# A fast parallel sparse polynomial GCD algorithm. 

Jiaxiong Hu and Michael Monagan<br>Department of Mathematics, Simon Fraser University


#### Abstract

We present a parallel GCD algorithm for sparse multivariate polynomials with integer coefficients. The algorithm combines a Kronecker substitution with a BenOr/Tiwari sparse interpolation modulo a smooth prime to determine the support of the GCD. We have implemented the algorithm in Cilk C for 31, 63 and 127 bit primes. We compare it with Maple and Magma's implementations of Zippel's GCD algorithm.


## 1 Introduction

Let $A$ and $B$ be two non-constant polynomials in $\mathbb{Z}\left[x_{0}, x_{1}, \ldots, x_{n}\right]$. Let $G=\operatorname{gcd}(A, B)$ and let $\bar{A}=A / G$ and $\bar{B}=B / G$. So $G$ is the greatest common divisor of $A$ and $B$ and the polynomials $\bar{A}$ and $\bar{B}$ are called the cofactors of $A$ and $B$.

In this paper we present a sparse GCD algorithm for computing $G$. We will compare it with Zippel's sparse GCD algorithm from [22]. Zippel's algorithm is the main GCD algorithm used by Maple, Magma and Mathematica for polynomials in $\mathbb{Z}\left[x_{0}, x_{1}, \ldots, x_{n}\right]$.

Let $A=\sum_{i=0}^{d A} a_{i} x_{0}^{i}, \quad B=\sum_{i=0}^{d B} b_{i} x_{0}^{i}$ and $G=\sum_{i=0}^{d G} c_{i} x_{0}^{i}$ where the coefficients $a_{i}, b_{i}$ and $c_{i}$ are polynomials in $\mathbb{Z}\left[x_{1}, \ldots, x_{n}\right]$ and $a_{d A}, b_{d B}$ and $c_{d G}$ are the leading coefficients of $A, B$ and $G$. Let $\mathrm{LC}(A)$ denote the leading coefficient of $A$. Let \# denote the number of terms in $A$ and let $\operatorname{Supp}(A)$ denote the support of $A$, that is, the set of monomials appearing in $A$.

We will assume $\operatorname{gcd}\left(a_{i}\right)=1$ and $\operatorname{gcd}\left(b_{i}\right)=1$, that is, the contents have already been computed and divided out. Let $\Gamma=\operatorname{gcd}\left(a_{d A}, b_{d B}\right)$. Since $L C(G) \mid L C(A)$ and $\operatorname{LC}(G) \mid L C(B)$ it must be that $\mathrm{LC}(G) \mid \Gamma$ thus $\Gamma=L C(G) \Delta$ for some polynomial $\Delta \in \mathbb{Z}\left[x_{1}, x_{2}, \ldots, x_{n}\right]$.

Example 1 For $G=x_{1} x_{0}^{2}+x_{2} x_{0}+3, \bar{A}=\left(x_{2}-x_{1}\right) x_{0}+x_{2}$, and $\bar{B}=\left(x_{2}-x_{1}\right) x_{0}+x_{1}+2$ we have $\# G=3, \operatorname{LC}(G)=x_{1}, \Gamma=x_{1}\left(x_{2}-x_{1}\right), \Delta=x_{2}-x_{1}$, and $\operatorname{Supp}(G)=\left\{x_{1} x_{0}^{2}, x_{2} x_{0}, 1\right\}$.

We provide an overview of the algorithm. 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 $\operatorname{gcd}\left(h_{i}\right)$ which is $\Delta$ and divide $H$ by $\Delta$ to obtain $G$.

We compute $H$ modulo a sequence of primes $p_{1}, p_{2}, \ldots$ and recover the integer coefficients of $H$ using Chinese remaindering. The use of Chinese remaindering is standard. Details may be found in $[4,8]$. Let $H_{1}$ be the result of computing $H \bmod p_{1}$. For the remaining primes we use the sparse interpolation approach of Zippel [22] which assumes $\operatorname{Supp}\left(H_{1}\right)=\operatorname{Supp}(H)$. From now on we focus on the computation of $H \bmod p_{1}$ where we compute $\operatorname{Supp}(H)$ which is the hard part of the sparse GCD algorithm.

To compute $H \bmod p$ the algorithm will pick a sequence of points $\beta_{1}, \beta_{2}, \ldots$ from $\mathbb{Z}_{p}^{n}$, compute monic images $g_{j}=\operatorname{gcd}\left(A\left(x_{0}, \beta_{j}\right), B\left(x_{0}, \beta_{j}\right) \in \mathbb{Z}_{p}\left[x_{0}\right]\right.$ of $G$, in parallel, then multiply $g_{j}$ by the scalar $\Gamma\left(\beta_{j}\right) \in \mathbb{Z}_{p}$. Then we use Kaltofen's modification [13] of the Ben-Or/Tiwari sparse interpolation [3] to interpolate each coefficient $h_{i}\left(x_{1}, \ldots, x_{n}\right)$ of $H$, in parallel, from the coefficients of the scaled images $\Gamma\left(\beta_{j}\right) \times g_{j}\left(x_{0}\right)$.

Let $t=\max \# h_{i}$ be the maximum number of terms of the coefficients of $H$. Let $d=$ $\max _{i=1}^{n} \operatorname{deg}_{x_{i}} H$ and let $D=\max _{i=0}^{d G} \operatorname{deg} h_{i}$. The cost of sparse interpolation algorithms is determined mainly by the number of points $\beta_{1}, \beta_{2}, \ldots$ needed which depends on $t, d$ and $D$. The cost also depends on the size of the prime $p$. Table 1 below presents data for several sparse interpolation algorithms.

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

|  | \#points | size of $p$ | for the benchmark problem |
| :--- | :--- | :--- | :--- |
| Zippel [1979] | $O(n d t)$ | $p>2 n d^{2} t^{2}$ | $=6.4 \times 10^{9}$. |
| BenOr/Tiwari [1988] | $O(t)$ | $p>p_{n}^{D}$ | $=19^{60}=5.3 \times 10^{77}$. |
| Monagan/Javadi [2010] | $O(n t)$ | $p>n D t^{2}$ | $=4.8 \times 10^{8}$. |
| Discrete Logs | $O(t)$ | $p>(d+1)^{n}=3.7 \times 10^{10}$. |  |

Table 1: Data for different sparse interpolation algorithms

Notes: the figure $O(n d t)$ for Zippel's algorithm is for the worst case. The average case (for random inputs) is $O(d t)$ points. Also, Kaltofen and Lee showed in [12] how to modify Zippel's algorithm so that it will work for primes much smaller than $2 n d^{2} t^{2}$.

The primary disadvantage of the Ben-Or/Tiwari algorithm is the size of the prime. In [10] Javadi and Monagan modify the Ben-Or/Tiwari algorithm to work for a prime of size $O\left(n D t^{2}\right)$ but using $O(n t)$ points.

The discrete logs method, first proposed by Murao and Fujise [16], is a modification of the Ben-Or/Tiwari algorithm which computes discrete logarithms in the cyclic group $\mathbb{Z}_{p}^{*}$. We use this method. We give details for it in Section 1.2. The advantage over the Ben-Or/Tiwari algorithm is that the prime size is $O(n \log d)$ bits instead of $O(D \log n)$ bits.

In the GCD algorithm, not all points $\beta_{j} \in \mathbb{Z}_{p}^{n}$ can be used. If $\operatorname{gcd}\left(\bar{A}\left(x_{0}, \beta_{j}\right), \bar{B}\left(x_{0}, \beta_{j}\right)\right) \neq 1$ then $\beta_{j}$ is said to be unlucky and the image $g_{j}\left(x_{0}\right)$ cannot be used to interpolate $H$. In Zippel's algorithm, where the $\beta_{j}$ are chosen at random from $\mathbb{Z}_{p}^{n}$, unlucky images, once identified, can simply be skipped. But this is not the case for the point sequence $\left(2^{j}, 3^{j}, \ldots, p_{n}^{j}\right)$ used by the Ben-Or/Tiwari algorithm and the point sequence $\left(\omega_{1}^{j}, \ldots, \omega_{n}^{j}\right)$ used by the discrete logs method, because these points are not random.

In Section 2, we modify the prime power sequence $\left(2^{j}, 3^{j}, \ldots, p_{n}^{j}\right)$ to avoid and handle unlucky evaluation points. Our modification for the discrete logarithm sequence increases the size of $p$ which negates some of its advantage over the prime power sequence. This led us to consider using a Kronecker substitution on $x_{1}, x_{2}, \ldots, x_{n}$ to map the GCD computation into a bivariate computation in $\mathbb{Z}_{p}\left[x_{0}, y\right]$. Some Kronecker substitutions result in all evaluation
points being unlucky so they cannot be used. We call these Kronecker substitutions unlucky. Our second contribution is to show that are only finitely many of them, and how to detect them quickly so that a larger Kronecker substitution 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 in $y$ is exponential in $n$. This prompted us to investigate the distribution of the unlucky evaluation points. Our third contribution is a result on the mean and variance of the number of unlucky evaluations.

In Section 3 we assemble a Monte-Carlo GCD algorithm which chooses $p$ and computes $H \bmod p$. We have implemented our algorithm in C and parallelized it using Cilk C. We did this initially for 31 bit primes then for 63 bit primes to handle more polynomials. The first timing results revealed that almost all the time (over 95\%) was spent in evaluating $A\left(x_{0}, \beta_{j}\right)$ and $B\left(x_{0}, \beta_{j}\right)$. Our fifth contribution is an algorithm that reduces the evaluation cost to $s=\# A+\# B$ multiplications per $\beta_{j}$, thus one multiplication per term. Furthermore, since most of the time is still spent in evaluation, we show how to parallelize the evaluations.

In Section 4 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 62,520 seconds, Magma dies with an internal error, and our new algorithm takes 5.9 seconds on 16 cores. We have also made a 127 bit prime implementation to be able to interpolate polynomials of higher degree and/or in more variables.

A remaining problem is the extra cost incurred when $\# \Delta>1$ since this likely increases $t$. We discuss some things we can do in the Conclusion.

The proofs in the paper make use of the Schwartz-Zippel Lemma and properties of the Sylvester resultant. We state these results here for later use.

The Sylvester resultant of two polynomials $A$ and $B$ in $x$, denoted $\operatorname{res}_{x}(A, B)$, is the determinant of Sylvester's matrix. We gather the following facts about it into Lemma 1 below. Note, in the Lemma $\operatorname{deg} A$ denotes the total degree of $A$.

Lemma 1 Let $F$ be a field and let $A$ and $B$ be polynomials in $F\left[x_{0}, x_{1}, \ldots, x_{n}\right]$ with positive degree in $x_{0}$. Let $R=\operatorname{res}_{x_{0}}(A, B)$. Then
(i) $R$ is a polynomial in $F\left[x_{1}, \ldots, x_{n}\right] \quad$ ( $x_{0}$ is eliminated),
(ii) $\operatorname{deg} R \leq \operatorname{deg} A \operatorname{deg} B \quad$ (Bezout bound).

If $\alpha \in F^{n}$ satisfies $\operatorname{deg}_{x_{0}} A\left(x_{0}, \alpha\right)=\operatorname{deg}_{x_{0}}(A)$ and $\operatorname{deg}_{x_{0}} B\left(x_{0}, \alpha\right)=\operatorname{deg}_{x_{0}}(B)$ then
(iii) $\operatorname{gcd}\left(A\left(x_{0}, \alpha\right), B\left(x_{0}, \alpha\right)\right) \neq 1 \Longleftrightarrow \operatorname{res}_{x_{0}}\left(A\left(x_{0}, \alpha\right), B\left(x_{0}, \alpha\right)\right)=0$ and
(iv) $\operatorname{res}_{x_{0}}\left(A\left(x_{0}, \alpha\right), B\left(x_{0}, \alpha\right)\right)=R(\alpha)$.

Proofs may be found in Ch. 3 and Ch. 6 of [5]. In particular the proof in Ch. 6 of [5] for (ii) for bivariate polynomials generalizes to the multivariate case. Note that the degree condition on $\alpha$ means that the dimension of Sylvester's matrix for $A$ and $B$ in $x_{0}$ is the same as for $A\left(x_{0}, \alpha\right)$ and $B\left(x_{0}, \alpha\right)$ which proves (iv).

Lemma 2 (Schwarz-Zippel [20, 22]). Let $F$ be a field and $f \in F\left[x_{1}, x_{2}, \ldots, x_{n}\right]$ be non-zero with total degree $d$ and let $S \subset F$. If $\beta$ is chosen at random from $S^{n}$ then $\operatorname{Prob}[f(\beta)=0] \leq$ $\frac{d}{|S|}$. In particular, if $F=\mathbb{Z}_{p}$ and $S=\mathbb{Z}_{p}$ then $\operatorname{Prob}[f(\beta)=0] \leq \frac{d}{p}$.

### 1.1 Ben-Or Tiwari Sparse Interpolation

Let $C\left(x_{1}, \ldots, x_{n}\right)=\sum_{i=1}^{t} a_{i} M_{i}$ where $a_{i} \in \mathbb{Z}$ and $M_{i}$ are monomials in $\left(x_{1}, \ldots, x_{n}\right)$. In our context, $C$ represents one of the coefficients of $H=\Delta G$ we wish to interpolate. Let $D=\operatorname{deg} C$ and let $d=\max _{i} \operatorname{deg}_{x_{i}} C$ and let $p_{n}$ denote the $n$ 'th prime. Let

$$
v_{j}=C\left(2^{j}, 3^{j}, 5^{j}, \ldots, p_{n}^{j}\right) \text { for } j=0,1, \ldots, 2 t-1
$$

The Ben-Or/Tiwari sparse interpolation algorithm [3] interpolates $C\left(x_{1}, x_{2}, \ldots, x_{n}\right)$ from the $2 t$ points $v_{j}$. In practice, $t$ is not known in advance so the algorithm needs to be modified to also determine $t$. We will discuss this later. Let $m_{i}=M_{i}\left(2,3,5, \ldots, p_{n}\right) \in \mathbb{Z}$ and let $\lambda(z)=\prod_{i=1}^{t}\left(z-m_{i}\right) \in \mathbb{Z}[z]$. The algorithm proceeds in 4 steps.

Step 1 Compute $\lambda(z)$ from $v_{j}$ using either the Berlekamp-Massey algorithm [14] or the Euclidean algorithm [2, 21].

Step 2 Compute the integer roots $m_{i}$ of $\lambda(z)$.
Step 3 Factor the integers $m_{i}$ using trial division by $2,3, \ldots, p_{n}$ from which we obtain $M_{i}$. For example, if $m_{i}=45000=2^{3} 3^{2} 5^{4}$ then $M_{i}=x_{1}{ }^{3} x_{2}{ }^{2} x_{3}{ }^{4}$.

Step 4 Solve the $t \times t$ linear system

$$
\left.\begin{array}{cccc}
{\left[\begin{array}{cccc}
1 & 1 & \ldots & 1 \\
m_{1} & m_{2} & \ldots & m_{t} \\
m_{1}{ }^{2} & m_{2}{ }^{2} & \ldots & m_{t}{ }^{2} \\
\vdots & \vdots & \vdots & \vdots \\
m_{1}{ }^{t-1} & m_{2}{ }^{t-1} & \ldots & m_{t}^{t-1}
\end{array}\right]} & {\left[\begin{array}{c}
a_{1} \\
a_{2} \\
a_{3} \\
\vdots \\
a_{t}
\end{array}\right]} & =  \tag{1}\\
\mathbf{V} & & = & \mathbf{b}
\end{array}\right]
$$

for the unknown coefficients $a_{i}$ in $C\left(x_{1}, \ldots, x_{n}\right)$. The matrix $V$ in (1) is a transposed Vandermonde matrix. The linear system $V a=b$ can be solved in $O\left(t^{2}\right)$ arithmetic operations (see [23]). Note, the master polynomial $P(Z)$ in [23] is $\lambda(z)$.

Notice that the largest integer in $\lambda(z)$ is the constant coefficient $\Pi_{i=1}^{t} m_{i}$ which is of size $O(\operatorname{tn} \log D)$ bits. Moreover, in [11], Kaltofen, Lakshman and Wiley noticed a severe expression swell occurs if either the Berlekamp-Massey algorithm or the Euclidean algorithm are used to compute $\lambda(z)$ over $\mathbb{Q}$. For our purposes, because we want to interpolate $H$ modulo a prime $p$, we will run steps 1,2 , and 4 modulo a prime $p$. Provided $p>\max _{i=1}^{t} m_{i} \leq p_{n}^{D}$ then the integers $m_{i}$ mod $p$ remain unique. The roots of $\lambda(z) \in \mathbb{Z}_{p}[z]$ can be found using Rabin's algorithm [19, 7] which has classical complexity $O\left(t^{2} \log p\right)$.

Steps 1, 2, and 4 may be accelerated with fast multiplication. Let $M(t)$ denote the cost of multiplying two polynomials of degree $t$ in $\mathbb{Z}_{p}[t]$. The fast Euclidean algorithm can be used to accelerate Step 1. It has complexity $O(M(t) \log t)$. See Ch. 11 of [7]. Computing the roots of $\lambda(z)$ in Step 2 can be done in $O(M(t) \log t \log p)$. See Ch 14 of [7]. Step 4 may be done in $O(M(t) \log t)$ using fast interpolation. See Ch 10 of [7].

### 1.2 Ben-Or/Tiwari using discrete logarithms in $\mathbb{Z}_{p}$

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 $O(D \log n)$. Murao and Fujise [16] were the first to use this method. We explain how as follows.

First we pick a prime $p$ of the form $p=q_{1} q_{2} q_{3} \ldots q_{n}+1$ with $q_{1}$ even, $\operatorname{gcd}\left(q_{i}, q_{j}\right)=1$ and $q_{i}>\operatorname{deg}_{x_{i}} C$ where $C\left(x_{1}, \ldots, x_{n}\right)$ is the polynomial we want to interpolate. 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 $p-1=q_{1} q_{2} \ldots q_{n}$. We set $\omega_{i}=\alpha^{(p-1) / q_{i}}$ so that $\omega_{i}^{q_{i}}=1$. In the Ben-Or/Tiwari algorithm we replace the points $\left(2^{j}, 3^{j}, \ldots, p_{n}^{j}\right)$ with $\left(\omega_{1}^{j}, \omega_{2}^{j}, \ldots, \omega_{n}^{j}\right)$. After Step 1 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 determine the degrees $d_{k}$ in Step 3 we first compute the discrete logarithm $x:=\log _{\alpha} m_{i}$, that is, solve $\alpha^{x} \equiv m_{i} \quad(\bmod p)$ for $0 \leq x<p-1$. We have

$$
\begin{align*}
x & =\log _{\alpha} m_{i}=\log _{\alpha} \prod_{k=1}^{n} \omega_{k}^{d_{k}}=\sum_{k=1}^{n} d_{k} \log _{\alpha} \omega_{k}=\sum_{k=1}^{n} d_{k} \frac{p-1}{q_{k}}  \tag{2}\\
& =d_{1}\left(q_{2} q_{3} \ldots q_{n}\right)+d_{2}\left(q_{1} q_{3} \ldots q_{n}\right)+\ldots+d_{n}\left(q_{1} q_{2} \ldots q_{n-1}\right) \tag{3}
\end{align*}
$$

We now solve $(3) \bmod q_{k}$ for $d_{k}$ for $1 \leq k \leq n$ to determine $M_{i}$. We obtain $d_{k}=x[(p-$ 1) $\left./ q_{k}\right]^{-1} \bmod q_{k}$. Note the condition $\operatorname{gcd}\left(q_{i}, q_{j}\right)=1$ ensures $\left(q_{1} \ldots q_{k-1} q_{k+1} \ldots q_{n}\right)$ is invertible $\bmod q_{k}$. Step 4 remains unchanged.
For $p=q_{1} q_{2} \ldots q_{n}+1$, a discrete logarithm can be computed in $O\left(\sum_{i=1}^{n} \sqrt{q}_{i}\right)$ multiplications in $\mathbb{Z}_{p}$ using the Pohlig-Helman algorithm [18]. Since the $q_{i} \sim d$ this leads to an $O(n \sqrt{d})$ cost. Kaltofen showed in [13] 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.

The algorithm can be modified to determine $t$ and hence output $\lambda(z)$ with high probability as follows. If we attempt to compute $\lambda(z)$ after $j=2,4,6, \ldots$, points, we will see $\operatorname{deg} \lambda(z)=$ $1,2,3, \ldots, t-1, t, t, t, \ldots$ with high probability. Thus we simply wait until the degree of $\lambda(z)$ does not change. This is first discussed by Kaltofen, Lee and Lobo in [12].

### 1.3 Bad and Unlucky Evaluation Points

Let $A$ and $B$ be non constant polynomials in $\mathbb{Z}\left[x_{1}, \ldots, x_{n}\right]\left[x_{0}\right]$ with $G=\operatorname{gcd}(A, B)$ and let $\bar{A}=A / G$ and $\bar{B}=B / G$. Let $p$ be prime such that $L C(A) L C(B) \bmod p \neq 0$.

Definition 1 Let $\alpha \in \mathbb{Z}_{p}^{n}$ and let $\bar{g}_{\alpha}(x)=\operatorname{gcd}(\bar{A}(x, \alpha), \bar{B}(x, \alpha))$. We say $\alpha$ is bad if $L C(A)(\alpha)=0$ or $L C(B)(\alpha)=0$ and $\alpha$ is unlucky if $\operatorname{deg} \bar{g}_{\alpha}(x)>0$.

Example 2 Let $G=\left(x_{1}-16\right) x_{0}+1, \bar{A}=x_{0}^{2}+\left(x_{1}-1\right)\left(x_{2}-9\right) x_{0}+1$ and $\bar{B}=x_{0}^{2}+1$. Then $L C(A)=L C(B)=x_{1}-16$ so $\left\{(16, \beta): \beta \in \mathbb{Z}_{p}\right\}$ are bad and $\left\{(1, \beta): \beta \in \mathbb{Z}_{p}\right\}$ and $\left\{(\beta, 9): \beta \in \mathbb{Z}_{p}\right\}$ are unlucky.

The algorithm cannot reconstruct $G$ using $g_{\alpha}=\operatorname{gcd}(A(x, \alpha), B(x, \alpha))$ if $\alpha$ is unlucky. Brown's idea in [4] to detect unlucky $\alpha$ is based on the following Lemma.

Lemma 3 Let $h_{\alpha}=G\left(x_{0}, \alpha\right) \bmod p$. If $\alpha$ is not bad then $h_{\alpha} \mid g_{\alpha}$ and $\operatorname{deg}_{x_{0}} g_{\alpha} \geq \operatorname{deg}_{x_{0}} G$.
For a proof of Lemma 3 see Lemma 7.3 of [8]. Brown only uses $\alpha$ which are not bad and the images $g_{\alpha}$ of least degree in $x_{0}$ to interpolate $G$. The following Lemma implies if the prime $p$ is large then unlucky evaluations points are rare.

Lemma $4 \operatorname{Prob}[\alpha$ is bad or unlucky $] \leq \frac{\operatorname{deg} A \operatorname{deg} B+\operatorname{deg} A+\operatorname{deg} B}{p-\operatorname{deg} A-\operatorname{deg} B}$.
Proof: $\operatorname{Prob}[\alpha$ is bad $]=\operatorname{Prob}[\operatorname{LC}(A)(\alpha) \mathrm{LC}(B)(\alpha)=0] \leq \frac{\operatorname{deg}(\operatorname{LC}(A))}{p}+\frac{\operatorname{deg}(\operatorname{LC}(B))}{p} \leq$ $\frac{\operatorname{deg} A+\operatorname{deg} B}{p}$. To determine $\operatorname{Prob}[\alpha$ is unlucky $\mid \alpha$ is not bad] we have

$$
\begin{aligned}
& \alpha \text { is unlucky } \Longleftrightarrow \\
& \Longleftrightarrow \\
& \Longleftrightarrow \\
& \Longleftrightarrow \\
& \operatorname{res}_{x}(\bar{A}(x, \alpha), \bar{B}(x, \alpha)) \neq 1 \text { (by definition) } \\
& \Longleftrightarrow R(\alpha)=0 \text { where } R=\operatorname{res}_{x_{0}}(\bar{A}, \bar{B})(\alpha \text { is not bad }) .
\end{aligned}
$$

Hence $\operatorname{Prob}[\alpha$ is unlucky $\mid \alpha$ is not bad $] \leq \frac{\operatorname{deg} R}{p-\operatorname{deg} A-\operatorname{deg} B}$ (by Schwartz-Zippel). Now the $\operatorname{Prob}[\alpha$ is bad or unlucky $] \leq \frac{\operatorname{deg} A+\operatorname{deg} B}{p}+\frac{\operatorname{deg} R}{p-\operatorname{deg} A-\operatorname{deg} B} \leq \frac{\operatorname{deg} R+\operatorname{deg} A+\operatorname{deg} B}{p-\operatorname{deg} A-\operatorname{deg} B} \leq \frac{\operatorname{deg} A \operatorname{deg} B+\operatorname{deg} A+\operatorname{deg} B}{p-\operatorname{deg} A-\operatorname{deg} B}$ by Lemma 1.

The following algorithm applies Lemma 3 to compute a lower bound $d$ for $\operatorname{deg}_{x_{i}} G$. Note, later in the paper when we use Algorithm DegBound, if it happens that $d>\operatorname{deg}_{x_{i}} G$ ( $\alpha$ is unlucky) then this won't affect the correctness of our algorithm, only the efficiency.

Algorithm DegreeBound $(A, B, i)$
Input: $A, B \in \mathbb{Z}\left[x_{0}, x_{1}, \ldots, x_{n}\right]$ with $\operatorname{deg} A>0$ and $\operatorname{deg} B>0$.
Output: $d \geq \operatorname{deg}_{x_{i}}(G)$ where $G=\operatorname{gcd}(A, B)$.
1 Set $L A=\mathrm{LC}\left(A, x_{i}\right)$ and $L B=\mathrm{LC}\left(B, x_{i}\right)$. So $L A, L B \in \mathbb{Z}\left[x_{0}, \ldots, x_{i-1}, x_{i+1}, \ldots, x_{n}\right]$.
2 Pick a prime $p \gg \operatorname{deg} A \operatorname{deg} B$ such that $L A \bmod p \neq 0$ and $L B \bmod p \neq 0$.
3 Pick $\alpha=\left(\alpha_{0}, \ldots, \alpha_{i-1}, \alpha_{i+1}, \ldots, \alpha_{n}\right) \in \mathbb{Z}_{p}^{n}$ at random until $L A(\alpha) L B(\alpha) \neq 0$.
4 Compute $a=A\left(\alpha_{0}, \ldots, \alpha_{i-1}, x_{i}, \alpha_{i+1}, \ldots, \alpha_{n}\right)$ and $b=B\left(\alpha_{0}, \ldots, \alpha_{i-1}, x_{i}, \alpha_{i+1}, \ldots, \alpha_{n}\right)$.
5 Compute $g=\operatorname{gcd}(a, b)$ in $\mathbb{Z}_{p}\left[x_{i}\right]$ using the Euclidean algorithm and output $d=\operatorname{deg}_{x_{i}} g$.

## 2 Unlucky evaluations and Kronecker substitutions

Consider again Example 2 where $G=\left(x_{1}-16\right) x_{0}+1, \bar{A}=x_{0}^{2}+\left(x_{1}-1\right)\left(x_{2}-9\right) x_{0}+1$ and $\bar{B}=x_{0}^{2}+1$. For the Ben-Or/Tiwari points $\alpha_{j}=\left(2^{j}, 3^{j}\right)$ for $0 \leq j<2 t$ 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 the first $s>0$ such that $2^{s}>p$ so that $\left(2^{s}, 3^{s}, \ldots, p_{n}^{s}\right) \bmod p$ is not fixed and use $\alpha_{j}=\left(2^{j}, 3^{j}, \ldots, p_{n}^{j}\right)$ for $s \leq j<2 t+s$. Step 1 works as before so we obtain $\lambda(z)$ from which we determine the $m_{i}$. To solve the shifted transposed Vandermonde system

$$
\begin{array}{cccc}
{\left[\begin{array}{cccc}
m_{1}^{s} & m_{2}^{s} & \ldots & m_{t}^{s} \\
m_{1}^{s+1} & m_{2}^{s+1} & \ldots & m_{t}{ }^{++1} \\
\vdots & \vdots & \vdots & \vdots \\
m_{1}{ }^{s+t-1} & m_{2}{ }^{s+t-1} & \ldots & m_{t}^{s+t-1}
\end{array}\right]} & {\left[\begin{array}{c}
c_{1} \\
c_{2} \\
\vdots \\
c_{t}
\end{array}\right]} & =\left[\begin{array}{c}
v_{s} \\
v_{s+1} \\
\vdots \\
v_{s+t-1}
\end{array}\right]
\end{array}
$$

we first solve the transposed Vandermonde system
as before to obtain $b=V^{-1} u$. Observe that the matrix $W=D V$ where $D$ is the $t$ by $t$ diagonal matrix with $D_{i, i}=m_{i}^{s}$. To solve $W c=u$ we have

$$
(D V) c=u \Longrightarrow c=\left(V^{-1} D^{-1}\right) u=D^{-1}\left(V^{-1} u\right)=D^{-1} b
$$

Thus $c_{i}=u m_{i}^{-s}$ and we can solve $W c=u$ in $O\left(t^{2}+t \log s\right)$ multiplications.
Referring again to Example 2, if we use the discrete logarithm evaluation points $\alpha_{j}=$ $\left(\omega_{1}^{j}, \omega_{2}^{j}\right)$ for $0 \leq j<2 t$ then $\alpha_{0}=(1,1)$ is unlucky and also, since $\omega_{1}^{q_{1}}=1$, all points $\alpha_{q_{1}}, \alpha_{2 q_{1}}, \alpha_{3 q_{1}}, \ldots$ are unlucky. The obvious fix is to shift the sequence to start at $j=1$ and pick $q_{i}>2 t$ is problematic because we don't know $t$ at this point. The following fix will work. Starting at $j=1$ if an unlucky evaluation point is encountered at $j=q_{i}$ for some $i$, then it is probably because $x_{i}=1$ is unlucky so let us restart the algorithm with a new prime $p=q_{1} q_{2} \ldots, q_{n}+1$ but with $q_{i}$ doubled in length. Repeating this if necessary, eventually this will terminate. But, because of the increase in the length of $p$ we were led to consider using a Kronecker substitution for the GCD problem.

## 3 Kronecker Substitutions

We propose to use a Kronecker substitution to map a multivariate polynomial GCD problem in $\mathbb{Z}\left[x_{0}, x_{1}, \ldots, x_{n}\right]$ into a bivariate GCD problem in $\mathbb{Z}[x, y]$. After making the Kronecker substitution, we need to interpolate $H(x, y)=\Delta(x, y) G(x, y)$ where $\operatorname{deg}_{y} H(x, y)$ will be exponential in $n$. To make discrete logarithms in $\mathbb{Z}_{p}$ feasible, we follow Kaltofen [13] and pick $p=2^{k} s+1>\operatorname{deg}_{y} H(x, y)$ with $s$ small.

Definition 2 Let $D$ be an integral domain and let $f$ be a non-zero polynomial in $D\left[x_{0}, x_{1}, \ldots, x_{n}\right]$. Let $r \in \mathbb{Z}^{n-1}$ with $r_{i}>0$. Let $K_{r}: D\left[x_{0}, x_{1}, \ldots, x_{n}\right] \rightarrow D[x, y]$ be the Kronecker substitution $K_{r}(f)=f\left(x, y, y^{r_{1}}, y^{r_{1} r_{2}}, \ldots, y^{r_{1} r_{2} \ldots r_{n-1}}\right)$.

Let $d_{i}=\operatorname{deg}_{x_{i}} f$ be the partial degrees of $f$ for $1 \leq i \leq n$. Observe that $K_{r}$ is invertible if $r_{i}>d_{i}$ for $1 \leq i \leq n-1$. Not all such Kronecker substitutions can be used, however, for the GCD problem. We consider an example.

Example 3 Consider the following GCD problem

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

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

We call such a Kronecker substitution unlucky. 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.

Definition 3 Let $K_{r}$ be a Kronecker substitution. We say $K_{r}$ is bad if $\operatorname{deg}_{x} K_{r}(A)<\operatorname{deg}_{x_{0}} A$ or $\operatorname{deg}_{x} K_{r}(B)<\operatorname{deg}_{x_{0}} B$ and $K_{r}$ is unlucky if $\operatorname{deg}_{x} \operatorname{gcd}\left(K_{r}(\bar{A}), K_{r}(\bar{B})\right)>0$.

Lemma 5 Let $f \in \mathbb{Z}\left[x_{1}, \ldots, x_{n}\right]$ be non-zero and $d_{i} \geq 0$ for $1 \leq i \leq n$. Let $X$ be the number of Kronecker substitutions from the sequence $r_{k}=\left[d_{1}+k, d_{2}+k, \ldots, d_{n-1}+k\right]$ for $k=$ $1,2,3, \ldots$ for which $K_{r}(f)=0$. Then $X \leq(n-1) \sqrt{2 \operatorname{deg} f}$.

Proof: $K_{r}(f)=0 \Longleftrightarrow f\left(y, y^{r_{1}}, y^{r_{1} r_{2}}, \ldots, y^{r_{1} r_{2} \ldots r_{n-1}}\right)=0$
$\Longleftrightarrow f \bmod \left\langle x_{1}-y, x_{2}-y^{r_{1}}, \ldots, x_{n}-y^{r_{1} r_{2} \ldots r_{n-1}}\right\rangle=0$
$\Longleftrightarrow f \bmod \left\langle x_{2}-x_{1}^{r_{1}}, x_{3}-x_{2}^{r_{2}}, \ldots, x_{n}-x_{n-1}^{r_{n-1}}\right\rangle=0$.
Thus $X$ is the number ideals $I=\left\langle x_{2}-x_{1}^{r_{1}}, \ldots, x_{n}-x_{n-1}^{r_{n-1}}\right\rangle$ for which $f \bmod I=0$ with $r_{i}=d_{i}+1, d_{i}+2, \ldots$ We prove that $X \leq(n-1) \sqrt{2 \operatorname{deg} f}$ 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 $f\left(x_{1}, x_{2}\right) \bmod \left\langle x_{2}-x_{1}^{r_{1}}\right\rangle=0 \Longrightarrow x_{2}-x_{1}^{r_{1}} \mid f$. Now $X$ is maximal when $d_{1}=0$ and $r_{1}=1,2,3, \ldots$ We have

$$
\sum_{r_{1}=1}^{X} r_{1} \leq \operatorname{deg} f \Longrightarrow X(X+1) / 2 \leq \operatorname{deg} f \Longrightarrow X<\sqrt{2 \operatorname{deg} f}
$$

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

$$
S \bmod I=0 \Longleftrightarrow S \bmod \left\langle x_{2}-x_{1}^{r_{1}}, \ldots, x_{n-2}-x_{n-1}^{r_{n-2}}\right\rangle=0
$$

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

Theorem 1 Let $A, B \in \mathbb{Z}\left[x_{0}, x_{1}, \ldots, x_{n}\right]$ be non-zero, $G=\operatorname{gcd}(A, B), \bar{A}=A / G$ and $B=\bar{B} / G$. Let $d_{i} \geq \operatorname{deg}_{x_{i}} G$ and let $X$ count the number of bad and unlucky Kronecker substitutions $K_{r_{k}}$ from the sequence $r_{k}=\left[d_{1}+k, d_{2}+k, \ldots, d_{n-1}+k\right]$ for $k=1,2,3, \ldots$. Then

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

Proof Let $L A$ be the leading coefficient of $A$ and $L B$ the leading coefficient of $B$ in $x_{0}$. Then $K_{r}$ is bad $\Longleftrightarrow K_{r}(L A)=0$ or $K_{r}(L B)=0$. Applying the previous lemma we have the number of bad Kronecker substitutions is

$$
\leq(n-1) \sqrt{2 \operatorname{deg} L A}+(n-1) \sqrt{2 \operatorname{deg} L B} \leq(n-1)(\sqrt{2 \operatorname{deg} A}+\sqrt{2 \operatorname{deg} B})
$$

Now let $R=\operatorname{res}_{x_{0}}(\bar{A}, \bar{B})$. We will assume $K_{r}$ is not bad.

$$
\begin{aligned}
K_{r} \text { is unlucky } & \Longleftrightarrow \operatorname{deg}_{x}\left(\operatorname{gcd}\left(K_{r}(\bar{A}), K_{r}(\bar{B})\right)>0\right. \\
& \Longleftrightarrow \operatorname{res}_{x}\left(K_{r}(\bar{A}), K_{r}(\bar{B})\right)=0 \\
& \Longleftrightarrow K_{r}\left(\operatorname{res}_{x}(\bar{A}, \bar{B})\right)=0 \\
& \Longleftrightarrow K_{r}(R)=0 \quad\left(K_{r} \text { is not bad }\right) .
\end{aligned}
$$

By Lemma 5, the number of unlucky Kronecker substitutions $\leq 2(n-1) \sqrt{2 \operatorname{deg} R} \leq(n-$ 1) $\sqrt{2 \operatorname{deg} A \operatorname{deg} B}$ by Lemma 1. Adding the two contributions proves the theorem.

In algorithm PGCD below, we identify an unlucky substitution as follows. After computing the first two monic images $g_{1}(x)$ and $g_{2}(x)$ in step 9 if both $\operatorname{deg}_{x} g_{1}>d_{0}$ and $\operatorname{deg}_{x} g_{2}>d_{0}$ then with high probability $K_{r}$ is unlucky so we try the next Kronecker substitution $r=\left[r_{1}+1, r_{2}+1, \ldots, r_{n-1}+1\right]$.

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 $G C D$ problem

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

in $\mathbb{Z}[x, y, z, w]$. We have $\Gamma=w y$ and $\Delta=y$. For $K(f)=f\left(x, y, y^{3}, y^{9}\right)$ we have $\operatorname{gcd}(K(A), K(B))$ $=K(G) \operatorname{gcd}\left(y^{10} x+y^{3}, y^{4} x+y^{9}\right)=\left(y^{9} x^{2}+y^{4}\right) y^{3}=y^{7}\left(y^{5} x^{2}+1\right)$.

One must not try to compute $\operatorname{gcd}(K(A), K(B))$ because the degree of the content of $\operatorname{gcd}(K(A), K(B))\left(y^{7}\right.$ 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 $g_{j}=\operatorname{gcd}\left(K(A)\left(x, \alpha^{j}\right), K(B)\left(x, \alpha^{j}\right)\right)$ any content is divided out, and when we scale by $K(\Gamma)\left(\alpha^{j}\right)$ and interpolate $y$ in $K(H)$ using our sparse interpolation, we recover any content. We obtain $K(H)=K(\Delta) K(G)=y^{10} x^{2}+y^{5}$, then invert $K$ to obtain $H=(y w) x^{2}+\left(y^{2} z\right)$.

### 3.1 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 also be very large.

Example 5 Consider the following $G C D$ problem

$$
\begin{aligned}
G & =x_{0}+x_{1}^{d}+x_{2}^{d}+\ldots+x_{n}^{d}, \\
\bar{A} & =x_{0}+x_{1}+\ldots+x_{n-1}+x_{n}, \text { and } \\
\bar{B} & =x_{0}+x_{1}+\ldots+x_{n-1}+1 .
\end{aligned}
$$

Using $r=[d+1, d+1, \ldots, d+1]$ we need $p>(d+1)^{n}$. But $R=\operatorname{res}_{x_{0}}(\bar{A}, \bar{B})=1-x_{n}$ and $K_{r}(R)=1-y^{r_{1} r_{2} \ldots r_{n-1}}=1-y^{(d+1)^{n-1}}$ which means there could be as many as $(d+1)^{n-1}$ unlucky evaluation points, that is, one in $d+1$.

To guarantee that we avoid unlucky evaluation points with high probability we would need to pick $p \gg \operatorname{deg}_{y} K_{r}(R)$ which could be much larger than what is needed to interpolate $K_{r}(H)$. But this upper bound based on the resultant is a worst case. This lead us to investigate what the expected number of unlucky evaluation points is. 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 2 Let $\mathbb{F}_{q}$ be a finite field with $q$ elements and $f=x^{l}+\sum_{i=0}^{l-1}\left(\sum_{j=0}^{d_{i}} a_{i j} y^{j}\right) x^{i}$ and $g=x^{m}+\sum_{i=0}^{m-1}\left(\sum_{j=0}^{e_{i}} b_{i j} y^{j}\right) x^{i}$ with $l \geq 1, m \geq 1$, and $a_{i j}, b_{i j} \in \mathbb{F}_{q}$. Let $X=\mid\left\{\alpha \in \mathbb{F}_{q}\right.$ : $\operatorname{gcd}(f(x, \alpha), g(x, \alpha)) \neq 1\} \mid$ be a random variable over all choices $a_{i j}, b_{i j} \in \mathbb{F}_{q}$. So $0 \leq X \leq q$ and for $f$ and $g$ not coprime in $\mathbb{F}_{q}[x, y]$ we have $X=q$. Then
(i) if $d_{i} \geq 0$ and $e_{i} \geq 0$ then $\mathrm{E}[X]=1$ and
(ii) if $d_{i} \geq 1$ and $e_{i} \geq 1$ then $\operatorname{Var}[X]=1-1 / q$.

Proof (i): 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 $C_{\beta}\left(c_{0}, \ldots, c_{d}\right)=\sum_{i=0}^{d} c_{i} \beta^{i}$. 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 $\operatorname{deg} a>0$ and $\operatorname{deg} b>0$ are coprime with probability $1-1 / q$ (see Ch 11 of Mullen and Panario [17]). 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 $\beta$ we obtain

$$
E[X]=\sum_{\beta \in \mathbb{F}_{q}} \operatorname{Prob}[\operatorname{gcd}(A(x, \beta), B(x, \beta)) \neq 1]=q\left(1-\left(1-\frac{1}{q}\right)\right)=1
$$

Proof of claim. Since $B=\left\{1, y-\beta,(y-\beta)^{2}, \ldots,(y-\beta)^{d}\right\}$ 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 $u_{0}, u_{1}, \ldots, u_{d} \in \mathbb{F}_{q}$. Since $C(\beta)=u_{0}$ it follows that all $q^{d}$ choices for $u_{1}, \ldots, u_{d}$ result in $C(\beta)=u_{0}$ hence $C$ is balanced.

That $E[X]=1$ was a surprise to us. We thought $E[X]$ would have a logarithmic dependence on $\operatorname{deg} f$ and $\operatorname{deg} g$. In light of Theorem 2 , when picking $p>\operatorname{deg}_{y}\left(K_{r}(H)\right)$ we will ignore the unlucky evaluation points, and, should the algorithm encounter unlucky evaluations, restart the algorithm with a larger prime.

## 4 GCD Algorithm

## Algorithm $\operatorname{PGCD}(A, B, \Gamma)$

Input $A=a_{l} x_{0}^{l}+\ldots+a_{0}$, and $B=b_{m} x_{0}^{m}+\ldots+b_{0}$ with $a_{i}, b_{i} \in \mathbb{Z}\left[x_{1}, \ldots x_{n}\right]$ and $\Gamma \in$ $\mathbb{Z}\left[x_{1}, \ldots, x_{n}\right]$ satisfying $\operatorname{gcd}\left(a_{i}\right)=1(A$ is primitive $)$ and $\operatorname{gcd}\left(b_{i}\right)=1(B$ is primitive $)$ and $\Gamma=\operatorname{gcd}\left(a_{l}, b_{m}\right)=L C(G) \times \Delta$ where $G=\operatorname{gcd}(A, B)$.
Output A prime $p$ and polynomial $H \in \mathbb{Z}_{p}\left[x_{0}, x_{1}, \ldots, x_{n}\right]$ satisfying $H=\Delta \times G \bmod p$ with probability at least $1-\frac{\operatorname{deg} A \operatorname{deg} B+\min (\operatorname{deg} A, \operatorname{deg} B)}{p-\operatorname{deg} A-\operatorname{deg} B}$.

Compute $d_{i}=\operatorname{DegBound}(A, B, i)$ for $0 \leq i \leq n-1$.
Initialize $r_{i}=1+\min \left(\operatorname{deg}_{x_{i}} A, \operatorname{deg}_{x_{i}} B, d_{i}+\operatorname{deg}_{x_{i}} \Gamma\right)$ for $1 \leq i \leq n-1$.
Kronecker-substitution:
1 Let $Y=\left(y, y^{r_{1}}, y^{r_{1} r_{2}}, \ldots, y^{r_{1} r_{2} \ldots r_{n-1}}\right)$ be the Kronecker substitution.
Set $K(A)=A(x, Y), K(B)=B(x, Y)$ and $K(\Gamma)=\Gamma(Y)$.
2 If $\operatorname{deg}_{x} K(A)<\operatorname{deg}_{x_{0}} A$ or $\operatorname{deg}_{x} K(B)<\operatorname{deg}_{x_{0}} B$ then this Kronecker substitution is bad. Set $r_{i}=r_{i}+1$ for $1 \leq i \leq n-1$ and goto Kronecker-substitution.

## Pick-a-Prime:

3 Pick a new prime $p>\prod_{i=1}^{n} r_{i}$ of the form $p=2^{k} s+1$ with $s$ small.
4 If $\operatorname{deg}_{x}(K(A) \bmod p)<\operatorname{deg}_{x} A$ or $\operatorname{deg}_{x}(K(B) \bmod p)<\operatorname{deg}_{x} B$ then the prime is bad so goto Pick-a-Prime.

5 Set shift $s=1$ and $j=0$ and compute a random generator $\alpha$ for $\mathbb{Z}_{p}^{*}$.

## Compute-next-image:

6 Set $j=j+1$. If $j=p-1$ then we've run out of evaluation points. This could happen if the number of terms of one of the coefficients of $\Delta G$ is dense. Goto Pick-a-Prime and increase the length of $p$ by 10 bits.

7 Compute $a_{i}=K(A)\left(x, \alpha^{j}\right) \bmod p$ and $b_{i}=K(B)\left(x, \alpha^{j}\right) \bmod p$. If $\operatorname{deg}_{x} a_{i}<\operatorname{deg}_{x_{0}} A$ or $\operatorname{deg}_{x} b_{i}<\operatorname{deg}_{x_{0}} B$ then $\alpha_{j}$ is bad so set $s=j$ and goto Compute-next-image.

8 Compute $g_{j}=\operatorname{gcd}\left(K(A)\left(x, \alpha^{j}\right), K(B)\left(x, \alpha^{j}\right)\right)$ using the Euclidean algorithm.
9 Case $\operatorname{deg} g_{j}<d_{0}$ : (the degree bound $d_{0}$ is wrong)
Set $d_{0}=\operatorname{deg} g_{j}, s=j$ and goto Compute-next-image.

10 Case $\operatorname{deg} g_{j}>d_{0}:\left(\alpha^{j}\right.$ is unlucky $)$
10a If this happens for $j=s$ and $j=s+1$ then the Kronecker substitution is very probably unlucky so set $r_{i}=r_{i}+1$ for $1 \leq i \leq n-1$ and goto Kroneckersubstitution.
10b If this is the 2nd unlucky evaluation then goto Pick-a-Prime and increase the length of $p$ by 10 bits. Otherwise set $s=j$ and goto Compute-next-image.

11 Case $\operatorname{deg} g_{j}=d_{0}$ : (we have a new image)
11a Scale the image: Set $g_{j}=K(\Gamma)\left(\alpha^{j}\right) g_{j} \bmod p$. If $s-j$ is even then goto compute-next-image - we need at least two new images for the next step.

11b Run the Berlekamp-Massey algorithm on the coefficients of the images $g_{s}, g_{s+1}, \ldots, g_{s+j}$ to obtain $\lambda_{i}(z)$ for $0 \leq i \leq d_{0}$.
11c If any $\lambda_{i}(z)$ changed from the previous step goto Compute-next-image.
11 d Compute the roots of each $\lambda_{i}(z)$. If any $\lambda_{i}(z)$ has fewer than $\operatorname{deg} \lambda_{i}(z)$ distinct roots goto Compute-next-image.
11e Complete the sparse interpolation to obtain polynomials $h_{i}(y) \in \mathbb{Z}_{p}[y]$.
Note, $s$ is the shift used for the shifted transposed Vandermonde systems. Set $H(x, y):=\sum_{i=0}^{d_{0}} h_{i}(y) x^{i}$ which we hope is equal to $\Delta(Y) G(x, Y)$.
11f Invert the Kronecker substitution to obtain $H\left(x_{0}, x_{1}, \ldots, x_{n}\right)$.
If $\operatorname{deg}_{x_{i}} H>\min \left(\operatorname{deg}_{x_{i}} A, \operatorname{deg}_{x_{i}} B, d_{i}+\operatorname{deg}_{x_{i}} \Gamma\right)$ for any $1 \leq i \leq n$ then $H \neq \Delta G$ so goto Compute-next-image.
11h Probabilistic check: Pick $\beta \in \mathbb{Z}_{p}^{n}$ at random until $\operatorname{deg} A\left(x_{0}, \beta\right)=\operatorname{deg}_{x_{0}} A$ and $\operatorname{deg} B\left(x_{0}, \beta\right)=\operatorname{deg}_{x_{0}} B$. Compute $g_{\beta}=\operatorname{gcd}\left(A\left(x_{0}, \beta\right), B\left(x_{0}, \beta\right)\right)$. If $H\left(x_{0}, \beta\right)=$ $\Gamma(\beta) g_{\beta}$ then output $(p, H)$. Otherwise either $t_{i}$ is wrong for some $i$ or $d_{0}>$ $\operatorname{deg}_{x_{0}} G$ or $\beta$ is unlucky. In all cases continue goto Compute-next-image.

To prove the claim on the output $(p, H)$ let $H=\sum_{i=0}^{d_{0}} h_{i} x_{0}^{i}$ and let $G=\sum_{i=0}^{d G} c_{i} x_{0}^{i}$. We will bound the probability that the algorithm outputs $H \neq \Delta G \bmod p$. Notice that if the algorithm outputs $H$ it must be that $\operatorname{deg}_{x_{0}} H=d_{0}=\operatorname{deg}_{x_{0}} g_{\beta}$. Now either $d_{0}>$ $d G$ or $d_{0}=d G$. If $d_{0}>d G$ then $H$ is wrong. Now $d_{0}>d G \Rightarrow \beta$ is unlucky thus $\operatorname{Prob}\left[d_{0}>d G\right] \leq \operatorname{Prob}[\beta$ is unlucky $]$ which is at most $\frac{\operatorname{deg} A \operatorname{deg} B}{p-\operatorname{deg} A-\operatorname{deg} B}$. If $d_{0}=d G$ then $H$ is output iff $h_{i}(\beta)=\Delta(\beta) c_{i}(\beta) \bmod p$ for $0 \leq i \leq d_{0}$. Let $f_{i}=h_{i}-\Delta c_{i} \bmod p$. $H \neq \Delta G$ implies $f_{i} \neq 0$ for at least one $i$, say $i=j$. The Schwartz-Zippel lemma implies $\operatorname{Prob}\left[f_{j}(\beta)=0\right] \leq \frac{\operatorname{deg} f_{i}}{p-\operatorname{deg} A-\operatorname{deg} B}$. Now the degree condition on $\operatorname{deg}_{x_{i}} H$ means the total degree $\operatorname{deg} f_{i} \leq \min (\operatorname{deg} A, \operatorname{deg} B)$ thus $\operatorname{Prob}\left[f_{j}(\beta)=0\right] \leq \frac{\min (\operatorname{deg} A, \operatorname{deg} B)}{p-\operatorname{deg} A-\operatorname{deg} B}$. Adding both probabilities $\operatorname{Prob}[H \neq \Delta G \bmod p] \leq \frac{\min (\operatorname{deg} A, \operatorname{deg} B)}{p-\operatorname{deg} A-\operatorname{deg} B}+\frac{\operatorname{deg} A \operatorname{deg} B}{p-\operatorname{deg} A-\operatorname{deg} B}$ and the result follows.

### 4.1 Determining t

Algorithm PGCD assumes in step 12b that if none of the $\lambda_{i}(z)$ changed then $(j-s+1) / 2=t$ but it could be that $(j-s+1) / 2<t$. Let $V_{r}=\left(v_{0}, v_{1}, \ldots, v_{2 r-1}\right)$ 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 PGCD, we determine the $t$ by computing $c(z)$ s on the input sequence $V_{r}$ for $r=1,2,3, \ldots$ 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 $[14,12]$ 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 [12], Kaltofen, Lee and Lobo proved (Theorem 3) that the BMA encounters the first zero discrepancy after $2 t$ points with probability at least

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

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

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

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 $f\left(\alpha^{j}\right)$ 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 | $69 z+1$ |
| 2 | $24 z^{2}+59 z+1$ |
| 3 | $24 z^{2}+59 z+1$ |
| 4 | $24 z^{2}+59 z+1$ |
| 5 | $70 z^{7}+42 z^{6}+6 z^{3}+64 z^{2}+34 z+1$ |
| 6 | $70 z^{7}+42 z^{6}+25 z^{5}+87 z^{4}+16 z^{3}+20 z^{2}+34 z+1$ |
| 7 | $z^{7}+67 z^{6}+95 z^{5}+2 z^{4}+16 z^{3}+20 z^{2}+34 z+1$ |
| 8 | $31 z^{8}+61 z^{7}+91 z^{6}+84 z^{5}+15 z^{4}+7 z^{3}+35 z^{2}+79 z+1$ |
| 9 | $31 z^{8}+61 z^{7}+91 z^{6}+84 z^{5}+15 z^{4}+7 z^{3}+35 z^{2}+79 z+1$ |

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, $\lambda(z)=z^{2} c(1 / z)=z^{2}+59 z+24$ 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. Furthermore, the recovered monomials may have degree higher than the degree of the input polynomial $f(y)$. Algorithm PGCD must check $H$ for monomials of too high degree in step 12e for the degree argument in the proof of the claim to be valid.

### 4.2 Evaluation

Let $A, B \in \mathbb{Z}_{p}\left[x_{0}, x_{1}, \ldots, x_{n}\right], s=\# A+\# B, d_{i}=\max _{i=1}^{n}\left(\operatorname{deg}_{x_{i}} A, \operatorname{deg}_{x_{i}} B\right)$ and $d=$ $\max _{i=1}^{n} d_{i}$. If we apply a Kronecker substitution and compute $K(A)=A\left(x, y, y^{r_{1}}, \ldots, y^{r_{1} r_{2} \ldots r_{n-1}}\right)$ with $r_{i}=d_{i}+1$, then $\operatorname{deg}_{y} K(A)<(d+1)^{n}$. To evaluate $K(A)(x, y)$ at $y=\alpha^{k}$ we can evaluate any monomial in $y$ in $K(A)$ in $O(n \log d)$ multiplications. Instead we first compute $\beta_{1}=\alpha^{k}$ and $\beta_{i}=\beta_{i-1}^{r_{i-1}}$ for $i=2,3, \ldots, n-1$. To compute $A\left(x, \beta_{1}, \ldots, \beta_{n}\right)$ and $B\left(x, \beta_{1}, \ldots, \beta_{n}\right)$ we first precompute $n$ tables of powers $1, \beta_{i}, \beta_{i}^{2}, \ldots, \beta_{i}^{d_{i}}$ for $1 \leq i \leq n$ using at most $n d$ multiplications. Now, for each term of $A$ and $B$ of the form $c x_{0}^{e_{0}} x_{1}^{e_{1}} \ldots x_{n}^{e_{n}}$ we can compute $c \times \beta_{1}^{e_{1}} \times \ldots \times \beta_{n}^{e_{n}}$ using the tables in $n$ multiplications. Hence we can evaluate $A\left(x, \alpha^{k}\right)$ and $B\left(x, \alpha^{k}\right)$ in at most $n d+n s$ multiplications in $\mathbb{Z}_{p}$. Thus for $T$ evaluation points $\alpha, \alpha^{2}, \ldots, \alpha^{T}$, the evaluation cost is $O(n T(d+s))$ multiplications.

When we first implemented algorithm PGCD we noticed that often over $95 \%$ of the time was spent evaluating the input polynomials $A$ and $B$ at the points $\alpha^{k}$. This happens when $\# G \ll \# A+\# B$. The following method uses the fact that for a monomial $M_{i}\left(x_{1}, x_{2}, \ldots, x_{n}\right)$

$$
M_{i}\left(\beta_{1}^{k}, \beta_{2}^{k}, \ldots, \beta_{n}^{k}\right)=M_{i}\left(\beta_{1}, \beta_{2}, \ldots, \beta_{n}\right)^{k}
$$

to reduce the total evaluation cost from $O(n T(s+d))$ multiplications in $\mathbb{Z}_{p}$ to $O(s T+n d+n s)$ thus one multiplication per term instead of $n$. Note, no sorting on $x_{0}$ is needed in step 4 b if the monomials in the input $A$ are are sorted in in lexicographical order with $x_{0}>x_{1}>\ldots>x_{n}$.

## Algorithm Evaluate.

Input Polynomial $A\left(x_{0}, x_{1}, \ldots, x_{n}\right)=\sum_{i=1}^{m} c_{i} x_{0}^{e_{i}} M_{i}\left(x_{1}, \ldots, x_{n}\right) \in \mathbb{Z}_{p}\left[x_{0}, \ldots, x_{n}\right]$
$T>0, \beta_{1}, \beta_{2}, \ldots, \beta_{n} \in \mathbb{Z}_{p}$, and integers $d_{1}, d_{2}, \ldots, d_{n}$ with $d_{i} \geq \operatorname{deg}_{x_{i}} A$.
Output $A\left(x_{0}, \beta_{1}^{k}, \ldots, \beta_{n}^{k}\right)$ for $1 \leq k \leq T$.
step 1 Create the vector $C=\left[c_{1}, c_{2}, \ldots, c_{m}\right] \in \mathbb{Z}_{p}^{m}$ of coefficients.
step 2 Compute tables of powers $\beta_{i}^{j}$ for $j=0,1, \ldots, d_{i}$ for $1 \leq i \leq n$.
step 3 Compute the vector $\Gamma=\left[M_{i}\left(\beta_{1}, \beta_{2}, \ldots, \beta_{n}\right)\right.$ for $\left.1 \leq i \leq m\right]$.
step 4 For $k=1,2, \ldots, T$ do
step 4a Compute the vector $C:=\left[C_{i} \times \Gamma_{i}\right.$ for $\left.1 \leq i \leq m\right]$.
step 4b Assemble $\sum_{i=1}^{m} C_{i} x_{0}^{e_{i}}=A\left(x_{0}, \beta_{1}^{k}, \ldots, \beta_{n}^{k}\right)$.
Algorithm Evaluate serializes evaluation on $k$. To parallelize it on $k$ for $N$ cores, we replace steps 3 and 4 with those below. To simplify notation, for vectors $u, v \in \mathbb{Z}_{p}^{m}$ let $u \otimes v=$ [ $u_{1} v_{1}, u_{2} v_{2}, \ldots, u_{m} v_{m}$ ] denote point-wise product.
step 3 Compute $\Gamma_{1}=\left[M_{i}\left(\beta_{1}, \beta_{2}, \ldots, \beta_{n}\right)\right.$ for $\left.1 \leq i \leq m\right]$ and $B_{1}=C \otimes \Gamma_{1}$.
step 4a For $k=2,3, \ldots, N$ do compute $\Gamma_{k}:=\Gamma_{1} \bigotimes \Gamma_{k-1}$ and $B_{k}:=\Gamma_{k} \bigotimes B_{k-1}$.
step 4b For $i=N, 2 N, 3 N, \ldots$ while $i<T$ do
For $k=1,2, \ldots, N$ while $i+k<T$ in parallel do
Compute $B_{k}:=B_{k} \otimes \Gamma_{N}$.
Assemble $\sum_{j=1}^{m} B_{k, j} x^{e_{j}} \in \mathbb{Z}_{p}[x]=A\left(x_{0}, \beta_{1}^{k}, \ldots, \beta_{n}^{k}\right)$.

Notice that once step 3 is completed, evaluation is reduced to a sequence of pointwise vector products. Each of these vector products could be parallelized in blocks of size $m / N$ for $N$ cores. In Cilk C, this is only effective, however, if the blocks are very large (at least $50,000)$ so that the work done in a block is much larger than the time it takes Cilk to create a task. For this reason, it is necessary to also parallelize on $k$.

If $N$ is large and $T$ is not much larger than $N$ step 4a will become a parallel bottleneck. We parallelized step 4a by first computing $\Gamma_{2}$ then computing $\Gamma_{3}$ and $\Gamma_{4}$ in parallel from $\Gamma_{1}$ and $\Gamma_{2}$, then computing $\Gamma_{5}, \Gamma_{6}, \Gamma_{7}, \Gamma_{8}$ in parallel from $\Gamma_{1}, \Gamma_{2}, \Gamma_{3}, \Gamma_{4}$, etc.

The space used is $N m$ words of memory for $\Gamma_{1}, \Gamma_{2}, \ldots, \Gamma_{N}$ and a further $N m$ words for $B_{1}, \ldots, B_{N}$. For inputs $A$ and $B$ both of size $10^{8}$ terms, on our 16 core machine ( $N=16$ ), we exceeded the soft memory limit of 32 gigabytes. We saved a factor of 2 in space by recycling the memory $\Gamma_{1}, \Gamma_{2}, \ldots, \Gamma_{N-1}$ for $B_{1}, B_{2}, \ldots, B_{N-1}$ in step 4 . We also decided to parallelize each pointwise multiplication $B_{k}=B_{k} \bigotimes \Gamma_{N}$ in two blocks of size $m / 2$ so we do $N / 2$ points at a time instead of $N$ in step 4 to save a further factor of 2 in space.

## 5 Benchmarks

We have implemented algorithm PGCD for 31, 63 and 127 bit primes in C. For 127 bit primes we use the 128 bit signed integer type __int128_t supported by the gcc compiler. We have parallelized parts of the algorithm using Cilk C. To assess how good it is, we have compared it with the implementations of Zippel's algorithm in Maple 2015 and Magma 2.21. For Maple we were able to determine the time spent computing $G$ modulo the first prime only in Zippel's algorithm. It is over $90 \%$ of the total GCD time. For Magma we could not do this so the Magma timings are for the entire GCD computation over $\mathbb{Z}$.

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 the maximum parallel speedup is a factor of $16 \times 2.2 / 3.0=11.7$.

For our first benchmark (see Table 2) 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 $100 d$ random terms for $G$ and 100 random terms for $\bar{A}$ and $\bar{B}$. The integer coefficients of $G, \bar{A}, \bar{B}$ were generated at random from $\left[0,2^{31}-1\right]$. 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 $1, x_{0}^{d}, x_{1}^{d}, \ldots, x_{6}^{d}$ so that $G$ is monic in all variables and $\Gamma=1$. Our GCD code uses the 62 bit prime $p=29 \times 2^{57}+1$. Maple uses the 32 bit prime $2^{32}-5$ for the first image in Zippel's algorithm. Magma is using 23.5 bit primes so that it can use floating point arithmetic.

In Table 2 column $d$ is the maximum degree of the terms of $G, A, B$ in each variable, column $t$ is the maximum number of terms of the coefficients of $G$ and column eval is the \%age of the time spent evaluating the inputs, that is computing $K(A)\left(x_{0}, \alpha^{j}\right)$ and $K(B)\left(x_{0}, \alpha^{j}\right)$ for $j=1,2, \ldots, T$ where $T$ is slightly more than $2 t$.

Our second benchmark (see Table 3) is for our 9 variable benchmark problem from Section 1 where the degree of $G, \bar{A}, \bar{B}$ is at most 20 in each variable. The terms are generated at random as before but restricted to have total degree at most 60 . The 62 bit prime is $29 \times 2^{59}+1$ and the 127 bit prime is $4085 \times 2^{115}+1$.

Tables 2 and 3 show that most of the time is in evaluation. They show a parallel speedup

|  |  | New GCD algorithm |  |  | Zippel's algorithm |  |  |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| $n$ | $d$ | $t$ | 1 core (eval) | 4 cores 16 cores | Maple | Magma |  |
| 5 | 10 | 114 | $0.66 \mathrm{~s}(72 \%)$ | 0.225 s | 0.096 s | 48.04 s | 6.97 s |
| 5 | 20 | 122 | $1.46 \mathrm{~s}(74 \%)$ | 0.473 s | 0.187 s | 185.70 s | 318.22 s |
| 5 | 50 | 121 | $3.38 \mathrm{~s}(72 \%)$ | 1.050 s | 0.397 s | 1525.80 s | $>10^{4} \mathrm{~s}$ |
| 5 | 100 | 102 | $7.01 \mathrm{~s}(73 \%)$ | 2.126 s | 0.760 s | 6018.23 s | $>10^{4} \mathrm{~s}$ |
| 5 | 200 | 111 | $15.03 \mathrm{~s}(74 \%)$ | 4.522 s | 1.629 s | NA | NA |
| 5 | 500 | 128 | $41.10 \mathrm{~s}(74 \%)$ | 12.389 s | 4.425 s | NA | NA |
| 8 | 5 | 89 | $0.35 \mathrm{~s}(70 \%)$ | 0.133 s | 0.065 s | 30.87 s | 2.39 s |
| 8 | 10 | 110 | $0.65 \mathrm{~s}(71 \%)$ | 0.221 s | 0.089 s | 138.41 s | 6.15 s |
| 8 | 20 | 114 | $1.44 \mathrm{~s}(72 \%)$ | 0.466 s | 0.176 s | 664.33 s | 63.49 s |
| 8 | 50 | 113 | $3.65 \mathrm{~s}(73 \%)$ | 1.102 s | 0.410 s | 6390.22 s | 1226.77 s |
| 8 | 100 | 121 | $7.32 \mathrm{~s}(73 \%)$ | 2.168 s | 0.792 s | NA | NA |
| 8 | 200 | need 127 bit primes |  |  |  |  | NA |

Table 2: Timings (seconds) for GCD problems.

|  | New (62 bit prime) |  | 126 bit prime | Maple | Magma |
| ---: | ---: | ---: | ---: | ---: | ---: |
| \#G \#A | 1 core (eval) 16 cores (eval) | 1 core (eval) |  |  |  |
| $10^{3} 10^{5}$ | $0.83 \mathrm{~s}(72 \%)$ | $0.137 \mathrm{~s}(61 \%)$ | $4.80 \mathrm{~s}(59 \%)$ | 341.9 s | 63.55 s |
| $10^{3}$ | $10^{6}$ | $6.73 \mathrm{~s}(88 \%)$ | $0.732 \mathrm{~s}(79 \%)$ | $32.22 \mathrm{~s}(93 \%)$ | 5553.5 s |
| $10^{4} 10^{6}$ | $61.23 \mathrm{~s}(91 \%)$ | $5.895 \mathrm{~s}(82 \%)$ | $248.23 \mathrm{~s}(70 \%)$ | 62520.1 s | FAIL |

Table 3: Timings (seconds) for large 9 variable problems (FAIL $=$ Internal error).
approaching a factor of 10 . This is good but not great. There is a parallel bottleneck in how we compute the $\lambda_{i}(z)$ polynomials. For $N=16$ cores, after generating a new batch of $N / 2=8$ images we run the Berlekamp-Massey algorithm to see if we have enough points to interpolate the next coefficient of $H$. The algorithm is serial and quadratic in $t$ which limits parallel speedup for $t=100$. We are not sure how to improve this.

In comparing the new algorithm with Maple's implementation of Zippel's algorithm, for $n=8, d=50$ in Table 2 we achieve a factor of $1748=6390.22 / 3.654$ speedup on 1 core. Since Zippel's algorithm uses $O(d t)$ points and our Ben-Or/Tiwari algorithm uses $2 t+O(1)$ points, we get a factor of $O(d)$ speedup because of this.

Our improved evaluation gives us a another factor of $n$ speedup over Maple's implementation of Zippel's algorithm. Another factor is the cost of multiplication in $\mathbb{Z}_{p}$. The reader should realize that the running time of algorithm PGCD is proportional to the cost of multiplication in $\mathbb{Z}_{p}$. Maple is using $\% \mathrm{p}$ to divide in C which generates a hardware division instruction which is much more expensive than a multiplication. We are using Roman Pearce's implementation of Möller and Granlund [15] which reduces division by $p$ to multiplications and other cheap operations.

## 6 Conclusion and Final Remarks

We have shown that a Kronecker substitution can be used to reduce a multivariate GCD computation to bivariate by using a discrete logs Ben-Or/Tiwari point sequence. Our parallel implementation is fast and practical. Several questions remain. The Ben-Or/Tiwari method requires $2 t+O(1)$ points. Can we use fewer points? Can we do anything when $\# \Delta \gg 1$ which increases $t$ ?

We note that if either $A$ or $B$ is monic in some $x_{i}$ then if we interchange $x_{0}$ with $x_{i}$ then $\Gamma=1$ hence $\Delta=1$. Similarly, if either $A$ or $B$ have a constant term, then we can reverse the coefficients of both $A$ and $B$ in $x_{0}$ so that $\Gamma=1$ and $\Delta=1$.

Our algorithm interpolates $H$ from univariate images in $\mathbb{Z}_{p}\left[x_{0}\right]$. If instead we interpolate $H$ from bivariate images in $\mathbb{Z}_{p}\left[x_{0}, x_{1}\right]$, this will likely reduce both $t$ and $\# \Delta$. For our benchmark problem this would reduce $t$ by a factor of 5 and the cost of the bivariate GCD computations in $\mathbb{Z}_{p}\left[x_{0}, x_{1}\right]$, if computed with Brown's dense GCD algorithm, would be negligible compared with the cost of evaluating $A$ and $B$. Although we have not implemented this we estimate a speedup of a factor of 3 on 16 cores.

For polynomials in more variables or higher degree algorithm PGCD may need primes larger than 127 bits. We cite the methods of Garg and Schost [6], Giesbrecht and Roche [9] and Arnold, Giesbrecht and Roche [1] which can use a smaller prime $p$ and would also use fewer than $2 t+O(1)$ evaluations. These methods would compute $a_{i}=K_{r}(A)(x, y), b_{i}=$ $K_{r}(B)(x, y)$ then $\left.g_{i}=\operatorname{gcd}\left(a_{i}, b_{i}\right)\right)$ all $\bmod \left\langle p, y^{q_{i}}-1\right\rangle$ for several primes $q_{i}$ and recover the exponents of $y$ using Chinese remaindering. The algorithms in $[6,9,1]$ differ in the size of $q_{i}$ and how they avoid or recover from exponent collisions. It is not clear whether this can be made to work for the GCD problem as these methods assume a division free evaluation. Computing $g_{i}$ requires division and $y=1$ may be bad or unlucky. It is also not clear whether these methods would be faster as they require $q_{i} \gg t$ which means computing $g_{i}$ will be expensive for large $t$.

## References

[1] A. Arnold, M. Giesbrecht and D. Roche. Faster Sparse Multivariate Polynomial Interpolation of Straight-Line Programs. arXiv:1412.4088[cs.SC], 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 Proceedings of STOC '20, ACM Press, pp. 301-309, 1988.
[4] W. S. Brown. On Euclid's Algorithm and the Computation of Polynomial Greatest Common Divisors. J. ACM 18 (1971), 478-504.
[5] D. Cox, J. Little, D. O'Shea. Ideals, Varities and Algorithms. Springer-Verlag, 1991.
[6] Sanchit Garg and Eric Schost. Interpolation of polynomials given by straight-line programs. J. Theor. Comp. Sci., 410: 2659-2662, June 2009.
[7] J. von zur Gathen and J. Gerhard. Modern Computer Algebra. Cambridge University Press, UK, 1999.
[8] K. O. Geddes, S. R. Czapor, and G. Labahn. Algorithms for Computer Algebra. Kluwer, 1992.
[9] Mark Giesbrecht and Daniel S. Roche. Diversification improves interpolation. In Proc. ISSAC '2011, ACM Press, pp. 123-130, 2011.
[10] Mahdi Javadi and Michael Monagan. Parallel Sparse Polynomial Interpolation over Finite Fields. In Proceedings of PASCO '2010, ACM Press, pp. 160-168, 2010.
[11] Erich Kaltoren, Y.N. Lakshamn and John-Michael Wiley. Modular Rational Sparse Multivariate Interpolation Algorithm. In Proc. ISSAC 1990, pp. 135-139, ACM Press, 1990.
[12] Erich Kaltofen, Wen-shin Lee, Austin 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.
[13] E. Kaltofen. Fifteen years after DSC and WLSS2 what parallel computations I do today. In Proceedings of PASCO '2010, ACM Press, pp. 10-17, 2010.
[14] J. L. Massey. Shift-register synthesis and bch decoding. IEEE Trans. on Information Theory, 15:122-127, 1969.
[15] Niels Möller and Torbjorn Granlund. Improved division by invariant integers IEEE Transactions on Computers, 60, 165-175, 2011.
[16] Hirokazu Murao and Tetsuro Fujise. Modular Algorithm for Sparse Multivariate Polynomial Interpolation and its Parallel Implementation. J. Symb. Cmpt. 21:377-396, 1996.
[17] Gary Mullen, Daniel Panario. Handbook of Finite Fields. CRC Press, 2013.
[18] S. Pohlig and M. Hellman. An improved algorithm for computing logarithms over GF ( $p$ ) and its cryptographic significance. IEEE Trans. on Information Theory, 24:106-110, 1978.
[19] Michael Rabin. Probabilistic algorithms in finite fields. SIAM J. Comput, 9:273-280, 1979.
[20] Jack Schwartz, Fast probabilistic algorithms for verification of polynomial identities. J. ACM, 27:701-717, 1980.
[21] 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).
[22] Richard Zippel. Probabilistic algorithms for sparse polynomials. In Proceedings of EUROSAM '79, pages 216-226. Springer-Verlag, 1979.
[23] Richard Zippel. Interpolating Polynomials from their Values. J. Symb Cmpt. 9, 3 (1990), 375-403.

