CONSTRUCTIVE RECOGNITION OF $\mathrm{SL}_3(q)$

F. LÜBECK, K. MAGAARD, AND E.A. O'BRIEN

Abstract. We present a constructive recognition algorithm for groups of Lie type $\mathrm{SL}_3(q)$. This is a necessary component for constructive recognition algorithms of quasisimple groups of Lie type.

1. Introduction

A major research topic over the past decade has been the development of efficient algorithms for the investigation of subgroups of $\mathrm{GL}_d(\mathbb{F}_q)$ where $\mathbb{F}_q$ is a finite field of size $q = p^e$. We refer the interested reader to O'Brien [20] for background related to this work, and to Carter [7] for concepts related to groups of Lie type.

Let $G = \langle S \rangle$ be a finite quasisimple group given by a finite set S of generators. We assume that we know the isomorphism type of the simple quotient of G. Let $H = \langle Y \rangle$ be a known quasisimple group where there exists an epimorphism $H \to G$. An algorithm for *constructive recognition* of G by (H, Y) constructs such an epimorphism which has the property that images and preimages can be computed explicitly. More precisely, a subset $X \subset G$ is constructed such that X generates G and there is a bijection $Y \to X$ which defines a homomorphism $H \to G$. Further, we assume that there are algorithms to write every $h \in H$ as an explicit product of elements in Y and every $g \in G$ as a product of elements in X. If we record how X is constructed from S, this solves the *word problem* for G: every $g \in G$ can be expressed as a product of the given generators in S.

Babai & Szemerédi [1] introduced the *black-box group* model, where group elements are represented by bit-strings of uniform length; the only group operations permissible are multiplication, inversion, and checking for equality with the identity element. A *black-box algorithm* is one which does not use specific features of the group representation, nor particulars of how group operations are performed; it can only

1991 *Mathematics Subject Classification.* 20C20, 20C40.

Key words and phrases. Black-box groups, constructive recognition.

Magaard was partially supported by NSA grant MDA-9049810020. O'Brien was partially supported by the Marsden Fund of New Zealand via grant UOA0412. We thank the referee for helpful criticism and suggestions.

use the operations listed above. Both permutation groups and matrix groups defined over finite fields are covered by this model.

Recently, Kantor & Seress [16] proved the following.

Theorem 1.1. *There is a black-box Las Vegas algorithm which, when given as input a perfect group $G \leq \mathrm{GL}_d(F_q)$ where $G/Z(G)$ is isomorphic to a classical simple group C of known characteristic, produces a constructive isomorphism $C \longmapsto G/Z$.*

Implementations of the algorithm for $\mathrm{PSL}_d(q)$ are available in both GAP [12] and MAGMA [4]. These algorithms do not run in time polynomial in the size of the input: their complexity involves q. A critical obstruction is the search for a p-singular element. If G is a group of Lie type defined over $\mathbb{F}_q$, then $\frac{2}{5q} < \rho(G) < \frac{5}{q}$, where $\rho(G)$ denotes the proportion of p-singular elements in G (see [13] for details). Hence a random search for a transvection, a vital component for the algorithms, needs $O(q)$ selections.

In ongoing work, Kantor & Magaard [17] are developing similar algorithms for the exceptional groups.

All existing constructive recognition algorithms of groups of Lie type of untwisted Lie rank ≥ 3 are recursive and rely on the ability to solve the word problem for the classical groups having untwisted Lie rank 2: namely, $\mathrm{SL}_3(q)$, $\mathrm{Sp}_4(q)$, $\Omega_4^{\pm}(q)$ and $\mathrm{SU}_3(q)$.

Of these, we argue that (P)$\mathrm{SL}_3(q)$ is the most critical case. Consider the following situation. Let G be a finite simple group of Lie type and rank ≥ 2, B a Borel subgroup, U the unipotent radical of B and R the long root group of U labelled by the highest root of the root system of G. If S is a conjugate of R we call S *opposite* to R if the group generated by R and S is isomorphic to $\mathrm{SL}_2(q)$. Let Ω be the set of G-conjugates of R which are opposite to R. That U acts *transitively* on Ω underpins existing algorithmic solutions of the word problem for these groups: if we can compute $u \in U$ with $S_1^u = S_2$, then we can solve the word problem for G. This concept of *effective transitivity* underpins the algorithms of [6], [16] and [17]. Now if moreover G is not unitary, symplectic or ${}^2F_4(q)$, then with high probability, R, S_1 and S_2 generate a subgroup isomorphic to $\mathrm{SL}_3(q)$ and hence u can be obtained inside $\mathrm{SL}_3(q)$. Effective transitivity is also used in [17] to construct the centralizer of a fundamental SL_2.

Recognizing that an effective solution to $\mathrm{SL}_3(q)$ requires the ability to work effectively with $\mathrm{SL}_2(q)$, Brooksbank & Kantor [6] identify that the ultimate obstruction to a polynomial-time algorithm for constructive recognition of the classical groups is $\mathrm{PSL}_2(q)$. Building on the work of [16], they produce black-box polynomial-time constructive recognition algorithms for $\mathrm{PSL}_d(q)$, subject to the availability of an *oracle* to recognize constructively a group having central quotient $\mathrm{PSL}_2(q)$.

A consequence of the work of Landazuri & Seitz [19] is that faithful projective representation of $\mathrm{PSL}_2(q)$ in cross characteristic have degree that is polynomial in q rather than in $\log q$. Hence the critical case is a matrix representation of $\mathrm{SL}_2(q)$ in defining characteristic. Conder & Leedham-Green [9] and Conder, Leedham-Green & O'Brien [10] provide an algorithm which constructively recognizes $\mathrm{SL}_2(q)$ as a linear group in defining characteristic in time polynomial in the size of the input, subject to the availability of a *discrete log oracle*.

We exploit the solution for $\mathrm{SL}_2(q)$ to obtain a new constructive recognition algorithm for $\mathrm{SL}_3(q)$. Our principal result is the following.

Theorem 1.2. *There is a black-box Las Vegas algorithm to constructively recognize a group G whose central quotient is* $\mathrm{PSL}_3(q)$.

The algorithm assumes that a recognition algorithm for groups with central quotient isomorphic to $\mathrm{PSL}_2(q)$ *is available and that we know the prime factors of $q^2 - 1$.*

Let χ be the cost of an invocation of the $(\mathrm{P})\mathrm{SL}_2(q)$ *recognition algorithm, ξ the cost of constructing a random element of G, and μ the cost of a group operation in G.*

The complexity of the algorithm to construct a new generating set X for G is $O(\chi \log q + (\mu \log q + \xi) \log\log q + \mu \log^2 q)$. In time $O(\chi + \mu \log q)$, we can obtain for $g \in G$ a word in X.

We prove this theorem by exhibiting an algorithm with the stated complexity. The algorithm to construct the new generating set requires $O(\log q)$ calls to the $\mathrm{SL}_2(q)$ oracle which is the same complexity as in [6]. The algorithm of [16, 3.6.3] has complexity $O(\xi q e + \mu q \log^2 q)$ where $q = p^e$.

As known group H, we use the standard copy of $\mathrm{SL}_3(q)$. Denote the standard right H-module by V. As generating set Y of H, we use a subset of its Steinberg generators [7, Theorem 12.1.1]. These are non-diagonal matrices of the form $I + N$, where N is a matrix with precisely one nonzero entry.

1.1. **An overview of the paper.** In Section 2 we record various results about the standard copy $H \cong \mathrm{SL}_3(q)$. These are used in later sections to find possible images of our chosen generators of H in the group G under investigation. Various lemmas have the additional hypothesis that $q \neq 2, 3, 4, 7$. The case $q = 7$ requires only a very minor modification of our algorithm, which we identify at the end of Section 3. The other cases are so small that they can be handled directly.

Moreover most of the modifications are only needed in Section 3. The labelling needs no modification and the straight line programs only need modification if $q = 2$.

In Section 3 we show how to find generators for a set of six root subgroups in G which are normalized by a single maximal torus. The root subgroups are then easy to parameterize and this yields image elements of our chosen set of generators of H and determines a homomorphism $\pi : H \to G$; this is explained in Section 4.

In Section 5 we give an algorithm to write an arbitrary $g \in G$ as a product of the images under π of the Steinberg generators of H. Hence, we can compute a preimage $\pi^{-1}(g)$. Applying this algorithm to the user-supplied generators of G, we can prove that π is an epimorphism.

The complexity of the algorithm is mainly determined by the complexity of an $\mathrm{SL}_2(q)$-recognition algorithm used in some of the steps. We repeatedly need to find random elements with specific orders: to find these, we assume only that we know the prime factors of $q^2 - 1$.

In Section 6 we discuss the complexity of the overall algorithm, and finally report on our implementations of the algorithm, which are publicly available in GAP [12] and MAGMA [4].

While our discussion primarily focuses on $\mathrm{SL}_3(q)$, we also identify those few modifications needed for $\mathrm{PSL}_3(q)$.

1.2. **A commentary on two algorithms.** Our principal result is a slight improvement over that of Brooksbank & Kantor [6]: our algorithm applies for $q \geq 7$ whereas that of [6] has the hypothesis that $q \geq 17$. Of potentially greater significance is that our algorithm is demonstrably practical. Its implementation is already a central component of the matrix recognition routines under development in GAP and MAGMA. Since it is significantly different from that of [6], we feel that it is imperative to give a complete and self-contained description of our algorithm.

We now explain the main differences between the two algorithms. Brooksbank & Kantor [6] construct a pair of opposite maximal parabolic subgroups of G, which intersect in a Levi factor L, and their unipotent radicals Q and $Q(\gamma)$. The unipotent radicals are elementary abelian groups of order q^2. Along the way they construct a maximal torus T of L and the six T-invariant root groups of G. We construct L, then T and finally L^x such that $T \subset L \cap L^x$. From this we produce the six T-invariant root groups of G. Once the root groups are found, their elements must be labelled by elements of $\mathbb{F}_q$. In [6] the label of a root element u is obtained by conjugating it into L' where an $\mathrm{SL}_2(q)$-oracle can then be used. We use commutators with other root elements to obtain an element of L' which determines the label of u. How do we write an element of the Borel subgroup as a product of root elements? We use commutators with fixed root elements, see Lemma 5.1, whereas [6] uses commutators with elements of T. It is here that field size is an issue. Finally, we handle effective conjugation, see Lemma 5.2, by producing two subgroups of order $q-1$ in B which must be B-conjugate. The conjugating element of B is found by means of a base change calculation in H. In [6] effective conjugation is handled inside the normalizer of Q and requires computations inside factor groups.

2. The action of H on its natural module

Let H be our standard copy of $\mathrm{SL}_3(\mathbb{F}_q)$, for $q = p^e$, acting as 3×3-matrices from the right on the standard module V.

In this section we characterize certain configurations of elements and subgroups of H up to conjugacy. Some of the following statements are not valid for $q < 8$; we usually note the exceptional cases to the stated results.

If $g, h \in H$, then $[g,h] := g^{-1}h^{-1}gh$; if $U, V \leq H$, then $[U,V]$ is the subgroup generated by the commutators of elements.

For $a, b, c \in \mathbb{F}_q$ let

$$Y_\alpha(a) = \begin{pmatrix} 1 & 0 & 0 \\ a & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad Y_{-\alpha}(a) = \begin{pmatrix} 1 & a & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad Y_\beta(b) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & b & 1 \end{pmatrix},$$

$$Y_{-\beta}(b) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix}, \quad Y_\gamma(c) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ c & 0 & 1 \end{pmatrix}, \quad Y_{-\gamma}(c) = \begin{pmatrix} 1 & 0 & c \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

We identify the indices $\Psi = \{\pm\alpha, \pm\beta, \pm\gamma\}$ with a root system via $\alpha = (1,-1,0)$, $\beta = (0,1,-1) \in \mathbb{R}^3$, $\gamma = \alpha + \beta$.

The sets $Y_\mu := \{Y_\mu(f) \mid f \in \mathbb{F}_q\}$, $\mu \in \Psi$, are root subgroups of H; they are normalized by the maximal torus consisting of diagonal matrices in H. Two root subgroups are *opposite* if they generate a subgroup isomorphic to $\mathrm{SL}_2(\mathbb{F}_q)$.

Let $f_0 = 1, f_1, \ldots, f_{e-1}$ be an $\mathbb{F}_p$-basis of $\mathbb{F}_q$.

Proposition 2.1. *The group $H = \mathrm{SL}_3(\mathbb{F}_q)$ is generated by the elements $Y_\mu(f_i)$, where $\mu \in \Psi$, and $0 \leq i < e$.*

Let $\tilde{H}$ be the group generated by symbols $y_\mu(f_i)$, for $\mu \in \Psi$, $0 \leq i < e$. subject to the following relations:

$$\begin{array}{rcl} y_\mu(f_i)^p & = & 1 \textit{ for } \mu \in \Psi, 0 \leq i < e; \\ [y_\mu(f_i), y_\mu(f_j)] & = & 1 \textit{ for } \mu \in \Psi, 0 \leq i < j < e; \\ [y_\mu(f_i), y_\nu(f_j)] & = & \left\{ \begin{array}{l} 1, \textit{ if } \mu + \nu \notin \Psi \cup \{0\}, \\ y_{\mu+\nu}(C_{\mu\nu} f_i f_j), \textit{ if } \mu+\nu \in \Psi. \end{array} \right. \end{array}$$

$C_{\mu\nu} = -1$ if $(\mu, \nu) \in \{(\alpha, \beta), (\beta, -\gamma), (\gamma, -\beta), (-\alpha, \gamma), (-\beta, -\alpha), (-\gamma, \alpha)\}$ and $C_{\mu\nu} = 1$ otherwise.

For $\rho \in \Psi$, $f \in \mathbb{F}_q$ we write in the last relation $y_\rho(f) := \prod_{i=0}^{e-1} y_\rho(f_i)^{k_i}$ if $f = \sum_{i=0}^{e-1} k_i f_i$ with integers k_i.

Then $y_\mu(f_i) \mapsto Y_\mu(f_i)$ defines an isomorphism $\tilde{H} \to H$.

Proof. The first statement and the observation that the map on generators defines a homomorphism $\tilde{H} \to H$ follows from simple calculations with matrices. That the relations are sufficient to yield an isomorphism is shown in [2, Theorem 4.2]. □

Lemma 2.2. *If $h \in H$ has order $q^2 - 1$, then $z = h^{q+1}$ is H-conjugate to $\begin{pmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ 0 & 0 & \lambda^{-2} \end{pmatrix}$ where λ is a generator of $\mathbb{F}_q^\times$.*

Remark: Observe that $\lambda^{-1}z$ is a pseudo-reflection for $q \neq 2, 4$.

Proof. Observe $y = h^{q-1}$ has order $q+1$ and so $dim([y, V]) = 2$ and $dim(C_V(y)) = 1$. Moreover, since y is a semisimple element of H, we have that $V = C_V(y) \oplus [y, V]$ and y acts irreducibly on $[y, V]$. Thus, by Schur's Lemma, z must act like a scalar, say λ, on $[y, V]$. Since $z \in H$, we see that z must act as the scalar λ^{-2} on $C_V(y)$. Since z has order $q - 1$, either λ or λ^{-2} has order $q - 1$. If q is odd, then the order of λ^{-2} is half the order of λ, so λ must have order $q - 1$ as claimed. If q is even, the squaring and the inverse maps are order preserving and hence both λ and λ^{-2} have order $q - 1$; again the claim follows. □

Lemma 2.3. *If z is as in Lemma 2.2, then z is a generator of the center of $C_H(z)$.*

Proof. We use Lemma 2.2. If $q \in \{2, 4\}$ then z generates the center of G. Otherwise $C_H(z)$ is conjugate to the subgroup of matrices of the form $\begin{pmatrix} a & b & 0 \\ c & d & 0 \\ 0 & 0 & e \end{pmatrix}$, where e is the inverse of the determinant of

$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$. The claim follows since $|z|$ is the order of the center of this subgroup. □

Remark: We observe that for $q \notin \{2,4\}$ the derived group of $C_H(z)$ is a *long root* $\mathrm{SL}_2(q)$. The natural homomorphism from H to $\mathrm{PSL}_3(q)$ is one-to-one on $[C_H(z), C_H(z)]$.

Lemma 2.4. *The proportion of elements of order q^2-1 in H is $\phi(q^2-1)/(2(q^2-1))$, where ϕ is the Euler function.*

Proof. Let $h \in H$ have order q^2-1. Then h has an eigenvalue $a \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q$. The other eigenvalues are a^q and $a^{-q-1} \in F_q$. Hence a is a generator of $\mathbb{F}_{q^2}^\times$ and h generates its centralizer in H. There are $\phi(q^2-1)/2$ conjugacy classes of such elements, and each of these classes has $|H|/(q^2-1)$ elements. □

Let $x \in H$ and let α be an eigenvalue of x on V. Let $E_{\alpha,x}$ denote the α-eigenspace of x.

Lemma 2.5. *Assume $q \notin \{2,3,4,7\}$. If z is as in Lemma 2.2 and s and t are conjugates of z such that $E_{\lambda^{-2},s}$ is not contained in $E_{\lambda,t}$, and $E_{\lambda^{-2},t}$ is not contained in $E_{\lambda,s}$ and $[s,t] \neq 1$, then $S := \langle s,t \rangle$ is conjugate to a subgroup of $C_H(z)$ containing $[C_H(z), C_H(z)]$ if $q \neq 11$. If $q = 11$, then with probability at least $1/2$, $S := \langle s,t \rangle$ is conjugate to a subgroup of $C_H(z)$ containing $[C_H(z), C_H(z)]$. For fixed s the proportion of conjugates t which fulfill the above conditions is at least $(q-3)/q$.*

Proof. By hypothesis $dim(E_{\lambda,s} \cap E_{\lambda,t}) = 1$, since two conjugates of z with the same 2-dimensional eigenspace commute. Hence S is contained in the stabilizer in H of $\langle w \rangle := E_{\lambda,s} \cap E_{\lambda,t}$.

We now claim that $E_{\lambda^{-2},s} \oplus E_{\lambda^{-2},t}$ is an S-invariant complement of $\langle w \rangle$ in V. To see this let $0 \neq v \in E_{\lambda^{-2},s}$. We already know that $V = E_{\lambda^{-2},t} \oplus E_{\lambda,t}$ and hence we can write $v = v_1 + v_2$ with $v_1 \in E_{\lambda^{-2},t}$ and $v_2 \in E_{\lambda,t}$. Now $vt = \lambda^{-2}v_1 + \lambda v_2$ and $E_{\lambda^{-2},s} \oplus E_{\lambda^{-2},t} = \langle v_1, v_2 \rangle$. Clearly $\langle v_1, v_2 \rangle$ is t-invariant and therefore so is $E_{\lambda^{-2},s} \oplus E_{\lambda^{-2},t}$. Reversing the roles of s and t shows that $E_{\lambda^{-2},s} \oplus E_{\lambda^{-2},t}$ is also s-invariant and thus S-invariant.

Hence $(E_{\lambda^{-2},s} \oplus E_{\lambda^{-2},t}) \oplus \langle w \rangle$ is an S-invariant decomposition of V, which shows that S is contained in a conjugate of $C_H(z)$. The order of the projection of s in $\mathrm{PGL}_2(q)$ is $q-1$ if q is not congruent to 1 mod 3 and $(q-1)/3$ otherwise. In fact the action of s^2 on the points of the projective space of the natural $\mathrm{GL}_2(q)$-module is equivalent to the action of $\begin{pmatrix} \lambda^3 & 0 \\ 0 & \lambda^{-3} \end{pmatrix}$. Thus the projection of s^2 into $\mathrm{PSL}_2(q)$ is an element of order $q-1$ respectively $(q-1)/3$. If $q \neq 2, 3, 4$ or 7, then the projection of s^2 lies in a unique subgroup of order $(q-1)/2$ of $\mathrm{PSL}_2(q)$. Moreover, our choice of s and t ensures that projections of s^2 and t^2 lie in distinct subgroups of order $(q-1)/2$ of $\mathrm{PSL}_2(q)$.

When $q \neq 11$, Dickson's theorem on maximal subgroups of $\mathrm{SL}_2(q)$ [14, Chapter II, §8] shows that the projections of s^2 and t^2 must generate $\mathrm{PSL}_2(q)$ and our claim follows. For $q = 11$, a direct calculation using structure constants shows that s^2 and t^2 generate $\mathrm{PSL}_2(q)$ with probability $\geq 1/2$.

We now prove the last claim of the theorem. For fixed s, the condition $E_{\lambda^{-2},t} \not\subset E_{\lambda,s}$ implies that an element conjugating s to t does not map an eigenvector in $E_{\lambda^{-2},s}$ to $E_{\lambda,s}$, so one has to avoid q^2-1 out of q^3-1 possible images. The same estimate holds for s and t interchanged. The condition $[s,t] \neq 1$ is not fulfilled if t is in the centralizer of s. The H-conjugacy class of s intersects this centralizer in two classes, containing 1 and $q(q^2-1)$ elements, respectively. The stated estimate follows from these numbers. □

Lemma 2.6. *The proportion of elements g of order q^2-1 in $\mathrm{GL}_2(q)$ is $\phi(q^2-1)/(2(q^2-1))$. Further g^{q+1} is a generator of $Z(\mathrm{GL}_2(q))$.*

Proof. Such a g of order q^2-1 has eigenvalues a and a^q for some $a \in \mathbb{F}_{q^2}$, and g generates its centralizer in $\mathrm{GL}_2(q)$. This establishes the claimed proportion of such elements.

Further g^{q+1} has eigenvalues $a^{q+1} = (a^q)^{q+1}$ and so is a scalar matrix of order $q-1$ which establishes the second claim. □

Lemma 2.7. *Assume $q \notin \{2, 4\}$. If $h \in H$ has order q^2-1, and z_1, z_2 are conjugate to h^{q+1} and $[z_1, z_2] = 1$ and $z_2 \notin \langle z_1 \rangle$ then there is a basis $\mathcal{B}$ with respect to which*

$$z_1 = \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ 0 & 0 & \lambda^{-2} \end{pmatrix}$$

and

$$z_2 = \begin{pmatrix} \lambda^{-2} & 0 & 0 \\ 0 & \lambda & 0 \\ 0 & 0 & \lambda \end{pmatrix}.$$

Proof. Clear, since z_1 and z_2 are commuting semisimple elements. □

Lemma 2.8. *Assume $q \notin \{2, 4\}$, and consider z_1, z_2 as in Lemma 2.7. The root groups $Y_{\pm\alpha}$ are the unique z_2-invariant root subgroups of $C_H(z_1)$ and $Y_{\pm\beta}$ are the unique z_1-invariant root subgroups of $C_H(z_2)$.*

Proof. A root subgroup in $\mathrm{GL}_2(q)$ is the Sylow p-subgroup of the centralizer of a vector in the natural module of $\mathrm{GL}_2(q)$. If the eigenvalues λ and λ^{-2} are distinct, then z_i can stabilize at most two 1-dimensional subspaces of the natural module: that is, at most two root subgroups of $C_H(z_{i+1})$ (index taken mod 2). Since we have already exhibited two z_i-invariant root groups in $C_H(z_{i+1})$, our claim follows. □

Lemma 2.9. *Assume $q > 3$. Let $t := \begin{pmatrix} \lambda & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \lambda^{-1} \end{pmatrix}$ where λ is a generator of $\mathbb{F}_q^\times$. Then t normalizes precisely the six root groups, Y_μ for $\mu \in \Phi$, of H.*

Proof. A root group (transvection group) of H is uniquely determined by a maximal flag of V: namely, a 1-space contained in a 2-space of V. A 2-space is t-invariant precisely if it is the sum of eigenspaces of t. The number of t-invariant 2-spaces is 3 and the number of t-invariant subspaces contained in a fixed 2-space is 2, giving a total of six invariant maximal flags. The claim follows. □

Lemma 2.10. *Let L denote the set of lower triangular matrices contained in H. Then $N_H(Y_\gamma) = L$.*

Proof. Easy calculation. □

Lemma 2.11. *Let Z be a root subgroup of H such that $S = \langle Y_\gamma, Z\rangle$ is isomorphic to $\mathrm{SL}_2(q)$. Let t be the element from Lemma 2.9 and $T = \langle t\rangle$. Let $x \in L$ and suppose that T^x normalizes Z. Then $Z = Y_{-\gamma}^x$.*

Proof. Lemma 2.10 implies that x normalizes Y_γ. Hence T^x normalizes Y_γ. From Lemma 2.9 we deduce that T^x normalizes six root groups, one of which is Y_γ. Moreover only one of the invariant root groups is opposite to Y_γ. By hypothesis Z is a T^x-invariant root group opposite to Y_γ. The same is true for $Y_{-\gamma}^x$. The claim follows by uniqueness. □

3. Finding a set of root subgroups in G

Recall that H is our standard copy of $\mathrm{SL}_3(q)$ and has root groups Y_μ for $\mu \in \Psi = \{\pm\alpha, \pm\beta, \pm\gamma\}$.

Let $G = \langle S\rangle$ be a black-box copy of $\mathrm{SL}_3(q)$ with $q \notin \{2, 3, 4, 7\}$. We comment on these exceptional cases below.

The first step of our constructive recognition algorithm for $\mathrm{SL}_3(q)$ is to produce six subgroups X_μ in G which will be the images of the subgroups Y_μ in H. We assume that we can constructively recognize groups isomorphic to $\mathrm{SL}_2(q)$.

The algorithm to produce the six subgroups is the following.

(1) Search G randomly for an element x of order $q^2 - 1$.

(2) Set $s = x^{q+1}$.

(3) Find a conjugate t of s such that $\langle s, t\rangle \cong \mathrm{GL}_2(q)$.

(4) Find an element y in $\langle s, t\rangle$ of order $q^2 - 1$ and set $g_\alpha = y^{q+1}$.

(5) Find another conjugate u of s such that $\langle u, g_\alpha\rangle \cong \mathrm{GL}_2(q)$.

(6) Find an element w in $\langle u, g_\alpha\rangle$ of order $q^2 - 1$ and set $g_\beta = w^{q+1}$.

(7) Consider $K = \langle s, t\rangle'$. Observe that $K = C_G(g_\alpha)'$. Establish a constructive isomorphism π_α from $\mathrm{SL}_2(q)$ to K.

(8) Observe that g_β is central in $\langle u, g_\alpha\rangle$, as is the $(q + 1)$-st power of any element of order $q^2 - 1$ in $\mathrm{GL}_2(q)$. Thus g_β centralizes g_α and hence it normalizes K.

    Modify π_α so that its images of the standard root subgroups consisting of lower and upper triangular matrices in $\mathrm{SL}_2(q)$ are g_β-invariant root subgroups of G. Label these root subgroups X_α and $X_{-\alpha}$ respectively. This step is considered in detail below.

(9) Consider $M = \langle u, g_\alpha\rangle'$. Observe that $M = C_G(g_\beta)'$. Establish a constructive isomorphism π_β from $\mathrm{SL}_2(q)$ to M.

(10) Modify π_β so that the images in G of the standard root subgroups of this $\mathrm{SL}_2(q)$ are g_α-invariant. Each root group will centralize exactly one of $\{X_\alpha, X_{-\alpha}\}$. By X_β we denote the g_α-invariant root group that commutes with $X_{-\alpha}$. We name the other $X_{-\beta}$.

(11) Set $X_\gamma = [X_\alpha, X_\beta]$ and $X_{-\gamma} = [X_{-\alpha}, X_{-\beta}]$.

Recall that our algorithm assumes the existence of an $\mathrm{SL}_2(q)$ oracle. If the input to this oracle is black-box, we employ the algorithm of [16] which has complexity involving q. If the input is a matrix representation in the defining characteristic, then the complexity of the algorithm in [10] involves $\log q$. The complexity of the $\mathrm{SL}_2(q)$ oracle influences how we compute the invariant root subgroups in steps (8) and (10), a task we now discuss.

If G is a black-box group, we search for random elements k of order $q + 1$ in K and repeatedly conjugate the images under π_α of the standard root subgroups in $\mathrm{SL}_2(q)$ with k until we find some which are invariant under g_β. If $p = 2$, each k yields all pairs of opposite root subgroups in K; otherwise k yields half of these pairs. The composite of π_α with conjugation by the appropriate power of k now defines the modified map. The proportion of elements of order $q + 1$ in K is at least $1/\log\log q$.

If $\mathrm{SL}_2(q)$ can be recognized with complexity smaller than $O(q)$, then we can modify π_α more efficiently as follows. Observe that the action of g_β on K can be pulled back to the standard $\mathrm{SL}_2(q)$. We compute the matrix $A' = (g_\beta^{-1}(A\pi_\alpha)g_\beta)\pi_\alpha^{-1} \in \mathrm{SL}_2(q)$ for each A in

$$\{\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}\}.$$

The action of g_β on K is a conjugation action within $\langle s, t\rangle \cong \mathrm{GL}_2(q) \supset K \cong \mathrm{SL}_2(q)$. Hence, there is a 2×2-matrix T whose conjugation action describes the pulled back action on the standard $\mathrm{SL}_2(q)$. This is well-defined up to a scalar in $\mathbb{F}_q$ and is easily computed from the linear equations $AT = TA'$ for the two pairs A, A' as above.

Now T has order dividing $q-1$ (since $|g_\beta| = q-1$) and so it is diagonalizable over $\mathbb{F}_q$. Let B be a matrix whose rows form a basis of eigenvectors of T. With respect to this basis, T is diagonal and so normalizes the corresponding standard root subgroups. We define $\pi'_\alpha : A \mapsto \pi_\alpha(B^{-1}AB)$ for $A \in \mathrm{SL}_2(q)$, and now replace π_α by π'_α. Then the images under the new π_α of the standard root subgroups in $\mathrm{SL}_2(q)$ are g_β-invariant.

The lemmas in Section 2 show that all searches in steps (1) to (6) will quickly be successful. With the various proportions given in the lemmas, it is easy to determine in each step the number of sample elements to consider to ensure success of the step with some prescribed probability. Lemma 2.6 establishes that g_α is a generator for the center of $\langle s, t\rangle \cong \mathrm{GL}_2(q)$ and g_β is a generator for the center of $\langle u, g_\alpha\rangle \cong \mathrm{GL}_2(q)$. Lemmas 2.7 and 2.8 guarantee that steps (7)–(11) produce the correct set of root subgroups.

Recall that Lemma 2.5 does not apply for $q = 7$. However, we can readily modify the algorithm to construct the root groups in this case: in steps (3) and (5), we construct $\mathrm{GL}_2(q)$ as the centralizer of the involution obtained by powering s and u respectively; see [5] for the relevant algorithm. The rest of this algorithm applies unchanged. As noted in the introduction, the remaining exceptional cases ($q = 2, 3, 4$) can be handled readily.

If $G \cong \mathrm{PSL}_3(q)$ and $q \equiv 1 \bmod 3$, then we modify this algorithm to search in step (1) for an element of order $(q^2-1)/3$. If $q \not\equiv 1 \bmod 3$, then the algorithm applies without modification.

4. Labelling the elements in the root subgroups of G

Our goal is to define a constructive epimorphism $\pi : H \longrightarrow G$. In particular we must define $Y_\mu(f_i)\pi$ for every root $\mu \in \Psi$ and every f_i, $0 \le i < e$, in our $\mathbb{F}_p$-basis of $\mathbb{F}_q$.

In Section 3, we identified subgroups X_μ of G which are the images of the Y_μs. We now parameterize each root subgroup by the additive group $\mathbb{F}_q$. The algorithm is the following.

(1) First label $X_{\pm\alpha}$ using the map π_α from step (8) in Section 3. Set $Y_\alpha(f_i)\pi = X_\alpha(f_i) = \begin{pmatrix} 1 & 0 \\ f_i & 1 \end{pmatrix} \pi_\alpha$ and $Y_{-\alpha}(f_i)\pi = X_{-\alpha}(f_i) = \begin{pmatrix} 1 & f_i \\ 0 & 1 \end{pmatrix} \pi_\alpha$.

(2) Now we can freely choose $X_\beta(1)$ in X_β, because the diagonal matrices in $N_{\mathrm{GL}_3(q)}(H)$ which centralize Y_α and $Y_{-\alpha}$ act transitively on Y_β. We use π_β from step (10) in Section 3 and choose $Y_\beta(1)\pi = X_\beta(1) = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \pi_\beta$. Then, using π_β we also know that
$$Y_{-\beta}(1)\pi = X_{-\beta}(1) = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \pi_\beta.$$

(3) Now the labelling for all root subgroups is uniquely determined by the relations given in Section 2.1:
$$Y_\gamma(f_i)\pi = X_\gamma(f_i) = [X_\beta(1), X_\alpha(f_i)],$$
$$Y_{-\gamma}(f_i)\pi = X_{-\gamma}(f_i) = [X_{-\alpha}(f_i), X_{-\beta}(1)];$$
the remaining elements of $X_{\pm\beta}$ are determined by:
$$Y_\beta(f_i)\pi = X_\beta(f_i) = [X_\gamma(1), X_{-\alpha}(f_i)],$$
$$Y_{-\beta}(f_i)\pi = X_{-\beta}(f_i) = [X_\alpha(f_i), X_{-\gamma}(1)].$$

This completes the definition of π. To establish that π is a homomorphism, we check that the $X_\mu(f_i)$ satisfy the Steinberg relations given in Lemma 2.1. If they do, then we can proceed.

If $g \in X_\mu$ for some μ, then we can determine its label using similar relations as in (3) above. An easily computed commutator of g with an appropriate $X_\nu(1)$ is in X_α or $X_{-\alpha}$. We determine its preimage under π_α in the standard root subgroup of $\mathrm{SL}_2(q)$ and read off its label.

5. The straight-line programs

In Section 4 we defined a homomorphism $\pi : H \longrightarrow G$ by defining for every root μ and every $t \in \mathbb{F}_q$ an element $X_\mu(t)$ of G. We now show how to express elements of G as products of the elements $X_\mu(t)$.

We can apply this algorithm to the user supplied generators S of G. If we find preimages for each, we know that π is surjective.

Define $n_\gamma(\lambda) := X_{-\gamma}(\lambda)X_\gamma(-\lambda^{-1})X_{-\gamma}(\lambda)$ and $h_\gamma(\lambda) := n_\gamma(\lambda)n_\gamma(-1)$; and similarly $h_\alpha(\lambda), h_\beta(\lambda)$.

Let B be the subgroup of G generated by the $h_\alpha(t)$s, the $h_\beta(t)$s and the groups X_μ, where $\mu \in \{\alpha, \beta, \gamma\}$ is a positive root. Then B is the standard Borel subgroup of G and the image under π of the set of lower triangular matrices L of H. Recall from Lemma 2.10 that $N_H(Y_\gamma) = L$ and thus $N_G(X_\gamma) = B$.

Let $g \in G$. We now demonstrate how to write g as a word in the $X_\mu(t)$.

1. If $\langle X_\gamma, X_{-\gamma}^g \rangle \cong \mathrm{SL}_2(q)$, set $x = 1$, else select random $x \in \pi(H)$ until $\langle X_\gamma, X_{-\gamma}^{gx} \rangle \cong \mathrm{SL}_2(q)$. Since x is constructed as a random word in the $X_\mu(t)$, we can record its factorisation. Now $X_{-\gamma}^{gx}$ is guaranteed to be opposite to X_γ.
2. B acts transitively on root groups of G lying opposite to X_γ. Hence we find $b \in B$ such that $X_{-\gamma}^{gxb} = X_{-\gamma}$. We describe how to do this in Lemma 5.2 below. Lemma 5.1 shows how to write $b \in B$ as a product of root elements.
3. The element gxb lies in $N_G(X_{-\gamma})$. Hence $n_\gamma(1)$ permutes X_γ and $X_{-\gamma}$. Thus $(gxb)^{n_\gamma(1)} = b_1$ lies in $N_G(X_\gamma) = B$. We use Lemma 5.1 to express b_1 as a product of root elements. Hence $g = b_1^{n_\gamma^{-1}(1)} b^{-1} x^{-1}$ is a known product of root elements.

We comment on step (1) in more detail. The probability that a selected x fails is $2d/(q^2+q+1)$, where $d = \gcd(3, q-1)$; this is the sum of the indexes of the maximal parabolics containing X_γ.

Kantor [15] and Cooperstein [11] show that two root subgroups of a group of Lie type either generate a root $\mathrm{SL}_2(q)$ or a nilpotent group, which in this case has class at most 2. Hence we decide if $S = \langle X_\gamma, X_{-\gamma}^{gx} \rangle$ is nilpotent of class 2 by computing certain commutators; if not, then S is a root $\mathrm{SL}_2(q)$, and we must now recognize it constructively in order to apply Lemma 5.2.

Lemma 5.1. *If $x \in B$, the standard Borel subgroup of G, then there is an algorithm to express x as a straight-line program in the Steinberg generators of G which requires identification of the labels of up to six root elements.*

Proof. Since $x \in B$, following [7, Theorem 5.3.3, Corollary 8.4.4] we can write

$$x = h_\alpha(t_\alpha) h_\beta(t_\beta) X_\alpha(s_\alpha) X_\beta(s_\beta) X_\gamma(s_\gamma).$$

We need to determine each of the five labels. Computing inside $B\pi^{-1} = L$ and mapping the result back to B, we observe that

- $X_\alpha(1)^x = X_\alpha(t_\beta^{-1} t_\alpha^2) X_\gamma(-s_\beta t_\beta^{-1} t_\alpha^2)$,
- $X_\beta(1)^x = X_\beta(t_\beta^2 t_\alpha^{-1}) X_\gamma(s_\alpha t_\beta^2 t_\alpha^{-1})$,
- $X_\gamma(1)^x = X_\gamma(t_\alpha t_\beta)$.

We now determine the label c of $X_\gamma(1)^x$ as explained at the end of Section 4. Observe that:

- $[X_\alpha(1)^x, X_\beta(-1)] = X_\gamma(t_\beta^{-1}t_\alpha^2)$; we compute this commutator and determine its label, say a.
- $[X_\alpha(-1), X_\beta(1)^x] = X_\gamma(t_\beta^2 t_\alpha^{-1})$; we compute this commutator and determine its label, say b.

Thus we learn that $a = t_\beta^{-1}t_\alpha^2$, $b = t_\beta^2 t_\alpha^{-1}$, $c = t_\alpha t_\beta$. To determine s_α and s_β, we construct explicitly other elements of X_γ, determine the labels, and so deduce the values. More precisely, observe that $X_\alpha(-a)X_\alpha(1)^x = X_\gamma(-s_\beta a) \in X_\gamma$. If d is the corresponding label, then $s_\beta = -d/a$. Similarly $X_\beta(-t_\beta^2 t_\alpha^{-1})X_\beta(1)^x = X_\gamma(s_\alpha t_\beta^2 t_\alpha^{-1})$. Hence we deduce that $s_\alpha = (s_\alpha t_\beta^2 t_\alpha^{-1})/(t_\beta^2 t_\alpha^{-1})$.

Observe that $t_\alpha^3 = ac$ and $t_\beta = c/t_\alpha$. Thus t_α is determined up to a cube root of ac: namely, up to an element of $Z(G)$. Once a choice for t_α has been made, t_β is uniquely determined. Let t_α be a cube root of ac; construct the corresponding $h_\alpha(t_\alpha), h_\beta(t_\beta)$. If $\pi(H) = \mathrm{PSL}_3(q)$ then every choice of t_α works. Otherwise we construct

$$h_\beta^{-1}(t_\beta)h_\alpha^{-1}(t_\alpha)xX_\beta(-s_\beta)X_\alpha(-s_\alpha)$$

and decide, by taking the p-th power and comparing to the identity element, if it is a member of X_γ. If so, we read off its label s_γ and so deduce the complete list of defining labels for x. □

Finally we describe how to construct the element b used in Step 2.

Lemma 5.2. *If $\langle X_\gamma, X_{-\gamma}^{gx}\rangle \cong \mathrm{SL}_2(q)$, then there is an algorithm to construct an element b in B such that $X_{-\gamma}^{gxb} = X_{-\gamma}$ which requires one* $\mathrm{SL}_2(q)$*-recognition and two parameterizations of elements in B using Lemma 5.1.*

Proof. In $\mathrm{SL}_2(q)$ each pair of opposite root subgroups determines a unique subgroup of order $q-1$ (a maximal torus) as intersection of their normalizers. For X_γ and $X_{-\gamma}$ this is $K = \langle h_\gamma(\lambda)\rangle \subset \langle X_\gamma, X_{-\gamma}\rangle$ for a generator λ of $\mathbb{F}_q^\times$. We use an $\mathrm{SL}_2(q)$-recognition algorithm to recognize a corresponding subgroup $K' = \langle h_\gamma(\lambda)'\rangle \subset \langle X_\gamma, X_{-\gamma}^{gx}\rangle$. Note that both K and K' are contained in B. By Lemma 2.11 every $b \in B$ that conjugates K' to K will conjugate $X_{-\gamma}^{gx}$ to $X_{-\gamma}$.

We now describe how to find such a $b \in B$. We first identify $h_\gamma(\lambda)' \in B$ using Lemma 5.1 and observe that $h_\gamma(\lambda)'$ is conjugate to $h_\gamma(\mu)$ for some $\mu \in \mathbb{F}_q^\times$. Now the eigenvalues of $h_\gamma(\lambda)'$ are $\mu, 1, \mu^{-1}$ since $h_\gamma(\lambda)'$

is contained in a long root $\mathrm{SL}_2(q)$. Thus

$$h_\gamma(\lambda)'\pi^{-1} = \begin{pmatrix} \mu & 0 & 0 \\ \rho & 1 & 0 \\ \tau & \sigma & \mu^{-1} \end{pmatrix}.$$

The eigenspaces of $h_\gamma(\lambda)'\pi^{-1}$ are spanned by the vectors

$$v_1 = (1,0,0), \quad v_2 = (\frac{\rho}{1-\mu}, 1, 0), \quad v_3 = (\frac{\rho\sigma + (\mu^{-1}-1)\tau}{(\mu^{-1}-\mu)(\mu^{-1}-1)}, \frac{\sigma}{\mu^{-1}-1}, 1)$$

respectively. Let $C \in H$ be the matrix having rows v_1, v_2, v_3. Now $(h_\gamma(\lambda)'\pi^{-1})^{C^{-1}}$ is a diagonal matrix contained in $\langle h_\gamma(\lambda)\rangle\pi^{-1}$. Hence $b = C^{-1}\pi$ is our desired matrix. To see this, observe that b is constructed so as to conjugate $\langle h_\gamma(\lambda)'\rangle$ to $\langle h_\gamma(\lambda)\rangle$ and so, by Lemma 2.11, it will conjugate $X_{-\gamma}^{gx}$ to $X_{-\gamma}$. This proves our lemma. □

Remark: Lemma 5.2 is one key component in the algorithms under development in [17] for constructively recognizing exceptional groups of Lie type of BN-pair rank at least 2. Compare step (3) with [16, 3.1.3, 3.3.2].

Remark: We can decide whether $H\pi$ is $\mathrm{SL}_3(q)$ or $\mathrm{PSL}_3(q)$ by evaluating $z\pi$ for $z \in Z(\mathrm{SL}_3(q))$. A generator for the center of H can be readily obtained as a word in the Y_μs: it is $h_\alpha(\omega)h_\beta(\omega^2)$ where ω is a primitive cube root of unity in $\mathbb{F}_q$ if such exists.

6. The complexity of the algorithm

Observe that the constructive isomorphism from H to G is set up once. Two principal components are needed to establish this isomorphism. The first outlined in Section 3 identifies the root groups in G. The second outlined in Section 4 labels an $\mathbb{F}_p$-basis for each root group.

Lemma 2.4 shows that the probability that a random element of G has order q^2-1 is $\phi(q^2-1)/2(q^2-1)$. Since this value is at least $1/\log\log q$ (see [18, §II.8]), we expect to make at most $O(\log\log q)$ random selections to find a suitable element in step (1) of Section 3.

If we know the prime factorisation of q^2-1, then testing if a group element has precisely this order takes at most $O(\mu \log q)$ time (by repeated squaring).

Seress [21, Theorem 2.3.9] describes a polynomial-time Monte Carlo algorithm to compute the derived group of a black-box group. For $\mathrm{SL}_3(q)$ it requires $O(\log^2 q)$ group operations.

There are two calls to the oracle to recognize constructively $\mathrm{SL}_2(q)$. Their cost depends on the chosen representation. If the input is the natural representation, then the complexity of the algorithm in [10] is

$O(\xi \log\log q)$; if the input is an absolutely irreducible group of $k \times k$ matrices in defining characteristic, then the complexity is $O(k^5 \log\log q)$. Recall that this algorithm assumes the availability of a discrete log oracle. Otherwise we employ the algorithm of [16] which has complexity $O(q)$.

In steps (8) and (10), we map two matrices from $\mathrm{SL}_2(q)$ to G. Evaluating π_α needs $O(\log q)$ group operations, each with cost μ; the complexity of evaluating π_α^{-1} is the same as that of an $\mathrm{SL}_2(q)$ recognition, and so has cost $O(\chi)$. Conjugation in G costs $O(1)$ group operations.

The labelling of the root groups requires construction of $O(\log q)$ straight-line programs of length $O(\log q)$ for elements of $(\mathrm{P})\mathrm{SL}_2(q)$, and the evaluation of their images in G. The construction of each straight-line program takes $O(\log q)$ field operations.

Each writing of an element of G as a word in its Steinberg generators requires us to recognize constructively a copy of $\mathrm{SL}_2(q)$. Further, we must identify the labels of individual elements of X_γ; we do this by constructing the preimage of conjugates of these elements in Y_α. Each of the two calls to the algorithm of Lemma 5.1 requires six such identifications.

Of course, the value of μ in the statement of Theorem 1.2 depends on the actual representation. For a matrix group of degree d defined over a finite field, we assume that field operations are carried out in constant time and so group operations can be performed in at most time $O(d^3)$.

Hence the complexity of the algorithm is that stated in Theorem 1.2.

7. Implementation and performance

Babai *et al.* [3] present a Monte Carlo polynomial-time algorithm to identify the non-abelian composition factor of a quasisimple black-box group of Lie type in known defining characteristic.

As a preprocessing step to our algorithm, we expect that this algorithm has been employed to conclude with high probability that the quasisimple input group G has $\mathrm{PSL}_3(q)$ as a composition factor.

In theory we are concerned with a group isomorphic to $\mathrm{PSL}_3(q)$, and which will be defined modulo scalars; but in practice we deal with linear groups, so we have a subgroup G of $\mathrm{GL}_3(q)$ that is isomorphic, modulo scalars, to $\mathrm{PSL}_3(q)$. We may also replace G by its derived group, so that G is isomorphic to $\mathrm{PSL}_3(q)$ or to $\mathrm{SL}_3(q)$.

Algorithms to generate random elements of a finite group are discussed in [21, pp. 26-30]. Our implementation uses the algorithm of

[8]; after an initial preprocessing stage, the cost of obtaining a random element is two group multiplications.

Recall that in step (3) of Section 3 we must decide if $K = \langle s, t \rangle \cong \mathrm{GL}_2(q)$. We first use the "naming" algorithms to decide if K contains $\mathrm{SL}_2(q)$ and then use (projective) orders or determinants of elements depending on the representation.

Implementations of the algorithm are publicly available in GAP [12] and MAGMA [4]. The latter uses O'Brien's implementation of the constructive recognition algorithm for $(\mathrm{P})\mathrm{SL}_2(q)$ [10]. The computations reported in Table 1 were carried out using MAGMA V2.11 on a Pentium IV 2.8 GHz processor. The input to the algorithm is a representation of $(\mathrm{P})\mathrm{SL}_3(p^e)$ given as a subgroup of $\mathrm{GL}_d(p^e)$. In the column entitled "Time", we list the CPU time in seconds needed to construct the homomorphism between the standard copy and the input representation.

TABLE 1. Performance of implementation for some groups

d	p	e	Time
3	5	4	0.1
10	5	4	0.3
45	5	4	23.0
3	7	10	1.2
15	7	10	36.0
3	11	8	1.9
15	11	8	18.9

REFERENCES

[1] László Babai and Endre Szemerédi. On the complexity of matrix group problems, I. In *Proc. 25th IEEE Sympos. Foundations Comp. Sci.*, pages 229–240, 1984. 1

[2] L. Babai, A.J. Goodman, W.M. Kantor, E.M. Luks and P.P. Pálfy. Short presentations for finite groups. *J. Algebra*, **194**, 79–112, 1997. 2

[3] 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(4), 383–401, 2002. 7

[4] W. Bosma, J. Cannon, and C. Playoust. The MAGMA algebra system I: The user language. *J. Symbolic Comput.*, 24, 235–265, 1997. 1, 1.1, 7

[5] J.N. Bray. An improved method of finding the centralizer of an involution, *Arch. Math. (Basel)* **74**, 241–245, 2000. 3

[6] Peter A. Brooksbank and William M. Kantor. On constructive recognition of a black box $\mathrm{PSL}(d, q)$. In *Groups and Computation, III (Columbus, OH, 1999)*, volume 8 of *Ohio State Univ. Math. Res. Inst. Publ.*, de Gruyter, Berlin, 95–111, 2001. 1, 1, 1.2

[7] Roger Carter. Simple groups of Lie Type. Wiley-Interscience, 1989. 1, 1, 5

[8] 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**, 4931–4948, 1995. 7

[9] Marston Conder and Charles R. Leedham-Green. Fast recognition of classical groups over large fields. In *Groups and Computation, III (Columbus, OH, 1999)*, volume 8 of *Ohio State Univ. Math. Res. Inst. Publ.*, de Gruyter, Berlin, 113–121, 2001. 1

[10] 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, 1203–1221, 2006. 1, 3, 6, 7

[11] Bruce N. Cooperstein. Subgroups of exceptional groups of Lie type generated by long root elements. II. Characteristic two. *J. Algebra* **70** (1981), no. 1, 283–298. 5

[12] The GAP Group. GAP – Groups, Algorithms, and Programming, Version 4.4.7 (http://www.gap-system.org), 2006. 1, 1.1, 7

[13] R.M. Guralnick and F. Lübeck, On p-singular elements in Chevalley groups in characteristic p. In *Groups and Computation, III (Columbus, OH, 1999)*, volume 8 of *Ohio State Univ. Math. Res. Inst. Publ.*, pages 169–182, de Gruyter, Berlin, 2001. 1

[14] B. Huppert. Endliche Gruppen I. Springer-Verlag, 1967. 2

[15] William M. Kantor. Subgroups of classical groups generated by long root elements. *Trans. Amer. Math. Soc.* **248** (1979), no. 2, 347–379. 5

[16] William M. Kantor and Ákos Seress. Black box classical groups. *Mem. Amer. Math. Soc.*, 149(708):viii+168, 2001. 1, 1, 1, 3, 5, 6

[17] W.M. Kantor and K. Magaard, Constructive recognition for exceptional groups of Lie type. Preprint. 1, 5

[18] D.S. Mitrinović, J. Sándor and B. Crstici. Handbook of Number Theory. *Mathematics and Its Applications* **351**, Kluwer Academic Publishers, 1996. 6

[19] V. Landazuri and G.M. Seitz. On the minimal degrees of projective representations of the finite Chevalley groups. *J. Algebra* 32, 418–443, 1974. 1

[20] E.A. O'Brien. Towards effective algorithms for linear groups. *Finite Geometries, Groups and Computation*, Walter de Gruyter, Berlin, 163–190, 2006. 1

[21] Ákos Seress. Permutation group algorithms. *Cambridge Tracts in Mathematics*, volume 152, Cambridge University Press, Cambridge, 2003.

6, 7

RWTH Aachen, Lehrstuhl D für Mathematik, Templergraben 64, 52062 Aachen, Germany.

Wayne State University, Detroit, MI 48202, USA.

Department of Mathematics, University of Auckland, Private Bag 92019, Auckland, New Zealand.

E-mail address: `frank.luebeck@math.rwth-aachen.de`

E-mail address: `kaym@math.wayne.edu`

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