

On the Complexity of Solving Bivariate Systems: the Case of Non-singular Solutions

Romain Lebreton

LIRMM, UMR 5506 CNRS
Université de Montpellier II
Montpellier, France
lebreton@lirmm.fr

Esmail Mehrabi

Computer Science Dept.
Western University
London, ON, Canada
emehrab@uwo.ca

Éric Schost

Computer Science Dept.
Western University
London, ON, Canada
eschost@uwo.ca

ABSTRACT

We give an algorithm for solving bivariate polynomial systems over either $k(T)[X, Y]$ or $\mathbb{Q}[X, Y]$ using a combination of lifting and modular composition techniques.

Categories and Subject Descriptors

I.1.2 [Computing Methodologies]: Symbolic and Algebraic Manipulation—*Algebraic algorithms*

Keywords

Bivariate polynomial systems; complexity.

1. INTRODUCTION AND MAIN RESULTS

We investigate the complexity of solving bivariate polynomial systems. This question is interesting in its own right, but it also plays an important role in many higher-level algorithms, such as computing the topology of plane and space curves [13, 8] or solving general polynomial systems [18].

Many recent contributions on this question discuss computing real solutions of bivariate systems with integer or rational coefficients [15, 12, 30, 4, 14], by a combination of symbolic elimination and real root isolation techniques. Our interest here is on complexity of the “symbolic” component of such algorithms. One of our main results says that we can solve bivariate systems with integer coefficients in essentially optimal time, at least for non-singular solutions.

Geometric description. Let \mathbb{A} be a domain, let \mathbb{K} be its field of fractions and let $\overline{\mathbb{K}}$ be an algebraic closure of \mathbb{K} .

Let X, Y be the coordinates and let $Z \subset \overline{\mathbb{K}}^2$ be a finite set defined over \mathbb{K} and of cardinality δ (so the defining ideal $I \subset \overline{\mathbb{K}}[X, Y]$ of Z is generated by polynomials in $\mathbb{K}[X, Y]$). To describe Z , one may use a Gröbner basis of I , say for the lexicographic order $Y > X$. Such bases can however be unwieldy (they may involve a large number of polynomials, making modular computations difficult). Triangular decompositions are an alternative for which this issue is alleviated.

Geometrically, performing a triangular decomposition of the defining ideal of Z amounts to writing Z as the disjoint union of finitely many *equiprojectable sets*. Let $\pi : \overline{\mathbb{K}}^2 \rightarrow \overline{\mathbb{K}}$ be the projection on the X -space given by $(x, y) \mapsto x$. To $p = (x, y)$ in Z , we associate the positive integer $N(Z, p)$ defined as the cardinality of the fiber $\pi^{-1}(x) \cap Z$: this is the number of points in Z lying above x . We say that Z is *equiprojectable* if there exists a positive integer n such that $N(Z, p) = n$ for all $p \in Z$ (see [10] for illustrations).

It is proved in [3] that Z is equiprojectable if and only if its defining ideal I admits a Gröbner basis for the lexicographic order $Y > X$ that is a *monic triangular set*, i.e. of the form $\mathbf{T} = (U(X), V(X, Y))$, with U and V monic in respectively X and Y and with coefficients in \mathbb{K} (that result holds over a perfect field, so it applies over $\overline{\mathbb{K}}$; the fact that I has generators in $\mathbb{K}[X, Y]$ implies that \mathbf{T} has coefficients in \mathbb{K}). The degree $m = \deg(U, X)$ is the cardinality of $\pi(Z)$, and the equalities $n = \deg(V, Y)$ and $\delta = mn$ hold; we will say that \mathbf{T} has *bidegree* (m, n) .

When Z is not equiprojectable, it can be decomposed into equiprojectable sets, usually in a non-unique manner. The *equiprojectable decomposition* [10] is a canonical way to do so: it decomposes Z into subsets Z_{n_1}, \dots, Z_{n_s} , where for all $i \in \{1, \dots, s\}$, Z_{n_i} is the set of all $p \in Z$ for which $N(Z, p) = n_i$. This decomposition is implicit in the Cerlienco-Murredu description of the lexicographic Gröbner basis of the defining ideal of Z [7]; it can also be derived from Lazard’s structure theorem for bivariate Gröbner bases [22].

If Z is defined over \mathbb{K} , then all Z_{n_i} are defined over \mathbb{K} as well, so they can be represented by monic triangular sets

$$\mathbf{T}_1 \left| \begin{array}{l} V_1(X, Y) \\ U_1(X) \end{array} \right. \quad \dots \quad \mathbf{T}_s \left| \begin{array}{l} V_s(X, Y) \\ U_s(X) \end{array} \right. \quad (1)$$

with coefficients in \mathbb{K} . If we let $m_i = |\pi(Z_{n_i})|$, then \mathbf{T}_i has bidegree (m_i, n_i) for all i , and $\sum_{i \leq s} m_i n_i = \delta$.

By abuse of notation, we will call the family of monic triangular sets $\mathcal{T} = (\mathbf{T}_1, \dots, \mathbf{T}_s)$ the *equiprojectable decomposition* of Z . If I is a radical ideal of $\mathbb{K}[X, Y]$ that remains radical in $\overline{\mathbb{K}}[X, Y]$, its zero-set Z is defined over \mathbb{K} ; then, we define the equiprojectable decomposition of I as that of Z .

Solving systems. Let now F and G be in $\mathbb{A}[X, Y]$. In this paper, we are interested in the set $Z(F, G)$ of *non-singular solutions* of the system $F = G = 0$, that is, the points (x, y) in $\overline{\mathbb{K}}^2$ such that $F(x, y) = G(x, y) = 0$ and $J(x, y) \neq 0$, where J is the Jacobian determinant of (F, G) . Remark that $Z(F, G)$ is a finite set, defined over \mathbb{K} ; if F and G have total degree at most d , then $Z(F, G)$ has cardinality $\delta \leq d^2$.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISSAC’13, June 26–29, 2013, Boston, Massachusetts, USA.
Copyright 2013 ACM 978-1-4503-2059-7/13/06 ...\$15.00.

For instance, for generic F and G , $Z(F, G)$ coincides with their whole zero-set $V(F, G)$, it is equiprojectable ($s = 1$), the corresponding triangular set $\mathbf{T} = \mathbf{T}_1$ takes the form $\mathbf{T} = (U(X), Y - \eta(X))$ and U is (up to a constant in \mathbb{K}) the resultant of F and G in Y .

Given F and G , our goal will be, up to a minor adjustment, to compute the triangular sets $\mathcal{T} = (\mathbf{T}_1, \dots, \mathbf{T}_s)$ that define the equiprojectable decomposition of $Z(F, G)$.

Representing these polynomials requires $O(d^2)$ elements of \mathbb{K} . We will show below that one can compute them using $O^-(d^3)$ operations in \mathbb{K} , where $O^-(\cdot)$ indicates the omission of logarithmic factors. It is a major open problem to compute \mathcal{T} in time $O^-(d^2)$, just like it is an open problem to compute the resultant of F and G in such a cost [16, Problem 11.10].

Size considerations. In this paper, we are mainly interested in a refinement of this situation to cases where \mathbb{A} is endowed with a “length” function; in such cases, the cost analysis must take this length into account. Rather than giving an axiomatic treatment, we will assume that we are in one of the following situations:

- $\mathbb{A} = k[T]$ and $\mathbb{K} = k(T)$, for a field k , where we use the length function $\lambda(a) = \deg(a)$, for $a \in \mathbb{A} - \{0\}$;
- $\mathbb{A} = \mathbb{Z}$ and $\mathbb{K} = \mathbb{Q}$, where we use the length function $\lambda(a) = \log(|a|)$, for $a \in \mathbb{A} - \{0\}$.

In both cases, the length of $a \in \mathbb{A}$ represents the amount of storage needed to represent it, in terms of elements of k , resp. bits. It will be useful to introduce a notion of length for polynomials with coefficients in \mathbb{K} : if P is such a polynomial, $\lambda(P)$ denotes the maximum of the lengths $\lambda(n_i)$ and $\lambda(d_i)$, where n_i and d_i are the numerators and denominators of the coefficients of P , when written in reduced form using a common denominator.

When $\mathbb{A} = k[T]$, we are studying the intersection of two surfaces in a 3-dimensional space with coordinates T, X, Y ; the output describes the solution curve for generic T .

In that case, write again $d = \max(\deg(F), \deg(G))$, as well as $\ell = \max(\lambda(F), \lambda(G))$. Then, the polynomials U_1, \dots, U_s in the equiprojectable decomposition (1) of $Z(F, G)$ are in $k(T)[X]$, and the sum of their degrees in X is at most d^2 . These polynomials are all factors of the resultant $\text{res}(F, G, Y)$, which implies that $\lambda(U_i)$ is at most $2d\ell$ for each i , so that representing them involves $O(d^3\ell)$ coefficients in k .

For the polynomials V_1, \dots, V_s , however, the bounds are worse: [11] proves that $\lambda(V_i)$ only admits a weaker bound of order $d^3\ell + d^4$, so they involve $O(d^5\ell + d^6)$ coefficients in k . Practice shows that these bounds are realistic: the polynomials V_i are usually much larger than the polynomials U_i . In order to resolve this issue, we will use the polynomials N_1, \dots, N_s defined by $N_i = U_i'V_i \bmod U_i$ for all i . Then, Theorem 2 from [11] combined with the bi-homogeneous Bézout bound shows that $\lambda(N_i) \leq 2d\ell + d^2$ for all i ; thus, storing these polynomials uses $O(d^3\ell + d^4)$ coefficients in k .

Entirely similar considerations apply in the case $\mathbb{A} = \mathbb{Z}$; in that case, Theorem 1 from [11] and an arithmetic Bézout theorem [21] prove that $\lambda(U_i) \leq 2d\ell + 24d^2$, and similarly for $\lambda(N_i)$, so $O(d^3\ell + d^4)$ bits are sufficient to store them.

We call *modified equiprojectable decomposition* of $Z(F, G)$ the set of polynomials $\mathcal{C} = (\mathbf{C}_1, \dots, \mathbf{C}_s)$, with $\mathbf{C}_i = (U_i, N_i)$. These are not monic triangular sets anymore (N_i is not monic in Y), but *regular chains* [2]. In the particular case

where $s = 1$ and $V = V_1$ has the form $V(X, Y) = Y - \eta(X)$, it coincides with the rational univariate representation [29].

Main results. Our main results are the following theorems, that give upper bounds on the cost of computing the modified equiprojectable decomposition. We start with the case $\mathbb{A} = k[T]$, where we count operations in k at unit cost. Our second result concerns the case $\mathbb{A} = \mathbb{Z}$; in this case, we measure the cost of our algorithm using bit operations.

In what follows, we let $M : \mathbb{N} \rightarrow \mathbb{N}$ be such that over any ring, univariate polynomials of degree less than d can be multiplied in $M(d)$ ring operations, under the super-linearity conditions of [16, Ch. 8]: using FFT techniques, we can take $M(d) \in O(d \log(d) \log \log(d))$. We also let ω be such that we can multiply $n \times n$ matrices using $O(n^\omega)$ ring operations, over any ring. The best known bound is $\omega < 2.38$ [33].

THEOREM 1. *Let k be a field and let F, G be in $k[T][X, Y]$, with $d = \max(\deg(F), \deg(G))$ and $\ell = \max(\lambda(F), \lambda(G))$. If k has characteristic at least $4d^2(6d^2 + 9d\ell)$, one can compute the modified equiprojectable decomposition of $Z(F, G)$ over $k(T)[X, Y]$ by a probabilistic algorithm with probability of success at least $1/2$, using*

$$O\left(M(d^2)M(d\ell + d^2)d^{(\omega-1)/2} \log(d\ell)\right) \subset O^-(d^{3.69\ell} + d^{4.69})$$

operations in k .

THEOREM 2. *Let $\varepsilon > 0$, let F, G be in $\mathbb{Z}[X, Y]$, and write $d = \max(\deg(F), \deg(G))$ and $\ell = \max(\lambda(F), \lambda(G))$. One can compute the modified equiprojectable decomposition of $Z(F, G)$ over $\mathbb{Q}[X, Y]$ by a probabilistic algorithm with probability of success at least $1/2$, using $O(d^{3+\varepsilon}\ell + d^{4+\varepsilon})$ bit operations.*

In both cases, one can easily obtain a cost of $O^-(d^4\ell + d^5)$ using modular methods: e.g., over $\mathbb{A} = k[T]$, solve the system at $O(d\ell + d^2)$ values of T , each of which costs $O^-(d^3)$ operations in k , and use rational function interpolation. Our main contribution is to show that this direct approach is sub-optimal; over $\mathbb{A} = \mathbb{Z}$, the cost of our algorithm almost matches the known upper bounds on the output size.

The structure of our algorithm is the same in both cases: we compute $Z(F, G)$ modulo an ideal \mathfrak{m} of \mathbb{A} , lift the result modulo a high power of \mathfrak{m} and reconstruct all rational function coefficients. This approach is similar to the algorithm of [10]; the key difference is in how we implement the lifting process. The result in [10] assumes that the input system is given by means of a straight-line program; here, we assume that the input is dense, and we rely on fast *modular composition techniques*.

Our results imply similar bounds for computing the resultant $R = \text{res}(F, G, Y)$, at least for systems without singular roots: one can reconstruct R from U_1, \dots, U_s , taking care if needed of the leading coefficients of F and G in Y ; we leave the details to the reader. The main challenge is to handle systems with multiplicities without affecting the complexity. We expect that deflation techniques will make this possible.

After a section of preliminaries, we give (Section 3) an algorithm to compute $Z(F, G)$ over an arbitrary field in time $O^-(d^3)$. Section 4 is devoted to computing normal forms modulo triangular sets by means of modular composition techniques; this is the key ingredient of the main algorithm given in Section 5. Section 6 presents experimental results.

2. PRELIMINARIES

2.1 Notation and basic results

In the introduction, we defined monic triangular sets with coefficients in a field. We will actually allow coefficients in a ring \mathbb{A} ; as in the introduction, all monic triangular sets will be bivariate, that is, in $\mathbb{A}[X, Y]$.

For positive integers m, n , $\mathbb{A}[X]_m$ denotes the set of all $F \in \mathbb{A}[X]$ such that $\deg(F) < m$, and $\mathbb{A}[X, Y]_{m, n}$ the set of all $F \in \mathbb{A}[X, Y]$ such that $\deg(F, X) < m$ and $\deg(F, Y) < n$. Using Kronecker's substitution, we can multiply polynomials in $\mathbb{A}[X, Y]_{m, n}$ in $O(M(mn))$ operations in \mathbb{A} .

For a monic triangular set \mathbf{T} in $\mathbb{A}[X, Y]$, the monicity assumption makes the notion of remainder modulo the ideal $\langle \mathbf{T} \rangle$ well-defined; if \mathbf{T} has bidegree (m, n) , then for any F in $\mathbb{A}[X, Y]$, the remainder $F \bmod \langle \mathbf{T} \rangle$ is in $\mathbb{A}[X, Y]_{m, n}$. In terms of complexity, we have the following result about computations with such a triangular set (see [25, 24]).

LEMMA 1. *Let \mathbf{T} be a monic triangular set in $\mathbb{A}[X, Y]$ of bidegree (m, n) . Then, given $F \in \mathbb{A}[X, Y]_{m', n'}$, with $m \leq m'$ and $n \leq n'$, one can compute $F \bmod \langle \mathbf{T} \rangle$ in $O(M(m'n'))$ operations in \mathbb{A} . Additions, resp. multiplications modulo $\langle \mathbf{T} \rangle$ can be done in $O(mn)$, resp. $O(M(mn))$ operations in \mathbb{A} .*

We continue with a result on polynomial matrix multiplication. The proof is the same as that of [5, Lemma 8], up to replacing univariate polynomials by bivariate ones. Remark that for such rectangular matrix multiplications, one could actually use an algorithm of Huang and Pan's [19], which features a slightly better exponent (for current values of ω).

LEMMA 2. *Let $\mathbf{M}_1, \mathbf{M}_2$ be matrices of sizes $(a \times b)$ and $(b \times c)$, with entries in $\mathbb{A}[X, Y]_{m, n}$. If $a = O(\ell^{1/2})$, $b = O(\ell^{1/2})$ and $c = O(\ell)$, one can compute $\mathbf{M} = \mathbf{M}_1 \mathbf{M}_2$ using $O(M(mn)\ell^{(\omega+1)/2})$ operations in \mathbb{A} .*

2.2 Chinese Remainder techniques

Let $\mathcal{T} = (\mathbf{T}_1, \dots, \mathbf{T}_s)$ be a family of monic triangular sets in $\mathbb{A}[X, Y]$, where \mathbb{A} is a ring. In such situations, we write $\langle \mathcal{T} \rangle = \langle \mathbf{T}_1 \rangle \cap \dots \cap \langle \mathbf{T}_s \rangle$; if \mathbb{A} is a field, we write $V(\mathcal{T}) = V(\mathbf{T}_1) \cup \dots \cup V(\mathbf{T}_s)$, where $V(\mathbf{T})$ denotes the zero-set of \mathbf{T} over the algebraic closure of \mathbb{A} .

Following [10], we say that \mathcal{T} is *non-critical* if for i in $\{1, \dots, s\}$, $F_i = U_1 \cdots U_{i-1} U_{i+1} \cdots U_s$ is invertible modulo U_i ; if \mathbb{A} is a field, this simply means that U_1, \dots, U_s are pairwise coprime. The family \mathcal{T} is a *non-critical decomposition* of an ideal I if \mathcal{T} is non-critical and $\langle \mathcal{T} \rangle = I$.

Let $\mathcal{T} = (\mathbf{T}_1, \dots, \mathbf{T}_s)$ be a non-critical family of triangular sets, with $\mathbf{T}_i = (U_i(X), V_i(X, Y))$ of bidegree (m_i, n_i) , and suppose that there exists n such that $n_i = n$ for all i ; let also $m = m_1 + \dots + m_s$. Under these conditions, the following lemma shows how to merge \mathcal{T} into a single monic triangular set \mathbf{T} of bidegree (m, n) . Because \mathbb{A} may not be a field, we assume that $\mathcal{R} = (R_1, \dots, R_s)$ is part of the input, with $R_i = 1/F_i \bmod U_i$; we call them the *cofactors* of \mathcal{T} .

LEMMA 3. *Given a non-critical family \mathcal{T} as above, under the assumption $n_i = n$ for all i , and given the cofactors \mathcal{R} , we can compute a monic triangular set \mathbf{T} of bidegree (m, n) such that $\langle \mathbf{T} \rangle = \langle \mathcal{T} \rangle$ using $(nM(m) \log(m))$ operations in \mathbb{A} .*

Given $F \in \mathbb{A}[X, Y]$ reduced modulo $\langle \mathbf{T} \rangle$, we can compute all polynomials $F_i = F \bmod \langle \mathbf{T}_i \rangle$ using $O(nM(m) \log(m))$ operations in \mathbb{A} .

PROOF. For $i = 1, \dots, s$, write $V_i = \sum_{j=0}^n v_{i,j} Y^j$, with all $v_{i,j}$ in $\mathbb{A}[X]$. Algorithm 10.22 in [16], where our polynomials R_i are written s_i , allows us to apply the Chinese Remainder Theorem, yielding v_0, \dots, v_n in $\mathbb{A}[X]$ such that $v_{i,j} = v_j \bmod U_i$ for all i, j . Since $v_{i,n} = 1$ for all i , $v_n = 1$ as well, so we let $U = U_1 \cdots U_s$, $V = \sum_{j=0}^n v_j Y^j$ and $\mathbf{T} = (U, V)$. Computing U takes $O(M(m) \log(m))$ by [16, Lemma 10.4] and computing V takes a total time of $O(nM(m) \log(m))$ by [16, Coro. 10.23].

To prove the second point, write $F = \sum_{j=0}^{n-1} f_j Y^j$, with all f_j in $\mathbb{A}[X]$. For $j = 0, \dots, n-1$, we apply the modular reduction algorithm of [16, Algo. 10.16] to compute $f_{1,j}, \dots, f_{s,j}$, with $f_{i,j} = f_j \bmod U_i$; we return $F_i = \sum_{j=0}^{n-1} f_{i,j} Y^j$, for $i = 1, \dots, s$. The total time is n times the cost of modular reduction, that is, $O(nM(m) \log(m))$. \square

COROLLARY 1. *Let \mathbb{K} be a field and let $\mathcal{T} = (\mathbf{T}_1, \dots, \mathbf{T}_s)$ be a non-critical family of monic triangular sets in $\mathbb{K}[X, Y]$, with $\mathbf{T}_i = (U_i, V_i)$ of bidegree (m_i, n_i) for all i . Suppose that the ideal $\langle \mathcal{T} \rangle \cdot \overline{\mathbb{K}}[X, Y]$ is radical. Then one can compute the equiprojectable decomposition of the ideal $\langle \mathcal{T} \rangle$ using $O(M(\delta) \log(\delta))$ operations in \mathbb{K} , with $\delta = \sum_{1 \leq i \leq s} m_i n_i$.*

PROOF. Partition \mathcal{T} in the classes of the equivalence relation where $(U, V) \equiv (U^*, V^*)$ if and only if $\deg(V, Y) = \deg(V^*, Y)$. Let $\mathcal{T}_1, \dots, \mathcal{T}_t$ be these classes; for $j \in \{1, \dots, t\}$, let $\mu_j = \sum_{(U, V) \in \mathcal{T}_j} \deg(U, X)$ and let ν_j be the common value of $\deg(V, Y)$ for $(U, V) \in \mathcal{T}_j$; then, $\sum_{1 \leq j \leq t} \mu_j \nu_j = \delta$.

For $j = 1, \dots, t$, let \mathbf{T}_j^* be the monic triangular set obtained by applying the previous lemma to \mathcal{T}_j . Since \mathbb{K} is a field, the cofactors \mathcal{R}_j are computed in time $O(M(\mu_j) \log(\mu_j))$ using [16, Algo. 10.18], so the total time for any fixed j is $O(\nu_j M(\mu_j) \log(\mu_j))$, which is $O(M(\nu_j \mu_j) \log(\nu_j \mu_j))$. Summing over all j , the total cost is seen to be $O(M(\delta) \log(\delta))$.

Since $\langle \mathcal{T} \rangle$ is radical in $\overline{\mathbb{K}}[X, Y]$, we deduce that for all $i \in \{1, \dots, s\}$, the zero-set Z_i of \mathbf{T}_i is equiprojectable, with fibers for the projection $\pi : \mathbb{K}^2 \rightarrow \overline{\mathbb{K}}$ all having cardinality n_i . Thus, the triangular sets $\mathbf{T}_1^*, \dots, \mathbf{T}_t^*$ form the equiprojectable decomposition of $\langle \mathcal{T} \rangle$. \square

2.3 Specialization properties

Consider a domain \mathbb{A} , its fraction field \mathbb{K} , and a maximal ideal $\mathfrak{m} \subset \mathbb{A}$ with residual field $k = \mathbb{A}/\mathfrak{m}$. Given F and G in $\mathbb{A}[X, Y]$, our goal here is to relate the equiprojectable decomposition of $Z(F, G)$ to that of $Z(F \bmod \mathfrak{m}, G \bmod \mathfrak{m})$, where the former is defined over \mathbb{K} and the latter over k .

The following results give quantitative estimates for ideals of "good reduction" in the two cases we are interested in, $\mathbb{A} = k[T]$ and $\mathbb{A} = \mathbb{Z}$; in both cases, we use the length function λ defined in the introduction. The case $\mathbb{A} = k[T]$, while not treated in [10], is actually the simpler, so we only sketch the proof; for $\mathbb{A} = \mathbb{Z}$, we can directly apply [10, Th. 1].

LEMMA 4. *Let F, G be in $k[T][X, Y]$ and let $\mathbf{T}_1, \dots, \mathbf{T}_s \subset k(T)[X, Y]$ be the equiprojectable decomposition of $Z(F, G)$. If $d = \max(\deg(F), \deg(G))$ and $\ell = \max(\lambda(F), \lambda(G))$, there exist $A \in k[T] - \{0\}$ of degree at most $2d^2(6d^2 + 9d\ell)$ and with the following property.*

If an element $t_0 \in k$ does not cancel A , then none of the denominators of the coefficients of $\mathbf{T}_1, \dots, \mathbf{T}_s$ vanishes at $T = t_0$ and their evaluation at $T = t_0$ forms the equiprojectable decomposition of $Z(F(t_0, X, Y), G(t_0, X, Y))$.

PROOF. The approach of [10, Section 3] still applies in this case, and shows that if an element $t_0 \in k$ satisfies three assumptions (denoted by H_1, H_2, H_3 in [10]), then the specialisation property holds. These properties imply the existence of a non-zero $A \in k[T]$ as claimed in the lemma; its degree can be bounded using the results of [31, 11]. \square

LEMMA 5. *Let F, G be in $\mathbb{Z}[X, Y]$ and let $\mathbf{T}_1, \dots, \mathbf{T}_s \subset \mathbb{Q}[X, Y]$ be the equiprojectable decomposition of $Z(F, G)$. If $d = \max(\deg(F), \deg(G))$ and $\ell = \max(\lambda(F), \lambda(G))$, there exists $A \in \mathbb{N} - \{0\}$, with $\lambda(A) \leq 8d^5(3\ell + 10\log(d) + 22)$ and with the following property.*

If a prime $p \in \mathbb{N}$ does not divide A , then none of the denominators of the coefficients of $\mathbf{T}_1, \dots, \mathbf{T}_s$ vanishes modulo p , and their reduction modulo p forms the modified equiprojectable decomposition of $Z(F \bmod p, G \bmod p)$.

3. A DIRECT ALGORITHM

In this section, we work over a field \mathbb{K} . We give an algorithm that takes as input F, G in $\mathbb{K}[X, Y]$ and computes the equiprojectable decomposition $\mathbf{T}_1, \dots, \mathbf{T}_s$ of $Z(F, G)$. If F and G have degree at most d , the running time is $O(d^3)$, that is, essentially the same as computing $\text{res}(F, G, Y)$ (we count all operations in \mathbb{K} at unit cost). This result is by no means surprising (a particular case appears in [23]) and certainly not enough to prove our main theorems. We will only use it as the initialization step of our lifting process.

The rest of this section is devoted to prove this proposition, following a few preliminaries.

PROPOSITION 1. *Let F, G be in $\mathbb{K}[X, Y]$, of total degree at most d . If the characteristic of \mathbb{K} is greater than $2d^2 + d + 1$, one can compute the equiprojectable decomposition of $Z(F, G)$ using $O(\mathbf{M}(d)\mathbf{M}(d^2)\log(d)^2)$ operations in \mathbb{K} .*

Regular GCDs and quotients. Let R be a nonzero, squarefree polynomial in $\mathbb{K}[X]$, and let F, G be in $\mathbb{K}[X, Y]$. A regular GCD of (F, G) modulo R is a non-critical decomposition of the ideal $\langle R, F, G \rangle$; a regular quotient of F by G modulo R is a non-critical decomposition of the ideal $\langle R, F \rangle : G$. If $\mathcal{F} = (\mathbf{T}_1, \dots, \mathbf{T}_s)$ is a regular GCD of (F, G) modulo R , with $\mathbf{T}_i = (U_i, V_i)$ for all i , and if F is monic in Y , then $\mathcal{F} = (\xi_1, \dots, \xi_s)$, with $\xi_i = (U_i, F/V_i \bmod U_i)$ for all i , is a regular quotient of F by G modulo R .

If F, G have degree at most d in Y , and R, F, G have degree at most m in X , then using the algorithm of [1], both operations can be done in time $O(\mathbf{M}(d)\mathbf{M}(m)\log(d)\log(m))$.

Radical computation. Regular quotients allow us to compute radicals. Let indeed $\mathbf{T} = (U, V)$ be a monic triangular set of bidegree (m, n) in $\mathbb{K}[X, Y]$, with U squarefree and with m and n less than the characteristic of \mathbb{K} ; we prove that $I = \langle U, V \rangle : \partial V / \partial Y$ is the radical of the ideal $\langle \mathbf{T} \rangle$.

Let I' be the extension of I in $\mathbb{K}[X, Y]$. Over \mathbb{K} , the assumption on m ensures that U is still squarefree, so the ideal $\langle U, V \rangle$ is the intersection of primary ideals of the form $\mathfrak{p}_i = \langle X - x_i, (Y - y_i)^{e_i} \rangle$, where $(x_i, y_i)_{1 \leq i \leq t}$ are the zeros of \mathbf{T} , and $e_i \in \mathbb{N}_{>0}$ is the multiplicity of the factor $Y - y_i$ in $V(x_i, Y)$. Then, I' is the intersection of the ideals $\mathfrak{p}_i : \partial V / \partial Y$, which we can rewrite as

$$I' = \bigcap_{1 \leq i \leq t} \langle X - x_i, (Y - y_i)^{e_i} \rangle : (e_i(Y - y_i)^{e_i - 1}).$$

The assumption on n implies that $e_i \neq 0$ in \mathbb{K} , so that I' is the intersection of the maximal ideals $\langle X - x_i, Y - y_i \rangle$; our claim is proved. As a consequence, under the above assumption on \mathbf{T} , we can compute a non-critical decomposition of the radical of $\langle \mathbf{T} \rangle$ in time $O(\mathbf{M}(n)\mathbf{M}(m)\log(n)\log(m))$.

After these preliminaries, we can turn to the algorithm. In what follows, J is the Jacobian determinant of (F, G) , $H = \gcd(F, G)$, $F^* = F/H$ and $G^* = G/H$. The idea is classical: we compute the resultant $R = \text{res}(F^*, G^*, Y)$, then a regular GCD of (F^*, G^*) modulo R ; make the result radical and finally we remove all points where J vanishes.

Step 0. We compute H, F^*, G^* as defined above. Corollary 11.9 in [16] gives an expected $O(d\mathbf{M}(d)\log(d))$ operations for computing H ; we briefly explain how to make it deterministic, up to an acceptable increase in running time (this is routine; some details are left to the reader).

Choosing $2d^2 + d + 1$ values x_i in \mathbb{K} , we compute $H_i = \gcd(F(x_i, Y), G(x_i, Y))$, $F_i^* = F(x_i, Y)/H_i$ and $G_i^* = G(x_i, Y)/H_i$. Lucky values of x_i are those where the leading coefficient of (say) F in Y and the resultant of (F^*, G^*) in Y are non-zero. We are sure to find at least $d^2 + 1$ of them; these will be those x_i 's where H_i has maximal degree. These are enough to reconstruct H, F^* and G^* by interpolation, hence a total running time of $O(\mathbf{M}(d)\mathbf{M}(d^2)\log(d))$.

Step 1. Compute the (nonzero) resultant $R = \text{res}(F^*, G^*, Y)$. Using Reischert's algorithm [28], this takes time $O(\mathbf{M}(d)\mathbf{M}(d^2)\log(d))$.

Step 2. Replace R by its squarefree part, which takes time $O(\mathbf{M}(d^2)\log(d))$. Note that $V(R, F^*, G^*) = V(F^*, G^*)$.

Step 3. Compute a regular GCD $\mathcal{F} = (\mathbf{T}_1, \dots, \mathbf{T}_s)$ of (F^*, G^*) modulo R , in time $O(\mathbf{M}(d)\mathbf{M}(d^2)\log(d)^2)$. Note that $V(\mathcal{F})$ is equal to $V(R, F^*, G^*)$, that is, $V(F^*, G^*)$.

Step 4. For $i = 1, \dots, s$, writing $\mathbf{T}_i = (U_i, V_i)$, compute a regular quotient of V_i by $\partial V_i / \partial Y$ modulo U_i .

Letting (m_i, n_i) be the bidegree of \mathbf{T}_i , the cost for each i is $O(\mathbf{M}(d)\mathbf{M}(m_i)\log(d)\log(m_i))$. Using the super-linearity of \mathbf{M} , the total is seen to be $O(\mathbf{M}(d)\mathbf{M}(d^2)\log(d)^2)$.

Let $\mathcal{S} = (\xi_1, \dots, \xi_t)$ be the union of all triangular sets obtained this way, with $\xi_i = (P_i, Q_i)$. Since d^2 is less than the characteristic of \mathbb{K} , this is also the case for all m_i and n_i . As a result, by the discussion above, $\langle \mathcal{S} \rangle$ is the defining ideal of $V(F^*, G^*)$; in particular, it is radical in $\mathbb{K}[X, Y]$.

Step 5. For $i = 1, \dots, t$, compute $J_i = J \bmod P_i$, where J is the Jacobian determinant of (F, G) . Using fast simultaneous modular reduction, this costs $O(d\mathbf{M}(d^2)\log(d))$.

Step 6. For $i = 1, \dots, t$, compute a regular quotient of Q_i by J_i modulo P_i ; again, the cost is $O(\mathbf{M}(d)\mathbf{M}(d^2)\log(d)^2)$. Let \mathcal{U} be the union of all resulting triangular sets; then, $\langle \mathcal{U} \rangle$ is the defining ideal of $V(F^*, G^*) - V(J)$, and one verifies that the latter set is $Z(F, G)$.

Step 7. Finally, apply the algorithm of Corollary 1 to \mathcal{U} to obtain the equiprojectable decomposition of $Z(F, G)$. Since $Z(F, G)$ has size at most d^2 , the cost is $O(\mathbf{M}(d^2)\log(d))$.

4. NORMAL FORM ALGORITHMS

We consider now the problem of reducing $F \in \mathbb{A}[X, Y]$ modulo several triangular sets. Our input is as follows:

- $\mathcal{F} = (\mathbf{T}_1, \dots, \mathbf{T}_s)$ is a non-critical family of monic triangular sets in $\mathbb{A}[X, Y]$, where $\mathbf{T}_i = (U_i, V_i)$ of bidegree (m_i, n_i) for all i and \mathbb{A} is a ring;

- $\mathcal{R} = (R_1, \dots, R_s)$ is the family of cofactors associated to \mathcal{T} , as in Subsection 2.2;
- F is in $\mathbb{A}[X, Y]$, of total degree less than d .

We make the following assumptions:

$$\sum_{i=1, \dots, s} m_i n_i \leq d^2 \quad \text{and} \quad n_i \leq d \text{ for all } i. \quad (\mathbf{H})$$

Then, the size of input and output is $\Theta(d^2)$ elements of \mathbb{A} .

Already for $s = 1$, in which case we write (m, n) instead of (m_1, n_1) , the difficulty of the problem can vary significantly: if both m and n are $O(d)$, Lemma 1 shows that the reduction can be done in optimal time $O(d^2)$; however, when $m \simeq d^2$ and $n \simeq 1$, that same lemma gives a sub-optimal $O(d^3)$.

In this section, we prove the following propositions, which give algorithms with better exponents. The first one applies over any ring \mathbb{A} ; it uses fast matrix multiplication to achieve an exponent $(\omega + 3)/2 \simeq 2.69$.

PROPOSITION 2. *Under assumption (H), there exists an algorithm that takes as input polynomials \mathcal{T} , \mathcal{R} and F as above and returns all $F \bmod \langle \mathbf{T}_i \rangle$, for i in $\{1, \dots, s\}$, using $O(M(d^2)d^{(\omega-1)/2} \log(d))$ operations in \mathbb{A} .*

When $\mathbb{A} = \mathbb{Z}/N\mathbb{Z}$, for some prime power N , better can be done in a *boolean* model: this second proposition shows that we can get arbitrarily close to linear time (in the boolean model, input and output sizes are $\Theta(d^2 \log(N))$).

PROPOSITION 3. *Under assumption (H), for any $\varepsilon > 0$, there exists an algorithm that takes as input a prime power N , and polynomials \mathcal{T} , \mathcal{R} and F as above, all with coefficients in $\mathbb{Z}/N\mathbb{Z}$, and returns all $F \bmod \langle \mathbf{T}_i \rangle$, for i in $\{1, \dots, s\}$, using $d^{2+\varepsilon} O(\log(N))$ bit operations.*

4.1 Reduction modulo one triangular set

We first discuss a simplified version of the problem, where we reduce F modulo a single monic triangular set. In other words, we take $s = 1$; then, we simply write $\mathbf{T} = (U, V)$ instead of \mathbf{T}_1 , and we denote its bidegree by (m, n) instead of (m_1, n_1) . The polynomial F is in $\mathbb{A}[X, Y]$, of degree less than d ; thus our assumptions are the following:

$$mn \leq d^2 \quad \text{and} \quad n \leq d. \quad (\mathbf{H}')$$

In [27], Poteaux and Schost give two algorithms computing $F \bmod \langle \mathbf{T} \rangle$. The first one, originating from [26, Ths. 4-6], applies only in a boolean model, when $\mathbb{A} = \mathbb{Z}/p\mathbb{Z}$ for a prime p . Only a small change is needed to make it work modulo a prime power $N = p^\ell$. In both cases, when the base ring, or field, is too small, we need to enlarge it, by adding elements whose differences are invertible. In our case, we extend the basering $\mathbb{Z}/N\mathbb{Z}$ by a polynomial that is irreducible modulo p (since if $x - x'$ is a unit modulo p , it is a unit modulo N). The analysis of [26, Ths. 4-6] remains valid with this minor modification, and yields the following result.

PROPOSITION 4. [26, 27] *Under assumption (H'), for any $\varepsilon > 0$, there exists an algorithm that takes as input a prime power N and F and \mathbf{T} as above, with coefficients in $\mathbb{Z}/N\mathbb{Z}$, and returns $F \bmod \langle \mathbf{T} \rangle$ using $d^{2+\varepsilon} O(\log(N))$ bit operations.*

Since in this case the input and output size is $\Theta(d^2 \log(N))$ bits, this algorithm is close to being optimal.

If we consider the question over an abstract ring \mathbb{A} , no quasi-optimal algorithm is known. Under assumption (H'), the second algorithm of [27] runs in time $O(d^{\omega+1})$; this is subquadratic in the size d^2 of the problem, but worse than $O(d^3)$. The following proposition gives an improved result.

PROPOSITION 5. *Under assumption (H'), there exists an algorithm that takes as input F and \mathbf{T} as above and returns $F \bmod \langle \mathbf{T} \rangle$ using $O(M(d^2)d^{(\omega-1)/2})$ operations in \mathbb{A} .*

The rest of this subsection is devoted to prove this proposition. As in [27], we use a baby steps / giant steps approach inspired by Brent and Kung's algorithm [6], but with a slightly more refined subdivision scheme.

Let thus F be in $\mathbb{A}[X, Y]$, of total degree less than d , and let $\mathbf{T} = (U, V)$ be of bidegree (m, n) . The steps of the algorithm are given below: they consist in computing some powers of Y modulo $\langle \mathbf{T} \rangle$ (baby steps, at Step 3), doing products of matrices with entries in $\mathbb{A}[X, Y]$ (Step 4), and concluding using Horner's scheme (giant steps, at Step 6).

Step 0. Replace F by $F \bmod U$; as a consequence, we can assume $F \in \mathbb{A}[X, Y]_{r,d}$, with $r = \min(d, m)$. For future use, note that $mn \leq rd$: if $r = d$, this is because $mn \leq d^2$. Else, $r = m$, and the claim follows from the fact that $n \leq d$.

We do d reductions of polynomials of degree less than d by a polynomial of degree m in $\mathbb{A}[X]$; this is free if $d < m$ and costs $O(dM(d))$ otherwise.

Step 1. Let $t = \lceil d/n \rceil - 1$ and write F as $F = F_0 + F_1 Y^n + \dots + F_t Y^{nt}$, with all F_i in $\mathbb{A}[X, Y]_{r,n}$.

Step 2. Let $\rho = \lfloor d/n^{1/2} \rfloor$ and $\mu = \lceil (t+1)/\rho \rceil$; note that since $d \geq n$, $\rho \geq 1$ so μ is well-defined. Furthermore, both ρ and μ are $O(d/n^{1/2})$ and $\rho\mu \geq t+1$. Set up the $(\mu \times \rho)$ matrix $\mathbf{M}_1 = [F_{i\rho+j}]_{0 \leq i < \mu, 0 \leq j < \rho}$ with entries in $\mathbb{A}[X, Y]_{r,n}$, where we set $F_k = 0$ for $k > t$.

Step 3. For $i = 0, \dots, \rho$, compute $\sigma_i = Y^{ni} \bmod \langle \mathbf{T} \rangle$. Cost: since $\deg(V, Y) = n$, $\sigma_1 = (Y^n \bmod \langle \mathbf{T} \rangle)$ is equal to $Y^n - V$, so computing it takes time $O(mn)$. Then, it takes time $O(\rho M(mn))$ to deduce all other σ_i 's.

Step 4. Let $\nu = \lceil m/r \rceil - 1$; for $i = 0, \dots, \rho - 1$, write $\sigma_i = \sigma_{i,0} + \sigma_{i,1} X^r + \dots + \sigma_{i,\nu} X^{r\nu}$, with all $\sigma_{i,j}$ in $\mathbb{A}[X, Y]_{r,n}$. Build the $\rho \times (\nu + 1)$ matrix $\mathbf{M}_2 = [\sigma_{i,j}]_{0 \leq i < \rho, 0 \leq j \leq \nu}$ and compute $\mathbf{M} = \mathbf{M}_1 \mathbf{M}_2$.

Cost: we have seen that $mn \leq rd$, so that $m/r \leq d/n$ and thus $\nu = O(d/n)$. Using the bounds on ρ, μ, ν and Lemma 2, we deduce that the cost is $O(M(rn)(e/n)^{(\omega+1)/2})$.

Step 5. Let $[m_{i,j}]_{0 \leq i < \mu, 0 \leq j \leq \nu}$ be the entries of \mathbf{M} , which are in $\mathbb{A}[X, Y]_{2r-1, 2n-1}$. For $i = 0, \dots, \mu - 1$, compute $G_i = m_{i,0} + m_{i,1} X^r + \dots + m_{i,\nu} X^{r\nu}$ and $H_i = G_i \bmod \langle \mathbf{T} \rangle$. Because $m_{i,j} = F_{i\rho} \sigma_{0,j} + F_{i\rho+1} \sigma_{1,j} + \dots + F_{(i+1)\rho-1} \sigma_{\rho-1,j}$, we deduce that $G_i = F_{i\rho} \sigma_0 + F_{i\rho+1} \sigma_1 + \dots + F_{(i+1)\rho-1} \sigma_{\rho-1}$. Since $\sigma_j = Y^{nj} \bmod \langle \mathbf{T} \rangle$ for all j , this proves that $H_i = F_{i\rho} + F_{i\rho+1} Y^n + \dots + F_{(i+1)\rho-1} Y^{n(\rho-1)} \bmod \langle \mathbf{T} \rangle$.

Computing a single G_i takes $O(r\nu)$ additions in \mathbb{A} , which is $O(mn)$ since $r\nu = O(m)$. By construction, G_i is in $\mathbb{A}[X, Y]_{r(\nu+2)-1, 2n-1}$; since $r(\nu+2) = O(m)$, Lemma 1 implies that reducing G_i to obtain H_i takes time $O(M(mn))$. The total for all G_i 's is $O(\mu M(mn))$.

Step 6. Compute $H = H_0 + H_1 \sigma_\rho + \dots + H_{\mu-1} \sigma_\rho^{\mu-1} \bmod \langle \mathbf{T} \rangle$ using Horner's scheme; the expression given above for the polynomials H_i implies that $H = F \bmod \langle \mathbf{T} \rangle$. Cost: $O(\rho)$ additions and multiplications modulo \mathbf{T} , each of which costs $O(M(mn))$ operations in \mathbb{A} .

Summary. Summing all contributions, we obtain

$$O\left(M(d)d + M(mn)(d/n)^{1/2} + M(rn)(d/n)^{(\omega+1)/2}\right).$$

The first two terms are easily seen to be $O(M(d^2)d^{1/2})$. To deal with the last term, note that $r \leq d$ implies $M(rn) \leq M(dn)$, and the super-linearity of M implies that $M(dn) \leq M(d^2)n/d$. Thus, the third term is $O(M(d^2)(d/n)^{(\omega-1)/2})$, which is $O(M(d^2)d^{(\omega-1)/2})$. Proposition 5 is proved.

4.2 Reduction modulo several triangular sets

We now prove Proposition 2 and 3. To simplify the presentation, we give details for the first result (in the algebraic model); the proof in the boolean model requires no notable modification (just use Proposition 4 instead of 5 below).

Let thus $\mathcal{T} = (\mathbf{T}_1, \dots, \mathbf{T}_s)$ be monic triangular sets of the form $\mathbf{T}_i = (U_i, V_i)$, with coefficients in \mathbb{A} and bidegrees (m_i, n_i) . We also assume that the cofactors $\mathcal{R} = (R_1, \dots, R_s)$ are given. Given F in $\mathbb{A}[X, Y]$ of degree less than d , we consider the question of reducing F modulo all $\langle \mathbf{T}_i \rangle$, under assumption (H). Our proof distinguishes three cases, from the most particular to the general case.

Identical n_i 's. Assume first that there exists n such that $n_i = n$ for all i . Writing $m = m_1 + \dots + m_s$, Lemma 3 shows that we can build a monic triangular set \mathbf{T} in $\mathbb{A}[X, Y]$ of bidegree (m, n) , such that the ideal $\langle \mathbf{T} \rangle$ is the intersection of all $\langle \mathbf{T}_i \rangle$, in time $O(nM(m) \log(m))$

To compute all $F \bmod \langle \mathbf{T}_i \rangle$, because (H) implies $mn \leq d^2$, we first compute $H = F \bmod \langle \mathbf{T} \rangle$ using Proposition 5, in time $O(M(d^2)d^{(\omega-1)/2})$. Then, we use Lemma 3 to reduce H modulo all $\langle \mathbf{T}_i \rangle$ in time $O(nM(m) \log(m))$. Since $nM(m)$ is $O(M(d^2))$, the cost of this step is negligible.

Similar n_i 's. We now relax the assumption that all n_i 's are the same; instead, we assume that there exists n such that $n_i \in \{n, \dots, 2n-1\}$ for all i ; as above, we write $m = m_1 + \dots + m_s$, and we introduce $n' = 2n-1$.

For $i = 1, \dots, s$, define $V_i^* = Y^{n'-n_i}V_i$ and $\mathbf{T}_i^* = (U_i, V_i^*)$, so that the \mathbf{T}_i^* 's are monic triangular sets of bidegrees (m_i, n') . These new triangular sets and F may not satisfy (H) anymore, but they will, provided we replace d by $d' = 2d$. Indeed, notice that the inequality $n' \leq 2n_i$ holds for all i ; using (H), this yields

$$\sum_{i=1, \dots, s} m_i n' \leq 2 \sum_{i=1, \dots, s} m_i n_i \leq 2d^2 \leq d'^2,$$

and similarly $n' \leq d'$. The algorithm in the previous paragraph then allows us to compute all $H_i = F \bmod \langle \mathbf{T}_i^* \rangle$, still in time $O(M(d^2)d^{(\omega-1)/2})$.

Then, for all i , we compute the remainder $H_i \bmod \langle \mathbf{T}_i \rangle$. Using Lemma 1, this can be done in time $O(M(m_i n))$ for each i , for a negligible total cost of $O(M(mn)) \subset O(M(d^2))$.

Arbitrary degrees. Finally, we drop all assumptions on the degrees n_i . Instead, we partition the set $S = \{1, \dots, s\}$ into S_0, \dots, S_κ such that i is in S_ℓ if and only if n_i is in $\{2^\ell, \dots, 2^{\ell+1}-1\}$. Because all n_i satisfy $n_i \leq d$, κ is in $O(\log(d))$. We write as usual $m = m_1 + \dots + m_s$.

We are going to apply the algorithm of the previous paragraph to all S_ℓ independently. Remark that if all \mathbf{T}_i and F satisfy assumption (H), the subset $\{\mathbf{T}_i \mid i \in S_\ell\}$ and F still satisfy this assumption. Let us thus fix $\ell \in \{0, \dots, \kappa\}$. The only thing that we need to take care of are the cofactors required for Chinese Remaindering. As input, we assumed

that we know all $R_i = 1/(U_1 \cdots U_{i-1} U_{i+1} \cdots U_s) \bmod U_i$; what we need are the inverses $R_{\ell, i} = 1/\prod_{i' \in S_\ell, i' \neq i} U_{i'} \bmod U_i$ for $i \in S_\ell$. These polynomials are computed easily: first, we form the product $B_\ell = \prod_{i \notin S_\ell} U_i$; using [16, Lemma 10.4], this takes $O(M(m) \log(m))$ operations in \mathbb{A} . Then, we reduce B_ℓ modulo all U_i , for $i \in S_\ell$, for the same amount of time as above. Finally, we obtain all $R_{\ell, i}$ as $R_{\ell, i} = R_i B_\ell \bmod U_i$; the time needed for these products is $O(M(m))$.

Once the polynomials $R_{\ell, i}$ are known, the algorithm above gives us $F \bmod \langle \mathbf{T}_i \rangle$, for $i \in S_\ell$, in time $O(M(d^2)d^{(\omega-1)/2})$; this dominates the cost of computing the polynomials $R_{\ell, i}$. Summing over ℓ finishes the proof of Prop. 2.

5. PROOF OF THE MAIN RESULTS

We assume here that we are one of the cases $\mathbb{A} = k[T]$ or $\mathbb{A} = \mathbb{Z}$ and we prove Theorems 1 and 2. Given F, G in $\mathbb{A}[X, Y]$ and writing as before $\mathcal{T} = (\mathbf{T}_1, \dots, \mathbf{T}_s)$ for the equiprojectable decomposition of $Z(F, G)$ and $\mathcal{C} = (\mathbf{C}_1, \dots, \mathbf{C}_s)$ for its modified version, we show here how to compute the latter.

The algorithm follows the template given in [10]: compute the equiprojectable decomposition modulo a randomly chosen maximal ideal \mathfrak{m} of \mathbb{A} , lift it modulo \mathfrak{m}^N , for N large enough, and reconstruct all rational fractions that appear as coefficients in \mathcal{C} from their expansion modulo \mathfrak{m}^N .

We suppose that $\lambda(F), \lambda(G) \leq \ell$ and $\deg(F), \deg(G) \leq d$, where λ is the length function defined in the introduction. For $i \leq s$, we write (m_i, n_i) for the bidegree of \mathbf{T}_i ; then, we have the upper bound $\sum_{i \leq s} m_i n_i \leq d^2$; besides, each n_i is at most d , so assumption (H) of Section 4 holds.

5.1 One lifting step

Here, \mathfrak{m} is a maximal ideal of \mathbb{A} ; we assume that none of the denominators of the coefficients of the polynomials in \mathcal{T} vanishes modulo \mathfrak{m} . Thus, for $N \geq 1$, we can define $\mathcal{T}_N = \mathcal{T} \bmod \mathfrak{m}^N$ by reducing all coefficients of $\mathcal{T} \bmod \mathfrak{m}^N$. Given \mathcal{T}_N , we show how to compute \mathcal{T}_{2N} ; this will be the core of our main algorithm. We start by describing some basic operations in $\mathbb{A}_N = \mathbb{A}/\mathfrak{m}^N$ (when $N = 1$, we also use the notation k to denote the residual field \mathbb{A}/\mathfrak{m}).

For complexity analyzes, we assume that $\mathbb{A} = k[T]$ and that \mathfrak{m} has the form $\langle T - t_0 \rangle$, for some t_0 in k ; we discuss the case $\mathbb{A} = \mathbb{Z}$ afterwards. Remark in particular that operations $(+, -, \times)$ in \mathbb{A}_N can be done in $O(M(N))$ operations in k .

Univariate inversion. Given Q monic of degree m in $\mathbb{A}_N[X]$ and $F \in \mathbb{A}_N[X]$ of degree less than m , consider the problem of computing $1/F$ in $\mathbb{A}_N[X]/\langle Q \rangle$, if it exists.

This is done by computing the inverse modulo \mathfrak{m} (*i.e.*, over $k[X]$) by an extended GCD algorithm and lifting it by Newton iteration [16, Ch. 9]. The first step uses $O(M(m) \log(m))$ operations in k , the second one $O(M(m)M(N))$.

Bivariate inversion. Given a monic triangular set \mathbf{T} in $\mathbb{A}_N[X, Y]$ of bidegree (m, n) and $F \in \mathbb{A}_N[X, Y]_{m, n}$, consider the computation of $1/F$ in $\mathbb{A}_N[X, Y]/\langle \mathbf{T} \rangle$, assuming it exists.

We use the same approach as above, but with bivariate computations. For inversion modulo \mathfrak{m} , we use [9, Prop. 6], which gives a cost $O(M(m)M(n) \log(m)^3 \log(n)^3)$. The lifting now takes $O(M(mn)M(N))$.

Lifting \mathcal{T}_N . We can now explain the main algorithm, called Lift in the next subsection. In what follows, we write $\mathcal{T}_N = (\mathbf{T}_{N,1}, \dots, \mathbf{T}_{N,s})$; all computations take place over \mathbb{A}_{2N} .

Step 0. First, as in the proof of Corollary 1, we compute the cofactors \mathcal{R}_N associated to \mathcal{T}_N using [16, Algo. 10.18]; this time, though, we work over the ring \mathbb{A}_{2N} . Steps 1 and 2 of that algorithm work over any ring; Step 3, which computes inverses modulo the polynomials $\mathbf{T}_{N,j}$, is dealt with using the remarks made above on univariate inversion. Because $\mathbf{T}_{N,j}$ has bidegree (m_j, n_j) for all j , with $\sum_{j \leq s} m_j n_j \leq d^2$, the total cost is $O(M(d^2)M(N) \log(d))$ operations in k .

Step 1. We will use formulas from [31] to lift from \mathcal{T}_N to \mathcal{T}_{2N} . First, we reduce the polynomials F, G and the entries of their Jacobian matrix J modulo \mathfrak{m}^{2N} ; as a result, we will now see these polynomials as elements of $\mathbb{A}_{2N}[X, Y]$.

Over $\mathbb{A} = k[T]$, the assumption that $\lambda(F), \lambda(G) \leq \ell$ means that F and G have degree at most ℓ in T ; we are reducing them modulo the polynomial $(T - t_0)^{2N}$. The time for one coefficient reduction is $O(M(\ell))$, since when $2N > \ell$, no work is needed. The total time is $O(d^2 M(\ell))$.

Step 2. We compute $F_{N,j} = F \bmod \langle \mathbf{T}_{N,j} \rangle$ over $\mathbb{A}_{2N}[X, Y]$ for all $j \in \{1, \dots, s\}$, as well as $G_{N,j} = G \bmod \langle \mathbf{T}_{N,j} \rangle$ and $J_{N,j} = J \bmod \langle \mathbf{T}_{N,j} \rangle$. This is the most costly part of the algorithm: because we know the cofactors \mathcal{R}_N associated to \mathcal{T}_N , and because assumption (H) of Section 4 is satisfied, Proposition 2 shows that one can compute all $F_{N,j}$ using $O(M(d^2)M(N)d^{(\omega-1)/2} \log(d))$ operations in k . The same holds for all $G_{N,j}$ and $J_{N,j}$.

Step 3. Finally, for all j , we compute the (2×2) Jacobian matrix $M_{N,j}$ of $\mathbf{T}_{N,j}$ in $\mathbb{A}_{2N}[X, Y]$ and the vector

$$\delta_{N,j} = M_{N,j} J_{N,j}^{-1} \begin{bmatrix} F_{N,j} \\ G_{N,j} \end{bmatrix} \text{ over } \mathbb{A}_{2N}[X, Y] / \langle \mathbf{T}_{N,j} \rangle.$$

Proposition 4 in [31] then proves that $\mathbf{T}_{2N,j} = \mathbf{T}_{N,j} + \delta_{N,j}^*$, where $\delta_{N,j}^*$ is the canonical preimage of $\delta_{N,j}$ over $\mathbb{A}_{2N}[X, Y]$.

The dominant cost is the inversion of the matrices $J_{N,j}$. By the remark above, the cost for a given j is $O(M(m_j)M(n_j) \log(m_j)^3 \log(n_j)^3 + M(m_j n_j)M(N))$; summing over j , this step is negligible compared to Step 2.

Summary. When $\mathbb{A} = k[T]$, the cost of deducing \mathcal{T}_{2N} from \mathcal{T}_N is $O(d^2 M(\ell) + M(d^2)M(N)d^{(\omega-1)/2} \log(d))$ operations in k , which is $O^-(d^2 \ell + d^{(\omega+3)/2} N)$.

When $\mathbb{A} = \mathbb{Z}$ and $\mathfrak{m} = \langle p \rangle$, for a prime p , the algorithm does not change, but the complexity analysis does. Using the fact that computations modulo p^r can be done in $O^-(\log(p^r))$ bit operations, and using Proposition 3, we obtain a cost of $d^{2+\varepsilon} O^-(\ell + N \log(p))$ bit operations, for any $\varepsilon > 0$.

5.2 Main algorithm

We will now analyze the main steps of the following algorithm, proving our main theorems. For simplicity, we suppose that $\mathbb{A} = k[T]$; the modifications for $\mathbb{A} = \mathbb{Z}$ follow.

Input: F, G in $\mathbb{A}[X, Y]$, $\mathfrak{m} \subset \mathbb{A}$, $\ell \in \mathbb{N}$, $d \in \mathbb{N}$
Output: $\mathcal{C} = (\mathbf{C}_1, \dots, \mathbf{C}_s)$
(1) $\mathcal{T}_1 \leftarrow Z(F \bmod \mathfrak{m}, G \bmod \mathfrak{m})$
(2) $i \leftarrow 1$
(3) **while** $\lambda(\mathfrak{m}^{2^i}) < 4d\ell + 48d^2$ **do**
 (3.a) $\mathcal{T}_{2^i} \leftarrow \text{Lift}(\mathcal{T}_{2^{i-1}}, F, G)$
 (3.b) $i \leftarrow i + 1$
end
(4) $\mathcal{C}_{2^{i-1}} \leftarrow \text{Convert}(\mathcal{T}_{2^{i-1}})$
(5) **return** $\text{RationalReconstruction}(\mathcal{C}_{2^{i-1}})$

Step 1. Over $\mathbb{A} = k[T]$, the maximal ideal \mathfrak{m} has the form $\mathfrak{m} = \langle T - t_0 \rangle$, for some $t_0 \in k$. Reducing F and G modulo \mathfrak{m} takes $O(d^2)$ operations in k by the plain algorithm.

We assume that t_0 is not a root of the polynomial A defined in Lemma 4. By assumption, the cardinality of k is at least twice the degree of A , so choosing t_0 at random, our assumption is satisfied with probability at least $1/2$.

We use the algorithm of Section 3 over k to compute the equiprojectable decomposition \mathcal{T}_1 of $Z(F \bmod \mathfrak{m}, G \bmod \mathfrak{m})$; under our assumption on t_0 , \mathcal{T}_1 coincides with \mathcal{T} mod \mathfrak{m} . This step takes $O(M(d)M(d^2) \log(d)^2)$ operations in k .

Step 3. We saw in the introduction that over either $\mathbb{A} = k[T]$ or $\mathbb{A} = \mathbb{Z}$, all polynomials U_i and N_i in \mathcal{C} satisfy $\lambda(U_i), \lambda(N_i) \leq 2d\ell + 24d^2$. In order to reconstruct the coefficients of these polynomials from their expansion modulo \mathfrak{m}^N , it is thus enough to ensure that $2(2d\ell + 24d^2) \leq \lambda(\mathfrak{m}^N)$; this accounts for the bound in the **while** loop. If we wanted to compute \mathcal{T} instead, the bound would be of order $d^3 \ell + d^4$.

Step 3.a. For each value of i , we call the algorithm described in the previous subsection; we saw that the cost is $O(d^2 M(\ell) + M(d^2)M(2^i)d^{(\omega-1)/2} \log(d))$ operations in k . The last value i_0 of the loop index is such that $2^{i_0} < 4d\ell + 48d^2 \leq 2^{i_0+1}$. We deduce the total running time:

$$O\left(d^2 M(\ell) \log(\ell) + M(d^2)M(d\ell + d^2)d^{(\omega-1)/2} \log(d)\right).$$

Step 4. We obtain $\mathcal{C} \bmod \mathfrak{m}^{2^{i_0}}$ from $\mathcal{T} \bmod \mathfrak{m}^{2^{i_0}}$ by applying subroutine **Convert**, which does the following: for all $i \leq s$, \mathbf{T}_i has the form (U_i, V_i) and $\mathbf{C}_i = (U_i, N_i)$, with $N_i = V_i U_i' \bmod U_i$, over the ring $\mathbb{A}_{2^{i_0}}[X, Y]$. The cost is negligible compared to that of the lifting.

Step 5. Finally, **RationalReconstruction** recovers the rational coefficients appearing in \mathcal{C} from their expansion modulo $\mathfrak{m}^{2^{i_0}}$ (the index i_0 was chosen such that this precision is sufficient). There are $O(d^2)$ coefficients, each of them having numerator and denominator of degree $O(d\ell + d^2)$, so the total time is $O(d^2 M(d\ell + d^2) \log(d\ell))$ operations in k .

Summary. Summing all previous costs, we see that the total time admits the upper bound claimed in Theorem 1,

$$O\left(M(d^2)M(d\ell + d^2)d^{(\omega-1)/2} \log(d\ell)\right).$$

Over $\mathbb{A} = \mathbb{Z}$, \mathfrak{m} is of the form $\langle p \rangle$, for a suitable p chosen as follows: let $B = 6 \cdot 8d^5(3\ell + 10 \log(d) + 22)$. Using [16, Th. 18.10], we can compute in time $O^-(\log(d\ell))$ an integer $p \in [B+1, \dots, 2B]$ such that with probability at least $1/2$, p is prime and does not divide the integer A of Lemma 5. We apply the same algorithm as above (in particular, since $p \geq B$, the computation modulo p will satisfy the requirement on the characteristic of the field $k = \mathbb{Z}/p\mathbb{Z}$ of Proposition 1).

Using the analysis in the previous subsection and the bounds on the bit-size of the output, it is straightforward to derive an upper bound of $d^{2+\varepsilon} O^-(d\ell + d^2)$ bit operations, for any $\varepsilon > 0$. Up to doubling ε , the polylogarithmic terms can be discarded, and we get the result of Theorem 2.

6. EXPERIMENTAL RESULTS

We report here on preliminary results obtained with an experimental implementation of our main algorithm in the case $\mathbb{A} = \mathbb{Z}$, based on Shoup's NTL [32]. Although Theorem 2 features the best complexity, it relies ultimately on an idea of Kedlaya and Umans' [20], and we are not aware

of an efficient implementation of it, nor do we know how to derive one. Instead, we used the baby steps / giant steps idea underlying Theorem 1, which applies over any ring.

Our prototype is limited to inputs with word-size coefficients, and handles only the generic case described in the introduction, with only one triangular set of the form $U(X), Y - \eta(X)$ in \mathcal{S} . We did implement some classical optimizations not described above in the lifting process, such as halving the precision needed for the Jacobian matrix [18, § 4.4]. In the size ranges below, we choose our prime p of about 50 bits (this agrees with the bound given in the previous section; also, in this generic case, it is easy to verify that such a prime is “lucky”). Our implementation does polynomial matrix multiplication with exponent $\omega = 3$. Nevertheless, this step was carefully implemented, using FFT techniques for evaluation / interpolation and fast multiplication of matrices modulo small primes.

We compare our results to a Chinese Remainder approach that computes the resultant and the last subresultant modulo many primes. NTL only computes resultants, so we used an implementation of the fast subresultant algorithm already used in [17] that mimics NTL’s built-in resultant implementation. We give timings for the two kinds of modular arithmetic supported by NTL, `ZZ_p` and `lzz_p`, for respectively “large” primes and word-size primes. The latter is usually faster, as confirmed below, but the former allows us to choose fewer but larger primes for modular computations, which may be advantageous.

The following table shows timings needed to compute the output modulo p^N , where p is a 50 bit prime, and N is a power of 2, using these various approaches; inputs are random dense polynomials, and correctness was verified by comparing that the results of all approaches agreed. On these examples, our lifting algorithm does better than our CRT-based resultant implementation. The next step in our implementation will be to confirm whether this is still the case when we lift the general position assumption.

degree	precision	Lifting	CRT, <code>ZZ_p</code>	CRT, <code>lzz_p</code>
100	32	295.67	1474.88	899.48
100	64	558.75	2949.76	1798.96
100	128	1241.4	5899.52	3597.93
120	32	421.78	2711.36	1990.40
120	64	774.14	5422.72	3980.80
120	128	1728.1	10845.44	7961.60
140	32	818.97	4902.24	2671.89
140	64	1486.35	9804.48	5343.79
140	128	3045.91	19608.96	10687.59
160	32	1072.1	7610.6	5293.64
160	64	1896.64	15221.2	10587.28
160	128	3958.17	30442.4	21174.56
180	32	1394.61	11121.48	6541.90
180	64	2399.61	22242.96	13097.57
180	128	4951.37	44485.92	26195.15

Acknowledgements. We acknowledge support from NSERC, the CRC program and ANR grant HPAC (ANR-11-BS02-013). We thank all reviewers for their remarks.

7. REFERENCES

- [1] C. J. Accettella, G. M. D. Corso, and G. Manzini. Inversion of two level circulant matrices over \mathbb{Z}_p . *Lin. Alg. Appl.*, 366:5 – 23, 2003.
- [2] P. Aubry, D. Lazard, and M. M. Maza. On the theories of triangular sets. *JSC*, 28(1,2):45–124, 1999.
- [3] P. Aubry and A. Valibouze. Using Galois ideals for computing relative resolvents. *JSC*, 30(6):635–651, 2000.
- [4] E. Berberich, P. Emeliyanenko, and M. Sagraloff. An elimination method for solving bivariate polynomial systems: Eliminating the usual drawbacks. In *ALLENEX*, pages 35–47. SIAM, 2011.
- [5] A. Bostan, C.-P. Jeannerod, and E. Schost. Solving structured linear systems with large displacement rank. *Theor. Comput. Sci.*, 407(1-3):155–181, 2008.
- [6] R. P. Brent and H. T. Kung. Fast algorithms for manipulating formal power series. *J. ACM*, 25(4):581–595, 1978.
- [7] L. Cerlienco and M. Mureddu. From algebraic sets to monomial linear bases by means of combinatorial algorithms. *Discrete Math.*, 139:73–87, 1995.
- [8] J. Cheng, S. Lazard, L. M. Peñaranda, M. Pouget, F. Rouillier, and E. P. Tsigaridas. On the topology of real algebraic plane curves. *Mathematics in Computer Science*, 4(1):113–137, 2010.
- [9] X. Dahan, X. Jin, M. Moreno Maza, and É. Schost. Change of order for regular chains in positive dimension. *Theor. Comput. Sci.*, 392(1–3):37–65, 2008.
- [10] X. Dahan, M. Moreno Maza, É. Schost, W. Wu, and Y. Xie. Lifting techniques for triangular decompositions. In *ISSAC’05*, pages 108–115. ACM, 2005.
- [11] X. Dahan and É. Schost. Sharp estimates for triangular sets. In *ISSAC*, pages 103–110. ACM, 2004.
- [12] D. Diochnos, I. Emiris, and E. Tsigaridas. On the asymptotic and practical complexity of solving bivariate systems over the reals. *JSC*, 44(7):818–835, 2009.
- [13] M. El Kahoui. Topology of real algebraic space curves. *J. Symb. Comput.*, 43(4):235–258, 2008.
- [14] P. Emeliyanenko and M. Sagraloff. On the complexity of solving a bivariate polynomial system. In *ISSAC’12*. ACM, 2012.
- [15] I. Z. Emiris and E. P. Tsigaridas. Real solving of bivariate polynomial systems. In *CASC*, pages 150–161. Springer, 2005.
- [16] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 1999.
- [17] P. Gaudry and É. Schost. Construction of secure random curves of genus 2 over prime fields. In *Eurocrypt’04*, volume 3027 of *LNCS*, pages 239–256. Springer, 2004.
- [18] M. Giusti, G. Lecerf, and B. Salvy. A Gröbner free alternative for polynomial system solving. *J. Comp.*, 17(1):154–211, 2001.
- [19] X. Huang and V. Pan. Fast rectangular matrix multiplication and applications. *J. Complexity*, 14(2):257–299, 1998.
- [20] K. S. Kedlaya and C. Umans. Fast polynomial factorization and modular composition. *SICOMP*, 40(6):1767–1802, 2011.
- [21] T. Krick, L. M. Pardo, and M. Sombra. Sharp estimates for the arithmetic Nullstellensatz. *Duke Math. J.*, 109:521–598, 2001.
- [22] D. Lazard. Ideal bases and primary decomposition: case of two variables. *J. Symbolic Comput.*, 1(3):261–270, 1985.
- [23] X. Li, M. Moreno Maza, and V. Pan. Computations modulo regular chains. In *ISSAC*, pages 239–246. ACM, 2009.
- [24] X. Li, M. Moreno Maza, and É. Schost. Fast arithmetic for triangular sets: from theory to practice. *JSC*, 44:891–907, 2009.
- [25] C. Pascal and E. Schost. Change of order for bivariate triangular sets. In *ISSAC’06*, pages 277–284. ACM, 2006.
- [26] A. Poteaux and É. Schost. Modular composition modulo triangular sets and applications. *Comput. Comp. (to appear)*.
- [27] A. Poteaux and É. Schost. On the complexity of computing with zero-dimensional triangular sets. *JSC (to appear)*.
- [28] D. Reischert. Asymptotically fast computation of subresultants. In *ISSAC’97*, pages 233–240. ACM, 1997.
- [29] F. Rouillier. Solving zero-dimensional systems through the rational univariate representation. *Appl. Algebra Engrg. Comm. Comput.*, 9(5):433–461, 1999.
- [30] F. Rouillier. On solving systems of bivariate polynomials. In *ICMS*, volume 6327 of *LNCS*, pages 100–104. Springer, 2010.
- [31] É. Schost. Computing parametric geometric resolutions. *Appl. Algebra Engrg. Comm. Comput.*, 13(5):349–393, 2003.
- [32] V. Shoup. A new polynomial factorization algorithm and its implementation. *JSC*, 20(4):363–397, 1995.
- [33] V. Vassilevska Williams. Multiplying matrices faster than Coppersmith-Winograd. *STOC ’12*, pages 887–898. ACM, 2012.