THE ESSENTIAL DIMENSION OF THE NORMALIZER OF A MAXIMAL TORUS IN THE PROJECTIVE LINEAR GROUP

AUREL MEYER† AND ZINOVY REICHSTEIN††

Abstract. Let p be a prime, k be a field of characteristic $\neq p$ and N be 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 $n \geq 1$.

Contents

1. Introduction

Let p be a prime, k be a field of characteristic $\neq p$ and N be the normalizer of a split maximal torus in the projective linear group PGL_n, for some integer n. The purpose of this paper is to compute the essential dimension $\mathrm{ed}_k(N; p)$ of N at p. For the definition of essential dimension of an algebraic group (and more generally, of a functor), we refer the reader to [Re_2], [BF], [BRV] or [Me]. As usual, if the reference to k is clear from the context, we will sometimes write ed in place of ed_k.

1991 *Mathematics Subject Classification.* 11E72, 20D15, 16K20.

Key words and phrases. Essential dimension, central simple algebra, character lattice, finite p-group, Galois cohomology.

† Aurel Meyer was partially supported by a University Graduate Fellowship at the University of British Columbia.

†† Z. Reichstein was partially supported by NSERC Discovery and Accelerator Supplement grants.

We begin by explaining why we are interested in the essential dimension of N. One of the central problems in the theory of essential dimension is to find the exact value of the essential dimension of the projective linear group PGL_n or equivalently, of the functor

$$H^1(\,*\,,\mathrm{PGL}_n)\colon K \mapsto \{\text{ degree } n \text{ central simple algebras } A/K, \\ \text{up to } K\text{-isomorphism }\},$$

where K is a field extension of k. This problem arises naturally 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 $\mathrm{ed}(\mathrm{PGL}_n) \le n^2$; see [Pr, 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,

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

(cf. [Re$_1$, Theorem 16.1(b)] or [RY, Theorem 8.6]), falls far below the best known upper bound,

$$\mathrm{ed}(\mathrm{PGL}_n) \le \begin{cases} \frac{(n-1)(n-2)}{2}, \text{ for every odd } n \ge 5 \text{ and} \\ n^2 - 3n + 1, \text{ for every } n \ge 4; \end{cases} \tag{1}$$

see [LR], [LRRS, Theorem 1.1], [Le, Proposition 1.6] and [FF].

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 $n = p_1^{r_1} \dots p_s^{r_s}$ 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

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

for every $i = 1, \dots, s$; cf. [Re$_2$, Proposition 9.8].

It is important to note that the proofs of the upper bounds (1) are not based on a direct analysis of the functor $H^1(\,*\,,\mathrm{PGL}_n)$. Instead, one works with the related functor

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

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. 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 $H^1(\,*\,,N) \to H^1(\,*\,,\mathrm{PGL}_n)$. The surjectivity of this morphism (which is a special case of a more general result of T. Springer, see [Se$_2$, III.4.3, Lemma 6]) leads to the inequalities

$$\mathrm{ed}(N) \ge \mathrm{ed}(\mathrm{PGL}_n) \text{ and } \mathrm{ed}(N; p) \ge \mathrm{ed}(\mathrm{PGL}_n; p)\,; \tag{2}$$

see [Me, Proposition 1.3], [BF, Lemma 1.9] or [Re$_2$, Proposition 4.3]. The inequalities (1) were, in fact, proved as upper bounds on $\mathrm{ed}(N)$; see [LRRS] and [Le]. 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 [LRRS] and [Le]. 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 $[n/p]$ denotes the integer part of n/p and p^e denotes 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, using rather different techniques. There is nothing about the methods we use that in any way guarantees that the lower bounds should match the upper bounds, thus yielding an exact value of $\mathrm{ed}(N;p)$. The fact that this happens, for any base field k of characteristic $\neq p$, 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 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 (2), 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

$$\mathrm{ed}_k(\mathrm{PGL}_p;p) = 2, \tag{3}$$

see [Re$_2$, Corollary 5.7] or [RY, 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; see [RS, Corollary 1.3]. The upper bound on the essential dimension of a crossed product given by [LRRS, Corollary 3.10] then yields the inequality

$$\mathrm{ed}(\mathrm{PGL}_{p^2}; p) \le p^2 + 1\,,$$

which is stronger than Corollary 1.2 for any $p \geq 3$. If $r \geq 3$ we do not know how close the true value of $\mathrm{ed}(\mathrm{PGL}_{p^r}; p)$ is to $\mathrm{ed}(N; p) = p^{2r-1} - p^r + 1$; in this case Corollary 1.2 gives the best currently known upper bound on $\mathrm{ed}(\mathrm{PGL}_{p^r}; p)$. We remark that, beyond the obvious inequality $\mathrm{ed}(\mathrm{PGL}_{p^r}; p) \le \mathrm{ed}(\mathrm{PGL}_{p^r})$, the relationship between $\mathrm{ed}(\mathrm{PGL}_{p^r}; p)$ and $\mathrm{ed}(\mathrm{PGL}_{p^r})$ is quite mysterious as well.

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

Acknowledgements

The authors are grateful to A. S. Merkurjev and J.-P. Tignol for helpful comments.

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 $x \in X(\overline{k})$ in general position.

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

Indeed, the "only if" implication 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 $\le \dim(X)$. □

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

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

$$\mathrm{ed}_k(G) = \mathrm{ed}_k(G; p) = d\,, \tag{4}$$

where G is a linear algebraic group defined over k.

All such assertions will be proved by the following 2-step procedure.

(i) Construct a generically free linear representation of G of dimension $d+\dim(G)$ defined over k. This implies that $\mathrm{ed}_k(G) \le d$; see [Re$_2$, Theorem 3.4] or [BF, Proposition 4.11].

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

Since clearly $\mathrm{ed}(G;p) \le \mathrm{ed}(G)$, the desired equality (4) 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 by "diagonalizable over k" we mean 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 1-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 1-dimensional character spaces. Let $\Lambda \subset X(D)$ be the set of characters (weights) of D which occur in this decomposition. Note that here $|\Lambda| \le \dim(V)$, 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. [Se$_1$, Section 8.1].

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

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 $w \in W(\overline{k})$ in general position, $\mathrm{Stab}_G(w) = D$; cf. Remark 2.1. Choosing v in general position in $V(\overline{k})$, we see that

$$\mathrm{Stab}_G(v, w) = \mathrm{Stab}_G(v) \cap \mathrm{Stab}_G(w) = \mathrm{Stab}_D(v) = \{1\}\,.$$

$\square$

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 (in fact, a *monomial*) 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, i.e., via

$$\sigma\colon v_\lambda \mapsto v_{\sigma(\lambda)}\,. \tag{5}$$

for any $\sigma \in F$ and any $\lambda \in \Lambda$. The diagonalizable group D-acts by the character λ on each 1-dimensional space $\mathrm{Span}(v_\lambda)$, i.e., via

$$t\colon v_\lambda \mapsto \lambda(t)v_\lambda \tag{6}$$

for any $t \in D$ and $\lambda \in \Lambda$. Extending (5) and (6) 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 [LR, Lemma 3.1] or [Le, 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 $\phi\colon \mathbb{Z}[\Lambda] \to X(D)$ 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* (5) *and* (6)*, 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 $\rho\colon D \to U$, which is dual to ϕ under the usual (anti-equivalence) Diag between finitely generated abelian groups and diagonalizable algebraic groups. Applying Diag to the exact sequence

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

of finitely generated abelian $\mathbb{Z}[F]$-modules we obtain an F-equivariant exact sequence

$$1 \longrightarrow N \longrightarrow D \xrightarrow{\rho} U \longrightarrow Q \longrightarrow 1\,,$$

of diagonalizable groups, where $U = \mathrm{Diag}(\mathbb{Z}[\Lambda])$, $N = \mathrm{Diag}(\mathrm{Coker}(\phi))$ and $Q = \mathrm{Diag}(\mathrm{Ker}(\phi))$; cf. [Ja, I 5.6] or [DG, 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 (i) the D-action on U is generically free, and (ii) the F-action on Q is generically free.

It is now easy to see that (i) is equivalent to (a) and (ii) is equivalent to (b); cf. Remark 2.1. □

4. Subgroups of prime-to-p index

Our starting point is the following lemma.

Lemma 4.1. *Let G' be a closed subgroup of a smooth algebraic group G defined over k. Assume that the index $[G : G']$ 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 [Me, 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

$$\mathrm{ed}(G;p) \geq \mathrm{ed}(H;p) + \dim(H) - \dim(G)\,; \tag{7}$$

for any prime p; see, [BRV, Lemma 2.2] or [Me, 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, i.e., 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, e.g., [Me, Proposition 1.3].

Let X be a G-torsor over K and X/G the quotient by the action of G'. For a field L/K and an L-point $\mathrm{Spec}(L) \to X/G'$ we construct a G'-torsor Y as the pullback

$$\begin{tikzcd}
Y \arrow[r] \arrow[d] & X \arrow[d] \\
\mathrm{Spec}(L) \arrow[r] & X/G' \arrow[d] \\
& \mathrm{Spec}(K)
\end{tikzcd}$$

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

$$\begin{tikzcd}
H^1(L,G') \arrow[rr] & & H^1(L,G) \\
{[Y]} \arrow[r, mapsto] & {[X]_L} & \\
& {[X]} \arrow[u, mapsto] & \\
& & H^1(K,G) \arrow[uuu]
\end{tikzcd}$$

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; cf. [DG, III 3.2.7]. Let K_s be the separable closure of K. X being 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, cf. [DG, 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, e.g., [Bo, 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

$$\dim_K K[X/G'] = \sum_{j=1}^{r} [L_j : K] \quad \text{ is prime to } p,$$

one of the fields L_j must be of degree prime to p over K. We now take $L = L_j$. □

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

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

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

The inequality $\mathrm{ed}(\mathrm{S}_n;p) \geq \mathrm{ed}_k(D;p) \geq [n/p]$ is well known; see, [BuR$_1$, Section 6], [BuR$_2$, Section 7], or [BF, 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 [Lö]. The proof is quite easy; however, since it has not previously appeared in print, we reproduce it below.

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

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

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

The S_m-invariant subset $\Lambda = \{\sigma_1^*, \dots, \sigma_m^*\}$ of $X(D)$ gives rise to the m-dimensional k-representation V_Λ of $D \rtimes \mathrm{S}_m$, 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 $T \rtimes \mathrm{S}_n$, 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 pth 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 pth 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 pth 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 $G = T \rtimes P_n$. Before we proceeding with the proof of the proposition, we recall that the character lattice $X(T)$ is naturally isomorphic to

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

where we identify the character

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

of $T = \mathbb{G}_m^n/\Delta$ with $(a_1, \dots, a_n) \in \mathbb{Z}^n$. Note that $(t_1, \dots, t_n)$ 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 $a_1 + \dots + a_n = 0$. An element σ of S_n (and in particular, of $P_n \subset \mathrm{S}_n$) acts on $\mathbf{a} = (a_1, \dots, a_n) \in X(T)$ by naturally permuting $a_1, \dots, a_n$.

For notational convenience, we will denote by $\mathbf{a}_{i,j}$ the element of $(a_1, \dots, a_n) \in X(T)$ 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). Step (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 $\sigma \in \mathrm{S}_n$ such that $\sigma(1) = 1$.

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

To construct V, let $\Lambda = \{\mathbf{a}_{1,i} \,|\, i = 2, \dots, n\}$ and let W be a $[n/p]$-dimensional faithful linear representation of P_n constructed in the proof of Corollary 4.2 (we may adjoin pth roots of unity to k if needed, see Remark 2.2). Applying Lemma 3.2(b), we see that $V = V_\Lambda \times W$ is generically free.

Step (ii): Since the natural projection $p\colon T \rtimes P_n \to P_n$ has a section, so does the map $p^*\colon H^1(K, T \rtimes P_n) \to H^1(K, P_n)$ 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 $\mathrm{ed}(P_n; p) = \mathrm{ed}(\mathrm{S}_n; p)$ by Lemma 4.1 and $\mathrm{ed}(\mathrm{S}_n; p) = [n/p]$ by Corollary 4.2. $\square$

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

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

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

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

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

It is easy to see that, in the terminology of [Me, 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; indeed, by the

Skolem-Noether theorem, any two embeddings of $L \otimes_K K'$ into $\mathrm{M}_n(K')$ are conjugate. By [Me, Proposition 1.3], we conclude that $\mathrm{ed}(N;p) \le \mathrm{ed}(\mathrm{S}_n;p)$. Combining this with Corollary 4.2 yields the desired inequality $\mathrm{ed}(N;p) \le [n/p]$. □

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

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

Let $\Lambda = \{\mathbf{a}_{1,2}, \dots, \mathbf{a}_{p-1,p}, \mathbf{a}_{p,1}\}$ and $V = V_\Lambda \times L$, where L is a 1-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(b) tells us that V is a generically free representation of $T \rtimes P_n$.

Step (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,

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

see (2) and (3). This completes the proof of Proposition 5.1(b) and of Theorem 1.1(b). □

6. Proof of Theorem 1.1 part (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 $r \ge 2$ and follow the strategy of Section 2. In this section we will carry out Step (i). That is, we will construct construct a generically free representation V of $T \rtimes P_n$ of dimension p^{2r-1}. This will show that $\mathrm{ed}(T \rtimes P_n) \le p^{2r-1} - p^r + 1$. 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 $1, 2, \dots, p^r$ into p "big blocks" $B_1, \dots, B_p$, where each B_i consists of the p^{r-1} consecutive integers $(i-1)p^{r-1} + 1, (i-1)p^{r-1} + 2, \dots, ip^{r-1}$.

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

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

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 $\phi\colon \mathbb{Z}[\Lambda] \to X(T)$ is faithful.

The elements $\mathbf{a}_{\alpha,\beta}$ clearly generate $X(T)$ as an abelian group, as α and β range over $1, 2, \dots, p^r$. 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 $j - i \equiv 3, \dots, p \pmod p$, 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

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

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 non-trivial. Indeed, Z_n does not fix the non-zero element

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

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

7. Theorem 1.1 part (c): The lower bound

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

$$\mathrm{ed}(N; p) \geq p^{2r-1} - p^r + 1 , \tag{8}$$

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

$$q := p^e, \text{ where } e \geq 1 \text{ if } p \text{ is odd and } e \geq 2 \text{ if } p = 2. \tag{9}$$

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

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

cf. [Me, Proposition 1.5(1)]. Thus for the purpose of proving (8) we may replace k by k'. In particular, we may assume that k' contains a primitive qth 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 (7) 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

$$\mathrm{ed}(T_{(q)} \rtimes P_n; p) \geq p^{2r-1}\,. \tag{10}$$

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 recent result of Karpenko and Merkurjev [KM].

Theorem 7.1. *Let G be a finite p-group and k be a field containing a primitive pth root of unity. Then $\mathrm{ed}_k(G;p) = \mathrm{ed}_k(G)$ = 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 qth root of unity and hence, a primitive pth root of unity. Hence, Theorem 7.1 applies in our situation. That is, in order to prove (10) it suffices to show that $T_{(q)} \rtimes P_n$ does not have a faithful linear representation of dimension $< p^{2r-1}$. Lemma 3.1 further reduces this representation-theoretic assertion to the combinatorial statement of Proposition 7.2 below.

Before stating Proposition 7.2 we recall that the character lattice of $T_{(q)} \simeq \mu_q^n/\mu_q$ is

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

where we identify the character

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

of $T_{(q)}$ with $(a_1, \dots, a_n) \in (\mathbb{Z}/q\mathbb{Z})^n$. Here $(t_1, \dots, t_n)$ stands for an element of μ_q^n, modulo the diagonally embedded μ_q, so the above character is well defined if and only if $a_1 + \dots + a_n = 0$ 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 of Proposition 7.2 will rely on the following special case of Nakayama's Lemma [AM, Proposition 2.8].

Lemma 7.3. *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, $\Lambda \backslash pM$, 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 non-trivial 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 (9) of q, we have proved the proposition for $r = 1$.

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

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

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

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

$$\mathbf{a} = (a_1, a_2, \dots, a_{p^r}) \mapsto \mathbf{s} = (s_1, \dots, s_{p^{r-1}})\,, \tag{11}$$

where $s_i = a_{(i-1)p+1} + a_{(i-1)p+2} + \ldots + a_{ip}$ is the sum of the entries of $\mathbf{a}$ in the ith 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 $\mathbf{s} = (s_1, \dots, s_{p^{r-1}})$ 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

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

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, i.e., if

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

We are now ready to prove the claim. Suppose $\mathbf{a} = (a_1, \dots, a_{p^r}) \in X_{p^r}$ lies in the preimage of $\mathbf{s} = (s_1, \dots, s_{p^{r-1}})$, as in (11). 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

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

$\mathbf{s}$ 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 $\mathbf{s}$ cannot have exactly one entry not divisible by p.) Thus $\mathbf{a}$ is non-scalar 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 $\mathbf{s}$ 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). □

8. Proof of Theorem 1.1 part (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

$$n = n_1 p^{e_1} + n_2 p^{e_2} + ... + n_u p^{e_u}\,, \tag{12}$$

of n, where $1 \le e = e_1 < e_2 < ... < e_u$, and $1 \le n_i < p$ for each i. Subdivide the integers $1, ..., n$ into $n_1 + ... + n_u$ blocks B^i_j of length p^{e_i}, for j ranging over $1, 2, ..., n_i$. 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^i_j.

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

Step (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 $\mathrm{ed}_k(T \rtimes P_n) \le p^{e_1}(n - p^{e_1})$. 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

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

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 $\phi\colon \mathbb{Z}[\Lambda] \to X(T)$ 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 non trivial element of $Z_n = (Z_{p^{e_1}})^{n_1} \times \cdots \times (Z_{p^{e_u}})^{n_u}$. We may assume that the first component of σ in the above direct product is non-trivial, and therefore σ permutes elements in the first block B_1^1 cyclically. Note that B_1^1 is of size at least p as $e = e_1 \geq 1$, and that we have at least 2 blocks. The second block is also of size $\geq p$ and if $p = 2$, at least of size 4 by (12). It follows from this that σ does not fix the non-zero element

$$\mathbf{a}_{1,p^e+1} - \mathbf{a}_{1,p^e+2} + \mathbf{a}_{2,p^e+2} - \mathbf{a}_{2,p^e+1}$$

which lies in the kernel of ϕ.

Step (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 $\geq p^{e_1}(n - p^{e_1})$. 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*

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

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 (12) of n, with $0 \leq e_1 < e_2 < \ldots < e_u$ 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 $j = 1, \dots, n/p$. 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,

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

and thus

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

This completes the proof of Lemma 8.1 and thus of parts (d) of Proposition 5.1 and of Theorem 1.1. □

References

[AM] M. F. Atiyah, I. G. Macdonald, Introduction to commutative algebra. Addison-Wesley Publishing Co., 1969.

[Bo] N. Bourbaki, Algebra II. Chapters 4–7, Elements of Mathematics, Translated from the 1981 French edition by P. M. Cohn and J. Howie, Springer-Verlag, Berlin, 2003.

[BRV] P. Brosnan, Z. Reichstein, A. Vistoli, *Essential dimension, spinor groups and quadratic forms*, Annals of Math., to appear. Preprint available at `http://annals.math.princeton.edu/issues/2008/FinalFiles/BrosnanReichsteinVistoliFinal.pdf`

[BF] G. Berhuy, G. Favi, *Essential dimension: a functorial point of view (after A. Merkurjev)*, Doc. Math. **8** (2003), 279–330.

[BuR$_1$] J. Buhler, Z. Reichstein, *On the essential dimension of a finite group*, Compositio Math. **106** (1997), no. 2, 159–179.

[BuR$_2$] J. Buhler, Z. Reichstein, *On Tschirnhaus transformations*, in Topics in number theory (University Park, PA, 1997), 127–142, Math. Appl. **467**, Kluwer Acad. Publ., Dordrecht, 1999.

[DG] M. Demazure, P. Gabriel, *Groupes algébriques. Tome I: Géométrie algébrique, généralités, groupes commutatifs*, Masson & Cie, Éditeur, Paris, 1970.

[FF] G. Favi, M. Florence, *Tori and essential dimension*, J. Algebra, **319** (2008), no. 9, 3885-3900.

[Ja] J. C. Jantzen, *Representations of algebraic groups*, Mathematical Surveys and Monographs 107, AMS, Providence, 2003.

[KM] N. A. Karpenko and A. S. Merkurjev, *Essential dimension of finite p-groups*, Inventiones Math., **172**, no. 3 (2008), pp. 491-508.

[Le] N. Lemire, *Essential dimension of algebraic groups and integral representations of Weyl groups*, Transform. Groups **9** (2004), no. 4, 337–379.

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

[LR] M. Lorenz, Z. Reichstein, *Lattices and parameter reduction in division algebras*, MSRI Preprint 2000-001, `http://www.msri.org/publications/preprints/online/2000-001.html`

[LRRS] M. Lorenz, Z. Reichstein, L. H. Rowen, D. J. Saltman, *Fields of definition for division algebras*, J. London Math. Soc. (2) **68** (2003), no. 3, 651–670.

[Me] A. Merkurjev, *Essential dimension*, to appear in Proceedings of the International Conference on the algebraic and arithmetic theory of quadratic forms (Chile 2007), Contemporary Mathematics, American Mathematical Society, Providence, RI. Preprint posted at `http://www.math.ucla.edu/%7Emerkurev/publicat.htm`

[Pr] C. Procesi, *Non-commutative affine rings*, Atti Acc. Naz. Lincei, S. VIII, v. VIII, fo. 6 (1967), 239–255.

[Re$_1$] Z. Reichstein, *On a theorem of Hermite and Joubert*, Canad. J. Math. **51** (1999), no. 1, 69–95.

[Re$_2$] Z. Reichstein, *On the notion of essential dimension for algebraic groups*, Transform. Groups **5** (2000), no. 3, 265–304.

[RY] Z. Reichstein and B. Youssin, *Essential dimensions of algebraic groups and a resolution theorem for G-varieties*, with an appendix by János Kollár and Endre Szabó, Canad. J. Math. **52** (2000), no. 5, 1018–1056,

[RS] L. H. Rowen, D. J. Saltman, *Prime-to-p extensions of division algebras*, Israel J. Math. **78** (1992), no. 2-3, 197–207.

[Se$_1$] J.-P. Serre, Linear representations of finite groups, Graduate Texts in Mathematics, **42**. Springer-Verlag, New York–Heidelberg, 1977.

[Se$_2$] J.-P. Serre, Galois Cohomology, Springer, 1997.

Department of Mathematics, University of British Columbia, Vancouver, BC V6T 1Z2, Canada