Fast constructive recognition
of black box orthogonal groups ☆

Peter A. Brooksbank [a,*], William M. Kantor [b]

[a] *Department of Mathematics, Bucknell University, Lewisburg, PA 17837, USA*
[b] *Department of Mathematics, University of Oregon, Eugene, OR 97403, USA*

Received 12 July 2005

Available online 22 March 2006

Communicated by Derek Holt

Dedicated to Charles Leedham-Green on the occasion of his 65th birthday

Abstract

We present an algorithm that constructively recognises when a given black box group is a nontrivial homomorphic image of the orthogonal group $\Omega^\varepsilon(d, q)$ for known ε, d and q. The algorithm runs in polynomial time assuming oracles for handling $\mathrm{SL}(2, q)$ subgroups and discrete logarithms in $\mathbb{F}_q^*$.
© 2006 Elsevier Inc. All rights reserved.

1. Introduction

For more than ten years a "Matrix Group Project" has been underway, whose goal is to produce algorithms to determine the structure of the group $G = \langle S \rangle$ generated by a set S of invertible matrices over a finite field (cf. [L-G,KS3]). Ultimately this reduces to problems involving the composition factors of G: determine them and provide algorithms for computing effectively within them. In order to handle quotient groups, this question has been generalised to the setting of "black box" groups, in which we only have the ability to calculate such things as products of elements of $G = \langle S \rangle$ within suitable time constraints.

The present paper is one of a series [CFL,KS1,BK,Br2] that deal with algorithmic questions concerning simple black box groups. While the emphasis of all of these papers has been on proving theoretical bounds on the timing and reliability of the algorithms they contain, work is also underway to provide practically efficient implementations that can be incorporated into the computer algebra systems MAGMA and GAP.

The algorithms presented in [KS1] solve the desired computational problems for all classical black box groups $G = \langle \mathcal{S} \rangle$, but they do not run in time polynomial in the input length. The groups $\mathrm{PSL}(d, q)$, $\mathrm{PSp}(d, q)$ and $\mathrm{PSU}(d, q)$ were dealt with in [BK,Br2], in polynomial time under additional computational assumptions. The present paper extends this result to the remaining classical groups: the orthogonal groups $\mathrm{P}\Omega^\varepsilon(d, q)$. Consequently, the class of "polynomial-time with oracle" constructive recognition algorithms now includes all classical groups (cf. [KS3]). Ideally we would have liked to present a uniform treatment of the classical groups, along the same lines as [KS1,Br1]. However, while the general architectures of the algorithms for the unitary and orthogonal groups are identical, the substantially different techniques used to handle various subproblems ultimately led to both [Br2] and the present paper. Moreover, awkward aspects of orthogonal groups make the solution of many subproblems more technical than their counterparts in the other classical groups.

In order to give statements of our results, we need some terminology and notation. A *black box group* is a group whose elements are encoded (not necessarily uniquely) using 0–1 strings of uniform length N, and which is equipped with an *oracle* (the "black box") that finds a string representing the product of two given elements, finds a string representing the inverse of a given element, and tests whether a given string represents the identity element of the group. If we are given G as $G = \langle \mathcal{S} \rangle$ for some subset $\mathcal{S}$ of G, then the *input length* is $N|\mathcal{S}|$.

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 say that a homomorphism $\Psi : H \to G$ is *effective* if there is a procedure (which may be deterministic or randomised) that computes $h\Psi \in G$ for any given $h \in H$, and also a procedure that computes some preimage for any given element of $H\Psi$.

For a prime power q, an $\mathrm{SL}(2, q)$-*oracle* is a deterministic algorithm which, for any input black box group G isomorphic either to $\mathrm{SL}(2, q)$ or $\mathrm{PSL}(2, q)$, produces an effective epimorphism $\Psi : \mathrm{SL}(2, q) \to G$. A $\mathrm{DLog}(\mathbb{F}_q^*)$-*oracle* is a deterministic algorithm that computes discrete logarithms in $\mathbb{F}_q^*$, given a generator of $\mathbb{F}_q^*$.

The following complexity parameters are used in our timing estimates:

μ: An upper bound on the time requirement for each group operation in G (i.e., the cost of using the black box).

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 [Ba] produces such elements in a black box group in time polynomial in the input length. We assume that $\xi \geqslant \mu|\mathcal{S}|$.

χ: An upper bound on the time requirement for each application of either of the hypothesised $\mathrm{SL}(2, q)$- or $\mathrm{DLog}(\mathbb{F}_q^*)$-oracles.

Our main result is as follows:

Theorem 1.1. *Suppose that a black box group $G = \langle S \rangle$ is a nontrivial homomorphic image of $\Omega^\varepsilon(d, q)$ for known ε, $d \geqslant 3$ and q; exclude the cases $\Omega^-(4, q)$ and $\Omega^-(6, q)$. Suppose further that $\mathrm{SL}(2, q)$- and $\mathrm{DLog}(\mathbb{F}_q^*)$-oracles are available. Then there is an*

$$O\big(d^3 \log d \log q \{d + \log^3 q\} + \mu d^3 \log^2 d + \xi d^2 \log d \log q + \chi d^2 \log d \log^2 q\big)$$

time Las Vegas algorithm which, with probability $> 1/2$, constructs an effective epimorphism $\Psi : \Omega^\varepsilon(d, q) \to G$. The routine finding the Ψ-image of any given element of $\Omega^\varepsilon(d, q)$ is deterministic and runs in $O(\mu d^2 \log q)$-time. The routine finding a preimage of any given element of G is Las Vegas, succeeding with probability $> 1/2$ in $O(\chi d^2 \log q)$-time.

See Section 2.2 for the definition of Las Vegas algorithms. In the theorem we assume that d is even if q is, as the odd-dimensional groups in characteristic 2 are symplectic groups and hence were already handled in [BK]. The cases $\Omega^-(4, q) \cong \mathrm{PSL}(2, q^2)$ and $\Omega^-(6, q)$ were excluded from the statement of the theorem because they are most easily handled as groups defined over $\mathbb{F}_{q^2}$: the 4-dimensional group $\Omega^-(4, q)$ using an $\mathrm{SL}(2, q^2)$-oracle, and the 6-dimensional group $\Omega^-(6, q)$ using the $\mathrm{SU}(4, q)$ algorithm presented in [Br2, 6.2]. (The latter hypothesises a discrete log oracle for a cyclic group of order $q + 1$ in addition to the oracles in the above theorem; see Section 5 for more about these groups.) Other well-known isomorphisms deal with the remaining low-dimensional cases: $\Omega(3, q) \cong \mathrm{PSL}(2, q)$ and $\Omega^+(4, q) \cong \mathrm{SL}(2, q) \circ \mathrm{SL}(2, q)$ are both handled using the $\mathrm{SL}(2, q)$-oracle; and $\Omega(5, q) \cong \mathrm{PSp}(4, q)$ is handled using the algorithm presented in [BK]. The group $\mathrm{P}\Omega^+(6, q)$ is recognised using the algorithm in [BK] for the isomorphic group $\mathrm{PSL}(4, q)$; see Section 2.2 below.

We prove somewhat more than is stated in the theorem, and thereby satisfy additional requirements of the Matrix Group Project [L-G]. For example, we construct a new generating S^* from the original generating set S having the property that, for any given $g \in G$, we can produce a preimage of g by writing a straight-line program from S^* to g and then evaluating the resulting straight-line program in $\Omega^\varepsilon(d, q)$ starting from $S^*\Psi^{-1}$. Keeping track of the steps used to construct S^* from S, we also *obtain straight-line programs from the set S to any given element g of G.*

The paper is organised as follows. In Section 2 we summarise the elementary properties of orthogonal groups needed for our algorithm, and also discuss the algorithmic background necessary for computation in black box groups. The main algorithm is presented in Sections 3 and 4. Constructive recognition algorithms naturally consist of two phases: a *preprocessing algorithm* that sets up a data structure defining a suitable epimorphism; and an *application algorithm* that uses this data structure to compute images and preimages of given elements. The preprocessing phase is handled in Section 3, and the application phase in Section 4. Section 5 indicates extensions and variations of Theorem 1.1. For example, in practice G is only *probably* a homomorphic image of the stated orthogonal group, but this is readily dealt with merely by assuming that G is, indeed, a homomorphic image: if the algorithm succeeds, then one checks that G satisfies a suitable presentation for the stated orthogonal group (cf. Section 5).

2. Preliminaries

In this section we summarise the required orthogonal group and algorithmic background (cf. [KS1,KL,Ta]).

2.1. Orthogonal groups

Let V be a vector space of dimension $d \geqslant 6$ over $\mathbb{F}_q$, where $q = p^k$, and let ϕ be a nondegenerate quadratic form on V. For $v, w \in V$, $(v, w) := \phi(v + w) - \phi(v) - \phi(w)$ is the associated symmetric bilinear form. A subspace $W \leqslant V$ is said to be: *totally singular* (t.s.) if $\phi(W) = 0$; and *nonsingular* if the restriction of ϕ to W is nondegenerate. The *Witt index* m of V is the dimension of a maximal t.s. subspace. A *hyperbolic line* is a nonsingular 2-space of Witt index 1.

For odd q, the quadratic form ϕ can be recovered from $(,)$ since $\phi(v) = (v, v)/2$. If d is odd, then (up to equivalence) ϕ is uniquely determined up to a scalar. We say that V has *type* $\varepsilon = 0, 1$ or -1 according as d is odd, d is even and V has maximal Witt index $m = d/2$, or d is even and V has minimal Witt index $m = (d - 2)/2$, respectively.

Lemma 2.1. *Let Λ_0 be a fixed hyperbolic line of the orthogonal space V of dimension $d \geqslant 5$. Then the following hold*:

(a) *With probability $> 1/4$, for a choice Λ of hyperbolic line $\langle \Lambda_0, \Lambda \rangle$ is a nonsingular 4-space of Witt index 2.*
(b) *With probability $> 1/4$, for a choice Σ of t.s. line $\langle \Lambda_0, \Sigma \rangle$ is a nondegenerate 4-space of Witt index 2.*

Proof. Let n_d^ε, h_d^ε and s_d^ε denote, respectively, the numbers of singular points, hyperbolic lines and t.s. lines in an orthogonal space of dimension d and type ε. Note that the number of nondegenerate 4-spaces of maximal Witt index that contain our fixed line Λ_0 is h_{d-2}^ε (the number of hyperbolic lines in the $(d - 2)$-space $\Lambda_0^\perp$). In (a), the number of suitable hyperbolic lines Λ in any such 4-space is $h_4^1 - (2q^2 - 1)$. Hence, noting that $h_d^\varepsilon = n_d^\varepsilon q^{d-2}/2$, we have

$$\mathrm{Prob}\big(\langle \Lambda_0, \Lambda \rangle \text{ is a 4-space of index 2}\big) = \big(h_{d-2}^\varepsilon / h_d^\varepsilon\big) \cdot \big(h_4^1 - 2q^2 + 1\big)$$
$$> \big(1/2q^4\big) \cdot \big(q^4/2\big) = 1/4.$$

In (b), the number suitable t.s. lines Σ in any such 4-space is $s_4^1 - 4 = 2(q - 1)$. Hence, noting that $s_d^\varepsilon = n_d^\varepsilon \cdot n_{d-2}^\varepsilon/(q + 1)$, we have

$$\mathrm{Prob}\big(\langle \Lambda_0, \Sigma \rangle \text{ is a 4-space of index 2}\big) = \big\{h_{d-2}^\varepsilon \cdot 2(q - 1)\big\}/s_d^\varepsilon$$
$$= q^{d-4}\big(q^2 - 1\big)/n_d^\varepsilon$$
$$> q^{d-4}(q - 1)\big(q^2 - 1\big)/q^{d-1} > 1/4. \quad \square$$

The group of all isometries of V is denoted $\mathrm{GO}^\varepsilon(V)$, where ε is the type of V. It is often convenient to represent elements of $\mathrm{GO}^\varepsilon(V)$ as matrices relative to a nice basis of V. Fix a generator, ρ, of the multiplicative group $\mathbb{F}_q^*$. Then a *standard basis* of V is an ordered basis of the form

$$\mathcal{B} = \begin{cases} e_1, \ldots, e_m, e_{-1}, \ldots, e_{-m} & \text{if } \varepsilon = 1, \\ e_1, \ldots, e_m, v_1, e_{-1}, \ldots, e_{-m} & \text{if } \varepsilon = 0, \text{ or} \\ e_1, \ldots, e_m, v_1, v_2, e_{-1}, \ldots, e_{-m} & \text{if } \varepsilon = -1, \end{cases} \tag{2.2}$$

where, for $i, j \in \{\pm 1, \ldots, \pm m\}$ and $s = 1, 2$,

$$\begin{aligned} &\phi(e_i) = 0 = (e_i, v_s) \quad \text{and} \quad (e_i, e_j) = \delta_{i,-j}, \\ &(v_1, v_2) = 1 \quad \text{and} \quad \phi(v_s) \neq 0 && \text{if } p = 2, \\ &(v_1, v_2) = 0 \quad \text{and} \quad (v_s, v_s) = 1 && \text{if } q \equiv 3 \pmod 4, \\ &(v_1, v_2) = 0, \qquad (v_1, v_1) = 1 \quad \text{and} \quad (v_2, v_2) = \rho && \text{if } q \equiv 1 \pmod 4. \end{aligned}$$

2.1.1. Point stabilisers and root groups

Let $\mathcal{B}$ be a standard basis for V and, for any $i \in \{\pm 1, \ldots, \pm m\}$, let $x_i = \langle e_i \rangle$. Then the point stabiliser $\Omega^\varepsilon(V)_{x_i}$ is a semidirect product

$$\Omega^\varepsilon(V)_{x_i} = Q(x_i) \rtimes \Omega^\varepsilon(V)_{x_i, x_{-i}}. \tag{2.3}$$

Here $Q(x_i) = O_p(\Omega^\varepsilon(V)_{x_i})$, of order q^{d-2}, is the natural module of the group $(\Omega^\varepsilon(V)_{x_i, x_{-i}})' \cong \Omega^\varepsilon(d-2, q)$, and consists of all isometries of the form

$$r_i(w): u \mapsto u + \big(u, w - \phi(w)e_i\big)e_i - (u, e_i)w, \quad w \in \langle e_i, e_{-i} \rangle^\perp. \tag{2.4}$$

Moreover,

$$\phi_{x_i}\big(r_i(w)\big) := \phi(w) \tag{2.5}$$

defines a nondegenerate quadratic form on $Q(x_i)$; the map $w \mapsto r_i(w)$ is an isometry $\langle e_i, e_{-i} \rangle^\perp \to Q(x_i)$, and

$$r_i(w)^g = r_i\big(w^g\big) \quad \text{whenever } g \in \big(\Omega^\varepsilon(V)_{x_i, x_{-i}}\big)'. \tag{2.6}$$

In particular,

$$r_i(w)^g = r_i\big(\rho^2 w\big) \quad \text{if } g = \mathrm{diag}\big(\rho^2, 1, \ldots, 1, \rho^{-2}, 1, \ldots, 1\big). \tag{2.7}$$

The $\Omega^\varepsilon(V)$-conjugates of $r_i(w)$ will be referred to generically in this paper as *root elements*; conjugates of the groups

$$R\big(\langle e_i, w \rangle\big) := \big\{ r_i(\lambda w) \mid \lambda \in \mathbb{F}_q \big\}, \tag{2.8}$$

of order q, will be called *root groups*. Note that this is not quite in accordance with the Lie-theoretic definition of root elements and root groups. For singular $w \neq 0$, however, conjugates of $r_i(w)$ are *long root elements*, and conjugates of $R(\langle e_i, w \rangle)$ are *long root groups*, as in the Lie-theoretic sense. In view of the isometry $\langle e_i, e_{-i} \rangle^{\perp} \rightarrow Q(x_i)$, long root groups correspond to singular points of V.

Lemma 2.9. *Let V be an orthogonal space with $d = \dim(V) \geqslant 6$. For any two distinct singular points $x, y \in V$, put $H := \langle Q(x), Q(y) \rangle \leqslant \Omega^{\varepsilon}(V)$.*

(a) *If x and y are not perpendicular, then $H = \Omega^{\varepsilon}(V)$.*
(b) *If x and y are perpendicular (so that $\Sigma = \langle x, y \rangle$ is totally singular), then the following hold:*
    1. *$H = U \rtimes \mathrm{SL}(\Sigma)$, where $U = O_p(H)$ is the subgroup of $\Omega^{\varepsilon}(V)$ that centralises each of the spaces Σ, $\Sigma^{\perp}/\Sigma$ and $V/\Sigma^{\perp}$; also $|U| = q^{2d-7}$.*
    2. *$Z(U) = [U, U]$ is the root group $R(\Sigma) = Q(x) \cap Q(y)$.*
    3. *With probability $(1 - 1/q^{2d-6})(1 - 1/q) > 0.49$, if $u_1, u_2 \in U$ then $1 \neq [u_1, u_2] \in R(\Sigma)$.*

Proof. (a) This is well known.

(b) Parts 1 and 2 follow from a matrix calculation; they may also be deduced using an argument similar to [KS1, Lemma 4.7(ii)]. For part 3, note that a choice u_1 is not in $Z(U)$ with probability at least $1 - 1/q^{2d-6}$. For such a u_1, it is straightforward to check that $C_U(u_1)$ has order q^{2d-8}; thus, with probability at least $1 - 1/q$, u_2 does not commute with u_1. □

2.1.2. Homomorphisms and isometries

For $n < d$, a subgroup of $\Omega^{\varepsilon}(d, q) = \Omega^{\varepsilon}(V)$ that induces $\Omega^{\varepsilon'}(n, q)$ on some nondegenerate n-space of V, and the identity on the orthogonal complement of this n-space, will be called a *naturally embedded $\Omega^{\varepsilon'}(n, q)$-subgroup*. For a homomorphic image G of $\Omega^{\varepsilon}(V)$, we call a subgroup $H \cong \Omega^{\varepsilon'}(n, q)$ of G *naturally embedded in G* if $H\Psi^{-1}$ is naturally embedded in $\Omega^{\varepsilon}(V)$ for some (not necessarily specified) epimorphism $\Psi : \Omega^{\varepsilon}(V) \rightarrow G$.

Suppose that $d > 6$, and let J be a naturally embedded $\Omega^{+}(6, q)$-subgroup of G. Furthermore, let $\Psi_J : \Omega^{+}(V_J) \rightarrow J$ be a fixed isomorphism, where V_J is a nondegenerate 6-space in V of Witt index 3. Let $x_+ \in V_J$ be any singular point, and put $Q_J := O_p(\Omega^{\varepsilon}(V_J)_{x_+})\Psi_J < J$.

Consider *any* epimorphism $\Phi : \Omega^{\varepsilon}(V) \rightarrow G$ extending Ψ_J. Let Q denote the group $Q(x_+)\Phi < G$ containing Q_J, where $Q(x_+) = O_p(\Omega^{\varepsilon}(V)_{x_+})$, and let L denote the group $(\Omega^{\varepsilon}(V)_{x_+, x_-})'\Phi < G$, where $x_- \in V_J$ is not perpendicular to x_+. (The groups Q and L do not depend on the choice of Φ.) By (2.5),

$$\phi_Q(u) := \phi_{x_+}\left(u\Phi^{-1}\right) \quad \text{for } u \in Q \tag{2.10}$$

defines a nondegenerate L-invariant quadratic form on Q, while

$$\phi_{\Psi_J}(u) := \phi_{x_+}\left(u\Psi_J^{-1}\right) \quad \text{for } u \in Q_J \tag{2.11}$$

defines a nondegenerate $L_J := L \cap J$-invariant quadratic form on $Q_J = Q \cap J$. Up to equivalence, any nondegenerate L-invariant quadratic form on Q is a scalar multiple of ϕ_Q, so *there is a unique such form extending* ϕ_{Ψ_J}. Thus, ϕ_Q *does not depend on the choice of* Φ *extending* Ψ_J.

Consequently, it makes sense to speak of an *isometry from the orthogonal space* $Q(x_+)$ (relative to ϕ_{x_+}) *to the orthogonal space* Q (relative to ϕ_Q) *independent of the particular choice of* Φ *used to define* ϕ_Q.

Proposition 2.12. *In the above notation, let* $f : Q(x_+) \to Q$ *be an isometry that coincides with* Ψ_J *on* $Q(x_+) \cap \Omega^+(V_J)$. *Then there is a unique epimorphism* $\Psi : \Omega^\varepsilon(V) \to G$ *extending both* f *and* Ψ_J.

Proof. *Uniqueness*: Let l' be any element of $\Omega^+(V_J)$ that interchanges $Q(x_+)$ and $Q(x_-)$. Since $\Omega^\varepsilon(V) = \langle Q(x_+) = Q(x_-) \rangle$, any epimorphism $\Omega^\varepsilon(V) \to G$ is determined by the image of the set $Q(x_+) \cup \{l'\}$. The uniqueness of Ψ now follows from the fact that f determines the image of $Q(x_+)$, while Ψ_J determines the image of l'.

Existence: First let Φ denote *any* epimorphism extending Ψ_J. Then Φ restricts on $Q(x_+)$ to an isometry $f_\Phi : Q(x_+) \to Q$ agreeing with Ψ_J (and hence also f) on $Q(x_+) \cap \Omega^+(V_J)$. If $w_1, \ldots, w_{d-6}$ is a standard basis of $V_J^\perp$, then the ff_Φ^{-1}-image of the set $\{r_1(w_i) \mid 1 \leqslant i \leqslant d-6\} \subset Q(x_+)$ is a set $\{r_1(w_i') \mid 1 \leqslant i \leqslant d-6\} \subset Q(x_+)$, where $w_1', \ldots, w_{d-6}'$ is another standard basis of $V_J^\perp$. Hence there is a matrix $C \in \mathrm{GO}^\varepsilon(d, q)$ inducing the identity on V_J and sending $w_i' \mapsto w_i$ for $1 \leqslant i \leqslant d-6$.

We claim that the epimorphism $\Psi : \Omega^\varepsilon(V) \to G$ sending $g' \mapsto (Cg'C^{-1})\Phi$ extends both Ψ_J and f, as desired. For $g' \in \Omega^+(V_J)$, since C is the identity on V_J, we have $Cg'C^{-1}\Phi = g'\Psi_J$ (since Φ extends Ψ_J). Next,

$$r_1(w_i) f \Psi^{-1} = C^{-1}\big(r_1(w_i) f \Phi^{-1}\big) C = r_1\big(w_i'\big)^C = r_1\big(w_i' C\big) = r_1(w_i).$$

Hence the restriction of Ψ to $Q(x_+)$ agrees with f. It follows that Ψ extends both Ψ_J and f, as claimed. $\quad\square$

2.2. Algorithmic preliminaries

We now summarise the relevant notions and procedures needed for computing with black box groups. The reader is referred to chapters 1 and 2 of [Se] for a detailed treatment of this topic.

2.2.1. Monte Carlo and Las Vegas algorithms

We will consider randomised algorithms. Such an algorithm is called *Monte Carlo* if the output may be incorrect, but an upper bound on the probability of that can be prescribed by the user; thus, there is always an output, but there is also an uncomfortable possibility of error. A randomised algorithm is called *Las Vegas* if the output is always correct, but there is a possibility that "failure" is output, and an upper bound on the probability of that can be prescribed by the user; no errors can occur, so this is more comforting.

2.2.2. Straight-line programs

Let X be a list of elements of a group G, and let $g \in G$. Informally, a *straight-line program* (SLP) *of length m from X to g* is a sequence $(g_1, \ldots, g_m)$ of group elements 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. Since we do not always want to compute or store each of the group elements g_i, we more formally define an SLP from X to g to be a sequence $(w_1, \ldots, w_m)$ such that, for each i, either w_i is a positive integer (representing the w_ith 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 enables us to construct SLPs inside one group and evaluate them in another.

2.2.3. Order properties

By a fundamental theorem of Zsigmondy [Zs], if p is a prime and $n \geqslant 2$, then there is a prime r dividing $p^n - 1$ but not $p^i - 1$ for $1 \leqslant i < n$, except when either $p = 2$ and $n = 6$, or $n = 2$ and p is a Mersenne prime. Such a prime r 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 ppd$^{\#}(p; n)$ if $n = 6$, $p = 2$ and $21 \mid j$; if $n = 2$, p is Mersenne and $4 \mid j$; or if j is divisible by a ppd of $p^n - 1$. If p is not a Fermat prime, we say that j is a ppd$^{\#}(p; 1)$ if j is not a power of 2; if p is a Fermat prime, we say that j is a ppd$^{\#}(p; 1)$ if $4 \mid j$. We call an element g of a group a ppd$^{\#}(p; n)$-element if $|g|$ is a ppd$^{\#}(p; n)$.

Computing the *exact* order of a given element g of a black box group G will not be necessary but, as was the case in [Br1,Br2,KS1], we will need to test whether g is a ppd$^{\#}(p; n)$-element for a given prime p and integer n:

Lemma 2.13. ([NP]; cf. [Br1, Lemma 3.1]) *Following a preprocessing computation requiring time $O(n^3 \log n \log^4 p)$, one can test whether or not a given element of a black-box group G has ppd$^{\#}(p; n)$-order in time $O(\mu n \log p)$.*

2.2.4. The $SL(2, q)$-oracle

Our algorithm assumes the availability of an oracle which, for any given black box group G isomorphic to $SL(2, q)$ for known q, returns an effective isomorphism $\Psi : SL(2, q) \to G$. In particular, the oracle constructs a new generating set $\mathcal{S}^*$ for G, and provides a procedure which, for any given $g \in G$, writes an SLP from $\mathcal{S}^*$ to g. We assume that $\chi \geqslant \mu \log q$ (the time required to evaluate an SLP of length $\log q$ within G).

Remark 2.14. All SLP procedures in fact take as input a string representing some element the black box group G (recall that the elements of G are not necessarily uniquely encoded). *We assume that the SLPs output by the $SL(2, q)$-oracle for strings representing the same element of G are identical.*

Remark 2.15. There will be situations in which the $SL(2, q)$-oracle is called for a group G, *which is only probably isomorphic to $SL(2, q)$. If it happens that $G \not\cong SL(2, q)$, then we* assume the oracle returns false.

2.2.5. Applying the oracle to quotient groups

We will need to handle the following situation. Let G be the given black box orthogonal group, and suppose we have (generators for) a black box subgroup H of G such that, for some $N \trianglelefteq H$, we have $H/N \cong \mathrm{SL}(2, q)$. We do not have generators for N, but there is an efficient deterministic test for membership in this subgroup. This is sufficient to enable us to compute with H/N as a black box group: the elements of H/N are the elements of H; group operations are performed exactly as they are in H; but the test for whether or not a given string represents the identity is performed instead using the membership test for N.

Thus, in this setting, we may apply the $\mathrm{SL}(2, q)$-oracle to the quotient group H/N. (For more about practical aspects of such applications in our specific setting, see Remark 4.7.) In particular, we obtain a set $\mathcal{S}_H^* \subset H$ such that $\langle \mathcal{S}_H^* \rangle N/N = H/N$, having the property that, for any given $h \in H$, an SLP can be obtained from $\mathcal{S}_H^*$ to an element $h_0 \in H$ such that $hN = h_0 N$.

In the situation described above, we shall wish to use the $\mathrm{SL}(2, q)$-oracle to construct (nearly) uniformly distributed random elements of the normal subgroup N. This is achieved as follows. Using the given generators for H, construct (nearly) uniformly distributed random elements $h \in H$. For each such element, the $\mathrm{SL}(2, q)$-oracle is then called within the quotient group H/N to write an SLP from $\mathcal{S}_H^*$ to $h_0 \in H$ with $hN = h_0 N$. The elements $h^{-1} h_0$ are now (nearly) uniformly distributed random elements of N. (For, by Remark 2.14, SLPs returned by the $\mathrm{SL}(2, q)$-oracle to elements of the same coset are identical.)

2.2.6. Algorithms for $\Omega^+(6, q)$

Our general algorithm makes essential use of subgroups isomorphic to $\Omega^+(6, q)$, and requires that we be able to recognise them constructively *as* 6-*dimensional groups*. The following is a version of Theorem 1.1 for the single case $d = 6$, $\varepsilon = 1$.

Lemma 2.16. *There is an* $O(\log q \{\xi + \chi \log q + \mu \log^2 q\})$-*time Las Vegas algorithm which, with probability* $> 3/4$, *when given a black box group* $G = \langle \mathcal{S} \rangle$, *known to be a nontrivial homomorphic image of* $\Omega^+(6, q)$, *and having available* $\mathrm{SL}(2, q)$- *and* $\mathrm{DLog}(\mathbb{F}_q^*)$-*oracles, constructs an effective epimorphism* $\Psi : \Omega^+(6, q) \to G$. *Moreover,*

(a) *there is an* $O(\mu \log q)$-*time deterministic algorithm that finds the image of any given element of* $\Omega^+(6, q)$; *and*
(b) *there is an* $O(\xi + \chi \log q + \mu \log^2 q)$-*time Las Vegas algorithm that finds a preimage of any given element of* G, *succeeding with probability* $\geqslant 1 - 1/128$.

Proof. Since $\mathrm{P}\Omega^+(6, q) \cong \mathrm{PSL}(4, q)$, we use the algorithm in [BK] to obtain an effective epimorphism $\Phi : G \to \mathrm{PSL}(4, q) = \mathrm{PSL}(W)$. Straightforward linear algebra provides an isomorphism $\theta : \mathrm{PSL}(W) \to \mathrm{P}\Omega^+(W \wedge W) = \Omega^+(6, q)$, where $W = \mathbb{F}_q^4$. Although computing the θ-image of any given element of $\mathrm{PSL}(W)$ is elementary, computing preimages requires a few comments.

Let $\perp$ refer to perpendicularity with respect to the usual dot product relative to the usual basis w_1, w_2, w_3, w_4 of W. For $1 \leqslant i \neq j \leqslant 4$, the group of $(\langle w_i \rangle, w_j^\perp)$-transvections of $\mathrm{PSL}(W)$ maps, under θ, to the long root group $R(w_i \wedge w_j^\perp)$ of $\mathrm{P}\Omega^+(6, q)$. Let $\mathcal{T}$ be the

union of $O(k)$-size generating sets for the 12 long root groups $R(w_i \wedge w_j^{\perp})$, obtained as the θ-images of generating sets of the 12 groups of $(\langle w_i \rangle, w_j^{\perp})$-transvections of W. To find $h\theta^{-1}$ for a given $h \in \mathrm{P}\Omega^+(6, q)$, we write an SLP from $\mathcal{T}$ to h using the deterministic algorithm in [Br1, Theorem 1.1], and then evaluate it from $\mathcal{T}\theta^{-1}$.

The map $\Phi\theta : G \to \mathrm{P}\Omega^+(6, q)$ is thus an effective epimorphism. The centre of G may now be constructed as follows using a standard short presentation for $\mathrm{P}\Omega^+(6, q)$ arising from elementary matrices [BGKLP]. Evaluate all of the relators from the appropriate elements of G. If the identity of G is obtained for all of them, then G is isomorphic to $\mathrm{P}\Omega^+(6, q)$; else at least one of the relations evaluates to a generator of $Z(G)$, namely to the element corresponding to $-1 \in \Omega^+(6, q)$. Now modify $\Phi\theta$ to produce an effective epimorphism $\Psi : \Omega^+(6, q) \to G$, in accordance with Theorem 1.1.

The stated complexity arises from the call to [BK, Theorem 1.1] (with $d = 4$ and $q > 17$) to construct an isomorphism $\Phi : G \to \mathrm{PSL}(4, q)$, which dominates the timing. That algorithm also accounts for the reliability stated in the lemma, and provides the routine for computing preimages. The routine for computing images is in [Br1, Theorem 1.1] (also see Section 4.1). □

3. Constructing an isomorphism

This section contains the heart of our algorithm: the *preprocessing phase* that sets up a data structure for an effective epimorphism from a concrete matrix group to a given black box group. Let $G = \langle \mathcal{S} \rangle$ be the given black box group, known to be a nontrivial homomorphic image of $\Omega^\varepsilon(d, q)$ for known $d, q = p^k$ and $\varepsilon \in \{-1, 0, 1\}$. In view of Lemma 2.16, and the discussion following Theorem 1.1, we may assume that $d \geqslant 7$. Our data structure will consist of the following components:

(a) A generating set $\mathcal{T}$ for the group of matrices $\Omega^\varepsilon(d, q)$.
(b) A new generating set $\mathcal{S}^*$ for G, whose elements are constructed using SLPs from the original generating set $\mathcal{S}$.
(c) A bijection $\mathcal{T} \to \mathcal{S}^*$ that extends to an epimorphism $\Omega^\varepsilon(d, q) \to G$. (This epimorphism will be our *effective* epimorphism; its effectiveness will be established following the *application phase* in Section 4.)
(d) A naturally embedded $\Omega^+(6, q)$-subgroup J of G.
(e) Several useful elements of J together with their preimages in $\Omega^\varepsilon(d, q)$.

3.1. Preliminary constructions for "large" fields

Throughout this subsection we will assume that $q \geqslant 16$. Analogous constructions for smaller field sizes will be given in Section 3.2.

Our first goal is to obtain generators for a naturally embedded $\Omega^+(6, q)$-subgroup J of G, together with an effective isomorphism $\Psi_J : \Omega^+(6, q) \to J$ (cf. (d) above). The isomorphism Ψ_J will be important for two reasons. First, it will allow us to construct useful elements of J to be used later in our algorithm (cf. (e)). Second, our target epimorphism $\Omega^\varepsilon(d, q) \to G$ will be constructed so as to extend Ψ_J.

Our other preliminary task is to construct a subgroup Q of G, of order q^{d-2}, corresponding to some subgroup $Q(x_i)$ of $\Omega^\varepsilon(d, q)$, as well as a certain complement to Q in $N_G(Q)'$ (cf. Section 2.1.1).

3.1.1. The elements σ and a

We begin by giving a Las Vegas procedure that takes as input the given black box group G and returns elements $a, \sigma \in G$ that act on the underlying module in a special manner. (Note that all $\mathrm{ppd}^\#(p;\ )$-tests in the following procedure are performed using Lemma 2.13.)

Procedure 3.1. Put $n := d - 3 - \varepsilon$. For up $96n$ choices $\tau \in G$, proceed as follows.

1. Test whether $\sigma := \tau^{q^2-1}$ has $\mathrm{ppd}^\#(p; nk)$-order. If so, set $b := \tau^{q^{n/2}+1}$.
2. If $\varepsilon = 1$, test whether b has $\mathrm{ppd}^\#(p; 2k)$-order; if so, set $a := b^{q+1}$. If $\varepsilon \neq 1$, set $a := b$.
3. Test whether a has $\mathrm{ppd}^\#(p; k)$-order.

If, for some choice τ, all of the tests performed in steps 1–3 are in the affirmative, then return the pair (σ, a). Report `failure` if all $96n$ choices fail at least one of the tests.

Lemma 3.2. *Procedure* 3.1 *is a Las Vegas algorithm that returns a* $\mathrm{ppd}^\#(p; nk)$ *element* σ, *and a* $\mathrm{ppd}^\#(p; k)$-*element* a *whose support in the natural module underlying G is a hyperbolic line upon which σ induces the identity. The procedure succeeds with probability* $> 1 - 1/e^3 > 7/8$ *and runs in* $O(d^3 \log d \log^4 q + d\{\xi + \mu d \log q\})$-*time.*

Proof. The routine is extracted from [Br1, Section 4.3.1], where such a procedure is given more generally for all classical groups: the correctness, timing and reliability estimates are taken directly from there. □

3.1.2. Constructing J

The following is a Las Vegas subroutine that takes as input the element a just constructed, together with the group G, and returns a (constructively recognised) naturally embedded $\Omega^+(6, q)$-subgroup.

Procedure 3.3. For up to 2^{16} choices $(g_1, g_2) \in G \times G$, proceed as follows:

1. Set $J := \langle a, a^{g_1}, a^{g_2} \rangle$.
2. Use Lemma 2.16 to test, constructively, whether $J \cong \Omega^+(6, q)$.

If we obtain a suitable J for some g_1, g_2, then return the effective isomorphism $\Psi_J : \Omega^+(6, q) \to J$ constructed in step 2. Report `failure` if none of the choices give rise to a suitable J.

Lemma 3.4. *Procedure* 3.3 *is a Las Vegas algorithm that returns a naturally embedded* $\Omega^+(6, q)$-*subgroup J, and effective isomorphism* $\Psi_J : \Omega^+(6, q) \to J$, *with probability* $> 1 - 1/e^3 > 7/8$, *in* $O(\log q\{\xi + \chi \log q + \mu \log^2 q\})$-*time.*

Proof. For a fixed pair g_1, g_2, $\langle a, a^{g_1}, a^{g_2} \rangle$ is a naturally embedded $\Omega^+(6, q)$-subgroup of G with probability $> 1/2^{14}$ [KS1, Lemma 4.12(ii)]. For a pair behaving in this manner, Lemma 2.16 produces a suitable isomorphism with probability $> 3/4$ (using a Las Vegas algorithm). Hence, all of our 2^{16} choices fail with probability $< (1 - 3/2^{16})^{2^{16}} < 1/e^3$, as claimed. The stated timing arises from the $O(1)$ calls to Lemma 2.16. $\quad\square$

Remark 3.5. Note that the requirement $q \geqslant 16$ is used for the first time in the proof of the lemma: [KS1, Lemma 4.12(ii)] applies only for such q.

3.1.3. Some elements of J

From now on, $\mathbb{F}_q$ will denote the field constructed along with the effective isomorphism $\Psi_J : \Omega^+(6, q) = \Omega^+(V_J) \to J$; ρ will denote a fixed generator of $\mathbb{F}_q^*$. Define a nondegenerate quadratic form ϕ, and associated symmetric form $(\ ,\ )$, on the row space $V := \mathbb{F}_q^d$ such that the usual basis of V is a standard basis $\mathcal{B}$ (see (2.2)). Matrices in $\Omega^+(6, q)$ will always be written relative to the standard basis $\{e_1, e_2, e_3, e_{-1}, e_{-2}, e_{-3}\}$ of V_J. We may also assume that $x_+ := \langle e_1 \rangle$ and $x_- := \langle e_{-1} \rangle$ are the 1-dimensional eigenspaces of $a\Psi_J^{-1}$. The subgroups $Q(x_\pm)$ of $\Omega^\varepsilon(d, q)$ consist of the transformations $r_{\pm 1}(w)$ given in (2.4). We identify these transformations with their corresponding matrices relative to $\mathcal{B}$.

Use Lemma 2.16(a) to construct all of the following elements of J in $O(\log q(\mu \log q))$-time:

(i) For $1 \leqslant j \leqslant k$, $r_{1j} := r_1(\rho^{j-1} e_2)\Psi_J$, $r_{2j} := r_1(\rho^{j-1} e_3)\Psi_J$, $r_{3j} := r_1(\rho^{j-1} e_{-2})\Psi_J$ and $r_{4j} := r_1(\rho^{j-1} e_{-3})\Psi_J$.

(ii) $l := l'\Psi_J$, where $l' \in \Omega^+(V_J)$ sends $e_{\pm i} \mapsto e_{\mp i}$ $(i = 1, 2)$ and $e_{\pm 3} \mapsto e_{\pm 3}$.

(iii) $h := h'\Psi_J$, where $h' \in \Omega^+(V_J)$ sends $e_1 \mapsto \rho^2 e_1$, $e_{-1} \mapsto \rho^{-2} e_{-1}$ and $v \mapsto v$ for $v \in \langle e_1, e_{-1} \rangle^\perp$. (N.B. h is an analogue of the element s used in [KS1, 4.4.2(v)]; by (2.7), h has order $(q-1)/(2, q-1)$, acting on Q as the scalar ρ^2.)

(iv) $h^+ := (h^+)'\Psi_J$, where $(h^+)' \in \Omega^+(V_J)$ induces an element of order $q - 1$ on x_+. (N.B. h^+ is an analogue of the element s^+ used in [KS1, 4.4.2(vi)].)

(v) A four-element generating set for $J_0 := \Omega^+(\langle e_{\pm 1}, e_{\pm 2} \rangle)\Psi_J$.

(vi) A four-element generating set for $D_J = \Omega^+(\langle e_{\pm 2}, e_{\pm 3} \rangle)\Psi_J$.

Here r_{ij} are long root elements of G, which for a fixed i generate a long root group (cf. (3.9)). Let

$$Q_J := \langle r_{ij} \mid 1 \leqslant i \leqslant 4, \ 1 \leqslant j \leqslant k \rangle = \left[Q(x_+) \cap \Omega^+(V_J) \right]\Psi_J. \tag{3.6}$$

Since $r_1(e_{\pm 2})^{l'} = r_{-1}(e_{\mp 2})$ and $r_1(e_{\pm 3})^{l'} = r_{-1}(e_{\pm 3})$ by matrix calculations,

$$Q_J^l = \left[Q(x_-) \cap \Omega^+(V_J) \right]\Psi_J. \tag{3.7}$$

Finally, since $\Omega^+(\langle e_{\pm 2}, e_{\pm 3} \rangle)$ is the identity on $\langle e_1, e_{-1} \rangle$,

$$N_J(Q_J)' = Q_J \rtimes D_J, \quad \text{and} \quad D_J \text{ normalises } Q_J^l. \tag{3.8}$$

3.1.4. Constructing Q and $N_G(Q)' = Q \rtimes L$

We now proceed to the second stage of our preliminary constructions. Set

$$
\begin{aligned}
r_{ij} &:= r_{1j}^{\sigma^{i-4}} \quad (5 \leqslant i \leqslant d-2 \text{ and } 1 \leqslant j \leqslant k), \\
R_i &:= \langle r_{ij} \mid 1 \leqslant j \leqslant k \rangle \quad (1 \leqslant i \leqslant d-2), \\
Q &:= \langle R_1, \ldots, R_{d-2} \rangle, \\
L &:= \langle D_J^{\sigma^i} \mid 0 \leqslant i \leqslant d-2 \rangle.
\end{aligned}
\tag{3.9}
$$

All of the elements r_{ij}, and the generators for L, are constructed in $O(\mu d \log q)$-time. Let Φ be any epimorphism $\Omega^\varepsilon(V) \to G$ that extends Ψ_J, let $\dot Q = Q(x_+)\Phi$, and let $\dot L = \Omega^\varepsilon(V)_{e_1,e_{-1}}\Phi$ (as in Section 2.1.2, $\dot Q$ and $\dot L$ do not depend on the choice of Φ). Then we have the following result.

Lemma 3.10. *With probability at least 7/8, $Q = \dot Q$ and $L = \dot L$.*

Proof. Let V' denote the orthogonal module underlying G, and recall the elements a and σ found by Procedure 3.1. Then $[V', a]$ is a hyperbolic line with singular points x'_+ and x'_- such that $\dot Q = Q(x'_+)$ and $\dot Q^l = Q(x'_-)$. Furthermore, σ is the identity on $[V', a]$, and hence induces on the $(d-2)$-space $\dot Q$ a transformation of $\mathrm{ppd}^\#(p; nk)$-order. Therefore, exactly as in [Br1, Section 4.4.2], $Q = \dot Q$ with the stated probability.

We assume now that $Q = \dot Q$, and turn to the proof of [KS1, Lemma 4.14] from the third paragraph onward. In that proof it is shown that a certain group G'_α (generators for which were obtained in a manner analogous to our construction of generators for the group $\langle Q, L \rangle$) is equal to $N_G(Q)' = (\Omega^\varepsilon(V')_{x_+})'$. Since each group $D_J^{\sigma^i}$ normalises both Q and Q^l, so does L, and hence $\langle Q, L \rangle = Q \rtimes L$. We now use L in place of the quotient group $H = G'_\alpha/Q$ in the proof of [KS1, Lemma 4.14] (which is shown there to be isomorphic to $\Omega(Q) \cong \Omega^\varepsilon(d-2, q)$) to see that L is indeed the desired complement to Q in $N_G(Q)'$. $\square$

Remark 3.11. The "prime" notation in the previous proof is used to distinguish between the vector space V' underlying the black box group G and the concrete space V that we constructed in Section 3.1.3. The space V' is used only within proofs; see also the proofs of Lemmas 3.22 and 4.9.

Remark 3.12. In the preceding lemma we see a fundamental difference between [KS1] and the present paper: since we did not have to handle some of the small field cases that caused a great deal of annoyance in [KS1], we were able to obtain a complement L without using the "effective transitivity" of the subgroup Q (cf. [KS1, Lemma 4.17]). This simplification is facilitated by the nature of the element σ we employ here, which, in its action on the module underlying G, induces the identity on the 2-dimensional support of a; the singular points of this hyperbolic line correspond to the groups $\dot Q$ and $\dot Q^l$. Hence, the σ-conjugates of D_J all normalise $\dot Q$ and $\dot Q^l$, unlike their counterparts in [KS1], which had to be modified

in order to have this property (cf. [KS1, Corollary 4.18]). Nevertheless, in Section 4 we will present a faster, randomised version of "effective transitivity."

Remark 3.13. At this stage of the algorithm, there appears to be no easy way to verify that Q and L are the desired subgroups $\dot{Q}$ and $\dot{L}$. Thus the main algorithm to date is Monte Carlo rather than Las Vegas. However, the uneasiness this causes will be removed following the execution of Procedure 3.21 in 3.3.3; that routine provides, as a by-product, a "correctness check" for our construction of $\dot{Q}$. Thus *we will henceforth assume that $Q = \dot{Q}$ and dispense with the "dot" notation.*

3.1.5. Total timing and reliability for Section 3.1 ("large" fields)

By Lemmas 3.2, 3.4 and 3.10, we see that all of the preliminary constructions are obtained with probability $> 1 - (1/e^3 + 1/e^3 + 1/8)$ in $O(d^3 \log d \log^4 q + \xi\{d + \log q\} + \chi \log^2 q + \mu \log q \{d + \log^2 q\})$-time.

3.2. Preliminary constructions for "small" fields

We next consider the case $q < 16$. Although timing considerations prevent us from using the entire algorithm in [KS1, Section 4] to handle these small fields, we can make use of some of its key subroutines to obtain analogues of the preliminary constructions in Section 3.1. We may assume that $d \geqslant 9$, since otherwise $q < 16$ and $d \leqslant 8$ so we can use brute force.

Remark 3.14. Due to the use of recursion, the reliability estimates of most of the randomised subroutines in [KS1] depend upon d. Our approach avoids recursion and hence this dependence on d: we are usually able to improve (by a factor of $\log d$) the timing of subroutines from [KS1]. In addition, the timing for any call to the $SL(2, q)$-oracle reduces, in this setting, to $O(\mu)$.

3.2.1. The elements t and σ

In Section 3.1, the commuting elements a and σ were fundamental to our preliminary constructions. For small q the situation is not quite so straightforward, but the hard work has already been done in [KS1, Section 4.2.1]. In each of the nine cases considered therein, a *long root element t is constructed*, together with an element τ fixing some point x_+ in the support of t. This is achieved using up to c_1 choices of element from G, and possibly up to c_2 constructive isomorphism tests for either $\Omega^+(4, q)$ or $\Omega^+(6, q)$. The integers c_1 and c_2 in each of the nine cases considered in [KS1, Section 4.2.1] are stated explicitly below, where we also indicate how to use τ to construct an element σ *fixing the point x_+.*

Case 1. $[\Omega^\varepsilon(d, q), q \leqslant 13$ odd]: $c_1 := 8q^2 d$; $\sigma := \tau^{p(q+1)}$.
Case 2. $[\Omega^+(2n, 2^k), n \geqslant 5$ odd, $k = 1, 2, 3]$: $c_1 := 16qd$; $\sigma := \tau^{p(q+1)}$.
Case 3. $[\Omega^-(2n, 2^k), n \geqslant 6$ even, $k = 1, 2, 3]$: $c_1 := 16qd$; $\sigma := \tau^{p(q+1)}$.
Case 4. $[\Omega^+(2n, 8), n \geqslant 6$ even]: $c_1 := 8q^2 d$; $c_2 := 2^{12}$; $\sigma := \tau^{(q-1)(q^2+1)} = \tau^{455}$.
Case 5. $[\Omega^-(2n, 8), n \geqslant 5$ odd]: $c_1 := 32d$; $c_2 := 2^{12}$; $\sigma := \tau^{q-1} = \tau^7$.
Case 6. $[\Omega^+(2n, 4), n \geqslant 6$ even]: $c_1 := 256d$; $c_2 := 2^{14}$; $\sigma := \tau^{51}$.

Case 7. $[\Omega^-(2n, 4), n \geqslant 5 \text{ odd}]$: $c_1 := 32d$; $c_2 := 2^{14}$; $\sigma := \tau^3$.
Case 8. $[\Omega^+(2n, 2), n \geqslant 6 \text{ even}]$: $c_1 := 8d$; $c_2 := 32$; $\sigma := \tau^5$.
Case 9. $[\Omega^-(2n, 2), n \geqslant 5 \text{ odd}]$: $c_1 := 40d$; $c_2 := 32$; $\sigma := \tau^5$.

The required element t is constructed from τ and σ in [KS1, Section 4.2.1].

Timing.　　The timing in each case is dominated by the $c_1 = O(d)$ choices of element from G (together with primitive prime divisor tests for each), and the $c_2 = O(1)$ constructive isomorphism tests. Hence, the elements σ, τ and t are obtained in $O(\xi d + \mu d^3)$-time (cf. [KS1, Section 4.2.1] and Remark 3.14 above).

Reliability.　　In [KS1, Section 4.2.1], c_1 and c_2 were chosen so that success is ensured with probability $\geqslant 1 - 1/8d^2$; here we crudely modified c_1 and c_2 so that we succeed with probability $\geqslant 1 - 1/8$ (cf. Remark 3.14).

3.2.2. Constructing J

As in [KS1, Section 4.2.2], choose up to 2^{21} triples $f_1, f_2, f_3 \in G$ in order to find, and constructively recognise, at least 2^6 subgroups $J_i = \langle t, t^{f_1}, t^{f_2}, t^{f_3} \rangle \cong \Omega^-(8, q)$ with probability $\geqslant 1 - 1/16$. (Note that $|\Omega^-(8, q)|$ is bounded.) For each i, use the resulting isomorphism $\Psi_i : \Omega^-(V_i) \to J_i$ to construct subgroups Q_{8i}, G_{8i} and D_i, where $G_{8i} = (\Omega^-(V_i)_{x_i})' \Psi_i$, $Q_{8i} = O_p(G_{8i})$ and

$$G_{8i} = Q_{8i} \rtimes D_i, \tag{3.15}$$

for some singular point $x_i \in V_i$. Furthermore, as in [KS1, Section 4.2.2], choose the x_i so that they all correspond to the same point x_+, fixed by σ, within the support of t. Finally, fix any i, and restrict Ψ_i to a suitable naturally embedded $\Omega^+(6, q)$ subgroup J of J_i containing t in order to obtain an isomorphism $\Psi_J : \Omega^+(6, q) \to J$.

Timing and reliability.　　The requisite number of groups Q_{8i} and G_{8i}, together with the distinguished isomorphism Ψ_J, are constructed, with probability $> 1 - 1/16$, in $O(\xi + \mu)$-time (cf. [KS1, Section 4.2.2] and Remark 3.14).

3.2.3. Some elements of J

We may again assume that elements of $\Omega^+(6, q)$ are written relative to a standard basis e_1, e_2, e_3, e_{-1}, e_{-2}, e_{-3}, and construct all of the elements and subgroups obtained in Section 3.1.3. We may further assume that t is contained in the long root group R_1.

3.2.4. Constructing $\dot{Q}$ and $\dot{L}$

Exactly as in [KS1, Section 4.3.1], use the 2^6 groups Q_{8i} and G_{8i}, together with the element τ, to (probably) construct the groups Q and $N_G(Q)'$. Use [KS1, Section 4.3.2] to (probably) construct a complement L to Q in $N_G(Q)'$. Then modify the generating set for Q to consist of L-conjugates of generators of R_1.

Timing and reliability. As in Section 3.1.4, with probability $> 1 - 1/8$, Q is the target group $\dot{Q}$ of order q^{d-2}, and L is the complement $\dot{L}$ to $\dot{Q}$ in $N_G(\dot{Q})'$ that normalises $\dot{Q}^l$. The constructions require $O(\mu d^3 \log^2 d)$-time (cf. [KS1, Section 4.3.2] and Remark 3.14 again).

Remark 3.16. Observe that [KS1, Section 43.2] uses the effective transitivity of Q [KS1, Lemma 4.17]. However, in the present setting there is no problem with the timing (cf. Remark 3.12) since q is bounded. *Once again, as in* Remark 3.13*, we assume that $Q = \dot{Q}$ and dispense with the "dot" notation.*

3.2.5. Total timing and reliability for Section 3.2

We obtain all of the elements and subgroups in Section 3.2 with probability $> 1 - (1/8 + 1/16 + 1/8)$ in $O(\xi d + \mu d^3 \log^2 d)$-time.

3.3. Algorithms for Q

Let $\Psi_J : \Omega^+(V_J) \to J$ be the isomorphism constructed by Procedure 3.3 or Section 3.2.2, and let Q and L be the groups constructed in (3.9) or Section 3.2.4. Recall that (with high probability) Q is the natural orthogonal module for L, where L corresponds to the group $\Omega^\varepsilon(V)_{e_1, e_{-1}} \cong \Omega^\varepsilon(d - 2, q)$. However, since we do not have the action of $\mathbb{F}_q$ on Q, we cannot yet regard Q *algorithmically* as an $\mathbb{F}_q$-module. The main construction in this section constitutes a step in this direction: in Section 3.3.3 we obtain a matrix A_Q representing a nondegenerate L-invariant quadratic form on Q.

3.3.1. Evaluating forms

In this subsection and the next we present two technical subroutines involving $\Omega^+(6, q)$-subgroups that are needed for our construction of the matrix A_Q.

Let $K = \langle \mathcal{S}_K \rangle$ be any given black box group isomorphic to $\Omega^+(6, q)$, and let $\Psi_K : \Omega^+(6, q) \to K$ be an effective isomorphism. Denote the invariant quadratic form on the underlying 6-space V_K by ϕ, fix a standard basis of V_K relative to ϕ, and let f_1, f_{-1} be a hyperbolic pair in this basis. Let Q_K denote the subgroup $O_p(\Omega^+(6, q)_{\langle f_1 \rangle}) \Psi_K$, and let $L_K := \Omega^+(6, q)_{f_1, f_{-1}} \Psi_K \cong \Omega^+(4, q)$. Then, as in (2.11), the natural $\Omega^+(6, q)_{f_1, f_{-1}}$-invariant quadratic form $\phi_{\langle f_1 \rangle}$ on $O_p(\Omega^+(6, q)_{\langle f_1 \rangle})$, together with the isomorphism Ψ_K, can be used to define a corresponding L_K-invariant form ϕ_{Ψ_K} effectively on Q_K.

Specifically, the following subroutine takes, as input, the effective isomorphism Ψ_K and any element $u \in Q_K$, and outputs the scalar $\phi_{\Psi_K}(u) \in \mathbb{F}_q$. (Of course the subroutine can then also compute values for the associated bilinear form $(\ ,\ )_{\Psi_K}$ via the equation $(u_1, u_2)_{\Psi_K} = \phi_{\Psi_K}(u_1 u_2) - \phi_{\Psi_K}(u_1) - \phi_{\Psi_K}(u_2)$.)

Subroutine A. Use Lemma 2.16 to compute the preimage

$$u' := u \Psi_K^{-1} \in O_p\big(\Omega^+(6, q)_{f_1}\big).$$

Compute the vector $f_{-1} u'$ relative to our standard basis, and return the f_1-component of this vector.

Lemma 3.17. *Subroutine* A *is a Las Vegas routine that returns* $\Psi_K(u)$, *with probability* $> 1 - 1/128$, *in* $O(\xi + \chi \log q + \mu \log^2 q)$-*time.*

Proof. By definition, $\phi_{\Psi_K}(u) = \phi_{\langle f_1 \rangle}(u') = \phi(w)$, where $r_1(w) = u'$ (see Eqs. (2.5) and (2.11)). The proof now follows from (2.4) by noting that $f_{-1} r_1(w) = f_{-1} + w + \phi(w) f_1$. The stated timing and reliability estimates are precisely those for Lemma 2.16. □

3.3.2. Matching forms

In the next section we will need to consider a second naturally embedded $\Omega^+(6, q)$-subgroup K of G such that $J_0 = J \cap K \cong \Omega^+(4, q)$; here K is given together with an effective isomorphism $\Psi_K : \Omega^+(6, q) \to K$. By Section 3.3.1, Ψ_J induces a nondegenerate $L \cap J$-invariant quadratic form ϕ_{Ψ_J} on $Q_J = Q \cap J$, while Ψ_K induces an $L \cap K$-invariant form on $Q_K = Q \cap K$. Then $Q_J \cap J_0$ is a hyperbolic line of Q_J (as well as of Q_K); nonsingular elements of Q_K are meaningful in terms of ϕ_{Ψ_K}.

In this setting the following subroutine takes as input J, K, Ψ_J, Ψ_K and a nonsingular element $u_0 \in Q \cap J_0$, and returns a scalar $\lambda \in \mathbb{F}_q$.

Subroutine B. Using Subroutine A in Section 3.3.1 twice (once with input Ψ_J and u_0, and once with input Ψ_K and u_0), compute $\alpha_1 := \phi_{\Psi_J}(u_0)$ and $\alpha_2 := \phi_{\Psi_K}(u_0)$. Return $\lambda := \alpha_1/\alpha_2 \in \mathbb{F}_q$.

Lemma 3.18. *Subroutine* B *is a Las Vegas routine which, with probability at least* $1 - 1/64$, *in* $O(\xi + \chi \log q + \mu \log^2 q)$-*time returns a scalar* λ *such that*

$$\phi_Q = \lambda \phi'_Q, \tag{3.19}$$

where ϕ_Q *and* ϕ'_Q *are* L-*invariant quadratic forms on* Q *that extend* ϕ_{Ψ_J} *and* ϕ_{Ψ_K}, *respectively.*

Proof. Since any two nondegenerate L-invariant quadratic forms on Q differ only by a scalar, there is a scalar behaving as in (3.19). Now use Lemma 3.17. □

3.3.3. The matrix A_Q

We do not yet know that Q is the correct group $\dot{Q}$, nor have we constructed an action of $\mathbb{F}_q$ turning Q into an $\mathbb{F}_q$-space. Thus, although there is a unique nondegenerate L-invariant quadratic form ϕ_Q on Q that extends the nondegenerate L_J-invariant quadratic form ϕ_{Ψ_J} on Q_J (defined, effectively, by Subroutine A), we do not yet have ϕ_Q available to work with. Nevertheless, we can now construct a lower triangular $((d - 2) \times (d - 2))$-matrix $A_Q = [\![\alpha_{ij}]\!]$, with entries in $\mathbb{F}_q$, that will allow us to specify the bilinear form $(\,,\,)_Q$ associated with ϕ_Q. Namely, we will have

$$\alpha_{ij} = (r_i, r_j)_Q \quad \text{for } 1 \leqslant i < j \leqslant d - 2, \tag{3.20}$$

where $r_i := r_{i1}$, $1 \leqslant i \leqslant d - 2$ (cf. (3.9)), will be a basis of Q once we know that $Q = \dot{Q}$ is an $\mathbb{F}_q$-space of dimension $d - 2$. More precisely, in the notation of 3.1.3 and (3.9), the following procedure takes as input

(1) the effective isomorphism $\Psi_J : \Omega^+(V_J) \to J$,
(2) the subgroup $J_0 \leqslant J$ (isomorphic to $\Omega^+(4, q)$) and $Q_0 := Q \cap J_0$,
(3) the elements r_i and $u_0 := r_1(e_2 + e_{-2})\Psi_J \in Q_0$, and
(4) the subgroup L,

and returns one of the following outputs:

(a) A matrix $A_Q = [\![\alpha_{ij}]\!]$ together with the report "$\boldsymbol{Q = \dot{Q}}$" (this is the desired output);
(b) The report "$\boldsymbol{Q \neq \dot{Q}}$" (this output occurs if the procedure runs as expected, but we failed to generate $\dot{Q}$ in Section 3.1.4 or 3.2.4); or
(c) The report "**failure**" (this output occurs if the procedure encounters bad luck with its random choices).

Procedure 3.21 *(construction of A_Q).*

1. Initialise $\alpha_{ij} := 0$ for all $i \geqslant j$. Next, for $1 \leqslant i < j \leqslant 4$, use Subroutine A to compute $\alpha_{ij} := (r_i, r_j)_{\Psi_J}$.
2. For all other pairs $1 \leqslant i < j \leqslant d - 2$, proceed as follows:
   <u>i</u> For at most $\lceil 12 \log(4d) \rceil$ choices $c \in L$, set

   $$K = K_{ij}(c) := \langle J_0, R_i^c, R_j^c \rangle$$

   for $J_0 \cong \Omega^+(4, q)$ in Section 3.1.3(v); use Lemma 2.16 to test whether $K \cong \Omega^+(6, q)$; if so, let $\Psi_K : \Omega^+(6, q) \to K$ denote the resulting effective isomorphism, and move to step <u>iii</u>.
   <u>ii</u> If no Ψ_K is found in step <u>i</u>, report "**failure**" and stop.
   <u>iii</u> Repeat Subroutine B at most $\lceil 2 \log(8d)/7 \rceil$ times (with input Ψ_J, Ψ_K, u_0) to find a scalar $\lambda = \lambda_{ij}(c)$.
   <u>iv</u> Repeat Subroutine A at most $\lceil \log(8d)/3 \rceil$ times to find $\beta_{ij} := (r_i^c, r_j^c)_{\Psi_K}$. Put $\alpha_{ij} := \lambda \beta_{ij}$.
3. Let $A_Q := [\![\alpha_{ij}]\!]$.
   If $A_Q + A_Q^{\mathrm{tr}}$ is nonsingular, return A_Q together with the report "$\boldsymbol{Q = \dot{Q}}$".
   If $A_Q + A_Q^{\mathrm{tr}}$ is singular, report "$\boldsymbol{Q \neq \dot{Q}}$" and stop.

Lemma 3.22. *Procedure* 3.21 *is a Las Vegas algorithm which, with probability at least* $15/16$, *in* $O(d^2 \log d \cdot \log q\{\xi + \chi \log q + \mu \log^2 q\} + d^3 \log q)$-*time returns an output of type* (a) *or* (b). *Moreover, if one of these outputs is returned, then it is guaranteed to be correct in the following sense*:

for a type (a) *output, $Q = \dot{Q}$ with probability* 1 *and* (3.20) *holds; and*
for a type (b) *output, $Q \neq \dot{Q}$ with probability* 1.

Proof. We consider correctness, reliability and timing of Procedure 3.21 separately. Recall that, since Remarks 3.13 and 3.16, we have been assuming that Q is the desired group $\dot{Q}$.

Correctness.   This principally involves showing that the scalars α_{ij} satisfy (3.20). This holds when $i = j$ since r_i is a long root element (cf. (2.8)). Since we want A_Q to be lower triangular, there is nothing to prove concerning the remaining α_{ij} initialised in step 1. Consider a fixed pair $1 \leqslant i < j \leqslant d - 2$ in step 2, and suppose that the procedure succeeds in producing the following:

2.<u>i</u>   an isomorphism $\Psi_K : \Omega^+(6, q) \to K = \langle J_0, R_i^c, R_j^c \rangle$ for some $c \in L$;
2.<u>iii</u> a scalar λ; and
2.<u>iv</u> a scalar α_{ij}.

In Lemma 3.18 we use the element u_0, which is nonsingular with respect to ϕ_Q (since Q_0 is isometric to the hyperbolic line $\langle e_2, e_{-2} \rangle$): λ in step 2.<u>iii</u> satisfies $\phi'_Q = \lambda \phi_Q$. Hence, $\alpha_{ij} = \lambda \beta_{ij} = \lambda (r_i^c, r_j^c)'_Q = \lambda (r_i, r_j)'_Q = (r_i, r_j)_Q$, as required.

Thus, we have obtained the correct values $(r_i, r_j)_Q$ for $1 \leqslant i < j \leqslant d - 2$. Moreover, $A_Q + A_Q^{\mathrm{tr}}$ is the matrix of $(,)_Q$ with respect to the $\mathbb{F}_q$-basis $r_1, \ldots, r_{d-2}$ of Q. In particular, this matrix is nonsingular.

It remains to show that, *if Q is not the desired group $\dot{Q}$, then $A_Q + A_Q^{\mathrm{tr}}$ is singular.* By 3.1.3(i), R_i, $1 \leqslant i \leqslant 4$, are $\mathbb{F}_q$-subspaces of the $\mathbb{F}_q$-space $\dot{Q}$, and hence so is Q by (3.9). Thus, $A_Q + A_Q^{\mathrm{tr}}$ is the matrix of a symmetric bilinear form defined using linearly dependent vectors $r_1, \ldots, r_{d-2}$ of that $\mathbb{F}_q$-space, and hence is, indeed, singular.

Thus, Q is the desired group $\dot{Q}$ precisely as the test in step 3 decides.

Reliability.   Observe that Lemma 2.16, Subroutines A and B are Las Vegas algorithms, so the present procedure is also Las Vegas.

Since an upper bound on the probability that $Q \neq \dot{Q}$ was already obtained in 3.1.4 and 3.2.4, and was included in the reliability estimates for Sections 3.1 and 3.2, *we may once again assume from now on that $Q = \dot{Q}$.*

Claim 1. *For a fixed pair $1 \leqslant i < j \leqslant d - 2$, a choice $c \in L$ produces $K = K_{ij}(c) \cong \Omega^+(6, q)$ with probability at least $1/4$.*

Let V' denote the $\mathbb{F}_q$-space underlying G. The group J_0 acts on V' with nondegenerate 4-dimensional support $[V', J_0]$, and Q_0 fixes a singular point x'_+ of $[V', J_0]$. Similarly, Q_0^l fixes the singular point $x'_- = (x'_+)^l \notin (x'_+)^\perp$ of $[V', J_0]$. Consider $V'_K = [V', K]$ and $Q_K = \langle Q_0, R_i^c, R_j^c \rangle$. Since V'_K is spanned by the spaces $[V', \langle J_0, R_i^c \rangle]$ and $[V', \langle J_0, R_j^c \rangle]$, each of dimension at most 5 (since x_+ is in the supports of J_0, R_i and R_i^c), we have $\dim V'_K \leqslant 6$.

By Lemma 2.1(a) or (b), with probability at least $1/4$, Q_K is a 4-space of Q of Witt index 2. In that case, by the definition of $Q(x_+)$ in Section 2.3, $[V', Q_K]$ is a 5-space with radical x'_+ and not containing x'_-. It follows that V'_K contains the distinct 5-spaces $[V', Q_K]$ and $[V', Q_K^l]$, so that $\dim V'_K \geqslant 6$. For such Q_K, moreover, since V'_K contains the hyperbolic line $\langle x'_+, x'_- \rangle$, and since Q_K is isometric to $V'_K / \langle x'_+, x'_- \rangle$, it follows that V'_K is nondegenerate of (maximal) Witt index 3.

Thus, with the stated probability, for a fixed choice $c \in L$ the group K is isomorphic to a subgroup of $\Omega^+(6, q)$. Claim 1 now follows by noting that, for such a c, the groups Q_K and Q_K^l generate a subgroup of K isomorphic to $\Omega^+(6, q)$.

Claim 2. *For a fixed pair $1 \leqslant i < j \leqslant d - 2$, steps 2.$\underline{i}$–2.$\underline{iv}$ produce α_{ij} with probability at least $1 - 1/8d^2$.*

For, if $c \in L$ gives rise to $K \cong \Omega^+(6, q)$, then Lemma 2.16 produces an effective isomorphism $\Psi_K : \Omega^+(6, q) \to K$ with probability $> 1/2$. Hence, by Claim 1, c produces a suitable Ψ_K with probability $> (1/4) \cdot (1/2)$. The probability that none of the $\lceil 12 \log(4d) \rceil$ choices c produces a suitable Ψ_K is therefore at most $(7/8)^{12 \log(4d)} < 1/(4d)^2$.

On the other hand, the repetitions of Subroutines A and B in steps 2.$\underline{iii}$ and 2.$\underline{iv}$ ensure that each of the scalars λ and α_{ij} is obtained with probability $> 1 - 1/(8d)^2 - 1/(8d)^2$. For, a single call to Lemma 3.17 succeeds with probability $\geqslant 1 - 1/128$; thus all of the calls fail with probability $\leqslant (1/2^7)^{\log(8d)/3} < (1/8d)^2$. Similarly, all of the calls to Lemma 3.18 fail with probability $\leqslant (1/8d)^2$.

Hence, steps 2.$\underline{i}$–$\underline{iv}$ compute the desired scalar α_{ij} with probability $> 1 - (1/16d^2 + 1/64d^2 + 1/64d^2) > 1 - 1/8d^2$, as claimed.

It now follows that the probability that at least one of the at most $d^2/2$ iterations of step 2 fails to produce the corresponding scalar α_{ij} for some i, j is less than $(d^2/2)(1/8d^2) = 1/16$. Thus the procedure computes *all* of the scalars α_{ij} with probability $> 15/16$.

Timing. The timing of each iteration of the main loop is dominated by step 2.$\underline{i}$, in which $O(\log d)$ calls to Lemma 2.16 are made. There are $O(d^2)$ iterations, so the procedure takes $O(d^2 \log d \cdot \log q \{\xi + \chi \log q + \mu \log^2 q\} + d^3 \log q)$-time, as stated. (The $d^3 \log q$ term is the time required to test whether $A_Q + A_Q^{\text{tr}}$ is nonsingular in step 3.) $\quad \square$

3.3.4. Defining Ψ on $Q(x_+)$

The following procedure takes as input the matrix A_Q returned by Procedure 3.21, and returns a bijection, f^*, between generating sets for the matrix group $Q(x_+)$ and the black box group Q. (Recall the definition of the elements r_{ij}: see (3.9) for $q \geqslant 16$, and Section 3.2.4 for $q < 16$.)

Procedure 3.23.

1. Initialise $w_1 := e_2$, $w_2 := e_{-2}$, $w_3 := e_3$, $w_4 := e_{-3}$.
2. Using linear algebra in the $(d - 6)$-dimensional orthogonal space $V_J^\perp \subset V = \mathbb{F}_q^d$, find a basis $w_5, \ldots, w_{d-2}$ of *singular* vectors having the property that $(w_i, w_j) = \alpha_{ij}$ for all $5 \leqslant i < j \leqslant d - 2$.
3. For $1 \leqslant i \leqslant d - 2$, $1 \leqslant j \leqslant k$, let $r'_{ij} := r_1(\rho^{j-1} w_i)$, and set

$$
\begin{aligned}
\mathcal{T}(x_+) &:= \{r'_{ij} \mid 1 \leqslant i \leqslant d - 2, \ 1 \leqslant j \leqslant k\}, \\
\mathcal{S}_Q^* &:= \{r_{ij} \mid 1 \leqslant i \leqslant d - 2, \ 1 \leqslant j \leqslant k\}.
\end{aligned}
\tag{3.24}
$$

4. Return the bijection $f^* : \mathcal{T}(x_+) \to \mathcal{S}_Q^*$ sending $r'_{ij} \mapsto r_{ij}$ for $1 \leqslant i \leqslant d - 2$, $1 \leqslant j \leqslant k$.

Only linear algebra and bookkeeping are required for the construction of $\mathcal{T}(x_+)$. The timing for Procedure 3.23 is therefore dominated by that of our earlier constructions. The following result demonstrates that the bijection f^* output by the procedure is induced by an epimorphism $\Omega^\varepsilon(V) \to G$ upon restriction; note, however, that the homomorphism $Q(x_+) \to Q$ defined by f^* is not yet an effective homomorphism.

Lemma 3.25. *Let* $f^* : \mathcal{T}(x_+) \to \mathcal{S}_Q^*$ *be the bijection output by Procedure* 3.23 *and let* $f^* \cup \Psi_J$ *denote the obvious bijection* $\mathcal{T}(x_+) \cup \Omega^+(V_J) \to \mathcal{S}_Q^* \cup J$. *Then there is a unique epimorphism* $\Psi : \Omega^\varepsilon(V) \to G$ *that restricts to* $f^* \cup \Psi_J$.

Proof. We will show that f^* extends (uniquely) to an isometry $f : Q(x_+) \to Q$ that coincides with Ψ_J on $Q(x_+) \cap \Omega^+(V_J)$; the result will then follow from Proposition 2.12. (Note that the quadratic form on $Q(x_+)$ is ϕ_{x_+}, defined as in (2.5) via the assignment $\phi_{x_+}(r_1(w)) := \phi(w)$ for $w \in \langle x_+, x_- \rangle^\perp$, and the quadratic form on Q is ϕ_Q, whose matrix was constructed in Procedure 3.21.)

By construction, the basis $w_1, \ldots, w_{d-2}$ satisfies $(w_i, w_j) = (r_i, r_j)_Q$ for $1 \leqslant i < j \leqslant d-2$ (cf. (3.20)). Since the w_i are singular, and the r_i are long root elements, we also have $\phi(w_i) = \phi_Q(r_i) = 0$. As the w_i and r_i are $\mathbb{F}_q$-bases of $\langle x_+, x_- \rangle^\perp$ and Q respectively, it follows that $w_i \mapsto r_i$ extends uniquely to an isometry $\langle x_+, x_- \rangle^\perp \to Q$. Thus the map $r_1(w_i) \mapsto r_i$ extends uniquely to an isometry $f : Q(x_+) \to Q$. Furthermore, the vectors $w_1, \ldots, w_4$ were selected so that $r_1(\rho^{j-1} w_i)\Psi_J = r_{ij}$ ($1 \leqslant i \leqslant 4, 1 \leqslant j \leqslant k$), so f coincides with Ψ_J on $Q(x_+) \cap \Omega^+(V_J)$. It now follows from Proposition 2.12 that there is a unique epimorphism $\Psi : \Omega^\varepsilon(V) \to G$ extending both Ψ_J and the map $r_1(w_i) \mapsto r_i$, for $1 \leqslant i \leqslant d-2$.

It remains to show that $r_{ij}'\Psi = r_{ij}$ for $1 \leqslant i \leqslant d-2$, $1 \leqslant j \leqslant k$ (that is, we must verify that the bijection f^* is consistent with the action of the field). This is clear for $1 \leqslant i \leqslant 4$, since f^* was constructed to agree with Ψ_J on $\mathcal{T}(x_+) \cap \Omega^+(V_J)$. For $i \geqslant 5$, by construction of the r_{ij}, there exists $\sigma_i \in L$ such that $r_{ij} = r_{1j}^{\sigma_i}$ ($5 \leqslant i \leqslant d-2$). Let σ_i' denote the unique element of the preimage $\sigma_i \Psi^{-1}$ that fixes the two vectors $e_1, e_{-1} \in V$. Then

$$r_1(w_i) = r_{i1}\Psi^{-1} = r_{11}^{\sigma_i}\Psi^{-1} = \left(r_{11}\Psi^{-1}\right)^{\sigma_i'} = r_1(w_1)^{\sigma_i'} = r_1\left(w_1^{\sigma_i'}\right),$$

so that $w_1^{\sigma_i'} = w_i$. Now

$$r_{ij}'\Psi = r_1\left(\rho^{j-1} w_i\right)\Psi = r_1\left(\rho^{j-1} w_1\right)^{\sigma_i'}\Psi = r_1\left(\rho^{j-1} w_1\right)\Psi^{\sigma_i'\Psi} = r_{1j}^{\sigma_i} = r_{ij}, \quad (3.26)$$

as desired. $\square$

3.4. The data structure for Ψ

We are finally in a position to describe our data structure.

3.4.1. Algorithms for the natural representation

Our strategy is to follow the algorithm presented in [Br1] for orthogonal groups in their natural representation. We therefore provide a brief commentary of the relevant parts of that algorithm, and indicate the modifications needed to make use of it in the present setting.

In [Br1, 4.4] a generating set, $\mathcal{T}(x_+)$, consisting of long root elements, is constructed for the group $Q(x_+)$. As the notation suggests, this generating set is an analogue of the one in (3.24), and has all of the properties of the latter set.

Then, in [Br1, 4.5], the set $\mathcal{T}(x_+)$ is replaced by a standard generating set $\Delta(x_+)$. More precisely, each element of a fixed, prescribed set $\Delta(x_+)$ is constructed using a short SLP from $\mathcal{T}(x_+)$. The routine to construct $\Delta(x_+)$ takes, in fact, *any* generating set $\mathcal{T}'$ for $Q(x_+)$, and produces SLPs of length $O(d \log q)$ from $\mathcal{T}'$ to the elements of $\Delta(x_+)$. Thus, although the set $\mathcal{T}(x_+)$ obtained here may differ from its analogue in [Br1], we can still construct *exactly the same set $\Delta(x_+)$ obtained in* [Br1, 4.5].

Finally, in [Br1, 4.6], $O(kd^2)$ elements of a "canonical" generating set $\mathcal{T}$ for $\Omega^\varepsilon(V)$ are constructed using short SLPs from the set $\Delta(x_+) \cup \Delta(x_+)^{l'}$. In our setting, by using the subroutines in [Br1, 4.4 through 4.6] we can construct an SLP of length $O(d \log q)$ from $\mathcal{T}(x_+) \cup \{l'\}$ to each element of the analogous set $\mathcal{T}$ constructed in [Br1].

Timing. The total timing for Sections 4.5 and 4.6 of [Br1] is $O(d^3 \log^2 q)$, and hence is dominated by the timing for our previous constructions.

3.4.2. New generators for G

By Lemma 3.25, the obvious bijection $\mathcal{T}(x_+) \cup \{l'\} \to \mathcal{S}_Q^* \cup \{l\}$ is the restriction of a unique epimorphism $\Psi : \Omega^\varepsilon(V) \to G$. In 3.4.1, an SLP was constructed from $\mathcal{T}(x_+) \cup \{l'\}$ to each $t' \in \mathcal{T}$; use the above bijection to find the image $t'\Psi$, and set

$$\mathcal{S}^* := \{t'\Psi \mid t' \in \mathcal{T}\}. \tag{3.27}$$

Timing. In fact we do not evaluate each SLP from scratch: this would require $O(\mu d \log q \cdot kd^2) = O(\mu d^3 \log^2 q)$-time. Instead, we construct the image of $\Delta(x_+)$ in G in $O(\mu kd \cdot d \log q) = O(\mu d^2 \log q)$-time. Then additional $O(\mu kd^2)$-time is required to construct all of the elements of $\mathcal{S}^*$. Hence, we construct $\mathcal{S}^*$ in $O(\mu d^2 \log q)$-time.

3.4.3. The data structure

The data structure returned by the preprocessing phase of the algorithm contains the following information (compare with the summary at the start of Section 3):

(a) The generating set $\mathcal{T}$ for $\Omega^\varepsilon(V)$.
(b) The subset $\mathcal{S}^*$ of G.
(c) The obvious bijection $\mathcal{T} \to \mathcal{S}^*$ (extending to the target epimorphism Ψ).
(d) The effective isomorphism $\Psi_J : \Omega^+(V_J) \to J$ (also extending to Ψ).
(e) The elements $l, h, h^+ \in J$ constructed in Section 3.1.3, together with their preimages $l', (h^+)' \in \Omega^+(V_J)$.

3.5. Total timing and reliability for Section 3

The failure probabilities of the subroutines in this section sum to less than $1/2$. The timing is dominated by Sections 3.1.1 and 3.3.3. Hence, our algorithm returns a suitable data structure, with probability $> 1/2$, in $O(d^3 \log d \log q \{d + \log^3 q\} + \xi d^2 \log d \log q + \chi d^2 \log d \log^2 q)$-time, as stated in Theorem 1.1.

4. Straight-line programs

The previous section dealt with the *preprocessing phase* of the algorithm: we produced a data structure that specifies an epimorphism, Ψ, from a matrix group $\Omega^\varepsilon(V)$ to our given black box group $G = \langle S \rangle$. This section deals with the *application phase* of the algorithm: we present routines that compute images and preimages under Ψ. It suffices to give algorithms to solve each of the following problems:

(SLP1) *Given $g' \in \Omega^\varepsilon(V)$, write an SLP of length $O(d^2 \log q)$ from $\mathcal{T}$ to g'.*
(SLP2) *Given $g \in G$, write an SLP of length $O(d^2 \log q)$ from $\mathcal{S}^*$ to g.*

For example, given $g' \in \Omega^\varepsilon(V)$, we take the SLP returned by **(SLP1)** and evaluate it from the set $\mathcal{S}^*$ (via the bijection $\mathcal{T} \to \mathcal{S}^*$) to obtain the image $g'\Psi$ in G.

4.1. An algorithm for (SLP1)

In [Br1, Section 5], a deterministic $O(d^3 \log q)$-time algorithm is presented, which writes an SLP of length $O(d^2 \log q)$ from essentially the same generating set $\mathcal{T}$ as the one constructed here to any given element of $\Omega^\varepsilon(d, q)$. This yields an $O(\mu d^2 \log q)$-time algorithm to find the Ψ-image of any given element of $\Omega^\varepsilon(d, q)$, as stated in Theorem 1.1.

4.1.1. Constructing $Z(G)$
Until now, we have assumed only that G is a homomorphic image of a known orthogonal group: we do not yet know whether G is $\Omega^\varepsilon(d, q)$ or $P\Omega^\varepsilon(d, q)$. Using **(SLP1)**, we can now decide which of those groups G is, by constructing $Z(G)$:

Note that $Z(\Omega^\varepsilon(d, q)) \neq 1$ if and only if $d \equiv 0 \pmod 4$ and $\varepsilon = 1$, or if $d \equiv 2 \pmod 4$ and $q \equiv \varepsilon \pmod 4$. In each of these cases, write an SLP from $\mathcal{T}$ to the matrix $z' = -I_d$, where I_d is the $d \times d$ identity matrix. Evaluate that SLP from $\mathcal{S}^*$ to obtain an element $z \in G$; then $Z(G) = \langle z \rangle$. If $z = 1$ then $G \cong P\Omega^\varepsilon(d, q)$. If $z \neq 1$ then, for convenience later on, replace $\mathcal{T}$ by $\mathcal{T} \cup \{-I_d\}$, $\mathcal{S}^*$ by $\mathcal{S}^* \cup \{z\}$, and extend the bijection accordingly.

Timing. **(SLP1)** writes an SLP of length $O(d^2 \log q)$ from $\mathcal{T}$ to $-I_d$ in $O(d^3 \log q)$-time; the evaluation of this SLP is then carried out in $O(\mu d^2 \log q)$-time. We consider the construction of $Z(G)$ to be part of the preprocessing phase, and its timing is dominated by that of Section 3.

Table 1

[KS1] algorithm	Present algorithm
V_L	$\langle e_1, e_{-1}\rangle^{\perp}$
"Q"	$Q(x_+)$
"L"	$\Omega^{\varepsilon}(V)_{e_1,e_{-1}} \cong \Omega^{\varepsilon}(\langle e_1, e_{-1}\rangle^{\perp})$
$\lambda_L^{-1}: \Omega^{\varepsilon}(V_L) \to L$	$\Psi\vert_{\Omega^{\varepsilon}(V)_{e_1,e_{-1}}}: \Omega^{\varepsilon}(\langle e_1, e_{-1}\rangle^{\perp}) \to L$
$[\lambda^{\bullet}: QL \to "Q""L"]$	$[(\Psi\vert_{\Omega^{\varepsilon}(V)_{e_1}})^{-1}: QL \to \Omega^{\varepsilon}(V)_{e_1}]$

4.2. An algorithm for (SLP2)

As in [KS1,Br2], it is much more difficult to write SLPs within the black box group G, so that significantly different techniques are required for problem **(SLP2)**. The algorithm we give is a modification of that in [KS1, Proposition 4.23]. Indeed, most of the work in this subsection involves adapting various subroutines from [KS1] so that they can be used in our setting, often refining them in order to satisfy our more stringent timing requirements.

4.2.1. Linear algebra in Q

It now becomes essential to regard Q *algorithmically* both as an $\mathbb{F}_q$- and $\mathbb{F}_p$-space. (Recall that we were able to finesse this issue in the preprocessing phase; notably in Section 3.3.4.) It is possible to approach this problem in a manner similar to that in which the unitary case was handled in [Br2, Section 4.4.3]. However, the more complicated structure of orthogonal groups and, ironically, the fact that Q is abelian here, renders this method rather involved. In the interests of exposition, we prefer to adapt routines from [KS1] that handle such computations. To do so we will need analogues of the following constructions.

(a) A subgroup $L \cong \Omega^{\varepsilon}(d - 2, q)$ of G, that is a complement to Q in $N_G(Q)'$ (cf. [KS1, Corollary 4.18]).
(b) An effective isomorphism $\lambda_L: L \to \Omega^{\varepsilon}(V_L)$, where V_L is an $\mathbb{F}_q$-space of dimension $d - 2$ (unfortunately this was obtained recursively in [KS1, 4.3.3]).
(c) The *existence* of an isomorphism $\lambda^{\bullet}: QL \to "Q""L"$ extending λ_L, for suitable matrix groups "Q" and "L" (cf. [KS1, Section 4.3.3] again).

Table 1 serves as a "dictionary," enabling the reader to pass back and forth between [KS1] and the present paper. We constructed generators for Q and a suitable complement L in Section 3.1.4. Although we do not yet have an isomorphism $\lambda_L: L \to \Omega^{\varepsilon}(d - 2, q)$ which is effective in *both* directions, in view of **(SLP1)** the restriction of Ψ to the subgroup $\Omega^{\varepsilon}(V)_{e_1,e_{-1}}$ is effective in one direction; fortunately this is all we will need. Finally, the restriction of Ψ to $\Omega^{\varepsilon}(V)_{e_1}$ confirms the existence of $\lambda^{\bullet}$ in (c). (The map $\lambda^{\bullet}$ and its analogue here are listed in square brackets in Table 1 because we only require their *existence* to have been established: we have not yet constructed an effective analogue of $\lambda^{\bullet}$.)

Lemma 4.1. *In deterministic* $O(\mu d^2 \log q)$*-time, a generating set*

$$B_p := \{t_{ij} \mid 1 \leqslant i \leqslant d - 2,\ 1 \leqslant j \leqslant k\} \tag{4.2}$$

for Q can be found such that the following hold:

(i) *For $1 \leqslant i \leqslant d - 2$, if $A_i := \langle t_{ij} \mid 1 \leqslant j \leqslant k \rangle$, then $|A_i| = q$ and $Q = A_1 \times \cdots \times A_{d-2}$.*
(ii) *In deterministic $O(\mu d)$-time, any given $u \in Q$ can be expressed in the form $u = \prod_{i=1}^{d-2} a_i$ with $a_i \in A_i$.*
(iii) *In deterministic $O(\chi d)$-time, any given $u \in Q$ can be expressed in the form $u = \prod_{i=1}^{d-2} \prod_{j=1}^{k} t_{ij}^{n_{ij}}$, where $0 \leqslant n_{ij} < p$. (We say that this expresses u "as an $\mathbb{F}_p$-vector relative to B_p.")*
(iv) *In deterministic $O(\chi d)$-time, one can write an SLP of length $O(d \log q)$ from B_p to any given $u \in Q$.*

Proof. The groups A_i in (i) are precisely those constructed in [KS1, Lemma 4.19], where (ii) was also proved. We now describe how the specific generating sets $\{t_{ij} \mid 1 \leqslant j \leqslant k\}$ for the A_i are constructed. We consider two cases, and refer to the corresponding cases in the proof of [KS1, Lemma 4.19] for the notation $W_1, W_2, \ldots$ that we use.

Case $p > 2$. Choose the 1-space W_1 so that it lies in the 4-space $\langle e_2, e_3, e_{-2}, e_{-3} \rangle = V_J \cap \langle e_1, e_{-1} \rangle^{\perp}$ (recall that $\langle e_1, e_{-1} \rangle^{\perp}$ plays the role of V_L). Find a matrix $c'_{d-2} \in \Omega^{\varepsilon}(\langle e_1, e_{-1} \rangle^{\perp})$ moving W_{d-2} to a 1-space Y_{d-2} within $\langle e_2, e_3, e_{-2}, e_{-3} \rangle$. Use **(SLP1)** to construct $c_{d-2} := c'_{d-2} \Psi \in L$. Fix vectors $v_1 \in W_1$ and $y \in Y_{d-2}$, and, for $1 \leqslant j \leqslant k$, use **(SLP1)** again to find each of the following elements of J:

$$t_{1j} := r_1(\rho^{j-1} v_1) \Psi \in A_1 \quad \text{and} \quad s_{d-2,j} := r_1(\rho^{j-1} y) \Psi.$$

Put $t_{d-2,j} := s_{d-2,j}^{c_{d-2}^{-1}} \in A_{d-2}$ and $v_{d-2} := y^{c'_{d-2}}$. Finally, for $2 \leqslant i \leqslant d - 3$, put $t_{ij} := t_{1j}^{c^{i-1}} \in A_j$, and $v_i := v_1^{c'^{i-1}}$ where $c \in L$ and c' are as defined in the proof of [KS1, Lemma 4.19].

Case $p = 2$. Write $d = 2n$ (in the present paper the letter "m" used in [KS1, Lemma 4.19] has been reserved for the Witt index of V). This time W_1 is a nondegenerate 2-space, which we choose so that it lies in the 4-space $\langle e_2, e_3, e_{-2}, e_{-3} \rangle$. Find a matrix $c'_{n-1} \in \Omega^{\varepsilon}(\langle e_1, e_{-1} \rangle^{\perp})$ moving the nondegenerate 2-space W_{n-1} to a 2-space Y_{n-1} within $\langle e_2, e_3, e_{-2}, e_{-3} \rangle$. Use **(SLP1)** to construct $c_{n-1} := c'_{n-1} \Psi \in L$. Relabel the groups constructed in [KS1, Lemma 4.19]: for $1 \leqslant l \leqslant n - 1$, put $A_{2l-1} := A_{lx}$, and $A_{2l} := A_{ly}$. Fix vectors $v_1, v_2 \in W_1$, $y_1, y_2 \in Y_{n-1}$, and, for $1 \leqslant j \leqslant k$, use **(SLP1)** to find each of the following elements of J:

$$t_{1j} := r_1(\rho^{j-1} v_1) \Psi \in A_1, \qquad t_{2j} := r_1(\rho^{j-1} v_2) \Psi \in A_2,$$
$$s_{d-3,j} := r_1(\rho^{j-1} y_1) \Psi, \qquad s_{d-2,j} := r_1(\rho^{j-1} y_2) \Psi.$$

Put $t_{d-3,j} := s_{d-3,j}^{c_{n-1}} \in A_{d-3}$, $t_{d-2,j} := s_{d-2,j}^{c_{n-1}} \in A_{d-2}$, $v_{d-3} := y_1^{c'_{n-1}}$ and $v_{d-2} := y_2^{c'_{n-1}}$. Finally, for $2 \leqslant l \leqslant n - 2$, put $t_{2l-1,j} := t_{1j}^{c^{l-1}} \in A_{2l-1}$, $t_{2l,j} := t_{2j}^{c^{l-1}} \in A_{2l}$, $v_{2l-1} := v_1^{c'^{l-1}}$ and $v_{2l} := v_2^{c'^{l-1}}$, where, once again, $c \in L$ and c' are as in [KS1, Lemma 4.19].

We can now prove (iii) of the lemma when $p > 2$, the case $p = 2$ being very similar. Let u be the given element of Q. Use (ii) to find elements $a_i \in A_i$ ($1 \leqslant i \leqslant d - 2$) such that $u = a_1 \ldots a_{d-2}$. For each $1 \leqslant i < d - 2$, use Lemma 2.16 (with Ψ_J) to write $a_i^{c^{1-i}} \in J$ as $a_i^{c^{1-i}} \prod_{j=1}^{k} t_{1j}^{n_{ij}}$. Then, for $1 \leqslant i < d - 2$,

$$a_i = \left(a_i^{c^{1-i}} \right)^{c^{i-1}} = \left(\prod_{j=1}^{k} t_{1j}^{n_{ij}} \right)^{c^{i-1}} = \prod_{j=1}^{k} (t_{1j}^{c^{i-1}})^{n_{ij}} = \prod_{j=1}^{k} t_{ij}^{n_{ij}},$$

as claimed. An $\mathbb{F}_p$-vector for a_{d-2} is computed similarly.

As presented above, the algorithm computes $O(d)$ preimages of elements of J using Lemma 2.16. However, all such computations are performed with elements that lie inside 1-dimensional subspaces of Q_J; hence each occurs within some $\Omega^+(4, q)$ subgroup of J. In view of the isomorphism $\Omega^+(4, q) \cong \mathrm{SL}(2, q) \circ \mathrm{SL}(2, q)$, Ψ_J can be restricted to a suitable $\Omega^+(4, q)$ subgroup, and then each $\mathbb{F}_p$-vector may be found using the $\mathrm{SL}(2, q)$-oracle in $O(\chi)$-time. The stated timing is that required for $O(d)$ such uses of the $\mathrm{SL}(2, q)$-oracle.

Now (iv) follows immediately from (iii). $\quad\square$

Parts (iii) and (iv) of the previous lemma permit us to compute effectively within Q both as an $\mathbb{F}_p$-space and an abstract group, respectively; our next lemma gives the algorithms necessary for computing with Q as $\mathbb{F}_q$-space.

Remark 4.3. In the above proof we constructed a basis $v_1, \ldots, v_{d-2}$ of $\langle x_+, x_- \rangle^{\perp}$ such that

$$r_1(v_i)\Psi = b_i := t_{i1} \quad \text{for } 1 \leqslant i \leqslant d - 2.$$

If we set

$$B := \{ b_i \mid 1 \leqslant i \leqslant d - 2 \}, \tag{4.4}$$

then the expression "write $u \in Q$ as an $\mathbb{F}_q$-vector relative to B" will mean "find $(\lambda_1, \ldots, \lambda_{d-2}) \in \mathbb{F}_q^{d-2}$ such that $u\Psi^{-1} = \prod_{i=1}^{d-2} r_1(\lambda_i v_i)$."

Lemma 4.5. *There are deterministic algorithms for each of the following.*

(i) *Given any $u \in Q$, write u as an $\mathbb{F}_q$-vector relative to B in $O(\chi d)$-time.*

(ii) *Given $g \in N_G(Q)$, find the $(d-2) \times (d-2)$ matrix $\tilde{g}$ representing the linear transformation induced by g on the $\mathbb{F}_q$-space Q in $O(\chi d^2)$-time.*

(iii) *Given any $u \in Q$, find $\phi_Q(u)$ in $O(\chi d)$-time, where ϕ_Q is the L-invariant quadratic form on Q represented by the matrix A_Q, constructed in Procedure 3.21.*

Proof. (i) Use Lemma 4.1(iii) to write u as an $\mathbb{F}_p$-vector relative to B_p using some $n_{ij} \in \mathbb{F}_p$. For $1 \leqslant i \leqslant d - 2$, put $\lambda_i := \sum_{j=1}^{k} n_{ij} \rho^{j-1}$. Return $(\lambda_1, \ldots, \lambda_{d-2}) \in \mathbb{F}_q^{d-2}$.

To see that this output is correct, note that each $t_{ij} \in B_p$ was constructed in the proof of Lemma 4.1 either as the image under Ψ_J of an element $r_1(\rho^{j-1}w)$ for some $w \in \langle e_1, e_{-1}\rangle^\perp$, or else as an explicit L-conjugate of such an element. Hence, just as in (3.26), if $b_i\Psi^{-1} = t_{i1}\Psi^{-1} = r_1(v_i)$ for fixed $1 \leqslant i \leqslant d-2$, then $t_{ij}\Psi^{-1} = r_1(\rho^{j-1}v_i)$ for $1 \leqslant j \leqslant k$.

(ii) This follows immediately from (i) (cf. [KS1, 4.4.3(5)]).

(iii) Recall that ϕ_Q can be evaluated via the $\phi_Q(u) = \phi_{x_+}(u\Psi^{-1})$ (cf. (2.10)). It follows that $\phi_Q(u)$ can easily be computed from the vector returned in (i) since we know all of the scalars $\phi(v_i)$ and (v_i, v_j). $\square$

4.2.2. Exploiting geometry

We have seen in Section 2.1.1 that each singular point x in an orthogonal module determines a subgroup $Q(x)$ of the corresponding orthogonal group. Thus the conjugates of the subgroup Q provide us with models of singular points within the given black box group G. The ability to distinguish between given points, determine whether two points are "perpendicular" and, under certain geometric constraints, conjugate one point to another, are fundamental tasks that are crucial to our algorithm.

Lemma 4.6. *There is an $O(\mu d)$-time deterministic algorithm, that decides whether or not two given "points" are equal.*

Proof. See [KS1, p. 72] (noting that the timing of this algorithm is $O(\mu d)$ rather than $O(\mu d \log d)$, since our group Q was generated using only $O(d)$ long root groups). $\square$

Although perpendicularity testing and conjugacy are also dealt with in [KS1, Section 4.3.2], we need new algorithms to bring these problems into polynomial time. Our methods will make essential use of the $SL(2, q)$-oracle, often giving rise to randomised (rather than deterministic) algorithms. We begin by describing a special use for the $SL(2, q)$-oracle.

Let $H = \langle S_H \rangle$ be a given subgroup of G such that $O_p(H)$ is a class 2 nilpotent group and $H/O_p(H)$ is isomorphic to $SL(2, q)$. Then there is a deterministic membership test for $O_p(H)$: for $a \in H$, we have $a \in O_p(H)$ if and only if $[a, a^t]$ commutes with a for all $t \in S_H$. The membership test requires at most $6|S_H|$ group operations in H. The next observation follows from the discussion in Section 2.2 concerning the use of the $SL(2, q)$-oracle with quotient groups.

(p-*Core*) *For any given black box group $H = \langle S_H \rangle$ such that $O_p(H)$ is a class 2 nilpotent group and $H/O_p(H) \cong SL(2, q)$, in $O(\chi)$-time one can use the $SL(2, q)$-oracle to construct an effective isomorphism $SL(2, q) \to H/O_p(H)$. Moreover, once such an effective isomorphism has been constructed, the oracle may be used to construct (nearly) uniformly distributed random elements of $O_p(H)$ in time $O(\chi)$ per element.*

Remark 4.7. In practice (for example, in the context of the Matrix Group Project) the $SL(2, q)$-oracle uses a vector space of characteristic p upon which the group $SL(2, q)$ acts;

see [CLO]. Assuming that $G < \mathrm{GL}(W)$ for a vector space W of characteristic p, a suitable module for $H/O_p(H)$ in application (**p-Core**) is obtained as follows: Find a section W_0 of W upon which H acts nontrivially and irreducibly; then $O_p(H)$ acts trivially on W_0, so that H acts as $\mathrm{SL}(2, q)$ on W_0 (cf. [KS3]).

We are now ready for our analogue of [KS1, Lemma 4.15]. We begin with a procedure that takes as input generators for two (distinct) groups Q^{g_1} and Q^{g_2} and decides whether or not they are perpendicular.

Procedure 4.8 *(perpendicularity test).*

1. Find a long root subgroup R_1 of Q^{g_1} that does not normalise Q^{g_2}.
2. Selecting one generator, r, from each of $d-2$ long root groups generating Q^{g_2}, use the $\mathrm{SL}(2, q)$-oracle to test whether $\langle R_1, r \rangle \cong \mathrm{SL}(2, q)$. When some choice r produces such a subgroup, use the $\mathrm{SL}(2, q)$-oracle to find an element $t \in \langle R_1, r \rangle$ such that $r \in R_1^t$.
3. Pick a generator u of Q^{g_1} not in R_1. Test whether $[u^t, Q^{g_2}] = 1$ by selecting one generator from each long root group generating Q^{g_2}. If so, report that Q^{g_1} and Q^{g_2} are perpendicular; else report that they are not perpendicular.

Lemma 4.9. *Procedure* 4.8 *is an* $O(\chi d + \mu d^2)$-*time deterministic algorithm, that decides whether or not two given distinct "points"* Q^{g_1} *and* Q^{g_2} *are perpendicular.*

Proof. As in the proofs of Lemmas 3.10 and 3.22, let V' denote the orthogonal space underlying G. For $i = 1, 2$, let x_i' denote the singular point of V' corresponding to Q^{g_i}, and let $\Sigma_1' = [V', R_1]$ be the t.s. line (containing x_1') corresponding to R_1. Then $x_2' \notin \Sigma_1'^{\perp}$ since R_1 does not normalise Q^{g_2}. One of the long root elements r generating Q^{g_2} satisfies $\Sigma_2' \cap \Sigma_1'^{\perp} = 0$, where $\Sigma_2' = [V, r]$. (For, if not, then each t.s. line corresponding to a generating long root group of Q^{g_2} contains a point in $\Sigma_1'^{\perp}$. Those points then span a subspace of the degenerate $(d-2)$-space $\Sigma_1'^{\perp}$, whereas they also project onto a spanning subset of the nondegenerate $(d-2)$-space $x_2'^{\perp}/x_2'$.)

Then $D := \langle R_1, r \rangle \cong \mathrm{SL}(2, q)$ for such an element r (cf. [KS1, 4.1.2(v)]). Here $[V', D]$ is a nondegenerate 4-space of Witt index 2; D fixes each of $q + 1$ t.s. lines and acts transitively on the remaining $q + 1$ t.s. lines of this 4-space, with Σ_1' and Σ_2' distinct members of the latter collection. An element $t \in D$ conjugating R_1 to the root group (in D) containing r moves Σ_1' to Σ_2'. Thus, if x_1' and x_2' are perpendicular, then t fixes the t.s. line $\Sigma = \langle x_1'^t, x_2' \rangle$ and hence sends $x_1' = \Sigma \cap \Sigma_1'$ to $x_2' = \Sigma \cap \Sigma_2'$; while if these points are not perpendicular then $x_1'^t$ is point of Σ_2' distinct from x_2'. The commutator test in step 3 distinguishes between those two possibilities.

Timing. A suitable R_1 not normalising Q^{g_2} is found in $O(\mu d^2)$-time by testing the equality of Q^{g_2} and $Q^{g_2 v}$ for one generator v from each generating long root group of Q^{g_1}; R_1 is the long root group containing such a v. Finding a generator r such that $\langle R_1, r \rangle \cong \mathrm{SL}(2, q)$ requires at most $d-2$ calls to the $\mathrm{SL}(2, q)$-oracle. $\quad \square$

The next procedure finds generators for the long root group $Q^{g_1} \cap Q^{g_2}$ when Q^{g_1} and Q^{g_2} are distinct and perpendicular. Recall that the element h in $\mathcal{S}^*$, introduced in 3.1.3(iii), acts on Q as a scalar that generates $\mathbb{F}_q$ as $\mathbb{F}_p$-space (by (2.7), its preimage h' in $\Omega^+(V_J)$ has this property with respect to $Q(x_+)$). In particular, if r is any root element of Q, then the normal closure $\langle r \rangle^{\langle h \rangle}$ is the root group containing r.

Procedure 4.10 *(intersection).*

1. Put $H := \langle Q^{g_1}, Q^{g_2} \rangle$.
2. Use (*p*-**Core**) to recognise constructively $H/O_p(H) \cong \mathrm{SL}(2, q)$, and to produce two random elements $u_1, u_2 \in O_p(H)$.
3. Put $z := [u_1, u_2]$: if $z \neq 1$, put $h_1 := h^{g_1}$ and return $\langle z^{h_1^i} \mid 0 \leqslant i < k \rangle$; else report `failure` and stop.

Lemma 4.11. *Procedure* 4.10 *is an* $O(\chi)$*-time Las Vegas algorithm that returns* $Q^{g_1} \cap Q^{g_2}$ *with probability* > 0.49.

Proof. It follows from Lemma 2.9(b) that $H = U \rtimes D$ with $U = O_p(H)$ of order q^{2d-7} and $U' = Z(U)$ the desired intersection $Q^{g_1} \cap Q^{g_2}$ (here $D \cong \mathrm{SL}(2, q)$). The correctness and stated reliability now follow directly from Lemma 2.9. (Note that if $z \neq 1$, then $\langle z \rangle^{\langle h_1 \rangle} = \langle z^{h_1^i} \mid 0 \leqslant i < k \rangle$ is the root group containing z.) $\square$

The following is a simple modification of [KS1, Lemma 4.16], bringing it into polynomial time.

Lemma 4.12. *In deterministic* $O(\mu d + \chi)$*-time, given a long root element* u *not normalising* Q^g*, a point* Q^w *can be found containing* u *and perpendicular to* Q^g.

Proof. Proceed exactly as in the proof of [KS1, Lemma 4.16]: test at most one element from each member of a generating set of long root groups for Q^g to find $a \in Q^g$ such that $[[a, u], u] \neq 1$; and let A denote the long root group containing a. Rather than list A, we instead use the fact that $\langle A, u \rangle \cong \mathrm{SL}(2, q)$, together with the $\mathrm{SL}(2, q)$-oracle, to find $b \in A^u$ such that $[u, a^b] = 1$. Now return $Q^w := Q^{gb}$. $\square$

If x is a singular point of an orthogonal space V, then the group $Q(x) < \Omega^\varepsilon(V)$ acts regularly on the set of singular points not perpendicular to x. The final procedure of this subsection is an algorithmic analogue of this transitivity within our black box group G. It takes as input generators for Q, and elements $g_1, g_2 \in G$ such that neither Q^{g_1} nor Q^{g_2} is perpendicular to Q, and outputs the unique element of Q conjugating Q^{g_1} to Q^{g_2} (cf. [KS1, Lemma 4.17]).

Procedure 4.13 *(transitivity of Q).* First test whether $Q^{g_1} = Q^{g_2}$ and, if so, return $u := 1 \in Q$. If Q^{g_1} and Q^{g_2} are distinct points, run Procedure 4.8 to decide whether or not Q^{g_1} and Q^{g_2} are perpendicular and go to the appropriate case below. (Recall that Pro-

cedure 4.8 constructs long root groups $R_i < Q^{g_i}$ with $\langle R_1, R_2 \rangle \cong \mathrm{SL}(2, q)$.) *All calls to Procedure 4.10 should be repeated up to 3 times.*

- **(Q^{g_1} and Q^{g_2} are not perpendicular)** Proceed exactly as in the proof of [KS1, Lemma 4.17] to reduce to the perpendicular case. (Note that Procedure 4.10 should be used in place of [KS1, Lemma 4.15(ii)], and that Lemma 4.12 should be used in place of [KS1, Lemma 4.16].)
- **(Q^{g_1} and Q^{g_2} are perpendicular)** Use Procedure 4.10 to find $Z := Q^{g_1} \cap Q^{g_2}$, and fix a generator $1 \neq z \in Z$. Use Procedure 4.10 again to find $Y := Q \cap Q^z$, and put $K := \langle R_1, R_2, Y \rangle$. Use (***p*-Core**) to recognise constructively $K/O_p(K) \cong \mathrm{SL}(2, q)$, and then to find $u \in Y$ conjugating R_1 to R_2 mod $O_p(K)$. Return u.

Lemma 4.14. *Procedure* 4.13 *is an* $O(\chi d + \mu d^2)$-*time Las Vegas algorithm that finds the unique* $u \in Q$ *such that* $Q^{g_1 u} = Q^{g_2}$ *with probability at least* 1/2.

Proof. Let x', x'_1, x'_2 denote the singular points of V' corresponding to Q, Q^{g_1}, Q^{g_2} respectively. The reduction to the perpendicular case is proved in [KS1, Lemma 4.17]. We may therefore assume that Q^{g_1} and Q^{g_2} are perpendicular, and hence that $\langle x'_1, x'_2 \rangle$ is totally singular. Then $z \in Z = Q^{g_1} \cap Q^{g_2}$ moves x' to a point of the 3-space $\langle x', x'_1, x'_2 \rangle$ perpendicular to x', and $w' = \langle x', x'^z \rangle \cap \langle x'_1, x'_2 \rangle$ is the radical of this 3-space. Now $Y = Q \cap Q^z = R(\langle x', w' \rangle)$, R_1 and R_2 induce three distinct transvection groups on $\langle x'_1, x'_2 \rangle$, fixing δ', x'_1 and x'_2 respectively. Hence there exists a unique $u \in Y$ moving R_1 to R_2 modulo $O_p(K)$; such u evidently moves x'_1 to x'_2, as desired.

There are at most three calls to Procedure 4.10, each of which is repeated up to three times in order to ensure success with probability at least $1 - (0.51)^3 > 0.867$. Hence the procedure finds the unique $u \in Q$ with probability at least $0.867^3 > 1/2$. The timing is dominated by that of Procedure 4.8. $\quad\square$

4.2.3. Straight-line programs from $\mathcal{S}^*$

We are finally in position to give our algorithm for **(SLP2)**. In fact, we merely give a commentary on the proof of [KS1, Proposition 4.23], indicating where subroutines given here are substituted for ones in [KS1], and giving a correspondingly revised timing estimate.

We are given $g \in G$. Find $y \in \{1, l\} \cup B^l$ such that Q^{gy} and Q are not perpendicular. (Our element l and generating set B for Q are suitable analogues of $j(\gamma)$ and $\mathcal{S}^*(2)$ appearing in [KS1, 4.2.2(iii)] and [KS1, 4.3.3], respectively.) Using Procedure 4.13 instead of [KS1, Lemma 4.17], find $u, v \in Q$ such that $gyul^{-1}v$ normalises both Q and Q^l, and replace g by $gyul^{-1}v$. (Repeat the randomised Procedure 4.13 twice, if necessary, for each of u and v.)

We have now reduced to the case where g normalises Q and Q^l. Fix a nonsingular vector $w \in Q$ and use Lemma 4.5(iii) to compute $\phi_Q(w)$, $\phi_Q(w^g)$ and $\phi_Q(w^{h^+})$, where $h^+ \in \mathcal{S}^*$ is the element constructed in 3.1.3(iv). Then, as in [KS1, Lemma 4.6(a)], $\phi_Q(w^g) = \lambda^2 \phi_Q(w)$ and $\phi_Q(w^{h^+}) = \zeta^2 \phi_Q(w)$ for nonzero scalars λ, ζ with ζ of order $q - 1$. Use $\mathrm{DLog}(\mathbb{F}_q^*)$ to find i such that $(\zeta^2)^i = \lambda^2$, and replace g by gh^{-i}.

Now g induces an isometry of Q. Use Lemma 4.5(ii) to find the matrix $\tilde{g}$ representing this isometry. If necessary, use Wall forms and the element h^+ once again to modify g so that $\tilde{g} \in \Omega^\varepsilon(Q)$.

Finally, write the unique matrix $g' \in \Omega^\varepsilon(V)_{e_1, e_{-1}}$ whose restriction to $\langle e_1, e_{-1} \rangle^\perp$ is $\tilde{g}$. Use **(SLP1)** to write an SLP of length $O(d^2 \log q)$ from $\mathcal{T}$ to g', and evaluate this SLP from $\mathcal{S}^*$ to obtain an element $g_0 \in G$ such that $g g_0^{-1} \in Z(G)$. In the cases where G is not simple, $\mathcal{S}^*$ contains a generator for $Z(G)$; hence the SLP is easily modified to give one to g, as required.

Timing.   The timing is dominated by the call to Lemma 4.5(ii) and the evaluation of an SLP of length $O(d^2 \log q)$ inside G. Hence, we obtain the desired SLP in $O(\chi d^2 + \mu d^2 \log q) = O(\chi d^2)$-time, as stated in Theorem 1.1.

Reliability.   Randomisation occurs only in finding the elements $u, v \in Q$ using Procedure 4.13. Repeating each such call ensures success with probability at least $(3/4) \cdot (3/4) > 1/2$.

5. Concluding remarks

5.1. Verifying a presentation

Our assumption in Theorem 1.1 was that the given black box group $G = \langle \mathcal{S} \rangle$ is *known* to be a homomorphic image of $\Omega^\varepsilon(d, q)$. Our constructive recognition algorithm is a Las Vegas algorithm only under this assumption.

In practice, however, it is unlikely that a user will know *with certainty* that a given $G = \langle \mathcal{S} \rangle$ is the image of a certain orthogonal group. There are two methods in print to determine that G is *probably* a homomorphic image of $\Omega^\varepsilon(d, q)$ for known ε, d and q ([KS2,BKPS]; compare [KS1, Section 7.2.1]). We make no claim as to what our algorithm will output if G is not what we think it is (although it can safely be stated that it will almost always fail). In such a situation, in order to be certain that a positive output Ψ from our procedure really is an epimorphism, one must construct and verify a suitable presentation for G. (There are additional reasons why a presentation for G is desirable; see [L-G].)

Thus, we will *assume* that G is indeed an epimorphic image of $\Omega^\varepsilon(V)$. Let $\Psi : \Omega^\varepsilon(V) \to G$ be an *alleged* epimorphism returned by the preprocessing phase of our algorithm. A presentation for G can be constructed and verified as follows (cf. [KS1, Section 7.2.1]).

Use [BGKLP] to write a "short" presentation $\langle X \mid R \rangle$ of $\Omega^\varepsilon(d, q)$ in $O(d^4 \log^2 q)$ time: a map $\varphi : X \to \Omega^\varepsilon(d, q)$ such that, if $F(X)$ is the free group with X and if $N = \langle R^{F(X)} \rangle$, then φ induces an isomorphism $\hat{\varphi} : F(X)/N \to \Omega^\varepsilon(d, q)$. Use **(SLP1)** to find $\mathcal{S}^{**} := X\varphi\Psi \subseteq G$. For each word $w(x_1, \dots) \in R$ (where $x_1, \dots \in X$), find $w(x_1\varphi\Psi, \dots)$ and test whether this is 1 in G; if so, then $\langle \mathcal{S}^{**} \rangle$ is a homomorphic image of $\Omega^\varepsilon(V)$, otherwise it is not a homomorphic image.

Finally, test that $G = \langle \mathcal{S}^{**} \rangle$ by verifying that $\mathcal{S} \subseteq \langle \mathcal{S}^{**} \rangle$ as follows. Use **(SLP2)** to find $\mathcal{S}\Psi^{-1} \subseteq \Omega^\varepsilon(V)$. Use **(SLP1)** to write short SLPs from $X\varphi$ to $\mathcal{S}\Psi^{-1}$, and evaluate these

SLPs from $X\varphi\Psi = S^{**}$ in order to (try to) obtain $\mathcal{S}$. (Of course, if $\mathcal{S} \subseteq \langle \mathcal{S}^{**} \rangle$ is false then G is not a homomorphic image of $\Omega^{\varepsilon}(V)$.)

The timing for such a verification is dominated by the construction of $\mathcal{S}^{**}$, which is obtained in $O(d^2 \log q \cdot \mu d^2 \log q) = O(\mu d^4 \log^2 q)$-time.

5.2. Other perfect central extensions of $\mathrm{P}\Omega^{\varepsilon}(V)$

We have focused on epimorphic images of $\Omega^{\varepsilon}(V)$. However, our algorithm applies with almost no change to homomorphic images of the corresponding spin group, or, more precisely, to all perfect central extensions G of $\mathrm{P}\Omega^{\varepsilon}(V)$. The only slight change has to do with finding $Z(G)$, but this is accomplished in almost the same manner as before.

5.3. Alternative algorithm for $\Omega^{-}(6, q)$

The case $\Omega^{-}(6, q)$ was excluded in Theorem 1.1. Recall that, for this group, we could simply switch to the equivalent case $\mathrm{SU}(4, q)$ and use [Br2]; but then we would need an additional discrete log oracle for cyclic groups of order $q + 1$. However, such an oracle can be sidestepped; essentially the same situation arose in [KS1, Section 4.6.3]. The algorithm presented there can be used here, replacing [KS1, Lemmas 4.15–4.17] by Procedure 4.8, Lemma 4.12 and Procedure 4.13, respectively. This allows us to obtain a data structure behaving as in Section 3. Then straight-line programs can be found as in Section 4.

Acknowledgment

References

[Ba] L. Babai, Local expansion of vertex-transitive graphs and random generation in finite groups, in: Proc. ACM Symp. on Theory of Computing, 1991, pp. 164–174.

[BGKLP] L. Babai, A.J. Goodman, W.M. Kantor, E.M. Luks, P.P. Pálfy, Short presentations for finite groups, J. Algebra 194 (1997) 79–112.

[BKPS] L. Babai, W.M. Kantor, P.P. Pálfy, Á. Seress, Black-box recognition of finite simple groups of Lie type by statistics of element orders, J. Group Theory 5 (2002) 383–401.

[Br1] P.A. Brooksbank, Constructive recognition of classical groups in their natural representation, J. Symbolic Comput. 35 (2003).

[Br2] P.A. Brooksbank, Fast constructive recognition of black box unitary groups, LMS J. Comput. Math. 6 (2003) 162–197.

[BK] P.A. Brooksbank, W.M. Kantor, On constructive recognition of a black box $\mathrm{PSL}(d, q)$, in: W.M. Kantor, Á. Seress (Eds.), Groups and Computation III, in: Ohio State Univ. Math. Res. Inst. Publ., vol. 8, de Gruyter, Berlin, 2001, pp. 95–111.

[CFL] G. Cooperman, L. Finkelstein, S. Linton, Recognizing $GL_n(2)$ in non-standard representation, in: L. Finkelstein, W.M. Kantor (Eds.), Groups and Computation II, Proceedings of a DIMACS Workshop, Amer. Math. Soc., 1997, pp. 85–100.

[CLO] M.D.E. Conder, C.R. Leedham-Green, E.A. O'Brien, Constructive recognition of PSL(2, q), Trans. Amer. Math. Soc. 358 (3) (2006) 1203–1221.

[KS1] W.M. Kantor, Á. Seress, Black box classical groups, Mem. Amer. Math. Soc. 149 (708) (2001).

[KS2] W.M. Kantor, Á. Seress, Prime power graphs for groups of Lie type, J. Algebra 247 (2002) 370–434.

[KS3] W.M. Kantor, Á. Seress, Computing with matrix groups, in: Groups, Combinatorics and Geometry, Durham, 2001, World Scientific, River Edge, NJ, 2003.

[KL] P. Kleidman, M. Liebeck, The Subgroup Structure of the Finite Classical Groups, London Math. Soc. Lecture Note Ser., vol. 129, Cambridge Univ. Press, 1990.

[L-G] C.R. Leedham-Green, The computational matrix group project, in: W.M. Kantor, Á. Seress (Eds.), Groups and Computation III, in: Ohio State Univ. Math. Res. Inst. Publ., vol. 8, de Gruyter, Berlin, 2001, pp. 229–247.

[NP] P.M. Neumann, C.E. Praeger, A recognition algorithm for special linear groups, Proc. London Math. Soc. (3) 65 (1992) 555–603.

[Se] Á. Seress, Permutation Group Algorithms, Cambridge Univ. Press, 2002.

[Ta] D.E. Taylor, The Geometry of the Classical Groups, Heldermann, Berlin, 1992.

[Zs] K. Zsigmondy, Zur Theorie der Potenzreste, Monatsh. Math. Phys. 3 (1892) 265–284.