Recognition of small dimensional representations of general linear groups

Kay Magaard, E.A. O'Brien and Ákos Seress

Abstract

Let G be isomorphic to a group H satisfying $\mathrm{SL}(d,q) \leq H \leq \mathrm{GL}(d,q)$ and let W be an irreducible $\mathbf{F}_q G$-module of dimension at most d^2. We present a Las Vegas polynomial-time algorithm which takes as input W and constructs a d-dimensional projective representation of G.

1 Introduction

A major research topic over the past decade has been the development of efficient algorithms for the investigation of subgroups of $\mathrm{GL}(d, \mathbf{F}_q)$ where $\mathbf{F}_q$ is a finite field of size $q = p^f$. We refer to the recent survey [15] for background related to this work.

A particular focus is the development of algorithms to construct an isomorphism between an arbitrary representation of a classical group and its "standard" (or natural) representation.

In 2001, Kantor and Seress [9] proved that there is a Las Vegas algorithm that, given as input an arbitrary permutation or (projective) matrix representation G of an almost simple classical group H of Lie type of known characteristic, constructs an isomorphism between G and the natural projective representation of H. Their algorithm also constructs a new "nice" generating set S for G such that any element can be reached efficiently from S by a short *straight-line program*: an efficiently stored group word on S that evaluates to g. (For a formal definition and discussion of their significance, see [18, p. 10].)

In this paper, we present efficient algorithms to construct such an isomorphism for a *projective matrix representation* of degree at most d^2 of the general linear groups having natural module of dimension d. In the natural module, a "nice" generating set can be constructed using the efficient algorithms of [5] or [11]. Hence this work supplements

This work was supported in part by by the NSA, the Marsden Fund of New Zealand, and the NSF. 2000 *Mathematics Subject Classification.* Primary 20C20, 20C40. We thank the referee for helpful comments and suggestions.

that of [9], providing fast polynomial time reduction for the most commonly occurring irreducible representations of general linear groups.

An additional motivation for our algorithm is the recent work of Ryba [17]. He presents a polynomial-time Las Vegas algorithm that, given as input an odd defining characteristic absolutely irreducible representation of a finite Chevalley group, constructs its action on the adjoint module. Since the adjoint module of $\mathrm{GL}(d,q)$ has dimension at most d^2-1, a combination of Ryba's algorithm and ours can be used to construct the action on the natural module.

A similar program has been carried out for the alternating and symmetric groups in [2] and [3]. The algorithm of [2] constructs an isomorphism between an arbitrary permutation or matrix representation of A_n or S_n and the natural permutation representation on n points; in [3] a specialized, faster algorithm does the same for the deleted permutation module, which is the smallest dimensional matrix representation of these groups.

2 Background and main result

We now consider in more detail our task. Let $\mathrm{SL}(d,q) \leq H \leq \mathrm{GL}(d,q)$ with $q = p^f$. Let V denote the natural module of H and V^* is its dual module. Define the Frobenius map $\delta : \mathrm{GL}(d,q) \mapsto \mathrm{GL}(d,q)$ by $(a_{i,j})^\delta = (a_{i,j}^p)$ for $(a_{i,j}) \in \mathrm{GL}(d,q)$.

Two representations ρ_1 and ρ_2 of H are *quasiequivalent* if there exists $\theta \in \mathrm{Aut}(H)$ such that ρ_1 is equivalent to $\theta\rho_2$.

Let H act on an irreducible $\mathbf{F}_q G$-module W of dimension at most d^2. For a discussion of such irreducible representations, see [12]. In particular, W is quasiequivalent to an irreducible section of $V \otimes V^{\delta^e}$, or $V^* \otimes V^{\delta^e}$ where $0 \leq e < f$.

The irreducible sections of $V \otimes V$ are the *symmetric* and *alternating* squares of V of dimension $d(d+1)/2$ and $d(d-1)/2$ respectively.

Consider $V^* \otimes V$ with basis $\{e_i \otimes e_j \mid 1 \leq i,j \leq d\}$ and let

$$w := \sum_{i=1}^{d} e_i \otimes e_i, \qquad U := \{\sum_{i,j} \alpha_{i,j} e_i \otimes e_j \mid \sum_{i=1}^{d} \alpha_{i,i} = 0\}, \qquad W_1 := U \cap \langle w \rangle.$$

The *adjoint module* of V is $W := U/W_1$. If $d \bmod p \equiv 0$ then W has dimension d^2-2, otherwise d^2-1.

The remaining irreducible representations of dimension at most d^2 are $V \otimes V^{\delta^e}$ and $V^* \otimes V^{\delta^e}$ where $0 < e < f$.

Our principal goal is an algorithm that, given as input such an irreducible representation W of H, constructs a d-dimensional projective representation of H.

Our algorithm assumes that we can construct random elements of a finite group G. Following the notation of [18, p. 24], an algorithm constructs an ε-uniformly distributed random element x of G if $(1-\varepsilon)/|G| < \mathrm{Prob}(x = g) < (1+\varepsilon)/|G|$ for all $g \in G$; if

$\varepsilon < 1/2$, then the algorithm constructs *nearly uniformly distributed* random elements of G. Babai [1] presents a black-box Monte Carlo algorithm to construct such elements in polynomial time. Another, more practical, option is the product replacement algorithm of Celler *et al.* [6], which also runs in polynomial time (see [16]). For a discussion of both algorithms, we refer the reader to [18, pp. 26-30].

Let ξ denote the cost of constructing a nearly uniformly distributed random element in a group G and let ρ_r denote the cost of a field operation in a finite field $\mathbf{F}_r$. Our main result is the following.

Theorem 2.1 *Let $d \geq 2$ and let $q = p^f$ be a prime power. Let $\mathrm{SL}(d,q) \leq H \leq \mathrm{GL}(d,q)$ where H has natural module V. Suppose that H is given as $G = \langle X \rangle$ acting irreducibly on a $\mathbf{F}_q$-vector space W of dimension $n \leq d^2$. Subject to Conjecture 4.9, given as input G, the value of d, and error probability $\epsilon > 0$, there is a polynomial-time Las Vegas algorithm that, with probability at least $1 - \epsilon$, sets up a data structure to construct the projective action of G on V in time $O(\xi d^2 \log q \log(1/\epsilon) + \rho_q d^9 \log^2 d \log^2 q)$. The time requirement to evaluate the image of $g \in G$ on V is $O(\xi + \rho_q d^8 \log q)$.*

We prove this theorem by exhibiting an algorithm with the stated complexity. We present the conjecture and evidence in its support in Section 4.3. Theorem 2.1 depends on Conjecture 4.9 *only* if W is the exterior square of V.

An $n \times n$ matrix over $\mathbf{F}_q$ requires $\Theta(d^4 \log q)$ space, so the running time of the algorithm, in terms of the input length N, is $O(\xi N + \rho_q N^{2.25} \log^2 N)$. We use the conventional estimate of $O(n^3)$ field operations for matrix multiplication; if it can be done in $O(n^\omega)$ field operations for some constant $\omega < 3$, then our algorithm runs in $O(\xi d^2 \log q + \rho_q d^{3+2\omega} \log^2 d \log^2 q)$ time.

In Section 3 we outline the basic algorithm common to all of the cases, and in Section 4 estimate the costs of common steps. We then study each representation in turn, and finally report briefly on an implementation of the algorithms in Magma [4].

3 The general strategy

Let $q = p^f$ be a prime power and let r be a prime. Recall from [14] that r is a *primitive prime divisor* of $q^d - 1$ if $r \mid q^d - 1$ but r does not divide $q^e - 1$ for $e < d$. We use the notation $\mathrm{ppd}(q;d)$ to describe such an r.

By a theorem of Zsigmondy (see [14]), $\mathrm{ppd}(q;d)$ primes exist for all q and d except when $q = 2, d = 6$ and q is a Mersenne prime, $d = 2$. To cover these exceptional cases, we call 9 a $\mathrm{ppd}(2;6)$ prime and 4 a $\mathrm{ppd}(q;2)$ prime for Mersenne prime q.

Recall that H has natural module V and is given as $G \leq \mathrm{GL}(W)$ where $W = \mathbf{F}_q^n$. Let $s \in H$ and assume that $|s|$ is divisible by some $\mathrm{ppd}(q;d)$. Hence s is a power of a Singer cycle and has d one-dimensional eigenspaces $\langle e_1 \rangle, \langle e_2 \rangle, \ldots, \langle e_d \rangle$ in $V \otimes \mathbf{F}_{q^d}$. Let $\sigma = \delta^f$ be the Frobenius map of $\mathrm{GL}(d, q^d)$ whose fixed points contain H. Thus σ centralizes $\langle s \rangle$ and so σ transitively permutes the eigenspaces of s acting on $V \otimes \mathbf{F}_{q^d}$.

Consequently, we can index the eigenspaces $\langle e_i \rangle$ of s and choose the eigenvectors e_i within the eigenspaces in such a way that $e_i^\sigma = e_{i+1}$ where the index is computed modulo d. If $e_1 s = \omega e_1$, then $e_i s = \omega^{q^{i-1}} e_i$ for $i \leq d$.

Our goal is to write the action A of an arbitrary $g \in G$ on $V \otimes \mathbf{F}_{q^d}$, in the basis $e_1, \ldots, e_d$. If W is either the alternating or symmetric square representation, then we also compute the base change matrix B between a particular $\mathbf{F}_q$-basis $b_1, \ldots, b_d$ and $e_1, \ldots, e_d$. Hence, we also learn the action BAB^{-1} of g on the natural module V. In the other cases, the action of G in a suitable $\mathbf{F}_q$-basis is recovered using the algorithm of [8] which has complexity $O(\rho_q d^5 \log^2 d)$.

We now summarise the algorithm `Decompose` to construct the projective action of G on V.

1. Find, by random search, $s \in G$ which satisfies the following:
   - if W is not the adjoint module, then s has n one-dimensional eigenspaces;
   - if W is the adjoint module, then s has $d^2 - d$ one-dimensional eigenspaces and an $n - (d^2 - d)$-dimensional eigenspace for the eigenvalue 1;
   - $|s|$ is divisible by some $\mathrm{ppd}(q; d)$ prime.
2. Construct a basis $\mathcal{B}_0$ consisting of eigenvectors for the action of s on $W \otimes \mathbf{F}_{q^d}$.
3. Label the elements of $\mathcal{B}_0$ by ordered pairs (i, j) with $1 \leq i, j \leq d$. The labels are found by discovering a sufficient number of algebraic dependencies among the eigenvalues. This labelling must be commensurate with the basis $e_1, \ldots, e_d$ and consistent with the action of σ.
4. From the eigenspace labelled with (i, j), compute the vector corresponding to $e_i \otimes e_j$.

Steps 1 to 4 create the data structure described in Theorem 2.1 and are applied *once*. To obtain the image of $g \in G$, we apply the following step.

5. First write g in the basis $\mathcal{B}_0$; then compute the action of g on $V \otimes \mathbf{F}_{q^d}$ in the basis $e_1, e_2, \ldots, e_d$; finally rewrite with respect to the basis $b_1, b_2, \ldots, b_d$ for the natural module V.

4 The common steps

We first discuss costs associated with the extension field $\mathbf{F}_{q^d}$. Since Step 1 is common to all representations, we next discuss it in detail and estimate its cost. To conclude this section, we discuss base change between the bases $b_1, \ldots, b_d$ of V and $e_1, \ldots, e_d$ of $V \otimes \mathbf{F}_{q^d}$, and consider properties of matrices written with respect to the latter basis.

4.1 Construction of the extension field

Since we work in $\mathbf{F}_{q^d}$, as a preprocessing step we construct that field.

Lemma 4.1 *The extension field $\mathbf{F}_{q^d}$ can be constructed by a Las Vegas algorithm, in $O(\rho_q d^3 \log^2 d \log q)$ time. The cost of a field operation in $\mathbf{F}_{q^d}$ is $O(\rho_q d^2 \log^2 d)$. Taking a square root of an element of $\mathbf{F}_{q^d}$ can be done in $O(\rho_q d^3 \log^2 d \log q)$ time.*

Proof: We construct $\mathbf{F}_{q^d}$ as an extension of $\mathbf{F}_q$. To do this, we search for monic polynomials of degree d from the polynomial ring $\mathbf{F}_q[x]$ until we find an irreducible $f(x)$. With probability at least $1/(d+1)$ (see [14]), a random polynomial of degree d defined over $\mathbf{F}_q[x]$ is irreducible. That a polynomial of degree d is irreducible over $\mathbf{F}_q[x]$ can be decided in $O(\rho_q d^2 \log d \log\log d \log q)$ time by a Las Vegas algorithm (see [19, 14.14]).

The elements of $\mathbf{F}_{q^d}$ are the residue classes of $\mathbf{F}_q[x]$ modulo $f(x)$, and they can be represented by the polynomials of degree at most $d-1$. We summmarize the cost of field operations in $\mathbf{F}_{q^d}$.

- $O(\rho_q d)$ for addition and $O(\rho_q d \log d \log\log d)$ for multiplication and division (see [19, 8.23, 9.6]);
- $O(\rho_q d^2 \log d \log\log d)$ for taking inverses; we compute the inverse of $a(x) \in \mathbf{F}_{q^d}$ by writing the greatest common divisor 1 of $f(x)$ and $a(x)$ in the form $1 = f(x)g(x) + a(x)h(x)$ by a Euclidean algorithm of length at most d, and then taking $h(x)$ as the inverse.
- Taking square roots of some $c \in \mathbf{F}_{q^d}$ can be done by factorizing the polynomial $x^2 - c$, in time $O(\rho_{q^d} \log q^d) = O(\rho_q d^3 \log^2 d \log q)$ (see [19, 14.14]). □

We use the following result, obtained in [2, 4.6] as an application of [19, 14.19].

Lemma 4.2 *The distinct linear factors of some $g(x) \in \mathbf{F}[x]$ of degree n can be computed by a Las Vegas algorithm, in $O(\rho_{\mathbf{F}} n \log^2 n \log(n|\mathbf{F}|) \log\log n)$ time.*

4.2 Step 1 of the general strategy: Finding s

We now discuss the search for random $s \in G$ which satisfies the conditions associated with Step 1 of `Decompose`. If an s satisfying the eigenvalue condition is found, then we check that $|s|$ is divisible by a $\mathrm{ppd}(q;d)$ prime as follows. If (d,q) equals $(6,2)$ or $(2,p)$ with p Mersenne, then define $m := 21$ and $m := p-1$, respectively; otherwise

$$m := q^d \prod_{j|d, j\neq d} \frac{d}{j}(q^j - 1). \tag{1}$$

Now $|s|$ is divisible by a $\mathrm{ppd}(q;d)$ prime if and only if $s^m \neq 1$. We decide this by raising the eigenvalues of s to the m-th power.

Lemma 4.3 *We can decide if $s \in G$ satisfies the conditions of Step 1 by a Las Vegas algorithm in $O(\rho_q d^6 \log q)$ time.*

Proof: The characteristic polynomial $c(x)$ of s can be computed using the algorithm of [10] in $O(\rho_q n^3 \log n) = O(\rho_q d^6 \log d)$ time. By Lemma 4.2, the distinct linear factors of $c(x)$ in $\mathbf{F}_{q^d}$ are obtained by a Las Vegas algorithm, in $O(\rho_{q^d} n \log^2 n \log(nq^d) \log\log n) = O((\rho_q d^2 \log^2 d)(d^2 \log^2 d)(\log d + d \log q)(\log d)) = O(\rho_q d^5 \log^5 d \log q)$ time. If W is the adjoint module, then the 1-eigenspace of s can be computed in $O(\rho_q d^6)$ time. Raising the eigenvalues of s to the power m in (1) takes $O(n\rho_{q^d} \log(q^{d^{3/2}})) = O(\rho_q d^{11/2} \log^2 d \log q)$ time, using the trivial upper estimate $2\sqrt{d}$ for the number of divisors of d. □

We now derive a sufficient condition to identify a suitable element of H, and use this condition in Lemma 4.5 to estimate the number of random elements processed in our search.

Theorem 4.4 *Suppose that $(d, q) \neq (3, 4)$ and that the order of $s \in H$ is a multiple of $(q^d - 1)/(q - 1)$. If W is not the adjoint module, then s has n distinct eigenvalues in $\mathbf{F}_{q^d}$. If W is the adjoint module, then 1 is an eigenvalue of s with eigenspace of dimension $n - (d^2 - d)$, and s has $d^2 - d$ other eigenvalues.*

Proof: Let α be a primitive element of $\mathbf{F}_{q^d}$. If the order of s is a multiple of $(q^d - 1)/(q - 1)$, then the eigenvalues of s in $V \otimes \mathbf{F}_{q^d}$ are $\omega, \omega^q, \ldots, \omega^{q^{d-1}}$, where $\omega = \alpha^k$ for some divisor k of $q - 1$.

If W is the symmetric square of V, then the eigenvalues of s in $W \otimes \mathbf{F}_{q^d}$ are $\omega^{q^{i-1}+q^{j-1}}$, for $1 \leq i \leq j \leq d$. Suppose that $\omega^{q^{i_1-1}+q^{j_1-1}} = \omega^{q^{i_2-1}+q^{j_2-1}}$ for some $i_1 \leq j_1, i_2 \leq j_2$. Then $\alpha^{k(q^{i_1-1}+q^{j_1-1})} = \alpha^{k(q^{i_2-1}+q^{j_2-1})}$. If the exponents on both sides are less than $q^d - 1$, then they must be equal. This implies $j_1 = j_2$, and so $i_1 = i_2$. If $k(q^{i_1-1} + q^{j_1-1}) \geq q^d - 1$, then $k = q - 1$ and $i_1 = j_1 = d$, so the only remaining possibility is $2(q - 1)q^{d-1} = q^d - 1 + (q - 1)(q^{i_2-1} + q^{j_2-1})$. This simplifies to $q^{d-1} = q^{d-2} + q^{d-3} + \cdots + q + 1 + q^{i_2-1} + q^{j_2-1}$. If $q = 2$, then we further simplify to $1 = q^{i_2-1} + q^{j_2-1}$, a contradiction. If $q \geq 3$, then $q = 3$ and $i_2 = j_2 = 1$, otherwise the right-hand-side of the last equation is not divisible by q. But this also leads to a contradiction.

If W is the alternating square, then the eigenvalues of s in $W \otimes \mathbf{F}_{q^d}$ are $\omega^{q^{i-1}+q^{j-1}}$, for $1 \leq i < j \leq d$. Since all occur as eigenvalues for the symmetric square, they are distinct.

If W is the adjoint module, then the eigenvalues of s in $W \otimes \mathbf{F}_{q^d}$, different from 1, are of the form $\omega^{q^{i-1}-q^{j-1}}$, for $1 \leq i \leq j \leq d$, $i \neq j$. If $\omega^{q^{i_1-1}-q^{j_1-1}} = \omega^{q^{i_2-1}-q^{j_2-1}}$ for some $i_1 \neq j_1, i_2 \neq j_2$, then $\alpha^{k(q^{i_1-1}+q^{j_2-1})} = \alpha^{k(q^{i_2-1}+q^{j_1-1})}$. As in the symmetric square case, the only solution of this equation implies that $j_2 \in \{i_2, j_1\}$. Since $j_2 \neq i_2$, we must have $j_2 = j_1$, and so $i_1 = i_2$.

Now consider the case $W = V \otimes V^\tau$. If $\tau = \delta^e$ and $0 < e < f$, then the eigenvalues of s in $W \otimes \mathbf{F}_{q^d}$ are $\omega^{q^{i-1}+p^e q^{j-1}}$, for $1 \le i, j \le d$. Suppose that $\omega^{q^{i_1-1}+p^e q^{j_1-1}} = \omega^{q^{i_2-1}+p^e q^{j_2-1}}$ for some $1 \le i_1, j_1, i_2, j_2 \le d$. Then

$$\alpha^{k(q^{i_1-1}+p^e q^{j_1-1})} = \alpha^{k(q^{i_2-1}+p^e q^{j_2-1})}. \tag{2}$$

If one of $i_1 = i_2$ and $j_1 = j_2$ holds, then clearly the other equality holds as well. If the exponents on both sides of (2) are equal, then $q^{i_1-1} - q^{i_2-1} = p^e(q^{j_2-1} - q^{j_1-1})$. If $j_1 \neq j_2$, then the exponent of p in the prime factorization of the right-hand-side of this last equation is equal to e mod f, but on the left-hand-side the exponent is 0 mod f, a contradiction. Hence $j_i = j_2$, and so $i_1 = i_2$. If $j_1, j_2 \le d-2$, then both exponents in (2) are at most $k(q^{d-1} + p^e q^{d-3}) < (q-1)(q^{d-1} + q^{d-2}) \le q^d - 1$, so the exponents must be equal. But this implies $i_1 = i_2$ and $j_1 = j_2$.

If $d \ge 5$ then, given any solution of (2), there exists $m \in \{0, 1, 2, 3, 4\}$ so that $j_k + m - 1 \le 2d - 2$ and $j_k + m - 1 \notin \{d-1, d\}$ holds for $k = 1, 2$. Hence, raising (2) to the q^m-th power and replacing the terms q^{i_k+m-1}, q^{j_k+m-1} by $q^{i_k+m-1-d}, q^{j_k+m-1-d}$, respectively, in the case these exponents are greater than d, we obtain a solution of (2) with $j_1, j_2 \le d-2$. Therefore, for the original j_1, j_2 we have $j_1 + m - 1 \equiv j_2 + m - 1 \pmod{d}$, implying $j_1 = j_2$. Similarly, if $3 \le d \le 4$ then, given any solution of (2), there exists $m \in \{0, 1, 2\}$ so that $j_k + m - 1 \le 2d - 2$ and $j_k + m - 1 \neq d$ holds for $k = 1, 2$. Hence it is enough to consider solutions of (2) with $j_1, j_2 \le d-1$ and at least one of the exponents of α on both sides of (2) is greater than $q^d - 1$. However, if $j_1 \le d-1$ then $k(q^{i_1-1} + p^e q^{j_1-1}) > q^d - 1$ is possible if and only if $k = q-1, i_1 = d, j_1 = d-1$, and again $k(q^{i_1-1} + p^e q^{j_1-1}) < 2(q^d - 1)$. Hence the only case to consider is $(q-1)(q^{d-1} + p^e q^{d-2}) = (q-1)(q^{i_2-1} + p^e q^{j_2-1}) + q^d - 1$. Here the left-hand-side is divisible by p, so we must have $i_2 = 1$ and $p = 2$. Thus our equation is equivalent to $2^e(q^{d-2} - q^{j_2-1}) = (q^{d-1} + q - 2)/(q-1)$. Using that $2^e \ge 2$ and $j_2 \le d-2$ (because $j_2 \neq j_1$), the left-hand-side of the last equation is greater than the right-hand-side, unless $d = 3, q = 4, j_2 = 1$. This is the exception identified in the statement of the lemma. The last subcase is $d = 2$. Since $i_1 \neq i_2, j_1 \neq j_2$, $\alpha^{k(1+p^e q)} = \alpha^{k(q+p^e)}$. But this implies that $k(1 + p^e q - q - p^e) = k(q-1)(p^e - 1)$ is a multiple of $q^2 - 1$, a contradiction since it is not divisible by a $\mathrm{ppd}(p; 2f)$ prime.

Finally, if $W = V^* \otimes V^\tau$, then the eigenvalues of s in $W \otimes \mathbf{F}_{q^d}$ are $\omega^{-q^{i-1}+p^e q^{j-1}}$, for $1 \le i, j \le d$. Suppose that $\omega^{-q^{i_1-1}+p^e q^{j_1-1}} = \omega^{-q^{i_2-1}+p^e q^{j_2-1}}$ for some $1 \le i_1, j_1, i_2, j_2 \le d$. Then $\alpha^{k(q^{i_2-1}+p^e q^{j_1-1})} = \alpha^{k(q^{i_1-1}+p^e q^{j_2-1})}$. As in the previous case, the only solution is $i_1 = i_2, j_1 = j_2$ if $(d, q) \neq (3, 4)$. □

As stated, Theorem 4.4 identifies a sufficient condition for an element of H to be suitable; its statement can be readily adapted to G.

Lemma 4.5 *The expected sample size for Step* 1 *is* $O(1/4d^2 \ln q)$ *and the expected running time is* $O(\xi d^2 \log q + \rho_q d^8 \log^2 q)$.

Proof: By [14], the probability that the order of a random $s \in H$ is divisible by a $\mathrm{ppd}(q;d)$ number is greater than $1/2d$. If the order of s is divisible by a $\mathrm{ppd}(q;d)$, then $C_H(s)$ is the intersection of H with the cyclic subgroup generated by a Singer cycle; further the order of s is largest if and only if $C_H(s) = \langle s \rangle$. The probability that this occurs is greater than $1/(2\ln q^d)$, since $\varphi(k) > k/(2\ln k)$ for all $k > 2$ (see [13, p. 227]). Hence the probability that a random $s \in H$ satisfies the order requirement of Theorem 4.4 is greater than $1/(4d^2 \ln q)$. Combining with Lemma 4.3, we obtain the statement of this lemma. □

4.3 The base change matrix

Let $s \in H$ have order a multiple of $(q^d-1)/(q-1)$. Consider a basis $b_1, b_2, \ldots, b_d$ of the natural module V of H in which s is represented by the matrix

$$S = \begin{pmatrix} 0 & 0 & \cdots & 0 & a_0 \\ 1 & 0 & \cdots & 0 & a_1 \\ 0 & 1 & \cdots & 0 & a_2 \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & 1 & a_{d-1} \end{pmatrix}. \tag{3}$$

This is the rational canonical form, or companion matrix, of the *left* action of s in V^*, but the b_i are row vectors, a basis for the *right* action of H on V. The characteristic polynomial of s is $x^d - \sum_{i=0}^{d-1} a_i x^i$, with the entries a_i in the last column of S.

Let s have eigenvalues $\omega, \omega^q, \ldots, \omega^{q^{d-1}}$ in $V \otimes \mathbf{F}_{q^d}$, and corresponding eigenvectors $e_1, e_2, \ldots, e_d$ satisfying $e_i^\sigma = e_{i+1}$. We now determine the base change matrix between the bases $b_1, b_2, \ldots, b_d$ of V and $e_1, e_2, \ldots, e_d$ of $V \otimes \mathbf{F}_{q^d}$, and obtain structural information about the matrix of an element of H in the basis $e_1, e_2, \ldots, e_d$.

Lemma 4.6 *The base change matrix between $b_1, b_2, \ldots, b_d$ and $e_1, e_2, \ldots, e_d$ has the form*

$$B = \begin{pmatrix} \mu & \mu\omega & \mu\omega^2 & \cdots & \mu\omega^{d-1} \\ \mu^q & (\mu\omega)^q & (\mu\omega^2)^q & \cdots & (\mu\omega^{d-1})^q \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \mu^{q^{d-1}} & (\mu\omega)^{q^{d-1}} & (\mu\omega^2)^{q^{d-1}} & \cdots & (\mu\omega^{d-1})^{q^{d-1}} \end{pmatrix}.$$

for some non-zero $\mu \in \mathbf{F}_{q^d}$.

Proof: Let $e_1 = (\alpha_1, \ldots, \alpha_d)$ in the basis $b_1, b_2, \ldots, b_d$. Then $e_1 s = \omega e_1$ implies that

$$(\alpha_1, \ldots, \alpha_d)S = (\omega\alpha_1, \ldots, \omega\alpha_d).$$

Also, by (3),

$$(\alpha_1, \ldots, \alpha_d)S = (\alpha_2, \ldots, \alpha_d, \beta),$$

with $\beta = \sum_{i=1}^{d} \alpha_i a_{i-1}$. Comparing the first $d-1$ corresponding entries in the two vectors on the right-hand-sides of these equations, we obtain $\alpha_i = \alpha_1 \omega^{i-1}$ for $2 \le i \le d$. Hence, with the notation $\mu := \alpha_1$, the first row of the base change matrix B is $(\mu, \mu\omega, \ldots, \mu\omega^{d-1})$. Since $e_i^\sigma = e_{i+1}$ for all i, the other rows of B can be obtained by taking the q-th power of entries in the previous row. □

Lemma 4.7 *Let $h \in H$ and let $A = (a_{ij})$ be the matrix of h in the basis $e_1, \ldots, e_d$. For $i, j \in \{1, \ldots, d\}$,*

$$a_{i+1,j+1} = a_{ij}^q$$

(where the index $d+1$ is interpreted as 1).

Proof: Let $c_1, c_2, \ldots, c_d$ be the column vectors of B^{-1}, where B is the base change matrix defined in Lemma 4.6. Then $e_1 c_1 = 1$ and for $2 \le j \le d$, $e_j c_1 = e_1^{\sigma^{j-1}} c_1 = 0$. Applying σ^{i-1} to these equations, we obtain $e_1^{\sigma^{i-1}} c_1^{\sigma^{i-1}} = e_i c_1^{\sigma^{i-1}} = 1$ and $e_j c_1^{\sigma^{i-1}} = 0$ for all $j \ne i$. Hence $c_i = c_1^{\sigma^{i-1}}$ for $2 \le i \le d$.

If M is the matrix of h in the basis $b_1, b_2, \ldots, b_d$, then its entries are in $\mathbf{F}_q$, and $A = BMB^{-1}$. Thus $a_{ij} = e_i M c_j$ and $a_{ij}^q = e_i^\sigma M^\sigma c_j^\sigma = e_{i+1} M c_{j+1} = a_{i+1,j+1}$. □

Lemma 4.8 *Let $h \in H$, and let $A = (a_{ij})$ be the matrix of h in the basis $e_1, \ldots, e_d$.*

(a) For $i, j \in \{1, \ldots, d\}$, $\mathrm{Prob}(a_{ij} = 0) < 4/q^d$. If $q \ge 3$ then $\mathrm{Prob}(a_{ij} = 0) < 2/q^d$.

(b) $\mathrm{Prob}(\text{all } a_{ij} \ne 0) > 5/8$.

Proof: Recall that the entries of A lie in $\mathbf{F}_{q^d}$, not $\mathbf{F}_q$.

(a) By Lemma 4.7, the first row of A determines uniquely the other rows of A, and $a_{ij} = 0$ if and only if $a_{1,j-i+1} = 0$ (if $j - i + 1 \le 0$ then define $a_{1,j-i+1}$ as $a_{1,j-i+1} := a_{1,d+j-i+1}$). There are q^{d^2-d} vectors of length d over $\mathbf{F}_{q^d}$ with a 0 entry in position $j - i + 1$. Not all of these vectors can occur as the first row of a matrix for $h \in \mathrm{GL}(d, q)$, so

$$\mathrm{Prob}(a_{1,j-i+1} = 0) < \frac{q^{d^2-d}}{|\mathrm{GL}(d,q)|} = \frac{1}{q^d} \prod_{k=1}^{d} \frac{1}{1 - 1/q^k}.$$

Observe that

$$\prod_{k=1}^{d} (1 - \frac{1}{q^k}) > (1 - \frac{1}{q})(1 - \sum_{k=2}^{d} \frac{1}{q^k}) > (1 - \frac{1}{q})(1 - \frac{1}{q^2} \frac{1}{1 - 1/q}) = \frac{q^2 - q - 1}{q^2}.$$

This last fraction is greater than $1/2$ if $q \ge 3$ and is equal to $1/4$ if $q = 2$, implying both claims of (a).

(b) By Lemma 4.7, it is enough to estimate the probability that the first row of A consists of all non-zero entries. By (a), if $q \geq 3$, then the probability that the first row of A contains a zero entry is less than $2d/q^d \leq 1/4$ (not considering the solvable case $(d, q) = (2, 3)$). Similarly, if $q = 2$ and $d \geq 6$, then the probability that the first row of A contains a zero entry is less than $4d/2^d \leq 3/8$. For $3 \leq d \leq 5$, we counted using GAP [7] the exact number of $h \in \mathrm{GL}(d, 2)$ with a zero entry in its matrix A. □

In Section 6.2, we require the following conjecture.

Conjecture 4.9 *Let h be a random element of H and let $A = (a_{ij})$ be the matrix of h in the basis $e_1, \ldots, e_d$. The probability that a principal 3×3 submatrix of A has determinant 0 is less than c/q^d for an absolute constant c, unless $d = 3m$ and the principal submatrix is from rows and columns $i, m+i, 2m+i$ for some i.*

In the exceptional case, the first row of the submatrix uniquely determines the other entries; in all other cases, there is an entry which is "independent" of the others. It appears that the value of this entry is roughly uniformly distributed in $\mathbf{F}_{q^d}^*$ and only one value makes the determinant 0. The conjecture and these observations are supported by experiment: we considered large random samples of elements from the alternating square representation of $\mathrm{GL}(d, q)$ for $q = 2, 3$ and $d \leq 9$.

4.4 Avoiding division by zero

Let $g \in G$, let $A = (a_{ij})$ be the matrix of g in the basis $e_1, \ldots, e_d$, and let $K = (\kappa_{ij,k\ell})$ be the matrix of g in the basis $\mathcal{B}_0$ obtained in Step 3 of `Decompose`.

In Step 4, we determine the constants associated with our choice of basis. In doing so, we perform arithmetic operations on the entries of K; some of these operations are not possible if certain entries of A are 0. Lemma 4.8 implies that, with high probability all entries of A are non-zero and so we can perform the computations.

In Step 5, we compute the action A of arbitrary $g \in G$ in the basis $e_1, \ldots, e_d$. The algorithms of Sections 5-7 may not work if A has zero entries; however, if we multiply g by a random $m \in G$, then both m and gm are uniformly distributed (but not independent) random elements of G. Hence, with probability at least 1/4, all entries in the $d \times d$ matrices of m and gm are non-zero. If so, then we compute the action of m and gm in the basis $e_1, \ldots, e_d$, and obtain the action of g as the ratio of these two matrices.

Summarizing, in Sections 5-7 we may assume that, for $g \in G$, *all entries of its matrix A are non-zero.* In Section 6.2 we assume Conjecture 4.9 which implies that *all of the 3×3 submatrices of A not having the exceptional indices have non-zero determinant.*

5 The symmetric square representation

In this case, q must be odd. Suppose that $G \leq \mathrm{GL}(W)$ where W is the symmetric square of V, and $s \in G$ was constructed as described in Step 1 of `Decompose`. We now discuss Steps 2 – 5.

5.1 Labelling the basis

Let $l_i = \omega^{q^{i-1}}$, for $1 \leq i \leq d$, be the eigenvalues of s in its action on $V \otimes \mathbf{F}_{q^d}$. The eigenspaces of s on W are $\langle e_{i,j} = e_i \otimes e_j \rangle$ for $1 \leq i \leq j \leq d$, and $e_{i,j}s = \omega^{q^{i-1}+q^{j-1}} e_{i,j}$.

We *know* the set L of products $l_{i,j} := l_i l_j$ for $1 \leq i \leq j \leq d$. We identify the indices (i,j) corresponding to every element of L, and choose a basis $\mathcal{B}_0 = \{f_{i,j}\}$, $f_{i,j} \in \langle e_{i,j} \rangle$, using the following procedure.

1. We construct the orbits $\Omega_1, \ldots, \Omega_k$ of eigenvalues under the Frobenius map σ. If d is odd, then there are $(d+1)/2$ orbits of size d; otherwise there are $d/2$ orbits of size d and one of size $d/2$.

2. We identify the orbit Ω_m of size d which satisfies the following test: for all distinct pairs $\alpha, \beta \in \Omega_m$, $\gamma = \alpha \cdot \beta$ is a square, and one of the square roots of γ is in Ω_j for some $j \neq m$.

3. Now we have identified that Ω_m is the orbit of $l_{1,1}$ under σ. We choose an arbitrary element of Ω_m as $l_{1,1}$ and label $l_{j,j} = l_{1,1}\sigma^{j-1}$.

4. For $i \neq j$, we evaluate $\gamma = l_{i,i}l_{j,j}$; record $l_{i,j}$ as the one of $\pm\sqrt{\gamma}$ in L.

5. From each orbit of eigenvalues Ω, we pick an arbitrary $l_{i,j} \in \Omega$ and compute its eigenspace $\langle e_{i,j} \rangle$. We choose the vector $f_{i,j} \in \langle e_{i,j} \rangle$ whose first non-zero coordinate is equal to 1.

6. For the other elements $l_{i,j}^{\sigma^r} \in \Omega$, we compute $f_{i+r,j+r} := f_{i,j}^{\sigma^r}$.

Since we identified $\omega^2 = l_{1,1}$, we can compute ω and the base change matrix B of Section 4.3. We may choose $\mu := 1$ as the $(1,1)$ entry of B.

Lemma 5.1 *The cost of this procedure is* $O(\rho_q d^9 \log^2 d \log q)$.

Observe that the largest exponent for d comes from Step 5, where we compute d eigenvectors at a cost $O(\rho_{q^d} n^3) = O(\rho_q d^8 \log^2 d)$ for each.

5.2 Determining the constants

The outcome of the last step is the following: for $1 \le i \le j \le d$, we now know

$$f_{i,j} = c_{i,j} e_{i,j}$$

for some $c_{i,j} \in \mathbf{F}_{q^d}$. We now describe a procedure to determine these constants.

Our choice of $f_{i,j}$ implies that $c_{i+1,j+1} = c_{i,j}^q$. Since we can multiply all basis vectors by the same scalar, we may *assume* that $c_{1,1} = 1$ (implying $c_{2,2} = c_{3,3} = \cdots = c_{d,d} = 1$). We have no further control over the $c_{i,j}$. Our procedure must compute their values.

We choose random $g \in G$ and compute the matrix $K = (\kappa_{ij,kl})$ representing the action of g on W with respect to the basis $\mathcal{B}_0 = \{f_{i,j}\}$. Let $A = (a_{ij})$ be the matrix of g in the basis $e_1, \ldots, e_d$. *A priori*, we do not know the entries of A, but we may assume that all entries of A are non-zero.

If $e_i g = \sum_{j=1}^d a_{ij} e_j$, then

$$(e_i \otimes e_j)g = (\sum_{s=1}^d a_{is} e_s) \otimes (\sum_{t=1}^d a_{jt} e_t)$$

and so

$$c_{ij} e_{ij} g = f_{ij} g = \sum \kappa_{ij,st} f_{st} = \sum \kappa_{ij,st} c_{st} e_{st}.$$

The basic equation for $\kappa_{ij,k\ell}$ is

$$\kappa_{ij,k\ell} \;=\; \frac{c_{ij}}{c_{k\ell}}(a_{ik}a_{j\ell} + a_{i\ell}a_{jk}(1-\delta_{kl})). \tag{4}$$

where $\delta_{kl} = 1$ if $k = l$, otherwise 0.

By choosing specific values of the indices, this equation allows us to readily deduce the following formulas.

(i) $\kappa_{ii,jj} = \frac{c_{ii}}{c_{jj}} a_{ij}^2 = a_{ij}^2$.

(ii) $\kappa_{ii,ii} = a_{ii}^2$.

Hence the values a_{ij} are determined up to a sign ambiguity.

(iii) $\kappa_{ij,jj} = \frac{c_{ij}}{c_{jj}} a_{ij} a_{jj} = c_{ij} a_{ij} a_{jj}$.

(iv) $\kappa_{ii,ij} = 2\frac{c_{ii}}{c_{ij}} a_{ii} a_{ij} = \frac{2}{c_{ij}} a_{ii} a_{ij}$.

We square (iii) and substitute (i), (ii) to get

(v) $c_{ij}^2 = \frac{\kappa_{ij,jj}^2}{\kappa_{jj,jj}\kappa_{ii,jj}} c_{jj} c_{ii} = \frac{\kappa_{ij,jj}^2}{\kappa_{jj,jj}\kappa_{ii,jj}}$.

Taking the product and ratio of (iv) and (iii), and using (i) and (v), we get

(vi) $a_{ii}a_{jj} = \frac{\kappa_{ij,jj}\kappa_{ii,ij}}{2\kappa_{ii,jj}}$.

(vii) $\frac{a_{ii}}{a_{jj}} = \frac{\kappa_{ii,ij}\kappa_{ij,jj}}{2\kappa_{jj,jj}\kappa_{ii,jj}}$.

Taking $j = k$, Equation (4) gives

$$\kappa_{ij,j\ell} = \frac{c_{ij}}{c_{j\ell}}(a_{ij}a_{j\ell} + a_{i\ell}a_{jj}).$$

Rearranging and multiplying by a_{ii}/a_{jj}, we get

(viii) $\frac{c_{ij}}{c_{j\ell}}a_{i\ell}a_{ii} = \frac{a_{ii}}{a_{jj}}(\kappa_{ij,j\ell} - c_{ij}^2\frac{a_{ij}}{c_{ij}}\frac{a_{j\ell}}{c_{j\ell}})$.

Using (iv)–(vii), the right-hand-side of (viii) can be expressed as a function of the $\kappa_{ij,k\ell}$.

We now describe the procedure to extract the c_{ij}.

1. If d is even, let $m = d/2$, else $m = d$.
2. We evaluate (viii) for $i = 1, j = 2, \ldots m, \ell = (2j-1)$ to obtain
$$\frac{c_{1j}}{c_{j,2j-1}}a_{1,2j-1}a_{11}$$
as a function of the $\kappa_{ij,k\ell}$.
3. Since $c_{1j}/c_{j,2j-1} = c_{1j}^{(1-q^{j-1})}$ is a power of c_{1j} with even exponent, by (v) we can express it as a function of the $\kappa_{ij,k\ell}$.
4. Thus we obtain $a_{1,2j-1}a_{11}$ as a function of the $\kappa_{ij,k\ell}$, and (iv) now yields the value of $c_{1,2j-1}$ without ambiguity.
5. If d is even, then we must compute the values $c_{1,2j}$. By (v), we know c_{12}^2, but c_{12} cannot be computed without ambiguity: writing $g \in G$ in the basis $e_1, \ldots, e_d$ or in $e_1, -e_2, e_3, \ldots, -e_d$ yields the same action on W. Hence we choose an arbitrary square root of c_{12}^2 from (v) to obtain c_{12}.
6. Evaluating (viii) with $i = 1$, $j = 2$, $\ell = 4, 6, \ldots, d$, we obtain $a_{1\ell}a_{11}$ as a function of c_{12}, $c_{2\ell} = c_{1,\ell-1}^q$, and the $\kappa_{ij,k\ell}$ without further ambiguity (apart from the sign of c_{12}).
7. Now (iv) yields the value of $c_{1\ell}$. The values c_{ij}, for $i > 1$, can be obtained as $c_{ij} = c_{1,j-i+1}^{q^{i-1}}$.

Lemma 5.2 *The cost of this procedure is* $O(\xi + \rho_q d^8 \log^2 d)$.

The most expensive part of this procedure is writing g with respect to the basis $\mathcal{B}_0 = \{f_{i,j}\}$.

5.3 Evaluating images

Assume we now wish to carry out Step 5 of `Decompose`: namely, we want to construct the projective image of an arbitrary $g \in G$.

We first compute the matrix $K = (\kappa_{ij,kl})$ representing the action of g on W with respect to the basis $\mathcal{B}_0 = \{f_{i,j}\}$. We then compute the a_{ij} in terms of the $\kappa_{rs,tu}$: for given i, j, we compute $a_{11}a_{ij} = (a_{11}/a_{ii})(a_{ii}a_{jj})$ using (vii) and (iv). Hence the a_{ij} values are recovered up to a scalar multiple a_{11}. By (ii), the value of a_{11} can be computed only up to a sign ambiguity.

Lemma 5.3 *The cost of this procedure is* $O(\rho_q d^8 \log^2 d)$.

If necessary, we compute the images of m and gm for random $m \in G$ as discussed in Section 4.4.

6 The alternating square representation

Suppose that $G \leq \mathrm{GL}(W)$ where W is the alternating square of V, and $s \in G$ is constructed as described in Step 1 of `Decompose`. We now discuss Steps 2 – 5.

6.1 Labelling the basis

Let $l_i = \omega^{q^{i-1}}$, for $1 \leq i \leq d$, be the eigenvalues of s in its action on $V \otimes \mathbf{F}_{q^d}$. A basis of W is the set of vectors $e_{i,j} = e_i \wedge e_j$ for $1 \leq i < j \leq d$, and $e_{i,j}s = \omega^{q^{i-1}+q^{j-1}} e_{i,j}$.

We *know* the set L of products $l_i l_j$ for $1 \leq i \leq j \leq d$. We identify the indices (i, j) corresponding to every element of L, and choose a basis $\mathcal{B}_0 = \{f_{i,j}\}$, $f_{i,j} \in \langle e_{i,j} \rangle$, using the following procedure.

1. We construct the orbits of the eigenvalues under the Frobenius map σ.
2. We choose an orbit of length d, and label an element of this orbit as $l_1 l_2$. Taking q-th powers determines $l_2 l_3$ and $l_3 l_4$.
3. We identify $l_1 l_4 = (l_1 l_2)(l_3 l_4)/(l_2 l_3)$, and by taking q-th powers we identify $l_2 l_5$ and $l_3 l_6$.
4. We identify $l_1 l_6 = (l_1 l_4)(l_3 l_6)/(l_3 l_4)$, and by taking q-th powers we identify $l_2 l_7$ and $l_3 l_8$.
5. In the same manner, we identify $l_1 l_{2i}$ for $i \leq d/2$. It remains to identify $l_1 l_{2i+1}$.
6. Assume $d = 2k + 1$. By taking q-th powers, we identify $l_2 l_{2k+1}$ and $l_3 l_1$. We compute $(l_1 l_3)/(l_1 l_2)^{q-1} = \omega^2$. From ω^2, we compute $\omega^{q^{i-1}+q^{j-1}}$ for all i, j. If these are not elements of L, then we select another L-orbit of length d and pick another candidate for $l_1 l_2$.

7. Assume $d = 2k$. We know that $(l_1 l_2)(l_3 l_4) = (l_1 l_3)^{q+1}$. We choose an orbit of length d in L and an element x of this orbit such that $x^{q+1} = (l_1 l_3)^{q+1}$. We compute $x/(l_1 l_2)^{q-1}$. If our choice of x is valid as $l_1 l_3$, then this last ratio is ω^2. From ω^2, we compute $\omega^{q^{i-1}+q^{j-1}}$ for all i, j. If these are not elements of L, then we select another L-orbit of length d and pick another candidate for $l_1 l_3$. The case $d = 4$ is exceptional: one orbit has size 4, the only other has size 2, and we choose x from the orbit of length 2.

8. From each orbit Ω of eigenvalues, we pick an arbitrary $l_{i,j} \in \Omega$ and compute its eigenspace $\langle e_{i,j} \rangle$. We choose the vector $f_{i,j} \in \langle e_{i,j} \rangle$ whose first non-zero coordinate is equal to 1.

9. For the other elements $l_{i,j}^{\sigma^r} \in \Omega$, we compute $f_{i+r,j+r} := f_{i,j}^{\sigma^r}$.

During this procedure, we computed ω^2, so we can compute the base change matrix B of Section 4.3.

Lemma 6.1 *The cost of this procedure is* $O(\rho_q d^9 \log^2 d \log q)$.

6.2 Determining the constants

The outcome of the last step is the following: for $1 \leq i < j \leq d$, we now know

$$f_{i,j} = c_{i,j} e_{i,j}$$

for some $c_{i,j} \in \mathbf{F}_{q^d}$. We now describe a procedure to determine these constants.

Our choice of $f_{i,j}$ implies that $c_{i+1,j+1} = c_{i,j}^q$. Since we can multiply all basis vectors by the same scalar, we may assume that $c_{1,2} = 1$ (implying $c_{2,3} = c_{3,4} = \cdots = c_{d-1,d} = 1$). We have no further control over the $c_{i,j}$. Our procedure must compute their values.

We choose random $g \in G$ and compute the matrix $K = (\kappa_{ij,kl})$ representing the action of g on W with respect to the basis $\{f_{i,j}\}$.

If $e_i g = \sum_{j=1}^d a_{ij} e_j$, then

$$(e_i \otimes e_j)g = (\sum_{s=1}^d a_{is} e_s) \wedge (\sum_{t=1}^d a_{jt} e_t)$$

and so

$$c_{ij} e_{ij} g = f_{ij} g = \sum \kappa_{ij,st} f_{st} = \sum \kappa_{ij,st} c_{st} e_{st}.$$

The basic equation for $\kappa_{ij,k\ell}$ is

$$\kappa_{ij,k\ell} = \frac{c_{ij}}{c_{k\ell}} (a_{ik} a_{j\ell} - a_{i\ell} a_{jk}). \tag{5}$$

Underpinning our procedure is the observation that, in dimension 3, the exterior square is the dual of the natural module. This means that for distinct triples i, j, k, the matrices

$$B_{ijk} = \begin{pmatrix} a_{ii} & \frac{c_{ik}}{c_{jk}}a_{ij} & \frac{c_{ij}}{c_{jk}}a_{ik} \\ \frac{c_{jk}}{c_{ik}}a_{ji} & a_{jj} & \frac{c_{ij}}{c_{ik}}a_{jk} \\ \frac{c_{jk}}{c_{ij}}a_{ki} & \frac{c_{ik}}{c_{ij}}a_{kj} & a_{kk} \end{pmatrix}, \quad C_{ijk} = \begin{pmatrix} \kappa_{jk,jk} & -\kappa_{ik,jk} & \kappa_{ij,jk} \\ -\kappa_{jk,ik} & \kappa_{ik,ik} & -\kappa_{ij,ik} \\ \kappa_{jk,ij} & -\kappa_{ik,ij} & \kappa_{ij,ij} \end{pmatrix} \tag{6}$$

satisfy $B_{ijk}C_{ijk} = \det(B_{ijk}) \cdot I$, because by (5) the entries of C_{ijk} are the appropriate 2×2 minors of B_{ijk}. This implies $\det(B_{ijk}C_{ijk}) = \det(B_{ijk}) \cdot \det(C_{ijk}) = \det(B_{ijk})^3$. Moreover, if C_{ijk} is invertible, then $B_{ijk} = \sqrt{C_{ijk}}C_{ijk}^{-1}$; so we can compute B_{ijk} up to a sign ambiguity. Conjecture 4.9 implies that we may assume that B_{ijk} is invertible unless $j = i + d/3$ and $k = i + 2d/3$.

We now describe the procedure to determine the constants $c_{i,j}$. Recall that we can assume $c_{12} = 1$ (and consequently $c_{i,i+1} = c_{12}^{q^{i-1}} = 1$ for all $i \leq d-1$).

1. For $k = 4, 6, \ldots, 2\lfloor d/2 \rfloor$, we compute B_{132} and B_{13k}. We choose the square root of determinants so that the common entries $B_{132}[1,1] = B_{13k}[1,1] = a_{11}$ and $B_{132}[2,2] = B_{13k}[2,2] = a_{33}$ are equal. Now $B_{132}[2,1] = a_{31}c_{32}/c_{12} = a_{31}$ and $B_{13k}[2,1] = a_{31}c_{3k}/c_{1k} = a_{31}c_{1,k-2}^{q^2}/c_{1k}$. Taking the ratio, we obtain $c_{1k} = c_{1,k-2}^{q^2}B_{132}[2,1]/B_{13k}[2,1]$ and so recursively we can compute $c_{14}, c_{16}, \ldots, c_{1,2\lfloor d/2 \rfloor}$. Since $c_{ij} = c_{i-1,j-1}^q$, we obtain all c_{ij} with $i < j$ and $j - i$ odd.

2. For $k = 5, 7, \ldots, 2\lceil d/2 \rceil - 1$, we compute B_{123} and B_{12k}, again choosing the square roots of determinants so that the common entries $B_{123}[1,1] = B_{12k}[1,1] = a_{11}$ and $B_{123}[2,2] = B_{12k}[2,2] = a_{22}$ are equal. Thus $B_{123}[2,1] = a_{21}c_{23}/c_{13} = a_{21}/c_{13}$ and $B_{12k}[2,1] = a_{21}c_{2k}/c_{1k} = a_{21}c_{1,k-1}^q/c_{1k}$. The ratio yields

$$c_{1k} = c_{13}c_{1,k-1}^q \frac{B_{123}[2,1]}{B_{12k}[2,1]}. \tag{7}$$

   Since $c_{1,k-1}$ is already known, we obtain c_{1k} as the value of c_{13} multiplied by a known quantity.

3. If d is odd, then $c_{13} = c_{1,d-1}^{q^2}$ is known, and we obtain all c_{1j} without ambiguity. The other c_{ij} are computed by the formula $c_{ij} = c_{i-1,j-1}^q$.

4. If d is even, then we compute B_{123} and B_{124}. Now (7), with the value $k = 4$, gives $c_{14} = c_{13}^{q+1}B_{123}[2,1]/B_{124}[2,1]$. Since c_{14} is already known, we obtain c_{13}^{q+1}. Using the formula $c_{ij} = c_{i-1,j-1}^q$, and the values of c_{1k} and c_{13}^{q+1}, we obtain c_{ij} with $i < j$ and $j - i$ even as $c_{ij} = c_{13}\overline{c_{ij}}$ if i is odd, and as $c_{ij} = {c_{13}}^{-1}\overline{c_{ij}}$ if i is even, where $\overline{c_{ij}}$ is a known element of $\mathbf{F}_{q^d}$.

   We do not compute the value for c_{13} because $(q+1)$-st roots cannot be computed in polynomial time. Instead we ignore both c_{13} and c_{13}^{-1}, and take $c_{ij} = \overline{c_{ij}}$ for all $i < j$ with $j - i$ even.

The justification is the following. Given $g \in G$, we want to compute the coefficients in the linear combination $e_m g = \sum_\ell a_{m\ell} e_\ell$. If we know c_{ij}, then $a_{m\ell}$ can be computed from an appropriate entry of some B_{ijk}. If we use $\overline{c_{ij}}$ in place of c_{ij}, then we obtain $c_{13} a_{m\ell}$ when both m and $m-\ell$ are odd; $a_{m\ell}/c_{13}$ when m is even and $m-\ell$ is odd; and $a_{m\ell}$ when $m-\ell$ is even. Thus, instead of $A = (a_{m\ell})$, we construct the conjugate of A by the diagonal matrix $\mathrm{Diag}(1, c_{13}, 1, c_{13}, \ldots, 1, c_{13})$. Equivalently, this is the action of g in the basis $e_1, e_2/c_{13}, e_3, e_4/c_{13}, \ldots, e_{d-1}, e_d/c_{13}$.

Lemma 6.2 *The cost of this procedure is* $O(\xi + \rho_q d^8 \log^2 d \log q)$.

6.3 Evaluating images

Given $g \in G$, we compute the matrix $K = (\kappa_{ij,kl})$ representing the action of g on W with respect to the basis $\mathcal{B}_0 = \{f_{i,j}\}$. We determine the first row of A by computing the matrices B_{12k} for $3 \le k \le d$. Following Lemma 4.7, the other entries of A can be obtained by taking q-th powers.

Lemma 6.3 *The cost of this procedure is* $O(\rho_q d^8 \log^2 d \log q)$.

If necessary, we compute the images of m and gm for random $m \in G$ as discussed in Section 4.4.

7 The adjoint representation

Let $V^* \otimes V$ have basis $B_0 := \{e_i \otimes e_j \mid 1 \le i, j \le d\}$. Recall from Section 2 our definition of the adjoint module $W := U/W_1$. Consider a basis for W which is the union of the set B_1 of $d^2 - d$ vectors $e_i \otimes e_j + W_1$, $i \ne j$, and a set B_2 of $d-1$ or $d-2$ vectors of the form $x + W_1$ for some $x \in \langle e_i \otimes e_i \mid 1 \le i \le d \rangle$. We can compute the subspaces $\langle e_i \otimes e_j + W_1 \rangle$; by choosing the vectors with first coordinates 1 from these subspaces, we construct B_1. Choosing the remaining basis vectors from the 1-eigenspace of s on W, we construct B_2.

For $g \in G$, let K_1 and K_2 be the matrices of g on $V^* \otimes V$ with basis B_0 and on W with basis $B_1 \cup B_2$, respectively. Independent of the particular choice of B_2, the $(d^2 - d) \times (d^2 - d)$ submatrices of K_1 and K_2, determined by the basis vectors $e_i \otimes e_j$ and $e_i \otimes e_j + W_1$, respectively, are identical.

7.1 Labelling the basis

Let $l_i = \omega^{q^{i-1}}$, for $1 \le i \le d$, be the eigenvalues of s in its action on $V \otimes \mathbf{F}_{q^d}$. The one-dimensional eigenspaces of s on W are $\langle e_{i,j} \rangle$ for $i \ne j$, and $e_{i,j} s = \omega^{q^{i-1} - q^{j-1}} e_{i,j}$.

We *know* the set L of products $l_{i,j} := l_i l_j^{-1}$ for $1 \le i, j \le d$, $i \ne j$. We identify the indices (i, j) corresponding to every element of L, and choose a basis for W using the following procedure.

1. We construct the $d-1$ orbits of elements of L under the Frobenius map σ.

2. We choose one of these orbits and declare an entry from this orbit as $l_{1,2}$. For $i \in \{2, \ldots, d-1\}$, we label $l_{i,i+1} = l_{i-1,i}^q$, and $l_{d,1} = l_{d-1,d}^q$.

3. For $k \in \{2, \ldots, d-1\}$, we perform the following:

   - We evaluate $\nu := l_{1,k} l_{1,2}^{q^{(k-1)}}$.
   - If $\nu \notin L$, then we choose a different orbit.
   - Otherwise we label ν as $l_{1,k+1}$.
   - For $j \in \{2, \ldots, d-k\}$, we identify $l_{j,j+k} := l_{j-1,j+k-1}^q$. Also we identify $l_{d+1-k,1} := l_{d-k,d}^q$.
   - For $j \in \{d+2-k, \ldots, d\}$, we identify $l_{j,j-d+k} := l_{j-1,j-d+k-1}^q$.

4. For each of the $d-1$ orbits on L, we pick a representative $l_{i,j}$ and compute its eigenspace $\langle e_{i,j} \rangle$. We choose the vector $f_{i,j} \in \langle e_{i,j} \rangle$ whose first non-zero coordinate is equal to 1.

   Since we can assume that the first coordinate of each e_i is 1, the vector $f_{i,j}$ corresponds *precisely* to $e_i \otimes e_j$, not just to a scalar multiple of it. For other eigenvalues $l_{i,j}^{\sigma^r}$, we compute $f_{i+r,j+r} := f_{i,j}^{\sigma^r}$. Let $B_1 := \{f_{i,j} \mid 1 \leq i \neq j \leq d\}$.

5. We compute the 1-eigenspace of s and choose an arbitrary basis B_2 for it. Then $B_1 \cup B_2$ is a basis of W.

Lemma 7.1 *The cost of this procedure is* $O(\rho_q d^9 \log^2 d \log q)$.

Since the exponent of ω in $l_{i,j}$ is a multiple of $q-1$, we cannot compute ω in polynomial time. Hence we cannot compute the base change matrix B of Section 4.3, but instead use the algorithm of [8] to perform the final base change.

7.2 Evaluating images

Given $g \in G$, we compute the matrix representing the action of g on W with respect to the basis $B_1 \cup B_2$. Let K_1 be the matrix of g in the basis B_0, let $K_2 = (\kappa_{ij,kl})$ be the matrix of g in the basis $B_1 \cup B_2$, let $A = (a_{ij})$ be the matrix of g in the basis $\{e_i\}$, and let $A^* = (a_{ij}^*)$ be the matrix of $g\varphi$ in the basis $\{e_i\}$. Here φ is a graph automorphism, namely an inverse transpose map, but it is taken with respect to the basis $b_1, \ldots, b_n$ defined in Section 4.3. Thus A and A^* are not inverse transposes of each other. The goal is to recover a scalar multiple of A. The basic equation in K_1 is

$$\kappa_{ij,kl} = a_{ik}^* a_{jl}. \tag{8}$$

The $(d^2-d)\times(d^2-d)$ submatrix of K_2 indexed by B_1 is the corresponding submatrix of K_1, so we may use (8) for $i \neq j$ and $k \neq \ell$.

To determine $a_{11}^* a_{ks}$ for any s, k, we use the following equations.

$$\begin{aligned}
\kappa_{1k,1s} &= a_{11}^* a_{ks} \quad k, s \geq 2 \\
\kappa_{21,2s} &= a_{22}^* a_{1s}, \quad s \neq 2 \\
\kappa_{2k,21} &= a_{22}^* a_{k1}, \quad k \neq 2 \\
\kappa_{31,32} &= a_{33}^* a_{12} \\
\kappa_{32,31} &= a_{33}^* a_{21} \\
a_{11}^*/a_{22}^* &= \kappa_{1j,1\ell}/\kappa_{2j,2\ell} \quad \text{for some } j, \ell \notin \{1,2\} \\
a_{11}^*/a_{33}^* &= \kappa_{1j,1\ell}/\kappa_{3j,3\ell} \quad \text{for some } j, \ell \notin \{1,3\}
\end{aligned}$$

Lemma 7.2 *The cost of this procedure is* $O(\rho_q d^8 \log^2 d)$.

If necessary, we compute the images of m and gm for random $m \in G$ as discussed in Section 4.4.

8 The other representations

We now discuss the outstanding cases: namely, $W = V \otimes V^\tau$ and $W = V^* \otimes V^\tau$. Recall that $H \leq \mathrm{GL}(d,q)$ with $q = p^f$. The operation τ is to take p^e-th powers of the entries in the matrices representing the group elements, for some fixed positive $e < f$, and so W is now irreducible.

As before, we can assume that the first coordinate of the basis vector e_1 is 1; since all other e_i are obtained by the Frobenius map σ, their first coordinate is also 1. Consequently, the first coordinates of the d^2 vectors $e_i \otimes e_j$ are 1. These d^2 vectors form a basis of W. Hence, we can compute the one-dimensional subspaces $\langle e_i \otimes e_j \rangle$ and, by picking the vectors with first coordinates 1 from these subspaces, we choose the vectors $e_i \otimes e_j$.

8.1 Labelling the basis when $W = V \otimes V^\tau$

Let $l_i = \omega^{q^{i-1}}$, for $1 \leq i \leq d$, be the eigenvalues of s in its action on $V \otimes \mathbf{F}_{q^d}$. The eigenspaces of s on W are $\langle e_{i,j} = e_i \otimes e_j \rangle$ for $1 \leq i, j \leq d$, and $e_{i,j}s = \omega^{q^i + q^j p^e} e_{i,j}$.

We *know* the set L of products $l_{i,j} := l_i (l_j)^{p^e}$ for $1 \leq i, j \leq d$. We identify the indices (i,j) corresponding to every element of L, and choose a basis $\mathcal{B}_0 = \{f_{i,j}\}$, $f_{i,j} \in \langle e_{i,j} \rangle$, using the following procedure.

1. We construct the orbits of eigenvalues under the Frobenius map σ.

2. We choose one of these orbits and declare an entry λ from this orbit as $l_{1,1}$.

3. If there exists an eigenvalue ν where $\nu^{1+p^e} = \lambda^{1+qp^e}$ and $\lambda^{q+1}\nu^{-1}$ is also an eigenvalue, then we identify $l_{1,2}$ as ν and $\lambda^{q+1}\nu^{-1}$ as $l_{2,1}$. Otherwise we choose a new orbit.

4. For $i \in \{2, \ldots, d\}$, we label $l_{i,i} = l_{i-1,i-1}^q$,

5. For $i \in \{2, \ldots, d-1\}$, we label $l_{i,i+1} = l_{i-1,i}^q$, and $l_{d,1} = l_{d-1,d}^q$.

6. For $i \in \{2, \ldots, d-1\}$, we label $l_{i+1,1} := l_{1,1}l_{i,i}l_{i+1,i+1}/l_{i,i+1}l_{1,i}$ and $l_{1,i+1} := l_{1,1}l_{i+1,i+1}/l_{i+1,1}$.

7. The remaining values $l_{i,j}$ can now be identified by taking q-th powers of the already labelled elements of their orbits.

8. For each orbit on L, we pick a representative $l_{i,j}$ and compute its eigenspace $\langle e_{i,j} \rangle$. We choose the vector $f_{i,j} \in \langle e_{i,j} \rangle$ whose first non-zero coordinate is equal to 1.

   Since we can assume that the first coordinate of each e_i is 1, the vector $f_{i,j}$ corresponds *precisely* to $e_i \otimes e_j$, not just to a scalar multiple of it. For other eigenvalues $l_{i,j}^{\sigma^r}$, we compute $f_{i+r,j+r} := f_{i,j}\sigma^r$.

Lemma 8.1 *The cost of this procedure is* $O(\rho_q d^9 \log^2 d \log q)$.

We use the algorithm of [8] to perform the final base change.

8.2 Evaluating images when $W = V \otimes V^\tau$

Given $g \in G$, we compute its $d^2 \times d^2$ matrix $K = (\kappa_{ij,kl})$ in the basis $\{e_i \otimes e_j\}$. Let $A = (a_{ij})$ be the matrix of g in the basis $\{e_i\}$. The goal is to recover a scalar multiple of A. The basic equation for $\kappa_{ij,k\ell}$ is

$$\kappa_{ij,k\ell} = a_{ik}a_{j\ell}^{p^e}.$$

We choose an arbitrary non-zero entry $\kappa_{i_0 j_0, k_0 \ell_0}$ in K. For this fixed j_0, ℓ_0, the matrix with (i,k) entry $\kappa_{ij_0,k\ell_0} = a_{ik}a_{j_0\ell_0}^{p^e}$ is a projective image of g.

Lemma 8.2 *The cost of this procedure is* $O(\rho_q d^8 \log^2 d)$.

8.3 Labelling the basis when $W = V^* \otimes V^\tau$

Labelling the basis for this case is essentially identical to that described in Section 8.1.

8.4 Evaluating images when $W = V^* \otimes V^\tau$

Given $g \in G$, we compute its $d^2 \times d^2$ matrix $K = (\kappa_{ij,kl})$ in the basis $\{e_i \otimes e_j\}$. Let $A = (a_{ij})$ be the matrix of g in the basis $\{e_i\}$ and let $A^* = (a_{ij}^*)$ be the matrix of $g\varphi$ in the basis $\{e_i\}$ for a graph automorphism φ. The goal is to recover a scalar multiple of A^*. The basic equation for $\kappa_{ij,k\ell}$ is

$$\kappa_{ij,k\ell} = a_{ik}^* a_{j\ell}^{p^e}.$$

We choose an arbitrary non-zero entry $\kappa_{i_0 j_0, k_0 \ell_0}$ in K. For this fixed j_0, ℓ_0, the matrix with (i,k) entry $\kappa_{ij_0,k\ell_0} = a_{ik}^* a_{j_0\ell_0}^{p^e}$ is a projective image of g.

Lemma 8.3 *The cost of this procedure is* $O(\rho_q d^8 \log^2 d)$.

9 Implementation and performance

We have implemented our algorithms in Magma. We use the algorithm of [6] to generate random elements and in Step 1 choose a sample of size $4d^2$.

Table 1: Performance of implementation for some groups

d	q	G	Setup	Image
5	7^{10}	Symmetric square	1.0	0.02
		Alternating square	0.3	0.01
		Adjoint	1.6	0.05
		$V \otimes V^\tau$	1.5	0.05
10	5^6	Symmetric square	15.7	0.4
		Alternating square	6.0	0.2
		Adjoint	34.5	1.1
		$V \otimes V^\tau$	73.6	1.0
15	3^2	Symmetric square	50.1	1.7
		Alternating square	23.4	1.3
		Adjoint	140.2	5.2
		$V \otimes V^\tau$	150.7	1.2

The computations reported in Table 1 were carried out using Magma V2.13 on a Pentium IV 2.8 GHz processor. Let G act as the named representation of $H := \mathrm{GL}(d,q)$. The first four steps of Algorithm `Decompose` provide the data structure of Theorem 2.1. Having computed this data, we can now compute the projective image

of $g \in G$. In the final two columns, we list the CPU times in seconds to set up the data structure, and to evaluate the image of a randomly chosen element of G.

References

[1] László Babai. Local expansion of vertex-transitive graphs and random generation in finite groups. *Theory of Computing*, (Los Angeles, 1991), pp. 164–174. Association for Computing Machinery, New York, 1991.

[2] Robert Beals, Charles R. Leedham-Green, Alice C. Niemeyer, Cheryl E. Praeger and Ákos Seress. A black-box group algorithm for recognizing finite symmetric and alternating groups I, *Trans. Amer. Math. Soc.* **355** (2003), 2097–2113.

[3] Robert Beals, Charles R. Leedham-Green, Alice C. Niemeyer, Cheryl E. Praeger and Ákos Seress. Constructive recognition of finite symmetric and alternating groups acting as matrix groups on their natural permutation modules, *J. Algebra* **292** (2005), 4–46.

[4] Wieb Bosma, John Cannon, and Catherine Playoust. The Magma algebra system I: The user language, *J. Symbolic Comput.* **24** (1997), 235–265.

[5] Peter A. Brooksbank. Constructive recognition of classical groups in their natural representation. *J. Symbolic Comput.* **35** (2003), 195–239.

[6] Frank Celler, Charles R. Leedham-Green, Scott H. Murray, Alice C. Niemeyer and E.A. O'Brien. Generating random elements of a finite group. *Comm. Algebra* **23** (1995), 4931–4948.

[7] The GAP Group. GAP – Groups, Algorithms, and Programming. Version 4.4.9; 2007. (http://www.gap-system.org)

[8] S.P. Glasby, C.R. Leedham-Green, and E.A. O'Brien. Writing projective representations over subfields. *J. Algebra* **295** (2006), 51–61.

[9] William M. Kantor and Ákos Seress. Black box classical groups. *Mem. Amer. Math. Soc.*, **149**, 2001.

[10] W. Keller-Gehrig. Fast algorithms for the characteristic polynomial. *Theoret. Comput. Sci.* **36** (1985), 309–317.

[11] C.R. Leedham-Green and E.A. O'Brien. Constructive recognition of classical groups in odd characteristic. Preprint 2007.

[12] Martin W. Liebeck. On the orders of maximal subgroups of the finite classical groups. *Proc. London Math. Soc.* (3), 50:426–446, 1985.

[13] M.B. Nathanson. *Elementary Methods in Number Theory.*

[14] P.M. Neumann and C.E. Praeger. A recognition algorithm for special linear groups. *Proc. London Math. Soc.* **65** (1992), 555–603.

[15] E.A. O'Brien. Towards effective algorithms for linear groups. *Finite Geometries, Groups and Computation*, (Colorado), pp. 163-190. De Gruyter, Berlin, 2006.

[16] Igor Pak. The product replacement algorithm is polynomial. In 41st *Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, 476–485, IEEE Comput. Soc. Press, Los Alamitos, CA, 2000.

[17] Alexander J.E. Ryba. Identification of matrix generators of a Chevalley group. *J. Algebra* **309** (2007), 484–496.

[18] Ákos Seress. *Permutation Group Algorithms*, volume 152 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2003.

[19] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, 2002.

School of Mathematics
University of Birmingham
Birmingham
B15 2TT, UK
k.magaard@bham.ac.uk

Department of Mathematics
University of Auckland
Auckland
New Zealand
obrien@math.auckland.ac.nz

Department of Mathematics
The Ohio State University
Columbus, OH 43210
USA
akos@math.ohio-state.edu