Factoring via Power Series Methods: Towards VP Closure in Positive Characteristic

A thesis Submitted
in Partial Fulfilment of the Requirements
for the Degree of
MASTER OF SCIENCE

by

Bhaskar Goyal

to the

School of Mathematical Sciences

National Institute of Science Education and Research

Bhubaneswar

2026

DECLARATION

I hereby declare that I am the sole author of this thesis in partial fulfilment of the requirements for a postgraduate degree from National Institute of Science Education and Research (NISER). I authorize NISER to lend this thesis to other institutions or individuals for the purpose of scholarly research.

.....
Signature of the Student
Date:

The thesis work reported in the thesis entitled *Factoring via Power Series Methods: Towards VP Closure in Positive Characteristic* was carried out **under our supervision**.

Signature of Supervisor

.....
School of Mathematical Sciences,
NISER, Bhubaneswar
Date:

Signature of Co-supervisor

.....
Department of Computer Science and Engineering,
IIT Kanpur
Date: 07-May'26

ABSTRACT

This report investigates the closure of algebraic complexity classes under factorization, specifically focusing on the long-standing open problem: whether the factors of a polynomial with circuit complexity s can be represented by circuits of size $poly(s)$ in positive characteristic. While this property is well-established for fields of characteristic zero [Kal87], the case of characteristic p remains unsolved for the closure property of the class VP.

We first present a rigorous proof for factor reconstruction in general fields using the power series techniques developed by Dutta, Saxena, and Sinhababu [DSS22], which utilize Newton iteration to handle factors of high multiplicity. Following this, we explore a new approach to the characteristic p problem by shifting the analysis from standard power series to the p-adic setting. This method attempts to bypass the traditional roadblocks of vanishing derivatives and p-th powers by lifting the computation to a domain where characteristic-zero techniques might be applied. However, we demonstrate that our approach ultimately fails to resolve the problem. This report provides a detailed analysis of the technical barriers that continue to make the positive characteristic case an elusive challenge in Algebraic Complexity Theory.

ACKNOWLEDGEMENTS

First and foremost, I would like to express my deepest gratitude to my primary supervisor, *Prof. Nitin Saxena*, for his unique mentorship, and the immense trust he placed in me. By giving me the opportunity to work on an open problem alongside him, he gave me an invaluable experience; one where I had full exploratory freedom, while also having the chance to learn from one of the best in the field. I was constantly amazed by his ability to quickly get to the heart of the problem; it would take me days to truly understand the insights he shared within minutes. I am also grateful to *Foram Lakhani*, who was my collaborator on this project. Not only did we spend countless hours working together, but she also taught me some of the much needed basics of arithmetic circuit complexity.

I am also grateful to *Dr. Deepak Dalai*, who made my visit to Kanpur possible by agreeing to be my secondary supervisor. I thank the UGCS committee of my department, and the academic section at NISER for letting me take a semester off to work on this project. I thank all the great teachers I have had at NISER, especially *Dr. Aritra Banik*, for making me fall in love with the subject, and for providing mentorship and support whenever I needed it.

But most of all, I am grateful to my family and friends, who have always been there for me. I cannot fathom what my life would have been without *Ghugniland* — every single one of you has been a source of joy and support, and I am so lucky to have you all in my life. I am grateful to my parents for their unwavering support and encouragement, and for always believing in me. It is their belief in me that has given me the confidence to pursue whatever I set my mind to, and I am forever grateful for that.

Contents

1 Preliminaries **1**
- 1.1 Introduction . . . 1
- 1.2 Algebraic Model of Computation and Complexity Classes . . . 2
- 1.3 The Factoring Problem in VP . . . 5
- 1.4 Organization of the Report . . . 6

2 Technical Tools for Circuit Factoring **8**
- 2.1 Notation . . . 8
- 2.2 Algebraic Primitives for Factoring . . . 9
  - 2.2.1 Schwartz-Zippel Lemma . . . 9
  - 2.2.2 Newton Iteration . . . 9
  - 2.2.3 Pre-processing . . . 11
- 2.3 Structural Results about Arithmetic Circuits . . . 14

3 Proof for VP closure **16**
- 3.1 Reduction of Factoring to Power Series Root Approximation . . . 16
- 3.2 Proof of the Main Theorem . . . 19
- 3.3 Fields that are not Algebraically Closed . . . 21
  - 3.3.1 Resultants and Norm . . . 22

4 Characteristic p **24**
- 4.1 Some p-adic analysis . . . 25
- 4.2 An Approach by Interpreting the circuit over $\hat{\mathbb{Q}}_p$ . . . 29
  - 4.2.1 Counterexample to Integral Power Series Solutions . . . 31
- 4.3 Conclusion . . . 32

List of Figures

1.1 An algebraic circuit for $f = x_1^2 x_2 + x_1 x_2^2$ using 2 multiplication gates and 1 addition gate. Labels in grey show the polynomial computed at each gate. Edges carry the value computed by their source gate. . . . 4

4.1 Newton polygon of $f(y) = (y - p)(y - p^2)(y - p^3)$. All four coefficient points lie on the lower convex hull. Each segment has horizontal length 1 and slope equal to the negative of the p-adic valuation of the corresponding root. 29

4.2 Newton polygon of $F(0, y) = y^{p^i} + g(0)^{p^i} + p\,E'(0, y)$, where $\deg_y F \geq p^i$. The bold blue segment is the *first segment* of the Newton polygon, running from $(0, 0)$ to $(p^i, 0)$ with slope 0 and horizontal length p^i. Points between 0 and p^i from $p\,E'(0, y)$ have $v_p \geq 1$ and lie strictly above the x-axis; points with zero coefficient are omitted ($v_p = +\infty$). The dashed blue line indicates the polygon may continue beyond p^i up to $\deg_y F$, with further points from E' floating above; this is irrelevant to the first-break analysis. By Lemma 4.6, $F(0, y)$ has exactly p^i roots in $\overline{\mathbb{Q}_p}$ with p-adic valuation 0. Since the first segment has length $p^i > 1$, these roots cannot be separated by valuation, obstructing the lifting step. 31

Chapter 1

Preliminaries

1.1 Introduction

How complex is a given polynomial? This is the fundamental question at the heart of Algebraic Complexity Theory. Instinctively, one might measure the complexity of a polynomial by its degree, its number of variables, or the number of monomials. However, these classical measures often fail to capture the actual computational effort required to evaluate the polynomial.

A polynomial with a vast number of monomials may be efficiently computable if it possesses an underlying structural simplicity. For example, consider the polynomial:

$$f(x_1, \ldots, x_n) = (x_1 + x_2 + \cdots + x_n)^{100}$$

When expanded, f consists of $\binom{n+99}{100}$ monomials—a number that grows rapidly with n. Yet, f is computationally "easy" to evaluate, requiring only $n - 1$ additions and a few multiplications via repeated squaring. Conversely, a polynomial with a sparse representation can be comparatively difficult to compute. The monomial $g(x_1, \ldots, x_n) = x_1 x_2 \cdots x_n$ contains only one term, but any sequence of multiplications computing it requires at least $n - 1$ operations.

This discrepancy necessitates a robust model of computation that accounts for the sequence of arithmetic operations rather than the final expanded form. To this end, we will introduce the *Arithmetic Circuit*, which provides the formal framework for measuring the complexity of polynomials in terms of the minimum number of

fundamental operations required for their construction.

While understanding how to compute a polynomial is the starting point, this report looks at the problem of factorization through this lens. If a polynomial is "easy" to compute, are its factors also "easy" to compute? This simple question is very important in Algebraic Complexity Theory.

Before addressing the complexity of factorization, we must first establish a rigorous mathematical framework to quantify "ease of computation". The following sections define the formal model and the complexity classes that form the landscape of this study.

1.2 Algebraic Model of Computation and Complexity Classes

We first formally introduce *Arithmetic Circuit* (or *Algebraic Circuit*).

Definition 1.1. *(Algebraic Circuit) An algebraic circuit C over a field $\mathbb{F}$ is a directed acyclic graph with the following nodes:*

- *The input gates, with in-degree* 0*, are labelled by variables like $x_1, x_2, \ldots, x_n$.*
- *The internal gates, with in-degree at least* 1*, are labelled by either $+$ or $\times$.*
- *A unique gate with outdegree* 0*, called the output gate.*
- *A constant gate, with in-degree* 0*, is labelled by a field element from $\mathbb{F}$.*

An algebraic circuit computes a polynomial in the following way: the polynomial computed at an input gate is the variable it is labelled with, the polynomial computed at a constant gate is the field element it is labelled with, and the polynomial computed at an internal gate is the result of applying the operation (addition or multiplication)

to the polynomials computed by its children. The polynomial computed by the circuit is the value at the output gate.

*The **size** of an algebraic circuit is the number of gates in the circuit. The **depth** of an algebraic circuit is the length of the longest path from an input gate to the output gate.*

Arithmetic circuits are arguably the most natural model to study the complexity of a polynomial; they describe how a polynomial can be computed from the variables and field elements via fundamental operations like addition and multiplication. The complexity of a polynomial f can then be defined as the *size of the smallest circuit* computing it.

Example 1.2 (An algebraic circuit for $x_1^2x_2 + x_1x_2^2$)**.** *Consider the polynomial*

$$f(x_1, x_2) \;=\; x_1^2x_2 + x_1x_2^2 \;=\; x_1x_2(x_1 + x_2).$$

A naive evaluation would compute $x_1^2x_2$ and $x_1x_2^2$ separately and add them, requiring five multiplications and one addition. The factored form $x_1x_2(x_1 + x_2)$ suggests a more efficient circuit using only two multiplications and one addition, which we now describe explicitly.

To allow for a fair comparison among polynomials, it is logical to describe their complexity in terms of their fundamental parameters, such as the number of variables n or the degree d. To this end, we utilize the notion of *polynomial families*, which are sequences $(f_n)_{n\in\mathbb{N}}$ where f_n is a multivariate polynomial over some field $\mathbb{F}$. We study the complexity of these families with respect to the growth in terms of the parameter n.

For example, the determinant of an $n \times n$ matrix is a polynomial of degree n in n^2 variables (the entries of the matrix). We refer to this as the *determinant*

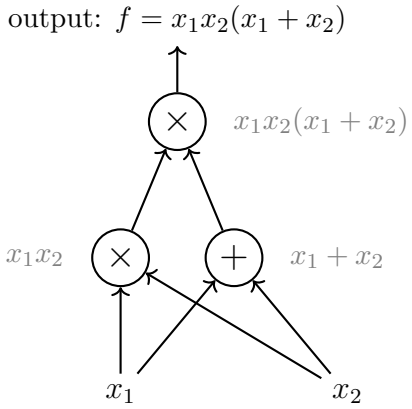

Figure 1.1: An algebraic circuit for $f = x_1^2x_2 + x_1x_2^2$ using 2 multiplication gates and 1 addition gate. Labels in grey show the polynomial computed at each gate. Edges carry the value computed by their source gate.

family, denoted by $\{Det_n\}_{n\in\mathbb{N}}$. *Algebraic Complexity Classes* consist of such families of polynomials and are used to categorize the asymptotic complexity of a polynomial family.

We now define the class of efficiently computable polynomial families, denoted by VP. This class is the centrepiece of this report and serves as the algebraic analogue of the class P in Boolean complexity theory.

Definition 1.3. *A family of polynomials $(f_n)_{n\in\mathbb{N}}$ is in* VP *if the degree of f_n is poly(n), and f_n can be computed by an algebraic circuit of size poly(n) for all large enough n.*

Example 1.4. *(elementary symmetric polynomials) The k-th elementary symmetric polynomial in n variables, denoted by $e_k(x_1, x_2, \ldots, x_n)$, is defined as the sum of all products of k distinct variables. For example, $e_2(x_1, x_2, x_3) = x_1x_2 + x_1x_3 + x_2x_3$. To see that the family of elementary symmetric polynomials is in* VP*, we can use the following recursive relation:*

$$e_k(x_1, x_2, \ldots, x_n) = e_k(x_1, x_2, \ldots, x_{n-1}) + x_n e_{k-1}(x_1, x_2, \ldots, x_{n-1}).$$

This relation allows us to compute e_k using a circuit of size $O(n^2)$, which is polynomial

in n. Therefore, the family of elementary symmetric polynomials is in VP.

1.3 The Factoring Problem in VP

A natural question to ask is: how robust are these complexity classes under different algebraic operations? Is the property of efficient computability inherited by the factors of a polynomial? Specifically, if a polynomial f can be computed by a circuit of size s, can its factors also be computed by circuits of size polynomial in s and the degree of f?

Remark 1.5. The polynomial $f = x_1x_2(x_1 + x_2)$ in Example 1.2 has two irreducible factors over any field: x_1x_2 and $x_1 + x_2$ (or equivalently the three linear factors x_1, x_2, $x_1 + x_2$ if we allow non-monic factorisation). Observe that each factor is computed as an intermediate value at a gate of Φ — namely g_2 computes x_1x_2 and g_1 computes x_1+x_2. This is of course specific to this example; in general the factors of a polynomial need not appear as intermediate values of a circuit computing it, and constructing circuits for the factors from a circuit for f is precisely the content of this report.

It is quite intuitive to expect that if a polynomial can be computed efficiently, then its factors should also be computable efficiently. If this were not the case, it would imply that there are polynomials that cannot be computed efficiently, but their product can be computed efficiently, which would be a rather strange phenomenon. However, proving this result is non-trivial.

The following definition formalizes this notion into the closure of a complexity class under factoring.

Definition 1.6. *A complexity class $\mathcal{C}$ is said to be* ***closed under factoring*** *if for every family of polynomials $(f_n)_{n\in\mathbb{N}}$ in $\mathcal{C}$, and for every factor g_n of f_n, the family $(g_n)_{n\in\mathbb{N}}$ is also in $\mathcal{C}$.*

In a series of influential papers, Kaltofen [Kal86, Kal87, Kal89, KT90] established that for characteristic zero fields (or fields of sufficiently large characteristic), the class VP is indeed closed under factoring. However, the question of whether VP is closed under factoring over fields of small positive characteristic remains open. This is a significant open problem in algebraic complexity theory, and this report is dedicated to this problem. We will explore the known results, the techniques used in the proofs, and the challenges that arise when trying to extend these results to fields of small characteristic.

We conclude this section by formally stating the closure theorem for VP for characteristic zero fields. We will prove this in chapter 3. Unless stated otherwise, $\mathbb{F}$ will denote an algebraically closed field of characteristic zero, and $\mathbf{x}$ will denote the set of variables $\{x_1, x_2, \ldots, x_n\}$.

Theorem 1.7 (Kaltofen [Kal87])**.** *Let $f \in \mathbb{F}[\mathbf{x}]$ be an n-variate, degree d polynomial that can be computed by a circuit of size s. Let g be a polynomial such that g divides f. Then g can be computed by a circuit of size at most poly(s, n, d).*

Open Question 1.8. Does Theorem 1.7 hold when $\mathbb{F} = \mathbb{F}_q$, a field of small positive characteristic?

1.4 Organization of the Report

In chapter 2, we will introduce the main tools and techniques that we will use in the proof of Theorem 1.7. This will include some structural results related to arithmetic circuits, along with some necessary algebraic preliminaries. In chapter 3, we will present the proof of Theorem 1.7 for characteristic zero fields via an approach from [DSS22]. Finally, in chapter 4, we will discuss the challenges and open problems related to extending this result to fields of small positive characteristic, and analyse a

proposed approach via p-adic methods.

Chapter 2

Technical Tools for Circuit Factoring

Before we move on to the actual proof of the closure of VP under factoring, we need to develop some machinery for getting our polynomial into a nice form from which we can compute factors easily. This chapter mostly contains simple mathematical lemmas that will be used in the proof of Theorem 1.7.

First, we fix some notation that we will use throughout this report.

2.1 Notation

- Lowercase bold letters like $\mathbf{x}, \mathbf{a}, \mathbf{b}, \boldsymbol{\alpha}$ will denote vectors like $x_1, x_2, ..., x_n$.
- $\mathbb{F}$ will denote the base field over which we are working. Unless otherwise specified, we will assume that $\mathbb{F}$ is algebraically closed and has characteristic 0.
- $\mathbb{F}[\mathbf{x}]$ will denote the ring of polynomials in the variables $\mathbf{x}$ over $\mathbb{F}$.
- $\mathbb{F}(\mathbf{x})$ will denote the field of rational functions in the variables $\mathbf{x}$ over $\mathbb{F}$.
- $\mathbb{F}[[\mathbf{x}]]$ will denote the ring of formal power series in the variables $\mathbf{x}$ over $\mathbb{F}$.
- $\langle \mathbf{x} \rangle^k$ will denote the ideal generated by all monomials of total degree at least k in the variables $\mathbf{x}$. So $\langle \mathbf{x} \rangle^k$ is the set of all polynomials with no monomial of degree less than k.
- $[n]$ will denote the set $\{1, 2, ..., n\}$.

- $Rad(f)$ will denote the radical of a polynomial f, which is the product of all distinct irreducible factors of f.
- $\partial_y f$ will denote the partial derivative of f with respect to y.

2.2 Algebraic Primitives for Factoring

2.2.1 Schwartz-Zippel Lemma

The Schwartz-Zippel lemma is a fundamental tool in algebraic complexity theory, particularly in the context of polynomial identity testing. It provides a probabilistic method for determining whether a given polynomial is identically zero. It is a purely mathematical result, and we will use it extensively in this chapter. We state it below:

Lemma 2.1 (Schwartz-Zippel Lemma [Zip79, Sch80])**.** *Let f be a non-zero polynomial in $\mathbb{F}[x_1, x_2, ..., x_n]$ of total degree d. Let S be a finite subset of $\mathbb{F}$ and let $\boldsymbol{\beta} \in_r S^n$ be a random point. Then*

$$\Pr[f(\boldsymbol{\beta}) = 0] \leq \frac{d}{|S|}.$$

The proof of this lemma is by induction on the number of variables. The base case, where f is a univariate polynomial, is clear from the fact that a non-zero univariate polynomial of degree d can have at most d roots. For the inductive step, we can write f as a univariate polynomial in x_1 with coefficients that are polynomials in the other variables. We can then apply the inductive hypothesis to these coefficient polynomials to bound the probability that they vanish at a random point, and use this to bound the probability that f vanishes at a random point.

2.2.2 Newton Iteration

The most important tool used throughout this thesis is classical Newton Iteration. It can be used to find the power series roots of a polynomial up to any degree of

precision. We state and prove it below:

Theorem 2.2 (Implicit Function Theorem [BCS13])**.** *Let $P(\mathbf{x}, y) \in \mathbb{F}(\mathbf{x})[y]$, $P'(\mathbf{x}, y) := \partial_y P(\mathbf{x}, y)$ and $\mu \in \mathbb{F}$ be such that $P(\mathbf{0}, \mu) = 0$ but $P'(\mathbf{0}, \mu) \neq 0$. Then there is a unique power series S such that $S(\mathbf{0}) = \mu$ and $P(\mathbf{0}, S) = 0$.*

Moreover, there exists a rational function $y_t, \forall t \geq 0$ such that

$$y_{t+1} = y_t - \frac{P(\mathbf{x}, y_t)}{P'(\mathbf{x}, y_t)} \text{ and } S \equiv y_t \bmod \langle x \rangle^{2^t} \text{ with } y_0 = \mu.$$

Proof. We give a proof by induction on t. The base case, $t = 0$ holds from the hypothesis itself, as $y_0 = \mu$. Suppose a y_t satisfying the induction hypothesis exists, and define $y_{t+1} = y_t - \frac{P(\mathbf{x}, y_t)}{P'(\mathbf{x}, y_t)}$.

We first need to show that y_{t+1} is well defined. We can see that $y_t \equiv y_{t-1} \bmod \langle x \rangle^{2^{t-1}}$, so $y_t(\mathbf{0}) = \mu$. Then $P'(\mathbf{x}, y_t)|_{x=\mathbf{0}} = P'(\mathbf{0}, \mu) \neq 0$. So $P'(\mathbf{x}, y_t)$ is invertible in the power series ring $\mathbb{F}[[\mathbf{x}]]$. So y_{t+1} is well defined.

Now we show that y_t is indeed a higher degree approximation of the power series root. We apply Taylor expansion on P:

$$\begin{aligned} P(\mathbf{x}, y_{t+1}) &= P\left(\mathbf{x}, y_t - \frac{P(\mathbf{x}, y_t)}{P'(\mathbf{x}, y_t)}\right) \\ &= P(\mathbf{x}, y_t) - P'(\mathbf{x}, y_t)\frac{P(\mathbf{x}, y_t)}{P'(\mathbf{x}, y_t)} + \frac{P''(\mathbf{x}, y_t)}{2!}\left(\frac{P(\mathbf{x}, y_t)}{P'(\mathbf{x}, y_t)}\right)^2 - ... \\ &\equiv 0 \bmod \langle x \rangle^{2^{t+1}} \end{aligned}$$

So $P(\mathbf{x}, y_t) \equiv 0 \bmod \langle \mathbf{x} \rangle^{2^{t+1}}$ and $y_{t+1} \equiv y_t \bmod \langle \mathbf{x} \rangle^{2^t}$.

The power series defined as $S \equiv y_t \bmod \langle \mathbf{x} \rangle^{2^t}, \forall t \geq 0$ will be the root of P, as if it is not the case, $P(\mathbf{x}, y_t) \neq 0 \bmod \langle \mathbf{x} \rangle^{2^t}$ for some t, yielding a contradiction. Uniqueness of this power series follows from the fact that μ is a simple root of $P(\mathbf{0}, y)$.

□

2.2.3 Pre-processing

To factor a general polynomial, we need to do several pre-processing steps to get it into a form where we can apply Theorem 2.2. We state and prove a few lemmas that will be essential in this.

Square Free Factorisation: Note how in the hypothesis for Theorem 2.2, we need to ensure that μ is a root with multiplicity one. A polynomial whose every root is of multiplicity one is said to be square-free. How can we check if a given polynomial is square free? The following lemma gives an easy criterion for this:

Lemma 2.3. *(Square Free Criterion) Let f be a polynomial in $\mathbb{F}(\mathbf{x})[y]$ with $\deg_y f \geq 1$. Then f is square-free if and only if f and $\partial_y f$ are co-prime.*

Proof. If f and $\partial_y f$ are co-prime, then $\gcd(f, \partial_y f) = 1$, implying that they have no common roots. If there is a root $g(\mathbf{x})$ of f multiplicity e greater than one, we must have that $(y - g(\mathbf{x}))^2 | f$. But then $g(\mathbf{x})$ is a common root of f and $\partial_y f$, which is a contradiction. So there is no root of multiplicity greater than one, i.e. f is square-free.

Conversely, assume that f is square-free. For the sake of contradiction, assume that $\gcd(f, \partial_y f)$ has degree at least one. Let g be an irreducible polynomial with degree at least one such that $g | \gcd(f, \partial_y f)$. (Such a g always exists due to our assumption). Let h be such that $f = gh$. Then $\partial_y f = g\partial_y h + h\partial_y g$. So $g | h\partial_y g$. Since g is irreducible, we get that $g | h$ ($(\deg \partial_y g) < \deg g$). Hence $g^2 | f$, which is a contradiction. Hence, $\gcd(f, \partial_y f) = 1$. □

Note that using this lemma, we can also find the square-free part of a polynomial f by simply dividing f by $\gcd(f, \partial_y f)$. This can be done using a polynomial number of gates using the structural results in section 2.3; namely (gcd computation, partial derivatives, and Strassen division gates).

Resultants, GCD, and Co-primality: To easily check whether two polynomials are co-prime, we need not compute their gcd explicitly. We can simply check whether the gcd is non-trivial or not by a tool called the resultant. We develop the theory of resultants in this section.

Suppose we have univariate polynomials $f = \sum_{i=1}^{l} f_i y^i$ and $g = \sum_{i=1}^{m} g_i y^i$ in $\mathbb{F}[y]$. Then consider the map T such that $T(u, v) = uf + vg$ for polynomials u, v of degrees less than m and l respectively. Then the matrix corresponding to the map T is given by:

$$Syl_y(f,g) := \begin{pmatrix} f_0 & 0 & \cdots & 0 & g_0 & 0 & \cdots & 0 \\ f_1 & f_0 & \cdots & 0 & g_1 & g_0 & \cdots & 0 \\ f_2 & f_1 & & \vdots & \vdots & g_1 & & \vdots \\ \vdots & \vdots & \ddots & f_0 & g_m & \vdots & \ddots & 0 \\ f_l & f_{l-1} & & f_1 & 0 & g_m & & g_0 \\ 0 & f_l & & f_2 & 0 & 0 & & g_1 \\ \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & f_l & 0 & 0 & \cdots & g_m \end{pmatrix}.$$

This is known as the Sylvester Matrix. We use this to define the resultant.

Definition 2.4. *(Resultant) Let $f(\mathbf{x})[y]$ and $g(\mathbf{x})[y]$ be multivariate polynomials. Then the resultant of f and g with respect to y is defined as the determinant of the Sylvester matrix:*

$$Res_y(f,g) = \det(Syl_y(f,g))$$

Remark 2.5. Here we view f and g as univariate polynomials in y over the field $\mathbf{F}(\mathbf{x})$ (i.e. f_i's and g_i's are in $\mathbf{F}(\mathbf{x})$). The resultant also lies in $\mathbf{F}(\mathbf{x})$, as it is just the determinant of a matrix over $\mathbf{F}(\mathbf{x})$.

Note that the determinant is non-zero iff the map T is injective. That is, $Res_y(f, g) = 0$ iff T has a non-trivial kernel.

Lemma 2.6. *(gcd and resultant) Let $f(\mathbf{x})[y]$ and $g(\mathbf{x})[y]$ be multivariate polynomials with positive degree in y. Then f and g have a common root iff $Res_y(f, g) = 0$.*

Proof. We have already seen that the map T that takes $(u, v) \to uf + vg$ for $\deg u < m, \deg v < l$ has a non-trivial kernel iff $Res_y(f, g) = 0$. This means that there exist non-trivial solutions u, v to the equation $uf = -vg$ iff $Res_y(f, g) = 0$. Since the polynomial ring is a UFD, this implies that $g|uf$, but since $\deg u < \deg g$, we get that f and g share a common factor. Hence, f and g share a common factor iff $Res_y(f, g) = 0$. □

We conclude this section with the following powerful lemma, which lets us project down to univariate polynomials while preserving their co-primality.

Lemma 2.7. *(co-primality) Let $f, g \in \mathbb{F}(\mathbf{x})[y]$ be co-prime polynomials with positive degree in y, and let $\boldsymbol{\beta} \in_r \mathbb{F}^n$. Then $f(\boldsymbol{\beta}, y)$ and $g(\boldsymbol{\beta}, y)$ are co-prime with positive degree in y.*

Proof. If $f = \sum_{i=1}^{l} f_i y^i$ and $g = \sum_{i=1}^{m} g_i y^i$, we take $f_l.g_m.Res_y(f, g)$ and evaluate it at $\mathbf{x} = \boldsymbol{\beta}$. By Lemma 2.1, we get that this is non-zero with a large probability. This implies, $Res_y(f(\boldsymbol{\beta}, y), g(\boldsymbol{\beta}, y)) \neq 0$. Then by Lemma 2.6, $f(\boldsymbol{\beta}, y)$ and $g(\boldsymbol{\beta}, y)$ are co-prime. □

Remark 2.8. In general, it is not true that $Res_y(f, g)(\boldsymbol{\beta}) = Res_y(f(\boldsymbol{\beta}, y), g(\boldsymbol{\beta}, y))$. It is possible that $f_l(\boldsymbol{\beta}) = 0$ or $g_m(\boldsymbol{\beta}) = 0$, making the y-degree lower, and changing the resultant. That is why we have multiplied f_l and g_m to the resultant before applying Lemma 2.1. This ensures that f_l and g_m are non-zero at $\boldsymbol{\beta}$, and hence $Res_y(f, g)(\boldsymbol{\beta}) = Res_y(f(\boldsymbol{\beta}, y), g(\boldsymbol{\beta}, y))$.

Transformation to Monic Polynomial In the case of multivariate polynomials, the meaning of the term 'monic' isn't as straightforward as it is in the univariate case. Let $f \in \mathbb{F}[\mathbf{x}, y]$ be a polynomial. Then f is said to be monic with respect to y if the coefficient of the highest degree term (wrt y) is in $\mathbb{F}$ i.e. it is a constant.

We can transform an arbitrary polynomial $f \in \mathbb{F}(\mathbf{x})$ to monic by applying a random linear transformation:

Lemma 2.9. *(monic transform) Let $f \in \mathbb{F}[\mathbf{x}]$ be a polynomial with total degree $d \geq 0$ and let $\alpha_i, \beta_i \in_r \mathbb{F}$ be random elements, for $i = 1, 2, ..., n$. Then $g(\mathbf{x}, y) := f(\mathbf{x}+\boldsymbol{\alpha}y+\boldsymbol{\beta})$ is monic wrt y, and the degree wrt y is d.*

Proof. Let $|\boldsymbol{\gamma}|$ denote $\sum_{i=1}^{n} \gamma_i$. Then the homogeneous part of degree d in f is $f_d := \sum_{|\boldsymbol{\gamma}|=d} c_{\boldsymbol{\gamma}} \mathbf{x}^{\boldsymbol{\gamma}}$. On applying the transform $x_i \to \alpha_i y + x_i + \beta_i$, we get the coefficient of y^d in g as $\sum_{|\boldsymbol{\gamma}|=d} c_{\boldsymbol{\gamma}} \boldsymbol{\alpha}^{\boldsymbol{\gamma}}$. We can use the Lemma 2.1 to argue that with a high probability f_d does not vanish on a random $\boldsymbol{\alpha}$. That this is the highest degree term wrt y is clear from the fact that the total degree of f is d. □

2.3 Structural Results about Arithmetic Circuits

In this section, we summarise some well-known structural results about arithmetic circuits. As these are standard folklore results in algebraic complexity, we state them without proof and provide references for the interested reader.

Homogeneous Part of a Polynomial: The homogeneous part of a polynomial f of degree d is the sum of all monomials in f that have total degree d. We denote the homogeneous part of f by $H_d[f]$. The following lemma shows that we can compute the homogeneous part of a polynomial efficiently given a circuit for the polynomial itself. For a proof, see [SY10].

Lemma 2.10. *(Homogenisation [Str73]) Let f be a polynomial that can be computed by a circuit of size s. Then there exists a circuit of size at most $O(r^2 s)$ computing $H_0[f], H_1[f], \ldots, H_r[f]$.*

As obvious, this will be used to truncate a power series root to a polynomial

of a certain degree, which we can then use to compute the factors of the original polynomial.

Division Gates: In the proof of Theorem 1.7, we will need to compute the power series root of a polynomial using Newton Iteration. This involves division of polynomials, which we can do efficiently using only addition and multiplication gates through the following lemma, which is from [Str73]. This actually gives a power series approximation of the inverse of a polynomial.

Lemma 2.11. *(Division Gates [Str73]) Let f, g be two degree-d polynomials that can be computed by circuits of size s and let $\mu \in \mathbb{F}$ be such that $g(\mathbf{0}) = \mu \neq 0$. Then $f/g \bmod \langle \mathbf{x} \rangle^{d+1}$ can be computed by a circuit of size polynomial in s and d.*

The proof of this lemma is based on the following simple identity in the power series ring $\mathbb{F}[[\mathbf{x}]]$, along with the previous homogenisation result:

$$\frac{1}{g} = \frac{1}{\mu} \cdot \frac{1}{1 - (1 - g/\mu)} = \frac{1}{\mu} \sum_{i=0}^{\infty} (1 - g/\mu)^i$$

.

Partial Derivatives: To use Newton Iteration, we need to compute partial derivatives of polynomials. The following lemma shows that we can do this efficiently given a circuit for the original polynomial. This result is from [Kal87].

Lemma 2.12. *(Partial Derivatives [Kal87]) Let f be a multivariate polynomial that can be computed by a circuit of size s. Then for any variable x_i, $\frac{\partial^k f}{\partial x_i^k}$ can also be computed by a circuit of size at most $O(k^2 s)$.*

Chapter 3

Proof for VP closure

In a series of influential works, Kaltofen [Kal86, Kal87, Kal89, KT90] showed that the class VP is closed under factoring. This chapter is dedicated to the proof of this theorem.

Theorem 1.7 (Kaltofen [Kal87])**.** *Let $f \in \mathbb{F}[\mathbf{x}]$ be an n-variate, degree d polynomial that can be computed by a circuit of size s. Let g be a polynomial such that g divides f. Then g can be computed by a circuit of size at most poly(s, n, d).*

The main ingredients used in the proof by Kaltofen are Hensel lifting and an effective version of Hilbert's irreducibility theorem. These techniques are rather mathematically involved, and we would like to present a simpler version in this chapter. We present a version based on [DSS22], which proceeds by reducing the problem of factoring polynomials to the problem of approximating power series roots of the polynomial. We will then show how these power series roots can be approximated using circuits.

3.1 Reduction of Factoring to Power Series Root Approximation

It starts with this simple observation for univariate polynomials: the polynomial $f(x) \in \mathbb{F}[x]$ has a linear factor $(x - a)$, iff $f(a) = 0$, i.e. a is a root of f. We can generalise this to the multivariate case, $(y - g(\mathbf{x}))$ is a factor of $f \in \mathbb{F}[\mathbf{x}, y]$ iff

$f(\mathbf{x}, g(\mathbf{x})) = 0$. This can be seen by viewing f as a univariate polynomial over the field $\mathbb{F}(\mathbf{x})$, i.e. $f \in \mathbb{F}(\mathbf{x})[y]$.

So if we can guarantee the existence of linear factors, finding the roots of the polynomials is enough to find the factors. Unfortunately, it is not the case for multivariate polynomials, even over an algebraically closed field. For example, $xy-1 \in \mathbb{C}[x, y]$ is an irreducible polynomial over a closed field which is not linear. However, there is a workaround: we can show that with a random shift on the variables, a polynomial can be written as a product of linear factors of the form $(y - g(\mathbf{x}))$ where g is in the power series ring $\mathbb{F}[[\mathbf{x}]]$. This is the key result of the paper and we state in the following theorem:

Theorem 3.1 (Power Series Complete Split, [DSS22])**.** *Let $f \in \mathbb{F}[\mathbf{x}]$ be a polynomial with $\deg(rad(f)) = d_0 > 0$. Let $\alpha_i, \beta_i \in_r \mathbb{F}$ and $\tau : x_i \to \alpha_i y + x_i + \beta_i, i \in [n]$, where y is a new variable. Then over $\mathbb{F}[[\mathbf{x}]]$, $f(\tau\mathbf{x}) = k.\prod_{i \in d_0}(y - g_i)^{\gamma_i}$, where $k \in \mathbb{F}^*, \gamma_i > 0$ and $g_i(\mathbf{0}) = \mu_i$. Moreover, μ_i's are distinct non-zero field elements.*

Proof. Let the irreducible factorisation of f be $f = \prod_{i \in [m]} f_i^{e_i}$. On applying τ, f and all its irreducible factors f_i become monic in y (Lemma 2.9). Since τ is invertible, all these monic factors $\tilde{f}_i := f_i(\tau\mathbf{x})$ remain irreducible.

Our goal is to start with a univariate factorisation and apply Newton Iteration. To do so, consider $\tilde{f}_i(\mathbf{0}, y) = f_i(\alpha y + \beta)$. We can show that $\tilde{f}_i(\mathbf{0}, y)$ and $\partial_y \tilde{f}_i(\mathbf{0}, y)$ are co-prime using Lemma 2.7. This implies that $\tilde{f}_i(\mathbf{0}, y)$ is square-free (Lemma 2.3), and hence its roots are simple. This will ensure that $\tilde{f}_i(\mathbf{x}, y)$ satisfies the hypothesis for applying Newton Iteration.

Since $\mathbb{F}$ is algebraically closed, $\tilde{f}_i(\mathbf{0}, y)$ (a univariate polynomial) splits into linear factors as $\tilde{f}_i = \prod_{j=1}^{\deg(f_i)}(y - \mu_{j,i})$ for distinct non-zero field elements $\mu_{j,i}$. For each of these, we can use Theorem 2.2 to build a power series root $g_{j,i} \in \mathbb{F}[[x]]$ with

$g_{j,i}(0) = \mu_{j,i}$. So we can write $\tilde{f_i} = \prod_{j=1}^{\deg(f_i)}(y - g_{j,i})$ with $g_{j,i}(0) = \mu_{j,i}$. So each f_i splits into linear factors over $\mathbb{F}[[\mathbf{x}]][y]$.

Since f_i's are co-prime, we can get that $\tilde{f_i}(\mathbf{0}, y)$ are mutually co-prime using Lemma 2.7. That is, $\mu_{j,i}$ are all distinct and there are $\sum_{i\in[m]} \deg(f_i) = d_0$ many of them. Hence, $f(\tau\mathbf{x})$ can be factored as

$$\prod_{i\in[m]} f_i(\tau\mathbf{x})^{e_i} = \prod_{i\in[d_0]} (y - g_i)^{\gamma_i}$$

for $\gamma_i > 0$ and $g_i(0)$ distinct field constants for all i.

□

Remark 3.2. Theorem 3.1 says that a polynomial $p \in \mathbb{F}[\mathbf{x}]$, after a random linear transformation, completely splits over $\mathbb{F}[[\mathbf{x}]]$. This shift is, in fact, necessary. For example, take $f(x, y) = y^2 - x \in \mathbb{F}(x)[y]$. Over the algebraic closure of $\mathbb{F}(\mathbf{x})$, it splits as $f(x, y) = (y - \sqrt{x})(y + \sqrt{x})$. However, $\sqrt{x}$ cannot be written as a power series over $\mathbb{F}$. But if we apply a random shift on x, for example, $f(x + 1, y)$, it has roots $\sqrt{x+1}$ and $-\sqrt{x+1}$ which can be written as power series using the Binomial Theorem.

The following is well known result in commutative algebra, that states that $\mathbb{F}[[\mathbf{x}]]$ is a unique factorisation domain.

Proposition 3.3. *([ZS75], chapter 7) The power series ring $\mathbb{F}[[\mathbf{x}]]$ is a Unique Factorisation Domain (UFD), and so is $\mathbb{F}[[\mathbf{x}]][y]$.*

A simple consequence of Proposition 3.3 and Theorem 3.1 is the following corollary which will help us recover the polynomial factors of f.

Corollary 3.4. *Suppose g is a polynomial factor of f. Let $f(\tau\mathbf{x}) = k.\prod_{i\in[d_o]}(y - g_i)^{\gamma_i}$ as in the previous theorem. As $g(\tau\mathbf{x})|f(\tau\mathbf{x})$, $g(\tau\mathbf{x}) = k'.\prod_{i\in[d_o]}(y - g_i)^{c_i}$ with $0 \le c_i \le \gamma_i$. Moreover, we can get g back by applying τ^{-1} on $g(\tau\mathbf{x})$.*

So how can we go from these power series roots to the polynomial factors of the polynomial? Corollary 3.4 hints at this: if g is a factor of f, some of the linear factors of $f(\tau\mathbf{x})$ must combine to produce $g(\tau\mathbf{x})$.

Suppose $g(\tau\mathbf{x}) = k'.\prod_{i\in[d_o]}(y-h_i)^{c_i}$, and let the degree of f be d. If g is a non-trivial factor, $\deg(g) \leq (d-1)$. If we know these power series roots h_i's, we can obtain g by taking these h_i's up to precision of $\deg(g)$ (which is less than d), do the corresponding multiplication, and truncate it again to $\deg(g)$. So the problem of finding factors of f is reduced to the problem of finding power series roots of f. And to upper bound the size of the circuits computing g, we only need to bound the circuit size for approximating h_i's up to the precision equal to the degree of g.

Note that to obtain the VP closure result, we only need to upper bound the size of the circuits computing g; we don't need to find exactly which of the power series roots of f combine to produce g. However, it will be needed if we wish to algorithmically compute g. That is beyond the scope of this thesis, but we refer the reader to [DSS22].

3.2 Proof of the Main Theorem

Once the problem of finding factors is reduced to power series root approximation, the proof for VP closure becomes fairly simple.

Theorem 1.7 (Kaltofen [Kal87])**.** *Let $f \in \mathbb{F}[\mathbf{x}]$ be an n-variate, degree d polynomial that can be computed by a circuit of size s. Let g be a polynomial such that g divides f. Then g can be computed by a circuit of size at most poly(s,n,d).*

Proof. Let $f(\mathbf{x}) = g^e(\mathbf{x})h(\mathbf{x})$, (without loss of generality) where g and h are co-prime, and let d_g be the degree of g. We take the linear shift τ as in Theorem 3.1 to ensure that the transformed polynomial $\tilde{f}(x,y) = f(\tau\mathbf{x})$ is monic, and it splits over $\mathbb{F}[[\mathbf{x}]]$. Also note that $\tilde{g}, \tilde{h}$ remain co-prime (Lemma 2.7). So using Theorem 3.1,

$\tilde{f} = k.\prod_{i\in d_0}(y-g_i)^{\gamma_i}$, with $g_i(\mathbf{0}) = \mu_i$ being distinct.

Since $\tilde{g}$ has multiplicity e, all the power series roots of $\tilde{g}$ must occur with multiplicity e. We cannot approximate a root of multiplicity greater than 1 using Newton iteration (NI), as the denominator becomes non-invertible. So we need to do some preprocessing.

But note that the square free part of $\tilde{f}$ contains all the power series roots of $\tilde{f}$, and hence g with multiplicity 1. This can be done by dividing (for now we construct a circuit that allows division gates; we will remove them later) $\tilde{f}$ by $\gcd(f, \partial_y \tilde{f})$ in poly size (see square free factorisation in subsection 2.2.3). We work with this polynomial (renamed to $\tilde{f}$ for convenience) to get the power series roots of $\tilde{g}$.

So let $\tilde{g}_i$ be a power series root of $\tilde{g}$, which has multiplicity 1 in the factorisation of $\tilde{f}$. We will show that it can be approximated up to degree d via a circuit of size $\text{poly}(s, d)$. We start with the base case $\tilde{g}_{i,1} = \tilde{g}_i(\mathbf{0}) = \mu_i \equiv \tilde{g} \bmod \langle \mathbf{x} \rangle$. This is just a field element, and hence has circuit size 1. On this, we iteratively apply the Newton Iteration formula $\tilde{g}_{i,k+1} = \tilde{g}_{i,k} - \frac{\tilde{g}_{i,k}(\mathbf{x},\tilde{g}_{i,k})}{\partial_y \tilde{g}_{i,k}(\mathbf{x},\tilde{g}_{i,k})} \bmod \langle \mathbf{x} \rangle^{2^k}$, where each $\tilde{g}_{i,k} \equiv \tilde{g}_i \bmod \langle \mathbf{x} \rangle^{2^t}$ for $1 \leq k \leq \log d$.

From the Newton iteration formula, along with efficient partial derivative computation (Lemma 2.12), it can be seen that if $\tilde{g}_{i,k}$ has circuit size s_k, then $\tilde{g}_{i,k+1}$ has circuit size $s_{k+1} = s_k + poly(s)$. After all the iterations, we get a circuit (containing division gates) of size $poly(s, \log d)$ for $\tilde{g}_{i,d}$. We can then use Strassen's division (Lemma 2.11) and homogenisation (Lemma 2.10) to get a circuit of size $poly(s, d)$ without any division gates. So up to a precision of degree d, all power series roots of $\tilde{g}$ can be computed in size $poly(s, d)$.

We can obtain the circuit for $\tilde{g} = \prod_i (y - \tilde{g}_i)$ by multiplying the circuits for $(y - \tilde{g}_{i,d})$ for each i and truncating (truncate lemma) up to $\deg(g)$. And then we can apply τ^{-1} to obtain g. The final circuit size is still $poly(s, d)$.

□

Remark 3.5. We again emphasize the fact that to show the bound on the size of the circuit computing $\tilde{g}$, we do not need to show exactly which power series roots of $\tilde{f}$ combine to produce g, nor do we need to show by what power the (simple) roots of the square free factorisation of $\tilde{f}$ were raised to produce $\tilde{g}$, since it is less than d.

From Theorem 1.7, it follows that the class VP is closed under factoring.

3.3 Fields that are not Algebraically Closed

Suppose we have a polynomial $f \in \mathbb{F}[\mathbf{x}]$ where $\mathbb{F}$ is not algebraically closed. It is not feasible to extend the base field to the algebraic closure of $\mathbb{F}$, as the splitting field of f can be very large. However, we can still apply the techniques of the previous section to get a factorisation of f over $\mathbb{F}(\alpha)$, for some root α of $f(\mathbf{0}, y)$ in the algebraic closure of $\mathbb{F}$. We can then use the resultant to get a factorisation of f over $\mathbb{F}$ in the following way.

Suppose we have constructed a circuit $C(\mathbf{x}, \alpha)$ that computes a good enough approximation of the power series root $\gamma(\mathbf{x}, \alpha)$ of $f(\mathbf{x}, y)$ in the extension field $\mathbb{F}(\alpha)$, where α is an algebraic element with minimal polynomial $A(z) \in \mathbb{F}[z]$ (usually $A(y)$ will just be $f(\mathbf{0}, y)$). To recover a factor of f that resides in the base field $\mathbb{F}[\mathbf{x}, y]$, we define the reconstructed polynomial $Q(\mathbf{x}, y)$ as the resultant of $A(z)$ and the shifted circuit representation:

$$Q(\mathbf{x}, y) = \mathrm{Res}_z\left(A(z), y - C(x, z)\right) \tag{3.1}$$

By the properties of the Resultant (explained in the next subsection), this is equivalent to:

$$Q(\mathbf{x}, y) = \prod_{i=1}^{d} (y - C(\mathbf{x}, \alpha_i)) \tag{3.2}$$

where $\alpha_1, \ldots, \alpha_d$ are the roots of $A(z)$ in its splitting field. This construction ensures that:

1. $Q(\mathbf{x}, y) \in \mathbb{F}[\mathbf{x}, y]$, as the resultant of two polynomials in $\mathbb{F}[\mathbf{x}, y, z]$ with respect to z eliminates z and keeps the coefficients in $\mathbb{F}$.

2. **Complexity:** Since the resultant is the determinant of a Sylvester matrix of size $O(d)$, and C is a circuit of size s, the complexity of the reconstructed factor $Q(\mathbf{x}, y)$ is $poly(s, d)$, preserving its membership in the class VP.

3. **Factorization:** If $C(\mathbf{x}, \alpha)$ approximates a power series root of $f(\mathbf{x}, y)$, then $Q(\mathbf{x}, y)$ will be a factor of f over $\mathbb{F}[\mathbf{x}, y]$ as it is the product of linear factors corresponding to the roots.

3.3.1 Resultants and Norm

To understand why the resultant computes the norm, we examine the Sylvester matrix through the lens of linear transformations. Given two polynomials $f, g \in \mathbb{F}[z]$ of degrees n and m respectively, the Sylvester matrix S represents the $\mathbb{F}$-linear map:

$$\begin{aligned} \phi : \mathcal{P}_{m-1} \times \mathcal{P}_{n-1} &\to \mathcal{P}_{n+m-1} \\ (u, v) &\mapsto uf + vg \end{aligned}$$

where $\mathcal{P}_k$ denotes the space of polynomials of degree at most k.

The determinant of this matrix, $\text{Res}_z(f, g)$, vanishes if and only if f and g share a common root. More deeply, this determinant is related to the map of multiplication by g in the quotient ring $\mathbb{A} = \mathbb{F}[z]/\langle f(z) \rangle$. Let $\alpha_1, \ldots, \alpha_n$ be the roots of $f(z)$ in its algebraic closure. In the extension field $\mathbb{F}(\alpha)$, the eigenvalues of the linear operator representing multiplication by g are exactly the values $\{g(\alpha_1), g(\alpha_2), \ldots, g(\alpha_n)\}$.

Recall that the determinant of any linear operator is equal to the product of its eigenvalues. Consequently, for a monic polynomial f, the resultant satisfies the following identity:

$$\mathrm{Res}_z(f, g) = \prod_{i=1}^{n} g(\alpha_i) \tag{3.3}$$

By definition, the *algebraic norm* of an element $g(\alpha)$ in the extension $\mathbb{F}[z]/\langle f(z)\rangle$ is the product of g evaluated at all conjugates of α. Thus, the Sylvester matrix provides a direct determinantal method to compute the norm:

$$\mathrm{Res}_z(f, g) = N_{\mathbb{F}(\alpha)/\mathbb{F}}(g(\alpha)) \tag{3.4}$$

Chapter 4

Characteristic p

In the previous chapters we saw the VP closure result for fields with characteristic 0. In fact, to make our lives easier, we were only dealing with algebraically closed fields. As we saw in the last section, we can extend the result efficiently to non-algebraically closed fields as well.

The problem, however, remains open for fields of small positive characteristic. In fact, there is one specific case which has remained unsolved in this.

Open Question 4.1. Let $\mathbb{F}_q$ be a field of characteristic p, and let $f = g^{p^e} \in \mathbb{F}_q[\mathbf{x}]$ be a polynomial in VP, for some polynomial $g \in \mathbb{F}_q[\mathbf{x}]$. Is g also in VP?

This is the central objective of this thesis. We will analyse a proposed approach based on p-adic numbers, and see why it doesn't work. Before this, we need to understand why the older techniques fail for this particular case. For now, suppose $f \in \mathbb{F}_q[\mathbf{x}, y]$. Then $\partial_y f = p^e g^{e-1} = 0$ (as we are in characteristic p). This means the term $f(\mathbf{0}, \mu)/\partial_y f(\mathbf{0}, \mu)$ used in the Newton iteration formula becomes undefined in $\mathbb{F}_q$ (Here μ is a root of $f(\mathbf{0}, y)$, as in the proof of Theorem 3.1), making it impossible to use Newton Iteration the way we did for other cases.

Our proposed approach to solve Open Question 4.1 is to lift the problem to a p-adic field, and then use the techniques we have developed for characteristic 0 fields. The idea is that if we can show that g is in VP over the p-adic field, and that all the computations are well defined mod p, then we can reduce it back to $\mathbb{F}_q$ and conclude that g is in VP over $\mathbb{F}_q$ as well. The p-adic field is the most natural characteristic-zero

extension of $\mathbb{F}_p$, because $\mathbb{F}_p$ is the residue field of the p-adic integers.

To discuss our proposed approach, we need to understand the basics of p-adic numbers. We will discuss that in the next section.

4.1 Some p-adic analysis

Here we introduce some of the basic definitions from p-adic analysis, to build up to Lemma 4.6. We have borrowed these from [Neu99].

Definition 4.2. *Let p be a prime number. A p-adic integer is a formal infinite series*

$$a_0 + a_1 p + a_2 p^2 + ...$$

where $0 \leq a_i < p$ for all $i = 0, 1, 2, ...$ The set of p-adic integers is denoted by $\mathbb{Z}_p$.

In analogy to the Laurent series, we can extend the domain to the formal series

$$a_{-m} p^{-m} + ... + a_{-1} p^{-1} + a_0 + a_1 p + a_2 p^2 + ...$$

where $m \in \mathbb{Z}$ and $0 \leq a_v < p$. We call such a series a p-adic number and we denote the set of p-adic numbers by $\mathbb{Q}_p$.

Definition 4.3. *The p-adic valuation of an integer $x \in \mathbb{Z}$, $v_p(x)$ is defined as follows:*

$$v_p(x) = \max\{r : p^r | x\}$$

By convention, $v_p(0) = \infty$. For a rational $a/b \in \mathbb{Q}$,

$$v_p(a/b) = v_p(a) - v_p(b)$$

Definition 4.4. *The p-adic norm of $x \in \mathbb{Q}$ is given by*

$$|x|_p = \begin{cases} p^{-v_p(x)} & x \neq 0 \\ p^{-\infty} = 0 & x = 0 \end{cases}$$

The completion of $\mathbb{Q}$ with respect to the p-adic norm is $\mathbb{Q}_p$. However, $\mathbb{Q}_p$ is not algebraically closed. The algebraic closure of $\mathbb{Q}_p$ is denoted by $\bar{\mathbb{Q}}_p$. It is known that $\bar{\mathbb{Q}}_p$ is not complete with respect to the p-adic norm. The completion of $\bar{\mathbb{Q}}_p$ with respect to the p-adic norm is denoted by $\mathbb{C}_p$.

Ring extension: Suppose we have a polynomial f in the field $\mathbb{F}_p$, and let $\mathbb{F}_{p^l}$ be an extension over $\mathbb{F}_p$ of degree l, containing a root of f. Note how taking mod p of the elements of $\mathbb{Z}_p$ gives us $\mathbb{F}_p$. That is, $\mathbb{Z}_p/\mathfrak{p}\mathbb{Z}_\mathfrak{p} \cong \mathbb{F}_p$. This is known as the *residue field* of $\mathbb{Z}_p$. We want an extension $\hat{\mathbb{Z}}_p$ of degree l over $\mathbb{Z}_p$ such that taking mod p gives us $\mathbb{F}_{p^l}$. This is called an *unramified extension* of $\mathbb{Z}_p$.

Suppose the degree of f is d. It is known that we can come up with a monic irreducible polynomial over $\mathbb{Z}_p$ of any required degree. Let g be such a polynomial of degree d. We can construct a lift $\tilde{g} \in \mathbb{Z}_p[x]$ of g, by viewing the coefficients $\{0, 1, 2, ..., (p-1)\}$ of g as the same elements in $\mathbb{Z}_p$. By Hensel's lemma, $\tilde{g}$ is also an irreducible polynomial in $\mathbb{Z}_p[x]$. Then $\hat{\mathbb{Z}}_p := \mathbb{Z}_p/(\tilde{g})$ is the required extension.

Valuation for Algebraic Extensions and Ramification: More generally, let K be an algebraic extension of $\mathbb{Q}_p$. The p-adic norm, v_p, extends naturally as $|\alpha|_p = |min_{\mathbb{Q}_p,\alpha}(0)|_p^{1/d}$ for any $\alpha \in K$, where $min_{\mathbb{Q}_p,\alpha}(x)$ denotes the minimal polynomial of α over $\mathbb{Q}_p$ and d denotes its degree. See [Bak07, Theorem 5.2].

The p-adic valuation v_p is a homomorphism from K^* (as a multiplicative group) to the additive group $\mathbb{Q}$. Its image is of the form $\frac{1}{e}\mathbb{Z}$ where e is a divisor of $n = [K : \mathbb{Q}_p]$. There exists an element $\pi \in K$, called uniformizer, that has $v_p(\pi) = 1/e$. Every $x \in K$ has a unique representation $\sum_{i \geq m} \alpha_i \pi^i$ where $m \in \mathbb{Z}$ and $\alpha_i \in S, f := n/e$. Here S is a set of representatives of the residue field (which is isomorphic to $\mathbb{F}_{p^f}$) in K. An extension is called *unramified*, if $\pi = p, e = 1$ and $f = n$.

The *integral extension of* $\mathbb{Z}_p$, denoted as $\mathcal{O}_K$, is a subring of K with series representation that do not have negative powers of π, i.e. $\sum_{i \geq 0} \alpha_i \pi^i$.

Newton Polygon and Valuation of the Roots: Our goal was to use techniques from p-adic analysis to ensure that all the computations we do (using characteristic 0 techniques) are well defined mod p. To do so, we develop a tool called the Newton Polygon, which gives us information about the roots of a polynomial in terms of the valuations of its coefficients. We will use this tool to show that the roots of $F(\mathbf{0}, y)$ are in $\hat{\mathbb{Z}}_p$, and we intended to show that we can then perform Newton Iteration to get a power series solution for $F(\mathbf{x}, y) = 0$ which is in $\hat{\mathbb{Z}}_p[[\bar{x}]]$. This would have ensured that all the computations we do using this power series solution are well defined mod p.

Definition 4.5. *For a univariate polynomial $p(x) = 1 + a_1X + a_2X^2 + ... + a_nX^n \in K[X]$, $a_n \neq 0$, we associate a point $(i, v(a_i)) \in \mathbb{R}^2$ for each $i = 1, 2, .., n$, ignoring the point (i, ∞) for $a_i = 0$. We take the lower convex envelope of the set of points $\{(0, 0), (1, v(a_1)), ..., (n, v(a_n))\}$. This produces a polygonal chain which we call the* ***Newton Polygon*** *of $p(x)$.*

The Newton polygon consists of line segments whose slopes are strictly increasing, and they contain information about the roots of the polynomial.

The following lemma, stated and proved in [Gou20, Proposition 7.4.6], provides the information about the valuations of the roots of a polynomial.

Lemma 4.6. [Gou20, Proposition 7.4.6] *Let $f(X) = 1 + a_1X + a_2X^2 + \cdots + a_nX^n \in K[X]$, and assume that the first break of the Newton polygon of $f(X)$ occurs at the point (i, mi). Then $f(X)$ has no roots with absolute value less than p^m and has exactly i roots (counting multiplicities, in $\mathbb{C}_p$) with absolute value p^m.*

Remark 4.7. The theorem says the first segment of the Newton polygon isolates roots with the smallest p-adic valuation. One can apply the theorem repeatedly to each subsequent segment to recover the valuations of all roots, but for our purposes the first break statement suffices.

Example 4.8 (Newton polygon and root valuations)**.** *Consider the polynomial constructed explicitly as:*

$$f(y) \;=\; (y-p)(y-p^2)(y-p^3).$$

Expanding:

$$\begin{aligned} f(y) &= y^3 - (p+p^2+p^3)\,y^2 + (p^3+p^4+p^5)\,y - p^6 \\ &= y^3 - p(1+p+p^2)\,y^2 + p^3(1+p+p^2)\,y - p^6. \end{aligned}$$

The coefficients are:

$$a_3 = 1, \quad a_2 = -p(1+p+p^2), \quad a_1 = p^3(1+p+p^2), \quad a_0 = -p^6.$$

With valuations:

$$v_p(a_0) = 6, \quad v_p(a_1) = 3, \quad v_p(a_2) = 1, \quad v_p(a_3) = 0.$$

We plot the points $(i, v_p(a_i))$ for $i = 0, 1, 2, 3$:
All four points lie on the lower convex hull, giving three segments. Applying Lemma 4.6 to the first segment: it runs from $(0, 6)$ to $(1, 3)$, has slope $\lambda = -3$, and horizontal length $k = 1$. The theorem predicts exactly one root of f with $v_p(\alpha) = 3$. Indeed the root is $\alpha = p^3$ and $v_p(p^3) = 3$.

Applying the theorem to each subsequent segment recovers all roots:

Segment	*Slope λ*	*Length k*	*Root ($v_p = -\lambda$)*
$(0,6) \to (1,3)$	-3	1	p^3 $(v_p = 3)$
$(1,3) \to (2,1)$	-2	1	p^2 $(v_p = 2)$
$(2,1) \to (3,0)$	-1	1	p $(v_p = 1)$

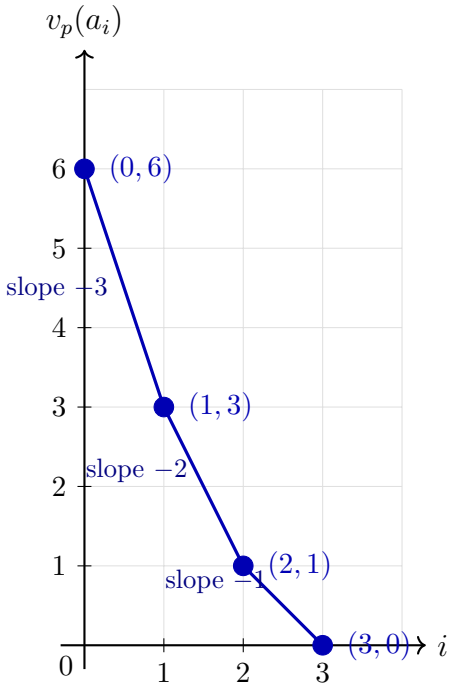

Figure 4.1: Newton polygon of $f(y) = (y-p)(y-p^2)(y-p^3)$. All four coefficient points lie on the lower convex hull. Each segment has horizontal length 1 and slope equal to the negative of the p-adic valuation of the corresponding root.

4.2 An Approach by Interpreting the circuit over $\hat{\mathbb{Q}}_p$

We now explain how we can interpret the circuit over $\mathbb{Q}_p$. Suppose we are given a circuit over $\mathbb{F}_p$ computing the polynomial $f(\mathbf{x}) = u^{p^i}(\mathbf{x})$ which is in VP. From here, we start interpreting it as a circuit over $\mathbb{Q}_p$, which leads to an addition of $pE(\mathbf{x})$ to f, for some polynomial E. We call this new polynomial $\tilde{f} = f + pE$. The addition of pE is the result of the fact that there are some terms computed by the circuit that vanish in $\mathbb{F}_p$ due to characteristic p. Note that $\tilde{f} \equiv f \bmod p$. We will be working with $\tilde{f}$ instead of f because we want to use the techniques we have developed for characteristic 0 fields, and $\tilde{f}$ is a polynomial over a characteristic 0 field.

Note that the depth reduction of Valiant et al. [VSBR83] implies that the $\tilde{f}$ can

be computed by a circuit of depth $O(\log r)$, and since the degree of a polynomial computed by a circuit of depth $O(\log r)$ is at most $2^{O(\log r)} = r^{O(1)}$, we can assume that the degree of E is at most $r^{O(1)}$. In other words, $\tilde{f}$ is in the class VP over $\mathbb{Z}_p$.

Preprocessing: Since we are now in a characteristic 0 field, we can apply Theorem 3.1. Let $F(\mathbf{x}, y) = \tilde{f}(\tau\mathbf{x})$, where τ is the map that sends $x_i \to \alpha_i y + x_i + \beta_i, i \in [n]$ for $\alpha_i, \beta_i \in_r \mathbb{Q}_p$ random. Following the notation in Theorem 3.1, let μ_i's be the roots of the univariate polynomial $F(\mathbf{0}, y)$.

We will show that all the μ_i's are in $\hat{\mathbb{Z}}_p$ using Lemma 4.6.

Remark 4.9. Note that Theorem 3.1 already gives a power series factorisation over $\mathbb{Q}_p$. However, we need to ensure that the power series solution we get is in $\hat{\mathbb{Z}}_p[[\mathbf{x}]]$, so that all the computations we do using this power series solution are well defined mod p. This is why we need to show that the roots of $F(\mathbf{0}, y)$ are in $\hat{\mathbb{Z}}_p$.

Lemma 4.10. *The roots of $F(\mathbf{0}, y)$ are in $\hat{\mathbb{Z}}_p$.*

Proof. Before we apply Lemma 4.6, we need to ensure $F(\mathbf{0}, y)$ is in the form as in the statement of the lemma. The preprocessing ensures that the constant term has valuation 0 and that $F(\mathbf{0}, y)$ is monic in y. Actually after the preprocessing we have $F(\mathbf{0}, y) = y^{p^i} + g(\mathbf{0})^{p^i} + pE'(\mathbf{0}, y)$, for some polynomials g, E' with constant term of g of valuation 0. We divide by the constant term to get $\tilde{F}$ which has the form required in Lemma 4.6. Note that this has not changed the valuation of any of the coefficients. The Newton polygon of $\tilde{F}(\mathbf{0}, y)$ (univariate polynomial over y) has the vertices $(0, 0), (p^i, 0)$ and (k, j) such that $j \geq 0$, for $k \neq 0, p^i$, so that the first break is at $(p^i, 0)$. By Lemma 4.6, $\tilde{F}(\mathbf{0}, y)$ has no roots of absolute value less than p^0, and hence all the roots are in $\hat{\mathbb{Z}}_p$.The roots of F and $\tilde{F}$ are the same and hence, the roots of F have absolute value equal to p^0.

□

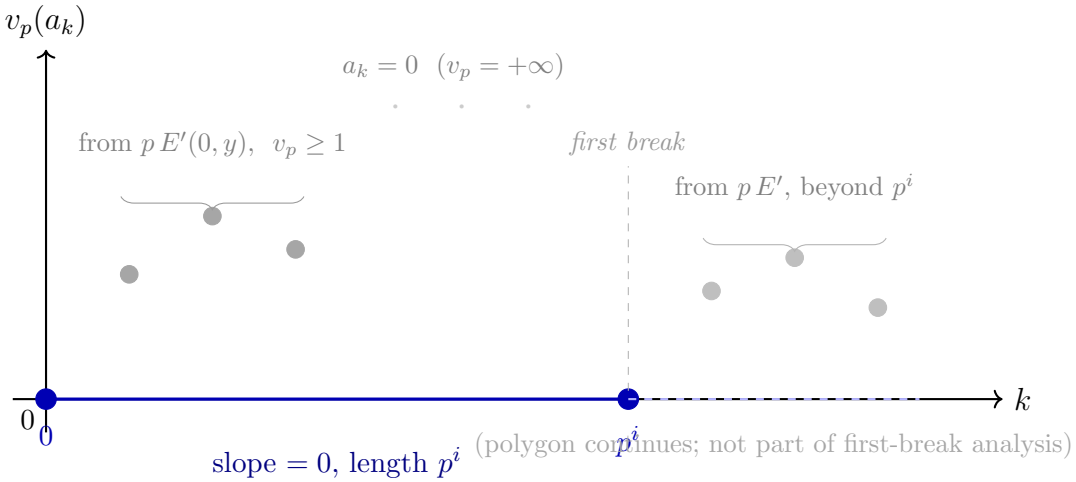

Figure 4.2: Newton polygon of $F(0,y) = y^{p^i} + g(0)^{p^i} + p\,E'(0,y)$, where $\deg_y F \geq p^i$. The bold blue segment is the *first segment* of the Newton polygon, running from $(0,0)$ to $(p^i,0)$ with slope 0 and horizontal length p^i. Points between 0 and p^i from $p\,E'(0,y)$ have $v_p \geq 1$ and lie strictly above the x-axis; points with zero coefficient are omitted ($v_p = +\infty$). The dashed blue line indicates the polygon may continue beyond p^i up to $\deg_y F$, with further points from E' floating above; this is irrelevant to the first-break analysis. By Lemma 4.6, $F(0,y)$ has exactly p^i roots in $\overline{\mathbb{Q}_p}$ with p-adic valuation 0. Since the first segment has length $p^i > 1$, these roots cannot be separated by valuation, obstructing the lifting step.

The idea was that we can use Lemma 4.10 as a base case, and use the Newton Iteration formula to inductively show that each coefficient of the power series solution is in $\hat{\mathbb{Z}}_p$.

Unfortunately, this is not true, and there is a simple counterexample to show that the coefficients of the power series solution are not in $\hat{\mathbb{Z}}_p$.

4.2.1 Counterexample to Integral Power Series Solutions

One might assume that if the univariate polynomial $F(0,y)$ has roots in $\mathbb{Z}_p$, then by Newton Iteration, one can always lift these to a power series root in $\mathbb{Z}_p[[x]]$. We provide a counterexample to show that this is not the case when the derivative at the root has a positive valuation.

Example 4.11. *(Counterexample to the power series roots being integral) Consider the polynomial* $F(x, y) = y^2 - (1+x) \in \mathbb{Z}_2[x, y]$. *At* $x = 0$, *the roots of* $F(0, y) = y^2 - 1$ *are* $y = \pm 1$, *which are clearly 2-adic integers in* $\mathbb{Z}_2$. *Note that* F *has the form as in Lemma 4.10.*

However, the power series solution for $F(x, y) = 0$ *is* $y = \sqrt{1+x}$. *Applying the generalized binomial expansion, we obtain:*

$$y = 1 + \frac{1}{2}x - \frac{1}{8}x^2 + \frac{1}{16}x^3 - \frac{5}{128}x^4 + \dots \tag{4.1}$$

While this is a valid formal power series over the characteristic zero field $\mathbb{Q}_2$, *the coefficients* c_n *for* $n \geq 1$ *do not belong to the ring of integers* $\mathbb{Z}_2$. *For instance,* $v_2(c_1) = -1$ *and* $v_2(c_2) = -3$.

The derivative $F'_y(0, 1) = 2$ *has a positive valuation, which forces the division by* p *during the construction of the series. Consequently, the coefficients have strictly negative valuations, meaning the power series* $y(x)$ *is not in* $\mathbb{Z}_2[[x]]$.

In the context of our VP closure problem, this implies that even if the constant terms of our circuit computations are well-defined modulo p, the higher-order terms in the power series are not well defined modulo p. Consequently, the resulting polynomial g cannot be recovered by simply taking the reduction modulo p of the characteristic-zero power series.

4.3 Conclusion

To our current understanding, the p-adic approach to factoring remains incomplete unless a mechanism is found to bridge the gap between $\mathbb{Q}_p$ and $\mathbb{F}_p$ roots without relying on the standard reduction map. Specifically, because the p-adic valuation of the coefficients can be strictly negative, the "mod p" operation is not well-defined for the resulting root.

Unless one can extract the $\mathbb{F}_p$-factorization from the $\mathbb{Q}_p$-data through a process that bypasses direct modular reduction, the p-adic lifting technique appears to fail for the inseparable case.

Bibliography

[Bak07] A. J. Baker. An Introduction to p-adic Numbers and p-adic Analysis. https://www.ams.org/open-math-notes/files/course-material/OMN-202003-110818-1-Course_notes-v2.pdf, 2007. Lecture notes, University of Glasgow.

[BCS13] Peter Bürgisser, Michael Clausen, and Amin Shokrollahi. *Algebraic Complexity Theory*, volume 315 of *Grundlehren der mathematischen Wissenschaften*. Springer Science & Business Media, Berlin, Heidelberg, 2013.

[DSS22] Pranjal Dutta, Nitin Saxena, and Amit Sinhababu. Discovering the Roots: Uniform Closure Results for Algebraic Classes Under Factoring. *J. ACM*, 69(3):18:1–18:39, 2022.

[Gou20] Fernando Q. Gouvêa. *p-adic numbers*. Springer Cham, 2020.

[Kal86] Erich L. Kaltofen. Uniform Closure Properties of P-Computable Functions. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing (STOC 1986)*, pages 330–337. ACM, 1986.

[Kal87] Erich Kaltofen. Single-Factor Hensel Lifting and Its Application to the Straight-Line Complexity of Certain Polynomials. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC 1987)*, page 443–452. Association for Computing Machinery, 1987.

[Kal89] Erich Kaltofen. Factorization of Polynomials Given by Straight-Line Programs. *Adv. Comput. Res.*, 5:375–412, 1989.

[KT90] Erich Kaltofen and Barry M. Trager. Computing with polynomials given by black boxes for their evaluations: greatest common divisors, factorization, separation of numerators and denominators. *J. Symbolic Comput.*, 9(3):301–320, 1990.

[Neu99] Jürgen Neukirch. *Algebraic Number Theory*. Springer Berlin, Heidelberg, 1999.

[Sch80] Jacob T. Schwartz. Fast Probabilistic Algorithms for Verification of Polynomial Identities. *Journal of the ACM (JACM)*, 27(4):701–717, 1980.

[Str73] Volker Strassen. Vermeidung von Divisionen. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 1973(264):184–202, Nov 1973.

[SY10] Amir Shpilka and Amir Yehudayoff. Arithmetic Circuits: A Survey of Recent Results and Open Questions. *Found. Trends Theor. Comput. Sci.*, 5(3–4):207–388, 2010.

[VSBR83] Leslie G. Valiant, Sven Skyum, Stuart Berkowitz, and Charles Rackoff. Fast Parallel Computation of Polynomials Using Few Processors. *SIAM Journal on Computing*, 12(4):641–644, 1983.

[Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Symbolic and Algebraic Computation, EUROSAM '79, An International Symposium on Symbolic and Algebraic Manipulation*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226, Berlin, Heidelberg, 1979. Springer-Verlag.

[ZS75] Oscar Zariski and Pierre Samuel. *Commutative Algebra*, volume 1 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, Heidelberg,

Berlin, 1975. With the cooperation of I. S. Cohen. Corrected reprinting of the 1958 edition.