The essential dimension
of the normalizer of a maximal torus
in the projective linear group

Aurel Meyer and Zinovy Reichstein

The essential dimension
of the normalizer of a maximal torus
in the projective linear group

Aurel Meyer and Zinovy Reichstein

Let p be a prime, k a field of characteristic $\neq p$ and N the normalizer of the maximal torus in the projective linear group PGL_n. We compute the exact value of the essential dimension $\mathrm{ed}_k(N; p)$ of N at p for every $n \geq 1$.

1. Introduction

Let k be a field, Fields$_k$ the category of field extensions K/k, and F a covariant functor from Fields$_k$ into the category of sets. As usual, for a field extension L/K, we denote the image of $a \in F(K)$ under the natural map $F(K) \to F(L)$ by a_L.

Given a field extension L/k, an object $a \in F(L)$ is said to *descend* to an intermediate field $k \subseteq K \subseteq L$ if a is in the image of the induced map $F(K) \to F(L)$. The *essential dimension* $\mathrm{ed}(a)$ of $a \in F(L)$ is the minimum of the transcendence degrees $\mathrm{trdeg}_k(K)$ taken over all fields $k \subseteq K \subseteq L$ such that a descends to K. The essential dimension $\mathrm{ed}(a; p)$ of a at a prime integer p is the minimum of $\mathrm{ed}(a_{L'})$, taken over all finite field extensions L'/L such that the degree $[L' : L]$ is prime to p.

The essential dimension $\mathrm{ed}(F)$ of the functor F (respectively, the essential dimension $\mathrm{ed}(F; p)$ of F at a prime p) is the supremum of $\mathrm{ed}(a)$ (respectively, of $\mathrm{ed}(a; p)$) taken over all $a \in F(L)$ and over all field extensions L/k. Informally speaking, the essential dimension of $a \in F(L)$ can be thought of as the minimal number of parameters one needs to define a, and $\mathrm{ed}(F)$ as the minimal number of parameters required to define any object in F.

An important example is the Galois cohomology functor $F_G = H^1(*, G)$ sending a field K/k to the set $H^1(K, G)$ of isomorphism classes of G-torsors over

Spec(K), in the fppf topology. Here G is an algebraic group defined over k. The essential dimension of this functor is a numerical invariant of G, which, informally speaking, measures the complexity of G-torsors over fields. This number is usually denoted by $\mathrm{ed}_k(G)$ or, if k is fixed throughout, simply by $\mathrm{ed}(G)$. The notion of essential dimension was originally introduced and has since been extensively studied in this context; see for example [Buhler and Reichstein 1997; Reichstein 2000; Reichstein and Youssin 2000; Lemire 2004; Chernousov and Serre 2006]. The theory of essential dimension of algebraic groups may be viewed as a natural extension of the theory of *special groups* initiated in [Serre 1958]. Over an algebraically closed field k special groups are precisely those of essential dimension 0, these groups were classified in [Grothendieck 1958]. The more general definition of essential dimension for a covariant functor given above is due to Merkurjev [Berhuy and Favi 2003; Merkurjev 2007].

The purpose of this paper is to compute the relative essential dimension $\mathrm{ed}(N; p)$, where N is the normalizer of the (split) maximal torus in the projective linear group PGL_n. Before proceeding to state our main result, we would like to explain why we are interested in the essential dimension of N.

We begin by recalling that elements of $H^1(K, G)$ can often be naturally identified with K-forms of a single "split" algebraic object over k. Here by an algebraic object we mean a tensor t defined on a finite-dimensional k-vector space V; the group $G \subset \mathrm{GL}(V)$ then naturally arises as the automorphism group of t [Serre 1997, Chapter III]. Two examples will be of primary interest in the sequel:

$$H^1(*, \mathrm{PGL}_n) : K \mapsto \left\{ \begin{array}{c} \text{degree } n \text{ central simple algebras } A/K, \\ \text{up to } K\text{-isomorphism} \end{array} \right\} \qquad (1)$$

and

$$H^1(*, N) : K \mapsto \left\{ K\text{-isomorphism classes of pairs } (A, L) \right\}, \qquad (2)$$

where K is a field extension of k, A is a degree n central simple algebra over K, L is a maximal étale subalgebra of A, and N is the normalizer of a split maximal torus in PGL_n, as above. For the functor (1) the split central simple k-algebra of degree n is M_n, its automorphism group is PGL_n. Similarly, in the case of the functor (2) the split pair (A, L) is $(\mathrm{M}_n, \mathrm{Diag}_n)$, where Diag_n denotes the subalgebra of diagonal matrices in $\mathrm{M}_n(k)$. The automorphism group of this split pair is N.

Computing the essential dimension of the projective linear group PGL_n, or equivalently, of the functor (1), is a fundamental problem in the theory of central simple algebras. To the best of our knowledge, it was first raised by C. Procesi, who showed (using different terminology) that $\mathrm{ed}(\mathrm{PGL}_n) \leq n^2$ [Procesi 1967, Theorem 2.1]. This problem and the related question of computing the relative essential dimension $\mathrm{ed}(\mathrm{PGL}_n; p)$ at a prime p remain largely open. The best currently

known lower bound [Reichstein 1999, Theorem 16.1(b); Reichstein and Youssin 2000, Theorem 8.6] is

$$\mathrm{ed}(\mathrm{PGL}_{p^r};\, p) \geq 2r,$$

and it falls far below the best known upper bound [Lorenz and Reichstein 2000; Lorenz et al. 2003, Theorem 1.1; Lemire 2004, Proposition 1.6; Favi and Florence 2008], given by

$$\mathrm{ed}(\mathrm{PGL}_n) \leq \begin{cases} \frac{1}{2}(n-1)(n-2) & \text{for every odd } n \geq 5, \\ n^2 - 3n + 1 & \text{for every } n \geq 4. \end{cases} \tag{3}$$

We remark that the primary decomposition theorem reduces the computation of $\mathrm{ed}(\mathrm{PGL}_n;\, p)$ to the case where n is a power of p. That is, if $n = p_1^{r_1} \ldots p_s^{r_s}$ then $\mathrm{ed}(\mathrm{PGL}_n;\, p_i) = \mathrm{ed}(\mathrm{PGL}_{p_i^{r_i}};\, p_i)$. The computation of $\mathrm{ed}(\mathrm{PGL}_n)$ also partially reduces to the prime power case, because

$$\mathrm{ed}(\mathrm{PGL}_{p_i^{r_i}}) \leq \mathrm{ed}(\mathrm{PGL}_n) \leq \mathrm{ed}(\mathrm{PGL}_{p_1^{r_1}}) + \cdots + \mathrm{ed}(\mathrm{PGL}_{p_s^{r_s}})$$

for every $i = 1, \ldots, s$ [Reichstein 2000, Proposition 9.8].

Note that the proofs of the upper bounds (3) are not based on a direct analysis of the functor $H^1(*, \mathrm{PGL}_n)$. Instead, one works with the related functor $H^1(*, N)$ of (2). This functor is often more accessible than $H^1(*, \mathrm{PGL}_n)$ because many of the standard constructions in the theory of central simple algebras depend on the choice of a maximal subfield L in a given central simple algebra A/K. Projecting a pair (A, L) to the first component, we obtain a surjective morphism of functors $H^1(*, N) \to H^1(*, \mathrm{PGL}_n)$, [Rowen 1980, Corollary 3.1.11]. The surjectivity of this morphism leads to the inequalities

$$\mathrm{ed}(N) \geq \mathrm{ed}(\mathrm{PGL}_n) \quad \text{and} \quad \mathrm{ed}(N;\, p) \geq \mathrm{ed}(\mathrm{PGL}_n;\, p); \tag{4}$$

see [Merkurjev 2007, Proposition 1.3], [Berhuy and Favi 2003, Lemma 1.9] or [Reichstein 2000, Proposition 4.3].

The inequalities (3) were, in fact, proved as upper bounds on $\mathrm{ed}(N)$ [Lorenz et al. 2003; Lemire 2004]. It is thus natural to try to determine the exact values of $\mathrm{ed}(N)$ and $\mathrm{ed}(N;\, p)$. In addition to being of independent interest, these numbers represent a limitation on the techniques used in [Lorenz et al. 2003] and [Lemire 2004]. This brings us to the main result of this paper.

Theorem 1.1. *Let N the normalizer of a maximal torus in the projective linear group PGL_n defined over a field k with $\mathrm{char}(k) \neq p$. Then:*

(a) $\mathrm{ed}_k(N;\, p) = [n/p]$, *if n is not divisible by p.*

(b) $\mathrm{ed}_k(N;\, p) = 2$, *if $n = p$.*

(c) $\mathrm{ed}_k(N;\, p) = n^2/p - n + 1$, *if $n = p^r$ for some $r \geq 2$.*

(d) $\mathrm{ed}_k(N;\, p) = p^e(n - p^e) - n + 1$, *in all other cases.*

Here $[n/p]$ is the integer part of n/p and p^e is the highest power of p dividing n.

In each part we will prove an upper bound and a lower bound on $\mathrm{ed}(N)$ separately. We do not have an a priori reason why the two should match, thus yielding an exact value of $\mathrm{ed}(N; p)$; the fact that this happens may be viewed as a lucky coincidence. We also remark that our proof of the upper bounds on $\mathrm{ed}_k(N; p)$ in part (c) and (d) does not use the assumption that $\mathrm{char}(k) \neq p$. These bounds are valid for every base field k.

As we mentioned above, the computation of $\mathrm{ed}(\mathrm{PGL}_n; p)$ reduces to the case where n is a power of p. A quick glance at the statement of Theorem 1.1 shows that the computation of $\mathrm{ed}(N; p)$ does not. On the other hand, the proof of part (c), where $n = p^r$ and $r \geq 2$, requires the most intricate arguments. Another reason for our special interest in part (c) is that it leads to a new upper bound on $\mathrm{ed}(\mathrm{PGL}_n; p)$. More precisely, combining the upper bound in part (c) with (4), and remembering that the upper bound in part (c) is valid for any the ground field k, we obtain the following inequality.

Corollary 1.2. *Let $n = p^r$ be a prime power. Then*

$$\mathrm{ed}_k(\mathrm{PGL}_n; p) \leq p^{2r-1} - p^r + 1$$

for any field k and for any $r \geq 2$. □

Corollary 1.2 fails for $r = 1$ because

$$\mathrm{ed}_k(\mathrm{PGL}_p; p) = 2; \tag{5}$$

see [Reichstein 2000, Corollary 5.7] or [Reichstein and Youssin 2000, Lemma 8.5.7]. For $r = 2$, Corollary 1.2 is valid but is not optimal. Indeed, in this case L. H. Rowen and D. J. Saltman showed that, after a prime-to-p extension L/K, every degree p^2 central simple algebra A/K becomes a $(\mathbb{Z}/p\mathbb{Z})^2$-crossed product [Rowen and Saltman 1992, Corollary 1.3]. The upper bound on the essential dimension of a crossed product given by [Lorenz et al. 2003, Corollary 3.10] then yields the inequality $\mathrm{ed}(\mathrm{PGL}_{p^2}; p) \leq p^2 + 1$, which is stronger than Corollary 1.2 for any $p \geq 3$. Merkurjev [2008] recently showed that in fact, $\mathrm{ed}_k(\mathrm{PGL}_{p^2}; p) = p^2 + 1$ for any field k of characteristic different from p. For $r \geq 3$ Corollary 1.2 gives the best currently known upper bound on $\mathrm{ed}(\mathrm{PGL}_{p^r}; p)$.

We remark that the inequalities of (4) have counterparts for algebraic groups other than PGL_n. Indeed, if G is a linear group defined over k, C is a Cartan subgroup of G and $N(C)$ is the normalizer of C then by a theorem of T. Springer the natural map $H^1(K, N(C)) \to H^1(K, G)$ is surjective for every perfect field extension K/k [Serre 1997, III.4.3, Lemma 6]. Consequently, $\mathrm{ed}_k(N(C)) \geq \mathrm{ed}_k(G)$ if $\mathrm{char}(k) = 0$ and $\mathrm{ed}_k(N(C); p) \geq \mathrm{ed}_k(G; p)$ if $\mathrm{char}(k) \neq p$; compare [Reichstein

2000, Proposition 4.3]. It would thus be of interest to prove an analogue of Theorem 1.1 in the more general setting, where N is the normalizer of a split maximal torus in an arbitrary simple (or semisimple) linear algebraic group G. The new technical difficulty one encounters in this more general setting is that the natural sequence

$$1 \to T \to N \to W \to 1,$$

may not split. Here T is a split maximal torus and $W = N/T$ is the Weyl group of G. The fact that this sequence splits for $G = \mathrm{PGL}_n$ is an important ingredient in our proof of the upper bound on $\mathrm{ed}(N; p)$.

A key ingredient in our proofs of the lower bounds in Theorem 1.1(c) and (d) is a recent theorem of Karpenko and Merkurjev [2008] on the essential dimension of a p-group, stated as Theorem 7.1 below. To the best of our knowledge, these lower bounds were not accessible by previous techniques. Corollary 1.2 and the other parts of Theorem 1.1 do not rely on the Karpenko–Merkurjev theorem.

2. A general strategy

Let G be an algebraic group defined over a field k. Recall that the action of G on an algebraic variety X defined over k is generically free if the stabilizer subgroup $\mathrm{Stab}_G(x)$ is trivial for $x \in X(\bar{k})$ in general position.

Remark 2.1. If G is a finite constant group and X is irreducible then the G-action on X is generically free if and only if it is faithful.

Indeed, the "only if" is obvious. Conversely, if the G-action on X is faithful then $\mathrm{Stab}_G(x) = \{1\}$ for any x outside of the closed subvariety $\bigcup_{1 \neq g \in G} X^{\langle g \rangle}$, whose dimension is at most $\dim X - 1$. □

Remark 2.2. Suppose k'/k is a field extension of degree prime to p. Then the essential dimension at p does not change if we replace k by k' [Merkurjev 2007, Proposition 1.5(2)]. This happens in particular if $\mathrm{char}(k) \neq p$ and k' is obtained from k by adjoining a primitive p-th root of unity. Thus in the course of proving Theorem 1.1 we may assume without loss of generality that k contains a primitive p-th root of unity.

In the sequel we will repeatedly encounter the following situation. Suppose we want to show that

$$\mathrm{ed}_k(G) = \mathrm{ed}_k(G; p) = d, \qquad (6)$$

where G is a linear algebraic group defined over k. All such assertions will be proved in two steps:

(i) Construct a generically free linear representation of G over k of dimension $d + \dim G$. This implies that $\mathrm{ed}_k(G) \leq d$; see [Reichstein 2000, Theorem 3.4] or [Berhuy and Favi 2003, Proposition 4.11].

(ii) Prove the lower bound $\mathrm{ed}_k(G; p) \geq d$.

Since clearly $\mathrm{ed}(G; p) \leq \mathrm{ed}(G)$, equality (6) follows from (i) and (ii).

The group G will always be of the form $G = D \rtimes F$, where D is diagonalizable and F is finite. In the next section we will recall some known facts about representations of such groups. This will help us in carrying out step (i) and, in the most interesting cases, step (ii) as well, via the Karpenko–Merkurjev Theorem 7.1.

3. Representation-theoretic preliminaries

We will work over a ground field k which remains fixed throughout. Suppose that a linear algebraic k-group G contains a diagonalizable (over k) group D and the quotient G/D is a constant finite group F. Here diagonalizable over k means that D is a subgroup of the split torus $\mathbb{G}_m^d$ defined over k or, equivalently, that every linear representation of D defined over k decomposes as a direct sum of one-dimensional subrepresentations.

Denote the group of (multiplicative) characters of D by $X(D)$. Note that since D is diagonalizable over k, every multiplicative character of D is defined over k. Consider a linear k-representation $G \to \mathrm{GL}(V)$. Restricting this representation to D, we decompose V into a direct sum of one-dimensional character spaces. Let $\Lambda \subset X(D)$ be the set of characters (weights) of D which occur in this decomposition. Note that here $|\Lambda| \leq \dim V$, and equality holds if and only if each character from Λ occurs in V with multiplicity 1. The finite group F acts on $X(D)$ and Λ is invariant under this action. Moreover, if the G-action (and hence, the D-action) on V is generically free then Λ generates $X(D)$ as an abelian group. In summary, we have proved the following lemma; cf. [Serre 1977, Section 8.1].

Lemma 3.1. *Suppose that every F-invariant generating set Λ of $X(D)$ contains at least d elements. If $G \to \mathrm{GL}(V)$ is a generically free k-representation of G then $\dim V \geq d$.* $\qquad\square$

As we explained in the previous section, we are interested in constructing low-dimensional generically free representations of G. In this section we will prove simple sufficient conditions for generic freeness for two particular families of representations.

Lemma 3.2. *Let W be a faithful representation of F and V be a representation of G whose restriction to D is generically free. Then $V \times W$ is a generically free representation of G.*

Here we view W as a representation of G via the natural projection

$$G \to G/D = F.$$

Proof. For $w \in W(\bar{k})$ in general position, we have $\mathrm{Stab}_G(w) = D$ by Remark 2.1. Choosing v in general position in $V(\bar{k})$, we see that $\mathrm{Stab}_G(v, w) = \mathrm{Stab}_G(v) \cap \mathrm{Stab}_G(w) = \mathrm{Stab}_D(v) = \{1\}$. $\qquad \square$

From now on we will assume that $G = D \rtimes F$ is the semidirect product of D and F. In this case, given an F-invariant generating set $\Lambda \subset X(D)$, we can construct a linear k-representation V_Λ of G so that each character from Λ occurs in V_Λ exactly once. To do this, we associate a basis element v_λ to each $\lambda \in \Lambda$. The finite group F acts on

$$V_\Lambda = \mathrm{Span}(v_\lambda \mid \lambda \in \Lambda)$$

by permuting these basis elements in the natural way, that is, via

$$\sigma : v_\lambda \mapsto v_{\sigma(\lambda)} \tag{7}$$

for any $\sigma \in F$ and any $\lambda \in \Lambda$. The diagonalizable group D-acts by the character λ on each one-dimensional space $\mathrm{Span}(v_\lambda)$, that is, via

$$t : v_\lambda \mapsto \lambda(t)v_\lambda \tag{8}$$

for any $t \in D$ and $\lambda \in \Lambda$. Extending (7) and (8) linearly to all of V_Λ, we obtain a linear representation $G = D \rtimes F \to \mathrm{GL}(V_\Lambda)$. Note that by our construction $\dim V_\Lambda = |\Lambda|$.

Our second criterion for generic freeness is a variant of [Lorenz and Reichstein 2000, Lemma 3.1] or [Lemire 2004, Proposition 2.1]. For the sake of completeness we outline a characteristic-free proof.

Lemma 3.3. *Let Λ be an F-invariant subset of $X(D)$ and $\phi : \mathbb{Z}[\Lambda] \to X(D)$ be the natural morphism of $\mathbb{Z}[F]$-modules, taking $\lambda \in \Lambda$ to itself. Let V_Λ be the linear representation of $G = D \rtimes F$ defined by (7) and (8), as above. The G-action on V_Λ is generically free if and only if*

(a) *Λ spans $X(D)$ (or equivalently, ϕ is surjective) and*

(b) *the F-action on $\mathrm{Ker} \, \phi$ is faithful.*

Proof. Let $U \simeq \mathbb{G}_m^n$ be the diagonal subgroup of $\mathrm{GL}(V_\Lambda)$, in the basis e_λ, where $\lambda \in \Lambda$. Here $n = |\Lambda| = \dim V_\Lambda$. The G-action on V induces an F-equivariant morphism $\rho : D \to U$, which is dual to ϕ under the usual (antiequivalence) Diag between finitely generated abelian groups and diagonalizable algebraic groups. Applying Diag to the exact sequence

$$(0) \longrightarrow \mathrm{Ker} \, \phi \longrightarrow \mathbb{Z}[\Lambda] \overset{\phi}{\longrightarrow} X(D) \longrightarrow \mathrm{Coker} \, \phi \longrightarrow (0)$$

of finitely generated abelian $\mathbb{Z}[F]$-modules and setting

$$U = \mathrm{Diag}(\mathbb{Z}[\Lambda]), \quad N = \mathrm{Diag}(\mathrm{Coker} \, \phi), \quad Q = \mathrm{Diag}(\mathrm{Ker} \, \phi),$$

we obtain an F-equivariant exact sequence

$$1 \longrightarrow N \longrightarrow D \overset{\rho}{\longrightarrow} U \longrightarrow Q \longrightarrow 1$$

of diagonalizable groups; see [Jantzen 2003, I 5.6] or [Demazure and Gabriel 1970, IV 1.1]. Since U is F-equivariantly isomorphic to a dense open subset of V, the G-action on V is generically free if and only if the G-action on U is generically free. On the other hand, the G-action on U is generically free if and only if the D-action on U is generically free and the F-action on Q is generically free. But the first of these conditions is equivalent to (a), while the second is equivalent to (b); see Remark 2.1. □

4. Subgroups of prime-to-p index

Lemma 4.1. *Let G' be a closed subgroup of a smooth algebraic group G defined over k. Assume that the index $[G : G'] := \dim_k k[G/G']$ is finite and prime to p. Then $\mathrm{ed}(G; p) = \mathrm{ed}(G'; p)$.*

In the case where G is finite a proof can be found in [Merkurjev 2007, Proposition 4.10]; the argument below proceeds along similar lines.

Proof. Recall that if G is a linear algebraic group and H is a closed subgroup then

$$\mathrm{ed}(G; p) \geq \mathrm{ed}(H; p) + \dim H - \dim G \qquad (9)$$

for any prime p; see [Brosnan et al. 2008, Lemma 2.2] or [Merkurjev 2007, Corollary 4.3]. Since $\dim G' = \dim G$, this yields $\mathrm{ed}(G; p) \geq \mathrm{ed}(G'; p)$.

To prove the opposite inequality, it suffices to show that for any field K/k the map $H^1(K, G') \to H^1(K, G)$ induced by the inclusion $G' \subset G$ is p-surjective, meaning that for every $\alpha \in H^1(K, G)$ there is a finite field extension L/K of degree prime to p such that α_L is in the image of $H^1(L, G') \to H^1(L, G)$; see for example [Merkurjev 2007, Proposition 1.3].

Let $X \to \mathrm{Spec}(K)$ be a G-torsor and X/G' be the natural quotient of X by the action of G'. Recall that X/G' is a K-form of G/G' and that it is constructed by descent [Serre 1962, 1.3.2]. Alternatively, X/G' may be viewed as the Galois twist of G/G' by X with respect to the natural G-action on G/G' [Milne 1980, p. 134].

For a field L/K and an L-point $\mathrm{Spec}(L) \to X/G'$ we construct a G'-torsor Y as the pullback

$$
\begin{array}{ccc}
Y & \longrightarrow & X \\
\downarrow & & \downarrow \\
\mathrm{Spec}(L) & \longrightarrow & X/G' \\
& & \downarrow \\
& & \mathrm{Spec}(K)
\end{array}
$$

In this situation $Y \times^{G'} G \cong X_L$ as G-torsors. Thus we have the natural diagram

$$
\begin{array}{ccc}
H^1(L, G') & \longrightarrow & H^1(L, G) \\
[Y] \longmapsto [X]_L & & \uparrow \\
\uparrow & & \\
[X] & & \\
& H^1(K, G) &
\end{array}
$$

where $[X]$ and $[Y]$ denote the classes of X and Y in $H^1(K, G)$ and $H^1(L, G')$, respectively. It remains to show the existence of such an L-point, with the degree $[L : K]$ prime to p.

Note that G/G' is affine, since G and G' are of the same dimension and hence $G/G' \cong (G/G^\circ)/(G'/G^\circ) = \operatorname{Spec} k[G/G^\circ]^{G'/G^\circ}$ where G° is the connected component of G (and G'). Furthermore G/G' is smooth [Demazure and Gabriel 1970, III 3.2.7]. Let K_s be a separable closure of K. Since X is a G-torsor, we have $X_{K_s} \cong G_{K_s}$ and $(X/G')_{K_s} \cong (G/G')_{K_s}$ which implies that X/G' is also affine [Demazure and Gabriel 1970, III 3.5.6 d)]. Thus, $K[X/G'] \otimes K_s \cong k[G/G'] \otimes K_s$ is reduced and its dimension $\dim_K K[X/G'] = [G : G']$ is not divisible by p by assumption.

Therefore $K[X/G']$ is étale or, equivalently, a product of separable field extensions of K

$$
K[X/G'] = L_1 \times \cdots \times L_r;
$$

see for example [Bourbaki 1990, V, Theorem 4]. For each L_j the projection $K[X/G'] \to L_j$ is an L_j-point of X/G' and since $\dim_K K[X/G'] = \sum_{j=1}^{r} [L_j : K]$ is prime to p, one of the fields L_j must be of degree prime to p over K. We now take $L = L_j$. $\qquad \square$

Corollary 4.2. *Suppose k is a field of characteristic $\neq p$. Then $\operatorname{ed}_k(S_n; p) = [n/p]$.*

Proof. Let $m = [n/p]$ and let $D \simeq (\mathbb{Z}/p\mathbb{Z})^m$ be the subgroup generated by the disjoint p-cycles

$$
\sigma_1 = (1, \ldots, p), \ldots, \sigma_m = \big((m-1)p + 1, \ldots, mp\big).
$$

The inequality $\operatorname{ed}(S_n; p) \geq \operatorname{ed}_k(D; p) \geq [n/p]$ is well known; see any of [Buhler and Reichstein 1997, Section 6; Buhler and Reichstein 1999, Section 7; Berhuy and Favi 2003, Proposition 3.7].

To the best of our knowledge, the opposite inequality was first noticed by J.-P. Serre (private communication, May 2005) and independently by R. Lötscher [Lötscher 2008]. The proof is quite easy. However, since it has not previously appeared in print, we reproduce it below.

The semidirect product $D \rtimes S_m$, where S_m permutes $\sigma_1, \ldots, \sigma_m$, embeds in S_n with index prime to p. By Lemma 4.1, $\mathrm{ed}_k(D \rtimes S_m; p) = \mathrm{ed}_k(S_n; p)$ and it suffices to show that $\mathrm{ed}_k(D \rtimes S_m) \leq [n/p]$. As we mentioned in Section 2, in order to prove this, it is enough to construct a generically free m-dimensional representation of $D \rtimes S_m$ defined over k. Moreover, by Remark 2.2 we may assume that $\zeta_p \in k$, where ζ_p denotes a primitive p-th root of unity.

To construct a generically free m-dimensional representation of $D \rtimes S_m$, let $\sigma_1^*, \ldots, \sigma_m^* \subset X(D)$ be the "basis" of D dual to $\sigma_1, \ldots, \sigma_m$. That is,

$$\sigma_i^*(\sigma_j) = \begin{cases} \zeta_p & \text{if } i = j, \\ 1 & \text{otherwise.} \end{cases}$$

The S_m-invariant subset $\Lambda = \{\sigma_1^*, \ldots, \sigma_m^*\}$ of $X(D)$ gives rise to the m-dimensional k-representation V_Λ of $D \rtimes S_m$, as in Section 3. An easy application of Lemma 3.3 shows that this representation is generically free. □

5. First reductions and proof of Theorem 1.1 parts (a) and (b)

Let $T \simeq \mathbb{G}_m^n / \Delta$ be the diagonal maximal torus in PGL_n, where $\Delta = \mathbb{G}_m$ is diagonally embedded into $\mathbb{G}_m^n$. Recall that the normalizer N of T is isomorphic to $T \rtimes S_n$, where we identify S_n with the subgroup of permutation matrices in PGL_n.

Let P_n be a Sylow p-subgroup of S_n. Lemma 4.1 tells us that

$$\mathrm{ed}_k(N; p) = \mathrm{ed}_k(T \rtimes P_n; p).$$

Note also that by Remark 2.2 we may assume without loss of generality that k contains a primitive p-th root of unity.

Thus in order to prove Theorem 1.1 it suffices to establish the following proposition.

Proposition 5.1. *Let $T \simeq \mathbb{G}_m^n / \Delta$, where $\Delta = \mathbb{G}_m$ is diagonally embedded into $\mathbb{G}_m^n$. Assume that k is of characteristic $\neq p$, containing a primitive p-th root of unity. Then*:

(a) $\mathrm{ed}_k(T \rtimes P_n) = \mathrm{ed}_k(T \rtimes P_n; p) = [n/p]$, *if n is not divisible by p.*

(b) $\mathrm{ed}_k(T \rtimes P_n) = \mathrm{ed}_k(T \rtimes P_n; p) = 2$, *if $n = p$.*

(c) $\mathrm{ed}_k(T \rtimes P_n) = \mathrm{ed}_k(T \rtimes P_n; p) = n^2/p - n + 1$, *if $n = p^r$ for some $r \geq 2$.*

(d) $\mathrm{ed}_k(T \rtimes P_n) = \mathrm{ed}_k(T \rtimes P_n; p) = p^e(n - p^e) - n + 1$, *in all other cases.*

Here P_n is a Sylow p-subgroup of S_n, $[n/p]$ is the integer part of n/p and p^e is the highest power of p dividing n.

The assumption that k contains a primitive p-th root of unity is only needed for the proof of the first equality in parts (a) and (b).

Our proof of each part of this proposition will be based on the strategy outlined in Section 2, with $G = T \rtimes P_n$. We start by recalling that the character lattice $X(T)$ is naturally isomorphic to

$$\{(a_1, \ldots, a_n) \in \mathbb{Z}^n \mid a_1 + \cdots + a_n = 0\},$$

where we identify the character

$$(t_1, \ldots, t_n) \to t_1^{a_1} \ldots t_n^{a_n}$$

of $T = \mathbb{G}_m^n / \Delta$ with $(a_1, \ldots, a_n) \in \mathbb{Z}^n$. Note that $(t_1, \ldots, t_n)$ is viewed as an element of $\mathbb{G}_m^n$ modulo the diagonal subgroup Δ, so the above character is well defined if and only if $a_1 + \cdots + a_n = 0$. An element σ of S_n (and in particular, of $P_n \subset S_n$) acts on $\mathbf{a} = (a_1, \ldots, a_n) \in X(T)$ by naturally permuting $a_1, \ldots, a_n$.

For notational convenience, we will denote by $\mathbf{a}_{i,j} = (a_1, \ldots, a_n) \in X(T)$ the element such that $a_i = 1$, $a_j = -1$ and $a_h = 0$ for every $h \neq i, j$.

We also recall that for $n = p^r$ the Sylow p-subgroup P_n of S_n can be described inductively as the wreath product

$$P_{p^r} \cong P_{p^{r-1}} \wr \mathbb{Z}/p \cong (P_{p^{r-1}})^p \rtimes \mathbb{Z}/p.$$

For general n, P_n is the direct product of certain P_{p^r}; see Section 8.

Proof of Proposition 5.1(a). (i) Since n is not divisible by p, we may assume that P_n is contained in S_{n-1}, where we identify S_{n-1} with the subgroup of S_n consisting of permutations $\sigma \in S_n$ such that $\sigma(1) = 1$.

We will now construct a generically free linear representation V of $T \rtimes P_n$ of dimension $n - 1 + [n/p]$. This will show that $\mathrm{ed}(T \rtimes P_n) \leq [n/p]$.

To construct V, let $\Lambda = \{\mathbf{a}_{1,i} \mid i = 2, \ldots, n\}$ and V_Λ be as in Section 3 and let W be an $[n/p]$-dimensional faithful linear representation of P_n constructed in the proof of Corollary 4.2. Applying Lemma 3.2, we see that $V = V_\Lambda \times W$ is generically free.

(ii) Since the natural projection $p : T \rtimes P_n \to P_n$ has a section, so does the map $p^* : H^1(K, T \rtimes P_n) \to H^1(K, P_n)$ of Galois cohomology sets. Hence, p^* is surjective for every field K/k. This implies that

$$\mathrm{ed}(T \rtimes P_n) \geq \mathrm{ed}(P_n; p) = [n/p].$$

Here $\mathrm{ed}(P_n; p) = \mathrm{ed}(S_n; p)$ by Lemma 4.1 and $\mathrm{ed}(S_n; p) = [n/p]$ by Corollary 4.2. $\qquad\square$

Remark 5.2. We will now outline a different and perhaps more conceptual proof of the upper bound $\mathrm{ed}(N; p) \leq [n/p]$ of Theorem 1.1(a). As we pointed out in the introduction, $\mathrm{ed}(N; p)$ is the essential dimension at p of the functor

$$H^1(*, N) : K \mapsto \{K\text{-isomorphism classes of pairs } (A, L)\},$$

where A is a degree n central simple algebra over K and L is a maximal étale subalgebra of A. Similarly, $\mathrm{ed}(S_n; p)$ is the essential dimension at p of the functor

$$H^1(*, S_n) : K \mapsto \{K\text{-isomorphism classes of } n\text{-dimensional étale algebras } L/K\}.$$

Let $\alpha : H^1(*, S_n) \to H^1(*, N)$ be the map taking an n-dimensional étale algebra L/K to $(\mathrm{End}_K(L), L)$. Here we embed L in $\mathrm{End}_K(L) \simeq M_n(K)$ via the regular action of L on itself.

It is easy to see that, in the terminology of [Merkurjev 2007, Section 1.3], α is p-surjective. That is, for any class (A, L) in $H^1(K, N)$ there exists a prime-to-p extension K'/K such that $(A \otimes_K K', L \otimes_K K')$ lies in the image of α. In fact, any K'/K of degree prime to p which splits A will do (such an extension exists because we are assuming that the degree n of A is not divisible by p). Indeed, by the Skolem–Noether theorem, any two embeddings of $L \otimes_K K'$ into $M_n(K')$ are conjugate. By [Merkurjev 2007, Proposition 1.3], we conclude that $\mathrm{ed}(N; p) \leq \mathrm{ed}(S_n; p)$. Combining this with Corollary 4.2 yields the desired inequality $\mathrm{ed}(N; p) \leq [n/p]$.
□

Proof of Proposition 5.1(b). Here $n = p$ and $P_n \simeq \mathbb{Z}/p$ is generated by the p-cycle $(1, 2, \ldots, n)$. We follow the strategy outlined in Section 2.

(i) To show that $\mathrm{ed}_k(T \rtimes P_n) \leq 2$, we construct a generically free k-representation of $T \rtimes P_n$ of dimension $2 + \dim(T \rtimes P_n) = n + 1$.

Let $\Lambda = \{\mathbf{a}_{1,2}, \ldots, \mathbf{a}_{p-1,p}, \mathbf{a}_{p,1}\}$ and $V = V_\Lambda \times L$, where L is a one-dimensional faithful representation of $P_n \simeq \mathbb{Z}/p$ and $T \rtimes P_n$ acts on L via the natural projection $T \rtimes P_n \to P_n$. Note that $\dim V = |\Lambda| + 1 = n + 1$. Since Λ generates $X(T)$, Lemma 3.2 tells us that V is a generically free representation of $T \rtimes P_n$.

(ii) Recall that $\mathrm{ed}_k(T \rtimes P_n; p) = \mathrm{ed}_k(N; p)$ by Lemma 4.1. On the other hand, as we mentioned in the introduction,

$$\mathrm{ed}_k(N; p) \geq \mathrm{ed}_k(\mathrm{PGL}_p; p) = 2;$$

see (4) and (5). This completes the proof of Proposition 5.1(b) and of Theorem 1.1(b).
□

6. Proof of Theorem 1.1(c): The upper bound

In the next two sections we will prove Proposition 5.1(c), and hence Theorem 1.1(c). We will assume that $n = p^r$ for some $r \geq 2$ and follow the strategy of Section 2. In this section we will carry out Step (i). That is, we will construct a generically free representation V of $T \rtimes P_n$ of dimension p^{2r-1}. This will show

that $\mathrm{ed}(T \rtimes P_n) \leq p^{2r-1} - p^r + 1$. Our V will be of the form V_Λ for a particular P_n-invariant $\Lambda \subset X(T)$, following the recipe of Section 3. Note that this construction (and thus the above inequality) will not require any assumption on the base field k.

For notational convenience, we will subdivide the integers $1, 2, \ldots, p^r$ into p big blocks $B_1, \ldots, B_p$, where each B_i consists of the p^{r-1} consecutive integers $(i-1)p^{r-1} + 1, \ (i-1)p^{r-1} + 2, \ldots, ip^{r-1}$.

We define $\Lambda \subset X(T)$ as the P_n-orbit of the element

$$\mathbf{a}_{1, p^{r-1}+1} = (\underbrace{1, 0, \ldots, 0}_{B_1}, \underbrace{-1, 0, \ldots, 0}_{B_2}, \underbrace{0, 0, \ldots, 0}_{B_3}, \ldots, \underbrace{0, 0, \ldots, 0}_{B_p})$$

in $X(T)$. Thus, Λ consists of elements $\mathbf{a}_{\alpha, \beta}$, subject to the condition that if α lies in the big block B_i then β has to lie in B_j, where $j - i \equiv 1$ modulo p. There are p^r choices for α. Once α is chosen, there are exactly p^{r-1} further choices for β. Thus

$$|\Lambda| = p^r \cdot p^{r-1} = p^{2r-1}.$$

As described in Section 3, we obtain a linear representation V_Λ of $T \rtimes P_n$ of the desired dimension

$$\dim V_\Lambda = |\Lambda| = p^{2r-1}.$$

It remains to prove that V_Λ is generically free. By Lemma 3.3 it suffices to show that

(i) Λ generates $X(T)$ as an abelian group and

(ii) the P_n action on the kernel of the natural morphism $\phi : \mathbb{Z}[\Lambda] \to X(T)$ is faithful.

The elements $\mathbf{a}_{\alpha, \beta}$ clearly generate $X(T)$ as an abelian group, as α and β range over $1, 2, \ldots, p^r$. Thus in order to prove (i) it suffices to show that $\mathrm{Span}_\mathbb{Z}(\Lambda)$ contains every element of this form. Suppose α lies in the big block B_i and β in B_j. If $j - i \equiv 1 \pmod{p}$, then $\mathbf{a}_{\alpha, \beta}$ lies in Λ and there is nothing to prove. If $j - i \equiv 2 \pmod{p}$ then choose some $\gamma \in B_{i+1}$ (where the subscript $i + 1$ should be viewed modulo p) and write

$$\mathbf{a}_{\alpha, \beta} = \mathbf{a}_{\alpha, \gamma} + \mathbf{a}_{\gamma, \beta}.$$

Since both terms on the right are in Λ, we see that in this case $\mathbf{a}_{\alpha, \beta} \in \mathrm{Span}_\mathbb{Z}(\Lambda)$. Using this argument recursively, we see that $\mathbf{a}_{\alpha, \beta}$ also lies in $\mathrm{Span}_\mathbb{Z}(\Lambda)$ if $j - i \equiv 3, \ldots, p \pmod{p}$, i.e., for all possible i and j. This proves (i).

To prove (ii), denote the kernel of ϕ by M. Since P_n is a finite p-group, every normal subgroup of P_n intersects the center of P_n, which we shall denote by Z_n. Thus it suffices to show that Z_n acts faithfully on M.

Recall that Z_n is the cyclic subgroup of P_n of order p generated by the product of disjoint p-cycles

$$\sigma_1 \ldots \sigma_{p^{r-1}} = (1 \ \cdots \ p)(p+1 \ \cdots \ 2p) \cdots (p^r - p + 1 \ \cdots \ p^r).$$

Since $|Z_n| = p$, it either acts faithfully on M or it acts trivially, so we only need to check that the Z_n-action on M is nontrivial. Indeed, Z_n does not fix the nonzero element

$$\mathbf{a}_{1, p^{r-1}+1} + \mathbf{a}_{p^{r-1}+1, 2p^{r-1}+1} + \cdots + \mathbf{a}_{(p-1)p^{r-1}+1, 1} \in \mathbb{Z}[\Lambda]$$

which lies in M. This proves the upper bound of Proposition 5.1(c) and Theorem 1.1(c). $\qquad\square$

7. Proof of Theorem 1.1(c): The lower bound

In this section we will continue to assume that $n = p^r$. We will show that

$$\mathrm{ed}(N; p) \geq p^{2r-1} - p^r + 1, \tag{10}$$

thus completing the proof of Proposition 5.1(c) and Theorem 1.1(c). Let

$$q := p^e, \text{ where } e \geq 1 \text{ if } p \text{ is odd and } e \geq 2 \text{ if } p = 2. \tag{11}$$

be a power of p. The specific choice of e will not be important in the sequel; in particular, the reader may assume that $q = p$ if p is odd and $q = 4$, if $p = 2$. Whatever e we choose, $q = p^e$ will remain unchanged for the rest of this section.

We now recall that if k'/k is a field extension then

$$\mathrm{ed}_k(N; p) \geq \mathrm{ed}_{k'}(N; p),$$

by [Merkurjev 2007, Proposition 1.5(1)]. Thus for the purpose of proving (10) we may replace k by k'. In particular, we may assume that k' contains a primitive q-th root of unity.

Let $T_{(q)} = \mu_q^n / \mu_q$ be the q-torsion subgroup of $T = \mathbb{G}_m^n / \Delta$. Applying the inequality (9) to $G = T \rtimes P_n$ and its finite subgroup $H = T_{(q)} \rtimes P_n$, we obtain

$$\mathrm{ed}(T \rtimes P_n; p) \geq \mathrm{ed}(T_{(q)} \rtimes P_n; p) - p^r + 1.$$

Thus it suffices to show that

$$\mathrm{ed}(T_{(q)} \rtimes P_n; p) \geq p^{2r-1}. \tag{12}$$

The advantage of replacing $T \rtimes P_n$ by $T_{(q)} \rtimes P_n$ is that $T_{(q)} \rtimes P_n$ is a finite p-group, so that we can apply the following result:

Theorem 7.1 [Karpenko and Merkurjev 2008]. *Let G be a finite p-group and k be a field containing a primitive p-th root of unity. Then* $\mathrm{ed}_k(G; p) = \mathrm{ed}_k(G)$ *equals the minimal value of* $\dim V$, *where V ranges over all faithful linear k-representations* $G \to \mathrm{GL}(V)$.

Now recall that we are assuming that k contains a primitive q-th root of unity and hence, a primitive p-th root of unity. Hence, Theorem 7.1 applies in our situation. That is, in order to prove (12) it suffices to show that $T_{(q)} \rtimes P_n$ does not have a faithful linear representation of dimension less than p^{2r-1}. Lemma 3.1 further reduces this representation-theoretic assertion to the combinatorial statement of Proposition 7.2 below. Before stating the proposition we recall that the character lattice of $T_{(q)} \simeq \mu_q^n / \mu_q$ is

$$X_n := \{(a_1, \ldots, a_n) \in (\mathbb{Z}/q\mathbb{Z})^n \mid a_1 + \cdots + a_n = 0 \text{ in } \mathbb{Z}/q\mathbb{Z}\},$$

where we identify the character

$$(t_1, \ldots, t_n) \to t_1^{a_1} \ldots t_n^{a_n}$$

of $T_{(q)}$ with $(a_1, \ldots, a_n) \in (\mathbb{Z}/q\mathbb{Z})^n$. Here $(t_1, \ldots, t_n)$ stands for an element of μ_q^n, modulo the diagonally embedded μ_q, so the character above is well defined if and only if $a_1 + \cdots + a_n = 0$ in $\mathbb{Z}/q\mathbb{Z}$. (This is completely analogous to our description of the character lattice of T in the previous section.) Note that X_n depends on the integer $q = p^e$, which we assume to be fixed throughout this section.

Proposition 7.2. *Let* $n = p^r$ *and* P_n *be a Sylow p-subgroup of* S_n. *If* Λ *is a* P_n-*invariant generating subset of* X_n *then* $|\Lambda| \geq p^{2r-1}$ *for any* $r \geq 1$.

Our proof relies on the following special case of Nakayama's Lemma:

Lemma 7.3 [Atiyah and Macdonald 1969, Proposition 2.8]. *Let* $q = p^e$ *be a prime power,* $M = (\mathbb{Z}/q\mathbb{Z})^d$ *and* Λ *be a generating subset of M (as an abelian group). If we remove from* Λ *all elements that lie in* pM, *the remaining set,* $\Lambda \setminus pM$, *will still generate M.* $\square$

Proof of Proposition 7.2. We argue by induction on r. For the base case, set $r = 1$. We need to show that $|\Lambda| \geq p$. Assume the contrary. In this case P_n is a cyclic p-group, and every nontrivial orbit of P_n has exactly p elements. Hence, $|\Lambda| < p$ is only possible if every element of Λ is fixed by P_n. Since we are assuming that Λ generates X_n as an abelian group, we conclude that P_n acts trivially on X_n. This can happen only if $p = q = 2$. Since these values are ruled out by our definition (11) of q, we have proved the proposition for $r = 1$.

In the previous section we subdivided the integers $1, 2, \ldots, p^r$ into p big blocks $B_1, \ldots, B_p$ of length p^{r-1}. Now we will now work with small blocks $b_1, \ldots, b_{p^{r-1}}$,

where b_j consists of the p consecutive integers

$$(j-1)p+1, \ (j-1)p+2, \ \ldots, \ jp.$$

We can identify $P_{p^{r-1}}$ with the subgroup of P_{p^r} that permutes the small blocks $b_1, \ldots, b_{p^{r-1}}$ without changing the order of the elements in each block.

For the induction step, assume $r \geq 2$ and consider the homomorphism $\Sigma : X_{p^r} \to X_{p^{r-1}}$ given by

$$\mathbf{a} = (a_1, a_2, \ldots, a_{p^r}) \mapsto \mathbf{s} = (s_1, \ldots, s_{p^{r-1}}), \tag{13}$$

where $s_i = a_{(i-1)p+1} + a_{(i-1)p+2} + \cdots + a_{ip}$ is the sum of the entries of $\mathbf{a}$ in the i-th small block b_i. Thus

(i) if Λ generates X_{p^r} then $\Sigma(\Lambda)$ generates $X_{p^{r-1}}$.

(ii) if Λ is a P_{p^r}-invariant subset of X_{p^r} then $\Sigma(\Lambda)$ is a $P_{p^{r-1}}$-invariant subset of $X_{p^{r-1}}$.

Let us remove from $\Sigma(\Lambda)$ all elements which lie in $pX_{p^{r-1}}$. The resulting set, $\Sigma(\Lambda) \setminus pX_{p^{r-1}}$, is clearly $P_{p^{r-1}}$-invariant. By Lemma 7.3 this set generates $X_{p^{r-1}}$. Thus by the induction assumption $|\Sigma(\Lambda) \setminus pX_{p^{r-1}}| \geq p^{2r-3}$.

We claim that the fiber of each element $\mathbf{s} = (s_1, \ldots, s_{p^{r-1}})$ in $\Sigma(\Lambda) \setminus pX_{p^{r-1}}$ has at least p^2 elements in Λ. If we can show this, then we will be able to conclude that

$$|\Lambda| \geq p^2 \cdot |\Sigma(\Lambda) \setminus pX_{p^{r-1}}| \geq p^2 \cdot p^{2r-3} = p^{2r-1},$$

thus completing the proof of Proposition 7.2.

Let σ_i be the single p-cycle, cyclically permuting the elements in the small block b_i. To prove the claim, note that the subgroup

$$\langle \sigma_i \mid i = 1, \ldots, p^{r-1} \rangle \simeq (\mathbb{Z}/p\mathbb{Z})^{p^{r-1}}$$

of P_n acts on each fiber of Σ.

To simplify the exposition in the argument to follow, we introduce the following bit of terminology. Let us say that $\mathbf{a} \in (\mathbb{Z}/q\mathbb{Z})^n$ is *scalar in the small block* b_i if all the entries of $\mathbf{a}$ in the block b_i are the same, that is, if

$$a_{(i-1)p+1} = a_{(i-1)p+2} = \cdots = a_{ip}.$$

We are now ready to prove the claim. Suppose $\mathbf{a} = (a_1, \ldots, a_{p^r}) \in X_{p^r}$ lies in the preimage of $\mathbf{s} = (s_1, \ldots, s_{p^{r-1}})$, as in (13). If $\mathbf{a}$ is scalar in the small block b_i then clearly

$$s_i = a_{(i-1)p+1} + a_{(i-1)p+2} + \cdots + a_{ip} \in p\mathbb{Z}/q\mathbb{Z}.$$

Since we are assuming that $\mathbf{s}$ lies in

$$\Sigma(\Lambda) \setminus pX_{p^{r-1}},$$

s must have at least two entries that are not divisible by p, say, s_i and s_j. (Recall that $s_1 + \cdots + s_{p^r} = 0$ in $\mathbb{Z}/q\mathbb{Z}$, so **s** cannot have exactly one entry not divisible by p.) Thus **a** is nonscalar in the small blocks b_i and b_j. Consequently, the elements $\sigma_i^\alpha \sigma_j^\beta (a)$ are distinct, as α and β range between 0 and $p-1$. All of these elements lie in the fiber of **s** under Σ. Therefore we conclude that this fiber contains at least p^2 distinct elements. This completes the proof of the claim and thus of Proposition 7.2, Proposition 5.1(c) and Theorem 1.1(c). $\qquad\square$

8. Proof of Theorem 1.1(d)

In this section we assume that n is divisible by p but is not a power of p. We will modify the arguments of the last two sections to show that

$$\mathrm{ed}(T \rtimes P_n) = \mathrm{ed}(T \rtimes P_n; p) = p^e(n - p^e) - n + 1,$$

where p^e is the highest power of p dividing n. This will complete the proof of Proposition 5.1 and thus of Theorem 1.1.

Write out the p-adic expansion

$$n = n_1 p^{e_1} + n_2 p^{e_2} + \cdots + n_u p^{e_u}, \tag{14}$$

of n, where $1 \le e = e_1 < e_2 < \cdots < e_u$, and $1 \le n_i < p$ for each i. Subdivide the integers $1, \ldots, n$ into $n_1 + \cdots + n_u$ blocks B_j^i of length p^{e_i}, for j ranging over $1, 2, \ldots, n_i$. By our assumption there are at least two such blocks. The Sylow subgroup P_n is a direct product

$$P_n = (P_{p^{e_1}})^{n_1} \times \cdots \times (P_{p^{e_u}})^{n_u}$$

where each $P_{p^{e_i}}$ acts on one of the blocks B_j^i.

Once again we will use the strategy outlined in Section 2.

(i) We will construct a generically free representation of $T \rtimes P_n$ of dimension $p^{e_1}(n - p^{e_1})$. This will prove the upper bound $\mathrm{ed}_k(T \rtimes P_n) \le p^{e_1}(n - p^{e_1}) - n + 1$. Note that this construction (and thus the above inequality) do not require any assumption on the field k.

To construct this representation, let $\Lambda \subset X(T)$ be the union of the P_n-orbits of the elements

$$\mathbf{a}_{1, j+1} \text{ where } j = p^{e_1}, \ldots, n_1 p^{e_1}, n_1 p^{e_1} + p^{e_2}, \ldots, n - p^{e_u},$$

i.e., the union of the P_n-orbits of elements of the form $(1, 0, \ldots, 0, -1, 0, \ldots, 0)$, where 1 appears in the first position of the first block and -1 appears in the first position of one of the other blocks. For $\mathbf{a}_{\alpha,\beta}$ in Λ there are p^{e_1} choices for α and $n - p^{e_1}$ choices for β. Thus

$$\dim V_\Lambda = |\Lambda| = p^{e_1}(n - p^{e_1}).$$

It is not difficult to see that Λ generates $X(T)$ as an abelian group. To conclude with Lemma 3.3 that V_Λ is a generically free representation of $T \rtimes P_n$, it remains to show that the P_n-action on the kernel of the natural morphism $\phi : \mathbb{Z}[\Lambda] \to X(T)$ is faithful when $e_1 \geq 1$. As in Section 6 we only need to check that the center Z_n of P_n acts faithfully on the kernel. Let σ be a nontrivial element of $Z_n = (Z_{p^{e_1}})^{n_1} \times \cdots \times (Z_{p^{e_u}})^{n_u}$, with each $Z_{p^{e_i}}$ cyclic of order p. Let h, h' be in the first block B_1^1 and l, l' in some other block B_i^j (there are at least two blocks each of size at least p). The element

$$\mathbf{a} = \mathbf{a}_{h,l} - \mathbf{a}_{h,l'} + \mathbf{a}_{h',l'} - \mathbf{a}_{h',l}$$

lies in the kernel of ϕ. To fix $\mathbf{a}$, σ must either (1) fix all h, h', l, l' or (2) $\sigma(h) = h', \sigma(h') = h$ and $\sigma(l) = l', \sigma(l') = l$. Since σ is nontrivial we may choose B_i^j such that (1) is not possible and if $p \neq 2$, (2) is not possible either. If $p = 2$, by (14), B_i^j is at least of size 4 and we can choose l, l' within B_i^j such that (2) does not hold. Therefore σ does not fix a nonzero element of the kernel of ϕ.

(ii) We now want to prove the lower bound,

$$\mathrm{ed}(T \rtimes P_n; p) \geq p^{e_1}(n - p^{e_1}) - n + 1.$$

Arguing as in Section 7 (and using the same notation, with $q = p$), it suffices to show that $\mathrm{ed}(T_{(p)} \rtimes P_n; p) \geq p^{e_1}(n - p^{e_1})$. By the Karpenko–Merkurjev Theorem 7.1 this is equivalent to showing that every faithful representation of $T_{(p)} \rtimes P_n$ has dimension at least $p^{e_1}(n - p^{e_1})$. By Lemma 3.1 it now suffices to prove the following lemma.

Lemma 8.1. *Let n be a positive integer, P_n be the Sylow subgroup of S_n, p^e be the highest power of p dividing n, and*

$$X_n := \big\{ (a_1, \ldots, a_n) \in (\mathbb{Z}/p\mathbb{Z})^n \mid a_1 + \cdots + a_n = 0 \text{ in } \mathbb{Z}/p\mathbb{Z} \big\}.$$

Then every P_n-invariant generating subset of X_n has at least $p^e(n - p^e)$ elements.

In the statement of the lemma we allow $e = 0$, to facilitate the induction argument. For the purpose of proving the lower bound in Proposition 5.1(d) we only need this lemma for $e \geq 1$.

Proof. Once again, we consider the p-adic expansion (14) of n with $0 \leq e_1 < e_2 < \cdots < e_u$ and $1 \leq n_i < p$. We may assume that n is not a power of p, since otherwise the lemma is vacuous.

We will argue by induction on $e = e_1$. For the base case, let $e_1 = 0$. Here the lemma is obvious: since X_n has rank $n - 1$, every generating set (P_n-invariant or not) has to have at least $n - 1$ elements.

For the induction step, we may suppose $e = e_1 \geq 1$; in particular, n is divisible by p. Define $\Sigma : X_n \to X_{n/p}$ by sending $(a_1, \ldots, a_n)$ to $(s_1, \ldots, s_{n/p})$, where

$$s_j = a_{(j-1)p+1} + \cdots + a_{jp}$$

for $j = 1, \ldots, n/p$. Arguing as in Section 7 we see that $\Sigma(\Lambda) \setminus pX_{n/p}$ is a $(P_{p^{e_1-1}})^{n_1} \times \cdots \times (P_{p^{e_u-1}})^{n_u}$-invariant generating subset of $X_{n/p}$ and that every

$$\mathbf{s} \in \Sigma(\Lambda) \setminus pX_{n/p}$$

has at least p^2 preimages in Λ. By the induction assumption,

$$|\Sigma(\Lambda) \setminus pX_{n/p}| \geq p^{e-1}\left(\frac{n}{p} - p^{e-1}\right)$$

and thus

$$|\Lambda| \geq p^2 \cdot p^{e-1}\left(\frac{n}{p} - p^{e-1}\right) = p^e(n - p^e)$$

This completes the proof of Lemma 8.1 and thus of parts (d) of Proposition 5.1 and of Theorem 1.1. $\qquad\qquad\square$

Acknowledgements

References

[Atiyah and Macdonald 1969] M. F. Atiyah and I. G. Macdonald, *Introduction to commutative algebra*, Addison-Wesley, Reading, MA, 1969. MR 39 #4129 Zbl 0175.03601

[Berhuy and Favi 2003] G. Berhuy and G. Favi, "Essential dimension: a functorial point of view (after A. Merkurjev)", *Doc. Math.* **8** (2003), 279–330. MR 2004m:11056 Zbl 1101.14324

[Bourbaki 1990] N. Bourbaki, *Algebra II: Chapters 4–7*, Springer, Berlin, 1990. MR 91h:00003 Zbl 0719.12001

[Brosnan et al. 2008] P. Brosnan, Z. Reichstein, and A. Vistoli, "Essential dimension, spinor groups and quadratic forms", preprint, 2008, http://www.math.ubc.ca/~reichst/pfister-numbers.pdf. To appear in *Annals of Math.*

[Buhler and Reichstein 1997] J. Buhler and Z. Reichstein, "On the essential dimension of a finite group", *Compositio Math.* **106**:2 (1997), 159–179. MR 98e:12004 Zbl 0905.12003

[Buhler and Reichstein 1999] J. Buhler and Z. Reichstein, "On Tschirnhaus transformations", pp. 127–142 in *Topics in number theory (University Park, PA, 1997)*, edited by D. Scott, Math. Appl. **467**, Kluwer Acad. Publ., Dordrecht, 1999. MR 2000b:12003 Zbl 0937.12001

[Chernousov and Serre 2006] V. Chernousov and J.-P. Serre, "Lower bounds for essential dimensions via orthogonal representations", *J. Algebra* **305**:2 (2006), 1055–1070. MR 2007i:20070 Zbl 05078318

[Demazure and Gabriel 1970] M. Demazure and P. Gabriel, *Groupes algébriques, I: Géométrie algébrique, généralités, groupes commutatifs*, Masson, 1970. MR 46 #1800 Zbl 0203.23401

[Favi and Florence 2008] G. Favi and M. Florence, "Tori and essential dimension", *J. Algebra* **319**:9 (2008), 3885–3900. MR 2009b:20085 Zbl 1141.14028

[Grothendieck 1958] A. Grothendieck, "Torsion homologique et sections rationnelles", in *Séminaire C. Chevalley, 2ème année: Anneaux de Chow et applications* (Exposé no. 5), Secrétariat de mathématique, ENS, Paris, 1958.

[Jantzen 2003] J. C. Jantzen, *Representations of algebraic groups*, 2nd. ed., Mathematical Surveys and Monographs **107**, American Mathematical Society, Providence, RI, 2003. MR 2004h:20061 Zbl 1034.20041

[Karpenko and Merkurjev 2008] N. A. Karpenko and A. S. Merkurjev, "Essential dimension of finite *p*-groups", *Invent. Math.* **172**:3 (2008), 491–508. MR 2009b:12009 Zbl 05279042

[Lemire 2004] N. Lemire, "Essential dimension of algebraic groups and integral representations of Weyl groups", *Transform. Groups* **9**:4 (2004), 337–379. MR 2005j:20056 Zbl 1076.14060

[Lorenz and Reichstein 2000] M. Lorenz and Z. Reichstein, "Lattices and parameter reduction in division algebras", preprint 2000-001, MSRI, 2000, Available at http://www.msri.org/publications/preprints/online/2000-001.html.

[Lorenz et al. 2003] M. Lorenz, Z. Reichstein, L. H. Rowen, and D. J. Saltman, "Fields of definition for division algebras", *J. London Math. Soc.* (2) **68**:3 (2003), 651–670. MR 2004j:16022 Zbl 1071.16012

[Lötscher 2008] R. Lötscher, "Application of multihomogeneous rational covariants to the determination of essential dimension of finite groups", preprint, 2008.

[Merkurjev 2007] A. S. Merkurjev, "Essential dimension", preprint, 2007, Available at http://www.math.ucla.edu/~merkurev/papers/review.dvi. To appear in *Algebraic and arithmetic theory of quadratic forms* (Llanquihue, Chile, 2007), Contemporary Mathematics, Amer. Math. Soc.

[Merkurjev 2008] A. S. Merkurjev, "Essential *p*-dimension of PGL(p^2)", preprint, 2008, Available at http://www.math.ucla.edu/~merkurev/papers/pgl.dvi.

[Milne 1980] J. S. Milne, *Étale cohomology*, Princeton Mathematical Series **33**, Princeton University Press, Princeton, NJ, 1980. MR 81j:14002 Zbl 0433.14012

[Procesi 1967] C. Procesi, "Non-commutative affine rings", *Atti Accad. Naz. Lincei Mem. Cl. Sci. Fis. Mat. Natur. Sez. I* (8) **8** (1967), 239–255. MR 37 #256 Zbl 0204.04802

[Reichstein 1999] Z. Reichstein, "On a theorem of Hermite and Joubert", *Canad. J. Math.* **51**:1 (1999), 69–95. MR 2000h:12008 Zbl 0942.12001

[Reichstein 2000] Z. Reichstein, "On the notion of essential dimension for algebraic groups", *Transform. Groups* **5**:3 (2000), 265–304. MR 2001j:20073 Zbl 0981.20033

[Reichstein and Youssin 2000] Z. Reichstein and B. Youssin, "Essential dimensions of algebraic groups and a resolution theorem for *G*-varieties", *Canad. J. Math.* **52**:5 (2000), 1018–1056. MR 2001k:14088 Zbl 1044.14023

[Rowen 1980] L. H. Rowen, *Polynomial identities in ring theory*, Pure and Applied Mathematics **84**, Academic Press, New York, 1980. MR 82a:16021 Zbl 0461.16001

[Rowen and Saltman 1992] L. H. Rowen and D. J. Saltman, "Prime-to-*p* extensions of division algebras", *Israel J. Math.* **78**:2-3 (1992), 197–207. MR 93k:12006 Zbl 0795.16012

[Serre 1958] J.-P. Serre, "Modules projectifs et espaces fibrés à fibre vectorielle", in *Séminaire P. Dubreil, M.-L. Dubreil-Jacotin et C. Pisot, 1957/58,* Fasc. 2, Exposé 23, Secrétariat mathématique, Paris, 1958. MR 31 #1277 Zbl 0132.41202

[Serre 1962] J.-P. Serre, "Cohomologie galoisienne des groupes algébriques linéaires", pp. 53–68 in *Colloq. Théorie des Groupes Algébriques* (Brussels, 1962), Librairie Universitaire, Louvain, 1962. MR 32 #4177 Zbl 0145.17501

[Serre 1977] J.-P. Serre, *Linear representations of finite groups*, Graduate Texts in Mathematics, Springer, New York, 1977. 42. MR 56 #8675 Zbl 0355.20006

[Serre 1997] J.-P. Serre, *Galois cohomology*, Springer, Berlin, 1997. MR 98g:12007 ZBL 0902. 12004

aurel@math.ubc.ca   *Department of Mathematics, University of British Columbia,*
                                     1984 Mathematics Road, Vancouver, BC V6T 1Z2, Canada
                                     http://www.math.ubc.ca/~aurel

reichst@math.ubc.ca   *Department of Mathematics, University of British Columbia,*
                                     1984 Mathematics Road, Vancouver, BC V6T 1Z2, Canada
                                     http://www.math.ubc.ca/~reichst