ESSENTIAL DIMENSION

Angelo Vistoli
Scuola Normale Superiore

Budapest, May 2008

Joint work with Patrick Brosnan and Zinovy Reichstein
University of British Columbia

Posted at http://arxiv.org/abs/math.AG/0701903

Let us start from an example. Fix a base field k. Let K be an extension of k, and consider a smooth projective curve C of genus g defined over K. The curve C will always be defined over an intermediate field $k \subseteq L \subseteq K$ whose transcendence degree over k is finite.

The *essential dimension* $\operatorname{ed} C$ is the minimum of $\operatorname{tr\,deg}_k L$, where L is a field of definition of C. In other words, $\operatorname{ed} C$ is the minimal number of independent parameters needed to write down C.

It is also natural to ask: what is the supremum of $\operatorname{ed} C$ for all curves of fixed genus g over all extensions K?

For example, if $g = 0$ then C is a conic in $\mathbb{P}^2_K$. After a change of coordinates, it is defined by an equation $ax^2 + by^2 + z^2 = 0$. Hence it is defined over $k(a, b)$, and $\operatorname{ed} C \leq 2$. It follows from Tsen's theorem that if a and b are independent variables, then $\operatorname{ed} C = 2$.

All fields will have characteristic 0.

Let k be a field, Fields_k the category of extensions of k.

Let $F\colon \mathrm{Fields}_k \to \mathrm{Sets}$ be a functor. If ξ is an object of some $F(K)$, a *field of definition* of ξ is an intermediate field $k \subseteq L \subseteq K$ such that ξ is in the image of $F(L) \to F(K)$.

Definition (Merkurjev)**.** The *essential dimension of* ξ, denoted by $\operatorname{ed}\xi$, is the least transcendence degree $\operatorname{tr}\deg_k L$ of a field of definition L of ξ.

The *essential dimension of* F, denoted by $\operatorname{ed} F$, is the supremum of the essential dimensions of all objects ξ of all $F(K)$.

The essential dimension $\operatorname{ed}\xi$ is finite, under weak hypothesis on F. But $\operatorname{ed} F$ could still be $+\infty$.

The previous question could be stated as: if F_g is the functor that associates with each extension $k \subseteq K$ the set of isomorphism classes of smooth projective curves of genus g, what is $\operatorname{ed} F_g$? In other words, how many independent variables do you need to write down a general curve of genus g? An easy argument using moduli spaces of curves reveals that $\operatorname{ed} F_g \geq 3g - 3$ for $g \geq 2$.

Theorem (BRV)**.**

$$\operatorname{ed} F_g = \begin{cases} 2 & \text{if } g = 0 \\ +\infty & \text{if } g = 1 \\ 5 & \text{if } g = 2 \\ 3g - 3 & \text{if } g \geq 3. \end{cases}$$

Merkurjev's definition generalizes the notion of *essential dimension of a group*, due to Buhler and Reichstein.

Definition (Buhler, Reichstein, reinterpreted through Merkurjev)**.** Let G be an algebraic group over k. The essential dimension of G, denoted by $\operatorname{ed} G$, is the essential dimension of the functor $\mathrm{H}^1(-, G)$, where $\mathrm{H}^1(K, G)$ is the set of isomorphism classes of G-torsors over K.

Assume that $\sigma\colon (k^n)^{\otimes r} \to (k^n)^{\otimes s}$ is a tensor on an n-dimensional k-vector space k^n, and G is the group of automorphisms of k^n preserving σ. Then G-torsors over K correspond to *twisted forms* of σ, that is, n-dimensional vector spaces V over K with a tensor $\tau\colon V^{\otimes r} \to V^{\otimes s}$ that becomes isomorphic to $(k^n, \sigma) \otimes_k \overline{K}$ over the algebraic closure $\overline{K}$ of K.

For example, GL_n-torsors correspond to n-dimensional vector spaces, which are all defined over k. Hence $\operatorname{ed} \mathrm{GL}_n = 0$. Also $\operatorname{ed} \mathrm{SL}_n = \operatorname{ed} \mathrm{Sp}_n = 0$.

Let $G = \mathrm{O}_n$. The group O_n is the automorphism group of the standard quadratic form $x_1^2 + \cdots + x_n^2$, which can be thought of as a tensor $(k^n)^{\otimes 2} \to k$. Then O_n-torsors correspond to non-degenerate quadratic forms on K of dimension n. Since every such quadratic form can be diagonalized in the form $a_1 x_1^2 + \cdots + a_n x_n^2$, it is defined over $k(a_1, \ldots, a_n)$. So $\operatorname{ed} \mathrm{O}_n \leq n$.

In fact, it was proved by Reichstein that $\operatorname{ed} \mathrm{O}_n = n$. Also, $\operatorname{ed} \mathrm{SO}_n = n - 1$.

In general finding lower bounds is much harder than finding upper bounds.

PGL_n is the automorphism group of $\mathbb{P}^n$ and of the matrix algebra M_n. Therefore PGL_n-torsors over K correspond to twisted forms of the multiplication tensor $\mathrm{M}_n^{\otimes 2} \to \mathrm{M}_n$, that is, to K-algebras A that become isomorphic to M_n over $\overline{K}$. These are the *central simple algebras* of degree n.

Main open problem: what is $\operatorname{ed}\mathrm{PGL}_n$?

Assume that k contains enough roots of 1. It is know that $\operatorname{ed}\mathrm{PGL}_2 = \operatorname{ed}\mathrm{PGL}_3 = 2$; this follows from the fact that central division algebras of degree 2 and 3 are *cyclic*. This is easy for degree 2; in degree 3 it is a theorem of Albert. A cyclic algebra of degree n over K has a presentation of the type $x^n = a$, $y^n = b$ and $yx = \omega xy$, where $a, b \in K^*$ and ω is a primitive n^{th} root of 1. Hence a cyclic algebra is defined over a field of the type $k(a, b)$, and has essential dimension at most 2.

When n is a prime larger than 3, it is only known (due to Lorenz, Reichstein, Rowen and Saltman) that

$$2 \leq \operatorname{ed} \mathrm{PGL}_n \leq \frac{(n-1)(n-2)}{2}.$$

Computing $\operatorname{ed} \mathrm{PGL}_n$ when n is a prime is an extremely interesting question, linked with the problem of cyclicity of division algebras of prime degree. If every division algebra of prime degree is cyclic, then $\operatorname{ed} \mathrm{PGL}_n = 2$. Most experts think that a generic division algebra of prime degree larger than 3 should not be cyclic. One way to show this would be to prove that $\operatorname{ed} \mathrm{PGL}_n > 2$ when n is a prime larger than 3. Unfortunately, our methods don't apply to this problem.

What was known about classical groups:

$$\operatorname{ed}\mathrm{GL}_n = \operatorname{ed}\mathrm{SL}_n = \operatorname{ed}\mathrm{Sp}_n = 0$$
$$\operatorname{ed}\mathrm{O}_n = n$$
$$\operatorname{ed}\mathrm{SO}_n = n-1$$
$$\operatorname{ed}\mathrm{PGL}_n \leq n^2 - n\,.$$

and some more assorted results about PGL_n.

How about spin groups?

Recall that Spin_n is the universal cover of the group SO_n. It is a double cover, hence there is a central extension

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

where $\mu_2 = \{\pm 1\}$ is the group of square roots of 1.

The following result is due to Chernousov–Serre and Reichstein–Youssin.

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

Furthermore $\operatorname{ed} \operatorname{Spin}_n$ had been computed by Rost for $n \leq 14$:

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

All this seemed to suggest that $\operatorname{ed} \operatorname{Spin}_n$ should be a slowly increasing function of n.

Assume that $\sqrt{-1} \in k$.

Theorem (BRV). *If n is not divisible by 4 and $n \geq 15$, then*

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

Theorem (BRV, Merkurjev). *If n is divisible by 4 and $n \geq 16$, call 2^k the largest power of 2 that divides n. Then*

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

The lower bound is due to Merkurjev, and improves on a previous lower bound due of BRV.

$$
\begin{aligned}
\operatorname{ed}\operatorname{Spin}_{15} &= 23\\
\operatorname{ed}\operatorname{Spin}_{16} &= 24\\
\operatorname{ed}\operatorname{Spin}_{17} &= 120\\
\operatorname{ed}\operatorname{Spin}_{18} &= 103\\
\operatorname{ed}\operatorname{Spin}_{19} &= 341\\
326 \leq \operatorname{ed}\operatorname{Spin}_{20} &\leq 342\\
\operatorname{ed}\operatorname{Spin}_{21} &= 814\\
\operatorname{ed}\operatorname{Spin}_{22} &= 793\\
\operatorname{ed}\operatorname{Spin}_{23} &= 1795\\
1780 \leq \operatorname{ed}\operatorname{Spin}_{24} &\leq 1796\\
\operatorname{ed}\operatorname{Spin}_{25} &= 3796
\end{aligned}
$$

This result can be applied to the theory of quadratic forms. Assume that $\sqrt{-1} \in k$. For each extension $k \subseteq K$, denote by h_K the hyperbolic quadratic form $x_1 x_2$ on K. Recall that the *Witt ring* $\mathrm{W}(K)$ is the set of isometry classes of non-degenerate quadratic forms on K, modulo the relation that identifies q and q' if $q \oplus h_K^r \simeq q' \oplus h_K^s$ for some $r, s \geq 0$. Addition is induced by direct product, multiplication by tensor product. There is a rank homomorphism $\mathrm{rk}\colon \mathrm{W}(K) \to \mathbb{Z}/2\mathbb{Z}$; its kernel is denoted by $\mathrm{I}(K)$.

Let $q = a_1 x_1^2 + \cdots + a_n x_n^2$. Then $[q] \in \mathrm{I}(K)$ if and only if n is even. $[q] \in \mathrm{I}(K)^2$ if and only if $[q] \in \mathrm{I}(K)$ and $[a_1 \ldots a_n] = 1 \in K^*/(K^*)^2$, or, equivalently, if q comes from an SO_n-torsor. It is known that $[q] \in \mathrm{I}(K)^3$ if and only if $[q] \in \mathrm{I}(K)^2$ and q comes from a Spin_n-torsor (we say that q *has a spin structure*).

For each $n \geq 0$, an *n-fold Pfister form* is a quadratic form of dimension 2^s of the form

$$\ll a_1, \ldots, a_s \gg \stackrel{\text{def}}{=} (x_1 + a_1 x_2^2) \otimes \cdots \otimes (x_1 + a_s x_2^2)$$

for some $(a_1, \ldots, a_s) \in (K^*)^s$. If $[q] \in \mathrm{W}(K)$, it is easy to see $[q] \in \mathrm{I}(K)^s$ if and only if $[q]$ a sum of classes of s-fold Pfister forms.

For each non-degenerate quadratic form q over K of dimension n such that $[q] \in \mathrm{I}(K)^s$, denote by $\mathrm{Pf}(s, q)$ the least integer N such that $[q] \in \mathrm{W}(K)$ can be written as a sum of N s-fold Pfister forms. For any $n > 0$, the *s-fold Pfister number* $\mathrm{Pf}_k(s, n)$ is the supremum of the $\mathrm{Pf}(s, q)$ taken over all extensions K of k and all q as above.

The following estimates are elementary:

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

Nothing is known about $\mathrm{Pf}_k(s,n)$ for $s > 3$; for all we know $\mathrm{Pf}_k(s,n)$ could be infinite. However, it is known that $\mathrm{Pf}_k(3,n)$ is finite for all $n > 0$. This follows from the existence of a "versal" Spin_n-torsor.

Theorem (BRV)**.**

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

The idea of the proof is the following. Suppose that q is an n-dimensional quadratic form over an extension K of k coming from a Spin_n-torsor with maximum essential dimension. Then if q comes from a form q' with a spin structure from $k \subseteq L \subseteq K$, the transcendence degree of L over k has to be very large. Assume that $[q]$ is a sum of N 3-fold Pfister forms. By the Witt cancellation theorem, $q \oplus h_K^r$ is a sum of N Pfister forms (unless N is really small, which is easy to exclude). Each 3-fold Pfister form is defined over an extension of k of transcendence degree at most 3. Then $q \oplus h_K^r$ comes from a form q_L with a spin structure over an extension $k \subseteq L$ of transcendence degree at most $3N$. By making a further small extension of L we may assume that q_L splits as $q' \oplus h_L^r$. Again by Witt's cancellation theorem, q comes from q'; hence $3N$ has to be large.

The proof of the result on spinor groups is based on the study of essential dimension of *gerbes*. Suppose that G is an algebraic group over a field k, and suppose that Z is a central subgrop of G. Set $H = G/Z$. Let P be an H-torsor over an K, let ∂P be the gerbe over K of liftings of P to G: if E is an extension of K, then $\partial P(E)$ is the category of G-torsors Π over $\operatorname{Spec} E$, with a G-equivariant morphism $\Pi \to P_E$. The gerbe ∂P is banded by Z; that is, the automorphism group of an object of $\partial P(E)$ is canonically isomorphic to $Z(E)$. We are interested in the essential dimension of ∂P over the field K, for the following reason.

Theorem.

$$\operatorname{ed} G \geq \operatorname{ed}(\partial P/K) - \dim H.$$

What can we say about essential dimension of gerbes banded by Z?

By a well known result of Grothendieck and Giraud equivalence classes of gerbes banded by Z are parametrized by $\mathrm{H}^2(K, Z)$. We are interested in the case $Z = \mu_n$. From the Kummer sequence

$$1 \longrightarrow \boldsymbol{\mu}_n \longrightarrow \mathbb{G}_{\mathrm{m}} \xrightarrow{\times n} \mathbb{G}_{\mathrm{m}} \longrightarrow 1$$

we get an exact sequence

$$0 = \mathrm{H}^1(K, \mathbb{G}_{\mathrm{m}}) \longrightarrow \mathrm{H}^2(K, \boldsymbol{\mu}_n) \longrightarrow \mathrm{H}^2(K, \mathbb{G}_{\mathrm{m}}) \xrightarrow{\times n} \mathrm{H}^2(K, \mathbb{G}_{\mathrm{m}}).$$

The group $\mathrm{H}^2(K, \mathbb{G}_{\mathrm{m}})$ is called the *Brauer group* of K, and is denoted by $\mathrm{Br}\, K$. If $\overline{K}$ is the algebraic closure of K and $\mathcal{G}$ is the Galois group of $\overline{K}$ over K, then $\mathrm{Br}\, K = \mathrm{H}^2(\mathcal{G}, \overline{K}^*)$. Thus $\mathrm{H}^2(K, \boldsymbol{\mu}_n)$ is the n-torsion part of $\mathrm{Br}\, K$.

A gerbe $\mathcal{X}$ banded by $\boldsymbol{\mu}_n$ has a class $[\mathcal{X}]$ in $\mathrm{Br}\, K$.

Each element of $\operatorname{Br} K$ comes from a PGL_m-torsor for some m, via the non-commutative boundary operator $\mathrm{H}^1(K, \mathrm{PGL}_m) \to \mathrm{H}^2(K, \mathbb{G}_\mathrm{m})$ coming from the sequence

$$1 \longrightarrow \mathbb{G}_\mathrm{m} \longrightarrow \mathrm{GL}_m \longrightarrow \mathrm{PGL}_m \longrightarrow 1.$$

The least m such that $\alpha \in \operatorname{Br} K$ is in the image of $\mathrm{H}^1(K, \mathrm{PGL}_m)$ is called the *index* of α, denoted by $\operatorname{ind} \alpha$.

Theorem. *Let $\mathcal{X}$ be a gerbe banded by $\boldsymbol{\mu}_n$, where n is a prime power larger than* 1*. Then* $\operatorname{ed} \mathcal{X}$ *equals the index* $\operatorname{ind} [\mathcal{X}]$ *in* $\operatorname{Br} K$*.*

If n is a power of a prime p, then the index is also a power of p. This allows to show that in several case $\operatorname{ed} \mathcal{X}$ is much larger than previously suspected.

The essential ingredient in the proof is a theorem of Karpenko on the canonical dimension of Brauer-Severi varieties.

These results can be applied to the sequence

$$1 \longrightarrow \mu_2 \longrightarrow \mathrm{Spin}_n \longrightarrow \mathrm{SO}_n \longrightarrow 1.$$

By associating each SO_n-torsor the class of its boundary gerbe ∂P, we obtain a non-abelian boundary map

$$\mathrm{H}^1(K, \mathrm{SO}_n) \longrightarrow \mathrm{H}^2(K, \mu_2) \subseteq \mathrm{Br}\, K.$$

The image of P in $\mathrm{Br}\, K$ is know as the *Hasse–Witt invariant* of the associated quadratic form. The two result above combine to give the following.

Theorem. *If P is an SO_n-torsor, then*

$$\mathrm{ed}\, \mathrm{Spin}_n \geq \mathrm{ind}[\partial P] - \frac{n(n-1)}{2}.$$

Theorem. *If P is an SO_n-torsor, then*

$$\operatorname{ed} \mathrm{Spin}_n \geq \operatorname{ind}[\partial P] - \frac{n(n-1)}{2}.$$

It is known that if P is a generic quadratic form of determinant 1, its Hasse–Witt invariant has index $2^{\lfloor \frac{n-1}{2} \rfloor}$. From this we obtain the inequality

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

How about the essential dimension of gerbes banded by μ_n, when n is not a prime power?

Let $\mathcal{X}$ be a gerbe banded by μ_n. Write the prime factor decomposition

$$\operatorname{ind}[\mathcal{X}] = p_1^{e_1} \dots p_r^{e_r}.$$

Then

$$\operatorname{ed} \mathcal{X} \leq p_1^{e_1} + \cdots + p_r^{e_r} - r + 1$$

Conjecturally, equality holds. This is equivalent to a conjecture of Colliot-Thélène, Karpenko and Merkurjev on the canonical dimension of Brauer–Severi schemes. They proved it for $\operatorname{ind}[\mathcal{X}] = 6$.