

A Fast Parallel Sparse Polynomial GCD Algorithm

Jiaxiong Hu, Michael Monagan
Department of Mathematics, Simon Fraser University
jha107@sfu.ca, mmonagan@sfu.ca

Abstract

We present a parallel GCD algorithm for sparse multivariate polynomials with integer coefficients. The algorithm combines a Kronecker substitution with a Ben-Or/Tiwari sparse interpolation modulo a smooth prime to determine the support of the GCD. We have implemented our algorithm in C for primes of various size and have parallelized it using Cilk C. We compare our implementation with Maple and Magma's serial implementations of Zippel's GCD algorithm.

1 Introduction

Let A and B be two polynomials in $\mathbb{Z}[x_0, x_1, \dots, x_n]$. In this paper we present a modular GCD algorithm for computing $G = \gcd(A, B)$ the greatest common divisor of A and B which is designed for sparse A and B . We will compare our algorithm with Zippel's sparse modular GCD algorithm from [47]. Zippel's algorithm is the main GCD algorithm currently used by the computer algebra systems Fermat, Maple, Magma and Mathematica for polynomials in $\mathbb{Z}[x_0, x_1, \dots, x_n]$.

Multivariate polynomial GCD computation was a central problem in Computer Algebra in the 1970's and 1980's. Whereas classical algorithms for polynomial multiplication and exact division are sufficient for many inputs, this is not the case for polynomial GCD computation. Euclid's algorithm, and variants of it such as the reduced PRS algorithm [8] and the subresultant PRS algorithm [6], produce an n dimensional intermediate expression swell where the intermediate polynomials are usually much larger than A, B and G . This renders these algorithms useless even for inputs of a very modest size.

There are four ideas in the literature that have been used to avoid or reduce this intermediate expression swell. The first was Brown's dense modular GCD algorithm [5] which interpolates x_1, x_2, \dots, x_n in G from univariate images in x_0 . The second was Moses and Yun's EZ-GCD algorithm [33] which uses multivariate Hensel lifting instead of interpolation, to recover x_1, x_2, \dots, x_n in G . The third is the heuristic GCD algorithm of Char, Geddes and Gonnet [7] which reduces a polynomial GCD problem to a single large integer GCD. The fourth by Sasaki and Suzuki [40] is to use truncated power-series for x_1, \dots, x_n when executing the Euclidean algorithm with x_0 the main variable.

Because polynomials in many variables in practice are often sparse, subsequent research focused on trying to improve the cost of the GCD algorithms when G is sparse. Zippel's landmark sparse polynomial interpolation method in [47] was developed to solve the sparse GCD problem. It is one of the earliest probabilistic algorithms. It interpolates x_1, x_2, \dots, x_n one variable at a time. Rayes, Wang and Weber [39] parallelized

parts of Zippel’s algorithm for shared memory computers. Kaltofen and Trager’s [24] “black box” GCD algorithm also uses sparse polynomial interpolation.

Wang’s EEZ-GCD algorithm [45] improved on the EZ-GCD algorithm by choosing zero evaluation points and Hensel lifting the variables one at a time to preserve sparsity. Wang’s algorithm was used as the default GCD algorithm in the Reduce, Maple, Magma, and Macsyma computer algebra systems. When the evaluation points cannot be zero, Wang’s algorithm fails to preserve sparsity. Several approaches to address this include the work of Kaltofen [23], Tsuji [44], and Monagan and Tuncer [35].

For the interested reader, Chapter 7 of [14] provides a description of the algorithms in [5, 45, 47, 7]. We now provide an overview of our GCD algorithm. Our GCD algorithm also uses sparse polynomial interpolation.

Let $A = \sum_{i=0}^{d_A} a_i x_0^i$, $B = \sum_{i=0}^{d_B} b_i x_0^i$ and $G = \sum_{i=0}^{d_G} c_i x_0^i$ where $d_A > 0$, $d_B > 0$ and the coefficients a_i , b_i and c_i are in $\mathbb{Z}[x_1, \dots, x_n]$. Our GCD algorithm first computes and removes contents, that is computes $\text{cont}(A, x_0) = \gcd(a_i)$ and $\text{cont}(B, x_0) = \gcd(b_i)$. These GCD computations in $\mathbb{Z}[x_1, x_2, \dots, x_n]$ are computed recursively.

Let $\bar{A} = A/G$ and $\bar{B} = B/G$ be the cofactors of A and B respectively. Let $\#A$ denote the number of terms in A and let $\text{Supp}(A)$ denote the set of monomials appearing in A . Let $\text{LC}(A)$ denote the leading coefficient of A taken in x_0 and let $\Gamma = \gcd(\text{LC}(A), \text{LC}(B)) = \gcd(a_{d_A}, b_{d_B})$. Since $\text{LC}(G) | \text{LC}(A)$ and $\text{LC}(G) | \text{LC}(B)$ it must be that $\text{LC}(G) | \Gamma$ thus $\Gamma = \text{LC}(G)\Delta$ for some polynomial $\Delta \in \mathbb{Z}[x_1, \dots, x_n]$.

Example 1. If $G = x_1 x_0^2 + x_2 x_0 + 3$, $\bar{A} = (x_2 - x_1)x_0 + x_2$ and $\bar{B} = (x_2 - x_1)x_0 + x_1 + 2$ we have $\#G = 3$, $\text{Supp}(G) = \{x_1 x_0^2, x_2 x_0, 1\}$, $\text{LC}(G) = x_1$, $\Gamma = x_1(x_2 - x_1)$, and $\Delta = x_2 - x_1$.

Let $H = \Delta \times G$ and $h_i = \Delta \times c_i$ so that $H = \sum_{i=0}^{d_G} h_i x_0^i$. Our algorithm will compute H not G . After computing H it must then compute $\text{cont}(H, x_0) = \gcd(h_i) = \Delta$ and divide H by Δ to obtain G . We compute H modulo a sequence of primes p_1, p_2, \dots , and recover the integer coefficients of H using Chinese remaindering. The use of Chinese remaindering is standard. Details may be found in [5, 14]. Let H_1 be the result of computing $H \bmod p_1$. For the remaining primes we use the sparse interpolation approach of Zippel [47, 48] which assumes $\text{Supp}(H_1) = \text{Supp}(H)$. Let us focus on the computation of $H \bmod p_1$.

To compute $H \bmod p$ the first prime p our algorithm will pick a sequence of evaluation points β_1, β_2, \dots from \mathbb{Z}_p^n , compute monic images

$$g_j := \gcd(A(x_0, \beta_j), B(x_0, \beta_j)) \in \mathbb{Z}_p[x_0]$$

of G then multiply g_j by the scalar $\Gamma(\beta_j) \in \mathbb{Z}_p$. Because the scaled image $\Gamma(\beta_j) \times g_j(x_0)$ is an image of a polynomial, H , we can use polynomial interpolation to interpolate each coefficient $h_i(x_1, \dots, x_n)$ of H from the coefficients of the scaled images.

Let $t = \max_{i=0}^{d_G} \#h_i$. The parameter t measures the sparsity of H . Let $d = \max_{i=1}^n \deg_{x_i} H$ and $D = \max_{i=0}^{d_G} \deg h_i$. The cost of sparse polynomial interpolation algorithms is determined mainly by the number of points β_1, β_2, \dots needed and the size of the prime p needed. These depend on n , t , d and D . Table 1 below presents data for several sparse polynomial interpolation algorithms. In Table 1 p_n denotes the n ’th prime which has size $O(\log n \log \log n)$ bits.

To get a sense of how large the prime needs to be for the different algorithms in Table 1 we include data for the following **benchmark problem**: Let G, \bar{A}, \bar{B} have nine variables ($n = 8$), have degree $d = 20$ in each variable, and have total degree $D = 60$

	#points	size of p	benchmark
Zippel [1979]	$O(ndt)$	$p > 2nd^2t^2$	$= 6.4 \times 10^9$
BenOr/Tiwari [1988]	$2t + O(1)$	$p > p_n^D$	$= 5.3 \times 10^{77}$
Monagan/Javadi [2010]	$O(nt)$	$p > nDt^2$	$= 4.8 \times 10^8$
Murao/Fujise [1996]	$2t + O(1)$	$p > (d+1)^n$	$= 3.7 \times 10^{10}$

Table 1: Some sparse interpolation algorithms

(to better reflect real problems). Let G have 10,000 terms with $t = 1000$. Let \bar{A} and \bar{B} have 100 terms so that $A = G\bar{A}$ and $B = G\bar{B}$ have about one million terms.

Notes on the sparse interpolation methods: Zippel’s sparse interpolation algorithm [47] is probabilistic. It was developed for polynomial GCD computation and implemented in Macsyma by Zippel. Kaltofen and Lee showed in [27] how to modify Zippel’s algorithm so that it will work effectively for primes much smaller than $2nd^2t^2$. The implementation of Zippel’s algorithm in Maple is due to de Kleine, Monagan and Wittkopf [29]. They use a different strategy to scale the univariate images $g_j(x_0)$.

The Ben-Or/Tiwari algorithm [3] is deterministic. The primary disadvantage of the Ben-Or/Tiwari algorithm is the size of the prime. In [22], Monagan and Javadi modify the Ben-Or/Tiwari algorithm to work for a smaller prime but using $O(nt)$ points. Murao and Fujise’s method [37] is a modification of the Ben-Or/Tiwari algorithm which computes discrete logarithms in the cyclic group \mathbb{Z}_p^* . We call this method the “discrete logs” method. We give details for it in Section 1.3. The advantage over the Ben-Or/Tiwari algorithm is that the prime size is $O(n \log d)$ bits instead of $O(D \log n \log \log n)$ bits.

Our goal is to design a GCD algorithm that recovers $H \bmod p$ using $O(t)$ points and a prime p of size $O(n \log d)$ bits. Also, for a multi-core computer with N cores we want to compute N of these $O(t)$ images in parallel. But, in a GCD algorithm that uses interpolation from values, not all evaluation points can be used. Let $\beta_j \in \mathbb{Z}_p^n$ be an evaluation point. If $\gcd(\bar{A}(x_0, \beta_j), \bar{B}(x_0, \beta_j)) \neq 1$ then β_j is said to be *unlucky* and this image cannot be used to interpolate H . Section 1.4 characterizes which evaluation points are unlucky and describes how they can be detected. In Zippel’s algorithm from [47], where the β_j are chosen at random from \mathbb{Z}_p^n , unlucky β_j , once identified, can simply be skipped. This is not the case for the geometric evaluation point sequences used by the Ben-Or/Tiwari algorithm and the discrete logs method. In Section 2, we consider whether these point sequences can be modified to handle unlucky evaluation points.

To reduce the probability of encountering unlucky evaluation points, the prime p may need to be larger than that shown in Table 1. Our modification for the discrete logarithm sequence increases the size of p which negates much of its advantage. This led us to consider using a Kronecker substitution K_r on x_1, x_2, \dots, x_n to map the GCD computation into a bivariate computation in $\mathbb{Z}_p[x_0, y]$. Some Kronecker substitutions result in all evaluation points being unlucky so they cannot be used. We call these Kronecker substitutions *unlucky*. In Section 2 we show (Theorem 1) that there are only finitely many of them and how to detect them so that another one may be tried.

If a Kronecker substitution is not unlucky there can still be many unlucky evaluation points because the degree of the resulting polynomials $K_r(A)$ and $K_r(B)$ in y is exponential in n . In order to avoid unlucky evaluation points one may simply choose the prime $p \gg \max(\deg_y(K_r A), \deg_y(K_r B))$, which is what we do for our “simplified” version of our GCD algorithm. But this may mean p is not a machine prime which will significantly increase the cost of all modular arithmetic in \mathbb{Z}_p as multi-precision

arithmetic is needed. However, it is well known in the computer algebra research community that unlucky evaluation points are infact rare. This prompted us to investigate the distribution of the unlucky evaluation points. Our next contribution (Theorem 3) is a result for the expected number of unlucky evaluations. This theorem justifies our “faster” version of our GCD algorithm which first tries a smaller prime.

In Section 3 we assemble a “Simplified Algorithm” which is a Las Vegas GCD algorithm. It first applies a Kronecker substitution to map the GCD computation into $\mathbb{Z}[x_0, y]$. It then chooses p randomly from a large set of smooth primes and computes $H \bmod p$ using sparse interpolation in y then uses further primes and Chinese remaindering to recover the integer coefficients in H . The algorithm chooses a Kronecker substitution large enough to be a priori not unlucky. It also assumes a term bound $\tau \geq \max \#h_i$ is given. These assumptions lead to a much simpler algorithm.

In Section 4, we relax the term bound requirement and we first try a Kronecker substitution just large enough to recover H . This complicates significantly the GCD algorithm. In Section 4 we present a heuristic GCD algorithm which we can prove always terminates and outputs $H \bmod p$. The heuristic algorithm will usually be much faster than the simplified algorithm but it can, in theory, fail several times before it finds a Kroenecker substitution K_r , a sufficiently large prime p , and evaluation points β_j which are all good.

We have implemented our algorithm in C and parallelized it using Cilk C [11]. We did this initially for 31 bit primes then for 63 bit primes and then for 127 bit primes to handle polynomials in more variables. The first timing results revealed that almost all the time was spent in evaluating $A(x_0, \beta_j)$ and $B(x_0, \beta_j)$ and not interpolating H . In Section 5.1 we describe an improvement for evaluation and how we parallelized it.

In Section 6 we compare our new algorithm with the C implementations of Zippel’s algorithm in Maple and Magma. The timing results are very promising. For our benchmark problem, Maple takes 22,111 seconds, Magma takes 1,611 seconds. and our new algorithm takes 4.67 seconds on 16 cores.

If $\#\Delta > 1$ then the number of terms in H may be (much) larger than the number of terms in G . Sections 5.2 and 5.3 describe two practical improvements, homogenization and interpolation of G from bivariate images, to reduce $\#\Delta$ and hence reduce t . The second improvement reduces the time for our benchmark problem from 4.67 seconds to 0.652 seconds on 16 cores. Both improvements would also benefit Zippel’s sparse GCD algorithm [47].

We conclude this Section by citing also the sparse interpolation methods of Garg & Schost [12], Giesbrecht & Roche [17] and Arnold, Giesbrecht and Roche [1] which can use a smaller prime and would also use fewer than $2t + O(1)$ evaluations. These methods would compute $a_i = K_r(A)(x, y)$, $b_i = K_r(B)(x, y)$ and $g_i = \gcd(a_i, b_i)$ all mod $\langle p, y^{q_i} - 1 \rangle$ for several primes q_i and recover the exponents of y in $K_r(H)$ using Chinese remaindering. The algorithms differ in the size of q_i and how they avoid and recover from exponent collisions modulo q_i . It is not clear whether this approach can work for the GCD problem as these methods assume a division free evaluation but computing g_i modulo $\langle p, y^{q_i} - 1 \rangle$ requires division and $y = 1$ may be bad or unlucky. These methods also require $q_i \gg t$ which means computing g_i modulo $\langle p, y^{q_i} - 1 \rangle$ will be expensive for large t .

1.1 Some notation and results

The proofs in the paper make use of properties of the Sylvester resultant, the Schwartz-Zippel Lemma and require bounds for the size of the integer coefficients appearing in

certain polynomials. We state these results here for later use.

Let $f = \sum_{i=1}^t a_i M_i$ where $a_i \in \mathbb{Z}$, $a_i \neq 0$, $t \geq 1$ and M_i is a monomial in n variables x_1, x_2, \dots, x_n . We denote by $\deg f$ the total degree of f , $\deg_{x_i} f$ the degree of f in x_i , and $\#f$ the number of terms of f . We need to bound the size of the integer coefficients of certain polynomials. For this purpose let $\|f\|_1 = \sum_{i=1}^t |a_i|$ be the one-norm of f and $\|f\| = \max_{i=1}^t |a_i|$ be the height of f . For a prime p , let ϕ_p denote the modular mapping $\phi_p(f) = f \pmod p$.

Lemma 1 (Schwartz-Zippel [41, 47]). *Let F be a field and $f \in F[x_1, x_2, \dots, x_n]$ be non-zero with total degree D and let $S \subset F$. If β is chosen at random from S^n then $\text{Prob}[f(\beta) = 0] \leq \frac{D}{|S|}$. Hence if $R = \{\beta | f(\beta) = 0\}$ then $|R| \leq D|S|^{n-1}$.*

Lemma 2. *Gelfond [15] Lemma II page 135. Let f be a polynomial in $\mathbb{Z}[x_1, x_2, \dots, x_n]$ and let d_i be the degree of f in x_i . If g is any factor of f over \mathbb{Z} then $\|g\| \leq e^{d_1 + d_2 + \dots + d_n} \|f\|$ where $e \approx 2.71828$.*

Let A be an $m \times m$ matrix with $A_{i,j} \in \mathbb{Z}$. Hadamard's bound $H(A)$ for $|\det(A)|$ is

$$|\det A| \leq \prod_{i=1}^m \sqrt{\sum_{j=1}^m A_{i,j}^2} = H(A).$$

Lemma 3. *Goldstein & Graham [19] Let A be an $m \times m$ matrix with entries $A_{i,j} \in \mathbb{Z}[y]$. Let B be the $m \times m$ integer matrix with $B_{i,j} = \|A_{i,j}\|_1$. Then $\|\det A\| \leq H(B)$.*

For polynomials $A = \sum_{i=0}^s a_i x_0^i$ and $B = \sum_{i=0}^t b_i x_0^i$, Sylvester's matrix is the following $s+t$ by $s+t$ matrix

$$S = \begin{bmatrix} a_s & 0 & & 0 & b_t & 0 & & 0 \\ a_{s-1} & a_s & & 0 & b_{t-1} & b_t & & 0 \\ \vdots & a_{s-1} & \ddots & 0 & \vdots & b_{t-1} & \ddots & 0 \\ a_1 & \vdots & & a_s & b_1 & \vdots & & b_t \\ a_0 & a_1 & & a_{s-1} & b_0 & b_1 & & b_{t-1} \\ 0 & a_0 & & \vdots & 0 & b_0 & & \vdots \\ 0 & 0 & \ddots & a_1 & 0 & 0 & \ddots & b_1 \\ 0 & 0 & & a_0 & 0 & 0 & & b_0 \end{bmatrix}. \quad (1)$$

where the coefficients of A are repeated in the first t columns and the coefficients of B are repeated in the last s columns. The Sylvester resultant of the polynomials A and B in x , denoted $\text{res}_x(A, B)$, is the determinant of Sylvester's matrix. We gather the following facts about it into Lemma 4 below.

Lemma 4. *Let D be any integral domain and let A and B be two polynomials in $D[x_0, x_1, \dots, x_n]$ with $s = \deg_{x_0} A > 0$ and $t = \deg_{x_0} B > 0$. Let $a_s = LC(A)$, $b_t = LC(B)$, $R = \text{res}_{x_0}(A, B)$, $\alpha \in D^n$ and p be a prime. Then*

- (i) R is a polynomial in $D[x_1, \dots, x_n]$,
- (ii) $\deg R \leq \deg A \deg B$ (Bezout bound) and
- (iii) $\deg_{x_i} R \leq t \deg_{x_i} A + s \deg_{x_i} B$ for $1 \leq i \leq n$.

If D is a field and $a_s(\alpha) \neq 0$ and $b_t(\alpha) \neq 0$ then

- (iv) $\text{res}_{x_0}(A(x_0, \alpha), B(x_0, \alpha)) = R(\alpha)$ and
(v) $\deg_{x_0} \gcd(A(x_0, \alpha), B(x_0, \alpha)) > 0 \iff \text{res}_{x_0}(A(x_0, \alpha), B(x_0, \alpha)) = 0$.

If $D = \mathbb{Z}$ and $\phi_p(a_s) \neq 0$ and $\phi_p(b_t) \neq 0$ then

- (vi) $\text{res}_{x_0}(\phi_p(A), \phi_p(B)) = \phi_p(R)$ and
(vii) $\deg_{x_0} \gcd(\phi_p(A), \phi_p(B)) > 0 \iff \text{res}_{x_0}(\phi_p(A), \phi_p(B)) = 0$.

Proofs of (i), (ii), (iv) and (v) may be found in Ch. 3 and Ch. 6 of [10]. In particular the proof in Ch. 6 of [10] for (ii) for bivariate polynomials generalizes to the multivariate case. Note that the condition on α that the leading coefficients a_s and b_t do not vanish means that the dimension of Sylvester's matrix for $A(x_0, \alpha)$ and $B(x_0, \alpha)$ is the same as that for A and B which proves (v). The same argument used to prove (iv) and (v) works for (vi) and (vii). To prove (iii) we have

$$\deg_{x_i} \det S \leq \sum_{c \in \text{columns}(S)} \max_{f \in c} \deg_{x_i} f = \sum_{j=1}^t \deg_{x_i} A + \sum_{j=1}^s \deg_{x_i} B.$$

1.2 Ben-Or Tiwari Sparse Interpolation

Let $C(x_1, \dots, x_n) = \sum_{i=1}^t a_i M_i$ where $a_i \in \mathbb{Z}$ and M_i are monomials in (x_1, \dots, x_n) . In our context, C represents one of the coefficients of $H = \Delta G$ we wish to interpolate. Let $D = \deg C$ and let $d = \max_{i=1}^n \deg_{x_i} C$ and let p_n denote the n 'th prime. Let

$$v_j = C(2^j, 3^j, 5^j, \dots, p_n^j) \text{ for } j = 0, 1, \dots, 2t - 1.$$

The Ben-Or/Tiwari sparse interpolation algorithm [3] interpolates $C(x_1, x_2, \dots, x_n)$ from the $2t$ points v_j . Let $m_i = M_i(2, 3, 5, \dots, p_n) \in \mathbb{Z}$ and let $\lambda(z) = \prod_{i=1}^t (z - m_i) \in \mathbb{Z}[z]$. The algorithm proceeds in 5 steps.

- 1 Compute $v_j = C(2^j, 3^j, 5^j, \dots, p_n^j)$ for $j = 0, 1, \dots, 2t - 1$.
- 2 Compute $\lambda(z)$ from v_j using the Berlekamp-Massey algorithm [30] or the Euclidean algorithm [2, 43].
- 3 Compute the integer roots m_i of $\lambda(z)$.
- 4 Factor the integers m_i using trial division by $2, 3, \dots, p_n$ from which we obtain M_i . For example, for $n = 3$, if $m_i = 45000 = 2^3 3^2 5^4$ then $M_i = x_1^3 x_2^2 x_3^4$.
- 5 Solve the following $t \times t$ linear system for the unknown coefficients a_i in $C(x_1, \dots, x_n)$.

$$V a = \begin{bmatrix} 1 & 1 & \dots & 1 \\ m_1 & m_2 & \dots & m_t \\ m_1^2 & m_2^2 & \dots & m_t^2 \\ \vdots & \vdots & \vdots & \vdots \\ m_1^{t-1} & m_2^{t-1} & \dots & m_t^{t-1} \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ \vdots \\ a_t \end{bmatrix} = \begin{bmatrix} v_0 \\ v_1 \\ v_2 \\ \vdots \\ v_{t-1} \end{bmatrix} = b \quad (2)$$

The matrix V above is a transposed Vandermonde matrix. Recall that

$$\det V = \det V^T = \prod_{1 \leq j < k \leq t} (m_j - m_k).$$

Since the monomial evaluations $m_i = M_i(2, 3, 4, \dots, p_n)$ are distinct it follows that $Va = b$ has a unique solution. The linear system $Va = b$ can be solved in $O(t^2)$ arithmetic operations (see [48]). Note, the master polynomial $P(Z)$ in [48] is $\lambda(z)$.

Notice that the largest integer coefficient in $\lambda(z)$ is the constant term $\prod_{i=1}^t m_i$ which is at most p_n^{Dt} hence of size $O(tD \log n \log \log n)$ bits. Moreover, in [25], Kaltofen, Lashman and Wiley noticed that a severe expression swell occurs if either the Berlekamp-Massey algorithm or the Euclidean algorithm is used to compute $\lambda(z)$ over \mathbb{Q} . For our purposes, because we want to interpolate H modulo a prime p , we run Steps 2, 3, and 5 modulo p . Provided we pick $p > \max_{i=1}^t m_i \leq p_n^D$ the integers m_i remain unique modulo p and we recover the monomials $M_i(x_1, \dots, x_n)$ in Step 4 and the linear system in Step 5 has a unique solution modulo p . For Step 3, the roots of $\lambda(z) \in \mathbb{Z}_p[z]$ can be found using Berlekamp's algorithm [4] which has classical complexity $O(t^2 \log p)$.

In [3], Ben-Or and Tiwari assume a sparse term bound $T \geq t$ is known, that is, we are given some T such that $t \leq T \ll (d+1)^n$ and in Step 1 we may compute $2T$ evaluations in parallel. In practice such a bound on t may not be known in advance so the algorithm needs to be modified to also determine t . For p sufficiently large, if we compute $\lambda(z)$ after $j = 2, 4, 6, \dots$ points, we will see $\deg \lambda(z) = 1, 2, 3, \dots, t-1, t, t, t, \dots$ with high probability. Thus we may simply wait until the degree of $\lambda(z)$ does not change. This problem is first discussed by Kaltofen, Lee and Lobo in [26]. We will return to this in Section 4.5.

Steps 2, 3, and 5 may be accelerated with fast multiplication. Let $M(t)$ denote the number of arithmetic operations in \mathbb{Z}_p to multiply two polynomials of degree t in $\mathbb{Z}_p[t]$. The fast Euclidean algorithm can be used to accelerate Step 2. It does $O(M(t) \log t)$ arithmetic operations in \mathbb{Z}_p . See Ch. 11 of [13]. Computing the roots of $\lambda(z)$ in Step 3 can be done in $O(M(t) \log t \log(pt))$. See Corollary 14.16 of [13]. Step 5 may be done in $O(M(t) \log t)$ using fast interpolation. See Ch 10 of [13]. Thus for large t the root finding step dominates. We cite Grenet, van der Hoeven and Lecerf [18] who show how a smooth FFT prime of the form $p = 2^k s + 1$, e.g. $p = 2^{30} 3 + 1$, may be used to eliminate the $\log p$ factor and how to use a Graeffe transform to eliminate a $\log t$ factor to reduce the cost of the root finding step to $O(3M(t) \log t)$ arithmetic operations in \mathbb{Z}_p . We summarize these complexity results in Table 2 below.

Step	Classical	Fast	smooth FFT p
2	$O(t^2)$	$O(M(t) \log t)$	$O(M(t) \log t)$
3	$O(t^2 \log p)$	$O(M(t) \log t \log(pt))$	$O(M(t) \log t)$
5	$O(t^2)$	$O(M(t) \log t)$	$O(M(t) \log t)$

Table 2: Number of arithmetic operations in \mathbb{Z}_p for t monomials.

1.3 Ben-Or/Tiwari with discrete logarithms

The discrete logarithm method modifies the Ben-Or/Tiwari algorithm so that the prime needed is a little larger than $(d+1)^n$ thus of size is $O(n \log d)$ bits instead of $O(D \log n \log \log n)$. Murao & Fujise [37] were the first to try this approach. Some practical aspects of it are discussed by van der Hoven and Lecerf in [21]. We explain how the method works.

To interpolate $C(x_1, \dots, x_n)$ we first pick a prime p of the form $p = q_1 q_2 q_3 \dots q_n + 1$ satisfying $2|q_1$, $q_i > \deg_{x_i} C$ and $\gcd(q_i, q_j) = 1$ for $1 \leq i < j \leq n$. Finding such primes is not difficult and we omit presenting an explicit algorithm here.

Next we pick a random primitive element $\alpha \in \mathbb{Z}_p$ which we can do using the partial factorization $p - 1 = q_1 q_2 \dots q_n$ (see [42]). We set $\omega_i = \alpha^{(p-1)/q_i}$ so that $\omega_i^{q_i} = 1$ and replace the evaluation points $(2^j, 3^j, \dots, p_n^j)$ with $(\omega_1^j, \omega_2^j, \dots, \omega_n^j)$. After Step 2 we factor $\lambda(z)$ in $\mathbb{Z}_p[z]$ to determine the m_i . If $M_i = \prod_{k=1}^n x_k^{d_k}$ we have $m_i = \prod_{k=1}^n \omega_k^{d_k}$. To compute d_k in Step 4 we compute the discrete logarithm $x := \log_\alpha m_i$, that is, we solve $\alpha^x \equiv m_i \pmod{p}$ for $0 \leq x < p - 1$. We have

$$x = \log_\alpha m_i = \log_\alpha \prod_{k=1}^n \omega_k^{d_k} = \sum_{k=1}^n d_k \frac{p-1}{q_k}. \quad (3)$$

Taking (3) mod q_k we obtain $d_k = x[(p-1)/q_k]^{-1} \pmod{q_k}$. Note the condition $\gcd(q_i, q_j) = 1$ ensures $(p-1)/q_k$ is invertible mod q_k . Step 5 remains unchanged.

For $p = q_1 q_2 \dots q_n + 1$, a discrete logarithm can be computed in $O(\sum_{i=1}^n e_i (\log p + \sqrt{p_i}))$ multiplications in \mathbb{Z}_p using the Pohlig-Helman algorithm where the factorization of $p-1 = \prod_{i=1}^m p_i^{e_i}$. See [38, 42]. Since the $q_i \sim d$ this leads to an $O(n\sqrt{d})$ cost. Kaltofen showed in [28] that this can be made polynomial in $\log d$ and n if one uses a Kronecker substitution to reduce multivariate interpolation to a univariate interpolation and uses a prime $p > (d+1)^n$ of the form $p = 2^k s + 1$ with s small.

We cite also the work of Giesbrecht, Labahn and Lee in [16] for numeric sparse interpolation. They pick complex roots of unity $\omega_1, \dots, \omega_n$ of relatively prime orders q_1, \dots, q_n . The discrete logarithm computation becomes a numerical logarithm.

1.4 Bad and Unlucky Evaluation Points

Let A and B be non constant polynomials in $\mathbb{Z}[x_0, \dots, x_n]$, $G = \gcd(A, B)$ and $\bar{A} = A/G$ and $\bar{B} = B/G$. Let p be prime such that $LC(A)LC(B) \pmod{p} \neq 0$.

Definition 1. Let $\alpha \in \mathbb{Z}_p^n$ and let $\bar{g}_\alpha(x) = \gcd(\bar{A}(x, \alpha), \bar{B}(x, \alpha))$. We say α is bad if $LC(A)(\alpha) = 0$ or $LC(B)(\alpha) = 0$ and α is unlucky if $\deg \bar{g}_\alpha(x) > 0$. If α is not bad and not unlucky we say α is good.

Example 2. Let $G = (x_1 - 16)x_0 + 1$, $\bar{A} = x_0^2 + 1$ and $\bar{B} = x_0^2 + (x_1 - 1)(x_2 - 9)x_0 + 1$. Then $LC(A) = LC(B) = x_1 - 16$ so $\{(16, \beta) : \beta \in \mathbb{Z}_p\}$ are bad and $\{(1, \beta) : \beta \in \mathbb{Z}_p\}$ and $\{(\beta, 9) : \beta \in \mathbb{Z}_p\}$ are unlucky.

Our GCD algorithm cannot reconstruct G using the image $g_\alpha(x) = \gcd(A(x, \alpha), B(x, \alpha))$ if α is unlucky. Brown's idea in [5] to detect unlucky α is based on the following Lemma.

Lemma 5. Let α and g_α be as above and $h_\alpha = G(x, \alpha) \pmod{p}$. If α is not bad then $h_\alpha | g_\alpha$ and $\deg_x g_\alpha \geq \deg_x G$.

For a proof of Lemma 5 see Lemma 7.3 of [14]. Brown only uses α which are not bad and the images $g_\alpha(x)$ of least degree to interpolate G . The following Lemma implies if the prime p is large then unlucky evaluations points are rare.

Lemma 6. If α is chosen at random from \mathbb{Z}_p^n then

$$\text{Prob}[\alpha \text{ is bad or unlucky}] \leq \frac{\deg AB + \deg A \deg B}{p}.$$

PROOF: Let b be the number of bad evaluation points and let r be the number of unlucky evaluation points that are not also bad. Let B denote the event α is bad and G denote the event α is not bad and U denote the event α is unlucky. Then

$$\begin{aligned} \text{Prob}[B \text{ or } U] &= \text{Prob}[B] + \text{Prob}[G \text{ and } U] \\ &= \text{Prob}[B] + \text{Prob}[G] \times \text{Prob}[U|G] \\ &= \frac{b}{p^n} + \left(1 - \frac{b}{p^n}\right) \frac{r}{p^n - b} = \frac{b}{p^n} + \frac{r}{p^n}. \end{aligned}$$

Now α is bad $\implies \text{LC}(A)(\alpha)\text{LC}(B)(\alpha) = 0 \implies \text{LC}(AB)(\alpha) = 0$. Applying Lemma 1 with $f = \text{LC}(AB)$ we have $b \leq \deg \text{LC}(AB)p^{n-1}$. Let $R = \text{res}_{x_0}(\bar{A}, \bar{B}) \in \mathbb{Z}_p[x_1, \dots, x_n]$. Now α is unlucky and not bad $\implies \deg \gcd(\bar{A}(x, \alpha), \bar{B}(x, \alpha)) > 0$ and $\text{LC}(\bar{A})(\alpha) \neq 0$ and $\text{LC}(\bar{B})(\alpha) \neq 0 \implies R(\alpha) = 0$ by Lemma 4 (iv) and (v). Applying Lemma 1 we have $r \leq \deg(R)p^{n-1}$. Substituting into the above we have

$$\text{Prob}[B \text{ or } U] \leq \frac{\deg \text{LC}(AB)}{p} + \frac{\deg R}{p} \leq \frac{\deg AB}{p} + \frac{\deg A \deg B}{p} \quad \square$$

The following algorithm applies Lemma 6 to compute an upper bound d for $\deg_{x_i} G$.

Algorithm DegreeBound(A, B, i)

Input: Non-zero $A, B \in \mathbb{Z}[x_0, x_1, \dots, x_n]$ and i satisfying $0 \leq i \leq n$.

Output: $d \geq \deg_{x_i}(G)$ where $G = \gcd(A, B)$.

- 1 Set $LA = \text{LC}(A, x_i)$ and $LB = \text{LC}(B, x_i)$.
So $LA, LB \in \mathbb{Z}[x_0, \dots, x_{i-1}, x_{i+1}, \dots, x_n]$.
- 2 Pick a prime $p \gg \deg A \deg B$ such that $LA \pmod p \neq 0$ and $LB \pmod p \neq 0$.
- 3 Pick $\alpha = (\alpha_0, \dots, \alpha_{i-1}, \alpha_{i+1}, \dots, \alpha_n) \in \mathbb{Z}_p^n$ at random until $LA(\alpha)LB(\alpha) \neq 0$.
- 4 Compute $a = A(\alpha_0, \dots, \alpha_{i-1}, x_i, \alpha_{i+1}, \dots, \alpha_n)$ and
 $b = B(\alpha_0, \dots, \alpha_{i-1}, x_i, \alpha_{i+1}, \dots, \alpha_n)$.
- 5 Compute $g = \gcd(a, b)$ in $\mathbb{Z}_p[x_i]$ using the Euclidean algorithm.
- 6 Output $d = \deg_{x_i} g$.

2 Kronecker Substitutions

Consider again Example 2 from Section 1.4 where $G = (x_1 - 16)x_0 + 1$, $\bar{A} = x_0^2 + 1$ and $\bar{B} = x_0^2 + (x_1 - 1)(x_2 - 9)x_0 + 1$. For the Ben-Or/Tiwari points $\alpha_j = (2^j, 3^j)$ for $0 \leq j < 2t$ observe that $\alpha_0 = (1, 1)$ and $\alpha_2 = (4, 9)$ are unlucky and $\alpha_4 = (16, 81)$ is bad. Since none of these points can be used to interpolate G we need to modify the Ben-Or/Tiwari point sequence. For the GCD problem, we want random evaluation points to avoid bad and unlucky points. The following fix works.

Pick $0 < s < p$ at random and use $\alpha_j = (2^{s+j}, 3^{s+j}, \dots, p_n^{s+j})$ for $0 \leq j < 2t$. Steps 1, 2 and 3 work as before. To solve the *shifted* transposed Vandermonde system

$$Wc = \begin{bmatrix} m_1^s & m_2^s & \dots & m_t^s \\ m_1^{s+1} & m_2^{s+1} & \dots & m_t^{s+1} \\ \vdots & \vdots & \vdots & \vdots \\ m_1^{s+t-1} & m_2^{s+t-1} & \dots & m_t^{s+t-1} \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_t \end{bmatrix} = \begin{bmatrix} v_s \\ v_{s+1} \\ \vdots \\ v_{s+t-1} \end{bmatrix} = u.$$

we first solve the transposed Vandermonde system

$$Vb = \begin{bmatrix} 1 & 1 & \dots & 1 \\ m_1 & m_2 & \dots & m_t \\ \vdots & \vdots & \vdots & \vdots \\ m_1^{t-1} & m_2^{t-1} & \dots & m_t^{t-1} \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_t \end{bmatrix} = \begin{bmatrix} v_s \\ v_{s+1} \\ \vdots \\ v_{s+t-1} \end{bmatrix} = u$$

as before to obtain $b = V^{-1}u$. Observe that the matrix $W = VD$ where D is the t by t diagonal matrix with $D_{i,i} = m_i^s$. Solving $Wc = u$ for c we have

$$c = W^{-1}u = (VD)^{-1}u = (D^{-1}V^{-1})u = D^{-1}(V^{-1}u) = D^{-1}b.$$

Thus $c_i = b_i m_i^{-s}$ and we can solve $Wc = u$ in $O(t^2 + t \log s)$ multiplications.

Referring again to Example 2, suppose we use the discrete logarithm evaluation points $\alpha_j = (\omega_1^j, \omega_2^j)$ for $0 \leq j < 2t$. Observe that $\alpha_0 = (1, 1)$ is unlucky and also, since $\omega_1^{q_1} = 1$, all $\alpha_{q_1}, \alpha_{2q_1}, \alpha_{3q_1}, \dots$ are unlucky. Shifting the sequence to start at $j = 1$ and picking $q_i > 2t$ to avoid unlucky evaluation points is problematic because for the GCD problem as t may be larger than $\max\{\#a_i, \#b_i\}$, or smaller; there is no way to know in advance. This difficulty led us to consider using a Kronecker substitution.

2.1 Kronecker Substitutions and GCDs

We propose to use a Kronecker substitution to map a multivariate polynomial GCD problem in $\mathbb{Z}[x_0, x_1, \dots, x_n]$ into a bivariate GCD problem in $\mathbb{Z}[x, y]$. After making the Kronecker substitution, we need to interpolate $H(x, y) = \Delta(y)G(x, y)$ where $\deg_y H(x, y)$ will be exponential in n . To make discrete logarithms in \mathbb{Z}_p feasible, we follow Kaltofen [28] and pick $p = 2^k s + 1 > \deg_y H(x, y)$ with s small.

Definition 2. Let D be an integral domain and let f be a polynomial in $D[x_0, x_1, \dots, x_n]$. Let $r \in \mathbb{Z}^{n-1}$ with $r_i > 0$. Let $K_r : D[x_0, x_1, \dots, x_n] \rightarrow D[x, y]$ be the Kronecker substitution $K_r(f) = f(x, y, y^{r_1}, y^{r_1 r_2}, \dots, y^{r_1 r_2 \dots r_{n-1}})$.

Let $d_i = \deg_{x_i} f$ be the partial degrees of f for $1 \leq i \leq n$. We note that if $r_i > d_i$ for $1 \leq i \leq n-1$ then K_r is invertible and therefore $K_r(f) = 0 \iff f = 0$. Not all such Kronecker substitutions can be used, however, for the GCD problem. We consider an example.

Example 3. Consider the following GCD problem

$$G = x + y + z, \quad \bar{A} = x^3 - yz, \quad \bar{B} = x^2 - y^2$$

in $\mathbb{Z}[x, y, z]$. Since $\deg_y G = 1$ the Kronecker substitution $K_r(G) = G(x, y, y^2)$ is invertible. But $\gcd(K_r(\bar{A}), K_r(\bar{B})) = \gcd(\bar{A}(x, y, y^2), \bar{B}(x, y, y^2)) = \gcd(x^3 - y^3, x^2 - y^2) = x - y$. If we proceed to interpolate the $\gcd(K_r(\bar{A}), K_r(\bar{B}))$ we will obtain $(x - y)K_r(G)$ in expanded form from which we cannot recover G .

We call such a Kronecker substitution unlucky. Lemma 7 below says that if the r_i are sufficiently large then K_r is a good Kronecker substitution. We will use such a Kronecker substitution in Algorithm MGCD in Section 4.4 – see Step 2. .

Definition 3. Let K_r be a Kronecker substitution. We say K_r is bad if $\deg_x K_r(A) < \deg_{x_0} A$ or $\deg_x K_r(B) < \deg_{x_0} B$ and K_r is unlucky if $\deg_x \gcd(K_r(\bar{A}), K_r(\bar{B})) > 0$. If K_r is not bad and not unlucky we say K_r is good.

Lemma 7. Let $A, B \in \mathbb{Z}[x_0, x_1, \dots, x_n]$ with $\deg_{x_0} A > 0$ and $\deg_{x_0} B > 0$ and let $D_i = \deg_{x_i} A \deg_{x_0} B + \deg_{x_i} B \deg_{x_0} A$. The Kronecker substitution K_r with $r_i > D_i$ for $1 \leq i \leq n-1$ is good.

Proof. Let $G = \gcd(A, B)$ and $\bar{A} = A/G$ and $\bar{B} = B/G$ and let $LC(A)$ and $LC(B)$ the the leading coefficients of A and B with respect to x_0 .

If we pick $r_i > \max(\deg_{x_i} LC(A), \deg_{x_i} LC(B))$ then $K_r(LC(A)) \neq 0$ and $K_r(LC(B)) \neq 0$ so K_r is not bad. The assumption that $\deg_{x_0} A > 0$ and $\deg_{x_0} B > 0$ implies

$$\max\{\deg_{x_i} LC(A), \deg_{x_i} LC(B)\} \leq D_i\}$$

so K_r with $r_i > D_i$ is not bad.

Let $R = \text{res}_{x_0}(\bar{A}, \bar{B})$. By Lemma 4(iii), we have

$$\deg_{x_i} R \leq \deg_{x_0} \bar{B} \deg_{x_i} \bar{A} + \deg_{x_0} \bar{A} \deg_{x_i} \bar{B}.$$

Since $\deg_{x_i} \bar{A} \leq \deg_{x_i} A$ and $\deg_{x_i} \bar{B} \leq \deg_{x_i} B$ for $0 \leq i \leq n$ we have

$$\deg_{x_i} R \leq \deg_{x_0} B \deg_{x_i} A + \deg_{x_0} A \deg_{x_i} B = D_i.$$

So if we pick $r_i > D_i$ then $K_r(R) \neq 0$, that is, K_r is not unlucky. Thus a Kronecker substitution K_r with the sequence $r_i > D_i$ for $1 \leq i \leq n-1$ is good. \square

In order to use a smaller prime we would like to use a Kronecker substitution K_r with $r_i = d_i + 1, d_i + 2, \dots$ if possible. Theorem 1 below tells us that the number of unlucky Kronecker substitutions is finite. To detect them we will also avoid bad Kronecker substitutions in an analogous way Brown did to detect unlucky evaluation points.

Proposition 1. Let $f \in \mathbb{Z}[x_1, \dots, x_n]$ be non-zero and $d_i = \deg(f, x_i)$ for $1 \leq i \leq n$. Let X be the number of Kronecker substitutions K_r such that $K_r(f) = 0$ where

$$r \in \{[d_1 + k, d_2 + k, \dots, d_{n-1} + k] \text{ for } k = 1, 2, 3, \dots\}$$

Then $X \leq (n-1)\sqrt{2 \deg f}$.

$$\begin{aligned} \text{PROOF: } K_r(f) = 0 &\iff f(y, y^{r_1}, y^{r_1 r_2}, \dots, y^{r_1 r_2 \dots r_{n-1}}) = 0 \\ &\iff f \bmod \langle x_1 - y, x_2 - y^{r_1}, \dots, x_n - y^{r_1 r_2 \dots r_{n-1}} \rangle = 0 \\ &\iff f \bmod \langle x_2 - x_1^{r_1}, x_3 - x_2^{r_2}, \dots, x_n - x_{n-1}^{r_{n-1}} \rangle = 0. \end{aligned}$$

Thus X is the number of ideals $I = \langle x_2 - x_1^{r_1}, \dots, x_n - x_{n-1}^{r_{n-1}} \rangle$ for which $f \bmod I = 0$ with $r_i = d_i + 1, d_i + 2, \dots$. We prove that $X \leq (n-1)\sqrt{2 \deg f}$ by induction on n .

If $n = 1$ then I is empty so $f \bmod I = f$ and hence $X = 0$ and the Lemma holds. For $n = 2$ we have $f(x_1, x_2) \bmod \langle x_2 - x_1^{r_1} \rangle = 0 \implies x_2 - x_1^{r_1} | f$. Now X is maximal when $d_1 = 0$ and $r_1 = 1, 2, 3, \dots$. We have

$$\sum_{r_1=1}^X r_1 \leq \deg f \implies X(X+1)/2 \leq \deg f \implies X < \sqrt{2 \deg f}.$$

For $n > 2$ we proceed as follows. Either $x_n - x_{n-1}^{r_{n-1}}|f$ or it doesn't. If not then the polynomial $S = f(x_1, \dots, x_{n-1}, x_{n-1}^{r_{n-1}})$ is non-zero. For the sub-case $x_n - x_{n-1}^{r_{n-1}}|f$ we obtain at most $\sqrt{2 \deg f}$ such factors of f using the previous argument. For the case $S \neq 0$ we have

$$S \bmod I = 0 \iff S \bmod \langle x_2 - x_1^{r_1}, \dots, x_{n-2} - x_{n-1}^{r_{n-2}} \rangle = 0$$

Notice that $\deg_{x_i} S = \deg_{x_i} f$ for $1 \leq i \leq n-2$. Hence, by induction on n , $X < (n-2)\sqrt{2 \deg f}$ for this case. Adding the number of unlucky Kronecker substitutions for both cases yields $X \leq (n-1)\sqrt{2 \deg f}$. \square

Theorem 1. *Let $A, B \in \mathbb{Z}[x_0, x_1, \dots, x_n]$ be non-zero, $G = \gcd(A, B)$, $\bar{A} = A/G$ and $\bar{B} = B/G$. Let $d_i \geq \deg_{x_i} G$. Let X be the number of Kronecker substitutions K_r where $r \in \{[d_1 + k, d_2 + k, \dots, d_{n-1} + k] \text{ for } k = 1, 2, 3, \dots\}$ which are bad and unlucky. Then*

$$X \leq \sqrt{2}(n-1) \left[\sqrt{\deg A} + \sqrt{\deg B} + \sqrt{\deg A \deg B} \right].$$

PROOF: Let $LA = LC(A)$ and $LB = LC(B)$ be the leading coefficients of A and B in x_0 . Then K_r is bad $\iff K_r(LA) = 0$ or $K_r(LB) = 0$. Applying Proposition 1, the number of bad Kronecker substitutions is at most

$$(n-1)(\sqrt{2 \deg LA} + \sqrt{2 \deg LB}) \leq (n-1)(\sqrt{2 \deg A} + \sqrt{2 \deg B}).$$

Now let $R = \text{res}_{x_0}(\bar{A}, \bar{B})$. We will assume K_r is not bad.

$$\begin{aligned} K_r \text{ is unlucky} &\iff \deg_x(\gcd(K_r(\bar{A}), K_r(\bar{B}))) > 0 \\ &\iff \text{res}_x(K_r(\bar{A}), K_r(\bar{B})) = 0 \\ (K_r \text{ is not bad and is a homomorphism}) &\iff K_r(\text{res}_x(\bar{A}, \bar{B})) = 0 \\ &\iff K_r(R) = 0 \end{aligned}$$

By Proposition 1, the number of unlucky Kronecker substitutions $\leq (n-1)\sqrt{2 \deg R} \leq (n-1)\sqrt{2 \deg A \deg B}$ by Lemma 4(ii). Adding the two contributions proves the theorem. \square

Theorem 1 tells us that the number of unlucky Kronecker substitutions is finite. Algorithm MGCD1 in Section 4, after identifying an unlucky Kronecker substitution will try the next Kronecker substitution $r = [r_1 + 1, r_2 + 1, \dots, r_{n-1} + 1]$.

It is still not obvious that a Kronecker substitution that is not unlucky can be used because it can create a content in y of exponential degree. The following example shows how we recover $H = \Delta G$ when this happens.

Example 4. *Consider the following GCD problem*

$$G = wx^2 + zy, \quad \bar{A} = ywx + z, \quad \bar{B} = yzx + w$$

in $\mathbb{Z}[x, y, z, w]$. We have $\Gamma = wy$ and $\Delta = y$. For $K_r(f) = f(x, y, y^3, y^9)$ we have

$$\gcd(K_r(A), K_r(B)) = K_r(G) \gcd(y^{10}x + y^3, y^4x + y^9) = (y^9x^2 + y^4)y^3 = y^7(y^5x^2 + 1).$$

One must not try to compute $\gcd(K_r(A), K_r(B))$ because the degree of the content of $\gcd(K_r(A), K_r(B))$ (y^7 in our example) can be exponential in n the number of variables and we cannot compute this efficiently using the Euclidean algorithm. The crucial observation is that if we compute **monic** images $g_j = \gcd(K_r(A)(x, \alpha^j), K_r(B)(x, \alpha^j))$ any content is divided out, and when we scale by $K_r(\Gamma)(\alpha^j)$ and interpolate y in $K_r(H)$ using sparse interpolation, we recover any content. We obtain $K_r(H) = K_r(\Delta)K_r(G) = y^{10}x^2 + y^5$, then invert K_r to obtain $H = (yw)x^2 + (y^2z)$.

2.2 Unlucky primes

Let A, B be polynomials in $\mathbb{Z}[x_0, x_1, \dots, x_n]$, $G = \gcd(A, B)$, $\bar{A} = A/G$ and $\bar{B} = B/G$. In the introduction we defined the polynomials $\Gamma = \gcd(LC(A), LC(B))$, $\Delta = \Gamma/LC(G)$ and $H = \Delta G$ where $LC(A)$, $LC(B)$ and $LC(G)$ are the leading coefficients of A , B and G in x_0 respectively.

Let $K_r : \mathbb{Z}[x_0, x_1, \dots, x_n] \rightarrow \mathbb{Z}[x, y]$ be a Kronecker substitution $K_r(f) = f(x, y, y^{r_1}, y^{r_1 r_2}, \dots, y^{r_1 r_2 \dots r_{n-1}})$ for some $r_i > 0$. Our GCD algorithm will compute $\gcd(K_r(A), K_r(B))$ modulo a prime p . Some primes cannot be used.

Example 5. Consider the following GCD problem in $\mathbb{Z}[x_0, x_1]$ where a and b are positive integers.

$$G = x_0 + b x_1 + 1, \quad \bar{A} = x_0 + x_1 + a, \quad \bar{B} = x_0 + x_1$$

In this example, $\Gamma = 1$ so $H = G$. Since there are only two variables the Kronecker substitution is $K_r(f) = f(x, y)$ hence $K_r(\bar{A}) = x + y + a$, $K_r(\bar{B}) = x + y$. Notice that $\gcd(K_r(\bar{A}), K_r(\bar{B})) = 1$ in $\mathbb{Z}[x, y]$, but $\gcd(\phi_p(K_r(\bar{A})), \phi_p(K_r(\bar{B}))) = x + y$ for any prime $p|a$. Like Brown's modular GCD algorithm in [5], our GCD algorithm must avoid these primes.

If our GCD algorithm were to choose primes from a pre-computed set of primes $S = \{p_1, p_2, \dots, p_N\}$ then notice that if we replace a in example 5 with $a = \prod_{i=1}^N p_i$ then every prime would be unlucky. To guarantee that our GCD algorithm will succeed on all inputs we will bound the number of primes that cannot be used and pick our prime from a sufficiently large set at random.

Because our algorithm will always choose $r_i > \deg_{x_i} H$, the Kronecker substitution K_r leaves the coefficients of H unchanged. Let p_{min} be the smallest prime in S . From Section 1, $H = \sum_{i=0}^{dG} h_i x_0^i$ with $t = \max(\#h_i)$, we have $\#H \leq (d+1)t$ hence if p is chosen at random from S then

$$\text{Prob}[\text{Supp}(\phi_p(H)) \neq \text{Supp}(H)] \leq \frac{(d+1)t \log_{p_{min}} \|H\|}{N}.$$

Theorem 2 below bounds $\|H\|$ from the inputs A and B .

Definition 4. Let p be a prime and let K_r be a Kronecker substitution. We say p is bad if $\deg_x \phi_p(K_r(A)) < \deg_x K_r(A)$ or $\deg_x \phi_p(K_r(B)) < \deg_x K_r(B)$ and p is unlucky if $\deg_x \gcd(\phi_p(K_r(A)), \phi_p(K_r(B))) > 0$. If p is not bad and not unlucky we say p is good.

Let $R = \text{res}_x(\bar{A}, \bar{B}) \in \mathbb{Z}[x_1, \dots, x_n]$ be the Sylvester resultant of \bar{A} and \bar{B} . Unlucky primes are characterized as follows; if p is not bad then Lemma 4(vii) implies p is unlucky $\iff \phi_p(K_r(R)) = 0$. Unlucky primes are detected using the same approach as described for unlucky evaluations in section 1.3 which requires that we also avoid bad primes. If p is bad or unlucky then p must divide the integer $M = \|K_r(LC(A))\| \cdot \|K_r(LC(B))\| \cdot \|K_r(R)\|$. Let $p_{min} = \min_{i=1}^N p_i$. Thus if p is chosen at random from S then

$$\text{Prob}[p \text{ is bad or unlucky}] \leq \frac{\log_{p_{min}} M}{N}.$$

Proposition 2. Let A be an $m \times m$ matrix with entries $A_{i,j} \in \mathbb{Z}[x_1, x_2, \dots, x_n]$ satisfying the term bound $\#A_{i,j} \leq t$, the degree bound $\deg_{x_k} A_{i,j} \leq d$ and the coefficient bound $\|A_{i,j}\| < h$ (for $1 \leq i, j \leq m$). Note if a term bound for $\#A_{i,j}$ is not known we may use $t = (1+d)^n$. Let $K_r : \mathbb{Z}[x_1, x_2, \dots, x_n] \rightarrow \mathbb{Z}[y]$ be the Kronecker map $K_r(f) = f(y, y^{r_1}, y^{r_1 r_2}, \dots, y^{r_1 r_2 \dots r_{n-1}})$ for $r_k > 0$ and let $B = K_r(A)$ be the $m \times m$

matrix of polynomials in $\mathbb{Z}[y]$ with $B_{i,j} = K_r(A_{i,j})$ for $1 \leq i, j \leq m$. Then

- (i) $\|\det A\| < m^{m/2} t^m h^m$ and
- (ii) $\|\det B\| < m^{m/2} t^m h^m$.

PROOF: To prove (i) let S be the $m \times m$ matrix of integers given by $S_{i,j} = \|A_{i,j}\|_1$. We claim $\|\det A\| \leq H(S)$ where $H(S)$ is Hadamard's bound on $|\det S|$. Then applying Hadamard's bound to S we have

$$H(S) = \prod_{i=1}^m \sqrt{\sum_{j=1}^m S_{i,j}^2} = \prod_{i=1}^m \sqrt{\sum_{j=1}^m \|A_{i,j}\|_1^2} < \prod_{i=1}^m \sqrt{m(th)^2} = m^{m/2} t^m h^m$$

which establishes (i).

To prove our claim let K_s be a Kronecker map with $s_i > md$ and let C be the $m \times m$ matrix with $C_{i,j} = K_s(A_{i,j})$. Notice that $\deg_{x_k} A_{i,j} \leq d$ implies $\deg_{x_k} \det A \leq md$ for $1 \leq k \leq n$. Thus $K_s(\det A)$ is a bijective map on the monomials of $\det A$ thus $K_s(\det A) = \det C$ which implies $\|\det A\| = \|\det C\|$. Now let W be the $m \times m$ matrix with $W_{i,j} = \|C_{i,j}\|_1$ and let $H(W)$ be Hadamard's bound on $|\det W|$. Then $\|\det C\| \leq H(W)$ by Lemma 3 and since K_s is bijective $S = W$ hence $H(S) = H(W)$. Therefore $\|\det A\| = \|\det C\| \leq H(W) = H(S)$ which proves the claim.

To prove (ii), let S and T be the $m \times m$ matrices of integers given by $S_{i,j} = \|A_{i,j}\|_1$ and $T_{i,j} = \|B_{i,j}\|_1$ for $1 \leq i, j \leq m$. From the claim in part (i) if $r_k > md$ we have $\|\det A\| = \|\det B\| \leq H(T) = H(S)$. Now if $r_k \leq md$ for any $1 \leq k \leq n-1$ then $K_r(\det A)$ is not necessarily one-to-one on the monomials in $\det A$. However, for all $r_k > 0$ we still have

$$\|K_r(A_{i,j})\|_1 \leq \|A_{i,j}\|_1 \quad \text{for } 1 \leq i, j \leq m$$

so that $T_{i,j} \leq S_{i,j}$ hence $H(T) \leq H(S)$. We have $\|\det B\| \leq H(T) \leq H(S)$ and (ii) follows. \square

Theorem 2. Let $A, B, G, \bar{A}, \bar{B}, \Delta, H$ be as given at the beginning of this section and let $R = \text{res}_{x_0}(\bar{A}, \bar{B})$. Suppose $A = \sum_{i=0}^{d_A} a_i(x_1, \dots, x_n)x_0^i$ and $B = \sum_{i=0}^{d_B} b_i(x_1, \dots, x_n)x_0^i$ satisfy $\deg A \leq d$, $\deg B \leq d$, $d_A > 0$, $d_B > 0$, $\|a_i\| < h$ and $\|b_i\| < h$. Let $K_r : \mathbb{Z}[x_0, x_1, \dots, x_n] \rightarrow \mathbb{Z}[x, y]$ be the Kronecker map $K_r(f) = f(x, y, y^{r_1}, y^{r_1 r_2}, \dots, y^{r_1 r_2 \dots r_{n-1}})$. If K_r is not bad, that is, $K_r(a_{d_A}) \neq 0$ and $K_r(a_{d_B}) \neq 0$, then

- (i) $\|K_r(LC(A))\| \leq (1+d)^n h$ and $\|K_r(LC(B))\| \leq (1+d)^n h$,
- (ii) $\|K_r(R)\| \leq m^{m/2} (1+d)^{nm} E^m$ and
- (iii) if $r_i > \deg_{x_i} H$ for $1 \leq i \leq n-1$ then $\|H\| \leq (1+d)^n E^2$

where $m = d_A + d_B$ and $E = e^{(n+1)d} h$.

PROOF: Since $LC(A) \in \mathbb{Z}[x_1, \dots, x_n]$ we have $\#LC(A) \leq (1+d)^n$ thus $\|K_r(LC(A))\| \leq (1+d)^n \|LC(A)\| \leq (1+d)^n h$. Using the same argument we have $\|K_r(LC(B))\| \leq (1+d)^n h$ which proves (i).

Let $\bar{A} = \sum_{i=0}^{d_{\bar{A}}} \bar{a}_i x_0^i$ and $\bar{B} = \sum_{i=0}^{d_{\bar{B}}} \bar{b}_i x_0^i$. Because $A = G\bar{A}$ and $B = G\bar{B}$, Lemma 2 implies $\|\bar{A}\| < E$ and $\|\bar{B}\| < E$. Let S be Sylvester's matrix formed from $K_r(\bar{a}_i)$ and $K_r(\bar{b}_i)$. Now $K_r(R) = \det S$ and S has dimension $d_{\bar{A}} + d_{\bar{B}} \leq d_A + d_B = m$. Applying Proposition 2 to S we have

$$\|K_r(R)\| = \|\det S\| \leq t^m E^m m^{m/2}$$

where $t = \max_{i,j} \#S_{i,j}$. Since $\bar{A}|A$ and $\bar{B}|B$ we have $\deg_{x_j} \bar{a}_i(x_1, \dots, x_n) \leq d$ and $\deg_{x_j} \bar{b}_i(x_1, \dots, x_n) \leq d$ thus $\#S_{i,j} \leq (1+d)^n$ and (ii) follows.

For (iii) since $G|A$ and $\Delta|LC(A)$, Lemma 2 implies $\|G\| < E$ and $\|\Delta\| < E$. Thus

$$\|H\| = \|\Delta G\| \leq \#\Delta \cdot \|\Delta\| \cdot \|G\| \leq (1+d)^n E^2. \quad \square$$

We remark that our definition for unlucky primes differs from Brown [5]. Brown's definition depends on the vector degree whereas ours depends only on the degree in x_0 the main variable. The following example illustrates the difference.

Example 6. Consider the following GCD problem and prime p .

$$G = x + y + 1, \quad \bar{A} = (y + p)x^2 + y^2, \quad \bar{B} = yx^3 + y + p.$$

We have $\gcd(\phi_p(\bar{A}), \phi_p(\bar{B})) = \gcd(yx + y^2, yx^2 + y) = y$. By definition 4, p is not unlucky but by Brown's definition, p is unlucky.

Our GCD algorithm in $\mathbb{Z}[x_0, x_1, \dots, x_n]$ only needs monic images in $\mathbb{Z}_p[x_0]$ to recover H whereas Brown needs monic images in $\mathbb{Z}_p[x_0, x_1, \dots, x_n]$ to recover G . A consequence of this is that our bound on the number of unlucky primes is much smaller than Brown's bound (see Theorems 1 and 2 of [5]). This is relevant because we also require p to be smooth.

2.3 The number of unlucky evaluation points

Even if the Kronecker substitution is not unlucky, after applying it to input polynomials A and B , because the degree in y may be very large, the number of bad and unlucky evaluation points may be very large.

Example 7. Consider the following GCD problem

$$G = x_0 + x_1^d + x_2^d + \dots + x_n^d, \quad \bar{A} = x_0 + x_1 + \dots + x_{n-1} + x_n^{d+1}, \quad \bar{B} = x_0 + x_1 + \dots + x_{n-1} + 1.$$

To recover G , if we use $r = [d+1, d+1, \dots, d+1]$ for x_1, x_2, \dots, x_{n-1} we need $p > (d+1)^n$. But $R = \text{res}_{x_0}(\bar{A}, \bar{B}) = 1 - x_n^{d+1}$ and $K_r(R) = 1 - (y^{r_1 r_2 \dots r_{n-1}})^{d+1} = 1 - y^{(d+1)^n}$ which means there could be as many as $(d+1)^n$ unlucky evaluation points. If $p = (d+1)^n + 1$, all evaluation points would be unlucky.

To guarantee that we avoid unlucky evaluation points with high probability we would need to pick $p \gg \deg K_r(R)$ which could be much larger than what is needed to interpolate $K_r(H)$. Algorithm MGCD in Section 4 does this. But this upper bound based on the resultant is a worst case bound. This lead us to investigate what the expected number of unlucky evaluation points is over all possible inputs of a given size. We ran an experiment. We computed all monic quadratic and cubic bivariate polynomials over small finite fields \mathbb{F}_q of size $q = 2, 3, 4, 5, 7, 8, 11$ and counted the number of unlucky evaluation points to find the following result.

Theorem 3. Let \mathbb{F}_q be a finite field with q elements and $f = x^l + \sum_{i=0}^{l-1} (\sum_{j=0}^{d_i} a_{ij} y^j) x^i$ and $g = x^m + \sum_{i=0}^{m-1} (\sum_{j=0}^{e_i} b_{ij} y^j) x^i$ with $l \geq 1$, $m \geq 1$, and $a_{ij}, b_{ij} \in \mathbb{F}_q$. Let $X = |\{\alpha \in \mathbb{F}_q : \gcd(f(x, \alpha), g(x, \alpha)) \neq 1\}|$ be a random variable over all choices $a_{ij}, b_{ij} \in \mathbb{F}_q$. So $0 \leq X \leq q$ and for f and g not coprime in $\mathbb{F}_q[x, y]$ we have $X = q$. If $d_i \geq 0$ and $e_i \geq 0$ then $\mathbb{E}[X] = 1$.

PROOF: Let $C(y) = \sum_{i=0}^d c_i y^i$ with $d \geq 0$ and $c_i \in \mathbb{F}_q$ and fix $\beta \in \mathbb{F}_q$. Consider the evaluation map $C_\beta : \mathbb{F}_q^{d+1} \rightarrow \mathbb{F}_q$ given by $C_\beta(c_0, \dots, c_d) = \sum_{i=0}^d c_i \beta^i$. We claim that C is balanced, that is, C maps q^d inputs to each element of \mathbb{F}_q . It follows that $f(x, \beta)$ is also balanced, that is, over all choices for $a_{i,j}$ each monic polynomial in $\mathbb{F}_q[x]$ of degree n is obtained equally often. Similarly for $g(x, \beta)$.

Recall that two univariate polynomials a, b in $\mathbb{F}_q[x]$ with degree $\deg a > 0$ and $\deg b > 0$ are coprime with probability $1 - 1/q$ (see Ch 11 of Mullen and Panario [36]). This is also true under the restriction that they are monic. Therefore $f(x, \beta)$ and $g(x, \beta)$ are coprime with probability $1 - 1/q$. Since we have q choices for β we obtain

$$E[X] = \sum_{\beta \in \mathbb{F}_q} \text{Prob}[\text{gcd}(A(x, \beta), B(x, \beta)) \neq 1] = q(1 - (1 - \frac{1}{q})) = 1.$$

Proof of claim. Since $B = \{1, y - \beta, (y - \beta)^2, \dots, (y - \beta)^d\}$ is a basis for polynomials of degree d we can write each $C(y) = \sum_{i=0}^d c_i y^i$ as $C(y) = u_0 + \sum_{i=1}^d u_i (y - \beta)^i$ for a unique choice of $u_0, u_1, \dots, u_d \in \mathbb{F}_q$. Since $C(\beta) = u_0$ it follows that all q^d choices for u_1, \dots, u_d result in $C(\beta) = u_0$ hence C is balanced. \square

That $E[X] = 1$ was a surprise to us. We thought $E[X]$ would have a logarithmic dependence on $\deg f$ and $\deg g$. This explains why unlucky evaluation points occur so rarely in practice and why codes work with much smaller primes than worst case bounds would require. In light of Theorem 3, our heuristic algorithm in Section 4 will first pick $p > \deg_y(K_r(H))$ and, should the algorithm encounter unlucky evaluations, restart the algorithm with a larger prime.

3 Simplified Algorithm

We now present our GCD algorithm. It consists of two parts: the main routine MGCD and the subroutine PGCD. PGCD computes the GCD modulo a prime and MGCD calls PGCD several times to obtain enough images to reconstruct the coefficients of the target polynomial H using Chinese Remaindering. In this section, we assume that we are given a term bound τ on the number of terms in the coefficients of target polynomial H , that is $\tau \geq \#h_i(x_1, x_2, \dots, x_n)$. We will also apply Lemma 7 to choose a Kronecker substitution that is a priori not bad and not unlucky. These assumptions will enable us to choose the prime p so that PGCD computes G modulo p with high probability. We will relax these assumptions in the next section. The algorithm will need to treat bad and unlucky primes and bad and unlucky evaluation points.

3.1 Bad and unlucky evaluations

In this section, the Kronecker substitution K_r is assumed to be good. We also assume that the prime p is good.

Proposition 3. *Let $d = \max\{\max\{\deg_{x_i} A, \deg_{x_i} B\}_{0 \leq i \leq n}\}$ and let $r_i = 2d^2 + 1$ for $1 \leq i \leq n$. Note $2d^2 + 1 \geq (\deg_{x_i} A \deg_{x_0} B + \deg_{x_i} B \deg_{x_0} A) + 1$. Then*

- (1) $\deg_y K_r(A) < (2d^2 + 1)^n$ and $\deg_y K_r(B) < (2d^2 + 1)^n$,
- (2) $\deg_y LC(K_r(A))(y) < (2d^2 + 1)^n$ and $\deg_y LC(K_r(B))(y) < (2d^2 + 1)^n$,
- (3) $\deg_y K_r(H) < (2d^2 + 1)^n$, and
- (4) $\deg_y K_r(R) < 2d(2d^2 + 1)^n$, where $K_r(R) = \text{res}_x(K_r(\bar{A}), K_r(\bar{B}))$.

Proof. For (1), after the Kronecker substitution, the exponent of $y \leq e_1 + e_2(2d^2 + 1) + \dots + e_n(2d^2 + 1)^{n-1}$, where e_i is the exponent of x_i and $e_i \leq d$ for all i . So $\deg_y K_r(A)$ and $\deg_y K_r(B)$ are bounded by

$$\begin{aligned} d + d(2d^2 + 1) + \dots + d(2d^2 + 1)^{n-1} &= d(1 + (2d^2 + 1) + \dots + (2d^2 + 1)^{n-1}) \\ &= d\left(1 + \frac{(2d^2 + 1)^n - (2d^2 + 1)}{(2d^2 + 1) - 1}\right) \\ &= \frac{2d^3}{2d^2} + \frac{d(2d^2 + 1)^n - d(2d^2 + 1)}{2d^2} \\ &= \frac{d(2d^2 + 1)^n - d}{2d^2} \\ &< (2d^2 + 1)^n. \end{aligned}$$

Property (2) follows from (1). For (3), recall that $\deg_y K_r(H) = \deg_y K_r(\Delta G)$. Since $\Delta = \gcd(\text{LC}(\bar{A}), \text{LC}(\bar{B}))$, we have

$$\begin{aligned} \deg_y K_r(\Delta G) &= \deg_y K_r(\Delta) + \deg_y K_r(G) \\ &\leq \min(\deg_y K_r(\text{LC}(\bar{A})), \deg_y K_r(\text{LC}(\bar{B}))) + \deg_r K_r(G) \\ &\leq \min(\deg_y K_r(\bar{A}), \deg_y K_r(\bar{B})) + \deg_r K_r(G) \\ &= \min(\deg_y K_r(A), \deg_y K_r(B)) < (2d^2 + 1)^n. \end{aligned}$$

For (4),

$$\deg_y K_r(R) \leq \deg_y K_r(\bar{A}) \deg_x K_r(\bar{B}) + \deg_y K_r(\bar{B}) \deg_x K_r(\bar{A}),$$

where $\deg_x K_r(\bar{A}) = \deg_{x_0} \bar{A} \leq \deg_{x_0} A \leq d$, $\deg_x K_r(\bar{B}) = \deg_{x_0} \bar{B} \leq \deg_{x_0} B \leq d$, and $\deg_y K_r(\bar{A}) \leq \deg_y K_r(A)$ and $\deg_y K_r(\bar{B}) \leq \deg_y K_r(B)$. So we have

$$\deg_y K_r(R) < d(2d^2 + 1)^n + d(2d^2 + 1)^n = 2d(2d^2 + 1)^n.$$

□

By proposition 3(1), a prime $p > (2d^2 + 1)^n$ is sufficient to recover the exponents for the Kronecker substitution. With the assumption that p is not bad and not unlucky, we have the following lemma.

Lemma 8. *Let p be a prime. If α is chosen at random from $[0, p - 1]$, then*

$$(i) \text{ Prob}[\alpha \text{ is bad}] < \frac{2(2d^2 + 1)^n}{p} \text{ and}$$

$$(ii) \text{ Prob}[\alpha \text{ is unlucky or } \alpha \text{ is bad}] < \frac{(2d + 2)(2d^2 + 1)^n}{p}.$$

Proof. $\text{Prob}[\alpha \text{ is bad}] = \text{Prob}[\text{LC}(K_r(A)(\alpha)) \text{LC}(K_r(B)(\alpha)) = 0]$
 $\leq \deg \text{LC}(K_r(AB))(y)/p < 2(2d^2 + 1)^n/p$. For (ii) from the proof of Lemma 6 we have this probability $\leq \deg K_r(\text{LC}(AB))/p + \deg K_r(R)/p$ where $R = \text{res}_{x_0}(\bar{A}, \bar{B})$. Applying proposition 3(1) and (4) we have the probability $< 2(2d^2 + 1)^n/p + 2d(2d^2 + 1)^n/p$ and the result follows. □

The probability that our algorithm does not encounter a bad or unlucky evaluation can be estimated as follows. Let U denote the bound of the number of bad and unlucky evaluation points and $\tau \geq \max_i \{\#h_i\}$. We need 2τ good consecutive evaluation points (a segment of length 2τ in the sequence $(1, \dots, p-1)$) to compute the feedback polynomial for h_i . We remark that our analysis here is similar to that of Comer, Kaltofen and Pernet [9] who considered numerical sparse interpolation in the presence of outliers.

Suppose α^k is a bad or unlucky evaluation point where $s \leq k < s + 2\tau - 1$ for any positive integer $s \in (0, p-1]$. Then every segment of length 2τ starting at α^i where $k - 2\tau + 1 \leq i \leq k$ includes the point α^k . Hence our algorithm fails to determine the correct feedback polynomial. The union of all segments including α^k has length $4\tau - 1$. We can not use every segment of length 2τ from $k - 2\tau + 1$ to $k + 2\tau - 1$ to construct the correct feedback polynomial. The worst case occurs when all bad and unlucky evaluation points, their corresponding segments of length $4\tau - 1$ do not overlap. Since there are at most U of them, we can not determine the correct feedback polynomials for at most $U(4\tau - 1)$ points. Note, this does not mean that all those points are bad or unlucky, there is only one bad or unlucky point in each segment of length 2τ . U is bounded by $2(2d^2 + 1)^n + 2d(2d^2 + 1)^n = (2d + 2)(2d^2 + 1)^n$.

Lemma 9. *Suppose p is good. Then*

$$\begin{aligned} & \text{Prob}[2\tau \text{ evaluation points fail to determine the feedback polynomial}] \\ & \leq \frac{4\tau U - U}{p-1} < \frac{4\tau U}{p-1} = \frac{4\tau(2d+2)(2d^2+1)^n}{p-1}. \end{aligned}$$

So if we choose a prime $p > 4X\tau(2d+2)(2d^2+1)^n$ for some positive number X , then the probability that PGCD fails is at most $\frac{1}{X}$.

We note that the choice of p in previous lemma implies $p > (2d^2+1)^n \geq \deg_y(K_r(H))$. So we can recover the exponents of y in H .

3.2 Bad and unlucky primes

Our goal here is to construct a set S of smooth primes, with $|S|$ large enough so if we choose a prime $p \in S$ at random, the probability that p is good is at least $\frac{1}{2}$. Recall that a prime p is said to be y -smooth if $q|p-1$ implies $q \leq y$. The choice of y affects the efficiency of discrete logarithm computation in \mathbb{Z}_p .

A bad prime must divide $\|LC(K_r(A))\|$ or $\|LC(K_r(B))\|$ and an unlucky prime must divide $\|Kr(R)\|$. Recall that in Section 2.2,

$$M = \|LC(K_r(A))\| \|LC(K_r(B))\| \|Kr(R)\|.$$

We want to construct a set $S = \{p_1, p_2, \dots, p_N\}$ of N smooth primes with each $p_i > 4\tau(2d+4)(2d^2+1)^n X$. If $p > 4\tau(2d+4)(2d^2+1)^n X$, then the probability that our algorithm fails to determine the feedback polynomial is $< \frac{1}{X}$. The size N of S can be estimated as follows. If

$$N = Y \lceil \log_{4X\tau(2d+4)(2d^2+1)^n} M \rceil > Y \log_{pmin} M,$$

where a bound for M is given by Theorem 2 (ii), $pmin = \min_{p_i \in S} p_i$ and $Y > 0$, Then

$$\text{Prob}[p \text{ is bad or unlucky}] \leq \frac{\log_{pmin} M}{N} < \frac{1}{Y}.$$

We construct the set S which consists of N y -smooth primes so that $\min_{p_i \in S} p_i > 4\tau X(2d+4)(2d^2+1)^n$ which is the constraint for the bad or unlucky evaluation case. We conclude with the following result.

Theorem 4. *Let S be constructed as just described. Let p be chosen at random from S , s be chosen at random from $0 < s \leq p-1$ and α_p be a random generator of \mathbb{Z}_p^* . Let $E = \{\alpha_p^{s+j} : 0 \leq j < 2\tau\}$ be 2τ consecutive evaluation points. For any $X > 0$ and $Y > 0$, we have*

$$\text{Prob}[p \text{ is good and } E \text{ are all good}] > (1 - \frac{1}{X})(1 - \frac{1}{Y}).$$

3.3 The Simplified GCD Algorithm

Let $S = \{p_1, p_2, \dots, p_N\}$ is the set of N primes constructed in the previous section. We've split our GCD algorithm into two subroutines, subroutine MGCD and PGCD. The main routine MGCD chooses a Kronecker substitution K_r and then chooses a prime p from S at random and calls PGCD to compute $K_r(H) \bmod p$.

Algorithm MGCD is a Las Vegas algorithm. The choice of S means that algorithm PGCD will compute $K_r(H) \bmod p$ with probability at least $(1 - \frac{1}{X})(1 - \frac{1}{Y})$. By taking $X = 4$ and $Y = 4$ this probability is at least $\frac{1}{2}$. The design of MGCD means that even with probability $\frac{1}{2}$, the expected number of calls to algorithm PGCD is linear in the minimum number of primes needed to recover H using Chinese remaindering, that is, we do not need to make the probability that algorithm PGCD computes $H \bmod p$ high for algorithm MGCD to be efficient.

Algorithm MGCD(A, B, τ)

Inputs $A, B \in \mathbb{Z}[x_0, x_1, \dots, x_n]$ and a term bound τ satisfying $n > 0$, A and B are primitive in x_0 , $\deg_{x_0} A > 0$, $\deg_{x_0} B > 0$ and $\tau \geq \max \#h_i$.

Output $G = \gcd(A, B)$.

- 1 Compute $\Gamma = \gcd(LC(A), LC(B))$ in $\mathbb{Z}[x_1, \dots, x_n]$.
- 2 Set $r_i = 1 + (\deg_{x_i} A \deg_{x_0} B + \deg_{x_i} B \deg_{x_0} A)$ for $1 \leq i < n$.
- 3 Let $Y = (y, y^{r_1}, y^{r_1 r_2}, \dots, y^{r_1 r_2 \dots r_{n-1}})$.
Set $K_r A = A(x, Y)$, $K_r B = B(x, Y)$ and $K_r \Gamma = \Gamma(Y)$.
- 4 Construct the set S of smooth primes according to Theorem 4 with $X = 4$ and $Y = 4$.
- 5 Set $\hat{H} = 0$, $M = 1$, $d_0 = \min(\deg_{x_0} A, \deg_{x_0} B)$.

LOOP: // Invariant: $d_0 \geq \deg_{x_0} H = \deg_{x_0} G$.

- 6 Call PGCD($K_r A, K_r B, K_r \Gamma, S, \tau, M$).
If PGCD outputs FAIL then **goto** LOOP.
Let p and $\hat{H}p = \sum_{i=0}^{dx} \hat{h}_i(y)x^i$ be the output of PGCD.
- 7 If $dx > d_0$ then either p is unlucky or all evaluation points were unlucky so **goto** LOOP.

8 If $dx < d_0$ then either this is the first image or all previous images in \widehat{H} were unlucky so set $d_0 = dx$, $\widehat{H} = Hp$, $M = p$ and **goto** LOOP.

Chinese-Remaindering

9 Set $Hold = \widehat{H}$. Solve $\{\widehat{H} \equiv Hold \pmod{M} \text{ and } \widehat{H} \equiv \widehat{H}p \pmod{p}\}$ for \widehat{H} . Set $M = M \times p$. If $\widehat{H} \neq Hold$ then **goto** LOOP.

Termination.

10 Set $\widetilde{H} = K_r^{-1}\widehat{H}(x, y)$ and let $\widetilde{H} = \sum_{i=0}^{d_0} \widetilde{c}_i x_0^i$ where $\widetilde{c}_i \in \mathbb{Z}[x_1, x_2, \dots, x_n]$.

11 Set $\widehat{G} = \widetilde{H} / \gcd(\widetilde{c}_0, \widetilde{c}_1, \dots, \widetilde{c}_{d_0})$ (\widehat{G} is the primitive part of \widetilde{H}).

12 If $\widehat{G}|A$ and $\widehat{G}|B$ then **output** \widehat{G} .

13 **goto** LOOP.

Algorithm PGCD($K_r A, K_r B, K_r \Gamma, S, \tau, M$)

Inputs $K_r A, K_r B \in \mathbb{Z}[x, y]$, $K_r \Gamma \in \mathbb{Z}[y]$, S a set of smooth primes, a term bound $\tau \geq \max \#h_i$ and M a positive integer.

Output With probability $\geq \frac{1}{2}$ a prime p and polynomial $Hp \in \mathbb{Z}_p[x, y]$ satisfying $Hp = K_r(H) \pmod{p}$ and p does not divide M .

1 Pick a prime p at random from S that is not bad and does not divide M .

2 Pick a random shift s such that $0 < s < p$ and any generator α for \mathbb{Z}_p^* .

Compute-and-scale-images:

3 For j from 0 to $2\tau - 1$ do

4 Compute $a_j = K_r A(x, \alpha^{s+j}) \pmod{p}$ and $b_j = K_r B(x, \alpha^{s+j}) \pmod{p}$.

5 If $\deg_x a_j < \deg_x K_r A$ or $\deg_x b_j < \deg_x K_r B$ then **output** FAIL (α^{s+j} is a bad evaluation point.)

6 Compute $g_j = \gcd(a_j, b_j) \in \mathbb{Z}_p[x]$ using the Euclidean algorithm and set $g_j = K_r \Gamma(\alpha^{s+j}) \times g_j \pmod{p}$.

End for loop.

7 Set $d_0 = \deg g_0(x)$. If $\deg g_j(x) \neq d_0$ for any $1 \leq j \leq 2\tau - 1$ **output** FAIL (unlucky evaluations).

Interpolate-coefficients:

8 For $i = 0$ to d_0 do

9 Run the Berlekamp-Massey algorithm on the coefficients of x^i in the images $g_0, g_1, \dots, g_{2\tau-1}$ to obtain $\lambda_i(z)$ and set $\tau_i = \deg \lambda_i(z)$.

10 Compute the roots m_j of each $\lambda_i(z)$ in \mathbb{Z}_p . If the number of distinct roots of $\lambda_i(z)$ is not equal τ_i then **output** FAIL (the feedback polynomial is wrong due to undetected unlucky evaluations.)

- 11 Set $e_k = \log_\alpha m_k$ for $1 \leq k \leq \tau_i$ and let $\sigma_i = \{y^{e_1}, y^{e_2}, \dots, y^{e_{\tau_i}}\}$.
- 12 Solve the τ_i by τ_i shifted transposed Vandermonde system

$$\left\{ \sum_{k=1}^{\tau_i} (\alpha^{s+j})^{e_k} u_k = \text{coefficient of } x^i \text{ in } g_j(x) \text{ for } 0 \leq j < \tau_i \right\}$$

modulo p for u and set $h_i(y) = \sum_{k=1}^{\tau_i} u_k y^{e_k}$. Note: $(\alpha^{s+j})^{e_k} = m_k^{s+j}$

End for loop.

- 13 Set $Hp := \sum_{i=0}^{d_0} h_i(y)x^i$ and **output** (p, Hp) .

We remark that we do not check for termination after each prime because computing the primitive part of \hat{H} or doing the trial divisions $\hat{G}|A$ and $\hat{G}|B$ in Step 12 could be more expensive than algorithm PGCD. Instead algorithm MGCD waits until the Chinese remaindering stabilizes in Step 9 before proceeding to test for termination.

Theorem 5. *Let $N = \log_{p_{\min}} \|2H\|$. So N primes in S are sufficient to recover the integer coefficients of H using Chinese remaindering. Let X be the number of calls that Algorithm MGCD makes to Algorithm PGCD. Then $E[X] \leq 2(N + 1)$.*

Proof. Because the Kronecker substitution K_r is not bad, and the primes p used in PGCD are not bad and the evaluations points $\{\alpha^{s+j} : 0 \leq j \leq 2\tau - 1\}$ used in PGCD are not bad, in Step 6 of Algorithm PGCD, $\deg g_j(x) \geq \deg_{x_0} G$ by Lemma 5. Therefore $d_0 \geq \deg_{x_0} H = \deg_{x_0} G$ throughout Algorithm MGCD and $\deg_x \hat{H} = \deg_{x_0} \hat{G} \geq \deg_{x_0} G$. Since A and B are primitive in x_0 , if $\hat{G}|A$ and $\hat{G}|B$ then it follows that $\hat{G} = G$, so if algorithm MGCD terminates, it outputs G .

To prove termination observe that Algorithm MGCD proceeds in two phases. In the first phase MGCD loops while $d_0 > \deg_{x_0} H$. In this phase no useful work is accomplished. Observe that the loops in PGCD are of fixed length 2τ and $d_0 + 1$ so PGCD always terminates and algorithm MGCD tries another prime. Because at least $3/4$ of the primes in S are good, and, for each prime, at least $3/4$ of the possible evaluation point sequences are good, eventually algorithm PGCD will choose a good prime and a good evaluation point sequence after which $d_0 = \deg_{x_0} H$.

In the second phase MGCD loops using images Hp with $\deg_x Hp = d_0$ to construct \hat{H} . Because the images $g_j(x)$ used satisfy $\deg_x g_j(x) = d_0 = \deg_{x_0} H$ and we scale them with $\Gamma(\alpha^{s+j})$, PGCD interpolates $Hp = H \pmod p$ thus we have modular images of H . Eventually $\hat{H} = H$ and the algorithm terminates.

Because the probability that the prime chosen from S is good is at least $3/4$ and the evaluations points α^{s+j} are all good is at least $3/4$, the probability that PGCD outputs a good image of H is at least $1/2$. Since we need N images of H to recover H and one more to stabilize (see Step 9), $E[X] \leq 2(N + 1)$ as claimed. \square

4 Faster Algorithm

In this section we consider the practical design of algorithms MGCD and PGCD. We make three improvements. Unfortunately, each improvement leads to a major complication.

4.1 Term Bounds

Recall that $H = \Delta G = \sum_{i=0}^{dG} h_i(x_1, \dots, x_n) x_0^i$. Algorithms MGCD and PGCD assume a term bound τ on $\#h_i(y)$. In practice, good term bounds are usually not available. For the GCD problem, one cannot even assume that $\#G \leq \min(\#A, \#B)$ so we must modify the algorithm to compute $t_i = \#h_i(y)$.

We will follow Kaltofen et. al. [26] which requires $2t_i + O(1)$ evaluation points to determine t_i with high probability. That is, we will loop calling the Berlekamp-Massey algorithm after 2, 4, 6, 8, \dots , evaluation points and wait until we get two consecutive zero discrepancies, equivalently, we wait until the output $\lambda_i(z)$ does not change. This means $\lambda_i(z)$ is correct with high probability when p is sufficiently large. We give details in Section 4.5. This loop will only terminate, however, if the sequence of points is generated by a polynomial and therein lies a problem.

Example 8. Consider the following GCD problem in $\mathbb{Z}[x, y]$. Let p be a prime and let

$$G = 1, \quad \bar{A} = (yx + 1)((y + 1)x + 2), \quad \bar{B} = (yx + 2)(y + p + 1)x + 2).$$

Observe that $LC(A) = y(y + 1)$, $LC(B) = y(y + p + 1)$, $\Gamma = y$ and $\gcd(A \bmod p, B \bmod p) = (y + 1)x + 2$ so p is unlucky.

Suppose we run algorithm PGCD with inputs $A = G\bar{A}$, $B = G\bar{B}$ and $\Gamma = y$ and suppose PGCD selects the prime p . Let $F(x, y) = x + \frac{2}{y+1}$. Algorithm PGCD will compute monic images $g_j(x) = F(x, \alpha^{s+j}) \bmod p$ which after scaling by $\Gamma = \alpha^{s+j}$ are images of $yx + \frac{2y}{y+1}$ which is not a polynomial in y . So the Berlekamp-Massey algorithm will likely not stabilize and algorithm PGCD will loop trying to compute $\lambda_0(z)$. The problem is that scaling by $\Gamma = y$ does not result in a polynomial. We note that the same problem may be caused by an unlucky Kronecker substitution.

Our solution is to scale with either $\Gamma = LC(A)$ or $\Gamma = LC(B)$, whichever has fewer terms. Then, assuming p is not bad, $LC(\gcd(A \bmod p, B \bmod p))$ must divide both $LC(A) \bmod p$ and $LC(B) \bmod p$ thus scaling $g_j(x)$ by $LC(A)(\alpha^{s+j}) \bmod p$ or $LC(B)(\alpha^{s+j}) \bmod p$ will always give an image of a polynomial. The downside of this solution is that it may increase $t_i = \#h_i(y)$.

Another difficulty caused by $\lambda_i(z)$ stabilizing too early is that the support σ_i of $K_r(h_i)$ computed in Step 11 of PGCD may be wrong. We consider an example.

Example 9. Consider the following GCD problem in $\mathbb{Z}[x, y]$. Let p and q be prime and

$$G = x + py + qy^2 + py^4, \quad \bar{A} = 1, \quad \bar{B} = 1.$$

Suppose MGCD chooses p first and suppose PGCD returns $x + qy^2 \bmod p$ so that $\sigma_0 = \{y^2\}$. Suppose MGCD chooses q next and suppose $\lambda_0(z)$ stabilizes too early and $\sigma_0 = \{y^3\}$ which is wrong. This could also be due to a missing term, for example, if $G = x + pqy + qy^2 + py^3$. If we combine these two images modulo p and q using Chinese remaindering to obtain \hat{H} of the form $x + \cdot y^2 + \cdot y^3$ we have a bad image in \hat{H} and we need somehow to detect it. Once detected, we do not want to restart the entire algorithm because we might be throwing away a lot of good images in \hat{H} . Our solution (see Steps 7–10 of algorithm MGCD1) is probabilistic. Suppose M is the product of primes used so far. Let β be chosen at random from $[0, p)$. For each prime $p|M$ we test if

$$h(x, \beta) | K_r(A)(x, \beta) \quad \text{and} \quad h(x, \beta) | K_r(B)(x, \beta) \quad \text{where} \quad h = \hat{H} \bmod p$$

If the test fails it means $h(x, y)$ is a bad image and we remove it from \hat{H} . By repeating this test we eventually detect all bad images in \hat{H} .

4.2 Using smaller primes

Another consideration is the size of the primes that we use. We have implemented our GCD algorithm for 63 bit primes and 127 bit primes. By choosing a Kronecker substitution that is a priori good, and requiring that the 2τ evaluation points are good, the primes in S must be greater than $4\tau(2d+2)(2d^2+1)^n$ where d bounds the degree of A and B in all variables. If instead we choose $r_i > \deg_{x_i} H$ then we will still be able to recover H from $K_r(H)$ but K_r may be unlucky.

Since $\deg_{x_i} H \leq \min(\deg_{x_i} A, \deg_{x_i} B) \leq d$, using $r_i = d+1$ we replace the factor $(2d^2+1)^n$ with $(d+1)^n$. We will detect if K_r is unlucky when $\deg g_j(x) > d_0$ by computing $d_0 = \text{DegreeBound}(A, B, 0)$ periodically (see Step 6 of MGCD1) so that eventually we obtain $d_0 = \deg_{x_0} G$ and can detect unlucky K_r . Once detected we will increase r_i by 1 to try a larger Kronecker substitution.

Recall that p is an unlucky prime if $p|R$ where $R = \text{res}_{x_0}(\bar{A}, \bar{B})$. Because the inputs A and B are primitive in x_0 it follows that the integer coefficients of \bar{A} and \bar{B} are relatively prime. Therefore, the integer coefficients of R are also likely to have a very small common factor like 2. Thus the expected number of unlucky primes is very close to 0. In Theorem 3 we showed that the expected number of unlucky evaluations is 1 hence instead of using $p > 4\tau(2d+2)(d+1)^n$ we first try a prime $p > 4(d+1)^n$. Should we encounter bad or unlucky evaluation points we will increase the length of p until we don't. This reduces the length of the primes for most inputs by at least a factor of 2.

Example 10. For our benchmark problem where $n = 8$, $d = 20$ and $\tau = 1000$ we have $\log_2[4\tau(2d+2)(2d^2+1)^n] = 94.5$ bits which precludes our using 63 bit primes. On the other hand $\log_2[4(d+1)^n] = 37.1$ bits, meaning a 63 bit prime is more than sufficient.

4.3 Using fewer evaluation points

Let $K_r(h_i) = \sum_{j=1}^{t_i} c_{ij}y^{e_{ij}}$ for some coefficients $c_{ij} \in \mathbb{Z}$ and exponents e_{ij} so that $\text{Supp}(K_r(h_i)) = \{y^{e_{ij}} : 1 \leq j \leq t_i\}$. Because of the size of the primes chosen by algorithm MGCD, it is likely that the first good image Hp computed by PGCD has the entire support of $K_r(H)$, that is, $\text{Supp}(\hat{h}_i) = \text{Supp}(K_r(h_i))$. Assuming this to be so, we can compute the next image of $K_r(H)$ modulo p using only t evaluations instead of $2t + O(1)$ as follows. We choose a prime p and compute $g_j(x)$ for $0 \leq j < t$ as before in PGCD. Assuming these t images are all good, one may solve the t_i by t_i shifted transposed Vandermonde systems

$$\left\{ \sum_{k=1}^{t_i} (\alpha^{s+j})^{e_{ik}} u_{ik} = \text{coefficient of } x^i \text{ in } g_j(x) \text{ for } 0 \leq j \leq \tau_i - 1 \right\}$$

for the unknown coefficients u_{ij} obtaining $Hp = \sum_{i=0}^{d_0} \sum_{j=0}^{t_i} u_{ij}y^{e_{ij}}$.

It is possible that the prime p used in PGCD may divide a coefficient c_{ij} in $K_r(H)$ in which case we will need to call PGCD again to compute more of the support of $K_r(H)$.

Definition 5. Let $f = \sum_{i=0}^d c_i y^{e_i}$ be a polynomial in $\mathbb{Z}[y]$. We say a prime p causes missing terms in f if p divides any coefficient c_i in f .

Our strategy to detect when $\text{Supp}(\hat{h}_i) \not\subset \text{Supp}(K_r(h_i))$ is probabilistic. We compute one more image $j = \tau_i$ and check that the solutions of the Vandermonde systems are consistent with this image. Thus we require $t+1$ evaluations instead of $2t + O(1)$. Once missing terms are detected, we call PGCD again to determine $\text{Supp}(K_r(h_i))$.

4.4 Algorithm MGCD1

We now present our algorithm as algorithm MGCD1 which calls subroutines PGCD1 and SGCD1. Like MGCD, MGCD1 loops calling PGCD1 to determine the $Hp = K_r(H) \bmod p$. Instead of calling PGCD1 for each prime, MGCD1 after PGCD1 returns an image Hp , MGCD1 assumes the support of $K_r(H)$ is now known and uses SGCD1 for the remaining images.

Algorithm MGCD1(A, B)

Inputs $A, B \in \mathbb{Z}[x_0, x_1, \dots, x_n]$ satisfying $n > 0$, A and B are primitive in x_0 , and $\deg_{x_0} A > 0$, $\deg_{x_0} B > 0$.

Output $G = \gcd(A, B)$.

- 1 If $\#LC(A) < \#LC(B)$ set $\Gamma = LC(B)$ else set $\Gamma = LC(A)$.
- 2 Call Algorithm DegreeBound(A, B, i) to get $d_i \geq \deg_{x_i} G$ for $0 \leq i \leq n$.
If $d_0 = 0$ **return 1**.
- 3 Set $r_i = \min(\deg_{x_i} A, \deg_{x_i} B, d_i + \deg_{x_i}(\Gamma))$ for $1 \leq i \leq n$.
Set $\delta = 1$.

Kronecker-Prime

- 4 Set $r_i = r_i + 1$ for $1 \leq i < n$. Let $Y = (y, y^{r_1}, y^{r_1 r_2}, \dots, y^{r_1 r_2 \dots r_{n-1}})$.
Set $K_r A = A(x, Y)$, $K_r B = B(x, Y)$ and $K_r \Gamma = \Gamma(Y)$.
If K_r is bad **goto** Kronecker-Prime otherwise set $\delta = \delta + 1$.

RESTART

- 5 Set $\widehat{H} = 0$, $M = 1$ and MissingTerms = true.
Set $\sigma_i = \phi$ and $\tau_i = 0$ for $0 \leq i \leq d_0$.

LOOP: // Invariant: $d_0 \geq \deg_{x_0} H$.

- 6 Compute $dx = \text{DegreeBound}(A, B, 0)$.
If $dx < d_0$ set $d_0 = dx$ and **goto** RESTART.
- 7 For each prime $p|M$ do // check current images
 - 8 Set $a = K_r A \bmod p$, $b = K_r B \bmod p$ and $h = \widehat{H} \bmod p$.
 - 9 Pick β from $[0, p - 1]$ at random.
 - 10 If $K_r \Gamma(\beta) \neq 0$ and either $h(x, \beta)$ does not divide $a(x, \beta)$ or does not divide $b(x, \beta)$ then h is wrong so set $M = M/p$ and $\widehat{H} = \widehat{H} \bmod M$ to remove this image.

End for loop.

If MissingTerms then // for first iteration

- 11 Pick a new smooth prime $p > 2^\delta \prod_{i=1}^n r_i$ that is not bad.
- 12 Call PGCD1($K_r A, K_r B, K_r \Gamma, d_0, \tau, r, p$).

13 If PGCD1 returned UNLUCKY($dmin$) set $d_0 = dmin$ and **goto** RESTART.
 If PGCD1 returned FAIL **goto** Kronecker-Prime.

14 Let $\widehat{H}p = \sum_{i=0}^{d_0} \widehat{h}_i(y)x^i$ be the output of PGCD1.
 Set MissingTerms = false, $\sigma_i := \sigma_i \cup \text{Supp}(\widehat{h}_i)$ and $\tau_i = |\sigma_i|$ for $0 \leq i \leq d_0$.

else

15 Pick a new prime $p > 2^\delta \prod_{i=1}^n r_i$ that is not bad.

16 Call SGCD1($K_rA, K_rB, K_r\Gamma, d_0, \sigma, \tau, p$).

17 If SGCD1 returned UNLUCKY($dmin$) set $d_0 = dmin$ and **goto** RESTART.
 If SGCD1 returned FAIL **goto** Kronecker-Prime.
 If SGCD1 returned MISSINGTERMS set $\delta = \delta + 1$, Missingterms = true and **goto** LOOP.

18 Let $\widehat{H}p = \sum_{i=0}^{d_0} \widehat{h}_i(y)x^i$ be the output of SGCD1.

End If

Chinese-Remaindering

19 Set $Hold = \widehat{H}$. Solve $\{\widehat{H} \equiv Hold \pmod{M} \text{ and } \widehat{H} \equiv \widehat{H}p \pmod{p}\}$ for \widehat{H} . Set $M = M \times p$. If $\widehat{H} \neq Hold$ then **goto** LOOP.

Termination.

20 Set $\widetilde{H} = K_r^{-1}\widehat{H}(x, y)$. Let $\widetilde{H} = \sum_{i=0}^{d_0} \widetilde{c}_i x_0^i$ where $\widetilde{c}_i \in \mathbb{Z}[x_1, \dots, x_n]$.

21 Set $\widehat{G} = \widetilde{H} / \text{gcd}(\widetilde{c}_0, \widetilde{c}_1, \dots, \widetilde{c}_{d_0})$ (\widehat{G} is the primitive part of \widetilde{H}).

22 If $\deg \widehat{G} \leq \deg A$ and $\deg \widehat{G} \leq \deg B$ and $\widehat{G}|A$ and $\widehat{G}|B$ then **return** \widehat{G} .

23 **goto** LOOP.

Algorithm PGCD1($K_rA, K_rB, K_r\Gamma, d_0, \tau, r, p$)

Inputs $K_rA, K_rB \in \mathbb{Z}[x, y]$ and $K_r\Gamma \in \mathbb{Z}[y]$, $d_0 \geq \deg_{x_0} G$ where $G = \text{gcd}(A, B)$, term bound estimates $\tau \in \mathbb{Z}^{d_0+1}$, $r \in \mathbb{Z}^n$, and a smooth prime p .

Output $Hp \in \mathbb{Z}_p[x, y]$ satisfying $Hp = K_r(H) \pmod{p}$ or FAIL or UNLUCKY($dmin$).

1 Pick a random shift $s \in \mathbb{Z}_p^*$ and any generator α for \mathbb{Z}_p^* .

2 Set $T = 0$.

LOOP

3 For j from $2T$ to $2T + 1$ do

4 Compute $a_j = K_rA(x, \alpha^{s+j}) \pmod{p}$ and $b_j = K_rB(x, \alpha^{s+j}) \pmod{p}$.

5 If $\deg_x a_j < \deg_x K_rA$ or $\deg_x b_j < \deg_x K_rB$ then return FAIL (α^{s+j} is a bad evaluation point.)

6 Compute $g_j = \gcd(a_i, b_i) \in \mathbb{Z}_p[x]$ using the Euclidean algorithm.
 Make g_j monic and set $g_j = K_r \Gamma(\alpha^{s+j}) \times g_j \pmod p$.

End for loop.

7 Set $dmin = \min \deg g_j(x)$ and $dmax = \max \deg g_j$ for $2T \leq j \leq 2T + 1$.
 If $dmin < d_0$ **output** UNLUCKY(dmin).
 If $dmax > d_0$ **output** FAIL.

8 Set $T = T + 1$.
 If $T < \#K_r \Gamma$ or $T < \max_{i=0}^{d_0} \tau_i$ **goto** LOOP.

9 For i from 0 to d_0 do

10 Run the Berlekamp-Massey algorithm on the coefficients of x^i in the images $g_0, g_1, \dots, g_{2T-1}$ to obtain $\lambda_i(z)$ and set $\tau_i = \deg \lambda_i(z)$. If either of the last two discrepancies were non-zero **goto** LOOP.

End for loop.

11 For i from 0 to d_0 do

12 Compute the roots m_k of $\lambda_i(z)$. If $\lambda_i(0) = 0$ or the number of distinct roots of $\lambda_i(z)$ is not equal τ_i then **goto** LOOP ($\lambda_i(z)$ stabilized too early)

13 Set $e_k = \log_{\alpha} m_k$ for $1 \leq k \leq \tau_i$ and let $\sigma_i = \{y^{e_1}, y^{e_2}, \dots, y^{e_{\tau_i}}\}$.
 If $e_k \geq \prod_{i=1}^n r_i$ then $e_k > \deg_y K_r(H)$ so **output** FAIL (either the $\lambda_i(z)$ stabilized too early or K_r or p or all evaluations are unlucky).

14 Solve the τ_i by τ_i shifted transposed Vandermonde system

$$\left\{ \sum_{k=1}^{\tau_i} (\alpha^{s+j})^{e_k} u_k = \text{coefficient of } x^i \text{ in } g_j(x) \text{ for } 0 \leq j < \tau_i \right\}$$

modulo p for u and set $\widehat{h}_i(y) = \sum_{k=1}^{\tau_i} u_k y^{e_k}$. Note: $(\alpha^{s+j})^{e_k} = m_k^{s+j}$.

End for loop.

15 Set $Hp = \sum_{i=0}^{d_0} \widehat{h}_i(y) x^i$ and **output** Hp .

The main **for** loop in Step 3 of algorithm PGCD1 evaluates $K_r A$ and $K_r B$ at α^{s+j} for $j = 2T$ and $j = 2T + 1$ in Step 4 and computes their gcd in Step 6, that is, it computes two images before running the Berlekamp-Massey algorithm in Step 10. In our parallel implementation of algorithm PGCD1, for a multi-core computer with $N > 1$ cores, we compute N images at a time in parallel. We discuss this in Section 5.1.

Algorithm SGCD1($K_r A, K_r B, K_r \Gamma, d_0, \sigma, \tau, p$)

Inputs $K_r A, K_r B \in \mathbb{Z}[x, y]$, $K_r \Gamma \in \mathbb{Z}[y]$, $d_0 \geq \deg_{x_0} G$ where $G = \gcd(A, B)$, supports σ_i for $K_r(h_i)$ and $\tau_i = |\sigma_i|$, a smooth prime p .

Output FAIL or UNLUCKY(dmin) or MISSINGTERMS or $Hp \in \mathbb{Z}_p[x, y]$ satisfying if $d_0 = \deg_{x_0} G$ and $\sigma_i = \text{Supp}(K_r(h_i))$ then $Hp = K_r(H) \pmod p$.

1 Pick a random shift s such that $0 < s < p$ and any generator α for \mathbb{Z}_p^* .

- 2 Set $T = \max_{i=1}^{d_0} \tau_i$.
- 3 For j from 0 to T do // includes 1 check point
 - 4 Compute $a_j = K_r A(x, \alpha^{s+j}) \bmod p$ and $b_j = K_r B(x, \alpha^{s+j}) \bmod p$.
 - 5 If $\deg_x a_j < \deg_x K_r A$ or $\deg_x b_j < \deg_{x_0} K_r B$ then **output** FAIL (α^{s+j} is a bad evaluation point.)
 - 6 Compute $g_j = \gcd(a_j, b_j) \in \mathbb{Z}_p[x]$ using the Euclidean algorithm.
Make g_j monic and set $g_j = K_r \Gamma(\alpha^{s+j}) \times g_j \bmod p$.
- End for loop.
- 6 Set $dmin = \min \deg g_j(x)$ and $dmax = \max \deg g_j$ for $0 \leq j \leq T$.
If $dmin < d_0$ **output** UNLUCKY($dmin$).
If $dmax > d_0$ **output** FAIL.
- 7 For i from 0 to d_0 do
 - 8 Let $\sigma_i = \{y^{e_1}, y^{e_2}, \dots, y^{e_{\tau_i}}\}$.
Solve the τ_i by τ_i shifted transposed Vandermonde system

$$\left\{ \sum_{k=1}^{\tau_i} u_k (\alpha^{s+j})^{e_k} = \text{coefficient of } x^i \text{ in } g_j(x) \text{ for } 0 \leq j \leq \tau_i - 1 \right\}$$
 - modulo p for u and set $\hat{h}_i(y) = \sum_{k=1}^{\tau_i} u_k y^{e_k}$. Note $(\alpha^{s+j})^{e_k} = m_k^{s+j}$.
 - 9 If $\hat{h}_i((\alpha)^{s+\tau_i}) \neq \text{coefficient of } x^i \text{ in } g_{\tau_i}$ then **output** MISSINGTERMS.
- End for loop.
- 10 Set $Hp = \sum_{i=0}^{d_0} \hat{h}_i(y) x^i$ and **output** Hp .

We prove that algorithm MGCD1 terminates and outputs $G = \gcd(A, B)$. We first observe that because MGCD1 avoids bad Kronecker substitutions and bad primes, and because the evaluation points α^{s+j} used in PGCD1 and SGCD1 are not bad, we have $K_r(\Gamma)(\alpha^{s+j}) \neq 0$ and $\deg g_j(x) \geq \deg_{x_0} G$ by Lemma 5. Hence $\deg_x \hat{H} = \deg_{x_0} \hat{G} \geq \deg_{x_0} G$. Therefore, if algorithm MGCD1 terminates, the conditions A and B are primitive and $\hat{G}|A$ and $\hat{G}|B$ imply $\hat{G} = G$.

To prove termination we observe that Algorithm MGCD1 proceeds in four phases. In the first phase MGCD1 loops while $d_0 > \deg_x K_r(H) = \deg_{x_0} G$. Because Γ is either $LC(A)$ or $LC(B)$, even if K_r or p or all evaluation points are unlucky, the scaled images in Step 6 of algorithm PGCD1 are images of a polynomial in $\mathbb{Z}[x, y]$ hence the $\lambda_i(z)$ polynomials must stabilize and algorithm PGCD1 always terminates.

Now if PGCD1 or SGCD1 output UNLUCKY($dmin$) then d_0 is decreased, otherwise, they output FAIL or MISSINGTERMS or an image Hp and MGCD1 executes Step 6 at the beginning of the main loop. Eventually the call to DegreeBound in Step 6 will set $d_0 = \deg_{x_0} G$ after which unlucky Kronecker substitutions, unlucky primes and unlucky evaluation points can be detected.

Suppose $d_0 = \deg_{x_0} G$ for the first time. In the second phase MGCD1 loops while PGCD1 outputs FAIL due to an unlucky Kronecker substitution or an unlucky prime

or bad or unlucky evaluation points or the Berlekamp-Massey algorithm stabilized too early. If PGCD1 outputs FAIL, since we don't know if this is due to an unlucky Kronecker substitution or an unlucky prime p , MGCD1 increases r_i by 1 and the size of p by 1 bit. Since there are only finitely many unlucky K_r , eventually K_r will be lucky. And since there are only finitely many unlucky primes, eventually p will be lucky. Finally, since we keep increasing the length of p , eventually p will be sufficiently large so that no bad or unlucky evaluations are encountered in PGCD1 and the Berlekamp-Massey algorithm does not stabilize too early. Then PGCD1 succeeds and outputs an image Hp with $\deg_x Hp = d_0 = \deg_{x_0} G$.

In the third phase MGCD1 loops while the $\sigma_i \not\supseteq K_r(h_i)$, that is, we don't yet have the support for all $K_r(h_i) \in \mathbb{Z}[y]$ either because of missing terms or because a $\lambda_i(z)$ polynomial stabilized too early in PGCD1, and went undetected.

We now prove that Step 9 of SGCD1 detects that $\sigma_i \not\supseteq \text{Supp}(K_r(h_i))$ with probability at least $\frac{3}{4}$ so that PGCD1 is called again in MGCD1.

Suppose $\sigma_i \not\supseteq \text{Supp}(K_r(h_i))$ for some i . Consider the first τ_i equations in Step 8 of SGCD1. We first argue that this linear system has a unique solution. Let $m_k = \alpha^{e_k}$ so that $(\alpha^{s+j})^{e_k} = m_k^{s+j}$. The coefficient matrix W of the linear system has entries

$$W_{jk} = m_k^{s+j-1} \text{ for } 1 \leq j \leq \tau_i \text{ and } 1 \leq k \leq \tau_i.$$

W is a shifted transposed Vandermonde matrix with determinant

$$\det W = m_1^s \times m_2^s \times \cdots \times m_{\tau_i}^s \times \prod_{1 \leq j < k \leq \tau_i} (m_j - m_k).$$

Since $m_k = \alpha^{e_k}$ we have $m_k \neq 0$ and since $p > \deg_y K_r(H)$ the m_k are distinct hence $\det W \neq 0$ and the linear system has a unique solution for u .

Let $E(y) = \phi_p(K_r(h_i)(y)) - \hat{h}_i(y)$ where $\hat{h}_i(y) = \sum_{k=1}^{\tau_i} u_k y^{e_k}$ is the polynomial in $\mathbb{Z}_p[y]$ computed in Step 8 of SGCD1. It satisfies $E(\alpha^{s+j}) = 0$ for $0 \leq j < \tau_i$. If $\sigma_i \not\supseteq \text{Supp}(K_r(h_i))$ then $E(y) \neq 0$ and algorithm SGCD1 tests for this in Step 9 when it checks if $E(\alpha^{s+\tau_i}) \neq 0$. It is possible, however, that $E(\alpha^{s+\tau_i}) = 0$. We bound the probability that this can happen.

Lemma 10. *If s is chosen at random from $[1, p-1]$ then*

$$\text{Prob}[E(\alpha^{s+\tau_i}) = 0] < \frac{1}{4}.$$

Proof. The condition in Step 13 of algorithm PGCD1 means $\deg \bar{h}_i(y) < \prod_{j=1}^n r_j$ hence $\deg_y(E) < \prod_{j=1}^n r_j$. Now s is chosen at random so $\alpha^{s+\tau_i}$ is random on $[1, p-1]$ therefore

$$\text{Prob}[E(\alpha^{s+\tau_i}) = 0] \leq \frac{\deg_y(E)}{p-1} < \frac{\prod_{j=1}^n r_j}{p-1}.$$

Since the primes in SGCD1 satisfy $p > 4 \prod_{j=1}^n r_j$ the result follows. \square

Thus eventually $\sigma_i \not\supseteq \text{Supp}(K_r(h_i))$ is detected in Step 9 of algorithm SGCD1. Because we cannot tell whether this is caused by missing terms or $\lambda_i(z)$ stabilizing too early and going undetected in Steps 12 and 13 of PGCD1, we increase the size of p by 1 bit in Step 17 so that with repeated calls to PGCD1, $\lambda_i(z)$ will eventually not stabilize early and we obtain $\sigma_i \supseteq \text{Supp}(K_r(h_i)) \pmod p$.

How many good images are needed before $\sigma_i \supseteq \text{Supp}(K_r(h_i))$ for all $0 \leq i \leq d_0$? Let p_{min} be the smallest prime used by algorithm PGCD1. Let $N = \lfloor \log_{p_{min}} \|K_r(H)\| \rfloor$. Since at most N primes $\geq p_{min}$ can divide any integer coefficient in $K_r(H)$ then $N+1$ good images from PGCD1 are sufficient to recover the support of $K_r(H)$.

In the fourth and final phase MGCD1 loops calling SGCD1 while $\widehat{H} \neq K_r(H)$. If SGCD1 outputs an image Hp then since $d_0 = \deg_{x_0} H$ and $\sigma_i \supseteq \text{Supp}(K_r(h_i))$ then Hp satisfies $Hp = H \pmod{p}$. The image is combined with previously computed images in \widehat{H} using Chinese remaindering. But as noted in example 9, \widehat{H} may contain a bad image. A bad image arises because either PGCD1 returns a bad image Hp because a $\lambda_i(z)$ stabilized too early or because SGCD1 uses a support with missing terms and fails to detect it.

Consider the prime p and polynomial $h(x, y)$ in Step 8 of MGCD1. Suppose $h(x, y)$ is a bad image, that is, $h \neq K_r(H) \pmod{p}$. We claim Steps 7–10 of MGCD1 detect this bad image with probability at least $1/2$ and since the test for a bad image is executed repeatedly in the main loop, algorithm MGCD1 eventually detects it and removes it hence eventually MGCD1 computes $K_r(H)$ and terminates with output G .

To prove the claim recall that $H = \Delta G$ and $LC(H) = \Gamma$. Because Step 8 of PGCD1 requires $T \geq \#K_r(\Gamma)$ this ensures algorithm PGCD1 always outputs Hp with $LC(Hp) = K_r(\Gamma) \pmod{p}$ hence $LC(h) = K_r(\Gamma) \pmod{p}$.

If $h = K_r(H) \pmod{p}$ and $K_r(\Gamma)(\beta) \neq 0$ then in Step 10 of MGCD1 $h(x, \beta)$ must divide $a(x, \beta)$ and divide $b(x, \beta)$ as $a(x, \beta) = K_r(A)(x, \beta)$ and $b(x, \beta) = K_r(B)(x, \beta)$. Now suppose $h \neq K_r(H) \pmod{p}$. Then Step 10 of MGCD1 fails to detect this bad image if $K_r(\Gamma)(\beta) \neq 0$ and $h(x, \beta) | a(x, \beta)$ and $h(x, \beta) | b(x, \beta)$ in $\mathbb{Z}_p[x]$. Since $\deg_x h = d_0 = \deg_x K_r(H)$ it must be that $h(x, \beta)$ is an associate of $K_r(H)(x, \beta)$. But since $LC(h) = K_r(\Gamma) \pmod{p} = LC(K_r(H)) \pmod{p}$ we have $h(x, \beta) = K_r(H)(x, \beta) \pmod{p}$. Let $E = h - K_r(H) \pmod{p}$. Therefore the test for a bad image h succeeds iff $K_r(\Gamma)(\beta) \neq 0$ and $E(x, \beta) \neq 0$. Lemma 11 below implies the test succeeds with probability at least $1/2$.

Lemma 11. *If β is chosen at random from $[0, p-1]$ then*

$$\text{Prob}[K_r(\Gamma)(\beta) \neq 0] \geq \frac{3}{4} \quad \text{and} \quad \text{Prob}[E(x, \beta) \neq 0] \geq \frac{3}{4}.$$

Proof. The primes p chosen in Step 15 of MGCD1 satisfy $p > 2^\delta \prod_{i=1}^n r_i$ with $\delta \geq 2$. Since $\deg_y K_r(\Gamma) < \prod_{i=1}^n r_i$ by Step 3 of MGCD1 then $\text{Prob}[K_r(\Gamma)(\beta) \pmod{p} = 0] \leq \frac{\deg_y(\Gamma)}{p} < \frac{1}{4}$. Since $\deg_y h < \prod_{i=1}^n r_i$ by Step 13 of PGCD 1 and since r_i is chosen in Step 3 of MGCD1 so that $r_i \geq \deg_{x_i} H$ we have $\deg_y K_r(H) < \prod_{i=1}^n r_i$. Hence $\text{Prob}[E(x, \beta) = 0] \leq \frac{\deg_y E}{p} < \frac{1}{4}$. \square

4.5 Determining t

Algorithm PGCD1 tests in Steps 9 and 10 if both of the last two discrepancies are 0 before it executes Step 11. But it is possible that in Step 11 $\tau_i < \#h_i$.

Let $V_r = (v_0, v_1, \dots, v_{2r-1})$ be a sequence where $r \geq 1$. The Berlekamp-Massey algorithm (BMA) with input V_r computes a feedback polynomial $c(z)$ which is the reciprocal of $\lambda(z)$ if $r = t$. In PGCD1, we determine the t by computing $c(z)$ s on the input sequence V_r for $r = 1, 2, 3, \dots$. If a $c(z)$ remains unchanged from the input V_k to the input V_{k+1} , then we conclude that this $c(z)$ is *stable* which implies that the last two

consecutive discrepancies are both zero, see [30, 26] for a definition of the discrepancy. However, it is possible that the degree of $c(z)$ on the input V_{k+2} might increase again. In [26], Kaltofen, Lee and Lobo proved (Theorem 3) that the BMA encounters the first zero discrepancy after $2t$ points with probability at least

$$1 - \frac{t(t+1)(2t+1)\deg(C)}{6|S|}$$

where S is the set of all possible evaluation points. Here is an example where we encounter a zero discrepancy before $2t$ points. Consider

$$f(y) = y^7 + 60y^6 + 40y^5 + 48y^4 + 23y^3 + 45y^2 + 75y + 55$$

over \mathbb{Z}_{101} with generator $\alpha = 93$. Since f has 8 terms, 16 points are required to determine the correct $\lambda(z)$ and two more for confirmation. We compute $f(\alpha^j)$ for $0 \leq j \leq 17$ and obtain $V_9 = (44, 95, 5, 51, 2, 72, 47, 44, 21, 59, 53, 29, 71, 39, 2, 27, 100, 20)$. We run the BMA on input V_r for $1 \leq r \leq 9$ and obtain feedback polynomials in the following table.

r	Output $c(z)$
1	$69z + 1$
2	$24z^2 + 59z + 1$
3	$24z^2 + 59z + 1$
4	$24z^2 + 59z + 1$
5	$70z^7 + 42z^6 + 6z^3 + 64z^2 + 34z + 1$
6	$70z^7 + 42z^6 + 25z^5 + 87z^4 + 16z^3 + 20z^2 + 34z + 1$
7	$z^7 + 67z^6 + 95z^5 + 2z^4 + 16z^3 + 20z^2 + 34z + 1$
8	$31z^8 + 61z^7 + 91z^6 + 84z^5 + 15z^4 + 7z^3 + 35z^2 + 79z + 1$
9	$31z^8 + 61z^7 + 91z^6 + 84z^5 + 15z^4 + 7z^3 + 35z^2 + 79z + 1$

The ninth call of the BMA confirms that the feedback polynomial returned by the eighth call is the desired one. But, by our design, the algorithm terminates at the third call because the feedback polynomial remains unchanged from the second call. It also remains unchanged for V_4 . In this case, $\lambda(z) = z^2c(1/z) = z^2 + 59z + 24$ has roots 56 and 87 which correspond to monomials y^4 and y^{20} since $\alpha^4 = 56$ and $\alpha^{20} = 87$. The example shows that we may encounter a stable feedback polynomial too early.

5 Implementation and Optimizations

5.1 Evaluation

Let $A, B \in \mathbb{Z}_p[x_0, x_1, \dots, x_n]$, $s = \#A + \#B$, and $d = \max_{i=1}^n d_i$ where $d_i = \max(\deg_{x_i} A, \deg_{x_i} B)$. If we use a Kronecker substitution

$$K(A) = A(x, y, y^{r_1}, \dots, y^{r_1 r_2 \dots r_{n-1}}) \text{ with } r_i = d_i + 1,$$

then $\deg_y K(A) < (d+1)^n$. Thus we can evaluate the s monomials in $K(A)(x, y)$ and $K(B)(x, y)$ at $y = \alpha^k$ in $O(sn \log d)$ multiplications. Instead we first compute $\beta_1 = \alpha^k$ and $\beta_{i+1} = \beta_i^{r_i}$ for $i = 1, 3, \dots, n-2$ then precompute n tables of powers $1, \beta_i, \beta_i^2, \dots, \beta_i^{d_i}$ for $1 \leq i \leq n$ using at most nd multiplications. Now, for each term in A and B of the form $cx_0^{e_0} x_1^{e_1} \dots x_n^{e_n}$ we compute $c \times \beta_1^{e_1} \times \dots \times \beta_n^{e_n}$ using the tables

in n multiplications. Hence we can evaluate $K(A)(x, \alpha^k)$ and $K(B)(x, \alpha^k)$ in at most $nd + ns$ multiplications. Thus for T evaluation points $\alpha, \alpha^2, \dots, \alpha^T$, the evaluation cost is $O(ndT + nsT)$ multiplications.

When we first implemented algorithm PGCD we noticed that often well over 95% of the time was spent evaluating the input polynomials A and B at the points α^k . This happens when $\#H \ll \#A + \#B$. The following method uses the fact that for a monomial $M_i(x_1, x_2, \dots, x_n)$

$$M_i(\beta_1^k, \beta_2^k, \dots, \beta_n^k) = M_i(\beta_1, \beta_2, \dots, \beta_n)^k$$

to reduce the total evaluation cost from $O(ndT + nsT)$ multiplications to $O(nd + ns + sT)$. Note, no sorting on x_0 is needed in Step 4b if the monomials in the input A are sorted on x_0 .

Algorithm Evaluate.

Input $A = \sum_{i=1}^m c_i x_0^{e_i} M_i(x_1, \dots, x_n) \in \mathbb{Z}_p[x_0, \dots, x_n]$, $T > 0$, $\beta_1, \beta_2, \dots, \beta_n \in \mathbb{Z}_p$, and integers d_1, d_2, \dots, d_n with $d_i \geq \deg_{x_i} A$.

Output $y_k = A(x_0, \beta_1^k, \dots, \beta_n^k)$ for $1 \leq k \leq T$.

- 1 Create the vector $C = [c_1, c_2, \dots, c_m] \in \mathbb{Z}_p^m$.
- 2 Compute $[\beta_i^j : j = 0, 1, \dots, d_i]$ for $1 \leq i \leq n$.
- 3 Compute $\Gamma = [M_i(\beta_1, \beta_2, \dots, \beta_n) : 1 \leq i \leq m]$.
- 4 For $k = 1, 2, \dots, T$ do
 - 4a Compute the vector $C := [C_i \times \Gamma_i \text{ for } 1 \leq i \leq m]$.
 - 4b Assemble $y_k = \sum_{i=1}^m C_i x_0^{e_i} = A(x_0, \beta_1^k, \dots, \beta_n^k)$.

The algorithm computes y_k as the matrix vector product.

$$\begin{bmatrix} \Gamma_1 & \Gamma_2 & \dots & \Gamma_m \\ \Gamma_1^2 & \Gamma_2^2 & \dots & \Gamma_m^2 \\ \vdots & \vdots & \ddots & \vdots \\ \Gamma_1^T & \Gamma_2^T & \dots & \Gamma_m^T \end{bmatrix} \begin{bmatrix} c_1 x_0^{e_1} \\ c_2 x_0^{e_2} \\ c_3 x_0^{e_3} \\ \vdots \\ c_m x_0^{e_m} \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_T \end{bmatrix}.$$

Even with this improvement evaluation still takes most of the time so we must parallelize it. Each evaluation of A could be parallelized in blocks of size m/N for N cores. In Cilk C, this is only effective, however, if the blocks are large enough (at least 50,000) so that the time for each block is much larger than the time it takes Cilk to create a task. For this reason, it is necessary to also parallelize on k . To parallelize on k for N cores, we multiply the previous N values of C in parallel by the vector

$$\Gamma_N = [M_i(\beta_1, \beta_2, \dots, \beta_n)^N : 1 \leq i \leq m]$$

Because most of the time is still in evaluation, we have considered the asymptotically fast method of van der Hoven and Lecerf [20] and how to parallelize it. For our evaluation problem it has complexity $O(nd + ns + s \log^2 T)$ which is better than our $O(nd + ns + sT)$ method for large T . In [34], Monagan and Wong implemented this method using 64 bit machine integers and in comparing it with our method used here, found the break even point to be around $T = 500$.

5.2 The non-monic case and homogenization.

Algorithm PGCD interpolates $H = \Delta G$ from scaled monic images $K(\Gamma)(\alpha^j)g_j(x)$ which are computed in Step 6. If the number of terms of Δ is m and $m > 1$ then it is likely that $\#H$ is greater than $\#G$, which means we need more evaluation points for sparse interpolation. For sparse inputs, this may increase t by a factor of m .

One such example occurs in multivariate polynomial factorization. Given a polynomial f in $\mathbb{Z}[x_0, x_1, \dots, x_n]$, factorization algorithms first identify and remove repeated factors by doing a square-free factorization. See Section 8.1 of [14]. The first Step of square-free factorization computes

$$g = \gcd(f, h = \frac{\partial f}{\partial x_0}).$$

Then we have $\Gamma = \gcd(LC(f), LC(h)) = \gcd(LC(f), dLC(f)) = LC(f)$ and $\Delta = LC(f)/LC(g)$ which can be a large polynomial.

Obviously, if either A or B is monic in x_i for some $i > 0$ then we may simply use x_i instead of x_0 as the main variable in our GCD algorithm so that $\#\Gamma = \#\Delta = 1$. Similarly, if either A or B have a constant term in any x_i , that is, $A = \sum_{j=0}^i a_j x_i^j$ and $B = \sum_{j=0}^i b_j x_i^j$ and either a_0 or b_0 are integers, then we can reverse the coefficients of both A and B in x_i so that again $\#\Gamma = \#\Delta = 1$. But many multivariate GCD problems in practice do not satisfy any of these conditions, for example, the polynomial $A = (x_1^2 + 1)x_0^2 + (x_1^2 + 2) = (x_0^2 + 1)x_1^2 + (x_0^2 + 2)$.

Suppose A or B has a constant term. We propose to exploit this by homogenizing A and B . Let f be a non-zero polynomial in $\mathbb{Z}[x_0, x_1, \dots, x_n]$ and

$$H_z(f) = f\left(\frac{x_0}{z}, \frac{x_1}{z}, \dots, \frac{x_n}{z}\right)z^{\deg f}$$

denote the homogenization of f in z . For example, $H_z(A) = 2z^4 + (x_0^2 + x_1^2)z^2 + x_0^2x_1^2$ which as a polynomial in z has an integer leading coefficient. We have the following properties of $H_z(f)$.

Lemma 12. *Let a and b be in $\mathbb{Z}[x_0, x_1, \dots, x_n]$. For non-zero a and b*

- (i) $H_z(a)$ is homogeneous in z, x_1, \dots, x_n of degree $\deg a$,
- (ii) $H_z(a)$ is invertible: if $f(z) = H_z(a)$ then $H_z^{-1}(f) = f(1) = a$,
- (iii) $H_z(ab) = H_z(a)H_z(b)$, and
- (iv) $H_z(\gcd(a, b)) = \gcd(H_z(a), H_z(b))$.

PROOF: To prove (i) let $M = x_0^{d_1} x_1^{d_2} \dots x_n^{d_n}$ be a monomial in a and let $d = \deg a$. Then

$$H_z(M) = z^d \frac{x_0^{d_0}}{z^{d_0}} \dots \frac{x_n^{d_n}}{z^{d_n}}.$$

Observe that since $d \geq d_0 + d_1 + \dots + d_n$ then $\deg_z(H_z(M)) \geq 0$ and $\deg H_z(M) = d$. Properties (ii) and (iii) follow easily from the definition of H_z . To prove (iv) let $g = \gcd(a, b)$. Then $a = g\bar{a}$ and $b = g\bar{b}$ for some \bar{a}, \bar{b} with $\gcd(\bar{a}, \bar{b}) = 1$. Now

$$\begin{aligned} \gcd(H_z(a), H_z(b)) &= \gcd(H_z(g\bar{a}), H_z(g\bar{b})) \\ &= \gcd(H_z(g)H_z(\bar{a}), H_z(g)H_z(\bar{b})) \text{ by (iii)} \\ &= H_z(g) \times \gcd(H_z(\bar{a}), H_z(\bar{b})) \text{ up to units.} \end{aligned}$$

Let $c(z) = \gcd(H_z(\bar{a}), H_z(\bar{b}))$ in $\mathbb{Z}[z, x_0, \dots, x_n]$. It suffices to prove that $\gcd(\bar{a}, \bar{b}) = 1$ implies $c(z)$ is a unit. Now $c(z) = \gcd(H_z(\bar{a}), H_z(\bar{b})) \Rightarrow c(z)|H_z(\bar{a})$ and $c(z)|H_z(\bar{b})$ which implies

$$H_z(\bar{a}) = c(z)q(z) \quad \text{and} \quad H_z(\bar{b}) = c(z)r(z)$$

for some $q, r \in \mathbb{Z}[z, x_0, \dots, x_n]$. Applying H^{-1} to these relations we get $\bar{a} = c(1)q(1)$ and $\bar{b} = c(1)r(1)$. Now $\gcd(\bar{a}, \bar{b}) = 1$ implies $c(1)$ is a unit and thus $q(1) = \pm\bar{a}$ and $r(1) = \pm\bar{b}$. We need to show that $c(z)$ is a unit. Let $d = \deg H_z(\bar{a})$. Since $\deg H_z(\bar{a}) = \deg \bar{a}$ by (i) and $q(1) = \pm\bar{a}$ then $\deg q(1) = d$ and hence $\deg q(z) \geq d$. Now since $H_z(\bar{a}) = c(z)q(z)$ it must be that $\deg c(z) = 0$ and $\deg q(z) = d$. Since $c(1) = \pm 1$ then $\deg c(z) = 0$ implies $c(z) = \pm 1$. \square .

Properties (iii) and (iv) mean we can compute $G = \gcd(A, B)$ using

$$G = H_z^{-1} \gcd(H_z(A), H_z(B)).$$

Notice also that homogenization preserves sparsity. To see why homogenization may help we consider an example.

Example 11. Let $G = x^2 + y + 1$, $\bar{A} = xy + x + y + 1 = (y+1)x + (y+1) = (x+1)y + (x+1)$ and $B = x^2y + xy^2 + x^2 + y^2 = (y+1)x^2 + y^2(x+1)$. Then $H_z(G) = z^2 + yz + x^2$, $H_z(\bar{A}) = z^2 + (x+y)z + xy$, and $H_z(B) = (x^2 + y^2)z + (x^2y + xy^2)$.

Notice in Example 11 that A and B are neither monic in x nor monic in y but since A has a constant term, $H_z(A)$ is monic in z . If we use x as x_0 in Algorithm PGCD then $\Gamma = \gcd(y+1, y+1) = y+1 = \Delta$ and we interpolate $H = \Delta G = (y+1)x^2 + (y^2 + 2y + 1)$ and $t = 3$. If we use y as x_0 in Algorithm PGCD then $\Gamma = \gcd(x+1, x+1) = x+1 = \Delta$ and we interpolate $H = \Delta G = (x+1)y + (x^3 + x^2 + x + 1)$ and $t = 4$. But if we use z as x_0 in Algorithm PGCD then $\Gamma = \gcd(1, x^2 + y^2) = 1$ hence $\Delta = 1$ and we interpolate $H_z(G) = z^2 + yz + x^2$ and $t = 1$.

In our implementation, if $\#\Gamma > 1$ and A or B has a constant term we always homogenize. There is, however, a cost to in homogenizing for the GCD problem, namely, we increase the number of variables to interpolate by 1 and we increase the cost of the univariate images in $\mathbb{Z}_p[z]$ if the degree increases. The degree may increase by up to a factor of $n+1$. For example, if $G = 1 + \prod_{i=0}^n x_i^{d-1}$, $\bar{A} = 1 + \prod_{i=0}^n x_i$ and $\bar{B} = 1 - \prod_{i=0}^n x_i$ then $\deg_{x_i} A = d = \deg_{x_i} B$ but $\deg_z H_z(A) = (n+1)d = \deg_z H_z(B)$. Homogenizing can also increase t when G has many terms of the same total degree.

5.3 Bivariate images

Recall that we interpolate $H = \sum_{i=0}^{dG} h_i(x_1, \dots, x_n)x_0^i$ where $H = \Delta G$. The number of evaluation points used by algorithm PGCD is $2t + O(1)$ where $t = \max_{i=0}^{dG} \#h_i$. Since the cost of our algorithm is multiplied by the number of evaluation points needed we can reduce the cost of algorithm PGCD if we can reduce t .

Algorithm PGCD interpolates H from univariate images in $\mathbb{Z}_p[x_0]$. If instead we interpolate H from bivariate images in $\mathbb{Z}_p[x_0, x_1]$, this will likely reduce t when $\#\Delta = 1$ and when $\#\Delta > 1$. For our benchmark problem, where $\Delta = 1$, $D = 60$ and $n = 8$, doing this reduces t from 1198 to 122 saving a factor of 9.8. On the other hand, we must now compute bivariate GCDs in $\mathbb{Z}_p[x_0, x_1]$ instead of univariate GCDs in $\mathbb{Z}_p[x_0]$ using the Euclidean algorithm. To decide whether this will lead to an overall gain, we need to know the cost of computing bivariate images and the likely reduction in t .

To compute a bivariate GCD in $\mathbb{Z}_p[x_0, x_1]$ we have implemented Brown's dense modular GCD algorithm from [5] which interpolates x_1 using univariate images in $\mathbb{Z}_p[x_0]$. If G is sparse, then for sufficiently large t and n , G is likely dense in x_0 and x_1 , so using a dense GCD algorithm is efficient. Let $d = \max(\deg A, \deg B)$ and let $s = \#A + \#B$. The complexity of Brown's algorithm is $O(d^3)$ arithmetic operations in \mathbb{Z}_p . In comparison, the complexity of Euclid's algorithm for computing a GCD in $\mathbb{Z}_p[x_0]$ is $O(d^2)$ so using bivariate images increases the cost of an image GCD by a factor of $O(d)$. But if the cost of a bivariate images is less than the cost of evaluating the inputs A and B , which using our evaluation algorithm from 5.1 is s multiplications in \mathbb{Z}_p , then the cost of the bivariate images will not increase the overall cost of the algorithm significantly. For our benchmark problem, $s = 2 \times 10^6$ and $d^3 = 40^3 = 64,000$, the cost of a bivariate image is negligible compared with the cost of an evaluation.

To estimate the gain made by using bivariate images instead of univariate images, let us write

$$H = \sum_{i=0}^{d_0} h_i(x_1, \dots, x_n) x_0^i = \sum_{i=0}^{d_0} \sum_{j=0}^{d_1} h_{ij}(x_2, \dots, x_n) x_0^i x_1^j$$

and define $t_1 = \max \#h_i$ and $t_2 = \max \#h_{ij}$. The ratio t_1/t_2 is reduction of the number of evaluation points needed by our algorithm. The maximum reduction in t occurs when the terms in H are distributed evenly over the coefficients of H in x_1 , that is, then $t_1/t_2 = 1 + d_1 = 1 + \deg_{x_1} \Delta + \deg_{x_1} G$. For some very sparse inputs, there is no gain. For example, for

$$H = x_0^d + x_1^d + x_2^d + \dots + x_n^d + 1$$

we have $t_1 = n$ and $t_2 = n - 1$ and the gain is negligible.

If H has total degree D and H is dense then the number of terms in $h_i(x_1, \dots, x_n)$ is $\binom{D-i+n}{n}$ which is a maximum for h_0 where $\#h_0 = \binom{D+n}{n}$. A conservative assumption is that $\#h_i$ is proportional to $\binom{n+D-i}{n}$ and similarly $\#h_{ij}$ is proportional to $\binom{n-1+D-(i+j)}{n-1}$. In this case, the reduction is a factor of

$$\frac{\#h_0}{\#h_{00}} = \binom{n+D}{n} / \binom{n-1+D}{n-1} = \frac{n+D}{n}.$$

For our benchmark problem where $n = 8$ and $D = 60$ this is $8.5 = \frac{68}{8}$. Table 4 in Section 6 shows that this idea is very effective for gcds with many terms. We remark that the use of bivariate images would also benefit an implementation of Zippel's sparse GCD algorithm from [47].

6 Benchmarks

We have implemented algorithm PGCD1 for 31, 63 and 127 bit primes in Cilk C. For 127 bit primes we use the 128 bit signed integer type `__int128_t` supported by the gcc compiler. We parallelized evaluation (see Section 5.1) and we interpolate the coefficients $h_i(y)$ in parallel in Step 11 of Algorithm PGCD1.

The new algorithm requires $2t + \delta$ images (evaluation points) for the first prime and $t + 1$ images for the remaining primes. The additional image ($t + 1$ images instead of t) is used to check if the support of H (see Step 9 of Algorithm SGCD1) obtained from the first prime is correct.

To assess how good our new algorithm is, we have compared it with the serial implementations of Zippel's algorithm in Maple 2016 and Magma V2.22. For Maple we

are able to determine the time spent computing G modulo the first prime in Zippel’s algorithm. It is typically over 99% of the total GCD time. The reason for this is that Zippel’s algorithm requires $O(ndt)$ images for the first prime but only t images for the remaining primes.

We also timed Maple’s implementation of Wang’s EEZ-GCD algorithm from [45, 46]. It was much slower than Zippel’s algorithm on these inputs so we have not included timings for it. Note, older versions of Maple and Magma both used the EEZ-GCD algorithm for multivariate polynomial GCD computation.

All timings were made on the gaby server in the CECM at Simon Fraser University. This machine has two Intel Xeon E-2660 8 core CPUs running at 3.0 GHz on one core and 2.2 GHz on 8 cores. Thus maximum parallel speedup is a factor of $16 \times 2.2/3.0 = 11.7$.

6.1 Benchmark 1

For our first benchmark (see Table 3) we created polynomials G, \bar{A} and \bar{B} in 6 variables ($n = 5$) and 9 variables ($n = 8$) of degree at most d in each variable. We generated $100d$ terms for G and 100 terms for \bar{A} and \bar{B} . That is, we hold t approximately fixed to test the dependence of the algorithms on d .

The integer coefficients of G, \bar{A}, \bar{B} were generated at random from $[0, 2^{31} - 1]$. The monomials in G, \bar{A} and \bar{B} were generated using random exponents from $[0, d - 1]$ for each variable. For G we included monomials 1 and x_0^d so that G is monic in x_0 and $\Gamma = 1$. The reason we make G monic for our benchmarks is because Maple and Magma use different strategies for handling the non-monic case. Maple and Magma code for generating the input polynomials is given in the Appendix.

Our new algorithm used the 62 bit prime $p = 29 \times 2^{57} + 1$. Maple used the 32 bit prime $2^{32} - 5$ for the first image in Zippel’s algorithm.

			New GCD algorithm		Zippel’s algorithm	
n	d	t	1 core (eval)	16 cores	Maple	Magma
5	5	110	0.29s (64%)	0.074s (3.9x)	3.57s	0.60s
5	10	114	0.62s (68%)	0.091s (6.8x)	48.04s	6.92s
5	20	122	1.32s (69%)	0.155s (8.5x)	185.70s	296.06s
5	50	121	3.48s (69%)	0.326s (10.7x)	1525.80s	$> 10^5 s$
5	100	123	7.08s (69%)	0.657s (10.8x)	6018.23s	NA
5	200	125	14.64s (71%)	1.287s (11.4x)	NA	NA
5	500	135	38.79s (71%)	3.397s (11.4x)	NA	NA
8	5	89	0.27s (61%)	0.065s (4.2x)	32.47s	2.28s
8	10	110	0.63s (65%)	0.098s (6.4x)	138.41s	7.33s
8	20	114	1.35s (66%)	0.163s (8.3x)	664.33s	78.77s
8	50	113	3.52s (66%)	0.336s (10.5x)	6390.22s	800.15s
8	100	121	7.43s (68%)	0.645s (11.5x)	NA	9124.73s

Table 3: Real times (seconds) for GCD problems.

In Table 3 column d is the maximum degree of the terms of G, \bar{A}, \bar{B} in each variable, column t is the maximum number of terms of the coefficients of G , that is, if $G = \sum_{i=0}^d g_i(x_1, \dots, x_n)x_0^i$, $t = \max_{i=0}^d \#g_i$. Timings are shown in seconds for the new algorithm for 1 core and 16 cores. For 1 core we show the %age of the time spent evaluating the inputs, that is computing $K(A)(x_0, \alpha^j)$ and $K(B)(x_0, \alpha^j)$ for $j = 1, 2, \dots, T$.

The parallel speedup on 16 cores is shown in parentheses.

Table 3 shows that most of the time in the new algorithm is in evaluation. It shows a parallel speedup approaching the maximum of 11.7 on this machine. There was a parallel bottleneck in how we computed the $\lambda_i(z)$ polynomials that limited parallel speedup to 10 on these benchmarks. For N cores, after generating a new batch of N images we used the Euclidean algorithm for Step 12b which is quadratic in the number of images j computed so far. To address this we now use an incremental version of the Berlekamp-Massey algorithm which is $O(Nj)$.

In comparing the new algorithm with Maple’s implementation of Zippel’s algorithm, for $n = 8, d = 50$ in Table 3 we achieve a speedup of a factor of $1815 = 6390.22/3.52$ on 1 core. Since Zippel’s algorithm uses $O(ndt)$ points and our Ben-Or/Tiwari algorithm uses $2t + O(1)$ points, we get a factor of $O(nd)$ speedup because of this.

6.2 Benchmark 2

Our second benchmark (see Table 4) is for 9 variables where the degree of G, \bar{A}, \bar{B} is at most 20 in each variable. The terms are generated at random as in Benchmark 1 but are restricted to have total degree at most 60. As in Benchmark 1, we make G monic in all variables so that $\Gamma = 1$. The reason we make G monic is because Maple and Magma use different strategies for handling the non-monic case. The row with $\#G = 10^4$ and $\#A = 10^6$ is our benchmark problem from Section 1. Maple and Magma code for generating the input polynomials is given in the Appendix.

We show two sets of timings for our new algorithm. The first set of timings under “New GCD: univariate images” is for interpolating G from univariate images in $\mathbb{Z}_p[x_0]$. The second set of timings under “New GCD: bivariate images” is for interpolating G from bivariate images in $\mathbb{Z}_p[x_1, x_1]$ as described in Section 5.3. Here $t = \max_{i,j} \#g_{i,j}$ where $G = \sum_{i,j} g_{i,j}(x_2, \dots, x_n)x_0^i x_1^j$. The faster timings for bivariate images is due to the reduction in t .

The timings reported for the new algorithm are for the first prime only. We used the 62 bit prime $29 \times 2^{57} + 1$ for both sets of timings. Although one prime is sufficient for these problems to recover H that is, no Chinese remaindering is needed, our algorithm uses an additional 63 bit prime to verify $H \bmod p_1 = H$. The time for the second prime is always less than 50% of the time for the first prime because it needs only $t + 1$ points instead of $2t + \delta$ points and it does not need to compute degree bounds.

Table 4 shows again that most of the time in the new algorithm is in evaluation. This is also true of Zippel’s algorithm and hence of Maple and Magma too. Because Maple uses random evaluation points, and not a power sequence, the cost of each evaluation in Maple is $O(n(\#A + \#B))$ multiplications instead of $\#A + \#B$ evaluations for the new algorithm. Also, Maple is using `% p` to divide in \mathbb{C} which generates a hardware division instruction which is much more expensive than a hardware multiplication instruction. We can also see that when $\#G$ becomes much larger than $\#\bar{A}$ and $\#\bar{B}$, the Magma timings decrease but the Maple timings increase. This is because Magma is also interpolating \bar{A} and \bar{B} whereas Maple interpolates G only. For the new algorithm, we are using Roman Pearce’s implementation of Möller and Granlund [32] which reduces division by p to two multiplications plus other cheap operations. Magma is doing something similar. It is using floating point primes (25 bits) so that it can multiply modulo p using floating point multiplications. This is one reason why Maple is slower than Magma.

#G #A	New GCD: univariate images			New GCD: bivariate images			Zippel's algorithm	
	t	1 core	16 cores	t	1 core	16 cores	Maple	Magma
10 ² 10 ⁵	13	0.14s (62%)	0.043	4	0.156s (65%)	0.032s	51.8s	7.87s
10 ³ 10 ⁵	113	0.59s (66%)	0.100s	13	0.306s (55%)	0.066s	210.9s	50.98s
10 ⁴ 10 ⁵	1197	7.32s (48%)	1.022s	118	2.299s (36%)	0.224s	7003.4s	10.70s
10 ² 10 ⁶	13	1.36s (70%)	0.167s	4	1.024s (60%)	0.164s	797.4s	50.31s
10 ³ 10 ⁶	130	5.70s (90%)	0.520s	14	1.713s (69%)	0.228s	2135.9s	216.81s
10 ⁴ 10 ⁶	1198	48.17s (87%)	4.673s	122	7.614s (75%)	0.685s	22111.6s	1784.02s
10 ⁵ 10 ⁶	11872	466.09s (82%)	45.83s	1115	80.04s (57%)	6.079s	NA	693.20s
10 ² 10 ⁷	11	12.37s (67%)	1.46s	3	10.69s (58%)	1.63s	NA	354.90s
10 ³ 10 ⁷	122	47.72s (91%)	4.470s	16	15.76s (71%)	2.09s	NA	1553.91s
10 ⁴ 10 ⁷	1212	429.61s (98%)	37.72s	122	57.23s (90%)	5.10s	NA	8334.93s
10 ⁵ 10 ⁷	11867	3705.4s (98%)	311.6s	1114	438.87s (90%)	34.4s	NA	72341.0s
10 ⁶ 10 ⁷	117508	47568.s (90%)	3835.9s	11002	4794.5s (83%)	346.1s	NA	NA
10 ² 10 ⁸	12	129.26s (69%)	15.86s	4	101.8s (60%)	16.88s	NA	NA
10 ³ 10 ⁸	121	522.14s (92%)	49.17s	17	150.0s (73%)	23.25s	NA	NA
10 ⁴ 10 ⁸	1184	4295.0s (99%)	412.69s	121	555.5s (89%)	78.76s	NA	NA
10 ⁵ 10 ⁸	11869	43551.s (99%)	3804.93s	1162	4417.7s (98%)	626.19s	NA	NA

Table 4: Timings (seconds) for 9 variable GCDs

7 Conclusion

We have shown that a Kronecker substitution can be used to reduce a multivariate GCD computation to a bivariate one, and how to interpolate the GCD using a discrete logs Ben-Or/Tiwari point sequence. The main design difficulty is how to handle unlucky Kronecker substitutions, unlucky primes, unlucky evaluation points, and missing terms.

We have implemented our algorithm in Cilk C for 32 bit, 64 bit and 128 bit primes. Our parallel algorithm is fast and practical. We observed that for input polynomials with many terms, evaluating the inputs was the bottleneck. We also presented two practical ways one can reduce the number of evaluation points needed.

Acknowledgement

We would like to acknowledge Adriano Arce for his discrete logarithm code, Alex Fan for his Chinese remaindering code, Alan Wong for integrating the bivariate GCD code and Roman Pearce for his 64 bit and 128 bit modular multiplication codes. We would also thank one of the anonymous referees for their many helpful suggestions.

References

- [1] A. Arnold, M. Giesbrecht and D. Roche. Faster Sparse Multivariate Polynomial Interpolation of Straight-Line Programs. [arXiv:1412.4088\[cs.SC\]](https://arxiv.org/abs/1412.4088), December 2014.
- [2] N. B. Atti, G. M. Diaz-Toca, and H. Lombardi. The Berlekamp-Massey algorithm revisited. *AAECC* **17** pp. 75–82, 2006.
- [3] M. Ben-Or and P. Tiwari. A deterministic algorithm for sparse multivariate polynomial interpolation. In *Proc. of STOC '88*, ACM Press, pp. 301–309, 1988.
- [4] Elwyn Berlekamp. Factoring polynomials over large finite fields. *Mathematics of Computation*, **24**(111) 713–735, 1970.
- [5] W. S. Brown. On Euclid’s Algorithm and the Computation of Polynomial Greatest Common Divisors. *J. ACM* **18**:478–504, 1971.
- [6] W. S. Brown and J. F. Traub. On Euclid’s Algorithm and the Theory of Subresultants. *J. ACM*, **18**:4, 505–514. ACM Press, 1971.
- [7] B. W. Char, K. O. Geddes, G. H. Gonnet. GCDHEU: Heuristic polynomial GCD algorithm based on integer GCD computation. *J. Symbolic Comp.* **7**:1, 31–48. Elsevier, 1989.
- [8] G. E. Collins. Subresultants and reduced polynomial remainder sequences. *J. ACM* **14**:1, 128–142. ACM Press, 1967.
- [9] Matthew Comer, Erich Kaltofen and Clement Pernet. Sparse Polynomial Interpolation and Berlekamp/Massey Algorithms that Correct Outlier Errors in Input Values. *Proceedings of ISSAC 2012*, pp. 138–145, ACM, 2012.
- [10] D. Cox, J. Little, D. O’Shea. *Ideals, Varieties and Algorithms*. Springer-Verlag, 1991.
- [11] Matteo Frigo, Charles E. Leiserson and Keith H. Randall. The Implementation of the Cilk-5 Multithreaded Language. *Proceedings of PLDI 1988*, ACM, 212–223, 1998.
- [12] Sanchit Garg and Eric Schost. Interpolation of polynomials given by straight-line programs. *J. Theor. Comp. Sci.*, **410**:2659–2662, June 2009.
- [13] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, UK, 1999.

- [14] K. O. Geddes, S. R. Czapor, and G. Labahn. *Algorithms for Computer Algebra*. Kluwer, 1992.
- [15] Gelfond, A. O. *Transcendental and Algebraic Numbers*, GITTL, Moscow, 1952; English translation by Leo F. Boron, Dover, New York, 1960.
- [16] Mark Giesbrecht, George Labahn and Wen-shin Lee. Symbolic-numeric sparse interpolation of multivariate polynomials. *J. Symbolic Computation* **44**:943–959, 2009.
- [17] Mark Giesbrecht and Daniel S. Roche. Diversification improves interpolation. In *Proc. ISSAC 2011*, ACM Press, pp. 123–130, 2011.
- [18] Bruno Grenet, Joris van der Hoeven and Grégoire Lecerf. Randomized Root Finding over Finite FFT-fields using Tangent Graeffe Transforms. *Proceedings of ISSAC 2015*, pp. 197–204, ACM, 2015.
- [19] Goldstein, A., Graham, G. A Hadamard-type bound on the coefficients of a determinant of polynomials. *SIAM Review* **1** 394–395, 1974.
- [20] Joris van der Hoven and Grégoire Lecerf. On the bit complexity of sparse polynomial multiplication. *J. Symb. Cmpt.* **50**:227–254, 2013.
- [21] Joris van der Hoven and Grégoire Lecerf. Sparse polynomial interpolation in practice. *CCA* **48**:187–191, September 2015.
- [22] Mahdi Javadi and Michael Monagan. Parallel Sparse Polynomial Interpolation over Finite Fields. In *Proc. of PASCO 2010*, ACM Press, pp. 160–168, 2010.
- [23] Erich Kaltofen. Sparse Hensel lifting. *Proceedings of EUROCAL '85*, LNCS **204**:4–17, Springer, 1985.
- [24] Erich Kaltofen and Barry Trager. Computing with polynomials given by black boxes for their evaluations: greatest common divisors, factorization, separation of numerators and denominators. *Proc. FOCS '88*, 96–105. IEEE, 1988.
- [25] E. Kaltofen, Y.N. Lakshman and J-M. Wiley. Modular Rational Sparse Multivariate Interpolation Algorithm. In *Proc. ISSAC 1990*, pp. 135–139, ACM Press, 1990.
- [26] E. Kaltofen, W. Lee, and A. Lobo. Early Termination in Ben-Or/Tiwari Sparse Interpolation and a Hybrid of Zippel’s algorithm. In *Proc. ISSAC 2000*, ACM Press, pp. 192–201, 2000.
- [27] Erich Kaltofen and Wen-shin Lee. Early termination in sparse interpolation algorithms. *J. Symbolic Computation* **36**:365–400, 2003.
- [28] E. Kaltofen. Fifteen years after DSC and WLSS2. In *Proc. of PASCO 2010*, ACM Press, pp. 10–17, 2010.
- [29] J. de Kleine, M. Monagan, A. Wittkopf. Algorithms for the Non-monic case of the Sparse Modular GCD Algorithm. *Proceedings of ISSAC '2005*, ACM Press, pp. 124–131, 2005.
- [30] J. L. Massey. Shift-register synthesis and BCH decoding. *IEEE Trans. on Information Theory*, 15:122–127, 1969.
- [31] Alfred Menezes, Paul van Oorschot and Scott Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1996. <http://cacr.uwaterloo.ca/hac/>
- [32] Niels Möller and Torbjorn Granlund. Improved division by invariant integers. *IEEE Trans. on Computers*, **60**:165–175, 2011.
- [33] Joel Moses and David Y.Y. Yun. The EZ GCD algorithm. *Proc. ACM '73*, 159–166. ACM Press, 1973.
- [34] Michael Monagan and Alan Wong. Fast parallel multi-point evaluation of sparse polynomials. *Proceedings of PASCO 2017*, ACM digital library, July 2017.
- [35] Michael Monagan and Baris Tuncer. Using Sparse Interpolation in Hensel Lifting. *Proceedings of CASC 2016*, LNCS **9890**:381–400, 2016.

- [36] Gary Mullen and Daniel Panario. *Handbook of Finite Fields*. CRC Press, 2013.
- [37] Hirokazu Murao and Tetsuro Fujise. Modular Algorithm for Sparse Multivariate Polynomial Interpolation and its Parallel Implementation. *J. Symb. Cmpt.* **21**:377–396, 1996.
- [38] S. Pohlig and M. Hellman. An improved algorithm for computing logarithms over $\text{GF}(p)$ and its cryptographic significance. *IEEE Trans. on Information Theory*, **24**:106–110, 1978.
- [39] M. O. Rayes, P. S. Wang and K. Weber. Parallelization fo the Sparse Modular GCD Algorithm for Multivariate Polynomials on Shared Memory Processors. *Proc. ISSAC '94*, 66–73. ACM Press, 1994.
- [40] Tateaki Sasaki and Masayuki Suzuki. Three New Algorithms for Multivariate Polynomial GCD. *J. Symbolic Computaton* **13**:395–411, 1992.
- [41] Jack Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, **27**:701–717, 1980.
- [42] Douglas Stinson. *Cryptography, Theory and Practice*, Chapman and Hall, 2006.
- [43] Y. Sugiyama, M. Kashara, S. Hirashawa and T. Namekawa. A Method for Solving Key Equation for Decoding Goppa Codes. *Information and Control* **27**:87–99, 1975.
- [44] Kuniaki Tsuji. An improved EZ-GCD algorithm for multivariate polynomials. *J. Symbolic Computation* **44**:99–100, 2009.
- [45] Paul Wang. The EEZ-GCD algorithm. *ACM SIGSAM Bulletin*, **14**(2): 50–60, 1980.
- [46] Paul S. Wang. An Improved Multivariate Polynomial Factoring Algorithm. *Mathematics of Computation* **32**(144): 1215–1231. American Math. Soc., 1978.
- [47] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Proc. of EUROSAM '79*, pp. 216–226. Springer-Verlag, 1979.
- [48] Richard Zippel. Interpolating Polynomials from their Values. *J. Symb Cmpt.* **9**:375–403, 1990.

Appendix

Our new code is written in MIT Cilk C. It may be downloaded from <http://www.cecm.sfu.ca/CAG/code/HuMonGCD> . There are two .zip archives, one for using univariate images (gcduni.zip), the other for using bivariate images (gcdbiv.zip). To use it one must use the old MIT Cilk C compiler.

Maple code for the 6 variable gcd benchmark is below.

```
kernelopts(numcpus=1); # run on one core
r := rand(2^31);
X := [u,v,w,x,y,z];
getpoly := proc(X,t,d) local i,e,x;
    e := rand(0..d);
    add( r()*mul(x^e(),x=X), i=1..t );
end;

infolevel[gcd] := 3; # to see output from Zippel's algorithm

for d in [5,10,20,50,100] do
    s := 100; t := 100*d;
    g := add(x^d,x=X) + r() + getpoly(X,t-7,d-1);
    abar := getpoly(X,s-1,d)+r();
    bbar := getpoly(X,s-1,d)+r();
```

```

a := expand(g*abar);
b := expand(g*bbar);
st := time(); h := gcd(a,b); gcdtime := time()-st;
printf("d=%d  time=%8.3f\n",d,gcdtime);
end do:

```

Magma code for the 6 variable gcd benchmark.

```

p := 2^31;
Z := IntegerRing();
P<u,v,w,x,y,z> := PolynomialRing(Z,6);

randpoly := function(d,t)
M := [ u^Random(0,d)*v^Random(0,d)*w^Random(0,d)
      *x^Random(0,d)*y^Random(0,d)*z^Random(0,d) : i in [1..t] ];
C := [ Random(p) : i in [1..t] ];
g := Polynomial(C,M);
return g;
end function;

for d in [5,10,20,50] do
s := 100; t := 100*d;
g := u^d+v^d+w^d+x^d+y^d+z^d + randpoly(d,t-7) + Random(p);
abar := randpoly(d+1,s-1) + Random(p);
bbar := randpoly(d+1,s-1) + Random(p);
a := g*abar;
b := g*bbar;
d; time h := Gcd(a,b);
end for;

```

Maple code for the 9 variable gcd benchmark.

```

kernelopts(numcpus=1); # run on one core
local D;
p := 2^31-1;
r := rand(p);
X := [x1,x2,x3,x4,x5,x6,x7,x8,x9];
getmon := proc(X,d,D) local e,m,x;
e := rand(d);
do m := mul( x^e(), x=X ); until degree(m)<=D;
end;
getpoly := proc(X,t,d,D) local i;
add( r()*getmon(X,d,D), i=1..t );
end;

infolevel[gcd] := 3; # to see output from Zippel's algorithm

d := 20; D := 60; N := 10^5;
t := 1;
while t<N/10 do
t := 10*t;
s := N/t;

```

```

g := x1^d + getpoly(X,t-2,d,D) + r();
abbar := x1^d + getpoly(X,s-2,d,D) + r();
bbar := x1^d + getpoly(X,s-2,d,D) + r();
a := expand(g*abbar);
b := expand(g*bbar);
st := time(): h := gcd(a,b); gcdtime := time()-st;
printf("s=%d t=%d gcdtime=%8.3f\n",s,t,gcdtime);
od:

```

Magma code for the 9 variable gcd benchmark.

```

p := 2013265921;
Z := IntegerRing();
P<x1,x2,x3,x4,x5,x6,x7,x8,x9> := PolynomialRing(Z,9);

function getmon(d,D)
M := x1^Random(0,d)*x2^Random(0,d)*x3^Random(0,d)*x4^Random(0,d)*
x5^Random(0,d)*x6^Random(0,d)*x7^Random(0,d)*x8^Random(0,d)*x9^Random(0,d);
if Degree(M) le D then return M; end if;
return getmon(d,D);
end function;

randpoly := function(d,D,t)
M := [ getmon(d,D) : i in [1..t] ];
C := [ Random(p) : i in [1..t] ];
g := Polynomial(C,M);
return g;
end function;

d := 20; D := 60; N := 100000;
t := 1;
while t lt N/10 do
t := 10*t;
s := N/t;
g := x1^d + randpoly(d,D,t-2) + Random(p);
printf "t=%o s=%o\n", t, s;
abbar := x1^d + randpoly(d,D,s-2) + Random(p);
bbar := x1^d + randpoly(d,D,s-2) + Random(p);
a := g*abbar;
b := g*bbar;
printf " #g=%o #a=%o #b=%o\n", #Coefficients(g), #Coefficients(a), #Coefficients(b);
time h := Gcd(a,b);
end while;

```