Effective black-box constructive recognition of classical groups

Heiko Dietrich, C. R. Leedham-Green, and E. A. O'Brien

Dedicated to the memory of Ákos Seress

ABSTRACT. We describe a black-box Las Vegas algorithm to construct standard generators for classical groups defined over finite fields. We assume that the field has size at least 4 and that oracles to solve certain problems are available. Subject to these assumptions, the algorithm runs in polynomial time. A practical implementation of our algorithm is distributed with the computer algebra system MAGMA.

1. Introduction

In [19, 26] we developed constructive recognition algorithms for the classical groups in their natural representation. These are well-analysed and efficient, both theoretically and in practice; our implementations are distributed with the computer algebra system MAGMA [9]. A core idea is to construct centralisers of involutions, and use these to construct, as subgroups of the input group, classical groups of smaller rank, so facilitating recursion. We now develop these ideas to obtain such algorithms for classical groups given as black-box groups.

Let $\tilde{G} \leq \mathrm{GL}_d(q)$ be a classical group in its natural representation, and let $G = \langle X \rangle$ be isomorphic to a central quotient of $\tilde{G}$, where X is a given generating set. A *constructive recognition* algorithm for G constructs a surjective homomorphism from $\tilde{G}$ to G, and for any given $g \in G$ constructs an element of its inverse image in $\tilde{G}$. We realise such an algorithm in two stages. For each classical group $\tilde{G}$, we define a specific ordered set of *standard generators* $\tilde{\mathcal{S}}$. The first task is to construct, as words in X, an ordered subset $\mathcal{S}$ of G that is the image of $\tilde{\mathcal{S}}$ under a surjective homomorphism from $\tilde{G}$ to G. The second task is to solve the *constructive membership problem* for G with respect to $\mathcal{S}$: namely, express $g \in G$ as a word in $\mathcal{S}$, and so as a word in X; we also solve the constructive membership for $\tilde{G}$ with respect to $\tilde{\mathcal{S}}$. Now the surjective homomorphism $\varphi\colon \tilde{G} \to G$ that maps $\tilde{\mathcal{S}}$ to $\mathcal{S}$ is *constructive*: $\tilde{g} \in \tilde{G}$ is written as a word $w(\tilde{\mathcal{S}})$ in $\tilde{\mathcal{S}}$, and its image $\varphi(\tilde{g})$ is $w(\mathcal{S})$. Similarly, we compute a preimage in $\tilde{G}$ of $g \in G$ under φ. In summary, we provide an algorithm to solve the first of these tasks; we discuss the second in Section 1.3.

Babai and Szemerédi [6] introduced the concept of a *black-box group*: group elements are represented by bit strings of uniform length, where more than one bit string may represent the same element. Three *oracles* are provided to supply the group-theoretic functions of multiplication, inversion, and checking for equality with the identity element. A *black-box* algorithm is one that uses only these oracles. A common assumption is that other oracles are available to perform certain tasks.

For an overview of the *Matrix Group Recognition Project*, to which this work contributes, see [37]. Much of the background and preliminaries needed for this paper are summarised in [19, 26, 37].

1.1. The groups and their standard copies.

Throughout, $\mathrm{GL}_d(q)$ is the group of invertible $d \times d$ matrices over the field $\mathrm{GF}(q)$. The groups under discussion are $\mathrm{SL}_d(q)$, $\mathrm{Sp}_d(q)$, $\mathrm{SU}_d(q)$, $\Omega_d^{\pm}(q)$, and $\Omega_d(q)$. We assume that $q \geq 4$, and $d \geq 3$ for the orthogonal groups. All of the groups are perfect, and with the exception of $\Omega_4^+(q)$, all are quasisimple.

The definition of these groups, except for the first, depends on the choice of a bilinear or quadratic form. Groups defined by two forms of the same type are conjugate in the corresponding general linear

Key words and phrases. classical groups, constructive recognition, black-box algorithms.

This work was supported in part by the Marsden Fund of New Zealand via grant UOA 105. Dietrich was supported by an ARC-DECRA Fellowship, project DE140100088. The last two authors were supported by GNSAGA-INdAM while this work was completed; we thank our hosts, Patrizia Longobardi and Mercede Maj of the University of Salerno, for their generous hospitality. We thank Damien Burns for early work on the odd characteristic case; we thank Jianbei An, Gerhard Hiss, and Martin Liebeck for helpful discussions; we thank the referee and editor for their comments.

group; the *standard copy* of a classical group is its unique conjugate which preserves a chosen *standard form*. Our standard forms and copies are described in detail in [19, 26]. The *standard generators* of a classical group $\tilde{G}$ satisfy a specific *standard presentation*. The latter is used to define standard generators for a (black-box) group G isomorphic to a central quotient of $\tilde{G}$: namely, a generating set of G satisfying this presentation.

We write $\mathrm{SX}_d(q)$ for a conjugate of one of the above groups in the natural representation; we call SL, SU, Sp, Ω, and $\Omega^\pm$ the *type* of the group.

Definition 1.1. The standard generators $\mathcal{S}(d, q, \mathrm{SX})$ of $\mathrm{SX}_d(q)$ are given in [19, Table 1] and [26, Tables 1 & 2], depending on whether q is even or odd.

The definition of the standard generators of $\mathrm{SX}_d(q)$ implies a *fixed* choice of primitive element for the underlying field. Observe that $\mathcal{S}(d, q, \mathrm{SX})$ has cardinality at most 8 and, with the exception of one element, the *cycle* of $\mathrm{SX}_d(q)$, all standard generators lie in naturally embedded subgroups $\mathrm{SX}_4(q)$ of $\mathrm{SX}_d(q)$. This observation is crucial since we construct $\mathcal{S}(d, q, \mathrm{SX})$ by a recursion to classical groups of smaller degree.

1.2. Main results. Let $G = \langle X \rangle$ be isomorphic to a central quotient of $\mathrm{SX}_d(q)$. We present and analyse a black-box Las Vegas algorithm that takes as input X, and the parameters (d, q, SX) of G, and outputs standard generators of G as words in X. All words are given as *straight-line programs* (SLPs) [42, p. 10] which may be regarded as efficiently stored group words in X.

The complexity of a black-box algorithm is measured in terms of the number of calls to the standard oracles for the black-box G. Let μ be an upper bound on the time required for each group operation.

Our algorithm assumes the existence of the following.

- An oracle $\mathcal{O}$ to compute the order of a given $g \in G$.
- An oracle Π to compute a given power of $g \in G$.
- An oracle ξ to construct a (nearly) uniformly distributed random element of G as an SLP in X.
- An oracle χ to recognise constructively (a central quotient of) $\mathrm{SL}_2(q)$.

We abuse notation by identifying the oracle with its cost. We ignore the cost of standard integer operations such as computing the greatest common divisor of two integers.

Our main result is the following theorem; it is proved in Sections 4–6. In Section 7 we discuss the complexity and the cost of realising the oracles.

Theorem 1.2. *Let $G = \langle X \rangle$ be a black-box group isomorphic to a central quotient of $\mathrm{SX}_d(q)$. Assume the availability of the oracles specified above. If $q \geq 4$, then there is a black-box Las Vegas polynomial-time algorithm which constructs, as* SLPs *in X, standard generators for G. The time required by the algorithm is $O(d \log d(\mu + \xi + \mathcal{O} + \Pi) + d((\chi + \mu) \log^2 q + \xi \log q \log \log q))$.*

With minor modifications, which we identify in Section 4, the algorithm works well for $q = 3$; our algorithm does not apply to $q = 2$.

1.3. Rewriting. A black-box algorithm, with complexity $O(d^2 q)$ group operations, to write an element of G as an SLP in the standard generators was developed by Ambrose *et al.* [1]; recently, Schneider extended this result to cover missing cases. We know of no polynomial-time algorithm to perform this task.

The implementation of our algorithm accepts as input either a permutation or linear representation of $\mathrm{SX}_d(q)$. If the input is an absolutely irreducible representation in defining characteristic, then we use the polynomial-time algorithm of Costi [18] to perform the rewriting task. Using standard algorithms for modules, we reduce arbitrary matrix representations in defining characteristic to this case. Recall, from [31], that a faithful linear or projective representation of a finite group of Lie type in cross characteristic has degree which is polynomial in q. Hence, all other input representations have size $O(q)$; so, in these cases, the extension to [1] runs in time polynomial in the size of the input.

1.4. Related work. Kantor & Seress [24] developed the first black-box constructive recognition algorithms for classical groups. The complexity of these algorithms involves a factor of q. By assuming the availability of the oracle χ, Brooksbank and Kantor [11–14] present algorithms with complexity polynomial in d, $\log q$, and the number of calls to χ.

These algorithms construct Steinberg generators for the group, so the generating set returned has size $O(d^2 \log q)$ and requires significant storage. In practical applications, when we use the methods of COMPOSITIONTREE [7], we work with groups having classical groups as homomorphic images and construct kernels to these homomorphisms; now a small fixed number of standard generators is useful.

Table 1 lists the principal contributors to the stated complexity of each the algorithms of [11–14] and also the comparable costs of our algorithm. In Section 7 we discuss the cost of these oracles, and our additional two, $\mathcal{O}$ and Π.

Algorithm	ξ	χ	μ
SL [11]	$d^2 \log q$	$d^3 \log d \log q$	$d^4 \log d \log^3 q$
SU [12]	$d^2 \log d$	$d^2 \log d \log q$	–
Ω^ϵ [13]	$d^2 \log d \log q$	$d^2 \log d \log^2 q$	$d^3 \log^2 d$
Sp [14]	$d + \log q$	1	$d^2 \log^2 q$
Ours	$d(\log d + \log q)$	$d \log^2 q$	$d(\log d + \log^2 q)$

TABLE 1. Coefficients of oracles in the complexity of the algorithms

2. Structure of the general algorithm

Our black-box algorithm follows the general approach of our algorithms for the natural representation described in [19, 26]. Let $G = \langle X \rangle$ be isomorphic to a central quotient of a classical group $\mathrm{SX}_d(q)$. We use a recursion to construct standard generators $\mathcal{S}_G$ of G as SLPs in X. The base cases of this recursion are discussed in Section 3.1; in the following, suppose that G is not a base case.

For odd q, find, by random search, an element of even order that powers to an involution $g \in G$ which corresponds to an element in $\tilde{G}$ with -1- and 1-eigenspaces of dimension $m \in [d/3, 2d/3]$ and $d - m$, respectively. In the centraliser of g in G, construct two commuting subgroups $H, K \leq G$ with $H \cong \mathrm{SX}_m(q)$ and $K \cong \mathrm{SX}_{d-m}(q)$. Using recursion, construct the standard generators $\mathcal{S}_H$ and $\mathcal{S}_K$ of H and K, respectively. With the exception of the cycle of G, all standard generators of G lie in $\mathcal{S}_H \cup \mathcal{S}_K$. The cycle of G is constructed by *gluing* the cycles in $\mathcal{S}_H$ and $\mathcal{S}_K$.

For even q, find, by random search, an element that powers to $g \in G$ which is the image of an element in $\tilde{G}$ with 1-eigenspace of dimension in $[2d/3, 5d/6]$, acting irreducibly on a complement. By taking g and a random conjugate h of g in G, construct $H = \langle g, h \rangle \leq G$ isomorphic to $\mathrm{SX}_m(q)$ with $m \in [d/3, 2d/3]$. Using recursion, construct the standard generators $\mathcal{S}_H$ of H and a specific involution $i \in H$. In $C_G(i)$, find $K \leq G$ which is isomorphic to $\mathrm{SX}_{d-m}(q)$ and commutes element-wise with H. By recursion, construct the standard generators $\mathcal{S}_K$ of K, and, finally, glue the cycles in $\mathcal{S}_H$ and $\mathcal{S}_K$.

To ensure that the algorithm is Las Vegas in the natural representation is easy: modulo a (known) base change, the standard generators returned are identical to those listed in [19, Table 1] and [26, Tables 1 & 2]. To establish this for the black-box algorithm is more challenging. That groups of Lie type have *short presentations* was first established by Guralnick *et al.* [22]; explicit short presentations, on our standard generators, for the classical groups appear in [27]. By evaluating the standard presentation of $\mathrm{SX}_d(q)$ in the output of our algorithm, $\mathcal{S}_G$, we verify the correctness of our result.

The main challenge in developing the black-box algorithm was to devise a strategy for gluing the cycles. Other difficulties arise in the construction of the two smaller subgroups for the recursion.

The remainder of the paper is as follows. In Section 3, we recall some preliminary results. In Sections 4 and 5, we describe the construction of the two smaller subgroups H and K for odd and even q, respectively. In Section 6, we discuss how to glue the cycles of H and K; this completes the construction of the standard generators of G. The complexity of our algorithm is discussed in Section 7. We comment on our implementation in Section 8.

3. Preliminaries

3.1. Base cases. If G is isomorphic to a (central quotient of a) classical group of small rank, then we treat it as a base case.

Definition 3.1. The *base cases* for even q are $\mathrm{SL}_d(q)$ with $d \leq 5$; $\mathrm{SU}_d(q)$ with $d \leq 7$; $\mathrm{Sp}_d(q)$ with $d \leq 6$; $\Omega_d^+(q)$ with $d \leq 8$; and $\Omega_d^-(q)$ with $d \leq 10$. The base cases for odd q are $\mathrm{SL}_d(q)$, $\mathrm{SU}_d(q)$, and $\mathrm{Sp}_d(q)$ with $d \leq 4$; $\Omega_d(q)$ with $d \leq 5$; $\Omega_d^\pm(q)$ with $d \leq 6$; and $\Omega_7(q)$ and $\Omega_8^\pm(q)$ with $q \equiv 3 \bmod 4$.

The next theorem follows from [11–14, 16, 30].

Theorem 3.2. *Let G be isomorphic to a central quotient of a base case group $\mathrm{SX}_d(q)$. There is a black-box Las Vegas algorithm that constructively recognises G. Subject to the existence of the relevant oracles identified in Section 1.2, the algorithm runs in time $O((\chi + \mu)\log^2 q + \xi \log q \log\log q)$.*

In practice, we sometimes employ algorithms other than those cited above to deal with base cases.

3.2. Automorphism groups of classical groups. The following facts are well-known, see [40, p. 192 & Proposition 13.11] and [21, Sec. 2.2, 2.5, 2.7].

Remark 3.3. a) The universal versions of the finite classical groups are $\mathrm{SL}_d(q)$, $\mathrm{SU}_d(q)$, $\mathrm{Sp}_{2n}(q)$, $\mathrm{Spin}_{2n}^\pm(q)$, and $\mathrm{Spin}_{2n+1}(q)$. If H is one of these, then $H/Z(H)$ is the adjoint version. If $H/Z(H)$ is simple, then $\mathrm{Aut}(H) \cong \mathrm{Aut}(H/Z(H))$; every automorphism of H can be written as a product of a graph, field, diagonal, and inner automorphism.

b) Let $H = \mathrm{SL}_d(q)$. Then $H/Z(H)$ is simple, the diagonal automorphisms of H are induced by conjugation with diagonal matrices in $\mathrm{GL}_d(q)$, and field automorphisms are induced by the usual Frobenius action on matrix entries. If $d = 2$, then there is no graph automorphism; if $d > 2$, then the graph automorphism is the inverse-transpose.

c) Let $H = \mathrm{SU}_d(q)$ and $d \geq 3$. Then $H/Z(H)$ is simple, the diagonal automorphisms are induced by conjugation with diagonal matrices in $\mathrm{GU}_d(q)$, and there are no graph automorphisms. Field automorphisms act on matrix entries. Recall that $\mathrm{SU}_2(q) \cong \mathrm{SL}_2(q)$.

d) Let $H = \mathrm{Sp}_d(q)$ and $d \geq 4$. Then $H/Z(H)$ is simple and field automorphisms act on matrix entries. If q is even, then H has no diagonal automorphisms; H has a non-trivial graph automorphism (of order 2) only if $d = 4$. If q is odd, then the diagonal automorphisms are induced by conjugation with elements of the conformal group, and H has no graph automorphism.

e) Let $H = \Omega_d^+(q)$ with q even and $d \geq 6$. Then H is simple, field automorphisms act on matrix entries, and there are no diagonal automorphisms. If $d \geq 6$ and $d \neq 8$, then $|\mathrm{Out}(H)| = 2e$ where $q = 2^e$; there is a graph automorphism of order 2, induced by conjugation by a certain permutation matrix, see [33, p. 194]. If $d = 8$, then $|\mathrm{Out}(H)| = 6e$ where $q = 2^e$; there are graph automorphisms of order 2 and 3. If $d = 4$, then $\Omega_4^+(q) = \mathrm{SL}_2(q) \times \mathrm{SL}_2(q)$. The graph automorphism swaps the two factors, and for each $\mathrm{SL}_2(q)$ there are field automorphisms; thus, $|\mathrm{Out}(\Omega_4^+(q))| = 2e^2$.

f) Let $H = \Omega_d^-(q)$ with q even and $d \geq 4$. Then H is simple and there are no graph or diagonal automorphisms. Field automorphisms are induced by the usual action on matrix entries followed by conjugation by some matrix in $\mathrm{GL}_d(q)$; thus, $|\mathrm{Out}(H)| = 2e$ where $q = 2^e$.

g) Let $H = \Omega_d(q)$ with both d and q odd. Then H is simple, field automorphisms act on matrix entries, there is no graph automorphism, and $|\mathrm{Out}(H)| = 2e$ where $q = p^e$; there is a diagonal automorphism of order 2.

h) Let $H = \Omega_d^{\pm}(q)$ with $d \geq 4$ even and q odd. If $H \neq \Omega_4^+(q)$, then $K = H/Z(H) = \mathrm{P}\Omega_d^{\pm}(q)$ is simple, and we can identify $\mathrm{Aut}(H)$ with $\mathrm{Aut}(K)$. The automorphisms of K are as for $\Omega_d^{\pm}(q)$ with q even, with two exceptions. There are diagonal automorphisms, and there is no graph automorphism of order 3. The automorphisms of $\Omega_4^+(q) = \mathrm{SL}_2(q) \circ \mathrm{SL}_2(q)$ are as for even q, with the exception that diagonal automorphisms exist.

For an integer m let 1_m be the $m \times m$ identity matrix.

Lemma 3.4. *Let $G = \mathrm{SX}_d(q)$ and let $H \cong \mathrm{SX}_m(q)$ with m even such that*

$$H = \begin{pmatrix} \mathrm{SX}_m(q) & 0 \\ 0 & 1_{d-m} \end{pmatrix} \leq G.$$

Suppose that either G and H have the same type, or q is even and H has type Ω^+ and G is orthogonal or symplectic. With some exceptions for $H \cong \Omega_4^+(q)$, and $H \cong \mathrm{Sp}_4(q)$ and $H \cong \Omega_8^+(q)$ with q even, every automorphism of H lifts to an automorphism of G.

PROOF. This follows from Remark 3.3; note that $\alpha \in \mathrm{Aut}(H)$ does not lift if its decomposition into an inner, diagonal, field, and graph automorphism contains a graph automorphism of $\mathrm{Sp}_4(q)$, a graph automorphism of $\Omega_8^+(q)$ of order 3, or a field automorphism of $\Omega_4^+(q)$ which acts differently on the two $\mathrm{SL}_2(q)$ factors. □

3.3. Involution centralisers. If G is a central quotient of $\mathrm{SX}_d(q)$, then the centraliser $C_G(i)$ of an involution $i \in G$ can be constructed using an algorithm of Bray [10]. If $g \in G$, then $[i, g]$ either has odd order $2k + 1$, in which case $g[i, g]^k$ commutes with i, or has even order $2k$, in which case both $[i, g]^k$ and $[i, g^{-1}]^k$ commute with i. If g is random among the elements of G for which $[i, g]$ has odd order, then $g[i, g]^k$ is random in $C_G(i)$, see [39, Theorem 11]. That such *Bray generators*, $g[i, g]^k$, of $C_G(i)$ can be constructed follows from the next theorem established in [28, 39].

Theorem 3.5. *There is a constant $c > 0$ such that if $i \in G$ is an involution and G is a central quotient of $\mathrm{SX}_d(q)$, then the proportion of $g \in G$ with $[i, g]$ of odd order is bounded below by c/d.*

To construct a Bray generator, we apply the order and power oracles to a random element.

3.4. Zsigmondy primes. Recall that if q is a prime-power and $l > 0$, then a (q, l)-*Zsigmondy prime* r is one that divides $q^l - 1$ but not $q^i - 1$ for $i < l$. Such primes exist, except for $(q, l) = (2, 6)$ and $(q, l) = (q, 2)$ with q a Mersenne prime. If an order oracle for $G \cong \mathrm{SX}_d(q)$ is available, then repeated computations of the form $\gcd(q^i - 1, |g|)$ yield all l and r such that r is a (q, l)-Zsigmondy prime dividing $|g|$. If a (q, l)-Zsigmondy prime divides $|g|$, then g is a *ppd(q, l) element.*

Every semisimple element in $G = \mathrm{SX}_d(q)$ lies in a maximal torus; the structure of these tori is known, see for example [36, Sec. 3]. If G is linear or unitary, then its maximal tori are isomorphic to

$$[(q^{e_1} - (-1)^{\varepsilon}) \times \ldots \times (q^{e_k} - (-1)^{\varepsilon})]/(q - (-1)^{\varepsilon}),$$

where $(e_1, \ldots, e_k)$ is a partition of d, each $q^e \pm 1$ denotes a cyclic group of that order, and $\varepsilon = 1$ (or -1) if G is linear (or unitary). If $G = \mathrm{Sp}_{2n}(q)$ or $G = \Omega_{2n+1}(q)$, then the maximal tori are

$(q^{e_1}+1)\times\ldots\times(q^{e_k}+1)\times(q^{f_1}-1)\times\ldots\times(q^{f_j}-1)$ where $(e_1^+,\ldots,e_k^+,f_1^-,\ldots,f_j^-)$ is a signed partition of n. The maximal tori for $\Omega_{2n}^{\pm}(q)$ are the same, with k even for Ω^+, and k odd for Ω^-. Observe that if C is cyclic of order n and p is a prime dividing n, then at least $1-1/p$ of all elements in C have order divisible by p. Hence, if T is a maximal torus containing a direct factor q^e-1 with $e>1$ and (q,e)-Zsigmondy primes exist, then the proportion of $\mathrm{ppd}(q,e)$ elements in T is at least $2/3$; a similar observation holds for q^e+1 and $\mathrm{ppd}(q,2e)$ elements.

We now summarise easy but important consequences of properties of $\mathrm{ppd}(q,e)$ elements as discussed in [35]; to obtain the stated proportions, using [36], we count the number of tori (up to conjugacy) with suitable direct factors.

Remark 3.6. a) A subgroup H of $\mathrm{SL}_d(q)$ is irreducible if H contains a $\mathrm{ppd}(q,d)$ element, or if it contains two elements g_1 and g_2 such that each g_j is a $\mathrm{ppd}(q,e_j)$ and $\mathrm{ppd}(q,d-e_j)$ element, where $e_j\le d-e_j$, and e_j does not divide $d-e_j$, and $\{e_1,d-e_1\}\ne\{e_2,d-e_2\}$. The analogous result holds for other classical groups. The proportion of such elements in $\mathrm{SX}_d(q)$ is $O(1/d)$.

b) A subgroup of an orthogonal or symplectic $\mathrm{SX}_d(q)$ with $d=2n$ does not preserve a quadratic form of $+$ type if it contains a $\mathrm{ppd}(q,d)$ element; it does not preserve a quadratic form of $-$ type if it contains a $\mathrm{ppd}(q,d-2)$ element of order not dividing $(q^{n-1}+1)(q-1)$. The proportion of such elements in $\mathrm{SX}_d(q)$ is $O(1/d)$.

c) A subgroup of $\mathrm{SL}_d(q)$ does not preserve a bilinear form if it contains a $\mathrm{ppd}(q,e)$ element with odd $e>d/2$; it does not preserve a sesquilinear form if it contains a $\mathrm{ppd}(q,e)$ element with even $e>d/2$. The proportion of such elements in $\mathrm{SL}_d(q)$ is $O(1/d)$.

4. Two smaller subgroups in odd characteristic

As outlined in Section 2, our algorithm for constructive recognition in the natural representation [26] carries over readily to a black-box algorithm, with the exception of gluing the cycles. We describe gluing in Section 6; here we comment on the construction of the subgroups used for the recursion.

Let G be isomorphic to a central quotient of $\tilde G=\mathrm{SX}_d(q)$ with $q>3$ odd. If $i\in\tilde G$ is an involution with ±1-eigenspaces $E_\pm$, then

$$C_{\tilde G}(i)=(\mathrm{GX}(E_+)\times\mathrm{GX}(E_-))\cap\tilde G,$$

where $\mathrm{GX}(E_\pm)$ is the general linear, general unitary, symplectic, or orthogonal group acting on $E_\pm$. If j is the image of i in G, then $C_G(j)$ is the image in G of $C_{\tilde G}(i)$, unless $\mathrm{GX}(E_+)\cong\mathrm{GX}(E_-)$, and the images of i and $-i$ in G are equal, in which case $C_G(j)$ is the image of $C_{\tilde G}(i)$ extended by the image of a 2-cycle that interchanges E_+ and E_-. In [26], we call i a *strong involution* if $d/3<\dim(E_-)\le 2d/3$. Here we allow $d/3\le\dim(E_-)\le 2d/3$ so that i is a strong involution if and only if $-i$ is; this has negligible side effects. An involution in G is *strong* if it is the image of a strong involution in $\tilde G$.

If $i\in G$ is a strong involution, then $C_G(i)''$, the second derived subgroup of $C_G(i)$, is isomorphic to a central quotient of $\mathrm{SX}_e(q)\times\mathrm{SX}_{d-e}(q)$ with $d/3\le e\le 2d/3$. We now describe how to construct these direct factors as subgroups of $C_G(i)$.

Theorem 4.1. *Let $G=\langle X\rangle$ be a central quotient of $\mathrm{SX}_d(q)$ for $d\ge 6$ and odd $q>3$. There exists a black-box Las Vegas algorithm to construct a strong involution $i\in G$, and to find generating sets for A_1 and A_2, where the generalised Fitting subgroup $F^*(C_G(i))=C_G(i)''$ is a central quotient of $\mathrm{SX}_e(q)\times\mathrm{SX}_{d-e}(q)$, and A_1 and A_2 are the images of $\mathrm{SX}_e(q)$ and $\mathrm{SX}_{d-e}(q)$. The algorithm also returns the names of these two classical groups. If G is orthogonal of $+$ type, then i is chosen such that A_1 and A_2 have $+$ type. The algorithm runs in time $O(d\log d(\mu+\xi+\mathcal{O}+\Pi))$.*

PROOF. The restriction on d ensures that $F^*(C_G(i))$ is a central quotient of the direct product of two perfect groups.

We prove the theorem by exhibiting an algorithm which has the claimed complexity. Suppose first that G is isomorphic to a central quotient of $\mathrm{SL}_d(q)$.

(1) By a random search, find $g \in G$ of even order; set $i = g^{|g|/2}$ and $S = \{g\}$.

(2) Construct three Bray generators of $C_G(i)$ and place them in S.

(3) Construct random elements of $\langle S\rangle$, looking for two elements that power to elements a_1 and a_2 satisfying the following two conditions: first, each a_j is a $\mathrm{ppd}(q, e_j)$ element and $e_1 + e_2 = d$; second, if b_j is a random $\langle S\rangle$-conjugate of a_j, then $\langle a_1, b_1\rangle$ and $\langle a_2, b_2\rangle$ commute. If $e_1 \notin [d/3, 2d/3]$, then i is not a strong involution and we return to Step (1). If after $O(d)$ trials no such elements are found, then repeat Step (2) and then (3).

(4) Set $T_1 = \{a_1, b_1\}$ and $T_2 = \{a_2, b_2\}$, and, to ease exposition, suppose $T_j \leq A_j$. For $g \in C_G(i)$ we check membership in A_1 and A_2 by checking commutativity with $\langle T_2\rangle$ and $\langle T_1\rangle$, respectively. We decompose $g \in \langle S\rangle$ as $g = g_1g_2g_3g_4$ where each g_j is a power of g of largest possible order such that $|g_1|, |g_2|, |g_3|$ are pairwise coprime and none divides $q - 1$, and $g_1 \in A_1$, $g_2 \in A_2$, $g_3 \notin A_1 \cup A_2$, and $|g_4|$ divides $q - 1$. Taking random $g \in \langle S\rangle$ and adding its component g_j to T_j for $j = 1, 2$, we seek *witnesses* (as in [35]) to establish that $\langle T_j\rangle = A_j$. If this fails, then repeat Steps (2)–(4), and continue. The presence of these witnesses, which are returned by the procedure, proves that the algorithm has terminated correctly.

We now supply further details for these steps, and assess the complexity of the algorithm.

$(1')$ By [29], a strong involution $i \in G$ is found after $O(\log d)$ repetitions of Step (1); thus, we expect to return to this step $O(\log d)$ times at a cost of $O(\log d(\xi + \mathcal{O} + \Pi))$.

$(2')$ A sample of $O(d)$ random elements yields a Bray generator. It is proved in [34, Corollary 1.2] that the probability that 3 random elements of a finite almost simple group K, conditional on them generating $K/F^*(K)$, generate K is greater than 139/150. As observed in [38, Theorem 4.1], the probability that $k + 1$ random elements of a finite abelian k-generator group generate the group is greater than $1/2.72$. Since $C_G(i)$ is an extension of a central quotient of $\mathrm{SL}_e(q) \times \mathrm{SL}_{d-e}(q)$ by a cyclic group (or by a dihedral group when $d = 2e$ and $i = -i$), the probability that $\langle S\rangle = C_G(i)$, with S as in Step (2), is bounded away from 0 by an absolute positive constant. In particular, the probability that S generates a group containing $F^*(C_G(i))$ is very high (observe S contains g as well as the Bray generators). The expected number of returns to Step (2) is $O(\log d)$, at the cost of $O(d \log d(\xi + \mu + \mathcal{O} + \Pi))$.

$(3')$ Using the notation of the theorem, we may assume that $e_1 = e$, so $e_2 = d - e$. Recall, from [35, Theorem 5.7], that the probability that an element of $\mathrm{SL}_f(q)$ is a $\mathrm{ppd}(k, q)$ element is approximately $1/k$ where $f/2 < k \leq f$. Thus the proportion of elements of $F^*(C_G(i))$ that power to a candidate for a_1 is approximately $1/e_1$, or $(1/e_1)(1-1/e_1)$ if $e_2 \geq e_1$, and similarly for a_2. If we find $a_j, b_j \in C_G(i)$ with the stated properties, then we can suppose $\langle a_j, b_j\rangle \leq A_j$; the probability of this being false is exponentially small. The total cost of Step (3) is as in Step (2); the factor of $\log d$ arises as we may have to return to this step for $O(\log d)$ involutions.

$(4')$ Elements of T_j have order coprime to $q - 1$, thus lie in a central quotient of $\mathrm{SL}_{e_j}(q)$. We first seek witnesses to show that $G_j = \langle T_j\rangle$ is not a central quotient of a classical group that preserves a form. If e_j is even, then we rule out the possibility that G_j is an image of a symplectic or orthogonal group by finding a $\mathrm{ppd}(q, k)$ element for some odd k greater than $e_j/2$; similarly, if q is a square, then we rule out the possibility that G_j is the image of the unitary group by finding a $\mathrm{ppd}(q, k)$ element for some even $k > e_j/2$. We are now in a position in which we can, in principle, apply the algorithm of [35]. That algorithm applies to a subgroup K of $\mathrm{SL}_n(q)$, in its natural representation, where K is known to act irreducibly, and to preserve no non-zero form. The algorithm seeks witnesses to prove

$K = \mathrm{SL}_n(q)$ by virtue of their orders. The witnesses are constructed by a random process, and the algorithm uses only ppd information. Here we have a central quotient of $\mathrm{SL}_{e_j}(q)$ rather than the group itself, but this does not harm the validity of the algorithm.

The remaining issue is that the elements that we place in the generating sets T_j do not approximate to a random distribution. However the probability that g_j, as in (4), is a $\mathrm{ppd}(q,k)$ element approximates closely to the probability that a random element of A_j is a $\mathrm{ppd}(q,k)$ element. For smaller values of k, the probability is slightly reduced because the chances that an element of $C_G(i)$ will map to an element of order a multiple of a given (q,k)-Zsigmondy prime in both components is slightly increased. Since the algorithm in [35] seeks $\mathrm{ppd}(q,k)$ elements for large values of k, this is not a problem. It needs $O(\log\log d)$ random elements to find the required witnesses, so Step (4) is asymptotically faster than Steps (2) and (3). Once these witnesses have been found (and witnesses for one factor all commute with the witnesses of the other), then we have proved that the algorithm has run correctly: based on element orders, the groups generated by these witnesses are not isomorphic to central quotients of *proper* subgroups of $\mathrm{SL}_{e_j}(q)$, thus, they must be central quotients of $\mathrm{SL}_{e_j}(q)$.

Now suppose that G is a central quotient of $\Omega^+_{2n}(q)$. New difficulties arise. Firstly, A_1 or A_2 may be an image of $\Omega^+_4(q)$; secondly, we must reject the involution i if its centraliser is a central quotient of two orthogonal groups of $-$ type; finally, we cannot choose the elements a_1 and a_2 to act irreducibly on the respective direct factors because such elements do not exist. The impact of the first is limited to a minor change in the associated statistics. The others we address by seeking elements with one of the following sets of properties.

(a) There exist even integers e_1 and e_2 with $e_1 + e_2 = 2n$, and integers $u_1 \neq v_1$ and $u_2 \neq v_2$, and elements a_1, a_2, b_1, and b_2 are found such that a_j has order the product of a (q,u_j)-Zsigmondy prime and a $(q, e_j - u_j)$-Zsigmondy prime, and b_j has order the product of a (q,v_j)-Zsigmondy prime and a $(q, e_j - v_j)$-Zsigmondy prime, and a_1 and b_1 both commute with a_2 and b_2, cf. Remark 3.6. These elements are sought by powering up random elements of $\langle S\rangle$. Again it is almost certain that a_1 and b_1 correspond to elements of one factor, and that a_2 and b_2 correspond to elements of the other. Also, a_1 and b_1, together, serve as irreducibility witnesses (as did a_1 alone in the special linear case), and also as witnesses to the fact that they generate a subgroup of $\Omega^+_{e_1}(q)$, as opposed to $\Omega^-_{e_1}(q)$; similarly for a_2 and b_2. Thus the algorithm proceeds as before.

(b) Elements are found that power to $\mathrm{ppd}(q, e_j)$ elements a_j, $j = 1, 2$, where $e_1 + e_2 = 2n$, and a_1 commutes with a_2 and a random conjugate of a_2. Now a_1 and a_2 are witnesses that $F^*(C_G(i))$ is a central factor of the direct product of two groups of type Ω^-, and the involution i is rejected. As pointed out in [26, Lemma 2.2], we fall into the Ω^- case if and only if both $q \equiv 3 \bmod 4$ and $e_j \equiv 2 \bmod 4$.

The proportion of elements of $\Omega^+_{e_1}(q)$ satisfying the order condition imposed in (a) is $O(\log d/d)$. However, if, in the notation of (a), either u_1 or $e_1 - u_1$ is small, then the probability that a random element of $\Omega^+_{e_2}(q)$ has order a multiple of this prime tends (slowly) to 1 as e_2 tends to infinity. But consider large d: if we just count the cases that arise when $e_1/3 \leq v_1 \leq 2e_1/3$, then the proportion of elements of $\Omega^+_{e_1}(q)$ of the appropriate order remains $O(\log d/d)$, and, because e_1 and e_2 are of comparable size, the probability that a random element of $\Omega^+_{e_2}(q)$ has order a multiple of one of the relevant primes is bounded away from zero by an absolute positive constant. The requisite proportions are given, to more accuracy than required here, in [26, Section 8].

The other groups are dealt with in the same style. □

The algorithm of Theorem 4.1 may be trivially extended to deal with smaller values of d, provided that i is chosen so that $F^*(C_G(i))$ is a central quotient of the direct product of two perfect groups.

In practice, the steps in this algorithm can run faster by applying various simple devices, such as using conjugation to generate new elements of the T_j. Theoretically, the most expensive step of the algorithm is (2): we must test $O(d)$ random elements to obtain a Bray generator of the involution.

Recall [5, Corollary 4.2]: if p is a prime and G is a finite simple classical group acting naturally on a projective space of dimension $d-1$, then the proportion of p-regular elements in G is at least $1/2d$.

Remark 4.2. In the gluing process, we deal with the following situation: the involution $i \in G$ is not strong and, using the previous notation, A_1 and A_2 are quotients of $\mathrm{SX}_e(q)$ and $\mathrm{SX}_{d-e}(q)$ with $e \leq 6$. In contrast to the above discussion, this time e is known, and we only want to construct A_1. We proceed as follows. The first step is to use a modification of Theorem 4.1 to construct $B \leq A_2$ with $C_{A_2}(B) \leq Z(A_2)$, for example, $B = A_2$. Since $e \leq 6$ is small, elements in B can in general be readily constructed by taking random Bray generators of $C_G(i)$ to the power $\exp(\mathrm{SX}_e(q))$, cf. [5, Corollary 4.2]. Observe that $h \in C_G(i)$, of order not dividing $|Z(A_2)|$, lies in A_1 if and only if $[h, b] = 1$ for every generator b of B. Using this, we find a non-central $h \in A_1$, and construct A_1 as the normal closure of h in $C_G(i)$ by applying the algorithm of [42, Theorem 2.3.9]; we use Remark 3.6 and [35] to verify the correctness of our computation.

Remark 4.3. The case $q = 3$ requires special care, here and in gluing (see Section 6). The principal reason is that one of the factors A_j may be soluble. In all other important respects, the algorithm is identical with that for larger odd q, and displays similar performance.

5. Two smaller subgroups in even characteristic

Throughout this section, let $q \neq 2$ be even and let $G = \langle X \rangle$. To simplify exposition, we assume that G is isomorphic to $\mathrm{SX}_d(q)$, and not to an arbitrary central quotient. We also assume that G is *not* a base case. Let $\varphi\colon G \to \tilde{G}$ be an (unknown) isomorphism to the standard copy $\tilde{G}$ of $\mathrm{SX}_d(q)$, with underlying field $\mathbb{F}$. The aim of this section is to construct, as SLPs in X, generators for commuting subgroups $H \cong \mathrm{SX}_m(q)$ and $K \cong \mathrm{SX}_{d-m}(q)$ of G, where, in general, $m \in [d/3, 2d/3]$.

5.1. Constructing the first subgroup.

In [19, Sec. 5], we devised an algorithm to construct $\tilde{H} \leq \tilde{G}$ with $\tilde{H} \cong \mathrm{SX}_m(q)$. In general, $m \in [d/3, 2d/3]$ is even, and $\tilde{H}$ has the same type as $\tilde{G}$; if $\tilde{G}$ is symplectic or orthogonal, then m is divisible by 4 and $\tilde{H}$ has type Ω^+.

We briefly recall this construction. By a random search, find $g \in \tilde{G}$ that powers to $h \in \tilde{G}$ which has a 1-eigenspace of dimension $e \in [2d/3, 5d/6]$ and acts irreducibly on a complement. A construction of $O(1)$ random elements of $\tilde{G}$ suffices to find u so that $\tilde{H} = \langle h, h^u \rangle \cong \mathrm{SX}_m(q)$ with $m = 2(d-e)$; more precisely, modulo a base change,

$$\tilde{H} = \begin{pmatrix} \mathrm{SX}_m(q) & 0 \\ 0 & 1_{d-m} \end{pmatrix} \leq \tilde{G}.$$

Motivated by that approach, we now develop a black-box algorithm to construct $H \leq G$ with $H \cong \mathrm{SX}_m(q)$ and $\varphi(H) = \tilde{H}$ as above. We seek $g \in G$ whose order is divisible by two Zsigmondy primes, say p and r satisfying (Z1) and (Z2) below, which witness that the image of $g^{|g|/p}$ in $\tilde{G}$, firstly, acts irreducibly on a subspace of dimension $i \in [d/6, d/3]$ and, secondly, acts trivially on a complement to this space.

If $G \cong \Omega_d^-(q)$, then we seek $H \cong \Omega_m^+(q)$ with $m \in \{d-4, d-6\}$ divisible by 4. While our algorithm is capable of constructing subgroups of other ranks, this restriction arises from gluing; we explain this in more detail in Remark 6.1.

We start with an easy observation.

Lemma 5.1. *Let $h \in \mathrm{GL}_d(q)$ have an e-dimensional 1-eigenspace. If h is a ppd$(q, d-e)$ element, then h acts irreducibly on a complement to its 1-eigenspace.*

We now describe the construction of H in detail for SL. Let $\varphi\colon G \to \tilde{G} = \mathrm{SL}_d(q)$. By a random search, find $g \in G$ such that

(Z1) $|g|$ is divisible by a (q,i)-Zsigmondy prime p with $i \in [d/6, d/3]$;

(Z2) $|g|$ is divisible by a (q,e)-Zsigmondy prime r with $i \nmid 2e$ and $e + 2i > d$.

We can also assume that g has odd order; otherwise, replace g by g^c where c is the smallest 2-power satisfying $c \geq 2d-2$; now g is semisimple and lies in some maximal torus of G.

Lemma 5.2. *If $g \in G \cong \mathrm{SL}_d(q)$ satisfies* (Z1) *and* (Z2)*, then the image of $h = g^{|g|/p}$ in $\tilde{G}$ has a 1-eigenspace of dimension $d - i \in [2d/3, 5d/6]$ and acts irreducibly on a complement.*

PROOF. Suppose that g lies in a maximal torus $S = (q^k - 1) \times S^*$ and k is divisible by both e and i. Since $e > d/3$, we know that $k \in \{e, 2e\}$; now $i \mid k$ yields a contradiction to $i \nmid 2e$. Thus, g must lie in a maximal torus $T = (q^j - 1) \times (q^f - 1) \times T^*$ with $i \mid j$ and $e \mid f$; as shown above, $i \nmid f$ and, by assumption, $p \nmid |T^*|$. Since $e + 2i > d$, it follows that $j = i$, hence $g \in T = (q^i - 1) \times (q^f - 1) \times T^*$ and $p \nmid |(q^f - 1) \times T^*|$. Thus, the image of $h = g^{|g|/p}$ in $\tilde{G}$ has a 1-eigenspace of dimension $d - i$ and acts irreducibly on the i-dimensional space associated with $q^i - 1$, see Lemma 5.1. □

Suppose we have found $g \in G$ satisfying (Z1) and (Z2), and set $h = g^{|g|/p}$. As outlined in [19], it follows from [41] that the construction of $O(1)$ random elements of G suffices to find u such that $H = \langle h, h^u \rangle \cong \mathrm{SL}_m(q)$, where $m = 2i \in [d/3, 2d/3]$. We could use [4] to verify that $H \cong \mathrm{SL}_m(q)$. More efficiently, we proceed as follows. First, we use Remark 3.6 and consider a sample of $O(d)$ random elements in H until we find witnesses that H does not preserve a bilinear or sesquilinear form, and it acts irreducibly on an m-dimensional space. Then we apply [35] as in Theorem 4.1 and seek witnesses that H is isomorphic to $\mathrm{SL}_m(q)$. If we cannot find these witnesses, then we construct another H; only $O(1)$ repetitions are required.

The strategy for unitary, symplectic, and orthogonal types is similar. One change is due to the different structure of maximal tori, which requires an adjustment of the Zsigmondy prime divisors we seek; for small d, this requires specialised techniques. A second is that the isomorphism type of $H \cong \mathrm{SX}_m(q)$ is not uniquely determined if G is orthogonal or symplectic; both $\Omega^+_m(q)$ and $\Omega^-_m(q)$ are possible. We use Remark 3.6 to detect one or the other. If we confirm $H \cong \Omega^-_m(q)$, then we constructively recognise H and replace H by $H^\star \leq H$ with $H^\star \cong \Omega^+_{m-4}(q)$; having constructively recognised H, we can write down generators for $H^\star$. (Alternatively, we could construct a new group until $H \cong \Omega^+_m(q)$.)

We now analyse the complexity of the resulting algorithm.

Lemma 5.3. *There is a black-box Las Vegas algorithm which takes as input $G \cong \mathrm{SX}_d(q)$, which is not a base case, and constructs $H \leq G$ with $H \cong \mathrm{SX}_m(q)$, admitting $K \leq C_G(H)$ with $K \cong \mathrm{SX}_{d-m}(q)$; in general, $m \in [d/3, 2d/3]$ is even. If G is linear or unitary, then so is H. In all other cases, H is of type Ω^+ and m is divisible by 4. If G has type Ω^-, then $m \in \{d-4, d-6\}$ is divisible by 4. The time required is $O(d(\xi + \mathcal{O}) + \Pi + \mu)$.*

PROOF. The correctness of the algorithm is established in [19, Sec. 5]; it remains to show that the construction of $O(1)$ random elements in G is sufficient to find $g \in G$ satisfying (Z1) and (Z2) above. (If G has type Ω^-, then we show that $O(d)$ random elements suffice.) We give the proof in detail for SL and Ω^-; the remaining cases are dealt with analogously. In the following, we assume that d is large enough so that all required Zsigmondy primes exist and intervals are non-empty. We use Remark 3.6 and [35] as in Theorem 4.1 to verify that the output of this algorithm, H, satisfies $H \cong \mathrm{SX}_m(q)$; this verification dominates the overall complexity.

First, suppose $G \cong \mathrm{SL}_d(q) = \tilde{G}$. Let $\hat{G}$ be a simply connected reductive algebraic group such that $\hat{G}^F = \tilde{G}$ for some Frobenius morphism F. Call $g \in \tilde{G}$ admissible if some power of g has odd order and satisfies (Z1) and (Z2) above. Let $A(\tilde{G})$ be the set of all admissible $g \in G$. By the multiplicative Jordan decomposition, every $g \in \tilde{G}$ can be written uniquely as $g = su$ where $s \in \tilde{G}$ is semisimple, $u \in \tilde{G}$ is unipotent, and $su = us$. Since $g = su \in A(\tilde{G})$ if and only if $s \in A(\tilde{G})$, and $A(\tilde{G})$ is invariant under conjugation, we can apply [36, Theorem 1.3] to estimate the proportion $|A(\tilde{G})|/|\tilde{G}|$. For convenience, we recall this result here.

Let $\mathbb{F}$ be the underlying field of $\tilde{G}$. Let $T_0 \leq \tilde{G}$ be an F-stable maximal torus with Weyl group W. The $\tilde{G}$-conjugacy classes of F-stable maximal tori in $\tilde{G}$ are in one-to-one correspondence with the F-conjugacy classes of W. For an F-conjugacy class C in W, denote by $T_C \leq \hat{G}$ a representative of the corresponding $\tilde{G}$-conjugacy class of maximal tori, and let $T_C^F = T_C \cap \tilde{G}$. It is proved in [36, Theorem 1.3] that

$$\frac{|A(\tilde{G})|}{|\tilde{G}|} = \sum\nolimits_C \frac{|C|}{|W|} \cdot \frac{|T_C^F \cap A(\tilde{G})|}{|T_C^F|}$$

where C runs over all F-conjugacy classes of W. Our strategy is to restrict to special classes C which allow us to determine lower bounds for $|C|/|W|$ and $|T_C^F \cap A(\tilde{G})|/|T_C^F|$, thus providing a lower bound for $|A(\tilde{G})|/|\tilde{G}|$.

If the type is SL, then $W = S_d$, the symmetric group of degree d, and the F-classes of W are the conjugacy classes of W, so parametrised by partitions of d. Let C be a conjugacy class of S_d corresponding to a partition $(i, e, \ldots)$ where $i \in [d/6, d/3]$ and $e > d/3$ with $e + 2i > d$ and $i \nmid 2e$; call such a class *admissible*. Note that $T_C^F = (q^i - 1) \times (q^e - 1) \times T^*$. Let p and r be (q, i)-Zsigmondy and (q, e)-Zsigmondy primes, respectively. The proportion of elements in T_C^F with order divisible by pr is at least $1/4$, thus, $|T_C^F \cap A(\tilde{G})|/|T_C^F| \geq 1/4$ for each such class C. In conclusion,

$$\frac{|A(\tilde{G})|}{|\tilde{G}|} \geq \frac{1}{4d!} \sum\nolimits_C |C|,$$

where C runs over all admissible classes.

Let $l = \lceil d/6 \rceil$ and $u = \lfloor d/3 \rfloor$. It remains to estimate the number $N(d)$ of elements of S_d in admissible classes, that is, elements of cycle type $(i, e, \ldots)$ with $i \in [l, u]$, $e + 2i > d$, and $i \nmid 2e$. For this, we run over $i \in [l, u]$ and $e \in [d - 2i + 1, d - i]$, and count how many elements of cycle type $(i, e, \ldots)$ exist. First, we show that there is no over-counting. Since $e > d/3 \geq i$ and $e + 2i > d$,

$$(\#) \quad d - e - i \in [0, \ldots, i - 1],$$

and e is the unique largest entry in the cycle decomposition. Now suppose we encounter cycle types $(i, e, j, \ldots)$ and $(j, e, i, \ldots)$ with $i, j \in [l, u]$, $e \in [d-2i+1, d-i] \cap [d-2j+1, d-j]$, and $i+j+e \leq d$. The latter, together with $(\#)$, implies $j \leq d - i - e \leq i - 1$, hence $j < i$. By symmetry, we get $i < j$, a contradiction. Thus, there is no over-counting.

We next determine the number $\tilde{N}(d)$ of elements of cycle type $(i, e, \ldots)$ with $i \in [l, u]$ and $e \in [d - 2i + 1, d - i]$ as

$$\begin{aligned} \tilde{N}(d) &= \sum\nolimits_{i=l}^{u} \sum\nolimits_{e=d-2i+1}^{d-i} \binom{d}{i} \binom{d-i}{e} (i-1)!(e-1)!(d-e-i)! \\ &= d! \sum\nolimits_{i=l}^{u} \sum\nolimits_{e=d-2i+1}^{d-i} \frac{1}{ie}. \end{aligned}$$

For each i, there are at most three $2e \in [2d - 4i + 2, \ldots, 2d - 2i]$ with $i \mid 2e$, hence ignoring the three largest summands in $\sum_{e=d-2i+1}^{d-i} \frac{1}{ie}$ yields a lower bound for $N(d)$; in summary,

$$N(d) \geq d! \sum\nolimits_{i=l}^{u} \frac{1}{i} \sum\nolimits_{e=d-2i+4}^{d-i} \frac{1}{e}.$$

There is an absolute constant $z_1 > 0$ such that for large enough d and all $i \in [d/6, d/3]$,

$$\sum\nolimits_{e=d-2i+4}^{d-i} \frac{1}{e} \;\;\geq\;\; \int_{d-2i+4}^{d-i} \frac{1}{x}\mathrm{d}x = \log\left(\frac{d-i}{d-2i+4}\right) > z_1.$$

Thus, there is an absolute constant $z_2 > 0$ with

$$N(d) \geq d!\, z_1 \sum\nolimits_{i=l}^{u} \frac{1}{i} \geq d!\, z_1 \log(u/l) \geq d!\, z_2.$$

Since $N(d) = \sum_C |C|$, where C runs over all admissible classes, there is an absolute constant $z_3 > 0$ with $|A(\tilde{G})|/|\tilde{G}| \geq \frac{1}{4d!}N(d) \geq z_3$, which proves the claim for SL. The types Sp, SU, and Ω^+ are dealt with analogously.

Now consider $G \cong \Omega_d^-(q) = \tilde{G}$ and write $\tilde{d} = d/2$. Suppose $\tilde{d}$ is odd, define $c = (\tilde{d}-3)/2$, and suppose d is large enough such that $c/2 > 6$. We want to find $g \in G$ which, in the natural representation, acts irreducibly on a space of dimension $\tilde{d}-3$ and as the identity on a complement. For this, we seek $g \in G$ with order divisible by a $(q, 2c)$-Zsigmondy prime p and by a (q, e)-Zsigmondy prime r with $e \in [c/2+4, c+3] \setminus \{c, 2c/3\}$. Note that $e \nmid 2c$; thus, if g is such an element, then it must lie in a maximal torus of G isomorphic to $T = (q^c+1) \times (q^f \pm 1) \times T^*$ such that $r \mid q^f \pm 1$ and $e \mid 2f$. Since $\tilde{d} - c - f < c$, the power $h = g^{|g|/p}$ must lie in the direct factor q^c+1; hence, in the natural representation, h acts irreducibly on a space of dimension $2c = \tilde{d}-3$ and as the identity on a complement.

The Weyl group W of $\Omega_d^-(q)$ has order $2^{\tilde{d}-1}\tilde{d}!$ and the F-conjugacy classes of W correspond to signed partitions $(b_1^-, \ldots, b_i^-, c_1^+, \ldots, c_j^+)$ of $\tilde{d}$ where i is odd. The associated maximal torus is

$$T_C^F \cong (q^{b_1^-}+1) \times \ldots \times (q^{b_i^-}+1) \times (q^{c_1^+}-1) \times \ldots \times (q^{c_j^+}-1).$$

If there exist z elements in $S_{\tilde{d}}$ of cycle type $(a_1, \ldots, a_k)$, then $2^{\tilde{d}-k}z$ elements of W correspond to each $(a_1^-, \ldots, a_i^-, a_{i+1}^+, \ldots, a_k^+)$ with i odd.

As before, let $c = (\tilde{d}-3)/2$ and $e \in [c/2+4, c+3] \setminus \{c, 2c/3\}$. Let $\mathcal{C}_e$ be the union of all F-classes corresponding to signed partitions $(c^+, e^-, \ldots)$. Note that $\tilde{d}-c-e < c/2$ and there are $\tilde{d}!/ce$ elements in $S_{\tilde{d}}$ of cycle type $(c, e, \ldots)$; we claim that there are $2^{\tilde{d}-3}\tilde{d}!/ce$ elements in W corresponding to signed partitions $(c^+, e^-, \ldots)$, that is, $|\mathcal{C}_e| = 2^{\tilde{d}-3}d!/ce$.

To prove the claim, let $\lambda = (u_1, \ldots, u_t)$ be a partition of $\tilde{d} - c - e$. Define $\pi(\lambda) = u_1 \ldots u_t$ and $\zeta(\lambda) = n!$, where n is the number of $u_i = 1$. Using this notation, $S_{\tilde{d}}$ contains $\tilde{d}!/ce\zeta(\lambda)\pi(\lambda)$ elements corresponding to the partition $(c, e, u_1, \ldots, u_t)$. Each such partition gives rise to 2^{t-1} signed partitions $(c^+, e^-, u_1^{\epsilon_1}, \ldots, u_t^{\epsilon_t})$ with an even number of $\epsilon_i = +$, and, as shown above, for each such signed partition there exist $2^{\tilde{d}-2-t}\tilde{d}!/ce\zeta(\lambda)\pi(\lambda)$ elements in W. Thus, each partition $(c, e, u_1, \ldots, u_t) \vdash \tilde{d}$ yields $2^{\tilde{d}-3}\tilde{d}!/ce\zeta(\lambda)\pi(\lambda)$ elements in W corresponding to signed partitions $(c^+, e^-, u_1^{\epsilon_1}, \ldots, u_t^{\epsilon_t})$, where $\lambda = (u_1, \ldots, u_t)$. Clearly, $|\mathcal{C}_e|$ is the sum of these numbers, running over all partitions $\lambda \vdash \tilde{d} - c - e$, thus

$$|\mathcal{C}_e| = \frac{2^{\tilde{d}-3}\tilde{d}!}{ce} \sum\nolimits_{\lambda \vdash \tilde{d}-e-c} \frac{1}{\zeta(\lambda)\pi(\lambda)} = \frac{2^{\tilde{d}-3}\tilde{d}!}{ce};$$

the last equation follows since $m! = |S_m| = \sum_{\lambda \vdash m} m!/\zeta(\lambda)\pi(\lambda)$ for every integer $m \geq 1$.

Recall that $|T_C^F \cap A(\tilde{G})|/|T_C^F| \geq 1/4$ for each F-class $C \in \mathcal{C}_e$. In conclusion, there is an absolute constant $z > 0$ such that

$$\frac{|A(\tilde{G})|}{|\tilde{G}|} \;\;\geq\;\; \frac{1}{2^{\tilde{d}+1}\tilde{d}!} \sum\nolimits_{e=\lceil c/2 \rceil+3}^{c+1} |\mathcal{C}_e| \geq \frac{1}{16} \sum\nolimits_{e=\lceil c/2 \rceil+3}^{c+1} \frac{1}{ce} > z/d;$$

recall that $e \in [c/2+2, c+3] \setminus \{c, 2c/3\}$, so we estimate the sum over all such e by running with e from $\lceil c/2 \rceil + 2$ to $c+1$. The case of even $\tilde{d}$ is dealt with analogously. □

5.2. Constructing the second subgroup. Let $G = \langle X \rangle$ be isomorphic to $\mathrm{SX}_d(q)$ and let $H \leq G$ be constructed as in Lemma 5.3. We now describe the construction of $K \leq G$ such that $K \cong \mathrm{SX}_{d-m}(q)$ and H commutes with K. The approach is to constructively recognise H, explicitly write down a suitable involution $i \in H$, and then to find K in $C_G(i)$. As a first step, we comment on the structure of $C_G(i)$.

5.2.1. *Involution centralisers.* The *corank* of an involution $i \in \tilde{G}$ is the rank of the matrix $i - 1_d$. The next theorem describes the structure of involution centralisers in $\tilde{G}$; it is a modification of [19, Theorem 6.1], and was proved by Aschbacher & Seitz [2].

Theorem 5.4. *Let $i \in \tilde{G}$ be an involution of corank $r \leq d/2$. There exists $c \in \mathrm{GL}_d(\mathbb{F})$ such that*

$$i^c = \begin{pmatrix} 1_r & 0 & 1_r \\ 0 & 1_{d-2r} & 0 \\ 0 & 0 & 1_r \end{pmatrix}$$

and the elements of $C_{\tilde{G}^c}(i^c)$ have upper block triangular form with diagonal blocks a, b, a, of degrees r, $d-2r$, and r, respectively. Consider the homomorphism

$$\psi\colon C_{\tilde{G}^c}(i^c) \to \mathrm{GL}_r(\mathbb{F}) \times \mathrm{GL}_{d-2r}(\mathbb{F}), \quad \begin{pmatrix} a & \star & \star \\ 0 & b & \star \\ 0 & 0 & a \end{pmatrix} \mapsto (a, b).$$

(i) *If $\tilde{G}$ is linear or unitary, then the image of ψ contains $A \times B$ with $A = \mathrm{SX}_r(q)$ and $B = \mathrm{SX}_{d-2r}(q)$, both of the same type as $\tilde{G}$.*
(ii) *If $\tilde{G}$ is symplectic, then the image of ψ is $A \times B$ with $B = \mathrm{Sp}_{d-2r}(q)$ and*

$$A = \mathrm{Sp}_r(q) \quad or \quad A = \begin{pmatrix} 1 & \star \\ 0 & \mathrm{Sp}_{r-1}(q) \end{pmatrix} \quad or \quad A = \begin{pmatrix} 1 & \star & \star \\ 0 & \mathrm{Sp}_{r-2}(q) & \star \\ 0 & 0 & 1 \end{pmatrix}.$$

(iii) *If $\tilde{G}$ is orthogonal, then the image of ψ is $A \times B$ with A as in (ii). If $A = \mathrm{Sp}_r(q)$, then $B' = \mathrm{SX}_{d-2r}(q)$ has the same type as $\tilde{G}$; if $A \neq \mathrm{Sp}_r(q)$, then $B = \mathrm{Sp}_{d-2r}(q)$.*

The *standard form* of i is i^c; note that, in general, i^c is not in the standard copy $\tilde{G}$. If $i \in \tilde{G}$ is an involution of specific corank $r \in \{2, \ldots, d/2\}$, then its image under every automorphism of $\tilde{G}$ has the same corank. (We remark that this actually holds for all $\mathrm{SX}_d(q)$, including base cases, with the exceptions of $\mathrm{Sp}_4(q)$ and $\Omega_8^+(q)$: these have graph automorphisms which change the corank of involutions.) This allows us to define the corank of an involution $i \in G$ via $\varphi\colon G \to \tilde{G}$.

Let $i \in \tilde{G}$ be an involution of corank $r < d/2$ in standard form; let ψ and A be as in Theorem 5.4. If $\tilde{G}$ is linear or unitary, then $C \leq C_{\tilde{G}}(i)$ is *sufficient* if $\psi(C) \geq \mathrm{SX}_r(q) \times \mathrm{SX}_{d-2r}(q)$. If $\tilde{G}$ is symplectic or orthogonal, then $C \leq C_{\tilde{G}}(i)$ is *sufficient* if $\psi(C)$ contains $\mathrm{SX}_{d-2r}(q)$ and the projection to the irreducible diagonal block of A contains $\mathrm{Sp}_r(q)$, $\mathrm{Sp}_{r-1}(q)$, and $\mathrm{Sp}_{r-2}(q)$, respectively. If $i \in G$ is an involution, then a sufficient subgroup of $C_G(i)$ is defined via the isomorphism $\varphi\colon G \to \tilde{G}$ followed by a conjugation.

Theorem 5.5. *Let i be an involution in $G = \langle X \rangle \cong \mathrm{SX}_d(q)$. There is a black-box Monte Carlo algorithm which constructs, as* SLPs *in X, a generating set for a sufficient subgroup of $C_G(i)$; it runs in time $O(d(\mu + \xi + \mathcal{O} + \Pi))$.*

PROOF. It suffices to consider $O(d)$ random elements to construct a Bray generator of $C_G(i)$. As explained in the proof of [19, Theorem 6.4], a constant number of Bray generators suffices to generate a sufficient subgroup. □

5.2.2. *The second subgroup.* Let $G = \langle X \rangle$ be isomorphic to $\mathrm{SX}_d(q)$ and let $H \leq G$ be constructed as in Lemma 5.3, hence H is isomorphic to $\mathrm{SL}_m(q)$, $\mathrm{SU}_m(q)$, or $\Omega_m^+(q)$. By construction, there exists an isomorphism $\varphi\colon G \to \tilde{G}$ to the standard copy of $\mathrm{SX}_d(q)$ such that

$$\tilde{H} = \varphi(H) = \begin{pmatrix} \mathrm{SX}_m(q) & 0 \\ 0 & 1_{d-m} \end{pmatrix}.$$

We now describe the construction of $K \leq G$ with $K \cong \mathrm{SX}_{d-m}(q)$ of the same type as G such that

$$\tilde{K} = \varphi(K) = \begin{pmatrix} 1_m & 0 \\ 0 & \mathrm{SX}_{d-m}(q) \end{pmatrix}.$$

By recursion, we construct standard generators $\mathcal{S}_H$ of H. Note that $\tilde{\mathcal{S}}_H = \varphi(\mathcal{S}_H)$ is an automorphic image of the standard generators $\mathcal{S}(m, q, \mathrm{SX})$ embedded in $\tilde{H}$, say $\alpha(\tilde{\mathcal{S}}_H) = \mathcal{S}(m, q, \mathrm{SX})$ with $\alpha \in \mathrm{Aut}(\tilde{H})$. If $\tilde{H}$ is linear or unitary, then so is $\tilde{G}$, hence α lifts to an automorphism of $\tilde{G}$, see Remark 3.3. If $\tilde{H}$ is orthogonal, then α lifts to an automorphism of $\tilde{G} \in \{\mathrm{Sp}_d(q), \Omega_d^{\pm}(q)\}$, with possible exceptions for $\tilde{H} \cong \Omega_m^+(q)$ with $m \in \{4, 8\}$, see Lemma 3.4. We comment on this case in Remark 5.8; for now, suppose that $H \not\cong \Omega_m^+(q)$ with $m \in \{4, 8\}$. Under these assumptions, φ can be modified by an automorphism of $\mathrm{SX}_d(q)$ so that we can assume that $\tilde{\mathcal{S}}_H = \mathcal{S}(m, q, \mathrm{SX})$.

We use $\mathcal{S}_H$ to construct $i, f \in H$ such that there exists a base change matrix $c \in \mathrm{GL}_d(\mathbb{F})$ with

$$\varphi(i)^c = \begin{pmatrix} 1_r & 1_r & 0 \\ 0 & 1_r & 0 \\ 0 & 0 & 1_{d-m} \end{pmatrix}, \quad \varphi(f)^c = \begin{pmatrix} u & 0 & 0 \\ 0 & u & 0 \\ 0 & 0 & 1_{d-m} \end{pmatrix}, \quad \text{and} \quad \varphi(C_H(i))^c = \begin{pmatrix} A & \star & 0 \\ 0 & A & 0 \\ 0 & 0 & 1_{d-m} \end{pmatrix}$$

where $r = m/2$, $A \cong \mathrm{SX}_r(q)$ acts irreducibly, and $u \in \mathrm{SX}_r(q)$ is fixed-point free of odd order. We then apply the next proposition to construct the required subgroup $K \leq C_G(i)$. To visualise the situation, we now assume that φ is chosen such that $\varphi(i)$ has standard form, and

$$\varphi(f) = \begin{pmatrix} u & 0 & \star \\ 0 & 1_{d-m} & 0 \\ 0 & 0 & u \end{pmatrix},$$

and

$$\varphi(C_G(i))' = C_{\tilde{G}}(\varphi(i))' = \begin{pmatrix} A & \star & \star \\ 0 & \mathrm{SX}_{d-m}(q) & \star \\ 0 & 0 & A \end{pmatrix};$$

as in Theorem 5.4, denote by ψ the projection $C_{\tilde{G}}(\varphi(i))' \to A \times \mathrm{SX}_{d-m}(q)$.

Proposition 5.6. *There is a black-box Las Vegas algorithm which, using the above notation, constructs from i and f a subgroup $K \leq C_G(i)$ with $K \cong \mathrm{SX}_{d-m}(q)$ and*

$$\tilde{K} = \varphi(K) = \begin{pmatrix} 1_r & 0 & 0 \\ 0 & \mathrm{SX}_{d-m}(q) & 0 \\ 0 & 0 & 1_r \end{pmatrix}.$$

The algorithm runs in time $O(d(\mu + \xi + \mathcal{O} + \Pi))$.

PROOF. Suppose first that G is not of type Ω^-; hence $m \in [d/3, 2d/3]$, and therefore m and $d - m$ are approximately equal. Using Theorem 5.5, we find a sufficient subgroup $C \leq C_G(i)$, and then construct K as a subgroup of C. Applying a simple modification of Theorem 4.1, we obtain $K_1 \leq C$ with

$$\varphi(K_1) = \begin{pmatrix} 1_r & \star & \star \\ 0 & \mathrm{SX}_{d-m}(q) & \star \\ 0 & 0 & 1_r \end{pmatrix}.$$

We stress that, by construction of i, the types and degrees of $\mathrm{SX}_{d-m}(q)$ and $A \cong \mathrm{SX}_r(q)$ are known; thus, modulo a normal 2-subgroup of $C_G(i)$, Theorem 4.1 is essentially applied to $\mathrm{SX}_r(q) \times \mathrm{SX}_{d-m}(q)$.

If $h \in K_1$ is random and $\varphi(h)$ has diagonal blocks $1_r, b, 1_r$, then, as seen in the proof of [19, Lem. 7.1], the element $k = (fh(ff^h)^{(|f|-1)/2})^2$ lies in K and $\varphi(k)$ has diagonal blocks $1_r, b^2, 1_r$. It is proved in [23] that an $O(1)$ random search in a perfect classical group suffices to find a generating set M such that $\{x^2 \mid x \in M\}$ generates the group; thus, collecting $O(1)$ elements k of this type suffices to generate $K \le C$ with $\varphi(K) = \tilde{K}$. We use Remark 3.6 and [35] to verify that $K \cong \mathrm{SX}_{d-m}(q)$; the rank $d-m$ is known by construction, and K has the same type as G.

We proceed analogously if G has type Ω^-; see Remark 5.7 for more details on the construction of a subgroup of *small* degree. $\square$

Remark 5.7. In the gluing process, we also use the following modification of Proposition 5.6. Let $i, f \in G$ be as above, so

$$\varphi(i) = \begin{pmatrix} 1_r & 0 & 0 \\ 0 & 1_{d-2r} & 0 \\ 0 & 0 & 1_r \end{pmatrix}, \quad \varphi(f) = \begin{pmatrix} u & 0 & * \\ 0 & 1_{d-2r} & 0 \\ 0 & 0 & u \end{pmatrix}, \quad \text{and} \quad \varphi(C_G(i))' = \begin{pmatrix} A & \star & \star \\ 0 & \mathrm{SX}_{d-2r}(q) & \star \\ 0 & 0 & A \end{pmatrix},$$

but, this time, $d - 2r \le 10$, and $A \le \mathrm{SX}_r(q)$ is as in Theorem 5.4(iii), containing a subgroup $\mathrm{Sp}_{r'}(q)$ with $r' \in \{r, r-1, r-2\}$. The degree r is known, and we want to construct $K \le G$ with $\varphi(K) = \mathrm{diag}(1_r, \mathrm{SX}_{d-2r}(q), 1_r)$; note that we also know r' by investigating Zsigmondy prime divisors of random elements in $C_G(i)$. We now proceed as in Remark 4.2. The first step is to use a modification of Theorem 4.1 to construct $B \le C_G(i)$ such that $\varphi(B)$ has diagonal blocks $\hat{B}, 1_{d-2r}, \hat{B}$ with $\hat{B} \le A$ and $C_A(\hat{B}) = 1$, for example, $\hat{B} = \mathrm{Sp}_{r'}(q)$ or $\Omega^{\pm}_{r'}(q)$. Since $d - 2r \le 10$ is small, elements in B can be readily constructed by taking random Bray generators of $C_G(i)$ to the power $\exp(\mathrm{SX}_{d-2r}(q))$. Let $K_1 \le C_G(i)$ be as in the proof of Proposition 5.6. Observe that $h \in C_G(i)$, of order not dividing $|Z(A)|$, lies in K_1 if and only if $[h, b]$ is a 2-element for every generator b of B. Having found such an h, we construct K_1 as the normal closure of h in $C_G(i)$. Finally, the subgroup $K \le K_1$ we seek is constructed as before.

Remark 5.8. We comment on the two exceptional cases $\Omega_4^+(q)$ and $\Omega_8^+(q)$.

a) Let $H \cong \Omega_4^+(q) = \mathrm{SL}_2(q) \times \mathrm{SL}_2(q)$. Using the above notation, $\alpha(\tilde{\mathcal{S}}_H) = \mathcal{S}(4, q, \Omega^+)$ for some $\alpha \in \mathrm{Aut}(\tilde{H})$. If α does not lift to an automorphism of $\tilde{G} = \mathrm{SX}_d(q)$, then, modulo automorphisms of $\tilde{H}$ that lift to $\tilde{G}$, it must be a field automorphism of H, acting differently on the two direct factors $\mathrm{SL}_2(q)$ of H. In other words, the semisimple elements $\delta, y \in \mathcal{S}(4, q, \Omega^+)$ are defined with respect to two different primitive elements of the underlying field $\mathrm{GF}(q)$ which are equal modulo applying a Frobenius automorphism. This has no impact on the above construction, and we obtain $K \cong \mathrm{SX}_{d-m}(q)$ as before. However, such $\delta, y \in \mathrm{SX}_d(q)$ cannot be used as the semisimple standard generators of $\mathrm{SX}_d(q)$; instead we must replace y by a suitable power y^{p^j} for some $j \in \{0, \ldots, e-1\}$ where $q = p^e$. We correct this when gluing the standard generators. For simplicity, in the remainder of this paper, we suppose that if $H \cong \Omega_4^+(q)$, then $\tilde{\mathcal{S}}_H = \mathcal{S}(4, q, \Omega^+)$. This remark also holds for odd q and $H \cong \Omega_4^+(q) = \mathrm{SL}_2(q) \circ \mathrm{SL}_2(q)$.

b) Let $H \cong \Omega_8^+(q)$. Using the above notation, $\alpha(\tilde{\mathcal{S}}_H) = \mathcal{S}(8, q, \Omega^+)$ for some $\alpha \in \mathrm{Aut}(\tilde{H})$. If α does not lift to an automorphism of $\tilde{G} = \mathrm{SX}_d(q)$, then, modulo automorphisms of $\tilde{H}$ that lift to $\tilde{G}$, it must be a graph automorphism γ of H having order 3. Such a graph automorphism may change the corank of an involution, and the above construction to obtain $K \cong \mathrm{SX}_{d-8}(q)$ will fail. We remedy the situation as follows. Having constructively recognised H, we can compute images under γ. Let $i_0 = i \in H$ be the involution we have constructed for Proposition 5.6, and define $i_j = \gamma^j(i)$ for $j = 1, 2$. In the natural representation, exactly one of i_0, i_1, i_2 has corank 4 and the other two have corank 2. The centraliser of each involution of corank 2 contains a subgroup $\Omega_{d-4}^+(q)$, which is not the case for a centraliser of an involution of corank 4. Thus, in the centralisers of these involutions, we look for an element that witnesses a subgroup $\Omega_{d-4}^+(q)$, for example, a $\mathrm{ppd}(q, d-6)$ element. We will find such witnesses with $O(d)$ trials for the two involutions of corank 2. If the remaining involution is i_1 or i_2, then we apply γ or γ^2 to the standard generators that we found for H. Using this strategy, we can assume that $\tilde{\mathcal{S}}_H = \mathcal{S}(8, q, \Omega^+)$. For simplicity, in the remainder of this paper, we suppose that if $H \cong \Omega_8^+(q)$, then $\tilde{\mathcal{S}}_H = \mathcal{S}(8, q, \Omega^+)$.

6. Gluing the cycles

Let $G = \langle X \rangle$ be isomorphic to $\mathrm{SX}_d(q)$ with q even or odd. Using the algorithms of the previous sections, we have constructed commuting subgroups $H \cong \mathrm{SX}_m(q)$ and $K \cong \mathrm{SX}_{d-m}(q)$ of G, such that there exists an isomorphism $\varphi\colon G \to \tilde{G}$ to the standard copy of $\mathrm{SX}_d(q)$ with

$$\tilde{H} = \varphi(H) = \left(\begin{smallmatrix} \mathrm{SX}_m(q) & 0 \\ 0 & 1_{d-m} \end{smallmatrix}\right) \quad \text{and} \quad \tilde{K} = \varphi(K) = \left(\begin{smallmatrix} 1_m & 0 \\ 0 & \mathrm{SX}_{d-m} \end{smallmatrix}\right).$$

The isomorphism φ is unknown, but we use it to visualise the situation. By recursion, we have constructed standard generators $\mathcal{S}_H$ and $\mathcal{S}_K$ for H and K, respectively. Recall that $\mathcal{S}_H \cup \mathcal{S}_K$ contains standard generators $\mathcal{S}_G$ of G, with the exception of the cycle v_G. In this section, we describe how to *glue* the cycles v_H and v_K of H and K, respectively, to obtain a suitable cycle v_G; this will complete the construction of the standard generators of G. For odd q, our approach follows that of [26]; for even q, we use a strategy different to that of [19], see Remark 6.1.

As outlined in Section 5.2.2, we can suppose that φ maps $\mathcal{S}_H$ onto the standard generators $\mathcal{S}(m, q, \mathrm{SX})$ of $\mathrm{SX}_m(q)$ embedded in $\tilde{H}$, that is, $\tilde{\mathcal{S}}_H = \varphi(\mathcal{S}_H) = \mathcal{S}(m, q, \mathrm{SX})$; note that Remark 5.8a) also applies to odd q. Now consider $\tilde{\mathcal{S}}_K = \varphi(\mathcal{S}_K) \subseteq \tilde{K}$, which is an automorphic image of the standard generators $\mathcal{S} = \mathcal{S}(d-m, q, \mathrm{SX})$ of $\mathrm{SX}_{d-m}(q)$, say $\tilde{\mathcal{S}}_K = \beta(\mathcal{S})$ with $\beta \in \mathrm{Aut}(\tilde{K})$. By Remark 3.3, we decompose $\beta = \beta_{\mathrm{g}} \circ \beta_{\mathrm{f}} \circ \beta_{\mathrm{d}} \circ \beta_{\mathrm{i}}$ into a graph, field, diagonal, and inner automorphism, respectively. In all cases, β_{i} and β_{d} lift to automorphisms of $\tilde{G}$ which fix $\tilde{\mathcal{S}}_H$ element-wise; therefore, we can suppose the following:

(i) $\varphi(\mathcal{S}_H) \subseteq \tilde{H}$ are the standard generators $\mathcal{S}(m, q, \mathrm{SX})$ of $\mathrm{SX}_m(q)$ embedded in $\tilde{H}$,
(ii) $\varphi(\mathcal{S}_K) \subseteq \tilde{K}$ are the standard generators $\mathcal{S}(d-m, q, \mathrm{SX})$ of $\mathrm{SX}_{d-m}(q)$ embedded in $\tilde{K}$, or the image of these under field and graph automorphisms of $\tilde{K}$.

If K is not isomorphic to $\mathrm{Sp}_4(q)$ or $\Omega_8^+(q)$ with q even, then we can also assume that (i) and (ii) hold with the roles of H and K interchanged. Note that if q is even, $d - m = 4$, and $K \cong \mathrm{Sp}_4(q)$ (which arises in our algorithms only for $d \in \{8, 12, 16\}$), then the graph automorphism of K does not lift to an automorphism of $\tilde{G} = \mathrm{Sp}_d(q)$; similarly for $\Omega_8^+(q)$. We comment on this in Section 6.3.

In Section 6.1, we describe the general strategy for gluing in $\mathrm{SL}_d(q)$ in both even and odd characteristic. In Section 6.2, we describe gluing in $G \cong \mathrm{SU}_d(q)$ with d odd and q even; this exemplifies the algorithm used for other types.

6.1. General strategy. Let $\tilde{G} = \mathrm{SL}_d(q)$, $\tilde{H} = \mathrm{SL}_m(q)$, and $\tilde{K} = \mathrm{SL}_{d-m}(q)$. Suppose $d = 2n$ is even and choose a (necessarily hyperbolic) basis $\{e_1, f_1, \ldots, e_n, f_n\}$ of the natural $\tilde{G}$-module such that (i) and (ii) hold. Write $m = 2z$. Recall that ω is a fixed primitive element of $\mathrm{GF}(q)$. By assumption, the cycles $v_H \in \mathcal{S}_H$ and $v_K \in \mathcal{S}_K$ satisfy

$$\varphi(v_H) = (e_1, e_2, \ldots, e_z)(f_1, f_2, \ldots, f_z) \quad \text{and} \quad \varphi(v_K) = (e_{z+1}, e_{z+2}, \ldots, e_n)(f_{z+1}, f_{z+2}, \ldots, f_n);$$

here $(e_1, e_2, \ldots, e_z)$ is the permutation mapping $e_1 \to e_2 \to \ldots \to e_z \to e_1$, and similarly for the other cycles. By (ii), this also holds if $\varphi(\mathcal{S}_K)$ is an image of $\mathcal{S}(d-m, q, \mathrm{SL})$ under field or inverse-transpose automorphisms.

Observe $v_G = v_K g v_H$ is a cycle in G where $g \in G$ is mapped to $\varphi(g) = (e_z, e_{z+1})(f_z, f_{z+1}) \in \tilde{G}$; indeed

$$\varphi(v_K)\varphi(g)\varphi(v_H) = (e_1, e_2, \ldots, e_n)(f_1, f_2, \ldots, f_n).$$

It remains to construct $g \in G$; in fact, we are only able to construct $g \in G$ such that $\varphi(g)$ maps e_z and f_z to ce_{z+1} and cf_{z+1}, respectively, for some unknown non-zero scalar c; such a *glue element* suffices. We find a suitable glue element in the centraliser $C_G(i)$ of a certain involution. This requires a case distinction.

For odd q, we use $\mathcal{S}_H$ and $\mathcal{S}_K$ to construct $A, B \le G$ with $\varphi(A) = \mathrm{diag}(1_{m-2}, \mathrm{SL}_2(q), 1_{d-m})$ and $\varphi(B) = \mathrm{diag}(1_m, \mathrm{SL}_2(q), 1_{d-m-2})$; let $\mathcal{S}_A$ and $\mathcal{S}_B$ be standard generators of A and B, respectively. The glue element g can now be found in $C_G(i)$ where $i \in A \times B$ is an involution with $\varphi(i) = \mathrm{diag}(1_{m-2}, -1_4, 1_{d-m-2})$. Using Theorem 4.1 and the algorithm of Remark 4.2, we extract from $C_G(i)$ the subgroup $K \le G$ with

$$\varphi(K) = \mathrm{diag}(1_{m-2}, \mathrm{SL}_4(q), 1_{d-m-2});$$

note that $g \in K$. We constructively recognise K and obtain an isomorphism $\psi\colon K \to \mathrm{SL}_4(q)$. Using $\psi(A), \psi(B) \le \mathrm{SL}_4(q)$, we can find a base change matrix $w \in \mathrm{SL}_4(q)$ such that $\psi(\mathcal{S}_A)^w$ and $\psi(\mathcal{S}_B)^w$ are the standard generators of $\mathrm{SL}_2(q)$, but $\psi(\mathcal{S}_B)^w$ may be twisted by field or graph automorphisms. Having constructively recognised K, we can find $g \in K$ such that $\psi(g)^w$ is the permutation matrix defined by $(1,3)(2,4)$. Thus, by construction, there is a scalar c such that $\varphi(g)$ maps e_z and f_z either to ce_{z+1} and cf_{z+1}, or to cf_{z+1} and ce_{z+1}, depending on whether inverse-transpose is involved in (ii) or not. In both cases, $v_G = v_K g v_H$ is a cycle of G: in the first case, choose

$$\{e_1, f_1, \dots, e_z, f_z, ce_{z+1}, cf_{z+1}, \dots, ce_n, cf_n\}$$

as a hyperbolic basis of $\tilde{G}$, so

$$\varphi(v_K)\varphi(g)\varphi(v_H) = (e_1, e_2, \dots, e_z, ce_{z+1}, \dots, ce_n)(f_1, f_2, \dots, e_z, cf_{z+1}, \dots, cf_n)$$

is a cycle of $\tilde{G}$. In the second case, choose

$$\{e_1, f_1, \dots, e_z, f_z, cf_{z+1}, ce_{z+1}, \dots, cf_n, ce_n\}$$

as a hyperbolic basis of $\tilde{G}$, so

$$\varphi(v_K)\varphi(g)\varphi(v_H) = (e_1, e_2, \dots, e_z, cf_{z+1}, \dots, cf_n)(f_1, f_2, \dots, e_z, ce_{z+1}, \dots, ce_n)$$

is a cycle of $\tilde{G}$.

For even q, the situation is more complicated. We use $\mathcal{S}_H$ and $\mathcal{S}_K$ to construct $A, B \le G$ as for odd q. We also construct $i_H \in H$ and $i_K \in K$ with $\varphi(i_H) = \mathrm{diag}(s, \dots, s, 1_{d-m+2})$ and $\varphi(i_K) = \mathrm{diag}(1_{m+2}, s, \dots, s)$ where $s = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$; it is possible that $\varphi(i_K)$ is the inverse-transpose of this element, but, as for odd q, this has no impact. Now $i = i_H i_K$ is an involution of corank $r = d/2 - 2$ and, modulo a base change,

$$\varphi(C_G(i)') = \begin{pmatrix} \mathrm{SL}_r(q) & \star & \star \\ 0 & \mathrm{SL}_4(q) & \star \\ 0 & 0 & \mathrm{SL}_r(q) \end{pmatrix}.$$

Using Theorem 4.1 and Remark 5.7, we construct $N \le C_G(i)$ corresponding to the middle block $\mathrm{SL}_4(q)$ of $\varphi(C_G(i))$; thus $N \cong \mathrm{SL}_4(q)$, and N contains A, B, and the glue element g. We now proceed as for odd q, construct the glue $g \in G$, and $v_G = v_K g v_H$; this completes the construction of $\mathcal{S}_G$. Note that our construction of i requires that m is even; similarly, for symplectic and orthogonal groups, we require that m is divisible by 4.

Remark 6.1. For even q and in the natural representation, we find the glue g in the centraliser of an involution i of corank 2, see [19]. More precisely, we construct g in

$$\begin{pmatrix} \mathrm{SL}_2(q) & 0 & \star \\ 0 & 1_{d-4} & \\ 0 & 0 & \mathrm{SL}_2(q) \end{pmatrix} \le \begin{pmatrix} \mathrm{SL}_2(q) & \star & \star \\ 0 & \mathrm{SL}_{d-4}(q) & \star \\ 0 & 0 & \mathrm{SL}_2(q) \end{pmatrix} = C_G(i)'.$$

We stress that $g \notin \mathrm{diag}(\mathrm{SL}_2(q), 1_{d-4}, \mathrm{SL}_2(q))$: to find g, we must inspect the top right block of matrices in $C_G(i)$. This approach does not work in the black-box situation: we cannot *see* this top right block, and we cannot align bases and write down the required glue element. As a consequence, as outlined above, for a black-box group we must find the glue element in a *clean middle block* of an involution centraliser, where we can align bases and write down the element we seek. Recall that the construction of the middle block of an involution centraliser $C_G(i)$ requires an element f *compatible* with i: namely, they interact as required by Proposition 5.6 (and Remark 5.7). For $G \cong \Omega_d^-(q)$, such compatible i and f do not exist for some values of d, so the glue element cannot be constructed in a middle block of an involution centraliser. To avoid this problem, we construct the first subgroup, $H \cong \Omega^+_m(q)$, with large rank $m \in \{d-4, d-6\}$ divisible by 4. The second subgroup $K \cong \Omega^-_{d-m}(q)$ now has small rank, and we find the glue in $\Omega^-_{d-m+4}(q)$ where $d-m+4 \leq 10$, see Section 6.3 for more details. This explains why, for Ω^-, we must construct a first subgroup H of large rank.

6.2. A detailed example: $\mathrm{SU}_d(q)$ **with** q **even.** We describe gluing in $G \cong \mathrm{SU}_d(q)$ with q even and $d = 2n+1$ odd. Let $\tilde{G} = \mathrm{SU}_d(q)$ have hyperbolic basis $\{e_1, f_1, \ldots, e_n, f_n, w\}$, and assume that (i) and (ii) hold. We now construct standard generators of G from $\mathcal{S}_H$ and $\mathcal{S}_K$ by gluing the cycles of $H \cong \mathrm{SU}_m(q)$ and $K \cong \mathrm{SU}_{d-m}(q)$. Denote by $\mathbb{F} = \mathrm{GF}(q^2)$ the underlying field of $\mathrm{SU}_d(q)$ with primitive element ω; let $\delta = \omega^{q+1}$ and write $m = 2z$. By assumption, the cycles $v_H \in \mathcal{S}_H$ and $v_K \in \mathcal{S}_K$ satisfy

$$\begin{aligned}\varphi(v_H) &= (e_1, e_2, \ldots, e_z)(f_1, f_2, \ldots, f_z) \text{ and}\\ \varphi(v_K) &= (e_{z+1}, e_{z+2}, \ldots, e_n)(f_{z+1}, f_{z+2}, \ldots, f_n, w).\end{aligned}$$

To construct the cycle for G, we compute $v_K g v_H$ where $g \in G$ with $\varphi(g) = (e_z, e_{z+1})(f_z, f_{z+1}) \in \tilde{G}$. We use $\mathcal{S}_H \cup \mathcal{S}_K$ to construct $i_H, f_H, s_H, t_H, \delta_H \in H$ with

$$\begin{aligned}\varphi(i_H) &= \mathrm{diag}(\left(\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right), \ldots, \left(\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right), 1_2, 1_{d-m}) \in \tilde{H},\\ \varphi(f_H) &= \mathrm{diag}(\delta, \delta^{-1}, \ldots, \delta, \delta^{-1}, 1_2, 1_{d-m}) \in \tilde{H},\\ \varphi(s_H) &= \mathrm{diag}(1_{m-2}, \left(\begin{smallmatrix}0&1\\1&0\end{smallmatrix}\right), 1_{d-m}) \in \tilde{H},\\ \varphi(t_H) &= \mathrm{diag}(1_{m-2}, \left(\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right), 1_{d-m}) \in \tilde{H},\\ \varphi(\delta_H) &= \mathrm{diag}(1_{m-2}, \delta, \delta^{-1}, 1_{d-m}) \in \tilde{H},\end{aligned}$$

and $i_K, f_K, s_K, t_K, \delta_K \in K$ such that

$$\begin{aligned}\varphi(i_K) &= \mathrm{diag}(1_m, 1_2, \left(\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right), \ldots, \left(\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right), 1) \in \tilde{K},\\ \varphi(f_K) &= \mathrm{diag}(1_m, 1_2, \delta, \delta^{-1}, \ldots, \delta, \delta^{-1}, 1) \in \tilde{K},\\ \varphi(s_K) &= \mathrm{diag}(1_m, \left(\begin{smallmatrix}0&1\\1&0\end{smallmatrix}\right), 1_{d-m-2}) \in \tilde{K},\\ \varphi(t_K) &= \mathrm{diag}(1_m, \left(\begin{smallmatrix}1&1\\0&1\end{smallmatrix}\right), 1_{d-m-2}) \in \tilde{K},\\ \varphi(\delta_K) &= \mathrm{diag}(1_m, \delta, \delta^{-1}, 1_{d-m-2}) \in \tilde{K},\end{aligned}$$

or they are images of these under field automorphisms, cf. (ii) above; note that $\mathrm{SU}_d(q)$ has no graph automorphism. Let

$$\iota = i_H i_K \quad \text{and} \quad f = f_H f_K,$$

so $\varphi(\iota)$ is an involution of corank $r = n-2$ in $\mathrm{GL}_d(\mathbb{F})$, and there exists a permutation matrix $b \in \mathrm{GL}_d(\mathbb{F})$ such that $\varphi(\iota)^b$ has standard form and $\varphi(f)^b$ remains a diagonal matrix; for simplicity, suppose that $b = 1$ in the following. Thus,

$$\tilde{C} = \begin{pmatrix} \mathrm{SU}_r(q) & * & * \\ 0 & \mathrm{SU}_5(q) & * \\ 0 & 0 & \mathrm{SU}_r(q) \end{pmatrix} \leq C_{\tilde{G}}(\varphi(\iota));$$

the underlying basis of $\tilde{C}$ is

$$\{e_1, e_2, \ldots, e_{z-1}, e_{z+2}, \ldots, e_n, e_z, f_z, e_{z+1}, f_{z+1}, w, f_1, f_2, \ldots, f_{z-1}, f_{z+2}, \ldots, f_n\}.$$

We use f and the algorithm of Remark 5.7 to construct $L \leq C_G(\iota)$ with

$$\tilde{L} = \varphi(L) = \begin{pmatrix} 1_r & 0 & 0 \\ 0 & \mathrm{SX}_5(q) & 0 \\ 0 & 0 & 1_r \end{pmatrix} \leq \tilde{C},$$

which contains the image $\varphi(g)$ of the glue $g \in G$ we seek.

Let N be the standard copy of $\mathrm{SU}_5(q)$. Using a base case algorithm, we construct an isomorphism

$$\psi\colon L \to N.$$

Let $s_1, t_1, \delta_1 \in N$ and $s_2, t_2, \delta_2 \in N$ be the images under ψ of $s_H, t_H, \delta_H \in H \cap L$ and $s_K, t_K, \delta_K \in K \cap L$, respectively. If $i \in \{1,2\}$, then $\{s_i, t_i, \delta_i\} \subseteq N$ generates a subgroup isomorphic to $\mathrm{SL}_2(q)$; we now define a basis $\{v_1, \dots, v_5\}$ by choosing and constructing $v_2 \in \mathrm{im}\,(t_1 - 1_4) \setminus \{0\}$, $v_1 = v_2 s_1$, $v_4 \in \mathrm{im}\,(t_2 - 1_4) \setminus \{0\}$, $v_3 = v_4 s_2$, and $v_5 \in (\mathrm{Eig}(s_1 t_1, 1) \cap \mathrm{Eig}(s_2 t_2, 1)) \setminus \{0\}$.

Lemma 6.2. *If b is the base change matrix to the basis $\{v_1, \dots, v_5\}$, then there exist $j_1, j_2 \in \mathbb{N}$ with*

$$s_1^b = \mathrm{diag}(\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right), 1_3), \quad t_1^b = \mathrm{diag}(\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right), 1_3), \quad \delta_1^b = \mathrm{diag}(\delta^{j_1}, \delta^{-j_1}, 1_3)$$

and

$$s_2^b = \mathrm{diag}(1_2, \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right), 1), \quad t_2^b = \mathrm{diag}(1_2, \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right), 1), \quad \delta_2^b = \mathrm{diag}(1_2, \delta^{j_2}, \delta^{-j_2}, 1).$$

There exist $c, c' \in \mathbb{F}$ such that $\{v_1, \dots, v_5\}$ corresponds to $\{e_z, f_z, ce_{z+1}, cf_{z+1}, c'w\}$.

PROOF. Denote by $\bar{s_i}, \bar{t_i}, \bar{\delta_i}$ the matrices displayed in the lemma. By definition, $\tilde{L}$ has block diagonal form $\mathrm{diag}(1_r, \mathrm{SU}_5(q), 1_r)$; let $\pi\colon \tilde{L} \to \mathrm{SU}_5(q)$ be the projection onto the middle block. By construction, $\alpha = \pi \circ \varphi \circ \psi^{-1}$ is an isomorphism $\mathrm{SU}_5(q) \to \mathrm{SU}_5(q)$ which maps s_i, t_i, and δ_i to $\bar{s_i}$, $\bar{t_i}$, and $(\bar{\delta_i})^{-j_i}$, respectively. Let κ be an inner automorphism, adjusting the hermitian form, such that $\alpha' = \alpha \circ \kappa^{-1}$ is an automorphism of $\mathrm{SU}_5(q)$. Then α' maps $\kappa(s_i)$, $\kappa(t_i)$, and $\kappa(\delta_i)$ to $\bar{s_i}$, $\bar{t_i}$, and $(\bar{\delta_i})^{-j_i}$, respectively.

By Remark 3.3, we can decompose $\alpha' = \alpha_{\mathrm{f}} \circ \alpha_{\mathrm{d}} \circ \alpha_{\mathrm{i}}$. Hence, $\beta = \alpha_{\mathrm{d}} \circ \alpha_{\mathrm{i}} \circ \kappa$ satisfies

$$\begin{aligned} \mathrm{diag}(\mathrm{SL}_2(q), 1_3) &= \langle \bar{s_1}, \bar{t_1}, \bar{\delta_1} \rangle = \alpha' \circ \kappa(\langle s_1, t_1, \delta_1 \rangle) = \beta(\langle s_1, t_1, \delta_1 \rangle) \quad \text{and} \\ \mathrm{diag}(1_2, \mathrm{SL}_2(q), 1_1) &= \langle \bar{s_2}, \bar{t_2}, \bar{\delta_2} \rangle = \alpha' \circ \kappa(\langle s_2, t_2, \delta_2 \rangle) = \beta(\langle s_2, t_2, \delta_2 \rangle). \end{aligned}$$

The outer automorphism group of $\mathrm{SL}_2(q) \cong \mathrm{SU}_2(q)$ is the group of field automorphisms; thus, modulo field automorphisms, each automorphic image of the standard generators of $\mathrm{SU}_2(q)$ is conjugate in $\mathrm{SU}_2(q)$ to $\mathcal{S}(2, q, \mathrm{SU})$. In summary, there exists an isomorphism γ, realised as conjugation by the base change matrix b defined in the lemma, such that $\gamma(s_i) = \bar{s_i}$, $\gamma(t_i) = \bar{t_i}$, and $\gamma(\delta_i) = \bar{\delta_i}$ for $i \in \{1,2\}$. $\square$

We now show how to construct the cycle of G, and thereby complete the construction of standard generators of G. Let $N = \mathrm{SU}_5(q)$ and let b and $s_i^b, t_i^b, \delta_i^b \in N^b$ be as in Lemma 6.2. The hermitian form preserved by N^b is

$$\mathrm{diag}(\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 & x \\ x & 0 \end{smallmatrix}\right), y)$$

for some $x, y \in \mathrm{GF}(q^2)$. It follows from [43, Theorem 7.1(iii)] that $x \in \mathrm{GF}(q)$, so the algorithm of [20] is used to find $s \in \mathrm{GF}(q^2)$ with $s^{q+1} = x^{-1}$. Write $t = s^{-1}$, so

$$h = \left(\begin{smallmatrix} 0 & 0 & s & 0 & 0 \\ 0 & 0 & 0 & s & 0 \\ t & 0 & 0 & 0 & 0 \\ 0 & t & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{smallmatrix}\right) \in N^b,$$

and we construct $g \in G$ with $\psi(g) = h^{b^{-1}} \in N$. Clearly, $\varphi(g)$ maps e_z and f_z to ce_{z+1} and cf_{z+1} for some $c \in \mathbb{F}$, and $v = v_K g v_H$ is a cycle for G.

Since d is odd, some subset $\mathcal{S}_1$ of the standard generators of G lies in H and some subset $\mathcal{S}_2$ of the standard generators lies in K. However, in $\tilde{G}$, our constructed cycle $\varphi(v)$ is not necessarily compatible with the underlying hyperbolic bases for $\varphi(\mathcal{S}_1)$ and $\varphi(\mathcal{S}_2)$. The solution is to redefine $\mathcal{S}_1$ as the image of a suitable subset of $\mathcal{S}_K$ under conjugation by $v^{-m/2}$. For this, we require that the conditions (i) and (ii) are formulated with H and K interchanged, so that $\varphi(\mathcal{S}_K) = \mathcal{S}(d-m, q, \mathrm{SU})$. This strategy requires $d - m \geq 5$; in particular, $d = 7$ must be treated separately.

6.3. Gluing in orthogonal and symplectic groups. If q is odd, then we glue as outlined in Section 6.1; we deal with forms as described in Section 6.2.

For even q, the situation is more complicated. If $G \cong \mathrm{SX}_d(q)$ is symplectic or orthogonal, then, by construction, $K \cong \mathrm{SX}_{d-m}(q)$ has the same type as G, and $H \cong \Omega_m^+(q)$ with m divisible by 4. We use the same approach to construct standard generators for G from those of H and K. If G has type different to Ω^+, then the non-cycle standard generators of G are those in K, and not in H. In this case, we require that (i) and (ii) hold with the roles of H and K interchanged. This poses some problems if K is isomorphic to $\mathrm{Sp}_4(q)$ or $\Omega_8^+(q)$; we comment on this below.

If G is symplectic, then we glue the cycles in $\mathrm{Sp}_6(q) \leq G$, extracted from an involution centraliser. In a variation of Lemma 6.2, we use $\Omega_4^+(q) \leq H \cap \mathrm{Sp}_6(q)$ and $\mathrm{SL}_2(q) \leq K \cap \mathrm{Sp}_6(q)$ to choose a basis which allows us to construct the glue element. If $H \cong \Omega_m^+(q)$ and $K \cong \mathrm{Sp}_4(q)$, which implies $d \in \{8, 12, 16\}$, then the standard generators $\mathcal{S}_K$ of K may correspond to the automorphic image of $\mathcal{S}(4, q, \mathrm{Sp})$ under a graph automorphism of K. In this case, we cannot assume that (i) and (ii) hold with the roles of H and K interchanged, and we cannot use $\mathcal{S}_K$ as a subset of the standard generators of G. We detect and correct this as in Remark 5.8b).

If G is orthogonal, then we glue the cycles in a subgroup $\Omega_k^{\pm}(q)$ of the same type as G, or in $\mathrm{Sp}_k(q)$, cf. Theorem 5.4(iii). Here we describe the case $\Omega_k^{\pm}(q)$; the other is similar (and can be avoided in practice). If $G \cong \Omega_d^+(q)$ with $d \equiv 2 \bmod 4$, then we glue the cycles in a subgroup $\Omega_6^+(q) \leq G$; we choose a suitable basis using $\Omega_4^+(q) \leq H$ and $D \leq K$ with $D \cong \langle \mathrm{diag}(\omega, \omega^{-1})\rangle$. If $d \equiv 0 \bmod 4$, then we glue in $\Omega_8^+(q) \leq G$, and we use $\Omega_4^+(q) \leq H$ and $\Omega_4^+(q) \leq K$ to align the basis. Working in this $\Omega_8^+(q)$, or if $K \cong \Omega_8^+(q)$, we face the same problem as in case Sp with $K \cong \mathrm{Sp}_4(q)$; we deal with it in the same manner. If $K \cong \Omega_4^+(q)$, then we may need to adjust the semisimple elements of $\mathcal{S}_K$, see Remark 5.8. If $G \cong \Omega_d^-(q)$, then $m \in \{d-4, d-6\}$ and m is divisible by 4. We glue in $\Omega_k^-(q) \leq G$ with $k \in \{8, 10\}$; we adjust the basis using K and $\Omega_4^+(q) \leq H$. Some standard generators of G lie in H, and some lie in K, thus we proceed as for SU in odd degree, see Section 6.2.

6.4. The cost. We now summarise the cost of the algorithm to glue the cycles. Observe that we must construct the centraliser of an involution; the algorithm also requires one base case call.

Lemma 6.3. *Given $H, K \leq G$, $\mathcal{S}_H \subseteq H$, and $\mathcal{S}_K \subseteq K$, the algorithm to construct standard generators $\mathcal{S}_G$ of G has complexity $O(d(\mu + \xi + \mathcal{O} + \Pi) + \mathcal{B})$ where $\mathcal{B} = (\chi + \mu)\log^2 q + \xi \log q \log\log q$ reflects the cost of recognising a single base case.*

7. Complexity of the algorithm

We now summarise the complexity of the algorithm. Suppose the input group G has parameters (d, q, SX). By recursion, we apply our main algorithm 2^i times to degree $d/2^i$ for $i \in \{0, 1, \ldots, j\}$ with $j \approx \log d$. In degree r, the time to construct two smaller subgroups is $O(r(\mu + \xi + \mathcal{O} + \Pi))$ and the time for gluing is $O(r(\mu + \xi + \mathcal{O} + \Pi) + \mathcal{B})$. Summing this up and adding the cost for the $O(d)$ base cases gives the complexity stated in Theorem 1.2.

While we present black-box algorithms, their primary application is to recognise absolutely irreducible classical groups in representations (other than the natural representation) in the defining characteristic. Our dual approach introduces complications with the complexity analysis.

One problem arises because efficient algorithms which solve certain problems for matrices or permutations are not available for black-box groups. We address this issue by introducing oracles for these various algorithms. The reader may assign to these oracles whatever timing estimates are appropriate to the context of interest.

Consider the order oracle $\mathcal{O}$. Since $|G|$ divides $|\mathrm{GL}_d(q)|$, the order of an element of G may be computed, in the black-box context, in time $O(\mu d^4 \log^2 q)$, given a prime factorisation of $|\mathrm{GL}_d(q)|$. But, given this factorisation, the order of an element of $\mathrm{GL}_d(q)$ may be computed with $O(d^3 \log d + d^2 \log d \log\log d \log q)$ field operations; see [26, Lemma 2.7]. Furthermore, if G is a matrix group in defining characteristic, it usually suffices to compute the pseudo-order of an element: this can be computed in the cited time *without* knowing the factorisation.

Next consider the power oracle Π. Using fast exponentiation, the complexity for Π is $O(\mu d \log q)$; but the algorithm in [26, Lemma 10.1] allows us to compute large powers of $g \in \mathrm{GL}_d(q)$ with $O(d^3 \log d + d^2 \log d \log\log d \log q)$ field operations.

Now consider the cost of χ, the oracle to recognise a central quotient of $\mathrm{SL}_2(q)$. The algorithm of [25] produces inverse isomorphisms between a black-box copy of $\mathrm{SL}_2(2^e)$ and the natural copy in time that is polynomial in e. Such an algorithm appears in [8] for $q \equiv 1 \bmod 4$; it is polynomial in $\log q$ and the square of the characteristic of $\mathrm{GF}(q)$. But the only known way of producing such an isomorphism with complexity that is polynomial in $\log q$ when the characteristic is not bounded is the algorithm of [16], which applies when the group is given as a matrix representation in the defining characteristic, and *assumes* a discrete logarithm oracle for $\mathrm{GF}(q)$. If the representation is not in defining characteristic, then the complexity involves q (but remains polynomial in the size of the input).

Babai [3] presented a Monte Carlo algorithm to construct in polynomial time independent nearly uniformly distributed random elements of a finite group. An alternative is the *product replacement algorithm* of Celler *et al.* [15].

A difficulty of another kind arises because our algorithm is recursive. We recurse to two classical groups, each of rank roughly half that of the parent group. In the matrix group case, these groups act faithfully (modulo a central subgroup) on a section of the given module that may also be about half the dimension of the given module, but will often be much smaller. This means that almost all the oracles, including μ, will now be replaced, in these recursive calls, by oracles that run much faster. A consequence is that the time spent in the recursive calls will, at the worst, multiply the complexity of the algorithm by a constant (depending on how much faster these oracles run in smaller cases).

It is difficult to produce a complexity analysis that allows for these complications. We content ourselves with giving the complexity of the main algorithm in three components as above, each being given in terms of oracles that may be used in the input group, and with no reference to the fact that they might be replaced, in subgroups, by faster oracles. The cost of the recursion is provably insignificant if the input is a matrix group in any characteristic, or a permutation group, when the oracles are replaced by faster oracles that apply in the recursive calls.

When the algorithms are applied to (absolutely irreducible) representations of classical groups in the defining characteristic, the complexity should be interpreted as a function of three variables, the dimension d of the natural representation of the group, the dimension n of the given representation, and the size q of the field. In this context, the complexity of our algorithms, for fixed q, is $O(d \log d\, n^3)$ if we assume that we can construct a random element of a group, given by a generating set of bounded size, with a bounded number of group multiplications. While evidence suggests that the algorithm of [15] achieves this, the provable performance is much worse.

8. Realisation and performance

Our implementation in MAGMA accepts as input a permutation or linear representation of $\mathrm{SX}_d(q)$. We use our implementations of [10, 14–16, 30]. We use Schneider's implementations of (the extension to) the algorithm of [1], and also of [18], to write an element of a classical group as an SLP in its standard generators. If $G \leq \mathrm{GL}_n(\mathbb{F})$ is an absolutely irreducible representation of $\mathrm{SX}_d(q)$, with $n \leq d^2$, and $\mathbb{F}$ and $\mathrm{GF}(q)$ have the same characteristic, then Corr's implementation of the Las Vegas algorithms of [17, 32] is used to construct the projective action of G on $\mathrm{GF}(q)^d$; then G can be constructively recognised by our algorithms of [19, 26]. To all individual base cases, we apply (our implementations of) specially designed base-case algorithms or COMPOSITIONTREE [7]. We observe that the latter also readily constructs standard generators for many representations of moderate dimension of $\mathrm{SX}_d(2)$ for $d \leq 20$.

In practice, black-box groups arise as permutation groups or linear groups. Once we construct the subgroup H (or K), we restrict to act on a faithful representation of a central quotient of H (or K) by taking its action on an irreducible section of the given module. All constructive recognition is performed on this faithful representation.

Table 2 displays runtimes of our MAGMA implementation to construct standard generators. All times are in rounded seconds and averaged over 5 runs; the computations were carried out using MAGMA V2.20-3 on a computer with a 2.9 GHz processor. As input we used $\mathrm{SX}_d(q)$ in both its natural and exterior square representations. We apply the algorithm described in this paper to these, ignoring the nature of the input representation.

Our implementation is the first that can construct standard generators for all classical groups in arbitrary matrix representations over all finite fields. It can readily be applied to representations of degree up to about 300. We observe that the runtime is often dominated by evaluations of SLPs.

	Natural representation				Exterior square			
group / q	2^5	2^8	3^4	3^6	2^5	2^8	3^4	3^6
$\mathrm{SL}_{14}(q)$	27	34	44	48	55	106	420	679
$\mathrm{SL}_{20}(q)$	49	64	83	93	413	581	637	946
$\mathrm{SU}_{14}(q)$	4	7	7	15	116	283	534	807
$\mathrm{SU}_{20}(q)$	12	26	15	37	711	977	678	1304
$\mathrm{Sp}_{14}(q)$	6	7	24	37	69	164	146	507
$\mathrm{Sp}_{20}(q)$	21	35	41	57	830	1208	1122	1560
$\Omega^+_{14}(q)$	5	6	97	168	76	411	366	826
$\Omega^+_{20}(q)$	16	26	163	230	297	893	456	995
$\Omega^-_{14}(q)$	7	13	115	124	306	656	564	734
$\Omega^-_{20}(q)$	16	24	208	220	657	879	750	1098
$\Omega_{13}(q)$	–	–	108	140	–	–	160	707
$\Omega_{19}(q)$	–	–	203	228	–	–	792	1103

TABLE 2. Runtimes for constructing standard generators

References

[1] S. Ambrose, S. H. Murray, C. E. Praeger, and C. Schneider. Constructive membership testing in black-box classical groups. Proceedings of The Third International Congress on Mathematical Software. *Lecture Notes in Computer Science*, **6327** (2010), 54–57.

[2] M. Aschbacher and G. M. Seitz. Involutions in Chevalley groups over fields of even order, *Nagoya Math. J.* **63** (1976), 1–91.

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

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

[5] L. Babai, P. Pálfy, and J. Saxl. On the number of p-regular elements in finite simple groups. LMS J. Comput. Math. **12** (2009), 82–119.

[6] László Babai and Endre Szemerédi. On the complexity of matrix group problems, I. In *Proc.* 25*th IEEE Sympos. Foundations Comp. Sci.*, pages 229–240, 1984.

[7] Henrik Bäärnhielm, Derek Holt, C.R. Leedham-Green, and E.A. O'Brien. A practical model for computation with matrix groups. *J. Symbolic Comput.* 2014.
[8] Alexandre Borovik and S. Yalçınkaya. Fifty shades of black. http://arxiv.org/abs/1308.2487.
[9] W. Bosma, J. Cannon, and C. Playoust. The MAGMA algebra system I: The user language. *J. Symbolic Comput.* **24** (1997), 235–265.
[10] J. N. Bray. An improved method of finding the centralizer of an involution. *Arch. Math. (Basel)* **74** (2000), 241–245.
[11] P. A. Brooksbank and W. M. Kantor. On constructive recognition of a black box $\mathrm{PSL}(d, q)$. In *Groups and Computation, III (Columbus, OH, 1999)*, pp. 95–111. Volume 8 of *Ohio State Univ. Math. Res. Inst. Publ.*, de Gruyter, Berlin, 2001.
[12] P. A. Brooksbank. Fast constructive recognition of black-box unitary groups. *LMS J. Comput. Math.* **6** (2003), 162–197.
[13] P. A. Brooksbank and W. M. Kantor. Fast constructive recognition of black box orthogonal groups. *J. Algebra* **300** (2006), 256–288.
[14] P. A. Brooksbank. Fast constructive recognition of black box symplectic groups. *J. Algebra* **320** (2008), 885–909.
[15] F. Celler, C. R. Leedham-Green, S. H. Murray, A. C. Niemeyer, and E. A. O'Brien. Generating random elements of a finite group. *Comm. Algebra* **23** (1995), 4931–4948.
[16] M. D. E. Conder, C. R. Leedham-Green, and E. A. O'Brien. Constructive recognition of $\mathrm{PSL}(2, q)$. *Trans. Amer. Math. Soc.* **358** (2006), 1203–1221.
[17] Brian Corr. Estimation and Computation with Matrices Over Finite Fields. PhD thesis, University of Western Australia, 2013.
[18] E. M. Costi. Constructive membership testing in classical groups. PhD thesis, Queen Mary, University of London, 2009.
[19] H. Dietrich, C. R. Leedham-Green, F. Lübeck, and E. A. O'Brien. Constructive recognition of classical groups in even characteristic. *J. Algebra* **391** (2013), 227–255.
[20] J. Doliskani and E. Schost. Taking roots over high extensions of finite fields. *Math. Comp.* **83** (2014) 435–446.
[21] D. Gorenstein, R. Lyons, and R. Solomon. The classification of finite simple groups, Number 3. Mathematical Surveys and Monographs, AMS, Providence, 1998.
[22] R.M. Guralnick, W.M. Kantor, M. Kassabov, and A. Lubotzky, Presentations of finite simple groups: a quantitative approach. *J. Amer. Math. Soc.* **21**, 711–774, 2008.
[23] W. M. Kantor and A. Lubotzky. The probability of generating a finite classical group. *Geom. Dedicata* **36** (1990), 67–87.
[24] W. M. Kantor and Á. Seress. Black box classical groups. *Mem. Amer. Math. Soc.* **149**, 2001.
[25] W. M. Kantor and M. Kassabov. Black box groups isomorphic to $\mathrm{PGL}(2, 2^e)$. *J. Algebra*, 2014.
[26] C. R. Leedham-Green and E. A. O'Brien. Constructive recognition of classical groups in odd characteristic. *J. Algebra*, **322** (2009), 833–881.
[27] C. R. Leedham-Green and E. A. O'Brien. Short presentations for classical groups. In preparation, 2014.
[28] Martin W. Liebeck. On products of involutions in finite classical groups of even characteristic. J. Algebra, 2014.
[29] Frank Lübeck, Alice C. Niemeyer, and Cheryl E. Praeger. Finding involutions in finite Lie type groups of odd characteristic. *J. Algebra* **321** (2009), 3397-3417.
[30] F. Lübeck, K. Magaard, and E. A. O'Brien. Constructive recognition of $\mathrm{SL}_3(q)$. *J. Algebra* **316** (2007), 619–633.
[31] Vicente Landazuri and Gary M. Seitz. On the minimal degrees of projective representations of the finite Chevalley groups. *J. Algebra* **32**, 418–443, 1974.
[32] Kay Magaard, E.A. O'Brien, and Ákos Seress. Recognition of small dimensional representations of general linear groups. *J. Aust. Math. Soc.* **85**, 229–250, 2008.
[33] G. Malle and D. Testerman. Linear algebraic groups and finite groups of Lie type. Cambridge Studies in Advanced Mathematics, 133. Cambridge University Press, Cambridge, 2011.
[34] Nina Menezes, Martyn Quick, and Colva Roney-Dougal. The probability of generating a finite simple group. *Israel J. Math.* (1) **198** (2013), 371–392.
[35] A. C. Niemeyer and C. E. Praeger. A recognition algorithm for classical groups over finite fields. *Proc. London Math. Soc.* **77** (1998), 117–169.
[36] A. C. Niemeyer and C. E. Praeger. Estimating proportions of elements in finite groups of Lie type. *J. Algebra* **324** (2010), 122–145.
[37] E. A. O'Brien. Algorithms for matrix groups. Groups St Andrews 2009 in Bath II, London Math. Soc. Lecture Note Series **388** (2011), 297–323.
[38] Igor Pak. On probability of generating a finite group. Preprint (1999), available at `math.ucla.edu/~pak/papers/sim.ps`.
[39] Christopher W. Parker and Robert A. Wilson. Recognising simplicity of black-box groups. *J. Algebra* **324** (2010), 885–915.
[40] C. Parker and P. Rowley. Symplectic amalgams. Springer Monographs in Mathematics. Springer-Verlag London, Ltd., London, 2002.

[41] C. E. Praeger, Á. Seress, and S. Yalçınkaya. Generation of finite classical groups by pairs of elements with large fixed point spaces. *J. Algebra*, 2014.

[42] Á. Seress. Permutation group algorithms. Volume 152 of Cambridge Tracts in Mathematics, Cambridge University Press, Cambridge, 2003.

[43] D. E. Taylor. The geometry of the classical groups. Sigma Series in Pure Mathematics, **9**. Heldermann Verlag, Berlin, 1992.

SCHOOL OF MATHEMATICAL SCIENCES, MONASH UNIVERSITY, MELBOURNE, VIC 3800, AUSTRALIA

E-mail address: heiko.dietrich@monash.edu

SCHOOL OF MATHEMATICAL SCIENCES, QUEEN MARY, UNIVERSITY OF LONDON, LONDON E1 4NS, UNITED KINGDOM

E-mail address: c.r.leedham-green@qmul.ac.uk

DEPARTMENT OF MATHEMATICS, PRIVATE BAG 92019, AUCKLAND, UNIVERSITY OF AUCKLAND, NEW ZEALAND

E-mail address: obrien@math.auckland.ac.nz