ESSENTIAL DIMENSION OF MODULI OF CURVES AND OTHER ALGEBRAIC STACKS

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

Abstract. In this paper we address questions of the following type. Let k be a base field and K/k be a field extension. Given a geometric object X over a field K (e.g. a smooth curve of genus g) what is the least transcendence degree of a field of definition of X over the base field k? In other words, how many independent parameters are needed to define X? To study these questions we introduce a notion of essential dimension for an algebraic stack. Using the resulting theory, we give a complete answer to the question above when the geometric objects X are smooth or stable curves.

Contents

2000 *Mathematics Subject Classification.* Primary 14A20, 14H10, 11E04.
[†]Supported in part by an NSERC discovery grant.
[‡]Supported in part by the PRIN Project "Geometria sulle varietà algebriche", financed by MIUR.

1. Introduction

This paper was motivated by the following question.

Question 1.1. Let k be a field and $g \geq 0$ be an integer. What is the smallest integer d such that for every field K/k, every smooth curve X of genus g defined over K descends to a subfield $k \subset K_0 \subset K$ with $\operatorname{tr\,deg}_k K_0 \leq d$?

Here by "X descends to K_0" we mean that the exists a curve X_0 over K_0 such that X is K-isomorphic to $X_0 \times_{\operatorname{Spec} K_0} \operatorname{Spec} K$.

In order to address this and related questions, we will introduce and study the notion of essential dimension for algebraic stacks; see §2. The essential dimension $\operatorname{ed} \mathcal{X}$ of a scheme $\mathcal{X}$ is simply the dimension of $\mathcal{X}$; on the other hand, the essential dimension of the classifying stack $\mathcal{B}_k G$ of an algebraic group G is the essential dimension of G in the usual sense; see [Rei00] or [BF03]. The notion of essential dimension of a stack is meant to bridge these two examples. The minimal integer d in Question 1.1 is the essential dimension of the moduli stack of smooth curves $\mathcal{M}_g$. We show that $\operatorname{ed} \mathcal{X}$ is finite for a broad class of algebraic stacks of finite type over a field; see Corollary 3.4. This class includes all Deligne–Mumford stacks and all quotient stacks of the form $\mathcal{X} = [X/G]$, where G a linear algebraic group.

Our main result is the following theorem.

Theorem 1.2. *Let $\mathcal{M}_{g,n}$ (respectively, $\overline{\mathcal{M}}_{g,n}$) be the stacks of n-pointed smooth (respectively, stable) algebraic curves of genus g over a field k of characteristic* 0. *Then*

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

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

In particular, the values of $\operatorname{ed} \mathcal{M}_{g,0} = \operatorname{ed} \mathcal{M}_g$ give a complete answer to Question 1.1.

Note that $3g-3+n$ is the dimension of the moduli space $\mathbf{M}_{g,n}$ in the stable range $2g-2+n > 0$ (and 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 proof of Theorem 1.2 for $(g,n) \neq (1,0)$ relies on two results of independent interest. One is the "Genericity Theorem" 6.1 which says that the essential dimension of a smooth integral Deligne–Mumford stack satisfying an appropriate separation hypothesis 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 1.2 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.

The second main ingredient in our proof of Theorem 1.2 is the following formula, which we use to compute the essential dimension of the generic gerbe.

Theorem 1.3. *Let $\mathcal{X}$ be a gerbe over a field K banded by a group G. Let $[\mathcal{X}] \in \mathrm{H}^2(K, G)$ be the Brauer class of $\mathcal{X}$.*
(a) *If $G = \mathbb{G}_m$ and $\operatorname{ind}[\mathcal{X}]$ is a prime power then* $\operatorname{ed} \mathcal{X} = \operatorname{ind}[\mathcal{X}] - 1$.
(b) *If $G = \boldsymbol{\mu}_{p^r}$, where p is a prime and $r \geq 1$, then* $\operatorname{ed} \mathcal{X} = \operatorname{ind}[\mathcal{X}]$.

Our proof of this theorem can be found in the preprint [BRV07, Section 7]. A similar argument was used by N. Karpenko and A. Merkurjev in the proof of [KM08, Theorem 3.1], which generalizes Theorem 1.3(b). For the sake of completeness, we include an alternative proof of Theorem 1.3 in §4.

Theorem 1.3 has a number of applications beyond Theorem 1.2. Some of these have already appeared in print. In particular, we used Theorem 1.3 to study the essential dimension of spinor groups in [BRV], N. Karpenko and A. Merkurjev [KM08] used it to study the essential dimension of finite p-groups, and A. Dhillon and N. Lemire [DL09] used it, in combination with the Genericity Theorem 6.1, to give an upper bound for the essential dimension of the moduli stack of SL_n-bundles over a projective curve. In this paper we will also use Theorem 1.3 (in combination with Theorems 6.1) to study the essential dimension of the stacks of hyperelliptic curves (Theorem 7.2) and principally polarized abelian varieties (Theorem 9.1).

In the case where $(g, n) = (1, 0)$ Theorem 1.2 requires a separate argument, which is carried out in §8. In this case Theorem 1.2 is a consequence of the fact that the group schemes of l^n-torsion points on a Tate curve has essential dimension l^n, where l is a prime.

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 J. Alper, K. Behrend, C.-L. Chai, D. Edidin, N. Fakhruddin, A. Merkurjev, B. Noohi, G. Pappas, M. Reid and B. Totaro for helpful conversations.

2. The essential dimension of a stack

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.

Definition 2.1. Let $a \in F(L)$, where L is an object of Fields_k. Suppose that $K \subset L$ is a field extension of k. We say that *a descends* to K or that K is a *field of definition* for a if a is in the image of the induced map $F(K) \to F(L)$.

The *essential dimension* $\operatorname{ed} a$ of $a \in F(L)$ is the minimum of the transcendence degrees $\operatorname{tr deg}_k K$ taken over all fields $k \subseteq K \subseteq L$ such that a descends to K.

The essential dimension $\operatorname{ed} F$ of the functor F is the supremum of $\operatorname{ed} a$ taken over all $a \in F(L)$ with L in Fields_k. We will write $\operatorname{ed} F = -\infty$ if F is the empty functor.

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

The following definition singles out a class of functors that is sufficiently broad to include most interesting examples, yet "geometric" enough to allow one to get a handle on their essential dimension.

Definition 2.2. Suppose $\mathcal{X}$ is an algebraic stack over k. The *essential dimension* $\operatorname{ed} \mathcal{X}$ of $\mathcal{X}$ is defined to be the essential dimension of the functor $F_{\mathcal{X}} \colon \mathrm{Fields}_k \to \mathrm{Sets}$ which sends a field L/k to the set of isomorphism classes of objects in $\mathcal{X}(L)$.[1]

As in Definition 2.1, we will write $\operatorname{ed}_k \mathcal{X}$ when we need to be specific about the dependence on the base field k. Similarly for $\operatorname{ed}_k \xi$, where ξ is an object of $F_{\mathcal{X}}$.

Example 2.3. Let G be an algebraic group defined over k and $\mathcal{X} = \mathcal{B}_k G$ be the classifying stack of G. Then $F_{\mathcal{X}}$ is the Galois cohomology functor sending K to the set $\mathrm{H}^1(K, G)$ of isomorphism classes of G-torsors over $\operatorname{Spec}(K)$, in the fppf topology. The essential dimension of this functor is a numerical invariant of G, which, roughly speaking, measures the complexity of G-torsors over fields. This number is usually denoted by $\operatorname{ed}_k G$ or (if k is fixed throughout) simply by $\operatorname{ed} G$; following this convention, we will often write $\operatorname{ed} G$ in place of $\operatorname{ed} \mathcal{B}_k G$. Essential dimension was originally introduced and has since been extensively studied in this context; see e.g., [BR97, Rei00, RY00, Kor00, Led02, JLY02, BF03, Lem04, CS06, Gar]. The more general Definition 2.1 is due to A. Merkurjev; see [BF03, Proposition 1.17].

Example 2.4. Let $\mathcal{X} = X$ be a scheme of finite type over a field k, and let $F_X \colon \mathrm{Fields}_k \to \mathrm{Sets}$ denote the functor given by $K \mapsto X(K)$. Then an easy argument due to Merkurjev shows that $\operatorname{ed} F_X = \dim X$; see [BF03, Proposition 1.17].

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

Example 2.5. Let $\mathcal{X} = \mathcal{M}_g$ be the stack of smooth algebraic curves of genus g. Then the functor $F_{\mathcal{X}}$ sends K to the set of isomorphism classes of n-pointed smooth algebraic curves of genus g over K. Question 1.1 asks about the essential dimension of this functor in the case where $n = 0$.

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

Example 2.6. 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 $[X/G]$. Recall that K-points of $[X/G]$ are by definition diagrams of the form

$$\begin{array}{ccc} T & \xrightarrow{\ \psi\ } & X \\ \big\downarrow{\scriptstyle \pi} & & \\ \operatorname{Spec}(K) & & \end{array} \tag{2.1}$$

where π is a G-torsor and ψ is a G-equivariant map. The functor $F_{[X/G]}$ associates to a field K/k the set of isomorphism classes of such diagrams.

In the case where G is a special group (recall that this means that every G-torsor over $\operatorname{Spec}(K)$ is split, for every field K/k) the essential dimension of $F_{[X/G]}$ has been previously studied in connection with the so-called "functor of orbits" $\mathbf{Orb}_{X,G}$ given by the formula

$$\mathbf{Orb}_{X,G}(K) \overset{\text{def}}{=} \text{set of } G(K)\text{-orbits in } X(K).$$

Indeed, if G is special, the functors $F_{[X/G]}$ and $\mathbf{Orb}_{X,G}$ are isomorphic; an isomorphism between them is given by sending an object (2.1) of $F_{[X/G]}$ to the $G(K)$-orbit of the point $\psi s\colon \operatorname{Spec}(K) \to X$, where $s\colon \operatorname{Spec}(K) \to T$ is a section of $\pi\colon T \to \operatorname{Spec}(K)$.

Of particular interest are the natural GL_n-actions on $\mathbb{A}^N$ = affine space of homogeneous polynomials of degree d in n variables and on $\mathbb{P}^{N-1}$ = projective space of degree d hypersurfaces in $\mathbb{P}^{n-1}$, where $N = \binom{n+d-1}{d}$ is the number of degree d monomials in n variables. For general n and d the essential dimension of the functor of orbits in these cases is not known. Partial results can be found [BF04] and [BR05, Sections 14-15]. Additional results in this setting will be featured in a forthcoming paper.

Remark 2.7. If the functor F in Definition 2.1 is limit-preserving, a condition satisfied in all cases of interest to us, then every element $a \in F(L)$ descends to a field $K \subset L$ that is finitely generated over k. Thus in this case $\operatorname{ed} a$ is finite. In particular, $\operatorname{ed}\xi$ is finite for every object $\xi \in \mathcal{X}(K)$ and every field extension K/k; the limit-preserving property in this case is proved in [LMB00, Proposition 4.18],

In §3 we will show that, in fact, $\operatorname{ed}\mathcal{X} < \infty$ for a broad class of algebraic stacks $\mathcal{X}$; cf. Corollary 3.4. On the other hand, there are interesting examples where $\operatorname{ed}\mathcal{X} < \infty$; see Theorem 1.2 or [BS08].

The following observation is a variant of [BF03, Proposition 1.5].

Proposition 2.8. *Let $\mathcal{X}$ be an algebraic stack over k, and let K be a field extension of k. Then $\operatorname{ed}_K \mathcal{X}_K \leq \operatorname{ed}_k \mathcal{X}$.*

Proof. If L/K is a field extension, then the natural morphism $\mathcal{X}_K(L) \to \mathcal{X}(L)$ is an equivalence. Suppose that M/k is a field of definition for an object ξ in $\mathcal{X}(L)$. Let N be a composite of M and K over k. Then N

is a field of definition for ξ, $\operatorname{tr deg}_K N \leq \operatorname{tr deg}_k M$, and the proposition follows. ♠

3. A fiber dimension theorem

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

Example 3.1. Let G be an algebraic group defined over k, and let $X \to Y$ be a G-equivariant morphism of k-algebraic spaces, locally of finite type and of relative dimension $\leq d$. Then the induced map of quotient stacks $[X/G] \to [Y/G]$ is representable, locally of finite type and of relative dimension $\leq d$.

The following result may be viewed as a partial generalization of the fiber dimension theorem (see [Har77, Exercise II.3.22 or Proposition III.9.5]) to the setting where schemes are replaced by stacks and dimension by essential dimension.

Theorem 3.2. *Let d be an integer $\mathcal{X} \to \mathcal{Y}$ be a representable k-morphism of algebraic stacks which is locally of finite type and of fiber dimension at most d. Let L/k be a field, $\xi \in X(L)$. Then*

(a) $\operatorname{ed}_k \xi \leq \operatorname{ed}_k f(\xi) + d$, *and*
(b) $\operatorname{ed}_k \mathcal{X} \leq \operatorname{ed}_k \mathcal{Y} + d$.

In particular, if $\operatorname{ed}_k \mathcal{Y}$ is finite, then so is $\operatorname{ed}_k \mathcal{X}$.

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

$$\begin{array}{ccc} \operatorname{Spec} L & \xrightarrow{\xi} & \mathcal{X} \\ \big\downarrow & & \big\downarrow{\scriptstyle f} \\ \operatorname{Spec} K & \xrightarrow{\eta} & \mathcal{Y} \end{array}$$

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

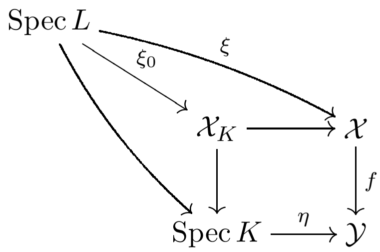

Moreover, ξ factors through $K(p)$, where p denotes the image of ξ_0 in $\mathcal{X}_K$. Since $\mathcal{X}_K$ has dimension at most d over K, we have $\operatorname{tr}\deg_K K(p) \leq d$. Therefore,

$$\operatorname{tr}\deg_k K(p) = \operatorname{tr}\deg_K K(p) + \operatorname{tr}\deg_k K \leq \operatorname{ed} f(\xi) + d$$

and part (a) follows.

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

Corollary 3.3. *Consider an action of an algebraic group G on an algebraic space X, defined over a field k. Assume X is locally of finite type over k. Then*

$$\operatorname{ed}_k G \geq \operatorname{ed}_k [X/G] - \dim X \, .$$

Proof. The natural G-equivariant map $X \to \operatorname{Spec} k$ gives rise to a map $[X/G] \to \mathcal{B}_k G$ of quotient stacks. This latter map is locally of finite type and of relative dimension $\leq \dim X$; see Example 3.1. Applying Theorem 3.2(b) to this map, we obtain the desired inequality. ♠

Corollary 3.4. (Finiteness of essential dimension) *Let $\mathcal{X}$ be an algebraic stack of finite type over k. Suppose that for any algebraically closed extension Ω of k and any object ξ of $\mathcal{X}(\Omega)$ the group scheme $\underline{\operatorname{Aut}}_\Omega(\xi) \to \operatorname{Spec} \Omega$ is affine. Then $\operatorname{ed}_k \mathcal{X} < \infty$.*

Note that Corollary 3.4 fails without the assumption that all the $\underline{\operatorname{Aut}}_\Omega(\xi)$ are affine. For example, by Theorem 1.2, $\operatorname{ed} \mathcal{M}_{1,0} = +\infty$.

Proof. We may assume without loss of generality that $\mathcal{X} = [X/G]$ is a quotient stack for some affine algebraic group G acting on an algebraic space X. Indeed, by a Theorem of Kresch [Kre99, Proposition 3.5.9] $\mathcal{X}$ is covered by quotient stacks $[X_i/G_i]$ of this form and hence, $\operatorname{ed} \mathcal{X} = \max_i \operatorname{ed} [X_i/G_i]$.

If $\mathcal{X} = [X/G]$ then by Corollary 3.3,

$$\operatorname{ed} [X/G] \leq \operatorname{ed}_k G + \dim(X) \, .$$

The desired conclusion now follows from the well-known fact that $\operatorname{ed}_k G < \infty$ for any affine algebraic group G; see [Rei00, Theorem 3.4] or [BF03, Proposition 4.11]. ♠

4. The essential dimension of a gerbe over a field

The goal of this section is to prove Theorem 1.3 stated in the Introduction. We proceed by briefly recalling some background material on gerbes from [Mil80, p. 144] and [Gir71, IV.3.1.1], and on canonical dimension from [KM06] and [BR05].

Gerbes. Let $\mathcal{X}$ be a gerbe defined over a field K *banded* by an abelian K-group scheme G. There is a notion of equivalence of gerbes banded by G; the set of equivalence classes is in a natural bijective correspondence with the group $\mathrm{H}^2(K, G)$. The identity element of $\mathrm{H}^2(K, G)$ corresponds to the class of the neutral gerbe $\mathcal{B}_K G$. 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 over K.

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

The *determination functor* $D_X \colon \mathrm{Fields}_K \to \mathrm{Sets}$ is defined as follows. For any field extension L/K, $D_X(L)$ is the empty set, if $X(L) = \emptyset$, and a set consisting of one element if $X(L) \neq \emptyset$. The natural map $D(L_1) \to D(L_2)$ is then uniquely determined for any $K \subset L_1 \subset L_2$. It is shown in [KM06] that if X be a complete regular K-variety then

$$\mathrm{cd}\, X = \mathrm{ed}\, D_X \,. \tag{4.1}$$

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

$$\mathrm{cd}\, X = p^a - 1\,; \tag{4.2}$$

see [KM06, Example 3.10] or [BR05, Theorem 11.4].

If $m = p_1^{a_1} \dots p_r^{a_r}$ is the prime decomposition of m then the class of X in $\mathrm{Br}\, L$ is the sum of classes $\alpha_1, \dots, \alpha_r$ whose indices are $p_1^{a_1}, \dots, p_r^{a_r}$. Denote by $X_1, \dots, X_r$ the Brauer–Severi varieties associated with $\alpha_1, \dots, \alpha_r$. It is easy to see that $K(X_1 \times \dots \times X_r)$ is a generic splitting field for X. Hence,

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

J.-L. Colliot-Thélène, N. Karpenko and A. Merkurjev [CTKM06] conjectured that equality holds, i.e.,

$$\mathrm{cd}\, X = p_1^{a_1} + \dots + p_r^{a_r} - r\,. \tag{4.3}$$

As we mentioned above, this in known to be true if m is a prime power (i.e., $r = 1$). Colliot-Thélène, Karpenko and Merkurjev also proved (4.3) for

$m = 6$; see [CTKM06, Theorem 1.3]. Their conjecture remains open for all other m.

Theorem 4.1. *Let d be an integer with $d > 1$. Let K be a field and $x \in \mathrm{H}^2(K, \boldsymbol{\mu}_d)$. Denote the image of x in $\mathrm{H}^2(K, \mathbb{G}_m)$ by y, the $\boldsymbol{\mu}_d$-gerbe associated with x by $\mathcal{X} \to \mathrm{Spec}(K)$, the $\mathbb{G}_m$-gerbe associated with y by $\mathcal{Y} \to \mathrm{Spec}(K)$, and the Brauer–Severi variety associated with y by P. Then*
(a) $\mathrm{ed}\,\mathcal{Y} = \mathrm{cd}\,P$ *and*
(b) $\mathrm{ed}\,\mathcal{X} = \mathrm{cd}\,P + 1$.
In particular, if the index of x is a prime power p^r then $\mathrm{ed}\,\mathcal{Y} = p^r - 1$ *and* $\mathrm{ed}\,\mathcal{X} = p^r$.

Proof. The last assertion follows from (a) and (b) by (4.2).

(a) The functor $F_{\mathcal{Y}}\colon \mathrm{Fields}_K \to \mathrm{Sets}$ sends a field L/K to the empty set, if $P(L) = \emptyset$, and to a set consisting of one point, if $P(L) \neq \emptyset$. In other words, F_Y is the determination functor D_P introduced above. The essential dimension of this functor is $\mathrm{cd}\,P$; see (4.1).

(b) First note that the natural map $\mathcal{X} \to \mathcal{Y}$ is of finite type and representable of relative dimension ≤ 1. By Theorem 3.2(b) we conclude that $\mathrm{ed}\,\mathcal{X} \leq \mathrm{ed}\,\mathcal{Y} + 1$. By part (a) it remains to prove the opposite inequality, $\mathrm{ed}\,\mathcal{X} \geq \mathrm{ed}\,\mathcal{Y} + 1$. We will do this by constructing an object α of $\mathcal{X}$ whose essential dimension is $\geq \mathrm{ed}\,\mathcal{Y} + 1$.

We will view $\mathcal{X}$ as a torsor for $\mathcal{B}_K\boldsymbol{\mu}_d$ in the following sense. There exist maps

$$\mathcal{X} \times \mathcal{B}_K\boldsymbol{\mu}_d \longrightarrow \mathcal{X}$$
$$\mathcal{X} \times \mathcal{X} \longrightarrow \mathcal{B}_K\boldsymbol{\mu}_d$$

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

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

$$\mathrm{ed}\,a = \mathrm{cd}\,P = \mathrm{ed}\,\mathcal{Y}, \tag{4.4}$$

where we view a as an object in $\mathcal{Y}$. Non-canonically lift $a\colon \mathrm{Spec}\,L \to \mathcal{Y}$ to a map $\mathrm{Spec}\,L \to \mathcal{X}$ (this can be done, because $\mathcal{X} \to \mathcal{Y}$ is a $\mathbb{G}_{\mathrm{m}}$-torsor). Let $\mathrm{Spec}\,L(t) \to \mathcal{B}_L\boldsymbol{\mu}_d$ denote the map classified by $(t) \in \mathrm{H}^1(L(t), \boldsymbol{\mu}_d) = L(t)^{\times}/L(t)^{\times d}$. Composing these two maps, we obtain an object

$$\alpha\colon\ \mathrm{Spec}\,L(t) \to \mathcal{X} \times \mathcal{B}_L\boldsymbol{\mu}_d \to \mathcal{X}.$$

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

$$\begin{array}{ccc} \operatorname{Spec} L(t) & \xrightarrow{\alpha} & \mathcal{X} \\ \downarrow & \overset{\beta}{\nearrow} & \\ \operatorname{Spec} M & & \end{array} \tag{4.5}$$

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

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

$$\begin{array}{ccc} \operatorname{Spec} L & \xrightarrow{a} & \mathcal{Y} \\ \downarrow & \nearrow & \\ \operatorname{Spec} M_0 & & . \end{array}$$

which shows that $\operatorname{ed} a \leq \operatorname{tr}\deg_K M_0 = \operatorname{tr}\deg_K M - 1 \leq \operatorname{ed}\mathcal{Y} - 1$, contradicting (4.4).

Case 2. Now suppose the restriction of ν to M is trivial. The map $\operatorname{Spec} L \to \mathcal{X}$ sets up an isomorphism $\mathcal{X}_L \cong \mathcal{B}_L\boldsymbol{\mu}_d$. The map $\operatorname{Spec} L(t) \to \mathcal{X}$ factors through $\mathcal{X}_L$ and thus induces a class in $\mathcal{B}_L\boldsymbol{\mu}_d(L(t)) = \mathrm{H}^1(L(t), \boldsymbol{\mu}_d)$. This class is (t). Tensoring the diagram (4.5) with L over K, we obtain

$$\begin{array}{ccc} \operatorname{Spec} L(t)\otimes L & \xrightarrow{\alpha} & \mathcal{X}_L \cong \mathcal{B}_L\boldsymbol{\mu}_d \\ \downarrow & \overset{\beta}{\nearrow} & \\ \operatorname{Spec} M\otimes L & & \end{array}$$

Recall that $L = K(P)$ is the function field of P. Since P is absolutely irreducible, the tensor products $L(t)\otimes L$ and $M\otimes L$ are fields. The map $\operatorname{Spec} M\otimes L \to \mathcal{B}_L\boldsymbol{\mu}_d$ is classified by some $m \in (M\otimes L)^\times/(M\otimes L)^{\times d} = \mathrm{H}^1(M\otimes L, \boldsymbol{\mu}_d)$. The image of m in $L(t)\otimes L$ would have to be equal to t modulo d-th powers. We will now derive a contradiction by comparing the valuations of m and t.

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

Corollary 4.2. *Let $1 \to Z \to G \to Q \to 1$ denote an extension of group schemes over a field k with Z central and isomorphic to (a) $\mathbb{G}_{\mathrm{m}}$ or (b) $\boldsymbol{\mu}_{p^r}$ for some prime p and some $r > 1$. Let $\operatorname{ind}(G, Z)$ as the maximal value of $\operatorname{ind}\bigl(\partial_K(t)\bigr)$ as K ranges over all field extensions of k and t ranges over all torsors in $\mathrm{H}^1(K, Q)$. If $\operatorname{ind}(G, Z)$ is a prime power (which is automatic in case (b)) then*

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

Proof. Choose $t \in H^1(K, Z)$ so that $\operatorname{ind}\bigl(\partial_K(t)\bigr)$ attains its maximal value, $\operatorname{ind}(G, Z)$. Let $X \to \operatorname{Spec}(K)$ be the Q-torsor representing t. Then G acts on X via the projection $G \to Q$, and $[X/G]$ is the Z-gerbe over $\operatorname{Spec}(K)$ corresponding to the class $\partial_K(t) \in \mathrm{H}^2(K, Z)$. By Theorem 1.3.

$$\operatorname{ed}[X/G] = \begin{cases} \operatorname{ind}\bigl(\partial_K(t)\bigr) - 1 \text{ in case (a),} \\ \operatorname{ind}\bigl(\partial_K(t)\bigr) \text{ in case (b).} \end{cases}$$

Since $\dim(X) = \dim(Q)$, applying Corollary 3.3 to the G-action on X, we obtain

$$\operatorname{ed}_K G_K \geq \begin{cases} (\operatorname{ind}(G, Z) - 1) - \dim Q = \operatorname{ind}(G, Z) - \dim(G) \text{ in case (a),} \\ \operatorname{ind}(G, Z) - \dim Q = \operatorname{ind}(G, Z) - \dim(G) \text{ in case (b).} \end{cases}$$

Since $\operatorname{ed}_k G \geq \operatorname{ed}_K G_K$ (see [BF03, Proposition 1.5] or our Proposition 2.8), the corollary follows. ♠

5. Gerbes over complete discrete valuation rings

In this section we prove two results on the structure of étale gerbes over complete discrete valuation rings that will be used in the proof of Theorem 6.1.

5.1. **Big and small étale sites.** Let S be a scheme. We let Sch/S denote the category of all schemes T equipped with a morphism to S. As in [SGA72], we equip Sch/S with the étale topology. Let $\text{ét}/S$ denote the full subcategory of Sch/S consisting of all schemes étale over S (also equipped with the étale topology). The site Sch/S is the *big étale site* and the category $\text{ét}/S$ is the *small étale site*. We let $S_{\text{Ét}}$ denote the category of sheaves on Sch/S and $S_{\text{ét}}$ the category of sheaves on $\text{ét}/S$. Since the obvious inclusion functor from the small to the big étale site is continuous, it induces a continuous morphism of sites $u\colon \text{ét}/S \to \mathrm{Sch}/S$ and, thus, a morphism $u\colon S_{\text{Ét}} \to S_{\text{ét}}$. Moreover, the adjunction morphism $F \to u_* u^* F$ is an isomorphism for F a sheaf in $S_{\text{ét}}$ [SGA72, VII.4.1]. We can therefore regard $S_{\text{ét}}$ as a full subcategory of $S_{\text{Ét}}$.

Definition 5.1. Let S be a scheme. An *étale gerbe* over S is a separated locally finitely presented Deligne–Mumford stack over S that is a gerbe in the étale topology.

Let $\mathcal{X} \to S$ be an étale gerbe over a scheme S. Then, by definition, there is an étale atlas, i.e., a morphism $U_0 \to \mathcal{X}$ where $U_0 \to S$ is surjective, étale and finitely presented over S. This atlas gives rise to a groupoid $\mathcal{G} \stackrel{\text{def}}{=} [U_1 \stackrel{\text{def}}{=} U_0 \times_{\mathcal{X}} U_0 \rightrightarrows U_0]$ in which each term is étale over S. Since $\mathcal{X}$ is the stackification of $\mathcal{G}$ which is a groupoid on the small étale site $S_{\text{ét}}$, it follows that $\mathcal{X} = u^*\mathcal{X}'$ for $\mathcal{X}'$ a gerbe on $S_{\text{ét}}$. In other words, we have the following proposition.

Proposition 5.2. *Let $\mathcal{X} \to S$ be an étale gerbe over a scheme S. Then there is a gerbe $\mathcal{X}'$ on $S_{\text{ét}}$ such that $\mathcal{X} = u^*\mathcal{X}'$.*

In fact, when S is the spectrum of a henselian discrete valuation ring, we can do better by using the following

Proposition 5.3. *Let $S = \operatorname{Spec} R$ with R a henselian discrete valuation ring. Let $f\colon T \to S$ be a surjective étale morphism. Then there is an open component T' of T such that $f_{|T'}\colon T' \to S$ is a finite étale morphism*

Proof. Let s denote the closed point of S. Since f is surjective, there exists a $t \in T$ such that $f(t) = s$. Since f is étale, f is quasi-finite at t (by [Gro67, 17.6.1]. Now, it follows from [Gro67, 18.5.11] that $T' \stackrel{\text{def}}{=} \operatorname{Spec} \mathcal{O}_{T,t}$ is an open component of T which is finite and étale. ♠

Now for a scheme S, let fét$/S$ denote the category of finite étale covers $T \to S$. We can consider fét$/S$ as a site in the obvious way. Then the inclusion morphism induces a continuous morphism of sites $v\colon \text{ét}/S \to \text{fét}/S$. If S is a henselian trait (i.e., the spectrum of a henselian discrete valuation ring) with closed point s, then the inclusion morphism $i\colon s \to S$ induces an equivalence of categories $i^*\colon \text{fét}/S \to \text{fét}/s$. Since the site fét$/s$ is equivalent to $s_{\text{ét}}$, this induces the *specialization morphism* $\operatorname{sp}\colon S_{\text{ét}} \to s_{\text{ét}}$, which is inverse to the inclusion morphism $i\colon s_{\text{ét}} \to S$; cf. [SGA73, p. 89]. Let $\tau = \operatorname{sp} \circ u\colon S_{\text{Ét}} \to s_{\text{ét}}$.

Corollary 5.4. *Let $\mathcal{X} \to S$ be an étale gerbe over a henselian trait S with closed point s. Then there is a gerbe $\mathcal{X}''$ over $s_{\text{ét}}$ such that $\mathcal{X} = \tau^*\mathcal{X}''$.*

Proof. Since $\mathcal{X} \to S$ is an étale gerbe, there is an étale atlas $X_0 \to S$ of $\mathcal{X}$. By Proposition 5.3 we may assume that X_0 such that X_0' is finite over S; and then $X_1 \stackrel{\text{def}}{=} X_0 \times_{\mathcal{X}} X_0$ is also finite, because $\mathcal{X}$ is separated, by hypothesis. Now the equivalence of categories $i^*\colon \text{fét}/S \to \text{fét}/s$ produces an gerbe $\mathcal{X}''$ over $s_{\text{ét}}$ such that $\mathcal{X} = \tau^*\mathcal{X}''$. ♠

5.2. **Group extensions and gerbes.** Let k be a field with separable closure $\overline{k}$ and absolute Galois group $G = \operatorname{Gal}(\overline{k}/k)$. Let

$$1 \longrightarrow F \xrightarrow{\ i\ } E \xrightarrow{\ p\ } G \longrightarrow 1 \tag{5.1}$$

be an extension of profinite groups with F finite and all maps continuous. From this data, we can construct a gerbe $\mathcal{X}_E$ over $(\operatorname{Spec} k)_{\text{ét}}$. To determine the gerbe it is enough to give its category of sections over $\operatorname{Spec} L$ where L/k

is a finite separable extension. Let $K = \{g \in G \mid g(\alpha) = \alpha, \alpha \in L\}$. Then the objects of the category $\mathcal{X}_E(L)$ are the solutions of the embedding problem given by (5.1). That is, an object of $\mathcal{X}_E(L)$ is a continuous homomorphism $\sigma\colon K \to E$ such that $p \circ \sigma(k) = k$ for $k \in K$. If $s_i\colon K \to E, i = 1, 2$ are two objects in $\mathcal{X}_E(L)$ then a morphism from s_1 to s_2 is an element $f \in F$ such that $fs_1f^{-1} = s_2$; cf. [DD99, p. 581].

By the results of Giraud [Gir71, Chapter VIII], it is easy to see that any gerbe $\mathcal{X} \to \operatorname{Spec} k$ with finite inertia arises from a sequence (5.1) as above. We explain how to get the extension: Given $\mathcal{X}$, we can find a separable Galois extension L/k and an object $\xi \in \mathcal{X}(L)$. This gives an extension of groups $\operatorname{Aut}_{\mathcal{X}}(\xi) \to \operatorname{Aut}_{\operatorname{Spec} k}(\operatorname{Spec} L) = \operatorname{Gal}(L/k)$. Pulling back this extension via the map $G = \operatorname{Gal}(k) \to \operatorname{Gal}(L/k)$ gives the desired extension (5.1).

Now, suppose that E is an extension of groups as in (5.1). Let L/k be a field extension, which is separable but not necessarily finite. Let $\overline{L}$ denote a fixed separable closure of L and let $\overline{k}$ denote the separable closure of k in $\overline{L}$. Then there is an obvious map $r\colon \operatorname{Gal}(\overline{L}/L) \to \operatorname{Gal}(\overline{k}/k)$. Let $u\colon (\operatorname{Spec} k)_{\text{Ét}} \to (\operatorname{Spec} k)_{\text{ét}}$ denote the functor of §5.1. Then $u^*\mathcal{X}(L)$ has the same description as in the case where L is a finite extension of k. In other words, we have the following proposition.

Proposition 5.5. *Let L/k be a separable extension and let $\mathcal{X}_E$ be the gerbe as above. Then the objects of the category $u^*\mathcal{X}_E(L)$ are the morphisms $s\colon \operatorname{Gal}(L) \to \operatorname{Gal}(k)$ making the following diagram commute.*

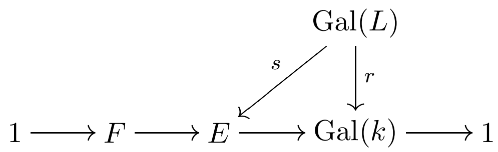

Moreover, if $s_i, i = 1, 2$ are two objects in $u^*\mathcal{X}_E(L)$, then the morphisms from s_1 to s_2 are the elements $f \in F$ such that $fs_1f^{-1} = s_2$.

5.3. **Splitting the inertia sequence.** We begin by recalling some results and notation from Serre's chapter in [GMS03].

Let A be a discrete valuation ring. Write $S = S_A$ for $\operatorname{Spec} A$, $s = s_A$ for the closed point in S and $\eta = \eta_A$ for the generic point. When A is the only discrete valuation ring under consideration, we suppress the subscripts. If A is henselian, then the choice of a separable closure $k(\overline{\eta})$ of $k(\eta)$ induces a separable closure of $k(s)$ and a map $\operatorname{Gal}(k(\overline{\eta})) \to \operatorname{Gal}(k(\overline{s}))$ between the absolute Galois groups. The kernel of this map is called the *inertia*, written as $I = I_A$. If $\operatorname{char} k(s) = p > 0$, then we set $I^{\mathrm{w}} = I^{\mathrm{w}}_A$ equal to the unique p-Sylow subgroup of I; otherwise we set $I^{\mathrm{w}} = \{1\}$. The group I^{w} is called the *wild inertia*. The group $I_{\mathrm{t}} = I_{A,t} \overset{\mathrm{def}}{=} I/I^{\mathrm{w}}$ is called the *tame inertia* and the group $\operatorname{Gal}(k(\eta))_{\mathrm{t}} \overset{\mathrm{def}}{=} \operatorname{Gal}(k(\eta))/I^{\mathrm{w}}$ is called the *tame Galois group*. We therefore have sequences

$$1 \longrightarrow I \longrightarrow \mathrm{Gal}(k(\eta)) \longrightarrow \mathrm{Gal}(k(s)) \longrightarrow 1 \quad \text{and} \tag{5.2}$$

$$1 \longrightarrow I_{\mathrm{t}} \longrightarrow \mathrm{Gal}(k(\eta))_{\mathrm{t}} \longrightarrow \mathrm{Gal}(k(s)) \longrightarrow 1\,, \tag{5.3}$$

called the *inertia exact sequence* and the *tame inertia exact sequence*, respectively.

For each prime l, set $\mathbb{Z}_l(1) = \varprojlim \boldsymbol{\mu}_{l^n}$ so that

$$\prod_{l \neq p} \mathbb{Z}_l(1) = \varprojlim_{p \nmid n} \boldsymbol{\mu}_n.$$

Then there is a canonical isomorphism $c\colon I_{\mathrm{t}} \to \prod_{l\neq p} \mathbb{Z}_l(1)$ [GMS03, p. 17]. To explain this isomorphism, let $g \in I_{\mathrm{t}}$ and let $\pi^{1/n}$ be an n-th root of a uniformizing parameter $\pi \in A$ with n not divisible by p. Then the image of $c(g)$ in $\boldsymbol{\mu}_n$ is $g(\pi^{1/n})/\pi^{1/n}$.

The following result extends Lemma 7.6 in [GMS03] from the case of complete to henselian discrete valuation rings.

Proposition 5.6. *Let A be a henselian discrete valuation ring. Then the sequence* (5.2) *is split.*

Proof. Because we need the ideas from the proof, we will repeat Serre's argument. Set $K = k(\eta)$ and $\overline{K} = k(\overline{\eta})$. Set $K_{\mathrm{t}} = \overline{K}^{I_{\mathrm{t}}}$: the maximal tamely ramified extension of K. Let π be a uniformizing parameter in A. Then, for each non-negative integer n not divisible by p, choose an n-th root π_n of π in $\overline{K}$ such that $\pi_{nm}^m = \pi_n$. Set $K_{\mathrm{ram}} \stackrel{\mathrm{def}}{=} K[\pi_n]_{(p\nmid n)}$. Then K_{ram} is totally and tamely ramified over K. Moreover any $K_{\mathrm{t}} = K_{\mathrm{ram}}K_{\mathrm{unr}}$. It follows that $\mathrm{Gal}(k(s))$ map be identified with the subgroup of elements $g \in \mathrm{Gal}(K)_{\mathrm{t}}$ fixing each of the π_n; cf. [Del80]. This splits the sequence (5.3).

Now, in [GMS03], Serre extends this splitting non-canonically to a splitting of (5.2) as follows. Since $k(s)$ has characteristic p, the p-cohomological dimension of $\mathrm{Gal}(k(s))$ is ≤ 1; see [Ser02]. Consequently, any homomorphism $\mathrm{Gal}(k(s)) \to \mathrm{Gal}(K)_{\mathrm{t}}$ can be lifted to $\mathrm{Gal}(K)$. ♠

While the splitting of (5.3) is not canonical, we need to know that it is possible to split two such sequences, associated to henselian discrete valuation rings $A \subseteq B$, in a compatible way.

Proposition 5.7. *Let $A \subseteq B$ be an extension of henselian discrete valuation rings, such that a uniformizing parameter for A is also a uniformizing parameter for B. Then it is possible to find maps $\sigma_B\colon \mathrm{Gal}(k(s_B)) \to \mathrm{Gal}(k(\eta_B))_{\mathrm{t}}$ (resp. $\sigma_A\colon \mathrm{Gal}(k(s_A)) \to \mathrm{Gal}(k(\eta_A))_{\mathrm{t}}$) splitting the tame inertia exact sequence* (5.3) *for B (resp. A) and such that the diagram*

$$\begin{array}{ccc} \mathrm{Gal}(k(s_B)) & \xrightarrow{\sigma_B} & \mathrm{Gal}(k(\eta_B))_{\mathrm{t}} \\ \downarrow & & \downarrow \\ \mathrm{Gal}(k(s_A)) & \xrightarrow{\sigma_A} & \mathrm{Gal}(k(\eta_A))_{\mathrm{t}}, \end{array}$$

with vertical morphisms given by restriction, commutes.

Proof. Let $\pi \in A$ be a uniformizing parameter for A, and hence for B. For each n not divisible by $p = \mathrm{char}(k(s_A))$, choose an n-th root π_n of π in $k(\overline{\eta}_B)$. Now, set $\sigma_B(k(s_B)) = \{g \in \mathrm{Gal}(k(\eta_B))_{\mathrm{t}} \colon g(\pi_n) = \pi_n \text{ for all } n\}$ and similarly for A. By the proof of Proposition 5.6, this defines splitting of the tame inertia sequences. Moreover, these splittings lift to splittings of the inertia exact sequence. ♠

Remark 5.8. By the proof of Proposition 5.6, the splittings σ_B and σ_A in Proposition 5.7 can be lifted to maps $\tilde{\sigma}_B \colon \mathrm{Gal}(k(s_B)) \to \mathrm{Gal}(k(\eta_B))$ (resp. $\tilde{\sigma}_A \colon \mathrm{Gal}(k(s_A)) \to \mathrm{Gal}(k(\eta_a)))$. However, since these liftings are non-canonical it is not clear that $\tilde{\sigma}_B$ and $\tilde{\sigma}_A$ can be chosen compatibly.

5.4. **Tame gerbes and splittings.** The following result is certainly well known; for the sake of completeness we supply a short proof.

Proposition 5.9. *Let $\mathcal{X} \to S$ be an étale gerbe over a henselian trait, with closed point s. Let $i \colon s \to S$ denote the inclusion of the closed point and* $\mathrm{sp} \colon S \to s$ *denote the specialization map. Then restriction map*

$$i^* \colon \mathcal{X}(S) \longrightarrow \mathcal{X}(s)$$

induces an equivalence of categories with quasi-inverse given by

$$\mathrm{sp}^* \colon \mathcal{X}(s) \longrightarrow \mathcal{X}(S)\,.$$

Proof. Since the composite $s \to S \xrightarrow{\mathrm{sp}} s$ is an auto-equivalence and $\mathcal{X}$ is obtained by pullback from $\mathcal{X}_s$, it suffices to show that the functor $i^* \colon \mathcal{X}(S) \to \mathcal{X}(s)$ is faithful. For this, suppose $\xi_i \colon S \to \mathcal{X}, i = 1, 2$ are two objects of $\mathcal{X}(S)$. Then the sheaf $\mathrm{Hom}(\xi_1, \xi_2)$ is étale over S. Since S is henselian, it follows that the sections of $\mathrm{Hom}(\xi_1, \xi_2)$ over S are isomorphic (via restriction) to the sections over s. Thus $i^* \colon \mathcal{X}(S) \to \mathcal{X}(s)$ is fully faithful. ♠

A Deligne-Mumford stack $\mathcal{X} \to S$ is *tame* if, for every geometric point $\xi \colon \mathrm{Spec}\,\Omega \to \mathcal{X}$, the order of the automorphism group $\mathrm{Aut}_{\mathrm{Spec}\,\Omega}(\xi)$ has order prime to the characteristic of Ω. For tame gerbes over a henselian discrete valuation ring, we have the following analogue of the splitting in Proposition 5.7.

Theorem 5.10. *Let $A \subseteq B$ be an extension of henselian traits, such that a uniformizing parameter for A is sent to a uniformizing parameter for B, and let $\mathcal{X}$ be a tame étale gerbe over A. Write $j_B \colon \{\eta_B\} \to B$ (resp. $j_A \colon \{\eta_A\} \to A$) for the inclusion of the generic points. Then it is possible to find functors $\sigma_B \colon \mathcal{X}(k(\eta_B)) \to \mathcal{X}(B)$ and $\sigma_A \colon \mathcal{X}(k(\eta_A)) \to \mathcal{X}(A)$ such that the diagram*

$$\begin{array}{ccccc}
\mathcal{X}(A) & \xrightarrow{j_A^*} & \mathcal{X}(k(\eta_A)) & \xrightarrow{\sigma_A} & \mathcal{X}(A) \\
\downarrow{\scriptstyle h^*} & & \downarrow{\scriptstyle h^*} & & \downarrow{\scriptstyle h^*} \\
\mathcal{X}(B) & \xrightarrow{j_B^*} & \mathcal{X}(k(\eta_B)) & \xrightarrow{\sigma_B} & \mathcal{X}(B)\,,
\end{array}$$

where $h\colon \operatorname{Spec} B \to \operatorname{Spec} A$ *is the morphism induced by the inclusion* $A \subseteq B$*, commutes (up to natural isomorphism) and the horizontal composites are isomorphic to the identity.*

Proof. Since $\mathcal{X}$ is an étale gerbe, there is an extension E as in (5.1) with $G = \operatorname{Gal}(k(s_A))$ such that $\mathcal{X}$ is the pull-back of $\mathcal{X}_E$ to the big étale site over S_A. Since $\mathcal{X}$ is tame, the band, i.e., the group F in (5.1), has order prime to $\operatorname{char} k(s_A)$.

Now, pick splittings σ_B and σ_A compatibly, as in Proposition 5.7.

We define a functor $\mathcal{X}(k(\eta_A)) \to \mathcal{X}(k(\eta_B))$ as follows. Using Proposition 5.5 we can identify $\mathcal{X}(k(\eta_B))$ with category of sections $s\colon \operatorname{Gal}(k\eta_B)) \to E$. Given such a section s, the tameness of E implies that $s(I^{\mathrm{w}}) = 1$. Therefore, s induces a map $\operatorname{Gal}(k(\eta_B)_{\mathrm{t}} \to \operatorname{Gal}(k(s_B)$ (which we also denote by the symbol s). Let $\sigma_B(s)$ denote the section $s \circ \sigma_B\colon \operatorname{Gal}(k(s_B)) \to \mathcal{X}_B$. This defines σ_B on the objects in $\mathcal{X}(k(\eta_B))$. If we define σ_A in the same way, it is clear that the diagram above commutes on objects. We define σ_B (resp. σ_A) on morphisms, by setting $\sigma_B(f) = f$ (and similarly for A). We leave the rest of the verification to the reader. ♠

5.5. **Genericity.**

Theorem 5.11. *Let R be a discrete valuation ring, $S = \operatorname{Spec}(R)$ and*

$$\mathcal{X} \longrightarrow S$$

a tame étale gerbe. Then $\operatorname{ed}_{k(s)} \mathcal{X}_s \le \operatorname{ed}_{k(\eta)} \mathcal{X}_\eta$, where s is the closed point of S and η is the generic point.

Proof. First note that we can assume that R is complete. If not, replace R with its completion at s. The field $k(s)$ does not change, but $k(\eta)$ is replaced by a field extension. By Proposition 2.8, the essential dimension of $\mathcal{X}_{k(\eta)}$ does not increase.

If R is equicharacteristic, then by Cohen's structure theorem, $R = k[\![t]\!]$ with $k = k(s)$. If not, denote by $W\big(k(s)\big)$ the unique complete discrete valuation ring with residue field $k(s)$ and uniformizing parameter p. This is called a Cohen ring of $k(s)$ in [Gro64, 19.8]. If $k(s)$ is perfect then $W\big(k(s)\big)$ is the ring of Witt vectors of $k(s)$, but this is not true in general, and $W\big(k(s)\big)$ is only determined up to a non-canonical isomorphism. By [Gro64, Théorème 19.8.6], there is a homomorphism $W\big(k(s)\big) \to R$ inducing the identity on $k(s)$; since $\mathcal{X}$ is pulled back from $k(s)$ via the specialization map, we can replace R with $W\big(k(s)\big)$.

Now suppose $b\colon \operatorname{Spec} L \to \mathcal{X}_s$ is a morphism from the spectrum of a field with $\operatorname{ed}_{k(s)} b = \operatorname{tr\,deg}_{k(s)} L = \operatorname{ed} \mathcal{X}_s$. (There is such a morphism because the $\operatorname{ed} \mathcal{X}_s$ is finite.) Set $B = L[\![t]\!]$ in the case that R is equicharacteristic and equal to $W(L)$ otherwise. In either case, B is a complete discrete valuation ring with residue field L. In the first case we have a canonical embedding $R = k[\![t]\!] \subseteq L[\![t]\!] = B$; in the second case, again by [Gro64, Théorème 19.8.6] (due to Cohen), we have a lifting $R = W\big(k(s)\big) \to W(L) = B$ of the embedding $k(s) \subseteq L$, which is easily seen to be injective. Therefore there is a unique morphism $\beta = b \circ \mathrm{sp}\colon S_B \to \mathcal{X}$ whose specialization to the closed point of B coincides with ξ.

Suppose there is a subfield M of $k(\eta_B)$ containing $k(\eta_R)$ such that the following hold

(1) The restriction $j_B^*\beta$ of β to $k(\eta_B)$ factors through M;
(2) $\operatorname{tr\,deg}_{k(\eta_R)} M < \operatorname{ed}_{k(s)} b$.

Complete M with respect to the discrete valuation induced from $k(\eta_B)$ and call the resulting complete discrete valuation ring A. It follows that there is a class α in $\mathcal{X}(k(\eta_A))$ whose restriction to $k(\eta_B)$ coincides with $j_B^*\beta$. But then, by Theorem 5.10, we have $\beta = h_*\sigma_A(\alpha)$. This implies that $b\colon \operatorname{Spec} L \to \mathcal{X}_s$ factors through the special fiber of A. Since the transcendence degree of $k(s_A)$ over $k(s)$ is less than $\operatorname{ed}_{k(s)} b$, this is a contradiction. ♠

Corollary 5.12. *Let R be an equicharacteristic complete discrete local ring and $\mathcal{X} \to \operatorname{Spec}(R)$ be a tame étale gerbe. Then*

$$\operatorname{ed}_{k(s)} \mathcal{X}_s = \operatorname{ed}_{k(\eta)} \mathcal{X}_\eta \,,$$

where s denotes the closed point of $\operatorname{Spec}(R)$ and η denotes the generic point.

Proof. Set $k = k(s)$. Since R is equicharacteristic, we have $R = k[\![t]\!]$ and $\mathcal{X}_{k(\eta)}$ is the pullback to $k(\eta)$ of $\mathcal{X}_{k(s)}$ via the inclusion of k in $k(\!(t)\!)$. Therefore $\operatorname{ed}_{k(s)} \mathcal{X}_{k(s)} \geq \operatorname{ed}_{k(\eta)} \mathcal{X}_{k(\eta)}$. The opposite inequality is given by Theorem 5.11. ♠

Theorem 5.13. *Suppose that $\mathcal{X}$ is an étale gerbe over a smooth scheme $\mathbf{X}$ locally of finite type over a field k of characteristic 0. Let K be an extension of k, $\xi \in \mathcal{X}(\operatorname{Spec} K)$. Then*

$$\operatorname{ed} \xi \leq \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})} + \dim \mathbf{X} - \operatorname{codim}_{\mathcal{X}} \xi.$$

Proof. We proceed by induction on $\operatorname{codim} \xi$. If $\operatorname{codim}_{\mathcal{X}} \xi = 0$, then the morphism $\xi\colon \operatorname{Spec} K \to \mathcal{X}$ is dominant. Hence ξ factors through $\mathcal{X}_{k(\mathbf{X})}$, and the result is obvious.

Assume $\operatorname{codim}_{\mathcal{X}} \xi > 0$. Call $\mathbf{Y}$ the closure of the image of $\operatorname{Spec} K$ in $\mathbf{X}$; then $\mathbf{Y}$ is generically smooth over $\operatorname{Spec} k$, because k is perfect. By restricting to a neighborhood of the generic point of $\mathbf{Y}$, we may assume that $\mathbf{Y}$ is contained in a smooth hypersurface $\mathbf{X}'$ of $\mathbf{X}$. Denote by $\mathcal{Y}$ and $\mathcal{X}'$ the inverse images in $\mathcal{X}$ of $\mathbf{Y}$ and $\mathcal{X}'$ respectively. Set $R = \mathcal{O}_{\mathbf{X},\mathbf{Y}}$ and call $\mathcal{X}$ the pullback of $\mathcal{X}$ to R. Then we can apply Theorem 5.11 to the gerbe $\mathcal{X}_R \to \operatorname{Spec} R$ and conclude that

$$\operatorname{ed}_{k(\mathbf{X}')} \mathcal{X}'_{k(\mathbf{X}')} \leq \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})}.$$

Using the inductive hypothesis we have

$$\begin{aligned}
\operatorname{ed} \xi &\leq \operatorname{ed}_{k(\mathbf{X}')} \mathcal{X}_{k(\mathbf{X}')} + \dim \mathbf{X}' - \operatorname{codim}_{\mathcal{X}'} \xi \\
&\leq \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})} + \dim \mathbf{X} - 1 - \operatorname{codim}_{\mathcal{X}'} \xi \\
&\leq \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})} + \dim \mathbf{X} - \operatorname{codim}_{\mathcal{X}} \xi.
\end{aligned}$$

♠

6. A genericity theorem for a smooth Deligne–Mumford stack

It is easy to see that Theorem 5.13 fails if $\mathcal{X}$ is not assumed to be a gerbe. In this section we will use Theorem 5.13 to prove the following weaker result for a wider class of Deligne–Mumford stacks.

Recall that a Deligne–Mumford stack $\mathcal{X}$ over a field k is *tame* if the order of the automorphism group of any object of $\mathcal{X}$ over an algebraically closed field is prime to the characteristic of k.

Theorem 6.1. *Let $\mathcal{X}$ be a smooth integral tame Deligne–Mumford stack locally of finite type over a perfect field k. Then*

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

Here the dimension of $\mathcal{X}$ is the dimension of the moduli space of any non-empty open substack of $\mathcal{X}$ with finite inertia.

Before proceeding with the proof, we record two immediate corollaries.

Corollary 6.2. *If $\mathcal{X}$ is as above and $\mathcal{U}$ is an open dense substack, then* $\operatorname{ed}_k\mathcal{M} = \operatorname{ed}_k\mathcal{U}$. ♠

Corollary 6.3. *If the conditions of the Theorem 6.1 are satisfied, and the generic object of $\mathcal{X}$ has no non-trivial automorphisms (i.e., $\mathcal{X}$ is an orbifold, in the topologists' terminology), then* $\operatorname{ed}_k\mathcal{X} = \dim\mathcal{X}$.

Proof. Here the generic gerbe $\mathcal{X}_K$ is a scheme, so $\operatorname{ed}_K\mathcal{X}_K = \dim\mathcal{X}$. ♠

Proof of Theorem 6.1. The inequality $\operatorname{ed}\mathcal{X} \geq \operatorname{ed}_{k(\mathbf{X})}\mathcal{X}_{k(\mathbf{X})} + \dim\mathcal{X}$ is obvious: so we only need to prove the inequality

$$\operatorname{ed}\xi \leq \operatorname{ed}_{k(\mathbf{X})}\mathcal{X}_{k(\mathbf{X})} + \dim\mathcal{X} \tag{6.1}$$

for any field extension L of k and any object ξ of $\mathcal{X}(L)$.

First of all, let us reduce the general result to the case that $\mathcal{X}$ has finite inertia. The reduction is immediate from the following lemma, that is essentially due to Keel and Mori.

Lemma 6.4 (Keel–Mori). *There exists an integral Deligne–Mumford stack with finite inertia $\mathcal{X}'$, together with an étale representable morphism of finite type $\mathcal{X}' \to \mathcal{X}$, and a factorization* $\operatorname{Spec} L \to \mathcal{X}' \to \mathcal{X}$ *of the morphism* $\operatorname{Spec} L \to \mathcal{X}$ *corresponding to ξ.*

Proof. We follow an argument due to Conrad. By [Con, Lemma 2.2] there exist

(i) an étale representable morphism $\mathcal{W} \to \mathcal{X}$ such that every morphism $\operatorname{Spec} L \to \mathcal{X}$, where L is a field, lifts to $\operatorname{Spec} L \to \mathcal{W}$, and
(ii) a finite flat representable map $Z \to \mathcal{W}$, where Z is a scheme.

Condition (ii) implies that $\mathcal{W}$ is a quotient of Z by a finite flat equivalence relation $Z \times_{\mathcal{W}} Z \rightrightarrows Z$, which in particular tells us that $\mathcal{W}$ has finite inertia. We can now take $\mathcal{X}'$ to be a connected component of $\mathcal{W}$ containing a lifting $\operatorname{Spec} L \to \mathcal{W}$ of $\operatorname{Spec} L \to \mathcal{X}$. ♠

Suppose that we have proved the inequality (6.1) whenever $\mathcal{X}$ has finite inertia. If denote by ξ' the object of $\mathcal{X}'$ corresponding to a lifting $\operatorname{Spec} L \to \mathcal{X}'$, we have

$$\operatorname{ed}\xi \leq \operatorname{ed}\xi' \leq \operatorname{ed}_{k(\mathbf{X}')} \mathcal{X}'_{k(\mathbf{X}')}.$$

On the other hand, the morphism $\mathcal{X}'_{k(\mathbf{X}')} \to \mathcal{X}_{k(\mathbf{X})}$ induced by the étale representable morphism $\mathcal{X}' \to \mathcal{X}$ is representable with fibers of dimension 0, hence

$$\operatorname{ed}_{k(\mathbf{X}')} \mathcal{X}'_{k(\mathbf{X}')} = \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}'_{k(\mathbf{X}')} \leq \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})}$$

by Theorem 3.2 (the first equality follows immediately from the fact that the extension $k(\mathbf{X}) \subseteq k(\mathbf{X}')$ is finite).

So, in order to prove the inequality (6.1) we may assume that $\mathcal{X}$ has finite inertia. Denote by $\mathbf{Y} \subseteq \mathbf{X}$ the closure of the image of the composite $\operatorname{Spec} L \to \mathcal{X} \to \mathbf{X}$, where $\operatorname{Spec} L \to \mathcal{X}$ corresponds to ξ, and call $\mathcal{Y}$ the reduced inverse image of $\mathbf{Y}$ in $\mathcal{X}$. Since k is perfect, $\mathcal{Y}$ is generically smooth; by restricting to a neighborhood of the generic point of $\mathbf{Y}$ we may assume that $\mathcal{Y}$ is smooth.

Denote by $\mathcal{N} \to \mathcal{Y}$ the normal bundle of $\mathcal{Y}$ in $\mathcal{X}$. Consider the deformation to the normal bundle $\phi\colon \mathcal{M} \to \mathbb{P}^1_k$ for the embedding $\mathcal{Y} \subseteq \mathcal{X}$. This is a smooth morphism such that $\phi^{-1}\mathbb{A}^1_k = \mathcal{X} \times_{\operatorname{Spec} k} \mathbb{A}^1_k$ and $\phi^{-1}(\infty) = \mathcal{N}$, obtained as an open substack of the blow-up of $\mathcal{X} \times_{\operatorname{Spec} k} \mathbb{P}^1_k$ along $\mathcal{Y} \times \{\infty\}$ (the well-known construction, explained for example in [Ful98, Chapter 5], generalizes immediately to algebraic stacks). Denote by $\mathcal{M}^0$ the open substack whose geometric points are the geometric points of $\mathcal{M}$ with stabilizer of minimal order (this is well defined because $\mathcal{M}$ has finite inertia).

We claim that $\mathcal{M}^0 \cap \mathcal{N} \neq \emptyset$. This would be evident if $\mathcal{X}$ where a quotient stack $[V/G]$, where G is a finite group of order not divisible by the characteristic of k, acting linearly on a vector space V, and $\mathcal{Y}$ where the of the form $[X/G]$, where W is a G-invariant linear subspace of V. However, étale locally on $\mathbf{X}$ every tame Deligne–Mumford stack if a quotient $[X/G]$, where G is a a finite group of order not divisible by the characteristic of k (see, e.g., [AV02, Lemma 2.2.3]). Since G is tame and X is smooth, it is well-known that étale-locally on $\mathbf{X}$, the stack $\mathcal{X}$ has the desired form, and this is enough to prove the claim.

Set $\mathcal{N}^0 \stackrel{\text{def}}{=} \mathcal{M}^0 \cap \mathcal{N}$. The object ξ corresponds to a dominant morphism $\operatorname{Spec} L \to \mathcal{Y}$. The pullback $\mathcal{N} \times_{\mathcal{Y}} \operatorname{Spec} K$ is a vector bundle V over $\operatorname{Spec} L$, and the inverse image $\mathcal{N}^0 \times_{\mathcal{Y}} \operatorname{Spec} L$ of $\mathcal{N}^0$ is not empty. We may assume that L is infinite; otherwise $\operatorname{ed}\xi = 0$ and there is nothing to prove. Assuming that L is infinite, $\mathcal{N}^0 \times_{\mathcal{Y}} \operatorname{Spec} L$ has an L-rational point, so there is a lifting $\operatorname{Spec} L \to \mathcal{N}^0$ of $\operatorname{Spec} L \to \mathcal{Y}$, corresponding to an object η of $\mathcal{N}^0(\operatorname{Spec} L)$. Clearly the essential dimension of ξ as an object of $\mathcal{X}$ is the same as its essential dimension as an object of $\mathcal{Y}$, and $\operatorname{ed}\xi \leq \operatorname{ed}\eta$. Let us apply Theorem 5.13 to the gerbe $\mathcal{M}^0$. The function field of the moduli space $\mathbf{M}$ of $\mathcal{M}$ is $k(\mathbf{X})(t)$, and its generic gerbe is $\mathcal{X}_{k(\mathbf{X})(t)}$; by Proposition 2.8, we have $\operatorname{ed}_{k(\mathbf{X})(t)} \mathcal{X}_{k(\mathbf{X})(t)} \leq \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})}$. The composite $\operatorname{Spec} L \to \mathcal{N}^0 \subseteq \mathcal{M}^0$

has codimension at least 1, hence we obtain

$$\begin{aligned} \operatorname{ed} \xi &< \operatorname{ed}_{k(\mathbf{X})(t)} \mathcal{X}_{k(\mathbf{X})(t)} + \dim \mathbf{M} \\ &\leq \operatorname{ed}_{k(\mathbf{X})} \mathcal{X}_{k(\mathbf{X})} + \dim \mathbf{X} + 1. \end{aligned}$$

This concludes the proof. ♠

The following examples show that Theorem 6.1 fails for more general algebraic stacks, including singular Deligne–Mumford stacks.

Examples 6.5.

(a) Let k be any field. Suppose $G \stackrel{\text{def}}{=} \mathbb{G}_{\mathrm{a}} \times \mathbb{G}_{\mathrm{a}}$ acts on $\mathbb{A}^3$ by the formula $(s,t)(x,y,z) = (x+sz, y+tz, z)$, and define $\mathcal{X} \stackrel{\text{def}}{=} [\mathbb{A}^3/G]$. 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 $[H/G] \simeq \mathbb{A}^2 \times \mathcal{B}_k G$; 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 $G \stackrel{\text{def}}{=} \boldsymbol{\mu}_r^n$ act on X via the formula

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

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 $[\{0\}/G] \simeq \mathcal{B}_k \boldsymbol{\mu}_r^n$, which has essential dimension n.

(c) The following example shows that Corollary 6.2 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

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

On the other hand, let Y be the locus of matrices of rank $n-1$, which is 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 a morphism $[Y/G] \to \mathbb{P}^{n-1}$. If L is an extension of $\mathbb{C}$, every L-valued point of $\mathbb{P}^{n-1}$ lifts to an L-valued point of Y: hence 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 6.6. Under what hypotheses the genericity theorem holds? Let $\mathcal{X} \to \operatorname{Spec} k$ be an integral algebraic stack. Using the results of [LMB00, Chapter 11], one can define the generic gerbe $\mathcal{X}_K \to \operatorname{Spec} K$ of $\mathcal{X}$, which is an fppf gerbe over a field of finite transcendence degree over k. What conditions on $\mathcal{X}$ ensure the equality

$$\operatorname{ed}_k \mathcal{X} = \operatorname{ed}_K \mathcal{X}_K + \operatorname{tr}\deg_k K\,?$$

Smoothness seems necessary; as we have seen above, there are counterexamples even for Deligne–Mumford stacks with very mild singularities. We think that the best result that one hope for is the following. Suppose that $\mathcal{X}$ is smooth with quasi-affine diagonal, and let $\xi \in \mathcal{X}(\operatorname{Spec} L)$ be a point. Assume that the automorphism group scheme of ξ over L is linearly reductive. Then $\operatorname{ed}\xi \leq \operatorname{ed}_K \mathcal{X}_K + \operatorname{tr}\deg_k K$. In particular, if all the automorphism groups are linearly reductive, then $\operatorname{ed}\mathcal{X} = \operatorname{ed}_K \mathcal{X}_K + \operatorname{tr}\deg_k K$.

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

Recall that the base field k is assumed to be of characteristic 0.

The assertion that $\operatorname{ed}\overline{\mathcal{M}}_{g,n} = \operatorname{ed}\mathcal{M}_{g,n}$ whenever $2g-2+n>0$ is an immediate consequence of Corollary 6.2. Moreover, if $g \geq 3$, or $g=2$ and $n \geq 1$, or $g=1$ and $n \geq 2$, then

$$\operatorname{ed}_k \mathcal{M}_{g,n} = \operatorname{ed}_k \overline{\mathcal{M}}_{g,n} = 3g-3+n.$$

Indeed, in all these cases the automorphism group of a generic object of $\mathcal{M}_{g,n}$ is trivial, so the generic gerbe is trivial, and $\operatorname{ed}\mathcal{M}_{g,n} = \dim\mathcal{M}_{g,n}$ by Corollary 6.3.

The remaining cases of Theorem 1.2, with the exception of $(g,n)=(1,0)$, are covered by the following proposition. The case where $(g,n)=(1,0)$ requires a separate argument which will be carried out in the next section.

Proposition 7.1. *(a)* $\operatorname{ed}\mathcal{M}_{0,1} = 2$*,*

(b) $\operatorname{ed}\mathcal{M}_{0,1} = \operatorname{ed}\mathcal{M}_{0,2} = 0$*,*

(c) $\operatorname{ed}\mathcal{M}_{1,1} = 2$*,*

(d) $\operatorname{ed}\mathcal{M}_{2,0} = 5$*.*

Proof. (a) Since $\mathcal{M}_{0,0} \simeq \mathcal{B}_k\mathrm{PGL}_2$, we have $\operatorname{ed}\mathcal{M}_{0,0} = \operatorname{ed}\mathrm{PGL}_2 = 2$, where the last inequality is proved in [Rei00, Lemma 9.4 (c)] (the argument there is valid for any field k of characteristic $\neq 2$).

Alternative proof of (a): 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.

(b) 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.

Alternative proof of (b): $\mathcal{M}_{0,2} = \mathcal{B}_k\mathbb{G}_{\mathrm{m}}$ and $\mathcal{M}_{0,1} = \mathcal{B}_k(\mathbb{G}_{\mathrm{m}} \ltimes \mathbb{G}_{\mathrm{a}})$, 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).

(c) 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 $\boldsymbol{\mu}_2$ and is neutral by [Sil86, Proposition 1.4 (c)], and so $\operatorname{ed}_k \mathcal{X} = \operatorname{ed}_{k(j)} \mathcal{X} + 1 = \operatorname{ed} \mathcal{B}_{k(j)}\boldsymbol{\mu}_2 + 1 = 2$.

(d) is a special case of Theorem 7.2 below, since $\mathcal{H}_2 = \mathcal{M}_{2,0}$. ♠

Theorem 7.2. *Let $\mathcal{H}_g$ denote the stack of hyperelliptic curves of genus $g > 1$ over a field k of characteristic* 0*. Then*

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

Proof. 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 $(\mathcal{H}_g)_K \overset{\text{def}}{=} \operatorname{Spec} K \times_{\mathbf{H}_g} \mathcal{H}_g$ the generic gerbe of $\mathcal{H}_g$. From Theorem 6.1 we have

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

so we need to show that $\operatorname{ed}_K (\mathcal{H}_g)_K$ 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.

Let $\mathcal{D}_g$ be the stack over k whose objects 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 that 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 $2g + 2 \geq 5$). Denote by $(P \to \operatorname{Spec} K, \Delta)$ the object of $\mathcal{D}_g(\operatorname{Spec} K)$ corresponding the generic point $\operatorname{Spec} K \to \mathbf{H}_g$. It is well-known that $P(K) = \emptyset$; we give a proof for lack of a suitable reference.

Let C be a conic without rational points defined over some extension L of k. Let V be the L-vector space $\mathrm{H}^0(C, \omega_{C/L}^{-(g+1)})$; denote the function field of V by $F = L(V)$. Then there is a tautological section σ of $\mathrm{H}^0(C_F, \omega_{C_F/F}^{-(g+1)}) = \mathrm{H}^0(C, \omega_{C/L}^{-(g+1)}) \otimes_L F$. Note that $C_F(F) = \emptyset$, because the extension $L \subseteq F$ is purely transcendental. The zero scheme of σ is a divisor on C_F that is étale over $\operatorname{Spec} F$, and defines a morphism $C_F \to \mathcal{D}_g$. This morphism is clearly dominant: so $K \subseteq F$, and $C_F = P \times_{\operatorname{Spec} L} \operatorname{Spec} F$. Since $C_F(F) = \emptyset$ we have $P(K) = \emptyset$, as claimed.

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 $\boldsymbol{\mu}_2$. 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 $\mathcal{B}_K\boldsymbol{\mu}_2$, whose essential dimension over K 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 $\mathrm{H}^2(K, \boldsymbol{\mu}_2)$ represents the image in $\mathrm{H}^2(K, \boldsymbol{\mu}_2)$ of the class $[P]$ in $\mathrm{H}^1(K, \mathrm{PGL}_2)$ under the non-abelian boundary map $\mathrm{H}^1(K, \mathrm{PGL}_2) \to \mathrm{H}^2(K, \boldsymbol{\mu}_2)$. According to Theorem 1.3 its essential dimension is the index of $[P]$, which equals 2. ♠

The results above apply to more than stable curves. Assume that we are in the stable range $2g - 2 + n > 0$. Denote by $\mathfrak{M}_{g,n}$ the stack of all reduced n-pointed local complete intersection curves of genus g. This is the algebraic stack over $\operatorname{Spec} k$ whose objects over a k-scheme T are finitely presented proper flat morphisms $C \to T$, where C is an algebraic space, whose geometric fibers are connected reduced local complete intersection curves of genus g, together with n sections $T \to C$ whose images are contained in the smooth locus of $C \to T$. We do not require the sections to be disjoint.

The stack $\mathfrak{M}_{g,n}$ contains $\mathcal{M}_{g,n}$ as an open substack. By standard results in deformation theory, every reduced local complete intersection curve is unobstructed, and is a limit of smooth curves. Furthermore there is no obstruction to extending the sections, since these map into the smooth locus. Therefore $\mathfrak{M}_{g,n}$ is smooth and connected, and $\mathcal{M}_{g,n}$ is dense in $\mathfrak{M}_{g,n}$. However, the stack $\mathfrak{M}_{g,n}$ is very large (it is certainly not of finite type), and in fact it is very easy to see that its essential dimension is infinite. However, consider the open substack $\mathfrak{M}_{g,n}^{\mathrm{fin}}$ consisting of objects whose automorphism group is finite. Then $\mathfrak{M}_{g,n}^{\mathrm{fin}}$ is a Deligne–Mumford stack, and Theorem 6.1 applies to it. Thus we get the following strengthened form of Theorem 1.2 (under the assumption that $2g - 2 + n > 0$).

Theorem 7.3. *If $2g - 2 + n > 0$ and the characteristic of k is 0, then*

$$\operatorname{ed} \mathfrak{M}_{g,n}^{\mathrm{fin}} = \begin{cases} 2 & \text{if } (g,n) = (1,1), \\ 5 & \text{if } (g,n) = (2,0), \\ 3g - 3 + n & \text{otherwise.} \end{cases}$$

It is not hard to show that $\mathfrak{M}_{g,n}^{\mathrm{fin}}$ does not have finite inertia.

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

In this section we will finish the proof of Theorem 1.2 by showing that $\operatorname{ed} \mathcal{M}_{1,0} = +\infty$.

We remark that the moduli stack $\mathcal{M}_{1,0}$ of genus 1 curves should not be confused with the moduli stack $\mathcal{M}_{1,1}$ of elliptic curves. The objects of

$\mathcal{M}_{1,0}$ are torsors for elliptic curves, where as the objects of $\mathcal{M}_{1,1}$ are elliptic curves themselves. The stack $\mathcal{M}_{1,1}$ is Deligne–Mumford and, as we saw in the last section, its essential dimension is 2. The stack $\mathcal{M}_{1,0}$ is not Deligne–Mumford, and we will now show that its essential dimension is ∞.

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

$$0 \longrightarrow \boldsymbol{\mu}_n \longrightarrow E[n] \longrightarrow \mathbb{Z}/n \longrightarrow 0.$$

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

Lemma 8.1. *Let $E = E_q$ be a Tate curve as above 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 $\boldsymbol{\mu}_{l^n} = \mathbb{Z}/l^n$. Let $L = K(t)$ and consider the class $(t) \in \mathrm{H}^1(L, \boldsymbol{\mu}_{l^n}) = L^*/(L^*)^n$.

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 $\operatorname{ind}(E[l^n], \boldsymbol{\mu}_{l^n}) \geq l^n$. Then, since $\dim \mathbb{Z}/l^n = 0$, Corollary 4.2 implies that $\operatorname{ed} \mathcal{B}_k E[l^n] \geq l^n$.

To see that $\operatorname{ed} \mathcal{B}_k E[l^n] \leq l^n$, note that $E[l^n]$ admits an l^n-dimensional generically free representation $V = \operatorname{Ind}_{\boldsymbol{\mu}_{l^n}}^{E[l^n]} \chi$, where $\chi\colon \boldsymbol{\mu}_{l^n} \to \mathbb{G}_m$ is the tautological character. Thus,

$$\operatorname{ed} \mathcal{B}_k E[l^n] \dim(V) \leq l^n\,,$$

as desired; see [BR97, Theorem 3.1] or [BF03, Proposition 4.11]. ♠

Theorem 8.2. *Let $E = E_q/K$ denote the Tate curve over a field K as above. Then* $\operatorname{ed} E = +\infty$.

Proof. For each prime power l^n, the morphism $\mathcal{B}_k E[l^n] \to \mathcal{B}_k E$ is representable of fiber dimension 1. By Theorem 3.2

$$\operatorname{ed} E \geq \operatorname{ed} \mathcal{B}_k E[l^n] = l^n - 1$$

for all n. ♠

Remark 8.3. It is shown in [BS08] that if A is an abelian variety over k and k is a number field then $\operatorname{ed}_k A = +\infty$. On the other hand, if $k = \mathbb{C}$ is the field of complex numbers then $\operatorname{ed}_{\mathbb{C}}(A) = 2\dim(A)$; see [Bro].

Now we can complete the proof of Theorem 1.2.

Theorem 8.4. *Let k be a field. Then* $\operatorname{ed}_k \mathcal{M}_{1,0} = +\infty$.

Proof. Set $F = k((t))$. It suffices to show that $\operatorname{ed}_F \mathcal{M}_{1,0} \otimes_k F$ is infinite. Consider the morphism $\mathcal{M}_{1,0} \to \mathcal{M}_{1,1}$ which sends a genus 1 curve to its Jacobian. 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:

$$\begin{array}{ccc} \mathcal{B}_k E & \longrightarrow & \mathcal{M}_{1,0} \otimes_k F \\ \downarrow & & \downarrow \\ \operatorname{Spec} F & \longrightarrow & \mathcal{M}_{1,1} \otimes_k F. \end{array}$$

It follows that the morphism $\mathcal{B}_k E \to \mathcal{M}_{1,0}$ is representable, with fibers of dimension ≤ 0. Applying Theorem 3.2 once again, we see that $+\infty = \operatorname{ed} \mathcal{B}_F E \leq \operatorname{ed} \mathcal{M}_{1,0} \otimes_k F$. ♠

9. On the essential dimension of the stack of principally polarized abelian varieties

Let g be a positive integer, and $\mathcal{A}_g$ be the stack of principally polarized abelian varieties of dimension g over a fixed field k of characteristic 0.This stack is well-known to be Deligne–Mumford, smooth and connected of dimension $g(g+1)/2$.

The proof of the following result was suggested to us by Miles Reid for part (a), and Najmuddin Fakhruddin for part (b).

Theorem 9.1.
(a) *If g is odd, then* $\operatorname{ed} \mathcal{A}_g = \frac{g(g+1)}{2} + 1$.
(b) *If g is even, write $g = 2^s h$, with h is odd. Then*

$$\frac{g(g+1)}{2} + 2 \leq \operatorname{ed} \mathcal{A}_g \leq \frac{g(g+1)}{2} + 2^{s+3}.$$

Moreover, $\operatorname{ed} \mathcal{A}_2 = 5$.

Proof. If $\mathbf{A}_g$ is the moduli space of $\mathcal{A}_g$, by the genericity theorem we only need to know the essential dimension of the generic gerbe $(\mathcal{A}_g)_{k(\mathbf{A}_g)}$. Since the only non-trivial automorphism of a generic abelian variety is the antipodal map $x \mapsto -x$, the gerbe $(\mathcal{A}_g)_{k(\mathbf{A}_g)}$ is banded by $\boldsymbol{\mu}_2$, so

$$\operatorname{ed} \mathcal{A}_g = \frac{g(g+1)}{2} + i_g\,,$$

where i_g is the index of its class of this gerbe in the Brauer group of $k(\mathbf{A}_g)$.

(a) Suppose that g is odd. If $g = 1$ then $\mathcal{A}_1 = \mathcal{M}_{1,1}$, and the statement follows from Theorem 1.2.

Assume $g > 1$. We have to show that the gerbe $(\mathcal{A}_g)_{k(\mathbf{A}_g)} \to \operatorname{Spec} k(\mathbf{A}_g)$ admits a section, for then $(\mathcal{A}_g)_{k(\mathbf{A}_g)} \simeq \mathcal{B}_{k(\mathbf{A}_g)}\boldsymbol{\mu}_2$, and $i_g = 1$.

Each abelian scheme $X \to S$ has a canonical action of $\boldsymbol{\mu}_2$ by the antipodal map; this embeds $\boldsymbol{\mu}_2$ into the center of the automorphism group scheme of $X \to S$. Let us denote by $\mathcal{B}_g$ the *rigidification* of $\mathcal{A}_g$ along $\boldsymbol{\mu}_2$, as described in [ACV03, 5.1]. There is a tautological morphism $\mathcal{A}_g \to \mathcal{B}_g$. If Ω is an algebraically closed extension of k, this morphism induces an isomorphism between the isomorphism classes in $\mathcal{B}_g(\operatorname{Spec}\Omega)$ and the isomorphism classes in $\mathcal{A}_g(\operatorname{Spec}\Omega)$. If $X \to \operatorname{Spec}\Omega$ is in $\mathcal{A}_g(\operatorname{Spec}\Omega)$, the automorphism group of its image in $\mathcal{B}_g(\operatorname{Spec}\Omega)$ is the quotient $\operatorname{Aut}_\Omega(X)/\boldsymbol{\mu}_2$. The stack $\mathcal{B}_g$ is a smooth integral Deligne–Mumford stack with moduli space $\mathbf{A}_g$.

A family $X \to S$ of principally polarized abelian varieties induces a family $X/\boldsymbol{\mu}_2 \to S$ of generalized singular Kummer varieties. Consider the tautological family of principally polarized abelian varieties $\mathcal{X}_g \to \mathcal{A}_g$. There is a stack $\mathcal{X}_g/\boldsymbol{\mu}_2 \to \mathcal{A}_g$ whose objects are pairs $(X \to S, S \to X/\boldsymbol{\mu}_2)$, where $X \to S$ is an object of $\mathcal{A}_g(S)$ and $S \to X/\boldsymbol{\mu}_2$ is a section of the generalized singular Kummer variety $X/\boldsymbol{\mu}_2 \to S$. Now, if $(X \to S, S \to X/\boldsymbol{\mu}_2)$ is an object of $\mathcal{X}_g/\boldsymbol{\mu}_2$, we can let $\boldsymbol{\mu}_2$ act on it by the antipodal map on $X \to S$, while we leave the section $S \to X/\boldsymbol{\mu}_2$ fixed; this defines an embedding of $\boldsymbol{\mu}_2$ into the center of the automorphism group of $(X \to S, S \to X/\boldsymbol{\mu}_2)$. Let $\mathcal{W}_g$ be the rigidification of $\mathcal{X}_g/\boldsymbol{\mu}_2$ along $\boldsymbol{\mu}_2$; there is a canonical representable morphism $\mathcal{W}_g \to \mathcal{B}_g$, whose geometric fibers are generalized Kummer varieties; the zero section $\mathcal{A}_g \to \mathcal{X}_g$ induces a section $\mathcal{B}_g \to \mathcal{W}_g$. The diagram

$$\begin{array}{ccc} \mathcal{X}_g/\boldsymbol{\mu}_2 & \longrightarrow & \mathcal{W}_g \\ \big\downarrow & & \big\downarrow \\ \mathcal{A}_g & \longrightarrow & \mathcal{B}_g \end{array}$$

is easily seen to be cartesian.

Suppose that L is an extension of k, and let X a g-dimensional abelian variety over L. Set $W \stackrel{\text{def}}{=} X/\boldsymbol{\mu}_2$, and call $\rho\colon X \to W$ the projection, $T \subseteq W$ the finite subset of codimension at least 3 consisting of singular points, and $w_0 \in T$ the image of the origin in X. The étale double cover

$$\rho\colon X \smallsetminus \rho^{-1}(T) \longrightarrow W \smallsetminus T$$

gives rise to an equivalence of categories between $\boldsymbol{\mu}_2$-equivariant invertible sheaves on X and invertible sheaves on $W \smallsetminus T$; the canonical sheaf $\omega_{W \smallsetminus T/L}$

corresponds to $\omega_{X/L}$. Choose an isomorphism between $\omega_{X/L}$ and $\mathcal{O}_X$ (as sheafs of $\mathcal{O}_X$-modules). This isomorphism becomes $\boldsymbol{\mu}_2$-equivariant if we let the generator -1 of $\boldsymbol{\mu}_2$ acts on $\mathcal{O}_X$ by multiplication by $(-1)^g = -1$. This gives a $\boldsymbol{\mu}_2$-equivariant isomorphism of $\omega_{X/L}^{\otimes 2}$ with $\mathcal{O}_X$, and an isomorphism $\pi_{W\smallsetminus T/L}^{\otimes 2} \simeq \mathcal{O}_{W\smallsetminus T}$. Hence, $X \smallsetminus \rho^{-1}(T)$ is the double cover of square roots in $\omega_{W\smallsetminus T/L}$ of the element corresponding to 1 under the isomorphism $\omega_{W\smallsetminus T/L}^{\otimes 2} \simeq \mathcal{O}_{W\smallsetminus T}$ above. Thus X can be recovered from the isomorphism $\omega_{W\smallsetminus T/L}^{\otimes 2} \simeq \mathcal{O}_{W\smallsetminus T}$ as the normalization of W in the function field of' the corresponding étale double cover.

Now, assume that W is a variety over an extension L of k, with a rational point $w_0 \in W(L)$, with the property that there exists an extension E/L and an isomorphism $W_E \overset{\text{def}}{=} W \times_{\operatorname{Spec} L} \operatorname{Spec} E \simeq Y/\boldsymbol{\mu}_2$, where Y is a g-dimensional abelian variety over E, such that w_0 corresponds to the image of the origin in $Y/\boldsymbol{\mu}_2$. Call T the singular locus of W; clearly T is finite, hence of codimension at least 3. Since $\mathrm{H}^0(W\smallsetminus T, \mathcal{O}^*) = L^*$, the Picard group of $W \smallsetminus T$ injects into the Picard group of $(W \smallsetminus T) \times_{\operatorname{Spec} L} \operatorname{Spec} E$. Hence $\omega_{W\smallsetminus T/L}^{\otimes 2}$ is isomorphic to $\mathcal{O}_{W\smallsetminus T}$. If we choose such an isomorphism, we get an étale double cover $U \to W \smallsetminus T$. Let X be the normalization of $W \smallsetminus T$ in the function field of U. We have a finite morphism $X \to W$, which is étale over $W \smallsetminus T$. The induced cover $X_E \to W_E$ is not necessarily isomorphic to $Y \to W_E$, because the pullback of the isomorphism $\omega_{W\smallsetminus T/L}^{\otimes 2} \simeq \mathcal{O}_{W\smallsetminus T}$ does not necessarily coincide with the isomorphism $\omega_{W_E\smallsetminus T_E/E}^{\otimes 2} \simeq \mathcal{O}_{W_E\smallsetminus T_E}$ coming from Y. However, the obstruction to these two covers being isomorphic lives in E^*/E^{*2}, hence these covers become isomorphic after making a further, at most quadratic, extension of E. It follows that $X \to W$ is ramified over w_0; hence we have a unique rational point of X lying over w_0. So X is a variety with a rational point that becomes an abelian variety after an extension of base field, and so it is itself an abelian variety. Then it is easy to see that $X/\boldsymbol{\mu}_2 = W$. In other words, if we have a variety with a rational point that becomes a generalized Kummer variety after a base change, it is itself a generalized Kummer variety.

Since the morphism $\mathcal{B}_g \to \mathbf{A}_g$ is generically an isomorphism, we have a lifting $\operatorname{Spec} k(\mathbf{A}_g) \to \mathcal{B}_g$; call w_0 the rational point on W coming from the section $\mathcal{B}_g \to \mathcal{W}_g$. Denote by W the fiber of $\mathcal{W}_g \to \mathcal{B}_g$. There exists an extension E of $k(\mathbf{A}_g)$ and a lifting $\operatorname{Spec} E \to \mathcal{A}_g$ of the composite $\operatorname{Spec} E \to \operatorname{Spec} L \to \mathcal{B}_g$. After base changing to E, the variety W becomes a generalized Kummer variety. As we have just seen, there exists an abelian variety X over $k(\mathbf{A}_g)$ such that $X/\boldsymbol{\mu}_2$ such that $X/\boldsymbol{\mu}_2 = W$. This variety is isomorphic over the algebraic closure $\overline{k(\mathbf{A}_g)}$ to the abelian variety coming from a lifting $\operatorname{Spec} \overline{k(\mathbf{A}_g)} \to \mathcal{A}_g$ of morphism $\operatorname{Spec} \overline{k(\mathbf{A}_g)} \to \mathbf{A}_g$; hence X acquires a principal polarization after base change to a finite extension of $k(\mathbf{A}_g)$. But since the Néron–Severi group of a generic principally polarized abelian variety is infinite cyclic, such a polarization is unique, and descends to a principal polarization over X. This yield a lifting $\operatorname{Spec} k(\mathbf{A}_g) \to \mathcal{A}_g$, and concludes the proof.

(b) Assume that g is even. We need to prove the inequalities

$$2 \le i_g \le 2^{s+3}.$$

Let us start from the inequality $i_g \ge 2$: we need to show that the gerbe $(\mathcal{A}_g)_{k(\mathbf{A}_g)}$ is not split. The restriction $\mathcal{H}_g \to \mathcal{A}_g$ of the Abel–Jacobi map is well known to be a locally closed embedding (the key point is that the antipodal map $x \mapsto -x$ on the Jacobian of a hyperelliptic curve is realized by the hyperelliptic involution); this induces a locally closed embedding of moduli spaces $\mathbf{H}_g \subseteq \mathbf{A}_g$. If $g = 2$, then this is an open embedding, so the generic gerbes of $\mathcal{H}_2 = \mathcal{M}_2$ and of $\mathcal{A}_2$ coincide, and $i_2 = 2$, as we have seen in the proof of Theorem 7.2. This proves the equality $\operatorname{ed} \mathcal{A}_2 = 5$.

When $g \ge 4$, consider the open subscheme $U \subseteq \mathcal{A}_g$ whose geometric points correspond to abelian varieties whose automorphism group is $\boldsymbol{\mu}_2$. Then U is smooth, and the restriction $\mathcal{A}_g\,|_U \to U$ is a gerbe banded by $\boldsymbol{\mu}_2$; furthermore, since the automorphism group of a generic hyperelliptic curve is $\boldsymbol{\mu}_2$, we have $U \cap \mathbf{H}_g \neq \emptyset$. Consider the class in the Brauer group of U coming from the class in $\mathrm{H}^2(U, \boldsymbol{\mu}_2)$ of $\mathcal{A}_g\,|_U$. Since the Brauer group of U injects in the Brauer group of its field of functions $k(\mathbf{A}_g)$, if the generic gerbe of $\mathcal{A}_g$ were split, the pullback of $\mathcal{A}_g$ along any morphism $\operatorname{Spec} L \to U$, where L is a field, would also be split. However, the pullback of $\mathcal{A}_g$ to the generic point $\operatorname{Spec} k(\mathbf{H}_g) \in U$, which is the generic gerbe of $\mathcal{H}_g$, is not split, as we have seen in the proof of Theorem 7.2. This gives a contradiction, and proves the inequality $i_g \ge 2$.

Let us finish the proof by showing the inequality $i_g \le 2^{s+3}$. Denote by $\mathcal{A}_g[3]$ the stack whose objects are families of principally polarized abelian varieties $X \to S$, together with a nowhere zero section $S \to X[3]$ of the 3-torsion group scheme $X[3] \to S$, and by $\mathbf{A}_g[3]$ its moduli space. It is very well known that $\mathbf{A}_g[3]$ is integral. The natural forgetful morphism $\mathcal{A}_g[3] \to \mathcal{A}_g$ is a representable étale cover of degree $3^{2g} - 1$; furthermore, the antipodal map $x \mapsto -x$ of an abelian variety does not leave any non-zero element of $X[3]$ invariant. Thus $\mathcal{A}_g[3]$ is generically a scheme, and the pullback of the generic gerbe of $\mathcal{A}_g$ to $\operatorname{Spec} \mathbf{A}_g[3]$ is split. This shows that i_g divides $3^{2g} - 1$. To finish the proof of part (b) we need to show that the highest power of 2 that divides $3^{2g} - 1$ is 2^{s+3}. Write $g = 2^s h$ as in the statement of part (b). Then

$$3^{2g} - 1 = 3^{2^{s+1}h} - 1 = (3^{2^{s+1}} - 1) \sum_{i=0}^{h-1} 3^{2^{s+1}i}.$$

The second factor is odd, so we may assume that $h = 1$. In other words, we need to show that the highest power of 2 dividing $3^{2^{s+1}} - 1$ is 2^{s+3}. This is easily done by induction on s. The base case, where $s = 0$ is evident, and the induction step follows from the formula

$$3^{2^{s+2}} - 1 = (3^{2^{s+1}} - 1)(3^{2^{s+1}} + 1)\,.$$

Here $3^{2^{s+1}} + 1 \equiv 2 \pmod 4$, so $3^{2^{s+1}} + 1$ contributes exactly one factor of 2 to the above product. ♠

Remark 9.2. It seems natural to conjecture that $i_g = 2$, hence $\operatorname{ed}\mathcal{A}_g = g(g+1)+2$, for any even g.

Question 9.3. What is the essential dimension of the stack of (unpolarized) abelian varieties of dimension g? This is not an algebraic stack, so the techniques of this paper do not apply.

References

[ACV03] Dan Abramovich, Alessio Corti, and Angelo Vistoli, *Twisted bundles and admissible covers*, Comm. Algebra **31** (2003), no. 8, 3547–3618, Special issue in honor of Steven L. Kleiman.

[AV02] Dan Abramovich and Angelo Vistoli, *Compactifying the space of stable maps*, J. Amer. Math. Soc. **15** (2002), no. 1, 27–75.

[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.

[BRV] Patrick Brosnan, Zinovy Reichstein, and Angelo Vistoli, *Essential dimension, spinor groups and quadratic forms*, `http://annals.math.princeton.edu/issues/2008/FinalFiles/BrosnanReichsteinVistoliFinal.pdf`.

[BRV07] ———, *Essential dimension and algebraic stacks*, 2007, `arXiv:math/0701903v1 [math.AG]`.

[BS08] Patrick Brosnan and Ramesh Sreekantan, *Essential dimension of abelian varieties over number fields*, preprint available at `http://www.math.ubc.ca/~brosnan/Papers/ab-number.pdf`, 2008.

[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, to appear, 2006.

[DD99] Pierre Dèbes and Jean-Claude Douai, *Gerbes and covers*, Comm. Algebra **27** (1999), no. 2, 577–594. MR MR1671938 (2000a:14017)

[Del80] Pierre Deligne, *La conjecture de Weil. II*, Inst. Hautes Études Sci. Publ. Math. (1980), no. 52, 137–252. MR MR601520 (83c:14017)

[DL09] Ajneet Dhillon and Nicole Lemire, *Upper bounds for the essential dimension of the moduli stack of* SL_n*-bundles over a curve*, 2009, preprint.

[Ful98] William Fulton, *Intersection theory*, second ed., Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics, vol. 2, Springer-Verlag, Berlin, 1998.

[Gar] Skip Garibaldi, *Cohomological invariants: exceptional groups and spin groups*, to appear in Memoirs of the AMS, preprint available at `arXiv:math.AG/0411424`.

[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.

[Gro64] Alexander Grothendieck, *Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. I*, Inst. Hautes Études Sci. Publ. Math. (1964), no. 20, 259.

[Gro67] A. Grothendieck, *Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas IV*, Inst. Hautes Études Sci. Publ. Math. (1967), no. 32, 361.

[Har77] Robin Hartshorne, *Algebraic geometry*, Graduate Texts in Mathematics, no. 52, Springer-Verlag, New York, 1977.

[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.

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

[KM08] ———, *Essential dimension of finite p-groups*, Invent. Math. **172** (2008), no. 3, 491–508. MR MR2393078 (2009b:12009)

[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.

[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.

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

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

[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.

[Ser02] Jean-Pierre Serre, *Galois cohomology*, English ed., Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2002, Translated from the French by Patrick Ion and revised by the author. MR MR1867431 (2002i:12004)

[SGA72] *Théorie des topos et cohomologie étale des schémas. Tome 2*, Lecture Notes in Mathematics, Vol. 270, Springer-Verlag, Berlin, 1972, Séminaire de Géométrie Algébrique du Bois-Marie 1963–1964 (SGA 4), Dirigé par M. Artin, A. Grothendieck et J. L. Verdier. Avec la collaboration de N. Bourbaki, P. Deligne et B. Saint-Donat. MR MR0354653 (50 #7131)

[SGA73] *Groupes de monodromie en géométrie algébrique. II*, Lecture Notes in Mathematics, Vol. 340, Springer-Verlag, Berlin, 1973, Séminaire de Géométrie

Algébrique du Bois-Marie 1967–1969 (SGA 7 II), Dirigé par P. Deligne et N. Katz. MR MR0354657 (50 #7135)

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

(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