CONSTRUCTIVE HOMOMORPHISMS FOR CLASSICAL GROUPS

SCOTT H. MURRAY AND COLVA M. RONEY-DOUGAL

Abstract. Let $\Omega \leq \mathrm{GL}(V)$ be a quasisimple classical group in its natural representation over a finite vector space V, and let $\Delta = \mathrm{N}_{\mathrm{GL}(V)}(\Omega)$. We construct the projection from Δ to Δ/Ω and provide fast, polynomial-time algorithms for computing the image of an element. Given a discrete logarithm oracle, we also represent Δ/Ω as a group with at most 3 generators and 6 relations. We then compute canonical representatives for the cosets of Ω. A key ingredient of our algorithms is a new, asymptotically fast method for constructing isometries between spaces with forms. Our results are useful for the matrix group recognition project, can be used to solve element conjugacy problems, and can improve algorithms to construct maximal subgroups.

1. Introduction

In this paper, we provide a variety of algorithms for classical groups. Fix a prime p and a power q of p, and let $u = 2$ for unitary groups and 1 otherwise. We consider groups $H \leq \mathrm{GL}_d(q^u)$ such that $\Omega \leq H \leq \Delta$, where Ω is a *quasisimple classical group* and $\Delta = \mathrm{N}_{\mathrm{GL}_d(q^u)}(\Omega)$ is the corresponding *conformal group* [KL90, Section 2.1]. Most of our algorithms are randomised Las Vegas in the sense of [Bab97]. We often need Las Vegas algorithms whose output is independent of the random choices made. In this case we call the output *canonical*.

The matrix group recognition project [LG01] seeks to compute efficiently composition series for matrix groups over finite fields. By finding a geometry preserved by the group, in the sense of Aschbacher's theorem [Asc84], a normal subgroup and its quotient can often be computed. This decomposition terminates on reaching groups that are almost simple, modulo their subgroup of scalar matrices. These groups are either classical groups in their natural representation (Aschbacher's class 8) or other almost simple groups (class 9). This paper provides algorithms for dealing with a group known to be in class 8. Algorithms to constructively recognise the quasisimple classical groups in their natural representation are known [Bro01, Bro03]. This paper presents efficient, practical reduction algorithms for the other class 8 groups.

Another motivation is constructing efficient algorithms for element conjugacy in classical groups H, when the dimension d is large. The fundamental problem is to determine if two elements are conjugate and, if so, provide a conjugating element. For the sake of memory efficiency, it makes sense to conjugate a single element to a canonical representative of its conjugacy class. Given a solution to this conjugacy problem for Δ [HM, Bri06], we can construct an algorithm to solve the element conjugacy algorithm in a group H between Ω and Δ, provided that we have *canonical* coset representatives for H/Ω. This, along with applications to the construction of maximal subgroups, are the primary motivations for the requirement that our algorithms give canonical solutions. See Section 4 for more details.

Date: September 10, 2010.

2000 *Mathematics Subject Classification.* Primary 20G40; 20H30, 20-04.

We would like to thank the anonymous referees, whose careful reading and helpful suggestions have significantly improved this paper. We also thank the Magma project at the University of Sydney, where some of the work was carried out. The second author would like to acknowledge the support of the Nuffield Foundation, and of EPSRC grant EP/C523229/1.

We give our timings in terms of elementary finite field operations: addition, negation, multiplication, and inversion. The number of field operations required by our algorithms is polynomial in d and $\log q$, except for some algorithms which require calls to a discrete logarithm oracle. We specify when this is the case, and count the number of calls to the oracle.

We consider multiplication of $d \times d$ matrices to take $O(d^\omega)$ field operations: for example, the standard method gives $\omega = 3$. For sufficiently large d (depending on the field size) Magma [BC07] uses the algorithm of [Str69] with $\omega = \log_2 7 + \epsilon$ for any $\epsilon > 0$: this gives a noticeable practical, as well as a theoretical, improvement.

A key algorithmic problem for classical groups is the construction of isometries between classical forms. We give a new method that is asymptotically faster than the method given in [HRD05].

Theorem 1.1. *Suppose we have two nondegenerate symplectic, unitary, or quadratic forms on the space $V = (\mathbb{F}_{q^u})^d$. We can determine if they are isometric, and find a canonical isometry between them, with a Las Vegas algorithm taking $O(d^\omega + d^2 \log^2 q)$ field operations.*

We now state our main theorem.

Theorem 1.2. *Let $\Omega \leq \mathrm{GL}_d(q^u)$ be a quasisimple classical group fixing a known classical form F, let $\Delta = \mathrm{N}_{\mathrm{GL}_d(q^u)}(\Omega)$, and let $G = \Delta/\Omega$.*

(1) *There is a deterministic algorithm which, on input F, constructs a finite presentation P_1 for G in $O(\log^2 q)$ field operations. There is a Las Vegas algorithm which constructs the image under the homomorphism $\Delta \to P_1$ of $g \in \Delta$ in $O(d^\omega + d^2 \log^2 q)$ field operations.*
(2) *There is a deterministic algorithm which, on input F, constructs a power-conjugate presentation P_2 for G with at most 3 generators and 6 relations in $O(\log^2 q)$ field operations. There is a Las Vegas algorithm which constructs the image under the homomorphism $\Delta \to P_2$ of $g \in \Delta$ in $O(d^\omega + d^2 \log^2 q)$ field operations, plus at most two calls to a discrete logarithm oracle for $\mathbb{F}_{q^2}$.*
(3) *There is a Las Vegas algorithm which, on input F and an element $g \in \Delta$, constructs a canonical representative of the coset Ωg in $O(d^\omega + d^2 \log^2 q)$ field operations.*

By the *type* of the form we mean one of: unitary, symplectic, orthogonal type +, orthgonal type −, orthogonal odd dimension. In Section 2 we define our canonical forms, and present algorithms for forms and classical groups, including proving Theorem 1.1. In Section 3 we prove Theorem 1.2. In Section 4 we present some applications, before concluding in Section 5 with some data on our implementations: our algorithms are now part of the standard release of Magma. The timings for our algorithms depend on the type of the form – in Theorems 1.1 and 1.2 we have given worst-case timings, but more detailed results are given below.

2. Groups and forms

In this section, we introduce some algorithms for classical forms and classical groups. We require that the output of each algorithm be *canonical*: for fixed input, every call to the algorithm gives the same output, even if the algorithm is randomised.

2.1. **Fields.** Let p be a prime and let q be a power of p. As is standard, we assume that $\mathbb{F}_q$ is constructed by adjoining a canonical root ξ of the Conway polynomial [JLPW95] to the prime field $\mathbb{F}_p$, so that ξ is the canonical primitive element of $\mathbb{F}_q$. See [Lüb] for a current list of the fields for which this assumption is valid. We let ζ be the canonical primitive element of $\mathbb{F}_{q^2}$, and recall that $\xi = \zeta^{q+1}$. Given a nonzero $\alpha \in \mathbb{F}_q$, the *discrete logarithm* $\log_\xi(\alpha)$ is the unique

$i = 0, 1, \ldots, q-2$ such that $\alpha = \xi^i$. We now show how to find canonical solutions to various equations over $\mathbb{F}_q$ or $\mathbb{F}_{q^2}$.

The next result is the main source of randomisation in our algorithms.

Theorem 2.1 ([GCL92, Theorem 8.12])**.** *A root in $\mathbb{F}_{q^2}$ for a quadratic polynomial with coefficients in $\mathbb{F}_q$ can be found by a Las Vegas algorithm in $O(\log q)$ field operations.*

Let $\mathbb{F}_q^\times$ denote the multiplicative group of $\mathbb{F}_q$ and let $\mathbb{F}_q^{\times 2}$ denote the set of squares in $\mathbb{F}_q^\times$. Every element of $\mathbb{F}_{q^2}$ can be written as $a_0 + a_1\zeta + \cdots + a_{m-1}\zeta^{m-1}$, where $p^m = q^2$ and $a_i \in \{0, \ldots, p-1\}$. Lexicographically ordering the coefficients induces an ordering on $\mathbb{F}_{q^2}$. We fix a *canonical root* of a quadratic equation by taking the smallest root with respect to this ordering on $\mathbb{F}_{q^2}$. Hence for $\alpha \in \mathbb{F}_q$ we can find a *canonical square root* $\sqrt{\alpha} \in \mathbb{F}_{q^2}$. For q even, α has a unique square root, equal to $\alpha^{q/2}$, so $\sqrt{\alpha}$ can be computed by a deterministic algorithm in $O(\log q)$ field operations. For $\alpha \in \mathbb{F}_q^\times$ with q odd, define $\iota(\alpha) = 0$ if $\alpha \in \mathbb{F}_q^{\times 2}$ and $\iota(\alpha) = 1$ otherwise. Since $\iota(\alpha) = 0$ if and only if $\alpha^{(q-1)/2} = 1$, there is a deterministic algorithm to determine $\iota(\alpha)$ which takes $O(\log q)$ field operations.

Canonical solutions for trace and norm equations are needed for the unitary groups.

Proposition 2.2. *Let $\alpha \in \mathbb{F}_q^\times$. There is a deterministic algorithm to find a canonical solution $\eta \in \mathbb{F}_{q^2}$ to the trace equation $\eta + \eta^q = \alpha$ which takes $O(1)$ field operations if q is odd, and $O(\log q)$ otherwise. There is a Las Vegas algorithm to find a canonical solution $\eta \in \mathbb{F}_{q^2}$ of the norm equation $\eta^{q+1} = \alpha$ which takes $O(\log q + \log^2 p)$ field operations.*

Proof. For the trace equation with q odd, $\eta = \alpha/2$. Otherwise, use the fact that $\alpha \mapsto \alpha^q$ is an $\mathbb{F}_q$-linear map. After we evaluate this map on an $\mathbb{F}_q$-basis of $\mathbb{F}_{q^2}$ deterministically in $O(\log q)$ field operations, the problem is reduced to two dimensional system of linear equations over $\mathbb{F}_q$. Since η exists by [Lan93, Theorem 6.3], it can now be found by linear algebra.

We construct a solution to the norm equation in three cases. If $\alpha \in \mathbb{F}_q^{\times 2}$, let $\eta := \sqrt{\alpha}$, then $\eta^{q+1} = \eta^2 = \alpha$. If $\alpha \notin \mathbb{F}_q^{\times 2}$ and $q \equiv 1 \pmod 4$, then $-1 \in \mathbb{F}_q^{\times 2}$, so $-\alpha \notin \mathbb{F}_q^{\times 2}$. Hence the polynomial $X^2 + \alpha$ is irreducible over $\mathbb{F}_q$, and its roots in $\mathbb{F}_{q^2}$ have norm α, which can be found by Theorem 2.1. If $\alpha \notin \mathbb{F}_q^{\times 2}$ and $q \equiv 3 \pmod 4$, then $-\alpha \in \mathbb{F}_q^{\times 2}$. Let $\beta = \sqrt{-\alpha}$ and write $p + 1 = 2^m s$ for s odd. Calculate $c \in \mathbb{F}_p$ in $O(\log^2 p)$ field operations by

$$c_1 := 0; \qquad c_{i+1} := \left(\frac{c_i + 1}{2}\right)^{\frac{p+1}{4}} \ (i = 1, \ldots, m-2); \qquad c := \left(\frac{c_{m-1} - 1}{2}\right)^{\frac{p+1}{4}}.$$

By [BGM93], the polynomial $g(X) = X^2 - 2cX - 1$ is irreducible over $\mathbb{F}_q$. Hence $-\alpha g(X/\beta) = X^2 - 2\beta c X + \alpha$ is also irreducible and its roots in $\mathbb{F}_{q^2}$ have norm α. □

The following elements are all used to compute with orthogonal groups.

Proposition 2.3.

1. *There is a deterministic algorithm to construct, on input an odd q, a canonical $\gamma \in \mathbb{F}_q^\times$ such that γ and $1 - 4\gamma$ are nonsquare. The algorithm takes $O(\log q)$ field operations.*
2. *There is a deterministic algorithm to construct, on input an even q, a canonical $\gamma \in \mathbb{F}_q^\times$ such that $X^2 + X + \gamma$ is irreducible over $\mathbb{F}_q$. The algorithm takes $O(\log^2 q)$ field operations.*
3. *There is a deterministic algorithm to construct, on input an odd q, a canonical $\nu \in \mathbb{F}_q^\times$ such that $1 + \nu^2$ is nonsquare. The algorithm takes $O(\log q)$ field operations.*

Proof. For (1), note that $\zeta + \zeta^q \neq 0$ (recall that ζ is the canonical primitive element in $\mathbb{F}_{q^2}$), as otherwise $\zeta^{q-1} = -1 = \zeta^{(q^2-1)/2}$. Set $\gamma = \xi(\zeta+\zeta^q)^{-2}$, then $\gamma \in \mathbb{F}_q$ because $\gamma^q = \gamma$. Also, $\gamma \notin \mathbb{F}_q^{\times 2}$ because $\xi \notin \mathbb{F}_q^{\times 2}$. Finally, $1 - 4\gamma = (\zeta - \zeta^q)^2(\zeta + \zeta^q)^{-2} \notin \mathbb{F}_q^{\times 2}$, since $(\zeta - \zeta^q)(\zeta + \zeta^q)^{-1} \notin \mathbb{F}_q$.

For (2), let $q = 2^m$. If m is odd, let $\gamma = 1$. Otherwise, let $m = 2^r s$ with s odd. Define a_i recursively: let $a_0 = 1$, and let a_{i+1} be the canonical root of $X^2 + X + a_i$ in $\mathbb{F}_q$. Define γ to be the first a_j for which $X^2 + X + a_j$ is irreducible, if any. Define $T : \mathbb{F}_q \to \mathbb{F}_q$ by $T(x) = x^2 + x$, and note that $T(a_i) = a_i^2 + a_i = a_{i-1}$ for $i \geq 1$. It is easy to show that $T^{2^i}(x) = x^{2^{2^i}} + x$ for all i. Now suppose $a = a_{2^r+1} \in \mathbb{F}_q$ exists. Then $T^{2^r+1}(a) = 1$, so $T^{2^{r+1}}(a) = T^{2^{r+1}-2^r-1}(1) = 0$, and so $a^{2^{2^{r+1}}} = a$. Hence $a \in \mathbb{F}_{2^{2^{r+1}}}$, which intersects $\mathbb{F}_q$ in $\mathbb{F}_{2^{2^r}}$. This implies that $a^{2^{2^r}} = a$, so $T^{2^r}(a) = 0$, which contradicts $T^{2^r+1}(a) = 1$. Therefore $j \leq 2^r \leq \log q$.

For (3), note that $4\zeta^{q+1}(\zeta - \zeta^q)^{-2} \in \mathbb{F}_q^{\times 2}$. Let $\nu = 2\zeta^{(q+1)/2}(\zeta - \zeta^q)^{-1} \in \mathbb{F}_q$ be its square root, then $1 + \nu^2 \notin \mathbb{F}_q^{\times 2}$. □

2.2. **Forms and Isometries.** In this subsection, we define our canonical forms, and present algorithms to construct isometries and similarities between forms.

Let $V = (\mathbb{F}_{q^u})^d$ and let $v_1, \ldots, v_d$ be the basis of V with $(v_i)_j = 1$ if $i = j$ and 0 otherwise. By $\mathrm{diag}(a_1, a_2, \ldots, a_d)$ we mean the $d \times d$ matrix with entry a_i in position (i, i) and 0 elsewhere. By $\mathrm{antidiag}(a_1, a_2, \ldots, a_d)$ we mean the $d \times d$ matrix with entry a_i in position $(i, d-i+1)$ and 0 elsewhere. By $A \oplus B$ we mean a block diagonal matrix, with blocks A and B along the main diagonal and 0 elsewhere. We denote the transpose of A by A^{Tr}.

The following results are standard and can be found in [BCS97, Chapter 16].

Theorem 2.4. *There are deterministic algorithms to find the row echelon form, the rank, the nullspace, or the determinant of a $d \times d$ matrix over $\mathbb{F}_q$. Each algorithm requires $O(d^\omega)$ field operations.*

We refer to [Tay92] or [Gro02] for basic terminology on classical forms. We fix the following notation: either β is a nondegenerate symplectic or unitary form over V; or Q is a nondegenerate quadratic form over V and β is its polar form, so that $2Q(v) = \beta(v, v)$. A vector v is *isotropic* if $\beta(v, v) = 0$ and *singular* if $Q(v) = 0$: note that if q is even and the form is quadratic then there can exist vectors that are isotropic but nonsingular. A vector is *anisotropic* if $Q(v) \neq 0$. The matrix of β is $F = (\beta(v_i, v_j))_{d \times d}$, and satisfies $\beta(u, v) = uFv^{\sigma\mathrm{Tr}}$, where σ is the field automorphism $x \mapsto x^q$ (nontrivial only in the unitary case). The matrix of Q is the upper triangular matrix $M = (m_{ij})_{d \times d}$ such that $Q(v) = vMv^{\mathrm{Tr}}$ for $v = (a_1, \ldots, a_d)$. If β is the polar form of Q, then $F = M + M^{\mathrm{Tr}}$ and F determines M if and only if q is odd. Forms β_1 and β_2 (or Q_1 and Q_2) are *isometric* if there exists an $A \in \mathrm{GL}_d(q^u)$ such that $\beta_1(u, v) = \beta_2(uA, vA)$ for all $u, v \in V$ (respectively, such that $Q_1(v) = Q_2(vA)$ for all $v \in V$). Forms β_1 and β_2 (or Q_1 and Q_2) are *similar* if there exists a $\lambda \in \mathbb{F}_{q^u}^\times$ such that β_1 is isometric to $\lambda\beta_2$ (respectively, such that Q_1 is isometric to λQ_2).

Definition 2.5 (Canonical classical forms). *We define the following canonical forms:*
Symplectic or even dimension unitary: *$d = 2m$ and V has basis $(e_1, \ldots, e_m, f_m, \ldots, f_1)$ with $\beta(e_i, e_j) = \beta(f_i, f_j) = 0$, $\beta(e_i, f_j) = \delta_{ij}$.*
Unitary, odd dimension: *$d = 2m + 1$ and V has basis $(e_1, \ldots, e_m, x, f_m, \ldots, f_1)$ with $\beta(e_i, e_j) = \beta(f_i, f_j) = \beta(e_i, x) = \beta(f_i, x) = 0$, $\beta(e_i, f_j) = \delta_{ij}$, $\beta(x, x) = 1$.*
Orthogonal, ∘ type: *$d = 2m + 1$ and V has basis $(e_1, \ldots, e_m, x, f_m, \ldots, f_1)$ with $Q^\circ(e_i) = Q^\circ(f_i) = \beta^\circ(e_i, e_j) = \beta^\circ(f_i, f_j) = \beta^\circ(e_i, x) = \beta^\circ(f_i, x) = 0$, $\beta^\circ(e_i, f_j) = \delta_{ij}$, $Q(x) = 1$.*
Orthogonal, + type: *$d = 2m$ and V has basis $(e_1, \ldots, e_m, f_m, \ldots, f_1)$ with $Q^+(e_i) = Q^+(f_j) =$*

$\beta^+(e_i,e_j) = \beta^+(f_i,f_j) = 0$ and $\beta^+(e_i,f_j) = \delta_{ij}$.
Orthogonal, − type: *$d = 2m+2$ and V has basis $(e_1,\dots,e_m,x,y,f_m,\dots,f_1)$ with $Q^-(e_i) = Q^-(f_j) = \beta^-(e_i,e_j) = \beta^-(f_i,f_j) = 0$, $\beta^-(e_i,f_j) = \delta_{ij}$, $\beta^-(a,b) = 0$ for $a \in \{e_i,f_j\}$, $b \in \{x,y\}$, $Q^-(x) = \beta^-(x,y) = 1$, $Q^-(y) = \gamma$, where γ is as in Proposition 2.3.*

It is well known (see for instance [Tay92]) that every nondegenerate quadratic, symplectic or unitary form over a finite field is similar to exactly one of the forms given in Definition 2.5. For odd dimension and characteristic, the two isometry classes of quadratic forms are similar. Otherwise, forms are similar if and only if they are isometric. The *discriminant* of Q is $\iota(\det(F))$. Two quadratic forms are isometric if and only if they have the same discriminant.

The following will be needed for constructing isometries and coset representatives. Unitary forms have an anisotropic vector whenever they are not identically zero, and quadratic forms have a nonsingular vector whenever they are not identically zero. However, symmetric forms may not have an anisotropic vector in even characteristic.

Lemma 2.6. *There is a deterministic algorithm which, on input a nonzero quadratic form, finds a canonical nonsingular vector v in $O(d^2)$ field operations. There is a deterministic algorithm which, on input a nonzero quadratic form in odd characteristic or a nonzero unitary form, finds a canonical anisotropic vector w in $O(d^2)$ field operations. There is a Las Vegas algorithm which, on input a nondegenerate quadratic form Q with q odd and $d \geq 2$, finds canonical nonsingular vectors u_1, u_2 such that $\iota(Q(u_1)) = 0$ and $\iota(Q(u_2)) = 1$ in $O(d^2 + \log q)$ field operations.*

Proof. We first discuss finding v or w. To find v, let $M = (m_{ij})$ be the matrix of the quadratic form. To find w, let M be the matrix of the polar form of Q or of the unitary form. To find v or w, now look for the smallest i such that $m_{ii} \neq 0$. If i exists, take $v = v_i$ or $w = v_i$. If none exists, let (i,j) be lexicographically minimal subject to $m_{ij} \neq 0$. Let $v = v_i + v_j$, and in the quadratic case let $w = v_i + v_j$ also. If M is unitary, let $w = v_i + \zeta v_j$, so that $\beta(v,v) = \zeta + \zeta^q$, which is nonzero as observed in the proof of Proposition 2.3(1).

To find u_1 and u_2, first choose v_1 nonsingular as above. Compute $v_1^\perp$ as the nullspace of the column vector Fv_1^{Tr} in $O(d^2)$ field operations, then recursively choose nonsingular $v_2 \in v_1^\perp$: note that $v_2 \notin \langle v_1 \rangle$ as v_1 is nonsingular. If possible, take $u_1 = v_i$ for square $Q(v_i)$ and $u_2 = v_j$ for nonsquare $Q(v_j)$. If this is not possible, then either the $Q(v_i)$ are both square, or both are nonsquare. Let $w = v_1 + \nu\sqrt{Q(v_1)/Q(v_2)}v_2$, where ν is as in Proposition 2.3. Then $Q(w) = (1+\nu^2)Q(v_1)$ and hence $\iota(Q(w)) = 1$ if and only if $\iota(Q(v_1)) = 0$, so let u_1 be one of w or v_1 and let u_2 be the other. □

Next we present the main technical ingredient of our isometry construction algorithm. We deal uniformly with symplectic, unitary and symmetric bilinear forms, and refer to the symplectic case as *case* S. We define the *initial k-block* of a matrix X to be the matrix consisting of the first k columns of the first k rows of X. For a matrix over $\mathbb{F}_{q^2}$, the map σ is the qth power map on matrix entries and so the application of σ takes $O(\log q)$ field operations for each entry. For a matrix X, we write X^* for $-X^{\mathrm{Tr}}$ in case S, for $X^{\sigma\mathrm{Tr}}$ in the unitary case, and for X^{Tr} in the orthogonal case. Furthermore, we write $X^\dagger$ for X^{Tr} in case S and for X^* otherwise. Let $a = \log q$ in the unitary case and 0 otherwise. If $SAS^\dagger = B$ we say that *S transforms A to B*. Note that we do not assume that our forms are nondegenerate, so symplectic forms can have odd dimension.

Theorem 2.7 (Diagonalise forms)**.** *Let A be the matrix of a (possibly degenerate) symmetric, unitary, or symplectic form over $\mathbb{F}_{q^u}$, where if q is even then the form is unitary or symplectic. There is a deterministic algorithm which, on input A, constructs a canonical $S \in \mathrm{GL}_d(q^u)$ such* *that $SAS^\dagger$ is diagonal, or block diagonal with blocks of size at most 2 in case S. The algorithm takes $O(d^\omega + d^2a)$ field operations, where a is $\log q$ in the unitary case and 0 otherwise.*

We prove the result via a sequence of lemmas.

Lemma 2.8. *Let A be a matrix of the form*

$$\begin{pmatrix} A_1 & 0 & A_2 \\ 0 & 0 & A_3 \\ A_2^* & A_3^* & A_4 \end{pmatrix},$$

where $A_1 \in \mathrm{GL}_k(q^u)$ for $1 \le k \le d-1$ (with k even in case S) and A_3 has $0 \le s < d-k$ rows. There is a deterministic algorithm which, on input A, constructs a canonical $S \in \mathrm{GL}_d(q^u)$ such that

$$SAS^\dagger = A_1 \oplus \begin{pmatrix} 0 & A_3 \\ A_3^* & A_5 \end{pmatrix}.$$

The algorithm takes $O(d^\omega + d^2a)$ field operations.

Proof. Let $S = \begin{pmatrix} I_k & 0 & 0 \\ 0 & I_s & 0 \\ -A_2^*A_1^{-1} & 0 & I_{d-k-s} \end{pmatrix}$. □

Lemma 2.9. *There is a deterministic algorithm which, on input $A \neq 0$, constructs a canonical $S \in \mathrm{GL}_d(q^u)$ such that $SAS^\dagger = A_1 \oplus 0$ with $A_1 \in \mathrm{GL}_k(q^u)$ for some $1 \le k \le d$ (with k even in case S). The algorithm takes $O(d^\omega)$ field operations.*

Proof. Let $S \in \mathrm{GL}_d(q^u)$ be such that SA is in row echelon form, constructed in $O(d^\omega)$ field operations by Theorem 2.4. Then

$$SAS^\dagger = \begin{pmatrix} X \\ 0 \end{pmatrix} S^\dagger = Y$$

for some matrix $X_{k\times d}$ with full row rank. Now, Y has its final $d-k$ rows all zero, and $Y = Y^*$. Thus the final $d-k$ columns of Y are all zero, and the initial k-block of Y is in $\mathrm{GL}_k(q^u)$. □

Lemma 2.10. *Let $d \equiv 0 \bmod 4$ in case S, and let d be even otherwise. There is a deterministic algorithm which, on input*

$$A = \begin{pmatrix} 0 & A_1 \\ A_1^* & A_2 \end{pmatrix}$$

with $A_1 \in \mathrm{GL}_{d/2}(q^u)$, constructs a canonical $S \in \mathrm{GL}_d(q^u)$ such that the initial $(d/2)$-block of $SAS^\dagger$ is invertible. The algorithm takes $O(d^\omega + d^2a)$ field operations.

Proof. First use Lemma 2.9 to construct $U \in \mathrm{GL}_{d/2}(q^u)$ in $O(d^\omega)$ such that $UA_2U^\dagger = A_3 \oplus 0$, with $A_3 \in \mathrm{GL}_k(q^u)$ for some $k \le d/2$ (and k even in case S). Construct $S_1 = (A_1U^\dagger)^{-1} \oplus U$ in $O(d^\omega + ad^2)$ field operations, then

$$B := S_1AS_1^\dagger = \begin{pmatrix} 0 & I_{d/2} \\ I_{d/2}^* & A_3 \oplus 0 \end{pmatrix}.$$

It is now routine to construct a canonical S_2 such that $S_2BS_2^\dagger$ has invertible initial $(d/2)$-block. □

Lemma 2.11. *Let l with $1 \le l \le d-1$ be given, with l even in case S. There is a deterministic algorithm which, on input an invertible matrix A, constructs a canonical $S \in \mathrm{GL}_d(q^u)$ such that the initial l-block of $SAS^\dagger$ is invertible. The algorithm takes $O(d^\omega + d^2a)$ field operations.*

Proof. If $l > 1$ then first construct a canonical permutation matrix S_1 transforming A to a matrix B whose initial l-block is not identically zero. If $l = 1$ and $a_{11} = 0$ then construct a canonical anisotropic vector v in $O(d^2)$ field operations, by Lemma 2.6, and let B be the form resulting from swapping this v with v_1. Let

$$B = \begin{pmatrix} B_1 & B_2 \\ B_2^* & B_3 \end{pmatrix},$$

where B_1 is $l \times l$. If B_1 is invertible, we are done. Otherwise, construct a matrix S_2 such that

$$C := S_2 B S_2^\dagger = \begin{pmatrix} C_1 \oplus 0 & C_2 \\ C_2^* & B_3 \end{pmatrix},$$

where $C_1 = C_1^* \in \mathrm{GL}_k(q^u)$ for some $k < l$ (with k even in case S). The matrix C can be computed in $O(d^\omega + ad^2)$ field operations by Lemma 2.9. Since C_1 is invertible, by Lemma 2.8 in $O(d^\omega + ad^2)$ field operations we construct a matrix S_3 such that

$$D := S_3 C S_3^\dagger = C_1 \oplus \begin{pmatrix} 0 & D_1 \\ D_1^* & D_2 \end{pmatrix},$$

where D_1 is $(l - k) \times (d - l)$. The fact that A and C_1 are both invertible implies that D_1 has full row rank, so construct a matrix $P \in \mathrm{GL}_{d-l}(q^u)$ in $O(d^\omega)$ field operations such that $D_1 P = (E_1 \ E_2)$ with $E_1 \in \mathrm{GL}_{l-k}(q^u)$. Let $S_4 := I_l \oplus P^\dagger$. Then

$$E := S_4 D S_4^\dagger = C_1 \oplus \begin{pmatrix} 0 & E_1 & E_2 \\ E_1^* & E_3 & E_4 \\ E_2^* & E_4^* & E_5 \end{pmatrix},$$

where E_3 is $(l - k) \times (l - k)$. By Lemma 2.10, in $O(d^\omega + ad^2)$ field operations we can construct a $2(l - k) \times 2(l - k)$ matrix M such that

$$M \begin{pmatrix} 0 & E_1 \\ E_1^* & E_3 \end{pmatrix} M^\dagger$$

has initial $(l - k)$-block invertible. Let $S_5 = I_k \oplus M \oplus I_{d-2l+k}$, then $S_5 E S_5^\dagger$ has invertible initial l-block. □

Proof of Theorem 2.7. If A is identically zero, there is nothing to do. Otherwise, by Lemma 2.9, in $O(d^\omega + d^2 a)$ field operations we can transform A to $S_1 A S_1^\dagger = A_1 \oplus 0$ with $A_1 \in \mathrm{GL}_r(q^u)$ for some $r \leq d$, with r even in case S. Then by Lemma 2.11, in $O(d^\omega + d^2 a)$ field operations we can construct a matrix S_2 transforming A_1 to a matrix A_2 whose initial k-block B_1 is invertible, where $k = 2\lfloor r/4 \rfloor$ in case S and $k = \lfloor r/2 \rfloor$ otherwise. Now by Lemma 2.8, in $O(d^\omega + d^2 a)$ field operations we can construct a matrix S_3 transforming A_2 to $B_1 \oplus C_1$, where $C_1 = C_1^* \in \mathrm{GL}_{r-k}(q^u)$. We now recurse on B_1 and C_1, stopping when we reach 2×2 matrices in case S or 1×1 matrices otherwise. The whole process completes in $O(d^\omega + d^2 a)$ field operations and produces canonical matrices at each step. □

We remark that the symmetric case of the above theorem is proved in [BCS97, Theorem 16.25], although we correct several minor errors in the proof.

Theorem 2.12 (Transform forms)**.** *Suppose we have two nondegenerate symplectic, unitary, or quadratic forms on the space $V = (\mathbb{F}_{q^u})^d$. We can determine if they are isometric, and find a canonical isometry between them, in $O(C)$ field operations, where C is given in Table 1. The algorithm used is deterministic for symplectic forms; otherwise it is Las Vegas.*

Table 1. Complexity for transforming forms

Form type	C
Symplectic	d^ω
Unitary	$d^\omega + d^2 \log q + d \log^2 p$
Quadratic, q odd	$d^\omega + d \log q$
Quadratic, q even	$d^\omega + d \log q + \log^2 q$

Proof. Note that it is enough to find an isometry or similarity from a given form to some fixed form. For quadratic forms we work at least initially with the polar form.

If the form is of unitary type, or the polar form of a quadratic form in odd characteristic, then use Theorem 2.7 to diagonalise the matrix of the form to $\operatorname{diag}(a_1, \ldots, a_d)$. In case S (resp. the form is the polar form of a quadratic form in even characteristic), then transform its matrix to a block diagonal matrix with 2×2 (and 1×1) blocks.

In the symplectic case, each 2×2 block is equal to $\operatorname{antidiag}(a, -a)$ for some $a \in \mathbb{F}_q^\times$. This is transformed to $\operatorname{antidiag}(1, -1)$ by $\operatorname{diag}(a^{-1}, 1)$.

In the unitary case, the form is transformed to I_d by $\operatorname{diag}(\alpha_1, \ldots, \alpha_d)$, where α_i is a canonical solution to $\alpha_i^{q+1} = a_i^{-1}$, using Proposition 2.2.

In the orthogonal case for q odd, if d is odd and the discriminant is nonsquare then let α be the first nonsquare entry, and multiply all entries by α^{-1} (we produce a similarity since $\alpha \neq 1$). In all orthogonal cases now transform all the square entries a_i to 1 by $\sqrt{a_i}^{-1}$ and the nonsquare entries a_i to the first nonsquare entry, μ, by $\sqrt{\mu/a_i}$. The entries μ are then changed in pairs to $\mu(1+\nu^2)$, using the fact that $\left(\begin{smallmatrix} 1 & \nu \\ -\nu & 1 \end{smallmatrix}\right)\left(\begin{smallmatrix} 1 & \nu \\ -\nu & 1 \end{smallmatrix}\right)^{\mathrm{Tr}} = (1+\nu^2)I_2$, where ν is as in Proposition 2.3. Each entry $\mu(1+\nu^2)$ can now be changed to 1, since $\mu(1+\nu^2) \in \mathbb{F}^{\times 2}$. If there is a single nonsquare entry remaining (so that d is even) then this is moved to the first row and transformed to ξ.

In the orthogonal case for q even, the way that we have transformed the polar form matrix F also makes the matrix M of the quadratic form block diagonal with blocks of size at most 2 (since F and M are identical above the diagonal). We now work with M. Since every element of $\mathbb{F}_q$ has a square root, we can convert every block in M to one of the forms (1), $\left(\begin{smallmatrix} 1 & a \\ 0 & 1 \end{smallmatrix}\right)$, or $\left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$. Note that a summand $\left(\begin{smallmatrix} 1 & a \\ 0 & 1 \end{smallmatrix}\right)$ must have $a \neq 0$, otherwise it would be degenerate and so Q would also be degenerate. This also shows that there is at most one summand (1).

Now consider a subform whose matrix is a pair of 2×2 blocks: $\left(\begin{smallmatrix} 1 & a \\ 0 & 1 \end{smallmatrix}\right) \oplus \left(\begin{smallmatrix} 1 & b \\ 0 & 1 \end{smallmatrix}\right)$ with respect to the basis u_1, u_2, u_3, u_4. Changing to the basis $u_1 + u_3, (u_1 + u_4)/b, u_1, bu_2 + a(u_3 + u_4)$, we get the form with matrix $\left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right) \oplus \left(\begin{smallmatrix} 1 & ab \\ 0 & b(a^2+b) \end{smallmatrix}\right)$. The second block can now be converted to $\left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$ or $\left(\begin{smallmatrix} 1 & a' \\ 0 & 1 \end{smallmatrix}\right)$ for some $a' \neq 0$ as above.

So we eventually get a direct sum of copies of $\left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$ together with at most one block of the form (1) or $\left(\begin{smallmatrix} 1 & a \\ 0 & 1 \end{smallmatrix}\right)$. If the polynomial $X^2 + X + a$ has a solution in $\mathbb{F}_q$, then $\left(\begin{smallmatrix} 1 & a \\ 0 & 1 \end{smallmatrix}\right)$ can be transformed to $\left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$, and otherwise it can be transformed to $\left(\begin{smallmatrix} 1 & 1 \\ 0 & \gamma \end{smallmatrix}\right)$. So we are done. □

Theorem 1.1 is just a simplified version of this result. Note that Theorems 1.1 and 2.12 apply unchanged to computing similarities rather than isometries.

2.3. **Groups.** Suppose β (or Q) is a nondegenerate form, as in the previous subsection. Then $\Delta := \mathrm{N}_{\mathrm{GL}_d(q^u)}(\Omega)$ consists of all similarities of the form with itself. The invariant group I consists of all isometries. We use notation from [KL90] for classical groups. For example, if β is a symplectic form, then $\Delta = \mathrm{CSp}_d(q, \beta)$; if β is the canonical symplectic form, then we abbreviate this to $\mathrm{CSp}_d(q)$.

Define $\tau : \Delta \to \mathbb{F}_{q^u}$ by $\beta(ux, vx) = \tau(x)\beta(u, v)$ for all $u, v \in V$. It is well known (see for example [KL90, Lemma 2.1.2]) that τ is a homomorphism with kernel I.

Lemma 2.13. *There is a deterministic algorithm which, on input $g \in \Delta$ and the matrix F of β, computes $\tau(g)$ in $O(d^2)$ field operations.*

Proof. Find w such that $wFv_1^{\mathrm{Tr}} \neq 0$ in $O(d)$ field operations. Then $\tau(g)$ is $\beta(wg, v_1 g)/\beta(w, v_1)$. □

For quadratic forms, the spinor norm is an epimorphism from the general orthogonal group $I = \mathrm{GO}_d(q, Q)$ to $\mathbb{F}_2^+$.

Definition 2.14 (Spinor norm)**.** *Let $g \in \mathrm{GL}(d, q)$ preserve the form Q.*

(1) *For q odd, let $U \leq V$ be the image of $I_d - g$ and define the bilinear form χ on U by $\chi(u, v) = 2\beta(w, v)$ where $w(I_d - g) = u$. The* spinor norm *of g is $\mathrm{sp}(g) = \iota(\det(\chi))$.*
(2) *For q even, the* spinor norm *of g is $\mathrm{sp}(g) = \mathrm{rank}(I_d + g) \bmod 2$.*

Our definition for odd q is from [Tay92], except for the factor of two which we include so the values of the spinor norm agree with [KL90, p.29]. We follow [KL90, Proposition 2.5.7] and define $\Omega_d(q, Q) := \mathrm{SO}_d(q, Q) \cap \ker(\mathrm{sp})$. What we call the spinor norm for even q is called the Dickson invariant by some authors.

Theorem 2.15. *There is a deterministic algorithm that, on input $g \in \mathrm{GO}_d(q, Q)$, computes $\mathrm{sp}(g)$. If q is even then the algorithm takes $O(d^\omega)$ field operations, otherwise it takes $O(d^\omega + \log q)$ field operations.*

Proof. If q is even, apply Theorem 2.4. If q is odd, compute the nullspace N of $a := I_d - g$ and find a matrix M whose rows are a basis to a complement of N in $O(d^\omega)$ field operations. Then the rows of Ma are a basis for the image of a. Calculate the form χ_g on Ma as $S = 2MF(Ma)^{\mathrm{Tr}}$ in $O(d^\omega)$ field operations. Finally, find $\iota(\det S)$. □

We finish this section with a discussion of reflections. Let $v \in V$ be nonsingular, so that $Q(v) \neq 0$. The *reflection* in v is the map $\mathrm{refl}_v : V \to V$, $u \mapsto u - \beta(u, v)v/Q(v)$.

Lemma 2.16. *Let Q be nondegenerate with polar form F, and let $u, v \in V$ be nonsingular.*

(1) *All reflections are elements of $\mathrm{GO}_d(q, Q)$, and have determinant -1 and order 2.*
(2) *For q even, $\mathrm{sp}(\mathrm{refl}_v) = 1$.*
(3) *For q odd, $\mathrm{sp}(\mathrm{refl}_v) = \iota(\beta(v, v))$.*
(4) *For q odd, $\Omega_d(q, Q)\,\mathrm{refl}_u = \Omega_d(q, Q)\,\mathrm{refl}_v$ if and only if $\iota(\beta(u, u)) = \iota(\beta(v, v))$.*

Proof. Parts (1) and (2) are well-known, and are easy exercises. For part (3), let $g = \mathrm{refl}_v$. Then $(I_d - g)$ has image $\langle v \rangle$, and maps $v \mapsto 2v$, so the matrix of χ_g is $(\beta(v, v))_{1\times 1}$. Part (4) follows from part (3) and the fact that sp is a homomorphism. □

Proposition 2.17. *Let Q be nondegenerate. For odd q and $d \geq 2$, there is a Las Vegas algorithm that constructs canonical reflections R_0, R_1 with $\mathrm{sp}(R_i) = i$ in $O(d^2 + \log q)$ field operations. For even q and $d \geq 2$, a canonical reflection R_0 can be constructed deterministically in $O(d^2)$ field operations.*

Proof. For q odd, by Lemma 2.6 we can find canonical vectors u_0, u_1 with $\iota(Q(u_i)) = i$. Note that $u_i F v_j^{\mathrm{Tr}}$ can be computed in $O(d)$ field operations for each j, as Fv_j is the jth row of F. Then row j of refl_{u_i} is $v_j - (u_i F v_j^{\mathrm{Tr}})Q(u_i)^{-1}u_i$. The case q even is similar. □

3. Constructive homomorphisms

In this section, for each type of classical group, we construct the quotient of the conformal group Δ by the quasisimple group Ω as a presentation in two ways. The first presentation has $O(q)$ generators, and a word for the image of an element of Δ can be found in polynomial time. The second presentation is polycyclic with at most four generators and at most six relations, but words for images can only be found using discrete logarithms. To our knowledge, for the orthogonal groups such presentations only exist in the literature for the projective groups [KL90, Sections 2.5–2.8]. Note that the first presentation has a constant number of generators and relations when considered as an FC-presentation in the sense of [CHM08]. We also compute canonical representatives for cosets of Ω, which are needed for the conjugacy problem in Section 4. Throughout this section we assume that Ω is quasisimple, which eliminates some small dimensional exceptional cases.

Our main result in this section is the following theorem.

Theorem 3.1. *Let $\Omega \leq \mathrm{GL}_d(q^u)$ be a quasisimple classical group fixing a known classical form, let $\Delta = \mathrm{N}_{\mathrm{GL}_d(q^u)}(\Omega)$ and let $G := \Delta/\Omega$. Let X be the matrix tranforming the canonical form to the given form (Theorem 1.1). Let X_i, $\mathcal{R}_i$, and C_i $(i = 1, 2)$ be defined as in Table 2.*

(0) *Δ is generated by Ω and X_0.*

(1) *$P_1 = \langle X_1 \mid \mathcal{R}_1\rangle$ is a presentation for G. The image of $g \in \Delta$ as a canonical word in P_1 can be computed in $O(C_1)$ field operations.*

(2) *$P_2 = \langle X_2 \mid \mathcal{R}_2\rangle$ is a polycyclic presentation for G. The image of $g \in \Delta$ as a canonical word in P_2 can be computed in $O(C_1)$ field operations plus C_2 discrete logarithms.*

(3) *A canonical representative of the coset Ωg, where $g \in \Delta$, can be computed in $O(C_3)$ field operations.*

For unitary and orthogonal groups, these algorithms are Las Vegas; in the other cases they are deterministic.

Note that Theorem 1.2 is just a simplified version of this result. The proof is straightforward in the linear and symplectic cases, and is similar to the unitary case.

Proof of Theorem 3.1, unitary case. Proof of (0): By [KL90, Table 2.1.C], $[\Delta : \Omega] = q^2 - 1$. The matrix $A(\lambda) \in \Delta$ for all $\lambda \in \mathbb{F}_{q^2}^\times$, as $A(\lambda)$ preserves the canonical unitary form up to scalars. The matrix $B(\lambda) \in \mathrm{GU}_d(q)$ for all $\lambda \in \mathbb{F}_{q^2}^\times$, as it preserves the canonical unitary form. The determinant of $B(\zeta)$ has order $q + 1$, so $B := \langle B(\lambda), \Omega\rangle/\Omega$ is cyclic of order $q + 1$. The τ map shows that $\langle A(\lambda), B\rangle/B$ is cyclic of order $q - 1$, so the result follows.

Proof of (1): First we check the presentation P_1. Since $A(\lambda)A(\mu) = A(\lambda\mu)$, we see that $a(\lambda)a(\mu) = a(\lambda\mu)$, and similarly $b(\lambda)b(\mu) = b(\lambda\mu)$. It follows from the proof of (0) that $b(\lambda)^{q+1} = 1$, and that some power of $a(\lambda)$ is a power of $b(\lambda)$. To show that $a(\lambda)^{q-1} = b(\lambda)^d$, note that $A(\lambda)^{q-1}B(\lambda)^{-d}$ has determinant 1.

We map $g \in \Delta$ to $a(\tau(g))b(\mu^{-d} \det(g)) \in P_1$, where μ is the canonical solution of $\mu^{q+1} = \tau(g)$. This is the correct image since it factors through det and τ correctly. Since $\tau(g)$ can be computed by a deterministic algorithm in $O(d^2)$ field operations by Lemma 2.13, and μ can be computed by a Las Vegas algorithm in $O(\log q + \log^2 p)$ field operations by Proposition 2.2, the result follows.

Proof of (2): It is clear that P_2 presents the same group as P_1. To write $g \in \Delta$ as a word in a and b, find the discrete logarithms of $\tau(g)$ and $\mu^{-d} \det(g)$.

Proof of (3): Use Theorem 1.1 to find X such that $\mathrm{SU}_d(q, \beta) = \mathrm{SU}_d(q)^X$. Take the coset representative of $g \in \Delta$ to be $(A(\tau(g))B(\mu^{-d} \det(g)))^X$. □

	$B(\lambda) = \left((\lambda^q) \oplus I_{d-2} \oplus (\lambda^{-1})\right)^X$	$b(\lambda)$	$a(\lambda)^{q-1} = b(\lambda)^d$ $[a(\lambda), b(\mu)]$	$+\log^2 p$	$b := b(\zeta)$	b^{q+1}, $a^{q-1} = b^d$		$+d\log^2 p$
$\Omega_d(q)$, q even	$R_0, C(\lambda) = \lambda^{q/2} I_d$	$r_0, c(\lambda)$	(3), $[r_0, c(\lambda)]$	$d^\omega + \log^2 q$	$r_0, c := c(\xi)$	(4), $[r_0, c]$, c^{q-1}	1	$d^\omega + \log^2 q$
$\Omega_d^\circ(q)$, d odd, q odd	R_0, R_1, $C(\lambda) = \left(\lambda^2 I_m \oplus (\lambda) \oplus I_m\right)^X$	r_0, r_1, $c(\lambda)$	(3), $[r_i, c(\lambda)]$, $c(-1) = r_0$	$d^\omega + \log q$	r_0, r_1, $c := c(\xi)$	(4), $[r_0, c]$, $[r_1, c]$, $c^{(q-1)/2}$	1	$d^\omega + d\log q$
$\Omega_d^+(q)$, d even, q odd	R_0, R_1, $C(\lambda) = (\lambda I_m \oplus I_m)^X$	r_0, r_1, $c(\lambda)$	(3), $r_i^{c(\lambda)} = r_{i+\iota(\lambda)}$	$d^\omega + d\log q$	r_0, r_1, $c := c(\xi)$	(4), $r_0^c = r_1$, $r_1^c = r_0$, $c^{q-1} = r_0$	1	$d^\omega + d\log q$
$\Omega_d^-(q)$, d even, q odd	R_0, R_1, $C(\lambda) = \left(\lambda^2 I_m \oplus \lambda I_2 \oplus I_m\right)^X$ $C_0^- = \left(\gamma I_m \oplus \left(\begin{smallmatrix} 0 & 1 \\ \gamma & 0 \end{smallmatrix}\right) \oplus I_m\right)^X$	r_0, r_1, $c_0, c(\lambda)$	(3), $[r_i, c(\lambda)]$, $c(-1) = r_0 r_1$ $[c_0, c(\lambda)], c_0^2 = c(\gamma)$, $r_i^{c_0} = r_{i+1}$	$d^\omega + d\log q$	r_0, r_1, $c := c(\sqrt{\xi\gamma^{-1}})c_0$	(4), $r_0^c = r_1$, $r_1^c = r_0$, $c^{q-1} = r_0$	1	$d^\omega + d\log q$

(1) The generators $R_0, R_1 \in X_0$ are defined as in Proposition 2.17. For the group $\Omega_d^-(q)$, we define γ as in Proposition 2.3.
(2) We define $a(\lambda) \in X_2$ to be the coset $\Omega A(\lambda)$, and similarly for $b(\lambda)$, r_0, r_1, $c(\lambda)$, c_0, for $\lambda, \mu \in \mathbb{F}_{q^u}^\times$ and $i \in \mathbb{F}_2^+$.
(3) The following relations are in $\mathcal{R}_1$ whenever the relevant generators are defined:
$a(\lambda)a(\mu) = a(\lambda\mu)$, $b(\lambda)b(\mu) = b(\lambda\mu)$, $c(\lambda)c(\mu) = c(\lambda\mu)$, $r_0^2 = r_1^2 = (r_0 r_1)^2 = 1$.
(4) The following relations are in $\mathcal{R}_2$ whenever the relevant generators are defined: $r_0^2 = r_1^2 = (r_0 r_1)^2 = 1$.

In the remainder of this section, we consider the orthogonal case. Since Ω is quasisimple by assumption, $d \geq 3$. If q is even, we also assume d is even, since in even characteristic the odd degree orthogonal groups are isomorphic to symplectic groups. For $\epsilon \in \{+, -, \circ\}$ we write $G = G^\epsilon(q) := \mathrm{CO}^\epsilon_d(q)/\Omega^\epsilon_d(q)$.

Our first result proves Theorem 3.1(0), and part of Theorem 3.1(1) for the orthogonal case.

Proposition 3.2. *The group* $\mathrm{CO}^\epsilon_d(q)$ *is generated by* $\Omega^\epsilon_d(q)$ *together with the generators* X_0 *in Table 2. Furthermore,* $P_1 = \langle X_1 | \mathcal{R}_1 \rangle$ *is a presentation for* $G^\epsilon_d(q)$.

Proof. It is easy to check that $C^\epsilon(\lambda) \in \mathrm{CO}^\epsilon_d(q)$ and $C^-_0 \in \mathrm{CO}^-_d(q)$. Note that $\tau(C^\epsilon(\lambda)) = \lambda^2$ when q is odd and ϵ is $\circ$ or $-$; whilst $\tau(C^\epsilon(\lambda)) = \lambda$ in all other cases. One may check that $\tau(C^-_0) = \gamma$.

The kernel of τ on $\mathrm{CO}^\epsilon_d(q)$ is $\mathrm{GO}^\epsilon_d(q)$, and its image is $\mathbb{F}^\times_q$ if d is even, and $\mathbb{F}^{\times 2}_q$ otherwise [KL90, §2.1]. For d odd, $\tau(C^\circ(\xi)) = \xi^2$ generates $\mathbb{F}^{\times 2}_q$. If ϵ is $+$ or q is even, then $\tau(C^\epsilon(\xi)) = \xi$ generates $\mathbb{F}^\times_q$. Finally, if ϵ is $-$ and q is odd, then $\tau(C^-(\xi)) = \xi^2$ and $\tau(C^-_0) = \gamma$ generate $\mathbb{F}^\times_q$, since γ is nonsquare. Since $\mathrm{GO}^\epsilon_d(q)$ is generated by $\Omega^\epsilon_d(q)$ and the reflections, $\mathrm{CO}^\epsilon_d(q)$ is generated by the given elements.

For q even or d odd, $G^\epsilon(q) = \langle r_0 \rangle \times \langle c(\xi) \rangle \cong \mathbb{F}^+_2 \times \mathbb{F}^\times_q$. For q odd, $G^+(q)$ is an extension of $\langle r_0, r_1 \rangle \cong (\mathbb{F}^+_2)^2$ by $\langle c(\xi) \rangle \cong \mathbb{F}^\times_q$, whilst $G^-(q)$ is an extension of $\langle r_0, r_1 \rangle \cong (\mathbb{F}^+_2)^2$ by $\langle c(\xi), c_1 \rangle \cong \mathbb{F}^\times_q$. Hence $G^\epsilon(q)$ has the same order as $\mathrm{CO}^\epsilon_d(q)/\Omega^\epsilon_d(q)$ [KL90, § 2.1]. It therefore suffices to show that the relations hold.

All relations involving only r_0 and r_1 hold because the quotient $\mathrm{GO}^\epsilon_d(q)/\Omega^\epsilon_d(q)$ is an elementary abelian 2-group. For the relations involving r_0 or r_1 conjugated by $c(\lambda)$ or c_0, note that $\mathrm{refl}^g_v = \mathrm{refl}_{vg}$ for $v \in V$ and $g \in \mathrm{CO}^\epsilon_d(q)$. For q even, all reflections are in the same coset of $\Omega^\pm_d(q)$, and so $r_0^{c(\lambda)} = r_0$. For q odd, $\iota(Q(vg)) = \iota(Q(v)) + \iota(\tau(g))$. For the relations involving products and powers of $c(\lambda)$ and c_0, one checks that $C^\epsilon(\lambda)C^\epsilon(\mu) = C^\epsilon(\lambda\mu)$ and so $c(\lambda)c(\mu) = c(\lambda\mu)$. Now, $C^\circ_{2m+1}(-1) = I_m \oplus (-1) \oplus I_m = \mathrm{refl}_x$, and since $Q^\circ(x) = 1$ we deduce that $c(-1) = r_0$. Finally, $C^-(\lambda)$ commutes with C^-_0; $(C^-_0)^2 = C^-(\gamma)$; and $C^-(-1) = I_m \oplus -I_2 \oplus I_m = \mathrm{refl}_x \,\mathrm{refl}_y$, so $c(-1) = r_0 r_1$. □

By setting $c = c(\xi)$, or $c = c(\sqrt{\xi\gamma^{-1}})c_0$ for q odd and $\epsilon = -$, we get presentations for the same groups with a bounded number of generators and relations.

Corollary 3.3. $P_2 = \langle X_2 | \mathcal{R}_2 \rangle$ *is a presentation for* $G^\epsilon_d(q)$.

We can now prove Theorem 3.1 for the orthogonal groups. If q is odd and Q is of $-$ type, we assume that the discrete log of γ has been precomputed in (2). We only give the case where q is odd, d is even, and Q is of $-$ type, as the other orthogonal cases are similar.

Proof of Theorem 3.1, orthogonal minus case. Proof of (0): This is immediate from Proposition 3.2.

Proof of (1): It is immediate from Proposition 3.2 that P_1 presents $G^\epsilon_d(q)$. For the homomorphism, we first find a canonical matrix X which tranforms the canonical form to F, in $O(d^\omega + d^2 \log q)$ field operations. We compute $\tau(g)$ in $O(d^2)$ field operations. If $\tau(g)$ is a square, we take $\lambda = \sqrt{\tau(g)}$, $z = c(\lambda)$ and $C = C^-(\lambda)$. Otherwise we take $\lambda = \sqrt{\tau(g)\gamma^{-1}}$, $z = c_0 c(\lambda)$, and $C = C^-_0 C^-(\lambda)$. We then let $h = g^{X^{-1}} C^{-1}$, find $a = \det(h)$ and $b = \mathrm{sp}(h)$ in $O(d^\omega + \log q)$ field operations. We map g to $r_0^{b'} r_1^b z$, where $b' = b$ if $a = 1$ and $b' = b + 1$ otherwise.

Proof of (2): It is immediate from Corollary 3.3 that P_2 presents $G^\epsilon_d(q)$. For the homomorphism, find $k = \log_{\xi\gamma} \lambda = \frac{\log \lambda}{\log \gamma + 1}$ with a discrete log call, and map g to $r_0^{b'} r_1^b c^k$.

Proof of (3): Write down R_0 and R_1 from Proposition 2.17 in $O(d^\omega + \log q)$ field operations, then the representative is $(R_0^{b'} R_1^b C)^X$. □

We finish with a special case, where our algorithms run faster.

Proposition 3.4. *Let Q be a nondegenerate quadratic form, and let $g \in \mathrm{GO}_d(q, Q)$. Then the image of g under the natural homomorphism to $\mathbb{F}_2^+$ (q even) or $(\mathbb{F}_2^+)^2$ (q odd) can be found by a deterministic algorithm in $O(d^\omega)$ field operations (q even) or $O(d^\omega + \log q)$ field operations (q odd) . A canonical coset representative for g can then be constructed by a deterministic algorithm in $O(d^2)$ field operations if q is even and, given ζ, by a Las Vegas algorithm in $O(d^\omega + \log q)$ field operations otherwise.*

4. Applications: conjugacy and maximal subgroups

Given a finite group G, the basic conjugacy problems are:

(1) find a set of canonical representatives of the conjugacy classes of G;
(2) given $x \in G$, find $g \in G$ such that x^g is a canonical class representative; and
(3) given a class representative x, find generators for $\mathrm{C}_G(x)$.

We conjugate to a class representative in problem 2, rather than designing an algorithm to conjugate arbitrary pairs of elements, because it reduces memory requirements. This way we need only work with a single element of the group, since the representative itself is implicit in the algorithm but does not usually need to be written down. This was our motivatation for the inclusion of canonical coset representatives in Theorem 3.1(3).

Suppose we can solve the element conjugacy problem in the group Δ. We briefly describe how to solve the same problem for groups G with $\Omega \leq G \leq \Delta$. This is a slight generalisation of the results of [Wal80], and is based on the following lemma.

Lemma 4.1. *Let Δ be a group, A a finite group, and $\phi : \Delta \to A$ an epimorphism. Let Ω be the kernel of ϕ. Suppose G is a group with $\Omega \leq G \trianglelefteq \Delta$. Given $g \in G$, the G-classes contained in g^Δ correspond to the elements of $A/\phi(\mathrm{C}_\Delta(g)G)$ under the map*

$$(g^h)^\Delta \mapsto \phi(\mathrm{C}_\Delta(g)Gh)$$

for h in Δ.

Proof. Clearly every G-class in g^Δ is of the form $(g^h)^G$ for some $h \in \Delta$. Now $(g^h)^G = (g^{h'})^G$ if and only if $g^{hg'} = g^{h'}$ for some $g' \in G$, that is, $hg'h'^{-1}$ is in $\mathrm{C}_\Delta(g)$ for some $g' \in G$. Since G is normal in Δ, this is equivalent to h being in $\mathrm{C}_\Delta(g)Gh'$, which means $\mathrm{C}_\Delta(g)Gh = \mathrm{C}_\Delta(g)Gh'$. Since $A/\phi(\mathrm{C}_\Delta(g)G)$ is naturally isomorphic to $\Delta/\mathrm{C}_\Delta(g)G$, we are done. □

Hence, in order to compute the classes in G from the classes in Δ, we need to know the images of centralisers under ϕ and we need representatives $h_a \in \phi^{-1}(a)$ for all $a \in A$. If G is not normal in Δ, we need to apply this lemma more than once: since Δ/Ω is soluble for classical groups Ω, every G with $\Omega \leq G \leq \Delta$ is subnormal in Δ.

Solving problem (1) is only possible for relatively small groups, but since Theorem 3.1(3) gives canonical *coset* representatives we can find canonical *class* representatives to solve problem (2) without first solving (1). Canonical class representatives also simplify the centraliser problem (3), and allow us to compare results between different runs of the algorithms. A detailed description of these algorithms is given in [HM].

An important application of Theorem 1.1 is to the construction of maximal subgroups of classical groups, as in [HRD05, HRD10]. When writing down generating matrices for a maximal subgroup, it is often convenient to construct initial matrices which preserve a form other than

TABLE 3. Spinor norm on $\mathrm{GO}^{\epsilon}_{d}(q, Q)$

		p					3^i			2^i				
Type	d	5	17	47	73	10000019	3^6	3^{11}	3^{16}	2^5	2^{10}	2^{20}	2^{40}	2^{80}
○	15	1	1	1	1	1	1	2	5					
	55	4	9	9	9	11	11	28	184					
	95	11	27	27	28	34	45	140	1083					
+	20	1	1	1	1	1	1	3	10	–	–	–	4	4
	60	4	11	10	11	13	13	38	246	–	1	12	60	78
	100	12	28	28	27	33	50	153	1408	2	7	57	311	413
−	20	1	1	2	1	1	1	3	10	–	–	–	4	3
	60	4	11	11	11	14	14	36	256	–	1	12	60	82
	100	11	28	27	26	33	48	148	1373	4	7	56	289	390

Magma's canonical classical form. We then conjugate the matrices so that they preserve the correct form. Since the isometry construction algorithm given in [HRD05] does not return the same conjugating matrix each time, different conjugates of the maximal subgroup are found each time it is constructed. Using Theorem 1.1, the *same* subgroup can now be constructed each time. This is not essential, but is often useful: for example when investigating containments between subgroups.

5. TIMINGS

In this section we present two tables of timings data for a Magma 2.14-9 [BC07] implementation of our algorithms. We tested our spinor norm algorithm on $\mathrm{GO}_d(q, Q)$ on all five cases: odd dimension and odd characteristic, and both types of form in even dimensions in both even and odd characteristic. In each case we computed the spinor norm of a random element of a random conjugate of the general orthogonal group.

Next we tested the canonical coset representative algorithms on all five cases. We took a random conjugate of the conformal orthogonal group, and then selected a random element. The time to find coset representatives for elements of the general orthogonal group lies between that taken to compute the spinor norm and to find coset representatives in the conformal orthogonal group.

The experiments were carried out on a 1.5 GHz PowerPC G4 processor. The machine has 1.25GB of RAM, but memory was not a factor. All times are given in milliseconds, and are the average of 50 trials; the symbol – indicates that the average time was less than 1 millisecond.

As we would expect, the time required grows extremely slowly with q, and somewhat more quickly with d. Far less time is required for even q than odd q. Notice however that the representation of the field is more significant than its size, as 3^{16} is only about four times larger than 10000019, yet the tests always take far longer.

REFERENCES

[Asc84] M. Aschbacher. On the maximal subgroups of the finite classical groups. *Invent. Math.*, 76(3):469–514, 1984.

[Bab97] László Babai. Randomization in group algorithms: conceptual questions. In *Groups and computation, II (New Brunswick, NJ, 1995)*, pages 1–17. Amer. Math. Soc., Providence, RI, 1997.

[BC07] W. Bosma and J.J. Cannon. *Handbook of Magma functions*. School of Mathematics and Statistics, University of Sydney, Sydney, 2.14 edition, 2007.

[BCS97] P. Bürgisser, M. Clausen, and M. A. Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1997.

Table 4. Coset representatives in $\mathrm{CO}_d^{\epsilon}(q,Q)$

		p					3^i							
Type	d	5	17	47	73	10000019	3^6	3^{11}	3^{16}	2^5	2^{10}	2^{20}	2^{40}	2^{80}
○	15	3	4	4	4	6	3	5	13					
	55	33	48	55	47	59	46	72	392					
	95	147	201	184	176	211	189	317	2342					
+	20	6	7	7	7	10	7	10	34	1	2	4	8	14
	60	46	62	68	65	77	76	148	936	17	18	26	124	170
	100	168	224	209	226	257	305	627	5645	49	67	127	553	629
−	20	7	9	9	9	11	153	15	40	1	1	3	7	11
	60	50	72	71	70	90	244	196	1168	14	12	25	131	154
	100	153	225	217	229	257	474	799	7969	71	60	119	553	736

[BGM93] I.F. Blake, S. Gao, and R.C. Mullin. Explicit factorization of $x^{2^k}+1$ over $\mathbf{F}_p$ with prime $p \equiv 3 \bmod 4$. *Appl. Algebra Engrg. Comm. Comput.*, 4(2):89–94, 1993.

[Bri06] John R. Britnell. Cyclic, separable and semisimple transformations in the finite conformal groups. *J. Group Theory*, 9(5):571–601, 2006.

[Bro01] P.A. Brooksbank. A constructive recognition algorithm for the matrix group $\Omega(d,q)$. In *Groups and computation, III*, volume 8 of *Ohio State Univ. Math. Res. Inst. Publ.*, pages 79–93. de Gruyter, Berlin, 2001.

[Bro03] P.A. Brooksbank. Constructive recognition of classical groups in their natural representation. *J. Symbolic Comput.*, 35(2):195–239, 2003.

[CHM08] Arjeh M. Cohen, Sergei Haller, and Scott H. Murray. Computing in unipotent and reductive algebraic groups. *LMS J. Comput. Math.*, 11:343–366, 2008.

[GCL92] K.O. Geddes, S.R. Czapor, and G. Labahn. *Algorithms for computer algebra*. Kluwer Academic Publishers, Boston, MA, 1992.

[Gro02] L.C. Grove. *Classical groups and geometric algebra*, volume 39 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2002.

[HM] Sergei Haller and Scott H. Murray. Computing conjugacy in finite classical groups. Unpublished.

[HRD05] D.F. Holt and C.M. Roney-Dougal. Constructing maximal subgroups of classical groups. *LMS J. Comput. Math.*, 8:46–79, 2005.

[HRD10] D.F. Holt and C.M. Roney-Dougal. Constructing maximal subgroups of orthogonal groups. *LMS J. Comput. Math.*, 2010. To appear.

[JLPW95] C. Jansen, K. Lux, R. Parker, and R. Wilson. *An Atlas of Brauer Characters*. Oxford University Press, Oxford, UK, 1995.

[KL90] P. Kleidman and M. Liebeck. *The subgroup structure of the finite classical groups*. Cambridge University Press, Cambridge, 1990.

[Lan93] Serge Lang. *Algebra*. Addison-Wesley Publishing Co., Reading, Mass., third edition, 1993.

[LG01] C.R. Leedham-Green. The computational matrix group project. In *Groups and computation, III*, volume 8 of *Ohio State Univ. Math. Res. Inst. Publ.*, pages 229–247. de Gruyter, Berlin, 2001.

[Lüb] F. Lübeck. `http://www.math.rwth-aachen.de/~Frank.Luebeck/data/ConwayPol`.

[Str69] V. Strassen. Gaussian elimination is not optimal. *Numer. Math.*, 13:354–356, 1969.

[Tay92] D.E. Taylor. *The geometry of the classical groups*. Heldermann Verlag, Berlin, 1992.

[Wal80] G. E. Wall. Conjugacy classes in projective and special linear groups. *Bull. Austral. Math. Soc.*, 22(3):339–364, 1980.

Department of Mathematics and Statistics, Building 11, University of Canberra, ACT, 2601, Australia

E-mail address: `murray@maths.usyd.edu.au`

School of Mathematics and Statistics, University of St Andrews, Fife KY16 9SS, UK.

E-mail address: `colva@mcs.st-and.ac.uk`