Lifting Techniques for Triangular Decompositions

Xavier Dahan
LIX, École polytechnique
91128 Palaiseau, France
dahan@lix.polytechnique.fr

Marc Moreno Maza
ORCCA, University of Western Ontario (UWO)
London, Ontario, Canada
moreno@orcca.on.ca

Érich Schost
LIX, École polytechnique
91128 Palaiseau, France
schost@lix.polytechnique.fr

Wenyuan Wu
ORCCA, UWO
wwu@orcca.on.ca

Yuzhen Xie
ORCCA, UWO
yxie@orcca.on.ca

ABSTRACT

We present lifting techniques for triangular decompositions of zero-dimensional varieties, that extend the range of the previous methods. We discuss complexity aspects, and report on a preliminary implementation. Our theoretical results are comforted by these experiments.

Categories and Subject Descriptors: I.I.2 [Computing Methodologies]: Symbolic and Algebraic Manipulation – *Algebraic Algorithms*

General Terms: Algorithms, experimentation, theory.

Keywords: Polynomial systems, triangular sets, Hensel lifting.

1. INTRODUCTION

Modular methods for computing polynomial GCDs and solving linear algebra problems have been well-developed for several decades, see [12] and the references therein. Without these methods, the range of problems accessible to symbolic computations would be dramatically limited. Such methods, in particular Hensel lifting, also apply to solving polynomial systems. Standard applications are the resolution of systems over $\mathbb{Q}$ after specialization at a prime, and over the rational function field $k(Y_1, \dots, Y_m)$ after specialization at a point $(y_1, \dots, y_m)$. These methods have already been put to use for Gröbner bases [26, 1] and primitive element representations, starting from [13], and refined notably in [14].

Triangular decompositions are well-suited to many practical problems: see some examples in [3, 11, 24]. In addition, these techniques are commonly used in differential algebra [4, 15]. Triangular decompositions of polynomial systems can be obtained by various algorithms [16, 18, 21] but none of them uses modular computations, restricting their practical efficiency. Our goal in this paper is to discuss such techniques, extending the preliminary results of [24].

Let us introduce the notation used below. If k is a perfect field (e.g., $\mathbb{Q}$ or a finite field), a *triangular set* is a family $T_1(X_1), T_2(X_1, X_2), \dots, T_n(X_1, \dots, X_n)$ in $k[X_1, \dots, X_n]$ which forms a reduced Gröbner basis for the lexicographic order $X_n > \cdots > X_1$ and generates a radical ideal (so T_i is monic in X_i). The notation $T^1, \dots, T^s$ denotes a family of s triangular sets, with $T^i = T^i_1, \dots, T^i_n$. Then, any 0-dimensional variety V can be represented by such a family, such that $I(V) = \cap_{i \le s} \langle T^i \rangle$ holds, and where $\langle T^i \rangle$ and $\langle T^{i'} \rangle$ are coprime ideals for $i \neq i'$; we call it a *triangular decomposition* of V. This decomposition is not unique: the different possibilities are obtained by suitably recombining the triangular sets describing the irreducible components of V.

In this paper, we consider 0-dimensional varieties defined over $\mathbb{Q}$. Let thus $F = F_1, \dots, F_n$ be a polynomial system in $\mathbb{Z}[X_1, \dots, X_n]$. Since we have in mind to apply Hensel lifting techniques, we will only consider the *simple roots* of F, that is, those where the Jacobian determinant J of F does not vanish. We write $Z(F)$ for this set of points; by the Jacobian criterion [10, Ch. 16], $Z(F)$ is finite, even though the whole zero-set of F, written $V(F)$, may have higher dimension.

Let us assume that we have at hand an oracle that, for any prime p, outputs a triangular decomposition of $Z(F \bmod p)$. Then for a prime p, a rough sketch of an Hensel lifting algorithm could be: (1) Compute a triangular decomposition $t^1, \dots, t^s$ of $Z(F \bmod p)$, and (2) Lift these triangular sets over $\mathbb{Q}$. However, without more precautions, this algorithm may fail to produce a correct answer. Indeed, extra factorizations or recombinations can occur modulo p. Thus, we have no guarantee that there exist triangular sets $T^1, \dots, T^s$ defined over $\mathbb{Q}$, that describe $Z(F)$, and with $t^1, \dots, t^s$ as modular images. Furthermore, if we assume no control over the modular resolution process, there is little hope of obtaining a quantification of primes p of "bad" reduction.

Consider for instance the variety $V \subset \mathbb{C}^2$ defined by the polynomials $326X_1 - 10X_2^6 + 51X_2^5 + 17X_2^4 + 306X_2^2 + 102X_2 + 34$ and $X_2^7 + 6X_2^4 + 2X_2^3 + 12$. For the order $X_2 > X_1$, the only possible description of V by triangular sets with rational coefficients corresponds to its irreducible decomposition, that is, $T^1 : (X_1 - 1,\ X_2^3 + 6)$ and $T^2 : (X_1^2 + 2,\ X_2^2 + X_1)$. Now, the following triangular sets describe the zeros of $(F \bmod 7)$, which are not the reduction modulo 7 of T^1 and T^2;

$$t^1 \left| \begin{array}{l} X_2^2 + 6X_2X_1^2 + 2X_2 + X_1 \\ X_1^3 + 6X_1^2 + 5X_1 + 2 \end{array} \right. \quad \text{and} \quad t^2 \left| \begin{array}{l} X_2 + 6 \\ X_1 + 6 \end{array} \right. ,$$

A lifting algorithm should discard t^1 and t^2, and replace them by the better choice t'^1 : $(X_1+6,\ X_2^3+6)$ and t'^2 : $(X_1^2+2,\ X_2^2+X_1)$, which are the reduction of T^1 and T^2 modulo 7. In [24], this difficulty was bypassed by restricting to *equiprojectable* varieties, *i.e.* varieties defined by a single triangular set, where no such ambiguity occurs. However, as this example shows, this assumption discards simple cases. Our main concern is to lift this limitation, thus extending these techniques to handle *triangular decompositions.*

Our answer consists in using a canonical decomposition of a 0-dimensional variety V, its *equiprojectable decomposition*, described as follows. Consider the map $\pi : V \subset \mathbb{A}^n(\overline{k}) \to \mathbb{A}^{n-1}(\overline{k})$ that forgets the last coordinate. To x in V, we associate $N(x) = \#\pi^{-1}(\pi(x))$, that is, the number of points lying in the same π-fiber as x. Then, we split V into the disjoint union $V_1 \cup \cdots \cup V_d$, where for all $i = 1, \dots, d$, V_i equals $N^{-1}(i)$, *i.e.*, the set of points $x \in V$ where $N(x) = i$. This splitting process is applied recursively to all $V_1, \dots, V_d$, taking into account the fibers of the successive projections $\mathbb{A}^n(\overline{k}) \to \mathbb{A}^i(\overline{k})$, for $i = n-1, \dots, 1$. In the end, we obtain a family of pairwise disjoint, equiprojectable varieties, whose reunion equals V, which form the *equiprojectable decomposition* of V. As requested, each of them is representable by a triangular set with coefficients in the definition field of V.

Looking back at the example, both $Z(F)$ and $Z(F \bmod 7)$ are described on the leftmost picture below (forgetting the actual coordinates of the points). Representing $Z(F)$ by T^1 and T^2, as well as $Z(F \bmod 7)$ by t'^1 and t'^2 amounts to grouping the points as on the central picture; this is the equiprojectable decomposition. The rightmost picture shows the description of $Z(F \bmod 7)$ by t^1 and t^2.

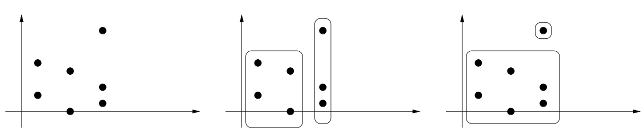

The above algorithm sketch is thus improved by applying lifting only after computing the equiprojectable decomposition of the modular output. Theorem 1 shows how to control the primes of bad reductions for the equiprojectable decomposition, thus overcoming the limitation that we pointed out previously. In what follows, the height of $x \in \mathbb{Z}$ is defined as $\operatorname{ht} x = \log |x|$; the height of $f \in \mathbb{Z}[X_1, \dots, X_n]$ is the maximum of the heights of its coefficients; that of $p/q \in \mathbb{Q}$, with $\gcd(p, q) = 1$, is $\max(\operatorname{ht} p, \operatorname{ht} q)$.

THEOREM 1. *Let $F_1, \dots, F_n$ have degree $\le d$ and height $\le h$. Let $T^1, \dots, T^s$ be the triangular description of the equiprojectable decomposition of $Z(F)$. There exists $A \in \mathbb{N} - \{0\}$, with $\operatorname{ht} A \le \mathfrak{a}(n, d, h)$, and, for $n \ge 2$,*

$$\mathfrak{a}(n, d, h) = 2n^2 d^{2n+1}(3h + 7\log(n+1) + 5n\log d + 10),$$

and with the following property. If a prime p does not divide A, then p cancels none of the denominators of the coefficients of $T^1, \dots, T^s$, and these triangular sets reduced mod p define the equiprojectable decomposition of $Z(F \bmod p)$.

Thus, the set of bad primes is finite and we have an explicit control on its size. Since we have to avoid some "discriminant locus", it is natural, and probably unavoidable, that the bound should involve the square of the Bézout number.

A second question is the coefficient size of the output. In what follows, we write $\deg V$ and $\operatorname{ht} V$ for the *degree* and *height* of a 0-dimensional variety V defined over $\mathbb{Q}$: the former denotes its number of points, and the later estimates its arithmetic complexity; see [17] and references therein for its definition. Let then $T^1, \dots, T^s$ be the triangular sets that describe the equiprojectable decomposition of $Z = Z(F)$. In [9], it is proved that all coefficients in $T^1, \dots, T^s$ have height in $O(n^{O(1)}(\deg Z + \operatorname{ht} Z)^2)$. However, better estimates are available, through the introduction of an alternative representation denoted by $N^1, \dots, N^s$. For $i \le s$, $N^i = N_1^i, \dots, N_n^i$ is obtained as follows. Let $D_1^i = 1$ and $N_1^i = T_1^i$. For $2 \le \ell \le n$ and $1 \le i \le s$, define

$$D_\ell^i = \prod_{1 \le j \le \ell-1} \frac{\partial T_j^i}{\partial X_j} \quad \text{and} \quad N_\ell^i = D_\ell^i T_\ell^i \bmod (T_1^i, \dots, T_{\ell-1}^i).$$

It is proved in [9] that all coefficients in $N^1, \dots, N^s$ have height in $O(n^{O(1)}(\deg Z + \operatorname{ht} Z))$. Since $T^1, \dots, T^s$ are easily recovered from $N^1, \dots, N^s$, our algorithm will compute the latter, their height bounds being the better.

Theorem 2 below states our main result regarding lifting techniques for triangular decompositions; in what follows, we say that an algorithm has a *quasi-linear* complexity in terms of some parameters if its complexity is linear in all of these parameters, up to polylogarithmic factors. We need the following assumptions:

- For any $C \in \mathbb{N}$, let $\Gamma(C)$ be the sets of primes in $[C+1, \dots, 2C]$. We assume the existence of an oracle O_1 which, for any $C \in \mathbb{N}$, outputs a random prime in $\Gamma(C)$, with the uniform distribution.
- We assume the existence of an oracle O_2, which, given a system F and a prime p, outputs the representation of the equiprojectable decomposition of $Z(F \bmod p)$ by means of triangular sets. We give in Section 2 an algorithm to convert any triangular decomposition of $Z(F \bmod p)$ to the equiprojectable one; its complexity analysis is subject of current research.
- For F as in Theorem 1, we write $\mathfrak{a}_F = \mathfrak{a}(n, d, h)$, $\mathfrak{h}_F = nd^n(h + 11\log(n+3))$ and $\mathfrak{b}_F = 5(\mathfrak{h}_F + 1)\log(2\mathfrak{h}_F + 1)$. The input system is given by a straight-line program of size L, with constants of height at most h_L.
- $\mathsf{C} \in \mathbb{N}$ is such that for any ring R, any $d \ge 1$ and monic $t \in R[X]$ of degree d, all operations $(+, -, \times)$ in $R[X]/t$ can be computed in $\mathsf{C} d \log d \log\log d$ operations in R [12, Ch. 8,9]. Then all operations $(+, -, \times)$ modulo a triangular set T in n variables can be done in quasi-linear complexity in C^n and $\deg V(T)$.

THEOREM 2. *Let $\varepsilon > 0$. There exists an algorithm which, given F, satisfying*

$$\frac{4\mathfrak{a}_F + 2\mathfrak{b}_F}{\varepsilon} + 1 < \frac{1}{2}\exp(2\mathfrak{h}_F + 1),$$

computes $N^1, \dots, N^s$ defined above. The algorithm uses two calls to O_1 with $C = 4\mathfrak{a}_F + 2\mathfrak{b}_F/\varepsilon$, two calls to O_2 with p in $[C+1, \dots, 2C]$, and its bit complexity is quasi-linear in $L, h_L, d, \log h, \mathsf{C}^n, \deg Z, (\deg Z + \operatorname{ht} Z), |\log \varepsilon|$. The algorithm is probabilistic, with success probability $\ge 1 - \varepsilon$.

To illustrate these estimates, suppose e.g. that we have $n = 10, d = 4, h = 100$, hence potentially 1048576 solutions; to ensure a success probability of 99%, the primes should have only about 20 decimal digits, hence can be generated without difficulty. Thus, even for such "large" systems, our results are quite manageable. Besides, computing the polynomials N^i instead of T^i enables us to benefit from their improved height bounds.

In the sequel, we use the following notation. For $n \in \mathbb{N}$, for $1 \le j \le i \le n$ and any field k, we denote $\pi^i_j : \mathbb{A}^i(\overline{k}) \to \mathbb{A}^j(\overline{k})$ the map $(x_1, \ldots, x_i) \mapsto (x_1, \ldots, x_j)$. The cardinality of a finite set G is written $\#G$.

2. SPLIT-AND-MERGE ALGORITHM

We start by reviewing the notion of equiprojectable decomposition of a 0-dimensional variety V, introduced in [8]. Then, in preparation for the modular algorithm of Section 4, we present an algorithm for computing this decomposition, given an arbitrary triangular decomposition of V. We call it *Split-and-Merge*, after its two phases: the *splitting* of what we call *critical pairs* (which is achieved by GCD computations) and the *merging* of what we call *solvable families* (which is performed by Chinese remaindering). The complexity analysis of the Split-and-Merge algorithm is work in progress [6]. From our preliminary study reported in [7], we believe that suitable improvements of the Split-and-Merge algorithm can run in quasi-linear time in the degree of V.

Let k be a perfect field and $\overline{k}$ one of its algebraic closures. Following [2], we first define the notion of equiprojectability.

Equiprojectable variety. Let $V \subset \mathbb{A}^n(\overline{k})$ be a 0-dimensional variety over k. For $1 \le i \le n$, the variety V is *equiprojectable* on $\pi^n_i(V)$ if all fibers of the restriction $\pi^n_i : V \to \pi^n_i(V)$ have the same cardinality. Then, for $1 \le i \le n$, V is *i-equiprojectable* if it is equiprojectable on all $\pi^n_j(V)$, $i \le j \le n$. Thus, any 0-dimensional variety is n-equiprojectable. Finally, V is *equiprojectable* if it is 1-equiprojectable. It is the case if and only if its defining ideal is generated by a triangular set $T_1, \ldots, T_n$ with coefficients in k. In this case, k being perfect, all fibers of the projection $\pi^n_i(V) \to \pi^n_{i-1}(V)$ share the same cardinality, which is the degree of T_i in X_i.

The variety V can be decomposed as the disjoint union of equiprojectable ones, in possibly several ways. Any such decomposition amounts to represent V as the disjoint union of the zeros of some triangular sets. The equiprojectable decomposition is a canonical way of doing so, defined by combinatorial means.

Equiprojectable decomposition. Let first W be a 0-dimensional variety in $\mathbb{A}^i(\overline{k})$, for some $1 \le i \le n$. For x in $\mathbb{A}^{i-1}(\overline{k})$, we define the preimage

$$\mu(x, W) = (\pi^i_{i-1})^{-1}(x) \cap W;$$

for any $d \ge 1$, we can then define

$$A(d, W) = \left\{ x \in W \mid \#\mu(\pi^i_{i-1}(x), W) = d \right\}.$$

Thus, x is in $A(d, W)$ if W contains exactly d points x' such that $\pi^i_{i-1}(x) = \pi^i_{i-1}(x')$ holds. Only finitely many of the $A(d, W)$ are not empty and the non-empty ones form a partition of W. Let $1 \le i \le n$. Writing $W = \pi^n_i(V)$, we define

$$B(i, d, V) = \{x \in V \mid \pi^n_i(x) \in A(d, W)\}.$$

Thus, $B(i, d, V)$ is the preimage of $A(d, W)$ in V, so these sets form a partition of V. If V is i-equiprojectable, then all $B(i, d, V)$ are $(i-1)$-equiprojectable. We then define inductively $B(V) = V$, and, for $1 < i \le n$, $B(d_i, \ldots, d_n, V) = B(i, d_i, B(d_{i+1}, \ldots, d_n, V))$. All $B(d_i, \ldots, d_n, V)$ are $(i-1)$-equiprojectable, only finitely many of them are not empty, and the non-empty ones form a partition of V.

The *equiprojectable decomposition* of V is its partition into the family of all non-empty $B(d_2, \ldots, d_n, V)$. All these sets being equiprojectable, they are defined by triangular sets. Note that we have not proved yet that the $B(d_2, \ldots, d_n, V)$ are defined over the same field as V. This will come as a by-product of the algorithms of this section. To do so, we introduce now the notions of *critical pair* and *solvable pair*.

Critical and solvable pairs. Let $T \neq T'$ be two triangular sets. The least integer ℓ such that $T_\ell \neq T'_\ell$ is called the *level* of the pair T, T'. If $\ell = 1$ we let $K_\ell = k$, otherwise we define $K_\ell = k[X_1, \ldots, X_{\ell-1}]/\langle T_1, \ldots, T_{\ell-1}\rangle$. Since a triangular set generates a radical ideal, the residue class ring K_ℓ is a direct product of fields. Therefore, every pair of univariate polynomials with coefficients in K_ℓ has a GCD in the sense of [22]. The pair T, T' is *critical* if T_ℓ and T'_ℓ are not relatively prime in $K_\ell[X_\ell]$. If T, T' is not critical, it is *certified* if $U, U' \in K_\ell[X_\ell]$ such that $UT_\ell + U'T'_\ell = 1$ are known. The pair T, T' is *solvable* if it is not critical and if for all $\ell < j \le n$ we have $\deg_{X_j} T_j = \deg_{X_j} T'_j$.

Introducing the notion of a certified solvable pair is motivated by efficiency considerations. Indeed, during the splitting step, solvable pairs are discovered. Then, during the merging step, the Bézout coefficients U, U' of these solvable pairs will be needed for Chinese Remaindering.

Solvable families. We extend the notion of *solvability* from a pair to a family of triangular sets. A family $\mathfrak{T}$ of triangular sets is *solvable* (resp. *certified solvable*) at level ℓ if every pair $\{T, T'\}$ of elements of $\mathfrak{T}$ is solvable (resp. certified solvable) of level ℓ.

The following proposition shows how to recombine such families. When this is the case, we say that all T in $\mathfrak{T}$ *divide* S. In what follows, we write $V(\mathfrak{T})$ for $\cup_{T \in \mathfrak{T}} V(T)$.

Proposition 1. *If $\mathfrak{T}$ is certified solvable at level ℓ, one can compute a triangular set S such that $V(S) = V(\mathfrak{T})$, using only multiplications in $K_\ell[X_\ell]$.*

Proof. First, we assume that $\mathfrak{T}$ consists of the pair $\{T, T'\}$. We construct S as follows. We set $S_i = T_i$ for $1 \le i < \ell$ and $S_\ell = T_\ell T'_\ell$. Let $\ell < i \le n$. For computing S_i, we see T_i and T'_i in $K_\ell[X_\ell][X_{\ell+1}, \ldots, X_i]$. We apply Chinese remaindering to the coefficients in T_i and T'_i of each monomial in $X_{\ell+1}, \ldots, X_i$ occurring in T_i or T'_i: since the Bézout coefficients U, U' for T_ℓ, T'_ℓ are known, this can be done using multiplications in $K_\ell[X_\ell]$ only. It follows from the Chinese Remaindering Theorem that the ideal $\langle S\rangle$ is equal to $\langle T\rangle \cap \langle T'\rangle$; for $i > \ell$, the equality $\deg_{X_i} T_i = \deg_{X_i} T'_i$ shows that S is monic in X_i, as requested.

Assume that $\mathfrak{T}$ consists of $s > 2$ triangular sets $T^1, \ldots, T^s$. First, we apply the case $s = 2$ to T^1, T^2, obtaining a triangular set $T^{1,2}$. Observe that every pair $T^{1,2}, T^j$, for $3 \le j \le s$, is solvable but not certified solvable: we obtain the requested Bézout coefficient by updating the known ones. Let us fix $3 \le j \le s$. Given $A_1, A_2, B_1, B_j, C_2, C_j \in K_\ell[X_\ell]$ such that $A_1 T^1_\ell + A_2 T^2_\ell = B_1 T^1_\ell + B_j T^j_\ell = C_2 T^2_\ell + C_j T^j_\ell = 1$ hold in $K_\ell[X_\ell]$, we let $\alpha = B_1 C_2 \mod T^j_\ell$ and $\beta = A_1 C_j T^1_\ell + A_2 B_j T^2_\ell \mod T^1_\ell T^2_\ell$. Then, $\alpha T^{1,2}_\ell + \beta T^j_\ell = 1$ in $K_\ell[X_\ell]$, as requested. Proceeding by induction ends the proof. $\square$

Splitting critical pairs. Let now V be a 0-dimensional variety over k. Proposition 3 below encapsulates the first part of the *Split-and-Merge* algorithm: given any triangular decomposition $\mathfrak{T}$ of V, it outputs another one, without critical pairs. We first describe the basic splitting step.

PROPOSITION 2. *Let $\mathfrak{T}$ be a triangular decomposition of V which contains critical pairs. Then one can compute a triangular decomposition $\mathsf{Split}(\mathfrak{T})$ of V which has cardinality larger than that of $\mathfrak{T}$.*

PROOF. Let T, T' be a critical pair of $\mathfrak{T}$ of level ℓ and let G be a GCD of T_ℓ, T'_ℓ in $K_\ell[X_\ell]$. First, assume that G is monic, in the sense of [22]; let Q and Q' be the quotients of T_ℓ and T'_ℓ by G in $K_\ell[X_\ell]$. We define the sets

$$\begin{aligned} A &= T_1, \ldots, T_{\ell-1}, G, T_{\ell+1}, \ldots, T_n, \\ B &= T_1, \ldots, T_{\ell-1}, Q, T_{\ell+1}, \ldots, T_n, \\ A' &= T_1, \ldots, T_{\ell-1}, G, T'_{\ell+1}, \ldots, T'_n, \\ B' &= T_1, \ldots, T_{\ell-1}, Q', T'_{\ell+1}, \ldots, T'_n. \end{aligned}$$

We let $\mathsf{Split}(\mathfrak{T}) = \{A, B, A', B'\}$, excluding the triangular sets defining the empty set. Since the pair T, T' is critical, $V(A)$ and $V(B)$ are non-empty. Since T_ℓ and T'_ℓ are not associate in $K_\ell[X_\ell]$, at least Q or Q' is not constant. Thus, $\mathsf{Split}(\mathfrak{T})$ has cardinality at least 3. Since $\langle T\rangle$ and $\langle T'\rangle$ are radical, if $Q \notin K_\ell$, G and Q are coprime in $K_\ell[X_\ell]$, so $V(T)$ is the disjoint union of $V(A)$ and $V(B)$. The same property holds for A' and B'. Thus, the proposition is proved.

Assume now that T_ℓ, T'_ℓ have no monic GCD in $K_\ell[X_\ell]$. Then, there exist triangular sets $C^1, \ldots, C^s, D^1 \ldots D^s$ such that $V(T)$ is the disjoint union of $V(C^1), \ldots, V(C^s)$, $V(T')$ is the disjoint union of $V(D^1), \ldots, V(D^s)$, at least one pair C^i, D^j is critical and C^i_ℓ, D^j_ℓ admits a monic GCD in $K_\ell[X_\ell]$. These triangular sets are obtained by the algorithms of [22] when computing a GCD of T_ℓ, T'_ℓ in $K_\ell[X_\ell]$. Then the results of the monic case prove the existence of $\mathsf{Split}(\mathfrak{T})$. □

PROPOSITION 3. *Let $\mathfrak{T}$ be a triangular decomposition of V. One can compute a triangular decomposition $\mathfrak{T}'$ of V with no critical pairs, and where each pair of triangular sets is certified.*

PROOF. Write $\mathfrak{T}_0 = \mathfrak{T}$, and define a sequence $\mathfrak{T}_i$ by $\mathfrak{T}_{i+1} = \mathsf{Split}(\mathfrak{T}_i)$, if $\mathfrak{T}_i$ contains critical pairs, and $\mathfrak{T}_{i+1} = \mathfrak{T}_i$ otherwise. Testing the presence of critical pairs is done by GCD computations, which yields the Bézout coefficients in case of coprimality. Let D be the number of irreducible components of V. Any family $\mathfrak{T}_i$ has cardinality at most D, so the sequence $\mathfrak{T}_i$ becomes stationary after at most D steps. □

Thus, we can now suppose that we have a triangular decomposition $\mathfrak{T}$ of V without critical pairs, and where every pair is certified, such as the one computed in Proposition 3. We describe the second part of the *Split-and-Merge* algorithm: merging solvable families in a suitable order, to obtain the equiprojectable decomposition of V.

For $0 \le \kappa \le n$, we say that $\mathfrak{T}$ satisfies property P_κ if for all $T, T' \in \mathfrak{T}$ the pair $\{T, T\}$ is certified, has level $\ell \le \kappa$ and for all $\kappa < i \le n$ satisfies $\deg_{X_i} T_i = \deg_{X_i} T'_i$. Observe that if $\mathsf{P}_0(\mathfrak{T})$ holds, then $\mathfrak{T}$ contains only one triangular set, and that the input family $\mathfrak{T}$ satisfies P_n.

The basic merging algorithm. Let $1 \le \kappa \le n$. We now define the procedure Merge_κ, which takes as input a family $\mathfrak{T}_\kappa$ of triangular sets which satisfies P_κ, and outputs several families of triangular sets, whose reunion defines the same set of points, and all of which satisfy $\mathsf{P}_{\kappa-1}$. First, we partition $\mathfrak{T}_\kappa$ using the equivalence relation $T \equiv T'$ if and only if $T_1, \ldots, T_{\kappa-1} = T'_1, \ldots, T'_{\kappa-1}$. Assumption P_κ shows that each equivalence class is certified and solvable of level κ. We then let $\mathfrak{S}^{(\kappa)}$ be the family of triangular sets obtained by applying Proposition 1 to each equivalence class.

LEMMA 1. *Let $S \neq S'$ in $\mathfrak{S}^{(\kappa)}$. The pair $\{S, S'\}$ is non-critical, certified, of level $\ell < \kappa$.*

PROOF. Let $T, T' \in \mathfrak{T}$, which respectively divide S and S'. Due to assumption P_κ, there exists $0 \le \ell \le \kappa$ such that $T_1, \ldots, T_{\ell-1} = T'_1, \ldots, T'_{\ell-1}$ and $(T_1, \ldots, T_\ell)$ and $(T'_1, \ldots, T'_\ell)$ have no common zero. Then, $\ell < \kappa$, since $T \not\equiv T'$. Thus, $T_1, \ldots, T_\ell = S_1, \ldots, S_\ell$ and $T'_1, \ldots, T'_\ell = S'_1, \ldots, S'_\ell$. Since $\{T, T'\}$ is certified of level $\ell < \kappa$, $\{S, S'\}$ is also. □

We partition $\mathfrak{S}^{(\kappa)}$ some more, into the classes of the equivalence relation $S \equiv' S'$ if and only if $\deg_{X_\kappa} S_\kappa = \deg_{X_\kappa} S'_\kappa$. Let $\mathfrak{S}^{(\kappa)}_1, \ldots, \mathfrak{S}^{(\kappa)}_\delta$ be the equivalence classes, indexed by the common degree in X_κ; we define $\mathsf{Merge}_\kappa(\mathfrak{T}_\kappa)$ as the data of all these equivalence classes.

LEMMA 2. *Each family $\mathfrak{S}^{(\kappa)}_d$ satisfies $\mathsf{P}_{\kappa-1}$.*

PROOF. Let $S \neq S'$ in $\mathfrak{S}^{(\kappa)}_d$, and let T, T' be as in the proof of Lemma 1; we now prove the degree estimate. For $\kappa < i \le n$, we have $\deg_{X_i} T_i = \deg_{X_i} S_i$ and $\deg_{X_i} T'_i = \deg_{X_i} S'_i$; assumption P_κ shows that $\deg_{X_i} S_i = \deg_{X_i} S'_i$ for $\kappa < i \le n$. Since $\deg_{X_\kappa} S_\kappa = \deg_{X_\kappa} S'_\kappa = d$, the lemma is proved. □

PROPOSITION 4. *$V(\mathfrak{S}^{(\kappa)}_d) = B(\kappa, d, V(\mathfrak{T}_\kappa))$ for all d.*

PROOF. We know that $V(\mathfrak{T}_\kappa)$ is the union of the $V(\mathfrak{S}^{(\kappa)}_d)$. Besides, both families $\{V(\mathfrak{S}^{(\kappa)}_d)\}$ and $\{B(\kappa, d, V(\mathfrak{T}))\}$ form a partition of $V(\mathfrak{T}_\kappa)$. Thus, it suffices to prove that for x in $V(\mathfrak{T}_\kappa)$, $x \in V(\mathfrak{S}^{(\kappa)}_d)$ implies that $\pi^n_\kappa(x) \in A(d, W)$, with $W = \pi^n_\kappa(V(\mathfrak{T}_\kappa))$. First, for S in $\mathfrak{S}^{(\kappa)}$, write $W_S = \pi^n_\kappa(S)$. Then Lemma 1 shows that the W_S form a partition of W, and that their images $\pi^\kappa_{\kappa-1}(W_S)$ are pairwise disjoint.

Let now $x \in V(\mathfrak{S}^{(\kappa)}_d)$ and $y = \pi^n_\kappa(x)$. There exists a unique $S \in \mathfrak{S}^{(\kappa)}$ such that $x \in V(S)$. The definition of $\mathfrak{S}^{(\kappa)}_d$ shows that there are exactly d points y' in W_S such that $\pi^\kappa_{\kappa-1}(y) = \pi^\kappa_{\kappa-1}(y')$. On the other hand, for any $y \in W_{S'}$, with $S' \neq S$, the above remark shows that $\pi^\kappa_{\kappa-1}(y) \neq \pi^\kappa_{\kappa-1}(y')$. Thus, there are exactly d points y' in W such that $\pi^\kappa_{\kappa-1}(y) = \pi^\kappa_{\kappa-1}(y')$; this concludes the proof. □

The main merging algorithm. We can now give the main algorithm. We start from a triangular decomposition $\mathfrak{T}$ of V without critical pairs, and where every pair is certified, so it satisfies P_n. Let us initially define $\mathfrak{T}_n = \{\mathfrak{T}\}$; note that $\mathfrak{T}_n$ is a *set of families of triangular sets*. Then, for $1 \le \kappa \le n$, assuming $\mathfrak{T}_\kappa$ is defined, we write $\mathfrak{T}_{\kappa-1} = \cup_{\mathfrak{U}^{(\kappa)} \in \mathfrak{T}_\kappa} \mathsf{Merge}_\kappa(\mathfrak{U}^{(\kappa)})$. Lemma 2 shows that this process is well-defined; note that each $\mathfrak{T}_\kappa$ is a set of families of triangular sets as well.

Let $\mathfrak{U}$ be a family of triangular sets in $\mathfrak{T}_0$. Then $\mathfrak{U}$ satisfies P_0, so by the remarks make previously, $\mathfrak{U}$ consists in a single triangular set. Proposition 4 then shows that the triangular sets in $\mathfrak{T}_0$ form the equiprojectable components of V.

3. PROOF OF THEOREM 1

In this section, we consider the simple solutions $Z(F)$ of a system $F = F_1, \dots, F_n$ in $\mathbb{Z}[X_1, \dots, X_n]$, that is, those where the Jacobian determinant J of F does not vanish. We prove that for all primes p but a finite number, the equiprojectable decomposition of $Z(F)$ reduces modulo p to that of $Z(F \bmod p)$. These results require to control the cardinality of the "specialization" of a variety at p. Such questions are easy to formulate using *primitive elements* and associated representations, which we now define as a preamble.

Primitive element descriptions. Let $W \subset \mathbb{C}^\ell$ be a 0-dimensional variety defined over $\mathbb{Q}$. Let Δ be a linear form in $\mathbb{Z}[X_1, \dots, X_\ell]$. Its *minimal polynomial* is the minimal polynomial $\mu \in \mathbb{Q}[T]$ of the multiplication endomorphism by Δ in $\mathbb{Q}[W]$; it is the squarefree part of $\Pi_{x \in W}(T - \Delta(x))$. Then Δ is a *primitive element* for W if the map $x \mapsto \Delta(x)$ is one-to-one on W. In this case, μ has degree $\deg W$ and $\mathbb{Q}[W]$ is isomorphic to the residue class ring $\mathbb{Q}[T]/\mu$. Writing $w_i \in \mathbb{Q}[T]$ for the image of X_i, we deduce that $\mu(T) = 0$ and $X_i = w_i(T)$, $1 \le i \le \ell$, form a parametrization of the points in W.

We will use quantitative estimates on the size of the coefficients in this representation, in terms of the degree and height of W. The following result is [5, Th. 2]; using the coefficient χ' leads to sharp height bound, as is the case for the polynomials N^i defined in the introduction.

Lemma 3. *Let h_Δ be an upper bound of the height of Δ, and $H_\Delta = \operatorname{ht} W + (\deg W) h_\Delta + (\deg W)\log(\ell+2) + (\ell+1)\log \deg W$. There exist $\chi, v_1, \dots, v_\ell$ in $\mathbb{Z}[T]$, such that $\chi, \chi', v_1, \dots, v_\ell$ have height at most H_Δ, μ_n equals χ divided by its leading coefficient, and $w_i = v_i/\chi' \bmod \chi$ for all i.*

Geometric considerations. Let now $Z = Z(F)$. For $1 \le i \le n$, let Δ_i be a linear form in $\mathbb{Z}[X_1, \dots, X_i]$ which is a primitive element for $\pi_i^n(Z)$, let $\mu_i \in \mathbb{Q}[T]$ be its minimal polynomial, and let $w_1, \dots, w_n \in \mathbb{Q}[T]$ be the parametrization of Z associated to Δ_n. Let finally p a prime. We first introduce assumptions on p (denoted by $\mathbf{H}_1$, $\mathbf{H}_2$, $\mathbf{H}_3$), that yield the conclusion of Theorem 1 in a series of lemmas; we then give quantitative estimates for these assumptions.

$\mathbf{H_1}$. The prime p divides no coefficients in $\mu_n, w_1, \dots, w_n$ and μ_n remains squarefree modulo p.

Let $\mathbb{F}_q$ be a finite extension of $\mathbb{F}_p$ such that ($\mu_n \bmod p$) splits in $\mathbb{F}_q$, let $\mathbb{Q}_q$ be the corresponding unramified extension of $\mathbb{Q}_p$ [20] and $\mathbb{Z}_q$ its ring of integers; then, μ_n splits in $\mathbb{Q}_q$, and has all its roots in $\mathbb{Z}_q$; thus, Z lies in $\mathbb{Z}_q^n$. Note that p divides no coefficient in $\mu_1, \dots, \mu_n$: the roots of μ_i are the values of Δ_i on $\pi_i^n(Z)$, so they are in $\mathbb{Z}_q$, hence the coefficients of μ_i are in $\mathbb{Z}_q \cap \mathbb{Q} = \mathbb{Z}_p$. The map $\mathbb{Z}_q \to \mathbb{F}_q$ of reduction modulo p extends to maps $a \in \mathbb{Z}_q^i \mapsto \overline{a} \in \mathbb{F}_q^i$ for all i. Given $A \subset \mathbb{Z}_q^i$, $\overline{A}$ is the set $\{\overline{a} \mid a \in A\}$. The same notation is used for the reduction of polynomials modulo p.

$\mathbf{H_2}$. All polynomials $\overline{\mu_i}$ are squarefree.

Lemma 4. *For $i \le n$, $\#\pi_i^n(Z)$ equals $\#\pi_i^n(\overline{Z})$.*

Proof. The inequality $\#\pi_i^n(\overline{Z}) \le \#\pi_i^n(Z)$ is obvious. By assumption $\mathbf{H_2}$, all values taken by Δ_i on $\pi_i^n(\overline{Z})$ are distinct, so $\#\pi_i^n(\overline{Z}) \ge \deg \mu_i = \#\pi_i^n(Z)$. $\square$

Lemma 5. *For all $d_2, \dots, d_n$, $B(d_2, \dots, d_n, \overline{Z})$ equals $\overline{B(d_2, \dots, d_n, Z)}$.*

Proof. We prove on $\ell = n+1, \dots, 2$ that for all $d_\ell, \dots, d_n$, $B(d_\ell, \dots, d_n, \overline{Z})$ equals $\overline{B(d_\ell, \dots, d_n, Z)}$; taking $\ell = 2$ gives the lemma. Since $B(X) = X$ for any variety X, this property holds for $\ell = n+1$. Assuming it for $B(d_{\ell+1}, \dots, d_n, Z)$, we prove it for $B(d_\ell, \dots, d_n, Z)$. Let $B = B(d_{\ell+1}, \dots, d_n, Z)$, $B_\ell = \pi_\ell^n(B)$ and $B_{\ell-1} = \pi_{\ell-1}^n(B)$; Lemma 4 implies that reduction modulo p is one-to-one on both B_ℓ and $B_{\ell-1}$. For y in $B_{\ell-1}$ and z in $\overline{B_{\ell-1}}$, we define

$$\mu(y) = (\pi_{\ell-1}^\ell)^{-1}(y) \cap B_\ell \quad \text{and} \quad \mu(z) = (\pi_{\ell-1}^\ell)^{-1}(z) \cap \overline{B_\ell}.$$

We first prove that $\mu(y)$ and $\mu(\overline{y})$ have the same cardinality for all y in $B_{\ell-1}$. To this effect, observe the equalities

$$\sum_{y \in B_{\ell-1}} \#\mu(y) = \#B_\ell, \qquad \sum_{z \in \overline{B_{\ell-1}}} \#\mu(z) = \#\overline{B_\ell}.$$

Let now y in $B_{\ell-1}$. Since $\overline{\mu(y)} \subset \mu(\overline{y})$, injectivity of the reduction mod p on B_ℓ implies that $\#\mu(y) \le \#\mu(\overline{y})$. Thus,

$$\#B_\ell = \sum_{y \in B_{\ell-1}} \#\mu(y) \le \sum_{y \in B_{\ell-1}} \#\mu(\overline{y}).$$

Injectivity of the reduction mod p on $B_{\ell-1}$ implies that

$$\sum_{y \in B_{\ell-1}} \#\mu(\overline{y}) = \sum_{z \in \overline{B_{\ell-1}}} \#\mu(z) = \#\overline{B_\ell}.$$

This sum equals $\#B_\ell$. Thus, all inequalities are equalities, giving our claim.

For x in B_ℓ, write $\nu(x) = \mu(\pi_{\ell-1}^\ell(x))$; define similarly $\nu(z)$ for z in $\overline{B_\ell}$. By the previous point, $\nu(x)$ and $\nu(\overline{x})$ have the same cardinality. Recalling from Section 2 that for $d \in \mathbb{N}$, we have defined $A(d, B_\ell)$ as the set $\{x \in B_\ell \mid \#\nu(x) = d\}$, and $A(d, \overline{B_\ell})$ as the set $\{z \in \overline{B_\ell} \mid \#\nu(z) = d\}$, one can see $\overline{A(d, B_\ell)} = A(d, \overline{B_\ell})$. To conclude, recall that by definition $\{x \in \overline{Z} \mid \pi_\ell^n(x) \in A(d, \pi_\ell^n(B(d_{\ell+1}, \dots, d_n, \overline{Z})))\} = B(d, d_{\ell+1}, \dots, d_n, \overline{Z})$. By the induction assumption, this equals $\{x \in \overline{Z} \mid \pi_\ell^n(x) \in A(d, \overline{B_\ell})\}$, and we have proved that this equals $\overline{\{x \in Z \mid \pi_\ell^n(x) \in A(d, B_\ell)\}}$. By definition, this is $\overline{B(d, d_\ell, \dots, d_n, Z)}$, which is what we wanted. $\square$

Lemma 6. *Let $T^1, \dots, T^s$ be the triangular sets that describe the equiprojectable decomposition of Z. Then p cancels no denominator in the coefficients of $T^1, \dots, T^s$, and the reduction of these triangular sets modulo p defines the equiprojectable decomposition of $\overline{Z}$.*

Proof. For $i \le s$, let $Z_i = Z(T^i)$. By Lemma 5, $\overline{Z_1}, \dots, \overline{Z_s}$ are the equiprojectable components of $\overline{Z}$. For $i \le s$, $\overline{Z_i}$ is described by a triangular set t^i with coefficients in $\mathbb{F}_p$. The coefficients of T^i are rational functions of the points in Z_i, given by interpolation formulas [9, §3]. With these formulas, Lemma 4 shows that all denominators are nonzero modulo p. The coefficients of t^i are obtained using the same formulas, using the coordinates of the points in $\overline{Z_i}$. Thus, $t^i = T^i \bmod p$. $\square$

$\mathbf{H_3}$. The Jacobian determinant of F vanishes nowhere on $\overline{Z}$.

Lemma 7. *The set $\overline{Z}$ equals $Z(\overline{F})$.*

Proof. First, we prove that $\overline{F}$ vanishes on $\overline{Z}$. Indeed, all F_i belong to the ideal generated by $I = (\mu_n, X_1 - w_1, \dots, X_n - w_n)$ in $\mathbb{Q}[T, X_1, \dots, X_n]$. Now, I is a Gröbner basis, so any F_i can be written in terms of I. Since p divides no denominator and no leading term in I, the division equality specializes modulo p, and $\overline{F}$ vanishes on $\overline{Z}$, as requested. Let then $Z' = Z(\overline{F})$. By Assumption $\mathbf{H_3}$, $\overline{Z} \subset Z'$, so it suffices to prove that $\#Z' \leq \#\overline{Z}$. Let $\mathbb{F}_r$ be a finite extension of $\mathbb{F}_p$ that contains the coordinates of all these points and let $\mathbb{Q}_r$ be the corresponding unramified extension of $\mathbb{Q}_p$. By Hensel's lemma, all points in Z' lift to pairwise distinct simple roots of F in $\mathbb{Q}_r^n$. Thus, $\#Z' \leq \#Z = \#\overline{Z}$. $\square$

Quantitative estimates. By Lemmas 6 and 7, assumptions $\mathbf{H_1}$, $\mathbf{H_2}$ and $\mathbf{H_3}$ imply Theorem 1. Thus, it suffices to give quantitative estimates for these assumptions. To this effect, we let D and H be upper bounds on the degrees and heights of the varieties $\pi_i^n(Z)$, h_Δ be an upper bound of the height of $\Delta_1, \ldots, \Delta_n$, and $H_\Delta = H + Dh_\Delta + D\log(n+2) + (n+1)\log D$.

Lemma 8. *There exists a in $\mathbb{N} - \{0\}$ such that if p does not divide a, $\mathbf{H_1}$ and $\mathbf{H_2}$ hold. Moreover a verifies:*

$$\mathrm{ht}\, a \leq n\left((2D-1)H_\Delta + (2D-1)\log(2D-1)\right).$$

Proof. Fix i in $1, \ldots, n$, and let $\chi, \chi', v_1, \ldots, v_i$ the polynomials associated to $\pi_i^n(Z)$ and Δ_i in Lemma 3; all of them have integer coefficients of height at most H_Δ. Let now a_i be the resultant of χ and χ'; by Hadamard's bound, $\mathrm{ht}\, a_i \leq (2D-1)H_\Delta + (2D-1)\log(2D-1)$. Suppose that p does not divide a_i. Then, χ keeps the same degree and remains squarefree modulo p. Furthermore, p divides no coefficient in any w_j, since all denominators in $1/\chi' \bmod \chi$ divide a_i. Thus, assumption $\mathbf{H_1}$ holds. Repeating this argument for all projections $\pi_i^n(Z)$, and taking $a = a_1 \cdots a_n$ gives assumption $\mathbf{H_2}$. $\square$

Lemma 9. *There exists a' in $\mathbb{N} - \{0\}$ such that if p does not divide aa', $\mathbf{H_1}$, $\mathbf{H_2}$ and $\mathbf{H_3}$ hold, and with $\mathrm{ht}\, a' \leq 2Dn(dH_\Delta + h + \log d + (d+1)D\log(n+1))$.*

Proof. Let $\chi, v_1, \ldots, v_n$ be associated to Δ_n as in Lemma 3, let J^h be the homogenization of J w.r.t. a new variable, and let $a' \in \mathbb{Z}$ be the resultant of $J^h(\chi', v_1, \ldots, v_n)$ and χ; then, $a' \neq 0$ by the definition of Z. The Jacobian determinant J has coefficients of height at most $n(h + \log d + (d+1)\log(n+1))$; estimating the height of the determinant of the Sylvester matrix of $J^h(\chi', v_1, \ldots, v_n)$ and χ yields the bound on $\mathrm{ht}\, a'$. Suppose now that p does not divide aa'. Then the degree of χ does not drop modulo p, and thus no root of $\overline{\chi}$ cancels $\overline{J^h(\chi', v_1, \ldots, v_n)}$. In other words, all points described by $\overline{\chi}(T) = 0$ and $\overline{\chi'}(T)X_i = \overline{v_i}(T)$, $1 \leq i \leq n$, are simple for $\overline{F}$. This set of points equals $\overline{Z}$, giving $\mathbf{H_3}$. $\square$

In view of Lemma 9, we prove Theorem 1 with $A = aa'$. By [23, Lemma 2.1], all Δ_i can be taken of height at most $h_\Delta = n(\log n + 2\log D) \leq n(\log n + 2n\log d)$. Using the arithmetic Bézout bound of [17], we get after simplifications that all H_Δ are bounded by $nd^n(h + 3\log(n+1) + 2n\log d + 3)$. The previous lemmas then give the upper bounds below, which finish proving Theorem 1 after a few simplifications.

$$\begin{aligned}\mathrm{ht}\, a &\leq 2nd^{2n}(h + 3\log(n+1) + 2n\log d + 7)\\ \mathrm{ht}\, a' &\leq 2n^2d^{2n+1}(2h + 4\log(n+1) + 3n\log d + 3).\end{aligned}$$

4. PROOF OF THEOREM 2

We now give the details of our lifting algorithm: given a polynomial system F, it outputs a triangular representation of its set of simple solutions $Z = Z(F)$, by means of the polynomials $N^1, \ldots, N^s$ defined in the introduction. First of all, we describe the required subroutines, freely using the notation of Theorem 2, and that preceding it. We do not give details of the complexity estimates for lack of space; they are similar to those of [24].

- `EquiprojDecomposition` takes as input a polynomial system F and outputs the equiprojectable decomposition of $Z(F)$, encoded by triangular sets. This routine is called here for systems defined over finite fields. For the experiments in the next section, we applied the triangularization algorithm of [21], followed by the Split-and-Merge algorithm of Section 2, modulo a prime. Studying the complexity of this task is left to the forthcoming [7]; hence, we consider this subroutine as an oracle here, which is called O_2 in Theorem 2.

- `Lift` applies the Hensel lifting algorithm of [24], but this time to a *family of triangular sets*, defined first modulo a prime p_1, to triangular sets defined modulo the successive powers $p_1^{2^\kappa}$. From [24], one easily sees that the κth lifting step has a bit complexity quasi-linear in $(L, h_L, \mathsf{C}^n, \sum_{i \leq s} \deg V(T^i), 2^\kappa, \log p_1)$, *i.e.* in $(L, h_L, \mathsf{C}^n, \deg Z, 2^\kappa, \log p_1)$.

- `Convert` computes the polynomials N^i starting from the polynomials T^i. Only multiplications modulo triangular sets are needed to perform this operation, so its complexity is negligible before that of `Lift`.

- `RationalReconstruction` does the following. Let $a = p/q \in \mathbb{Q}$, and $m \in \mathbb{N}$ with $\gcd(q, m) = 1$. If $\mathrm{ht}\, m \geq 2\mathrm{ht}\, a + 1$, given $a \bmod m$, this routine outputs a. If $\mathrm{ht}\, m < 2\mathrm{ht}\, a + 1$, the output may be undefined, or differ from a. We extend this notation to the reconstruction of all coefficients of a family of triangular sets. Using the fast Euclidean algorithm [12, Ch 5,11], its complexity is negligible before that of `Lift`.

- We do not consider the cost of prime number generation. We see them as input here; formally, in Theorem 2, this is handled by calls to oracle O_1.

Computing a triangular decomposition by lifting techniques

Input: The system F, primes p_1, p_2
Output: The polynomials $N^1, \ldots, N^s$.

$T^{1,0}, \ldots, T^{s,0} \leftarrow$ `EquiprojDecomposition`$(Z(F \bmod p_1))$
$u^1, \ldots, u^{s'} \leftarrow$ `EquiprojDecomposition`$(Z(F \bmod p_2))$
$m^1, \ldots, m^{s'} \leftarrow$ `Convert`$(u^1, \ldots, u^{s'})$
$\kappa \leftarrow 1$
while not(Stop) do
  $T^{1,\kappa}, \ldots, T^{s,\kappa} \leftarrow$ `Lift`$(T^{1,\kappa-1}, \ldots, T^{s,\kappa-1}) \bmod p_1^{2^\kappa}$
  $N^{1,\kappa}, \ldots, N^{s,\kappa} \leftarrow$ `Convert`$(T^{1,\kappa}, \ldots, T^{s,\kappa})$
  $N_{\mathbb{Q}}^{1,\kappa}, \ldots, N_{\mathbb{Q}}^{s,\kappa} \leftarrow$ `RationalReconstruction`$(N^{1,\kappa}, \ldots, N^{s,\kappa})$
  `Stop` $\leftarrow \{m^1, \ldots, m^{s'}\}$ `Equals` $\{N_{\mathbb{Q}}^{1,\kappa}, \ldots, N_{\mathbb{Q}}^{s,\kappa}\} \bmod p_2$
  $\kappa \leftarrow \kappa + 1$
end while
return $N_{\mathbb{Q}}^{1,\kappa-1}, \ldots, N_{\mathbb{Q}}^{s,\kappa-1}$

We still use the notation and assumption of Theorem 2. From [9, Th. 1], all coefficients of $N^1, \ldots, N^s$ have height in $n^{O(1)}(\deg Z + \mathrm{ht}\, Z)$, which can explicitly be bounded by $\mathfrak{h}_F$. For $p_1 \le \exp(2\mathfrak{h}_F + 1)$, define

$$\mathfrak{d} = \mathfrak{d}(p_1) = \left\lceil \log_2 \left(\frac{2\mathfrak{h}_F + 1}{\log p_1} \right) \right\rceil .$$

Then, $p_1^{2^{\mathfrak{d}(p_1)}}$ has height at least $2\mathfrak{h}_F + 1$. In view of the prerequisites for rational reconstruction, $\mathfrak{d}(p_1)$ bounds the number of lifting steps. From an intrinsic viewpoint, at the last lifting step, 2^κ is in $O(n^{O(1)}(\deg Z + \mathrm{ht}\, Z))$.

Suppose that p_1 does not divide the integer A of Theorem 1. Then, Hensel lifting computes approximations $T^{1,\kappa}, \dots, T^{s,\kappa} = T^1, \dots, T^s$ modulo $p_1^{2^\kappa}$. At the κth lifting step, let $N^{1,\kappa}, \dots, N^{s,\kappa}$ be the output of `Convert` applied to $T^{1,\kappa}, \dots, T^{s,\kappa}$, computed modulo $p_1^{2^\kappa}$; let $N_{\mathbb{Q}}^{1,\kappa}, \dots, N_{\mathbb{Q}}^{s,\kappa}$ be the same polynomials after rational number reconstruction, if possible. By construction, they have rational coefficients of height at most $2^{\kappa-1} \log p_1$. Supposing that p_2 does not divide the integer A of Theorem 1, failure occurs only if for some κ in $0, \dots, \mathfrak{d} - 1$, and some j in $1, \dots, s$, $N_{\mathbb{Q}}^{j,\kappa}$ and N^j differ, but coincide modulo p_2. For this to happen, p_2 must divide some non-zero number of height at most $\mathfrak{h}_F + 2^{\kappa-1} \log p_1 + 1$. Taking all κ into account, this shows that for any prime p_1, there exists a non-zero integer B_{p_1} such that $\mathrm{ht}\, B_{p_1} \le (\mathfrak{h}_F + 1)\mathfrak{d} + 2^{\mathfrak{d}} \log p_1$, and if p_2 does not divide B_{p_1}, the lifting algorithm succeeds. One checks that the above bound can be simplified into $\mathrm{ht}\, B_{p_1} \le \mathfrak{b}_F$.

Let $C \in \mathbb{N}$ be such that

$$C = \left\lceil \frac{4\mathfrak{a}_F + 2\mathfrak{b}_F}{\varepsilon} \right\rceil, \quad \text{so that } C \le \frac{1}{2} \exp(2\mathfrak{h}_F + 1);$$

let Γ be the set of pairs of primes in $[C+1, \dots, 2C]^2$ and γ be the number of primes in $C+1, \dots, 2C$; note that $\gamma \ge C/(2\log C)$ and that $\#\Gamma = \gamma^2$. The upper bound on C shows that all primes p less than $2C$ satisfy the requested inequality $\log p \le 2\mathfrak{h}_F + 1$. We can then estimate how many choices of (p_1, p_2) in Γ lead to failure. There are at most $\mathfrak{a}_F / \log C$ primes p_1 in $C+1, \dots, 2C$ which divide the integer A of Theorem 1, discriminating at most $\gamma \mathfrak{a}_F / \log C$ pairs (p_1, p_2). For any other value of p_1, there are at most $(\mathfrak{a}_F + \mathfrak{b}_F)/\log C$ choices of p_2 which divide A and B_{p_1}. This discriminates at most $\gamma(\mathfrak{a}_F + \mathfrak{b}_F)/\log C$ pairs (p_1, p_2). Thus the number of choices in Γ leading to failure is at most $\gamma(2\mathfrak{a}_F + \mathfrak{b}_F)/\log C$. The lower bound on γ shows that if (p_1, p_2) is chosen randomly with uniform probability in Γ, the probability that it leads to failure is at most

$$\frac{\gamma(2\mathfrak{a}_F + \mathfrak{b}_F)}{\#\Gamma \log C} = \frac{\gamma(2\mathfrak{a}_F + \mathfrak{b}_F)}{\gamma^2 \log C} = \frac{2\mathfrak{a}_F + \mathfrak{b}_F}{\gamma \log C} \le \frac{4\mathfrak{a}_F + 2\mathfrak{b}_F}{C},$$

which is at most ε, as requested.

To estimate the complexity of this algorithm, note that since we double the precision at each lifting step, the cost of the last lifting step dominates. From the previous discussion, the number of bit operations cost at the last step is quasi-linear in $(L, h_L, \mathsf{C}^n, \deg Z, 2^\kappa, \log p_1)$. The previous estimates show that at this step, 2^κ is in $O(n^{O(1)}(\deg Z + \mathrm{ht}\, Z))$, whereas $\log p_1$ is quasi-linear in $|\log \varepsilon|, \log h, d, \log n$. Putting all these estimates ends the proof of Theorem 2.

5. EXPERIMENTATION

We realized a first MAPLE 9.5 implementation of our modular algorithm on top of the `RegularChains` library [19]. Tests on benchmark systems [25] reveal its strong features,

Sys	Name	n	d	h	$\mathfrak{h}$
1	Cyclohexane	3	4	3	4395
2	Fee_1	4	4	2	24464
3	fabfaux	3	3	13	2647
4	geneig	6	3	2	116587
5	eco6	6	3	0	105718
6	Weispfenning-94	3	5	0	7392
7	Issac97	4	2	2	1511
8	dessin-2	10	2	7	358048
9	eco7	7	3	0	387754
10	Methan61	10	2	16	450313
11	Reimer-4	4	5	1	55246
12	Uteshev-Bikker	4	3	3	7813
13	gametwo5	5	4	8	159192
14	chemkin	13	3	11	850088102

Table 1: Features of the polynomial systems

Sys	p_1	$\mathfrak{d}$	a	C_a
1	4423	7	2	15
2	24499	8	4	70
3	2671	7	5	110
4	116663	10	5	162
5	105761	10	3	40
6	7433	7	3	31
7	1549	6	5	102
8	358079	11	7	711
9	387799	11	4	89
10	450367	11	6	362
11	55313	9	2	19
12	7841	7	5	125
13	159223	10	-	-
14	850088191	18	-	-

Table 2: Data for the modular algorithm

Sys	Δ_p	E_p	Lift	Total	Mem.	Output size
1	1	0.3	2	7	5	243
2	3	1	9	20	6	4157
3	8	0.4	6	22	7	5855
4	5	1	5	18	6	4757
5	12	1.5	6	35	6	2555
6	16	1.5	11	43	7	3282
7	66	0.4	4	133	8	4653
8	47	9	232	427	13	122902
9	1515	9	35	2873	11	9916
10	2292	6	82	4686	25	50476
11	3507	1	9	5569	38	2621
12	4879	2	22	8796	63	12870
13	∞	-	-	-	-	-
14	-	-	-	-	fail	-

Table 3: Results from our modular algorithm

Sys	`Triang.`	Mem.	Size	`gsolve`	Mem.	Size
1	0.4	4	169	0.2	3	239
2	2	6	1680	504	18	34375
3	512	275	6250	1041	34	27624
4	2.5	4	743	-	fail	-
5	5	5	3134	9	5	2236
6	3000	250	2695	4950	66	34932
7	-	fail	-	1050	31	31115
8	-	fail	-	-	error	-
9	1593	18	55592	-	fail	-
10	∞	-	-	-	fail	-
11	-	fail	-	-	fail	-
12	-	fail	-	-	fail	-
13	-	fail	-	∞	-	-
14	-	fail	-	-	fail	-

Table 4: Results from `Triangularize` and `gsolve`

compared with two other MAPLE solvers, `Triangularize`, from the `RegularChains` library, and `gsolve`, from the `Groebner` library. Remark that the triangular decompositions modulo a prime, that are needed in our algorithm, are performed by `Triangularize`. This function is a generic code:

essentially the same code is used over $\mathbb{Z}$ and modulo a prime. Thus, `Triangularize` is not optimized for modular computations.

Our computations are done on a 2799 MHz Pentium 4. For the time being our implementation handles square systems that generate radical ideals. We compare our algorithm called `TriangularizeModular` with `gsolve` and `Triangularize`;

For each benchmark system, Table 1 lists the numbers n, d, h $\mathfrak{h}$ and Table 2 lists the prime p_1, the *a priori* and actual number of lifting steps ($\mathfrak{d}$ and a) and the maximal height of the output coefficients (C_a). Table 3 gives the time of one call to `Triangularize` modulo p_1 (Δ_p), the equiprojectable decomposition (E_p), and the lifting (Lift.) in seconds — the first two steps correspond to the "oracle calls" O_2 mentioned in Theorem 2, which will be studied in [6]. Table 3 gives also the total time, the total memory usage and output size for `TriangularizeModular`, whereas Table 4 gives that data for `Triangularize` and `gsolve`.

The maximum time is set up to 10800 seconds; we set the probability of success to be at least 90%.

`TriangularizeModular` solves 12 of the 14 test systems before the timeout, while `Triangularize` succeeds with 7 and `gsolve` with 6. Among most of the problems which `gsolve` can solve, `TriangularizeModular` shows less time consumed, less memory usage, and smaller output size. Noticeably, quite a few of the large systems can be solved by `TriangularizeModular` with time extension: system 13 is solved in 18745 seconds. Another interesting system is Pinchon-1 (from the FRISCO project), for which $n = 29, d = 16, h = 20, \mathfrak{h} = 1409536095e + 29$, which we solve in 64109 seconds. Both `Triangularize` and `gsolve` fail these problems due to memory allocation failure. Our modular method demonstrates its efficiency in reducing the size of the intermediate computations, whence its ability to solve difficult problems.

We observed that for every test system, for which E_p can be computed, the Hensel lifting always succeeds, *i.e.* the equiprojectable decomposition over $\mathbb{Q}$ can be reconstructed from E_p. In addition, `TriangularizeModular` failed `chemkin` at the Δ_p phase rather than at the lifting stage. Furthermore, the time consumed in the equiprojectable decomposition and the Hensel lifting is rather insignificant comparing with that in triangular decomposition modulo a prime. For every tested example the Hensel lifting achieves its final goal in less steps than the theoretical bound. In addition, the primes derived from our theoretical bounds are of quite moderate size, even on large examples.

6. CONCLUSIONS

We have presented a modular algorithm for triangular decompositions of 0-dimensional varieties over $\mathbb{Q}$ and have demonstrated the feasibility of Hensel lifting in computing triangular decompositions of non-equiprojectable varieties. Experiments show the capacity of this approach to improve the practical efficiency of triangular decomposition.

By far, the bottleneck is the modular triangularization phase. This is quite encouraging, since it is the part for which we relied on generic, non-optimized code. The next step is to extend these techniques to specialize variables as well during the modular phase, following the approach initiated in [13] for primitive element representations, and treat systems of positive dimension.

Acknowledgment

The authors are thankful to François Lemaire (Université de Lille 1, France) for his support with the `RegularChains` library. *Merci, François !*

7. REFERENCES

[1] E. A. Arnold. Modular algorithms for computing Gröbner bases. *J. Symb. Comp.*, 35(4):403–419, 2003.

[2] P. Aubry and A. Valibouze. Using Galois ideals for computing relative resolvents. *J. Symb. Comp.*, 30(6):635–651, 2000.

[3] F. Boulier, L. Denis-Vidal, T. Henin, and F. Lemaire. Lépisme. In *ICPSS*, pages 23–27. University of Paris 6, France, 2004.

[4] F. Boulier and F. Lemaire. Computing canonical representatives of regular differential ideals. In *ISSAC 2000*, pages 37–46. ACM Press, 2000.

[5] X. Dahan. Borne de hauteur (polynomiale) sur les coefficients d'une représentation triangulaire d'une variété zéro-dimensionnelle présentant des symétries. Master's thesis, École Polytechnique, 2003.

[6] X. Dahan, M. Moreno Maza, É. Schost, W. Wu, and Y. Xie. The complexity of the Split-and-Merge algorithm. In preparation.

[7] X. Dahan, M. Moreno Maza, É. Schost, W. Wu, and Y. Xie. On the complexity of the D5 principle. Preprint.

[8] X. Dahan, M. Moreno Maza, É. Schost, W. Wu, and Y. Xie. Equiprojectable decompositions of zero-dimensional varieties. In *ICPSS*, pages 69–71. University of Paris 6, France, 2004.

[9] X. Dahan and É. Schost. Sharp estimates for triangular sets. In *ISSAC 04*, pages 103–110. ACM Press, 2004.

[10] D. Eisenbud. *Commutative algebra*, volume 150 of *GTM*. Springer-Verlag, 1995.

[11] M.V. Foursov and M. Moreno Maza. On computer-assisted classification of coupled integrable equations. *J. Symb. Comp.*, 33:647–660, 2002.

[12] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 1999.

[13] M. Giusti, J. Heintz, J. E. Morais, and L. M. Pardo. When polynomial equation systems can be solved fast? In *AAECC-11*, pages 205–231. Springer, 1995.

[14] M. Giusti, G. Lecerf, and B. Salvy. A Gröbner free alternative for polynomial system solving. *J. Complexity*, 17(1):154–211, 2001.

[15] É. Hubert. Notes on triangular sets and triangulation-decomposition algorithms. In *Symbolic and Numerical Scientific Computations*, volume 2630 of *LNCS*, pages 1–39. Springer, 2003.

[16] M. Kalkbrener. A generalized euclidean algorithm for computing triangular representations of algebraic varieties. *J. Symb. Comp.*, 15:143–167, 1993.

[17] T. Krick, L. M. Pardo, and M. Sombra. Sharp estimates for the arithmetic Nullstellensatz. *Duke Math. J.*, 109(3):521–598, 2001.

[18] D. Lazard. Solving zero-dimensional algebraic systems. *J. Symb. Comp.*, 13:117–133, 1992.

[19] F. Lemaire, M. Moreno Maza, and Y. Xie. The `RegularChains` library. In *Maple 10*, Maplesoft, Canada. To appear.

[20] P. J. McCarthy. *Algebraic extensions of fields*. Dover, New York, 1991.

[21] M. Moreno Maza. On triangular decompositions of algebraic varieties. Technical Report 4/99, NAG, UK, Presented at the MEGA-2000 Conference, Bath, UK. http://www.csd.uwo.ca/∼moreno.

[22] M. Moreno Maza and R. Rioboo. Polynomial gcd computations over towers of algebraic extensions. In *Proc. AAECC-11*, pages 365–382. Springer, 1995.

[23] F. Rouillier. Solving zero-dimensional systems through the rational univariate representation. *AAECC*, 9:433–461, 1999.

[24] É. Schost. Complexity results for triangular sets. *J. Symb. Comp.*, 36(3-4):555–594, 2003.

[25] The symbolicdata project, 2000–2002. http://www.SymbolicData.org.

[26] W. Trinks. On improving approximate results of Buchberger's algorithm by Newton's method. In *EUROCAL 85*, volume 203 of *LNCS*, pages 608–611. Springer, 1985.