Constructive membership in black-box groups

P.E. Holmes, S.A. Linton, E.A. O'Brien, A.J.E. Ryba and R.A. Wilson

Abstract

We present an algorithm to reduce the constructive membership problem for a black-box group G to three instances of the same problem for involution centralisers in G. If G is a simple group of Lie type in odd characteristic, then this reduction can be performed in (Monte Carlo) polynomial time.

1 Introduction

A vital component of many group-theoretic algorithms is an efficient solution of the *constructive membership problem* which may be defined as follows: given a finite group $G = \langle X \rangle$, and $g \in G$, express g as a straight-line program in X.

One may intuitively think of a *straight-line program* (SLP) for g as an efficiently stored group word on X that evaluates to g. For a formal definition, we refer the reader to [30, p. 10]. While the length of a word in a given generating set constructed in m multiplications and inversions can increase exponentially with m, the length of the corresponding SLP is *linear* in m. Babai & Szemerédi [5] prove that every element of G has an SLP of length at most $O(\log^2 |G|)$ in every generating set.

The concept of a *black-box group* was also introduced in [5]. In this model, 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. Permutation groups, groups of words with a confluent rewriting system, and matrix groups defined over finite fields are covered by this model. Over the past decade, a major research project, initiated by Babai and Beals, seeks to develop polynomial-time algorithms to determine the abstract group-theoretic structure of a black-box group. We refer the reader to [7] for an excellent account of this work.

Seress [30, p. 17] defines a *black-box algorithm* as one which does not use specific features of the group representation, nor particulars of how group operations are performed; it can only

We thank the referee for extensive and most helpful commentary. O'Brien was partially supported by the Marsden Fund of New Zealand via grant UOA412.

use the operations listed above. However, a common assumption is that *oracles* are available to perform certain tasks: for example an *order oracle* to compute the order of an arbitrary element. Babai & Beals [7] prove if the primes dividing the order of a black-box group are known, then the order of an element can be computed in polynomial time.

Many of the algorithms developed for black-box groups rely on random selections. Babai [4] presents a black-box Monte Carlo algorithm to construct in polynomial time nearly uniformly distributed random elements of a finite group. An alternative is the *product replacement algorithm* of Celler *et al.* [15]. That this also runs in polynomial time was established by Pak [28]. For a discussion of both algorithms, we refer the reader to [30, pp. 26–30].

In this paper, we show that the constructive membership problem in a black-box group G with order oracle can be reduced to three instances of the same problem for involution centralisers in G. Our *reduction algorithm* applies to all such groups. However, if G has no non-central involutions, then the algorithm is not effective; even if it is successful, the reduction may run in time exponential in the size of the input. We prove that the reduction algorithm runs in Monte Carlo polynomial time for the finite simple groups of Lie type defined over fields of odd characteristic.

We establish some notation. If the elements of a black-box group G are represented by bit-strings of uniform length n, then n is the *encoding length* of G and $|G| \leq 2^n$. If G also has Lie rank r and is defined over a field of size q, then $r = O(\sqrt{n})$ and $\log q = O(n)$. Let μ, ξ and ρ denote the costs of a group operation, constructing a random element of G, and an order oracle respectively.

Our principal result is the following.

Theorem 1 *Let G be a black-box group having an encoding of length n and equipped with an order oracle. There is a black-box Monte Carlo algorithm which reduces the constructive membership problem for G to three instances of the same problem for involution centralisers of G. Let $\varepsilon > 0$ denote the probability that the algorithm fails. If G is a simple group of Lie type defined over a field of odd characteristic, then this reduction algorithm is polynomial and can be carried out in time $O(n^{3/2}(\xi + \rho)\log(1/\varepsilon) + n\mu)$.*

Theorem 1 appears not to be true for groups of Lie type defined over fields of even characteristic. In particular, a key component of its proof is Theorem 8, which guarantees the abundance of elements of even order. But the corresponding result does not hold in even characteristic: now most elements are regular semisimple and have odd order, and the proportion of elements of even order is $O(1/q)$. Hence the complexity of the reduction algorithm in these cases is at least linear in q, and so is not polynomial in the size of the input.

Our reduction algorithm, `Reduction`, can readily be embedded into a constructive membership algorithm, `SLPViaCentralisers`, which we present in Section 2. A critical decision is how to solve each instance of the constructive membership problem for an involution centraliser. These can be solved either by a recursive call to `SLPViaCentralisers` or to an arbitrary constructive membership algorithm. However, our analysis of the cost of `Reduction` applies only to simple groups of Lie type in odd characteristic. The fundamental difficulty in producing

an analysis of `SLPViaCentralisers` is that it appears to require knowledge of the composition factors of a black-box group. These are not known to be computable in polynomial time. The best results in this direction are those of [7].

We can however control to some extent the Lie ranks of the non-abelian composition factors of the three involution centralisers. In particular we prove the following.

Theorem 2 *Let G be a simple group of Lie type and rank r, defined over a field of odd characteristic, having a black-box encoding of length n, and equipped with an order oracle. Let δ be a constant where $2/3 < \delta < 1$. If r is sufficiently large, then, at the cost of $O(n^2)$ random selections, we can choose the three involutions for* `Reduction` *so that the Lie ranks of the non-abelian composition factors of their centralisers are at most δr.*

Our principal motivation was a practical algorithm for constructive membership testing. As we demonstrate, our algorithm works well in practice, and often succeeds in cases where other constructive membership algorithms fail. The significance of Theorem 2 is that it allows us to direct the algorithm to choose involutions with relatively small centralisers, which ensures that (in practice) `SLPViaCentralisers` completes as quickly as possible. If the obstructions to a fully recursive algorithm could be overcome, then Theorem 2 could also be used to bound the depth of that recursion to $O(\log r)$, and the total number of recursive calls to a polynomial in r.

Black-box algorithms for constructive membership of the alternating groups have been developed by Beals *et al.* [9]. In various works, Brooksbank, Kantor, and Seress have also developed black-box algorithms for the classical groups; see, for example, [14] and [21]. These algorithms also compute *constructive isomorphisms* between the input group G and a "standard" (or natural) representation of G. Such an isomorphism is not a natural by-product of our work. Ambrose *et al.* [2] develop another general framework for membership testing in black-box groups.

Constructive membership in a permutation group can be decided by constructing a *base and strong generating set* (BSGS), a concept introduced by Sims [31]. For an analysis of the algorithm, see [18] or [30, p. 64]. For a discussion of practical algorithms to decide constructive membership in a soluble group described by a polycyclic presentation, see [32, Chapter 8].

Of course, the use of involution centralisers to obtain insight into group structure is not a new concept. As is well known, they played a fundamental role in the classification of finite simple groups. They were used extensively in early computations with sporadic groups; see [25] for a survey. Altseimer & Borovik [1] used them as a central component of an algorithm to distinguish between $\mathrm{PSp}_{2r}(q)$ and $\Omega_{2r+1}(q)$. Both Borovik [10] and Parker & Wilson [29] consider them in the general context of black-box groups.

The structure of the paper is as follows. In Section 2 we present a constructive membership algorithm which incorporates our reduction algorithm. In Section 3 we present and analyse an algorithm to construct the centraliser of an involution. In Sections 4 and 5 we prove Theorems 1 and 2. In Section 6 we report on a practical implementation in MAGMA [11] of the constructive membership algorithm for quasisimple linear groups.

2 The constructive membership algorithm

Our Monte Carlo constructive membership algorithm, `SLPViaCentralisers`, solves a slightly more general problem than that stated in the introduction. It takes as input a black-box group G equipped with an order oracle, a subgroup H of G, and $g \in G$. If the algorithm concludes that $g \in H$, then it returns an SLP for g in the generators of H, else it returns `false`. The algorithm is the following.

1. Find $h \in H$ with $|gh| = 2\ell$. Now define $z = (gh)^\ell$.
2. Find an involution $x \in H$ with $|xz| = 2m$. Now define $y = (xz)^m$.
3. Construct $X = C_H(x)$.
4. Solve the constructive membership problem for y in X.
5. Construct $Y = C_H(y)$.
6. Solve the constructive membership problem for z in Y.
7. Construct $Z = C_H(z)$.
8. Solve the constructive membership problem for gh in Z.
9. Compute and return an SLP for g.

In practice, we may wish to select h and x carefully, so that the involutions have the property identified in Theorem 2; we consider this in Section 5.

We now more precisely specify `Reduction`: it constructs the involutions x, y, z and their centralisers in H.

We make the following observations.

(a) Each instance of the constructive membership problem for an involution centraliser could be solved by a recursive call to `SLPViaCentralisers` or to a different algorithm. If any one of the constructive membership tests reports `false`, then `SLPViaCentralisers` terminates, returning `false`. If `SLPViaCentralisers` is called recursively, then it must also handle *base cases*: those groups where `Reduction` is not effective.

(b) Observe that $\langle x, z \rangle$ is D_{2m} having central involution $y = (xz)^m$. Hence y is in the centraliser of x and z is in the centraliser of y.

(c) It is easy to deduce that the method is constructive. After Step 1, we know an SLP w_h for h in the generators of H. After Step 2 we similarly know an SLP w_x for x. In Step 3 we record SLPs for the generators of X, and so the call in Step 4 will return an SLP for y. Similarly, in Step 5, we record SLPs for the generators of Y and so in Step 6 obtain an SLP w_z for z. Finally in Step 7 we record SLPs for the generators of Z; so in Step 8 we find an SLP w_{gh} for gh and hence an SLP $w_g = w_{gh} w_h^{-1}$ for g.

(d) If g is an involution, then we can choose h to be 1_H so that $z = g$ and hence we avoid Steps 1, 7 and 8. Both y and z are involutions; if `SLPViaCentralisers` is called recursively, then this remark applies to the subproblems solved at Steps 4 and 6.

The costs of both `Reduction` and `SLPViaCentralisers` depend on three central tasks:

- choose a suitable involution;
- given an involution, construct its centraliser;
- solve the constructive membership problem in this centraliser.

We consider these in detail in the remainder of the paper.

3 Constructing an involution centraliser

The centraliser of an involution in a black-box group having an order oracle can be constructed using an algorithm of Bray [12].

Theorem 3 *[12] If x is an involution in a group H, and w is an arbitrary element of H, then $[x, w]$ either has odd order $2k + 1$, in which case $w[x, w]^k$ commutes with x, or has even order $2k$, in which case both $[x, w]^k$ and $[x, w^{-1}]^k$ commute with x.*

Proof. In the first case $xw[x, w]^k = wx[x, w]^{k+1} = wx[x, w]^{-k} = w[x, w]^k x$ since x is an involution; in the second case $x[x, w^{\pm 1}]^k = x[x, w^{\pm 1}]^{-k} = [x, w^{\pm 1}]^{-k} x$. □

This theorem is used to convert a supply of independent nearly uniformly distributed random elements of H into a supply of elements of $C_H(x)$. While these are not, in general, nearly uniformly distributed, we have the following result (due to Richard Parker).

Theorem 4 *[12] With the above notation, if w is uniformly distributed among the elements of the group for which $[x, w]$ has odd order, then $w[x, w]^k$ is uniformly distributed among the elements of the centraliser of x.*

Proof. If $w' = yw$, where $y \in C_H(x)$, then $[x, w'] = [x, w]$ so that $w'[x, w']^k = yw[x, w]^k$; so each element of $C_H(x)$ occurs exactly once as w runs through any coset of $C_H(x)$ in H. □

Thus if the odd order case occurs sufficiently often (with probability at least a positive rational function of the input size), then we can construct nearly uniformly distributed random elements of the involution centraliser in Monte Carlo polynomial time. Of course, in practice, we can also use the output of the even-order case to obtain a generating set for the centraliser more rapidly.

We now restrict our attention to groups of Lie type over fields of odd characteristic. Here, the structure of the involution centralisers is well-known; see, for example, [19, Table 4.5.1].

Parker & Wilson [29] prove the following for classical groups.

Theorem 5 *There is an absolute constant $c > 0$ such that if H is a finite simple classical group of Lie rank r defined over a field of odd characteristic, and x is an involution in H, then $[x, h]$ has odd order for at least a proportion c/r of the elements $h \in H$.*

They also prove the following result for the exceptional groups.

Theorem 6 *There is an absolute constant $c > 0$ such that if H is a finite simple exceptional group, defined over a field of odd characteristic, and x is an involution in H, then $[x, h]$ has odd order for at least a proportion c of the elements $h \in H$.*

We now analyse the cost of constructing an involution centraliser C by generating elements of C using Theorem 3.

Theorem 7 *Let H be a simple group of Lie type defined over a field of odd characteristic, having a black-box encoding of length n and equipped with an order oracle. The centraliser in H of an involution can be computed in time $O(\sqrt{n}(\xi + \rho)\log(1/\varepsilon) + \mu n)$ with probability of success at least $1 - \varepsilon$, for positive ε.*

Proof. By Theorems 5 and 6, we need $O(\sqrt{n})$ random elements to find a commutator of odd order. The probability that two random elements of a cyclic group G generate G is

$$\prod(1 - \frac{1}{p^2}) > \frac{6}{\pi^2},$$

where the product is over all primes p dividing the order of G. The structure of the involution centralisers [19, Table 4.5.1] and the work of Liebeck & Shalev [23, Theorem] now imply that a constant number of elements generates the centraliser of an involution with arbitrarily high probability. These generators are obtained as powers of elements, each in time $O(n)$ group operations, using the standard doubling algorithm. □

4 Finding the involutions

Let G be a simple group of Lie type in odd characteristic and Lie rank r. Our analysis of `Reduction` assumes that G and its subgroup H coincide. Recall that `Reduction` constructs three involutions in G by powering up elements of even order. We need to estimate the size of the random samples required to obtain these elements. Observe that (since $G = H$) the involutions z and x are powers of random elements of even order, but $y = (xz)^m$ is obtained as a power of their product and so y is not a random element.

Parker & Wilson [29] prove the following.

Theorem 8 *There exists a constant $c > 0$ such that for every simple group G of Lie type in odd characteristic, of Lie rank r, and every conjugacy class C of involutions of G, the proportion of elements of G having a power in C is at least c/r^3.*

Indeed, they show that for the symplectic and orthogonal groups, this proportion is at least c/r^2.

Theorem 9 *Let G be a simple group of Lie type defined over a field of odd characteristic, having at least two conjugacy classes of involutions, and a black-box encoding of length n. In time $O(n^{3/2}(\xi+\rho)\log(1/\varepsilon)+n\mu)$ we can construct the three involutions x, y, z, with probability of success at least $1-\varepsilon$, for positive ε.*

Proof. Theorem 5.2 of [20] implies that at least $1/4$ of the elements of G have even order. Hence we obtain z with probability at least $1-\varepsilon$ by selecting at most $O(\log(1/\varepsilon))$ elements h. Now we need to obtain an involution x such that xz has even order. A sufficient condition for this is that x and z are in different conjugacy classes. Since G has at least two classes of involutions, and by Theorem 8 the proportion of elements of G which power into any given class of involutions is at least $c/n^{3/2}$, it follows that, with probability at least $c/n^{3/2}$, the involutions x and z are in different conjugacy classes. Thus we need at most $O(n^{3/2}\log(1/\varepsilon))$ random elements before we find one where xz has even order. Powering up to construct the involution takes time at most $O(\mu n)$. □

We now prove a similar result for groups having a unique class of involutions.

Theorem 10 *Let G be a simple group of Lie type defined over a field of odd characteristic, having a unique class of involutions, and a black-box encoding of length n. In time $O((\xi+\rho)\log(1/\varepsilon)+n\mu)$ we can construct the three involutions x, y, z, with probability of success at least $1-\varepsilon$, for positive ε.*

Proof. We deduce from [19, Table 4.5.1] that the relevant groups are $\mathrm{PSL}_2(q)$, $\mathrm{PSL}_3(q)$, $\mathrm{PSU}_3(q)$, $\mathrm{PSU}_4(q)$ for $q \equiv 3 \bmod 8$, $\mathrm{PSL}_4(q)$ for $q \equiv 5 \bmod 8$, $G_2(q)$, ${}^2G_2(q)$, and ${}^3D_4(q)$.

The Baer–Suzuki theorem [3, 39.6] implies that there exist two conjugates of an involution whose product has even order. We show that the proportion of pairs of involutions whose product has even order is at least a positive constant.

We illustrate the method of proof with the example of $G = {}^2G_2(q)$. Observe that $|G| = (q^3+1)(q-1)q^3$ and the involution centraliser has order $q(q^2-1)$. Hence the number of involutions is $a = q^2(q^2-q+1)$. We want to count the number of pairs of involutions whose product has even order greater than 2 and dividing $q+1$. Since the dihedral group they generate lies in the centraliser $2 \times \mathrm{PSL}_2(q)$ of another involution, its normaliser is contained in $H = 2 \times D_{q+1}$. The number of conjugates of H is $b = q^3(q-1)(q^2-q+1)/2$. The number of such pairs in H is $c = (q+1)(q-3)/4$. Moreover, no pair is in two distinct conjugates of H. Hence the desired proportion is

$$bc/a^2 = (q-3)(q^2-1)/8q(q^2-q+1) > (1-3/q)/8 \geq 1/9$$

since $q \geq 27$.

The other cases are similar. Since we need only an asymptotic result, we may assume that q is large and consider only the leading terms of the various polynomials in q which arise.

In the case of $\mathrm{PSL}_2(q)$ we look in D_{q-1} or D_{q+1} according as $q \equiv 1$ or $3 \bmod 4$. The number of such subgroups is of the order of $q^2/2$, and in each subgroup the number of pairs of involutions generating a suitable subgroup is at least of the order of $q^2/8$. But the total number of involutions is of the order of $q^2/2$, so the proportion of pairs whose product has even order is at least of the order of $(q^2/2)(q^2/8)/(q^2/2)^2 = 1/4$.

In $\mathrm{PSL}_3(q)$ our two involutions negate a common 1-space, and we can work in $\mathrm{GL}_2(q)$ instead. Similarly in $\mathrm{PSU}_3(q)$, we may work in $\mathrm{GU}_2(q)$. In $G_2(q)$ the involution centraliser is $2{\cdot}(\mathrm{PSL}_2(q) \times \mathrm{PSL}_2(q)).2$, and there is a dihedral group $D_{2(q^2-1)}$ which has index 2 in its normaliser. Therefore there are approximately q^8 involutions and $q^{12}/4$ such dihedral groups, each containing at least of the order of $q^4/2$ pairs of involutions whose product is regular semisimple of even order. Thus the desired proportion is at least of the order of $1/8$ in this case.

Both $\mathrm{PSL}_4(q)$ and $\mathrm{PSU}_4(q)$ are most easily treated as orthogonal groups, so we work in $\mathrm{SO}^+_6(q)$ or $\mathrm{SO}^-_6(q)$ and use the embedding of $\mathrm{GL}_2(q)$ into $\mathrm{SO}^+_4(q)$, as in Theorem 14. Indeed, Theorem 14 proves a more precise version of the result in these cases. The case of ${}^3D_4(q)$ is similar to $G_2(q)$: we take the involution centraliser $2{\cdot}(\mathrm{PSL}_2(q) \times \mathrm{PSL}_2(q^3)).2$ and the dihedral group $D_{2(q-1)(q^3+1)}$ inside it. □

Theorems 7, 9 and 10 now imply Theorem 1.

5 Prescribing the conjugacy classes of the involutions

Let G be a simple group of Lie type in odd characteristic and Lie rank r, having a black-box encoding of length n. In `Reduction` we construct three involutions in G by powering up elements of even order. If we simply choose the involutions as powers of random elements of even order, then each centraliser may have a composition factor of Lie rank $r-1$. We now prove Theorem 2: we can *choose* our involutions so that the non-abelian composition factors of their centralisers have Lie rank at most a proper fraction of r.

Theorem 8 implies that at a cost of at most $O(r^3)$ random selections we can choose precisely the conjugacy class of both x and z. In the proof of Theorem 2, we discuss how to identify the class. We now consider the choice of $y = (xz)^m$ in more detail. In particular, we consider the case where y lies in a conjugacy class of involutions whose eigenspaces on the natural module have a prescribed dimension.

We first prove a preliminary lemma.

Lemma 11 *Let p be an odd prime, $k \geq 2$, and let C be the (unique) subgroup of order $p^k + 1$ in the multiplicative group of the field F of order p^{2k}. Then the proportion of elements of C which lie in a proper subfield of F is at most $1/(2p-1)$.*

Proof. We need only consider subfields of order $p^{2k/\ell}$ where ℓ is prime, so that the number of elements of C lying in the subfield is $h := \gcd(p^k+1, p^{2k/\ell}-1)$. If $\ell = 2$, clearly $h = 2$. If ℓ is odd, then $p^{2k/\ell} - 1 = (p^{k/\ell}+1)(p^{k/\ell}-1)$, and

$$\begin{aligned} p^k + 1 &= (p^{k/\ell}+1)(p^{k-k/\ell} - p^{k-2k/\ell} + \cdots - p^{k/\ell} + 1) \\ &= (p^{k/\ell}+1)((p^{k/\ell}-1)(p^{k-2k/\ell} + \cdots + p^{k/\ell}) + 1) \end{aligned}$$

so $h = p^{k/\ell} + 1$. Therefore (counting ± 1 only once) the number of elements in C which lie in proper subfields is at most

$$2 + \sum_{\ell | k} (p^{k/\ell} - 1)$$

where ℓ is an odd prime. If there are at least two odd primes dividing k, then

$$\sum p^{k/\ell} < \sum_{m=0}^{k-2} p^m = (p^{k-1}-1)/(p-1) \leq \frac{1}{2}(p^{k-1}-1),$$

and so the stated proportion is at most

$$\frac{\frac{1}{2}(p^{k-1}-1)}{p^k+1} < \frac{1}{2p}.$$

If there is a unique odd prime dividing k, then the stated proportion is at most

$$\frac{p^{k/3}+1}{p^k+1} < \frac{1}{2p}.$$

If k is a power of 2, then the stated proportion is

$$\frac{2}{p^k+1} \leq \frac{1}{2p-1}.$$

□

In fact the same argument shows more.

Lemma 12 *If q is an odd prime power and $k \geq 2$, then the proportion of elements in C_{q^k+1} that are regular semisimple in $\mathrm{GL}_{2k}(q)$ is at least $1 - 1/(2q-1) \geq 4/5$.*

We also need the following order estimates for classical groups extracted from [29].

Lemma 13 *If $q \geq 3$ is a prime power, then*

(i) $\frac{1}{2}q^{d^2} \leq |\mathrm{GL}_d(q)| \leq q^{d^2}$;

(ii) $\frac{1}{2}q^{d^2} \leq |\mathrm{GU}_d(q)| \leq 2q^{d^2}$;

(iii) $\frac{1}{2}q^{d(d+1)/2} \leq |\mathrm{Sp}_d(q)| \leq q^{d(d+1)/2}$;

(iv) $\frac{1}{2}q^{d(d-1)/2} \leq |\mathrm{SO}_d(q)| \leq 2q^{d(d-1)/2}$.

An involution $g \in \mathrm{GL}_d(q)$ has *type* -1^a1^b if its -1-eigenspace in the natural module has dimension a and $a+b=d$. We now prove the main result of this section.

Theorem 14 *Let G be a quasisimple classical group in its natural representation of degree d, defined over a field of odd characteristic. Suppose that $d > 4k$ for positive k. Let x be an involution in G of type $-1^{2k}1^{d-2k}$. There exists a constant $c > 0$ such that, with probability at least c/k, the product of two random conjugates of x powers up to an involution of type $-1^{4k}1^{d-4k}$.*

Proof. We prove this result for $\mathrm{SL}_d(q)$ by looking inside the normaliser of a Singer cycle (namely, a cyclic subgroup of order $q^{4k}-1$) in $\mathrm{GL}_{4k}(q)$. In $\mathrm{SU}_d(q)$ we look at the normaliser of a Singer cycle in $\mathrm{GL}_{2k}(q^2)$, and in the symplectic groups we look at a Singer cycle in $\mathrm{GL}_{2k}(q)$. The orthogonal groups, as usual, are a little more complicated.

The normaliser of a Singer cycle in $\mathrm{GL}_{2k}(q)$ contains a subgroup $C_{q^{2k}-1}{:}C_2$, whose centre has order q^k-1. There are involutions of type -1^k1^k inverting the subgroup of order q^k+1. (All this can be seen already in the subgroup $\mathrm{GL}_2(q^k)$.) By Lemma 12, there are at least $(\frac{4}{5})^2q^{2k}$ pairs of involutions whose product is a regular semisimple element in this C_{q^k+1}, and in at least half of these cases the product has even order. For brevity call such pairs of involutions *good*. There are at least $\frac{1}{2}(\frac{4}{5})^2q^{2k} > \frac{1}{4}q^{2k}$ good pairs of involutions in the normaliser of a particular cyclic group of order q^k+1.

Now we estimate the numbers of these tori, and the numbers of pairs of involutions in the given conjugacy class, in order to estimate the proportion of these pairs which are good: namely those whose product powers up into the desired conjugacy class of involutions.

First look at the case $\mathrm{SL}_d(q)$. We embed $\mathrm{GL}_{4k}(q)$ naturally in $\mathrm{SL}_d(q)$, for $d > 4k$, and observe that the normaliser of $C_{q^{2k}+1}$ in $\mathrm{SL}_d(q)$ is $\mathrm{SL}_{d-4k}(q).C_{q^{4k}-1}.C_{4k}$, so the number of such tori is

$$\frac{|\mathrm{SL}_d(q)|}{|\mathrm{SL}_{d-4k}(q)|(q^{4k}-1).4k} \geq \frac{1}{8k}q^{4k(2d-4k-1)},$$

and the number of good pairs of involutions is at least $\frac{1}{32k}q^{8k(d-2k)}$. Similarly an involution of type $-1^{2k}1^{d-2k}$ has centraliser $\mathrm{SL}_{2k}(q).\mathrm{GL}_{d-2k}(q)$ so the total number of involutions from this conjugacy class is

$$\frac{|\mathrm{SL}_d(q)|}{|\mathrm{SL}_{2k}(q)||\mathrm{GL}_{d-2k}(q)|} \leq 4q^{4k(d-2k)}.$$

Hence the proportion of good pairs of involutions is at least $1/(2^9k)$.

Next consider the unitary groups. In this case we embed $\mathrm{GL}_{2k}(q^2)$ into $\mathrm{GU}_{4k}(q)$ and thence into $\mathrm{SU}_d(q)$, for $d > 4k$. The centraliser of an involution of type $-1^{2k}1^{d-2k}$ is $\mathrm{SU}_{2k}(q).\mathrm{GU}_{d-2k}(q)$. Hence there are at most

$$\frac{|\mathrm{SU}_d(q)|}{|\mathrm{SU}_{2k}(q)||\mathrm{GU}_{d-2k}(q)|} \leq 8q^{4k(d-2k)}$$

such involutions. The order of the normaliser of $C_{q^{2k}+1}$ is $(q^{4k}-1).|\mathrm{SU}_{d-4k}(q)|.4k$ so there are at least

$$\frac{|\mathrm{SU}_d(q)|}{|\mathrm{SU}_{d-4k}(q)|.(q^{4k}-1).4k} \geq \frac{1}{16k} q^{4k(2d-4k-1)}$$

such groups, each with at least $q^{4k}/4$ good pairs of involutions. Thus the proportion of good pairs of involutions from this class is at least $1/(2^{12}k)$.

Next consider the symplectic groups. We embed $\mathrm{GL}_{2k}(q)$ into $\mathrm{Sp}_{4k}(q)$ and thence into $\mathrm{Sp}_d(q)$, for $d > 4k$. The centraliser of an involution of type $-1^{2k}1^{d-2k}$ is $\mathrm{Sp}_{2k}(q) \times \mathrm{Sp}_{d-2k}(q)$ so the number of such involutions is at most

$$\frac{|\mathrm{Sp}_d(q)|}{|\mathrm{Sp}_{2k}(q)||\mathrm{Sp}_{d-2k}(q)|} \leq 4q^{2k(d-2k)}.$$

The normaliser of C_{q^k+1} is $\mathrm{GU}_2(q^k).C_{2k} \times \mathrm{Sp}_{d-4k}(q)$, so the number of such cyclic groups is

$$\frac{|\mathrm{Sp}_d(q)|}{2k|\mathrm{GU}_2(q^k)||\mathrm{Sp}_{d-4k}(q)|} \geq \frac{1}{8k} q^{2k(2d-4k-1)}.$$

Thus the proportion of good pairs of involutions from this class is at least $1/(2^9k)$.

Finally consider the orthogonal groups. The number of conjugates of an involution of type $-1^{2k}1^{d-2k}$ is at most

$$\frac{|\mathrm{O}_d(q)|}{|\mathrm{O}_{2k}(q)||\mathrm{O}_{d-2k}(q)|} \leq 4q^{2k(d-2k)}.$$

Now consider $\mathrm{O}_4(q^k) < \mathrm{O}_{4k}(q)$. Independent of the sign of this orthogonal group, it contains a dihedral group $D_{2(q^{2k}-1)}$. The normaliser of the corresponding cyclic group of order $q^{2k}-1$ in $\mathrm{O}_d(q)$ is $D_{2(q^{2k}-1)}.C_k \times \mathrm{O}_{d-4k}(q)$. Hence the number of conjugates of this dihedral group is at least

$$\frac{|\mathrm{O}_d(q)|}{2k.(q^{2k}-1)|\mathrm{O}_{d-4k}(q)|} \geq \frac{1}{8k} q^{4k(d-2k-1)}.$$

To complete the argument, we must estimate the number of elements in the cyclic group of order $C_{q^{2k}-1}$ which are regular semisimple in $\mathrm{O}_{4k}(q)$. Again, this is bounded below by a positive constant times q^{2k} and so the powers of q cancel as required. □

Corollary 15 *The same result holds for simple classical groups.*

Proof. Working modulo scalars has no effect on the above argument. □

Theorem 8 implies that in Step 1 of the algorithm we need at most $O(r^3)$ trials to find an involution in a particular conjugacy class. We now show, by allowing a range of dimensions for the eigenspace, that we can reduce the cost of this step to $O(r^2)$.

Lemma 16 *Let $0 < \kappa < \lambda < 1$ and let G be a simple classical group with natural module of dimension $d > 2/(\lambda - \kappa)$ defined over a field of odd characteristic. Then the proportion of elements of G which power to an involution whose -1-eigenspace on the natural module has dimension in the range κd to λd is at least c/d^2 for some constant c depending on κ and λ.*

Proof. If r is the Lie rank of G, then d is $r+1$ or $2r$ or $2r+1$. Therefore by Theorem 8, for each eigenspace dimension the proportion is at least a constant times d^{-3}. Since $(\lambda - \kappa)d > 2$, there exists at least one even integer in the range $(\kappa d, \lambda d)$. Hence there is at least one conjugacy class of involutions with -1-eigenspace dimension in this range, and indeed the number of possible dimensions is at least a constant times d. □

Proof of Theorem 2. We may assume $r > 8$, so exceptional groups of Lie type do not arise. Let d denote the dimension of the natural module for G. We choose δ' in the range $(2/3, \delta)$ and now define $\kappa = 1 - \delta'$ and $\lambda = \delta'/2$. Since $\delta' > 2/3$ we have $\kappa < \lambda$. If $d > 2/(\lambda - \kappa) = 4/(3\delta' - 2)$, then Lemma 16 implies that in $O(n)$ attempts we can choose the involution x and (its conjugate) z to have -1-eigenspace of dimension $2k$ where $(1 - \delta')d < 2k < \delta' d/2$. Theorem 14 implies that among a sample of $O(\sqrt{n})$ random conjugates of x, we find two whose product powers up to an involution y of type $-1^{4k}1^{d-4k}$.

It is easy to deduce that

$$(1 - \delta')d < d - 4k < \delta' d$$

$$(1 - \delta')d < 2k < \delta' d.$$

Hence both eigenspaces for each of the three involutions have dimension less than $\delta' d$.

Recall that the non-abelian composition factors of a centraliser of such an involution in a classical group are of the same classical type (linear, unitary, symplectic or orthogonal) in smaller Lie rank. Further r is $d-1$ (linear, unitary), or $(d-1)/2$ or $d/2$ (orthogonal, symplectic). Hence, for sufficiently large r, the Lie ranks of the non-abelian composition factors are at most δr.

For each involution in the sample, we must also identify its conjugacy class. We construct its centraliser C using Theorem 7; using the algorithm of [6], we construct the last term of the derived series of C, which is a product of at most two semisimple groups; we construct its composition factors using the algorithm of [7, Claim 5.3]; we identify the defining characteristic using the algorithm of [24]; finally, we name the composition factors using the algorithm of [8]. All of these steps can be performed using a sample of at most $O(n)$ elements. □

One might hope that Theorem 2 would allow us to bound by $\log r$ the depth of the recursion tree arising in a recursive application of `SLPViaCentralisers` to a simple group of Lie rank r. However, Theorem 2 does *not* apply to the centralisers, since they need not be simple. Such a result appears to require the ability to construct the composition factors of the centralisers: then both the number of recursive calls and the number of base cases could be bounded by a polynomial in r.

6 A practical realisation for matrix groups

We now consider how the constructive membership algorithm can be realised in practice for quasisimple linear groups.

We implemented a version of `SLPViaCentralisers` which solves the constructive membership problem in each centraliser, by constructing its composition factors using the *composition tree* algorithm (see [27]). If the Lie rank of a composition factor is too large for a direct solution of the membership problem, then we recursively apply `SLPViaCentralisers` to this factor.

Recall that a necessary component of `Reduction` is an order oracle. Celler & Leedham-Green [16] present an algorithm to determine the order of $g \in \mathrm{GL}(d, q)$. While it requires the factorisation of certain large integers, a variation can, as discussed in [27], in polynomial time determine a multiple of the order. From this multiple, we can determine if the element has *even order* and if so, construct an involution. Knowledge of a multiple of the order also suffices for Theorem 3. In practice, we use projective orders so that we can work in the simple group.

If the input group is classical in its natural representation, then we can determine the type of an involution directly.

6.1 Applications to sporadic groups

The original application of `Reduction` was as one step in the classification of conjugacy classes of subgroups of $E_7(5)$ isomorphic to the Rudvalis sporadic simple group (see [22]).

As we earlier observed, its performance depends critically on the proportions of elements gh (respectively xz) which power up to involutions of each class. For the sporadic groups, these proportions are constants which can be calculated from the character tables. In some cases, these proportions are zero: for certain choices of involution classes $\mathcal{C}_1$, $\mathcal{C}_2$, $\mathcal{C}_3$, there are no elements $x \in \mathcal{C}_1$, $z \in \mathcal{C}_2$, with xz powering to an element in $\mathcal{C}_3$. Hence we do not have a completely free choice of these three classes. In practice, we choose x, y and z all to be elements of the largest class of involutions, in which case it turns out that the probabilities are all positive. Indeed, as can be deduced from the character tables, the probability that gh powers to an element in this class is at least $5/64 = 0.078125$, while the probability that xz powers to an element in this class is at least $6181967/148341375 \approx 0.041674$.

For each sporadic group, we can calculate explicitly the proportion of $[x, g]$ which have odd order. For every class of involutions x this proportion is always greater than 17%, and therefore Bray's algorithm to construct an involution centraliser C completes rapidly. We now construct its composition tree and solve the membership problem for C directly.

Table 1 records some data supporting our claim that the algorithm works well for the sporadic groups.

6.2 Implementation and performance

`SLPViaCentralisers` is implemented in MAGMA. One motivation for its development is to solve the constructive membership problem for composition factors of matrix groups. The input to our implementation is an irreducible representation of a group of Lie type in odd defining characteristic, or a sporadic group.

`Reduction` constructs (at most) three involution centralisers. A composition tree is constructed for each centraliser, whose leaves are its composition factors. For each factor, we may generate further calls to `SLPViaCentralisers` until we construct a base case. Alternatively, if the factor is sufficiently small, we invoke the Schreier–Sims algorithm [31] (or its variations) to solve the problem.

Our implementation uses the following components:

- the product replacement algorithm [15] to generate random elements;
- the algorithm of Celler & Leedham-Green [16] to determine the order of an element;
- the algorithm of Liebeck & O'Brien [24] to determine the defining characteristic of a group of Lie type;
- the algorithms of Babai *et al.* [8] to identify a simple group of Lie type in known defining characteristic;
- the algorithm of Niemeyer & Praeger [26] to identify a classical group in its natural representation.
- the algorithm of Conder *et al.* [17] to solve the constructive membership problem for $\mathrm{SL}_2(q) \cong \mathrm{SU}_2(q) \cong \mathrm{Sp}_2(q)$; $\Omega_3(q) \cong \mathrm{PSL}_2(q)$; $\Omega_4^-(q) \cong \mathrm{PSL}_2(q^2)$; and $\Omega_4^+(q) \cong \mathrm{SL}_2(q) \circ \mathrm{SL}_2(q)$.

A variation of Theorem 3 allows us to decide constructively if two involutions x and y are conjugate in a group H. We construct random conjugates x_i of x, until we find $x_i y$ with odd order $2k+1$, say. In the dihedral group $D_{4k+2} = \langle x_i, y \rangle$, we can see that $(yx_i)^k$ conjugates x_i to y. If two random conjugates of x have a high enough probability of having a product of odd order, this provides an effective method. Moreover, it is constructive in the sense that it provides $h \in H$ such that $x^h = y$, and hence $C_H(x)^h = C_H(y)$.

We exploit this observation in our implementation. If we repeatedly test for membership in the same group, then we store the chosen involutions and their associated composition trees; as a preliminary step in a new membership test, we decide if the new involutions are conjugate to the known ones; if so, we do not need to construct a new composition tree.

Our constructive membership algorithm is competitive with the standard BSGS machinery for matrix groups of "moderate" dimension. If the matrix group has no subgroup of reasonable index, then our algorithm is currently the only practical approach. For example, the largest proper

Name	d	q	Time
J_4	112	2	8.5
$SL_{20}(5)$	20	5	10.0
$G_2(3^5)$	7	3^5	0.9
Ly	111	5	85.0
Th	248	3	2210
$Sp_{10}(9)$	10	9	3.1
$\Omega_{12}^+(7)$	12	7	2.1

Table 1: Performance of implementation for a sample of groups

subgroup of J_4 has index about 10^8; our algorithm readily succeeds in the 112-dimensional representation over $\mathrm{GF}(2)$.

In Table 1, we report on the application of `SLPViaCentralisers` to some of the larger sporadic groups and to groups of Lie type. The strategy of Theorem 2 to direct the choice of involution works well. For example, in $\mathrm{SL}_{20}(5)$, the three involutions chosen have types $-1^8 1^{12}, -1^{12}1^8$, and $-1^{10}1^{10}$. A further recursion then reduces to $d \leq 6$, and an invocation of a Schreier–Sims algorithm now completes the task. None of these examples completed using the existing machinery in MAGMA V2.12 on a Pentium IV 2.8 GHz processor with 2GB of RAM. The input to the algorithm is an irreducible subgroup of $\mathrm{GL}_d(q)$. In the column entitled "Time", we list the CPU time in seconds (averaged over three runs) needed to solve the constructive membership problem for a random element of the group.

References

[1] Christine Altseimer and Alexandre V. Borovik. Probabilistic recognition of orthogonal and symplectic groups. In *Groups and Computation, III (Columbus, OH, 1999)*, volume 8 of *Ohio State Univ. Math. Res. Inst. Publ.*, pages 1–20. de Gruyter, Berlin, 2001.

[2] Sophie Ambrose, Max Neunhöffer, Cheryl E. Praeger and Csaba Schneider. Generalised sifting in black-box groups, *London Math. Soc. J. Comput. Math.* **8**, 217–250, 2005.

[3] M. Aschbacher. Finite Group Theory, Cambridge Studies in Advanced Mathematics 10, Second Edition, 2000.

[4] 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.

[5] 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.

[6] László Babai, Gene Cooperman, Larry Finkelstein, Eugene Luks, and Ákos Seress. Fast Monte Carlo algorithms for permutation groups. 23rd Symposium on the Theory of Computing (New Orleans, LA, 1991). *J. Comput. System Sci.* **50** (1995), no. 2, 296–308.

[7] László Babai and Robert Beals. A polynomial-time theory of black box groups. I. Groups St. Andrews 1997 in Bath, I, 30–64, London Math. Soc. Lecture Note Ser., 260, Cambridge Univ. Press, Cambridge, 1999.

[8] László Babai, William M. Kantor, Péter P. Pálfy and Ákos Seress. Black box recognition of finite simple groups of Lie type by statistics of element orders, *J. Group Theory* **5** (2002), 383–401.

[9] Robert Beals, Charles R. Leedham-Green, Alice C. Niemeyer, Cheryl E. Praeger, and Ákos Seress. A black-box group algorithm for recognizing finite symmetric and alternating groups. I, Trans. Amer. Math. Soc. **355**, no. 5, 2097–2113, 2003.

[10] A.V. Borovik. Centralisers of involutions in black box groups. In *Computational and statistical group theory (Las Vegas, NV/Hoboken, NJ, 2001)*, 7–20, *Contemp. Math.*, **298**, Amer. Math. Soc., Providence, RI, 2002.

[11] Wieb Bosma, John Cannon, and Catherine Playoust. The MAGMA algebra system I: The user language. *J. Symbolic Comput.*, **24**, 235–265, 1997.

[12] John N. Bray. An improved method for generating the centralizer of an involution. *Arch. Math. (Basel)* **74** (2000), 241–245.

[13] 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.*, pages 95–111, Berlin, 2001. de Gruyter.

[14] P.A. Brooksbank and W.M. Kantor. Fast constructive recognition of black-box orthogonal groups. *J. Algebra* **300** (2006), 256–288.

[15] Frank Celler, Charles R. Leedham-Green, Scott H. Murray, Alice C. Niemeyer and E.A. O'Brien. Generating random elements of a finite group. *Comm. Algebra*, **23** (1995), 4931–4948.

[16] Frank Celler and C.R. Leedham-Green. Calculating the order of an invertible matrix. In *Groups and Computation II*, volume 28 of *Amer. Math. Soc. DIMACS Series*, pages 55–60. (DIMACS, 1995), 1997.

[17] 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.

[18] Merrick Furst, John Hopcroft, and Eugene Luks. Polynomial-time algorithms for permutation groups. 21st Annual Symposium on Foundations of Computer Science (Syracuse, N.Y., 1980), pp. 36–1, IEEE, New York, 1980.

[19] Daniel Gorenstein, Richard Lyons, and Ronald Solomon. The classification of the finite simple groups. Number 3. Part I, American Mathematical Society, Providence, RI, 1998.

[20] I.M. Isaacs, W.M. Kantor and N. Spaltenstein. On the probability that a group element is p-singular. *J. Algebra* **176** (1995), 139–181.

[21] William M. Kantor and Ákos Seress. Black box classical groups. *Mem. Amer. Math. Soc.*, **149**, 2001.

[22] P.B. Kleidman, U. Meierfrankenfeld and A.J.E. Ryba. $Ru < E_7(5)$. *Comm. Algebra* **28**, 3555–3583, 2000.

[23] Martin W. Liebeck and Aner Shalev. The probability of generating a finite simple group. *Geom. Ded.* **56** (1995), 103–113.

[24] Martin W. Liebeck and E.A. O'Brien. Finding the characteristic of a group of Lie type. *J. Lond. Math. Soc.* **75**, 741–754, 2007.

[25] S.A. Linton. The art and science of computing in large groups. *Computational Algebra and Number Theory* (Sydney, 1992), pp. 91–109, 1995. Kluwer Academic Publishers, Dordrecht.

[26] A.C. Niemeyer and C.E. Praeger. A recognition algorithm for classical groups over finite fields. *Proc. London Math. Soc.* **77** (1998), 117–169.

[27] E.A. O'Brien. Towards effective algorithms for linear groups. *Finite Geometries, Groups and Computation*, (Colorado), pp. 163-190. De Gruyter, Berlin, 2006.

[28] Igor Pak. The product replacement algorithm is polynomial. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, 476–485, IEEE Comput. Soc. Press, Los Alamitos, CA, 2000.

[29] Christopher W. Parker and Robert A. Wilson. Recognising simplicity of black-box groups. Preprint, 2007.

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

[31] Charles C. Sims. Computational methods in the study of permutation groups. In *Computational problems in abstract algebra*, pages 169–183, Oxford, 1970. (Oxford, 1967), Pergamon Press.

[32] Charles C. Sims. *Computation with finitely presented groups*. Cambridge University Press, 1994.

P.E. Holmes, Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, United Kingdom.

S.A. Linton, Centre for Interdisciplinary Research in Computational Algebra, University of St Andrews, St Andrews, Fife KY16 9SS, United Kingdom.

E.A. O'Brien, Department of Mathematics, University of Auckland, Auckland, New Zealand.

A.J.E. Ryba, Department of Computer Science, City University of New York, Flushing, NY 11367, USA.

R.A. Wilson, School of Mathematical Sciences, Queen Mary, University of London, London E1 4NS, United Kingdom.

Last updated February 7, 2008