Journal of Algebra 421 (2015) 16–26

Contents lists available at ScienceDirect

Journal of Algebra

www.elsevier.com/locate/jalgebra

Black box groups isomorphic to $\mathrm{PGL}(2, 2^e)$ ☆

William M. Kantor[a,b,*], Martin Kassabov[c]

[a] University of Oregon, Eugene, OR 97403, United States
[b] Northeastern University, Boston, MA 02115, United States
[c] Cornell University, Ithaca, NY 14853, United States

ARTICLE INFO

Article history:
Received 13 September 2013
Available online 8 September 2014
Communicated by Charles Leedham-Green

MSC:
primary 20D06
secondary 20B40, 68Q25

Keywords:
Computational group theory
Black box groups
Constructive recognition

ABSTRACT

A deterministic polynomial-time algorithm constructs an isomorphism between $\mathrm{PGL}(2, 2^e)$ and a black box group to which it is isomorphic.

Dedicated to the memory of Ákos Seress

1. Introduction

This note contains a polynomial-time algorithm for *recognizing* a black box group that is isomorphic to $\mathrm{PGL}(2, 2^e)$ by constructing such an isomorphism. The existence of a fast algorithm of this type has been open for a number of years. The standard

☆ This research was supported in part by NSA grant MDA-9049810020 and NSF grants DMS 0900932 and DMS 1303117.

* Corresponding author.
E-mail addresses: kantor@uoregon.edu (W.M. Kantor), kassabov@math.cornell.edu (M. Kassabov).

http://dx.doi.org/10.1016/j.jalgebra.2014.08.014
0021-8693/

way around this existence problem is based on a lovely idea of Leedham-Green [4,5], which avoids black box groups, instead focusing on groups represented on finite vector spaces and using a Discrete Log Oracle (this idea applies to all nonzero characteristics). Unfortunately, no black box version of our result is known in odd characteristic.

For the required background concerning black box groups, including the timing parameter and straight-line programs (SLPs), see [6,7]. Let μ be an upper bound on the time required for each group operation in $G = \langle \mathcal{S} \rangle$. We will assume that $|\mathcal{S}|$ is small and hence suppress it in our timing estimates (not suppressing it would multiply various times by $|\mathcal{S}|$).

The following appears to be the first deterministic polynomial-time black box recognition algorithm:

Theorem 1.

(i) *There is an $O(\mu e^3 \log e)$-time algorithm which, given a black box group $G = \langle \mathcal{S} \rangle$ isomorphic to $\mathrm{PGL}(2, q)$ with $q = 2^e$, finds 3-element generating sets $\hat{\mathcal{S}}$ and $\mathcal{S}^*$ of $\mathrm{PGL}(2, q)$ and G, respectively, and a bijection $\Psi : \hat{\mathcal{S}} \to \mathcal{S}^*$ inducing an isomorphism $\Psi : \mathrm{PGL}(2, q) \to G$.*

(ii) *In $O(\mu e^3)$ time an SLP of length $O(e)$ can be found from $\hat{\mathcal{S}}$ to any given element of $\mathrm{PGL}(2, q)$, and in $O(\mu e^3)$ time an SLP of length $O(e)$ can be found from $\mathcal{S}^*$ to any given element of G.*
In particular, the isomorphism Ψ is effective: the image of any given element of $\mathrm{PGL}(2, q)$ and the preimage of any given element of G can be found in $O(\mu e^3)$ time.

More precisely, *$\hat{\mathcal{S}}$ will essentially be $\left\{ \left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right), \mathrm{diag}(s, s^{-1}) \right\} \subset \mathrm{SL}(2, 2^e)$, where $\mathbb{F}_{2^e} = \mathbb{F}_2[s]$ and the minimal polynomial for s has been computed* (cf. **8** and **10**). (We identify $\mathrm{PGL}(2, 2^e)$ and $\mathrm{SL}(2, 2^e)$.) If desired, one can switch to $\mathbb{F}_2[t]$ where t is a root of another irreducible polynomial.

Our new and crucial observation is Proposition 4, which can be viewed as providing involution-producing formulas in a black box group isomorphic to $\mathrm{PGL}(2, 2^e)$ (cf. Remark 5). This is then used to produce first a Borel subgroup and a field $\mathfrak{F} \cong \mathbb{F}_q$, and then a group isomorphism. Also essential is Proposition 9, which recovers the entries of a matrix using the matrices for two given noncommuting involutions. An unexpected feature of this result is that any given element of G can be obtained quickly using our new generators (Theorem 1(ii)). A Monte Carlo algorithm related to the theorem appears in [2, Theorem 3.1].

Our field calculations all take place inside $\mathfrak{F}$, hence essentially "inside" G, which explains the timing in both parts of Theorem 1(ii). No Discrete Logs are involved, unless a user needs to describe the field using a generator of the multiplicative group. We have not tried to optimize the timing of our algorithm, for example by using fast multiplication in $\mathbb{F}_2[s]$ or fast computation of minimal polynomials, and we expect that careful optimization can reduce the time from essentially cubic to essentially quadratic in e.

2. Preliminaries

The following easy fact is critical:

Lemma 2. *If a, b, c, d are distinct points of the projective line over a field F, then there is one and only one involution in* $\mathrm{PGL}(2, F)$ *acting as* $(a, c)(b, d) \dots$.

Proof. By transitivity we may assume that the points are $\langle(1, 0)\rangle, \langle(1, 1)\rangle, \langle(0, 1)\rangle$, $\langle(s, 1)\rangle$, and then $\left(\begin{smallmatrix} 0 & 1 \\ s & 0 \end{smallmatrix}\right)$ behaves as required. □

Lemma 3. *Let $h \in G = \mathrm{PGL}(2, q)$ with $q > 2$ a power of a prime p not dividing $|h|$. Let $g \in G$ be such that $[h, h^g]^p \neq 1$. Then there is a unique involution $z \in G$ conjugating h to h^g.*

Proof. If h fixes two points a, b of the projective line, then the hypothesis $[h, h^g]^p \neq 1$ states that $|\{a, b, c, d\}| = 4$ for $\{c, d\} := \{a, b\}^g$. By the preceding lemma there is a unique involution $y \in G$ acting as $(a, c)(b, d) \dots$. Then $h^y \in \langle h^g \rangle$ since the stabilizer of c and d is cyclic, so that $h^y = (h^g)^\epsilon$ for $\epsilon = \pm 1$ since $|\mathrm{N}_G(\langle h \rangle) : \langle h \rangle| = 2$.

There is also an involution u acting as $(a, b)(c, d) \dots$, and hence commuting with y and satisfying $h^{yu} = (h^g)^{-\epsilon}$. Then y or yu is the unique involution that conjugates h to h^g.

If h does not fix any points, embed G into $\mathrm{PGL}(2, q^2)$ and let σ be the involutory field automorphism. Then h fixes two points of the larger projective line. We have already seen that there is a unique involution $z \in \mathrm{PGL}(2, q^2)$ such that $h^z = h^g$. Since z^σ also has this property, it follows that $z^\sigma = z$ is in G. □

Proposition 4. *Let $G = \mathrm{PGL}(2, q)$ with $q > 2$ even, let $2k + 1 = q^2 - 1$ denote the odd part of $|G|$, and let $1 \neq h \in G$ have odd order. For $g \in G$ with $[h, h^g] \neq 1$ either $[h, h^g]$ or $(hh^g)^k h$ is an involution.*

Proof. If $[h, h^g]$ is not an involution it has odd order. For z in the preceding lemma, $u := hz$ has odd order (as $u^2 = hh^z = hh^g \neq 1$). Then $(hh^g)^k h = (hh^z)^k h = (hzhz)^k hz \cdot z = (hz)^{2k+1} z = z$. □

Remark 5. This result deterministically produces many involutions in a black box group isomorphic to $G = \mathrm{PGL}(2, 2^e)$. The unavailability of such a deterministic or probabilistic result has long been the obstacle to a polynomial-time recognition algorithm for G. It has been folklore for several years that an involution u would lead to a Las Vegas algorithm based on the following idea: find $U := \mathrm{C}_G(u)$ using [3,1]; use a random element of U to produce a random element b of $\mathrm{N}_G(U)$ (cf. "lifting" below in **3**); turn U into a field generated by a single element corresponding to b; and finally make the algorithm Las Vegas by verifying a standard presentation of G.

We will need a standard elementary fact:

Lemma 6. *If z, z' are involutions in a group and $|zz'|$ divides $2k+1$, then $(zz')^{k+1}$ conjugates z to z'.*

3. The square root of a matrix

We will use square roots of 2×2 matrices over $F \cong \mathbb{F}_{2^e}$, $e > 1$. This already occurred in Proposition 4: $(hh^g)^{-k}$ is the square root of hh^g. Such square roots involve very elementary calculations (finiteness is not even required). We need to exclude the matrices having trace 0, which are precisely those of order dividing 2.

Lemma 7. *If $h = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ with $\mathrm{Tr}(h) = a + d \neq 0$, then*

$$\sqrt{h} = \frac{1}{\sqrt{a+d}} \begin{pmatrix} a+1 & b \\ c & d+1 \end{pmatrix}$$

is the unique matrix whose square is h.

Proof. Since $h^2 + \mathrm{Tr}(h)h + I = 0$ by Cayley–Hamilton, $\frac{1}{\sqrt{\mathrm{Tr}(h)}}(h+I)$ is a square root of h. It is unique: h has odd order > 1, so that any square root of h also has odd order and hence has to generate the same cyclic subgroup as h. □

For $u = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$, $g \in \mathrm{SL}(2, F)$, $|F| = q = 2^e$, and $q^2 - 1 = 2k + 1$ as in Proposition 4, write

$$\mathbb{B}(g) := \left(uu^g\right)^{k+1} g^{-1} \quad \text{if } \left[u, u^g\right] \neq 1. \tag{1}$$

Since $\sqrt{h} = h^{k+1}$ in the lemma, this definition of the partial function $\mathbb{B}$ is based on [3,1] and can be used for elements of either group in Theorem 1.

Lemma 8. *If $g = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ with $[u, u^g] \neq 1$ then $b \neq 0$ and $\mathbb{B}(g) = \begin{pmatrix} 1 & 0 \\ 1+\frac{a+d}{b} & 1 \end{pmatrix}$.*

Proof. Calculate $h := uu^g$ and use the preceding lemma for $\sqrt{h}$:

$$\mathbb{B}(g) = \begin{pmatrix} a & b \\ \frac{ab+a^2+1}{b} & a+b \end{pmatrix} \begin{pmatrix} d & b \\ c & a \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 1 + \frac{a+d}{b} & 1 \end{pmatrix}. \qquad \square$$

In particular, $\mathbb{B}(g) = u$ if and only if $g^2 = 1$. Let $\mathbb{B}_{21}(g)$ denote the 2, 1-entry of $\mathbb{B}(g)$. Starting with knowledge of the entries of u, we can use $\mathbb{B}_{21}$ to find the entries of most matrices:

Proposition 9. *Let $r = \begin{pmatrix} 0 & \lambda \\ \lambda^{-1} & 0 \end{pmatrix}$ with $\lambda \in F^*$, and let N be the normalizer of the maximal torus containing ru. If $g \in \mathrm{SL}(2, F) \backslash N$, then λ and the entries of g can be calculated (by* *formulas given below) using rational expressions in the square roots of the field elements* $\mathbb{B}_{21}(g_1 g g_2)$, $g_1, g_2 \in \{1, u, r\}^3$, *for which* $[u, u^{g_1 g g_2}] \neq 1$.

In particular, if $\mathrm{SL}(2, F) = \langle \mathcal{S} \rangle$ *then* F *is generated by the set*

$$\left\{\mathbb{B}_{21}(g_1 g g_2) \;\middle|\; g \in \mathcal{S} \cup \mathcal{S}^2,\ g_1, g_2 \in \{1, u, r\}^3,\ \textit{for which } \left[u, u^{g_1 g g_2}\right] \neq 1\right\}.$$

Proof. We need to determine the entries of $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. By the preceding lemma,

$$\mathbb{B}_{21}(g) = 1 + A := 1 + \frac{a+d}{b}, \qquad \mathbb{B}_{21}(gr) = 1 + B := 1 + \frac{\lambda^{-2} b + c}{a},$$
$$\mathbb{B}_{21}(rg) = 1 + C := 1 + \frac{\lambda^{-2} b + c}{d} \quad \text{and} \quad \mathbb{B}_{21}(rgr) = 1 + D := 1 + \frac{a+d}{\lambda^2 c}, \tag{2}$$

where $\mathbb{B}$ is not defined in those cases where the denominator is zero. Note that $\mathrm{Tr}(g) = \mathrm{Tr}(rgr) = a + d$ and $\mathrm{Tr}(rg) = \mathrm{Tr}(gr) = \lambda(\lambda^{-2} b + c)$.

If all of the elements A, B, C and D are defined and nonzero (which is equivalent to $abcd(a+d)(\lambda^{-2} b + c) \neq 0$), then the relation

$$ABCD\lambda^2 = (A + D)(B + C) \tag{3}$$

determines λ. Moreover, since $ad + bc = 1$,

$$a = \frac{ACD}{\Delta}, \qquad b = \frac{D(B+C)}{\Delta}, \qquad c = \frac{\lambda^{-2} A(B+C)}{\Delta} \quad \text{and} \quad d = \frac{ABD}{\Delta},$$

where $\Delta := \sqrt{\lambda^{-2} AD(B+C)(A+B+C+D)} \neq 0$.

If only three of A, B, C, D are defined, then g, gr, rg or rgr has 1, 2-entry 0, so consider the possibility that A is not defined but B, C and D are defined and nonzero (which is equivalent to $acd \neq 0$, $b = 0$ and $a \neq d$). Then

$$BCD\lambda^2 = B + C \tag{4}$$

determines λ. Moreover, since $ad = 1$,

$$a = \frac{C}{\Delta}, \qquad b = 0, \qquad c = \frac{BC}{\Delta} \quad \text{and} \quad d = \frac{B}{\Delta},$$

where $\Delta := \sqrt{BC} \neq 0$. Thus, we can recover λ and the entries of g. We still must consider the possibility that $acd \neq 0$, $b = 0$ and $a = d$. Then $ad + bc = 1$ implies that $a = d = 1$. If $c \neq 1, \lambda^{-2}$ then all entries of urg are nonzero and $\mathrm{Tr}(urg)\,\mathrm{Tr}(r \cdot urg) \neq 0$, so we can find λ and the entries of urg and hence of g by the previous paragraph. If $c = 1$ then $g = u \in N$. If $c = \lambda^{-2} \neq 1$ then $g' := urgu$ satisfies $\mathrm{Tr}(g')\,\mathrm{Tr}(rg') \neq 0$ and hence has been dealt with already.

If only two of the entries of g are not zero, then gu has only one zero entry, and we can recover the entries of g unless $g \in \{1, r\} \subset N$.

Finally, we need to deal with the case where A, B, C and D are all defined and at least one is 0, so that $abcd \neq 0$ and either $a = d$ or $b = \lambda^2 c$. Replacing g by gr if necessary, we may assume that $a = d$. Then $\mathrm{Tr}(gu) = b \neq 0$ and $\mathrm{Tr}(r \cdot gu) = \lambda(c + a) + \lambda^{-1} b$. If $\mathrm{Tr}(r \cdot gu) \neq 0$ then the previous cases determine λ and the entries of gu, hence also of g. On the other hand, if $\mathrm{Tr}(r \cdot gu) = 0$ then $a = c + \lambda^{-2} b$, which is precisely the condition that the involution g inverts ru and hence lies in N. (Note that it is not surprising that it is not possible to recover λ if $g \in N \backslash \{1, u\}$, since then $\mathbb{B}(g) \in N \cap \mathrm{C}_G(u) = \{1, u\}$ by (1).)

For the last statement of the proposition, note that the entries of $\mathcal{S}$ generate F. We have already computed the entries of every generator $s' \in \mathcal{S} \backslash N$, while we can obtain the entries of any $s \in \mathcal{S} \cap N$ from the entries of s' and ss'. (In fact we do not need $\mathcal{S}^2$, only $\mathcal{S} \cup s'\mathcal{S}$ for a single $s' \in \mathcal{S} \backslash N$.) □

Remark 10. The calculation of $\mathbb{B}(g)$ in Lemma 8 was unexpectedly simple. Although it will be less simple in larger groups having more elements of even order, such a calculation might occasionally be useful in order to speed up the recovery of the entries of a matrix from various values of $\mathbb{B}$ (as in Proposition 9 and **7** below).

4. Proof of Theorem 1

We are given a black box group $G = \langle \mathcal{S} \rangle \cong \mathrm{PGL}(2, q)$. We may assume that $q = 2^e > 2$. We proceed in several steps, each of which begins with a hint of its content.

1 *(Noncommuting involutions u, r).* If every element of $\mathcal{S}$ is an involution then two of them do not commute and we obtain noncommuting involutions u, r. If some $1 \neq h \in \mathcal{S}$ has odd order, then $\mathrm{C}_G(h)$ is cyclic and not normal in G, and hence $h^g \notin \mathrm{C}_G(h)$ for some $g \in \mathcal{S}$, so that Proposition 4 produces an involution u. Then some $s \in \mathcal{S}$ does not normalize $\mathrm{C}_G(u)$, so that $r := u^s \notin \mathrm{C}_G(u)$. (Time: $O(\mu e)$ using SLPs of length $O(e)$.)

2 *(The field: motivation).* In order to define a field we need to understand consequences of the assumed isomorphism $\hat{G} := \mathrm{PGL}(2, F) \to G$, where $F \cong \mathbb{F}_q$. (In **3** we will define a specific model of $\mathbb{F}_q$.)

Let $\hat{u} := \left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right)$, $\hat{U} := \mathrm{C}_{\hat{G}}(\hat{u})$, $\hat{B} := \mathrm{N}_{\hat{G}}(\hat{U})$ and $\hat{B}^* := \hat{B}/\hat{U}$. Writing matrices modulo scalar matrices, $\hat{B}^*$ consists of cosets $\tilde{s}\hat{U} = \left(\begin{smallmatrix} s & 0 \\ * & 1 \end{smallmatrix}\right)$ that can be viewed as field elements. For such cosets $\tilde{s}\hat{U}$ and $\tilde{t}\hat{U}$, their product in F^* corresponds to the group operation on $\hat{B}^*$, while field-addition occurs via $\hat{u}^{\tilde{s}+\tilde{t}} = \hat{u}^{\tilde{s}} \hat{u}^{\tilde{t}}$ when $\tilde{s}\hat{U} \neq \tilde{t}\hat{U}$.

Each $\hat{x} \in \hat{U}$ has the form $\hat{x} = \hat{X}(t) := \left(\begin{smallmatrix} 1 & 0 \\ t & 1 \end{smallmatrix}\right)$, $t \in F$. *If* $\hat{r} = \left(\begin{smallmatrix} 0 & \lambda \\ \lambda^{-1} & 0 \end{smallmatrix}\right)$ *with* $\lambda \in F^*$ (as in Proposition 9) *and* $t \neq 0$, *then* (using $2k + 1 = q^2 - 1$)

$$\mathfrak{b}(\hat{x}) := \left(\hat{u}\hat{u}^{\hat{r}}\right)^{k+1} \left(\hat{u}^{\hat{r}}\hat{x}\right)^{k+1} \hat{U} = \begin{pmatrix} \sqrt{t} & 0 \\ 0 & \sqrt{t}^{-1} \end{pmatrix} \hat{U} \in \hat{B}/\hat{U}, \tag{5}$$

which is independent of r. For, $\hat{u}\hat{u}^{\hat{r}}$ and $\hat{u}^{\hat{r}}\hat{x}$ have odd order (since $\hat{u}^{\hat{r}} \notin \hat{U}$). Then $(\hat{u}\hat{u}^{\hat{r}})^{k+1}(\hat{u}^{\hat{r}}\hat{x})^{k+1}$ conjugates $\hat{u}$ to $\hat{x}$ (by Lemma 6), hence normalizes $\hat{U}$ and so lies in $\hat{B}$. A simple matrix calculation in $\hat{B}$ produces all matrices conjugating $\hat{u}$ to $\hat{x}$, as stated in (5).

3 *(The field $\mathfrak{F}$).* We define what amounts to a "black box field" $\mathfrak{F}$ obtained from the involution u. Define $U := \mathrm{C}_G(u)$, $B := \mathrm{N}_G(U)$ and $B^* := B/U$. Then $\mathfrak{F}$ arises from the pair consisting of U and B^* as additive and multiplicative groups, respectively; as a set $\mathfrak{F}$ will be $B^* \cup \{0\}$ with 0 treated in the obvious ways. We do not yet have these as constructed groups.

- If $g \in G$, then $g \in U \Leftrightarrow [u, g] = 1$, and $g \in B \Leftrightarrow [u, u^g] = 1$.
- If $t_1U, t_2U \in B^*$ then $t_1U = t_2U \Leftrightarrow (t_1t_2^{-1})^2 = 1$.
- *Lifting elements of U to B/U*: If $1 \neq x \in U$, then uu^r and u^rx have odd order (since $u^r \notin U$), so that (as in (5))

$$\mathfrak{b}(x) := \left(uu^r\right)^{k+1}\left(u^rx\right)^{k+1}U \quad \text{conjugates } u \text{ to } x. \tag{6}$$

  Then $\mathfrak{b}(x)$ normalizes U and so lies in B/U. (Time: $O(\mu e)$ using SLPs of length $O(e)$.)
- *Labeling $U\backslash\{1\}$ using $\mathfrak{F}^* = B^*$*: If $tU \in B^*$ then $X(tU) := u^t \in U\backslash\{1\}$ defines the inverse of the map $\mathfrak{b} : U\backslash\{1\} \to B^*$ in (6). Also let $X(0) := 1$.
- *Field multiplication*: This is inherited from B/U and hence from G. (Time: μ.)
- *Field addition*: $t_1U + t_2U = tU = \mathfrak{b}(u^{t_1}u^{t_2})$ for distinct t_1U, t_2U, so that $u^t = u^{t_1}u^{t_2}$. (Time: $O(\mu e)$.)

All of the above imitated **2**. We emphasize that the field depends entirely on u (which uniquely determines U and B): it does not depend on r, which was used only to obtain elements of $B^* = B/U$. We also emphasize that nonzero field elements are cosets: whenever we write a nonzero field element we are implicitly also writing (and storing) a coset representative.

Field-theoretic calculations are postponed to Section 5.

4 *(Generators of $\mathfrak{F}$). The elements $\mathfrak{b}(\mathbb{B}(g_1gg_2))$ with $g \in \mathcal{S} \cup \mathcal{S}^2, g_1, g_2 \in \{1, u, r\}^3$ and $[u, u^{g_1gg_2}] \neq 1$, generate $\mathfrak{F} \cong \mathbb{F}_q$.* Here the partial function $\mathbb{B}$ is defined as in (1), $\mathfrak{b}$ was defined in (6), and "generate" means that the stated elements lie in no proper subfield. (Time: $O(\mu e)$ to find these elements using SLPs of length $O(e)$.)

For, there is an isomorphism $\Psi : \mathrm{PGL}(2, \mathfrak{F}) \to G$. By a basis change of $\mathfrak{F}^2$ we may assume that Ψ sends $\hat{u} := \left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right) \mapsto u$ and $\hat{r} := \left(\begin{smallmatrix} 0 & \lambda^{-1} \\ \lambda & 0 \end{smallmatrix}\right) \mapsto r$ for some uniquely determined $\lambda \in \mathfrak{F}$. Since the preimages of $\mathcal{S}$ cannot all lie in a proper subfield of $\mathfrak{F}$, the elements $\mathfrak{b}(\mathbb{B}(g_1gg_1))$ arising from the last part of Proposition 9 do not lie in a proper subfield of $\mathfrak{F}$.

Remark 11. The isomorphism Ψ sends $\hat{u} \mapsto u$ and induces an isomorphism from the usual Borel subgroup $\hat{B}$ of $\hat{G}$ to B. The latter isomorphism is uniquely determined up to conjugation and field automorphisms. Hence, we may assume in our arguments that $\hat{X}(t) \mapsto X(t)$. (We emphasize *in our arguments* since we have yet to construct such an isomorphism, as required in the theorem.)

We will abbreviate $\Psi' := \Psi^{-1}$ and $\hat{g} := \Psi'(g)$.

Remark 12 *(The projective line).* We also obtain the projective line $\mathfrak{P} = \mathfrak{F} \cup \{\infty\}$. The fact that we can obtain the action of any $g \in G$ on it motivated parts of this paper. It suffices to show how to compute 0^g and ∞^g. For, any $x \in \mathfrak{F}^*$ is represented by an element $b \in B$, so that $x^{b^{-1}u} = 0$ by definition and hence $x^g = 0^{ubg}$.

By the proof of Proposition 9, if $\mathbb{B}(g)$, $\mathbb{B}(gr)$, $\mathbb{B}(rg)$ and $\mathbb{B}(rgr)$ are all defined and are not equal to u (so that $g \notin N$), then

$$\begin{aligned} 0^g &= \mathfrak{b}\big(u\mathbb{B}(g)\big)\mathfrak{b}\big(u\mathbb{B}(gr)\big)\big(\mathfrak{b}\big(\mathbb{B}(g)\mathbb{B}(rgr)\big)\big)^{-1}, \\ \infty^g &= \mathfrak{b}\big(u\mathbb{B}(g)\big)\mathfrak{b}\big(u\mathbb{B}(gr)\big)\big(\mathfrak{b}\big(\mathbb{B}(rg)\mathbb{B}(gr)\big)\big)^{-1}. \end{aligned}$$

5 *(*PGL(2, 2) *and r').* By the proof of Proposition 9, there is an element $g \in \{1, r, u\}^3\mathcal{S}\{1, r, u\}^3$ such that three or four of the elements $\mathbb{B}(g)$, $\mathbb{B}(gr)$, $\mathbb{B}(rg)$, $\mathbb{B}(rgr)$ are defined, and each of those is not equal to u.

If all four behave this way, calculate

$$\tau := \mathfrak{b}\big(\mathbb{B}(g)\mathbb{B}(rgr)\big)\mathfrak{b}\big(\mathbb{B}(rg)\mathbb{B}(gr)\big)\big(\mathfrak{b}\big(\mathbb{B}(g)u\big)\mathfrak{b}\big(\mathbb{B}(rg)u\big)\mathfrak{b}\big(\mathbb{B}(gr)u\big)\mathfrak{b}\big(\mathbb{B}(rgr)u\big)\big)^{-1}.$$

Let $\tau = tU$, $t \in B$. Then uu^{tru} has order 3 and $r' := u^{tru}$ is an involution, so that $\langle u, r'\rangle \cong \mathrm{PGL}(2,2)$. (Time: $O(\mu e)$.)

For, the definitions of $\mathbb{B}$ and $\mathfrak{b}$ (in (1), (5) and (6)) imply that

$$\hat{u}^{\Psi'(\mathfrak{b}(\mathbb{B}(g)\mathbb{B}(rgr)))} = \Psi'\big(u^{\mathfrak{b}(\mathbb{B}(g)\mathbb{B}(rgr))}\big) = \Psi'\big(\mathbb{B}(g)\mathbb{B}(rgr)\big) = \mathbb{B}(\hat{g})\mathbb{B}(\hat{r}\hat{g}\hat{r}),$$

and hence that $\Psi'(\mathfrak{b}(\mathbb{B}(g)\mathbb{B}(rgr))) = \mathfrak{b}(\mathbb{B}(\hat{g})\mathbb{B}(\hat{r}\hat{g}\hat{r}))$; and similarly for the other terms in the definition of τ. Using (2), (3) and (5), we obtain

$$\begin{aligned} \Psi'(tU) &= \mathfrak{b}\begin{pmatrix} 1 & 0 \\ A+D & 1 \end{pmatrix}\mathfrak{b}\begin{pmatrix} 1 & 0 \\ B+C & 1 \end{pmatrix} \\ &\quad\times \left(\mathfrak{b}\begin{pmatrix} 1 & 0 \\ A & 1 \end{pmatrix}\mathfrak{b}\begin{pmatrix} 1 & 0 \\ B & 1 \end{pmatrix}\mathfrak{b}\begin{pmatrix} 1 & 0 \\ C & 1 \end{pmatrix}\mathfrak{b}\begin{pmatrix} 1 & 0 \\ D & 1 \end{pmatrix}\right)^{-1} \\ &= \begin{pmatrix} \sqrt{\frac{(A+D)(B+C)}{ABCD}} & 0 \\ 0 & \sqrt{\frac{(A+D)(B+C)}{ABCD}}^{-1} \end{pmatrix}\hat{U} = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}\hat{U}. \end{aligned}$$

Since $\hat{r} = \begin{pmatrix} 0 & \lambda^{-1} \\ \lambda & 0 \end{pmatrix}$, it follows that $\widehat{r'} = \Psi'(r') = \hat{u}^{\Psi'(tU)\hat{r}\hat{u}} = \hat{u}^{\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)\left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right)} = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, so that $\langle u, r'\rangle \cong \langle \hat{u}, \widehat{r'}\rangle \cong \mathrm{PGL}(2,2)$.

Similarly, assume that exactly three of the elements $\mathbb{B}(g)$, $\mathbb{B}(gr)$, $\mathbb{B}(rg)$, $\mathbb{B}(rgr)$ are defined and not equal to u. *If $\mathbb{B}(g)$ is not defined, then (4) can be used as above in place of (3), this time together with the field element*

$$\tau := \mathfrak{b}\big(\mathbb{B}(rg)\mathbb{B}(gr)\big)\Big(\mathfrak{b}\big(\mathbb{B}(rg)u\big)\mathfrak{b}\big(\mathbb{B}(gr)u\big)\mathfrak{b}\big(\mathbb{B}(rgr)u\big)\Big)^{-1},$$

letting $\tau = tU$ and $r' = u^{tru}$ in order to obtain $\langle u, r'\rangle \cong \mathrm{PGL}(2,2)$. The remaining cases are handled by replacing g by rg, gr or rgr.

Now we can introduce some standard elements of $\hat{G}$ and G for all $t \in \mathfrak{F}^*$:

$$\hat{n}(t) := \hat{X}(t)\hat{X}\big(t^{-1}\big)^{\widehat{r'}}\hat{X}(t) \quad \text{and} \quad \hat{h}(t) := \hat{n}(t)\hat{n}(1),$$
$$n(t) := X(t)X\big(t^{-1}\big)^{r'}X(t) \quad \text{and} \quad h(t) := n(t)n(1).$$

Then Ψ sends $\widehat{r'} \to r'$ and $\hat{h}(t) \mapsto h(t)$.

6. Proof of Theorem 1(i). In **10** we will express $\mathfrak{F}$ in the form $\mathbb{F}_2[s]$. Since the isomorphism Ψ sends $\hat{X}(1) \mapsto X(1)$, $\widehat{r'} \mapsto r'$ and $\hat{h}(s) \mapsto h(s)$, the map $\hat{\mathcal{S}} := \{\hat{X}(1), \widehat{r'}, \hat{h}(s)\} \to \mathcal{S}^* := \{X(1), r', h(s)\}$ determines Ψ.

The stated time is dominated by **10**. $\square$

7. Proof of Theorem 1(ii). In order to handle elements of $\hat{G}$, first consider $\hat{X}(t)$, $t \in \mathfrak{F}$. In $O(\mu e^3)$ time write t as $\sum_0^{e-1} a_i s^i$ with $a_i \in \mathbb{F}_2$ (cf. **9**), and then find an SLP from $\hat{\mathcal{S}}$ to $\hat{X}(t) = \prod_i \hat{X}(s^i)^{a_i}$ of length $O(e)$. Now it is easy to use the definitions at the end of **5**, together with straightforward row reduction, to find an SLP of length $O(e)$ from $\hat{\mathcal{S}}$ to any given element in $\hat{G} = \hat{B} \cup \hat{B}\hat{r}\hat{B}$, where $\hat{B} = \hat{U}\{\hat{h}(t) \mid t \in \mathfrak{F}^*\}$. This takes $O(\mu e^3)$ time (dominated by **9**).

Now consider $g \in G$. We must find an SLP from $\mathcal{S}^*$ to g. Assume that $g \notin \mathrm{N}_G(\langle u, r\rangle)$. For each $g_1, g_2 \in \{1, u, r\}^3$ such that $[u, u^{g_1 g g_2}] \neq 1$, find $\mathbb{B}(g_1 g g_2) \in U$ and $\mathbb{B}(\hat{g}_1\hat{g}\hat{g}_2) = \Psi'(\mathbb{B}(g_1 g g_2)) \in \hat{U}$, and then use Proposition 9 to find the entries of the matrix $\hat{g} = \Psi'(g)$. Use the preceding paragraph to find an SLP from $\hat{\mathcal{S}}$ to $\hat{g}$ of length $O(e)$. Then the Ψ-image of that SLP is the required SLP from $\mathcal{S}^*$ to g. Finding $\hat{g}$ takes $O(\mu e)$ time, finding an SLP to $\hat{g}$ takes $O(\mu e^3)$ time, and applying Ψ to the members of the SLP takes $O(e)$ time.

If $g \in \mathrm{N}_G(\langle u, r\rangle)$ then $g = gh(s)^{-1} \cdot h(s)$ with $h(s) \in \mathcal{S}^*$, and we can find an SLP from $\mathcal{S}^*$ to $gh(s)^{-1} \notin \mathrm{N}_G(\langle u, r\rangle)$ in the required time. $\square$

5. Field computations

The proof of Theorem 1 depends on the black box field $\mathfrak{F}$ in **3** that is a model of the field $\mathbb{F}_q$ inside the black box group G. Using it we will emulate some classical field-theoretic algorithms. Recall that each addition in $\mathfrak{F}$ takes $O(\mu e)$ time while each multiplication takes $O(\mu)$ time.

8 *(Preliminaries).*

- (i) Given $s \in \mathfrak{F}^*$, find m such that $\mathbb{F}_2[s] \cong \mathbb{F}_{2^m}$ in $O(\mu e)$ time: find the first $m \in \{1, \dots, e\}$ such that $s^{2^m} = s$. In particular, we can test whether $\mathbb{F}_2[s] = \mathfrak{F}$.
- (ii) For s and m in (i), the minimal polynomial of s over $\mathbb{F}_2$ is $\prod_0^{m-1}(x - s^{2^i})$, found in $O(\mu e^3)$ time. This yields an isomorphism $\mathbb{F}_2[x]/(f) \to \mathbb{F}_2[s]$ induced by $x + (f) \mapsto s$.
- (iii) The trace map $\mathrm{Tr}: \mathfrak{F} \to \mathbb{F}_2$, defined by $\mathrm{Tr}(t) := \sum_0^{e-1} t^{2^i}$, is calculated in $O(\mu e^2)$ time using $u^{\mathrm{Tr}(t)} = \prod_0^{e-1} u^{t^{2^i}}$.
- (iv) In **9** we will use linear equations over $\mathbb{F}_2$. While this $\mathbb{F}_2$ is contained in $\mathfrak{F}$, with its rather slow field operations, we can work with a standard model of $\mathbb{F}_2$ with much faster field operations.

9 *(Linear algebra).* Let $s \in \mathfrak{F}$ with $\mathfrak{F} = \mathbb{F}_2[s]$. *In $O(\mu e^3)$ time, given $t \in \mathfrak{F}$ we can find the unique $x_i \in \mathbb{F}_2$ such that $t = \sum_0^{e-1} x_i s^i$.*

First find all $2e - 1$ traces $a_{ij} := \mathrm{Tr}(s^{i+j}) \in \mathbb{F}_2$ and all $\mathrm{Tr}(ts^j) \in \mathbb{F}_2$, $0 \le i, j < e$, in $O(\mu e^3)$ time. Since $\mathrm{Tr}(ts^j) = \sum_{i=0}^{e-1} x_i \mathrm{Tr}(s^{i+j})$, we obtain

$$\mathrm{Tr}\bigl(ts^j\bigr) = \sum_{i=0}^{e-1} a_{ij} x_i, \quad 0 \le j < e.$$

In $O(e^3)$ time solve these linear equations over $\mathbb{F}_2$.

10 *(Field generator). In $O(\mu e^3 \log e)$ time we can find s such that $\mathfrak{F} = \mathbb{F}_2[s]$.*

By **4**, $\mathbb{F}_{2^e} \cong \mathfrak{F} = \mathbb{F}_2[\alpha_1, \dots, \alpha_n]$ with $n = O(1)$ (we are ignoring $|\mathcal{S}|$).

1. Factor $e = \prod_i q_i$ into powers $q_i = p_i^{k_i}$ of $O(\log e)$ different primes p_i. (Time: $O(e)$.)
2. Find the degree of each α_j over $\mathbb{F}_2$. (Time: $O(\mu e^2)$ as in **8**(i).)
3. For each i find α_{j_i} such that q_i divides $m_i := [\mathbb{F}_2[\alpha_{j_i}] : \mathbb{F}_2]$ (this exists since the α_j generate $\mathfrak{F}$).
4. For each i compute the polynomial $f_i := \prod_{t=0}^{(m_i/q_i)-1}(x - \alpha_{j_i}^{2^{tq_i}})$, whose coefficients generate a field $\mathbb{F}_{2^{q_i}}$. (Time: $O(\mu e^3 \log e)$.)
5. For each i find a coefficient c_i of f_i such that $\mathbb{F}_{2^{q_i}} = \mathbb{F}_2[c_i]$ (this exists since q_i is a prime power). (Time: $O(\mu e^2 \log e)$, testing all coefficients using **8**(i).)
6. Output $s := \prod_i c_i$. Then $\mathbb{F}_q = \mathbb{F}_2[s]$ since the groups $\mathbb{F}_2[c_i]^*$ have pairwise relatively prime orders, so that each c_i is a power of s. (Time: $O(\mu \log e)$.)

At this point we have completed the requirements made after the statement of Theorem 1: we have obtained $\mathfrak{F}$ as $\mathbb{F}_2[s]$, and we can find the minimal polynomial of s (cf. **8**(ii)).

6. Odd characteristic

We conclude with remarks about a case we have not been able to handle: $G = \mathrm{PSL}(2,q)$ with $q = p^e$ for an odd prime p. Since about a quarter of all group elements have even order, it is easy to (probabilistically) find an involution. The problem is to find an element of order p.

There is an analogue of Proposition 4 that might not be entirely useless, though we do not see how to use it. That result can be imitated given an element h such that $|hh^g|$ is a factor of the odd integer $(q \pm 1)/2$; if the involution t in Lemma 3 is in $\mathrm{PSL}(2,q)$, then we obtain that involution as before.

There is an analogue of Lemma 2 involving elements of order p acting as $g = (a, c, \ldots)(b, d, \ldots)\ldots$: given distinct a, b, for about half of all pairs c, d there is such an element g (in fact, two of them), and then $G_{ab}^g = G_{cd}$. However, we do not see how to use this fact.

Acknowledgments

We are grateful to Peter Brooksbank for helpful suggestions and for implementing our algorithm in Magma. We thank the referee for his statement of Lemma 7.

References

[1] A. Borovik, Centralisers of involutions in black box groups, in: R. Gilman, et al. (Eds.), Computational and Statistical Group Theory, in: Contemp. Math., vol. 298, AMS, Providence, RI, 2002, pp. 7–20.
[2] A. Borovik, Ş. Yalçınkaya, Fifty shades of black, preprint, arXiv:1308.2487v1.
[3] J.N. Bray, An improved method for generating the centralizer of an involution, Arch. Math. 74 (2000) 241–245.
[4] M.D.E. Conder, C.R. Leedham-Green, Fast recognition of classical groups over large fields, in: W.M. Kantor, Á. Seress (Eds.), Groups and Computation III, in: Ohio State Univ. Math. Res. Inst. Publ., vol. 8, de Gruyter, Berlin–New York, 2001, pp. 113–121.
[5] M.D.E. Conder, C.R. Leedham-Green, E.A. O'Brien, Constructive recognition of $\mathrm{PSL}(2,q)$, Trans. Amer. Math. Soc. 358 (2006) 1203–1221.
[6] W.M. Kantor, Á. Seress, Black box classical groups, Mem. Amer. Math. Soc. 149 (2001), No. 708.
[7] Á. Seress, Permutation Group Algorithms, Cambridge U. Press, Cambridge, 2002.