Recognizing simplicity of black-box groups and the frequency of p-singular elements in affine groups

László Babai and Aner Shalev

Abstract. We consider the asymptotic complexity of manipulating matrix groups over finite fields. The question is, given a matrix group G by a list of generators, what can we say in polynomial time about the structure of G?

While considerable progress has been made recently in identifying the nonabelian composition factors of a matrix group, the fundamental question of recognizing the simplicity of a non-abelian matrix group remains open.

For the purposes of polynomial time computation, we reduce this problem to the following "affine case:"

> Let A be an elementary abelian p-group which is a non-central minimal normal subgroup of a finite group G. Assume further that the quotient G/A is a finite simple group S of Lie type of characteristic p.
>
> The algorithmic goal is to distinguish the simple group S from the non-simple group G. Both S and G are given as "black-box groups of characteristic p."

The reduction to this basic problem involves a large number of recent techniques and results which we review along the way.

We address the affine case for the groups $S = \mathrm{PSL}(2, q)$ where $q = p^f$, p prime. We show that $\mathrm{PSL}(2, p)$ can be recognized in Monte Carlo polynomial time among all black-box groups of known characteristic. The situation for $f \geq 2$ seems much harder; we exhibit challenging open cases for every $f \geq 2$ when q is large.

Along with $S = \mathrm{PSL}(2, p)$, the positive result holds for all simple groups S of Lie type of characteristic p with the property that in every nontrivial S-module in characteristic p, every element of S has a fixed point. We call these groups "unisingular." Very recently, Guralnick, Saxl, and Tiep classified these groups. They found that a large number of classes of simple groups are unisingular, including certain classes of linear and unitary groups, all symplectic groups over an odd prime field, all orthogonal groups of odd degree over an odd prime field, many orthogonal groups of even degree, and a number of classes of exceptional groups.

1. Introduction

Let G be a finite group given by a list of generators. Our goal is to design efficient algorithms to obtain structural information about G. In this paper we take "efficient" to mean "Monte Carlo polynomial time." A Monte Carlo algorithm uses randomization and produces a correct output with probability $\geq 1-\epsilon$ where $\epsilon > 0$ is specified by the user. The running time of a polynomial time Monte Carlo algorithm is $O(|\log \epsilon| \cdot n^c)$ where c is a constant and n is the bit-length of the input. Throughout the paper we shall state the results under the tacit assumption that ϵ is a constant; multiplying by $|\log \epsilon|$ will adjust the timing estimates for variable ϵ.

1.1. Black-box groups. Typically G will be a quotient of a matrix group over a finite field, but the algorithms considered here work in the more general context of "black-box groups." The concept of black-box groups, introduced in [BaS] (1984), is critical to our work.

In a *black-box group*, a "black box" performs the group operations on codewords representing the group elements (multiplication, inverting, deciding whether a given string represents the identity.) Each codeword is a $(0,1)$-string of the same length n, called the *encoding length*. Not every string encodes a group element, and several strings may encode the same element. A black-box group is "given" if in addition to the oracles for group operations, a list of generators is given. If the input is a group with k generators and the encoding length is n then the length of the input is kn and correspondingly a polynomial-time algorithm is allowed to run for $(kn)^{O(1)}$ steps.

The rudiments of polynomial-time computation in black-box groups have been worked out in [BaCFLS]. The results include Monte Carlo polynomial-time algorithms for normal closure, the commutator subgroup, deciding solvability, nilpotence and deciding whether or not G is a p-group.

A detailed polynomial-time theory of black-box groups was given in [BaB], our standard reference. For the precise statement of the main results, the concept of "black-box groups of characteristic p" [BaB] is required; we give this definition in Section 4.1. Let it suffice to mention here that this concept includes all matrix groups over a field of characteristic p as well as quotients of such groups by recognizable normal subgroups. Membership testing in the normal subgroup is needed in order to decide what strings represent the identity in the quotient.

1.2. Random sampling. Most of the algorithms for matrix groups and black-box groups assume that we can choose uniformly distributed random elements from the group G. This assumption is not justified, but "nearly uniform"

random sampling can be achieved (every element has $(1 \pm \epsilon)/|G|$ probability to be chosen). For the purposes of the polynomial-time theory, it is shown in [Ba1] that nearly uniform random sampling can be performed in Monte Carlo polynomial time, even with exponentially small ϵ. For practical purposes, variants of the "product replacement" heuristic seem to perform very well [CLMNO]. For concepts and caveats regarding randomized algorithms for groups we refer to [Ba2].

In this paper, when using the term "pick an element at random from the (nonempty, finite) set Ω," we mean (nearly) uniform selection from Ω. For the analysis we think of the sampling as uniform, but "nearly uniform" would work just as well.

1.3. Center, outer automorphisms. A recent result from [BaPS] states that for every prime r and every nonabelian finite simple group S, a significant fraction of the elements of S is r-regular (of order not divisible by r) (see Theorem 4.9). This result was used in [BaB] to split a direct product of simple groups into its factors, a key step toward finding the nonabelian composition factors. In this paper, we give two additional applications: finding the center in a central extension of a simple group (Section 4.7) and analysing an algorithm by C. Leedham-Green to test membership in a simple normal subgroup (Section 4.4). The most important special case of the latter situation is when $S \leq G \leq \mathrm{Aut}(S)$ for some simple group S. In this case S can be constructed as the stable derivative (last member of the derived chain) of G (since, by Schreier's hypothesis, G/S is solvable); we are now also able to test elements of G for membership in S. In fact we show that any simple normal subgroup with a known set of generators is recognizable (Theorem 4.6).

1.4. Nonabelian composition factors. The main result of [BaB] is a polynomial-time algorithm to find a black-box representation of characteristic p for all non-abelian composition factors of a black-box group of characteristic p. Combining this with recent results on the statistical recognition of finite simple groups ([AltB], [KS2], [BaKPS]), we are now able to name the nonabelian composition factors:

Theorem 1.1 ([BaB, AltB, KS2, BaKPS]). *Given a black-box group G of known characteristic, the standard names of all nonabelian composition factors of G can be computed in Monte Carlo polynomial time.*

In particular, this result applies to all matrix groups over finite fields. (See Section 4.3.)

2. Recognizing simplicity: the main results

In spite of all this progress, the basic question, to *recognize the simplicity of a matrix group (or a black-box group of given characteristic)* remains open. This is the main question addressed in this paper.

Since by Theorem 1.1, every simple group can be recognized *among all simple groups*, the recognition of *simplicity* is equivalent recognizing each simple group (among all groups). We give a more precise formulation to this statement.

Theorem 2.1. *The following four computational tasks are equivalent in Monte Carlo polynomial time. In each case, a black box group G of known characteristic p is given.*

(a) Decide whether or not G is simple ("recognizing simplicity").

(b) Given the standard name of a simple group S, decide whether or not $G \cong S$ ("recognizing the simple group S").

(c) Same as (b) except that S is restricted to being a simple group of Lie type of charcateristic p.

(d) ("Affine case"). Assume that S is a simple group Lie type of characteristic p and G is known to have a normal subgroup A such that $G/A \cong S$ and either $A = 1$ or A is a non-central minimal normal subgroup of G and A an elementary abelian p-group. Decide whether or not $A = 1$.

The equivalence of (a) and (b) follows from Theorem 1.1. It is also clear that (d) reduces to (c) and (c) reduces to (b), each being a subproblem of the preceding one.

The principal result of this paper is a reduction of (b) (general case) to (d) ("affine case"). Full details of this reduction are given in Section 4.8.

The purpose of this reduction is twofold. First, it helps in solving the problem for certain classes of simple groups.

In particular, we show, using a monumental recent work by Kantor and Seress [KS1], how to recognize a simple group of Lie type (i. e., to distinguish such a group from *all* other groups, not only from other simple groups), assuming the characteristic is known and the field of definition is "tiny" (Section 4.6).

We also exhibit the first classes of Lie-type simple groups over *large fields* which can be recognized. We state the result after a definition.

Definition 2.2. Let G be a finite group. We say that G is *unisingular in characteristic p* if every nontrivial G-module M of characteristic p has the property that every element of G has a nonzero fixed point in its action on M.

In other words, G is unisingular in characteristic p if 1 is an eigenvalue of the action of every element of G on any nontrivial module of characteristic p.

We shall be particularly interested in the case when G is simple of Lie type of characteristic p. In such a case we shall simply call G "unisingular," tacitly referring to the same characteristic as the characteristic of the field of definition of G.

Theorem 2.3. *If S is a unisingular simple group of Lie type of characteristic p, then S can be recognized in Monte Carlo polynomial time among all black box groups of characteristic p.*

The definition of "black box groups of characteristic p" will be given in Section 4.1. We emphasize that Theorem 2.3 refers to recognition *among all groups*, not only among simple groups. For further discussion of the concept of recognition we refer to Section 4.2.

We complete the proof of Theorem 2.3 in Section 4.9 (Theorem 4.19).

We actually prove that the result holds more generally for those S for which every nontrivial S-module M of characteristic p has the property that a *non-negligible fraction* of the elements has a nonzero fixed point in its action on M. The running time will be inversely proportional to this "non-negligible fraction."

We prove that the groups $\mathrm{PSL}(2,p)$ are unisingular (Section 3.5), providing the first examples of simple groups over large fields which we can recognize. The power of Theorem 2.3 has very recently been greatly extended by Guralnick, Saxl and Tiep, who, answering the authors' question, classified all unisingular groups. We state their result [GST].

Theorem 2.4 (Guralnick, Saxl, Tiep). *Let S be a finite simple group of Lie type of characteristic p, defined over the field $GF(q)$, where $q = p^f$. Then S is unisingular if and only if S is one of the following:*

(a) $\mathrm{PSL}_n(p)$ *with* $n|p-1$*;*

(b) $\mathrm{PSU}_n(p)$ *with* $n|p+1$*;*

(c) $\mathrm{PSp}_{2n}(p)$ *with p odd;*

(d) $\Omega_{2n+1}(p)$ *with p odd;*

(e) $\mathrm{P}\Omega^{\epsilon}_{2n}(p)$ *with* $\epsilon = \pm$*, p odd, and* $\epsilon = (-1)^{n(p-1)/2}$*;*

(f) ${}^2\mathrm{G}_2(q)$, $\mathrm{F}_4(q)$, ${}^2\mathrm{F}_4(q)$, $\mathrm{E}_8(q)$, *q arbitrary;*

(g) $\mathrm{G}_2(q)$ *with* $q = p^f$ *and* $p \geq 3$*;*

(h) $\mathrm{E}_6(p)$ *with* $3|(p-1)$*;*

(i) ${}^2\mathrm{E}_6(p)$ *with* $3|(p+1)$*;*

(j) $\mathrm{E}_7(p)$ *with* p *odd.*

As a consequence, *we can recognize all these groups* in Monte Carlo polynomial time (among black box groups of characteristic p).

The second purpose of the reduction to the affine case was that the reduction allowed us to pinpoint where the main difficulty may lie. As an example, we offer as a challenge to distinguish $\mathrm{PSL}(2,p^2)$ from certain affine extension of $\mathrm{PSL}(2,p^2)$ by an elementary abelian group of order p^8 (Section 3.6).

3. The proportion of p-singular elements in affine extensions

In this section we describe general observations regarding the affine case, providing the basis for our analysis of the recognition of unisingular groups. We also show where the same idea fails for $\mathrm{PSL}(2,p^f)$, $f \geq 2$.

3.1. The proportion of p-singular elements. An element $g \in G$ is called *p-singular* if its order is a multiple of p; otherwise g is *p-regular*.

Let $\rho_p(G)$ denote the proportion of p-singular elements in a finite group G.

A good supply of p-singular elements is very helpful in recognizing finite simple groups of Lie type of characteristic p (cf. 4.14). The trouble is, if the field of definition is large, such elements become very scarce and difficult to find.

For instance, for $\mathrm{PSL}(2,q)$, we have the following proportion.

Proposition 3.1. $\rho_p(\mathrm{PSL}(2,q)) = 2/q$ *for odd* q *and* $\rho_p(\mathrm{PSL}(2,q)) = 1/q$ *for even* q. □

Recent work by Á. Bereczky [Ber] on linear groups and on certain classes of exceptional groups suggested that $\rho_p(S) \leq c/q$ may hold for all finite simple groups S of Lie type of characteristic p, defined over $\mathbf{F}_q$ (with a small absolute constant c). This was confirmed by independent work by Guralnick and Lübeck [GL].

Theorem 3.2 (Guralnick–Lübeck). *Let S be a finite simple group of Lie type of characteristic p defined over the field $\mathbf{F}_q$ ($q = p^f$). Then*

$$\rho'_p(S) < \frac{3}{q-1} + \frac{2}{(q-1)^2}, \tag{1}$$

where $\rho'_p(S)$ is the proportion of elements which commute with at least one element of order p.

Note that $\rho_p(S) < \rho'_p(S)$.

3.2. A naive algorithm for the affine case. It is evident that for $A \triangleleft G$ we have

$$\rho_p(G) \geq \rho_p(G/A).$$

We wish to estimate the quantity $\rho_p(G) - \rho_p(G/A)$.

If this quantity is "non-negligible" then a naive Monte Carlo algorithm will distinguish G from G/A: by sampling elements and checking for p-singularity, we statistically estimate ρ_p. The following statement gives the sample size for statistically significant estimation.

Proposition 3.3. *(a) We can statistically distinguish G from G/A by sampling*

$$\frac{2|\ln \epsilon|}{(\rho_p(G) - \rho_p(G/A))^2}$$

elements, where ϵ is the margin of error.

(b) If $\rho_p(G) \geq 1-\delta$ and $\rho_p(G/A) \leq \delta$ where $\delta < 1/4$ then we can statistically distinguish G from G/A by sampling

$$\frac{2|\ln \epsilon|}{|\ln(4\delta)|}$$

random elements. In other words, by sampling m random elements, we reduce the probability of error in our statistical inference below $(4\delta)^{m/2}$.

Proof. (a) Suppose the X_i are independent random $(0,1)$-variables with $\Pr(X_i = 1) = \rho$ and $\Pr(X_i = 0) = 1 - \rho$. Then for $\vartheta > 0$, the probability that $\rho - (1/m)\sum_{i=1}^m X_i \geq \vartheta$ is less than $\exp(-2m\vartheta^2)$ and the same holds for the probability that $\rho-(1/m)\sum_{i=1}^m X_i \leq -\vartheta$. This follows from a Chernoff bound stated as Theorem A4 in [AloS, p. 235].

We sample m elements from the given group; X_i will be the indicator whether or not the i-th element selected is p-singular. Let $\mu = (\rho_p(G) + \rho_p(G/A))/2$. If the number of p-singular elements selected is less than μm then we decide that the group is G/A, otherwise we declare it to be G. Setting $\vartheta = (\rho_p(G) - \rho_p(G/A))/2$, we see by Chernoff's bound that the probability that we make an erroneous decision is less than $\exp(-2m\vartheta^2)$ regardless of which of the two groups we are sampling.

This quantity is $\leq \epsilon$ for the stated value of m.

(b) This part is elementary. Suppose the X_i are independent random $(0,1)$-variables with $\Pr(X_i = 1) \le \delta$ and $\Pr(X_i = 0) = 1 - \Pr(X_i = 1)$. Then the probability that $(1/m)\sum_{i=1}^{m} X_i \ge 1/2$ is less than

$$\binom{m}{m/2}\delta^{m/2} < (4\delta)^{m/2}.$$

We again sample m elements from the given group; Y_i will be the indicator whether or not the i-th element selected is p-singular. If the number of p-singular elements selected is less than $m/2$ then we decide that the group is G/A, else it is G. Setting $Y_i = X_i$ if the group is G/a and $Y_i = 1 - X_i$ if the group is G we see by the above estimate that the probability that we make an erroneous decision is less than $(4\delta)^{m/2}$.

This quantity is $\le \epsilon$ for the stated value of m. □

We shall show that if $S = \mathrm{PSL}(2,p)$ then $\rho_p(G) - \rho_p(S) > 1 - 3/p$ regardless of the extension. Therefore $\mathrm{PSL}(2,p)$ can be efficiently distinguished from its affine extensions by sampling $O(|\log\epsilon|/\log p)$ elements (Section 3.5). More precisely, by sampling m elements we can distinguish these two groups with error probability less than $(8/p)^{m/2}$ (using $\delta = 2/p$ in part (b) of Proposition 3.3, cf. (Section 3.5)).

On the other hand, we find that for $S = \mathrm{PSL}(2,p^f)$, $f \ge 2$, there exists an affine extension G with $|A| = p^{4f}$ such that $\rho_p(G) - \rho_p(S) \le p^{1-f}$. The challenge is to distinguish G and S (as black-box groups of characteristic p) efficiently (for large p) (Section 3.6).

3.3. Split extensions suffice. The following observation saves us a lot of headache by allowing us to ignore non-split extensions.

Proposition 3.4. *Let A be an abelian p-group, S a finite group, and $1 \to A \to G \to S \to 1$ be an extension. Then the number of p-singular elements in G depends only on the S-module A_S and not on the specific extension.*

Proof. It suffices to show that the number of *p-regular* elements of G depends only on A_S. Let $\varphi : G \to S$ denote the natural epimorphism. Clearly, if $g \in G$ is p-regular then so is $\varphi(g) \in S$. It suffices to prove that for every p-regular $s \in S$, the number of p-regular elements in the coset $\varphi^{-1}(s)$ is determined by A_S.

Let $g \in \varphi^{-1}(s)$ and let $H = \langle A, g\rangle$. Now $A \lhd H$ and $|H/A|$ is the order of s, which is relatively prime to the order of A. Therefore $1 \to A \to H \to \langle s\rangle \to 1$ is a split extension, uniquely determined by the action of s on A. In particular, the distribution of orders of elements in the coset Ag is determined by the action of s on A. □

Remark 3.5. The proof shows that for every p-regular element $s \in S$, the number of p-singular elements in $\varphi^{-1}(s)$ is determined by the action of s on A.

3.4. Unisingular action of p-regular elements. Let A be a group and s an automorphism of A. By a *fixed point* of s we mean a nonidentity element of A fixed by s. If such an element exists, we call s *unisingular*. In the case that A is a vector space and s a linear transformation, s is unisingular exactly if 1 is an eigenvalue of s, i. e., $s-1$ is singular, hence the term.

The opposite of "unisingular" is "fixed-point-free."

Next we show that the proportion of p-regular elements in $\varphi^{-1}(s)$ critically depends on whether or not s is unisingular (in its action on A).

Proposition 3.6. *Let A be an abelian p-group, S a finite group, and $1 \to A \to G \to S \to 1$ be an extension. Let $s \in S$ be p-regular.*

(a) If s is fixed-point-free (in its action on A) then all elements of $\varphi^{-1}(s)$ are p-regular. (This case is bad for the naive algorithm.)

(b) If s is unisingular (not fixed-point-free) then at most a $1/p$ fraction of the elements of $\varphi^{-1}(s)$ is p-regular. (Good case.)

Proof. As in the proof of Proposition 3.4, let $\varphi : G \to S$ denote the natural epimorphism, let $g \in \varphi^{-1}(s)$ and let $H = \langle A, g\rangle$. As before, the extension $1 \to A \to H \to \langle s\rangle \to 1$ is split, so the elements of $\varphi^{-1}(s)$ can be written as affine transformations of the form $\alpha = \alpha(s,a) : x \mapsto sx + a$ $(a, x \in A)$ where sx denotes the image of $x \in A$ under the action of s (corresponding to conjugation by any member of the coset $\varphi^{-1}(s)$). (We write the operation in A additively.) Then $\alpha^i : x \mapsto s^i x + \beta(s,i)a$ where $\beta(s,i) = s^{i-1} + s^{i-2} + \cdots + s + 1$.

Let k be the order of s. Then $\alpha^k : x \mapsto x + \beta(s,k)a$ and therefore $\alpha^{kj} : x \mapsto x + j \cdot \beta(s,k)a$. It follows that the order of α is of the form kp^ℓ $(\ell \geq 0)$.

Now in case (a) we note that $\beta(s,k) = (s^k - 1)(s-1)^{-1} = 0$ and therefore the order of α is k regardless of the choice of a.

In case (b), $\beta(s,k) \neq 0$ since $\beta(s,k)b = kb \neq 0$ for $b \in \ker(s-1)$, $b \neq 0$. Therefore $\ker \beta(s,k)$ is a proper subgroup of A. Consequently the proportion of those $a \in A$ for which $\beta(s,k)a = 0$ is at most $1/p$. For all other $a \in A$ we have $\beta(s,k)a \neq 0$ and therefore the order of α is divisible by kp. □

Notation 3.7. Let S be a finite group and A an abelian p-group with an S-action. We use A_S to denote A as an S-module. Let $\lambda(A_S)$ denote the proportion in S of those p-regular elements which act as unisingular (not fixed-point-free) transformations on A.

The following is evident from the foregoing.

Corollary 3.8.

$$\lambda(A_S)(1-1/p) \le \rho_p(G) - \rho_p(S) \le \lambda(A_S).$$

Remark 3.9. Corollary 3.8 tells us that the naive algorithm succeeds if $\lambda(A_S)$ is not too small.

Observation 3.10. *A Lie-type simple groups S of characteristic p is unisingular if and only if*

$$\lambda(A_S) = 1 - \rho_p(S)$$

holds for every nontrivial S-module A_S of characteristic p.

Combining this observation with the Guralnick–Lübeck bound on $\rho_p(S)$ (Theorem 3.2), we obtain the following key inequality on $\lambda(A_S)$ for unisingular groups:

Corollary 3.11. *If S is a unisingular simple group of Lie type characteristic p over $\mathbf{F}_q$ then*

$$\lambda(A_S) > 1 - \frac{3}{q-1} - \frac{2}{(q-1)^2} \tag{2}$$

holds for every nontrivial S-module A_S of characteristic p.

This inequality implies that the naive algorithm performs very well on unisingular groups over large fields. We remark that for small fields, the work of Kantor and Seress quoted in Section 4.6 solves the recognition problem in a stronger sense (constructive recognition). However, no polynomial-time methods have previously been known to handle the cases when the field of definition is large.

3.5. $\mathrm{PSL}(2,p)$ **is unisingular.** In this section we prove that the naive algorithm works for the case when $S = \mathrm{PSL}(2,p)$, $p \ge 5$. In fact we prove that in this case, for any S-module A_S, *all* elements of S are unisingular. In the terminology of Definition 2.2, this means that $\mathrm{PSL}(2,p)$ is a *unisingular group.*

Theorem 3.12. *The groups* $\mathrm{PSL}(2,p)$ *(p an odd prime) are unisingular.*

Proof. Let A be an S-module where $S = \mathrm{PSL}(2,p)$. We need to prove that every element of S is unisingular (in its action on A).

Without loss of generality we may assume that A_S is irreducible. (Indeed, restriction to an irreducible quotient module can only decrease the number of unisingular elements.)

The irreducible modules of $\mathrm{SL}(2,p^f)$ in characteristic p have been characterized by Brauer and Nesbitt [BrN]. For the case $k=1$ we have the following modules M_t, $t=0,\dots,p-1$:

> M_t is the space of homogeneous polynomials of degree t in the two variables $\{x,y\}$. The action of $\mathrm{SL}(2,p)$ is defined by the linear substitutions corresponding to the elements of $\mathrm{SL}(2,p)$. (This is the t-th symmetric tensor power of the natural 2-dimensional module M_1.) Note that $\dim M_t = t+1$.

The irreducible modules of $\mathrm{PSL}(2,p)$ are those irreducible modules of $\mathrm{SL}(2,p)$ on which the center acts trivially, i. e., M_{2t}, $t=0,\dots,(p-1)/2$.

Let M_t' be the module obtained from M_t by extending the field of definition ($\mathbf{F}_p$) to its algebraic closure $\mathcal{F}_p$. This does not affect unisingularity.

Let s be an element of $\mathrm{SL}(2,p)$ with eigenvalues λ and $1/\lambda$ (in $\mathcal{F}_p$).

If $\lambda=\pm1$ then, after a suitable change of variables, we may assume that s acts on M_1' by $x\mapsto\pm x$. Therefore the element $x^{2t}\in M_{2t}'$ is a fixed point of s, hence s is unisingular.

If $\lambda\neq\pm1$ then $\lambda\neq1/\lambda$ and therefore s is diagonalizable over $\mathcal{F}_p$. After a suitable change of bases, s acts on M_1' by $x\mapsto\lambda x$ and $y\mapsto(1/\lambda)y$. Therefore the element $x^ty^t\in M_{2t}'$ is a fixed point of s, hence s is unisingular. □

Corollary 3.13. *Let A be an elementary abelian p-group, $S=\mathrm{PSL}(2,p)$, $p\geq5$, and let $1\to A\to G\to S\to1$ be an extension. Then $\rho_p(G)-\rho_p(S)>1-(3/p)$.*

Proof. Indeed, by Theorem 3.12, $\lambda(A_S)=1-\rho_p(S)=1-2/p$, and therefore by Corollary 3.8,

$$\rho(G)-\rho(S)\geq(1-2/p)(1-1/p)>1-3/p. \qquad \square$$

3.6. Case $\mathrm{PSL}(2,p^f)$, $f\geq2$: challenge. For unisingular groups over large fields, $\lambda(A_S)$ is close to 1 for all nontrivial S-modules of characteristic p (Corollary 3.11). Unfortunately nothing like this holds for $S=\mathrm{PSL}(2,p^f)$, $f\geq2$.

Proposition 3.14. *Let $S=\mathrm{PSL}(2,q)$ where $q=p^f$, $f\geq2$. Then there exists an irreducible S-module A_S of dimension 4 over $\mathbf{F}_q$ such that $\lambda(A_S)\leq p/q$.*

Proof. We are again guided by the Brauer–Nesbitt Theorem. In the general case ($q=p^f$), the theorem says that the irreducible modules in characteristic p for $\mathrm{SL}(2,q)$ are the tensor products of the form

$$V_0\otimes V_1^{\sigma}\otimes\cdots\otimes V_{f-1}^{\sigma^{f-1}},$$

where σ is the Frobenius automorphism of $\mathbf{F}_q$ and the V_i are modules described in the proof of Theorem 3.12.

Let now $V_0 = V_1 = M_1$ be the natural module for $\mathrm{SL}(2,q)$ (2-dimensional over $\mathbf{F}_q$). We can represent the 4-dimensional module $A = V_0 \otimes V_1^\sigma$ as a space of polynomials in the variables x_0, y_0, x_1, y_1, spanned by the four polynomials $z_0 z_1$ where $z_i \in \{x_i, y_i\}$.

Let now s be a p-regular element of S with eigenvalues $\lambda^{\pm 1}$. Noting that s is diagonalizable over $\mathcal{F}_p$, we see, that, after a suitable change of bases, s acts on M_1' by $x \mapsto \lambda x$ and $y \mapsto (1/\lambda)y$. It follows that $z_0 z_1 \mapsto \lambda^{\pm 1 \pm p} z_0 z_1$, therefore the four eigenvalues of the s-action on A are $\lambda^{\pm 1 \pm p}$. If the action of s on A is unisingular, one of these eigenvalues must be 1, therefore either $\lambda^{p-1} = 1$ or $\lambda^{p+1} = 1$. Standard calculations show that fewer than a p/q fraction of the elements of $\mathrm{SL}(2,q)$ satisfy this condition.

Since $-id \in \mathrm{SL}(2,q)$ acts trivially on A, we can view A as a module for $\mathrm{PSL}(2,q)$. The proportion of unisingular elements is clearly the same. □

Corollary 3.15. *Let $S = \mathrm{PSL}(2,q)$ where $q = p^f$, $f \geq 2$. Then S has an extension $1 \to A \to G \to S \to 1$ by the elementary abelian group A of order $|A| = q^4 = p^{4f}$ such that the proportion of p-regular elements in G is*

$$\rho(G) < (p+2)/q. \tag{3}$$

Proof. Combine Corollary 3.8 with the fact that $\rho(S) \leq 2/q$. □

4. Simplicity

4.1. Black-box groups of characteristic p. Following [BaB, Sec. 8.2], we say that H is a *black-box group of characteristic p* if H is a black-box group of some encoding length n and H is isomorphic to a section (quotient of subgroup) of $\mathrm{GL}(d,p)$ where $d = \lfloor n/\log p \rfloor$. When we say that such a group H is *given*, we tacitly assume that p is known.

The significance of this concept is that it is tailor-made to treat black-box algorithms for matrix groups, retaining the exact amount of information required for polynomial-time analysis.

Among the many advantages of working with this concept, let us highlight one.

Algorithms for matrix groups and for black-box groups often assume that we are able to determine the order of group elements. This, however, is only true if we possess a superset of the prime divisors of the order of the group. For $\mathrm{GL}(n,p)$, finding such a superset is, alas, as difficult as factoring integers. (Cf. [BaB, Sec. 8].)

This unjustifiable assumption seems difficult to avoid in the context of black-box groups. But for matrix groups as well as for black-box groups of a given characteristic, many significant applications survive without assuming a

superset of prime divisors of the order of the group; an explicitly defined small set of "pretend-primes" suffices instead. (See Remark 4.10 for details.)

4.2. Recognition within a given class of groups. Let $\mathcal{G}$ be a class of black-box groups. We say that a procedure *recognizes a group G within the class $\mathcal{G}$* if given a black-box group $H \in \mathcal{G}$, the procedure will accept H if $H \cong G$ and reject if $H \ncong G$.

The following result was proved recently:

Theorem 4.1. *Every finite simple group of a given characteristic can be recognized among finite simple groups of the same characteristic (given as black-box groups of that characteristic) by a polynomial-time Monte Carlo algorithm.*

This result appears in [BaKPS] and is based on work by Altseimer, Borovik [AltB] and Kantor and Seress [KS2]. Its content is that given two nonisomorphic finite simple groups of the same (known) characteristic, a polynomial time statistical procedure can tell them apart.

We emphasize that for this procedure we need to know that *both* groups are simple.

The key question we wish to address is whether the assumption of simplicity of the members of the class $\mathcal{G}$ can be dropped. (We continue to assume that G itself is simple.)

4.3. Naming the composition factors. Finding the orders of composition factors of a cyclic group is as hard as factoring integers, which is not expected to be doable in polynomial time.

We are faring much better regarding nonabelian composition factors.

First we state the main result of [BaB].

Theorem 4.2 ([BaB]). *Given a black-box group G of characteristic r, one can construct, in Monte-Carlo polynomial time, black-box-group representations of characteristic r of all nonabelian composition factors of G. Moreover, if a composition factor L is a sporadic group, alternating group, or a simple group of Lie type of characteristic other than r, then a permutation representation of L can be constructed in polynomial time.*

(This is a combination of Theorems 1.2 and 8.6 of [BaB].)

Remark 4.3. Given a permutation representation of L, we can switch to a standard representation via an explicit isomorphism in deterministic polynomial time (Kantor [Ka]). Thus we learn the standard name of L in polynomial time (and virtually everything else we may wish to know about L).

Combined with Theorem 4.1, this result yields the following powerful corollary.

Corollary 4.4. *Given a black-box group G of given characteristic, one can compute, in Monte Carlo polynomial time, the standard names of all nonabelian composition factors of G.* □

In particular, this result applies to matrix groups over finite fields.

Corollary 4.5. *Given a matrix group over a finite field, one can compute, in Monte Carlo polynomial time, the standard names of all nonabelian composition factors of G.* □

4.4. Membership testing in simple normal subgroups. The following result provides a useful tool in many contexts. It settles the "Outer automorphism problem" (recognizing S within G where $S \leq G \leq \mathrm{Aut}\,(S)$), stated as Problem 10.3 in [BaB].

Theorem 4.6. *Let S be a nonabelian simple group. Assume that S is a normal subgroup of a black-box group G of given characteristic. Given G and S, we can test membership in S of any $g \in G$ in Monte Carlo polynomial time.*

The following simple algorithm was communicated to us by Charles Leedham-Green. Our contribution is its analysis.
Algorithm (C. Leedham-Green)

> pick a reasonable number of random elements $s_i \in S$
> compute the g. c. d. of the orders of the elements gs_i
> **if** the g. c. d. is 1 **then** output "member"
> **else** output "probably not member"

Remark 4.7. We shall see that $O(\sqrt{n}\log n)$ is a "reasonable number," where n is the encoding length of the black-box group H.

Remark 4.8. This algorithm assumes that we can compute the orders of elements. See Remark 4.10 on how to avoid this assumption.

The analysis of the algorithm is based on the following result.

Theorem 4.9 ([BaPS]). *Let S be a finite simple group and r a prime number. Then at least a c/d fraction of the elements of S is r-regular where $c > 0$ is an absolute constant and $d = d(S)$ is defined as follows. For the alternating group A_t we set $d(A_t) = \sqrt{t}$, for a classical simple group S we define $d(S)$ to*

be the dimension of the projective space on which S acts; for all other simple groups S, $d(S) = 1$. □

(For the alternating groups, the weaker bound c/t was quoted in [BaB]. This would suffice for the polynomial-time claim but not for the claim stated in Remark 4.7.)

Analysis of the algorithm.

It is obvious that the order of g mod S will divide the order of each gs_i so if $g \notin S$ then the output will always be correct ("probably not member").

Let us now assume $g \in S$. In this case, each gs_i is a random member of S. Let $p_1, \ldots, p_k$ be the set of primes dividing the order of S; clearly, $k < n$ (in fact, $k < n/\log n$) where n is the encoding length of our black box group.

Suppose now that we chose m random elements s_i. By Theorem 4.9, the probability that a particular p_j divides the orders of each gs_i is less than $(1 - c/d)^m < \exp(-cm/d)$. The probability that the g. c. d. of these orders is greater than 1 is therefore less than $k \exp(-cm/d) \leq \epsilon$ assuming $m \geq (d/c)\log(k/\epsilon)$.

Note that $d = O(\sqrt{\log |S|}) = O(\sqrt{n})$; therefore sampling $m = O(\sqrt{n}\log n)$ random elements s_i suffices for any constant margin of error. □

Remark 4.10. As remarked in Section 4.1, we cannot compute the order of the elements in polynomial time. Hence, strictly speaking, we did not complete the proof of Theorem 4.6. This can be fixed by copying the methods of [BaB, Section 8.4] where a notion of "pseudo-order" with respect to a set of "pretend-primes" is introduced. The "pretend-primes" are the "small" primes ($\leq n$) plus a set of numbers obtained as cyclotomic polynomials and semi-cyclotomic polynomials of the characteristic p of our black-box group. (These numbers occur as the orders of certain maximal tori in simple groups of characteristic p.) More precisely, we divide these numbers by any small prime divisors; this way we end up with a set of pairwise relatively prime "pretend-primes."

The "pseudo-order" is the smallest multiple of the order which factors as a product of our "pretend-primes." The set of pretend-primes needed for computation in a black-box group of given characteristic is computable in polynomial time, and so is the pseudo-order of elements with respect to a given set of pretend-primes.

If we apply the above algorithm to black-box groups of known characteristic, we can replace "order" by "pseudo-order" of an element. As in [BaB, Section 8.4], this transition requires a more detailed version of Theorem [BaPS] (see [BaB, Theorem 8.7]). (Note: omit the word "cyclic" from the statement of [BaB, Theorem 8.7]. This correction does not affect any of the known applications of the theorem.)

4.5. Reduction to perfect, solvable-by-simple groups. A group G is perfect if $G' = G$. It is *solvable-by-simple* if it has a solvable normal subgroup N such that the quotient G/N is simple nonabelian. If S is a nonabelian simple group then we say that G is a solvable-by-simple extension of S or simply G is solvable-by-S if G has a solvable normal subgroup N such that $G/N \cong S$.

Theorem 4.11. *Let S be a finite simple group of Lie type of characteristic p. If S can be recognized in Monte-Carlo polynomial time among its perfect, solvable-by-simple extensions, given as black-box groups of characteristic p, then S can be recognized in Monte-Carlo polynomial time among all finite groups given as black-box groups of any known characteristic (not necessarily p and p need not be known).*

Proof. Let G be a black-box group of characteristic r; we wish to decide whether or not $G \cong S$, where S is also given as a black-box group of some (known) characteristic (not necessarily r or p).

Applying Corollary 4.4 to the given representation of S, we find the standard name of S and learn the value p. Applying Corollary 4.4 to G we are able to refute the hypothesis $G \cong S$ unless G has exactly one nonabelian composition factor, L, and this composition factor is a simple group of Lie type of characteristic p and r simultaneously. Since a finite number of groups can be treated as sporadic (using ad hoc methods), we infer that $r = p$. We may further assume that $L \cong S$. The question that remains is to decide under these conditions whether or not G is simple.

Let H be the stable derivative of G, i. e., the smallest normal subgroup of G such that G/H is solvable. This can be found in Monte Carlo polynomial time [BaCFLS].

Clearly, H is perfect and solvable-by-simple, given as a black-box group of the right characteristic. Therefore, by assumption, we are able to decide whether or not H is simple. We may assume it is, so $H \cong S$.

Now an application of Theorem 4.6 completes the proof: we pick some random elements of G and test whether or not they belong to H. If any of them does not, then $G \neq H$; otherwise we can "bet" that $G = H$ (the probability of error is at most 2^{-m} where m is the number of random elements selected). □

Remark 4.12. We did not estimate the overall probability of error in this algorithm; 2^{-m} is only the error probability attributable to our failure to catch an element outside the proper subgroup H. All other Monte Carlo subroutines used contribute to the error (most notably the routine to recognize the standard names of the composition factors (Corollary 4.4), and on a more fundamental level, the routines used for generating nearly uniform random elements and normal closures). These routines must run long enough to make their contribution to the error probability negligible. Exponentially small error

probabilities are achievable in polynomial time by the definition of polynomial-time Monte Carlo algorithms (cf. [Ba2]).

4.6. Recognizing classical groups over tiny fields. "Tiny" is a well-defined technical term: an input parameter q is *tiny* if it is given in *unary* (q tally marks). Since we measure the complexity of an algorithm as a function of the length of the input, a polynomial-time algorithm with a tiny field of order q as input is allowed to take time $q^{O(1)}$, as opposed to $(\log q)^{O(1)}$ under the ordinary (binary) encoding of the field order.

Theorem 4.13. *Every finite simple group of Lie type of a given characteristic over a tiny field can be recognized among all finite groups, given as black-box groups of any characteristic, by a polynomial-time Monte Carlo algorithm.*

The proof is based on the recent monumental work by Kantor and Seress [KS1] which proves the following main result.

Theorem 4.14 ([KS1]). *Let $\mathcal{C}_p$ denote the class of classical finite simple groups of characteristic p. Then one can, in Monte-Carlo polynomial time, constructively recognize black-box members of $\mathcal{C}_p$ within $\mathcal{C}_p$ assuming the field of definition is tiny.* □

In this result, p is an input variable, not assumed to be constant.

"Constructive recognition" means that an explicit isomorphism is constructed with a standard matrix representation of the group in question. As pointed out in [KS1], it follows, using Steinberg's presentations, that a presentation in terms of generators and relations can also be constructed in Monte Carlo polynomial time. This observation immediately upgrades the Monte Carlo algorithm of Theorem 4.14 to a *Las Vegas algorithm*, i. e., a Monte Carlo algorithm which does not err but is allowed to report failure with a small probability. (If the presentation does not work out as expected, report failure; else we have proof that the group is what we believe it is.) This conceptual upgrade is especially significant if we generate random elements using unproven heuristics (as we always do in practice). (See [Ba2] for a discussion of related issues.)

Unfortunately, a Las Vegas upgrade of Theorem 4.13 would seem to require entirely new ideas. Kantor and Seress [KS1] express pessimism about the possibility of such upgrade (in a paragraph following their Corollary 1.8) even when both groups are simple classical and one of them is over a tiny field. In fact, even a Las Vegas separation (proof of non-isomorphism) of $\mathrm{PSL}(2m, p)$ and $\mathrm{PSL}(2, p^m)$ in time $m^{O(1)}$ seems to be out of reach, where $\mathrm{PSL}(2, p^m)$ is given as a black-box group of characteristic p, while $\mathrm{PSL}(2m, p)$ is given explicitly in its natural projective representation and p is a fixed small prime.

The trouble is that we are unable to recover the geometry behind a linear group over a very large field, and in particular we are unable to construct a presentation of $\mathrm{PSL}(2, p^m)$, given as a black-box group, when m is moderately large.

Even bigger problems seem to loom when looking for a Las Vegas comparison of the unitary groups $\mathrm{PSU}(3, p^m)$ and $\mathrm{PSU}(3m, p)$: *short presentations* are not even known to exist for $\mathrm{PSU}(3, p^m)$ (even for fixed small p), let alone constructed efficiently. (A "short presentation" of a group G is a presentation that can be stated using only $(\log |G|)^{O(1)}$ symbols. In the case of $\mathrm{PSU}(3, p^m)$ this would mean length $(m \log p)^{O(1)}$.) Short presentations are known to exist for all finite simple groups with the possible exception of the rank-1 twisted Lie-type groups, viz. $\mathrm{PSU}(3, q)$, ${}^2B_2(q)$ (Suzuki) and ${}^2G_2(q)$ (Ree) [BaGKLP].

Proof of Theorem 4.13. According to Theorem 4.11, we may assume that G is solvable-by-S, where S is a finite simple group of Lie type of characteristic p over a tiny field $\mathbf{F}_q$ and G is given as a black-box group of characteristic p; the task is to decide whether or not $G \cong S$.

First we dispose of the case when S is an exceptional group (this case is not considered in [KS1]). Indeed, in this case $|S| < q^{248}$ so the order of S is polynomially bounded as a function of the length of the input (because $\mathbf{F}_q$ is *tiny*), and we can decide in polynomial time whether or not $|G| > |S|$ by exhaustive search. (The permutation group methods of [BaB] reviewed in Section 4.8 will considerably reduce the revolting exponent 248 but substantial further work will be required to produce a managable exponent.)

Let us now assume that S is classical and work from Theorem 4.14.

Let R denote the solvable radical of G (largest solvable normal subgroup). Membership in R can be decided in Monte Carlo polynomial time (check solvability of the normal closure of the given element, [BaCFLS]). Note that $G/R \cong S$.

All we need to do is decide whether or not $R = 1$. This is done by the standard method of "sifting" the defining relations of S as follows: let $\varphi : G \to S$ be the epimorphism to which our explicit $G/R \to S$ isomorphism lifts. Let $S = \langle s_i : i \le a \,|\, R_j : j \le b \rangle$; let $g_i \in G$ be a lifting of s_i to G. Let G be given by a set $\{h_\ell : \ell \le f\}$ of generators. Let $t_\ell = \varphi(h_\ell)$. Express t_ℓ as a word $w_\ell(\dots, s_i, \dots)$. Then R is the normal closure of the elements $R_j(\dots, g_i, \dots)$ and the elements $h_\ell^{-1} w_\ell(\dots, g_i, \dots)$. □

4.7. Central-by-simple groups

Theorem 4.15. *Let G be a black-box group of characteristic p. Assume $G/Z(G)$ is simple nonabelian. Then we can find $Z(G)$ in Monte Carlo polynomial time.*

Proof. The algorithm will again be based on Theorem 4.9.

For simplicity, we again pretend that we know a superset of primes dividing the order of G and so we can compute the order of elements. (We commented on this assumption in Section 4.1. We can avoid this assumption with the same technique as for Theorem 4.6. For the details, see Remark 4.10.)

Let r be a prime and let R be a Sylow r-subgroup of $Z(G)$. We shall collect generators for R in a set $\mathcal{S}$.

Algorithm

initialize the set $\mathcal{S} := \emptyset$
repeat a reasonable number of times:
  pick a random element $g \in G$
  write the order of g as $r^\ell b$ where r does not divide b
  let $h = g^b$
  if $h \in Z(G)$ **then** add h to $\mathcal{S}$

Remark 4.16. We shall see that $O(n^{3/2})$ is a "reasonable number," where n is the encoding length of the black-box group G.

Analysis. Let $S = G/Z(G)$ and let $\varphi : G \to S$ be the natural epimorphism.

Let us imagine that the random choice of elements of G is performed in two stages: first we pick a random element $s \in S$, and then a random member $g \in \varphi^{-1}(s)$. Clearly, if both choices are uniform, so is their composition.

By Theorem 4.9, with reasonable frequency we are likely to pick an r-regular $s \in S$. If this is the case, it is easy to see that the resulting element h will be a uniformly distributed random member of R. Indeed, let T denote the r'-Hall subgroup of $Z(G)$ (comprising all r-regular elements of $Z(G)$). Let further $g_0 \in \varphi^{-1}(s)$. Then $g = g_0 z_1 z_2$ where z_1 is a random member of T and z_2 is a random member of R. Now $h = g^b$ has order r^ℓ; on the other hand, s is r-regular; therefore $h \in Z(G)$. It follows that $h \in R$. The same is true for $h_0 := g_0^b$. Therefore $z_1^b z_2^b \in R$; consequently, $z_1^b = 1$. Finally, z_2^b is uniformly distributed over R, therefore so is h.

To estimate what is a "reasonable number" of repetitions, let $d = d(S)$ be the quantity defined in Theorem 4.9. $C_1 d \log |R|)$ repetitions are likely to turn out $\geq C_2 \log |R|$ random elements of R; for $C_2 > 1$, these are likely to generate R. We can use the crude estimates $d < \sqrt{n}$ and $\log |R| < n$ to obtain an $O(n^{3/2})$ bound on the number of repetitions (n is the encoding length). □

4.8. Reduction to perfect, p-by-simple groups. In this section we further narrow the obstacle exhibited in Theorem 4.11.

A group G is *p-by-simple* if it has a normal p-subgroup N such that the quotient G/N is simple nonabelian. If S is a nonabelian simple group then we

say that G is a p-by-simple extension of S if G is p-by-simple and $G/N \cong S$. We say that G is *non-central minimal p-by-simple* if N is a non-central minimal normal subgroup of G.

Theorem 4.17. *Let S be a finite simple group of Lie type of characteristic p. If S can be distinguished in Monte-Carlo polynomial time from its non-central minimal p-by-simple extensions, given as black-box groups of characteristic p, then S can be recognized in Monte-Carlo polynomial time among all finite groups given as black-box groups of any characteristic (not necessarily p and p need not be known).*

Proof. In view of Theorem 4.11, we may assume that we are given a perfect, solvable-by-simple group G such that $G/N \cong S$ where N is the solvable radical of G; and moreover that G is given as a black-box group of characteristic p. The task is to decide whether or not $N = 1$.

We note that membership in N is decidable in Monte Carlo polynomial time (check solvability of the normal closure). As soon as a nonidentity element of N is found, we are done.

First we perform the algorithm "PERM(G, m)" of [BaB, Theorem 6.1] with a judiciously chosen value m. ($m = n$ will suffice.) This algorithm carries the following guarantee:

Theorem 4.18 ([BaB]). *Assume G is a black-box group of given characteristic. Assume further that G has a nontrivial permutation representation of degree $\leq m$. Then Algorithm* PERM(G, m) *either returns a nonidentity element which belongs to a proper normal subgroup of G or returns a faithful permutation representation of degree $\leq m^c$. The algorithm works in Monte Carlo time $(m + n)^{O(1)}$.* □

If the algorithm produces a permutation representation, the fact that it is faithful can be verified in deterministic polynomial time. In this case we know virtually everything about G from the theory of permutation group algorithms (cf. [Lu3]). Specifically, an algorithm of Luks can be used to decide simplicity [Lu1].

If PERM fails to produce either of the desired objects, we perform the algorithm of Theorem 4.15. With each element h generated, we check whether $h \in N$.

If no element of N is found, we pretend that either $N = 1$ or N is an elementary abelian p-group which is non-central, minimal normal in G, and run the postulated algorithm to decide whether or not $N = 1$, checking along the way every element generated for membership in N. If none is found, we decide that $N = 1$, i. e., G is simple.

Analysis. We conclude that $N \neq 1$ only if we find a nonidentity element of N, so if $N = 1$ then our decision is certainly correct (except for possible errors made by the Monte Carlo algorithm for normal closures).

Suppose now that $N \neq 1$ and let M be a proper subgroup of N, maximal with respect to the condition that $M \lhd G$. Then N/M is a minimal normal subgroup of G/M and therefore it is an elementary abelian r-group for some r; let $|N/M| = r^k$. Note that $k \leq \log_2 |G| \leq n$.

If $N/M = Z(G/M)$ then it is easy to see that the second phase of the algorithm will find a nonidentity element of N. Indeed, pretend we compute in G/M. The only difference occurs when some element $u \in G$ is tested for being the identity. If $u = 1$, we report that u is the indentity and continue. If $u \in N$ and $u \neq 1$ then we terminate: we found what we wanted. If $u \notin N$, we report that $u \neq 1$ and continue. As long as no nonidentity element of N is encountered, the computations in G and in G/M proceed in the exact same steps.

Assume now that N/M is not central in G/M. Now S acts faithfully (via conjugation) on N/M.

If $r \neq p$ then this implies that S has a permutation representation of small degree ($\leq k^c \leq n^c$ for a small constant c). This follows from a theorem of Landazuri and Seitz [LaS] (cf. [BaB, Theorem 4.1]). But then, the first phase of the algorithm must have produced the required result.

Finally if $r = p$ then the last phase of the algorithm succeeds for a reason similar to why the algorithm to find the center worked when pretending to compute over G/M. □

4.9. Recognizing unisingular simple groups. Finally we are in the position to tackle some classes of Lie-type simple groups over large fields.

Theorem 4.19. *If S is a finite simple group of Lie type of characteristic p and S is unisingular then S can be recognized in Monte Carlo polynomial time among all black-box groups of known characteristic (not necessarily p and p need not be known).*

Proof. Let q be the order of the field of definition of S. For $q < 13$ we refer to the work of Kantor and Seress for groups over tiny fields and specifically to Theorem 4.13 above. Henceforth we assume $q \geq 13$. As a consequence, we have $\rho_p(S) < 1/3$. Indeed, by the quoted result of Guralnick and Lübeck (Theorem 3.2), we have

$$\rho_p(S) \leq \rho'_p(S) < \frac{3}{q-1} + \frac{2}{(q-1)^2} < \frac{1}{3}.$$

By Theorem 4.17, the proof of Theorem 4.19 is reduced to the following situation: we are given a black-box group G of characteristic p known to have a

normal subgroup N such that $G/N \cong S$ and either $N = 1$ or N is a non-central minimal normal p-subgroup of G.

The task is to decide whether or not $N = 1$.

Let us call $g \in G$ a "witness" of the assertion "$N \neq 1$" if g is p-singular (in G) but gN/N is p-regular (in S). The algorithm tries to find such a witness.

If a witness is found, we declare "$N \neq 1$." Otherwise we declare "probably $N = 1$."

Algorithm

repeat a reasonable number of times:
pick a random element $g \in G$
let b be the (pseudo)order of g
if $p|b$ **then** let $h = g^{b/p}$
if the normal closure $\langle h^G \rangle$ is solvable **then** output h, **halt**
end(repeat)
output "probably $N = 1$."

Assuming the normal closure and solvability routines do not fail, no error can occur if $N = 1$. Assume now that $N \neq 1$. Then the probability that a random $g \in G$ is a witness is $\rho_p(G) - \rho_p(S)$. Therefore the probability that we fail to find a witness is $(1 - \rho_p(G) + \rho_p(S))^m$ where m is the number of repetitions.

Combining Corollary 3.8 and Observation 3.10 we obtain that for unisingular simple groups, this quantity is always less than $(\rho_p(S) + 1/p)^m$. Under our assumption of $q \geq 13$ we have $\rho_p(S) < 1/3$ and therefore the probability of failing to detect $N \neq 1$ is less than $(5/6)^m$. So a modest constant number of repetitions is "reasonable."

For large p we can use the better estimate $(4/q + 1/p)^m$ to see that in that case, already $m = 1$ suffices with large probability. □

Remark 4.20. Note that this algorithm not only decides whether or not $N = 1$ but it is also finds a nontrivial element of N if $N \neq 1$. This demonstrates a special case of a general principle discovered by Beals [Be, Lemma 5.2]: *any* black-box algorithm which distinguishes a simple group S from its extension G ($G/N \cong S$, $N \neq 1$) can be modified to find a nontrivial element of a proper normal subgroup of G. (The Monte Carlo complexity of the modified algorithm is polynomially bounded in the input length and the complexity of the original algorithm.)

5. Open questions

1. *(p-core problem)* The big open question is to decide whether or not $O_p(G) = 1$ where $O_p(G)$ is the p-core of G (the largest normal p-subgroup) and G is a matrix group over $\mathbf{F}_p$ or more generally, a black-box group of characteristic p (Section 4.1). – Is the p-core problem any easier for matrix groups than for black-box groups of characteristic p?

2. *(Affine case)* The special case of the p-core problem described in the Abstract deserves particular attention. Its solution would resolve the simplicity problem and would constitute a major step toward solving the general p-core problem.

3. *(Non-unisingular groups.)* The question is, how bad are those simple groups S of Lie-type of characteristic p which are not unisingular. The measure of badness is the quantity

$$\lambda(S) = \min\ \lambda(A_S) \tag{4}$$

where the minimum is taken over all irreducible S-modlues A_S of characteristic p. The smaller this value, the worse the group. The running time of our naive algorithm for the affine case is proportional to $1/\lambda(S)$. (Recall that $\lambda(A_S)$ is the proportion in S of those p-regular elements which act as unisingular transformations on A. We called a linear transformation *unisingular* if 1 is among its eigenvalues. We called the group S unisingular if all elements of S act as unisingular transformations on all nontrivial S-modules.)

Let $\mathbf{F}_q$ be the field of definition of S. We are concerned with the cases of large q.

For unisingular S, we have $\lambda(S) = 1 - O(1/q)$, i. e., $\lambda(S)$ is close to 1 (see equation (2)).

On the other hand, for $S = \mathrm{PSL}(2, q)$, $q = p^f$ $(f \geq 2)$, Proposition 3.14 asserts that

$$\lambda(S) \leq p/q. \tag{5}$$

Does this upper bound generalize to all simple groups of Lie type that are not unisingular? Guralnick, Saxl, and Tiep suggest [GST] that a bound of the form

$$\lambda(S) \leq cpd^C/q \tag{6}$$

might hold, where d the Lie-rank, p is the characteristic, and $q = p^f$ is the order of the field of definition; c and C are absolute constants. Such a bound would show that the naive algorithm does not work for any of the non-unisingular groups with $f \geq 2$ when q is large.

The estimation of $\lambda(S)$ for the non-unisingular simple groups with $f = 1$ ($q = p$) is a further relevant question. The groups $\mathrm{PSL}(3, p)$ with $p \not\equiv 1 \pmod 3$ would be first class to consider.

4. Recognize (test membership in) the derived subgroup G' of a matrix group, or more generally, of a black box group of characteristic p.

Acknowledgments

The first named author wishes to thank Jonathan Alperin, George Glauberman, Charles Leedham-Green, Péter P. Pálfy and Ákos Seress for valuable discussions. The authors are indebted to Bill Kantor and Ákos Seress for their meticulous editorial work and in particular, for catching a number of inaccuracies in the paper. The authors are also indebted to Bill and Ákos for the opportunity to present this work at the OSU meeting and an update at the May 2000 meeting in Milan. Last but not least, we wish to thank Bob Guralnick, Jan Saxl and Pham Huu Tiep for providing a quick and complete solution, within days of the Milan meeting, to the problem of characterizing the unisingular simple groups of characteristic p. Their fast work allowed a last-minute update of the paper before being sent to print.

References

[AloS] N. Alon, J. H. Spencer: *The Probabilistic Method.* Wiley 1992.

[AltB] C. Altseimer, A. V. Borovik: Probabilistic recognition of orthogonal and symplectic groups. *In:* [G&C3] (this volume)

[Ba1] L. Babai: Local expansion of vertex-transitive graphs and random generation in finite groups. *In: Proc. 23rd ACM Symp. on Theory of Computing (STOC'91)*, 1991, pp. 164–174.

[Ba2] L. Babai: Randomization in group algorithms: conceptual questions. *In:* [G&C2], pp. 1–16.

[BaB] L. Babai, R. Beals: A polynomial-time theory of black-box groups I. In: *Groups St Andrews 1997 in Bath, I* (C.M. Campbell, E. F. Robertson, N. Ruskuc, G. C. Smith, eds.), London Math. Soc. Lect. Notes 260, Cambr. U. Press, 1999, pp. 30–64.

[BaCFLS] L. Babai, G. Cooperman, L. Finkelstein, E. M. Luks, Á. Seress: Fast Monte Carlo algorithms for permutation groups. *J. Computer and System Sci.* **50** (1995), 296–307.

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

[BaKPS] L. Babai, W. M. Kantor, P. P. Pálfy, Á. Seress: Black-box recognition of finite simple groups of Lie type by statistics of element orders. Manuscript, in preparation.

[BaPS] L. Babai, P. P. Pálfy, J. Saxl: On the number of p-regular elements in simple groups. In preparation.

[BaS] L. Babai, E. Szemerédi: On the complexity of matrix group problems I. *In: Proc. 25th IEEE Symp. Found. Comp. Sci. (FOCS'84)*, 1984, pp. 229–240.

[Be] R. Beals: Towards polynomial time algorithms for matrix groups. *In:* [G&C2], pp. 31–54.

[Ber] Á. Bereczky: On the proportion of p-singular elements in finite simple groups of Lie type of characteristic p. Private communication, 1999.

[BrN] R. Brauer, C. Nesbitt: On the modular characters of groups. *Annals of Math.* **42** (1941) 556–590.

[CLMNO] F. Celler, C. R. Leedham-Green, S. Murray, A. Niemeyer and E. A. O'Brien: Generating random elements of a finite group. *Comm. Alg.* **23** (1995), 4931–4948.

[G&C1] *Groups and Computation.* Proc. 1991 DIMACS Workshop. (L. Finkelstein and W. M. Kantor, eds.) DIMACS Ser. in Discr. Math. and Theor. Comp. Sci. Vol 11, A. M. S. 1993.

[G&C2] *Groups and Computation II.* Proc. 1995 DIMACS Workshop. (L. Finkelstein and W. M. Kantor, eds.) DIMACS Ser. in Discr. Math. and Theor. Comp. Sci. Vol 28, A. M. S. 1997.

[G&C3] *Groups and Computation III.* Proc. 1999 Workshop at the Ohio State University (W. M. Kantor and Ákos Seress, eds.) OSU Mathematical Research Institute Publications, deGruyter, Berlin – New York, 2000 (this volume)

[GL] R. M. Guralnick, F. Lübeck: On p-singular elements in Chevalley groups in characteristic p. *In:* [G&C3] (this volume)

[GST] R. M. Guralnick, Jan Saxl, Pham Huu Tiep: *private communication,* May 2000.

[Ka] W. M. Kantor: Sylow's theorem in polynomial time. *J. Computer Sys. Sci.* **30** (1985), 359–394.

[KS1] W. M. Kantor, Á. Seress: Black box classical groups. *Memoirs of the A. M. S.*, to appear.

[KS2] W. M. Kantor, Á. Seress: Prime power graphs for groups of Lie type. In preparation.

[LaS] V. Landazuri, G. M. Seitz: On the minimal degrees of projective representations of the finite Chevalley groups. *J. Algebra* **32** (1974), pp.

418–443.

[Lu1] E. M. Luks: Computing the composition factors of a permutation group in polynomial time. *Combinatorica* **7** (1987), pp. 87–99.

[Lu2] E. M. Luks: Computing in solvable matrix groups. *In: Proc. 33rd IEEE Symp. Found. Comp. Sci. (FOCS'92)*, 1992, pp. 111–120.

[Lu3] E. M. Luks: Permutation groups and polynomial-time computation. *In:* [G&C1], pp. 139–175.

Department of Computer Science
University of Chicago
Chicago, IL 60637
E-mail: laci@cs.uchicago.edu
Partially supported by NSF grant CCR-9732205
and
Department of Mathematics
Hebrew University
Jerusalem
E-mail: shalev@math.huji.ac.il
Partially supported by the Israel Science Foundation