

FAST CONSTRUCTIVE RECOGNITION OF BLACK-BOX UNITARY GROUPS

PETER A. BROOKSBANK

Abstract

In this paper, the author presents a new algorithm to recognise, constructively, when a given black-box group is a homomorphic image of the unitary group $SU(d, q)$ for known d and q . The algorithm runs in polynomial time, assuming the existence of oracles for handling $SL(2, q)$ subgroups, and for computing discrete logarithms in cyclic groups of order $q \pm 1$.

1. Introduction

A *black-box group* G is a group whose elements are encoded as 0–1 strings of uniform length N , and whose group operations are performed by an oracle (the ‘black box’). Given strings representing $g, h \in G$, the black box can compute strings representing gh and g^{-1} , and can decide whether $g = h$. Black-box groups are important because of their great generality; matrix groups and permutation groups are both examples of such groups.

Let H be a concrete group (such as a group of matrices or permutations) and let $G = \langle \mathcal{S} \rangle$ be a given black-box group. We will say that a homomorphism $\Psi: H \rightarrow G$ is *effective* if there are procedures that compute $h\Psi \in G$ for any given $h \in H$, and also $g\Psi^{-1} \in H$ for any given $g \in H\Psi$. If each of those procedures runs in time $O(f)$, then we say that Ψ is *$O(f)$ -effective*. A *black-box constructive recognition algorithm* for H is an algorithm that, for any input black-box group G , constructs an $O(f)$ -effective epimorphism $\Psi: H \rightarrow G$ whenever G is a nontrivial homomorphic image of H . The efficiency of such an algorithm is measured both in terms of the cost of setting up the effective epimorphism, and also in terms of f , the cost of each application.

There are currently two basic approaches to computing efficiently with matrix groups. One of these, led by Leedham-Green [16], is a geometric approach based on Aschbacher’s classification of subgroups of $GL(d, q)$; see [1]. The other, due to Kantor and Seress [14], is a purely black-box approach based on the work of Babai and Beals [3, 4]. The constructive recognition of quasisimple groups is a fundamental ingredient in both approaches.

1.1. *History of constructive recognition*

The breakthrough paper in the recognition of black-box simple groups, by Cooperman, Finkelstein and Linton [10], deals with the groups $PSL(d, 2)$. The result was later extended by those authors in joint work with Bratus to $PSL(d, q)$ for $d > 3$ and $q > 4$. In [13], Kantor and Seress give black-box constructive recognition algorithms for all quasisimple classical groups. A drawback to the algorithms of Kantor–Seress and Bratus *et al.* is that their running times are not polynomial in the input length; they contain small factors of q (the size of the

This research was supported, in part, by the National Science Foundation
 Received 14 April 2003, revised 6 May 2003; published 13 October 2003.
 2000 Mathematics Subject Classification 20C20, 20C34, 20D08
 © 2003, Peter A. Brooksbank

field), whereas N may contain only factors of $\log q$. We note, however, that the alternating and symmetric groups can be handled in polynomial time (see, for example, [5]).

The important role of $\text{SL}(2, q)$ -subgroups in the study of the finite simple groups of Lie type is well documented. This is reflected algorithmically in a development in black-box recognition [8], where the $\text{PSL}(d, q)$ - and $\text{PSp}(d, q)$ -algorithms in [13] are modified to obtain algorithms whose running times are polynomial if one assumes the availability of an oracle for handling $\text{SL}(2, q)$ -subgroups. Recently, Conder, Leedham-Green and O'Brien have made significant advances in the efficient treatment of $\text{SL}(2, q)$ (see, for example, [9]), suggesting that algorithms that admit the use of such an oracle will have considerable practical value.

We remark that, because of the size of the groups concerned, all known constructive recognition algorithms for classical groups employ randomized algorithms rather than the more-traditional deterministic ones. A randomized algorithm is called *Monte Carlo* if the output of the algorithm is incorrect with probability less than $1/2$; higher reliability can be achieved by repetition and majority vote. *Las Vegas* algorithms form a subclass of Monte Carlo algorithms: here a positive output is guaranteed to be correct, but failure may be reported (with probability less than $1/2$) if a suitable output has not been determined after a prescribed time.

1.2. Statement of results

In this paper we present a new Las Vegas black-box constructive recognition algorithm for unitary groups, which hypothesizes oracles for handling two-dimensional subgroups and for computing discrete logarithms. Replacing the oracles with known procedures for the computations that they perform, our algorithm still has improved running time over existing algorithms (see Section 1.4) and it runs in polynomial time, assuming the oracles. A similar treatment of the orthogonal groups $\text{P}\Omega^\epsilon(d, q)$ is in preparation. Hence, in view of [8], ‘polynomial time with oracle’ black-box constructive recognition algorithms will soon be available for all the classical groups. Furthermore, the author is currently implementing the unitary group algorithm in GAP4; see [11].

There are numerous stages in our algorithm where new techniques were developed in order to satisfy the more stringent timing goals that we have set. We highlight two, as follows.

(a) *The construction, within G , of the natural module for a $(d - 2)$ -dimensional subgroup L of G , together with an L -invariant form on that module (see Subsection 4.3.2).* This construction enables us to avoid the recursive approach taken in [13], and leads to our improved timing. We note that an L -invariant form is also constructed in [13], but only after the recursive call has been made.

(b) *The improved algorithms for $\text{SU}(3, q)$ and $\text{SU}(4, q)$ in Section 6.* The complexity of the algorithms presented in [13, 6.6.1 and 6.6.2] to handle these cases is some distance from our benchmark of efficiency, and a substantially different approach was required. The resulting algorithms are quite subtle, and take up a significant portion of the paper. Our main result can be stated as follows (see Section 1.3 for a description of the complexity parameters ξ , μ and χ).

THEOREM 1.1. *Suppose that there is an oracle available that constructively recognises any black box $\text{SL}(2, q)$, and another that computes discrete logarithms in cyclic groups of order $q \pm 1$. Then there is an*

$$O(d^2 \log d \{\xi + \chi \log q + d \log^4 q\})\text{-time}$$

Las Vegas black-box constructive recognition algorithm for the special unitary group $\text{SU}(d, q)$.

Moreover, if the black-box group $G = \langle \mathcal{S} \rangle$ is a homomorphic image of $\text{SU}(d, q)$ for known d and q , then the algorithm produces an

$$O(\xi + \chi\{d^2 + \log q\} + d^5 \log^2 q)$$

effective epimorphism $\Psi: \text{SU}(d, q) \rightarrow G$.

We will see that, as a byproduct of the procedure to compute $g\Psi^{-1}$, we obtain the following result, which is of vital importance in applications of constructive recognition algorithms.

THEOREM 1.2. *Let $G = \langle \mathcal{S} \rangle$ be a constructively recognised homomorphic image of $\text{SU}(d, q)$, and let $g \in G$ be given. Then there is a Las Vegas $O(\xi + \chi\{d^2 + \log q\})$ -time algorithm to write a straight-line program from \mathcal{S} to g .*

The term ‘straight-line program from \mathcal{S} ’ will be defined in Section 3.1, but can be thought of informally as a ‘word in \mathcal{S} ’.

1.3. Complexity parameters

The following are descriptions of the complexity parameters that we use.

μ : an upper bound on the time required for each group operation in G (using the ‘black box’). Evidently, $\mu \geq N$, but a reasonable upper bound on μ depends on the actual representation of the black-box group. An important example is when G is represented as a group of $n \times n$ matrices over a field of size r . In this situation, $\mu = O(n^3 \log r)$, assuming that field operations can be carried out in $O(\log r)$ time.

ξ : an upper bound on the time requirement, per element, for the construction of independent, (nearly) uniformly distributed random elements of G . A fundamental result of Babai [2] produces such elements in a black-box group (see [13, 2.2.2] for further discussion).

χ : an upper bound on the time requirement for each application of the hypothesised $\text{SL}(2, q)$ -oracle, and for each use of the discrete logarithm oracle. We assume that $\chi \geq \mu \log q$ (see Section 3.3).

1.4. Timing comparisons

We now compare the running time of our algorithm with that of the algorithm in [13, Section 6] for the case $G = \text{PSU}(d, q)$. For this purpose, we use the running time

$$\chi = O(\xi q \log q + \mu q \log^2 q)$$

of the $\text{PSL}(2, q)$ algorithm in [13, 3.6.1]. The timing in Theorem 1.1 is then dominated by the χ term, giving

$$O(\xi d^2 q \log d \log q + \mu d^2 q \log d \log^2 q). \quad (1)$$

On the other hand, the timing for the $\text{PSU}(d, q)$ algorithm in [13, 6.6.3] is

$$O(\xi \{d q^2 \log d + d^2 q \log d\} + \mu \{d^2 q^3 \log^2 q + d^4 q^2 \log d \log q\}). \quad (2)$$

We first compare the coefficient of μ in (1) and (2). For large d , there are at least a factor of d^2 fewer group operations required in our algorithm, and for large q , there are roughly a factor of q^2 fewer group operations required. Next, we compare the coefficient of ξ .

For small q , the two algorithms choose roughly the same number of random elements, while our algorithm makes a factor of q fewer choices when q is large.

In the most important practical settings (namely when the given representation of the unitary group is in the correct characteristic), we will be able to incorporate a more efficient $\mathrm{SL}(2, q)$ -routine than in [13, 3.6.1] (such as the one being developed by Leedham-Green and O'Brien), and we expect our algorithm to perform much better in such settings.

2. Unitary group preliminaries

In this section we summarise the properties of unitary groups that we will need; the reader is referred to [15] for a more thorough discussion of classical groups. Let \mathbb{F}_{q^2} be the field containing q^2 elements for some prime power $q = p^k$, and let V be a vector space of dimension $d \geq 3$ over \mathbb{F}_{q^2} . Let $\lambda \mapsto \bar{\lambda} = \lambda^q$ denote the involutory automorphism of \mathbb{F}_{q^2} , and let \mathbb{F}_q denote the fixed subfield of this automorphism. Let ρ denote a fixed generator of $\mathbb{F}_{q^2}^*$ so that $\zeta := \rho^{q+1}$ is a generator of \mathbb{F}_q^* . Let \mathbb{F}_p denote the prime subfield of \mathbb{F}_{q^2} , so that \mathbb{F}_{q^2} is a degree- $2k$ extension of \mathbb{F}_p . Finally let $(,)$ denote a non-degenerate hermitian form on V preserved by the elements of $\mathrm{SU}(V)$.

2.1. Standard bases

Define an indexing list I as follows

$$I = \begin{cases} 1, 2, \dots, m, -1, -2, \dots, -m, & \text{if } d \text{ is even,} \\ 1, 2, \dots, m, 0, -1, -2, \dots, -m, & \text{if } d \text{ is odd.} \end{cases} \quad (3)$$

A *standard basis* of V is a basis of the form $\mathcal{B} = (e_i)_{i \in I}$, where $(e_i, e_j) = \delta_{i, -j}$ (in particular, e_i is nonsingular if and only if d is odd and $i = 0$). Unless otherwise stated, we will always write elements of $\mathrm{SU}(V)$ as matrices relative to such a basis.

2.2. Stabilisers of isotropic points

Let e be an isotropic vector of V , and let x be the point (1-space) $\langle e \rangle$. Then the subgroup $\mathrm{SU}(V)_x$ fixing x has a normal subgroup

$$T(x) = \{u \mapsto u + \lambda(u, e)e \mid \lambda + \bar{\lambda} = 0\} \cong \mathbb{F}_q^+.$$

$\mathrm{SU}(V)$ conjugates of $T(x)$ are called either (*unitary*) *transvection groups* or *long root groups*. If $x \neq y$ are isotropic points of V , then either:

- (a) x and y are perpendicular and $\langle T(x), T(y) \rangle \cong T(x) \times T(y)$; or
- (b) $y \notin x^\perp$ and $\langle T(x), T(y) \rangle \cong \mathrm{SU}(2, q)$. (In this case, $V = \langle x, y \rangle \perp \langle x, y \rangle^\perp$, and $\langle T(x), T(y) \rangle$ induces $\mathrm{SU}(2, q)$ on the first summand and 1 on the second.)

The stabiliser $\mathrm{SU}(V)_x$ also contains a larger normal subgroup, $Q(x)$, containing $T(x)$ as its centre, and consisting of those isometries that induce 1 on x and on x^\perp/x . Let $\mathcal{B} = (e_i)_{i \in I}$ be a standard basis for V and, for $i \in I \setminus \{0\}$, let $r_i(w, \lambda)$ denote the isometry

$$r_i(w, \lambda): u \mapsto u + (u, w - \lambda e_i)e_i - (u, e_i)w, \quad (4)$$

where $w \in \langle e_i, e_{-i} \rangle^\perp$ and $\lambda + \bar{\lambda} = (w, w)$. Then, if $x = \langle e_1 \rangle$, we have

$$Q(x) = \{r_1(w, \lambda) \mid w \in \langle e_1, e_{-1} \rangle^\perp, \lambda + \bar{\lambda} = (w, w)\} = O_p(\mathrm{SU}(V)_x),$$

the largest normal p -subgroup of $SU(V)_x$. Note that the subgroup $\{r_1(0, \lambda) \mid \lambda + \bar{\lambda} = 0\}$ is the transvection group $T(x)$. If we exclude the case $d = 3$, as well as the cases $d = 4$, $q \leq 3$, then

$$SU(V)_x = Q(x) \rtimes SU(V)_{x,y},$$

and

$$(SU(V)_{x,y})' = SU(V)_{e_1, e_{-1}} \cong (SU(V)_x)' / Q(x) \cong SU(d - 2, q).$$

The following lemma summarises the elementary properties of $Q(x)$, which can easily be verified by direct calculation.

LEMMA 2.1. (i) $r_1(w, \lambda)^g = r_1(wg, \lambda)$ for all $g \in SU(V)_{e_1, e_{-1}}$.

(ii) $r_1(w, \lambda)r_1(w', \lambda') = r_1(w + w', \lambda + \lambda' + (w, w'))$.

(iii) $[r_1(w, \lambda), r_1(w', \lambda')] = r_1(0, (w, w') - (w', w))$.

(iv) $Z(Q(x)) = T(x)$ and $Q(x)/T(x)$ is elementary abelian.

(v) $Q(x)$ acts regularly on the set of isotropic points not in x^\perp .

Property (i) describes the action of $SU(V)_{e_1, e_{-1}} \cong SU(d - 2, q)$ on the quotient group $Q(x)/T(x)$. The next result (which follows easily from the properties in Lemma 2.1) states that $Q(x)/T(x)$ is, in fact, the natural module of $SU(V)_{e_1, e_{-1}}$.

LEMMA 2.2. Let $s \in GU(V) \setminus SU(V)$ sending $e_1 \mapsto \bar{\rho}e_1$, $e_{-1} \mapsto \rho^{-1}e_{-1}$ and $u \mapsto u$ for all $u \in \langle e_1, e_{-1} \rangle^\perp$. Then the following statements hold.

(i) The conjugation action of s turns $Q(x)/T(x)$ into an \mathbb{F}_{q^2} -space, and

$$(r_1(w, \lambda)T(x), r_1(w', \lambda')T(x))_{Q(x)/T(x)} := (w, w')$$

defines an $SU(V)_{e_1, e_{-1}}$ -invariant hermitian form on $Q(x)/T(x)$.

(ii) The map $\psi_x: \langle e_1, e_{-1} \rangle^\perp \rightarrow Q(x)$ sending $w \mapsto r_1(w, \gamma(w, w))$ for some fixed $\gamma = 1 - \bar{\gamma}$ is linear and the induced map $w \mapsto (w\psi_x)T(x)$ is an isometry $\langle e_1, e_{-1} \rangle^\perp \rightarrow Q(x)/T(x)$.

A conjugate of the group $R(e_1, w) = \langle r_1(\rho^i w, 0) \mid 0 \leq i < 2k - 1 \rangle$, where $0 \neq w \in \langle e_1, e_{-1} \rangle^\perp$ is isotropic, is called a *short root group* of G . In view of Lemma 2.2(ii), we have the following correspondence:

$$\left\{ \begin{array}{c} \text{short root} \\ \text{subgroups of } Q(x) \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{c} \text{isotropic} \\ \text{points of } \langle x, y \rangle^\perp \end{array} \right\}.$$

In particular, if $R < Q(x)$ is a short root group, then $RT(x)/T(x)$ is an isotropic point of the $(d - 2)$ -space $Q(x)/T(x)$.

LEMMA 2.3. Let $g \in SU(V)_{x,y}$ send $e_1 \mapsto \bar{\lambda}^{-1}e_1$ and $e_{-1} \mapsto \lambda e_{-1}$, and let g_\bullet be the restriction of g to the $(d - 2)$ -space $\langle e_1, e_{-1} \rangle^\perp$. Then the following statements hold.

(i) g induces by conjugation on $Q(x)/T(x)$ the transformation $\tilde{g} = g_\bullet / \lambda$.

(ii) If λ is a generator of $(\mathbb{F}_{q^2})^*$, then $\det(\tilde{g}) = \bar{\lambda} / \lambda^{d-1}$ generates the cyclic subgroup of $(\mathbb{F}_{q^2})^*$ consisting of the determinants of all linear transformations induced on $Q(x)/T(x)$ by elements of $SU(V)_{x,y}$.

(iii) g induces the scalar $\lambda\bar{\lambda} \in (\mathbb{F}_q)^*$ on the \mathbb{F}_q -space $T(x)$.

Proof. Statements (i) and (ii) follow from [13, Lemma 6.7] and its proof, while Statement (iii) is an easy matrix calculation. \square

2.3. Primitive prime divisors

By a fundamental theorem of Zsigmondy [20], if p is a prime and $n \geq 2$, then there is a prime dividing $p^n - 1$ but not $p^i - 1$ for $1 \leq i < n$, except when either $p = 2$ and $n = 6$, or $n = 2$ and p is a Mersenne prime. Such a prime is called a *primitive prime divisor* (ppd) of $p^n - 1$.

For $n > 1$, we call an integer $j > 1$ dividing $p^n - 1$ a $\text{ppd}^\#(p; n)$ if:

- $n = 6, p = 2$ and $21 \mid j$;
- $n = 2, p$ is Mersenne and $4 \mid j$; or
- j is divisible by a ppd of $p^n - 1$.

If p is not Fermat, then we say that j is a $\text{ppd}^\#(p; 1)$ if j is not a power of 2; if p is Fermat, j is a $\text{ppd}^\#(p; 1)$ if $4 \mid j$. We call an element g of a group G a $\text{ppd}^\#(p; n)$ -element if $|g|$ is a $\text{ppd}^\#(p; n)$. We also say that g is a $\text{ppd}^\#(p; n_1) \cdot \text{ppd}^\#(p; n_2)$ -element if $|g|$ is both a $\text{ppd}^\#(p; n_1)$ and a $\text{ppd}^\#(p; n_2)$.

Certain ppd elements are highly abundant in the classical groups, and are useful because of their action on the underlying vector space, and also because of the subgroups that they can be used to generate.

LEMMA 2.4 (See [13, 6.1.5]). *If $d \geq 3$ is odd, then the following statements hold.*

(i) $\text{GU}(d, p^k)$ contains elements of order $p^{kd} + 1$ and each is irreducible on V . If τ is a $\text{ppd}^\#(p; 2kd)$ -element of $\text{GU}(d, p^k)$, then $|\tau|$ divides $p^{kd} + 1$.

(ii) *The proportion of $\text{ppd}^\#(p; 2kd)$ elements of $\text{GU}(d, p^k)$ is at least $1/4d$.*

LEMMA 2.5 (See [13, Lemma 2.7]). *If $g \in \text{GL}(d, p^k)$ has $\text{ppd}^\#(p; kd)$ -order, then $\{vg^i \mid 0 \leq i < kd\}$ is a $\text{GF}(p)$ -basis of $\text{GF}(p^k)^d$ and $\{vg^i \mid 0 \leq i < d\}$ is a $\text{GF}(p^k)$ -basis of $\text{GF}(p^k)^d$ for any nonzero $v \in \text{GF}(p^k)^d$.*

2.4. Probabilistic generation

We now state some results that we will use to help verify the correctness and reliability claims of various subroutines of our algorithm. Recall that $q = p^k$.

LEMMA 2.6 (See [13, Lemma 3.8(ii)]). *Two elements of $\text{SL}(2, q)$, of the same $\text{ppd}^\#(p; 2k)$ order, generate $\text{SL}(2, q)$ with probability greater than 0.55.*

LEMMA 2.7. *If $d \geq 4$ and $q \geq 8$ then, with probability greater than $1/2$, three unitary transvection groups of $\text{SU}(d, q)$ generate a subgroup J inducing $\text{SU}(3, q)$ on the nonsingular 3-space $[V, J]$ and 1 on $[V, J]^\perp$.*

Proof. As in [13, Lemma 3.7], J acts irreducibly on the nonsingular 3-space $[V, J]$ with probability at least $(1 - 1/q)^4 \geq (7/8)^4 > 1/2$. The result now follows from [13, 6.1.4], if we note that there are no proper, irreducible subgroups of $\text{SU}(3, q)$ that are generated by transvection groups. \square

LEMMA 2.8. *Let $d \geq 5, q \geq 16$ and let $a \in \text{SU}(d, q) = \text{SU}(V)$ be an element of $\text{ppd}^\#(p; k)$ - or $\text{ppd}^\#(p; k/2) \cdot \text{ppd}^\#(p; k)$ -order, according as k is odd or even respectively, having two-dimensional nonsingular support $[V, a]$ on V . If b is a random conjugate of a and $W = [V, \langle a, b \rangle]$ then, with probability at least $1/32$, W is a nonsingular 4-space and $\langle a, b \rangle$ induces $\text{SU}(4, q)$ on W and 1 on W^\perp .*

Proof. As in [13, Lemma 5.10(v)], with probability at least $1/32$, $\langle a, b \rangle$ acts irreducibly on W . We now refer to [12, Theorem 5.7], for a catalogue of subgroups of $SU(4, q)$. In each case, there are no proper irreducible subgroups generated by two elements of the stated order having two-dimensional nonsingular support. \square

3. Algorithmic preliminaries

In this section, we develop the algorithmic background necessary for computing with black-box groups.

3.1. Straight-line programs

Let G be a group, let X be a list of elements of G , and let $g \in G$. A *straight-line program (SLP) of length m from X to g* is a sequence of group elements (g_1, \dots, g_m) such that $g_m = g$ and, for each i , one of the following holds: $g_i \in X$; or $g_i = g_j^{-1}$ for some $j < i$; or $g_i = g_j g_k$ for some $j, k < i$. SLPs can be thought of as space-efficient words. Indeed, since we do not always want to compute and/or store each of the group elements g_i in the sequence, we instead define an SLP from X to g to be a sequence (w_1, \dots, w_m) such that, for each i , either w_i is a positive integer (representing the w_i th element of X), or $w_i = (j, -1)$ for some $j < i$ (representing w_j^{-1}), or $w_i = (j, k)$ for some $j, k < i$ (representing $w_j w_k$), such that if each expression in the sequence is evaluated in the obvious way, then the value of w_m is g . This more abstract definition also enables us to construct SLPs inside one group, and evaluate them in another.

3.2. Orders of elements and primitive prime divisors

Computing the *exact* order of a given element g of a black-box group G will not be necessary, but we will need to detect certain properties of $|g|$ so that we may determine whether or not g acts in a prescribed way on the natural module.

For a given integer n , we can compute g^n in time $O(\mu \log n)$ by repeated squaring, using the binary expansion of n . It follows that we can decide whether $|g|$ divides n , by comparing g^n and 1 inside G . In Lemma 2.4(ii), we saw that certain primitive prime divisor elements occur often in unitary groups. The following result, due to Neumann and Praeger [17], states that one can efficiently test whether a given black-box group element is a $\text{ppd}^\#(p; n)$ -element for a specified prime p and positive integer n , using a deterministic algorithm.

LEMMA 3.1 (See [6, Lemma 3.1]). *Following a preprocessing computation requiring time $O(n^3 \log n \log^4 p)$, one can test whether or not given elements of G have $\text{ppd}^\#(p; n)$ -order in time $O(\mu n \log p)$ per element.*

3.3. Oracles

As in [8], we assume the availability of deterministic algorithms (oracles) to perform certain computational tasks.

The $SL(2, q)$ -oracle. We hypothesise a black-box constructive recognition algorithm for $SL(2, q)$. In practice, such an algorithm will presumably be Las Vegas (such as [13, 3.6.1], for example). The architecture of our main algorithm will not be affected by the presence of a ‘randomised oracle’, but we will need to be somewhat more careful with the reliability estimates of subroutines that use it.

Discrete logarithms We also assume the availability of a *discrete logarithm oracle* for $C_{q\pm 1}$. Specifically, we hypothesise the following.

DLO−: Given a generator ζ of $\text{GF}(q)^*$ and $\lambda \in \text{GF}(q)^*$, one can find the integer $0 \leq n < q - 1$ such that $\zeta^n = \lambda$.

DLO+: Given a generator ρ of $\text{GF}(q^2)^*$ and $\lambda \in \langle \rho^{q-1} \rangle$, one can find the integer $0 \leq n < q + 1$ such that $(\rho^{q-1})^n = \lambda$.

The $\text{SL}(2, q)$ -algorithms in [9] also assume a DLO−. Assuming a DLO+ is natural, given that $\text{GU}(d, q) / \text{SU}(d, q) \cong C_{q+1}$.

3.3.1. An application of the oracles

Let $\text{SL}(2, q) \cong L \leq G$, let $\Psi: \text{SL}(2, q) \rightarrow L$ be the effective isomorphism provided by the $\text{SL}(2, q)$ -oracle, and let T be the transvection group of L generated by

$$\left\{ \begin{pmatrix} 1 & 0 \\ \zeta^i & 1 \end{pmatrix} \Psi \mid 0 \leq i < k \right\}.$$

Given any $t \in T$, one can use the $\text{SL}(2, q)$ -oracle to find $t\Psi^{-1} = \begin{pmatrix} 1 & 0 \\ \lambda & 1 \end{pmatrix}$ for some $\lambda \in \mathbb{F}_{q^2}$. Using DLO−, one can then find the integer $0 \leq n < q - 1$ such that $\lambda = \zeta^n$. This leads to the following useful observation (recall that $T \cong \mathbb{F}_q^+$). Given $g, h \in N_G(T)$ (but not necessarily in L) inducing scalar transformations of T of the same order, one can replace h with a power of itself to ensure that g and h induce the same scalar transformation.

3.3.2. The parameter χ

As stated in Section 1.3, χ denotes the cost of constructing an effective isomorphism $\Psi: \text{SL}(2, q) \rightarrow L$, and the cost of a single call to DLO− or to DLO+. It also denotes the cost of finding $g\Psi^{-1}$ for any given $g \in L$ or $g'\Psi$ for any given $g' \in \text{SL}(2, q)$. In view of the latter, we assume that $\chi \geq \mu \log q$ (the time required to evaluate an SLP of length $O(\log q)$ inside L).

3.4. Derived subgroups

We will need an efficient algorithm to compute derived subgroups.

LEMMA 3.2. *Given a subgroup $A = \langle \mathcal{A}_A \rangle$ of a black-box group G , where G is a homomorphic image of $\text{SU}(d, q)$ for known d and q , there is a Monte Carlo algorithm that, with probability greater than $1/2$, and in time $O(\mu |\mathcal{A}_A| d^4 \log^2 q)$, computes a generating set of size $O(d^2 \log q)$ for the derived subgroup A' of A .*

Proof. The result follows immediately from [19, Theorem 2.3.12], using $H = K = A$, $\delta = 1/2$ and the crude estimate $l_G = O(d^2 \log q)$. \square

3.5. The natural module

In [6, 3.2], algorithms are presented to solve each of the following.

Trace(β): Given $\beta \in \mathbb{F}_q$, find $\alpha \in \mathbb{F}_{q^2}^*$ such that $\alpha + \bar{\alpha} = \beta$.

OrthogonalComplement(U, V): Given nonsingular $U \leq V$ of dimension r , find the orthogonal complement U^\perp (of dimension $d - r$) of U in V .

StandardBasis(U): Given nonsingular $U \leq V$, find a standard basis $(e_i)_{i \in J}$, where J is a suitable indexing list, as defined in Section 2.1.

Timing and reliability. As each of these functions is explicitly used precisely once in our algorithm, we note only that the Las Vegas algorithm for `StandardBasis` employs the other two functions and successfully finds a standard basis of U , with probability greater than $15/16$, in time $O(d \log d \log q \{d^3 + \log^2 q\})$.

4. The main algorithm

Let $G = \langle \mathcal{J} \rangle$ be a given black-box group, known to be a nontrivial homomorphic image of $SU(d, q)$ for some q and $d \geq 5$. In this section we present an algorithm to construct a data structure that will be used in Section 5 to construct an effective epimorphism $\Psi: SU(d, q) \rightarrow G$. The data structure will consist of the following:

- (a) a subset $\mathcal{T} \subset SU(d, q)$ (In particular, we will construct a field \mathbb{F}_{q^2} of order q^2 and an \mathbb{F}_{q^2} -space V of dimension d , and the elements of $SU(d, q)$ will be matrices relative to some standard basis \mathcal{B} of V);
- (b) a subset $\mathcal{J}^* \subset G$ (whose elements are constructed using SLPs from \mathcal{J});
- (c) a bijection $\mathcal{T} \rightarrow \mathcal{J}^*$ extending to an epimorphism $\Psi: SU(d, q) \rightarrow G$; and
- (d) for each $2 \leq i \leq m = \lfloor d/2 \rfloor$, an $SU(4, q)$ -subgroup J_i of G and, when d is odd, an $SU(3, q)$ -subgroup J_0 of G . Each J_i will be equipped with an appropriate effective isomorphism.

The case $d = 2$ is covered by the oracle assumption, since $SU(2, q) \cong SL(2, q)$. The cases $d = 3, 4$ require significantly different techniques, and are dealt with separately in Section 6. However, we will soon need to make calls to those low-dimensional cases in the course of our main algorithm.

Presentation of the algorithm. The algorithm is quite long and, in places, somewhat technical. We therefore break it up into a number of manageable pieces. Subsections are used to divide the algorithm into its primary components. Subsubsections are used to present small subroutines, and each comes with its own timing statement and, wherever applicable, its own reliability estimate. If a subroutine is also to be used independently outside the main algorithm, we will usually package it into a lemma. Finally, if a subroutine (or its correctness proof) is itself quite detailed, we will designate it as a numbered procedure, and give a separate correctness proof in addition to timing and reliability estimates.

Small fields. In view of the algorithm in [13] for $PSU(d, q)$, our primary concern in this paper is with large fields. (Indeed, our first constructions will require that q be sufficiently large.) It is tempting, therefore, simply to use the algorithm in [13] for small q . However, the use of recursion in that algorithm would compromise the running time of ours. This obstacle notwithstanding, we are able to employ certain constructions from [13] for $q < 16$, and then rejoin our main algorithm at a later stage to complete the recognition. More precisely: in Section 4.1, we assume that $q \geq 16$; in Section 4.2 we summarise the constructions that we require from the algorithm in [13, Section 6], for $q < 16$; and from Section 4.3 onward, each subroutine works for all field sizes.

4.1. The subgroup J

A naturally embedded $SU(e, q)$ -subgroup of $SU(d, q)$ is one that induces $SU(e, q)$ on some nonsingular e -space of $\mathbb{F}_{q^2}^d$, and is the identity on the orthogonal complement of that

e -space. We will say that a subgroup $H \leq G$ is a *naturally embedded* $SU(e, q)$ -subgroup of G if, for any epimorphism $\Psi: SU(d, q) \rightarrow G$, the preimage $H\Psi^{-1}$ of H is a naturally embedded $SU(e, q)$ -subgroup of $SU(d, q)$. Our first step will be to find generators for a naturally embedded $SU(4, q)$ -subgroup J . This is achieved in two stages, following a strategy similar to that employed in [13, 6.2.1 and 6.2.2], although more closely resembling [8, 2.3] (for $PSL(d, q)$) and [13, 4.2.1, Case 4] (for $P\Omega(d, q)$).

In the first stage we search for an element τ , whose action on the underlying module decomposes the space into the perpendicular sum of a 2-space and a $(d - 2)$ -space. One important difference between the type of element employed here and that used in [13] is that they occur with differing frequency. We require only a polynomial number of random choices in order to obtain a suitable τ with high probability, in contrast to [13, 6.2.1] (compare also [8, 2.3]).

In the second stage, we kill off the action of τ on the $(d - 2)$ -space by raising it to a suitable power, thereby obtaining an element a having two-dimensional support. With high probability, the element a , together with a random conjugate, generates our $SU(4, q)$ subgroup.

4.1.1. The elements τ and a

Let $q = p^k \geq 16$, and define an odd integer n as follows:

$$n := \begin{cases} d - 2, & \text{if } d \text{ is odd,} \\ d - 3, & \text{if } d \text{ is even.} \end{cases}$$

We now present a Las Vegas algorithm to construct:

- (a) an element τ of $\text{ppd}^\#(p; 2nk) \cdot \text{ppd}^\#(p; k) \cdot \text{ppd}^\#(p; k/2)$ - or $\text{ppd}^\#(p; 2nk) \cdot \text{ppd}^\#(p; k)$ -order, according as k is even or odd respectively; and
- (b) an element $a \in \langle \tau \rangle$ of $\text{ppd}^\#(p; k/2) \cdot \text{ppd}^\#(p; k)$ or $\text{ppd}^\#(p; k)$ -order whose support in the natural module underlying G is a nonsingular 2-space upon which τ^{q^2-1} induces 1.

PROCEDURE 4.1. For each of at most $32n$ choices of element $\tau \in G$, proceed as follows. Use Lemma 3.1 to test whether τ has $\text{ppd}^\#(p; 2nk)$ -order. If so, set $a := \tau^{q^n+1}$ and use Lemma 3.1 again to test whether a has $\text{ppd}^\#(p; k)$ -order when k is odd or $\text{ppd}^\#(p; k/2) \cdot \text{ppd}^\#(p; k)$ -order when k is even. If so, then return the pair (τ, a) .

Correctness. Let V denote the underlying d -space upon which G acts naturally. Let (τ, a) be any pair returned by the procedure. Then τ and a have the orders stated in (a) and (b), respectively. If d is odd, then τ also preserves a decomposition $V = V_2 \perp V_{d-2}$ of V into perpendicular nonsingular i -spaces V_i ($i = 2, d - 2$) and acts irreducibly on V_{d-2} . By Lemma 2.4(i), $\tau^{q^{d-2}+1}$ centralises V_{d-2} . If d is even, then τ preserves a decomposition $V = V_2 \perp V_1 \perp V_{d-3}$ for nonsingular subspaces V_i ($i = 1, 2, d - 3$) and, since $n = d - 3$ is odd, τ^{q^n+1} centralises the $(d - 2)$ -space $\langle V_1, V_{d-3} \rangle$. In each case the nonsingular 2-space $V_2 = [V, a]$ is centralised by τ^{q^2-1} .

Reliability. Recall that $q \geq 16$, and observe that $(q - 1, q^n + 1) \leq 2$. Hence a fixed-choice τ is returned by the procedure if it has $\text{ppd}^\#(p; 2nk) \cdot \text{ppd}^\#(p; k)$ - or $\text{ppd}^\#(p; 2nk) \cdot \text{ppd}^\#(p; k) \cdot \text{ppd}^\#(p; k/2)$ -order that is also divisible by 8 if $k = 2$ and p is Mersenne, or if $k = 1$ and p is Fermat.

We claim that there are at least $|G|/16n$ elements of G of the stated ppd order (our estimates are crude, and the additional divisibility requirement in the Mersenne and Fermat cases does not affect our lower bound). Assuming that this is the case, the procedure will fail to return a pair (τ, a) with probability no more than $\{(1 - 1/(16n))^{16n}\}^2 < 1/e^2$.

To count the number of suitable elements, we consider the cases d even and d odd separately. We assume that $G \cong \text{SU}(V)$, since the proportion of suitable elements will not decrease by taking central quotients. First, suppose that d is odd. There are

$$\frac{q^{2d-4}(q^{d-1} - 1)(q^d + 1)}{(q + 1)(q^2 - 1)}$$

hyperbolic lines V_2 in V , each of which gives rise to a decomposition $V = V_2 \perp V_{d-2}$ of V . By Lemma 2.4(ii), the number of elements of G of $\text{ppd}^\#(p; k) \cdot \text{ppd}^\#(p; 2k(d-2))$ -order, preserving a fixed decomposition $V = V_2 \perp V_{d-2}$, is at least

$$\frac{|\text{SU}(2, q)|}{4} \cdot \frac{|\text{GU}(d-2, q)|}{4(d-2)}.$$

Hence, the proportion of desired elements in G is at least

$$\frac{|\text{SU}(2, q)|}{4} \cdot \frac{|\text{GU}(d-2, q)|}{4(d-2)} \cdot \frac{q^{2d-4}(q^{d-1} - 1)(q^d + 1)}{(q + 1)(q^2 - 1)} = \frac{|\text{SU}(d, q)|}{16(d-2)}.$$

Next, let d be even. There are

$$\frac{q^{2d-4}(q^{d-1} - 1)(q^d + 1)}{(q + 1)(q^2 - 1)} \cdot \frac{q^{d-3}(q^{d-2} - 1)}{q + 1}$$

pairs (V_2, V_1) , where V_2 is a hyperbolic line and $V_1 \leq V_2^\perp$ is nonsingular. By Lemma 2.4(ii), the number of elements of G of $\text{ppd}^\#(p; k) \cdot \text{ppd}^\#(p; 2k(d-3))$ -order, preserving a fixed decomposition $V_2 \perp V_1 \perp V_{d-3}$, is at least

$$\frac{|\text{GU}(2, q)|}{4} \cdot \frac{|\text{GU}(d-2, q)|}{4(d-3)}.$$

In this case, the proportion of suitable elements is at least

$$\frac{|\text{GU}(2, q)|}{4} \cdot \frac{|\text{GU}(d-3, q)|}{4(d-3)} \cdot \frac{q^{2d-4}q^{d-3}(q^d - 1)(d^{d-1} + 1)(q^{d-2} - 2)}{(q^2 - 1)(q + 1)^2} = \frac{|\text{SU}(d, q)|}{16(d-3)}.$$

Timing. Constructing $O(d)$ random elements of G and testing ppd properties for each using Lemma 3.1 takes $O(d^3 \log d \log^4 q + d\{\xi + \mu d \log q\})$ time.

4.1.2. Constructing J

Choose up to 128 conjugates b of a and, for each one, set $J := \langle a, b \rangle$. Use Section 6.2 to test whether $J \cong \text{SU}(4, q)$ and, if so, to construct a field $\mathbb{F}_{q^2} = \text{GF}(q^2)$, a 4-space $V_J = (\mathbb{F}_{q^2})^4$ and an effective isomorphism $\Psi_J: \text{SU}(V_J) \rightarrow J$.

Reliability. For a fixed-choice b , by Lemma 2.8, $\langle a, b \rangle$ is a naturally embedded $\text{SU}(4, q)$ -subgroup of G with probability greater than $1/32$. For such b , the algorithm in Section 6.2 will construct the desired isomorphism with probability greater than $1/2$. Hence, all of our 128 choices fail with probability less than $(1 - 1/64)^{128} < 1/e^2$.

Timing. The total time required for the various calls to 6.2 is $O(\xi + \chi \log q)$.

REMARK 4.2. A naturally embedded $SU(4, q)$ -subgroup is also constructed in [13, 6.2.2], but using transvections rather than ppd elements. From a practical point of view, our method for constructing J is more desirable than its counterpart in [13, 6.2.2], since we expect to make fewer calls to the $SU(4, q)$ routine. Indeed, a random pair (a, b) succeeds with probability greater than $1/2^6$, whereas the lower bound for the probability that a random selection in [13] succeeds is $1/2^{10}$.

4.1.3. The natural module V

We have now constructed a field \mathbb{F}_{q^2} , a 4-space $V_J = (\mathbb{F}_{q^2})^4$, and an effective isomorphism $\Psi_J: SU(V_J) \rightarrow J$. We may assume that elements of $SU(V_J)$ are 4×4 matrices written relative to a standard basis $(\tilde{e}_i)_{i \in I_J}$ of V_J , where $I_J = \{\pm 1, \pm 2\}$. Let \mathbb{F}_p denote the prime field of \mathbb{F}_{q^2} , and let \mathbb{F}_q be the subfield fixed by the automorphism $\alpha \mapsto \bar{\alpha} = \alpha^q$.

Set $V := (\mathbb{F}_{q^2})^d$, and define a non-degenerate hermitian form $(,)$ on V by designating the usual basis of V to be a standard basis $\mathcal{B} = (e_i)_{i \in I}$. Then $\tilde{e}_i \mapsto e_i$ for $i \in I_J$ extends \mathbb{F}_{q^2} -linearly to an isometry $V_J \rightarrow \langle e_1, e_2, e_{-1}, e_{-2} \rangle$, and induces an embedding $SU(V_J) \rightarrow SU(V)$. Since J is a naturally embedded $SU(4, q)$ -subgroup of G , we view Ψ_J as being defined on the subspace $\langle e_1, e_2, e_{-1}, e_{-2} \rangle$ of V . Using a change of basis, we may assume that $x = \langle e_1 \rangle$ and $y = \langle e_{-1} \rangle$ are the one-dimensional eigenspaces of $a\Psi_J^{-1}$. Let ρ be a generator of $\mathbb{F}_{q^2}^*$, set $\zeta := \rho\bar{\rho}$ (a generator of \mathbb{F}_q^*) and use 3.5 to find $0 \neq \delta = -\bar{\delta}$ by setting $\delta := \text{Trace}(0)$. The field elements ρ , δ and ζ will be fixed for the remainder of the algorithm.

4.1.4. Some elements of J

Identify the transformation $r_i(w, \lambda)$, defined in (4), with its matrix in $SU(V_J)$ relative to \mathcal{B} . Use Section 5.1 with $d = 4$ to construct each of the following elements of J :

- (i) for $1 \leq s \leq k$, $t_i := r_1(0, \zeta^{i-1}\delta)\Psi_J$;
- (ii) for $1 \leq j \leq 2k$, $r_{1j} := r_1(\rho^{j-1}e_2, 0)\Psi_J$ and $r_{2j} := r_1(\rho^{j-1}e_{-2}, 0)\Psi_J$;
- (iii) $l := l'\Psi_J$, where $l' \in SU(V_J)$ sending $e_{\pm i} \mapsto e_{\mp i}$ for $i = 1, 2$; and
- (iv) $\sigma := \sigma'\Psi_J$, where $\sigma' \in SU(V_J)$ sending $e_1 \mapsto \bar{\rho}^{-1}e_1$, $e_{-1} \mapsto \rho e_{-1}$.

Timing. All of the elements in (i) – (iv) are obtained in $O(\mu k \log q)$ time.

4.2. Small fields

We now break off from the main algorithm to describe how the various constructions that we have obtained so far for $q \geq 16$ are obtained for smaller field sizes. We will use some of the methods developed in [13] for this purpose, observing that a subroutine whose running time contains an explicit factor of q is no longer a problem.

(a) As in [13, 6.2.1], make up to $24qd$ choices to find, with probability greater than $1 - 1/e^3$, an element τ of $p \cdot \text{ppd}^\#(p; 2k(d-2))$ - or $p \cdot \text{ppd}^\#(p; k(d-2)) \cdot \text{ppd}^\#(p; k(d-2)/2)$ -order according as d is odd or even respectively. Set $z := q^{2d-4} - 1$ and $t := \tau^z$.

(b) As in [13, 6.2.2], make up to 2^{16} choices of tuples of conjugates of t to find, with probability greater than $1 - 1/e^3$, $J \cong SU(4, q)$ if $q > 2$ or $J \cong SU(6, 2)$ if $q = 2$, and obtain an effective isomorphism Ψ_J from the appropriate group of matrices to J . If $q = 2$,

restrict Ψ_J to any $SU(4, 2)$ -subgroup whose image under Ψ contains t , and replace J with this image.

(c) Change the standard basis within the 4-space so that the centre of $t\Psi_J^{-1}$ is spanned by the first basis vector, and use Ψ_J to construct each of the elements listed in 4.1.4(i)-(iv).

Timing and reliability. The steps (a) through (c) are successfully carried out, with probability $1 - (2/e^3 + 1/5) > 0.7$, in time $O(\xi d + \mu d^2 + \chi)$.

4.3. The subgroup Q and the unitary module Q/T

The group $T := \langle t_i \mid 1 \leq i \leq k \rangle$ is the image in G of the transvection group $T(x) = T((e_1))$ of $SU(V)$, so that $Q = O_p(N_G(T))$ is the image in G of the subgroup $Q(x)$ of $SU(V)$. Effective computation with the group Q is at the heart of our algorithm, and this immediate goal will be our focus for the next two subsections. In Section 4.3.1 we obtain a probable generating set for Q . It follows from Lemma 2.2 that Q/T has the structure of a unitary $(d - 2)$ -space as an $N_G(T)/Q$ -module. In 4.3.2, we construct a non-degenerate $N_G(T)/Q$ -invariant hermitian form on Q/T . Later, in 4.4, we will show that we can compute efficiently with Q by constructing a nicer set of generators and giving a routine that writes an SLP from this set to any given $u \in Q$. In fact, this will provide the tools needed to define our target epimorphism $\Psi : SU(V) \rightarrow G$ effectively on the group Q . The important consequence of the constructions in Sections 4.3 and 4.4 is that we will be able to avoid the recursive approach taken in [13] by essentially following the algorithm in [6] for recognising $SU(d, q)$ in its natural representation.

4.3.1. Constructing Q

Recall the element τ constructed in 4.1.1 for $q \geq 16$ or in 4.2(a) for $q < 16$. Replace τ with $\tau^{(q^2-1)}$ for $q \geq 16$, or with τ^p for $q < 16$. Then $\tau \in (N_G(T) \cap N_G(T^l))'$ induces on Q/T a transformation of $\text{ppd}^\#(p; 2k(d - 3))$ - or $\text{ppd}^\#(p; 2k(d - 2))$ -order. In the former case, τ preserves a decomposition of Q/T consisting of a nonsingular point and its orthogonal complement. Recall the elements $r_{1j}, r_{2j} \in O_p(N_J(T)) < Q$ of 4.1.4(ii). Now put

$$r_{ij} := (r_{2j})^{\tau^{(i-2)}} \quad \text{for } 3 \leq i \leq d - 2, 1 \leq j \leq 2k, \quad (5)$$

and return the generating set

$$\mathcal{R}_Q^* := \{r_{ij} \mid 1 \leq i \leq d - 2, 1 \leq j \leq 2k\} \subset Q. \quad (6)$$

Reliability. We claim that \mathcal{R}_Q^* generates Q with probability at least $4/5$. Indeed, by Lemma 2.5, the vectors $r_{21}T, (r_{21}T)^\tau, \dots, (r_{21}T)^{\tau^{d-4}}$ span a $(d - 3)$ -space of Q/T and, with at least the stated probability, the vector $r_{11}T$ is not in this $(d - 3)$ -space. The claim now follows by noting that T is the Frattini subgroup of Q .

Timing. The time required to construct all of the conjugates r_{ij} is $O(\mu kd)$.

REMARK 4.3. Suppose that $\Psi : SU(V) \rightarrow G$ is any epimorphism extending Ψ_J . Let τ'_i denote the transformation $\tau^{(i-2)}\Psi^{-1} \in SU(V)_{e_1, e_{-1}}$. Let $w_1 = e_2$ and $w_2 = e_{-2}$, and let $w_i = w_2\tau'_i$ for $3 \leq i \leq d - 2$. Then, by Subsection 4.1.4(ii) and Lemma 2.1(i), we have

$$r_{ij}\Psi^{-1} = r_1(\rho^{j-1}w_i, 0), \quad (7)$$

for $1 \leq i \leq d - 2$ and $1 \leq j \leq 2k$.

4.3.2. A form on Q/T

We next present a Las Vegas algorithm to construct a non-degenerate $N_G(T)/Q$ -invariant hermitian form $(\cdot, \cdot)_{Q/T}$ on Q/T . Via equation (26), the isomorphism Ψ_J induces an $N_J(T)/Q_J$ -invariant form $(\cdot, \cdot)_{\Psi_J}$ on Q_J/T . Since our target epimorphism $\Psi: \text{SU}(V) \rightarrow G$ will extend Ψ_J , we construct the unique such form $(\cdot, \cdot)_{Q/T}$ extending $(\cdot, \cdot)_{\Psi_J}$.

For $1 \leq i \leq d-2$, denote the long root element r_{i1} simply by r_i . Since the vectors $r_i T$ (probably) span Q/T as \mathbb{F}_{q^2} -space, it will suffice to compute each of the $(d-2)^2$ scalars $(r_i T, r_j T)_{Q/T}$. For each pair $i < j$ for which that scalar is nonzero, our strategy will be to construct an $\text{SU}(4, q)$ -subgroup K_{ij} of G containing r_i and r_j such that $K_{ij} \cap J = \langle T, T^l \rangle \cong \text{SL}(2, q)$, for the element $l \in J$ of 4.1.4(iii). The scalar $(r_i T, r_j T)_{Q/T}$ will then be computed inside K_{ij} .

PROCEDURE 4.4. For $1 \leq i \leq d-2$, set $R_i := \langle r_{ij} \mid 1 \leq j \leq 2k \rangle$ a short root group of G containing r_i , and initialise $(r_i T, r_i T)_{Q/T} := 0$.

For $1 \leq i < j \leq d-2$, proceed as follows. If $[r_i, r_j] = 1 = [r_{i2}, r_j]$, set $(r_i T, r_j T)_{Q/T} = (r_j T, r_i T)_{Q/T} := 0$. Otherwise, perform the following steps.

1. Set $K = K_{ij} := \langle T, T^l, R_i, R_j \rangle$.
2. Use the techniques in Section 6.2 at most $3 \lceil \log d \rceil$ times to construct an effective isomorphism $\Phi: \text{SU}(4, q) \rightarrow K$ with high probability. If successful with at least one of those calls, go to Step 4.
3. Report `failure` if no Φ is found in Step 2.
4. Use Proposition 6.14 to modify Φ so that $(\cdot, \cdot)_\Phi$ extends to $(\cdot, \cdot)_{Q/T}$. Use equation (26) to compute the scalar $\alpha := (r_i T, r_j T)_\Phi$. Set $(r_i T, r_j T)_{Q/T} := \alpha$ and $(r_j T, r_i T)_{Q/T} := \bar{\alpha}$. Return $[(r_i T, r_j T)_{Q/T}]$ if this matrix is nonsingular; otherwise report that we failed to generate Q .

Correctness. Let $\Psi: \text{SU}(V) \rightarrow G$ denote any epimorphism extending Ψ_J . In view of (7), we have $r_i = r_1(w_i, 0)\Psi$, $r_{i2} = r_1(\rho w_i, 0)\Psi$ and $r_j = r_1(w_j, 0)\Psi$ for some (unknown) $w_i, w_j \in \langle e_1, e_{-1} \rangle^\perp$. It follows directly from Lemma 2.1(iii) that $[r_i, r_j] = [r_{i2}, r_j] = 1$ if and only if $(w_i, w_j) = 0$. Hence our initial commutator test detects precisely when $(r_i T, r_j T)_{Q/T} = 0$.

We next claim that if $(r_i T, r_j T)_{Q/T} \neq 0$, then $K = K_{ij}$ is a naturally embedded $\text{SU}(4, q)$ -subgroup of G . Recall the isotropic points $x = \langle e_1 \rangle$ and $y = \langle e_{-1} \rangle$ defined in 4.1.3. Since $K\Psi^{-1}$ is generated by the transvection groups $T(x) = T\Psi^{-1}$ and $T(y) = T^l\Psi^{-1}$ together with the root groups $R(w_i) = \{r_1(\lambda w_i, 0) \mid \lambda \in \mathbb{F}_{q^2}\} = R_i\Psi^{-1}$ and $R(w_j) = \{r_1(\lambda w_j, 0) \mid \lambda \in \mathbb{F}_{q^2}\} = R_j\Psi^{-1}$, it follows that $K\Psi^{-1}$ induces a subgroup of $\text{SU}(4, q)$ on the nonsingular 4-space $W = \langle e_1, e_{-1}, w_i, w_j \rangle$ and centralises W^\perp . By [12, Lemma 5.7], it suffices to show that $K(W) = K\Psi^{-1}$ is irreducible on W .

If a 1-space $z \in W$ is not perpendicular to w_l for $l = i$ or j , then $R(w_l)$ moves z . Similarly, $T(x)$ moves z if z is not perpendicular to x , and $T(y)$ moves z if z is not perpendicular to y . Hence $K(W)$ fixes no 1-space. Let $U \leq W$ be $K(W)$ -invariant, of dimension at least 2. Since $R(w_i)$ moves w_j within $\langle w_j, e_1 \rangle$, U contains a vector w not perpendicular to e_1 or not perpendicular to e_{-1} . In the former case, $T(x)$ moves w within $\langle w, e_1 \rangle$ so that $e_1 \in U$, while in the latter case $T(y)$ moves w within $\langle w, e_{-1} \rangle$ so that $e_{-1} \in U$. In either case, the action of $\langle T(x), T(y) \rangle$ ensures that $\langle e_1, e_{-1} \rangle \leq U$. But then, for $l = 1, 2$, $R(w_l)$ moves e_{-1} within $\langle e_{-1}, w_l \rangle$ so that $w_l \in U$. It follows that $U = W$, and hence that $K(W)$ acts irreducibly on W , as claimed.

We have now shown that when the initial commutator test fails, the resulting group K is a naturally embedded $SU(4, q)$ -subgroup. In particular, `failure` is reported by the procedure (that is, Step (iii) is reached) only if bad luck occurs with random choices in our numerous calls to 6.2. If Φ is found in Step (ii), we now have the necessary data to apply Proposition 6.14 and modify Φ as needed. Equation (26) then gives a deterministic routine to compute the scalar $(r_i T, r_j T)_\Phi = (r_i T, r_j T)_{Q/T}$.

Finally, the determinant of $[(r_i T, r_j T)_{Q/T}]$ is zero if and only if the $r_i T$ do not span Q/T as \mathbb{F}_{q^2} -space. For example, using additive notation, if $r_{d-2} T + \sum_{i=1}^{d-3} \alpha_i (r_i T) = 0$, then adding α_i times row i to row $d-2$ for each $1 \leq i \leq d-3$ makes the latter 0. Hence the procedure detects when the elements r_{ij} fail to generate Q , thereby upgrading the Monte Carlo construction of Q to Las Vegas.

Reliability. For fixed $i < j$, if $K_{ij} \cong SU(4, q)$, then a single call to 6.2 succeeds with probability greater than $1/2$. It follows that at least one of the $3\lceil \log d \rceil$ calls in Step (ii) produces a suitable Ψ_K with probability greater than $1 - 1/d^3$. Hence `failure` is reported for at least one of the less than $d^2/2$ pairs $i < j$ with probability less than $(d^2/2)/d^3 = 1/(2d) \leq 1/10$.

Timing. $O(d^2 \log d \{\xi + \chi \log q\})$ is required for the $O(d^2 \log d)$ calls to 6.2.

4.4. Constructing Ψ on Q

At the present stage of the algorithm, having successfully run Procedure 4.4, we may assume that $Q = \langle \mathcal{S}_Q^* \rangle$. Recall that our ultimate goal is to construct an epimorphism $\Psi: SU(V) \rightarrow G$ extending Ψ_J . Note that $Q = Q(x)\Psi$ for any such Ψ , where $x = \langle e_1 \rangle$ and $Q(x) = O_p(SU(V)_x)$; in 4.4 we construct a particular Ψ effectively on $Q(x)$. This is achieved in three stages, as follows.

- (a) Construct a generating set \mathcal{T}_Q of $Q(x)$ and a bijection $\mathcal{T}_Q \rightarrow \mathcal{S}_Q^*$ that extends to an isomorphism $Q(x) \rightarrow Q$ agreeing with Ψ_J on $Q_J(x)$; see Subsection 4.4.1.
- (b) Construct from \mathcal{T}_Q a ‘standard’ generating set $\Delta(x)$ of $Q(x)$, and then use the bijection in Stage (a) to obtain its image Δ^* in Q ; see Subsection 4.4.2.
- (c) Constructively recognise $\lceil (d-2)/2 \rceil$ subgroups J_i containing T such that the subgroups $Q_i := O_p(N_{J_i}(T))$ generate Q ; Lemma 4.8 then uses the Q_i to handle computations within Q ; see Subsection 4.4.3.

We will also see that, once an epimorphism Ψ extending Ψ_J is defined on $Q(x) \rightarrow Q$, it is uniquely determined (Proposition 4.5).

4.4.1. The set \mathcal{T}_Q

Label the usual basis of $\mathbb{F}_{q^2}^{d-2}$ with the elements r_1^*, \dots, r_{d-2}^* (that is, $r_1^* = (1, 0, \dots, 0)$ and so on). Equip $\mathbb{F}_{q^2}^{d-2}$ with the hermitian form defined by the matrix $[[\alpha_{ij}]]$ relative to r_1^*, \dots, r_{d-2}^* . Then $\mathbb{F}_{q^2}^{d-2}$ is a unitary space which we loosely associate with Q/T via the isometry $r_i^* \mapsto r_i T$. Note that r_1^*, r_2^* is a hyperbolic pair. Use 3.5 to compute $W := \langle r_1^*, r_2^* \rangle^\perp$ relative to $[[\alpha_{ij}]]$, and also to obtain a standard basis of W . Insert r_1^* and r_2^* in the appropriate positions to obtain a standard basis of $\mathbb{F}_{q^2}^{d-2}$ and set A to be the $(d-2) \times (d-2)$ matrix consisting of the vectors in this standard basis.

Recall that elements of $Q(x)$ have the form $r_1(w, \lambda)$ for $\lambda \in \mathbb{F}_{q^2}$ and $w \in \langle e_1, e_{-1} \rangle^\perp$; we denote the element $r_1(w, \lambda)T(x)$ of $Q(x)/T(x)$ simply by $\bar{r}(w)$. We next wish to find

vectors $w_i \in \langle e_1, e_{-1} \rangle^\perp$ such that $\bar{r}(w_i) \mapsto r_i T$ extends to an isometry $Q(x)/T(x) \rightarrow Q/T$. Note that the i th row of A^{-1} expresses r_i^* as an \mathbb{F}_{q^2} -vector relative to our standard basis of $\mathbb{F}_{q^2}^{d-2}$, and hence $r_i T$ as an \mathbb{F}_{q^2} -linear combination of *some* standard basis of Q/T (which we have not yet constructed). Pad A^{-1} to obtain a $(d-2) \times d$ matrix C by inserting 0s in each row at the positions corresponding to the basis vectors e_1 and e_{-1} . Denote the rows of C by w_1, \dots, w_{d-2} , interpreted as row vectors of V relative to \mathcal{B} . Then it is clear that $\bar{r}(w_i) \mapsto w_i \mapsto r_i^* \mapsto r_i T$ extends to an isometry $Q(x)/T(x) \rightarrow Q/T$. An immediate consequence of our constructions thus far is the following proposition.

PROPOSITION 4.5. *There is a unique epimorphism $\Psi: \text{SU}(V) \rightarrow G$ such that:*

- (i) Ψ extends the isomorphism $\Psi_J: \text{SU}(V_J) \rightarrow J$ constructed in Subsection 4.1.2, and
- (ii) the restriction of Ψ to $Q(x)$ induces the isometry $\bar{r}(w_i) \mapsto r_i T$.

Proof. Consider any epimorphism $\Psi: \text{SU}(V) \rightarrow G$ extending Ψ_J . Since $T(x) \leq \text{SU}(V_J)$, the image of $Q(x)/T(x)$ under Ψ determines the image of $Q(x)$. Since $l' \in \text{SU}(V_J)$ (defined in 4.1.4(iii)), the image of $Q(x)$ under Ψ determines the image of $Q(x)^{l'} = Q(y)$. The uniqueness of an epimorphism Ψ satisfying conditions (i) and (ii) now follows by the fact that $\text{SU}(V) = \langle Q(x), Q(y) \rangle$. It suffices then to demonstrate the existence of such a Ψ .

Fix an epimorphism $\Psi_0: \text{SU}(V) \rightarrow G$ extending Ψ_J . Then Ψ_0 induces an isometry $(Q_J(x)/T(x))^\perp \rightarrow (Q_J/T)^\perp$; let $w'_i \in V_J^\perp$ be such that $\bar{r}(w'_i)\Psi_0 = r_i T$ for $3 \leq i \leq d-2$. Then $w_i \mapsto w'_i$ extends to an isometry of V_J^\perp . Let $C \in \text{GU}(V_J^\perp) \cong \text{GU}(d-4, q)$ send $w_i \mapsto w'_i$ for $3 \leq i \leq d-2$. Let γ denote the automorphism of $\text{SU}(V)$ induced under conjugation with C . Then $\Psi := \gamma \circ \Psi_0: \text{SU}(V) \rightarrow G$ sending $M \mapsto (M^\gamma)\Psi_0$ satisfies conditions (i) and (ii) because Ψ and Ψ_0 agree on $\text{SU}(V_J)$ (since C is the identity on V_J), and

$$\bar{r}(w_i)\Psi = (\bar{r}(w_i)^\gamma)\Psi_0 = \bar{r}(w_i C)\Psi_0 = \bar{r}(w'_i)\Psi_0 = r_i$$

for $3 \leq i \leq d-2$. The second equality follows from Lemma 2.1(i). □

In view of Proposition 4.5, the vectors w_i that we have computed are precisely those vectors appearing in (7) for the unique epimorphism Ψ that we will construct. Set $w_1 := e_2$, $w_2 := e_{-2}$ and, for $1 \leq i \leq d-2$ and $1 \leq j \leq 2k$, $1 \leq s \leq k$, set $r'_{ij} := r_1(\rho^{j-1}w_i, 0)$ and $t'_s := r_1(0, \zeta^{s-1}\delta)$. Recalling the elements t_s from 4.1.4(i), set

$$\mathcal{R}_Q^* := \{t_s, r_{ij} \mid 1 \leq s \leq k, 1 \leq i \leq d-2, 1 \leq j \leq 2k\},$$

and

$$\mathcal{T}_Q := \{t'_s, r'_{ij} \mid 1 \leq s \leq k, 1 \leq i \leq d-2, 1 \leq j \leq 2k\}.$$

Then the bijection $\mathcal{T}_Q \rightarrow \mathcal{R}_Q^*$, sending

$$r'_{ij} \mapsto r_{ij} \quad \text{and} \quad t'_s \mapsto t_s, \tag{8}$$

extends to an isomorphism $Q(x) \rightarrow Q$, which in turn extends to the unique epimorphism $\Psi: \text{SU}(V) \rightarrow G$ of Proposition 4.5.

Timing and reliability. The time required for all of the computations in Subsection 4.4.1 is dominated by the timing stated in 3.5 for computing the standard basis of $\mathbb{F}_{q^2}^{d-2}$. That computation is successfully carried out, with probability greater than $15/16$, in time $O(d \log d \log q \{d + \log^2 q\} + d^4 \log q)$.

4.4.2. The set $\Delta(x)$

We next wish to construct a ‘standard’ generating set $\Delta(x)$ for $\mathcal{Q}(x)$. The matrices comprising $\Delta(x)$ will be precisely those that, relative to a suitable standard basis, comprise the set $\Delta(x)$ constructed in [6, 4.5]. This will be important later, as we plan to use $\Delta(x)$ to complete the construction of our data structure by calling upon algorithms presented in [6]. For $1 \leq i \leq d-2$, $1 \leq j \leq 2k$, set

$$s'_{ij} := r_1(\rho^{j-1}e_i, \lambda_{ij}), \tag{9}$$

where the scalar λ_{ij} satisfies

$$\lambda_{ij} = \begin{cases} v_j, & \text{if } d \text{ is odd and } i = 0, \text{ for some } v_j \text{ such that } v_j + \overline{v_j} = \zeta^{j-1}, \\ 0, & \text{otherwise.} \end{cases}$$

Now define

$$\Delta(x) := \{t'_i, s'_{ij} \mid 1 \leq s \leq k, i \in I', 1 \leq j \leq 2k\}.$$

The following procedure constructs each element of s'_{ij} of $\Delta(x)$ using SLPs from \mathcal{T}_Q (if d is odd, it therefore computes specific values for the scalars v_j). It then evaluates those SLPs from \mathcal{S}_Q^* to obtain the image, Δ^* , of $\Delta(x)$ in Q . Recall the matrix A constructed in 4.4.1 and, replacing the ordered index list $1, \dots, d-2$ with I' , denote the i th row of A by A_i for $i \in I'$.

PROCEDURE 4.6. For each $i \in I'$ and $1 \leq j \leq 2k$, proceed as follows.

1. Compute the vector $\rho^{j-1}A_i = (\alpha_1, \dots, \alpha_{d-2}) \in (\mathbb{F}_{q^2})^{d-2}$.
2. For $1 \leq u \leq d-2$ and $1 \leq v \leq 2k$, find integers $0 \leq a_{uv} < p$ such that $\alpha_u = \sum_{v=1}^{2k} a_{uv}\rho^{v-1}$ and hence write an SLP σ_{ij} of length $O(d \log q)$ from \mathcal{T}_Q to the element

$$\prod_{u=1}^{d-2} \prod_{v=1}^{2k} (r'_{uv})^{a_{uv}} = r_1(\rho^{j-1}e_i, \lambda_{ij}) \in \mathcal{Q}(x),$$

for some $\lambda_{ij} + \overline{\lambda_{ij}} = \zeta^{j-1}(e_i, e_i)$.

3. If d is even or if $i \neq 0$, then proceed as follows.

(i) Use linear algebra in the \mathbb{F}_p -space \mathbb{F}_q to find integers $0 \leq n_s < p$ ($1 \leq s \leq k$) such that $-\lambda_{ij} = \delta \sum_{s=1}^k n_s \zeta^{s-1}$, and hence write an SLP τ_{ij} of length $O(\log q)$ from $\{t'_1, \dots, t'_k\}$ to

$$r_1(0, -\lambda_{ij}) = \prod_{s=1}^k (t'_s)^{n_s}.$$

(ii) Replace σ_{ij} with the concatenation of σ_{ij} and τ_{ij} .

Else (if d is odd and $i = 0$), set $v_j := \lambda_{ij}$.

4. Use the bijection $\mathcal{T}_Q \rightarrow \mathcal{S}_Q^*$, defined in (8), to evaluate the SLP σ_{ij} from \mathcal{S}_Q^* , and denote the resulting element of Q by s_{ij} .

If d is odd, use the scalars v_j to determine the set $\Delta(x)$. Return $\Delta(x)$ together with its image $\Delta^* := \{t_s, s_{ij} \mid 1 \leq s \leq k, i \in I', 1 \leq j \leq 2k\}$ in Q .

Correctness. It follows immediately from the construction of A and the definition of r'_{ij} and t'_s that the SLP σ_{ij} evaluates from \mathcal{T}_Q to an element s'_{ij} as defined in (9).

Timing. $O(\mu d^2 \log^2 q)$ time is needed to evaluate all of the SLPs σ_{ij} from \mathcal{S}_Q^* .

REMARK 4.7. The bijection $\Delta(x) \rightarrow \Delta^*$ sending

$$t'_s \mapsto t_s \quad \text{and} \quad s'_{ij} \mapsto s_{ij}, \quad (10)$$

and the bijection $\mathcal{T}_Q \rightarrow \mathcal{S}_Q^*$ of (8) extend to the same isomorphism $Q(x) \rightarrow Q$. It is the sets $\Delta(x)$ and Δ^* that we will use to construct Ψ effectively in the remainder of the algorithm, so we now discard the sets \mathcal{T}_Q and \mathcal{S}_Q^* .

4.4.3. The subgroups J_i

In order to compute effectively with Q , we next construct some low-dimensional subgroups J_i of G . First some notation is required. Define

$$B' := (s_{i1}T)_{i \in I'}, \quad (11)$$

a standard \mathbb{F}_{q^2} -basis of Q/T with respect to $(\cdot, \cdot)_{Q/T}$, and set

$$B'_p := (s_{ij}T)_{i \in I', 1 \leq j \leq 2k}, \quad (12)$$

an \mathbb{F}_p -basis of Q/T . For $i \in I'$, set $S_i := \langle s_{ij} \mid 1 \leq j \leq 2k \rangle$ and, for $2 \leq i \leq m$, set

$$J_i := \langle T, T^l, S_i, S_{-i} \rangle,$$

where $l \in J$ was constructed in 4.1.4(iii). Then J_i is a naturally embedded $\text{SU}(4, q)$ -subgroup of G containing $\langle T, T^l \rangle$ (note that $J_2 = J$). If d is odd, set

$$J_0 := \langle T, T^l, S_0 \rangle \cong \text{SU}(3, q).$$

Observe that, for $i = 2, \dots, m$ (d even) or $i = 0, 2, \dots, m$ (d odd),

$$Q_i := O_p(N_{J_i}(T)) = \langle S_i, S_{-i} \rangle,$$

so that $(Q_i/T, Q_j/T)_{Q/T} \neq 0$ if and only if $i = j$ (Lemma 4.8). Set $\Psi_2 := \Psi_J$ and, for $3 \leq i \leq m$, use 6.2 at most $5 \lceil \log m \rceil$ times to (probably) construct an effective isomorphism

$$\Psi_i: \text{SU}(4, q) \rightarrow J_i. \quad (13)$$

A single call to 6.2 succeeds with probability greater than $1/2$, so we successfully construct *all* of the Ψ_i with probability greater than $1 - m/2^{5 \log m} > 1 - 1/81$. If d is odd, use 6.1 at most eight times to construct, with probability greater than $1 - 1/2^8$, an effective isomorphism

$$\Psi_0: \text{SU}(3, q) \rightarrow J_0. \quad (14)$$

Timing and reliability. All of the Ψ_i are obtained, with probability greater than $1 - (1/81 + 1/2^8) > 0.98$, in $O(d \log d \{\xi + \chi \log q\})$ time.

LEMMA 4.8. *Equipped with the isomorphisms Ψ_i of (13) and (14), for any given $u \in Q$, there are $O(\chi d)$ -time deterministic algorithms for each of the following.*

(i) *For $i \in I'$, $1 \leq j \leq 2k$, find integers $0 \leq a_{ij} < p$ such that uT has vector $(a_{21}, a_{22}, \dots, a_{-m, 2k})$ relative to the \mathbb{F}_p -basis B'_p of (12). That is,*

$$uT = \prod_{i \in I'} \prod_{j=1}^{2k} (s_{ij}T)^{a_{ij}}.$$

(ii) *For $i \in I'$, find scalars $\alpha_i \in \mathbb{F}_{q^2}$ such that uT has vector $(\alpha_2, \dots, \alpha_{-m})$ relative to the \mathbb{F}_{q^2} -basis B' of (11).*

(iii) *Find an SLP of length $O(d \log q)$ from Δ^* to u .*

Proof. Since the groups Q_i/T are pairwise perpendicular, we use Lemma 6.10(ii) for $2 \leq i \leq m$ (in the general setting (GS2) described before that lemma) to write the projection of uT on Q_i/T along $(Q_i/T)^\perp$ as an \mathbb{F}_p -vector $(a_{i1}, \dots, a_{i,2k}, a_{-i1}, \dots, a_{-i,2k})$ relative to $\mathcal{B}'_p \cap (Q_i/T)$. If d is even, this proves Part (i). If d is odd, set

$$w := \prod_{i \in I' \setminus \{0\}} \prod_{j=1}^{2k} s_{ij}^{\alpha_{ij}}$$

and $w_0 := uw^{-1} \in Q_0$. Now use Lemma 6.1(ii) to find the integers a_{0j} representing the vector w_0T relative to $\mathcal{B}'_p \cap (Q_0/T)$.

For Part (ii), in view of (9) and (10), set

$$\alpha_i := \sum_{j=1}^{2k} a_{ij} \rho^{j-1}$$

for each $i \in I'$.

For Part (iii), set

$$w := \prod_{i \in I'} \prod_{j=1}^{2k} s_{ij}^{a_{ij}},$$

so that $t := uw^{-1} \in T$. Now use Ψ_J to write an SLP of length $O(\log q)$ from $\{t_1, \dots, t_k\} \subset \Delta^*$ to t . \square

4.5. A data structure for G

We are finally in a position to complete our constructive recognition of G . Recall the elements $l' \in \text{SU}(V_J)$ and $l \in J$ defined in 4.1.4(iii), and set $\Delta(y) := \Delta(x)^{l'}$, a generating set for $Q(x)^{l'} = Q(y)$. Extend the bijection $\Delta(x) \rightarrow \Delta^*$ to a bijection

$$\Delta(x) \cup \Delta(y) \rightarrow \Delta^* \cup \Delta^{*l} \tag{15}$$

in the obvious way. As in the proof of Proposition 4.5, (15) extends to a unique epimorphism $\Psi: \text{SU}(V) \rightarrow G$.

Associated with the standard basis $\mathcal{B} = (e_i)_{i \in I}$ of V there are two ‘opposite’ t.i. subspaces $E^+ = \langle e_1, \dots, e_m \rangle$ and $E^- = \langle e_{-1}, \dots, e_{-m} \rangle$. In [6, Section 4.6], an algorithm is presented that constructs nice generating sets for the subgroups $\text{SU}(V)_{E^+, E^-}$, $O_p(\text{SU}(V)_{E^+})$ and $O_p(\text{SU}(V)_{E^-})$ using short SLPs from precisely the sets $\Delta(x)$ and $\Delta(y)$ constructed here. Their union is a generating set \mathcal{T} for $\text{SU}(V)$, of size $O(kd^2)$. Using (15), evaluate each of those SLPs from the set $\Delta^* \cup \Delta^{*l}$ to obtain the image, \mathcal{S}^* , of \mathcal{T} in G .

Next, recall the elements $\sigma' \in \text{SU}(V_J)$ and $\sigma \in J$ constructed in 4.1.4(iv). Add l' and σ' to \mathcal{T} , add l and σ to \mathcal{S}^* , and extend the bijection $\mathcal{T} \rightarrow \mathcal{S}^*$ in the obvious way. Our data structure for G consists of the latter bijection together with the $\lfloor (d-2)/2 \rfloor$ effective isomorphisms $\Psi_j: \text{SU}(4, q) \rightarrow J_i$ defined in (13) and, if d is odd, the effective isomorphism $\Psi_0: \text{SU}(3, q) \rightarrow J_0$ defined in (14).

Timing. The time required to write and evaluate all SLPs is $O(\mu kd^2)$.

4.6. Total timing and reliability for the main algorithm

The failure probabilities of the randomised subroutines of the main algorithm described in 4.1 through 4.5 sum to at most $1/e^2 + 1/e^2 + 1/5 + 1/10 + 1/16 + 1/81 < 2/3$. The timing

is dominated by 4.1.1 and 4.3.2. Hence our algorithm returns a suitable data structure, with probability greater than $1/3$, in time $O(d^2 \log d\{\xi + \chi \log q + d \log^4 q\})$. Note that higher reliability can be achieved by repeating the various randomised subroutines suitably many times.

5. Straight-line programs

In the previous section we constructed a bijection $\mathcal{T} \rightarrow \mathcal{S}^*$ that extends to an epimorphism $\Psi: \text{SU}(V) \rightarrow G$. In this section, we will make Ψ effective (and hence prove Theorem 1.1 for $d \geq 5$) by presenting algorithms to solve the following problems.

Given $g' \in \text{SU}(V)$, write an SLP of length $O(d^2 \log q)$ from \mathcal{T} to g' . (A1)

Given $g \in G$, write an SLP of length $O(d^2 \log q)$ from \mathcal{S}^* to g . (A2)

One can then easily compute images and preimages under Ψ . For example, given $g' \in \text{SU}(V)$, use the algorithm for solving problem (A1) to write an SLP of length $O(d^2 \log q)$ from \mathcal{T} to g' . Evaluate that SLP from \mathcal{S}^* using the bijection $\mathcal{T} \rightarrow \mathcal{S}^*$, and define $g = g'\Psi$ to be the resulting element of G . The preimage of any given $g \in G$ is similarly obtained using (A2) in place of (A1).

5.1. The algorithm for (A1)

In [6, Section 5], a deterministic $O(d^3 \log q)$ time algorithm is presented to solve problem (A1) using exactly the same set \mathcal{T} as the one that we have constructed here.

REMARK 5.1. We can use (A1) to construct $Z(G)$ as follows. If $j = \gcd(q + 1, d) > 1$, set $\alpha := \rho^{(q^2-1)/j}$ and $z' := \text{diag}(\alpha, \dots, \alpha) \in Z(\text{SU}(V))$. Write an SLP of length $O(d^2 \log q)$ from \mathcal{T} to z' and evaluate this SLP from \mathcal{S}^* to obtain an element $z \in Z(G)$. If $z = 1$, then we now know that $G = \text{PSU}(d, q)$; otherwise, add z' to \mathcal{T} and z to \mathcal{S}^* .

Timing. Each application of (A1) takes $O(d^3 \log q)$ time. An additional $O(\mu d^2 \log q)$ time is required to compute an image of any given $g' \in \text{SU}(V)$ in G . In particular, given a suitable data structure for G , $Z(G)$ is constructed in $O(\mu d^2 \log q)$ time.

5.2. The algorithm for (A2)

Our timing goals appear to dictate the use of a randomised algorithm to find an SLP from \mathcal{S}^* to any given $g \in G$, in contrast to [13, Proposition 6.15]. The procedure given here applies for $d \geq 4$; the three-dimensional version will be presented in 6.1. We proceed along the same lines as [13, 6.4], to develop the algorithmic properties of Q . The four-dimensional version of the following result is given in 6.2 (see Lemma 6.12).

LEMMA 5.2. *Let $d \geq 5$ and $g \in N_G(T)$ be given. Then, in deterministic $O(\chi d^2)$ time, one can determine the $(d - 2) \times (d - 2)$ matrix \tilde{g} representing the linear transformation induced by g on the $(d - 2)$ space Q/T relative to the \mathbb{F}_{q^2} -basis B' for Q/T of (11).*

Proof. For each $sT \in B'$, use Lemma 4.8(ii) to write $(sT)^g \in Q/T$ as an \mathbb{F}_{q^2} -vector relative to B' . Return the matrix whose rows are those $d - 2$ vectors. \square

By Lemma 2.1(v), the group $Q(x)$ acts regularly on the set of isotropic points of V not perpendicular to x . Within the black-box group G , this says that Q acts regularly on the set of transvection groups opposite T . The next lemma is the algorithmic version of this transitivity.

LEMMA 5.3. For $d \geq 4$, there is an $O(\xi + \chi \log q)$ -time Las Vegas algorithm that, with probability greater than $1/2$, finds the unique element of Q conjugating any given T^{g_1} opposite T to any other given T^{g_2} opposite T .

Proof. Fix $1 \neq t \in T$. Choose a random element $w \in Q$, and proceed as follows.

1. Replace g_2 with $g_2 w$.
2. Set $K := \langle T, T^{g_1}, T^{g_2} \rangle$ and use the techniques of Section 6.1 to test whether $K \cong \text{SU}(3, q)$.
3. If the test in Step 2 succeeds, apply Lemma 6.5 (the three-dimensional deterministic version of this lemma) to K, T^{g_1} and T^{g_2} to construct $u \in O_p(N_K(T))$ conjugating T^{g_1} to T^{g_2} inside K and return uw^{-1} .

Correctness. It suffices to show that, with sufficiently high probability, K is a naturally embedded $\text{SU}(3, q)$ -subgroup of G ; for then $u \in O_p(N_K(T)) \leq Q$, as required. Let x and y denote the centres of the transvection groups T and T^{g_1} respectively in the d -space underlying G . Each of the more than $[(q^2)^{d-2} - (q^2)^{d-2.5}]/(q^2 - 1)$ nonsingular points z_\bullet of $\langle x, y \rangle^\perp$ determines a nonsingular 3-space $\langle x, y, z_\bullet \rangle$ containing $q^3 - q$ isotropic points not perpendicular to x and not on $\langle x, y \rangle$. It follows that, of the $|(T^{g_1})^Q| = |Q| = q \cdot (q^2)^{d-2}$ isotropic points z of G not perpendicular to x , at least

$$(q^3 - q) \cdot \frac{(q^2)^{d-2} - (q^2)^{d-2.5}}{q^2 - 1} = q[(q^2)^{d-2} - (q^2)^{d-2-1/2}]$$

of them give rise to nonsingular 3-spaces $\langle x, y, z \rangle$. Furthermore, if z denotes the centre of T^{g_2} and $\langle x, y, z \rangle$ is nonsingular, then, as in the proof of Lemma 2.7, $K = \langle T, T^{g_1}, T^{g_2} \rangle$ is a naturally embedded $\text{SU}(3, q)$ subgroup. Hence, our choice $w \in Q$ gives rise to a suitable group K with probability greater than $q[(q^2)^{d-2} - (q^2)^{d-2.5}]/|Q| = 1 - 1/q > 2/3$.

Reliability. For $K \cong \text{SU}(3, q)$, the constructive test in 6.1 succeeds with probability greater than $3/4$. Hence, our procedure finds $u \in Q$ with probability greater than $(2/3)(3/4) = 1/2$.

Timing. The stated timing arises from the call to 6.1. □

REMARK 5.4. Comparing Lemma 5.3 with [13, Lemma 6.9], we avoid factors of q in the running time at the expense of using a randomised algorithm.

We next present our algorithm to solve problem (A2); it works for $d \geq 4$.

PROCEDURE 5.5. Suppose first that $d \geq 5$. Recall the elements $l, \sigma \in J$ (defined in 4.1.4) and $s_{ij} \in \Delta^*$ of our generating set \mathcal{S}^* , and consider the following routine.

1. Fix $1 \neq t \in T$ and perform the following steps.
 - (i) Find $s \in \mathcal{S}^*$ such that T^{gs} is not perpendicular to T as follows: if $[t, t^s] \neq 1$, set $s := 1$; if $[t^l, t^s] \neq 1$, set $s := l^{-1}$; otherwise, set $s := s_{i1}$ for the first $i \in I'$ such that $[t, t^{s_{i1}}] \neq 1$.
 - (ii) Apply Lemma 5.3 at most twice to (probably) find the unique element $w \in Q$ such that $T^{gs w} = T^l$ (hence $gswl^{-1} \in N_G(T)$). Apply Lemma 4.8(iii) to write an SLP from $\Delta^* \subset \mathcal{S}^*$ to w .
 - (iii) As in Step (ii), (probably) find the unique element $u \in Q$ such that $T^{l gswl^{-1} u} = T^l$ and write an SLP from $\Delta^* \subset \mathcal{S}^*$ to u .

(iv) Replace g with $gswl^{-1}u$. [Now g normalises T and T^l .]

2. Use Ψ_J to find the scalar $\lambda \in \mathbb{F}_q^*$ induced by g on the transvection group T . Use the DLO− (see 3.3) to find the integer $0 \leq n_1 < q - 1$ such that $\zeta^{n_1} = \lambda$, and replace g with $g\sigma^{-n_1}$. [Now g centralises T and T^l .]

3. Use Lemma 5.2 to find $(d - 2) \times (d - 2)$ matrices \tilde{g} and $\tilde{\sigma}$, representing the linear transformations induced by g and σ , respectively, on Q/T relative to B' . Use the DLO+ to find the integer $0 \leq n_2 < q + 1$ such that $(\det(\tilde{\sigma}^{q-1}))^{n_2} = \det(\tilde{g})$, and replace g with $g\sigma^{-n_2}$ and \tilde{g} with $\tilde{g}\tilde{\sigma}^{-n_2}$.

4. Write down the unique matrix g' of $SU(V)_{e_1, e_{-1}}$ inducing \tilde{g} on $Q(x)/T(x)$ and use (A1) to write an SLP of length $O(d^2 \log q)$ from \mathcal{T} to g' . Evaluate that SLP from \mathcal{S}^* to obtain an element $g_0 \in gZ(G)$. List the (no more than d) elements of $Z(G) = \langle z \rangle$ and then write an SLP of length $O(\log d)$ from $\{z\}$ to gg_0^{-1} and concatenate it with the one to g' to obtain the desired SLP to g .

In the case $d = 4$, proceed exactly as above, replacing Lemma 4.8(iii) with Lemma 6.10(iii) and Lemma 5.2 with Lemma 6.12.

Correctness. We first claim that $g \in N_G(T) \cap N_G(T^l)$ at the end of Stage 1. It suffices to show that the element s obtained in Step (i) behaves as stated. We may assume that $[t, t^g] = 1 = [t^l, t^{g^l}]$. For any $h \in G$, let h' denote $h\Psi^{-1}$; then $xg' \in \langle x, y \rangle^\perp$. There is at least one vector in $\mathcal{B} \cap \langle e_1, e_{-1} \rangle^\perp$ that is not perpendicular to xg' ; let $i \in I'$ be the index corresponding to such a vector. Then s'_{i1} moves xg' inside the line $\langle y, xg' \rangle$. Hence $x.(gs_{i1})' \in y^\perp \setminus \langle x, y \rangle^\perp$ so that $x.(gs_{i1})' \notin x^\perp$. For such i , $T^{gs_{i1}}$ is not perpendicular to T , as required.

Next, by Lemma 2.3(iii), σ induces the scalar $\rho\bar{\rho} = \zeta$ on T . It follows that g centralises T (and hence T^l) at the end of Stage 2. By Lemma 2.3(ii), $\det(\tilde{\sigma})$ generates the cyclic subgroup of $\mathbb{F}_{q^2}^*$ consisting of the determinants of all transformations induced on Q/T by elements of $g \in N_G(T) \cap N_G(T^l)$. Since $C_{\langle \sigma \rangle}(T) = \langle \sigma^{q-1} \rangle$, the DLO+ succeeds in finding the integer n_2 in Stage 3. Hence $\tilde{g} \in SU(d - 2, q)$ at the end of Stage 3. The correctness of Stage 4 is now clear: since gg_0^{-1} centralises T and T^l and induces 1 on Q/T , it follows that $gg_0^{-1} \in Z(G)$.

Reliability. The only randomised subroutine is Lemma 5.3 occurring in Stage 1, Steps (ii) and (iii). Each application of that lemma succeeds with probability greater than $1/2$ so that, applying this lemma again in each of Steps (ii) and (iii) if necessary, we ensure that the desired element of Q is obtained with probability greater than $3/4$. Hence, we find both u and w with probability greater than $1/2$.

Timing. The timing is dominated by $O(\xi + \chi[d^2 + \log q])$ for the calls to Lemmas 5.2 and 5.3.

REMARK 5.6. An additional $O(d^5 \log^2 q)$ (for $O(d^2 \log q)$ matrix multiplications) is required to evaluate the SLP, obtained in Procedure 5.5, from \mathcal{T} . Hence the total time to find $g\Psi^{-1}$ for any given $g \in G$ is

$$O(\xi + \chi[d^2 + \log q] + d^5 \log^2 q).$$

It follows that Ψ is $O(\xi + \chi\{d^2 + \log q\} + d^5 \log^2 q)$ -effective as stated in Theorem 1.1.

6. Low-dimensional cases

We have so far proved Theorem 1.1 for $d \geq 5$, presuming the existence of algorithms to handle the cases $d = 3, 4$. The algorithms in [13, 6.6.1 and 6.6.2] for the latter cases cannot be used here, since, once again, their running times contain factors of q . In this final section we present new constructive recognition algorithms for $d = 3, 4$ that satisfy our more stringent timing goals. We assume throughout this section that $q \geq 16$; for smaller q , the algorithms in [13, 6.6.1 and 6.6.2] may be used.

6.1. $SU(3, q)$

Let $G = \langle \mathcal{S} \rangle$ be a nontrivial homomorphic image of $SU(3, q)$. We present an $O(\xi + \chi \log q)$ -time Las Vegas algorithm to produce a data structure giving rise to an $O(\chi)$ -effective epimorphism $\Psi: SU(3, q) \rightarrow G$.

6.1.1. A subgroup $L \circ \langle z \rangle$

Exactly as in [13, 6.6.1], with probability greater than $1 - 1/2^8$, find an element $a \in G$ of $\text{ppd}^\#(p; 2k)$ -order dividing $(q + 1)/(2, q + 1)$. Then, for a single choice of G -conjugate b of a , we have $\langle a, b \rangle \cong \text{SL}(2, q) \circ \langle z \rangle$ with probability greater than $1/2$, where z has $\text{ppd}^\#(p; 2k)$ -order dividing $q + 1$. For up to sixteen choices of b , test whether or not this is the case as follows: use Lemma 3.2 to compute probable generators for $L := \langle a, b' \rangle$; use the $\text{SL}(2, q)$ -oracle to test whether $L \cong \text{SL}(2, q)$ and, if so, to obtain an effective isomorphism $\Psi_L: \text{SL}(2, q) \rightarrow L$. Since each call to Lemma 3.2 successfully produces generators for L with probability greater than $1/2$, at least one of our sixteen choices will produce a suitable L with probability greater than $1 - 1/2^8$.

Unlike the algorithm [13] that we have been following thus far, we wish to construct and use the element z such that $\langle a, b \rangle = L \circ \langle z \rangle$. For each of three distinct transvection groups T_i ($i = 1, 2, 3$) of L , proceed as follows: compute T_i^a ; write down a matrix $c' \in \text{SL}(2, q)$ sending $T_i^a \Psi_L^{-1}$ to $T_i \Psi_L^{-1}$ for $i = 1, 2, 3$; set $c := c' \Psi_L$ and $z := ac$. It is clear that z centralises L since it fixes the three distinct transvection groups T_i . Also, since a induces a $\text{ppd}^\#(p; 2k)$ -scalar on a nonsingular point not in the support of L , it follows that z induces the same scalar on this point.

Reliability. This is greater than $1 - 1/2^8 - 1/2^8 = 1 - 1/2^7$.

Timing. $O(\xi + \chi + \mu \log^2 q)$ is required for the various choices of element in G , calls to the $\text{SL}(2, q)$ -oracle, and uses of Lemma 3.2 (with $d = 3$).

6.1.2. Constructing Q

Use the $\text{SL}(2, q)$ -oracle to construct a transvection group T of L , an element $l \in L \setminus N_L(T)$ and $h \in L$, of order $q - 1$, normalising T and T^l . Choose $g \in G$ and set $L_1 := \langle T, T^{lg} \rangle$. Use the $\text{SL}(2, q)$ -oracle to test whether $L_1 \cong \text{SL}(2, q)$ and, if so, to construct an element h_1 of order $q - 1$ normalising T and T^{lg} . Then $u := [h, h_1]$ is the identity on the 1-space corresponding to T in the underlying module, so that $u \in Q = O_p(N_G(T))$. Since z centralises T and induces an irreducible transformation of Q/T (regarded as \mathbb{F}_p -space), $u \in Q \setminus T$ if and only if $u^z \neq u$. If $u \notin T$, then $Q = \langle u^{z^i} \mid 0 \leq i < 2k \rangle$ (since T is the Frattini subgroup of Q), so we have a $2k$ -element generating set for Q .

Reliability. We claim that $u \in Q \setminus T$ if and only if L and L_1 are distinct. Assuming that this is the case, we generate Q with probability $1 - (q + 1)/(q^3 + 1) > 0.99$. Note that Q/T acts regularly on the set of $\text{SL}(2, q)$ -subgroups of G containing T . Hence, for $\text{SL}(2, q)$ -subgroups L and L_1 , the unique $w \in Q$ conjugating $\langle h \rangle$ to $\langle h_1 \rangle$ is in T if and only if $L = L_1$. The claim now follows.

Timing. We obtain the $2k$ generators for Q in $O(\xi + \chi + \mu k)$ time.

6.1.3. Constructing \mathbb{F}_{q^2}

The elements z and h induce automorphisms \tilde{z} and \tilde{h} , respectively, of Q/T . Since the factor group $\text{PGU}(3, q)_x/Q(x) \cong \text{GF}(q^2)^*$, and z has $\text{ppd}^\#(p; 2k)$ -order, it follows that

$$\mathbb{F}_{q^2} := \mathbb{F}_p\text{-linear span of } \{\tilde{z}^i \mid 0 \leq i < 2k\}$$

is a field of order q^2 . Since h induces a transformation of Q/T of order $q - 1$, \tilde{h} is a generator of the field

$$\mathbb{F}_q := \mathbb{F}_p\text{-linear span of } \{\tilde{h}^j \mid 0 \leq j < k\} \subset \mathbb{F}_{q^2},$$

namely the subfield of \mathbb{F}_{q^2} fixed by the automorphism $x \mapsto x^q = \bar{x}$. In order that we may compute effectively with \mathbb{F}_{q^2} and \mathbb{F}_q , we shall need to compute the minimal polynomial of \tilde{z} as a transformation on the \mathbb{F}_p -space Q/T . We first need the following result, which is essentially Lemma 4.8 for $d = 3$. However, we need to prove it in a more general setting (GS1) than is required here, in order that it may be used in the proofs of Lemmas 4.8 and 6.10.

Suppose that our input group G is a naturally embedded $\text{SU}(3, q)$ -subgroup of a black-box unitary group H of dimension $n \geq 3$. Let $Q_H = O_p(N_H(T))$ and let $(\cdot, \cdot)_{Q_H/T}$ denote a non-degenerate $N_H(T)'/Q_H$ -invariant hermitian form on Q_H/T . In particular, Q/T is a nonsingular point of Q_H/T . (GS1)

LEMMA 6.1. *Let $u \in Q$ be the element constructed in 6.1.2, fix $1 \neq t \in T$ and, for $1 \leq i \leq k$, set $t_i := t^{h^{i-1}}$, $u_i := u^{h^{i-1}}$ and $u_{k+i} := (u^z)^{h^{i-1}}$. Then the following hold.*

- (i) $B'_p := (u_i T)_{i=1}^{2k}$ is an \mathbb{F}_p -basis for Q/T .
- (ii) $\mathcal{S}_Q^* := \{t_s, u_i \mid 1 \leq s \leq k, 1 \leq i \leq 2k\}$ generates Q .
- (iii) In (GS1), there is a deterministic $O(\chi)$ -time algorithm which, for any given $y \in Q_H$, writes the projection of yT on Q/T along $(Q/T)^\perp$ as an \mathbb{F}_p -vector relative to B'_p .
- (iv) There is a deterministic $O(\chi + \mu \log q)$ -time algorithm which, for any given $w \in Q$, writes an SLP of length $O(\log q)$ from \mathcal{S}_Q^* to w .

Proof. Statement (i) is clear, since \tilde{h} is a scalar transformation of Q/T of order $q - 1$ and uT and $u^z T$ are in different 1-spaces of Q/T viewed as \mathbb{F}_q -space.

Statement (ii) follows immediately from Statement (i).

For Statement (iii), precompute $r_i = [u_1, u_{k+i}] \in T$ and $s_i = [u_{k+1}, u_i] \in T$ for $1 \leq i \leq k$. Use the $\text{SL}(2, q)$ -oracle to find \mathbb{F}_p -vectors (a_1, \dots, a_k) and (a_{k+1}, \dots, a_{2k}) , representing $[u_{k+1}, y] \in T$ relative to s_1, \dots, s_k , and $[u_1, y] \in T$ relative to r_1, \dots, r_k , respectively, and return the vector $(a_1, \dots, a_k, a_{k+1}, \dots, a_{2k})$.

To see that this is the desired vector, write $yT \in Q_H/T$ additively as

$$yT = (yT)_\parallel + (yT)_\perp, \quad \text{where } (yT)_\parallel \in Q/T \text{ and } (yT)_\perp \in (Q/T)^\perp \leq Q_H/T.$$

Then

$$[uT, yT] = [uT, (yT)_\parallel], \quad \text{for all } uT \in Q/T.$$

Hence, if $(yT)_\parallel = \sum_{i=1}^{2k} b_i (u_i T)$ for $b_i \in \mathbb{F}_p$, then

$$[u_1, y] = \prod_{i=1}^k [u_1, u_{k+i}]^{b_{k+i}} = \prod_{i=1}^k r_i^{b_{k+i}}$$

and

$$[u_{k+1}, y] = \prod_{i=1}^k [u_{k+1}, u_i]^{b_i} = \prod_{i=1}^k s_i^{b_i}.$$

It follows immediately that $a_i = b_i$ for $1 \leq i \leq 2k$. The stated timing is clear.

Finally, for Statement (iv), use Statement (iii) to write wT as an \mathbb{F}_p -vector (a_1, \dots, a_{2k}) relative to B'_p . Set

$$w_0 := \prod_{i=1}^{2k} u_i^{a_i},$$

and then use the $\text{SL}(2, q)$ -oracle to write an SLP from $\{t_1, \dots, t_k\}$ to $w w_0^{-1} \in T$. Together with the a_i we easily obtain the desired SLP to w . The computation of w_0 takes $O(\mu \log q)$ time. \square

COROLLARY 6.2. *In deterministic $O(\chi k)$ time, one can find the minimal polynomial*

$$f(x) = x^{2k} - \sum_{i=1}^{2k} a_i x^{i-1}$$

of \tilde{z} over \mathbb{F}_p .

Proof. Fix $w \in Q \setminus T$ and, for $0 \leq i \leq 2k$, apply Lemma 6.1(ii) to express the vector $(wT)^{\tilde{z}^i}$ as an \mathbb{F}_p -vector z_i relative to B'_p . Using linear algebra in $(\mathbb{F}_p)^{2k}$, find integers a_1, \dots, a_{2k} such that $z_{2k} = a_1 z_0 + \dots + a_{2k} z_{2k-1}$, and set

$$f(x) := x^{2k} - \sum_{i=1}^{2k} a_i x^{i-1}.$$

The minimal polynomial $g(x)$ of \tilde{z} has degree $2k$ and, since $f(\tilde{z})$ acts trivially on Q/T , it follows that $f(x) = g(x)$. \square

The natural module. Now that we have the minimal polynomial of \tilde{z} , we can compute effectively with the field \mathbb{F}_{q^2} . In particular, we can construct the 3-space upon which G acts naturally.

Set $V := (\mathbb{F}_{q^2})^3$, and equip V with a hermitian form by designating the usual basis to be a standard basis $\mathcal{B} = e_1, v, e_{-1}$ relative to this form. In particular, $x := \langle e_1 \rangle$ and $y := \langle e_{-1} \rangle$ are isotropic. We will write elements of $\text{SU}(V)$ as matrices relative to \mathcal{B} .

Timing. The timing for 6.1.3 is $O(\chi k)$, dominated by the use of Corollary 6.2.

6.1.4. Recognising Q

We have assumed the existence of an epimorphism $\mathrm{SU}(V) \rightarrow G$, and we may therefore assume the existence of an epimorphism

$$\Psi: \mathrm{SU}(V) \rightarrow G \text{ sending } T(x) \rightarrow T \text{ and } T(y) \rightarrow T^l. \quad (16)$$

Eventually, we will construct an epimorphism Ψ satisfying (16). We proceed now, as we did in Subsection 4.4 of the main algorithm, by constructing such a map effectively from

$$\underline{Q}(x) = \left\{ \left(\begin{array}{ccc} 1 & 0 & 0 \\ \lambda & 1 & 0 \\ \nu & -\lambda^q & 1 \end{array} \right) \mid \nu + \nu^q + \lambda^{q+1} = 0 \right\}$$

to Q . We first make some observations, which hold for *any* epimorphism Ψ satisfying (16).

Since $h \in L$ induces the automorphism \tilde{h} of Q/T , it follows that

$$h' := h\Psi^{-1} = \mathrm{diag}(\tilde{h}, 1, 1/\tilde{h}). \quad (17)$$

That is, $h\Psi^{-1}$ is independent of the choice of Ψ satisfying (16). We do not yet have sufficient data to compute the preimage of $z\Psi^{-1} \in \mathrm{SU}(V)$ but, since \mathbb{F}_{q^2} was constructed using automorphisms of Q/T , we can determine the automorphism of $\underline{Q}(x)$ that it induces. Let α denote the automorphism of Q induced under conjugation by z . Then Ψ induces an isomorphism $\varphi: \underline{Q}(x) \rtimes A \rightarrow Q \rtimes \langle \alpha \rangle$ for some group of automorphisms A of $Q(x)$. Denoting $\alpha\varphi^{-1}$ by α' , by definition of α we have

$$\left(\begin{array}{ccc} 1 & 0 & 0 \\ \lambda & 1 & 0 \\ \nu & -\lambda^q & 1 \end{array} \right)^{\alpha'} = \left(\begin{array}{ccc} 1 & 0 & 0 \\ \lambda\tilde{z} & 1 & 0 \\ \nu & -(\lambda\tilde{z})^q & 1 \end{array} \right). \quad (18)$$

That is, α' is also independent of Ψ . Furthermore we can use (18) to compute the image under α' of any given element of $\underline{Q}(x)$ in $O(\mu)$ time.

We are now ready to determine a specific Ψ satisfying (16). Fix $u'_0 \in \underline{Q}(x) \setminus T(x)$ and $r \in Q \setminus T$. Since $\mathrm{GU}(V)_{x,y}$ is transitive on the nonzero vectors of \mathbb{F}_{q^2} -space $\underline{Q}(x)/T(x)$, there exists *some* epimorphism Ψ , satisfying (16), that induces a map $\varphi: \underline{Q}(x) \rtimes A \rightarrow Q \rtimes \langle \alpha \rangle$ sending $u'_0 T(x) \mapsto rT$. Using Procedure 6.3 below, we will demonstrate that such a φ is unique. We first make some observations.

(a) For any $w' \in \underline{Q}(x)$, by (17), we have $(w'^{h'})\varphi = (w'\varphi)^h$.

(b) Let $f(x) = x^{2k} - \sum_{i=1}^{2k} a_i x^{i-1}$ be the minimal polynomial of α on Q/T (and hence of α' on $\underline{Q}(x)/T(x)$) constructed in Corollary 6.2. Then

$$s' := (u'_0)^{\alpha'^{2k}} \left[\prod_{i=1}^{2k} \left((u')^{\alpha'^i} \right)^{a_i} \right]^{-1} \in T(x). \quad (19)$$

(c) For any $w \in Q$, define

$$s_w := w^{\alpha^{2k}} \left[\prod_{i=1}^{2k} (w^{\alpha^{i-1}})^{a_i} \right]^{-1} \in T; \quad (20)$$

then $u'_0\varphi = w$ implies that $s'\varphi = s_w$.

(d) Since f is irreducible, $s_{rt} = s_r t^{1-\sum_{i=1}^{2k} a_i} = s_r t^{f(1)} \neq s_r$ for each $t \in T$; and $s' = s_{rt}\varphi^{-1}$ if and only if

$$(t\varphi^{-1})^{f(1)} = s'(s_r\varphi^{-1})^{-1}. \quad (21)$$

The following procedure finds the unique element $t_0 \in T$ such that $\alpha' \mapsto \alpha$, $u'_0 \mapsto rt_0$ extends to an isomorphism $Q(x) \rtimes \langle \alpha' \rangle \rightarrow Q \rtimes \langle \alpha \rangle$ (and hence to our isomorphism φ).

PROCEDURE 6.3. Perform each of the following steps.

1. Use (18) to compute $1 \neq t'_1 := [u'_0, u_0^{\alpha'}] \in T(x)$ and set $t_1 := [r, r^\alpha] \in T$; for $2 \leq i \leq k$, set

$$t'_i := t'_1 h^{i-1} \quad \text{and} \quad t_i := t_1^{h^{i-1}}.$$

2. Use (20) to construct s_r in $O(\mu \log q)$ time. Use the $SL(2, q)$ -oracle to write an SLP from $\{t_1, \dots, t_k\}$ to s_r and then evaluate it from $\{t'_1, \dots, t'_k\}$ to find $s'_r := s_r \varphi^{-1}$.

3. For the element $s' \in T(x)$, defined in (19), compute

$$s'(s'_r)^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ v & 0 & 1 \end{pmatrix} \in T(x)$$

for some $v \in \mathbb{F}_{q^2}$ such that $v + v^q = 0$. Then equation (21) becomes

$$t'(\zeta)^{f(1)} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \zeta & 0 & 1 \end{pmatrix}^{f(1)} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ v & 0 & 1 \end{pmatrix}$$

for unknown $\zeta \in \mathbb{F}_{q^2}$ with $\zeta + \zeta^q = 0$.

4. Set $\zeta := v/f(1)$ to obtain the desired solution $t'_0 := t'(\zeta)$.

5. Use linear algebra to write an SLP from $\{t'_1, \dots, t'_k\}$ to t'_0 and evaluate it from $\{t_1, \dots, t_k\}$ to find $t_0 = t'_0 \varphi \in T$.

Return $u_0 := rt_0$.

Correctness. In Step 1, $[r, r^\alpha] = [rt, (rt)^\alpha]$ for all $t \in T$; it therefore follows that $t_1 = t'_1 \varphi \in T$. Also, by observation (a), $t'_i = t_i \varphi^{-1}$. That u_0 behaves as stated now follows from observations (a)–(d).

Timing. The computation of s_r and the various $SL(2, q)$ -oracle calls require a time of $O(\chi + \mu \log q)$.

REMARK 6.4. For $1 \leq i \leq k$, set $u_i := u_0^{h^{i-1}}$, $u_{k+i} := u_0^{zh^{i-1}}$, $u'_i := u_0^{h^{i-1}}$, and $u'_{k+i} := u_0^{\alpha' h^{i-1}}$, and redefine

$$\begin{aligned} \mathcal{T}_Q &:= \{t'_s, u'_i \mid 1 \leq s \leq k, 1 \leq i \leq 2k\}; \\ \mathcal{S}_Q^* &:= \{t_s, u_i \mid 1 \leq s \leq k, 1 \leq i \leq 2k\}. \end{aligned} \tag{22}$$

Then the bijection $\mathcal{T}_Q \rightarrow \mathcal{S}_Q^*$ sending $t'_s \mapsto t_s$ and $u'_i \mapsto u_i$ extends to φ . Also, the routine presented for Lemma 6.1(iv) goes through without change for our modified set \mathcal{S}_Q^* , so φ is now *effectively* defined on $Q(x)$.

6.1.5. Uniqueness of Ψ

The isomorphism φ is induced by some epimorphism Ψ satisfying (16). We will soon see that such a Ψ is uniquely determined, and it is this unique Ψ that we will construct. We first need a version of Lemma 5.3 for $d = 3$ (note that the following is a deterministic algorithm, unlike our general transitivity lemma).

LEMMA 6.5. *The unique element of Q conjugating a given $T^{g_1} \neq T$ to another given $T^{g_2} \neq T$ can be found in deterministic $O(\chi)$ time.*

Proof. We use the following routine.

1. For $i = 1, 2$, use the $\text{SL}(2, q)$ -oracle to construct an element $h_i \in \langle T, T^{g_i} \rangle \cong \text{SL}(2, q)$ normalising T and T^{g_i} , where $|h_i| = 2$ if q is odd, and $|h_i| = q - 1$ if q is even.

1'. If q is even, use DLO- twice inside $\langle T, T^l \rangle$, as in 3.3.1, to arrange for h, h_i ($i = 1, 2$) to induce the same scalar $\tilde{h}^2 \in \mathbb{F}_q$ on T . Then use the DLO- to find the integer n such that $\tilde{h}^n(1 - \tilde{h}) = 1$.

2. Set $u := (h_1 h_2)^{(p+1)/2}$ if q is odd, or $u := (h_1^{-1} h_2)^{h_1^n}$ if q is even.

3. Use the $\text{SL}(2, q)$ -oracle again inside $\langle T, T^{g_2} \rangle$ to find the unique $t \in T$ such that $T^{g_1 u t} = T^{g_2}$.

Return the element ut .

Correctness. It suffices to show that $\langle T, T_1^u \rangle = \langle T, T_2 \rangle$, as required in the last step of the procedure. Let w denote the unique element of Q such that $\langle h_1 \rangle^w = \langle h_2 \rangle$. If $p = 2$, then (since h_1 and h_2 induce the same scalar on T) it follows that $h_1^w = h_2$. Note that w acts on the 2-space Q/T via $v \mapsto v + c$ for some $c \in Q/T$. We also have $h_1 : v \mapsto \tilde{h}v$, so that $h_2 = h_1^w : v \mapsto \tilde{h}v + c(1 - \tilde{h})$. Then $u = (h_1^{-1} h_2)^{h_1^n} : v \mapsto v + c$, so that $uT = wT$, as required. The case $p > 2$ is similar but easier.

Timing. This is dominated by the various $\text{SL}(2, q)$ -oracle calls. □

For an epimorphism Ψ satisfying (16) and inducing $\varphi : Q(x) \rtimes \langle \alpha' \rangle \rightarrow Q \rtimes \langle \alpha \rangle$, the following routine finds the unique preimage $g\Psi^{-1}$, modulo scalars, of any given $g \in G$.

PROCEDURE 6.6. Set $u := u_1 \in Q$ and $u' := u'_1 \in Q(x)$. Set $T_0 := T, T_1 := T^l, T_2 := T_1^u$ and $T_3 := T_1^{u^z}$. Also set $x_0 := x, x_1 := y, x_2 := x_1.u'$ and $x_3 := x_1.u'^z$, and proceed as follows.

1. For $0 \leq i \leq 3$, compute T_i^g and test whether $T_i^g = T_j$ by testing $[t_i^g, t_j] = 1$ for a single generator t_i of T_i and t_j of T_j .

(i) If $T_i^g = T_j$, set $y_i := x_j$.

(ii) If $T_i^g \neq T_j$, use Lemma 6.5 to find the unique $r_i \in Q$ such that $T_1^{r_i} = T_i^g$. Use Lemma 6.1(iv) to write an SLP of length $O(\log q)$ from \mathcal{S}_Q^* to r_i and evaluate it from \mathcal{T}_Q to obtain $r'_i \in Q(x)$. Now set $y_i := x_1 r'_i$.

2. Compute the unique matrix in $\text{PSU}(V)$ sending x_i to y_i for $0 \leq i \leq 3$.

Correctness. It is clear, from the construction of the points y_i , that $g\Psi^{-1}$ sends x_i to y_i . That there is a unique matrix in $\text{PSU}(V)$ with this property follows from the fact that the groups T_0, T_1, T_2 and T_3 represent points in ‘general position’.

Timing. This is dominated by $O(\chi)$ for the calls to Lemmas 6.1 and 6.5.

PROPOSITION 6.7. *There is a unique epimorphism $\Psi : \text{SU}(V) \rightarrow G$ satisfying (16) and inducing $\varphi : Q(x) \rtimes \langle \alpha' \rangle \rightarrow Q \rtimes \langle \alpha \rangle$.*

Proof. Apply Procedure 6.6 to the element $l \in L$ to obtain $A_l \in \text{SU}(V)$, which is $l\Psi^{-1}$ up to scalar, and write it relative to our standard basis \mathcal{B} . Then there is a unique element

$$l' = \begin{pmatrix} 0 & 0 & \delta^q \\ 0 & 1 & 0 \\ 1/\delta & 0 & 0 \end{pmatrix} \in A_l Z(\text{SU}(V)) \quad (23)$$

fixing the basis vector v . Since our target epimorphism Ψ maps $\text{SU}(\langle e_1, e_{-1} \rangle)$ to L , it follows that $l'\Psi = l$. Finally, since Ψ is now determined on $Q(x)$ and $Q(x)l'$, it follows that Ψ is unique. \square

6.1.6. A data structure for G

For the element l' in (23), set $\mathcal{T} := \mathcal{T}_Q \cup \mathcal{T}_Q^{l'}$ and $\mathcal{S}^* := \mathcal{S}_Q^* \cup (\mathcal{S}_Q^*)^{l'}$, where \mathcal{T}_Q and \mathcal{S}_Q^* are defined in (22). Extend the bijection $\mathcal{T}_Q \rightarrow \mathcal{S}_Q^*$ in the obvious way to obtain a bijection $\mathcal{T} \rightarrow \mathcal{S}^*$. This bijection uniquely determines Ψ but, in order to construct Ψ effectively, we need a version of Procedure 5.5 to solve problem (A2) for $d = 3$. Note, again, that the following routine is deterministic, in contrast to the algorithm for the general case.

PROCEDURE 6.8. Use Procedure 6.6 to compute $g\Psi^{-1}$ up to scalar. Next, use Section 5, (A1), to write an SLP from \mathcal{T} modulo scalars to that matrix, and evaluate this SLP from \mathcal{S}^* to obtain an element $h \in G$. Then $gh^{-1} \in Z(G)$, and we modify the SLP in the obvious way to obtain an SLP to g .

Timing. This is dominated by $O(\chi)$ for the calls to Procedure 6.6.

Total timing and reliability for 6.1. A suitable data structure for G is found, with probability at least $1 - 1/2^7 - 0.01 > 0.95$, in time $O(\xi + \chi \log q + \mu \log^2 q) = O(\xi + \chi \log q)$.

6.2. $\text{SU}(4, q)$

Let $G = \langle \mathcal{S} \rangle$ be a nontrivial homomorphic image of $\text{SU}(4, q)$. In addition to a general-purpose algorithm, we also consider a special case that arises in the main algorithm for $d > 4$ in 4.3.2, where the output isomorphism is required to have some additional properties (see 6.2.7).

6.2.1. The subgroup L

We begin along the same lines as the algorithm in [6, 6.4], for recognising $\text{SU}(4, q)$ in its natural representation. Exactly as in [6, 6.4.1], with probability greater than $1 - 1/2^4$, find an element τ of $\text{ppd}^\#(p; 2k) \cdot \text{ppd}^\#(p; 6k)$ -order, divisible also by 8 if $k = 1$ and p is Mersenne. Set $a := \tau^{2(q^2 - q + 1)}$. Choose up to sixteen conjugates $b = a^g$, and for each one: use Lemma 3.2 to find $L := \langle a, b \rangle'$, use the $\text{SL}(2, q)$ -oracle to test whether $L \cong \text{SL}(2, q)$ and, if so, to construct an effective isomorphism $\Psi_L: \text{SL}(2, q) \rightarrow L$.

Reliability. As in [6, 6.3.1], with probability greater than $1/2$ for a single choice b , $\langle a, b \rangle' \cong \text{SL}(2, q)$. For such b , Lemma 3.2 correctly computes $L = \langle a, b \rangle'$ with probability greater than $1/2$. It follows that at least one of the sixteen choices b will succeed with probability greater than $1 - 1/2^4$. Hence we obtain a suitable L with probability greater than $1 - 1/2^4 - 1/2^4 = 7/8$.

Timing. All of the constructions are obtained in $O(\xi + \chi + \mu \log^2 q)$, as in 6.1.1.

6.2.2. *The subgroups J_i*

Use the $\text{SL}(2, q)$ -oracle to construct a transvection group T of L , and let $l \in L \setminus N_L(T)$. We now use the approach taken in [6, 6.4.4], to construct distinct $\text{SU}(3, q)$ -subgroups J_1 and J_2 of G . For $i = 1, 2$, proceed as follows for each of up to sixteen choices $g_i \in G$. Set $J_i := \langle T, T^l, T^{g_i} \rangle$. Use 6.1 to test whether $J_i \cong \text{SU}(3, q)$ and, if so, to construct an effective isomorphism $\Psi_i: \text{SU}(3, q) \rightarrow J_i$. Use Procedure 6.8 to test whether or not $t^{g_1} \in J_2$ and, if not, return J_1, J_2, Ψ_1 and Ψ_2 .

Reliability. For $i = 1, 2$, by Lemma 2.7, $J_i \cong \text{SU}(3, q)$ with probability greater than $1/2$. If $J_i \cong \text{SU}(3, q)$, the routine in 6.1 succeeds with probability greater than 0.95. Hence, a single choice of conjugate produces an isomorphism Ψ_i with probability greater than 0.475. Since two nonsingular 3-spaces inside a 4-space are equal with probability less than $1/q^2 < 0.01$, we obtain distinct $\text{SU}(3, q)$ -subgroups with probability greater than $1 - (2(0.525)^{16} + 0.01) > 0.95$.

Timing. $O(\xi + \chi \log q)$ is required for no more than 32 calls to 6.1.

6.2.3. *The subgroup L_0*

Our approach to constructing Ψ will involve modifying the isomorphisms Ψ_1 and Ψ_2 that we have just obtained so that, when viewed as maps on suitable 3-spaces of a 4-space, they extend simultaneously to a unique epimorphism $\Psi: \text{SU}(4, q) \rightarrow G$. In order to carry out the necessary modifications, we need some additional constructions.

For $i = 1, 2$, let $Q_i = O_p(N_{J_i}(T))$. Use Ψ_1 to construct $r, s \in Q_1$ such that rT and sT are nonsingular vectors spanning Q_1/T as \mathbb{F}_q -space and use Ψ_2 to construct $1 \neq u \in Q_2$. The next procedure constructs an $\text{SL}(2, q)$ subgroup L_0 centralising L , as well as an element $z \in L_0$ such that $(Q_2/T)^z$ is perpendicular to Q_1/T in Q/T (that is, such that $[Q_1, Q_2^z] = 1$).

PROCEDURE 6.9. If $[r, u] = 1 = [s, u]$, then report failure; otherwise, proceed as follows.

1. For $i = 1, 2$, use Ψ_i to find an element $\sigma_i \in C_{J_i}(L)$ of order $q + 1$.
2. [Repeat this step up to four times.] Set $A_0 := \langle \sigma_1, \sigma_2 \rangle$ and use Lemma 3.2 to find a probable set of generators for $L_0 := A_0'$. Use the $\text{SL}(2, q)$ -oracle to test whether $L_0 \cong \text{SL}(2, q)$ and, if so, to obtain an effective isomorphism $\Psi_0: \text{SL}(2, q) \rightarrow L_0$.
3. Fix distinct transvection groups S_j of L_0 for $j \in \{\infty, 0, 1\}$. For $i = 1, 2$, use the transvection groups $S_j^{\sigma_i}$ ($j = \infty, 0, 1$), together with Ψ_0 , to find a 2×2 matrix $\tilde{\sigma}_i$ representing the element of $\text{PGL}(2, q)$ induced by σ_i on the 2-space underlying L_0 . (See [13, p. 47], for example.)
4. For $i = 1, 2$, find the eigenvectors $\langle w_i \rangle, \langle w'_i \rangle$ of $\tilde{\sigma}_i$ over \mathbb{F}_{q^2} .
5. Let M_1 and M_2 be elements of $\text{SL}(2, q)$ sending $\langle w_2 \rangle \mapsto \langle w_1 \rangle$ and $\langle w_2 \rangle \mapsto \langle w'_1 \rangle$, respectively, and set $z_1 := M_1 \Psi_0$ and $z_2 := M_2 \Psi_0$.
6. If $[r, u^{z_1}] = 1 = [s, u^{z_1}]$, set $z := z_1$; otherwise, set $z := z_2$.

Return L_0 and z .

Correctness. Let $V = (\mathbb{F}_{q^2})^4$ and let $\Psi: \text{SU}(V) \rightarrow G$ denote any epimorphism. Then, for some nonsingular $v_1, v_2 \in V$ and $\lambda_r, \lambda_s, \lambda_u \in \mathbb{F}_{q^2}$, we have

$$r = r_1(v_1, \lambda_r)\Psi, \quad s = r_1(\rho_s v_1, \lambda_s)\Psi \quad \text{and} \quad u = r_1(v_2, \lambda_u)\Psi,$$

with $\rho_s \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q$ (see (4)). It follows immediately from Lemma 2.1(iii) that $[r, u] = 1 = [s, u]$ if and only if $(v_1, v_2) = 0$.

Assume that either $[r, u] \neq 1$ or $[s, u] \neq 1$ (so that $(v_1, v_2) \neq 0$). Then, in Step 2, $A_0 \leq C_G(L)$ and, for $i = 1, 2$, the element $\sigma'_i = \sigma_i \Psi^{-1}$ induces a $\text{ppd}^\#(p; 2k)$ -element on $\langle v_1, v_2 \rangle$ (of the same order for both $i = 1$ and $i = 2$) fixing $\langle v_i \rangle$ and $\langle v'_i \rangle = \langle v_i \rangle^\perp \cap \langle v_1, v_2 \rangle$. Since $(v_1, v_2) \neq 0$, it follows that A_0 acts irreducibly on $\langle v_1, v_2 \rangle$, and hence that $A'_0 = L_0 \cong \text{SL}(2, q)$.

For $i = 1, 2$, $\langle v_i \rangle$ and $\langle v'_i \rangle$ are the only nonsingular points in $\langle v_1, v_2 \rangle$ fixed by σ'_i (it acts nontrivially on $\langle v_i \rangle$ and as 1 on $\langle v'_i \rangle$). Observe that, since we have only the projective action of σ'_i (via the action of σ_i on the transvection groups of L_0 in Step 3), we cannot distinguish between $\langle v_i \rangle$ and $\langle v'_i \rangle$. Nevertheless, either $\langle v_2 \rangle M_1 = \langle v'_1 \rangle$ or $\langle v_2 \rangle M_2 = \langle v'_1 \rangle$; Step 6 decides which is the case, and defines z to be z_1 or z_2 in $L_0 \leq C_G(L)$ accordingly.

Reliability. $(v_1, v_2) = 0$ with probability $1/(q^2 - q - 1) < 1/100$. The only part of the algorithm that is randomised is Step 2, where Lemma 3.2 produces generators for A'_0 with probability greater than $1/2$. The subsequent $\text{SL}(2, q)$ -oracle call then determines whether or not we have A'_0 . Hence, we will obtain a suitable group L_0 in at least one of the four repetitions of Step 2 with probability greater than $1 - 1/2^4$. The procedure therefore returns L_0 and z with probability greater than $1 - 1/2^4 - 0.01 > 0.9$.

Timing. $O(\xi + \chi + \mu \log^2 q)$ is needed for Lemma 3.2 and the $\text{SL}(2, q)$ -oracle calls.

6.2.4. Recognising Q

Replace J_2 with J_2^z and Ψ_2 with $\Psi_2 \circ z$, defined effectively, for $j \in J_2$ and $j' \in \text{SU}(3, q)$, via the equations

$$j'(\Psi_2 \circ z) = (j' \Psi_2)^z \quad \text{and} \quad j(\Psi_2 \circ z)^{-1} = j^{z^{-1}} \Psi_2^{-1}.$$

We may assume that we have chosen standard bases e_i, v_i, e_{-i} for our two 3-spaces in such a way that each element of $Q_i \Psi_i^{-1}$ has the form $r_1(vv_i, \lambda)$ for $v, \lambda \in \mathbb{F}_{q^2}$ (see (4)). For $i \in \{1, 2\}$ and $1 \leq j \leq 2k$, use Ψ_i to construct the generator

$$u_{2k(i-1)+j} := r_1(\rho^{j-1}v_i, (\rho\bar{\rho})^{j+1})\Psi_i$$

of Q_i . Use the $\text{SL}(2, q)$ -oracle to obtain a generating set t_1, \dots, t_k for T , and set

$$\mathcal{S}_Q^* := \{t_s, u_j \mid 1 \leq s \leq k, 1 \leq j \leq 4k\}, \quad (24)$$

a generating set for $Q = O_p(N_G(T))$. Recall that our modification of Ψ_2 ensures that Q_1/T and Q_2/T are perpendicular nonsingular points of Q/T . As in Lemma 6.1, we need to consider Q/T in a more general setting, as follows.

Suppose that our input group G is a naturally embedded $\text{SU}(4, q)$ -subgroup of a black-box unitary group H of dimension $n \geq 4$. Let $Q_H = O_p(N_H(T))$ and let $(\cdot, \cdot)_{Q_H/T}$ denote a non-degenerate $N_H(T)'/Q_H$ -invariant hermitian form on Q_H/T . In particular, Q/T is a hyperbolic line of Q_H/T . (GS2)

LEMMA 6.10. (i) $B'_p := (u_i T)_{i=1}^{4k}$ is an \mathbb{F}_p -basis for Q/T .

(ii) In (GS2) there is a deterministic $O(\chi)$ -time algorithm that, for any given $y \in Q_H$, writes the projection of yT on Q/T along $(Q/T)^\perp$ as an \mathbb{F}_p -vector relative to B'_p .

(iii) There is a deterministic $O(\chi + \mu \log q)$ -time algorithm that, for any given $w \in Q$, writes an SLP of length $O(\log q)$ from \mathcal{S}_Q^* to w .

Proof. Statement (i) is clear, and Statement (iii) follows from Statement (ii) in the same way that Lemma 6.1(iv) followed from Lemma 6.1(iii). For Statement (ii), apply Lemma 6.1(iii) twice (using Ψ_i and J_i for $i = 1, 2$) to find the projections of yT on Q_1/T along $(Q_1/T)^\perp$ relative to $u_1T, \dots, u_{2k}T$, and on Q_2/T along $(Q_2/T)^\perp$ relative to $u_{2k+1}T, \dots, u_{4k}T$. Since Q_1/T and Q_2/T are perpendicular, the concatenation of those two vectors is the desired \mathbb{F}_p -vector. \square

6.2.5. Extending Ψ_1 and Ψ_2

Set $V := (\mathbb{F}_{q^2})^4$ and define a hermitian form on V by designating the usual basis of V to be $\mathcal{B} = e_1, v_1, v_2, e_{-1}$, where, for $i = 1, 2$, the basis $\mathcal{B}_i = e_1, v_i, e_{-1}$ is a standard basis for the nonsingular 3-space V_i that it spans. We have assumed that there exists $\Psi: \text{SU}(V) \rightarrow G$, and we may suppose further that $J_i\Psi^{-1} = \text{SU}(V_i)$ for $i = 1, 2$. Viewing Ψ_1 as a map on $\text{SU}(V_1)$, we may assume that Ψ extends Ψ_1 . Also, viewing Ψ_2 as a map on $\text{SU}(V_2)$, there exists $\alpha \in \text{Aut}(\text{SU}(3, q))$ such that the following property holds, where $\alpha \circ \Psi_2: \text{SU}(V_2) \rightarrow J_2$ sends $j' \mapsto j'^\alpha \Psi_2$:

there is an epimorphism $\Psi: \text{SU}(V) \rightarrow G$ extending Ψ_1 and $\alpha \circ \Psi_2$. (25)

Any such automorphism α factors as $\theta\gamma$, where θ is induced by an automorphism of \mathbb{F}_{q^2} and γ is induced under conjugation by an element of $\text{GU}(3, q)$. For α satisfying (25) the θ in any such factorisation is unique. The following procedure calculates the field automorphism inducing that θ .

PROCEDURE 6.11. Set $\sigma' := \text{diag}(1/\bar{\rho}, \bar{\rho}/\rho, \rho)$ and $\sigma := \sigma'\Psi_1 \in J_1$.

1. Fix $u \in Q_2$ and use Procedure 6.8 to find, relative to a suitable standard basis, $u\Psi_2^{-1} = r_1(v, v)$ and $u^\sigma\Psi_2^{-1} = r_1(v', v')$.

2. Express $v' = \lambda v$ and find the integer $0 \leq n < 2k$ such that $\lambda^{p^n} = 1/\rho$.

Set θ to be the automorphism $\lambda \mapsto \lambda^{-p^n}$.

Correctness. Note first that σ acts on Q/T , fixing the 1-spaces Q_1/T and Q_2/T . Furthermore, by Lemma 2.3(i), σ induces the scalar $1/\rho$ on Q_2/T . Hence, if Ψ is any epimorphism extending Ψ_1 and $1 \neq u \in Q_2$, then $u\Psi^{-1} = r_1(w, *)$ and $u^\sigma\Psi^{-1} = r_1(w', *)$, where $w' = (1/\rho)w$. Let Φ denote the restriction of Ψ to $\text{SU}(V_2)$. Then, if Φ_θ denotes the isomorphism $\theta: \text{SU}(V_2) \rightarrow J_2$ sending $A \mapsto A^\theta\Phi$ for some field automorphism θ , we have

$$u\Phi_\theta^{-1} = (u\Phi^{-1})^{\theta^{-1}} = r_1(w^{\theta^{-1}}, *) \quad \text{and} \quad u^\sigma\Phi_\theta^{-1} = r_1(w'^{\theta^{-1}}, *).$$

It follows that $w'^{\theta^{-1}} = (1/\rho)^{\theta^{-1}}w^{\theta^{-1}}$. Computed using Ψ_2 , the scalar induced by σ on Q_2/T is λ . We therefore require the unique field automorphism θ such that $\lambda^{\theta^{-1}} = 1/\rho$. That is precisely the automorphism that the procedure constructs.

Timing. This is dominated by $O(\chi)$, for the call to Procedure 6.8.

Change of basis. Apply the field automorphism θ to all of the generating matrices for $\text{SU}(V_2)$. We next find a matrix $C \in \text{GU}(3, q)$ such that, if $\gamma \in \text{Aut}(\text{SU}(3, q))$ represents the automorphism induced under conjugation by C , then $\gamma \circ \Psi_2$ satisfies property (25).

Find a hyperbolic pair f_1, f_{-1} in the support of $L\Psi_2^{-1}$, and a vector $w \in \langle f_1, f_{-1} \rangle^\perp$ such that $(w, w) = 1$ having the property that, if matrices in $\text{SU}(V_2)$ are written relative

to f_1, w, f_{-1} , then $x\Psi_1^{-1} = x\Psi_2^{-1}$ for all generators x of L . Clearly, this is a necessary condition for Ψ_2 to satisfy, in order for there to exist an extension of Ψ_1 and Ψ_2 . Also, since $C_{\text{GU}(V)}(\text{SU}(V_1))$ is transitive on the set of vectors $v \in \langle v_2 \rangle$ such that $(v, v) = 1$, we may choose any $w \in \langle f_1, f_{-1} \rangle^\perp$ subject to $(w, w) = 1$.

Let $C \in \text{GU}(3, q)$ map our given standard basis for $\text{SU}(V_2)$ to the standard basis f_1, w, f_{-1} , let γ denote the automorphism of $\text{SU}(V_2)$ induced under conjugation with C , and replace Ψ_2 with $\gamma \circ \Psi_2$.

6.2.6. A data structure for G

We have shown that, as maps on $\text{SU}(V_1)$ and $\text{SU}(V_2)$ relative to $\mathcal{B} = e_1, v_1, v_2, e_{-1}$, the isomorphism $\Psi_1: \text{SU}(3, q) \rightarrow J_1$ and the modified $\Psi_2: \text{SU}(3, q) \rightarrow J_2$ can be simultaneously extended to an epimorphism $\Psi: \text{SU}(V) \rightarrow G$. We note, once again, that since the images of the groups $Q(\langle e_1 \rangle)$ and $Q(\langle e_{-1} \rangle)$ are determined, such an extension is unique.

Recall that we intend to use Procedure 5.5 to write SLPs to given elements of G in the case $d = 4$. Accordingly, we now construct a suitable generating set \mathcal{S}^* for G , which can be used by the SLP algorithm for the general case. First, however, we need a four-dimensional version of Lemma 5.2.

LEMMA 6.12. *In deterministic $O(\chi)$ time, given any $g \in N_G(T)$, one can find the 2×2 matrix \tilde{g} representing the linear transformation induced by g on Q/T relative to the \mathbb{F}_{q^2} -basis $u_1T, u_{2k+1}T$.*

Proof. This follows easily from Lemma 6.10(ii) in the same way that Lemma 5.2 followed from Lemma 4.8(ii). □

Let σ' and $\sigma \in J_1$ be as in Procedure 6.11, let $l' \in \text{SU}(V_1)$ interchanging $\langle e_1 \rangle$ and $\langle e_{-1} \rangle$, and set $l := l'\Psi_1 \in J_1$. Recall the generating set \mathcal{S}_Q^* defined in (24), and use Ψ_1 and Ψ_2 to construct

$$\mathcal{T}_Q := \{t_s\Psi_1^{-1}, u_i\Psi_1^{-1} \mid 1 \leq s \leq k, 1 \leq i \leq 2k\} \cup \{u_i\Psi_2^{-1} \mid 2k+1 \leq i \leq 4k\},$$

the preimage of \mathcal{S}_Q^* under Ψ . Let \mathcal{S}_0^* be the generating set for $L_0 = G_G(L)'$ returned by the $\text{SL}(2, q)$ -oracle in Step 2 of Procedure 6.9. Observe that, for $s \in L_0$, $s\Psi^{-1} = \text{diag}(1, \tilde{s}, 1)$, where \tilde{s} is the 2×2 matrix induced by s on Q/T relative to $u_1T, u_{2k+1}T$. In $O(k\chi)$ time, use Lemma 6.12 to construct $\mathcal{T}_0 := \mathcal{S}_0^*\Psi^{-1}$. Finally, set

$$\begin{aligned} \mathcal{T} &:= \{\sigma', l'\} \cup \mathcal{T}_Q \cup \mathcal{T}_0, \quad \text{and} \\ \mathcal{S}^* &:= \{\sigma, l\} \cup \mathcal{S}_Q^* \cup \mathcal{S}_0^*, \end{aligned}$$

and define a bijection $\mathcal{T} \rightarrow \mathcal{S}^*$ in the obvious way.

Total timing and reliability for 6.2. A suitable data structure for G is found, with probability $1 - (1/8 + 0.05 + 0.1) > 0.725$, in time $O(\xi + \chi \log q)$.

6.2.7. A special case

The construction of a hermitian form on Q/T , presented in Procedure 4.4, requires that we be able to modify an effective isomorphism to obtain one with a certain specific property (see Proposition 6.14). The algorithmic setting in which that need arises is summarised as follows.

(a) We have naturally embedded $SU(4, q)$ -subgroups K and J of a black-box group G , where G is a homomorphic image of $SU(n, q)$ for some $n \geq 5$.

(b) There are transvection groups $T, T_1 \leq K$ such that $\langle T, T_1 \rangle = K \cap J \cong SL(2, q)$.

(c) We have generators for the group $Q = O_p(N_G(T))$ of order q^{2n-1} , as well as generators for $Q_J = Q \cap J$ and $Q_K = Q \cap K$, each of order q^3 .

(d) We have an effective isomorphism $\Psi_J : SU(4, q) \rightarrow J$.

We first demonstrate how any effective isomorphism $\Phi : SU(4, q) \rightarrow K$ defines a hermitian form on a certain elementary abelian section of K .

LEMMA 6.13. *Let $\Phi : SU(4, q) \rightarrow K$ be an effective isomorphism. Let e_1, e_2, e_{-1} and e_{-2} denote the standard basis relative to which the matrices of $SU(4, q)$ are written, and set $S := T(\langle e_1 \rangle)\Phi$ and $U := O_p(N_K(S))$. Suppose, for $u_1, u_2 \in U$, that $u_i\Phi^{-1} = r_1(w_i, \lambda_i)$ for $w_i \in \langle e_2, e_{-2} \rangle$ and $\lambda_i \in \mathbb{F}_{q^2}$. Then the assignment*

$$(u_1S, u_2S)_\Phi := (w_1, w_2), \tag{26}$$

defines an $N_K(S)/U$ -invariant non-degenerate hermitian form on U/T .

Proof. This follows easily from Lemma 2.2(i). □

In our setting we have two copies of $SU(4, q)$: one for J relative to a standard basis $(e_i)_{i \in I'}$, and one for K relative to a standard basis $(e'_i)_{i \in I'}$. Using a basis change if necessary, we may assume that Ψ_J maps $T(\langle e_1 \rangle) \rightarrow T$ and $T(\langle e_{-1} \rangle) \rightarrow T_1$, while Ψ_K maps $T(\langle e'_1 \rangle) \rightarrow T$ and $T(\langle e'_{-1} \rangle) \rightarrow T_1$. Hence we may use (26) to construct forms $(,)_{\Psi_J}$ and $(,)_\Phi$ on Q_J/T and Q_K/T respectively. Now there exists a unique $N_G(T)/Q$ -invariant hermitian form on Q/T extending $(,)_{\Psi_J}$, and another extending $(,)_\Phi$; there is no reason to suppose that they are the same form. However, we prove the following proposition.

PROPOSITION 6.14. *In the present setting, one can modify Φ so that $(,)_{\Psi_J}$ and $(,)_\Phi$ extend to the same $N_G(T)/Q$ -invariant hermitian form on Q/T .*

Proof. We may view the two copies of $SU(4, q)$ as naturally embedded subgroups of $SU(n, q)$ for $n \geq 5$. Specifically, let $(e_i)_{i \in I'}$ and $(e'_i)_{i \in I'}$ span nonsingular 4-spaces W and W' of an n -space V with $e_1 = e'_1$ and $e_{-1} = e'_{-1}$, so that $\Psi_J : SU(W) \rightarrow J$ and $\Phi : SU(W') \rightarrow K$. Clearly, there exists some automorphism $\alpha \in \text{Aut}(SU(4, q))$ such that the isomorphism $\Phi_\alpha : SU(W') \rightarrow K$ sending $A \mapsto A^\alpha\Phi \in K$ satisfies the following condition.

$$\text{There exists } \Psi : SU(V) \rightarrow G \text{ extending both } \Psi_J \text{ and } \Phi_\alpha. \tag{27}$$

As in 6.2.5, each $\alpha \in \text{Aut}(SU(4, q))$ can be written as $\alpha = \theta\gamma$, where θ is induced by an automorphism of \mathbb{F}_{q^2} and γ is induced under conjugation with an element of $GU(4, q)$. Once again, we first find the unique θ so that $\alpha = \theta\gamma$ satisfies (27).

Finding θ . Let $\sigma' \in SU(V_J)$ induce the scalar $1/\bar{\rho}$ on $\langle e_1 \rangle$ and ρ on $\langle e_{-1} \rangle$. Then $\sigma := \sigma'\Psi_J \in J$ has order $q^2 - 1$ and normalises T and T_1 . By Lemma 2.3(i), σ induces on Q/T a transformation $\tilde{\sigma}$ fixing a nonsingular 1-space U/T of Q_J/T and inducing the scalar $1/\rho$ on the $(d-3)$ -space $(U/T)^\perp$. Use Ψ_J to construct generators u_1, \dots, u_{2k} for U , and also generators w_1, \dots, w_{2k} for the subgroup $U_\perp \leq Q_J$, of order q^3 , such that U_\perp/T is perpendicular to U/T . Next, as in the proof of Lemma 6.10(ii), find the projection of rT on Q_J/T relative to $u_1T, \dots, u_{2k}T, w_1T, \dots, w_{2k}T$ for each generator r of Q_K . Using elementary linear algebra, find an \mathbb{F}_p -linear combination of those $4k$ vectors representing an element of U_\perp/T , and hence construct an element $u \in Q_K$ such that $uT \in U_\perp/T$.

Thus $\tilde{\sigma}$ induces the scalar $1/\rho$ on the \mathbb{F}_{q^2} -space spanned by uT , and we may use a similar method to that employed in Procedure 6.11 to compute θ .

Modifying Φ . Unlike the procedure in 6.2.5, we do not need to find Φ_α satisfying (27); we need only to modify Φ in such a way that $(,)_\Phi = (,)_{\Phi_\alpha}$ on Q_K/T for such a Φ_α . Apply θ to each element of our generating set for $SU(V_K)$ (hence modifying Φ). Then there exists $\gamma \in \text{Aut}(SU(4, q))$, induced under conjugation with some $C \in \text{GU}(4, q)$, such that Φ_γ satisfies (27). Clearly, any such γ satisfies the property

$$x\Phi_\gamma^{-1} = x\Psi_J^{-1} \text{ as matrices, for all } x \in \langle T, T^l \rangle. \quad (28)$$

If γ' is any automorphism induced under conjugation with an element of $\text{GU}(4, q)$ (say C') that satisfies (28), then $D := C'C^{-1}$ centralises $SU(\langle e_1, e_{-1} \rangle)$. Hence, if D_\bullet denotes the restriction of D to $\langle e'_2, e'_{-2} \rangle$, then D induces the transformation $w \mapsto \lambda w D_\bullet$ on $Q(\langle e_1 \rangle)/T(\langle e_1 \rangle)$ for some $\lambda \in \mathbb{F}_{q^2}$ with $\lambda^{q+1} = 1$. In particular, we have $(,)_{\Phi_{\gamma'}} = \lambda \bar{\lambda} (,)_{\Phi_\gamma} = (,)_{\Phi_\gamma}$ on Q_K/T . Hence we have reduced the problem to finding any automorphism γ satisfying (28). One can find a C giving rise to such γ via an elementary matrix calculation, and we replace Φ with Φ_γ . This completes the proof of the proposition. \square

Acknowledgements. This paper is based on part of the author's Ph.D. dissertation [7], completed under the direction of William Kantor. The author would like to thank Bill Kantor for the many stimulating discussions and helpful suggestions that have been invaluable in the preparation of this paper.

References

1. M. ASCHBACHER 'On the maximal subgroups of the finite classical groups', *Invent. Math.* 76 (1984) 469–514. 162
2. L. BABAI, 'Local expansion of vertex-transitive graphs and random generation in finite groups', *Proc. ACM Symp. on Theory of Computing* (1991) 164–174. 164
3. L. BABAI and R. BEALS, 'A polynomial-time theory of black box groups I', *Groups St Andrews 1997 in Bath, I*, London Math. Soc. Lecture Note Ser. 260 (ed. C. M. Campbell, E. F. Robertson, N. Ruskuc, and G. C. Smith, Cambridge University Press, 1999) 30–64. 162
4. R. BEALS and L. BABAI, 'Las Vegas algorithms for matrix groups', *Proc. IEEE Symp. Found. Comp. Sci.* (1993) 427–436. 162
5. R. BEALS, C. R. LEEDHAM-GREEN, A. C. NIEMEYER, C. E. PRAEGER and Á. SERESS, 'A black box algorithm for recognising finite symmetric and alternating groups, I', *Trans. Amer. Math. Soc.* 355 (2003) 2097–2113 (electronic). 163
6. P. A. BROOKSBANK, 'Constructive recognition of classical groups in their natural representation', *J. Symbolic Comput.* 35 (2003) 195–239. 168, 169, 174, 178, 180, 181, 190, 191
7. P. A. BROOKSBANK, 'Constructive recognition of the finite simple classical groups', Ph. D. thesis, Univeristy of Oregon, 2001. 196
8. P. A. BROOKSBANK and W. M. KANTOR, 'On constructive recognition of a black box $\text{PSL}(d, q)$ ', *Groups and Computation III*, Ohio State Univ. Math. Res. Inst. Publ. 8, (ed. W. M. Kantor and Á. Seress, Walter de Gruyter, Berlin/New York, 2001) 95–111. 163, 168, 171

9. M. CONDER and C. R. LEEDHAM-GREEN, ‘Fast recognition of classical groups over large fields’, *Groups and computation III*, Ohio State Univ. Math. Res. Inst. Publ. 8, (ed. W. M. Kantor and Á. Seress, Walter de Gruyter, Berlin/New York, 2001) 113–121. 163, 169
10. G. COOPERMAN, L. FINKELSTEIN and S. LINTON, ‘Constructive recognition of a black box group isomorphic to $GL(n, 2)$ ’, *Groups and computation II*, Proceedings of a DIMACS Workshop (ed. L. Finkelstein and W. M. Kantor, Amer. Math. Soc., Providence, RI, 1997) 85–100. 162
11. THE GAP GROUP, ‘Groups, algorithms, and programming’, Version 4.2 (The GAP Group, Aachen, St Andrews); <http://www-gap.dcs.st-and.ac.uk/gap>. 163
12. W. M. KANTOR and R. A. LIEBLER, ‘The rank 3 permutation representations of the finite classical groups’, *Trans. Amer. Math. Soc.* 271 (1982) 1–71. 168, 175
13. W. M. KANTOR and Á. SERESS, ‘Black box classical groups’, *Mem. Amer. Math. Soc.* 149 (2001). 162, 163, 164, 165, 166, 167, 168, 170, 171, 173, 174, 181, 182, 184, 191
14. W. M. KANTOR and Á. SERESS, ‘Computing with matrix groups’, *Groups, combinatorics and geometry (Durham 2001)* (ed. A. A. Ivanov, M. W. Liebeck and J. Saxl, World Scientific, 2003). 162
15. P. B. KLEIDMAN and M. W. LIEBECK, *The subgroup structure of the finite classical groups*, London Math. Soc. Lecture Note Ser. 129 (Cambridge Univ. Press, 1990). 165
16. C. R. LEEDHAM-GREEN, ‘The Computational Matrix Group Project’, *Groups and computation III*, Ohio State Univ. Math. Res. Inst. Publ. 8 (ed. W. M. Kantor and Á. Seress, Walter de Gruyter, Berlin/New York, 2001) 229–247. 162
17. P. M. NEUMANN and C. E. PRAEGER, ‘A recognition algorithm for special linear groups’, *Proc. London Math. Soc.* (3) 65 (1992) 555–603. 168
18. A. C. NIEMEYER and C. E. PRAEGER, ‘A recognition algorithm for classical groups over finite fields’, *Proc. London Math. Soc.* (1) 77 (1998) 117–169.
19. Á. SERESS, *Permutation group algorithms*, Cambridge Tracts in Math. 152 (Cambridge University Press, 2003). 169
20. K. ZSIGMONDY, ‘Zur Theorie der Potenzreste’, *Monatsh. Math. Phys.* 3 (1892) 265–284. 167

Peter A. Brooksbank brooksbank@math.ohio-state.edu
<http://www.math.ohio-state.edu/~brooksbank>

Department of Mathematics
The Ohio State University
231 W. 18th Ave.
Columbus OH 43210
USA