

On the Complexity and Parallel Implementation of Hensel’s Lemma and Weierstrass Preparation

Alexander Brandt, Marc Moreno Maza

Department of Computer Science, The University of Western Ontario,
London, Canada
abrandt5@uwo.ca, moreno@csd.uwo.ca

Abstract

Hensel’s lemma, combined with repeated applications of Weierstrass preparation theorem, allows for the factorization of polynomials with multivariate power series coefficients. We present a complexity analysis for this method and leverage those results to guide the load-balancing of a parallel implementation to concurrently update all factors. In particular, the factorization creates a *pipeline* where the terms of degree k of the first factor are computed simultaneously with the terms of degree $k - 1$ of the second factor, etc. An implementation challenge is the inherent irregularity of computational work between factors, as our complexity analysis reveals. Additional resource utilization and load-balancing is achieved through the parallelization of Weierstrass preparation. Experimental results show the efficacy of this mixed parallel scheme, achieving up to $9\times$ speedup on 12 cores.

Keywords: Formal power series · Weierstrass preparation · Hensel’s lemma · Hensel factorization · Parallel processing · Parallel pipeline

1 Introduction

Factorization via Hensel’s lemma, or simply Hensel factorization, provides a mechanism for factorizing univariate polynomials with multivariate power series coefficients. In particular, for a monic square-free multivariate polynomial in (X_1, \dots, X_n, Y) , one can compute its roots with respect to Y as power series in (X_1, \dots, X_n) . For a bivariate polynomial in (X_1, Y) , the classical Newton-Puiseux method is known to compute the polynomial’s roots with respect to Y as univariate Puiseux series in X_1 . The transition from power series to Puiseux series arises from handling the non-monic case.

The *Hensel-Sasaki Construction* or *Extended Hensel Construction* (EHC) was proposed in [13] as an efficient alternative to the Newton-Puiseux method for the case of univariate coefficients. In the same paper, an extension of the Hensel-Sasaki construction was proposed for multivariate coefficients. In [1], EHC was improved in terms of complexity estimates and practical implementation.

In this paper we present a method of Hensel factorization based on repeated applications of Weierstrass preparation theorem. This method proceeds by a *lazy evaluation* scheme, meaning that more power series terms can be computed on demand without having to restart the computation. This contrasts with a *truncated* implementation where only terms up to a pre-determined degree are computed. The complexity of this proposed method is also measured.

Denote by $M(n)$ a polynomial multiplication time [16, Ch. 8], that is, the number of operations sufficient to multiply two polynomials of degree n . Let $f \in \mathbb{K}[[X_1]][Y]$ have degree d_y in Y and total degree d , where \mathbb{K} is algebraically closed. Our method computes the first k terms of all factors of f within $\mathcal{O}(d_y^3 k + d_y^2 k^2)$ operations in \mathbb{K} . Moreover, we conjecture in Section 4 that our method can achieve $\mathcal{O}(d_y^3 k + d_y^2 M(k) \log k)$ operations in \mathbb{K} through relaxed algorithms [15]. The Hensel-Sasaki Construction of [1] computes all factors in $\mathcal{O}(d^3 M(d) + k^2 d M(d))$. Kung and Traub show that, over the complex numbers \mathbb{C} , the Newton-Puiseux method can be performed in $\mathcal{O}(d^2 k M(k))$ (resp. $\mathcal{O}(d^2 M(k))$) operations in \mathbb{C} using a linear lifting scheme (resp. quadratic lifting scheme) [10]. This complexity is lowered to $\mathcal{O}(d^2 k)$ by Chudnovsky and Chudnovsky in [6].

Nonetheless, the formulation of EHC in [1] is shown to be practically much more efficient than the method of Kung and Traub. Further, our serial implementation of lazy Hensel factorization has already been shown in [5] to be orders of magnitude faster than the implementation of EHC in [1]. We now extend our work by considering parallel processing techniques to further improve the performance of our lazy Hensel factorization. Our complexity analysis is very precise, giving the number of operations required to update each factor individually. This information helps guide our parallel implementation.

In Hensel factorization, computing power series terms of each factor relies on the computed terms of the previous factor. In particular, the output of one Weierstrass preparation becomes the input to another. These successive dependencies naturally lead to a parallel *pipeline* or chain of *producer-consumer* pairs. Within numerical linear algebra, parallel pipelines have already been employed in parallel implementations of singular value decomposition [9], LU decomposition, and Gaussian elimination [12]. Meanwhile, to the best of our knowledge, the only use of parallel pipeline in symbolic computation is [3], which examines a parallel implementation of triangular decomposition of polynomial systems.

However, in our case, work reduces with each application of Weierstrass, resulting in inherent load-balancing issues where each stage of the pipeline requires less work, limiting throughput. To overcome this challenge, we first make use of our complexity estimates to dynamically estimate the work required to update each factor. Second, we compose parallel processing schemes by applying the celebrated map-reduce pattern within Weierstrass preparation, and thus within a stage of the pipeline. Assigning multiple threads to a single pipeline stage is thus used to improve load-balancing and increase throughput. Experimental results show this composition is effective, leading to parallel speedup of up to $9\times$ on a 12-core machine.

The remainder of this paper is organized as follows. Section 2 reviews mathematical background and notations. Further background on our lazy power series of [5] is presented in Section 3. Algorithms and complexity analyses of Weierstrass preparation and Hensel factorization are given in Section 4. Section 5 describes

the parallel processing schemes of Weierstrass preparation and Hensel factorization. The practical use of our complexity estimates for dynamic scheduling is also discussed. Finally, Section 6 presents experimental data.

2 Background

We take this section to present basic concepts and notation of multivariate power series and univariate polynomials over power series (UPoPS). Further, we present constructive proofs for the theorems of Weierstrass preparation and Hensel's lemma for univariate polynomials over power series. Algorithms for Weierstrass preparation and factorization of UPoPS via Hensel's lemma are adapted from these proofs and later presented in Section 4.1 and Section 4.2, respectively. Further introductory details may be found in the book of G. Fischer [7].

2.1 Power Series and Univariate Polynomials over Power Series

Let \mathbb{K} be an algebraically closed field. We denote by $\mathbb{K}[[X_1, \dots, X_n]]$ the ring of formal power series with coefficients in \mathbb{K} and with variables X_1, \dots, X_n .

Let $f = \sum_{e \in \mathbb{N}^n} a_e X^e$ be a formal power series, where $a_e \in \mathbb{K}$, $X^e = X_1^{e_1} \dots X_n^{e_n}$, $(e_1, \dots, e_n) \in \mathbb{N}^n$, and $|e| = e_1 + \dots + e_n$. Let $k \in \mathbb{N}$. The *homogeneous part* and *polynomial part* of f in degree k are denoted by $f_{(k)}$ and $f^{(k)}$, and are defined by $f_{(k)} = \sum_{|e|=k} a_e X^e$ and $f^{(k)} = \sum_{i \leq k} f_{(i)}$. The *order* of a formal power series f , denoted by $\text{ord}(f)$, is defined as $\min\{i \mid f_{(i)} \neq 0\}$, if $f \neq 0$, and as ∞ otherwise.

Recall several properties regarding power series. First, $\mathbb{K}[[X_1, \dots, X_n]]$ is an integral domain. Second, the set $\mathcal{M} = \{f \in \mathbb{K}[[X_1, \dots, X_n]] \mid \text{ord}(f) \geq 1\}$ is the only maximal ideal of $\mathbb{K}[[X_1, \dots, X_n]]$. Third, for all $k \in \mathbb{N}$, we have $\mathcal{M}^k = \{f \in \mathbb{K}[[X_1, \dots, X_n]] \mid \text{ord}(f) \geq k\}$. Note that for $n = 0$ we have $\mathcal{M} = \langle 0 \rangle$. Further, note that $f_{(k)} \in \mathcal{M}^k \setminus \mathcal{M}^{k+1}$ and $f_{(0)} \in \mathbb{K}$. Fourth, a unit $u \in \mathbb{K}[[X_1, \dots, X_n]]$ has $\text{ord}(u) = 0$ or, equivalently, $u \notin \mathcal{M}$.

Let $f, g, h, p \in \mathbb{K}[[X_1, \dots, X_n]]$. The *sum* and *difference* $f = g \pm h$ is given by $\sum_{k \in \mathbb{N}} (g_{(k)} \pm h_{(k)})$. The product $p = gh$ is given by $\sum_{k \in \mathbb{N}} (\sum_{i+j=k} g_{(i)} h_{(j)})$. Notice that these formulas naturally suggest a *lazy evaluation* scheme, where the result of an arithmetic operation can be incrementally computed for increasing *precision* or, equivalently, for homogeneous parts of increasing degree. A power series f is said to be known to precision $k \in \mathbb{N}$, when $f_{(i)}$ is known for all $0 \leq i \leq k$. Such an update function, parameterized by k , for addition or subtraction is simply $f_{(k)} = g_{(k)} \pm h_{(k)}$; an update function for multiplication is $p_{(k)} = \sum_{i=0}^k g_{(i)} h_{(k-i)}$. Lazy schemes for power series are discussed further in Section 3. From these update formulas, the following observation follows.

Observation 1 (power series arithmetic) *Let $f, g, h, p \in \mathbb{K}[[X_1]]$ with $f = g \pm h$ and $p = gh$. $f_{(k)} = g_{(k)} \pm h_{(k)}$ can be computed in 1 operation in \mathbb{K} . $p_{(k)} = \sum_{i=0}^k g_{(i)} h_{(k-i)}$ can be computed in $2k - 1$ operations in \mathbb{K} .*

Now, let $f, g \in \mathbb{A}[Y]$ be univariate polynomials over power series where, $\mathbb{A} = \mathbb{K}[[X_1, \dots, X_n]]$. Writing $f = \sum_{i=0}^d a_i Y^i$, for $a_i \in \mathbb{A}$, we have that the degree of

f (denoted $\deg(f, Y)$ or simply $\deg(f)$) is d . Note that arithmetic operations for UPoPS are easily inherited from the arithmetic of its power series coefficients. A UPoPS is said to be known up to precision k if each of its power series coefficients are known up to precision k . A UPoPS f is said to be *general (in Y) of order j* if $f \bmod \mathcal{M}[Y]$ has order j when viewed as a power series. That is, for $f = \sum_{i=0}^d a_i Y^i$, $a_i \in \mathcal{M}$ for $0 \leq i < j$

2.2 Weierstrass Preparation Theorem and Factorization via Hensel's Lemma

The Weierstrass Preparation Theorem (WPT) is fundamentally a theorem regarding factorization. In the context of analytic functions, WPT implies that any analytic function resembles a polynomial in the vicinity of the origin. Generally, WPT can be stated for power series over power series, i.e. for the power series $\mathbb{K}[[X_1, \dots, X_n]][[Y]] = \mathbb{A}[[Y]]$. This can be used to prove that \mathbb{A} is both a unique factorization domain and a Noetherian ring. See [5] for a proof of WPT in the case of $\mathbb{A}[[Y]]$. Here, it is sufficient to state the theorem for UPoPS.

We begin with a simple lemma which serves as the basis of our eventual proof of WPT and our implementation.

Lemma 1 *Let $f, g, h \in \mathbb{K}[[X_1, \dots, X_n]]$ such that $f = gh$. Let $f_i = f_{(i)}, g_i = g_{(i)}, h_i = h_{(i)}$. If $f_0 = 0$ and $h_0 \neq 0$, then g_k is uniquely determined by f_1, \dots, f_k and h_0, \dots, h_{k-1}*

PROOF. We proceed by induction on k . Since $f_0 = g_0 h_0 = 0$ and $h_0 \neq 0$ both hold, the statement holds for $k = 0$. Now let $k > 0$, assuming the hypothesis holds for $k - 1$. To determine g_k it is sufficient to expand $f = gh$ modulo \mathcal{M}^{k+1} : $f_1 + f_2 + \dots + f_k = g_1 h_0 + (g_1 h_1 + g_2 h_0) + \dots + (g_1 h_{k-1} + \dots + g_{k-1} h_1 + g_k h_0)$; and, recalling $h_0 \in \mathbb{K} \setminus \{0\}$, we have $g_k = {}^{1/h_0}(f_k - g_1 h_{k-1} - \dots - g_{k-1} h_1)$. \square

Theorem 1 (Weierstrass Preparation Theorem) *Let $f = \sum_{i=0}^{d+m} a_i Y^i \in \mathbb{K}[[X_1, \dots, X_n]][[Y]]$ where $d \geq 0$ is the smallest integer such that $a_d \notin \mathcal{M}$ and $0 \leq m \in \mathbb{N}$. Assume that $f \not\equiv 0 \bmod \mathcal{M}[Y]$. Then, there exists a unique pair p, α satisfying the following:*

- (i) $f = p\alpha$,
- (ii) α is an invertible element of $\mathbb{K}[[X_1, \dots, X_n]][[Y]]$,
- (iii) p is a monic polynomial of degree d ,
- (iv) writing $p = Y^d + b_{d-1}Y^{d-1} + \dots + b_1Y + b_0$, we have $b_{d-1}, \dots, b_0 \in \mathcal{M}$.

PROOF. If $n = 0$, writing $f = \alpha Y^d$ with $\alpha = \sum_{i=0}^m a_{i+d} Y^i$ proves the existence of the decomposition. Now, assume $n \geq 1$. Write $\alpha = \sum_{i=0}^m c_i Y^i$, with $c_i \in \mathbb{K}[[X_1, \dots, X_n]]$. We will determine $b_0, \dots, b_{d-1}, c_0, \dots, c_m$ modulo successive powers of \mathcal{M} . Since we require α to be a unit, $c_0 \notin \mathcal{M}$. By definition, a_0, \dots, a_{d-1} are all 0 mod \mathcal{M} . Equating coefficients in $f = p\alpha$ we have:

$$\begin{aligned}
a_0 &= b_0 c_0 \\
a_1 &= b_0 c_1 + b_1 c_0 \\
&\vdots \\
a_{d-1} &= b_0 c_{d-1} + b_1 c_{d-2} + \cdots + b_{d-2} c_1 + b_{d-1} c_0 \\
a_d &= b_0 c_d + b_1 c_{d-1} + \cdots + b_{d-1} c_1 + c_0 \\
&\vdots \\
a_{d+m-1} &= b_{d-1} c_m + c_{m-1} \\
a_{d+m} &= c_m
\end{aligned} \tag{1}$$

and thus b_0, \dots, b_{d-1} are also all 0 mod \mathcal{M} . Then, $c_i \equiv a_{d+i}$ mod \mathcal{M} for all $0 \leq i \leq m$. All coefficients have thus been determined mod \mathcal{M} . Let $k > 0$ be an integer. Assume, inductively, that all $b_0, \dots, b_{d-1}, c_0, \dots, c_m$ have been determined mod \mathcal{M}^k , we will now determine them mod \mathcal{M}^{k+1} .

It follows from Lemma 1 that b_0 can be determined mod \mathcal{M}^{k+1} from the equation $a_0 = b_0 c_0$. Consider now the second equation. Since b_0 is known mod \mathcal{M}^{k+1} , and $b_0 \in \mathcal{M}$, the product $b_0 c_1$ is also known mod \mathcal{M}^{k+1} . Then, we can determine b_1 using Lemma 1 and the formula $a_1 - b_0 c_1 = b_1 c_0$. This procedure follows for b_2, \dots, b_{d-1} . With b_0, \dots, b_{d-1} known mod \mathcal{M}^{k+1} each c_0, \dots, c_m can be determined mod \mathcal{M}^{k+1} from the last $m + 1$ equations. \square

One requirement of Weierstrass Preparation Theorem is that $f \not\equiv 0 \pmod{\mathcal{M}[Y]}$. That is to say, f cannot vanish at $(X_1, \dots, X_n) = (0, \dots, 0)$ and, specifically, f is general of order $d = \deg(p)$. However, one can always apply a suitable linear change in coordinates to meet this requirement. See Algorithm 2 in Section 4. Weierstrass preparation provides a mechanism for factorizing a UPoPS into two factors. One may then apply WPT several times to fully factorize a UPoPS. The existence of such a factorization is given by Hensel's lemma for UPoPS.

Theorem 2 (Hensel's Lemma) *Let $f = Y^d + \sum_{i=0}^{d-1} a_i Y^i$ be a monic polynomial with $a_i \in \mathbb{K}[[X_1, \dots, X_n]]$. Let $\bar{f} = f(0, \dots, 0, Y) = (Y - c_1)^{d_1} (Y - c_2)^{d_2} \cdots (Y - c_r)^{d_r}$ for $c_1, \dots, c_r \in \mathbb{K}$ and positive integers d_1, \dots, d_r . Then, there exists $f_1, \dots, f_r \in \mathbb{K}[[X_1, \dots, X_n]][Y]$, all monic in Y , such that:*

- (i) $f = f_1 \cdots f_r$,
- (ii) $\deg(f_i, Y) = d_i$ for $1 \leq i \leq r$, and
- (iii) $\bar{f}_i = (Y - c_i)^{d_i}$ for $1 \leq i \leq r$.

PROOF. We proceed by induction on r . For $r = 1$, $d_1 = d$ and we have $f_1 = f$, where f_1 has all the required properties. Now assume $r > 1$. A change of coordinates in Y , sends c_r to 0. Define $g(X_1, \dots, X_n, Y) = f(X_1, \dots, X_n, Y + c_r) = (Y + c_r)^d + a_{d-1}(Y + c_r)^{d-1} + \cdots + a_0$. By construction, g is general of order d_r and WPT can be applied to obtain $g = p\alpha$ with p being of degree d_r and $\bar{p} = Y^{d_r}$. Reversing the change of coordinates we set $f_r = p(Y - c_r)$ and $f^* = \alpha(Y - c_r)$, and we have $f = f^* f_r$. f_r is a monic polynomial of degree d_r in Y with $\bar{f}_r = (Y - c_r)^{d_r}$. Moreover, we have $\bar{f}^* = (Y - c_1)^{d_1} (Y - c_2)^{d_2} \cdots (Y - c_{r-1})^{d_{r-1}}$. The inductive hypothesis applied to f^* implies the existence of f_1, \dots, f_{r-1} . \square

2.3 Parallel Patterns

We are concerned with *thread-level parallelism*, where multiple threads of execution within a single process enable concurrent processing. Our parallel implementation employs several so-called *parallel patterns*—algorithmic structures and organizations for efficient parallel processing. We review a few patterns: *map*, *producer-consumer*, and *pipeline*. See [11] for a detailed discussion.

2.3.1 Map

The map pattern applies a function to each item in a collection, simultaneously executing the function on each independent data item. Often, the application of a map produces a new collection with the same shape as the input collection. Alternatively, the map pattern may modify each data item in place or, when combined with the *reduce* pattern, produce a single data item. The reduce pattern combines data items pair-wise using some *combiner* function.

When data items to be processed outnumber available threads, the map pattern can be applied block-wise, where the data collection is (evenly) partitioned and each thread assigned a partition rather than a single data item.

Where a **for** loop has independent iterations, the map pattern is easily applied to execute each iteration of the loop concurrently. Due to this ubiquity, the map pattern is often implicit with such parallel for loops simply being labelled **parallel_for**. In this way, the number of threads to use and the partitioning of the data collection can be a dynamic property of the algorithm.

2.3.2 Producer-Consumer and Asynchronous Generators

The producer-consumer pattern describes two functions connected by a queue. The producer creates data items, pushing them to the queue, meanwhile the consumer processes data items, pulling them from the queue. Where both the creation of data items and their processing requires substantial computational work, producer and consumer may operate concurrently, with the queue facilitating inter-thread communication.

A *generator* or *iterator* is a special kind of co-routine function which **yields** data elements one at a time, rather than many together as a collection; see, e.g. [14, Ch. 8]. Combining the producer-consumer pattern with generators allows for an *asynchronous generator*, where the generator function is the producer and the calling function is the consumer. The intermediary queue allows the generator to produce items meanwhile the calling function processes them.

2.3.3 Pipeline

The pipeline pattern is a sequence of stages, where the output of one stage is used as the input to another. Two consecutive stages form a producer-consumer pair, with internal stages being both a consumer and a producer. Concurrency arises where each stage of the pipeline may be executed in parallel. Moreover, the pipeline pattern allows for earlier data items to flow from one stage to the next without waiting for later items to become available.

In terms of the latency of processing a single data item, a pipeline does not improve upon its serial counterpart. Rather, a parallel pipeline improves throughput, the amount of data that can be processed in a given amount of time. Throughput is limited by the slowest stage of a pipeline, and thus special care must be given to ensure each stage of the pipeline runs in nearly equal time.

A pipeline may be implicitly and dynamically created where an asynchronous generator consumes data from another asynchronous generator. The number of asynchronous generator calls, and thus the number of stages in the pipeline, can be dynamic to fit the needs of the application at runtime.

3 Lazy Power Series

As we have seen in Section 2.1, certain arithmetic operations on power series naturally lead to a lazy evaluation scheme. In this scheme, homogeneous parts of a power series are computed one at a time for increasing degree, as requested. Our implementation of lazy power series, lazy Weierstrass preparation, and lazy Hensel factorization is detailed in [5]. For the remainder of this paper, it is sufficient to understand that lazy power series rely on the following three principles:

- (i) the availability of an update function to compute the homogeneous part of a particular degree;
- (ii) the capturing of parameters required for that update function; and
- (iii) the storage of previously computed homogeneous parts.

Where a power series is constructed from arithmetic operations on other power series, the latter may be called the *ancestors* of the former. For example, the power series $f = gh$ has ancestors g and h and an update function $f_{(k)} = \sum_{i=0}^k g_{(i)}h_{(k-i)}$. In implementation, and in the algorithms which follow in this paper, we can thus augment a power series with: (i) its current precision; (ii) references to its ancestors, if any; and (iii) a reference to its update function.

Under this scheme, we make three remarks. Firstly, a power series can be lazily constructed using essentially no work. The initialization of a lazy power series only requires specifying the appropriate update function and storing references to its ancestors. Secondly, specifying an update function and the ancestors of a power series is sufficient for defining and computing that power series. Thirdly, when updating a particular power series, its ancestors can automatically and recursively be updated as necessary using their own update functions.

Hence, in the algorithms which are presented in the following sections, it is sufficient to define the update function for a power series. For example, Algorithm 1 simultaneously updates p and α as produced from a Weierstrass preparation. Further, operations on power series should be understood to be only the initialization of a power series, with no terms of the power series yet computed; e.g., Algorithm 3 for Hensel factorization.

4 Algorithms and Complexity

In this section we present algorithms for Weierstrass preparation and Hensel factorization adapted from their constructive proofs; see Theorem 1 and Theorem 2, respectively. For each algorithm we analyze its complexity. These results will guide our eventual parallel variations described later in Section 5.

4.1 Weierstrass Preparation

From the proof of Weierstrass preparation (Theorem 1), we derive WEIERSTRASSUPDATE (Algorithm 1). That proof proceeds modulo increasing powers of the maximal ideal \mathcal{M} , which is equivalent to computing homogeneous parts of increasing degree, just as required for our lazy power series. For an application of Weierstrass preparation producing p and α , this WEIERSTRASSUPDATE acts as the update function for p and α , updating both simultaneously.

By rearranging the terms of the first d equations of (1) and applying Lemma 1 we obtain “phase 1” of WEIERSTRASSUPDATE, where each coefficient of p is updated. By rearranging the terms of the next $m + 1$ equations of (1) we obtain “phase 2” of WEIERSTRASSUPDATE, where each coefficient of α is updated. From Algorithm 1, it is then routine to show the following two observations, which together lead to Theorem 3.

Algorithm 1 WEIERSTRASSUPDATE(k, f, p, α)

Input: $f = \sum_{i=0}^{d+m} a_i Y^i$, $p = Y^d + \sum_{i=0}^{d-1} b_i Y^i$, $\alpha = \sum_{i=0}^m c_i Y^i$, $a_i, b_i, c_i \in \mathbb{K}[[X_1, \dots, X_n]]$ satisfying Theorem 1, with $b_0, \dots, b_{d-1}, c_0, \dots, c_m$ known modulo \mathcal{M}^k , \mathcal{M} the maximal ideal of $\mathbb{K}[[X_1, \dots, X_n]]$.

Output: $b_0, \dots, b_{d-1}, c_0, \dots, c_m$ known modulo \mathcal{M}^{k+1} , updated in-place.

```

    ▷ phase one
1: for  $i = 0$  to  $d - 1$  do
2:    $F_{i(k)} := a_{i(k)}$ 
3:   if  $i \leq m$  then
4:     for  $j = 0$  to  $i - 1$  do
5:        $F_{i(k)} := F_{i(k)} - (b_j c_{i-j})_{(k)}$ 
6:   else
7:     for  $j = 0$  to  $m - 1$  do
8:        $F_{i(k)} := F_{i(k)} - (b_{i+j-m} c_{m-j})_{(k)}$ 
9:    $s := 0$ 
10:  for  $j = 1$  to  $k - 1$  do
11:     $s := s + b_{i(k-j)} \times c_{0(j)}$ 
12:     $b_{i(k)} := (F_{i(k)} - s) / c_{0(0)}$ 

    ▷ phase two
13:  $c_{m(k)} := a_{d+m(k)}$ 
14: for  $i = 1$  to  $m$  do
15:   if  $i \leq d$  then
16:      $c_{m-i(k)} := a_{d+m-i(k)} - \sum_{j=1}^i (b_{d-j} c_{m-i+j})_{(k)}$ 
17:   else
18:      $c_{m-i(k)} := a_{d+m-i(k)} - \sum_{j=1}^d (b_{d-j} c_{m-i+j})_{(k)}$ 

```

Observation 2 (Weierstrass phase 1 complexity) For WEIERSTRASSUPDATE over $\mathbb{K}[[X_1]][Y]$, computing the homogeneous part of degree k for b_i with $0 \leq i < d$ requires $2ki + 2k - 1$ operations in \mathbb{K} if $i \leq m$, or $2km + 2k - 1$ operations in \mathbb{K} if $i > m$.

Observation 3 (Weierstrass phase 2 complexity) For WEIERSTRASSUPDATE over $\mathbb{K}[[X_1]][Y]$, computing the homogeneous part of degree k for c_{m-i} , $0 \leq i < m$, requires $2ki$ operations in \mathbb{K} if $i \leq d$, or $2kd$ operations in \mathbb{K} if $i > d$.

Theorem 3 (Weierstrass preparation complexity) For Weierstrass preparation producing $f = p\alpha$ with $f, p, \alpha \in \mathbb{K}[[X_1]][Y]$, with $\deg(p) = d$, and $\deg(\alpha) = m$, computing p and α up to precision k requires $dmk^2 + dk^2 + dmk$ operations in \mathbb{K} .

PROOF. Let i be the index of a coefficient of p or α . As a first step, consider the cost of computing the homogeneous part of degree k of each coefficient of p and α . First consider $i < t = \min(d, m)$, that is, $b_0, \dots, b_t, c_0, \dots, c_t$. From Observation 2 we have that computing the k th homogeneous part alone requires $2ki + 2k - 1$ operations in \mathbb{K} for each b_i . From Observation 3 we have $2ki$ operations for each c_i . For $0 \leq i < t$, this yields a total of $2kt^2 + 2kt - t$. Next, we have three cases: (a) $t = d = m$, (b) $m = t < i < d$, or (c) $d = t < i < m$. In case (a) there is no additional work. In case (b), phase 1 contributes an additional $(d-m) * (2km + 2k - 1)$ operations and phase 2 contributes no additional operations. In case (c), phase 2 contributes an additional $(m-d) * (2kd)$ operations and phase 1 contributes no additional operations. In all cases, the total number of operations to update p and α from precision $k-1$ to precision k is $2dmk + 2dk - d$. Finally, to compute p and α up to precision k requires $dmk^2 + dk^2 + dmk$ operations in \mathbb{K} . \square

A useful consideration is when the input to Weierstrass preparation is monic. This necessarily makes α monic, and the overall complexity of Weierstrass preparation is reduced. This case arises for each application of Weierstrass preparation in Hensel factorization. The following corollary proves this, following Theorem 3.

Corollary 1 (Weierstrass preparation complexity for monic input) For an application of Weierstrass preparation producing $f = p\alpha$ with $f, p, \alpha \in \mathbb{K}[[X_1]][Y]$, f monic in Y , $\deg(p) = d$ and $\deg(\alpha) = m$, computing p and α up to precision k requires $dmk^2 + dmk$ operations in \mathbb{K} .

PROOF. If f is monic then α is necessarily monic and $c_m = 1$. For $i \geq m$ we save computing $(b_{i-m}c_m)_{(k)}$ for the update of $b_{i(k)}$. For $1 \leq i \leq d$ we save computing $(b_{d-j}c_{m-i+j})_{(k)}$ for $j = i$ for the update of each $c_{m-i(k)}$. First, consider updating p and α from precision $k-1$ to precision k . Let $t = \min(d, m)$. We have three cases: (a) $t = d = m$, (b) $m = t < i < d$, or (c) $d = t < i < m$. In case (a) we save $d(2k-1)$ operations in phase 2, as compared to case (a) from the proof of Theorem 3. In case (b) we save $(d-m)(2k-1)$ operations in phase 1 and $m(2k-1)$ operations in phase 2. In case (c) we save $d(2k-1)$ operations in phase 2. In all cases we save a total of $d(2k-1)$ operations, resulting in $2dmk$ operations in \mathbb{K} to update p and α from precision $k-1$ to precision k . Finally, to compute p and α up to precision k requires $dmk^2 + dmk$ operations in \mathbb{K} . \square

4.2 Hensel Factorization

Before we begin Hensel factorization, we will first see how to perform a translation, or Taylor shift, by lazy evaluation. For $f = \sum_{i=0}^d a_i Y^i \in \mathbb{K}[[X_1, \dots, X_n]][Y]$ and $c \in \mathbb{K}$, computing $f(Y+c)$ begins by pre-computing the coefficients of the binomial

Algorithm 2 TAYLORSHIFTUPDATE(k, f, \mathbf{S}, i)

Input: For $f = \sum_{j=0}^d a_j Y^j$, $g = f(Y + c) = \sum_{j=0}^d b_j Y^j$, obtain the homogeneous part of degree k for b_i . $\mathbf{S} \in \mathbb{K}^{(d+1) \times (d+1)}$ is a lower triangular matrix of coefficients of $(Y + c)^j$ for $j = 0, \dots, d$,
Output: $b_{i(k)}$, the homogeneous part of degree k of b_i .
1: $b_{i(k)} := 0$
2: **for** $\ell = i$ **to** d **do**
3: $j := \ell + 1 - i$
4: $b_{i(k)} := b_{i(k)} + S_{\ell+1, j} \times a_{\ell(k)}$
5: **return** $b_{i(k)}$

expansions $(Y + c)^j$ for $0 \leq j \leq d$. These coefficients are stored in a matrix \mathbf{S} . Then, each coefficient of $f(Y + c) = \sum_{i=0}^d b_i Y^i$ is a linear combination of the coefficients of f scaled by the appropriate elements of \mathbf{S} . Since those elements of \mathbf{S} are only elements of \mathbb{K} , this linear combination does not change the degree and, for some integer k , $b_{i(k)}$ relies only on $a_{\ell(k)}$ for $i \leq \ell \leq d$. Computing $b_{i(k)}$ from this linear combination is described as TAYLORSHIFTUPDATE in Algorithm 2; and its complexity is easily found and stated as Observation 4.

Observation 4 (Taylor shift complexity) For a UPoPS $f = \sum_{i=0}^d a_i Y^i \in \mathbb{K}[[X_1]][Y]$, computing the homogeneous part of degree k for all coefficients of the shifted UPoPS $f(Y + c)$ requires $d^2 + 2d + 1$ operations in \mathbb{K} .

Having defined update functions for Weierstrass preparation and Taylor shift, lazy Hensel factorization is immediate. Hensel factorization requires no additional update function, only the initialization of the appropriate chain of ancestors. Algorithm 3 shows this initialization of the factors and their ancestors through repeated applications of Taylor shift and Weierstrass preparation. Note that factors are

Algorithm 3 HENSELFACTORIZATION(f)

Input: $f = Y^d + \sum_{i=0}^{d-1} a_i Y^i$, $a_i \in \mathbb{K}[[X_1, \dots, X_n]]$.
Output: f_1, \dots, f_r satisfying Theorem 2.
1: $\hat{f} = f(0, \dots, 0, Y)$
2: $(c_1, \dots, c_r), (d_1, \dots, d_r) :=$ roots and their multiplicities of \hat{f}
3: $c_1, \dots, c_r :=$ SORT($[c_1, \dots, c_r]$) by increasing multiplicity ▷ see Theorem 5
4: $\hat{f}_1 = \hat{f}$
5: **for** $i = 1$ **to** $r - 1$ **do**
6: $g_i := \hat{f}_i(Y + c_i)$
7: $p_i, \alpha_i :=$ WeierstrassPreparation(g)
8: $\hat{f}_i := p_i(Y - c_i)$
9: $\hat{f}_{i+1} := \alpha_i(Y - c_i)$
10: $f_r := \hat{f}_r$
11: **return** f_1, \dots, f_r




Figure 1: The ancestor chain for the Hensel factorization $f = f_1 f_2 f_3 f_4$. Updating f_1 requires updating g_1, p_1, α_1 ; then updating f_2 requires updating $\hat{f}_2, g_2, p_2, \alpha_2$; then updating f_3 requires updating $\hat{f}_3, g_3, p_3, \alpha_3$; then updating f_4 requires only its own Taylor shift. These groupings form the eventual stages of the Hensel pipeline (Algorithm 8).

sorted by increasing degree; this enables better load-balancing in an eventual parallel implementation, see Theorem 5 and Algorithms 7–9. Fig. 1 shows the chain of ancestors created for the factorization $f = f_1 f_2 f_3 f_4$. Moreover, the figure shows the grouping of ancestors required to update each factor; the complexity of which is shown in Theorem 5. First, however, we analyze the complexity of HENSELFACTORIZATION for the common case where each factor has degree 1.

Theorem 4 (Hensel factorization complexity for simple roots) *Applying HENSELFACTORIZATION on $f \in \mathbb{K}[[X_1]][Y]$, where $\deg(f) = d$, with all resulting factors having degree 1, and updating each factor to precision k , requires $2/3 d^3 k + 1/2 d^2 k^2 + 5/2 d^2 k - 1/2 dk^2 + 35/6 dk - 9k$ operations in \mathbb{K} .*

PROOF. For each factor except the last, HENSELFACTORIZATION requires one Taylor shift, one Weierstrass preparation, and two more Taylor shifts. For the first factor we have that the first Taylor shift is of degree d , the Weierstrass preparation produces p_1 and α_1 of degree 1 and $d - 1$, respectively, and then the two Taylor shifts are of degree 1 and $d - 1$. This pattern continues for each factor but the last. f_d is obtained from the shifted α_{d-1} . The result is: a shift of degree $d - i + 1$ for $i = 1, \dots, d - 1$ (for each \hat{f}_i), $d - 1$ shifts of degree 1 (for each p_i), and a shift of degree $d - i$ for $i = 1, \dots, d - 1$ (for each α_i). From Observation 4, obtaining a Taylor shift of degree d' to precision k requires $d'^2 k + 2d'k + k$ operations in \mathbb{K} . Summing over each group of Taylor shifts gives, respectively, $k(1/3 d^3 + 3/2 d^2 + 13/6 d - 4)$, $4k(d - 1)$, and $k(1/3 d^3 + 1/2 d^2 + 1/6 d - 1)$, for a total of $k(2/3 d^3 + 2d^2 + 19/3 d - 9)$ operations in \mathbb{K} .

The remaining operations arise from the repeated Weierstrass preparations. For i from 1 to $d - 1$ we apply Weierstrass preparations to produce p_i, α_i pairs of respective degree 1, $d - i$. From Corollary 1 we have that each such Weierstrass preparation requires $(d - i)k^2 + (d - i)k$ operations in \mathbb{K} . Summing over $i = 1, \dots, d - 1$ yields $1/2(d^2 k^2 + d^2 k - dk^2 - dk)$. Finally, combining this with the previous Taylor shift costs leads to the desired result. \square

Theorem 5 (Hensel factorization complexity per factor) *Let \hat{d}_i be the degree of \hat{f}_i during the HENSELFACTORIZATION algorithm applied to $f \in \mathbb{K}[[X_1]][Y]$, $\deg(f) = d$. To update the first factor with degree d_1 to precision k requires $d_1 \hat{d}_2 k^2 + d^2 k + d_1 dk + 2d_1 k + 2dk + 2k$ operations in \mathbb{K} . To update the factor of degree d_i , for $1 < i < r$ to precision k requires $d_i \hat{d}_{i+1} k^2 + 2\hat{d}_i^2 k + d_i \hat{d}_i k + 2d_i k + 4\hat{d}_i k + 3k$ operations in \mathbb{K} . To update the final factor of degree d_r to precision k requires $d_r^2 k + 2d_r k + k$ operations in \mathbb{K} .*

PROOF. Updating the first factor produced by HENSELFACTORIZATION requires one Taylor shift of degree d , one Weierstrass preparation producing p_1 and α_1 of degree d_1 and $\hat{d}_2 = d - d_1$, and one Taylor shift of degree d_1 to obtain f_1 from p . From Observation 4 and Corollary 1 we have that the Taylor shifts require $k(d^2 + 2d + 1) + k(d_1^2 + 2d_1 + 1)$ operations in \mathbb{K} and the Weierstrass preparation requires $d_1(d - d_1)k^2 + d_1(d - d_1)k$ operations in \mathbb{K} . The total cost counted as operations in \mathbb{K} is thus $d_1 \hat{d}_2 k^2 + d^2 k + d_1 dk + 2d_1 k + 2dk + 2k$.

Updating each following factor, besides the last, requires one Taylor shift of degree \hat{d}_i to update \hat{f}_i from α_{i-1} , one Taylor shift of degree \hat{d}_i to update g_i from

\widehat{f}_i , one Weierstrass preparation to obtain p_i and α_i of degree d_i and $\widehat{d}_{i+1} = \widehat{d}_i - d_i$, and one Taylor shift of degree d_i to obtain f_i from p_i . The Taylor shifts require $2k(\widehat{d}_i^2 + 2\widehat{d}_i + 1) + k(d_i^2 + 2d_i + 1)$ operations in \mathbb{K} . The Weierstrass preparation requires $d_i(\widehat{d}_i - d_i)k^2 + d_i(\widehat{d}_i - d_i)k$ operations in \mathbb{K} . The total cost counted as operations in \mathbb{K} is thus $d_i\widehat{d}_{i+1}k^2 + 2\widehat{d}_i^2k + d_i\widehat{d}_ik + 2d_ik + 4\widehat{d}_ik + 3k$.

Finally, updating the last factor to precision k requires a single Taylor shift of degree d_r costing $d_r^2k + 2d_rk + k$ operations in \mathbb{K} . \square

Corollary 2 (Hensel factorization complexity per factor, per iteration) *Let \widehat{d}_i be the degree of \widehat{f}_i during the HENSELFACTORIZATION algorithm applied to $f \in \mathbb{K}[[X_1]][Y]$, $\deg(f) = d$. Computing the k th homogeneous part of f_1 with degree d_1 requires $2d_1\widehat{d}_2k + d_1^2 + d^2 + 2d_1 + 2d + 2$ operations in \mathbb{K} . Computing the k th homogeneous part of f_i of degree d_i , $1 < i < r$, requires $2d_i\widehat{d}_{i+1}k + d_i^2 + 2\widehat{d}_i^2 + 4\widehat{d}_i + 2d_i + 3$ operations in \mathbb{K} . Computing the k th homogeneous part of f_r with degree d_r requires $d_r^2 + 2d_r + 1$ operations in \mathbb{K} .*

PROOF. Follows directly from Observation 4, Corollary 1, and Theorem 5. \square

Corollary 3 (Hensel factorization complexity) *For $f \in \mathbb{K}[[X_1]][Y]$, where $\deg(f) = d$ and f factorizes as $f_1 \cdots f_r$, the factors can be computed up to precision k within $\mathcal{O}(d^3k + d^2k^2)$ operations in \mathbb{K} .*

PROOF. Let f_1, \dots, f_r have respective degrees d_1, \dots, d_r . Let $\widehat{d}_i = \sum_{j=i}^r d_j$ (thus $\widehat{d}_1 = d$ and $\widehat{d}_r = d_r$). From Theorem 5, each f_i , $1 \leq i < r$ requires $\mathcal{O}(d_i\widehat{d}_{i+1}k^2 + \widehat{d}_i^2k)$ operations in \mathbb{K} to be updated to precision k (or $\mathcal{O}(d_r^2k)$ for f_r). We have $\sum_{i=1}^{r-1} d_i\widehat{d}_{i+1} \leq \sum_{i=1}^{r-1} d_id < d^2$ and $\sum_{i=1}^r \widehat{d}_i^2 \leq \sum_{i=1}^r d^2 = rd^2 \leq d^3$. Hence, all factors can be updated to precision k within $\mathcal{O}(d^3k + d^2k^2)$ operations in \mathbb{K} . \square

Corollary 3 shows that the two dominant terms in the cost of computing a Hensel factorization of a UPoPS of degree d , up to precision k , are d^3k and d^2k^2 . From the proof of Theorem 5, the former term arises from the cost of the Taylor shifts in Y , meanwhile, the latter term arises from the (polynomial) multiplication of homogeneous parts in Weierstrass preparation. This observation then leads to the following conjecture. Recall that $M(n)$ denotes a polynomial multiplication time [16, Ch. 8]. From [15], relaxed algorithms, which improve the performance of lazy evaluation schemes, can be used to compute a power series product in $\mathbb{K}[[X_1]]$ up to precision k within $\mathcal{O}(M(k) \log k)$ operations in \mathbb{K} .

Conjecture 1 *Let $f \in \mathbb{K}[[X_1]][Y]$ factorize as $f_1 \cdots f_r$ using HENSELFACTORIZATION. Let $\deg(f) = d$. Updating the factors f_1, \dots, f_r to precision k using relaxed algorithms requires at most $\mathcal{O}(d^3k + d^2M(k) \log k)$ operations in \mathbb{K} .*

Comparatively, the Hensel-Sasaki Construction requires at most $\mathcal{O}(d^3M(d) + dM(d)k^2)$ operations in \mathbb{K} to compute the first k terms of all factors of $f \in \mathbb{K}[[X_1, Y]]$, where f has total degree d [1]. The method of Kung and Traub [10], requires $\mathcal{O}(d^2kM(k))$ (using linear lifting) or $\mathcal{O}(d^2M(k))$ (using quadratic lifting). Already,

Algorithm 4 UPDATETODEGPARALLEL(k, f, t)

Input: A positive integer k , $f \in \mathbb{K}[[X_1, \dots, X_n]]$ known to at least precision $k - 1$. If f has ancestors, it is the result of a binary operation. A positive integer t for the number of threads to use.

Output: f is updated to precision k , in place.

```
1: if  $f_{(k)}$  already computed then
2:   | return
3:    $g, h := \text{FIRSTANCESTOR}(f), \text{SECONDANCESTOR}(f)$ 
4:   UPDATETODEGPARALLEL( $k, g, t$ );
5:   UPDATETODEGPARALLEL( $k, h, t$ );
6:   if  $f$  is a product then
7:      $\mathcal{V} = [0, \dots, 0]$  ▷ 0-indexed list of size  $t$ 
8:     parallel for  $j = 0$  to  $t - 1$ 
9:       | for  $i = j^{k/t}$  to  $(j+1)^{k/t} - 1$  while  $i \leq k$  do
10:        |  $\mathcal{V}[j] := \mathcal{V}[j] + g_{(i)}h_{(k-i)}$ 
11:      $f_{(k)} := \sum_{j=0}^{t-1} \mathcal{V}[j]$  ▷ reduce
12:   else if  $f$  is a  $p$  from a Weierstrass preparation then
13:     | WEIERSTRASSPHASE1PARALLEL( $k, g, f, h, \text{WEIERSTRASSDATA}(f), t$ )
14:   else if  $f$  is an  $\alpha$  from a Weierstrass preparation then
15:     | WEIERSTRASSPHASE2PARALLEL( $k, g, h, f, t$ )
16:   else
17:     | UPDATETODEG( $k, f$ )
```

Corollary 3—where $d = \deg(f, Y)$ —shows that our Hensel factorization is an improvement on Hensel-Sasaki (d^2k^2 versus $dM(d)k^2$) and an improvement on Kung and Traub’s method with linear lifting. If Conjecture 1 is true, then Hensel factorization can be within a factor of $\log k$ of Kung and Traub’s method with quadratic lifting. Nonetheless, this conjecture is highly encouraging where $k \gg d$. Proving this conjecture is left to future work.

5 Parallel Algorithms

Section 4 presented lazy algorithms for Weierstrass preparation, Taylor shift, and Hensel factorization. It also presented complexity estimates for those algorithms. Those estimates will soon be used to help dynamically distribute hardware resources (threads) in a parallel variation of Hensel factorization; in particular, a Hensel factorization pipeline where each pipeline stage updates one or more factors, see Algorithms 7–9. But first, we will examine parallel processing techniques for Weierstrass preparation.

5.1 Parallel Algorithms for Weierstrass Preparation

Algorithm 1 shows that p and α from a Weierstrass preparation can be updated in two phases: p in phase 1, and α in phase 2. Ultimately, these updates rely on the computation of the homogeneous part of some power series product. Algorithm 4 presents a simple map-reduce pattern (see Section 2.3) for computing such a homogeneous part. Moreover, this algorithm is designed such that, recursively, all ancestors of a power series product are also updated using parallelism.

Using the notation of Algorithm 1, recall that, e.g., $F_i = a_i - \sum_{j=0}^{i-1} (b_j c_{i-j})$, for $i \leq m$. Using lazy power series arithmetic, this entire formula can be encoded by a chain of ancestors, and one simply needs to update F_i to trigger a cascade of updates through its ancestors. In particular, using Algorithm 4, the homogeneous part of each product $b_j c_{i-j}$ is recursively computed using map-reduce. Similarly, Lemma 1 can be implemented using map-reduce (see Algorithm 5) to replace Lines

Algorithm 5 LEMMAFORWEIERSTRASS(k, f, g, h, t)

Input: $f, g, h \in \mathbb{K}[[X_1, \dots, X_n]]$ such that $f = gh$, $f_{(0)} = 0$, $h_{(0)} \neq 0$, f known to precision k , and g, h known to precision $k - 1$. $t \geq 1$ the number of threads to use.

Output: $g_{(k)}$.
1: $\mathcal{V} = [0, \dots, 0]$ ▷ 0-indexed list of size t
2: **parallel.for** $j = 0$ to $t - 1$
3: | **for** $i = jk/t + 1$ to $(j+1)k/t$ **while** $i < k$ **do**
4: | | $\mathcal{V}[j] := \mathcal{V}[j] + g_{(k-i)}h_{(i)}$
5: **end for**
6: **return** $(f_{(k)} - \sum_{j=0}^{t-1} \mathcal{V}[j]) / h_{(0)}$

Algorithm 6 WEIERSTRASSPHASE2PARALLEL(k, f, p, α, t)

Input: $f = \sum_{i=0}^{d+m} a_i Y^i$, $p = Y^d + \sum_{i=0}^{d-1} b_i Y^i$, $\alpha = \sum_{i=0}^m c_i Y^i$, $a_i, b_i, c_i \in \mathbb{K}[[X_1, \dots, X_n]]$ satisfying Theorem 1. b_0, \dots, b_{d-1} known modulo \mathcal{M}^{k+1} , c_0, \dots, c_m known modulo \mathcal{M}^k , for \mathcal{M} the maximal ideal of $\mathbb{K}[[X_1, \dots, X_n]]$. $t \geq 1$ for the number of threads to use.

Output: c_0, \dots, c_m known modulo \mathcal{M}^{k+1} , updated in-place.
1: $work := 0$
2: **for** $i = 1$ to m **do** ▷ estimate work using Observation 3, Corollary 1
3: | **if** $i \leq d$ **then** $work := work + (i - 1)$
4: | **else** $work := work + d$
5: $t' := 1$; $targ := work / t$
6: $work := 0$; $j := 1$
7: $\mathcal{I} := [-1, 0, \dots, 0]$ ▷ 0-indexed list of size $t + 1$
8: **for** $i = 1$ to m **do**
9: | **if** $i \leq d$ **then** $work := work + (i - 1)$
10: | **else** $work := work + d$
11: | **if** $work \geq targ$ **then**
12: | | $\mathcal{I}[j] := i$; $work := 0$; $j := j + 1$
13: **if** $j \leq t$ **and** $t' < 2$ **then** ▷ work did not distribute evenly; try again with $t/2$ threads
14: | $t = t / 2$; $t' = 2$
15: | **goto** Line 6
16: **else if** $j \leq t$ **then** ▷ still not evenly distributed, use all threads in UPDATETODEGPARALLEL
17: | $\mathcal{I}[1] = m$; $t' = 2t$; $t = 1$
18: **parallel.for** $\ell = 1$ to t
19: | **for** $i = \mathcal{I}[\ell - 1] + 1$ to $\mathcal{I}[\ell]$ **do**
20: | | UPDATETODEGPARALLEL(k, c_{m-i}, t')

9–12 of Algorithm 1. Phase 1 of Weierstrass, say WEIERSTRASSPHASE1PARALLEL, thus reduces to a loop over i from 0 to $d - 1$, calling Algorithm 4 to update F_i to precision k , and calling Algorithm 5 to compute $b_{i(k)}$.

Algorithm 4 uses several simple subroutines: FIRSTANCESTOR gets the first ancestor of a power series, SECONDANCESTOR gets the second ancestor of a power series, and UPDATETODEG calls the update function of a lazy power series to ensure its precision is at least k ; see Section 3.

Now consider phase 2 of WEIERSTRASSUPDATE. Notice that computing the homogeneous part of degree k for c_{m-i} , $0 \leq i \leq m$ only requires each c_{m-i} to be known up to precision $k - 1$, since each $b_j \in \mathcal{M}$ for $0 \leq j < d$. This implies that the phase 2 **for** loop of WEIERSTRASSUPDATE has independent iterations. We can thus apply the map pattern directly to this loop itself, rather than relying on the map-reduce pattern to compute the homogeneous part of a product power series. However, consider the following two facts: the cost of computing each c_{m-i} is different (Observation 3), and, for a certain number of available threads t , it may be impossible to partition the iterations of the loop into t partitions of equal work. Yet, partitioning the loop itself is preferred for greater parallelism.

Hence, for phase 2, a dynamic decision is made to either apply the map pattern to the loop over c_{m-i} , or to apply the map pattern within UPDATETODEGPARALLEL

Algorithm 7 HENSELPIPELINESTAGE(k, f_i, t, GEN)

Input: An positive integer k , $f_i = Y^{d_i} + \sum_{i=0}^{d_i-1} a_i Y^i$, $a_i \in \mathbb{K}[[X_1, \dots, X_n]]$. A positive integer t the number of threads to use within this stage. GEN an generator for the previous pipeline stage.

Output: a sequence of integers j signalling f_i is known to precision j . This sequence ends with k .

```

1:  $p := \text{PRECISION}(f_i)$  ▷ get the current precision of  $f_i$ 
2: do
3:    $k' := \text{GEN}()$  ▷ A blocking function call until GEN yields
4:   for  $j = p$  to  $k'$  do
5:      $\text{UPDATE\_TODEGPARALLEL}(j, f_i, t)$ 
6:     yield  $j$ 
7:    $p := k'$ 
8: while  $k' < k$ 

```

for each c_{m-i} , or both. This decision process is detailed in Algorithm 6, where t partitions of equal work try to be found to apply the map pattern to only the loop itself. If unsuccessful, $t/2$ partitions of equal work try to be found with 2 threads given to each partition for UPDATE_TODEGPARALLEL. If that, too, is unsuccessful, then each c_{m-i} is updated one at a time using UPDATE_TODEGPARALLEL and the total number of threads t .

5.2 Parallel Algorithms for Hensel Factorization

Let $f = f_1 \cdots f_r$ be a Hensel factorization where the factors have respective degrees d_1, \dots, d_r . From Algorithm 3 and Figure 1, we have already seen that the repeated applications of Taylor shift and Weierstrass preparation naturally form a chain of ancestors, and thus a pipeline. Using the notation of Algorithm 3, updating f_1 requires updating g_1, p_1, α_1 . Then, updating f_2 requires updating $\hat{f}_2, g_2, p_2, \alpha_2$, and so on. These groups easily form stages of a pipeline, where updating f_1 to degree $k-1$ is a prerequisite for updating f_2 to degree $k-1$. Moreover, meanwhile f_2 is being updated to degree $k-1$, f_1 can simultaneously be updated to degree k . Such a pattern holds for all successive factors.

Algorithms 7 and 8 show how the factors of a Hensel factorization can all be simultaneously updated to degree k using asynchronous generators, forming the so-called *Hensel pipeline*. Algorithm 7 shows a single pipeline stage as an asynchronous generator, which itself consumes data from another asynchronous generator—just as expected from the pipeline pattern. Algorithm 8 shows the creation, and joining in sequence, of those generators. The key feature of these algorithms is that a generator (say, stage i) produces a sequence of integers (j) which signals to the consumer (stage $i+1$) that the previous factor has been computed up to precision j and thus the required data is available to update its own factor to precision j .

Notice that Algorithm 8 still follows our lazy evaluation scheme. Indeed, the factors are updated all at once up to precision k , starting from their current precision. However, for optimal performance, the updates should be applied for large increases in precision, rather than repeatedly increasing precision by one.

Further considering performance, Theorem 5 showed that the expected cost for updating each factor from a Hensel factorization is different. In particular, for $\hat{d}_i = \sum_{j=i}^r d_j$, updating factor f_i scales as $d_i \hat{d}_{i+1} k^2$. The work for each stage of the proposed pipeline is unequal and the pipeline is unlikely to achieve good parallel speedup. However, Corollary 2 shows that the work ratios between stages do not change for increasing k , and thus a static scheduling scheme is sufficient.

Algorithm 8 HENSELFACTORIZATIONPIPELINE($k, \mathcal{F}, \mathcal{T}$)

Input: A positive integer k , $\mathcal{F} = \{f_1, \dots, f_r\}$, the output of HENSELFACTORIZATION. $\mathcal{T} \in \mathbb{Z}^r$ a 0-indexed list of the number of threads to use in each stage, $\mathcal{T}[r-1] > 0$.
Output: f_1, \dots, f_r updated in-place to precision k .

```
1: GEN := () → {yield k}                                ▷ An anonymous function asynchronous generator
2: for i = 0 to r - 1 do
3:   if  $\mathcal{T}[i] > 0$  then
4:     ▷ Capture HENSELPIPELINESTAGE( $k, f_{i+1}, \mathcal{T}[i], \text{GEN}$ ) as a
      function object, passing the previous GEN as input
     GEN := ASYNCGENERATOR(HENSELPIPELINESTAGE,  $k, f_{i+1}, \mathcal{T}[i], \text{GEN}$ )
5: do
6:   |  $k' := \text{GEN}()$                                     ▷ ensure last stage completes before returning
7: while  $k' < k$ 
```

Algorithm 9 DISTRIBUTERESOURCESHENSEL(\mathcal{F}, t_{tot})

Input: $\mathcal{F} = \{f_1, \dots, f_r\}$ the output of HENSELFACTORIZATION. $t_{tot} > 1$ the total number of threads.
Output: \mathcal{T} , a list of size r , where $\mathcal{T}[i]$ is the number of threads to use for updating f_{i+1} . The number of positive entries in \mathcal{T} determines the number of pipeline stages. $\mathcal{T}[i] = 0$ encodes that f_{i+1} should be computed within the same stage as f_{i+2} .

```
1:  $\mathcal{T} := [0, \dots, 0, 1]$ ;  $t := t_{tot} - 1$                 ▷  $\mathcal{T}[r-1] = 1$  ensures last factor gets updated
2:  $d := \sum_{i=1}^r \deg(f_i)$ 
3:  $\mathcal{W} := [0, \dots, 0]$                                     ▷ A 0-indexed list of size  $r$ 
4: for i = 1 to r - 1 do
5:   |  $\mathcal{W}[i-1] := \deg(f_i)(d - \deg(f_i))$                 ▷ Estimate work by Theorem 5,  $d_i \widehat{d}_{i+1}$ 
6:   |  $d := d - \deg(f_i)$ 
7:  $totalWork := \sum_{i=0}^{r-1} \mathcal{W}[i]$ 
8:  $ratio := 0$ ;  $targ := 1 / t$ 
9: for i = 0 to r do
10:  |  $ratio := ratio + (\mathcal{W}[i] / totalWork)$ 
11:  | if  $ratio \geq targ$  then
12:    |  $\mathcal{T}[i] := \text{ROUND}(ratio \cdot t)$ ;  $ratio := 0$ 
13:  $t := t_{tot} - \sum_{i=0}^{r-1} \mathcal{T}[i]$                             ▷ Give any excess threads to the earlier stages
14: for i = 0 to r - 1 while  $t > 0$  do
15:   |  $\mathcal{T}[i] := \mathcal{T}[i] + 1$ ;  $t := t - 1$ 
16: return  $\mathcal{T}$ 
```

Notice that Algorithm 7 takes a parameter t for the number of threads to use. As we have seen in Section 5.1, the update for Weierstrass preparation can also be done in parallel. Consequently, each stage of the Hensel pipeline is augmented to further exploit such parallelism. We have thus composed the two pattern schemes, applying map-reduce within each stage of the parallel pipeline. This composition serves to load-balance the pipeline. For example, the first stage may be given t_1 threads and the second stage given t_2 threads, with $t_1 > t_2$, so that the two stages may execute in nearly equal time.

To further encourage load-balancing, each stage of the pipeline need not update a single factor, but rather a group of successive factors. Algorithm 9 applies Theorem 5 to attempt to load-balance each stage s of the pipeline through assigning a certain number of threads t_s and a certain group of factors f_{s_1}, \dots, f_{s_2} to it. The goal is for $\sum_{i=s_1}^{s_2} d_i \widehat{d}_{i+1} / t_s$ to be roughly equal for each stage.

6 Experimentation and Discussion

The previous section introduced parallel schemes for Weierstrass preparation and Hensel factorization based on the composition of the map-reduce and pipeline parallel patterns. Our lazy power series, along with those parallel schemes have been implemented in C/C++ as part of the Basic Polynomial Algebra Subprograms

(BPAS) library [2]. These parallel algorithms are implemented using generic support for task parallelism, thread pools, and asynchronous generators, also provided in the BPAS library. The details of this parallel support are discussed in [3] and [4].

Our experimentation was collected on a compute node running Ubuntu 18.04.4 with two Intel Xeon X5650 processors, each with 6 cores (12 cores total) at 2.67 GHz, and a 12x4GB DDR3 memory configuration at 1.33 GHz. BPAS was compiled using GMP 6.1.2 [8]. We work over \mathbb{Q} since these examples do not require algebraic numbers. All data shown is an average of 3 trials.

We begin by evaluating Weierstrass preparation for two families of examples:

$$(i) \quad u_r = \sum_{i=2}^r (X_1^2 + X_2 + i)Y^i + (X_1^2 + X_2)Y + X_1^2 + X_1X_2 + X_2^2$$

$$(ii) \quad v_r = \sum_{i=\lceil r/2 \rceil}^r (X_1^2 + X_2 + i)Y^i + \sum_{i=1}^{\lceil r/2 \rceil - 1} (X_1^2 + X_2)Y^i + X_1^2 + X_1X_2 + X_2^2$$

Applying Weierstrass preparation to u_r results in p with degree 2. Applying Weierstrass preparation to v_r results in p with degree $\lceil r/2 \rceil$. Fig. 2 summarizes the resulting execution times and parallel speedups. Generally, speedup increases with increasing degree in Y and increasing precision computed.

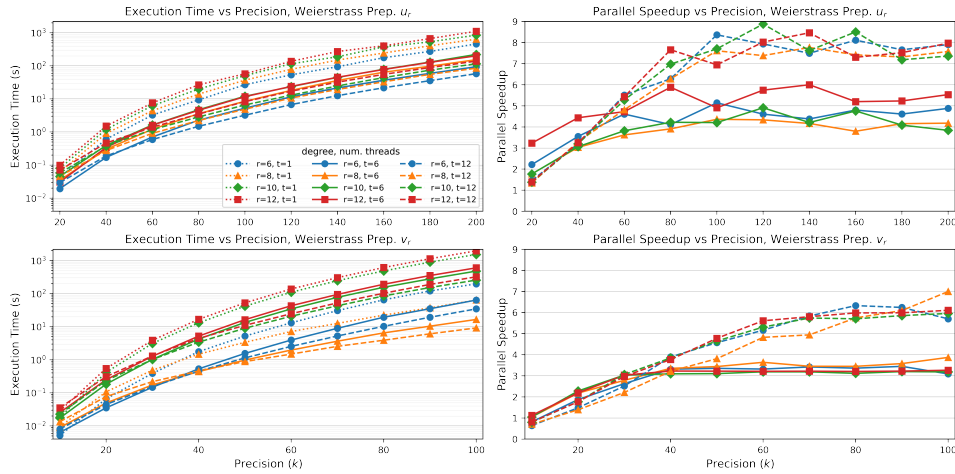


Figure 2: Comparing Weierstrass preparation of u_r and v_r for $r \in \{6, 8, 10, 12\}$ using 1, 6, and 12 threads. First column: execution time of u_r and v_r ; second column: parallel speedup of u_r and v_r . Profiling of v_6 shows that its exceptional relative performance is attributed to remarkably good branch prediction.

Within Weierstrass preparation, recall that parallelism arises in our implementation in two main ways: computing summations of products of homogeneous parts (the `parallel_for` loops in `UPDATETODEGPARALLEL` and `LEMMAFORWEIERSTRASS`), and the `parallel_for` loop over updating c_{m-i} in `WEIERSTRASSPHASE2PARALLEL`. The former has an inherent limitation: computing a multivariate product with one operand of low degree and one operand of high degree is much easier than computing one where both operands are of moderate degree. Evenly partitioning the iterations of the loop thus does not result in even work per thread. This is evident in comparing the parallel speedup between u_r and v_r ; the former, with higher degree in α , relies less on parallelism coming from those products. Better partitioning is thus needed and is left to future work.

We evaluate our parallel Hensel factorization using three families of problems:

$$(i) \ x_r = \prod_{i=1}^r (Y - i) + X_1(Y^3 + Y)$$

$$(ii) \ y_r = \prod_{i=1}^r (Y - i)^i + X_1(Y^3 + Y)$$

$$(iii) \ z_r = \prod_{i=1}^r (Y + X_1 + X_2 - i) + X_1 X_2 (Y^3 + Y)$$

Denote by f_1, \dots, f_r the factors of any one of these UPoPS. Let f_i be the factor with i as the root of f_i .

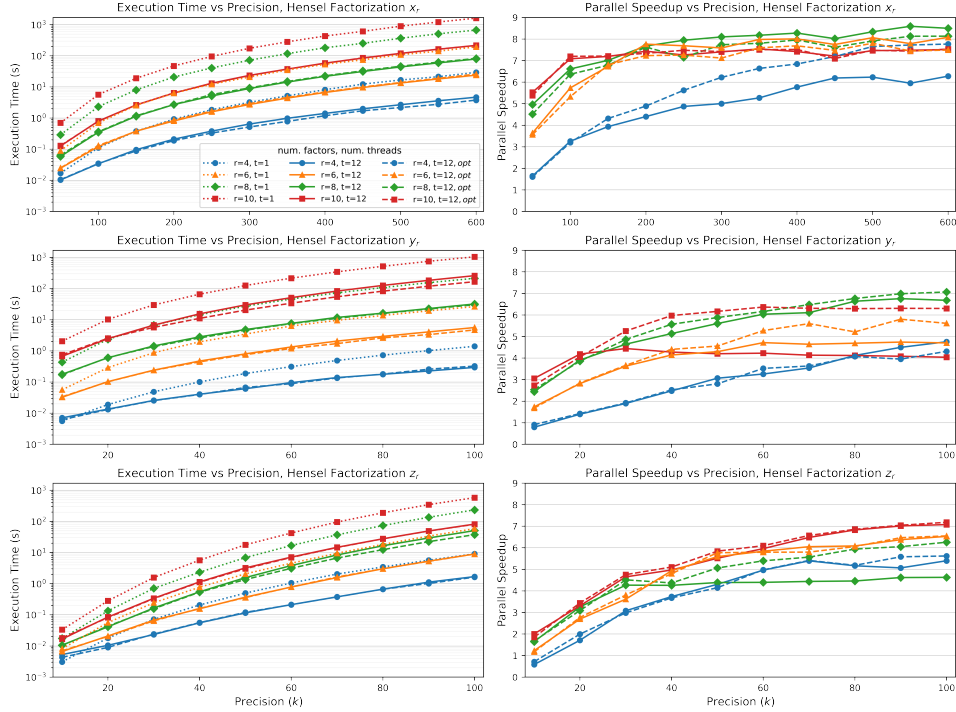


Figure 3: Comparing parallel Hensel factorization for x_r , y_r , and z_r for $r \in \{4, 6, 8, 10\}$. First column: execution time; second column: parallel speedup. For $t = 12$ resource distribution is determined by Algorithm 9; for $t = 12, opt$ serial execution time, rather than complexity estimates, estimates work in Algorithm 9.

Despite the inherent challenges of irregular parallelism arising from stages with unequal work, the composition of parallel patterns allows for load-balancing between stages and the overall pipeline to achieve relatively good parallel speed-up. Fig. 3 summarizes these results while Table 1 presents the execution time per factor (or stage, in parallel). Generally speaking, potential parallelism increases with increasing degree and increasing precision.

The distribution of discrete threads to a discrete number of pipeline stages is a challenge; a perfect distribution requires a fractional number of threads per stage. Nonetheless, in addition to the distribution technique presented in Algorithm 9, we can also examine hand-chosen assignments of threads to stages. One can first determine the time required to update each factor in serial, say for some small k ,

and then use that execution time as the work estimates in Algorithm 9, rather than using the complexity estimates. This latter technique is depicted in Fig. 3 as *opt* and in Table 1 as Time-est. threads. This technique is still not perfect, again because of the discrete nature of threads, and the imperfect parallelization of computing summations of products of homogeneous parts.

	factor	serial time (s)	shift time (s)	Complexity- Est. threads	parallel time (s)	wait time (s)	Time-est. threads	parallel time (s)	wait time (s)	
x_4	$k = 600$	f_1	18.1989	0.0012	6	4.5380	0.0000	7	3.5941	0.0000
		f_2	6.6681	0.0666	4	4.5566	0.8530	3	3.6105	0.6163
		f_3	3.4335	0.0274	1	4.5748	1.0855	0	-	-
		f_4	0.0009	0.0009	1	4.5750	4.5707	2	3.6257	1.4170
	<i>totals</i>	28.3014	0.0961	12	4.5750	6.5092	12	3.6257	2.0333	
y_4	$k = 100$	f_1	0.4216	0.0003	3	0.1846	0.0000	4	0.1819	0.0000
		f_2	0.5122	0.0427	5	0.2759	0.0003	4	0.3080	0.0001
		f_3	0.4586	0.0315	3	0.2842	0.0183	0	-	-
		f_4	0.0049	0.0048	1	0.2844	0.2780	4	0.3144	0.0154
	<i>totals</i>	1.3973	0.0793	12	0.2844	0.2963	12	0.3144	0.0155	
z_4	$k = 100$	f_1	5.2455	0.0018	6	1.5263	0.0000	7	1.3376	0.0000
		f_2	2.5414	0.0300	4	1.5865	0.2061	3	1.4854	0.0005
		f_3	1.2525	0.0151	1	1.6504	0.1893	0	-	-
		f_4	0.0018	0.0018	1	1.6506	1.6473	2	1.5208	0.7155
	<i>totals</i>	9.0412	0.0487	12	1.6506	2.0427	12	1.5208	0.7160	

Table 1: Times for updating each factor within the Hensel pipeline. Complexity-estimated threads use complexity estimates to estimate work within Algorithm 9; time-estimated threads use the serial execution time to estimate work and distribute threads. Wait time indicates the stage spent waiting on the previous.

In future, we must consider several important factors to improve performance. Relaxed algorithms should give better complexity and performance. For parallelism, better partitioning schemes for the map-reduce pattern within Weierstrass preparation should be considered. Finally, for the Hensel pipeline, more analysis is needed to optimize the scheduling and resource distribution, particularly considering coefficient sizes and the multivariate case.

Acknowledgements

The authors would like to thank NSERC of Canada (award CGSD3-535362-2019).

References

- [1] Parisa Alvandi, Masoud Ataei, Mahsa Kazemi, and Marc Moreno Maza. On the extended hensel construction and its application to the computation of real limit points. *J. Symb. Comput.*, 98:120–162, 2020.
- [2] M. Asadi, A. Brandt, C. Chen, S. Covanov, M. Kazemi, F. Mansouri, D. Mohajerani, R. H. C. Moir, M. Moreno Maza, D. Talaashrafi, Linxiao Wang, Ning Xie, and Yuzhen Xie. Basic Polynomial Algebra Subprograms (BPAS), 2021. www.bpaslib.org.
- [3] Mohammadali Asadi, Alexander Brandt, Robert H. C. Moir, Marc Moreno Maza, and Yuzhen Xie. On the parallelization of triangular decompositions. In

- ISSAC '20: International Symposium on Symbolic and Algebraic Computation*, pages 22–29. ACM, 2020.
- [4] Mohammadali Asadi, Alexander Brandt, Robert H. C. Moir, Marc Moreno Maza, and Yuzhen Xie. Parallelization of triangular decompositions: Techniques and implementation. *J. Symb. Comput.*, 2021. (to appear).
 - [5] Alexander Brandt, Mahsa Kazemi, and Marc Moreno Maza. Power series arithmetic with the BPAS library. In *Computer Algebra in Scientific Computing (CASC) 2020*, volume 12291 of *LNCS*, pages 108–128. Springer, 2020.
 - [6] David V Chudnovsky and Gregory V Chudnovsky. On expansion of algebraic functions in power and puiseux series, i. *Journal of Complexity*, 2(4):271–294, 1986.
 - [7] G. Fischer. *Plane Algebraic Curves*. AMS, 2001.
 - [8] Torbjørn Granlund et al. *GNU MP 6.0 Multiple Precision Arithmetic Library*. Samurai Media Limited, 2015.
 - [9] Azzam Haidar, Jakub Kurzak, and Piotr Luszczek. An improved parallel singular value algorithm and its implementation for multicore hardware. In *International Conference for High Performance Computing, Networking, Storage and Analysis, SC'13*, pages 90:1–90:12. ACM, 2013.
 - [10] H. T. Kung and Joseph F. Traub. All algebraic functions can be computed fast. *J. ACM*, 25(2):245–260, 1978.
 - [11] M. McCool, J. Reinders, and A. Robison. *Structured parallel programming: patterns for efficient computation*. Elsevier, 2012.
 - [12] Panagiotis D. Michailidis and Konstantinos G. Margaritis. Parallel direct methods for solving the system of linear equations with pipelining on a multicore using openmp. *J. Comput. Appl. Math.*, 236(3):326–341, 2011.
 - [13] T. Sasaki and F. Kako. Solving multivariate algebraic equation by Hensel construction. *Japan J. Indust. and Appl. Math.*, 1999.
 - [14] Michael L. Scott. *Programming Language Pragmatics (3. ed.)*. Academic Press, 2009.
 - [15] Joris van der Hoeven. Relax, but don't be too lazy. *J. Symb. Comput.*, 34(6):479–542, 2002.
 - [16] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, NY, USA, 2 edition, 2003.