Sub-quadratic time for Riemann-Roch spaces. The case of smooth divisors over nodal plane projective curves

Simon Abelard, Alain Couvreur, Grégoire Lecerf

► To cite this version:

Simon Abelard, Alain Couvreur, Grégoire Lecerf. Sub-quadratic time for Riemann-Roch spaces. The case of smooth divisors over nodal plane projective curves. ISSAC 2020 - 45th International Symposium on Symbolic and Algebraic Computation, Jul 2020, Kalamata, Greece. pp.14-21, 10.1145/3373207.3404053 . hal-02477371

HAL Id: hal-02477371

<https://hal.inria.fr/hal-02477371>

Submitted on 13 Feb 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sub-quadratic time for Riemann–Roch spaces

Case of smooth divisors over nodal plane projective curves

Simon Abelard^{1,3}, Alain Couvreur^{2,1,4}, and Grégoire Lecerf^{1,5}

¹Laboratoire d’informatique de l’École polytechnique (LIX, UMR 7161)

CNRS, École polytechnique, Institut Polytechnique de Paris
1, rue Honoré d’Estienne d’Orves
Bâtiment Alan Turing, CS35003
91120 Palaiseau, France

²Inria

³simon.abelard@lix.polytechnique.fr

⁴alain.couvreur@inria.fr

⁵gregoire.lecerf@lix.polytechnique.fr

Preliminary version, February 12, 2020

Abstract

We revisit the seminal Brill–Noether algorithm in the rather generic situation of smooth divisors over a nodal plane projective curve. Our approach takes advantage of fast algorithms for polynomials and structured matrices. We reach sub-quadratic time for computing a basis of a Riemann–Roch space. This improves upon previously known complexity bounds.

1 Introduction

Let $\mathbb{K}$ be an effective field and let $\bar{\mathbb{K}}$ denote an algebraic closure of $\mathbb{K}$. Here “effective” means that we can perform arithmetic operations and zero tests in $\mathbb{K}$. The projective space of dimension 2 over $\bar{\mathbb{K}}$ is written $\mathbb{P}^2$. The input projective curve $\mathcal{C}$ in $\mathbb{P}^2$ is given by its defining equation $Q(X, Y, Z) = 0$, where $Q \in \mathbb{K}[X, Y, Z]$ is homogeneous of total degree $\delta \geq 1$. This paper modifies the variant of the Brill–Noether algorithm proposed in [19] so as to reach sharper complexity bounds.

1.1 Hypotheses

Until the end of the paper, $\mathbb{K}$ is a sufficiently large field with the following restriction:
 $\mathbb{K}$ -H $\mathbb{K}$ is either finite or has characteristic zero, and is therefore a perfect field.

We will assume that the following hypotheses hold for $\mathcal{C}$:

$\mathcal{C}$ -H₁ Q is absolutely irreducible, that is irreducible over $\bar{\mathbb{K}}$;

$\mathcal{C}$ -H₂ $\mathcal{C}$ is nodal: each germ of curve at a singular point splits into two smooth germs with distinct tangent spaces. The number of singular points is written r , and the *nodal divisor*, written E is the symbolic sum of the singular points.

Let us recall that absolute irreducibility can be tested efficiently by means of the algorithms designed in [1]. For the second hypothesis it suffices to check that the Hessian of Q is non-degenerate at each singular point. The restriction on the type of singularities involves simplifications in the Brill–Noether algorithm [17, 7]: basically the desingularization of $\mathcal{C}$ is immediate, and the adjoint divisor simply writes from the singular locus. The last hypothesis necessary to our algorithm concerns the input divisor D , for which we want a basis of the Riemann–Roch space, written $\mathcal{L}(D)$: **D-H** The input divisor D is *smooth* and defined over $\mathbb{K}$, which means that its support is made of regular points of $\mathcal{C}$.

We will decompose a divisor D into $D = D_+ - D_-$, where D_+ and D_- are *positive* (also called *effective*) divisors. When $\deg D_+ < \deg D_-$, $\mathcal{L}(D)$ is (0) so we can freely assume that $\deg D_+ \geq \deg D_-$. The above hypotheses are essentially present in [19]: $\mathbb{K}$ -H is slightly more restrictive in order to simplify complexity analyses.

1.2 Notation

For complexity analyses we focus on an algebraic model over $\mathbb{K}$ (typically computation trees), so we count the number of arithmetic operations and zero tests performed by the algorithms. Over finite fields, we use Turing machines with sufficiently many but finite number of tapes. In order to simplify the presentation of complexity bounds, we use the *soft-Oh* notation: $f(n) \in \tilde{O}(g(n))$ means that $f(n) = g(n) \log_2^{O(1)}(g(n) + 3)$; see [5, chapter 25, section 7]. The vector space of polynomials of degree $< n$ in $\mathbb{K}[X]$ is written $\mathbb{K}[X]_{<n}$. For integer and polynomial arithmetic we content ourselves with softly linear cost bounds [5].

The constant ω denotes a real value between 2 and 3 such that two $n \times n$ matrices over a commutative ring can be multiplied with $O(n^\omega)$ operations; $\omega < 2.3728639$ [18]. The constant ϖ is an other real value between 1.5 and $(\omega + 1)/2$ such that the product of a $n \times \sqrt{n}$ matrix by a $\sqrt{n} \times \sqrt{n}$ matrix takes $O(n^\varpi)$ operations; $\varpi < 1.667$ [15, Theorem 10.1].

Given $M \in \text{GL}_3(\mathbb{K})$ and $P \in \mathbb{K}[X, Y, Z]$ we denote by $(P \circ M)(X, Y, Z)$ the polynomial $P(M \cdot (X, Y, Z)^\top)$.

1.3 Our contributions

The present paper is essentially based on the variant of the Brill–Noether algorithm designed in [19]. Our first result is the improvement of complexity bounds for the arithmetic of smooth divisors. A second contribution concerns the opportune use of structured linear algebra algorithms: we reformulate the Riemann–Roch problem in terms of modules of relations of rank $\leq \delta$, and compute bases thanks to the recent fast algorithm due to Neiger [23].

We represent the Riemann–Roch space

$$\mathcal{L}(D) := \{h \in \mathbb{K}(\mathcal{C}) : (h) \geq -D\} \cup \{0\}$$

by a *basis generator*, that is made of $M \in \text{GL}_3(\mathbb{K})$, an integer $l \leq \delta$, non-zero homogeneous polynomials $H, G_1, \dots, G_l$ in $\mathbb{K}[X, Y, Z]$ of respective total degrees d and $d_i \leq d$, such that:

- $\deg_Y(Q \circ M) = \delta$, $\deg_Y H \leq \delta - 1$, $\deg_Y G_i \leq \delta - 1$ for $i = 1, \dots, l$.
- The supports of $M^{-1}(D_+)$, $M^{-1}(D_-)$, $M^{-1}(E)$, and the solutions of $Q \circ M = H = 0$ are in the affine chart $Z = 1$.
- $\left(\frac{X^j G_i}{H}\right) \circ M^{-1}$ with $0 \leq j \leq d - d_i$ and $1 \leq i \leq l$ form a basis of $\mathcal{L}(D)$.

The actual vector-space basis of $\mathcal{L}(D)$ can be recovered in softly linear time from the basis generators, the latter being a more compact representation.

Theorem 1. *Under hypotheses $\mathbb{K}$ -H, $\mathcal{C}$ -H₁, $\mathcal{C}$ -H₂, with d defined below in (8), given a primitive element representation (see section 3.1) of D satisfying D -H, and assuming $|\mathbb{K}| \geq \max(\delta^4, 6(\delta d)^2)$, a basis generator of $\mathcal{L}(D)$ can be computed by a probabilistic algorithm of type Las Vegas with an expected number of $\tilde{O}\left((\delta^2 + \deg D_+)^{\frac{\omega+1}{2}}\right)$ field operations in characteristic zero or $> \max(\delta(\delta - 1), \delta d)$, or $\tilde{O}\left((\delta^2 + \deg D_+)^{\frac{\omega+1}{2}} \log q + (\delta^2 + \deg D_+) \log^2 q\right)$ bit operations if $\mathbb{K} = \mathbb{F}_q$.*

Our third contribution, central to this theorem, is a sharp degree bound d for H and the G_i ; namely (8). Such a bound is not supplied in [19] when $r > 0$, so when $\mathcal{C}$ is not smooth additional assumptions are required in [19, section 2].

1.4 Related work

Riemann–Roch spaces have various applications in applied algebra, number theory and cryptography (*e.g.* arithmetic in Jacobians of curves). Computing bases for these spaces is also pivotal to design geometric codes, where the encoding algorithm consists in evaluating a basis of a certain Riemann–Roch space at points of an algebraic curve. Currently in practice, algebraic curves used in coding theory are mostly limited to cases for which such bases are already known, so for the sake of diversity we aim to handle more general curves and divisors.

Algorithms and implementations for Riemann–Roch spaces have been thoroughly investigated over the past decades. To focus on the most recent contributions, we mention: Hess’ algorithm [9] that is implemented within the computer algebra software MAGMA, and Khuri–Makdisi’s approach [16] that is dedicated to group operations in Jacobians of genus- g curves in time $O(g^{\omega+\epsilon})$, where ϵ can be any positive number. More recently, Le Gluher and Spaenlehauer [19] revisited the Brill–Noether approach for smooth divisors D on a nodal curve, and obtained the complexity bound $O(\max(\delta^2, \deg D_+)^{\omega})$, yet under the aforementioned restriction on D . For conciseness we refer to [19] for further references.

2 Preliminaries

In order to obtain the aforementioned complexity bound for Riemann–Roch spaces, we rely on structured linear algebra algorithms that will be presented in section 4.1, and on modular composition and elimination, that are the purposes of this section.

2.1 Bivariate modular composition

At present time no algorithm with softly linear time is known for bivariate modular compositions over a general field $\mathbb{K}$. For practical purposes we appeal to a variant of the Paterson–Stockmeyer evaluation scheme designed by Nüsken and Ziegler [25]. We need a slight extension to express the complexity bound in terms of the degree of the modulus.

Algorithm 1

Input: $P \in \mathbb{K}[X, Y]$ of total degree n , $\chi \in \mathbb{K}[Y]$, $u \in \mathbb{K}[Y]_{<\deg \chi}$.

Output: $P(u(Y), Y) \bmod \chi(Y)$.

1. Let $p := \lfloor \sqrt{n} \rfloor$ and $q := \lceil n/p \rceil$.
2. For $i = 0, \dots, p-1$ do:
 - (a) Compute $u^i \bmod \chi$ and segment it into $M_{0,i}(Y) + M_{1,i}(Y)Y^n + \dots + M_{l-1,i}(Y)Y^{(l-1)n}$, with $\deg M_{j,i}(Y) < n$ and $l := \lceil \deg \chi / n \rceil$. This yields an $l \times p$ matrix $M \in \mathbb{K}[Y]^{l \times p}$.

- (b) For $j = 0, \dots, q-1$, let $N_{i,j}(Y) := P_{i+pj}^X(Y)$, where P_{i+pj}^X represents the coefficient of degree $i + pj$ of P regarded in $\mathbb{K}[Y][X]$. This yields a $p \times q$ matrix $N \in \mathbb{K}[Y]^{p \times q}$ of degree $\leq n$.
3. Compute the matrix product $R := MN$.
4. For $j = 0, \dots, q-1$, let $v_j(Y) := R_{0,j}(Y) + R_{1,j}(Y)Y^n + \dots + R_{l-1,j}(Y)Y^{(l-1)n} \bmod \chi(Y)$.
5. Return $\sum_{j=0}^{q-1} v_j u^{pj} \bmod \chi$.

Lemma 2. *Algorithm 1 is correct and takes $\tilde{O}\left(n^{\frac{\omega-1}{2}} \left(\deg \chi + n^{\frac{3}{2}}\right)\right)$ operations in $\mathbb{K}$.*

Proof. By construction, we have $v_j = P_{pj}^X + P_{1+pj}^X u + \dots + P_{p-1+pj}^X u^{p-1}$ for $j = 0, \dots, q-1$, whence

$$P(u(Y), Y) = \sum_{j=0}^{q-1} v_j u^{pj} \bmod \chi(Y).$$

This proves the correctness of the algorithm. Step 2a requires $O(p) = O(\sqrt{n})$ multiplications modulo χ . Step 3 costs

$$\tilde{O}\left(n^{\frac{\omega}{2}+1} \left\lceil \frac{l}{p} \right\rceil\right) = \tilde{O}\left(n^{\frac{\omega}{2}+1} \left(\frac{\frac{\deg \chi}{n} + 1}{n^{\frac{1}{2}}} + 1\right)\right) = \tilde{O}\left(n^{\frac{\omega-1}{2}} \left(\deg \chi + n^{\frac{3}{2}}\right)\right).$$

Step 5 involves $O(q) = O(\sqrt{n})$ multiplications and additions modulo χ , using Horner's method. $\square$

2.2 Primitive element representation

A *primitive element representation* of a set $\mathcal{E}$ of points in the affine plane $\mathbb{A}^2$ is the data of:

- (λ, μ) in $\bar{\mathbb{K}}^2$ such that the linear form $\lambda X + \mu Y$ *separates* the points in $\mathcal{E}$. This means that the form takes different values at different points of $\mathcal{E}$.
- A polynomial θ in $\bar{\mathbb{K}}[S]$ whose roots are the values of $\lambda X + \mu Y$ at the points of $\mathcal{E}$, that is

$$\theta(S) := \prod_{(x,y) \in \mathcal{E}} (S - (\lambda x + \mu y)).$$

So θ is monic and separable of degree $|\mathcal{E}|$.

- Polynomials u and v in $\bar{\mathbb{K}}[S]$ of degrees $< |\mathcal{E}|$ such that

$$\mathcal{E} = \{(u(\zeta), v(\zeta)) : \theta(\zeta) = 0\}.$$

Notice that such a representation is uniquely determined by (λ, μ) . If $(\lambda, \mu) \in \mathbb{K}^2$ and if $\theta, u, v \in \mathbb{K}[S]$, then the primitive element representation is said to be *defined over $\mathbb{K}$* .

If the annihilator ideal of $\mathcal{E}$ is generated by polynomials with coefficients in $\mathbb{K}$, then a primitive element representation does not necessarily exist over $\mathbb{K}$. However any value of λ/μ outside

$$\left\{ \frac{y_1 - y_2}{x_1 - x_2} : (x_1, y_1) \in \mathcal{E}, (x_2, y_2) \in \mathcal{E}, (x_1, y_1) \neq (x_2, y_2), x_1 \neq x_2 \right\}, \quad (1)$$

yields a primitive element. Therefore, a sufficient condition to ensure that such a primitive element exists is $|\mathbb{K}| > \binom{|\mathcal{E}|}{2}$. Otherwise, λ/μ needs to be taken in an algebraic extension of $\mathbb{K}$. We will not discuss these usual technical details but will make precise the conditions on the cardinality of $\mathbb{K}$ within each sub-algorithm. For the sake of complexity it will be convenient to change the variables X and Y linearly, so we recall the following lemma.

Lemma 3. (For instance [12, Proposition 9]) *If $F \in \mathbb{K}[X, Y, Z]$ is homogeneous of degree n , if $|\mathbb{K}| \geq n + 1$, and if $M \in \text{GL}_3(\mathbb{K})$, then we can compute $F \circ M$ with $\tilde{O}(n^2)$ operations in $\mathbb{K}$.*

2.3 Change of primitive element

Changing primitive elements mostly reduces to computing characteristic polynomials in $\mathbb{K}[S]/(\theta(S))$. This task has received a lot of attention in computer algebra, but so far no general algorithm is known with nearly linear time; for instance see [6] about the existing literature. For our present purposes it seems reasonable to appeal to the known complexity exponent ϖ : recall that univariate modular composition in degree n takes $O(n^\varpi)$ field operations.

Lemma 4. *Given a primitive element representation of $\mathcal{E}$ over $\mathbb{K}$ by $\lambda X + \mu Y$, and given $(\tilde{\lambda}, \tilde{\mu}) \in \mathbb{K}^2$, we can test if $\tilde{\lambda}X + \tilde{\mu}Y$ is primitive for $\mathcal{E}$, and, if so, compute the corresponding representation of $\mathcal{E}$, along with $w(S) \in \mathbb{K}[S]_{<|\mathcal{E}|}$ such that*

$$\begin{aligned} \mathbb{K}[S]/(\theta(S)) &\cong \mathbb{K}[S]/(\tilde{\theta}(S)) \\ S &\mapsto w(S) \\ \tilde{\lambda}u(S) + \tilde{\mu}v(S) &\mapsto S, \end{aligned}$$

is an isomorphism, with $O(|\mathcal{E}|^\varpi)$ field operations in characteristic zero or $> |\mathcal{E}|$, or $|\mathcal{E}|^\varpi \tilde{O}(\log q) + \tilde{O}(|\mathcal{E}| \log^2 q)$ bit operations if $\mathbb{K} = \mathbb{F}_q$.

Proof. Let Tr denote the trace map of $\mathbb{K}[S]/(\theta(S))$ and let $\tilde{\theta}$ be the characteristic polynomial of the multiplication endomorphism by $\tilde{\lambda}u(S) + \tilde{\mu}v(S)$ in this algebra. Le Verrier's method consists in computing $\text{Tr}((\tilde{\lambda}u + \tilde{\mu}v)^i)$ for $i = 1, \dots, |\mathcal{E}| - 1$. This task being dual to modular composition, it takes $O(|\mathcal{E}|^\varpi)$ operations in $\mathbb{K}$. Then the generating series $\tau(z) := \sum_{i \geq 0} \text{Tr}((\tilde{\lambda}u + \tilde{\mu}v)^i) z^i$ satisfies the Newton–Girard formula

$$-\frac{\nu'(z)}{\nu(z)} = \tau(z) + O(z^{|\mathcal{E}|}), \quad (2)$$

where $\nu(z) := z^{|\mathcal{E}|} \tilde{\theta}(1/z)$ is the reciprocal of $\tilde{\theta}$. Therefore ν is recovered with $\tilde{O}(|\mathcal{E}|)$ operations in characteristic zero or $> |\mathcal{E}|$. Testing if $\tilde{\lambda}X + \tilde{\mu}Y$ is primitive is equivalent to testing if $\tilde{\theta}$ is squarefree, that takes $\tilde{O}(|\mathcal{E}|)$ field operations in characteristic zero or $> |\mathcal{E}|$, or $\tilde{O}(|\mathcal{E}| \log^2 q)$ bit operations if $\mathbb{K} = \mathbb{F}_q$; for instance see [20]. By a deformation argument we further recover w up to a constant cost factor; see [6, section 2.6].

In positive characteristic, the integration of (2) is more tedious in general, but in the special case $\mathbb{K} = \mathbb{F}_q$, it is possible with $\tilde{O}(|\mathcal{E}| \log q)$ bit operations; see [6, Proposition 3]. $\square$

2.4 Curve intersection

For computing principal divisors on curves we will appeal to the following lemma, based on polynomial resultants. The technique is known so the proof is voluntarily concise. Details can be found in [3, section 4] or [12, section 5].

Lemma 5. *Given F of total degree m and G of total degree $\leq n$ in $\mathbb{K}[X, Y]$ such that $m \leq n$, F has degree m in Y , and the solutions of $F = G = 0$ is a finite set $\mathcal{E}$. We can check if X is primitive for $\mathcal{E}$, and compute a partition of $\mathcal{E} =: \mathcal{E}_1 \cup \dots \cup \mathcal{E}_s$, where $\mathcal{E}_i$ contains points with the same known intersection multiplicity m_i , with $\tilde{O}(nm^2 + n \deg_Y G)$ field operations in characteristic zero or $> mn$, or $\tilde{O}(((mn)^\varpi + n \deg_Y G) \log q + mn \log^2 q)$ bit operations if $\mathbb{K} = \mathbb{F}_q$.*

Proof. The remainder H of G by F regarded in $\mathbb{K}[X][Y]$ can be computed with $\tilde{O}(n \deg_Y G)$ operations in $\mathbb{K}$, so we obtain $\chi(X) := \text{Res}_Y(F(X, Y), H(X, Y))$ with cost $\tilde{O}(nm^2)$ by [21, Corollary 31]. Since χ has degree $\leq mn$, the squarefree decomposition $\chi =: \theta_1^{m_1} \dots \theta_s^{m_s}$ contributes to $\tilde{O}(mn)$ field operations in characteristic zero or $> mn$, or to $\tilde{O}(mn \log^2 q)$ bit operations over $\mathbb{F}_q$.

After fast multi-remaindering of F and H by $\theta_1, \dots, \theta_s$ [5, chapter 10], the directed evaluation paradigm [11, 2] yields a decomposition $\theta_i =: \theta_{i,1} \cdots \theta_{i,s_i}$, and bivariate polynomials $Q_{i,j}(X, Y)$ such that

$$Q_{i,j}(\zeta, Y) = \gcd(F(\zeta, Y), H(\zeta, Y)) \quad (3)$$

with $\deg_X Q_{i,j} < \deg \theta_{i,j}$, for all $\theta_{i,j}(\zeta) = 0$, $j = 1, \dots, s_i$, $i = 1, \dots, s$. It turns out that X is a primitive element of $\mathcal{E}$ if, and only if, each $Q_{i,j}(\zeta, Y)$ is a power of a degree 1 polynomial $Y - v_{i,j}(\zeta)$ with $\deg v_{i,j} < \deg \theta_{i,j}$. In this case, the representation $\theta_i(X) = Y - v_j(X) = 0$ of $\mathcal{E}_i$ is deduced in softly linear time by Chinese remaindering; details will be given in Lemma 10 below for a slightly more general situation. This takes $\tilde{O}(nm^2)$ operations in $\mathbb{K}$ when the characteristic p is zero or $> mn$. Otherwise when $p > 0$, the gcd (3) is required to be a power (coprime to p) of $Y^{p^{t_{i,j}}} - w_{i,j}(\zeta)$. We compute $A := X^{p^{t_{i,j}}} \bmod \theta_{i,j}(X)$ with bit cost $\tilde{O}(\deg \theta_{i,j} \log m \log q)$ since $p^{t_{i,j}} = O(\log m)$. Computing the characteristic polynomial $\tilde{\theta}_{i,j}$ of A and the expression $X = B(A)$ as in the proof of Lemma 4 involves bit cost $\tilde{O}((\deg \theta_{i,j})^\varpi \log q + \deg \theta_{i,j} \log^2 q)$. By modular composition, we deduce $Y^{p^{t_{i,j}}} - (w_{i,j} \circ B)(\zeta^{p^{t_{i,j}}})$ with $\tilde{O}((\deg \theta_{i,j})^\varpi \log q)$ bit cost. After the extraction of p -th roots, the latter expression finally becomes $(Y - v_{i,j}(\zeta))^{p^{t_{i,j}}}$, with further $\tilde{O}(\deg \theta_{i,j} \log^2 q)$ bit operations.

By [3, Proposition 2.7], $\chi(X)$ is the characteristic polynomial of X in $\mathbb{K}[X, Y]/(F(X, Y), G(X, Y))$, so m_i is the intersection multiplicity of the points represented by $\theta_i(X) = Y - v_i(X) = 0$. $\square$

3 Divisor

This section gathers complexity results for basic operations on smooth divisors of $\mathcal{C}$.

3.1 Primitive element representation

A *smooth positive* divisor D of $\mathcal{C}$ is a multi-set of smooth points of $\mathcal{C}$. The underlying set of points $\mathcal{E} = \{P_1, \dots, P_s\}$ is called the *support* of the divisor, and it is customary to write D as the formal sum $D = m_1 P_1 + \dots + m_s P_s$, where $m_i > 0$ is the *multiplicity* of P_i in D . Up to a linear change of variables we may assume that the support of D is in the affine chart $Z = 1$. In this case, a *primitive element* of D is a linear form $\lambda X + \mu Y$ that separates its support and satisfies the additional conditions:

$$\left| \begin{array}{cc} \frac{\partial Q}{\partial X}(P_i) & \frac{\partial Q}{\partial Y}(P_i) \\ \lambda & \mu \end{array} \right| \neq 0 \text{ for } i = 1, \dots, s, \quad (4)$$

where $Q := Q^h(X, Y, 1)$. Since the P_i are smooth on $\mathcal{C}$, if $\mu \neq 0$, the latter condition is equivalent to requiring that λ/μ is outside the set

$$\left\{ \frac{\frac{\partial Q}{\partial X}(u(\zeta), v(\zeta))}{\frac{\partial Q}{\partial Y}(u(\zeta), v(\zeta))} : \chi(\zeta) = 0, \frac{\partial Q}{\partial Y}(u(\zeta), v(\zeta)) \neq 0 \right\}. \quad (5)$$

If λ/μ is outside the sets (1) and (5) then it is primitive for D . The sum of the cardinalities of these two sets is $\leq \binom{\deg D + 1}{2}$, where $\deg D := m_1 + \dots + m_s$. Consequently, as soon as $|\mathbb{K}| > \binom{\deg D + 1}{2}$, primitive elements can be found in $\mathbb{K}$.

Proposition 6. *Given a smooth positive divisor $D = m_1 P_1 + \dots + m_s P_s$ whose support is in the affine chart $Z = 1$, and given a primitive element $\lambda X + \mu Y$ for D , there exist unique polynomials χ , u , and v in $\bar{\mathbb{K}}[S]$ with the following properties:*

Div-H₀ χ is monic of degree $\deg D$, and u , v have degrees $< \deg D$,

Div-H₁ $Q(u(S), v(S)) = 0 \bmod \chi(S)$,

Div-H₂ $\lambda u(S) + \mu v(S) = S$,

Div-H₃ $\mu \frac{\partial Q}{\partial X}(u(S), v(S)) - \lambda \frac{\partial Q}{\partial Y}(u(S), v(S))$ is coprime to $\chi(S)$.

Proof. We write χ_0 , u_0 , and v_0 for the primitive element representation of the support $\mathcal{E}$, so we have $Q(u_0(S), v_0(S)) = 0 \bmod \chi_0(S)$ and $\lambda u_0(S) + \mu v_0(S) = S$. The hypothesis (4) means that any point (x, y) of the divisor is a regular root of the map

$$\Xi: \begin{pmatrix} X \\ Y \end{pmatrix} \mapsto \begin{pmatrix} Q(X, Y) \\ \lambda X + \mu Y \end{pmatrix}.$$

For $n \geq 0$ we appeal to the following Newton iteration based on Ξ :

$$\begin{pmatrix} u_{n+1} \\ v_{n+1} \end{pmatrix} := \begin{pmatrix} u_n \\ v_n \end{pmatrix} - D\Xi(u_n, v_n)^{-1} \Xi(u_n, v_n) \bmod \chi_0^{2^{n+1}}.$$

It follows that (u_n, v_n) is the unique root of Ξ modulo $\chi_0^{2^n}$ that coincides to (u_0, v_0) modulo χ_0 . When 2^n is strictly larger than the largest multiplicity in D , we set

$$\chi(S) := (S - \lambda x_1 - \mu y_1)^{m_1} \cdots (S - \lambda x_r - \mu y_r)^{m_r},$$

where (x_i, y_i) denotes the coordinates of P_i , and then $u := u_n \bmod \chi$ and $v := v_n \bmod \chi$. Since χ divides $\chi_0^{2^n}$ the required properties are satisfied.

The uniqueness follows from the one of the lifted roots of Ξ , since conditions Div-H₀ to Div-H₃ imply that χ_0 , $u \bmod \chi_0$ and $v \bmod \chi_0$ constitute the primitive element representation of $\mathcal{E}$; that means $u_0 = u \bmod \chi_0$ and $v_0 = v \bmod \chi_0$. $\square$

A smooth positive divisor D as above will be represented by λ, μ, χ, u, v along with $\nabla Q(u, v) \bmod \theta$, where θ denotes the squarefree part of χ , and ∇Q represents the *gradient* of Q .

3.2 Lifting a divisor

We analyze the complexity of the Newton iteration seen in the proof of Proposition 6.

Lemma 7. *Let D be a smooth positive divisor parametrized by $\lambda X + \mu Y$. The representation of $2D$ by $\lambda X + \mu Y$ can be obtained with $\tilde{O}\left(\delta^{\frac{\omega}{2}+1} + (\deg D)^{\frac{\omega+2}{3}}\right)$ operations in $\mathbb{K}$.*

Proof. Let χ, u, v represent D , so $\Xi(u(S), v(S)) = 0 \bmod \chi(S)$. We can use the Newton iteration to obtain

$$\begin{pmatrix} \tilde{u}(S) \\ \tilde{v}(S) \end{pmatrix} := \begin{pmatrix} u(S) \\ v(S) \end{pmatrix} - D\Xi(u(S), v(S))^{-1} \cdot \Xi(u(S), v(S)) \bmod \chi(S)^2,$$

that yields $\Xi(\tilde{u}(S), \tilde{v}(S)) = 0 \bmod \chi(S)^2$. The evaluations of Q and of its partial derivatives at $(u(S), v(S))$ modulo $\chi(S)^2$ take

$$\tilde{O}\left(\delta^{\frac{\omega}{2}+1} + (\deg \chi)^{\frac{\omega+2}{3}}\right)$$

operations in $\mathbb{K}$ by Lemma 2. The inverse of the determinant of $D\Xi(u(S), v(S))$ contributes to $\tilde{O}(\deg \chi)$. $\square$

3.3 Nodal divisor

The *nodal divisor* of $\mathcal{C}$, written E , will be given by a primitive element representation $\lambda_E, \mu_E, \chi_E, u_E, v_E$ of the set of singular points of $\mathcal{C}$. In the terminology of the Brill–Noether algorithm, E plays the role of the *adjoint divisor* of $\mathcal{C}$. Since E only depends on $\mathcal{C}$, it might be regarded as a precomputation. Yet for the computation of a single Riemann–Roch space it is fair to take its cost into account. A probabilistic method is summarized in the next proposition. The hypothesis on $|\mathbb{K}|$ is flexible: in fact throughout the paper we have given priority to simple bounds that ensure (conditional) probabilities of success roughly about $1/2$ in the randomized sub-algorithms.

Proposition 8. *Assume $|\mathbb{K}| \geq \delta^4$. Given Q satisfying $\mathcal{C}\text{-}H_1$, we can check if $\mathcal{C}\text{-}H_2$ holds, compute $M \in \text{GL}_3(\mathbb{K})$ such that $Q \circ M$ has degree δ in Y and its singular locus lies in the chart $Z = 1$, and get a primitive element representation of $M^{-1}(E)$, with a probabilistic algorithm of type Las Vegas that takes an expected number of $\tilde{O}(\delta^3)$ field operations in characteristic zero or $> \delta(\delta - 1)$, or $\tilde{O}(\delta^{2\varpi} \log q + \delta^2 \log^2 q)$ bit operations if $\mathbb{K} = \mathbb{F}_q$.*

Proof. By taking α, β at random we easily find values in $\mathbb{K}$ such that $Q(X + \alpha Y, Y, Z + \beta Y)/Q(\alpha, 1, \beta)$ is monic in Y . The running time is $\tilde{O}(\delta^2)$ when using Lemma 3, since it suffices to ensure $Q(\alpha, 1, \beta) \neq 0$, and thanks to the Schwartz–Zippel lemma [5, Lemma 6.44] the expected number of trials is $O(1)$. From now we assume that Q is monic in Y .

The resultant $R(X, Z) := \text{Res}_Y \left(Q, \frac{\partial Q}{\partial Y} \right)$ is homogeneous of degree $\delta(\delta - 1)$. So up to replacing Z by $Z + \gamma X$ in Q , we can further assume that R has degree $\delta(\delta - 1)$ in X with high probability. In particular the solution set $\mathcal{E}$ of $Q = \frac{\partial Q}{\partial Y} = 0$ lies in the chart $Z = 1$. Let $X + \mu Y$ be a candidate primitive element for it. Then, we replace X by $X - \mu Y$ in Q , so X is finally expected to be primitive. Lemma 5 applies to $\frac{\partial Q}{\partial Y}$ and Q : γ and μ are suitable if, and only if, $R(X, 1)$ has degree $\delta(\delta - 1)$, that equals the number of solutions counted with multiplicities. Then we recover a parametrization $\theta(X) = Y - v(X) = 0$ of $\mathcal{E}$ via usual Chinese remaindering. The parametrization of E is deduced from

$$\begin{aligned} \theta_E(X) &:= \theta(X) / \gcd \left(\frac{\partial Q}{\partial X}(X, v(X)), \theta(X) \right) \\ v_E(E) &:= \theta(X) \bmod \theta_E(X). \end{aligned}$$

$\mathcal{C}\text{-}H_2$ holds if and only if the Hessian of Q has full rank at the singular points, that can be checked with further $\tilde{O} \left(\delta^{\frac{\omega+3}{2}} \right)$ operations in $\mathbb{K}$ thanks to Lemma 2. $\square$

3.4 Decomposition of a divisor

For performing arithmetic operations on divisors efficiently we decompose them, operate on components, and recompose them. For a divisor D defined over $\mathbb{K}$ there exists a unique *equal multiplicity decomposition* written $\sum_{i=1}^s m_i D_i$, where:

- the D_i are positive, defined over $\mathbb{K}$, and made of simple points,
- and the m_i are pairwise distinct.

Decompositions and recompositions can be computed fast, as summarized in the following lemmas.

Lemma 9. *The equal multiplicity decomposition of a smooth positive divisor D over $\mathbb{K}$ takes $\tilde{O}(\deg D)$ fields operations in characteristic zero or $> \deg D$, or $\tilde{O}(\deg D \log^2 q)$ bit operations if $\mathbb{K} = \mathbb{F}_q$.*

Proof. Let $\lambda X + \mu Y$, χ , u , v represent D as above. We compute the squarefree factorization of χ into $\theta_1^{m_1} \cdots \theta_s^{m_s}$, with $\tilde{O}(\deg D)$ field operations in characteristic zero or $> \deg D$, and with $\tilde{O}(\deg D \log^2 q)$ bit operations if $\mathbb{K} = \mathbb{F}_q$. So D writes as $m_1 D_1 + \cdots + m_s D_s$, where D_i is parametrized by $\lambda X + \mu Y$, $\chi_i := \theta_i^{m_i}$, $u_i := u \bmod \chi_i$, and $v_i := v \bmod \chi_i$ and $\nabla Q(u_i, v_i) \bmod \theta_i := \nabla Q(u, v) \bmod \theta_i$. Using fast multi-remaindering, this takes $\tilde{O}(\deg D)$ operations in $\mathbb{K}$; see [5, chapter 10]. $\square$

Lemma 10. *Let $D_1, \dots, D_s$ be smooth positive divisors over $\mathbb{K}$, with disjoint supports, and parametrized by the same primitive element $\lambda X + \mu Y$. If $\lambda X + \mu Y$ is primitive for the sum $D := D_1 + \cdots + D_s$, then its representation can be obtained with $\tilde{O}(\deg D)$ operations in $\mathbb{K}$.*

Proof. Let χ_i, u_i, v_i represent D_i , and let θ_i denote the squarefree part of χ_i . By assumption, the χ_i are pairwise coprime. Then $\chi := \chi_1 \cdots \chi_s$ can be computed with $\tilde{O}(\deg D)$ operations in $\mathbb{K}$; see [5, chapter 10]. Since u , v and $\nabla Q(u, v) \bmod \theta$ satisfy $u_i = u \bmod \chi_i$, $v_i = v \bmod \chi_i$, $\nabla Q(u_i, v_i) = \nabla Q(u, v) \bmod \theta_i$ for $i = 1, \dots, s$, they can be obtained via Chinese remaindering with $\tilde{O}(\deg D)$ operations in $\mathbb{K}$. $\square$

3.5 Change of primitive element

Assume that $D := m(P_1 + \cdots + P_s)$, so $\chi = \theta^m$ with θ separable of degree s . Consider

$$\begin{aligned} \Gamma : \quad \mathbb{K}[S]/(\chi(S)) &\cong (\mathbb{K}[Z]/(\theta(Z)))[[T - Z]]/(T - Z)^m \\ S &\mapsto T. \end{aligned} \quad (6)$$

Both directions of this isomorphism can be computed in softly linear time, namely $\tilde{O}(\deg D)$; see [10, section 4.2]. In fact, $(\Gamma(u), \Gamma(v))$ can be regarded as the simultaneous power series expansions of Q at $P_1, \dots, P_s$ with precision m . In order to change the primitive element for D , we first examine what happens to the underlying support, and then change the representations in the power series expansions.

Lemma 11. *Given $D = m(P_1 + \cdots + P_s)$ over $\mathbb{K}$ parametrized by $\lambda X + \mu Y$, and given $(\tilde{\lambda}, \tilde{\mu}) \in \mathbb{K}^2$, we can test if $\tilde{\lambda}X + \tilde{\mu}Y$ is primitive for D , and, if so, compute the corresponding representation, with $O((\deg D)^\varpi)$ field operations in characteristic zero or $> \deg D$, or $(\deg D)^\varpi \tilde{O}(\log q) + \tilde{O}(\deg D \log^2 q)$ bit operations if $\mathbb{K} = \mathbb{F}_q$.*

Proof. First it is checked that

$$\begin{vmatrix} \frac{\partial Q}{\partial X}(u, v) & \frac{\partial Q}{\partial X}(u, v) \\ \tilde{\lambda} & \tilde{\mu} \end{vmatrix}$$

is invertible modulo θ . If so, we change the primitive element for the support of D by means of Lemma 4:

$$\begin{aligned} \Phi : \quad \mathbb{K}[S]/(\theta(S)) &\cong \mathbb{K}[S]/(\tilde{\theta}(S)) \\ S &\mapsto w(S). \end{aligned}$$

That takes $O((\deg \theta)^\varpi)$ operations in $\mathbb{K}$ in characteristic zero or $> \deg \theta$, or $(\deg \theta)^\varpi \tilde{O}(\log q) + \tilde{O}(\deg \theta \log^2 q)$ bit operations if $\mathbb{K} = \mathbb{F}_q$. We convert D to local representation and get the following diagram:

$$\begin{array}{ccc} \Gamma : \quad \mathbb{K}[S]/(\chi(S)) & \rightarrow & (\mathbb{K}[Z]/(\theta(Z)))[[T - Z]]/(T - Z)^m \\ & & \downarrow \text{coefficient-wise extension of } \Phi \\ \tilde{\Gamma} : \quad \mathbb{K}[S]/(\tilde{\chi}(S)) & \rightarrow & (\mathbb{K}[Z]/(\tilde{\theta}(Z)))[[T - Z]]/(T - Z)^m, \end{array}$$

where $\tilde{\chi}(S) := \tilde{\theta}(S)^m$. The parametrization of D in terms of $\tilde{\lambda}X + \tilde{\mu}Y$ is $\tilde{u}(S) := \tilde{\Gamma}^{-1}(\Phi(\Gamma(u(S))))$ and $\tilde{v}(S) := \tilde{\Gamma}^{-1}(\Phi(\Gamma(v(S))))$. That incurs $O(m)$ compositions modulo $\tilde{\theta}$, that is $O(m(\deg \tilde{\theta})^\varpi) = O((\deg D)^\varpi)$. Finally $\nabla Q(\tilde{u}, \tilde{v}) \bmod \tilde{\theta}$ involves two compositions modulo $\tilde{\theta}$. $\square$

Proposition 12. *Given a smooth positive divisor D parametrized by $\lambda X + \mu Y$, and given $(\tilde{\lambda}, \tilde{\mu}) \in \mathbb{K}^2$, we can test if $\tilde{\lambda}X + \tilde{\mu}Y$ is primitive for D , and, if so, compute the corresponding representation with $O((\deg D)^\varpi)$ field operations in characteristic zero or $> \deg D$, or $(\deg D)^\varpi \tilde{O}(\log q) + \tilde{O}(\deg D \log^2 q)$ bit operations if $\mathbb{K} = \mathbb{F}_q$.*

Proof. We compute the equal multiplicity decomposition of $D = m_1 D_1 + \dots + m_s D_s$ as in Lemma 9. For each separable factor D_i of multiplicity m_i , we try to compute the primitive element representation $\tilde{\chi}_i, \tilde{u}_i, \tilde{v}_i$ of the support of D_i for $\tilde{\lambda}X + \tilde{\mu}Y$, by Lemma 11. If it fails then $\tilde{\lambda}X + \tilde{\mu}Y$ cannot be primitive for D . In order to check that $\tilde{\lambda}X + \tilde{\mu}Y$ is finally primitive for D it remains to verify that the squarefree parts of the $\tilde{\chi}_i$ are coprime. Then we may glue the representations of the $m_i D_i$ via Lemma 10. $\square$

3.6 Sum of divisors

Gathering tools presented above we obtain efficient sums and subtractions for divisors.

Proposition 13. *Given two smooth positive divisors D_1 and D_2 such that $|\mathbb{K}| \geq (\deg D_1 + \deg D_2)^2$, the sum $D := D_1 + D_2$ can be computed with a probabilistic algorithm of type Las Vegas that takes an expected*

$$\tilde{O}\left(\delta^{\frac{\varpi}{2}+1} + (\deg D)^{\frac{\varpi+1}{2}}\right)$$

field operations in characteristic zero or $> \deg D$, and

$$\tilde{O}\left(\left(\delta^{\frac{\varpi}{2}+1} + (\deg D)^{\frac{\varpi+1}{2}}\right) \log q + \deg D \log^2 q\right)$$

bit operations if $\mathbb{K} = \mathbb{F}_q$.

Proof. First, a common primitive element $\lambda X + \mu Y$ is found at random for D_1 and D_2 with an expected $O(\deg D^\varpi)$ field operations in characteristic zero or $> \deg D$, or $(\deg D)^\varpi \tilde{O}(\log q) + \tilde{O}(\deg D \log^2 q)$ bit operations if $\mathbb{K} = \mathbb{F}_q$, by Proposition 12. The number of trials is $O(1)$ thanks to the assumption on $|\mathbb{K}|$.

We split D_i into $\hat{D}_i + \check{D}_i$ for $i = 1, 2$ such that $\hat{D}_1$ and $\hat{D}_2$ have the same support $\mathcal{E}$, itself disjoint from $\check{D}_1$ and $\check{D}_2$. Let $\hat{\chi}_i, \hat{u}_i, \hat{v}_i$ denote the parametrization of $\hat{D}_i$ for $i = 1, 2$. Let w_1 and w_2 be the cofactors in the Bézout relation $\gcd(\hat{\chi}_1, \hat{\chi}_2) = w_1 \hat{\chi}_1 + w_2 \hat{\chi}_2$, then

$$\begin{aligned} \tilde{\chi}_3 &:= \text{lcm}(\hat{\chi}_1, \hat{\chi}_2), \\ \tilde{u}_3 &:= \hat{u}_1 w_2 (\hat{\chi}_2 / \gcd(\hat{\chi}_1, \hat{\chi}_2)) + \hat{u}_2 w_1 (\hat{\chi}_1 / \gcd(\hat{\chi}_1, \hat{\chi}_2)) \bmod \tilde{\chi}_3, \\ \tilde{v}_3 &:= \hat{v}_1 w_2 (\hat{\chi}_2 / \gcd(\hat{\chi}_1, \hat{\chi}_2)) + \hat{v}_2 w_1 (\hat{\chi}_1 / \gcd(\hat{\chi}_1, \hat{\chi}_2)) \bmod \tilde{\chi}_3, \end{aligned}$$

is the parametrization of the divisor of support $\mathcal{E}$ where the multiplicity of a point P in it is the maximum of the multiplicities of P in $\hat{D}_1$ and $\hat{D}_2$. Therefore the parametrization of $D_3 := \hat{D}_1 + \hat{D}_2$ is deduced by means of a single lifting step, that costs

$$\tilde{O}\left(\delta^{\frac{\varpi}{2}+1} + (\deg D_3)^{\frac{\varpi+2}{3}}\right)$$

by Lemma 7. Glueing $\check{D}_1 + \check{D}_2 + D_3$ takes softly linear time by Lemma 10. $\square$

Proposition 14. *Given two smooth positive divisors D_1 and D_2 by their primitive element representations, and such that $|\mathbb{K}| \geq (\deg D_1 + \deg D_2)^2$, a representation of $[D_1 - D_2]_+$ can be computed with a probabilistic algorithm of type Las Vegas that takes an expected number of*

$$\tilde{O}((\deg D_1)^\varpi + (\deg D_2)^\varpi)$$

field operations in characteristic zero or $> \deg D_1 + \deg D_2$, or

$$\tilde{O}(((\deg D_1)^\varpi + (\deg D_2)^\varpi) \log q + (\deg D_1 + \deg D_2) \log^2 q)$$

bit operations if $\mathbb{K} = \mathbb{F}_q$.

Proof. First, a common primitive element $\lambda X + \mu Y$ is found for D_1 and D_2 as is the proof of the latter proposition, and let χ_i, u_i, v_i denote the parametrization of D_i for $i = 1, 2$. The parametrization of $[D_1 - D_2]_+$ is $\chi := \chi_1 / \gcd(\chi_1, \chi_2)$, $u = u_1 \bmod \chi$, $v := v_1 \bmod \chi$. $\square$

4 Riemann–Roch space

We are now ready to revisit the Brill–Noether strategy. For the mathematical aspects of the proofs below, we refer the reader to [19]. The main improvements upon [19] concern fast structured linear algebra, and the extension to any smooth input divisor D .

4.1 Shifted Popov form

Let M denote a $m \times n$ matrix with entries in $\mathbb{K}[X]$, and let us consider a vector $\mathbf{s} := (s_1, \dots, s_n) \in \mathbb{Z}^n$ called a *shift* for the degrees. The $\mathbf{s}$ -degree of a row vector $\mathbf{a} = (a_1, \dots, a_n)$ in $\mathbb{K}[X]^n$ is defined as $\deg_{\mathbf{s}} \mathbf{a} := \max(a_1 + s_1, \dots, a_n + s_n)$. If $\mathbf{a}$ is non-zero then the *pivot index* of $\mathbf{a}$ is the largest index i where the latter maximum is attained. The entry a_i is called the *pivot*, and its degree is the *pivot degree*. If $\mathbf{a}$ is zero then its pivot index is set to zero. The matrix M is in *Popov form* if the following properties are satisfied:

- The positive pivot indices of the rows of M are in increasing order;
- The pivots of the rows of M are monic;
- The pivots of M have a degree strictly larger than the other elements in their column.

When $m = n$ and M is nonsingular then its pivots are the diagonal elements. In this case, M satisfies the “predictable degree” property:

Lemma 15. *If $\mathbf{b} = (b_1, \dots, b_n) := \mathbf{a}M$, then $\deg b_i + s_i = d_i + \deg a_i$ for $i = 1, \dots, n$, where d_i denotes the $\mathbf{s}$ -degree of the i -th row of M .*

The “naive algorithm” for Popov forms takes $\tilde{O}(mnr(\deg M)^2)$ operations in $\mathbb{K}$, where r is the rank of M and when $\mathbf{s}$ is zero [22, Theorem 7.1]. The current best bounds are $\tilde{O}(m^{\omega-1}nd)$ for a $m \times n$ matrix with $m \leq n$ [24, 23].

Proposition 16. *Assume that M is square, nonsingular, and in Popov form as above. Then, the elements of $\mathbf{s}$ -degree $\leq d$ in the $\mathbb{K}[X]$ -module generated by the rows $M_1, \dots, M_n$ of M form a $\mathbb{K}$ -vector space of basis $X^j M_i$, $j = 0, \dots, d - d_i$, $i = 1, \dots, n$.*

Proof. A $\mathbb{K}$ -relation between the elements of the candidate basis leads to a $\mathbb{K}[X]$ -relation between the rows of M . According to the assumptions, such a proper relation cannot exist so the candidate basis is free over $\mathbb{K}$. An element $X^j M_i$ with $j = 0, \dots, d - d_i$ and $i = 1, \dots, n$ satisfies $\mathbf{s}\text{-deg}(X^j M_i) \leq j + d_i \leq d$. Conversely let $\mathbf{b}$ be a $\mathbb{K}[X]$ -combination $\mathbf{a}M$ of the rows of M of $\mathbf{s}$ -degree $\leq d$. Lemma 15 implies that $\deg a_i = \deg b_i + s_i - d_i \leq d - d_i$. $\square$

4.2 Bivariate interpolation

Considering Y as the “main variable”, the subset of polynomials in $\mathbb{K}[X][Y]$ that vanish at a given set of points is a free $\mathbb{K}[X]$ -module. This motivates the definition of *basis generator* of a vector subspace of polynomials of degree $\leq d$ in $\mathbb{K}[X, Y]$: this is a set $F_1, \dots, F_n$ of polynomials such that $X^j F_i$, $j = 0, \dots, d - \deg F_i$, $i = 1, \dots, n$ form a vector basis; n is called the *rank*.

Proposition 17. *Let D be a smooth positive divisor and let $d \geq 1$. Assume that D and E are parametrized by X and write $\chi(X) = Y - v(X) = 0$ for the corresponding parametrization of D . Then, there exists a basis generator of rank $\leq \delta$ of polynomials $F \in \mathbb{K}[X, Y]$ of degree $\leq d$ such that $F(X, v(X)) = 0 \bmod \chi(X)$, $F(X, v_E(X)) = 0 \bmod \chi_E(X)$, and $\deg_Y F < \delta$. It takes $\tilde{O}(\min(d, \delta)^{\omega-1}(r + \deg D))$ operations in $\mathbb{K}$; recall that $r = \deg E$.*

Proof. Let $n := \min(d, \delta - 1)$. The parametrization of E in this context is $\chi_E(X) = Y - v_E(X) = 0$. In softly linear time we compute $a_i = v^i \bmod \chi$ and $b_i = v_E^i \bmod \chi_E$, for $i = 0, \dots, n$ and then consider the $\mathbb{K}[X]$ -module

$$\mathcal{M} := \{(f_0, \dots, f_n) \in \mathbb{K}[X]^{n+1} : f_0 a_0 + \dots + f_n a_n = 0 \bmod \chi \text{ and } f_0 b_0 + \dots + f_n b_n = 0 \bmod \chi_E\}. \quad (7)$$

Since $\mathbb{K}[X]$ is principal, $\mathcal{M}$ is a free module of rank $n + 1$, because it contains $(0, \dots, 0, \chi, 0, \dots, 0)$ with χ at position i , for all $i = 0, \dots, n$.

Using [23, Theorem 1.4] with $\mathbf{s} := (d, d - 1, \dots, d - n)$, the nonsingular matrix in $\mathbf{s}$ -Popov form whose rows are a basis of $\mathcal{M}$ can be computed with $\tilde{O}(n^{\omega-1}(\deg \chi + \deg \chi_E))$ operations in $\mathbb{K}$. $\square$

4.3 Denominator

Let $Q_1, \dots, Q_r$ denote the singular points of $\mathcal{C}$, let $\mathcal{C}' \rightarrow \mathcal{C}$ be the desingularization map for $\mathcal{C}$, and for any $i \in \{1, \dots, r\}$ let $Q_{i,1}$ and $Q_{i,2}$ represent the points of $\mathcal{C}'$ above Q_i . In the sequel we set

$$d := \left\lceil \frac{(\delta - 1)(\delta - 2) + \deg D_+}{\delta} \right\rceil. \quad (8)$$

Lemma 18. *There exists a non-zero homogeneous polynomial H in $\bar{\mathbb{K}}[X, Y, Z]$ of degree $\leq d$ such that Q does not divide H , $(H)_0 \geq D_+$, $(H)_0 \geq E$, and the intersection multiplicities of H at the singular points of $\mathcal{C}$ are 2 (recall that $(H)_0 \geq E$ means “ H is adjoint to $\mathcal{C}$ ”).*

Proof. Let us fix a homogeneous polynomial $L \in \mathbb{K}[X, Y, Z]$ of degree 1, and set $\bar{D} := D_+ + \sum_{j=1}^r (Q_{j,1} + Q_{j,2})$. Since $\mathcal{C}$ has degree δ with only ordinary singularities, its genus is $g = \frac{(\delta-1)(\delta-2)}{2} - r$, and we have $\deg(L)_0 = \delta$, so the hypothesis on d means $\deg(d(L)_0 - \bar{D}) \geq 2g$. The Riemann–Roch theorem [4, Corollary 3] thus implies that

$$\dim(\mathcal{L}(d(L)_0 - \bar{D} - Q_{i,j})) = \dim(\mathcal{L}(d(L)_0 - \bar{D})) - 1$$

holds for all $(i, j) \in \{1, \dots, r\} \times \{1, 2\}$.

So far we have proved that for any $(i, j) \in \{1, \dots, r\} \times \{1, 2\}$ there exists a function $h_{i,j} \in \mathcal{L}(d(L)_0 - \bar{D})$ that is not contained in $\mathcal{L}(d(L)_0 - \bar{D} - Q_{i,j})$. By [7, Théorème 2.7.1] (or [8, Théorème 2.5]), $h_{i,j}$ has a rational function representation of the form $\frac{H_{i,j}}{L^d}$, where $H_{i,j} \in \bar{\mathbb{K}}[X, Y, Z]$ is homogeneous of degree d and is not divisible by Q . In other words $(H_{i,j})_0 \geq D_+$, $(H_{i,j})_0 \geq E$, and the intersection multiplicity of $H_{i,j}$ at $Q_{i,j}$ is 2.

Let $\alpha_{i,j}$ for $i = 1, \dots, r$ and $j = 1, 2$ be parameters in $\bar{\mathbb{K}}$, and consider $H := \sum_{i=1}^r \sum_{j=1}^2 \alpha_{i,j} H_{i,j}$. By construction $(H)_0 \geq D_+$ and $(H)_0 \geq E$ hold, and the resultant $\text{Res}_Y(Q, H)$ is a non-zero polynomial in all the $\alpha_{i,j}$ when regarded in $\mathbb{K}[\alpha_{1,1}, \dots, \alpha_{r,2}][X]$. Let $T_{i,j}$ denote a tangent vector at the image in $\mathcal{C}$ of the germ of curve of $\mathcal{C}'$ at $Q_{i,j}$. Regarded in $\bar{\mathbb{K}}[\alpha_{1,1}, \dots, \alpha_{r,2}]$ the polynomial $\prod_{i=1}^r \prod_{j=1}^2 (T_{i,j} \cdot \nabla H(Q_i))$ is non-zero. Consequently almost all choices of the $\alpha_{i,j}$ yield H with the required properties. $\square$

Algorithm 2

Input: $Q \in \mathbb{K}[X, Y, Z]$, E , a smooth divisor D on $\mathcal{C}$.

Output: $M \in \text{GL}_3(\mathbb{K})$, $H \in \mathbb{K}[X, Y]$ of degree $\leq d$ such that $\deg_Y H < \delta$, $(H)_0 \geq M^{-1}(D_+)$, $(H)_0 \geq M^{-1}(E)$, and with intersection multiplicities exactly 2 at the singular points of $M^{-1}(\mathcal{C})$; and $(H)_0 - 2M^{-1}(E)$ for which X is primitive.

Assumptions: $\mathcal{C}$ -H₁, $\mathcal{C}$ -H₂, $\deg_Y Q = \delta$; the supports of E and D are in the chart $Z = 1$.

1. Take α, β at random in $\mathbb{K}$ and set

$$M := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \beta & 0 & 1 \end{pmatrix}^{-1} \begin{pmatrix} 1 & \alpha & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}^{-1}.$$

2. If $\deg_Y(Q \circ M) \neq \delta$ then go to step 1.
3. If the supports of $M^{-1}(E)$ and $M^{-1}(D)$ are not in the chart $Z = 1$ then go to step 1.
4. If X is primitive for $M^{-1}(E)$ and $M^{-1}(D_+ + D_-)$, then compute its primitive element representation as in section 2.2. Otherwise go to step 1.
5. Let d be as in (8). Compute a basis generator $H_1, \dots, H_l$ of the polynomials H satisfying $\deg H \leq d$, $\deg_Y H < \delta$, $(H)_0 \geq M^{-1}(D_+)$, $(H)_0 \geq M^{-1}(E)$.
6. Set $H(X, Y) := \sum_{i=1}^l \alpha_i(X) H_i(X, Y)$ with $\alpha_i(X) \in \mathbb{K}[X]_{\leq d - \deg H_i}$ taken at random.
7. Compute the intersection of $H(X, Y) = 0$ and $(Q \circ M)(X, Y, 1) = 0$. If the cardinality of the solution set is not δd counting multiplicities, or does not admit X as a primitive element then go to step 1.
8. If the multiplicities of H at the singular points of $Q \circ M$ are not 2, then go to step 6.
9. Return M , H and $(H)_0 - 2M^{-1}(E)$.

Proposition 19. Assume $|\mathbb{K}| > 6(\delta d)^2$. Algorithm 2 is correct and takes an expected $\tilde{O}\left((\delta^2 + \deg D_+)^{\frac{\omega+1}{2}}\right)$ operations in characteristic 0 or $> \delta d$, or $\tilde{O}\left((\delta^2 + \deg D_+)^{\frac{\omega+1}{2}} \log q + (\delta^2 + \deg D_+) \log^2 q\right)$ bit operations when $\mathbb{K} = \mathbb{F}_q$.

Proof. By Lemma 18 there exists $\bar{H} \in \bar{\mathbb{K}}[X, Y, Z]$ of degree d not divisible by Q , such that $(\bar{H})_0 \geq D_+$, $(\bar{H})_0 \geq E$, and with intersection multiplicities exactly 2 at the singular points of $\mathcal{C}$. Let $\mathcal{E}$ denote the set of the zeros of $\bar{H}$ on $\mathcal{C}$. Since Q is monic in Y , $(0 : 0 : 1) \notin \mathcal{E}$. Consequently for all but finite number of values of β the set $M(\mathcal{E})$ is in the chart $Z = 1$. On the other hand for almost all values of α , the form X is primitive for $M^{-1}(\mathcal{E})$ and $\deg_Y(Q \circ M) = \delta$ holds. Then, $(\bar{H} \circ M) \text{rem}_Y(Q \circ M)$ belongs to the $\bar{\mathbb{K}}$ extension of the polynomial space computed in step 5. Consequently the algorithm finishes with a correct output for almost all values of $\alpha, \beta, \alpha_1, \dots, \alpha_n$ over $\bar{\mathbb{K}}$.

Let us now estimate the probabilities involved by random choices over $\mathbb{K}$. The coefficient of Y^δ in $Q \circ M$ is a non-zero polynomial of degree $\leq 2\delta$ in α, β . By the Schwartz-Zippel lemma, $\deg_Y(Q \circ M) = \delta$ fails with probability $\leq \frac{2\delta}{6(\delta d)^2} < \frac{1}{2}$. Assuming that step 2 succeeds, then the

probability of step 3 failing is

$$\leq \frac{r + \deg D_+ + \deg D_-}{6(d\delta)^2} \leq \frac{(\delta-1)(\delta-2) + 4\deg D_+}{6(d\delta)^2} \leq \frac{2\delta d}{6(d\delta)^2} \leq \frac{1}{2}.$$

Once step 3 has succeeded then β is properly fixed, step 4 requires that X be primitive. Using (1), this fails with probability

$$\leq \frac{\binom{r + \deg D_+ + \deg D_- + 1}{2}}{6(d\delta)^2} \leq \frac{2(\delta d)^2 + \delta d}{6(d\delta)^2} \leq \frac{1}{2}.$$

The coefficient of $X^{\delta d}$ in $R(X, Z) := \text{Res}_Y(Q \circ M, Z^d H(X/Z, Y/Z))$ is a non-zero homogeneous polynomial of degree $\leq \delta$ in the coefficients of $\alpha_1, \dots, \alpha_n$. In addition the discriminant of the separable part of $R(X, 1)$ is non-zero of degree $\leq 2\delta^2 d$ in the coefficients of $\alpha_1, \dots, \alpha_n$. Thus, the probability that step 7 fails is $\leq \frac{2\delta^2 d + \delta}{6(d\delta)^2} \leq \frac{1}{2}$.

Let $T_{i,j}$ denote a tangent vector at the image in $\mathcal{C}$ of the germ of curve of $\mathcal{C}'$ at $Q_{i,j}$. The polynomial $\prod_{i=1}^r \prod_{j=1}^2 (T_{i,j} \cdot \nabla H(Q_i))$ is non-zero of total degree $2r$ in the coefficients of $\alpha_1, \dots, \alpha_n$. By the Schwartz–Zippel lemma, the probability that step 8 fails given that all the previous steps succeeded is $\leq \frac{2r}{6(d\delta)^2} \leq \frac{1}{2}$. Consequently, the expected number of times the algorithm returns to step 1 or step 6 is $O(1)$.

Assume that $\mathbb{K}$ has characteristic 0 or $> d\delta$. Step 2 takes softly linear time by Lemma 3. Step 4 contributes to $\tilde{O}(r^\varpi + (\deg D_+)^\varpi)$ by Proposition 12. Step 5 takes $\tilde{O}(\min(d, \delta)^{\omega-1}(r + \deg D_+)) = \tilde{O}(\delta^{\omega-1}(r + \deg D_+))$ by Proposition 17. Step 6 contributes to $\tilde{O}(\delta d)$. Step 7 is done via Lemma 5 with $\tilde{O}(d\delta^2)$ operations in $\mathbb{K}$.

In step 8 since X is primitive for $(H)_0$ and E and since the intersection multiplicities in $(H)_0$ are known, we can conveniently check whether the points in E have intersection multiplicity 2. And if so we deduce a primitive element representation of $(H)_0 - 2M^{-1}(E)$ in softly linear time. The total complexity bound is obtained by summing the cost of each step, thanks to $r = O(\delta^2)$ and $\delta^{\omega-1}(\delta^2 + \deg D_+) = O\left((\delta^2 + \deg D_+)^{\frac{\omega+1}{2}}\right)$. The same kind of analysis applies over $\mathbb{F}_q$, and is left to the reader. $\square$

The value of d defined in (8) guarantees that a denominator H can be found in degree $\leq d$. In favorable cases, smaller values for d are possible: in fact, when $\deg D_+ = O(\delta^2)$ and $r = 0$ the degree bound d used in [19] is sharper.

4.4 Riemann–Roch space

Once we have obtained a common denominator H as above for $\mathcal{L}(D)$, we focus on the numerators, as follows.

Algorithm 3

Input: $Q \in \mathbb{K}[X, Y, Z]$ of degree δ , E , and a smooth divisor D on $\mathcal{C}$.

Output: a basis generator of rank $\leq \delta$ of $\mathcal{L}(D)$.

Assumptions: $\mathcal{C}$ -H₁, $\mathcal{C}$ -H₂, $\deg_Y Q = \delta$; the supports of E and D are in the chart $Z = 1$.

1. Compute M , H and $D_{\text{res}} := (H)_0 - 2M^{-1}(E)$ by means of Algorithm 2.
2. Compute $D_{\text{num}} = M^{-1}(D_-) + (D_{\text{res}} - M^{-1}(D_+))$.
3. Compute a basis generator $G_1, \dots, G_l$ of the vector space of polynomials in $\mathbb{K}[X, Y]$ of degree $\leq d$ such that $(G)_0 \geq D_{\text{num}}$ and $(G)_0 \geq M^{-1}(E)$.
4. Return M and $G_1, \dots, G_l$.

Proposition 20. Assume that $|K| \geq 6(\delta d)^2$. Then, Algorithm 3 is correct and takes an expected $\tilde{O}\left((\delta^2 + \deg D_+)^{\frac{\omega+1}{2}}\right)$ field operations in characteristic 0 or $> \delta d$, or

$$\tilde{O}\left((\delta^2 + \deg D_+)^{\frac{\omega+1}{2}} \log q + (\delta^2 + \deg D_+) \log^2 q\right)$$

bit operations when $\mathbb{K} = \mathbb{F}_q$.

Proof. Combination of Propositions 13, 14, 17 and 19. $\square$

Proof. (Proof of Theorem 1) First we use Proposition 8. Once the resulting change of variables is applied to Q and E , Proposition 20 yields the claimed complexity bounds. $\square$

In the case $\mathbb{K} = \mathbb{F}_q$, most of the auxiliary routines get closer to optimality in theory: for bivariate composition, bounds *à la* Kedlaya–Umans are quasi-linear. Unfortunately they have not led to efficient practical implementations so far; see [14, section 8]. For curve intersections, better complexity bounds also exist, but under genericity assumptions; see [13, 26]. If assumptions could be dropped then the complexity bound of Theorem 1 would become $\tilde{O}(\delta^{\omega-1}(r + \deg D_+) \log q + (\delta^2 + \deg D_+)^{1+\epsilon} \log q + (\delta^2 + \deg D_+) \log^2 q)$ bit operations, for any fixed $\epsilon > 0$. The bottleneck would be structured linear algebra underlying Proposition 17. If ω were further proved to be close to 2, then our algorithm would be close to optimal in terms of the size of the output whenever $r = O(\deg D_+)$.

Acknowledgements. This paper is part of a project that has received funding from the French “Agence de l’Innovation de Défense”. We are grateful to Vincent Neiger for helpful discussions.

References

- [1] G. Chèze and G. Lecerf. Lifting and recombination techniques for absolute factorization. *J. Complexity*, 23(3):380–420, 2007.
- [2] X. Dahan, M. Moreno Maza, É. Schost, and Yuzhen Xie. On the complexity of the D5 principle. In J.-G. Dumas, editor, *Proceedings of Transgressive Computing 2006: a conference in honor of Jean Della Dora*, pages 149–168. U. J. Fourier, Grenoble, France, 2006.
- [3] C. Durvy and G. Lecerf. A concise proof of the Kronecker polynomial system solver from scratch. *Expo. Math.*, 26(2), 2007.
- [4] W. Fulton. *Algebraic Curves – An Introduction to Algebraic Geometry*. Addison-Wesley, 1989.
- [5] J. von zur Gathen and J. Gerhard. *Modern computer algebra*. Cambridge University Press, New York, 3rd edition, 2013.
- [6] B. Grenet, J. van der Hoeven, and G. Lecerf. Deterministic root finding over finite fields using Graeffe transforms. *Appl. Algebra Engrg. Comm. Comput.*, 27(3):237–257, 2016.
- [7] G. Haché. *Construction Effective des Codes Géométriques*. PhD thesis, Université Paris 6, 1996.
- [8] G. Haché. *L’algorithme de Brill-Noether appliqué aux courbes réduites*. Rapport de recherche n° 1998-01, Laboratoire d’Arithmétique, de Calcul formel et d’Optimisation ESA - CNRS 6090, Université de Limoges, France, 1998. https://www.unilim.fr/laco/rapports/1998/R1998_01.pdf.

- [9] F. Hess. Computing Riemann–Roch spaces in algebraic function fields and related topics. *J. Symbolic Comput.*, 33(4):425–445, 2002.
- [10] J. van der Hoeven and G. Lecerf. Composition modulo powers of polynomials. In *Proceedings of the 2017 ACM on International Symposium on Symbolic and Algebraic Computation*, ISSAC ’17, pages 445–452. New York, NY, USA, 2017. ACM.
- [11] J. van der Hoeven and G. Lecerf. Directed evaluation. Technical Report, HAL, 2018. <https://hal.archives-ouvertes.fr/hal-01966428>.
- [12] J. van der Hoeven and G. Lecerf. On the complexity exponent of polynomial system solving. Technical Report, HAL, 2018. <http://hal.archives-ouvertes.fr/hal-01848572>.
- [13] J. van der Hoeven and G. Lecerf. Fast computation of generic bivariate resultants. Technical Report, HAL, 2019. <https://hal.archives-ouvertes.fr/hal-02080426>.
- [14] J. van der Hoeven and G. Lecerf. Fast multivariate multi-point evaluation revisited. *J. Complexity*, 56:101405, 2020.
- [15] Xiaohan Huang and V. Y. Pan. Fast rectangular matrix multiplication and applications. *J. Complexity*, 14(2):257–299, 1998.
- [16] K. Khuri-Makdisi. Asymptotically fast group operations on Jacobians of general curves. *Math. Comp.*, 76(260):2213–2239, 2007.
- [17] D. Le Brigand and J.-J. Risler. Algorithme de Brill–Noether et codes de Goppa. *Bulletin de la société mathématique de France*, 116(2):231–253, 1988.
- [18] F. Le Gall. Powers of tensors and fast matrix multiplication. In K. Nabeshima, editor, *ISSAC’14: International Symposium on Symbolic and Algebraic Computation*, pages 296–303. New York, NY, USA, 2014. ACM.
- [19] A. Le Gluher and P.-J. Spaenlehauer. A fast randomized geometric algorithm for computing Riemann–Roch spaces. *Math. Comp.*, 2019. <https://doi.org/10.1090/mcom/3517>.
- [20] G. Lecerf. Fast separable factorization and applications. *Appl. Algebra Engrg. Comm. Comput.*, 19(2):135–160, 2008.
- [21] G. Lecerf. On the complexity of the Lickteig–Roy subresultant algorithm. *J. Symbolic Comput.*, 92:243–268, 2019.
- [22] T. Mulders and A. Storjohann. On lattice reduction for polynomial matrices. *J. Symbolic Comput.*, 35(4):377–401, 2003.
- [23] V. Neiger. Fast computation of shifted Popov forms of polynomial matrices via systems of modular polynomial equations. In *Proceedings of the ACM on International Symposium on Symbolic and Algebraic Computation*, ISSAC ’16, pages 365–372. New York, NY, USA, 2016. ACM.
- [24] V. Neiger, J. Rosenkilde, and G. Solomatov. Computing Popov and Hermite forms of rectangular polynomial matrices. In *Proceedings of the 2018 ACM International Symposium on Symbolic and Algebraic Computation*, ISSAC ’18, pages 295–302. New York, NY, USA, 2018. ACM.

- [25] M. Nüsken and M. Ziegler. Fast multipoint evaluation of bivariate polynomials. In S. Albers and T. Radzik, editors, *Algorithms – ESA 2004. 12th Annual European Symposium, Bergen, Norway, September 14-17, 2004*, volume 3221 of *Lect. Notes Comput. Sci.*, pages 544–555. Springer Berlin Heidelberg, 2004.
- [26] G. Villard. On computing the resultant of generic bivariate polynomials. In *Proceedings of the 2018 ACM International Symposium on Symbolic and Algebraic Computation, ISSAC '18*, pages 391–398. New York, NY, USA, 2018. ACM.