ESSENTIAL DIMENSION AND ALGEBRAIC STACKS I

PATRICK BROSNAN[†], ZINOVY REICHSTEIN[†], AND ANGELO VISTOLI[‡]

Abstract. Essential dimension is a numerical invariant of an algebraic group G which reflects the complexity of G-torsors over a field K. In the past 10 years it has been studied by many authors in a variety of contexts. In this paper we extend this notion to algebraic stacks. As an application of the resulting theory we obtain new results about the essential dimension of certain algebraic (and, in particular, finite) groups which occur as central extensions. In particular, we show that the essential dimension of the spinor group Spin_n grows exponentially with n, improving previous lower bounds of Chernousov-Serre and Reichstein-Youssin. In the last section we apply the result on spinor groups to show that quadratic forms with trivial discriminant and Hasse-Witt invariant are more complex, in high dimensions, than previously expected.

Contents

1. Introduction

Let k be a field. We will write Fields_k for the category of field extensions K/k. Let $F\colon \mathrm{Fields}_k \to \mathrm{Sets}$ be a covariant functor.

2000 *Mathematics Subject Classification.* Primary 14A20, 20G15, 11E04.

[†] Supported in part by an NSERC discovery grants and by PIMS Collaborative Research Group in Algebraic geometry, cohomology and representation theory.

[‡] Supported in part by the PRIN Project "Geometria sulle varietà algebriche", financed by MIUR.

Definition 1.1. Let $a \in F(L)$, where L is an object of Fields_k. We say that *a descends* 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* $\operatorname{ed}(a)$ of $a \in F(L)$ is the minimum of the transcendence degrees $\operatorname{tr deg}_k K$ taken over all fields $k \subseteq K \subseteq L$ such that a descends to K.

The essential dimension $\operatorname{ed}(a; p)$ of a at a prime integer p is the minimum of $\operatorname{ed}(a_{L'})$, taken over all finite field extensions L'/L such that the degree $[L' : L]$ is prime to p.

The essential dimension $\operatorname{ed} F$ of the functor F (respectively, the essential dimension $\operatorname{ed}(F; p)$ of F at a prime p) is the supremum of $\operatorname{ed}(a)$ (respectively, of $\operatorname{ed}(a; p)$) taken over all $a \in F(L)$ with L in Fields_k.

These notions are relative to the base field k. To emphasize this, we will sometimes write $\operatorname{ed}_k(a)$, $\operatorname{ed}_k(a; p)$, $\operatorname{ed}_k F$ or $\operatorname{ed}_k(F; p)$ instead of $\operatorname{ed}(a)$, $\operatorname{ed}(a; p)$, $\operatorname{ed} F$ or $\operatorname{ed}(F; p)$, respectively.

If the functor F is limit-preserving, a condition that is satisfied in all cases that interest us, every element $a \in F(L)$ has a field of definition K that is finitely generated over k, so $\operatorname{ed}(a)$ is finite; cf. Remark 2.5. On the other hand, $\operatorname{ed} F$ may be infinite even in cases of interest.

The following example describes a class of functors which will play a key role in the sequel.

Example 1.2. Let G be an algebraic group. Consider the Galois cohomology functor $F = \mathrm{H}^1(*, G)$ sending K to the set $\mathrm{H}^1(K, G)$ of isomorphism classes of G-torsors over $\operatorname{Spec}(K)$, in the fppf topology. The essential dimension of this functor is a numerical invariant of G, which, roughly speaking, measures the complexity of G-torsors over fields. This number is usually denoted by $\operatorname{ed}_k G$ or (if k is fixed throughout) simply by $\operatorname{ed} G$. Essential dimension was originally introduced and has since been extensively studied in this context; see [BR97, Rei00, RY00, Kor00, Led02, JLY02, BF03, Lem04, CS06, Gar]. The more general Definition 1.1 is due to A. Merkurjev; see [BF03, Proposition 1.17].

In this paper we define the notion of essential dimension for algebraic stacks and apply it in the "classical" setting of Example 1.2. The stacks that play a special role in these applications are gerbes banded by the multiplicative group μ_n; see §§3-4. Before proceeding to state out main results, we remark that the notion of essential dimension introduced in this paper also leads to interesting results for other types of stacks. In particular, in the forthcoming paper [BRV07] we will compute the essential dimension of the moduli stack $\mathcal{M}_{g,n}$ of smooth curves of genus g for every $g \geq 0$. We also note that a related (but not equivalent) notion of *arithmetic dimension* has been studied by C. O'Neil [O'N05, O'N].

We will now describe the main results of this paper. Let

$$1 \longrightarrow Z \longrightarrow G \longrightarrow Q \longrightarrow 1 \tag{1.3}$$

denote an exact sequence of algebraic groups over a field k with Z central and isomorphic to μ_n for some integer $n > 1$. For every field extension K/k the sequence (1.3) induces a connecting map $\partial_K \colon \mathrm{H}^1(K, Q) \to \mathrm{H}^2(K, Z)$. We define $\operatorname{ind}(G, Z)$ as the maximal value of $\operatorname{ind}\bigl(\partial_K(t)\bigr)$, as K ranges over all field extensions of k and t ranges over all torsors in $\mathrm{H}^1(K, Q)$. (Note that $\operatorname{ind}(G, Z)$ does not depend on the choice of the isomorphism $Z \simeq \mu_n$.) In §5 we will prove the following inequality.

Theorem 1.4. *Let G be an extension as in* (1.3). *Assume that n a power of a prime integer p. Then* $\operatorname{ed}(G; p) \geq \operatorname{ind}(G, Z) - \dim G$.

In §6 we will use Theorem 1.4 to give an alternative proof of a recent theorem of Florence about the essential dimension of a cyclic group (in a slightly strengthened form). We will then use Florence's theorem to settle a particular case of a conjecture of Ledet [Led02, Section 3], relating the essential dimensions of the cyclic group C_n and the dihedral group D_n (n odd); see Corollary 6.2.

In §§7–9 we continue to study essential dimensions of finite p-groups. Let G be a finite abstract group. We write $\operatorname{ed}_k G$ (respectively, $\operatorname{ed}_k(G; p)$) for the essential dimension (respectively, for the essential dimension at p) of the constant group G_k over the field k. Let $\exp(G)$ denote the exponent of G and let $\mathrm{C}(G)$ denote the center of G.

Theorem 1.5. *Let G be a p-group whose commutator $[G, G]$ is central and cyclic. Then*

$$\operatorname{ed}_k(G; p) = \operatorname{ed}_k G = \sqrt{|G/\mathrm{C}(G)|} + \operatorname{rank} \mathrm{C}(G) - 1\,.$$

for any base field k of characteristic $\neq p$ which contains a primitive root of unity of degree $\exp(G)$.

Note that with the above hypotheses, $|G/\mathrm{C}(G)|$ is a complete square. In the case where G is abelian we recover the identity $\operatorname{ed} G = \operatorname{rank}(G)$; see [BR97, Theorem6.1]. For most finite groups G the best previously known lower bounds on $\operatorname{ed} G$ were of the form

$$\operatorname{ed} G \geq \operatorname{rank}(A)\,, \tag{1.6}$$

where A was taken to be an abelian subgroup A of G of maximal rank. Theorem 1.5 represents a substantial improvement over these bounds. For example, if G is a non-abelian group of order p^3 and k contains a primitive root of unity of degree p^2 then Theorem 1.5 tells us that $\operatorname{ed} G = p$, while (1.6) yields only $\operatorname{ed} G \geq 2$.

In §9 we record several interesting corollaries of Theorem 1.5. In particular, we show that $\operatorname{ed}(G; p) \geq p$ for any non-abelian p-group G; see Corollary 9.3. We also answer a question of Jensen, Ledet and Yui [JLY02, p.204] by giving an example of a finite group G with a normal subgroup N such that $\operatorname{ed}(G/N) > \operatorname{ed} G$; see Corollary 9.6.

Another consequence of Theorem 1.5 is the following new bound on $\operatorname{ed}\operatorname{Spin}_n$. Here by Spin_n we will mean the totally split form of the spin group in dimension n over a field k.

Theorem 1.7. *Suppose k is a field of characteristic $\neq 2$, and that $\sqrt{-1} \in k$. If n is not divisible by 4 then*

$$2^{\lfloor (n-1)/2 \rfloor} - \frac{n(n-1)}{2} \leq \operatorname{ed}(\operatorname{Spin}_n; 2) \leq \operatorname{ed}(\operatorname{Spin}_n) \leq 2^{\lfloor (n-1)/2 \rfloor} .$$

If n is divisible by 4 then

$$2^{\lfloor (n-1)/2 \rfloor} - \frac{n(n-1)}{2} + 1 \leq \operatorname{ed}(\operatorname{Spin}_n; 2) \leq \operatorname{ed}\operatorname{Spin}_n \leq 2^{\lfloor (n-1)/2 \rfloor} + 1.$$

Theorem 1.7 is proved in §10, where we also prove similar estimates for the essential dimensions of pin and half-spin groups.

The lower bound of Theorem 1.7 was surprising to us because previously the best known lower bound, due of V. Chernousov and J.–P. Serre [CS06], was

$$(1.8) \qquad \operatorname{ed}\operatorname{Spin}_n \geq \begin{cases} \lfloor n/2 \rfloor + 1 & \text{if } n \geq 7 \text{ and } n \equiv 1, 0 \text{ or } -1 \pmod 8 \\ \lfloor n/2 \rfloor & \text{for all other } n \geq 11. \end{cases}$$

(The first line is due to B. Youssin and the second author in the case that $\operatorname{char} k = 0$ [RY00].) Moreover, in low dimensions, M. Rost [Ros99] (cf. also [Gar]) computed the following table of exact values:

$$\begin{array}{cccc} \operatorname{ed}\operatorname{Spin}_3 = 0 & \operatorname{ed}\operatorname{Spin}_4 = 0 & \operatorname{ed}\operatorname{Spin}_5 = 0 & \operatorname{ed}\operatorname{Spin}_6 = 0 \\ \operatorname{ed}\operatorname{Spin}_7 = 4 & \operatorname{ed}\operatorname{Spin}_8 = 5 & \operatorname{ed}\operatorname{Spin}_9 = 5 & \operatorname{ed}\operatorname{Spin}_{10} = 4 \\ \operatorname{ed}\operatorname{Spin}_{11} = 5 & \operatorname{ed}\operatorname{Spin}_{12} = 6 & \operatorname{ed}\operatorname{Spin}_{13} = 6 & \operatorname{ed}\operatorname{Spin}_{14} = 7. \end{array}$$

Taken together these results seemed to suggest that $\operatorname{ed}\operatorname{Spin}_n$ should be a slowly increasing function of n and gave no hint of its exponential growth.

The computation of $\operatorname{ed}\operatorname{Spin}_n$ gives an example of a split, semisimple, connected linear algebraic group whose essential dimension exceeds its dimension. (Note that for a semisimple adjoint group G, $\operatorname{ed} G \leq \dim G$; cf. Example 10.10.) Since $\operatorname{ed}\operatorname{SO}_n = n - 1$ for every $n \geq 3$ (cf. [Rei00, Theorem 10.4]), it also gives an example of a split, semisimple, connected linear algebraic group G with a central subgroup Z such that $\operatorname{ed} G > \operatorname{ed} G/Z$.

In the last section we follow a suggestion of A. Merkurjev and B. Totaro to apply our results on $\operatorname{ed}\operatorname{Spin}_n$ to a problem in the theory of quadratic forms. Let K be a field of characteristic different from 2 containing a square root of -1, $\mathrm{W}(K)$ be the Witt ring of K and $I(K)$ be the ideal of classes of even-dimensional forms in $\mathrm{W}(K)$. It is well known that if q is a non-degenerate n-dimensional quadratic form whose class $[q]$ in $\mathrm{W}(K)$ lies in $I^a(K)$, then $[q]$ can be expressed as the class a sum of a-fold Pfister forms. It is natural to ask how many Pfister form are needed. When $a = 1$ or 2 it is easy to see that n Pfister forms always suffice; see Proposition 11.1. We prove the following result, which shows that the situation is quite different when $a = 3$.

Theorem 1.9. *Let k be a field of characteristic different from 2 and n an even positive integer. Then there is a field extension K/k and a class $[q] \in I^3(K)$ represented by an n-dimensional quadratic form q/K such that $[q]$ cannot be written as the sum of fewer than*

$$\frac{2^{(n+4)/4} - n - 2}{7}$$

3*-fold Pfister forms over K.*

Finally we remark that the stack-theoretic approach to computing $\operatorname{ed} G$ developed in this paper has been recently extended and refined by Karpenko and Merkurjev [KM07]. Their main result is a general formula for the essential dimension of any finite p-group over any field k containing a primitive pth root of unity. This formula may be viewed as a common generalization of our Theorems 1.5 and 6.1.

Notation. Throughout this paper, a *variety* over a field k will be a geometrically integral separated scheme of finite type over k. Cohomology groups $\mathrm{H}^i(X, G)$, where X is a scheme and G is an abelian group scheme over X, will be taken with respect to the fppf topology unless otherwise specified. We will write μ_n for the group scheme of n-th roots of unity. If k is a field whose characteristic is prime to n, ζ_n will denote a primitive n-th root of unity in the algebraic closure of k. (Using the axiom of choice, we choose the ζ_n once and for all.) For typographical reasons, we sometimes write C_n for the cyclic group $\mathbb{Z}/n$.

Acknowledgments. We would like to thank the Banff International Research Station in Banff, Alberta (BIRS) for providing the inspiring meeting place where this work was started. We are grateful to K. Behrend, C.-L. Chai, D. Edidin, N. Fakhruddin, M. Florence, A. Merkurjev, B. Noohi, G. Pappas, D. Saltman and B. Totaro for helpful conversations.

2. The essential dimension of a stack: definition and first properties

We now return to the general setting of Definition 1.1. As we mentioned in Example 1.2 most of the existing theory is specific to the Galois cohomology functors $F = \mathrm{H}^1(*, G)$ for various algebraic k-groups G. On the other hand, many naturally arising functors where the essential dimension is of interest are not of the form $F(K) = \mathrm{H}^1(K, G)$ for any algebraic group G. Two such examples are given below. In this section we identify a class of functors which is sufficiently broad to include most such examples, yet "geometric" enough to allow one to get a handle on their essential dimension. These functors are isomorphism classes of objects of an algebraic stack; see Definition 2.3 below. As we shall see in the sequel, these "new" functors turns out to be very useful even if one is only interested in the “classical” setting of Example 1.2.

Example 2.1. Let X/k be a scheme of finite type over a field k, and let $F_X\colon \mathrm{Fields}_k \to \mathrm{Sets}$ denote the functor given by $K \mapsto X(K)$. Then an easy argument due to Merkurjev shows that $\operatorname{ed} F_X = \dim X$; see [BF03, Proposition 1.17].

In fact, this equality remains true for any algebraic space X. Indeed, an algebraic space X has a stratification by schemes X_i. Any K-point $\eta\colon \operatorname{Spec} K \to X$ must land in one of the X_i. Thus $\operatorname{ed} X = \max \operatorname{ed} X_i = \dim X$. ♠

Example 2.2. Let Curves_g be the functor that associates to a field K/k the set of isomorphism classes of smooth algebraic curves of genus g over K. The essential dimension of this functor is computed in [BRV07].

We are now ready to give the main definition of this section.

Definition 2.3. Suppose $\mathcal{X}$ is an algebraic stack over k. The *essential dimension* $\operatorname{ed} \mathcal{X}$ of $\mathcal{X}$ (respectively, the essential dimension $\operatorname{ed}(\mathcal{X}; p)$ of $\mathcal{X}$ at a prime integer p) is the essential dimension (respectively, the essential dimension at p) of the functor $F_{\mathcal{X}}\colon \mathrm{Fields}_k \to \mathrm{Sets}$ which sends a field L/k to the set of isomorphism classes of objects in $\mathcal{X}(L)$.[1]

As in Definition 1.1, we will write $\operatorname{ed}(\mathcal{X}/k)$, or $\operatorname{ed}_k \mathcal{X}$ when we need to be specific about the dependence on the base field k. Similarly for $\operatorname{ed}(\xi/k)$ or $\operatorname{ed}_k \xi$, where ξ is an object of $F_{\mathcal{X}}$, and for $\operatorname{ed}(\mathcal{X}/k; p)$, $\operatorname{ed}_k(\mathcal{X}; p)$, $\operatorname{ed}(\xi/k; p)$, $\operatorname{ed}_k(\xi; p)$.

All of the examples we have considered so far may be viewed as special cases of Definition 2.3. If $\mathcal{X}$ is a scheme of finite type (or an algebraic space), we recover Example 2.1. At the other extreme, if $\mathcal{X} = \mathcal{B}G$ is the classifying stack of an agebraic group G defined over k (so that $\mathcal{B}G(T)$ is the category of G-torsors on T), we recover Example 1.2. If $\mathcal{X} = \mathcal{M}_g$ is the stack of smooth algebraic curves of genus g, we recover Example 2.2.

Remark 2.4. If G is an algebraic group, we will often write $\operatorname{ed} G$ for $\operatorname{ed} \mathcal{B}G$. That is, we will write $\operatorname{ed} G$ for the essential dimension of the stack $\mathcal{B}G$ and not the essential dimension of the scheme underlying G. We do this to conform to the, now standard, notation introduced at the beginning of this paper (and in Example 1.2). Of course, by Example 2.1, the essential dimension of the underlying scheme is $\dim G$.

In fact, the reader may notice that we prefer to write $\operatorname{ed} \mathcal{B}G$ earlier in the paper, where we are working in a general stack-theoretic setting and $\operatorname{ed} G$ later on, where we are primarily concerned with essential dimensions of algebraic groups.

[1] In the literature the functor $F_{\mathcal{X}}$ is sometimes denoted by $\widehat{\mathcal{X}}$ or $\overline{\mathcal{X}}$.

Remark 2.5. $\operatorname{ed}\mathcal{X}$ takes values in the range $\{\pm\infty\}\cup\mathbb{Z}_{\geq 0}$ (with $-\infty$ occurring if and only if $\mathcal{X}$ is empty). On the other hand, $\operatorname{ed}\xi$ is finite for every object $\xi \in \mathcal{X}(K)$ and every field extension K/k.

Indeed, any field K/k can be written as a filtered direct limit $K = \operatorname{colim}_I K_i$ of its subfields K_i of finite transcendence degree. Since $\mathcal{X}$ is limit preserving (cf. [LMB00, Proposition 4.18]), ξ lies in the essential image of $\mathcal{X}(K_i) \to \mathcal{X}(K)$ for some $i \in I$ and thus $\operatorname{ed}\xi \leq \operatorname{tr}\operatorname{deg}_k(K_i) < \infty$.

We now recall Definitions (3.9) and (3.10) from [LMB00]. A morphism $f\colon \mathcal{X} \to \mathcal{Y}$ of algebraic stacks (over k) is said to be *representable* if, for every k-morphism $T \to \mathcal{Y}$, where T is an affine k-scheme, the fiber product $\mathcal{X}\times_{\mathcal{Y}}T$ is representable by a scheme over T. A representable morphism $f\colon \mathcal{X} \to \mathcal{Y}$ is said to be *locally of finite type and of fiber dimension* $\leq d$ if the projection $\mathcal{X} \times_{\mathcal{Y}} T \to T$ is locally of finite type over T and every fiber has dimension $\leq d$.

Proposition 2.6. *Let d be an integer $\mathcal{X} \to \mathcal{Y}$ be a representable k-morphism of algebraic stacks which is locally of finite type and of fiber dimension at most d. Let L/k be a field, $\xi \in X(L)$ and p be a prime integer. Then*

(a) $\operatorname{ed}_k \xi \leq \operatorname{ed}_k(f(\xi)) + d$,

(b) $\operatorname{ed}_k(\xi; p) \leq \operatorname{ed}_k(f(\xi); p) + d$,

(c) $\operatorname{ed}_k \mathcal{X} \leq \operatorname{ed}_k \mathcal{Y} + d$,

(d) $\operatorname{ed}_k(\mathcal{X}; p) \leq \operatorname{ed}_k(\mathcal{Y}; p) + d$.

Proof. (a) By the definition of $\operatorname{ed}_k(f(\xi))$ we can find an intermediate field $k \subset K \subset L$ and a morphism $\eta\colon \operatorname{Spec} K \to Y$ such that $\operatorname{tr}\operatorname{deg}_k L \leq \operatorname{ed} f(\xi)$ and the following diagram commutes.

$$\begin{tikzcd}
\operatorname{Spec} L \arrow[r, "\xi"] \arrow[d] & \mathcal{X} \arrow[d, "f"] \\
\operatorname{Spec} K \arrow[r, "\eta"] & \mathcal{Y}
\end{tikzcd}$$

Let $\mathcal{X}_K \overset{\text{def}}{=} \mathcal{X} \times_{\mathcal{Y}} \operatorname{Spec} K$. By the hypothesis, $\mathcal{X}_K$ is an algebraic space, locally of finite type over K and of relative dimension at most d. By the commutativity of the above diagram, the morphism $\xi\colon \operatorname{Spec} L \to \mathcal{X}$ factors through $\mathcal{X}_L$:

$$\begin{tikzcd}
\operatorname{Spec} L \arrow[dr, "\xi_0"] \arrow[drr, bend left, "\xi"] \arrow[ddr, bend right] & & \\
& \mathcal{X}_K \arrow[r] \arrow[d] & \mathcal{X} \arrow[d, "f"] \\
& \operatorname{Spec} K \arrow[r, "\eta"] & \mathcal{Y}
\end{tikzcd}$$

Let p denote the image of ξ_0 in $\mathcal{X}_K$. Since $\mathcal{X}_K$ has dimension at most d, we have $\operatorname{tr\,deg}_k k(p) \leq d$. Therefore $\operatorname{tr\,deg}_k k(p) \leq \operatorname{ed}(f(\xi)) + d$. Since ξ factors through $\operatorname{Spec} k(p)$, part (a) follows.

(b) Let L'/L be a field extension. By part (a), $\operatorname{ed}_k(\xi_{L'}) \leq \operatorname{ed}_k(f(\xi_{L'}))+d$. Taking the minimum over all prime-to p extensions L'/L, we obtain the desired inequality, $\operatorname{ed}_k(\xi; p) \leq \operatorname{ed}_k(f(\xi); p) + d$.

(c) follow from (a) and (d) follows from (b) by taking the maximum on both sides over all L/k and all $\xi \in \mathcal{X}(L)$. ♠

Corollary 2.7. *Let G be an algebraic group defined over k and H be a closed subgroup of G. Then*

(a) $\operatorname{ed} H \leq \operatorname{ed} G + \dim G - \dim H$ *and*

(b) $\operatorname{ed}(H; p) \leq \operatorname{ed}(G; p) + \dim G - \dim H$.

More generally, suppose G is acting on an algebraic space X (over k). Then

(c) $\operatorname{ed}\,[X/H] \leq \operatorname{ed}\,[X/G] + \dim G - \dim H$ *and*

(d) $\operatorname{ed}([X/H]; p) \leq \operatorname{ed}([X/G]; p) + \dim G - \dim H$.

Here $[X/G]$ and $[X/H]$ denote the quotient stacks for the actions of G and H on X.

Recall that the objects in $[X/G](K)$ are, by definition, diagrams of the form

$$\begin{tikzcd} T \arrow[r, "f"] \arrow[d, "\pi"] & X \\ \operatorname{Spec}(K), & \end{tikzcd}$$

where π is a G-torsor and f is a G-equivariant morphism.

Proof. The natural morphism $[X/H] \to [X/G]$ of quotient stacks is easily seen to be representable, of finite type and with fibers of dimension $d = \dim G - \dim H$. Applying Proposition 2.6(c) and (d) to this morphism, we obtain the inequalities (c) and (d) respectively. Specializing X to a point (i.e., to $\operatorname{Spec}(k)$) with trivial G-action, we obtain (a) and (b). ♠

In general, the inequality of Proposition 2.6 only goes in one direction. However, equality holds in the following important special case. We will say that a morphism $f\colon \mathcal{X} \to \mathcal{Y}$ of stacks is *isotropic* if for every field extension K of k and every object η of $\mathcal{Y}(K)$ there exists an object ξ of $\mathcal{X}(K)$ such that $f(\xi)$ is isomorphic to η.

Proposition 2.8. *Let $f\colon \mathcal{X} \to \mathcal{Y}$ be an isotropic morphism of stacks fibered over* Points_k. *Then* $\operatorname{ed} \mathcal{X} \geq \operatorname{ed} \mathcal{Y}$ *and* $\operatorname{ed}(\mathcal{X}; p) \geq \operatorname{ed}(\mathcal{Y}; p)$ *for every prime p.*

Proof. Both inequalities are immediate consequences of the following observation.

Let K be an extension of k and η an object of $\mathcal{Y}(K)$. If ξ is an object of $\mathcal{X}(K)$ such that $f(\xi)$ is isomorphic to η, then clearly a field of definition for ξ is also a field of definition for η. Thus $\operatorname{ed}\xi \geq \operatorname{ed}\eta$. ♠

Proposition 2.8 can be used to deduce the following well-known inequality. See [BF03] for another proof.

Theorem 2.9. *Let G be a linear algebraic group over a field k admitting a generically free representation on a vector space V. Then*

$$\operatorname{ed}\mathcal{B}G \leq \dim V - \dim G.$$

Proof. Let U denote a dense G-stable Zariski open subscheme of V on which G acts freely. Then $[U/G]$ is an algebraic space of dimension $\dim V - \dim G$ and the map $[U/G] \to \mathcal{B}G$ is representable and isotropic. ♠

3. Preliminaries on gerbes and canonical dimension

Let $\mathcal{X}$ be a gerbe defined over a field K *banded* by an abelian K-group scheme G. For background material on gerbes we refer the reader to [Mil80, p. 144] and [Gir71, IV.3.1.1].

There is a natural notion of equivalence of gerbes banded by G; the set of equivalence classes is in a natural bijective correspondence with the group $\mathrm{H}^2(K, G)$. Given a gerbe $\mathcal{X}$ banded by G, we write $[\mathcal{X}]$ for its class in $\mathrm{H}^2(K, G)$. The identity is the class $[\mathcal{B}_K G]$ of the neutral gerbe $\mathcal{B}_K G$.

We remark that it makes sense to talk about a gerbe banded by G, where G is not necessarily abelian, but we will not need to work in this more general setting, which makes the definition considerably more involved.

Let K be a field and let $\mathbb{G}_m$ denote the multiplicative group scheme over K. Recall that the group $\mathrm{H}^2(K, \mathbb{G}_m)$ is canonically isomorphic to the Brauer group $\mathrm{Br}(K)$ of Brauer equivalence classes of central simple algebras (CSAs) over K. By Wedderburn's structure theorem, any CSA over K isomorphic to the matrix algebra $\mathrm{M}_n(D)$ for D a division algebra over K which is unique up to isomorphism. Moreover, if A and B are two Brauer equivalent CSAs, the division algebras D and E corresponding to A and B respectively are isomorphic. For a class $[A] \in \mathrm{Br}(K)$, the *index* of A is $\sqrt{\dim_K D}$.

Let $\alpha \in \mathrm{H}^2(K, \mu_n)$, where n is a positive integer prime to $\operatorname{char} K$. We define the *index* $\operatorname{ind}\alpha$ to be the index of the image on α under the composition

$$\mathrm{H}^2(K, \mu_n) \hookrightarrow \mathrm{H}^2(K, \mathbb{G}_m) \xrightarrow{\cong} \mathrm{Br}(K).$$

Note that the index of α is the smallest integer d such that α is in the image of the (injective) connecting homomorphism

$$\partial \colon \mathrm{H}^1(K, \mathrm{PGL}_d) \longrightarrow \mathrm{Br}(K). \tag{3.1}$$

arising from the short exact sequence

$$1 \longrightarrow \mathbb{G}_m \longrightarrow \mathrm{GL}_d \longrightarrow \mathrm{PGL}_d \longrightarrow 1. \tag{3.2}$$

The *exponent* $\operatorname{ord}([A]))$ of a class $[A] \in \operatorname{Br} K$ is defined to be its order in the Brauer group. The exponent $\operatorname{ord}([A]))$ always divides the index $\operatorname{ind}([A]))$ [Her68, Theorem 4.4.5].

In the next section, we address the problem of computing the essential dimension of a μ_d-gerbe over a field K. Our computation will rely, in a key way, on the notion of canonical dimension introduced in [BR05] and [KM06], which we shall now recall.

Let X be a smooth projective variety defined over a field K. We say that L/K is a *splitting field* for X if $X(K) \neq \emptyset$. A splitting field L/K is called *generic* if for every splitting field L_0/K there exists a K-place $L \to L_0$. The canonical dimension of X is defined as the minimal value of $\operatorname{tr deg}_K(L)$, where L/K ranges over all generic splitting fields. Note that the function field $L = K(X)$ is a generic splitting field of X; see [KM06, Lemma 4.1]. In particular, generic splitting fields exist and $\operatorname{cdim}(X)$ is finite.

If X is a K-scheme and p is a prime integer, we also recall the relative variant of this notion from [KM06]. We say that L/K is a *p-generic splitting field* for X if for every splitting field L_0/K there is a prime-to-p extension L_1/L and a K-place $L \to L_1$. The minimal value of $\operatorname{tr deg}_K(L')$, as L' ranges over all p-generic splitting fields, is called the p-canonical dimension of X and is denoted by $\operatorname{cdim}_p(X)$. (Here, as usual, by a prime-to-p extension we mean a finite field extension whose degree is prime to p.)

Lemma 3.3. *If L/K is a p-generic extension for X then*

(a) any intermediate extension M/K, where $K \subset M \subset L$ is also p-generic, and

(b) any finite prime-to-p extension L'/K is also p-generic.

Proof. (a) Let L_0/K be a splitting field for X. Since L is generic, there is a prime-to-p extension L_1/L_0 and a K-place $L \to L_1$. Restricting this place to M, we obtain a desired place $M \to L_1$.

Part (b) is an immediate consequence of [KM06, Lemma 3.2]. ♠

The *determination functor* $D_X : \operatorname{Fields}_K \to \operatorname{Sets}$ is defined as follows: D associates to a field L/K the empty set, if $X(L) = \emptyset$, and a set consisting of one point, which we will denote by $a(L)$, if $X(L) \neq \emptyset$. The natural map $D(L_1) \to D(L_2)$ is then uniquely determined for any $K \subset L_1 \subset L_2$.

Lemma 3.4. *Let X be a complete regular K-variety then*

(a) $\operatorname{cdim}(X) = \operatorname{ed}(D_X)$.

(b) $\operatorname{cdim}_p(X) = \operatorname{ed}(D_X; p)$.

Proof. (a) Let L be a generic splitting field for X. By [KM06, Lemma 2.1], any subfield of L containing F is also a generic splitting field. Therefore $\operatorname{ed} D_X \geq \operatorname{ed} a(L) \geq \operatorname{cdim}(X)$.

To prove the opposite inequality, it suffices to show that $\operatorname{ed} a(F) \leq \operatorname{tr deg}_K L$ for any splitting field F/K and any generic splitting field L/K of X. In other words, given F and L as above, we want to construct an intermediate splitting field $K \subset F_0 \subset F$ with $\operatorname{tr}\deg_K(F_0) \leq \operatorname{tr}\deg_K(L)$. By the definition of a generic splitting field, there is a place $L \to F$. That is, there is a valuation ring R with quotient field L and a local homomorphism $f\colon R \to F$. We claim that the residue field $F_0 = R/m$, where m is the maximal ideal of R, has the desired properties. Clearly $K \subset F_0 \subset F$, and $\operatorname{tr}\deg_K(F_0) \leq \operatorname{tr}\deg_K(R) = \operatorname{tr}\deg_K(L)$, so we only need to check that F_0 is a splitting field for X.

Since L is itself a splitting field, we have a diagram

$$\begin{tikzcd}
\operatorname{Spec} L \arrow[r] \arrow[d] & X \arrow[d] \\
\operatorname{Spec} R \arrow[r] & \operatorname{Spec} K.
\end{tikzcd}$$

By the valuative crition, we obtain a map $\operatorname{Spec} R \to X$ and hence, a map $\operatorname{Spec} F_0 = \operatorname{Spec} R/m \to X$. Thus F_0 is splitting field for X, as claimed.

(b) Let L be a p-generic splitting field for X. Lemma 3.3 tells us that if L'/L is a prime-to-p extension and $K \subset M \subset L'$ is an intermediate splitting field then M is again p-generic. Thus $\operatorname{ed}(a(L);p) =$ minimal value of $\operatorname{tr}\deg_K(M) \geq \operatorname{cdim}_p(X)$.

To prove the opposite inequality, we need to show that $\operatorname{ed}(a(F);p) \leq \operatorname{tr}\deg_K(L)$ for any splitting field F/K and any p-generic splitting field L/K. Indeed, after replacing F by a prime-to-p-extension F', we obtain a K-place $L \to F'$. Now the same argument as in part (a) shows that there is an intermediate splitting subfield $K \subset F_0 \subset F'$ such that $\operatorname{tr}\deg_K(F_0) \leq \operatorname{tr}\deg_K(L)$. Thus

$$\operatorname{ed}(a(F);p) \leq \operatorname{ed}(a(F')) \leq \operatorname{tr}\deg_K(F_0) \leq \operatorname{tr}\deg_K(L)\,,$$

as claimed. ♠

Of particular interest to us will be the case where X is a Brauer-Severi variety over K. Let m be the index of X. If $m = p^a$ is a prime power then

$$\operatorname{cdim}(X) = \operatorname{cdim}_p(X) = p^a - 1\,; \tag{3.5}$$

see [KM06, Example 3.10] or [BR05, Theorem 11.4]. If the highest power of p dividing m is p^a then

$$\operatorname{cdim}_p(X) = p^a - 1\ ; \tag{3.6}$$

see [KM06, Example 5.10]. On the other hand, if m is divisible by more than one prime, $\operatorname{cdim}(X)$ is not known in general. Suppose $m = p_1^{a_1} \dots p_r^{a_r}$. Then the class of P in $\operatorname{Br} L$ is the sum of classes $\alpha_1, \dots, \alpha_r$ whose indices are $p_1^{a_1}, \dots, p_r^{a_r}$. Denote by $X_1, \dots, X_r$ the Brauer–Severi varieties with classes $\alpha_1, \dots, \alpha_r$. It is easy to see that $K(X_1 \times \dots \times X_r)$ is a generic splitting field for X. Hence,

$$\operatorname{cdim}(X) \leq \dim(X_1 \times \dots \times X_r) = p_1^{a_1} + \dots + p_r^{a_r} - r\,.$$

In [CTKM06], Colliot-Thélène, Karpenko and Merkurjev conjectured that equality holds, i.e.,

$$\operatorname{cdim}(X) = p_1^{a_1} + \cdots + p_r^{a_r} - r \tag{3.7}$$

for all $m \geq 2$. As we mentioned above, this in known to be true of m is a prime power (i.e., $r = 1$). Colliot-Thélène, Karpenko and Merkurjev also proved (3.7) for $m = 6$; see [CTKM06, Theorem 1.3].

4. The essential dimension of a μ_d-gerbe

We are now ready to proceed with our main theorem on gerbes.

Theorem 4.1. *Let d be an integer with $d > 1$. Let K be a field and $x \in \mathrm{H}^2(K, \mu_d)$, where $\operatorname{char} K$ does not divide d. Denote the image of x in $\mathrm{H}^2(K, \mathbb{G}_m)$ by y, the μ_d-gerbe associated to x by $\mathcal{X}$, the $\mathbb{G}_m$-gerbe associated to y by $\mathcal{Y}$, and the Brauer–Severi variety associated to y by P. Then*

(a) $\operatorname{ed}(\mathcal{Y}) = \operatorname{cdim} P$ *and (b)* $\operatorname{ed} \mathcal{X} = \operatorname{cdim} P + 1$.

Moreover, if the index of x is a prime power p^r then

(c) $\operatorname{ed}(\mathcal{Y}; p) = \operatorname{ed}(\mathcal{Y}) = p^r - 1$ *and (d)* $\operatorname{ed}(\mathcal{X}; p) = \operatorname{ed} \mathcal{X} = p^r$.

Proof. The functor $F_{\mathcal{Y}} : \mathrm{Fields}_K \to \mathrm{Sets}$ sends a field L/K to the empty set, if $P(L) = \emptyset$, and to a set consisting of one point, if $P(L) \neq \emptyset$. In other words, F_Y is the determination functor D_P introduced in the previous section. Thus $\operatorname{ed}(\mathcal{Y}) = \operatorname{ed}(D_P)$ and $\operatorname{ed}(\mathcal{Y}; p) = \operatorname{ed}(D_P; p)$. On the other hand, by Lemma 3.4, $\operatorname{ed}(D_P) = \operatorname{cdim}(P)$ and $\operatorname{ed}(D_P; p) = \operatorname{cdim}_p(P)$. This, in combination with (3.5), proves parts (a) and (c).

(b) First note that the natural map $\mathcal{X} \to \mathcal{Y}$ is finite type and representable of relative dimension 1. By Proposition 2.6 this implies that $\operatorname{ed} \mathcal{X} \leq \operatorname{ed}(\mathcal{Y}) + 1$. By part (a) it remains to prove the opposite inequality, $\operatorname{ed} \mathcal{X} \geq \operatorname{ed}(\mathcal{Y}) + 1$. We will do this by constructing an object α of $\mathcal{X}$ whose essential dimension is $\geq \operatorname{ed}(\mathcal{Y}) + 1$.

We will view $\mathcal{X}$ as a torsor for $\mathcal{B}\mu_d$ in the following sense: One has maps

$$\mathcal{X} \times \mathcal{B}\mu_d \to \mathcal{X}$$
$$\mathcal{X} \times \mathcal{X} \to \mathcal{B}\mu_d$$

satisfying various compatibilities, where the first map is the action of $\mathcal{B}\mu_d$ on $\mathcal{X}$ and the second map is the "difference" of two objects of $\mathcal{X}$. For the definition and a discussion of the properties of these maps, see [Gir71, Chapter IV, Sections 2.3, 2.4 and 3.3]. (Note that, in the notation of Giraud's book, $\mathcal{X} \wedge \mathcal{B}\mu_d \cong \mathcal{X}$ and the action operation above arises from the map $\mathcal{X} \times \mathcal{B}\mu_d \to \mathcal{X} \wedge \mathcal{B}\mu_d$ given in Chapter IV, Proposition 2.4.1. The "difference" operation — which we will not use here — arises similarly from the fact that, in Giraud's notation, $\mathrm{HOM}(\mathcal{X}, \mathcal{X}) \cong \mathcal{B}\mu_d$.)

Let $L = K(P)$ be the function field of P. Since L splits P, we have a natural map $a : \operatorname{Spec} L \to \mathcal{Y}$. Moreover since L is a generic splitting field for P,

$$\operatorname{ed}(a) = \operatorname{cdim}(P) = \operatorname{ed}(\mathcal{Y}), \tag{4.2}$$

$$\operatorname{ed}(a;p) = \operatorname{cdim}_p(P) = \operatorname{ed}(\mathcal{Y};p). \tag{4.3}$$

where we view a as an object in Y. Non-canonically lift $a : \operatorname{Spec} L \to \mathcal{Y}$ to a map $\operatorname{Spec} L \to \mathcal{X}$ using the fact that $\mathcal{X} \to \mathcal{Y}$ is isotropic. Let $\operatorname{Spec} L(t) \to \mathcal{B}\mu_d$ denote the map classified by $(t) \in \mathrm{H}^1(L(t), \mu_d) = L(t)^\times / L(t)^{\times d}$. Composing these two maps, we obtain an object

$$\alpha : \operatorname{Spec} L(t) \to \mathcal{X} \times \mathcal{B}\mu_d \to \mathcal{X}.$$

in $\mathcal{X}(L(t))$. Our goal is to prove that $\operatorname{ed}(\alpha) \geq \operatorname{ed}(\mathcal{Y}) + 1$. In other words, given a diagram of the form

$$\begin{tikzcd} \operatorname{Spec} L(t) \arrow[r, "\alpha"] \arrow[d] & \mathcal{X} \\ \operatorname{Spec} M \arrow[ur, "\beta"] & \end{tikzcd} \tag{4.4}$$

where $K \subset M \subset L$ is an intermediate field, we want to show that $\operatorname{tr}\deg_K(M) \geq \operatorname{ed}(\mathcal{Y})+1$. Assume the contrary: there is a diagram as above with $\operatorname{tr}\deg_K(M) \leq \operatorname{ed}(\mathcal{Y})$. Let $\nu \colon L(t)^* \to \mathbb{Z}$ be the usual discrete valuation corresponding to t and consider two cases.

Case 1. Suppose the restriction $\nu_{|M}$ of ν to M is non-trivial. Let M_0 denote the residue field of ν and $M_{\geq 0}$ denote the valuation ring. Since $\operatorname{Spec} M \to \mathcal{X} \to \mathcal{Y}$, there exists an M-point of P. Then by the valuative criterion of properness for P, there exists an $M_{\geq 0}$-point and thus an M_0-point of P. Passing to residue fields, we obtain the diagram

$$\begin{tikzcd} \operatorname{Spec} L \arrow[r, "a"] \arrow[d] & \mathcal{Y} \\ \operatorname{Spec} M_0 \arrow[ur] & \cdot \end{tikzcd}$$

which shows that $\operatorname{ed}(a) \leq \operatorname{tr}\deg_K M_0 = \operatorname{tr}\deg_K M - 1 \leq \operatorname{ed}(\mathcal{Y}) - 1$, contradicting (4.2).

Case 2. Now suppose the restriction of ν to M is trivial. The map $\operatorname{Spec} L \to \mathcal{X}$ sets up an isomorphism $\mathcal{X}_L \cong \mathcal{B}_L\mu_d$. The map $\operatorname{Spec} L(t) \to \mathcal{X}$ factors through $\mathcal{X}_L$ and thus induces a class in $\mathcal{B}\mu_d(L(t)) = \mathrm{H}^1(L(t), \mu_d)$. This class is (t). Tensoring the diagram (4.4) with L over K, we obtain

$$\begin{tikzcd} \operatorname{Spec} L(t) \otimes L \arrow[r, "\alpha"] \arrow[d] & \mathcal{X}_L \cong \mathcal{B}_L\mu_d \\ \operatorname{Spec} M \otimes L \arrow[ur, "\beta"] & \end{tikzcd}$$

Recall that $L = K(P)$ is the function field of P. Since P is absolutely irreducible, the tensor products $L(t) \otimes L$ and $M \otimes L$ are fields. The map

$\operatorname{Spec} M \otimes L \to \mathcal{B}_L\mu_d$ is classified by some $m \in (M \otimes L)^\times/(M \otimes L)^{\times d} = \mathrm{H}^1(M \otimes L, \mu_d)$. The image of m in $L(t) \otimes L$ would have to be equal to t modulo d-th powers. We will now derive a contradiction by comparing the valuations of m and t.

To apply the valuation to m, we lift ν from $L(t)$ to $L(t) \otimes L$. That is, we define ν_L as the valuation on $L(t) \otimes L = (L \otimes L)(t)$ corresponding to t. Since $\nu_L(t) = \nu(t) = 1$, we conclude that $\nu_L(m) \equiv 1 \pmod d$. This shows that ν_L is not trivial on $M \otimes L$ and thus ν is not trivial on M, contradicting our assumption. This contradiction completes the proof of part (b).

The proof of part (d) proceeds along similar lines, except for a small complication in prime characteristic, which is resolved by Lemma 4.7 below. For the sake of completeness we outline the argument.

Once again, applying Proposition 2.6 to the representable projection map $\mathcal{X} \to \mathcal{Y}$ of relative dimension 1, we obtain the inequality $\operatorname{ed}(\mathcal{X}; p) \le \operatorname{ed}(\mathcal{Y}; p) + 1$. By part (c) it remains to show that $\operatorname{ed}(\mathcal{X}; p) \ge \operatorname{ed}(\mathcal{Y}; p) + 1$. We will do this by showing that

$$\operatorname{ed}(\alpha; p) \ge \operatorname{ed}(\mathcal{Y}; p) + 1\,, \tag{4.5}$$

where $\alpha\colon \operatorname{Spec} L(t) \to \mathcal{X}$ is the same object in $\mathcal{X}(L(t))$ we constructed in the proof of part (b). Here we continue to denote the function field of P by $L = K(P)$. Once again, assume the contrary: there exists a finite field extension $L'/L(t)$ of degree prime to p, such that $\operatorname{ed}(\alpha_{L'}) \le \operatorname{ed}(\mathcal{Y}; p)$. Let L'_s be the separable closure of $L(t)$ in L'. By Lemma 4.7 $\operatorname{ed}(\alpha_{L'_s}) = \operatorname{ed}(\alpha_{L'})$. Hence, after replacing L' by L'_s, we may assume that L' is a finite separable extension of $L(t)$ of degree prime to p. The assumption that $\operatorname{ed}(\alpha_{L'}) \le \operatorname{ed}(\mathcal{Y}; p)$ means, by definition, that there exists a commutative diagram

$$\begin{tikzcd}
\operatorname{Spec}(L') \arrow[dd, "\text{prime-to-}p"'] \arrow[dr] & \\
& \operatorname{Spec} M \arrow[d, "\beta"] \\
\operatorname{Spec} L(t) \arrow[r, "\alpha"] & \mathcal{X}
\end{tikzcd} \tag{4.6}$$

where $K \subset M \subset L'$ is an intermediate field and $\operatorname{tr\,deg}_K(M) \le \operatorname{ed}(\mathcal{Y}; p)$.

Let $\nu\colon L(t)^* \to \mathbb{Z}$ be the discrete valuation corresponding to t and $\nu'_1, \dots, \nu'_r$ be the liftings of ν to L'. Since L'/L is separable, [Lan65, Proposition XII.18] tells us that

$$e_1 f_1 + \cdots + e_r f_r = [L' : L]\,,$$

where e_i and f_i are, respectively, the ramification index and the residue class degree of ν'_i. Since $[L' : L]$ is prime to p, at least one of the valuations $\nu'_1, \dots, \nu'_r$, say $\nu'\colon (L')^* \to \frac{1}{e}\mathbb{Z}$, has the property that its ramification index e and residue class degree f are both prime to p.

Once again we consider two cases. If the restriction of ν' to M is non-trivial then after passing to residue fields and arguing as in Case 1 of the proof of part (b), we arrive at the diagram

$$\begin{tikzcd}
\mathrm{Spec}(L_0') \arrow[dd, "\text{degree } f"'] \arrow[dr] & & \\
& \mathrm{Spec}\, M_0 \arrow[dr] & \\
\mathrm{Spec}\, L \arrow[rr] & & \mathcal{Y}
\end{tikzcd}$$

which shows that $\mathrm{ed}(a;p) \leq \mathrm{tr\,deg}_K M_0 = \mathrm{tr\,deg}_K M - 1 \leq \mathrm{ed}(\mathcal{Y};p) - 1$, contradicting (4.3).

If the restriction of ν' to M is trivial then the induced valuation on $M \otimes_K L'$ is trivial, and we argue as in Case 2 in the proof of part (b). Tensoring the diagram 4.6 with L over K and identifying X_L with $\mathcal{B}_L\mu_d$ via our chosen map $\phi\colon \mathrm{Spec}(L) \to \mathcal{X}$, we obtain the diagram

$$\begin{tikzcd}
\mathrm{Spec}(L' \otimes_K L) \arrow[dd, "\text{prime-to-}p"'] \arrow[dr] & & \\
& \mathrm{Spec}\, M \otimes_K L \arrow[dr, "\beta \otimes_K \phi"] & \\
\mathrm{Spec}\, L(t) \otimes_K L \arrow[rr, "\alpha \otimes_K \phi"] & & \mathcal{B}_L\mu_d .
\end{tikzcd}$$

Since L is absolutely irreducible, $L' \otimes_K L$, $M \otimes_K L$ and $L(t) \otimes_K L$ are fields. The map $\mathrm{Spec}\, M \otimes_K L \to \mathcal{B}_L\mu_d$ gives rise to an $m \in (M \otimes L)^{\times}/(M \otimes L)^{\times d} = \mathrm{H}^1(M \otimes L, \mu_d)$ whose image in $L' \otimes_K L$ is equal to t, modulo d-th powers.

Lifting μ from L' to $L' \otimes_K L$ (so that ν'_L is trivial on $1 \otimes_K L$), we see that

$$\nu'_L(m) - \nu'_L(t) = \nu'_L(m) - 1 \in d\frac{1}{e}\mathbb{Z}\,.$$

Since d and e are relatively prime, this shows that $\nu'_L(m) \neq 0$. We conclude that ν'_L is non-trivial on $M \otimes_K L$. Consequently, ν' is non-trivial on M, contradicting our assumption. This contradiction completes the proof of part (d) (modulo Lemma 4.7 below). ♠

Lemma 4.7. *Let F be a functor from the category of field extensions of K to the category of sets.*

(a) Suppose that the natural map $F(A) \to F(B)$ is bijective for every algebraic inseparable field extension B/A. Then $\mathrm{ed}(a) = \mathrm{ed}(a_B)$ for every algebraic inseparable field extension B/A and every $a \in F(A)$.

(b) Let $\mathcal{X}$ be a gerbe over a field K banded by μ_d, A a field over K, and B/A an algebraic inseparable field extension. Suppose d is not divisible by $\mathrm{char}\, K$. Then $\mathrm{ed}(a) = \mathrm{ed}(a_B)$ for every $a \in \mathcal{X}(A)$.

Proof. (a) The inequality $\operatorname{ed}(a) \geq \operatorname{ed}(a_B)$ follows immediately from Definition 1.1. To prove the opposite inequality it suffices to show that if a_B descends to an intermediate field M, where $K \subset M \subset B$, then a descends to an intermediate field M_0, where $K \subset M_0 \subset A$ and $\operatorname{tr}\deg_K(M_0) \leq \operatorname{tr}\deg_K(M)$. To prove this assertion, set $M_0 = M \cap A$ and $p = \operatorname{char} K$. Then for every $m \in M$ there is a power of p, say $e = p^r$ such that m^e lies in both A and M (and hence, in M_0). This shows that M is algebraic and inseparable over M_0; cf. [Lan65, VII.7]. Since M_0 is a subfield of M, the inequality $\operatorname{tr}\deg_K(M_0) \leq \operatorname{tr}\deg_K(M)$ is obvious. (In fact, equality holds, since M is algebraic over M_0.) The commutative diagram

$$\begin{tikzcd}
F(M) \arrow[r] & F(B) \ni a_B \\
F(M_0) \arrow[r] \arrow[u, "\simeq"] & F(A) \in a \arrow[u, "\simeq"]
\end{tikzcd}$$

shows that a descends to M_0.

(b) By part (a) it suffices to show that the natural map $\mathcal{X}(A) \to \mathcal{X}(B)$ is an isomorphism for every algebraic inseparable extension B/A. If A doesn't split the class $[\mathcal{X}] \in \mathrm{H}^1(K, \mu_d)$ then neither does B, so, $\mathcal{X}(A) = \mathcal{X}(B) = \emptyset$, and there is nothing to prove. If A splits $[\mathcal{X}]$, i.e., there is a map $\operatorname{Spec}(A) \to \mathcal{X}$ then we can use this map to trivialize $\mathcal{X}$. The map $\mathcal{X}(A) \to \mathcal{X}(B)$ can then be identified with the natural map $\mathrm{H}^1(A, \mu_d) \to \mathrm{H}^1(B, \mu_d)$ or equivalently,

$$A^*/(A^*)^d \to B^*/(B^*)^d\,.$$

We want to show that this map is an isomorphism. To do this, we recall that for every $y \in B^*$ there exists an $r \geq 0$ such that y^{p^r} lies in A. Since p and d are relatively prime, both injectivity and surjectivity of the above map follow. ♠

In view of Theorem 4.1(a), Conjecture (3.7) of Colliot-Thélène, Karpenko and Merkurjev be restated in the language of essential dimension of gerbes as follows.

Conjecture 4.8. *If $\mathcal{X}$ is a gerbe banded by μ_n over a field K, let $p_1^{a_1} \dots p_r^{a_r}$ be the decomposition into prime factors of the index of the class of $\mathcal{X}$ in the Brauer group of K. Then*

$$\operatorname{ed} \mathcal{X} = p_1^{a_1} + \dots + p_r^{a_r} - r + 1.$$

When the index is 6 this follows from [CTKM06, Theorem 1.3].

In view of the fact that the conjecture holds for $r = 1$, it can also be rephrased as follows: if m and n are relatively prime positive integers, $\mathcal{X}$ and $\mathcal{Y}$ are gerbes banded by μ_m and μ_n, then

$$\operatorname{ed}(\mathcal{X} \times \mathcal{Y}) = \operatorname{ed} \mathcal{X} + \operatorname{ed} \mathcal{Y} - 1.$$

Back to the language of canonical dimension, one could ask the following more general question. Let X and Y be smooth projective varieties over a field K. Assume that there are no rational functions $X \dashrightarrow Y$ or $Y \dashrightarrow X$. Then is it true that $\operatorname{cdim}(X \times Y) = \operatorname{cdim}(X) + \operatorname{cdim}(Y)$? A positive answer to this question would imply Conjecture 4.8.

5. Proof of Theorem 1.4

In this section k is a field and p is a prime number not equal to $\operatorname{char} k$. We begin with some preliminary facts.

Proposition 5.1. *Let U be an integral algebraic space locally of finite type over k with function field $K \overset{\text{def}}{=} k(U)$, and let $f\colon \mathcal{X} \to U$ be a stack over U. Let $\mathcal{X}_K$ denote the pullback of $\mathcal{X}$ to* $\operatorname{Spec} K$*. Then,*

$$\operatorname{ed} \mathcal{X} \geq \operatorname{ed}(\mathcal{X}_K/K) + \dim U;$$
$$\operatorname{ed}(\mathcal{X}; p) \geq \operatorname{ed}(\mathcal{X}_K/K; p) + \dim U.$$

Proof. Let $j : \mathcal{X}_K \to \mathcal{X}$ denote the inclusion. If $\xi : \operatorname{Spec} L \to \mathcal{X}_K$ be a morphism, then it is easy to see that $\operatorname{ed}(\xi/K) = \operatorname{ed}(j \circ \xi/k) + \dim U$. This implies the first inequality of the proposition directly. The second inequality then follows by taking the infinum of $\operatorname{ed}(\xi_M/K)$ for all prime-to-p extensions M/L. ♠

Let $\mathcal{X}$ be a locally noetherian stack over a field k with presentation $P\colon X \to \mathcal{X}$. Recall that the *dimension of $\mathcal{X}$ at a point* $\xi\colon \operatorname{Spec} K \to \mathcal{X}$ is given by $\dim_x(X) - \dim_x P$ where x is an arbitrary point of X lying over ξ [LMB00, (11.14)]. Let $\mathcal{Y}$ be stack-theoretic closure of the image of ξ; that is, the intersection of all the closed substacks $\mathcal{Y}_i$ such that $\xi^{-1}(\mathcal{Y}_i) = \operatorname{Spec} K$. The morphism ξ factors uniquely through $\mathcal{Y} \subseteq \mathcal{X}$. We defined *the dimension of the point* ξ to be the dimension of the stack $\mathcal{Y}$ at the point $\operatorname{Spec} K \to \mathcal{Y}$.

Proposition 5.2. *Let $\mathcal{X} \to \mathcal{Y}$ be a morphism of algebraic stacks over a field k. Let K/k be a field extension and let $y\colon$ $\operatorname{Spec} K \to \mathcal{Y}$ be a point of dimension $d \in \mathbb{Z}$. Let $\mathcal{X}_K \overset{\text{def}}{=} \mathcal{X} \times_{\mathcal{Y}} \operatorname{Spec} K$. Then*

$$\operatorname{ed}(\mathcal{X}_K/K) \leq \operatorname{ed}(\mathcal{X}/k) - d;$$
$$\operatorname{ed}(\mathcal{X}_K/K; p) \leq \operatorname{ed}(\mathcal{X}/k; p) - d.$$

Proof. By [LMB00, Theorem 11.5], $\mathcal{Y}$ is the disjoint union of a finite family of locally closed, reduced substacks $\mathcal{Y}_i$ such that each $\mathcal{Y}_i$ is an fppf gerbe over an algebraic space Y_i with structural morphism $A_i\colon \mathcal{Y}_i \to Y_i$. We can therefore replace $\mathcal{Y}$ by one of the $\mathcal{Y}_i$ and assume that $\mathcal{Y}$ is an fppf gerbe over an algebraic space Y. Without loss of generality, we can assume that Y is an integral affine scheme of finite type over k.

Let p be the image of y in Y. Since $\mathcal{Y}$ is limit-preserving, we can find an integral affine scheme U equipped with a morphism $i\colon U \to \mathcal{Y}$ and a dominant morphism $j\colon \operatorname{Spec} K \to U$ such that y is equivalent to $i \circ j$. We can also assume that the composition $U \to \mathcal{Y} \to Y$ is dominant.

Since $\mathcal{Y}$ is a gerbe over Y, it follows that $U \to \mathcal{Y}$ is representable of fiber dimension at most $\dim U - d$. Now, form the following diagram with Cartesian squares.

$$\begin{tikzcd}
\mathcal{X}_K \arrow[r] \arrow[d] & \operatorname{Spec} K \arrow[d] \\
\mathcal{X}_U \arrow[r] \arrow[d] & U \arrow[d] \\
\mathcal{X} \arrow[r] & \mathcal{Y}
\end{tikzcd}$$

Since the vertical maps in the lower square are finite type and representable of fiber dimension at most $\dim U - d$, it follows from Propositions 2.6 and 5.1 that

$$\begin{aligned}
\operatorname{ed}(X_K/K) &\leq \operatorname{ed}(\mathcal{X}_{k(U)}/k(U)) \\
&\leq \operatorname{ed} \mathcal{X}_U - \dim U \\
&\leq \operatorname{ed} \mathcal{X} + \dim U - d + \dim U \\
&\leq \operatorname{ed} \mathcal{X} - d.
\end{aligned}$$

♠

Similarly, $\operatorname{ed}(\mathcal{X}_K/K; p) \leq \operatorname{ed}(\mathcal{X}; p) - d$,

We now proceed with the proof of Theorem 1.4, which we restate for the convenience of the reader. Let

$$1 \longrightarrow Z \longrightarrow G \longrightarrow Q \longrightarrow 1 \tag{5.3}$$

denote an extension of groups over a field k, with Z central and isomorphic to μ_n for some integer $n > 1$. Recall from the introduction, that we defined $\operatorname{ind}(G, Z)$ as the maximal value of the index $\operatorname{ind}\big(\partial_K(t)\big)$, as K ranges over all field extensions of k and t ranges over all torsors in $\mathrm{H}^1(K, Q)$.

Theorem 5.4. *Let G be an extension as in* (5.3)*. Assume that $n = p^r$ for some non-negative integer r. Then $\operatorname{ed}_k(G; p) \geq \operatorname{ind}(G, Z) - \dim G$.*

Proof. Let K/k be a field extension and let $t : \operatorname{Spec} K \to \mathcal{B}Q$ be a Q-torsor over $\operatorname{Spec} K$. The dimension of $\mathcal{B}Q$ at the point t is $-\dim Q = -\dim G$. Let $\mathcal{X}$ denote the pull-back in the following diagram.

$$\begin{tikzcd}
\mathcal{X} \arrow[r] \arrow[d] & \operatorname{Spec} K \arrow[d, "t"] \\
\mathcal{B}G \arrow[r] & \mathcal{B}Q
\end{tikzcd}$$

By Proposition 5.2, $\operatorname{ed}(\mathcal{X}/K; p) \leq \operatorname{ed}(\mathcal{B}G/k; p) + \dim G$. On the other hand, since $\mathcal{B}G$ is a gerbe banded by Z over $\mathcal{B}Q$, $\mathcal{X}$ is a gerbe banded by Z over $\operatorname{Spec} K$. Therefore, by Theorem 4.1(d), $\operatorname{ed}(\mathcal{X}/K; p) = \operatorname{ind} \partial_K(t)$. By substitution, $\operatorname{ind} \partial_K(t) - \dim G \leq \operatorname{ed}_k(\mathcal{B}G; p)$. Since this inequality holds for all field extensions K/k and all Q-torsors t over K, the theorem follows. ♠

Remark 5.5. Suppose G is a simple algebraic group whose center Z is cyclic. It is tempting to apply Theorem 1.4 to the natural sequence

$$1 \longrightarrow Z \longrightarrow G \longrightarrow G^{\mathrm{ad}} \longrightarrow 1$$

where the adjoint group G^{ad} is G/Z. Given a torsor $t \in \mathrm{H}^1(K, G^{\mathrm{ad}})$, the central simple algebra representing $\partial_K(t) \in \mathrm{H}^2(K, Z)$ is called *the Tits algebra* of t. The values of the index of the Tits algebra were studied in [Tit92], where this index is denoted by $b(X)$ (for a group of type X) and its possible values are listed on p. 1133. A quick look at this table reveals that for most types these indices are smaller than $\dim G$, so that the bound of Theorem 1.4 becomes vacuous. The only exception are groups of types B and D, in which case Theorem 1.4 does indeed, give interesting bounds; cf. Remark 10.7.

Remark 5.6. If $n \geq 2$ is not necessarily a prime power, then the above argument shows that

$$\mathrm{ed}_k G \geq \mathrm{cdim}(\partial_K(t)) - \dim G \,.$$

Here by $\mathrm{cdim}(\partial_K(t))$ we mean the canonical dimension of the Brauer-Severi variety representing the class of $\partial_K(t)$ in $\mathrm{H}^2(K, \mu_n)$. Assuming that Conjecture 4.8 holds, we obtain the following (conjectural) inequality: if $\mathrm{ind}(G, Z) = p_1^{a_1} \dots p_r^{a_r}$ is the prime factorization of $\mathrm{ind}(G, Z)$ then

$$\mathrm{ed}_k G \geq p_1^{a_i} + \dots + p_r^{a_r} - r + 1 - \dim G \,.$$

As we remarked at the end of §4, Conjecture 4.8 (and thus the above inequality) is known to be true if $\mathrm{ind}(G, Z)$ is a prime power (i.e., $r = 1$) or $\mathrm{ind}(G, Z) = 6$.

6. Florence's theorem

In this section we give an alternative proof of a recent theorem of Florence (in a slightly strengthened form) by combining the lower bound of Theorem 1.4 with the Brauer–Rowen Theorem. The idea to use the Brauer–Rowen theorem in this context is due to Florence. Thus while our proof is different from the one in [Flo], it is not entirely independent.

Theorem 6.1 (M. Florence [Flo]). *Let p be a prime, k a field of characteristic $\neq p$. Suppose $\zeta_{p^n} \in k$ but $\zeta_{p^{n+1}} \notin k$ for some integer $n \geq 1$. Moreover, if $p = 2$ and $n = 1$, assume also that $k(\zeta_4) \neq k(\zeta_8)$. Then*

$$\mathrm{ed}_k(C_{p^m}; p) = \mathrm{ed}_k C_{p^m} = \begin{cases} p^{m-n} & \text{if } n < m, \\ 1 & \text{if } n \geq m. \end{cases}$$

Before proceeding with the proof, we remark that, in general, the value of $\mathrm{ed}_k C_{p^m}$ is not known if $\zeta_p \notin k$. On the other hand, for the sake of computing $\mathrm{ed}_k(C_{p^m}; p)$, we may replace k by $k(\zeta_p)$ and then apply Theorem 6.1. Indeed, since $[k(\zeta_p) : k]$ is prime to p, $\mathrm{ed}_k(C_{p^m}; p) = \mathrm{ed}_{k(\zeta_p)}(C_{p^m}; p)$.

Proof. If $m \leq n$, then $C_{p^m} = \mu_{p^m}$. Therefore $\operatorname{ed} C_{p^m} = 1$ (see [BF03, Example 2.3]) and hence, $\operatorname{ed}(C_{p^m}; p) = 1$ as well. We can therefore restrict our attention to the case where $n < m$.

We first show that $\operatorname{ed} C_{p^m} \leq p^{m-n}$. To do this, pick a faithful character $\chi\colon C_{p^n} \to \mathbb{G}_m$ defined over K and set $V \overset{\text{def}}{=} \operatorname{ind}_{C_{p^n}}^{C_{p^m}} \chi$. A simple calculation shows that V is faithful, thus, V is generically free since C_{p^m} is finite. By Theorem 2.9, it follows that $\operatorname{ed} C_{p^m} \leq \dim V = p^{m-n}$.

It remains to show that $\operatorname{ed}(C_{p^m}; p) \geq p^{m-n}$. By Theorem 1.4 it suffices to show that $\operatorname{ind}(C_{p^m}, C_{p^n}) \geq p^{m-n}$. To establish this inequality, we will view the representation V as a homomorphism $\rho\colon C_{p^m} \to \mathrm{GL}(V)$ of algebraic groups. Let $\pi : \mathrm{GL}(V) \to \mathrm{PGL}(V)$ denote the obvious projection and note that the kernel of $\pi \circ \rho$ is exactly C_{p^n}. It follows that we have a commutative diagram

$$\begin{tikzcd}
0 \arrow[r] & C_{p^n} \arrow[r] \arrow[d] & C_{p^m} \arrow[r, "\rho"] \arrow[d] & C_{p^{m-n}} \arrow[r] \arrow[d, "\iota"] & 1 \\
1 \arrow[r] & \mathbb{G}_m \arrow[r] & \mathrm{GL}(V) \arrow[r, "\pi"] & \mathrm{PGL}(V) \arrow[r] & 1
\end{tikzcd}$$

where the rows are exact and the columns are injective.

Let K/k be a field extension and let $t \in \mathrm{H}^1(K, C_{p^m})$ be a torsor. Let $\iota_*\colon \mathrm{H}^1(K, C_{p^{m-n}}) \to \mathrm{H}^1\big(K, \mathrm{PGL}(V)\big)$ denote the map induced by ι. Then, from the commutativity of the above diagram (and the injectivity of the columns), it follows that $\operatorname{ind}_K(t)$ is the index of the CSA $\iota_*(t)$.

We claim that there is a field extension K/k and a $t \in \mathrm{H}^1(K, C_{p^m})$ such that $\iota_*(t)$ is a division algebra. From this it will easily follow that $\operatorname{ind}_K(t) = \dim V = p^{m-n}$. In fact, we will take $t \in \mathrm{H}^1(K, C_{p^m})$ to be a versal element, constructed as follows. Let $L = K(x_1, \dots, x_{p^{m-n}})$ denote the field obtained by adjoining p^{m-n} independent variables to K, and let $C_{p^{m-n}}$ act on L by cyclically permuting the variables, i.e., $a \cdot x_i = x_{i+a} \pmod{p^{m-n}}$ for any $a \in C_{p^{m-n}}$. Let $K = L^{C_{p^{m-n}}}$. Then L/K defines a $C_{p^{m-n}}$-torsor t over K.

The CSA corresponding to $\iota_*(t) \in \mathrm{H}^1(K, \mathrm{PGL}_n)$ is the Brauer-Rowen algebra R_{p^n, p^m, p^m}; cf. [Row88, §7.3]. The fact that this algebra has index n is a variant of a theorem of Brauer; for a proof see [Flo, Theorem 2.17] or (for $k = \mathbb{Q}(\zeta_p)$ only) [Row88, Theorem 7.3.8]. ♠

Let D_n be the dihedral group of order $2n$. Ledet [Led02, Section 3] conjectured that if n is odd then $\operatorname{ed}_k C_n = \operatorname{ed}_k D_n$ over any field k of characteristic zero. As an application of Theorem 6.1 we will now prove Ledet's conjecture in the case where $n = p^r$ is a prime power and k contains a primitive pth root of unity.

Corollary 6.2. *Let p be an odd prime, $m \geq 1$ be an integer, and k be a field containing a primitive pth root of unity. Then $\operatorname{ed}_k D_{p^m} = \operatorname{ed}_k C_{p^m}$.*

Proof. If $\zeta_{p^m} \in k$ then we know that

$$\operatorname{ed}_k C_{p^m} = \operatorname{ed}_k D_{p^m} = 1\,;$$

see the proof of [BR97, Theorem 6.2]. Thus we may assume $\zeta_{p^m} \notin k$. Let n be the largest integer such that $\zeta_{p^n} \in k$. By our assumption $1 \leq n \leq m-1$. Since $C_{p^m} \subset D_{p^m}$, we have $\operatorname{ed}_k D_{p^m} \geq \operatorname{ed}_k C_{p^m}$.

To prove the opposite inequality, note that $D_{p^m} \simeq C_{p^m} \rtimes C_2$ has a subgroup isomorphic to $D_{p^n} = C_{p^n} \rtimes C_2$ of index p^{m-n}. As we pointed out above, $\zeta_{p^n} \in k$ implies $\operatorname{ed}_k D_{p^n} = 1$. Thus

$$\operatorname{ed}_k D_{p^m} \leq (\operatorname{ed}_k D_{p^n}) \cdot [D_{p^m} : D_{p^n}] = 1 \cdot p^{m-n} = \operatorname{ed} C_{p^m}\,.$$

where the inequality is given by [Led02, Section 3] and the last equality by Theorem 6.1. ♠

7. A theorem about the essential dimension of a finite group

The following proposition, extending [BR97, Theorem 5.3] and [Kan06, Theorem 4.5], will be used in the proof of Theorem 1.5 in the next section.

Theorem 7.1. *Let G be a finite group, H be a central cyclic subgroup of G, and $\chi\colon G \to k^*$ be a character of G whose restriction to H is faithful. Assume* $\operatorname{char} k$ *does not divide $|G|$. If H is maximal among central cyclic subgroups of G then*

(a) $\operatorname{ed}_k G = \operatorname{ed}_k(G/H) + 1$.

(b) Moreover, if G is a p-group then $\operatorname{ed}_k(G;p) = \operatorname{ed}_k(G/H;p) + 1$.

Proof. Let $\phi\colon G \to G/H \hookrightarrow \mathrm{GL}(V)$ be a faithful representation of G/H. Then $\phi \oplus \chi\colon G \to \mathrm{GL}(V \times \mathbb{A}^1)$ is a faithful representation of G.

(a) To prove the inequality $\operatorname{ed} G \leq \operatorname{ed}(G/H) + 1$, recall that $\operatorname{ed}(G/H)$ is the minimal dimension of a faithful G/H-variety X which admits a dominant G/H-equivariant rational map $f\colon V \dashrightarrow X$; see [BR97]. Then

$$f \times \mathrm{id}\colon V \times \mathbb{A}^1 \dashrightarrow X \times \mathbb{A}^1$$

is a dominant rational map of G-varieties. Thus

$$\operatorname{ed} G \leq \dim(X \times \mathbb{A}^1) = \operatorname{ed}(G/H) + 1\,.$$

We will now prove the opposite inequality, $\operatorname{ed} G \geq \operatorname{ed}(G/H) + 1$. Choose a dominant rational map

$$V \times \mathbb{A}^1 \dashrightarrow Y$$

of faithful G-varieties such that $\dim Y = \operatorname{ed} G$. Equivalently, if $K = k(V)$, we choose G-invariant subfield $F \subset K(t)$ such that the G-action on F is faithful and $\operatorname{tr}\deg_k(F) = \operatorname{ed} G$. (Here, of course, $F = k(Y)$.)

Let $\nu\colon K(t) \to \mathbb{Z}$ be the natural valuation associated to t (and trivial on K). Restricting ν to F we obtain a discrete valuation $\nu_{|F}$ of F. Denote the residue field for this valuation by F_0.

Claim: (i) $\nu_{|F}$ is non-trivial, and (ii) the inertia group of $\nu_{|F}$ (i.e., the subgroup of G that acts trivially on F_0) is precisely H.

To prove (i), assume the contrary. Then there is a natural G-equivariant embedding of F in the residue field $K(t)_0 = K$. Since H acts faithfully on F and trivially on K, this is impossible.

To prove (ii), let $H_F \subset G$ be the inertia group for $\nu_{|F}$. By our construction H is the inertia group for the valuation ν on $K(t)$. Hence, $H \subset H_F$. On the other hand, since $\operatorname{char} k$ is prime to $|G|$, H_F is cyclic (see [Ser79, Corollaries 2 and 3 to Proposition IV.7]) and central (see [Ser79, Propositioni IV.10]). By our assumption H is maximal among central cyclic subgroups of G; thus $H = H_F$, proving (ii).

We are now ready to complete the proof of the inequality $\operatorname{ed} G \geq \operatorname{ed}(G/H)+1$ (and thus of part (a)). By (i) the G-equivariant inclusion $F \hookrightarrow K(t)$ induces a G/H-equivariant inclusion $F_0 \hookrightarrow K(t)_0 = K$ of residue fields. In other words, we have a rational dominant G/H-equivariant map $V \dashrightarrow Y_0$ where $k(Y_0) = F_0$. Moreover, by (ii) G/H acts faithfully on F_0 (or equivalently, on Y_0). Since V is a faithful linear representation of G/H, we have $\operatorname{ed}(G/H) \leq \dim Y_0$. On the other hand,

$$\dim Y_0 = \operatorname{tr}\deg_k(F_0) = \operatorname{tr}\deg_k(F) - 1 = \dim Y - 1 \leq \operatorname{ed} G - 1\,.$$

and the desired inequality $\operatorname{ed}(G/H) \leq \operatorname{ed} G - 1$ follows.

(b) To prove the inequality $\operatorname{ed}(G;p) \leq \operatorname{ed}(G/H;p)+1$, recall that $\operatorname{ed}(G/H)$ is the minimal dimension of a faithful G/H-variety X which admits a diagram

$$\begin{tikzcd}
V' \arrow[d, dashed, "\text{prime-to-}p"'] \arrow[dr, dashed] & \\
V & X\,,
\end{tikzcd}$$

of dominant rational G/H-equivariant maps, where X is a generically free G-variety, $\dim V' = \dim V$ and $[k(V') : k(V)]$ is prime to p. (Note that, in the definition of $\operatorname{ed}(G;p)$ we require that G should transitively permute the connected components of V' (and of X). However, if G is a p-group and the degree of the cover $V' \to V$ is prime to p, a simple counting argument shows that G has to fix an irreducible component of V'. Consequently, in this situation V' and X are necessarily irreducible.) Multiplying every variety in the above diagram by $\mathbb{A}^1$ (on which G acts via the character χ), we obtain the diagram

$$\begin{tikzcd}
V' \times \mathbb{A}^1 \arrow[d, dashed, "\text{prime-to-}p"'] \arrow[dr, dashed] & \\
V \times \mathbb{A}^1 & X \times \mathbb{A}^1\,,
\end{tikzcd}$$

of dominant rational map of faithful G-varieties, which shows that

$$\operatorname{ed}(G;p) \leq \dim(X \times \mathbb{A}^1) = \operatorname{ed}(G/H;p) + 1\,.$$

To prove the opposite inequality, $\mathrm{ed}(G;p) \leq \mathrm{ed}(G/H;p) + 1$, choose a diagram

$$\begin{tikzcd} W \arrow[d, dashed, "\text{prime-to-}p"'] \arrow[dr, dashed] & \\ V \times \mathbb{A}^1 & Y\,, \end{tikzcd}$$

of rational maps of faithful G-varieties, where $[k(W) : k(V \times \mathbb{A}^1)]$ is finite and prime to p and $\dim Y = \mathrm{ed}_k(G;p)$. Equivalently, we have a diagram of field extensions

$$\begin{tikzcd} L \arrow[d, dash, "\text{prime-to-}p"'] \arrow[dr, dash] & \\ K(t) & F\,, \end{tikzcd}$$

where K, $K(t)$, L and F are the function fields of V, $V \times \mathbb{A}^1$, W and Y, respectively.

First observe that after replacing L by the separable closure L_s of $K(t)$ in L and F by $F \cap L_s$, we may assume without loss of generality that L is separable over $K(t)$. Indeed, the G-action on L and F descends to (and remains faithful on) L_s and $F \cap L_s$. Moreover, some power of every $x \in F$ lies in $F \cap L_s$; hence, F is algebraic over $F \cap L_s$, i.e., the two have the same transcendence degree over k.

From now on we will assume that L is separable over $K(t)$ (or equivalently, W is separable over $V \times \mathbb{A}^1$). Let ν be the natural discrete valuation on $K(t)$, by the degree in t (the same valuation we considered in part (a)). We claim that there exists a G-invariant lifting η of ν to L whose residue degree (i.e., the degree of the induced extension L_0/K of the residue fields) in prime to p.

Clearly the inertia group $H_L \subset G$ of L (with respect to η) is contained in the inertia group of $K(t)$, which is equal to H by our construction. On the other hand, the cyclic p-group H/H_L faithfully acts on the extension L_0/K whose degree is prime to p. This is only possible if $H/H_L = \{1\}$, i.e., $H = H_L$. Thus if the above claim is established, we can use the same argument as in part (a) (with $K(t)$, ν replaced by L, η) to complete the proof of the inequality $\mathrm{ed}(G;p) \leq \mathrm{ed}(G/H;p) + 1$.

To prove the claim, denote the liftings of ν on $K(t)$ to L by $\eta_1, \dots, \eta_r$. As usual, we will denote the ramification index of η_i by e_i and the residue degree by f_i. By [Lan65, Proposition XII.18]

$$e_1 f_1 + \cdots + e_r f_r = [L : K(t)]$$

or, equivalently,

$$\sum_{e,f \geq 1} |V_{e,f}| ef = [L : K(t)]$$

where $V_{e,f}$ be the subset of $\{\eta_1, \dots, \eta_n\}$ consisting of those valuations η_i with $e_i = e$ and $f_i = f$. (Note that $V_{e,f} = \emptyset$ for all but finitely many pairs (e, f).)

Thus there exist integers $e, f \geq 1$ such that e, f, and $|V_{e,f}|$ are all prime to p. The p-group G permutes the elements of $V_{e,f}$; since $|V_{e,f}|$ is prime to p, G must fixed a valuation $\eta \in V_{e,f}$. This is the valuation we were looking for. (Note that the residue degree f of η is prime to p by our construction.) This completes the proof of the claim and thus of part (b). ♠

Corollary 7.2. *Let G be a finite p-group and S be a direct summand of* $\mathrm{C}(G)$ *such that* $[G, G] \cap S = \{1\}$*. If* $\zeta_e \in k$*, where e is the exponent of S then* $\mathrm{ed}_k\, G = \mathrm{ed}_k(G/S) + \mathrm{rank}(S)$*.*

Remark 7.3. The following elementary observation will be useful in the sequel.

Let G be a finite group and S be a normal subgroup such that $[G, G] \cap S = \{1\}$*. Then* $\mathrm{C}(G/S) = \mathrm{C}(G)/S$*.*

To prove this assertion, suppose $g \in G$ projects to a central element of G/S. Then for every $x \in G$ the commutator $c = gxg^{-1}x^{-1}$ lies in both $[G, G]$ and S. Hence, $c = 1$, and thus $g \in \mathrm{C}(G)$, as desired. ♠

Proof of Corollary 7.2. Suppose $\mathrm{C}(G)$ as $S \oplus T$. Write $S = H_1 \oplus \cdots \oplus H_r$ as a direct sum of cyclic groups, where $r = \mathrm{rank}(S)$. We argue by induction on r. When $r = 0$, i.e., $S = \{1\}$, there is nothing to prove. For the induction step, note that by Remark 7.3,

$$S/H_r \simeq H_1 \oplus \cdots \oplus H_{r-1}$$

is a direct summand of $\mathrm{C}(G/H_r) = S/H_r \oplus T$. Thus by the induction assumption $\mathrm{ed}(G/H_r) = \mathrm{ed}(G/S) + r - 1$. On the other hand, by Theorem 7.1 $\mathrm{ed}\, G = \mathrm{ed}(G/H_r) + 1$, and the corollary follows. ♠

8. Proof of Theorem 1.5

Let Z be a maximal cyclic subgroup of $\mathrm{C}(G)$ containing $[G, G]$. Then $\mathrm{C}(G) = Z \oplus S$ for some central subgroup W of G. By Corollary 7.2,

$$\mathrm{ed}\, G = \mathrm{ed}(G/S) + \mathrm{rank}(S) = \mathrm{ed}(G/S) + \mathrm{rank}\ \mathrm{C}(G) - 1\,.$$

By Remark 7.3, $|\mathrm{C}(G/S)| = |\mathrm{C}(G)|/|S|$. Thus the quantity $|G/\mathrm{C}(G)|$ does not change when we replace G by G/S. We conclude that it suffices to prove Theorem 1.5 for G/S. In other words, we may assume without loss of generality that the center $Z = \mathrm{C}(G)$ of G is cyclic. In this case the theorem reduces to the identity

$$\mathrm{ed}(G; p) = \mathrm{ed}\, G = \sqrt{|G/Z|}\,.$$

To prove this identity it suffices to establish Lemma 8.1 below. Indeed, part (a) of this lemma tells us that $\mathrm{ed}(G; p) \geq \sqrt{|G/Z|}$ (see Theorem 1.4), and part (b) implies the inequality, $\mathrm{ed}\, G \leq \sqrt{|G/Z|}$ (see Theorem 2.9), and Theorem 1.5 follows.

Lemma 8.1. *Let G be a p-group such that the center $Z = \mathrm{C}(G)$ is cyclic and the quotient group $A \overset{\text{def}}{=} G/Z$ is abelian. (or equivalently, $[G, G] \subset Z$). Then*

(a) $\operatorname{ind}(G, Z) \geq \sqrt{|A|}$*, and*

(b) G has a faithful linear representation of dimension $\sqrt{|A|}$.

Proof. We will use additive notation for the groups Z and A, multiplicative for G. In this situation we can define a skew-symmetric bilinear form $\omega\colon A \times A \to Z$ by

$$\omega(a_1, a_2) = g_1 g_2 g_1^{-1} g_2^{-1},$$

where $a_i = g_i$, modulo Z, for $i = 1, 2$. (Note that $\omega(a_1, a_2)$ is independent of the choice of g_1 and g_2.) Clearly g lies in $\mathrm{C}(G)$ if and only if its image a lies in the kernel of ω; i.e., $\omega(a, b) = 0$ for every $b \in A$. Since we are assuming that $\mathrm{C}(G) = Z$, we conclude that the kernel of ω is trivial; i.e., ω is a *symplectic form* on A. It is well known (see for example [TA86, §3.1]) that the order of A, which equals the order of $G/\mathrm{C}(G)$, is a complete square.

Fix a generator z of Z. We recall the basic result on the structure of a symplectic form ω on a finite abelian group A (the proof is easy; it can be found, e.g., in [Wal63, §3.1] or [TA86, §7.1]). There exist elements a_1, …, a_{2r} in A and positive integers d_1, …, d_r with the following properties.

(a) d_i divides d_{i-1} for each $i = 2, \ldots, r$, and $d_r > 1$.
(b) Let i be an integer between 1 and r. If A_i denotes the subgroup of A generated by a_i and a_{r+i}, then there exists an isomorphism $A_i \simeq (\mathbb{Z}/d_i\mathbb{Z})^2$ such that a_i corresponds to (1,0) and a_{r+i} to $(0, 1)$.
(c) The subgroups A_i are pairwise orthogonal with respect to ω.
(d) $\omega(a_i, a_{r+i}) = z^{n/d_i} \in Z$.
(e) $A = A_1 \oplus \cdots \oplus A_r$.

Then the order of A is $d_1^2 \ldots d_r^2$, hence $\sqrt{|A|} = d_1 \ldots d_r$.

Let G_i be the inverse image of A_i in G; note that G_i commutes with G_j for any $i \neq j$.

Let u_1, …, u_{2r} be indeterminates, and set $K \overset{\text{def}}{=} k(u_1, \ldots, u_{2r})$. Identify Z with μ_n by sending z into ζ_n. Consider the boundary map

$$\partial_i\colon \mathrm{H}^1(K, A_i) \longrightarrow \mathrm{H}^2(K, Z)$$

obtained from the exact sequence

$$1 \longrightarrow Z \longrightarrow G_i \longrightarrow A_i \longrightarrow 1.$$

Claim. There exists a class $\xi_i \in \mathrm{H}^1(K, A_i)$ such that $\partial_i \xi_i$ is the class of the cyclic algebra $(u_i, u_{r+i})_{d_i}$ in $\operatorname{Br} K$.

To see that part (a) follows from the claim, consider the commutative diagram

$$\begin{tikzcd}
1 \arrow[r] & Z^r \arrow[r] \arrow[d, "m"] & \prod_i G_i \arrow[r] \arrow[d] & \prod_i A_i \arrow[r] \arrow[d] & 1 \\
1 \arrow[r] & Z \arrow[r] & G \arrow[r] & A \arrow[r] & 1
\end{tikzcd}$$

in which m is defined by the formula $m(z_1, \dots, z_r) = z_1 \dots z_r$, and the homomorphism $\prod_i G_i \to G$ is induced by the inclusions $G_i \subseteq G$. This yields a commutative diagram

$$\begin{tikzcd}
\prod_i \mathrm{H}^1(K, A_i) \arrow[r, "\prod_i \partial_i"] \arrow[d] & \mathrm{H}^2(K, Z)^r \arrow[d, "m_*"] \\
\mathrm{H}^1(K, A) \arrow[r, "\partial"] & \mathrm{H}^2(K, Z)
\end{tikzcd}$$

in which the map m_* is given by $m_*(\alpha_1, \dots, \alpha_r) = \alpha_1 \dots \alpha_r$. So, if $\xi \in \mathrm{H}^1(K, A)$ is the image of $(\xi_1, \dots, \xi_r)$ then $\partial\xi$ is the class of the product

$$(u_1, u_{r+1})_{d_1} \otimes_K (u_2, u_{r+2})_{d_2} \otimes_K \cdots \otimes_K (u_r, u_{2r})_{d_r},$$

whose index is $d_1 \dots d_r$. Hence $\operatorname{ind}(G, Z) \geq d_1 \dots d_r = \sqrt{|A|}$, as needed.

We now proceed with the proof of the claim. Choose a power of p, call it d, that is divisible by the order of Z and by the order of each a_i. Consider the group $\Lambda(d)$ defined by the presentation

$$\langle x_1, x_2, y \mid x_1^d = x_2^d = y^d = 1,\ x_1x_2 = yx_2x_1,\ x_1y = yx_1,\ x_2y = yx_2 \rangle.$$

Call $\rho_i \colon \Lambda(d) \to G_i$ the homomorphism obtained by sending x_1 to a_i, x_2 to a_{r+i}, and y to $z^{n/d_i} = \omega(a_i, a_{r+i})$.

Let ζ_d be a primitive d-th root of 1 in k such that $\zeta_n = \zeta_d^{n/d}$. The subgroup $\langle y \rangle$ in $\Lambda(d)$ is cyclic of order d; we fix the isomorphism $\langle y \rangle \simeq \mu_d$ so that y corresponds to ζ_d. The restriction of ρ_i to $\langle y \rangle \to Z$ corresponds to the homomorphism $\mu_d \to \mu_n$ defined by $\alpha \mapsto \alpha^{d/d_i}$. We have a commutative diagram

$$\begin{tikzcd}
1 \arrow[r] & \mu_d \arrow[r] \arrow[d, swap, "\substack{\alpha \\ \downarrow \\ \alpha^{d/d_i}}"] & \Lambda(d) \arrow[r] \arrow[d, "\rho_i"] & (\mathbb{Z}/d\mathbb{Z})^2 \arrow[r] \arrow[d] & 1 \\
1 \arrow[r] & \mu_n \arrow[r] & G_i \arrow[r] & A_i \arrow[r] & 1.
\end{tikzcd}$$

We have $\mathrm{H}^1\big(K, (\mathbb{Z}/d\mathbb{Z})^2\big) = (K^*/K^{*d})^2$. According to [Vel00, Example 7.2], the image of the element $(u_i, u_{r+i}) \in \mathrm{H}^1\big(K, (\mathbb{Z}/d\mathbb{Z})^2\big)$ is the cyclic algebra $(u_i, u_{r+i})_d$; hence, if ξ_i is the image in $\mathrm{H}^1(K, A_i)$ of (u_i, u_{r+i}), the image of ξ_i in $\mathrm{H}^2(K, \mu_d)$ is the algebra $(u_i, u_{r+i})_d^{\otimes d/d_i}$, which is equivalent to $(u_i, u_{r+i})_{d_i}$. This concludes the proof of Lemma 8.1(a).

We now turn to the proof of Lemma 8.1(b). Suppose $|\mathrm{C}(G)| = p^h$ and $|A| = p^{2m}$; we want to construct a faithful representation of G of dimension p^m. By [TA86, §3.1] A contains a Lagrangian subgroup L of order p^m. Denote by H the inverse image of L in G; then H is an abelian subgroup of G of order p^{h+m}. Since $\zeta_{p^e} \in k$ we can embed Z in k^* and extend this embedding to a homomorphism $\chi\colon H \to k^*$. We claim that the representation $\rho\colon G \to \mathrm{GL}_{p^m}$ induced by χ is faithful.

It is enough to show that $\rho(g) \neq \mathrm{id}$ for any $g \in G$ of order p, or, equivalently, that $\rho\mid_{\langle g\rangle}$ is non-trivial for any such g. If $s \in G$ consider the subgroup $H_s \overset{\text{def}}{=} s\langle g\rangle s^{-1} \cap H$ of H, which is embedded in $\langle g\rangle$ via the homomorphism $x \mapsto s^{-1}xs$. By Mackey's formula ([Ser77, §7.3]), $\rho\mid_{\langle g\rangle}$ contains all the representations of $\langle g\rangle$ induced by the restrictions $\chi\mid_{H_s}$ via the embedding above.

If $g \notin H$ then $H_1 = \langle g\rangle \cap H = \{1\}$: we take $s = 1$, and we see that $\rho\mid_{\langle g\rangle}$ contains a copy of the regular representation of $\langle g\rangle$, which is obviously non-trivial.

Assume $g \in H$. Then $H_s = \langle sgs^{-1}\rangle$ for any $s \in G$; it is enough to prove that $\chi(sgs^{-1}) \neq 1$ for some $s \in G$. If $\chi(g) \neq 1$ then we take $s = 1$. Otherwise $\chi(g) = 1$; in this case $g \notin \mathrm{C}(G)$, because $\chi\mid_{\mathrm{C}(G)}\colon \mathrm{C}(G) \to k^*$ is injective. Hence the image $\overline{g}$ of g in A is different from 0, and we can find $s \in G$ such that $\omega(\overline{s}, \overline{g}) \neq 1$. Then

$$\chi(sgs^{-1}) = \chi\bigl(\omega(\overline{s},\overline{g})g\bigr) = \chi\bigl(\omega(\overline{s},\overline{g})\bigr)\chi(g) = \chi\bigl(\omega(\overline{s},\overline{g})\bigr) \neq 1.$$

This concludes the of Lemma 8.1 and thus of Theorem 1.5. ♠

9. More on p-groups

In this section we will discuss some consequences of Theorem 1.5.

Example 9.1. Recall that a p-group G is called *extra-special* if its center Z is cyclic of order p, and the quotient G/Z is elementary abelian. The order of an extra special p-group G is an odd power of p; the exponent of G is either p or p^2; cf. [Rob96, pp. 145–146]. Note that every non-abelian group of order p^3 is extra-special. For extra-special p-groups Theorem 1.5 reduces to the following.

Let G be an extra-special p-group of order p^{2m+1}. Assume that the characteristic of k is different from p, that $\zeta_p \in k$, and $\zeta_{p^2} \in k$ if the exponent of G is p^2. Then $\mathrm{ed}(G;p) = \mathrm{ed}\, G = p^m$.

Example 9.2. Let p be an odd prime and $G = C_{p^r} \ltimes C_{p^s}$ be the natural semidirect product of cyclic groups of order p^r and p^s (in other words, C_{p^s} is identified with the unique subgroup of $C_{p^r}^*$ of order p^s). If $s \leq r/2$ then

$$\mathrm{ed}_k(C_{p^r} \ltimes C_{p^s}; p) = \mathrm{ed}_k(C_{p^r} \ltimes C_{p^s}) = p^s\,,$$

for any field k containing a primitive pth root of unity ζ_p.

Proof. The center $\mathrm{C}(G)$ of G is the (unique) subgroup of C_{p^r} of order p^s. Since $s \leq r/2$, this subgroup is central. Thus, if $\zeta_{p^r} \in k$, the equality $\operatorname{ed}_k G = p^s$ is an immediate consequence of Theorem 1.5. But we are only assuming that $\zeta_p \in k$, so Theorem 1.5 only tells us that $\operatorname{ed}_k G \geq p^s$. To prove the opposite inequality, we argue as follows. Let F be the prime subfield of k. By [Led02, Corollary to Proposition 2], $\operatorname{ed}_{F(\zeta_p)}(G) \leq p^s$. Since we are assuming that $F(\zeta_p) \subset k$, we conclude that $\operatorname{ed}_k G \leq p^s$ as well. ♠

Corollary 9.3. *Suppose k is a base field of characteristic $\neq p$. If G is a non-abelian finite p-group then* $\operatorname{ed}(G;p) \geq p$.

Proof. Assume the contrary: let G be a non-abelian p-group of smallest possible order such that $\operatorname{ed} G < p$. Since G has a non-trivial center, there exists a cyclic central subgroup $Z \subset G$. The short exact sequence

$$1 \longrightarrow Z \longrightarrow G \longrightarrow G/Z \longrightarrow 1 \tag{9.4}$$

gives rise to the exact sequence of pointed sets

$$\mathrm{H}^1(K,G) \longrightarrow \mathrm{H}^1(K,G/Z) \xrightarrow{\partial_K} \mathrm{H}^2(K,Z)$$

for any field extension K of our base field k. We will now consider two cases.

Case 1. Suppose the map $\mathrm{H}^1(K,G) \to \mathrm{H}^1(K,G/Z)$ is not surjective for some K/k. Then ∂_K is non-trivial, and Theorem 1.4 tells us that $\operatorname{ed}(G;p) \geq p$, a contradiction.

Case 2. Suppose the map $\mathrm{H}^1(K,G) \to \mathrm{H}^1(K,G/Z)$ is surjective for every K/k. Then the morphism $\mathcal{B}G \to \mathcal{B}(G/Z)$ is isotropic, and Proposition 2.8 implies that $p > \operatorname{ed}(G;p) \geq \operatorname{ed}(G/Z;p)$. By the minimality of G, the group G/Z has to be abelian. Consequently, $[G,G] \subset Z$ is cyclic and central in G. Since G is non-abelian, $|G/\mathrm{C}(G)| \geq p^2$. Theorem 1.5 now tells us that

$$\operatorname{ed}(G;p) = \sqrt{|G|/|\mathrm{C}(G)|} + \operatorname{rank}\,\mathrm{C}(G) - 1 \geq p\,,$$

a contradiction. ♠

We will conclude this section by answering the following question of Jensen, Ledet and Yui [JLY02, p. 204].

Question 9.5. Let G be a finite group and N be a normal subgroup. Is it true that $\operatorname{ed} G \geq \operatorname{ed}(G/N)$?

The inequality $\operatorname{ed} G \geq \operatorname{ed}(G/N)$ is known to hold in many cases; cf., e.g., Theorem 7.1. We will now show that it does not hold in general, even if H is assumed to be central.

Corollary 9.6. *For every real number $\lambda > 0$ there exists a finite p-group G, with a central subgroup $H \subset G$ such that* $\operatorname{ed}(G/H) > \lambda \operatorname{ed} G$.

Proof. Let Γ be a non-abelian group of order p^3. The center of Γ has order p; denote it by C. The center of $\Gamma^n = \Gamma \times \cdots \times \Gamma$ (n times) is then C^n.

Let H_n be the subgroup of C^n consisting of n-tuples $(c_1, \dots, c_n)$ such that $c_1 \dots c_n = 1$. Clearly

$$\operatorname{ed} \Gamma^n \leq n \cdot \operatorname{ed} \Gamma = np\,;$$

see Example 9.1.

On the other hand, Γ^n/H_n, is easily seen to be extra-special of order p^{2n+1}, so $\operatorname{ed}(\Gamma^n/H_n) = p^n$, again by Example 9.1. Setting $G = \Gamma^n$ and $H = H_n$, we see that the desired inequality $\operatorname{ed}(G/H) > \lambda \operatorname{ed} G$, holds for suitably large n. ♠

10. Spinor groups

In this section we will prove Theorem 1.7 stated in the introduction.

As usual, we will write $\langle a_1, \dots, a_n \rangle$ for the rank n-quadratic form q given by $q(x_1, \dots, x_n) = \sum_{i=1}^n a_i x_i^2$. Set h to be the standard hyperbolic quadratic form given by $h(x, y) = xy$. (Thus $h \cong \langle 1, -1 \rangle$). For each $n \geq 0$ define

$$h_n = \begin{cases} h_n^{\oplus n/2}, & \text{if } n \text{ is even,} \\ h_n^{\oplus (n-1/2)} \oplus \langle 1 \rangle, & \text{if } n \text{ is odd.} \end{cases} \tag{10.1}$$

Set $\operatorname{Spin}_n = \operatorname{Spin}(h_n)$; this is the totally split spin group which appears in the statement of Theorem 1.7. We also denote the totally split orthogonal and special orthogonal groups by $\mathrm{O}_n \overset{\text{def}}{=} \mathrm{O}(h_n)$ and $\mathrm{SO}_n \overset{\text{def}}{=} \mathrm{SO}(h_n)$.

Now, one of the hypotheses of Theorem 1.7 is that $\zeta_4 \in k$. Therefore we can write Spin_n as $\operatorname{Spin}(q)$, where

$$q(x_1, \dots, x_n) = -(x_1^2 + \dots + x_n^2).$$

Consider the subgroup $\Gamma_n \subseteq \mathrm{SO}_n$ consisting of diagonal matrices, which is isomorphic to μ_2^{n-1}. Call G_n the inverse image of Γ_n in Spin_n. It is a constant group scheme over k. Denote by μ_2 the kernel of the homomorphism $\operatorname{Spin}_n \to \mathrm{SO}_n$.

Lemma 10.2. *Every* Spin_n*-torsor over an extension* K *of* k *admits reduction of structure to* G_n*; i.e., the natural map* $\mathrm{H}^1(K, G_n) \to \mathrm{H}^1(K, \operatorname{Spin}_n)$ *is surjective for any field extension* K/k*.*

Proof. Let $P \to \operatorname{Spec} K$ be a Spin_n-torsor: we are claiming that the K-scheme P/G_n has a rational point. We have $P/G_n = (P/\mu_2)/\Gamma_n$. However $P/\mu_2 \to \operatorname{Spec} K$ is the SO_n torsor associated with $P \to \operatorname{Spec} K$, and every SO_n-torsor admits reduction of structure group to Γ_n. ♠

This means that the natural morphism $\mathcal{B}G_n \to \mathcal{B}\operatorname{Spin}_n$ is isotropic; so from Corollary 2.7 and Proposition 2.8 we get the bounds

$$\operatorname{ed} G_n - \dim \operatorname{Spin}_n \leq \operatorname{ed} \operatorname{Spin}_n \leq \operatorname{ed} G_n; \tag{10.3}$$

cf. also [BF03, Lemma 1.9]. Of course $\dim \operatorname{Spin}_n = n(n-1)/2$; we need to compute $\operatorname{ed} G_n$. The structure of G_n is well understood; in particular, it is very clearly described in [Woo89]. Recall that Spin_n is a subgroup of the group of units in the Clifford algebra A_n of the quadratic form $-(x_1^2 + \dots +$

x_n^2). The algebra A_n is generated by elements $e_1, \dots, e_n$, with relations $e_i^2 = -1$ and $e_ie_j + e_je_i = 0$ for all $i \neq j$. The element e_i is in Pin_n, and image of e_i in O_n is the diagonal matrix with -1 as the i-th diagonal entry, and 1 as all the other diagonal entries. The kernel of the homomorphism $\mathrm{Pin}_n \to \mathrm{O}_n$ is $\{\pm 1\}$. (For background material on the theory of Clifford algebras and spin modules, we refer the reader to [Che54], [FH91, §20.2] or [Ada96]).

For any $I \subseteq \{1, \dots, n\}$ write $I = \{i_1, \dots, i_r\}$ with $i_1 < i_2 < \dots < i_r$ and set $e_I \overset{\mathrm{def}}{=} e_{i_1} \dots e_{i_r}$. The group G_n consists of the elements of A_n of the form $\pm e_I$, where $I \subseteq \{1, \dots, n\}$ has an even number of elements. The element -1 is central, and the commutator $[e_I, e_J]$ is given by

$$[e_I, e_J] = (-1)^{|I \cap J|}.$$

It is clear from this description that G_n is a 2-group of order 2^n, the commutator $[G_n, G_n] = \{\pm 1\}$ is cyclic, and the center $\mathrm{C}(G)$ is given by

$$\mathrm{C}(G_n) = \begin{cases} \{\pm 1\} \simeq \mathbb{Z}/2\mathbb{Z}, \text{ if } n \text{ is odd,} \\ \{\pm 1, \pm e_{\{1,\dots,n\}}\} \simeq \mathbb{Z}/4\mathbb{Z}, \text{ if } n \equiv 2 \pmod 4, \\ \{\pm 1, \pm e_{\{1,\dots,n\}}\} \simeq \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}, \text{ if } n \text{ is divisible by } 4. \end{cases}$$

Theorem 1.5 now tells us that

$$\mathrm{ed}(G_n) = \begin{cases} 2^{(n-1)/2}, \text{ if } n \text{ is odd,} \\ 2^{(n-2)/2}, \text{ if } n \equiv 2 \pmod 4, \\ 2^{(n-2)/2} + 1, \text{ if } n \text{ is divisible by } 4. \end{cases}$$

Substituting this into (10.3), we obtain the bounds of Theorem 1.7. ♠

Remark 10.4. The same argument can be applied to Pin groups. (For the definition of Pin groups, see, e.g., [Ada96].) Here we replace G_n the inverse image G'_n of the diagonal subgroup of O_n in Pin_n. One easily checks that

$$\mathrm{C}(G'_n) \simeq \begin{cases} \mathbb{Z}/2\mathbb{Z}, \text{ if } n \text{ is even,} \\ \mathbb{Z}/4\mathbb{Z}, \text{ if } n \equiv 1 \pmod 4, \\ \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}, \text{ if } n \equiv 3 \pmod 4). \end{cases}$$

This yields the following bounds on the essential dimensions of Pin_n:

$$2^{\lfloor n/2 \rfloor} - \frac{n(n-1)}{2} \leq \mathrm{ed}_k \mathrm{Pin}_n \leq 2^{\lfloor n/2 \rfloor}, \quad \text{if } n \not\equiv 1 \pmod 4,$$

$$2^{\lfloor n/2 \rfloor} - \frac{n(n-1)}{2} + 1 \leq \mathrm{ed}_k \mathrm{Pin}_n \leq 2^{\lfloor n/2 \rfloor} + 1, \quad \text{if } n \equiv 1 \pmod 4.$$

for any field k of characteristic $\neq 2$ containing ζ_4.

Remark 10.5. When $n \leq 14$ the lower bound of Theorem 1.7 is negative and the upper bound is much larger than the true value of $\mathrm{ed}\,\mathrm{Spin}_n$. For $n = 15$ and 16 our inequalities yield

$$23 \leq \mathrm{ed}\,\mathrm{Spin}_{15} \leq 128$$

and

$$9 \leq \operatorname{ed}\operatorname{Spin}_{16} \leq 129.$$

When $n = 16$ our lower bound coincides with the lower bound (1.8) of Reichstein–Youssin and Chernousov–Serre, while for $n = 15$ it is substantially larger. When $n \geq 17$ the exponential part of the lower bound takes over, the growth becomes fast and the gap between the lower bound and the upper bound proportionally small. For values of n close to 15 our estimates are quite imprecise; it would be interesting to improve them.

Remark 10.6. The lower bounds in Theorem 1.7 hold over any field of characteristic different from 2 (and for any form of the Spin group).

On the other hand, if we do not assume that $\zeta_4 \in k$, we have the following slightly weaker upper bound

$$\operatorname{ed}\operatorname{Spin}_n \leq 2^{\lfloor (n-1)/2 \rfloor} + n - 1$$

for the totally split form of the spin group in dimension n, over k. To prove this inequality, we observe that a generically free representation of Spin_n can be constructed by taking a spin, or half-spin, representation V of Spin_n of dimension $2^{\lfloor (n-1)/2 \rfloor}$, and adding a generically free representation W of SO_n. Since the essential dimension of SO_{n-1} is $n-1$ over any field of characteristic different from 2, there is an SO_n-equivariant dominant rational map $f\colon W \dashrightarrow X$, where $\dim X = \dim \mathrm{SO}_n + n - 1$. Now $\mathrm{id} \times f\colon V \times W \dashrightarrow V \times X$ is a Spin_n-equivariant dominant rational map $V \times W$. Consequently,

$$\begin{aligned}\operatorname{ed}\operatorname{Spin}_n &\geq \dim(V \times X) - \dim \operatorname{Spin}_n \\ &= 2^{\lfloor (n-1)/2 \rfloor} + \dim \mathrm{SO}_n + n - 1 - \dim \operatorname{Spin}_n \\ &= 2^{\lfloor (n-1)/2 \rfloor} + n - 1\,,\end{aligned}$$

as claimed.

Remark 10.7. It is natural to ask whether the inequality

$$\operatorname{ed}\operatorname{Spin}_n \geq 2^{\lfloor (n-1)/2 \rfloor} - \frac{n(n-1)}{2}$$

can be proved by a direct application of Theorem 1.7 to the exact sequence

$$1 \longrightarrow \mu_2 \longrightarrow \operatorname{Spin}_n \longrightarrow \mathrm{SO}_n \longrightarrow 1 \tag{10.8}$$

without considering the finite subgroup G_n of Spin_n. The answer is "yes."

Indeed, consider the associated coboundary map

$$\mathrm{H}^1(K, \mathrm{SO}_m) \xrightarrow{\partial_K} \mathrm{H}^2(K, \mu_2).$$

A class in $\mathrm{H}^1(K, \mathrm{SO}_m)$ is represented by a m-dimensional quadratic form q of discriminant 1 defined over K. The class of $\partial_K(q) \in \mathrm{H}^2(K, \mu_2)$ is then the Hasse-Witt invariant of q; following Lam [Lam73], we will denote it by $c(q)$. (Note that since we are assuming that -1 is a square in k, the Hasse invariant and the Witt invariant coincide; see [Lam73, Proposition V.3.20].)

Our goal is thus to show that for every $n \geq 1$ there exists a quadratic form q_n of dimension n and discriminant 1 such that $c(q_n)$ has index $2^{\lfloor (n-1)/2 \rfloor}$.

If n is even this is proved in [Mer91, Lemma 5]. (Note that in this case $c(q) \in \mathrm{H}^2(K, \mu_2)$ is the class of the Clifford algebra of q.) If $n = 2r + 1$ is odd, set $K = k(a_1, b_1, a_2, b_2, \dots)$, where $a_1, b_1, a_2, b_2, \dots$ are independent variables, and define q_{2r+1} recursively by

$$q_1 = \langle 1 \rangle \text{ and } q_{2r+1} = \langle a_r, b_r \rangle \oplus a_r b_r \rangle \otimes q_{2r-1} \rangle.$$

One easily sees by induction on r that every q_{2r+1} has discriminant 1. Moreover, a direct computation using basic properties of the Hasse-Witt invariant (see, e.g., [Lam73, V.3.15 and 3.16]) shows that $c(q_{2r+1})$ is the class of the tensor product $(a_1, b_1)_2 \otimes_K \dots \otimes_K (a_r, b_r)_2$ of quaternion algebras. This class has index 2^r, as claimed. ♠

In summary, this approach recovers the lower bound of Theorem 1.7 in the case where n is not divisible by 4. In the case where n is divisible by 4 Theorem 1.7 gives a slightly stronger lower bound.

To conclude this section, we will now prove similar bounds on the essential dimensions on half-spin groups. We begin with the following simple corollary of [CGR06, Theorem 1.1], which appears to have been previously overlooked.

Lemma 10.9. *Let G be a closed (but not necessarily connected) subgroup of GL_n defined over a field k. Assume that $\operatorname{char} k = 0$ and either k is algebraically closed or G is connected. Then $\operatorname{ed} G \leq n$.*

Proof. By [CGR06, Theorem 1.1], there exists a finite k-subgroup $S \subseteq G$ such that every G-torsor over $\operatorname{Spec} K$, admits reduction of structure to S, for any field extension K/k. This means that the morphism $\mathcal{B}S \to \mathcal{B}G$ is isotropic and thus

$$\operatorname{ed} G \leq \operatorname{ed} S\,;$$

cf. also [BF03, Lemma 1.9]. The restriction of the representation $G \subseteq \mathrm{GL}_n$ to S is faithful; since S is a finite group over a field of characteristic zero, any faithful representation of S is necessarily generically free. By Theorem 2.9, $\operatorname{ed} S \leq n$, and hence, $\operatorname{ed} G \leq n$, as claimed. ♠

Example 10.10. Suppose G/k satisfies one of the conditions of Lemma 10.9 and the centralizer $C_G(G^0)$ of the connected component of G is trivial. Then the adjoint representation of G is faithful and Lemma 10.9 tells us that $\operatorname{ed} G \leq \dim G$. In particular, this inequality is valid for every connected semisimple adjoint group G. (In the case of simple adjoint groups, a stronger bound is given by [Lem04, Theorem 1.3].)

We are now ready to proceed with our bounds on the essential dimension of half-spin groups. Recall that the *half-spin group* HSpin_n is defined, for every n divisible by 4, as $\mathrm{Spin}_n / \langle \eta \rangle$, where η is an element of the center of Spin_n different from -1. (There are two such elements, but the resulting quotients are isomorphic.)

Theorem 10.11. *(a) Suppose k is a field of characteristic* 0 *and $\zeta_4 \in k$. Then*

$$2^{(n-2)/2} - \frac{n(n-1)}{2} \leq \operatorname{ed} \mathrm{HSpin}_n \leq 2^{(n-2)/2}$$

for any positive integer n divisible by 4.

The conditions that $\operatorname{char} k = 0$ and $\zeta_4 \in k$ are used only in the proof of the upper bound. The lower bound of Theorem 10.11 remains valid for any base field k of characteristic $\neq 2$.

Proof. The group HSpin_n contains $G_n/\langle\eta\rangle \simeq G_{n-1}$, which is an extra-special group of order 2^{n-1}. By Example 9.1 $\operatorname{ed}(G_n/\langle\eta\rangle) = 2^{(n-2)/2}$ and thus

$$\operatorname{ed} \mathrm{HSpin}_n \geq \operatorname{ed}(G_n/\langle\eta\rangle) - \dim \mathrm{HSpin}_n = 2^{(n-2)/2} - \frac{n(n-1)}{2},$$

as in the proof of Theorem 1.7.

For the upper bound notice that one of the two half-spin representations of Spin_n descends to HSpin_n, and is a faithful representation of HSpin_n of dimension $2^{(n-2)/2}$. The upper bound now follows from Lemma 10.9 ♠

11. Pfister numbers

Let k be a field of characteristic not equal to 2 and write $\mathrm{W}(k)$ for the Witt ring of k; see [Lam73, Chapter 2]. Let $I = I(k)$ denote the ideal of all even dimensional forms in the Witt ring. Then, for any integer $a > 0$, I^a is generated as an abelian group by the a-fold Pfister forms [Lam73, Proposition 1.2].

Let q be a quadratic form of rank $n > 0$ whose class $[q]$ in $\mathrm{W}(k)$ lies in I^a for $a > 0$. Define the *a-Pfister number* of q to be the minimum number r appearing in a representation

$$q = \sum_{i=1}^{r} \pm p_i$$

with the p_i being a-fold Pfister forms. The *(a, n)-Pfister number* $\mathrm{Pf}_k(a, n)$ is the supremum of the a-Pfister number of q taken over all field extensions K/k and all n-dimensional forms q such that $[q] \in I^a(K)$.

We have the following easy (and probably well-known) result.

Proposition 11.1. *Let k be a field of characteristic not equal to* 2 *and let n be a positive even integer.*

(a) $\mathrm{Pf}_k(1, n) \leq n$.
(b) $\mathrm{Pf}_k(2, n) \leq n - 2$.

Proof. (a) Immediate from the identity

$$\langle a_1, a_2\rangle = \langle 1, a_1\rangle - \langle 1, -a_2\rangle = \langle\!\langle -a_1\rangle\!\rangle - \langle\!\langle a_2\rangle\!\rangle$$

in the Witt ring.

(b) Let $q = \langle a_1, \dots, a_n \rangle$ be an n-dimensional quadratic form over K. Recall that $q \in I^2(K)$ iff n is even and $d_{\pm}(q) = 1$, modulo $(K^*)^2$ [Lam73, Corollary II.2.2]. Here $d_{\pm}(q)$ is the *signed determinant* given by $(-1)^{n(n-1)/2} d(q)$ where $d(q) = \prod_{i=1}^{n} a_n$ is the determinant [Lam73, p. 38].

To explain how to write q as a sum of $n-2$ Pfister forms, we will temporarily assume that $\zeta_4 \in K$. In this case we may assume that $a_1 \dots a_n = 1$. Since $\langle a, a \rangle$ is hyperbolic for every $a \in K^*$, we see that $q = \langle a_1, \dots, a_n \rangle$ is Witt equivalent to

$$\ll a_2, a_1 \gg \oplus \ll a_3, a_1 a_2 \gg \oplus \dots \oplus \ll a_{n-1}, a_1 \dots a_{n-2} \gg \ .$$

By inserting appropriate powers of -1, we can modify this formula so that it remains valid even if we do not assume that $\zeta_4 \in K$, as follows:

$$q = \langle a_1, \dots, a_n \rangle \simeq \sum_{i=2}^{n} (-1)^i \ll (-1)^{i+1} a_i, (-1)^{i(i-1)/2+1} a_1 \dots a_{i-1} \gg \quad \spadesuit$$

We do not have an explicit upper bound on $\mathrm{Pf}_k(3, n)$; however, we do know that $\mathrm{Pf}_k(3, n)$ is finite for any k and any n. To explain this, let us recall that $I^3(K)$ is the set of all classes $[q] \in \mathrm{W}(K)$ such that q has even dimension, trivial signed determinant and trivial Hasse-Witt invariant [KMRT98].

Let n be a positive integer. Let q be a non-degenerate n-dimensional quadratic form over K whose whose signed determinant is 1. The class of q in $\mathrm{H}^1(K, \mathrm{O}_n)$ lies in $\mathrm{H}^1(K, \mathrm{SO}_n)$. We say that *$q$ admits a spin structure* if its class is in the image of $\mathrm{H}^1(K, \mathrm{Spin}_n)$ into $\mathrm{H}^1(K, \mathrm{SO}_n)$. As pointed out in Remark 10.7, the obstruction to admitting a spin structure is the Hasse-Witt invariant $c(q)$. Thus, the forms in I^3 are exactly the even dimensional forms admitting a spin structure. The following result was suggested to us by Merkurjev and Totaro.

Proposition 11.2. *Let k be a field of characteristic different from* 2*. Then* $\mathrm{Pf}_k(3, n)$ *is finite.*

Sketch of proof. Let E be a versal torsor for Spin_n over a field extension L/k; cf. [GMS03, Section I.V]. Let q_L be the quadratic form over L corresponding to E under the map $\mathrm{H}^1(L, \mathrm{Spin}_n) \to \mathrm{H}^1(L, \mathrm{O}_n)$. The 3-Pfister number of q_L is then an upper bound for the 3-Pfister number of any n-dimensional form in I^3 over any field extension K/k. ♠

Remark 11.3. For $a > 3$ the finiteness of $\mathrm{Pf}_k(a, n)$ is an open problem.

The goal of this section is to prove Theorem 1.9 stated in the introduction, which says that

$$\mathrm{Pf}_k(3, n) \geq \frac{2^{(n+4)/4} - n - 2}{7}$$

for any field k be a field of characteristic different from 2 and any positive even integer n. Since this is a lower bound on $\mathrm{Pf}_k(3, n)$, we may assume, without loss of generality that k contains ζ_4 in the proof. To simplify matters, this assumption will be in force for the remainder of this section.

For each extension K of k, denote by $\mathrm{T}_n(K)$ the image of $\mathrm{H}^1(K, \mathrm{Spin}_n)$ in $\mathrm{H}^1(K, \mathrm{SO}_n)$. We will view T_n as a functor $\mathrm{Fields}_k \to \mathrm{Sets}$. The essential dimension of this functor is closely related to the essential dimension of Spin_n.

Lemma 11.4. $\operatorname{ed}\mathrm{Spin}_n - 1 \leq \operatorname{ed}\mathrm{T}_n \leq \operatorname{ed}\mathrm{Spin}_n$.

Proof. In the language of [BF03, Definition 1.12], we have a fibration of functors

$$\mathrm{H}^1(\_, \mu_2) \rightsquigarrow \mathrm{H}^1(\_, \mathrm{Spin}_n) \longrightarrow \mathrm{T}_n(\_).$$

The first inequality then follows from [BF03, Proposition 1.13] and the second follows from Proposition 2.8 (or [BF03, Lemma 1.9]). ♠

Lemma 11.5. *Let q and q' be non-degenerate even-dimensional quadratic forms over K. Suppose that q admits a spin structure. Then $q \oplus q'$ admits a spin structure if and only if q' admits a spin structure.*

Proof. Immediate from the fact that $I^3(K)$ is an ideal of $W(K)$. ♠

Let h_K be the standard 2-dimensional hyperbolic form $h_K(x,y) = xy$ over an extension K of k. For each n-dimensional quadratic form $q \in I^3(K)$, let $\operatorname{ed}_n(q)$ denote the essential dimension of the class of q in $\mathrm{T}_n(K)$.

Lemma 11.6. *Let q be an n-dimensional quadratic form over K whose class in $W(K)$ lies in $I^3(K)$. Then for any positive integer s*

$$\operatorname{ed}_{n+2s}(h_K^{\oplus s} \oplus q) \geq \operatorname{ed}_n(q) - \frac{s(s+2n-1)}{2}.$$

Proof. Set $m \overset{\text{def}}{=} \operatorname{ed}_{n+2s}(h_K^{\oplus s} \oplus q)$. Let F be a field of definition of $h_K^{\oplus s} \oplus q$ of transcendence degree m, and let $\widetilde{q}$ be an $(n+2s)$-dimensional quadratic form with a spin structure over F such that $\widetilde{q}_K$ is K-isomorphic to $h_K^{\oplus s} \oplus q$. Let X be the Grassmannian of s-dimensional subspaces of F^{n+2s} which are totally isotropic with respect to $\widetilde{q}$; the dimension of X is precisely $s(s+2n-1)/2$.

The variety X has a rational point over K; hence there exists an intermediate extension $F \subseteq E \subseteq K$ such that $\operatorname{tr\,deg}_F E \leq s(s+2n-1)/2$, with the property that $\widetilde{q}_E$ has a totally isotropic subspace of dimension s. Then $\widetilde{q}_E$ splits as $h_E^s \oplus q'$. By Witt's Cancellation Theorem, q'_K is K-isomorphic to q; hence $\operatorname{ed}_n(q) \leq m + s(s+2n-1)/2$, as claimed. ♠

Proof of Theorem 1.9. If $n \leq 10$ then the statement is vacuous, because then $2^{(n+4)/4} - n - 2 \leq 0$, so we assume that $n \geq 12$. We may also assume without loss of generality that $\zeta_4 \in k$. In this case $\mathrm{W}(K)$ is a $\mathbb{Z}/2$-vector space; it follows that the 3-Pfister number of a form q is the smallest r appearing in an expression

$$q = \sum_{i=1}^{r} \ll a_i, b_i, c_i \gg.$$

in $\mathrm{W}(K)$. Choose an n-dimensional form q such that $[q] \in I^3(K)$ and $\mathrm{ed}_n(q) = \mathrm{ed}\,\mathrm{T}_n$. (Here we view q as an object in $T_n(K)$.) Suppose that q is equivalent in the Witt ring to $\sum_{1=1}^{r} \ll a_i, b_i, c_i \gg$.

Let us write a Pfister form $\ll a, b, c \gg$ as

$$\ll a, b, c \gg = \langle 1 \rangle \oplus \ll a, b, c \gg_0,$$

where

$$\ll a, b, c \gg_0 \overset{\text{def}}{=} \langle a_i, b_i, c_i, a_i b_i, a_i c_i, b_i c_i, a_i b_i c_i \rangle.$$

Set

$$\phi \overset{\text{def}}{=} \sum_{1=1}^{r} \ll a_i, b_i, c_i \gg_0$$

if r is even, and

$$\phi \overset{\text{def}}{=} \langle 1 \rangle \oplus \sum_{1=1}^{r} \ll a_i, b_i, c_i \gg_0$$

if r is odd. Then q is equivalent to ϕ in the Witt ring, and hence, $\phi \in I^3(K)$. The dimension of ϕ is $7r$ or $7r + 1$, depending on the parity of r.

We claim that $n < 7r$. Indeed, assume the contrary. Then the dimension of q is less than of equal to the dimension of ϕ, so q is isomorphic to a form of type $h_K^s \oplus \phi$. Thus

$$\frac{3n}{7} \geq 3r \geq \mathrm{ed}_n(q) = \mathrm{ed}\,\mathrm{T}_n \overset{\text{by Lemma 11.4}}{\geq} \mathrm{ed}\,\mathrm{Spin}_n - 1\,.$$

The resulting inequality fails for every even $n \geq 12$ because, for such n, $\mathrm{ed}\,\mathrm{Spin}_n \geq n/2$; see (1.8).

So we may assume that $7r \geq n$; then there is an isomorphism between the quadratic forms ϕ and a form of the type $h_K^{\oplus s} \oplus q$. By comparing dimensions we get the equality $7r = n + 2s$ when r is even, and $7r + 1 = n + 2s$ when r is odd. The essential dimension of the form ϕ as an element of $\mathrm{T}_{7r}(K)$ or $\mathrm{T}_{7r+1}(K)$ is at most $3r$, while Lemma 11.6 tells us that this essential dimension is at least $\mathrm{ed}(q) - s(s + 2n - 1)/2$. From this, Lemma 11.4 and Theorem 1.7 we obtain the chain of inequalities

$$\begin{aligned} 3r &\geq \mathrm{ed}_n(q) - \frac{s(s+2n-1)}{2} \\ &= \mathrm{ed}\,\mathrm{T}_n - \frac{s(s+2n-1)}{2} \\ &\geq \mathrm{ed}\,\mathrm{Spin}_n - 1 - \frac{s(s+2n-1)}{2} \\ &\geq 2^{(n-2)/2} - \frac{n(n-1)}{2} - 1 - \frac{s(s+2n-1)}{2}. \end{aligned} \tag{11.7}$$

Now suppose r is even. Substituting $s = (7r - n)/2$ into the inequality (11.7), we obtain

$$\frac{49r^2 + (14n + 10)r - 2^{(n+4)/2} - n^2 + 2n - 8}{8} \geq 0.$$

We interpret the left hand side as a quadratic polynomial in r. The constant term of this polynomial is negative for all $n \geq 8$; hence this polynomial has one positive real root and one negative real root. Denote the positive root by r_+. The above inquality is then equivalent to $r \geq r_+$. By the quadratic formula

$$r_+ = \frac{\sqrt{49 \cdot 2^{(n+4)/2} + 168n - 367} - (7n+5)}{49} \geq \frac{2^{(n+4)/4} - n - 2}{7}.$$

This completes the proof of Theorem 1.9 when r is even. If r is odd then substituting $s = (7r+1-n)/2$ into (11.7), we obtain an analogous quadratic inequality, whose positive root is

$$r_+ = \frac{\sqrt{49 \cdot 2^{(n+4)/2} + 168n - 199} - (7n+12)}{49} \geq \frac{2^{(n+4)/4} - n - 2}{7},$$

and Theorem 1.9 follows. ♠

b

References

[Ada96] J. F. Adams, *Lectures on exceptional Lie groups*, Chicago Lectures in Mathematics, University of Chicago Press, Chicago, IL, 1996, With a foreword by J. Peter May, Edited by Zafer Mahmud and Mamoru Mimura. MR MR1428422 (98b:22001)

[BF03] Grégory Berhuy and Giordano Favi, *Essential dimension: a functorial point of view (after A. Merkurjev)*, Doc. Math. **8** (2003), 279–330 (electronic).

[BR97] J. Buhler and Z. Reichstein, *On the essential dimension of a finite group*, Compositio Math. **106** (1997), no. 2, 159–179.

[BR05] G. Berhuy and Z. Reichstein, *On the notion of canonical dimension for algebraic groups*, Adv. Math. **198** (2005), no. 1, 128–171.

[BRV07] Patrick Brosnan, Zinovy Reichstein, and Angelo Vistoli, *Essential dimension and algebraic stacks*, 2007, `arXiv:math/0701903v1 [math.AG]`.

[CGR06] V. Chernousov, Ph. Gille, and Z. Reichstein, *Resolving G-torsors by abelian base extensions*, J. Algebra **296** (2006), no. 2, 561–581.

[Che54] Claude C. Chevalley, *The algebraic theory of spinors*, Columbia University Press, New York, 1954.

[CS06] Vladimir Chernousov and Jean-Pierre Serre, *Lower bounds for essential dimensions via orthogonal representations*, J. Algebra **305** (2006), no. 2, 1055–1070.

[CTKM06] Jean-Louis Colliot-Thélène, Nikita A. Karpenko, and Alexander S. Merkurjev, *Rational surfaces and canonical dimension of* PGL_6, to appear, 2006.

[FH91] William Fulton and Joe Harris, *Representation theory*, Graduate Texts in Mathematics, vol. 129, Springer-Verlag, New York, 1991, A first course, Readings in Mathematics.

[Flo] Mathieu Florence, *On the essential dimension of cyclic p-groups*, to appear, preprint available on the author's home page.

[Gar] Skip Garibaldi, *Cohomological invariants: exceptional groups and spin groups*, to appear.

[Gir71] Jean Giraud, *Cohomologie non abélienne*, Springer-Verlag, Berlin, 1971, Die Grundlehren der mathematischen Wissenschaften, Band 179.

[GMS03] Skip Garibaldi, Alexander Merkurjev, and Jean-Pierre Serre, *Cohomological invariants in Galois cohomology*, University Lecture Series, vol. 28, American Mathematical Society, Providence, RI, 2003.

[Her68] I. N. Herstein, *Noncommutative rings*, The Carus Mathematical Monographs, No. 15, Published by The Mathematical Association of America, 1968.

[JLY02] Christian U. Jensen, Arne Ledet, and Noriko Yui, *Generic polynomials*, Mathematical Sciences Research Institute Publications, vol. 45, Cambridge University Press, Cambridge, 2002, Constructive aspects of the inverse Galois problem.

[Kan06] Ming-Chang Kang, *Essential dimensions of finite groups*, `http://www.arxiv.org/abs/math.AG/0611673`, 2006.

[KM06] Nikita A. Karpenko and Alexander S. Merkurjev, *Canonical p-dimension of algebraic groups*, Adv. Math. **205** (2006), no. 2, 410–433.

[KM07] ______, *Essential dimension of finite groups*, preprint available at `http://www.mathematik.uni-bielefeld.de/lag/man/263.html`, 2007.

[KMRT98] Max-Albert Knus, Alexander Merkurjev, Markus Rost, and Jean-Pierre Tignol, *The book of involutions*, American Mathematical Society Colloquium Publications, vol. 44, American Mathematical Society, Providence, RI, 1998, With a preface in French by J. Tits.

[Kor00] V. È. Kordonskiĭ, *On the essential dimension and Serre's conjecture II for exceptional groups*, Mat. Zametki **68** (2000), no. 4, 539–547.

[Lam73] T. Y. Lam, *The algebraic theory of quadratic forms*, W. A. Benjamin, Inc., Reading, Mass., 1973, Mathematics Lecture Note Series.

[Lan65] Serge Lang, *Algebra*, Addison-Wesley Publishing Co., Inc., Reading, Mass., 1965. MR MR0197234 (33 #5416)

[Led02] Arne Ledet, *On the essential dimension of some semi-direct products*, Canad. Math. Bull. **45** (2002), no. 3, 422–427.

[Lem04] N. Lemire, *Essential dimension of algebraic groups and integral representations of Weyl groups*, Transform. Groups **9** (2004), no. 4, 337–379.

[LMB00] Gérard Laumon and Laurent Moret-Bailly, *Champs algébriques*, Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge., vol. 39, Springer-Verlag, Berlin, 2000.

[Mer91] A. S. Merkur′ev, *Simple algebras and quadratic forms*, Izv. Akad. Nauk SSSR Ser. Mat. **55** (1991), no. 1, 218–224.

[Mil80] James S. Milne, *Étale cohomology*, Princeton Mathematical Series, vol. 33, Princeton University Press, Princeton, N.J., 1980. MR MR559531 (81j:14002)

[O'N] Catherine O'Neil, *Sampling spaces and arithmetic dimension*, To appear in special volume in honor of Serge Lang to be published by Springer-Verlag.

[O'N05] ______, *Models of some genus one curves with applications to descent*, J. Number Theory **112** (2005), no. 2, 369–385.

[Rei00] Z. Reichstein, *On the notion of essential dimension for algebraic groups*, Transform. Groups **5** (2000), no. 3, 265–304.

[Rob96] Derek J. S. Robinson, *A course in the theory of groups*, second ed., Graduate Texts in Mathematics, vol. 80, Springer-Verlag, New York, 1996.

[Ros99] Markus Rost, *On the galois cohomology of* Spin(14), `http://www.mathematik.uni-bielefeld.de/~rost/spin-14.html`, 1999.

[Row88] Louis H. Rowen, *Ring theory. Vol. II*, Pure and Applied Mathematics, vol. 128, Academic Press Inc., Boston, MA, 1988.

[RY00] Zinovy Reichstein and Boris Youssin, *Essential dimensions of algebraic groups and a resolution theorem for G-varieties*, Canad. J. Math. **52** (2000), no. 5, 1018–1056, With an appendix by János Kollár and Endre Szabó.

[Ser77] Jean-Pierre Serre, *Linear representations of finite groups*, Springer-Verlag, New York, 1977, Translated from the second French edition by Leonard L. Scott, Graduate Texts in Mathematics, Vol. 42.

[Ser79] ______, *Local fields*, Graduate Texts in Mathematics, vol. 67, Springer-Verlag, New York, 1979, Translated from the French by Marvin Jay Greenberg.

[TA86] J.-P. Tignol and S. A. Amitsur, *Symplectic modules*, Israel J. Math. **54** (1986), no. 3, 266–290.

[Tit92] Jacques Tits, *Sur les degrés des extensions de corps déployant les groupes algébriques simples*, C. R. Acad. Sci. Paris Sér. I Math. **315** (1992), no. 11, 1131–1138.

[Vel00] Montserrat Vela, *Explicit solutions of Galois embedding problems by means of generalized Clifford algebras*, J. Symbolic Comput. **30** (2000), no. 6, 811–842, Algorithmic methods in Galois theory.

[Wal63] C. T. C. Wall, *Quadratic forms on finite groups, and related topics*, Topology **2** (1963), 281–298.

[Woo89] Jay A. Wood, *Spinor groups and algebraic coding theory*, J. Combin. Theory Ser. A **51** (1989), no. 2, 277–313.

(Brosnan, Reichstein) Department of Mathematics, The University of British Columbia, 1984 Mathematics Road, Vancouver, B.C., Canada V6T 1Z2

(Vistoli) Scuola Normale Superiore, Piazza dei Cavalieri 7, 56126 Pisa, Italy
E-mail address, Brosnan: brosnan@math.ubc.ca
E-mail address, Reichstein: reichst@math.ubc.ca
E-mail address, Vistoli: angelo.vistoli@sns.it