Computing in Picard groups of projective curves over finite fields

Peter Bruin
Essen, 10 November 2009

Abstract

Following the work of K. Khuri-Makdisi (see [3] and [4]), I will describe a way of representing a smooth projective curve over a field, and of divisors on it, that allows fast computation of group operations in the Picard group. This is especially interesting in the case of modular curves, where such a representation can be computed from spaces of modular forms. If the base field is finite, there are additional operations such as choosing uniformly random elements of the Picard group and computing Frobenius maps, Frey–Rück pairings and (for modular curves) Hecke operators. I will explain, based in part on work of J.-M. Couveignes [1], how these operations can be done efficiently in the setting of Khuri-Makdisi's representation of the curve.

1. Representing curves, divisors and elements of the Picard group

Notation. Throughout, X denotes a complete, smooth and geometrically connected curve of genus g over a field k, and $\mathcal{L}$ denotes a line bundle on X of degree

$$d \geq 2g + 1.$$

The assumption on d implies that $\mathcal{L}$ is very ample, i.e. it defines a closed immersion

$$X \to \mathbf{P}\Gamma(X, \mathcal{L})$$

of varieties over k. (Here $\mathbf{P}V$ denotes the projective space over k of hyperplanes in the k-vector space V.) Points on X can therefore be described by hyperplanes; a point $P \in X(k)$ corresponds to the hyperplane $\Gamma(X, \mathcal{L}(-P))$ in $\Gamma(X, \mathcal{L})$.

The same assumption implies that the natural multiplication maps

$$\Gamma(X, \mathcal{L}^i) \otimes_k \Gamma(X, \mathcal{L}^j) \to \Gamma(X, \mathcal{L}^{i+j}).$$

are surjective, or in other words that the embedding given by $\mathcal{L}$ is *projectively normal*. (This is a classical theorem of Castelnuovo, Mattuck and Mumford.) Equivalently, the projective coordinate ring S_X of X can be expressed as

$$S_X = \bigoplus_{i \geq 0} \Gamma(X, \mathcal{L}^i).$$

We assume that we know the spaces $\Gamma(X, \mathcal{L}^i)$ and the multiplication maps between them for $i = 0, 1, \ldots, h$, where h is some sufficiently large positive integer. Giving these data is equivalent to giving the k-algebra

$$S_X^{(h)} = S_X/(\text{homogeneous elements of degree} > h).$$

The most straightforward way to give the k-algebra structure is to write down matrices with respect to certain k-bases of the $\Gamma(X, \mathcal{L}^i)$. However, a more efficient representation is to choose trivialisations of $\mathcal{L}$ at sufficiently many points of X (i.e. so many that a section of $\mathcal{L}^h$ is determined by its values at these points), so that the multiplication maps can be computed pointwise.

Example. Let X be a modular curve, say for concreteness $X_0(n)$ or $X_1(n)$ with $n \geq 5$. Then a suitable line bundle $\mathcal{L}$ is the line bundle ω^2 of modular forms of weight 2 on X. The spaces $\Gamma(X, \mathcal{L}^i)$ can be computed using modular symbols. We can express the modular forms as q-expansions, and the multiplication is simply multiplication of power series. Another possibility is to compute the values of the modular forms at sufficiently many rational points (with respect to some fixed trivialisation).

Effective divisors on X can be represented in the following way. Let D be an effective divisor on X, and let a be a positive integer such that

$$ad - \deg D \geq 2g + 1.$$

Then D is uniquely determined by the subspace

$$\Gamma(X, \mathcal{L}^a(-D)) \subset \Gamma(X, \mathcal{L}^a)$$

of codimension d.

Using the representation of divisors just described, Khuri-Makdisi has given the (asymptotically) fastest algorithms known to date (measured in operations in the field k) for computing with divisors on general curves. An important source of efficiency is that inside the algorithms, the space $\Gamma(X, \mathcal{L}^a)$ is replaced at strategical moments by a small subset consisting of sections whose common divisor equals D.

Khuri-Makdisi's algorithms allow us to add, subtract and intersect divisors of sufficiently small degree, and to test whether a given subspace of $\Gamma(X, \mathcal{L}^i)$ is of the form $\Gamma(X, \mathcal{L}^i(-D))$ for some effective divisor D, in time $O((\deg \mathcal{L})^{3+\epsilon})$. Using fast linear algebra, this can be improved to $O((\deg \mathcal{L})^{2.376})$. The algorithms are based on the following two results.

Lemma 1.1. *Let $\mathcal{M}$ and $\mathcal{N}$ be line bundles on X whose degrees are at least $2g + 1$. Then the canonical k-linear map*

$$\Gamma(X, \mathcal{M}) \otimes_k \Gamma(X, \mathcal{N}) \longrightarrow \Gamma(X, \mathcal{M} \otimes_{\mathcal{O}_X} \mathcal{N})$$

is surjective.

Lemma 1.2. *Let $\mathcal{M}$ and $\mathcal{N}$ be line bundles on X such that $\mathcal{N}$ is generated by global sections, and let D be any effective divisor on X. Then the inclusion*

$$\Gamma(X, \mathcal{M}(-D)) \subset \bigl\{ s \in \Gamma(X, \mathcal{M}) \bigm| s\Gamma(X, \mathcal{N}) \subset \Gamma(X, \mathcal{M} \otimes_{\mathcal{O}_X} \mathcal{N}(-D)) \bigr\}$$

is an equality.

We now describe how to represent elements of the group

$$\operatorname{Pic}^0 X = \{\text{line bundles of degree } 0\}/\cong.$$

The isomorphism class of a line bundle $\mathcal{M}$ of degree 0 is represented by an effective divisor D of degree d such that

$$\mathcal{M} \cong \mathcal{L}(-D).$$

This divisor is in turn represented as the subspace

$$\Gamma(X, \mathcal{L}^2(-D)) \subset \Gamma(X, \mathcal{L}^2).$$

Remark. Let $\operatorname{Gr}^d \Gamma(X, \mathcal{L}^2)$ denote the Grassmann variety classifying subspaces of codimension d in $\Gamma(X, \mathcal{L}^2)$. Then we are in fact considering the morphisms

$$\operatorname{Pic}^0_{X/k} \longleftarrow \operatorname{Sym}^d X \longrightarrow \operatorname{Gr}^d \Gamma(X, \mathcal{L}^2)$$

of varieties over k, where the left arrow (surjective) sends a divisor D of degree d to the isomorphism class of $\mathcal{L}(-D)$ and the right arrow (injective) sends D to the subspace $\Gamma(X, \mathcal{L}^2(-D))$ of $\Gamma(X, \mathcal{L}^2)$.

Remark. By taking $d \geq 2g + 1$ and representing D as a subspace of $\Gamma(X, \mathcal{L}^2)$, we have restricted ourselves to a description of Khuri-Makdisi's *medium model* of $\operatorname{Pic}^0 X$. In the *large* model, D is represented as the subspace $\Gamma(X, \mathcal{L}^3(-D))$ of $\Gamma(X, \mathcal{L}^3)$. The advantage of this is that the group operations are conceptually the easiest to implement. In the *small model*, D is also represented as a subspace of $\Gamma(X, \mathcal{L}^3)$, but $\mathcal{L}$ can be taken of degree less than $2g + 1$. The advantage of the small model is that the operations require asymptotically the least number of field operations.

2. Basic operations in the Picard group

There are three basic operations.

- *zero test*: given a subspace of codimension d in $\Gamma(X, \mathcal{L}^2)$, decide whether it represents $0 \in \operatorname{Pic}^0 X$.
- *zero element*: output a subspace of codimension $\deg \mathcal{L}$ in $\Gamma(X, \mathcal{L}^2)$ representing $0 \in \operatorname{Pic}^0 X$.
- *addflip*: given two subspaces of $\Gamma(X, \mathcal{L}^2)$ representing elements $x, y \in \operatorname{Pic}^0 X$, compute a subspace of $\Gamma(X, \mathcal{L}^2)$ representing $-x - y$.

From the "addflip" operation, one immediately gets negation ($-x = -x - 0$), addition ($x + y = -(-x - y)$) and subtraction ($x - y = -(-x) - y$). Testing equality of two elements x and y can be done by computing $x - y$ and testing whether the result equals zero.

Remark. There is also a "membership test", but we will not need it. As for the question how to get non-zero elements in the first place, this depends on the application. If k is finite, we can generate random elements to play with.

Algorithm 2.1 (Zero test). *Let $x \in \operatorname{Pic}^0 X$. Given the k-algebra $S_X^{(2)}$ and a subspace of the form $\Gamma(X, \mathcal{L}^2(-D))$ of $\Gamma(X, \mathcal{L}^2)$ representing x, this algorithm decides whether $x = 0$.*

1. *Compute the space*

$$\Gamma(X, \mathcal{L}(-D)) = \bigl\{ s \in \Gamma(X, \mathcal{L}) \bigm| s\Gamma(X, \mathcal{L}) \subset \Gamma(X, \mathcal{L}^2(-D)) \bigr\}.$$

(The equality holds by Lemma 1.2.)

2. *Output 'false' if $\Gamma(X, \mathcal{L}(-D)) = 0$, and 'true' if $\Gamma(X, \mathcal{L}(-D))$ is of dimension 1.*

Algorithm 2.2 (Zero element). *Given the k-algebra $S_X^{(2)}$, this algorithm outputs a subspace $\Gamma(X, \mathcal{L}^2(-D))$ of $\Gamma(X, \mathcal{L}^2))$ representing $0 \in \operatorname{Pic} X$.*

1. *Choose any non-zero element $s \in \Gamma(X, \mathcal{L})$, and let D denote its divisor.*
2. *Output the space $\Gamma(X, \mathcal{L}^2(-D)) = s\Gamma(X, \mathcal{L})$.*

Algorithm 2.3 (Addflip). *Let $x, y \in \operatorname{Pic}^0 X$. Given the k-algebra $S_X^{(5)}$ and subspaces of the form $\Gamma(X, \mathcal{L}^2(-D))$ and $\Gamma(X, \mathcal{L}^2(-E))$ of $\Gamma(X, \mathcal{L}^2)$ representing x and y, this algorithm outputs a subspace $\Gamma(X, \mathcal{L}^2(-F))$ representing $-x - y$.*

1. *Compute the space $\Gamma(X, \mathcal{L}^4(-D - E))$ as the product of $\Gamma(X, \mathcal{L}^2(-D))$ and $\Gamma(X, \mathcal{L}^2(-E))$. (This is possible by Lemma 1.1.)*
2. *Compute the space*

$$\Gamma(X, \mathcal{L}^3(-D - E)) = \bigl\{ s \in \Gamma(X, \mathcal{L}^3) \bigm| s\Gamma(X, \mathcal{L}) \subset \Gamma(X, \mathcal{L}^4(-D - E)) . \bigr\}$$

3. *Choose any non-zero $s \in \Gamma(X, \mathcal{L}^3(-D - E))$, and write $F = \operatorname{div} s$.*
4. *Compute the space*

$$\Gamma(X, \mathcal{L}^5(-D - E - F)) = s\Gamma(X, \mathcal{L}^2).$$

5. *Output the space*

$$\Gamma(X, \mathcal{L}^2(-F)) = \bigl\{ s \in \Gamma(X, \mathcal{L}^2) \bigm| \\ s\Gamma(X, \mathcal{L}^3(-D - E)) \subset \Gamma(X, \mathcal{L}^5(-D - E - F)) \bigr\}.$$

3. Functoriality for finite morphisms

Consider a finite morphism

$$f\colon X \to Y$$

between curves. Such a morphism gives rise to group homomorphisms

$$f^*\colon \operatorname{Div} Y \to \operatorname{Div} X \quad \text{and} \quad f_*\colon \operatorname{Div} X \to \operatorname{Div} Y.$$

These maps are defined on prime divisors as follows. If Q is a prime divisor on X, we have

$$f^*Q = \sum_{P\colon f(P)=Q} e(P)P,$$

where P runs over the prime divisors on X lying over Q and $e(P)$ is the ramification index at P. If P is a prime divisor on X, we have

$$f_*P = [k(P) : k(f(P))]f(P),$$

where $k(P)$ and $k(f(P))$ are the residue fields. The maps f^* and f_* induce two morphisms

$$\operatorname{Pic} f\colon \operatorname{Pic}^0 Y \to \operatorname{Pic}^0 X \quad \text{and} \quad \operatorname{Alb} f\colon \operatorname{Pic}^0 X \to \operatorname{Pic}^0 Y,$$

called the Picard and Albanese maps.

Suppose X and Y are given as above by projective embeddings via line bundles $\mathcal{L}_X$ and $\mathcal{L}_Y$, and the morphism f is induced by a graded k-algebra homomorphism

$$f^\#\colon S_Y \to S_X$$

between the projective coordinate rings of X and Y. In this case we have a canonical isomorphism

$$f^*\mathcal{L}_Y \cong \mathcal{L}_X$$

and induced injective maps

$$\Gamma(X, \mathcal{L}_X^i) \to \Gamma(Y, \mathcal{L}_Y^i).$$

Suppose the k-algebras $S_X^{(4)}$ and $S_Y^{(4)}$ are given together with the homomorphism $f^\#$. Then we can compute

$$\operatorname{Pic} f\colon \operatorname{Pic}^0 Y \to \operatorname{Pic}^0 X.$$

If we know the above spaces for $i \le 6$ and we know in addition a rational point of X, *and* we can factor polynomials over k, then we can also compute

$$\operatorname{Alb} f\colon \operatorname{Pic}^0 X \to \operatorname{Pic}^0 Y.$$

Let us briefly explain the algorithm for computing the Picard map; the algorithm for computing $\operatorname{Alb} f$ is more complicated. Let y be an element of $\operatorname{Pic}^0 Y$, represented by the space $\Gamma(Y, \mathcal{L}_Y^2(-E)$ for some effective divisor E with $\deg E = \deg \mathcal{L}_Y$. We first compute the image of $\Gamma(Y, \mathcal{L}_Y^2(-E))$ under the inclusion

$$\Gamma(Y, \mathcal{L}_Y^2) \longrightarrow \Gamma(X, \mathcal{L}_X^2).$$

This is a basepoint-free subspace W of $\Gamma(X, \mathcal{L}_X^2(-f^*E))$. We then reconstruct the full space $\Gamma(X, \mathcal{L}_X^2(-f^*E))$ by first computing $\Gamma(X, \mathcal{L}_X^4(-f^*E))$ as the product of W and $\Gamma(X, \mathcal{L}^2)$ (we don't prove why this works) and then computing $\Gamma(X, \mathcal{L}_X^2(-f^*E))$ by

$$\Gamma(X, \mathcal{L}_X^2(-f^*E)) = \bigl\{s \in \Gamma(X, \mathcal{L}_X^2) \bigm| \\ s\Gamma(X, \mathcal{L}_X^2) \subseteq \Gamma(X, \mathcal{L}_X^4(-f^*E))\bigr\}.$$

4. Curves over finite fields

From now on we now restrict ourselves to the case where k is a finite field with q elements. We write

$$J = \operatorname{Pic}^0_{X/k}$$

for the Jacobian variety of X over k; we have an isomorphism

$$\operatorname{Pic}^0 X \xrightarrow{\sim} J(k)$$

(This follows for example from the fact that the Brauer group of k is trivial.)

The zeta function of X is the power series

$$\begin{aligned} \mathrm{Z}_X &= \sum_{n=1}^{\infty} (\# \operatorname{Eff}^n X) t^n \\ &= \prod_{d=1}^{\infty} \frac{1}{(1-t^d)^{\# \operatorname{PDiv}^d X}} \\ &= \frac{L_X}{(1-t)(1-qt)}. \end{aligned}$$

Here $\operatorname{Eff}^n X$ denotes the set of effective divisors of degree n, and $\operatorname{PDiv}^d X$ denotes the set of prime divisors of degree d.

There are several operations that are particular to the case of finite fields:

- the Frobenius map on divisors on X and the Frobenius endomorphism of J;
- choosing uniformly random divisors on X and elements of $J(k)$;
- computing Frey–Rück pairings on $J(k)$;
- computing Kummer morphisms on $J(k)$.

Using all of the above, it is also possible to choose uniformly random l-torsion points of $J(k)$, where l is a prime number different from the characteristic of k. This technique was described by Couveignes [1] in the situation where the curve is represented using a non-constant morphism to the plane, together with data describing the various branches over the singularities of the image.

5. Choosing random divisors

We will now describe how to choose uniformly random divisors on X. The first step is to choose random prime divisors of a given degree n. This is done by intersecting X with a hypersurface of sufficiently large degree and decomposing the resulting divisor into prime divisors. With probability proportional to the number of distinct prime divisors of degree n occurring in this decomposition, we choose one of them.

Let $\mathcal{S}^a$ denote the set of hypersurface sections of degree a on X. If D is an element of $\mathcal{S}^a$ (or more generally an effective divisor), we write $\operatorname{Irr}^n D$ for the set of prime divisors of degree n in the support of D. Then we have a diagram

$$\begin{tikzcd} \displaystyle\coprod_{D \in \mathcal{S}^a} \operatorname{Irr}^n D \arrow[r] \arrow[d] & \operatorname{PDiv}^n X \\ \mathcal{S}^a, & \end{tikzcd}$$

where the horizontal arrow sends (D, P) to P and the vertical arrow sends (D, P) to D. If the inequality

$$ad - n \geq 2g - 1$$

is satisfied, then all fibres of the horizontal map has the same cardinality, since the fibre above P is in bijection with the set of lines in $\Gamma(X, \mathcal{L}^a(-P))$, and this space has dimension $1 - g + ad - n$ for all P by Riemann–Roch.

Algorithm 5.1 (Choosing random prime divisors). *Let n and a be positive integers such that*

$$ad - n \geq 2g + 1$$

Given the k-algebra $S_X^{(2a+2)}$, this algorithm outputs a uniformly random prime divisor of degree n on X as a subspace of $\Gamma(X, \mathcal{L}^a)$.

1. *Choose a uniformly random non-zero element $s \in \Gamma(X, \mathcal{L}^a)$, and let D denote the divisor of s. (This D is a uniformly random hypersurface section of degree a of X.)*
2. *Compute the set $\operatorname{Irr}^n D$ of (reduced) irreducible components of D of degree n over k. (We do not explain how this works; we need to be able to factor polynomials over k.)*
3. *With probability $\frac{\# \operatorname{Irr}^n D}{\lfloor ad/n \rfloor}$, output a uniformly random element $P \in \operatorname{Irr}^n D$ and stop.*
4. *Go to step 1.*

We can now describe how to choose uniformly random effective divisors of a given degree n. For this we need the zeta function of X, which was introduced above. We sketch a method that can be found in C. Diem's Habilitationsschrift [2]. If D is an effective divisor of degree n, its *decomposition type* is the vector of integers $(a_1, \ldots, a_n)$, where a_i is the number of prime divisors of degree i occurring in D (counted with multiplicities). Given the zeta function of X, we know the distribution of $(a_1, \ldots, a_n)$. We can generate a random decomposition type according to this distribution, and then generate a uniformly random element from the set of effective divisors with this decomposition type.

This algorithm can be used to choose uniformly random elements of $J(k) = \operatorname{Pic}^0 X$. Since all fibres of the map

$$\begin{aligned} \operatorname{Div}^d X &\to \operatorname{Pic}^0 X \\ D &\mapsto [\mathcal{L}(-D)] \end{aligned}$$

are isomorphic, we can choose a uniformly random element of $\operatorname{Pic}^0 X$ by choosing a random divisor D of degree d and interpreting it as the isomorphism class of $\mathcal{L}(-D)$.

6. The Kummer map and the Frey–Rück pairing

Let n be a positive integer, and suppose all the n-torsion points of J are defined over k. Taking Galois cohomology of the short exact sequence

$$0 \longrightarrow J[n] \longrightarrow J \xrightarrow{n} J \longrightarrow 0$$

and using that

$$\mathrm{H}^1(k, J[n]) \cong J[n](k),$$

we obtain the *Kummer isomorphism*

$$\begin{aligned} J(k)/nJ(k) &\xrightarrow{\sim} J[n](k) \\ x &\longmapsto \mathrm{F}_q(y) - y \quad \text{where } x = ny. \end{aligned}$$

A similar cohomological construction, but now in the case where k contains a primitive root of unity (i.e. $n \mid q - 1$), gives a non-degenerate pairing

$$\begin{aligned} J[n](k) \times J(k)/nJ(k) &\longrightarrow \mu_n(k) \\ (x, y) &\longmapsto [x, y]_n \end{aligned}$$

called the *Frey–Rück pairing*. One possible definition is the following. Let $x \in J[n](k)$ be represented by a line bundle $\mathcal{M}$ of degree 0 such that $\mathcal{M}^n$ is trivial, and let $y \in J(k)$ be represented by a divisor $E = E^+ - E^-$ of degree 0. We choose trivialisations

$$s \colon \mathcal{M}^n \xrightarrow{\sim} \mathcal{O}_X$$

and

$$t^{\pm} \colon k \xrightarrow{\sim} \mathrm{N}_{E^{\pm}/k}\mathcal{M}.$$

Then the composed isomorphism

$$k \xrightarrow[\sim]{(t^+)^n} (\mathrm{N}_{E^+/k}\mathcal{M})^n \cong \mathrm{N}_{E^+/k}(\mathcal{M}^n) \xrightarrow[\sim]{\mathrm{N}_{E^+/k}s} k \xrightarrow[\sim]{\mathrm{N}_{E^-/k}s^{-1}} \mathrm{N}_{E^-/k}(\mathcal{M}^n) \cong (\mathrm{N}_{E^-/k}\mathcal{M})^n \xrightarrow[\sim]{(t^-)^{-n}} k$$

equals multiplication by $[x, y]_n$.

Under the assumption that we know the zeta function of X, the Kummer isomorphism and the Frey–Rück pairing can be computed, and they can be used to generate uniformly random points of $J[l]$, where l is any prime number different from the characteristic of k.

References

[1] J.-M. Couveignes, Linearizing torsion classes in the Picard group of algebraic curves over finite fields. *Journal of Algebra* **321** (2009), 2085–2118.

[2] C. Diem, *On arithmetic and the discrete logarithm problem in class groups of curves*. Habilitationsschrift, Universität Leipzig, 2008.

[3] K. Khuri-Makdisi, Linear algebra algorithms for divisors on an algebraic curve. *Mathematics of Computation* **73** (2004), no. 245, 333–357.
Available online: `http://arxiv.org/abs/math.NT/0105182`.

[4] K. Khuri-Makdisi, Asymptotically fast group operations on Jacobians of general curves. *Mathematics of Computation* **76** (2007), no. 260, 2213–2239.
Available online: `http://arxiv.org/abs/math.NT/0409209`.