ESSENTIAL DIMENSION AND ALGEBRAIC STACKS

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

Abstract. We define and study the essential dimension of an algebraic stack. We compute the essential dimension of the stacks $\mathcal{M}_{g,n}$ and $\overline{\mathcal{M}}_{g,n}$ of smooth, or stable, n-pointed curves of genus g. We also prove a general lower bound for the essential dimension of algebraic groups with a non-trivial center. Using this, we find new exponential lower bounds for the essential dimension of spin groups and new formulas for the essential dimension of some finite p-groups. Finally, we apply the lower bound for spin groups to the theory of the Witt ring of quadratic forms over a field k.

Contents

1. Introduction 1
2. Generalities 7
3. Quotient stacks and finiteness 15
4. The essential dimension of a smooth Deligne–Mumford stack 17
5. Gerbes 22
6. Canonical dimension of smooth proper varieties 24
7. The essential dimension of a gerbe 25
8. The essential dimension of $\mathcal{M}_{g,n}$ for $(g,n) \neq (1,0)$ 28
9. Central extensions 30
10. Tate curves and the essential dimension of $\mathcal{M}_{1,0}$ 31
11. Essential dimension of p-groups I 33
12. Essential dimension of p-groups II 37
13. Spinor groups 39
14. Essential dimension of cyclic p-groups 43
15. Pfister numbers 45
References 49

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, 14H10.

[†]Supported in part by an NSERC discovery grant.

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

Definition 1.1. Let $a \in F(L)$ for L an object of Fields_k. A *field of definition* for a is an intermediate field $k \subseteq K \subseteq L$ such that a is in the image of the induced function $F(K) \to F(L)$.

The *essential dimension* $\operatorname{ed} a$ of a (with respect to L) is the minimum of the transcendence degrees $\operatorname{tr}\deg_k K$ taken over all fields of definition of a.

The *essential dimension* $\operatorname{ed} F$ *of the functor* F is the supremum of $\operatorname{ed} a$ taken over all $a \in F(L)$ with L in Fields_k.

Note that in Definition 1.1 the essential dimension of a depends on the field L. We write $\operatorname{ed} a$ instead of $\operatorname{ed}(a, L)$ to simplify the notation.

Remark 1.2. 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. On the other hand, $\operatorname{ed} F$ may be infinite even in cases of interest (see for example Theorem 1.8).

Example 1.3. Let G be an algebraic group. Consider the Galois cohomology functor $\mathrm{H}^1(*, G)$ sending K to the set $\mathrm{H}^1(K, G)$ of isomorphism classes of G-torsors over $\mathrm{Spec}(K)$. The essential dimension of this functor is a numerical invariant of G, usually denoted by $\operatorname{ed} G$. Essential dimension was originally introduced (in [BR97, Rei00]) and has since been extensively studied in this context; see, e.g., [RY00, Kor00, Led02, JLY02, BF03, Lem04, CS06, Gar06].

Definition 1.1 is due to A. Merkurjev, as is the following observation; cf. [BF03, Proposition 1.17].

Example 1.4. 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 $\operatorname{ed} F_X = \dim X$.

Note that the same is true if X is an algebraic space (see Proposition 2.15).

Many interesting naturally arising functors are not of the form discussed in Examples 1.3 or 1.4. One such example is the functor $\mathrm{Curves}_{g,n}$ that sends K into the set of isomorphism classes of n-pointed smooth algebraic curves of genus g over K. When $2g - 2 + n > 0$ this functor has a well-known extension $\overline{\mathrm{Curves}}_{g,n}$ that sends K into the set of isomorphism classes of n-pointed stable algebraic curves of genus g over K. Much of his paper was motivated by the following question.

Question 1.5. What are $\operatorname{ed} \mathrm{Curves}_{g,n}$ and $\operatorname{ed} \overline{\mathrm{Curves}}_{g,n}$?

Our starting point is the following definition.

Definition 1.6. Suppose $\mathcal{X}$ is an algebraic stack over k. The *essential dimension* of $\mathcal{X}$ is the essential dimension of the functor $F_{\mathcal{X}} \colon \mathrm{Fields}_k \to \mathrm{Sets}$ which sends a field L/k to the isomorphism classes of objects in $\mathcal{X}(L)$. We write $\operatorname{ed} \mathcal{X}$ for the *essential dimension* of the stack $\mathcal{X}$.

Note that all of the examples above may be viewed as special cases of 1.6. If $\mathcal{X}$ is a scheme of finite type (or an algebraic space), we recover Example 1.4. If $\mathcal{X} = \mathcal{B}G$, the classifying stack of G such that $\mathcal{B}G(T)$ is the category of G-torsors on T, we recover Example 1.3. Finally, Question 1.5 asks for the values of $\operatorname{ed}\mathcal{M}_{g,n}$ and $\operatorname{ed}\overline{\mathcal{M}}_{g,n}$, where $\mathcal{M}_{g,n}$ and $\overline{\mathcal{M}}_{g,n}$ are the stacks of n-pointed smooth, or stable, algebraic curves of genus g over a field k.

Remark 1.7. 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 described in Example 1.3. Of course, by Example 1.4, the essential dimension of the underlying scheme is $\dim G$.

In this paper we develop the theory of essential dimension for algebraic stacks. As a first application of this theory, we give the following answer to Question 1.5.

Theorem 1.8. *Assume that the characteristic of k is 0. Then*

$$\operatorname{ed}\mathrm{Curves}_{g,n} = \operatorname{ed}\mathcal{M}_{g,n} = \begin{cases} 2 & \text{if } (g,n) = (0,0) \text{ or } (1,1); \\ 0 & \text{if } (g,n) = (0,1) \text{ or } (0,2); \\ +\infty & \text{if } (g,n) = (1,0); \\ 5 & \text{if } (g,n) = (2,0); \\ 3g-3+n & \text{otherwise.} \end{cases}$$

Moreover for $2g-2+n>0$ we have $\operatorname{ed}\overline{\mathcal{M}}_{g,n} = \operatorname{ed}\mathcal{M}_{g,n}$.

Notice that $3g-3+n$ is the dimension of the moduli space $\mathrm{M}_{g,n}$ in the stable range $2g-2+n>0$ (or the dimension of the stack in all cases); the dimension of the moduli space represents an obvious lower bound for the essential dimension of a stack. The first four cases are precisely the ones where a generic object in $\mathcal{M}_{g,n}$ has non-trivial automorphisms, and the case $(g,n) = (1,0)$, is the only one where the automorphism group scheme of an object of $\mathcal{M}_{g,n}$ is not affine.

Our stack-theoretic formalism turns out to be useful even for studying the essential dimension of algebraic groups in the classical setting of Example 1.3. Our key result in this direction is Theorem 1.10 below.

Let

$$1 \longrightarrow Z \longrightarrow G \longrightarrow Q \longrightarrow 1 \tag{1.9}$$

denote an extension of group schemes over a field k with Z central and isomorphic to μ_n for some integer $n>1$. For every extension K of k the sequence (1.9) induces a connecting homomorphism $\partial_K\colon \mathrm{H}^1(K,Q) \to \mathrm{H}^2(K,Z)$. We define $\operatorname{ind}(G,Z)$ as the maximal value of $\operatorname{ind}\big(\partial_K(t)\big)$ as K ranges over all field extensions of k and t ranges over all torsors in $\mathrm{H}^1(K,Q)$.

(Note that $\operatorname{ind}(G, Z)$ does not depend on the choice of the isomorphism $Z \simeq \mu_n$.)

Theorem 1.10. *Let G be an extension as in* (1.9)*. Assume that n is a prime power. Then* $\operatorname{ed} G \geq \operatorname{ind}(G, Z) - \dim Q$.

Let G be a finite abstract group. We write $\operatorname{ed}_k G$ for the essential dimension of the constant group scheme G_k over the field k. Let $\exp G$ denote the exponent of G and let $\mathrm{C}(G)$ denote the center of G. One of the main consequences of Theorem 1.10 is the following result about the essential dimension of finite p-groups.

Theorem 1.11. *Let G be a p-group whose commutator $[G, G]$ is central and cyclic. Then*

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

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 $\operatorname{ed}(G) = \operatorname{rank}(G)$; cf. [RY00, Example 7.4]. For most finite groups G the best previously known lower bounds on $\operatorname{ed}(G)$ were of the form

$$\operatorname{ed}(G) \geq \operatorname{rank}(A)\,, \tag{1.12}$$

where A was taken to be an abelian subgroup A of G of maximal rank. Theorem 1.11 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 1.11 tells us that $\operatorname{ed}(G) = p$, while (1.12) yields only $\operatorname{ed}(G) \geq 2$.

Theorem 1.11 has a number of interesting consequences. One of them is that $\operatorname{ed}(G) \geq p$ for any non-abelian p-group G; see Corollary 12.3. Another is the following new bound on $\operatorname{ed} \operatorname{Spin}_n$. Here by Spin_n we will mean the totally split form of the the spin group in dimension n over a field k.

Theorem 1.13. *Suppose k is a field of characteristic $\neq 2$, and that $\sqrt{-1} \in k$. If n is not divisible by 4 then*

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

If n is divisible by 4 then

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

The lower bound in this theorem was surprising to us because previously the best known lower bound was the following result due of V. Chernousov and J.–P. Serre [CS06].

$$\operatorname{ed} \operatorname{Spin}_n \geq \begin{cases} \lfloor n/2 \rfloor + 1 & \text{if } n \geq 7 \text{ and } n \equiv 1,\, 0 \text{ or } -1 \pmod 8 \\ \lfloor n/2 \rfloor & \text{for all other } n \geq 11. \end{cases} \tag{1.14}$$

(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 [Gar06]) 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.

Note that the computation of $\operatorname{ed}\operatorname{Spin}_n$ gives an example of a split, simple, connected linear algebraic group whose essential dimension exceeds its dimension. (Note that for a simple adjoint group G, $\operatorname{ed}(G) \leq \dim(G)$; cf. Example 13.10.) It also gives an example of a split, semi-simple, connected linear algebraic group G with a central subgroup Z such that $\operatorname{ed} G > \operatorname{ed} G/Z$. This is because $\operatorname{ed}\operatorname{SO}_n = n - 1$ for $n \geq 3$; cf. [Rei00, Theorem 10.4].

Finally 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, and let $\mathrm{W}(K)$ be the Witt ring of K. Call $I(K)$ the augmentation ideal 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 a natural to ask how many form are needed. When $a = 1$ or $a = 2$ is easy to see that n Pfister forms suffice; see Proposition 15.1. We prove the following result.

Theorem 1.15. *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.

Description of contents. The rest of this paper is structured as follows.

§2 contains general results on essential dimension of algebraic stacks, which are used systematically in the rest of the paper.

§3 contains a discussion of essential dimension of quotient stacks; here we mostly rephrase known facts in our language. At the end of the section, we show finiteness of the essential dimension for a large class of algebraic stacks of finite type over a field. This class includes all Deligne–Mumford stacks and all quotient stacks of the form $[X/G]$ for G a linear algebraic group.

In §4 we prove Theorem 4.1 about essential dimension of smooth integral Deligne–Mumford stacks satisfying an appropriate separation hypothesis; it states that the essential dimension of such a stack is the sum of its dimension and the essential dimension of its generic gerbe. This somewhat

surprising result implies that the essential dimension of a non-empty open substack equals the essential dimension of the stack. In particular, it proves Theorem 1.8 in the cases where a general curve in $\mathcal{M}_{g,n}$ has no non-trivial automorphisms. It also brings into relief the important role played by gerbes in this theory.

Our main result on gerbes is Theorem 5.4, stated in §5 and proved in §6 and §7. It says that the essential dimension of a gerbe banded by μ_n, where n is a prime power, equals the index of its class in the Brauer group. Our proof is geometric: we link the essential dimension of a gerbe banded by μ_n with the canonical dimension of the associated Brauer–Severi variety, and use a result of Karpenko on the canonical dimension of Brauer–Severi varieties of prime-power index.

In §8 we use Theorems 4.1 and 5.4 to compute the essential dimensions of stacks of hyperelliptic curves. This and some special arguments complete the proof of Theorem 1.8, except for the statement that $\operatorname{ed}\mathcal{M}_{1,0} = +\infty$.

Theorem 5.4 is used again in §9, where we prove Theorem 1.10. The rest of the paper is dedicated to applications of this result.

In §10 we complete the proof of Theorem 1.8 by showing that $\operatorname{ed}\mathcal{M}_{1,0} = +\infty$. This is achieved by applying Theorem 1.10 to the group schemes of l^n-torsion points on the Tate curves, where l is a prime.

§11 contains our results on p-groups. We also answer a question of Jensen, Ledet and Yui by giving an example of a finite group G with a normal subgroup N such that $\operatorname{ed}(G/N) > \operatorname{ed} G$ [JLY02, p.204].

Theorem 1.13 is proved in §13, along with similar estimates for the essential dimensions of pin and half-spin groups.

Theorem 1.10 can also be applied to cyclic group over small fields. Little was known about the essential dimension of a cyclic group over $\mathbb{Q}$ until recently, when an important preprint [Flo06] of M. Florence appeared, computing the essential dimension of a cyclic group of order p^m, where p is a prime, over a field containing a primitive p-th root of 1; this implies that $\operatorname{ed}_{\mathbb{Q}}(\mathbb{Z}/p^m) \geq p^{m-1}$. In §14 we we recover this result as a consequence of Theorem 1.10 by making use of the Brauer-Rowen algebra, an idea we learned from [Flo06]. As a corollary of Florence's theorem we prove 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). We show that C_n and D_n have the same essential dimension if n is a prime power and k contains a primitive p-th root of 1.

Section §15 contains our application of the results on $\operatorname{ed}\operatorname{Spin}_n$ to the theory of quadratic forms. In particular, we prove Theorem 1.15.

Remark 1.16. One interesting issue that we do not address in this paper is the subject of the essential dimension at a prime p. For groups, this is defined in [RY00, Definition 6.3]. There is an obvious generalization to stacks which we leave to the reader to formulate. We hope that the main results of this paper would remain valid for the concept of essential dimension at a

prime p (with certain obvious alterations). In particular, we think it is very likely that (in the notation of [RY00]), $\operatorname{ed}(G;p)$ is given by the formula in Theorem 1.11 and $\operatorname{ed}(\operatorname{Spin}_n;2)$ is bounded by the formulas in Theorem 1.13. However, we have not checked this in detail.

Notation. In the paper, a *variety* over a field k will be a geometrically integral separated scheme of finite type over k. Cohomology groups $\mathrm{H}^i(T,\mathcal{F})$ will be taken with respect to the fppf topology unless otherwise specified.

As explained in Remark 1.7, we will write $\operatorname{ed} G$ for $\operatorname{ed} \mathcal{B}G$ and use these notations interchangeably. 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 $\operatorname{ed} G$ towards the end where we are primarily concerned with essential dimensions of algebraic groups.

We write μ_n for the groups scheme of n-th roots of unity. If k is a field, we write ζ_n for a primitive n-th root of unity in the algebraic closure of k. (Using the axiom of choice, we choose one 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, A. Merkurjev, B. Noohi, G. Pappas, D. Saltman and B. Totaro for very useful conversations. We would also like to thank M. Florence for sending us a copy of his preprint on the essential dimension of $\mathbb{Z}/p^n$ over $\mathbb{Q}(\mu_p)$.

2. Generalities

We begin by reformulating Definition 1.1 in the language of fibered categories. For this notion, we refer the reader to the Definition 3.1 of [Vis05].

For a field k, let $\mathrm{Points}_k \stackrel{\text{def}}{=} \mathrm{Fields}_k{}^{\mathrm{op}}$. We study categories $\mathcal{X}$ which are fibered over Points_k; these are a generalization of functors from Points_k to Sets. (Clearly $\mathcal{X}$ is fibered over Points_k if and only if $\mathcal{X}$ is *cofibered over* Fields_k ([Gro63, 6.10]), but we prefer to work with fibered categories.)

Definition 2.1. If ξ is an object of $\mathcal{X}(K)$, where K is an extension of k, a *field of definition* of ξ is an intermediate field $k \subseteq F \subseteq K$, such that ξ is in the essential image of the pullback functor $\mathcal{X}(F) \to \mathcal{X}(K)$.

Definition 2.2. Let $\mathcal{X}$ be a category fibered over Points_k. If K is an extension of k and ξ is an object of $\mathcal{X}(K)$, the *essential dimension of* ξ, written $\operatorname{ed}\xi$, is the least transcendence degree over k of a field of definition of ξ.

The *essential dimension of* $\mathcal{X}$, denoted by $\operatorname{ed}\mathcal{X}$, is the supremum of the essential dimension of all objects ξ in $\mathcal{X}(K)$ for all extensions K of k.

These notions are obviously relative to the base field k. (See Remark 2.4.) We will write $\operatorname{ed}(\xi/k)$ (resp. $\operatorname{ed}(\mathcal{X}/k)$) when we need to be specific about the dependence on the base field.

Note that $\operatorname{ed}\mathcal{X}$ takes values in the range $\{\pm\infty\}\cup\mathbb{Z}_{\geq 0}$, with $-\infty$ occurring if and only if $\mathcal{X}$ is empty.

Remark 2.3. With every functor $F\colon \mathrm{Fields}_k \to \mathrm{Sets}$, one can canonically associate a category $\mathcal{X}_F$ fibered over k (see [Vis05, Proposition 3.26]). It is an easy exercise in unravelling the definitions to see that $\operatorname{ed}\mathcal{X}_F$ as defined in 2.2 is equal to $\operatorname{ed} F$ as defined in 1.1.

Furthermore, given a fibered category $\mathcal{X} \to \mathrm{Aff}_k$ we get a functor

$$\overline{\mathcal{X}}\colon \mathrm{Fields}_k \longrightarrow \mathrm{Sets}$$

sending a field K into the set of isomorphism classes in $\mathcal{X}(\operatorname{Spec} K)$. It also straightforward to see that $\operatorname{ed}\mathcal{X}$ equals $\operatorname{ed}\overline{\mathcal{X}}$ as defined in 1.1.

Remark 2.4. Let L be an extension of a field K of transcendence degree d, and let $\mathcal{X} \to \mathrm{Points}_L$ be a fibered category. Then the composite of $\mathcal{X} \to \mathrm{Points}_L$ with the obvious functor $\mathrm{Points}_L \to \mathrm{Points}_K$ makes $\mathcal{X}$ into a category fibered over Points_K. We have

$$\operatorname{ed}(\mathcal{X}/K) = \operatorname{ed}(\mathcal{X}/L) + d.$$

The idea is that if F is an extension of K and ξ is an object of $\mathcal{X}(F)$, the image of ξ in Points_L defines an embedding of L in F; and every field of definition of ξ over K must contain L.

2.5. If ξ is an object in $\mathcal{X}(L)$ and $k \subset F \subset L$, then we say that ξ *descends to* F if F is a field of definition for ξ. The map $\operatorname{Spec} L \to \operatorname{Spec} F$ is then called a *compression* of ξ. Note that the identity morphism $\mathrm{id}_{\operatorname{Spec} L}$ is always a compression.

A compression $\operatorname{Spec} L \to \operatorname{Spec} F$ is called a *deflation* if $\operatorname{tr}\deg_k F < \operatorname{tr}\deg_k L$. If a deflation exists, we will say that ξ is *deflatable*. If ξ is not deflatable, then it is called *undeflatable*. To show that $\operatorname{ed}\mathcal{X} \geq d$, it suffices to exhibit an undeflatable object $\xi \in \mathcal{X}(L)$ with $\operatorname{tr}\deg_k L = d$.

We will call an undeflatable object $\xi \in \mathcal{X}(L)$ *maximally undeflatable* if $\operatorname{ed}\xi = \operatorname{ed}\mathcal{X}$. Clearly, for $d \in [0,\infty)$, we have $\operatorname{ed}(\mathcal{X}/k) = d$ if and only if there is a maximal undeflatable object $\xi \in \mathcal{X}(L)$ with $\operatorname{tr}\deg_k L = d$.

2.6. For a scheme S, we follow Laumon & Moret-Bailly [LMB00] in letting Aff_S denote the category of affine schemes over S. If $S = \operatorname{Spec} k$, then this category, which we denote by Aff_k, is equivalent to the category opposite to the category k-Alg of k-algebras. We equip Aff_S with the étale topology, and, by default, all notions of sheaves and stacks involving Aff_S are with respect to this topology.

2.7. A *stack* over a scheme S will mean a stack over Aff_S. That is, a stack over S is a category $\mathcal{X}$ fibered over Aff_S satisfying Definition 4.6 of [Vis05]. If $\mathcal{X}$ is a category fibered over Aff_k, then the restriction $\widetilde{\mathcal{X}}$ of $\mathcal{X}$ to Points_k via the obvious functor $\mathrm{Points}_k \to \mathrm{Aff}_k$ is a category fibered over Points_k. We write $\operatorname{ed}\mathcal{X} \overset{\text{def}}{=} \operatorname{ed}\widetilde{\mathcal{X}}$. This defines the notion of the essential dimension of a stack.

2.8. We use [LMB00, Definition 4.1] as our definition of an algebraic stack. That is, by an *algebraic stack* over a scheme S, we will mean a stack $\mathcal{X}$ in groupoids over Aff_S satisfying

(1) The diagonal morphism $\Delta\colon \mathcal{X} \to \mathcal{X} \times_S \mathcal{X}$ is representable, separated and quasi-compact,
(2) there is an algebraic space X over S and a morphism $X \to \mathcal{X}$ which is smooth and surjective.

We will make heavy use of the notion of gerbe. Let us recall that a category $\mathcal{X}$ fibered in groupoids over the category Aff_K is an *fppf gerbe* if the following conditions are satisfied.

(1) $\mathcal{X}$ is a stack with respect to the fppf topology.
(2) There exists a field extension K' of K such that $\mathcal{X}(\operatorname{Spec} K')$ is not empty.
(3) Given an affine scheme S over K and two objects ξ and η in $\mathcal{X}(S)$, there exists an fppf cover $\{S_i \to S\}$ such that the pullbacks ξ_{S_i} and η_{S_i} are isomorphic in $\mathcal{X}(S_i)$ for all i.

A gerbe is called *neutral* if $\mathcal{X}(\operatorname{Spec} K)$ is not empty.

We have the following easy observation.

Proposition 2.9. *Let $\mathcal{X}$ be an algebraic stack over a field k, and let $\mathcal{X}_{\mathrm{red}}$ denote the reduced substack* [LMB00, Lemma 4.10]. *Then* $\operatorname{ed} \mathcal{X}_{\mathrm{red}} = \operatorname{ed} \mathcal{X}$.

Proof. For every field K over k, the morphism $\mathcal{X}_{\mathrm{red}} \to \mathcal{X}$ induces an equivalence of categories $\mathcal{X}_{\mathrm{red}}(K) \to \mathcal{X}(K)$. ♠

A category $\mathcal{X}$ fibered over Points_k (resp. over Aff_k) is *limit preserving* if, whenever $K = \operatorname{colim} K_i$ is a filtered direct limit of fields (resp. k-algebras), $\operatorname{colim} X(\operatorname{Spec} K_i) \to X(K)$ is an equivalence of categories (see [Art74, p. 167]). Note that an algebraic stack (viewed as a category fibered over Aff_k) is limit preserving [LMB00, Proposition 4.18]. The property of being limit preserving provides the most basic instance of finiteness of essential dimension.

Proposition 2.10. *If $\mathcal{X}$ is a limit-preserving category fibered over* Points_k*, then any object ξ of $\mathcal{X}$ has finite essential dimension.*

Proof. Let K be a field. We can write it as a filtered direct limit $K = \operatorname{colim}_I K_i$ of all of its subfields K_i of finite transcendence degree. Since $\mathcal{X}$ is limit preserving, every object ξ of $X(K)$ must be in the essential image of $\mathcal{X}(K_i) \to \mathcal{X}(K)$ for some $i \in I$. Therefore $\operatorname{ed} \xi < \infty$. ♠

Let K be an extension of k; there is a tautological functor $\mathrm{Points}_K \to \mathrm{Points}_k$. We denote by $\mathcal{X}_K$ the pullback of $\mathcal{X}$ to K. This means the following. An object of $\mathcal{X}_K$ is a pair (ξ, E), where E is an extension of K, and ξ is an object of $\mathcal{X}(E)$ mapping to E in Points_k. An arrow from (ξ, E) to (ξ', E') is an embedding $E' \hookrightarrow E$ (corresponding to an arrow $E \to E'$ in Points_k) preserving K, and an arrow $\xi' \to \xi$ in $\mathcal{X}$ mapping to this embedding. The category $\mathcal{X}_K$ is fibered over Points_K.

Proposition 2.11. *Suppose that k is a field, K is an extension of k and L is an extension of K.*

(a) *Let $\mathcal{X}$ be a category fibered over* Points_k. *If ξ is an object of $\mathcal{X}(K)$ and ξ_L denotes its pullback in $\mathcal{X}(L)$, then*

$$\operatorname{ed}\xi_L \leq \operatorname{ed}\xi.$$

(b) *Let $\mathcal{X}$ be a limit-preserving fibered category over* Aff_k. *If ξ is an object of $\mathcal{X}$ and L is contained in a purely transcendental extension of K, then*

$$\operatorname{ed}\xi_L = \operatorname{ed}\xi.$$

Proof. Part (a) is obvious.

For part (b), let E be a purely transcendental extension of k containing K: then $\operatorname{ed}\xi_E \leq \operatorname{ed}\xi_L \leq \operatorname{ed}\xi$ by part (a), so it is enough to prove that $\operatorname{ed}\xi_E = \operatorname{ed}\xi$. So we can assume that L is purely transcendental over k.

We need to prove the inequality $\operatorname{ed}\xi_L \geq \operatorname{ed}\xi$. Let B be a transcendence basis of L over K. Let $E \subseteq L$ be a finitely generated subfield with $\operatorname{tr}\deg_k E = \operatorname{ed}\xi_L$, with an object ξ_E whose image in $\mathcal{X}(L)$ is isomorphic to ξ_L; then if $E \subseteq L' \subseteq L$ is an intermediate field and $\xi_{L'}$ is the image of ξ_E in $\mathcal{X}(L')$, we have $\operatorname{ed}\xi_{L'} = \operatorname{ed}\xi_L$. Since L is the directed limit of subfields of the form $K'(S)$, where $K' \subseteq K$ is a field of definition of ξ which is finitely generated over k and $S \subseteq B$ is a finite subset, after enlarging L' we can find such and L' of the form $K'(S)$. Let $\xi_{K'}$ be an object of $\mathcal{X}(K')$ whose image in $\mathcal{X}(K)$ is isomorphic to ξ. The image of $\xi_{K'}$ in $\xi_{L'}$ is not necessarily isomorphic to $\xi_{L'}$, but it will become so after enlarging K' and S. Since we have $\operatorname{ed}\xi_{K'} \geq \operatorname{ed}\xi$ and $\operatorname{ed}\xi_{L'} = \operatorname{ed}\xi_L$, we may substitute K' and L' for K and L and assume that K and L are finitely generated over k.

We may also assume that K is infinite, because otherwise K is finite over k, we have $\operatorname{ed}\xi = 0$ and the inequality is obvious. Again because $\mathcal{X}$ is limit-preserving, there will be an affine integral scheme U of finite type over k, with quotient field K, and an object ξ_U in $\mathcal{X}(U)$, whose image in $\mathcal{X}(K)$ is isomorphic to ξ. Let $\{x_1, \dots, x_n\}$ be a transcendence basis for L over K. There will exist an open affine subscheme V of $U \times \mathbb{A}^n_k$ and a dominant morphism $V \to W$ onto an integral affine scheme W which is of finite type of dimension $\operatorname{ed}\xi_L$ over k, together with an object ξ_W whose pullback in $\mathcal{X}(V)$ is isomorphic to the pullback of ξ_U along the first projection $\mathrm{pr}_1 \colon W \to V$. Since the fraction field of U is infinite, there will exist a non-empty open subscheme $U' \subseteq U$ and a section $U' \to W$ of $\mathrm{pr}_1 \colon W \to V$. From this we we see that the restriction $\xi_{U'}$ of ξ_U to U' is isomorphic to the pullback of ξ_W to $\mathcal{X}(U')$. If V' denotes the closure of the image of U' into V, we get an object $\xi_{V'}$ of $\mathcal{X}(V')$ whose image in $\mathcal{X}(U')$ is isomorphic to $\xi_{U'}$. Hence $k(V') \subseteq k(U') = K$ is a field of definition of ξ; since $\dim V' \leq \dim V = \operatorname{ed}\xi_L$ we conclude that $\operatorname{ed}\xi \leq \operatorname{ed}\xi_L$. ♠

The following observation is a variant of [BF03, Proposition 1.5]. We will use it repeatedly in the sequel.

Proposition 2.12. *Let $\mathcal{X}$ be a category fibered over* Points_k*, and let K be an extension of k. Then* $\mathrm{ed}(\mathcal{X}_K/K) \leq \mathrm{ed}(\mathcal{X}/k)$.

Proof. If L/K is a field extension, then the natural morphism $\mathcal{X}_K(L) \to \mathcal{X}(L)$ is an equivalence. Suppose than M/k is a field of definition for an object ξ in $\mathcal{X}(L)$. Then any field N containing both M and K is a field of definition for ξ. Since there is a field N/K with $\operatorname{tr}\deg_K N \leq \operatorname{tr}\deg_k M$ containing both M and K (any composite of M and K), we can find a field of definition for ξ as an object in $\mathcal{X}_K$ of transcendence degree $\leq \mathrm{ed}(\mathcal{X}/k)$. ♠

Remark 2.13. The proof shows the following: if L is an extension of K and ξ is an object in some $\mathcal{X}_K(L)$, call η the image of ξ in $\mathcal{X}$. Then $\operatorname{ed}\xi \leq \operatorname{ed}\eta$.

In some cases of interest we can arrange for equality Proposition 2.12.

Proposition 2.14. *Let $\mathcal{X}$ be a limit-preserving category fibered over* Aff_k*. Suppose one of the following conditions holds:*

(1) k is algebraically closed.
(2) K/k is purely transcendental and k is infinite,

Then $\mathrm{ed}(\mathcal{X}_K/K) = \mathrm{ed}(\mathcal{X}/k)$.

Proof. It is easy to see that, since $\mathcal{X}$ is limit-preserving, we can assume that K is finitely generated over k. Thus, $\operatorname{tr}\deg_k K < \infty$.

So pick $\xi \in \mathcal{X}(l)$ an undeflatable object for some field extension l/k with $\operatorname{tr}\deg_k l = n < \infty$. Again, since $\mathcal{X}$ is limit-preserving, we can assume that l is finitely generated over k. Set $L \stackrel{\text{def}}{=} l \otimes_k K$. Note that L is a field of transcendence degree n under either hypotheses (1) or (2). Write η for the restriction of ξ to L via the obvious map $\operatorname{Spec} L \to \operatorname{Spec} l$. We claim that, in either case (1) or (2), η is undeflatable over K.

To show this, assume that η is deflatable over K. Then there is a intermediate field $K \subset R \subsetneq L$ and an object $\gamma \in \mathcal{X}(R)$ such that the restriction of γ to L is η. Moreover $\operatorname{tr}\deg_K R < \operatorname{tr}\deg_K L$. Pick affine schemes U and V of finite type over k such that $k(U) = l$ and $k(V) = K$. In case (2) we can and will assume that $V = \mathbb{A}^m$. Since $\mathcal{X}$ is limit-preserving, we can find an affine scheme Z of finite type over k such that $k(Z) = R$ and γ is the restriction of some object $\widetilde{\gamma} \in \mathcal{X}(Z)$. Shrinking Z if necessary, we can assume that the rational map from Z to V inducing the inclusion of K into R is a morphism. We can also find a Zariski dense open $W \subset U \times V$ and a V-morphism $f\colon W \to Z$ inducing the inclusion of R into L.

For each point $v \in V$, write W_v (resp. Z_v) for the fiber of the map $Z \to V$ (resp. the map $p_2\colon W \to V$). Since $R \subsetneq L$, $\dim W > \dim Z$. It follows that there is a Zariski dense open subscheme M of V such that $\dim W_v > \dim Z_v$ for all $v \in M$.

Suppose v is a closed point in M such that $k(v) = k$. Then ξ is the restriction of γ to l via the map $\operatorname{Spec} l \to W_v \to Z_v$. It follows that ξ is deflatable over k, which is a contradiction.

Note that, in either case (1) or case (2), $M(k) \neq \emptyset$. Therefore the contradiction is always obtained. It follows that η is undeflatable over L.

Now, we are free to pick $\xi \in \mathcal{X}(l)$ with $\operatorname{ed}(\xi/k) = \operatorname{ed}(\mathcal{X}/k)$. Then, in either case (1) or (2), we have that $\operatorname{ed}(\mathcal{X}/k) = \operatorname{tr}\deg_k l = \operatorname{tr}\deg_K L = \operatorname{ed}(\mathcal{X}/L)$. This completes the proof of the statement. ♠

We will need the following generalization of Example 1.4.

Proposition 2.15. *The essential dimension of an algebraic space locally of finite type over k equals its dimension.*

Proof. Indeed, in this case 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 $\operatorname{ed} X = \max \operatorname{ed} X_i = \dim X$. ♠

2.16. If X is an algebraic space over an algebraic space S, then the category of arrows $T \to X$ where T is an object in Aff_S is fibered over Aff_S. It is equivalent to the fibered category $\mathcal{X}_{\mathrm{h}_X}$ arising from the functor $\mathrm{h}_X\colon \mathrm{Aff}_S \to \mathrm{Sets}$ given by $\mathrm{h}_X(T) = \mathrm{Mor}_S(T, X)$ via [Vis05, Proposition 3.26].

A category $\mathcal{X}$ fibered over Aff_S is said to be *representable* by an algebraic space if there is an algebraic space X over S and an equivalence of categories between $\mathcal{X}$ and $\mathcal{X}_{\mathrm{h}_X}$. We will follow the standard practice of identifying an algebraic space X with its corresponding representable stack $\mathcal{X}_{\mathrm{h}_X}$. This is permissible by Yoneda's lemma.

A morphism $f\colon \mathcal{X} \to \mathcal{Y}$ of categories fibered over S is said to be representable if, for every algebraic space T over $\mathcal{Y}$, the fiber product $\mathcal{X} \times_{\mathcal{Y}} T$ is representable as a category fibered over Aff_T.

Let d be an integer, and let k be a field. A morphism $f\colon \mathcal{X} \to \mathcal{Y}$ of categories fibered over Aff_k is said to be *representable of fiber dimension at most d* if, for every map $T \to \mathcal{Y}$ from an algebraic space, the fibered product $\mathcal{X} \times_{\mathcal{Y}} T$ is an algebraic space locally of finite type over T with fibers of relative dimension $\leq d$.

Proposition 2.17. *Let $\mathcal{X}$ and $\mathcal{Y}$ be fibered categories over k. Let d be an integer, and assume that there exists a morphism $\mathcal{X} \to \mathcal{Y}$ that is represented by morphisms locally of finite algebraic spaces, with fiber dimension at most d. Then* $\operatorname{ed}(\mathcal{X}/k) \leq \operatorname{ed}(\mathcal{Y}/k) + d$.

Proof. Let K be a field over k and let $x\colon \operatorname{Spec} K \to \mathcal{X}$ be an object of $\mathcal{X}(K)$. Then $f \circ x\colon \operatorname{Spec} K \to \mathcal{Y}$ is an object of $\mathcal{Y}(K)$, and we can find a field L with a morphism $y\colon \operatorname{Spec} L \to Y$ such that $k \subset L \subset K$, $\operatorname{tr}\deg_k L \leq \operatorname{ed} \mathcal{Y}$ and the following diagram commutes.

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

Let $\mathcal{X}_L \stackrel{\text{def}}{=} \mathcal{X} \times_{\mathcal{Y}} \operatorname{Spec} L$. By the hypothesis, $\mathcal{X}_L$ is an algebraic space, locally of finite type over L and of relative dimension at most d. By the commutativity of the above diagram, the morphism $x\colon \operatorname{Spec} K \to \mathcal{X}$ factors through $\mathcal{X}_L$. Let p denote the image of x is $\mathcal{X}_L$. Since $\mathcal{X}_L$ has dimension at most d, we have $\operatorname{tr}\deg_k k(p) \leq d$. Therefore $\operatorname{tr}\deg_k k(p) \leq \operatorname{ed}\mathcal{Y} + d$. Since x factors through $\operatorname{Spec} k(p)$ the result follows. ♠

Remark 2.18. The proof of Proposition 2.17 clearly shows the following: For any field K/k and any $\xi \in \mathcal{X}(K)$, $\operatorname{ed}\xi \leq \operatorname{ed} f(\xi) + d$.

The following simple observation will be used often in this paper.

Proposition 2.19. *Let U be an integral algebraic space locally of finite type over k with function field $K \stackrel{\text{def}}{=} k(U)$, 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*

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

Proof. If $\operatorname{Spec} L \to \mathcal{X}_K$ is maximally undeflatable over K, then the morphism $\operatorname{Spec} L \to \mathcal{X}$ obtained by composing with the canonical morphism $\mathcal{X}_K \to \mathcal{X}$ is maximally undeflatable over k. ♠

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 2.20. *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 $\mathcal{X}_K \stackrel{\text{def}}{=} \mathcal{X} \times_{\mathcal{Y}} \operatorname{Spec} K$. Then*

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

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 X_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 ξ 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 representable of fiber dimension at most $\dim U - d$,

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

♠

In general, the inequality of Proposition 2.17 only goes in one direction. However, in important special cases we can obtain an inequality in the reverse direction.

2.21. We will say that a morphism $f\colon \mathcal{X} \to \mathcal{Y}$ of categories fibered over Points_k is *isotropic* if for every 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 2.22. *Let $f\colon \mathcal{X} \to \mathcal{Y}$ be an isotropic morphism of categories fibered over* Points_k*. Then* $\operatorname{ed}\mathcal{X} \geq \operatorname{ed}\mathcal{Y}$.

Proof. 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 a field of definition for ξ is also a field of definition for η. ♠

Remark 2.23. One obvious example of an isotropic morphism is the total space of a vector bundle over a Deligne-Mumford stack. Any open substack of a vector bundle which is dense in every fiber is also isotropic.

We will use the following proposition.

Proposition 2.24. *Let $\mathcal{X}$ and $\mathcal{Y}$ be categories fibered over* Points_k*. Then*

$$\operatorname{ed}(\mathcal{X} \times_{\mathrm{Points}_k} \mathcal{Y}) \leq \operatorname{ed}\mathcal{X} + \operatorname{ed}\mathcal{Y}.$$

Proof. This is equivalent to Lemma 1.11 of [BF03]. The proof is immediate: if (ξ, η) is an object in some $(\mathcal{X} \times \mathcal{Y})(K)$, then $k \subseteq F \subseteq K$ is a field of definition for ξ with $\operatorname{tr}\deg_k F \leq \operatorname{ed}\mathcal{X}$ and $k \subseteq L \subseteq K$ is a field of definition for η with $\operatorname{tr}\deg_k L \leq \operatorname{ed}\mathcal{Y}$, then the subfield of K generated by F and L is a field of definition for (ξ, η), of transcendence degree at most $\operatorname{tr}\deg_k F + \operatorname{tr}\deg_k L \leq \operatorname{ed}\mathcal{X} + \operatorname{ed}\mathcal{Y}$. ♠

Remark 2.25. The inequality in Proposition 2.24 is often strict. For example, let $k = \mathbb{C}$, $\mathcal{X} = \mathcal{B}\mu_2$ and $\mathcal{Y} = \mathcal{B}\mu_3$. Then $\mathcal{X} \times_{\mathrm{Points}_k} \mathcal{Y} = \mathcal{B}\mu_6$. However, we have $\operatorname{ed}\mu_n = 1$ for all integers $n > 1$ by [BR97, Theorem 5.3].

3. Quotient stacks and finiteness

Suppose a linear algebraic group G is acting on an algebraic space X over a field k. We shall write $[X/G]$ for the quotient stack $[X/G]$. The functor $F_{[X/G]}$ associates to a field K/k the set isomorphism classes of diagrams

$$\begin{tikzcd} T \arrow[r, "\psi"] \arrow[d, "\pi"] & X \\ \operatorname{Spec}(K) & \end{tikzcd} \tag{3.1}$$

where π is a G-torsor and ψ is a G-equivariant map.

If G is an algebraic group over k, then $\mathcal{B}G \stackrel{\text{def}}{=} [\operatorname{Spec} k/G]$. The functor $F_{\mathcal{B}G}$ is equal to the functor $K \mapsto \mathrm{H}^1(K, G)$ sending K to the isomorphism classes of G-torsors over K.

Remark 3.2. As noted in the introduction, for G an algebraic group, the essential dimension $\operatorname{ed} \mathcal{B}G$ is equal to the essential dimension of Example 1.3 classically denoted by $\operatorname{ed} G$. To prevent confusion, we remind the reader that we will use the notations $\operatorname{ed} \mathcal{B}G$ and $\operatorname{ed} G$ interchangeably (as in Remark 1.7).

Proposition 3.3. *Let $G \to \operatorname{Spec} K$ be an algebraic group acting on an algebraic space X over K and let H be a closed subgroup of G. Then*

$$\operatorname{ed}[X/H] \leq \operatorname{ed}[X/G] + \dim G - \dim H.$$

Proof. The obvious morphism $[X/H] \to [X/G]$ has fibers of dimension $\dim G - \dim H$, so this is a consequence of Proposition 2.17. ♠

Lemma 3.4. *Suppose a linear algebraic group H is acting on an algebraic space X. If H is a subgroup of another linear algebraic group G then the quotient stacks $[X/H]$ and $[X *_H G/G]$ are isomorphic.*

Proof. Here $X *_H G$ is the quotient of $X \times G$ by the H action given by $h(x, g) = (xh^{-1}, hg)$. This fact is standard but it is as easy to prove it as it is to look for a reference.

Note that, when H acts freely on X, the quotients X/H and $X *_H G/G$ are both algebraic spaces

Let E be an object in $[X/H]$, i.e, an H-torsor over a k-scheme S equipped with an H-equivariant map to X. We associate to E the G-torsor $E *_H G$ equipped with its natural morphism to the algebraic space $X *_H G$.

On the other hand, suppose F is an object in $[X *_H G/G]$, i.e., a G-torsor over a k-scheme S equipped with a G-equivariant map to the algebraic space $X *_H G$. Consider the H-equivariant map $i\colon X \to X *_H G$ given by $x \mapsto (x, 1)$. We associate to F the H-torsor E over S defined by the pull-back diagram

$$\begin{tikzcd} E \arrow[r] \arrow[d] & X \arrow[d, "i"] \\ F \arrow[r] & X *_H G. \end{tikzcd}$$

It is not difficult to see that these operations give an equivalence of categories between $[X/H]$ and $[X *_H G/G]$. ♠

Now let $F^{\mathrm{spl}}_{[X/G]}$ be the subfunctor of $F_{[X/G]}$ defined as follows. For any field K/k, $F^{\mathrm{spl}}_{[X/G]}(K)$ consists of diagrams (3.1), where $\pi\colon T \to \operatorname{Spec}(K)$ is a split torsor.

Following [BF04], we define the *functor of orbits* $\mathbf{Orb}_{X,G}$ by $\mathbf{Orb}_{X,G}(K)$ = set of $G(K)$-orbits in $X(K)$.

Lemma 3.5. *The functors $F^{\mathrm{spl}}_{[X/G]}$ and $\mathbf{Orb}_{X,G}$ are isomorphic.*

Proof. Recall that a torsor $\pi\colon T \to \operatorname{Spec}(K)$ is split if and only if there exists a section $s\colon \operatorname{Spec}(K) \to T$.

Now we associate the $G(K)$-orbit of the K-point $\psi s\colon \operatorname{Spec}(K) \to X$ to the object (3.1) of $F^{\mathrm{spl}}_{[X/G]}$. Note that while the K-point $\psi s\colon \operatorname{Spec}(K) \to X$ depends on the choice of s, its $G(K)$-orbit does not, since any other section s' of π can be obtained from s by translating by an element of $G(K)$. Thus we have defined a map $\mathbf{Orb}_{X,G}(K) \to F^{\mathrm{spl}}_{[X/G]}(K)$ for each K/k; it is easy to see that these maps give rise to a morphism of functors

$$\mathbf{Orb}_{X,G} \longrightarrow F^{\mathrm{spl}}_{[X/G]}\,. \tag{3.6}$$

To construct the inverse map, note that a K-point $p\colon \operatorname{Spec}(K) \to X$ of X, gives rise to a G-equivariant morphism ψ from the split torsor $T = G \times \operatorname{Spec}(K)$ to X defined by $\psi\colon (g,x) \mapsto g \cdot x$. This morphism represents an object in $F^{\mathrm{spl}}_{[X/G]}(K)$. Translating $p \in X(K)$ by $g \in G(K)$ modifies ψ by composing it with an automorphism of T given by translation by g:

$$\begin{tikzcd} T \arrow[r, "\times g"] \arrow[dr, "\pi"'] & T \arrow[r, "\psi"] \arrow[d, "\pi"] & X \\ & \operatorname{Spec}(K) & \end{tikzcd}$$

It is now easy to see that the resulting map $F^{\mathrm{spl}}_{[X/G]} \to \mathbf{Orb}_{X,G}$ is a morphism of functors, inverse to (3.6). ♠

Recall that a linear algebraic group G/k is called *special* if every G-torsor over $\operatorname{Spec}(K)$ is split, for every field K/k.

Corollary 3.7. *Consider the action of a special linear algebraic group G/k on an algebraic space X locally of finite type k. Then*

(a) *The functors $F_{[X/G]}$ and $\mathbf{Orb}_{X,G}$ are isomorphic.*
(b) $\operatorname{ed}[X/G] \le \dim X$.

Proof. (a) Since G is special, $F^{\mathrm{spl}}_{[X/G]} = F_{[X/G]}$. Now apply Lemma 3.5.

(b) Let F_X be the functor $K \to X(K)$. Then sending a point $p \in X(K)$ to its $G(K)$-orbit induces a surjective morphism of functors $F_X \to \mathbf{Orb}_{X,G}$. Hence,

$$\begin{aligned} \operatorname{ed}[X/G] &= \operatorname{ed}\mathbf{Orb}_{X,G} \\ &\leq \operatorname{ed} F_X \\ &= \dim(X). \end{aligned}$$

♠

Corollary 3.8. *Let G/k be a linear algebraic group and let X/k be an algebraic space, locally of finite type over k equipped with a G-action. Then* $\operatorname{ed}[X/G] < \infty$.

Proof. Let $\rho\colon G \to \mathrm{GL}_r$ be an embedding and $Y = X *_G \mathrm{GL}_r$. By Lemma 3.4 the stacks $[X/G]$ and $[Y/\mathrm{GL}_r]$ are isomorphic. Since GL_r is special, Corollary 3.7 tells us that $\operatorname{ed}(X/G) = \operatorname{ed}(Y/\mathrm{GL}_r) \leq \dim Y < \infty$. ♠

Another consequence of Proposition 2.22 is the following "classical" theorem (see [BF03] for another proof).

Theorem 3.9. *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. ♠

Finiteness. The main theorem on finiteness of essential dimension is now an easy corollary our study of quotient stacks and of a result of A. Kresch.

Theorem 3.10. *Let $\mathcal{X}$ be an algebraic stack of finite type over k. If for any algebraically closed extension Ω of k and any object ξ of $\mathcal{X}(\Omega)$ the group scheme $\underline{\mathrm{Aut}}_\Omega(\xi) \to \operatorname{Spec}\Omega$ is affine, then* $\operatorname{ed}(\mathcal{X}/k) < \infty$.

Proof. By a Theorem of Kresch [Kre99, Proposition 3.5.9] $\mathcal{X}$ is covered by quotient stacks $[X_i/G_i]$. By Corollary 3.8, $\operatorname{ed}\mathcal{X} = \max_i \operatorname{ed}[X_i/G_i] < \infty$. ♠

Theorem 3.10 does not hold without the assumption that all the $\underline{\mathrm{Aut}}_\Omega(\xi)$ are affine. For example, by Theorem 1.8, $\operatorname{ed}\mathcal{M}_{1,0} = +\infty$. The proof of this will be given in §10, and we will also see (Theorem 10.2) that $\operatorname{ed}\mathcal{B}E = +\infty$ if E is the Tate elliptic curve over the power series field $\mathbb{C}(\!(t)\!)$.

4. The essential dimension of a smooth Deligne–Mumford stack

The goal of this section is to prove the following theorem which allows us, in several of the most interesting cases, to reduce the calculation of the essential dimension of a stack to that of the essential dimension of a gerbe over a field.

Recall that if $\mathcal{X}$ is an algebraic stack over a base scheme S, then $\mathcal{X}$ is said to be *separated over* S when the diagonal morphism $\Delta : \mathcal{X} \to \mathcal{X} \times_S \mathcal{X}$ is proper. In the case of a Deligne–Mumford stack, the diagonal morphism is always quasi-finite. So the diagonal morphism of a separated Deligne–Mumford stack is finite.

Recall also that the *inertia stack* $\mathcal{I}_{\mathcal{X}} \to \mathcal{X}$ is the fibered product

$$\mathcal{X} \times_{\mathcal{X}\times\mathcal{X}} \mathcal{X}$$

mapping to $\mathcal{X}$ via the second projection (with both maps $\mathcal{X} \to \mathcal{X} \times \mathcal{X}$ given by the diagonal). The inertia stack is a group stack, and represents the functors of isomorphisms of objects: that is, it is equivalent to the obvious fibered category over S whose objects are pairs (ξ, α), where ξ is an object over some morphism $T \to S$, and α is an automorphism of ξ in $\mathcal{X}(T)$.

We say that $\mathcal{X}$ has *finite inertia* when $\mathcal{I}_{\mathcal{X}}$ is finite over $\mathcal{X}$. A separated Deligne–Mumford stack has finite inertia; however, having finite inertia is a weaker condition than being separated. For example, when X is a scheme the inertia stack is the identity $X = X$, so X always has finite inertia, even when it is not separated.

By a result of Keel and Mori ([KM97], see also [Con]) an algebraic stack locally of finite type over $\operatorname{Spec} k$ with finite inertia has a moduli algebraic space $\mathbf{X}$, which is also locally of finite type over $\operatorname{Spec} k$. The morphism $\mathcal{X} \to \mathbf{X}$ is proper.

Theorem 4.1. *Let k be a field of characteristic 0, $\mathcal{X}$ a smooth connected Deligne–Mumford stack with finite inertia, locally of finite type over $\operatorname{Spec} k$. Let $\mathbf{X}$ the moduli space of $\mathcal{X}$, K the field of rational functions on $\mathbf{X}$. Denote by $\mathcal{X}_K$ the fibered product $\operatorname{Spec} K \times_{\mathbf{X}} \mathcal{X}$. Then we have*

$$\operatorname{ed}(\mathcal{X}/k) = \dim \mathbf{X} + \operatorname{ed}(\mathcal{X}_K/K).$$

Corollary 4.2. *If $\mathcal{X}$ is as above and $\mathcal{U}$ is an open dense substack, then $\operatorname{ed}(\mathcal{M}/k) = \operatorname{ed}(\mathcal{U}/k)$.*

Corollary 4.3. *If the conditions of the theorem are satisfied, and the generic object of $\mathcal{X}$ has no non-trivial automorphisms ($\mathcal{X}$ is an orbifold, in the topologists' terminology), then $\operatorname{ed}(\mathcal{X}/k) = \dim \mathbf{X}$.*

Corollary 4.4. *Assume that k has characteristic 0. If $g \geq 3$, or $g = 2$ and $n \geq 1$, or $g = 1$ and $n \geq 2$, then*

$$\operatorname{ed}(\mathcal{M}_{g,n}/k) = \operatorname{ed}\big(\overline{\mathcal{M}}_{g,n}/k\big) = 3g - 3 + n.$$

Proof. In all these case the automorphism group of a generic object of $\mathcal{M}_{g,n}$ is trivial, so the generic gerbe is trivial, and $\operatorname{ed} \mathcal{M}_{g,n} = \dim \mathcal{M}_{g,n}$. Similarly for $\overline{\mathcal{M}}_{g,n}$. ♠

Proof of Theorem 4.1. The equality $\operatorname{ed}(\mathcal{X}/k) \geq \dim \mathbf{X} + \operatorname{ed}(\mathcal{X}_K/K)$ is clear. Let us prove the opposite inequality.

Let F be an extension of k and ξ an object in $\mathcal{X}(F)$, corresponding to a morphism $\xi\colon \operatorname{Spec} F \to \mathcal{X}$. We need to show that the essential dimension

of ξ is less than or equal to $\dim \mathbf{X} + \operatorname{ed}(\mathcal{X}_K/K)$. Of course we can assume that F is infinite, otherwise $\operatorname{ed}\xi$ would be 0, in which case we are done.

We may also assume that $\mathbf{X}$ is an affine scheme. If it is not so, by [Knu71, II, Theorem 6.4] the composite $\operatorname{Spec} F \xrightarrow{\xi} \mathcal{X} \to \mathbf{X}$ admits a factorization $\operatorname{Spec} F \to U \to \mathbf{X}$, where U is an affine scheme and the morphism $U \to \mathbf{X}$ is étale. By substituting $\mathcal{X}$ with the pullback $\operatorname{Spec} U \times_{\mathbf{X}} \mathcal{X}$ the dimension stays the same, while the essential dimension of the generic gerbe can not increase.

We proceed by induction on the codimension in $\mathbf{X}$ of the closure of the image of the composite $\operatorname{Spec} F \xrightarrow{\xi} \mathcal{X} \to \mathbf{X}$. If this codimension is 0, then $\xi\colon \operatorname{Spec} F \to \mathcal{X}$ factors though $\mathcal{X}_K$, in which case the inequality is obvious. So we can assume that this codimension is positive, that is, the composite $\operatorname{Spec} F \xrightarrow{\xi} \mathcal{X} \to \mathbf{X}$ is not dominant.

Claim 4.5. There a morphism $\operatorname{Spec} F[\![t]\!] \to \mathcal{X}$, such that its restriction $\operatorname{Spec} F \subseteq \operatorname{Spec} F[\![t]\!] \to \mathcal{X}$ is isomorphic to ξ, and such that the image of the composite $\operatorname{Spec} F[\![t]\!] \to \mathcal{X} \to \mathbf{X}$ consists of two distinct points.

By Schlessinger's theorem, there exists a local complete noetherian F-algebra A with residue field F and is a formal versal deformation of ξ defined on A. Since $\mathcal{X}$ is not obstructed the ring A is a power series ring $F[\![t_1, \dots, t_m]\!]$. The composite $\operatorname{Spec} A \to \mathcal{X} \to \mathbf{X}$ is dominant; since F is infinite, if $a_1, \dots, a_m$ are general elements of F and the homomorphism $A \to F[\![t]\!]$ is defined by sending t_i to $a_i t$, the composite $\operatorname{Spec} F[\![t]\!] \to \operatorname{Spec} A \to \mathcal{X}$ has the required properties.

Claim 4.6. There exists a complete discrete valuation subring $R \subseteq F[\![t]\!]$ and fraction field $L \subseteq F(\!(t)\!)$, such that the following properties hold:

(a) $t \in R$,
(b) the residue field of R has transcendence degree over k at most equal to $\dim \mathbf{X} + \operatorname{ed}(\mathcal{X}_K/K)$, and
(c) The composite morphism $\operatorname{Spec} F(\!(t)\!) \subseteq \operatorname{Spec} F[\![t]\!] \to \mathcal{X}$ factors through $\operatorname{Spec} L$.

Consider the composite $\operatorname{Spec} F(\!(t)\!) \subseteq \operatorname{Spec} F[\![t]\!] \to \mathcal{X}$. The closure of its image in $\mathbf{X}$ has a codimension that is less than the codimension of the closure of $\operatorname{Spec} F$; hence by induction hypothesis there exists an intermediate field $k \subseteq E \subseteq F(\!(t)\!)$ such that $\operatorname{Spec} F(\!(t)\!) \to \mathcal{X}$ factors through $\operatorname{Spec} E$, and $\operatorname{tr}\deg_k L$ is at most $\dim \mathbf{X} + \operatorname{ed}(\mathcal{X}_K/K)$.

Set $L \stackrel{\text{def}}{=} E(t)$ and $R \stackrel{\text{def}}{=} L \cap F[\![t]\!]$; clearly R is a discrete valuation subring in L containing t. We claim that the transcendence degree of the residue field $R/\mathfrak{m}_R$ over k is less than

$$\operatorname{tr}\deg_k L \leq \dim \mathbf{X} + \operatorname{ed}(\mathcal{X}_K/K) + 1.$$

This is elementary: if $s_1, \dots, s_r$ are elements of R whose images in $R/\mathfrak{m}_R$ are algebraically independent over k, then it is easy to check that $s_1, \dots, s_r, t$

are algebraically independent over k. Thus R and L satisfy all the conditions of the claim, except completeness. By completing R we prove the claim.

Claim 4.7. The morphism $\operatorname{Spec} F[\![t]\!] \to \mathcal{X}$ factors through $\operatorname{Spec} R$.

This claim implies that $\xi\colon \operatorname{Spec} F \to \mathcal{X}$ factors through $\operatorname{Spec}(R/\mathfrak{m}_R)$, which shows that

$$\begin{aligned} \operatorname{ed}\xi &\leq \operatorname{tr}\deg\bigl((R/\mathfrak{m}_R)/k\bigr) \\ &\leq \dim \mathbf{X} + \operatorname{ed}(\mathcal{X}_K/K), \end{aligned}$$

thus proving the Theorem.

To prove the claim, let us first show that the morphism $\operatorname{Spec} F[\![t]\!] \to \mathbf{X}$ factors through $\operatorname{Spec} R$. This is trivial, since $\mathbf{X}$ is an affine scheme: if $\mathbf{X} = \operatorname{Spec} A$, the homomorphism $A \to F(\!(t)\!)$ corresponding to the composite $\operatorname{Spec} F(\!(t)\!) \subseteq \operatorname{Spec} F[\![t]\!] \to \mathcal{X} \to \mathbf{X}$ factors through $F[\![t]\!]$ and also through L, so its image in contained in $L \cap F[\![t]\!] = R$.

Now denote by $\mathcal{X}_R$ the normalization of the reduced pullback $(\operatorname{Spec} R \times_{\mathbf{X}} \mathcal{X})_{\mathrm{red}}$; by a well known theorem of Nagata, stating that the normalization of a complete local integral domain is finite, this is finite over $(\operatorname{Spec} R \times_{\mathbf{X}} \mathcal{X})_{\mathrm{red}}$. The restriction of $\mathcal{X}_R$ to $\operatorname{Spec} L \subseteq \operatorname{Spec} R$ coincides with $(\operatorname{Spec} L \times_{\mathbf{X}} \mathcal{X})_{\mathrm{red}}$; hence the morphism $\operatorname{Spec} L \to \mathcal{X}$ yields a morphism $\operatorname{Spec} L \to \mathcal{X}_R$. Thus the moduli space $\mathbf{X}_R$ of $\mathcal{X}_R$, which is integral and finite over $\operatorname{Spec} R$, admits a section over $\operatorname{Spec} L$; hence $\mathbf{X}_R = \operatorname{Spec} R$.

Since $\operatorname{Spec} F[\![t]\!]$ is normal and the morphism $\operatorname{Spec} F[\![t]\!] \to (\operatorname{Spec} R \times_{\mathbf{X}} \mathcal{X})_{\mathrm{red}}$ induced by $\operatorname{Spec} F[\![t]\!] \to \mathcal{X}$ is dominant, the morphism $\operatorname{Spec} F[\![t]\!] \to \mathcal{X}$ factors through $\mathcal{X}_R$.

Let $X_0 \to \mathcal{X}_R$ be an étale map, where X_0 is a scheme. Since R is complete, hence henselian, X_0 contains a component of the form $\operatorname{Spec} R_0$, where R_0 is a discrete valuation ring which is a finite extension of R. This component dominates $\mathcal{X}_R$, so we can assume $X_0 = \operatorname{Spec} R_0$. Set $X_1 = X_0 \times_{\mathcal{X}_R} X_0$; we have that $X_1 = \operatorname{Spec} R_1$, where R_1 is a product of discrete valuation rings. The stack $\mathcal{X}_R$ has a presentation $X_1 \rightrightarrows X_0$. If we set $\operatorname{Spec} F[\![t]\!] \times_{\mathcal{X}_R} \operatorname{Spec} R_0 = \operatorname{Spec} A$, we have that A is a product of discrete valuation rings, each of them an étale extension of $F[\![t]\!]$; this implies that the image of $t \in R$ in A is a uniformizing parameter in all of the factors of A; and this implies that t is also a uniformizing parameter in R_0. So R_0 is étale over R, because the characteristic of the base field is 0.

Now consider the morphism $\operatorname{Spec} L \to \mathcal{X}_R$: set $\operatorname{Spec} L_0 = \operatorname{Spec} L \times_{\mathcal{X}_R} X_0$ and $\operatorname{Spec} L_1 = \operatorname{Spec} L \times_{\mathcal{X}_R} X_1 = \operatorname{Spec}(L_0 \otimes_L L_0)$. Set also $S_i = R_i \otimes_R F[\![t]\!]$ and $M_i = L_i \otimes_L F(\!(t)\!)$ for $i = 0$ or 1. The ring M_i is a product of fields, S_i a product of discrete valuation rings.

The pullback of the commutative diagram

$$\begin{tikzcd}
\operatorname{Spec} L_1 \arrow[r, shift left] \arrow[r, shift right] \arrow[d] & \operatorname{Spec} L_0 \arrow[r] \arrow[d] & \operatorname{Spec} L \arrow[d] \arrow[dr] & \\
\operatorname{Spec} R_1 \arrow[r, shift left] \arrow[r, shift right] & \operatorname{Spec} R_0 \arrow[r] & \mathcal{X}_R \arrow[r] & \operatorname{Spec} R
\end{tikzcd}$$

via the morphism $\operatorname{Spec} F[\![t]\!] \to \operatorname{Spec} R$ is isomorphic to the diagram

$$\begin{tikzcd}
\operatorname{Spec} M_1 \arrow[r, shift left] \arrow[r, shift right] \arrow[d] & \operatorname{Spec} M_0 \arrow[r] \arrow[d] & \operatorname{Spec} F(\!(t)\!) \arrow[d] \arrow[dr] & \\
\operatorname{Spec} S_1 \arrow[r, shift left] \arrow[r, shift right] & \operatorname{Spec} S_0 \arrow[r] & \mathcal{X}_{F[\![t]\!]} \arrow[r] & \operatorname{Spec} F[\![t]\!].
\end{tikzcd}$$

But the morphism $\operatorname{Spec} F(\!(t)\!) \to \mathcal{X}_{F[\![t]\!]}$ extends to a morphism $\operatorname{Spec} F[\![t]\!] \to \mathcal{X}_{F[\![t]\!]}$; and this implies that M_0 is unramified over $F(\!(t)\!)$ with respect to the canonical valuation of $F(\!(t)\!)$. Hence L_0 is unramified over L; if we denote by T_i the normalization of R_i in L_i we have that T_0 is étale over R, and $T_1 = T_0 \otimes_R T_0$. The diagram

$$\begin{tikzcd}
\operatorname{Spec} L_1 \arrow[r, shift left] \arrow[r, shift right] \arrow[d] & \operatorname{Spec} L_0 \arrow[d] \\
\operatorname{Spec} R_1 \arrow[r, shift left] \arrow[r, shift right] & \operatorname{Spec} R_0
\end{tikzcd}$$

extends to a diagram

$$\begin{tikzcd}
\operatorname{Spec} T_1 \arrow[r, shift left] \arrow[r, shift right] \arrow[d] & \operatorname{Spec} T_0 \arrow[d] \\
\operatorname{Spec} R_1 \arrow[r, shift left] \arrow[r, shift right] & \operatorname{Spec} R_0
\end{tikzcd}$$

which defines the descent data for a morphism $\operatorname{Spec} R \to \mathcal{X}_R$ extending $\operatorname{Spec} L \to \mathcal{X}_R$. This proves the theorem. ♠

Remark 4.8. The stack $\mathcal{X}_K$ that appears in the statement of Theorem 4.1 can be defined in much greater generality. Let $\mathcal{X}$ be a locally noetherian integral algebraic stack. It is easy to see that all dominant maps $\operatorname{Spec} K \to \mathcal{X}$ are equivalent, in the sense of [LMB00, Definition 5.2], hence they define a point of $\mathcal{X}$, called the *generic point* of $\mathcal{X}$. Then the stack $\mathcal{X}_K$ of Theorem 4.1 is the gerbe of $\mathcal{X}$ at its generic point, in the sense of [LMB00, §11.1]. This is naturally called the *generic gerbe* of $\mathcal{X}$.

Theorem 4.1 is false in general for algebraic stacks, and also for singular Deligne–Mumford stacks.

Examples 4.9.

(a) Let k be any field. Let $G \overset{\text{def}}{=} \mathbb{G}_{\mathrm{a}} \times \mathbb{G}_{\mathrm{a}}$ act on $\mathbb{A}^3$ be the formula $(s,t)(x,y,z) = (x+sz, y+tz, z)$, and define $\mathcal{X} \overset{\text{def}}{=} [\mathbb{A}^3/G]$. Let $H \subseteq \mathbb{A}^3$ be the hyperplane defined by the equation $z = 0$. Then $\mathcal{X}$ is the union

of the open substack $[(\mathbb{A}^3 \smallsetminus H)/G] \simeq \mathbb{A}^1 \smallsetminus \{0\}$ and the closed substack $[H/G] \simeq \mathbb{A}^2 \times \mathcal{B}G$; hence its essential dimension is 2, its generic gerbe is trivial, and its dimension is 1.

(b) Let r and n be integers, $r > 1$. Assume that the characteristic of k is prime to r. Let $X \subseteq \mathbb{A}^n$ be the hypersurface defined by the equation $x_1^r + \cdots + x_n^r = 0$. Let $G \overset{\text{def}}{=} \boldsymbol{\mu}_r^n$ act on X via the formula

$$(s_1, \dots, s_n)(x_1, \dots, x_n) = (s_1 x_1, \dots, s_n x_n).$$

Set $\mathcal{X} = [X/G]$. Then $\mathcal{X}$ is the union of $[(X \smallsetminus \{0\})/G]$, which is a quasi-projective scheme of dimension $n-1$, and $[\{0\}/G] \simeq \mathcal{B}\boldsymbol{\mu}_r^n$, which has essential dimension n.

(c) The following example shows that Corollary 4.2 fails even for quotient stacks of the form $[X/G]$, where X is a complex affine variety and G is a connected complex reductive linear algebraic group.

Consider the action of $G = \mathrm{GL}_n$ on the affine space X of all $n \times n$-matrices by multiplication on the left. Since G has a dense orbit, and the stabilizer of a non-singular matrix in X is trivial, we have

$$\operatorname{ed}(\text{generic point of } [X/G]) = 0.$$

On the other hand, let Y be the locus of matrices of rank $n-1$, which forms a locally closed subscheme of X. There is a surjective GL_n-equivariant morphism $Y \to \mathbb{P}^{n-1}$, sending a matrix A into its kernel, which induces an isotropic morphism $[Y/G] \to \mathbb{P}^{n-1}$. Hence by Proposition 2.19 we have

$$\operatorname{ed}[X/G] \geq \operatorname{ed}[Y/G] \geq n-1.$$

It is not hard to see that the essential dimension of $[X/G]$ is the maximum of all the dimensions of Grassmannians of r planes in $\mathbb{C}^n$, which is $n^2/4$ if n is even, and $(n^2-1)/4$ if n is odd.

Question 4.10. One could ask to what class of curves we may apply Theorem 4.1. More specifically, let $\mathfrak{M}_g$ be the stack over $\operatorname{Spec} k$, whose objects over a k-scheme S are flat proper finitely presented maps $C \to S$, whose geometric fibers are reduced locally complete intersection irreducible curves of geometric genus g. The stack $\mathfrak{M}_g$ is an irreducible locally finitely presented smooth stack over $\operatorname{Spec} k$, but it is not Deligne–Mumford. It is not of finite type, either, and it is easy to see that $\operatorname{ed} \mathfrak{M}_g = \infty$ for all g.

However, $\mathfrak{M}_g$ will contain the largest open substack $\widetilde{\mathfrak{M}}_g$ with finite inertia: it follows from Theorem 4.1 that $\operatorname{ed} \widetilde{\mathfrak{M}}_g = \operatorname{ed} \mathcal{M}_g$. Is there a good description of $\widetilde{\mathfrak{M}}_g$? Does it contain all curves with finite automorphism groups?

5. Gerbes

In this section, we address the problem of computing the essential dimension a gerbe over a field K. The general problem seems difficult; however

we do have a formula in the case where the gerbe is banded by μ_{p^n} for p a prime.

Let G be a sheaf of abelian groups in the category Aff_K. A gerbe $\mathcal{X}$ over K is said to be *banded by* G if for any affine K-scheme S and any object ξ of $\mathcal{X}$ there is an isomorphism of group schemes $G_S \simeq \underline{\mathrm{Aut}}_S(\xi)$, which is compatible with pullbacks, in the obvious sense. (Here $\underline{\mathrm{Aut}}_S(\xi)$ denotes the group scheme of automorphisms of ξ over S.) A gerbe banded by G is neutral if and only if it is equivalent to the classifying stack $\mathcal{B}_K G$.

More generally, [Gir71] contains a notion of gerbe banded by G when G is not abelian; but we do not need the added generality, which makes the definition considerably more involved.

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

5.1. 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 n denote a non-negative integer and $\alpha \in \mathrm{H}^2(K, \mu_n)$. We define the *index* $\operatorname{ind} \alpha$ to be the index of the image on α under the composition

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

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

$$\partial\colon \mathrm{H}^1(K, \mathrm{PGL}_d) \longrightarrow \mathrm{H}^2(K, \mu_d) \tag{5.2}$$

arising from the short-exact sequence

$$1 \longrightarrow \mu_d \longrightarrow \mathrm{SL}_d \longrightarrow \mathrm{PGL}_d \longrightarrow 1. \tag{5.3}$$

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

The next two sections will be devoted to the proof of the following result.

Theorem 5.4. *Let $\mathcal{X}$ be a gerbe over a field K banded by μ_{p^m} for p a prime and k a positive integer. Then*

$$\operatorname{ed} \mathcal{X} = \operatorname{ind} [\mathcal{X}].$$

6. Canonical dimension of smooth proper varieties

The following lemma is well known.

Lemma 6.1. *Let X, Y and Z be varieties over K. Assume that Y is smooth and Z is proper. If there exist rational maps $\alpha\colon X \dashrightarrow Y$ and $\beta\colon Y \dashrightarrow Z$, then there exist a rational map $\gamma\colon X \dashrightarrow Z$.*

Proof. Immediate from Nishimura's lemma; cf. [Nis55] or [RY00, Proposition A.6]. ♠

Definition 6.2. Two smooth proper varieties X and Y are e-*equivalent* (or simply *equivalent*) if there exist rational maps $X \dashrightarrow Y$ and $Y \dashrightarrow X$.

From Lemma 6.1 above, this is in fact an equivalence relation.

Definition 6.3. If X and Y are smooth proper varieties over K, let $\mathrm{e}(X,Y)$ denote the least dimension of the closure of the image of a rational map $X \dashrightarrow Y$. We set $\mathrm{e}(X,Y) = +\infty$ if there are no rational maps $X \dashrightarrow Y$ and define $\mathrm{e}(X) = \mathrm{e}(X,X)$.

The integer $\mathrm{e}(X)$ has been introduced in [KM06], and is called *the canonical dimension of X*. (In the case where X is a Brauer-Severi variety of dimension $n-1$, this number coincides with the canonical dimension of the class of X in $H^1(K, \mathrm{PGL}_n)$, as defined in [BR05].)

Lemma 6.4. *Let X, X', Y and Y' be smooth proper varieties over K, such that X is equivalent to X' and Y is equivalent to Y'. Then $\mathrm{e}(X,Y) = \mathrm{e}(X',Y')$.*

Proof. Let $f\colon X \dashrightarrow Y$ be a rational map such that $\dim V = \mathrm{e}(X,Y)$, where $V \subseteq Y$ is the closure of the image of f. From Lemma 6.1, there exists a rational map $X' \dashrightarrow V$; by composing this with the embedding $V \subseteq Y$, we see that there a rational map $X' \dashrightarrow Y$ whose image has dimension at most $\dim V = \mathrm{e}(X,Y)$. Hence $\mathrm{e}(X',Y) \leq \mathrm{e}(X,Y)$.

On the other hand, from the same lemma we see that there a rational map $V \to Y'$; this can be composed with the dominant rational map $X \dashrightarrow V$ to obtain a rational map $X \to Y'$ whose image has dimension at most $\dim V = \mathrm{e}(X,Y)$. Hence $\mathrm{e}(X,Y') \leq \mathrm{e}(X,Y)$. We deduce that

$$\mathrm{e}(X',Y') \leq \mathrm{e}(X',Y) \leq \mathrm{e}(X,Y);$$

by symmetry, $\mathrm{e}(X,Y) = \mathrm{e}(X',Y')$. ♠

Corollary 6.5. *If X and Y are equivalent smooth proper varieties over K, then*

$$\mathrm{e}(X) = \mathrm{e}(Y) = \mathrm{e}(X,Y).$$

When we have resolution of singularities for varieties over K, then $\mathrm{e}(X)$ can also be defined as the least dimension of a smooth proper variety in the equivalence class of X.

7. The essential dimension of a gerbe

Let K be a field and let n be an integer with $n > 1$. Let $\mathcal{X} \to \operatorname{Spec} K$ be a gerbe banded by μ_n with index d. Let $P \to \operatorname{Spec} K$ be a Brauer–Severi variety of dimension $d-1$ whose class in $\mathrm{H}^1(K, \mathrm{PGL}_d)$ maps to $[\mathcal{X}] \in \mathrm{H}^2(K, \mu_n) \subseteq \operatorname{Br} K$ under the connecting homomorphism $\mathrm{H}^1(K, \mathrm{PGL}_d) \to \operatorname{Br} K$.

Theorem 7.1. $\operatorname{ed} \mathcal{X} = \mathrm{e}(P) + 1$.

Before proving the theorem we will prove the following easy corollaries.

Corollary 7.2. *If $\mathcal{X}_1$ and $\mathcal{X}_2$ are gerbes over $\operatorname{Spec} K$ banded by μ_{n_1} and μ_{n_2} respectively whose cohomology classes in $\operatorname{Br} K$ are the same, then $\operatorname{ed} \mathcal{X}_1 = \operatorname{ed} \mathcal{X}_2$.*

Proof. Clear. ♠

Corollary 7.3. *Theorem 5.4 holds. That is, if $n = p^m$, with $m > 0$, then $\operatorname{ed} \mathcal{X} = \operatorname{ind} [\mathcal{X}]$.*

Proof. If the index d is 1, then $\mathcal{X}$ is neutral. Thus $\mathcal{X} = \mathcal{B}\mu_n$ with $n > 1$. By [BF03, Example 2.3], $\operatorname{ed} \mathcal{B}\mu_n = 1$.

Assume then that $d > 1$. Then the class of $\mathcal{X}$ in $\operatorname{Br} K$ is also represented by a gerbe banded by μ_d. By Corollary 7.2, we can substitute this gerbe for $\mathcal{X}$, and assume $d = n$.

The class of $\mathcal{X}$ in $\mathrm{H}^2(K, \mu_d)$ comes from a division algebra of degree d. If P is the associated Brauer–Severi variety then $\mathrm{e}(P) = d - 1$ by a theorem of Karpenko [Kar00, Theorem 2.1] (see also [Mer03, §7.2]). ♠

Proof of Theorem 7.1. Since the exponent of ∂P divides n, there exists an invertible sheaf Λ on P whose degree when base changed to $P_{\overline{K}} \cong \mathbb{P}^{d-1}_{\overline{K}}$ is n. For each K scheme T, we denote by Λ_T the pullback of Λ to $P_T \stackrel{\text{def}}{=} P \times_K T$.

The gerbe $\mathcal{X}$ is equivalent to the gerbe whose sections over a K-scheme T consist of pairs (L, λ) where L is a an invertible sheaf on P_T, and λ is an isomorphism of sheaves of $\mathcal{O}_{P_T}$-modules between $L^{\otimes n}$ and Λ_T. We may, therefore, substitute this gerbe for $\mathcal{X}$ in proving Theorem 7.1.

Let $P^\vee$ denote the dual Brauer–Severi variety. Since P splits over $k(P^\vee)$ and $P^\vee$ splits over $K(P)$, P and $P^\vee$ are in the same e-equivalence class. Thus, $\mathrm{e}(P) = \mathrm{e}(P^\vee)$. Each point $\xi \in P^\vee$ gives, by definition, a hypersurface of degree 1 in $P_{k(\xi)}$, which we denote by H_ξ.

Claim. $\operatorname{ed} \mathcal{X} \leq \mathrm{e}(P) + 1$.

Proof. Let F be an extension of K, and let (L, λ) be a class in $\mathcal{X}(\operatorname{Spec} F)$. The degree of the pullback of L to $P_{\overline{F}} \cong \mathbb{P}^{d-1}_{\overline{F}}$ is 1. So $\mathrm{H}^0(P_F, L)$ is an n-dimensional vector space over F. Choose a non-zero section of L and let $H \subset P_F$ denote its divisor. Then H gives a morphism $\operatorname{Spec} F \to P^\vee$. We know that there exists a rational map $P^\vee \dashrightarrow P^\vee$ whose image has dimension $\mathrm{e}(P^\vee) = \mathrm{e}(P)$; call all V the closure of its image. Since there is

a rational map $P^\vee \dashrightarrow V$ and $P^\vee(F) \neq \emptyset$, we also have $V(F) \neq \emptyset$. Chose a morphism $\operatorname{Spec} F \to P^\vee$ whose image is contained in V, and call its image ξ. The transcendence degree of $K(\xi)$ over K is at most $\mathrm{e}(P)$. Then $k(\xi) \subseteq F$ and the pullback of $\mathcal{O}_{P_{k(\xi)}}(H_\xi)$ to P_F is $\mathcal{O}_{P_F}(H) \simeq L$. Fix an isomorphism of invertible sheaves $\mu\colon \mathcal{O}_{P_{k(\xi)}}(H_\xi)^{\otimes n} \simeq \Lambda_{k(\xi)}$: the pullback of μ gives an isomorphism of invertible sheaves $L^{\otimes n} \simeq \Lambda_F$, which will differ from λ by an element $a \in F^*$. Then (L, λ) is clearly defined over the field $k(\xi)(a)$, whose transcendence degree over K is at most $\mathrm{e}(P) + 1$. ♠

Now we prove that $\operatorname{ed} \mathcal{X} \geq e(P) + 1$.

Let $S \subseteq P$ be a divisor in the linear system of Λ; when pulled back to $P_{\overline{K}} \simeq \mathbb{P}^{d-1}_{\overline{K}}$, the hypersurface $S_{\overline{K}}$ has degree n. If F is an extension of K, we can determine an element of $\mathcal{X}(\operatorname{Spec} F)$ by specifying a hyperplane $H \subseteq P_F$ and a rational function $u \in k(P_F)$ whose divisor is $S_F - nH$ (here, as in what follows, we will write S_F to indicate the pullback of S to P_F): the line bundle is $\mathcal{O}_{P_F}(H)$, and the isomorphism

$$\mathcal{O}_{P_F}(H)^{\otimes n} = \mathcal{O}_{P_F}(nH) \simeq \mathcal{O}_{P_F}(S_F) = \Lambda_F$$

is given by multiplication by u. Every element of $\mathcal{X}(\operatorname{Spec} F)$ is isomorphic to one arising in this way. The hyperplane H and the rational function u are not unique, but it is easy to see that the class of u in $k(P_F)^*/k(P_F)^{*n}$ is uniquely determined by the element of $\mathcal{X}(\operatorname{Spec} F)$. This gives us an invariant, which is functorial in F.

Let $P^\vee \dashrightarrow P^\vee$ be a rational function whose image has dimension $\mathrm{e}(P)$, and call V the closure of its image in $P^\vee$. The generic point $\operatorname{Spec} k(V) \to V \subseteq P^\vee$ gives us a rational point ξ of $P^\vee_{k(V)}$, corresponding to a hyperplane $H_\xi \subseteq P_{k(V)}$; let u be a rational function on $P_{k(V)}$ whose divisor is $S_{k(V)} - nH_\xi$. Consider the element $\alpha \in \mathcal{X}\big(\operatorname{Spec} k(P)(t)\big)$ determined by the rational function tu, whose divisor is $S_{k(V)(t)} - nH_\xi$. We claim that α can not come from an extension of K whose transcendence degree is less than $\mathrm{e}(P) + 1$.

Let F be a subfield of $k(V)(t)$ containing K such that α is defined over F; we need to show that the transcendence degree of F over K is $\mathrm{e}(P) + 1$. Let Z be an integral variety over K with quotient field F. Since $\mathcal{X}$ has a section on $F = k(Z)$ there exists a hyperplane $H \subseteq P_{k(Z)}$, which gives a rational map $\psi\colon Z \dashrightarrow P^\vee$. By blowing up the base locus we can assume that ψ is a morphism.

Consider the rational map $\phi\colon V \times_K \mathbb{P}^1 \dashrightarrow Z$ corresponding to the embedding $k(Z) \subseteq k(V)(t)$. The class of tu in $k(P_{k(V)(t)})^*/k(P_{k(V)(t)})^{*n}$ comes from $k(P_{k(Z)})^*/k(P_{k(Z)})^{*n}$; hence there exist rational functions $v \in k(P_{k(Z)})^*$ and $w \in k(V_{k(P)(t)})^*$ such that

$$tu = w^n(\mathrm{id}_P \times \phi)^* v \in k(P_{k(V)(t)}) = k(P \times V \times \mathbb{P}^1).$$

Now, the valuation of tu along the divisor $D \overset{\text{def}}{=} P \times V \times \{0\} \subseteq P \times V \times \mathbb{P}^1$ is 1; since $n > 1$, the valuation of $(\mathrm{id}_P \times \phi)^* v$ at D can not be 0. Hence D can not dominate $P \times Z$ under the map $\mathrm{id}_P \times \phi$, or, equivalently, $V \times \{0\}$

can not dominate Z under ϕ. The restriction of ϕ gives a regular function $V = V \times \{0\} \to Z$. The composite

$$V \xrightarrow{\phi} Z \xrightarrow{\psi} P^\vee,$$

which is well defined because ψ is a morphism, has image closure of dimension less than $\dim Z$: hence, if $\dim Z \leq \mathrm{e}(P) = \mathrm{e}(P^\vee)$, by composing with the dominant rational map $P^\vee \dashrightarrow V$ we obtain a rational map $P^\vee \dashrightarrow P^\vee$ with image closure less than $\mathrm{e}(P^\vee)$, a contradiction. So the dimension of Z, which equals the transcendence degree of F over K, is $\mathrm{e}(P) + 1$, as claimed.

This completes the proof of Theorem 7.1 and, thus, the proof of Theorem 5.4. ♠

When the index of P is not a prime power, the essential dimension of P is smaller than the index. In fact, let m be the index of P, and consider the prime decomposition $m = p_1^{a_1} \dots p_r^{a_r}$. Then the class of P in $\operatorname{Br} K$ is the product of the classes $\alpha_1, \dots, \alpha_r$ of indices $p_1^{a_1}, \dots, p_r^{a_r}$. If $P_1, \dots, P_r$ are Brauer–Severi varieties with classes $\alpha_1, \dots, \alpha_r$, then the splitting fields of P are exactly the fields that split all of the P_i; hence $P_1 \times \dots \times P_r$ has a point over $k(P)$, and P has a point over $k(P_1 \times \dots \times P_r)$. So P and $P_1 \times \dots \times P_r$ are equivalent, and we have

$$\begin{aligned} \mathrm{e}(P) &= \mathrm{e}(P_1 \times \dots \times P_r) \\ &\leq p_1^{a_1} + \dots + p_r^{a_r} - r. \end{aligned}$$

In [CTKM06], Colliot-Thélène, Karpenko and Merkurjev conjecture that equality always holds and prove it for $m = 6$. This can be reformulated in the language of essential dimension of gerbes in the following fashion.

Conjecture 7.4. *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 this holds for $r = 1$, the conjecture can 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

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

Or, 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 $\mathrm{e}(X \times Y) = \mathrm{e}(X) + \mathrm{e}(Y)$? A positive answer to this question would imply the conjecture above.

8. The essential dimension of $\mathcal{M}_{g,n}$ for $(g,n) \neq (1,0)$

In this section we complete the proof of Theorem 1.8 when $(g,n) \neq (1,0)$. By Corollary 4.4, it will suffice to compute $\operatorname{ed}\mathcal{M}_{0,0}$, $\operatorname{ed}\mathcal{M}_{0,1}$ and $\operatorname{ed}\mathcal{M}_{0,2}$, $\operatorname{ed}\mathcal{M}_{1,1}$ and $\operatorname{ed}\mathcal{M}_{2,0}$.

It is easy to see that $\operatorname{ed}\mathcal{M}_{0,1} = \operatorname{ed}\mathcal{M}_{0,2} = 0$. Indeed, a smooth curve C of genus 0 with one or two rational points over an extension K of k is isomorphic to $(\mathbb{P}^1_k, 0)$ or $(\mathbb{P}^1_k, 0, \infty)$, hence it is defined over k. Alternatively, $\mathcal{M}_{0,2} = \mathcal{B}\mathbb{G}_{\mathrm{m}}$ and $\mathcal{M}_{0,1} = \mathcal{B}(\mathbb{G}_{\mathrm{m}} \ltimes \mathbb{G}_{\mathrm{a}})$, and the groups $\mathbb{G}_{\mathrm{m}}$ and $\mathbb{G}_{\mathrm{m}} \ltimes \mathbb{G}_{\mathrm{a}}$ are special (and hence have essential dimension 0).

We will now consider the remaining cases, starting with $\mathcal{M}_{0,0}$.

Since $\mathcal{M}_{0,0} \simeq \mathcal{B}\mathrm{PGL}_2$, the fact that $\operatorname{ed}\mathcal{M}_{0,0} = 2$ is classical. We recall the following result.

Proposition 8.1. *Let k be an algebraically closed field of characteristic 0. Then $\operatorname{ed}\mathrm{PGL}_n = 2$ for $n = 2, 3$ and 6.*

This is [Rei00, Lemma 9.4 (c)]. However, note that the proof does not really require the field k to be algebraically closed of characteristic 0: it goes through whenever $\operatorname{char} k$ does not divide n and k contains a primitive n-th root of unity. Thus we have the following.

Corollary 8.2. *Let k be a field of characteristic not equal to 2. Then $\operatorname{ed}\mathcal{M}_{0,0} = \operatorname{ed}\mathcal{B}\mathrm{PGL}_2 = 2$.*

This can also be proved directly very easily: the inequality $\operatorname{ed}\mathcal{M}_{0,0} \leq 2$ holds because every smooth curve of genus 0 over a field K is a conic in $\mathbb{P}^2_K$, and can be defined by an equation of the type $ax^2 + by^2 + x^2 = 0$ for some a, $b \in K$, hence is defined over $k(a,b)$. The opposite inequality follows from Tsen's theorem.

We will now proceed to compute the essential dimension of $\mathcal{M}_{1,1}$.

Proposition 8.3. *If k is a field of characteristic not equal to 2 or 3, then*

$$\operatorname{ed}\mathcal{M}_{1,1} = 2.$$

Proof. Every elliptic curves over a field K can be written as a cubic in $\mathbb{P}^2_K$ with equation $yz = x^3 + axz^2 + bz^3$, so it is defined over $k(a,b)$. Hence $\operatorname{ed}\mathcal{M}_{1,1} \leq 2$.

Let $\mathcal{M}_{1,1} \to \mathbb{A}^1_k$ denote the map given by the j-invariant and let $\mathcal{X}$ denote the pull-back of $\mathcal{M}_{1,1}$ to the generic point $\operatorname{Spec} k(j)$ of $\mathbb{A}^1$. Then $\mathcal{X}$ is banded by μ_2 and neutral by [Sil86, Proposition 1.4 (c)], and so $\operatorname{ed}\mathcal{X} = \operatorname{ed}\mathcal{B}_{k(j)}\mu_2 = 1$. This implies what we want. ♠

Ir remains to compute the essential dimension of $\mathcal{M}_{2,0}$. The equality $\operatorname{ed}\mathcal{M}_{2,0} = 5$ is a special case of the following more general result.

Theorem 8.4. *Let $\mathcal{H}_g$ denote the stack of hyperelliptic curves of genus $g > 1$ over a field k of characteristic 0 and let $\overline{\mathcal{H}}_g$ denote its closure in $\overline{\mathcal{M}}_g$.*

Then

$$\operatorname{ed}\mathcal{H}_g = \operatorname{ed}\overline{\mathcal{H}}_g = \begin{cases} 2g & \textit{if } g \textit{ is odd,} \\ 2g+1 & \textit{if } g \textit{ is even.} \end{cases}$$

Since $\mathcal{H}_2 = \mathcal{M}_{2,0}$ and $\overline{\mathcal{H}}_g = \overline{\mathcal{M}}_2$, Theorem 8.4 completes the proof of every case of Theorem 1.8, except for the fact that $\operatorname{ed}\mathcal{M}_{1,0} = +\infty$.

Proof. The closure $\overline{\mathcal{H}}_g$ is well known to be smooth, so by Corollary 4.2 it is enough to prove the statement about $\mathcal{H}_g$. Denote by $\mathbf{H}_g$ the moduli space of $\mathcal{H}_g$; the dimension of $\mathcal{H}_g$ is $2g-1$. Let K be the field of rational functions on $\mathbf{H}_g$, and denote by $(\mathcal{H}_g)_K \stackrel{\text{def}}{=} \operatorname{Spec} K \times_{\mathbf{H}_g} \mathcal{H}_g$ the generic gerbe of $\mathcal{H}_g$. From Theorem 4.1 we have

$$\operatorname{ed}\mathcal{H}_g = 2g - 1 + \operatorname{ed}\big((\mathcal{H}_g)_K/K\big),$$

so we need to show that $\operatorname{ed}\big((\mathcal{H}_g)_K/K\big)$ is 1 if g is odd, 2 if g is even.

For this we need some standard facts about stacks of hyperelliptic curves, which we recall.

Call $\mathcal{D}_g$ the stack over K whose object over a K-scheme S are pairs $(P \to S, \Delta)$, where $P \to S$ is a conic bundle (that is, a Brauer–Severi scheme of relative dimension 1), and $\Delta \subseteq P$ is a Cartier divisor which is étale of degree $2g+2$ over S. Every family $\pi\colon C \to S$ in $\mathcal{H}(S)$ comes with a unique flat morphism $C \to P$ of degree 2, where $P \to S$ is a smooth conic bundle; denote by $\Delta \subseteq P$ its ramification locus. Sending $\pi\colon C \to S$ to $(P \to S, \Delta)$ gives a morphism $\mathcal{H}_g \to \mathcal{D}_g$. Recall the usual description of ramified double covers: if we split $\pi_*\mathcal{O}_C$ as $\mathcal{O}_P \oplus L$, where L is the part of trace 0, then multiplication yields an isomorphism $L^{\otimes 2} \simeq \mathcal{O}_P(-\Delta)$. Conversely, given an object $(P \to S, \Delta)$ of $\mathcal{D}_g(S)$ and a line bundle L on P, with an isomorphism $L^{\otimes 2} \simeq \mathcal{O}_P(-\Delta)$, the direct sum $\mathcal{O}_P \oplus L$ has an algebra structure, whose relative spectrum is a smooth curve $C \to S$ with a flat map $C \to P$ of degree 2.

The morphism $\mathcal{H}_g \to \mathbf{H}_g$ factors through $\mathcal{D}_g$, and the morphism $\mathcal{D}_g \to \mathbf{H}_g$ is an isomorphism over the non-empty locus of divisors on a curve of genus 0 with no non-trivial automorphisms (this is non-empty because $g \geq 2$, hence $2g+2 \geq 5$). Call $(P \to \operatorname{Spec} K, \Delta)$ object of $\mathcal{D}_g(\operatorname{Spec} K)$ corresponding the generic point $\operatorname{Spec} K \to \mathbf{H}_g$. It is well known, and easy to show, that $P(K) = \emptyset$. By the description above, the gerbe $(\mathcal{H}_g)_K$ is the stack of square roots of $\mathcal{O}_P(-\Delta)$, which is banded by μ_2. When g is odd then there exists a line bundle of degree $g+1$ on P, whose square is isomorphic to $\mathcal{O}_P(-\Delta)$; this gives a section of $(\mathcal{H}_g)_K$, which is therefore isomorphic to $\mathcal{B}_K\mu_2$, whose essential dimension over μ_2 is 1. If g is even then such a section does not exist, and the stack is isomorphic to the stack of square roots of $\omega_{P/K}$, whose class in $\mathrm{H}^2(K, \mu_2)$ represents the image in $\mathrm{H}^2(K, \mu_2)$ of the class $[P]$ in $\mathrm{H}^1(K, \mathrm{PGL}_2)$ under the non-abelian boundary map $\mathrm{H}^1(K, \mathrm{PGL}_2) \to \mathrm{H}^2(K, \mu_2)$. According to Theorem 5.4 its essential dimension is the index of $[P]$, which equals 2. ♠

9. Central extensions

The rest of this paper will rely on our analysis of the following situation which we recall from the introduction (1.9).

Let

$$1 \longrightarrow Z \longrightarrow G \longrightarrow Q \longrightarrow 1 \tag{9.1}$$

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

We are now going to prove Theorem 1.10 from the introduction which we restate for the convenience of the reader.

Theorem 9.2. *Let G be an extension as in* (9.1)*. Assume that n is a prime power. Then*

$$\operatorname{ed}(\mathcal{B}G/k) \geq \operatorname{ind}(G, Z) - \dim Q.$$

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 $-\dim Q$. 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 2.20, $\operatorname{ed}(\mathcal{X}/K) \leq \operatorname{ed}(\mathcal{B}G/k) + \dim Q$. 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 5.4, $\operatorname{ed}(\mathcal{X}/K) = \operatorname{ind} \partial_K(t)$. By substitution, $\operatorname{ind} \partial_K(t) - \dim Q \leq \operatorname{ed}(\mathcal{B}G/k)$. Since this inequality holds for all field extensions K/k and all Q-torsors t over K, the result follows. ♠

Remark 9.3. An affirmative answer to Conjecture 7.4 would yield an inequality similar to the one in 9.2 without the assumption that n is a prime power: Let $\operatorname{ind}(G, Z) = \prod p_i^{a_i}$ be the prime factorization of $\operatorname{ind}(G, Z)$. The conjecture would imply that

$$\operatorname{ed}(\mathcal{B}G/k) \geq 1 - \dim Q + \sum (p_i^{a_i} - 1).$$

As remarked in §7, the conjecture is a theorem in the case that the index is 6 ([CTKM06, Theorem Theorem 1.3]). We therefore have that

$$\operatorname{ed}(\mathcal{B}G/k) \geq 4 - \dim Q\,.$$

Remark 9.4. Suppose G is a simple algebraic group whose center Z is cyclic. It is tempting to apply Theorem 9.2 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 $t \in H^1(K, G^{\mathrm{ad}})$, the central simple algebra representing $\partial_K(t) \in H^2(K, Z)$ is called *the Tits*

algebra of t. The possible values of the index of the Tits algebra were studied in [Tit92], where it is denoted by $b(X)$ (for 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 9.2 becomes vacuous. The only exception are groups of types B and D, in which case Theorem 9.2 does indeed, give interesting bounds; cf. Remark 13.7.

10. Tate curves and the essential dimension of $\mathcal{M}_{1,0}$

Our first application of Theorem 9.2 is to finish the proof of Theorem 1.8 from the introduction by showing that $\operatorname{ed}\mathcal{M}_{1,0} = +\infty$.

Note that by $\mathcal{M}_{1,0}$ we mean the moduli stack of genus 1 curves, not the moduli stack $\mathcal{M}_{1,1}$ of elliptic curves (which is Deligne-Mumford). The objects of $\mathcal{M}_{1,0}$ are torsors for elliptic curves as opposed to the elliptic curves which appear as the objects of $\mathcal{M}_{1,1}$. We will now see that these torsors are what causes the essential dimension to be infinite.

10.1. Let R be a complete discrete valuation ring with function field K and uniformizing parameter q. For simplicity, we will assume that $\operatorname{char} K = 0$. Let $E = E_q/K$ denote the Tate curve over K [Sil86, §4]. This is an elliptic curve over K with the property that, for every finite field extension L/K, $E(L) \cong L^*/q^{\mathbb{Z}}$. It follows that the kernel $E[n]$ of multiplication by an integer $n > 0$ fits canonically into a short exact sequence

$$0 \longrightarrow \mu_n \longrightarrow E[n] \longrightarrow \mathbb{Z}/n \longrightarrow 0.$$

Let $\partial\colon \mathrm{H}^0(K, \mathbb{Z}/n) \to \mathrm{H}^1(K, \mu_n)$ denote the connecting homomorphism. Then it is well-known (and easy to see) that $\partial(1) = q \in \mathrm{H}^1(K, \mu_n) \cong K^*/(K^*)^n$.

Theorem 10.2. *Let $E = E_q/K$ denote the Tate curve over a field K as in* (10.1)*. Then*

$$\operatorname{ed} E = +\infty.$$

Theorem 10.2 is an immediate consequence of the following statement.

Lemma 10.3. *Let $E = E_q$ be a Tate curve as in* (10.1) *and let l be a prime integer not equal to* $\operatorname{char} R/q$*. Then, for any integer $n > 0$,*

$$\operatorname{ed} E[l^n] = l^n.$$

Proof. We first show that $\operatorname{ed} E[l^n] \geq l^n$.

Let $R' \stackrel{\text{def}}{=} R[1^{1/l^n}]$ with fraction field $K' = K[1^{1/l^n}]$. Since l is prime to the residue characteristic, R' is a complete discrete valuation ring, and the Tate curve E_q/K' is the pullback to K' of E_q/K. Since $\operatorname{ed}(E_q/K') \leq \operatorname{ed}(E_q/K)$, it suffices to prove the lemma with K' replacing K. In other words, it suffices to prove the lemma under the assumption that K contains the l^n-th roots of unity.

In that case, we can pick a primitive l^n-th root of unity ζ and write $\mu_{l^n} = \mathbb{Z}/l^n$. Let $L = K(t)$ and consider the class $(t) \in \mathrm{H}^1(L, \mu_{l^n}) = L^*/(L^*)^n$.

It is not difficult to see that

$$\partial_K(t) = q \cup (t).$$

It is also not difficult to see that the order of $q \cup (t)$ is l^n (as the map $\alpha \mapsto \alpha \cup (t)$ is injective by cohomological purity). Therefore $\operatorname{ind}(q \cup (t)) = l^n$. It follows that $\operatorname{ind}(E[l^n], \mu_{l^n}) \geq l^n$. Then, since $\dim \mathbb{Z}/l^n = 0$, Theorem 9.2 implies that $\operatorname{ed} \mathcal{B}E[l^n] \geq l^n$.

To see that $\operatorname{ed} \mathcal{B}E[l^n] \leq l^n$, note that $E[l^n]$ admits an l^n-dimensional generically free representation $V = \operatorname{Ind}_{\mu_{l^n}}^{E[l^n]} \chi$ where $\chi \colon \mu_{l^n} \to \mathbb{G}_m$ is the tautological character. Thus, by Theorem 3.9, we have the desired inequality. ♠

Proof of Theorem 10.2. For each prime power l^n, the morphism $\mathcal{B}E[l^n] \to \mathcal{B}E$ is representable of fiber dimension 1. We therefore have

$$\begin{aligned} \operatorname{ed} E &\geq \operatorname{ed} \mathcal{B}E[l^n] \\ &= l^n - 1 \end{aligned}$$

for all n. ♠

G. Pappas pointed out the following corollary.

Corollary 10.4. *Let E be a curve over a number field K. Assume that there is at least one prime $\mathfrak{p}$ of K where E has semistable bad reduction. Then* $\operatorname{ed} E = +\infty$.

It seems reasonable to make the following guess.

Conjecture 10.5. *If E is an elliptic curve over a number field, then* $\operatorname{ed} E = +\infty$.

Remark 10.6. Note, however, that, if A is a d-dimension complex abelian variety, then $\operatorname{ed} A = 2d$; see [Bro].

Now we can complete the proof of Theorem 1.8.

Theorem 10.7. *Let k be a field. Then* $\operatorname{ed}(\mathcal{M}_{1,0}/k) = +\infty$.

Proof. Consider the morphism $\mathcal{M}_{1,0} \to \mathcal{M}_{1,1}$ which sends a genus 1 curve to its Jacobian. Let $F = k(\!(t)\!)$ and let E denote the Tate elliptic curve over F, which is classified by a morphism $\operatorname{Spec} F \to \mathcal{M}_{1,1}$. We have a Cartesian diagram:

$$\begin{tikzcd} \mathcal{B}E \arrow[r] \arrow[d] & \mathcal{M}_{1,0} \arrow[d] \\ \operatorname{Spec} F \arrow[r] & \mathcal{M}_{1,1}. \end{tikzcd} \tag{10.8}$$

It follows that $+\infty = \operatorname{ed} \mathcal{B}E \leq \operatorname{ed} \mathcal{M}_{1,0}$. ♠

11. Essential dimension of p-groups I

The goal of this section is to prove the following theorem from the introduction (Theorem 1.11).

Theorem 11.1. *Let G be a p-group whose commutator $[G,G]$ is central and cyclic. Then*

(a) *We have*
$$\operatorname{ed}_k G \geq \sqrt{|G/\mathrm{C}(G)|} + \operatorname{rank} \mathrm{C}(G) - 1$$
for any base field k of characteristic $\neq p$.

(b) *Moreover, assume that k contains a primitive root of unity of degree $\exp(G)$. Then G has a faithful representation of degree*
$$\sqrt{|G/\mathrm{C}(G)|} + \operatorname{rank} \mathrm{C}(G) - 1\,.$$

Theorem 1.11 is an immediate consequence of this result, since part (b) implies $\operatorname{ed}_k G \leq \sqrt{|G/\mathrm{C}(G)|} + \operatorname{rank} \mathrm{C}(G) - 1$; cf. e.g., Theorem 3.9. Our proof of part (a) will rely on the following lemma.

Lemma 11.2. *Let G be a finite group and H be a central cyclic subgroup. Assume that there exists a character $\chi\colon G \to k^*$ whose restriction to H is faithful. Then*

(1) $\operatorname{ed}(G) \geq \operatorname{ed}(G/H)$.

(2) Moreover, if H has prime order and is not properly contained in another central cyclic subgroup of G then $\operatorname{ed}(G) = \operatorname{ed}(H) + 1$.

Proof. Part (2) is proved in [BR97, Theorem 5.3] in characteristic zero and in [Kan06, Theorem 4.5] in prime characteristic.

To prove (1), let $\phi\colon G \to G/H \hookrightarrow \mathrm{GL}(V)$ be a faithful representation of G/H. Then $\phi \oplus \chi\colon G \to \mathrm{GL}(V \times k)$ is a faithful representation of G. Denote the class of the G-action on $V \times k$ by $\alpha \in H^1(K,G)$, where $K = k(V \oplus k)^G$. Let β be the image of α in $H^1(K, G/H)$. Then β is given by the induced action of G/H on
$$(V \times k)/H \simeq V \times k\,.$$
Here the quotient map $V \times k \to V \times k$ is given by $(v,x) \to (v, x^d)$, where $d = |H|$. This shows that induced action of G/H on $(V \times k)/H$ is again linear. Hence, β is a versal G/H-torsor; cf. [GMS03, Example 5.4] or [BR97, Theorem 3.1]. We conclude that
$$\operatorname{ed}(G) = \operatorname{ed}(\alpha) \geq \operatorname{ed}(\beta) = \operatorname{ed}(G/H)\,,$$
as claimed. ♠

We now proceed with the proof of Theorem 11.1. We begin with the following reduction.

Lemma 11.3. *In the course of proving Theorem 11.1 we may assume without loss of generality that the center $\mathrm{C}(G)$ is cyclic.*

Proof. Let Z be a maximal cyclic subgroup of $\mathrm{C}(G)$ containing $[G,G]$. Then $\mathrm{C}(G) = Z \oplus W$ for some central subgroup W of G. Note that $\mathrm{rank}(W) = \mathrm{rank}\ \mathrm{C}(Z) - 1$. Moreover, W projects isomorphically into the abelian subgroup G/Z. In particular, if H is a cyclic subgroup of W then after composing this projection with a suitable character of G/Z, we obtain a character $\chi_H \colon G \to k^*$, which is faithful on H. This means that Lemma 11.2 can be used to compare the essential dimensions of G and G/H.

We will now argue by induction on $|W|$. If $|W| = 1$, we are done. For the induction step we will choose a cyclic subgroup $H \subset W$ (in a way, to be specified below) and assume that parts (a) and (b) of Theorem 11.1 hold for $\overline{G} = G/H$. Our goal will then be to prove that they also hold for G.

It is easy to see that g_1 and g_2 commute in G if and only if their images $\overline{g_1}$ and $\overline{g_2}$ commute in $\overline{G}$. In particular, $\mathrm{C}(\overline{G}) = \mathrm{C}(G)/H \simeq Z \oplus W/H$. (Here Z projects isomorphically to a cyclic subgroup $\overline{Z}$ of $\overline{G} = G/H$, and we are identifying Z with $\overline{Z}$). Consequently,

$$|\overline{G}/\mathrm{C}(\overline{G})| = |G/\mathrm{C}(G)| \text{ and } \mathrm{rank}\ \mathrm{C}(\overline{G}) = \mathrm{rank}(W/H) + 1. \tag{11.4}$$

(a) We choose H to be a subgroup of prime order in W. By induction assumption,

$$\mathrm{ed}(\overline{G}) \geq \sqrt{|\overline{G}/\mathrm{C}(\overline{G})|} + \mathrm{rank}(\mathrm{C}(\overline{G})) - 1\,.$$

We will now consider two cases.

Case 1. H is properly contained in another cyclic subgroup of W. In this case $\mathrm{rank}(W/H) = \mathrm{rank}(W)$ and by (11.4)

$$\mathrm{rank}\ \mathrm{C}(\overline{G}) = \mathrm{rank}(W/H) + 1 = \mathrm{rank}(W) + 1 = \mathrm{rank}\ C(G)\,.$$

By Lemma 11.2(1), $\mathrm{ed}(G) \geq \mathrm{ed}(\overline{G})$ and thus

$$\begin{aligned}\mathrm{ed}(G) &\geq \sqrt{|\overline{G}/\mathrm{C}(\overline{G})|} + \mathrm{rank}(\mathrm{C}(\overline{G})) - 1\\ &= \sqrt{|G|/|\mathrm{C}(G)|} + \mathrm{rank}\ \mathrm{C}(G) - 1\,,\end{aligned}$$

as desired.

Case 2. H is not properly contained in any cyclic subgroup of W. In this case $\mathrm{rank}(W/H) = \mathrm{rank}(W) - 1$ and by (11.4)

$$\mathrm{rank}\ \mathrm{C}(\overline{G}) = \mathrm{rank}(W/H) + 1 = \mathrm{rank}(W) = \mathrm{rank}\ \mathrm{C}(G) - 1\,.$$

By Lemma 11.2(2), $\mathrm{ed}(G) = \mathrm{ed}(\overline{G}) + 1$ and thus

$$\begin{aligned}\mathrm{ed}(G) &= \mathrm{ed}(\overline{G}) + 1\\ &= \sqrt{|\overline{G}/\mathrm{C}(\overline{G})|} + \mathrm{rank}(\mathrm{C}(\overline{G}))\\ &= \sqrt{|G|/|\mathrm{C}(G)|} + (\mathrm{rank}(\mathrm{C}(G)) - 1)\,.\end{aligned}$$

(b) Here we choose H to be a maximal central cyclic subgroup of W (not necessarily of prime order). By our induction assumption, G/H has a

representation $\rho\colon G/H \hookrightarrow \mathrm{GL}(V)$ of dimension

$$\sqrt{|\overline{G}/\mathrm{C}(\overline{G})|} + \mathrm{rank}(\mathrm{C}(\overline{G})) - 1\,.$$

Then $\rho \oplus \chi_H$ is a faithful representation of G of dimension

$$\sqrt{|\overline{G}/\mathrm{C}(\overline{G})|} + \mathrm{rank}\ \mathrm{C}(\overline{G}) = \sqrt{|G|/|\mathrm{C}(G)|} + \mathrm{rank}(\mathrm{C}(G)) - 1\,;$$

cf. (11.4). This shows that Theorem 11.1 holds for G. ♠

Proof of Theorem 11.1(a). By Lemma 11.3 we may assume that $\mathrm{C}(G)$ is cyclic. Set $Z = \mathrm{C}(G)$. Since we are assuming that $[G, G] \subset Z$, the quotient $A \stackrel{\text{def}}{=} G/Z$ is abelian.

In this case Theorem 1.11 reduces to $\mathrm{ed}(G) \geq \sqrt{|A|}$. By Theorem 9.2 it suffices to show that

$$\mathrm{ind}(G, Z) \geq \sqrt{|A|}\,.$$

so we will now direct our attention towards computing $\mathrm{ind}(G, Z)$.

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, $\ldots$, a_{2r} in A and positive integers d_1, $\ldots$, 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 $\overline{A}_i \simeq (\mathbb{Z}/d_i\mathbb{Z})^2$ 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, $\ldots$, u_{2r} be indeterminates, and set $K \stackrel{\text{def}}{=} k(u_1, \ldots, u_{2r})$. 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 Theorem 11.1(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 such that $m(\alpha_1, \dots, \alpha_r) = \alpha_1 \dots \alpha_r$. So, if $\xi \in \mathrm{H}^1(K, A)$ is the image of $(\xi_1, \dots, \xi_r)$, we have that $\partial\xi$ 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 $\operatorname{ind}(G, Z) \geq d_1 \dots d_r = \sqrt{|A/K|}$, as needed.

Now we prove 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

$$\langle x_1, x_2, y \mid x_1^d = x_2^d = y^d = 1,\ x_1 x_2 = y x_2 x_1,\ x_1 y = y x_1,\ x_2 y = y x_2 \rangle.$$

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

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

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 Theorem 11.1(a). ♠

Proof of Theorem 11.1(b). By Lemma 11.3 we may assume that $\mathrm{C}(G) = Z$ is cyclic. In this case Theorem 11.1 asserts that G has a faithful representation of degree $\sqrt{|A|}$.

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 $H_s \stackrel{\mathrm{def}}{=} s\langle g\rangle s^{-1} \cap H$ 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 we have

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

This concludes the of Theorem 1.11(b). ♠

12. Essential dimension of p-groups II

We begin with several simple illustrations of Theorem 1.11.

Example 12.1. 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 1.11 reduces to the following.

Let G be an extra-special p-group of order p^{2m+1}. Assume that 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 $\operatorname{ed} G = p^m$.

Example 12.2. 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

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

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

Proof. Here $\mathrm{C}(G)$ is the (unique) subgroup of C_{p^r} of order p^s. If $s \leq r/2$, this subgroup is central. Thus, if ζ_{p^r}, the equality $\mathrm{ed}_k(G) = p^s$ is an immediate consequence of Theorem 1.11. Since we are only assuming that $\zeta_p \in k$, Theorem 1.11 only tells us that $\mathrm{ed}_k(G) \geq p^s$. To prove the opposite inequality, we argue as follows. Let F be the prime subfield of k. By [Led02, Corollary to Proposition 2], $\mathrm{ed}_{F(\zeta_p)}(G) \leq p^s$. Since we are assuming that $F(\zeta_p) \subset k$, we conclude that $\mathrm{ed}_k(G) \leq p^s$ as well. ♠

Corollary 12.3. *Suppose k is a base field of characteristic $\neq p$. If G is a non-abelian finite p-group then* $\mathrm{ed}\, G \geq p$.

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

$$1 \longrightarrow Z \longrightarrow G \longrightarrow G/Z \longrightarrow 1 \tag{12.4}$$

give 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 9.2 tells us that $\mathrm{ed}\, G \geq p$, 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 2.22 implies that $p > \mathrm{ed}\, G \geq \mathrm{ed}(G/Z)$. 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 1.11 now tells us that

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

a contradiction. ♠

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

Question 12.5. Let G be a finite group and N be a normal subgroup. Is it true that $\mathrm{ed}\, G \geq \mathrm{ed}(G/N)$?

The inequality $\mathrm{ed}(G) \geq \mathrm{ed}(G/N)$ is known to hold in many cases (cf., e.g., Lemma 11.2). We will now show that it does not hold in general, even if H is assumed to be central.

Corollary 12.6. *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 $(c_1, \ldots, c_n)$ such that $c_1 \ldots c_n = 1$. Clearly

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

see Example 12.1.

On the other hand, G_n/H_n, is easily seen to be extra-special of order p^{2n+1}, so $\operatorname{ed}(G_n/H_n) = p^n$, again by Example 12.1. Hence by taking n sufficiently large we finish the proof. ♠

13. Spinor groups

In this section we will prove Theorem 1.13 stated in the introduction. For an introduction to the structure of this group and the theory of Clifford algebras and spin modules, we refer the reader to [FH91, §20.2] and [Che54].

Recall that by Spin_n we mean the totally split form of the spin group in dimension n over a field of characteristic not equal to 2. Let us be explicitly about this.

Write $\langle a_1, \ldots, a_n \rangle$ for the rank n-quadratic form q given by $q(x_1, \ldots, x_n) = \sum_{i=1}^n a_i x_i^2$. 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

$$h_n = \begin{cases} h_n^{\oplus n/2}, & \text{if } n \text{ is even,} \\ h_n^{\oplus (n-1/2)} \oplus \langle 1 \rangle, & \text{if } n \text{ is odd.} \end{cases} \tag{13.1}$$

Set $\mathrm{O}_n \stackrel{\text{def}}{=} \mathrm{O}(h_n), \mathrm{SO}_n \stackrel{\text{def}}{=} \mathrm{SO}(h_n)$, and $\mathrm{Spin}_n = \mathrm{Spin}(h_n)$. These are all totally split groups.

Now, one of the hypotheses of Theorem 1.13 is that $\zeta_4 \in k$. Therefore we can write Spin_n as $\mathrm{Spin}(q)$, where

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

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 $\mathrm{Spin}_n \to \mathrm{SO}_n$.

Lemma 13.2. *Every* Spin_n*-torsor over an extension K of k admits a reduction of structure group to G_n.*

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 has a reduction of structure group to Γ_n. ♠

This means that the natural morphism $\mathcal{B}G_n \to \mathcal{B}\mathrm{Spin}_n$ is isotropic; so from Propositions 3.3 and 2.22 we get the bounds

$$\operatorname{ed} G_n - \dim \mathrm{Spin}_n \leq \operatorname{ed} \mathrm{Spin}_n \leq \operatorname{ed} G_n. \tag{13.3}$$

Of course $\dim \mathrm{Spin}_n = n(n-1)/2$; we need to compute $\operatorname{ed} G_n$. The structure of G_n is well understood (besides the references cited above, it is also very clearly described in [Woo89]). The group scheme Spin_n is a subgroup scheme of the group scheme of units in the Clifford algebra A_n of the quadratic form $-(x_1^2 + \cdots + x_n^2)$. The algebra A_n is generated by elements $e_1, \ldots, e_n$, with relations $e_i^2 = -1$ and $e_i e_j + e_j e_i = 0$ 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 any $I \subseteq \{1, \ldots, n\}$ write $I = \{i_1, \ldots, i_r\}$ with $i_1 < i_2 < \cdots < i_r$ and set $e_I \stackrel{\text{def}}{=} e_{i_1} \ldots e_{i_r}$. The group G_n consists of the elements of A_n of the form $\pm e_I$, where $I \subseteq \{1, \ldots, n\}$ 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

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

Theorem 1.11 now tells us that

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

Substituting this into (13.3), we obtain the bounds of Theorem 1.13. ♠

Remark 13.4. The same argument, with G_n replaced by the inverse image of the diagonal subgroup of O_n in Pin_n, yields the following bounds on the essential dimensions of Pin groups (over a field k satisfying the assumptions of Theorem 1.13):

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

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

Remark 13.5. When $n \leq 14$ the lower bound of Theorem 1.13 is negative and the upper bound is much larger than the true value of $\operatorname{ed} \mathrm{Spin}_n$. For $n = 15$ and 16 our inequalities yield

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

and

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

When $n = 16$ our lower bound coincides with the lower bound (1.14) 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 13.6. By Proposition 2.12, the lower bounds in the theorem hold for 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 get the slightly weaker upper bound

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

for the totally split form of the spin group in dimension n. 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-compression $f\colon W \dashrightarrow X$, 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-compression of $V \times W$. Consequently,

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

as claimed.

Remark 13.7. It is natural to ask whether the inequality

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

can be proved by a direct application of Theorem 1.13 to the exact sequence

$$1 \longrightarrow \mu_2 \longrightarrow \mathrm{Spin}_n \longrightarrow \mathrm{SO}_n \longrightarrow 1 \tag{13.8}$$

without considering the finite subgroup G_n of Spin_n. The answer is "yes." Note, however, that we see no way of obtaining the slightly stronger lower-bound which we obtain when n is divisible by 4.

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 $\partial_K(q) \in H^2(K, \mu_2)$ 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 $c(q) \in H^2(K, \mu_2)$ is the class of the Clifford algebra of q.) If $n = 2m+1$ is odd, set $K = k(a_1, b_1, \ldots, a_m, b_m)$, where $a_1, b_1, \ldots, a_m, b_m$ are independent variables, and define q_n recursively by

$$q_3 = \langle a_1, b_1, a_1 b_1 \rangle \text{ and } q_{n+2} = \langle a_n b_n \rangle \otimes q_n \oplus \langle a_n, b_n \rangle.$$

A direct computation using basic properties of the Hasse-Witt invariant (see, e.g., [Lam73, Section V.3]) shows that $c(q_{2m+1})$ is the class of the product $(a_1, b_1)_2 \otimes_K \cdots \otimes_K (a_m, b_m)_2$ of quaternion algebras. This class has index 2^m, as claimed. ♠

Any lower bound on $\operatorname{ed} \operatorname{Spin}_n$ obtained in this way (i.e., directly from Theorem 9.2) will necessarily be of the form $\operatorname{ed} \operatorname{Spin}_n \geq 2^m - \frac{n(n-1)}{2}$ for some integer $m \geq 0$. Thus, while this approach recovers the lower bound of Theorem 1.13 if n is not divisible by 4, it cannot be used to do so if n is divisible by 4.

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 13.9. *Let G be a closed (but not necessarily connected) subgroup of GL_n defined over a field k. Assume that one of the following conditions holds.*

(1) k is algebraically closed of characteristic 0.
(2) k has characteristic 0 and G is connected.
(3) G is connected and reductive.

Then $\operatorname{ed} G \leq n$.

Proof. According to [CGR06, Theorem 1.1], there exists a finite subgroup scheme $S \subseteq G$ such that every G-torsor over $\operatorname{Spec} K$, where K is an extension of G, has a reduction of structure groups to S. Hence the morphism $\mathcal{B}S \to \mathcal{B}G$ is isotropic, so $\operatorname{ed} G \leq \operatorname{ed} S$. But the restriction of the representation $G \subseteq \mathrm{GL}_n$ is generically free, hence $\operatorname{ed} S \leq n$. ♠

Example 13.10. Suppose G satisfies one of the conditions (a), (b) or (c) of Lemma 13.9 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 13.9 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 13.11. *Suppose k is a field of characteristic $\neq 2$ and $\zeta_4 \in k$. Let n be a positive integer divisible by 4. Then*

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

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 12.1 $\operatorname{ed}(G_n/\langle\eta\rangle) = 2^{(n-2)/2}$ and thus

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

as in the proof of Theorem 1.13.

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 13.9 ♠

14. Essential dimension of cyclic p-groups

In this section, we are going to prove the following theorem due to M. Florence. In the sequel p will denote a prime, different from the characteristic of our base field k, and ζ_d will denote a primitive dth root of unity in $\overline{k}$. Recall that we have set $C_n \stackrel{\text{def}}{=} \mathbb{Z}/n\mathbb{Z}$.

Theorem 14.1 (M. Florence [Flo06])**.** *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*

$$\operatorname{ed} C_{p^m} = \begin{cases} p^{m-n} & \text{if } n < m, \\ 1 & \text{if } n \geq m. \end{cases}$$

for any integer $m \geq 1$,

This theorem was independently obtained by us in the case where $n \geq \lfloor (m+1)/2 \rfloor$; in particular, for $m = 2$. However, our proof of the stronger result given by Theorem 14.1, will rely on an idea of M. Florence, in combination with the lower bound of Theorem 9.2.

Proof. If $m \leq n$, then $C_{p^m} = \mu_{p^m}$. Therefore $\operatorname{ed} C_{p^m} = 1$ by [BF03, Example 2.3]. We can therefore restrict our attention to the case $n < m$.

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 $V \stackrel{\text{def}}{=} \operatorname{Ind}_{C_{p^n}}^{C_{p^m}} \chi$. A simple calculation shows that V is faithful, thus, V is generically free since C_{p^m} is finite. By Theorem 3.9, it follows that $\operatorname{ed} C_{p^m} \leq \dim V = p^{m-n}$.

To show that $\operatorname{ed} C_{p^m} \geq p^{m-n}$, 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 $\mathrm{ind}_K(t)$ is the index of of the CSA $\iota_*(t)$.

We claim that there is a t such that $\iota_*(t)$ is a division algebra. From this it will clearly follow that $\mathrm{ind}_K(t) = \dim V = p^{m-n}$.

In fact, this $t \in \mathrm{H}^1(K, C_{p^m})$ is simply the "generic" one. (This is the part of the argument that we learned from Florence's preprint [Flo06].) Namely, 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 permuting the variables in the obvious way ($k \cdot x_i = x_{i+k} \pmod{p^{m-n}}$). Let $F = L^{C_{p^{m-n}}}$. Then L/K defines a $C_{p^{m-n}}$-torsor t over K.

In the case where $k = \mathbb{Q}(\zeta_{p^n})$ the torsor $\iota_*(t)$ is the "generic" algebra R_{p^n,p^m,p^m} of [Row88, §7.3]. By a theorem of Brauer (see Theorem [Row88, Theorem 7.3.8]) it is a division algebra. A similar argument (due to M. Florence) shows that the same is true if $\mathbb{Q}(\zeta_{p^n})$ is replaced by our field k (satisfying the assumptions of Theorem 14.1). ♠

Remark 14.2. Suppose $\zeta_6 \in k$. Note that in this case $[k(\zeta_{36} : k]$ always divides 6. Assume $[k(\zeta_{36}) : k] = 6$. (This occurs, for example, if $k = \mathbb{Q}(\zeta_6)$.) We claim that $\mathrm{ed}_k\, C_{36} \geq 4$.

Indeed, by Remark 9.3 it suffices to show that $\mathrm{ind}(C_{36}, C_6) = 6$. Let K/k be a field extension and consider the boundary map

$$\partial_K\colon K^*/(K^*)^6 = \mathrm{H}^1(K, C_6) \longrightarrow \mathrm{H}^2(K, C_6)$$

induced by the exact sequence $1 \to C_6 \to C_{36} \to C_6 \to 1$. By [Vel00, Theorem 7.1], ∂_K sends (a) to the class of the cyclic algebra $(a, \zeta_6)_6$. The index of this cyclic algebra clearly divides 6. Taking $K = k(a)$, where a is an independent variable over k, and applying Wedderburn's criterion (cf. e.g., [Pie82, Corollary 15.1d]), we conclude that in this case the cyclic algebra $(a, \zeta_6)_6$ has index 6. Thus $\mathrm{ind}(C_{36}, C_6) = 6$, as claimed.

A similar argument shows that if Conjecture 7.4 is valid for $n = p_1^{a_1} \dots p_r^{a_r}$ then $\mathrm{ed}_k(C_n) \geq p_1^{a_1} + \dots + p_r^{a_r} - r + 1$.

Let D_n be the dihedral group of order $2n$. Ledet [Led02, Section 3] conjectured that if n is odd then $\mathrm{ed}\, C_n = \mathrm{ed}\, D_n$ over any field k of characteristic

zero. We will now prove this conjecture in the case where $n = p^r$ is a prime power and k contains a primitive pth root of unity.

Corollary 14.3. *Let p be an odd prime 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 s be the largest integer n such that $\zeta_{p^n} \in k$. By our assumption $1 \leq s \leq m-1$. By Theorem 14.1

$$\operatorname{ed}_k C_{p^m} = p^{m-n}\,.$$

Since $C_{p^m} \subset D_{p^m}$, we clearly have $\operatorname{ed}_k D_{p^m} \geq \operatorname{ed}_k C_{p^m}$. Thus we only need to show that

$$\operatorname{ed}_k D_{p^m} \leq p^{m-n}\,.$$

To prove this 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}. Since k contains ζ_{p^n}, D_{p^n} has essential dimension 1 over k. Thus, by [Led02, Section 3],

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

This completes the proof of Corollary 14.3. ♠

15. 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 15.1. *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) If n is even $\langle a_1, \dots, a_n\rangle = \sum_{i=1}^n (-1)^i \ll -a_i \gg$.

(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

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

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:

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

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 13.7, 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 15.2. *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 $(3, n)$-Pfister number of q_L is then an upper bound for the $(3, n)$-Pfister number of any form over any field extension K/k. $\spadesuit$

Remark 15.3. For $a > 3$ the finiteness of $\mathrm{Pf}_k(a, n)$ is an open problem.

The main theorem in this section is a lower bound for $\mathrm{Pf}_k(3, n)$ stated as Theorem 1.15 in the Introduction. We restate it here for the reader's convenience.

Theorem 15.4. *Let k be a field of characteristic different from 2 and let n be an even positive integer. Then*

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

For each extension K of k, denote by $\mathrm{T}_n(K)$ the image of $\mathrm{H}^1(K, \mathrm{Spin}_n)$ into $\mathrm{H}^1(K, \mathrm{SO}_n)$; we get a functor $\mathrm{T}_n\colon \mathrm{Fields}_k \to \mathrm{Sets}$. The essential dimension of this functor is closely related to the essential dimension of Spin_n.

Lemma 15.5. $\operatorname{ed}\mathrm{Spin}_n - 1 \leq \operatorname{ed}\mathrm{T}_n \leq \operatorname{ed}\mathrm{Spin}_n$.

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

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

The first inequality then follows from [BF03, Proposition 1.13] and the second follows from Proposition 2.22. ♠

Lemma 15.6. *Let q and q' be non-degenerate 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 identity $c(q \oplus q') = c(q) + c(q')$ [Lam73, V.3.15]. ♠

Let h_K be the standard 2-dimensional hyperbolic form $h_K(x,y) = xy$ over an extension K of k discussed at the beginning of §13. For each n-dimensional quadratic form q admitting a spin structure over K, denote $\operatorname{ed}_n(q)$ the essential dimension of the class of q in $\mathrm{T}_n(K)$.

Lemma 15.7. *Let q be an dimensional quadratic form over K admitting a spin structure, and let s be a positive integer. Then we have*

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

Proof. Set $m \stackrel{\text{def}}{=} \operatorname{ed}_{n+2s}(h_K^{\oplus s} \oplus q)$; let F be a field of definition of $h_K^{\oplus s} \oplus q$ of transcendence degree m, and let $\widetilde{q}$ be a quadratic form with a spin structure on F whose base change to K is 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, the base change of q' to K is isomorphic to q; hence $\operatorname{ed}_n(q) \leq m + s(s+2n-1)/2$, as claimed. ♠

Proof of Theorem 15.4. 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 $\mathrm{W}(K)$.

We will take an n-dimensional form q with a spin structure such that $\mathrm{ed}_n(q) = \mathrm{ed}\,\mathrm{T}_n$. Suppose that q is equivalent in the Witt ring to a form of the type $\sum_{1=1}^{r} \ll a_i, b_i, c_i \gg$.

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

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

Set

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

if r is even, and

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

if r is odd. Then q is equivalent to ϕ in the Witt ring, and ϕ has a spin structure. The dimension of ϕ is $7r$ or $7r+1$, according to the parity of r.

We claim that $n < 7r$. If not, then the dimension of q is at most equal to the dimension of ϕ, so q is isomorphic to a form of type $h_K^s \oplus \phi$. By Lemma 15.6 and Theorem 1.13 we get the inequalities

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

The resulting inequality fails for every even $n \geq 12$ because, for such n, $\mathrm{ed}\,\mathrm{Spin}_n \geq n/2$; see (1.14).

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 by Lemma 15.7 we have that this essential dimension is at least $\mathrm{ed}_n(q) - s(s + 2n - 1)/2$. From this, Lemma 15.5 and

Theorem 1.13 we have the chain of inequalities

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

Now assume that r is even. We substitute the value $s = (7r-n)/2$ in the resulting inequality, obtaining

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

We interpret this as a quadratic inequality in r. The constant term of the polynomial is negative for all $n \geq 8$; hence if r_0 is the positive root, the equality is equivalent to $r \geq r_0$. By the quadratic formula

$$\begin{aligned} r_0 &= \frac{\sqrt{49 \cdot 2^{(n+4)/2} + 168n - 367} - (7n+5)}{49} \\ &\geq \frac{2^{(n+4)/4} - n - 2}{7}. \end{aligned}$$

This completes the proof of Theorem 15.4 when r is even. The calculations when r is odd are analogous: using the substitution $s = (7r+1-n)/2$ we obtain the root

$$\begin{aligned} r_0 &= \frac{\sqrt{49 \cdot 2^{(n+4)/2} + 168n - 199} - (7n+12)}{49} \\ &\geq \frac{2^{(n+4)/4} - n - 2}{7}. \end{aligned}$$

♠

References

[Art74] M. Artin, *Versal deformations and algebraic stacks*, Invent. Math. **27** (1974), 165–189.

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

[BF04] ———, *Essential dimension of cubics*, J. Algebra **278** (2004), no. 1, 199–216.

[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.

[Bro] P. Brosnan, *The essential dimension of a g-dimensional complex abelian variety is $2g$*, to appear in Transformation Groups.

[CGR06] V. Chernousov, P. 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.

[Con] Brian Conrad, *Keel–Mori theorem via stacks*, `http://www.math.lsa.umich.edu/~bdconrad/papers/coarsespace.pdf`.

[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, LAGRS preprint server, `http://www.math.uni-bielefeld.de/lag/`, 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.

[Flo06] Mathieu Florence, *On the essential dimension of cyclic p-groups*, LAGRS preprint server, `http://www.math.uni-bielefeld.de/lag/`, 2006.

[Gar06] Skip Garibaldi, *Cohomological invariants: exceptional groups and spin groups*, `http://www.mathcs.emory.edu/~skip/lens-ci/lens-ci.html`, 2006.

[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.

[Gro63] Alexander Grothendieck, *Revêtements étales et groupe fondamental. Fasc. II: Exposés 6, 8 à 11*, Séminaire de Géométrie Algébrique, vol. 1960/61, Institut des Hautes Études Scientifiques, Paris, 1963.

[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.

[Kar00] Nikita A. Karpenko, *On anisotropy of orthogonal involutions*, J. Ramanujan Math. Soc. **15** (2000), no. 1, 1–22.

[KM97] Seán Keel and Shigefumi Mori, *Quotients by groupoids*, Ann. of Math. (2) **145** (1997), no. 1, 193–213.

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

[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.

[Knu71] Donald Knutson, *Algebraic spaces*, Springer-Verlag, Berlin, 1971, Lecture Notes in Mathematics, Vol. 203.

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

[Kre99] Andrew Kresch, *Cycle groups for Artin stacks*, Invent. Math. **138** (1999), no. 3, 495–536.

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

[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.

[Mer03] Alexander Merkurjev, *Steenrod operations and degree formulas*, J. Reine Angew. Math. **565** (2003), 13–26.

[Nis55] Hajime Nishimura, *Some remarks on rational points*, Mem. Coll. Sci. Univ. Kyoto. Ser. A. Math. **29** (1955), 189–192. MR MR0095851 (20 #2349)

[Pie82] Richard S. Pierce, *Associative algebras*, Graduate Texts in Mathematics, vol. 88, Springer-Verlag, New York, 1982, , Studies in the History of Modern Science, 9.

[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.

[Sil86] Joseph H. Silverman, *The arithmetic of elliptic curves*, Graduate Texts in Mathematics, vol. 106, Springer-Verlag, New York, 1986.

[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.

[Vis05] Angelo Vistoli, *Grothendieck topologies, fibered categories and descent theory*, Fundamental algebraic geometry, Math. Surveys Monogr., vol. 123, Amer. Math. Soc., Providence, RI, 2005, pp. 1–104.

[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