



**This electronic thesis or dissertation has been
downloaded from Explore Bristol Research,
<http://research-information.bristol.ac.uk>**

Author:

Van Leeuwen, Barry

Title:

Number Field Sieve with provable complexity

General rights

Access to the thesis is subject to the Creative Commons Attribution - NonCommercial-No Derivatives 4.0 International Public License. A copy of this may be found at <https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>. This license sets out your rights and the restrictions that apply to your access to the thesis so it is important you read this before proceeding.

Take down policy

Some pages of this thesis may have been removed for copyright restrictions prior to having it been deposited in Explore Bristol Research. However, if you have discovered material within the thesis that you consider to be unlawful e.g. breaches of copyright (either yours or that of a third party) or any other law, including but not limited to those relating to patent, trademark, confidentiality, data protection, obscenity, defamation, libel, then please contact collections-metadata@bristol.ac.uk and include the following information in your message:

- Your contact details
- Bibliographic details for the item, including a URL
- An outline nature of the complaint

Your claim will be investigated and, where appropriate, the item in question will be removed from public view as soon as possible.

Number Field Sieve with provable complexity

Barry van Leeuwen

Supervisor: Dr. A.R. Booker

Chair: Dr. T. Dokchitser (University of Bristol)

Examiners: Dr. J. Bober (University of Bristol)

Dr. S. Siksek (University of Warwick)

A dissertation submitted to the

University of Bristol

in accordance with the requirements for award of the degree of

Master of Science by Research in Mathematics

at the

Faculty of Science

School of Mathematics

June 2, 2020

Word count: 28557

Abstract

In this thesis we give an in-depth introduction to the General Number Field Sieve, as it was used by Buhler, Lenstra, and Pomerance, [17], before looking at one of the modern developments of this algorithm: A randomized version with provable complexity. This version was posited in 2017 by Lee and Venkatesan, [14], and will be preceded by ample material from both algebraic and analytic number theory, Galois theory, and probability theory.

Dedication and Acknowledgements

I want to thank Dr. Andrew Booker, who as my supervisor managed to find what I needed even though it may not have been what I wanted. I also want to thank Dr. James Milne, Dr. Florian Bouyer, and Dr. Lynne Walling for providing some of the material used. I also want to thank Dr. Dan Fretwell, who helped me find my footing when I just started (which feels very long ago) and my mother, Cokky van Leeuwen, who with her sorcery managed to find typos that I could not.

To my wife, Sarah van Leeuwen, who with her continued support and motivation made possible what I thought impossible.

Declaration

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

Signed: Barry van Leeuwen

Date: June 2, 2020

Contents

Abstract	1
Contents	5
Notation	6
1 Introduction	7
2 Algebraic fundamentals	9
2.1 Field extensions	9
2.2 Algebraic number fields and Dedekind domains	12
2.3 Galois theory and the Frobenius element	17
2.4 Norms, discriminants, and ramification	20
2.4.1 Norms and Traces	20
2.4.2 Discriminants and Minkowski's bound	23
2.4.3 Relativity and Transitivity	24
2.5 Dirichlet's Unit Theorem	26
3 The General Number Field Sieve	28
3.1 Introducing the algorithm	29
3.2 The algorithm Explained	31
3.2.1 Choosing a polynomial	31
3.2.2 Sieving over \mathbb{Z}	33
3.2.3 Sieving over $\mathbb{Z}[\alpha]$	35
3.3 Dealing with obstructions	43
3.4 Computational efficiency	49
4 Preparing for randomness	52
4.1 Zeta function and the Prime Number Theorem	52

4.2	Generalizing the Riemann zeta function	54
4.2.1	Dirichlet characters	54
4.2.2	Dirichlet \mathcal{L} -functions	58
4.2.3	Dedekind zeta functions and Siegel zeroes	60
4.2.4	Siegel zeroes	62
4.3	Arithmetic progressions	63
4.3.1	Extensions and Refinements	64
4.3.2	Smooth numbers on Arithmetic Progressions	65
4.4	Probability measures and moments	68
5	The Randomized Number Field Sieve	72
5.1	Introducing the algorithm	72
5.2	Randomizing the polynomial	75
5.2.1	Uniform distribution of the polynomial	76
5.2.2	Uniform Distribution of Smooth Numbers	83
5.2.3	Proof of theorem 5.2	89
5.3	Algebraic obstructions Vol. 2: The congruence	94
5.3.1	Quadratic characters	95
5.3.2	Proof of lemma 5.19	97
5.3.3	Proof of lemma 5.23	104
5.3.4	Proof of theorem 5.3	107
5.4	Computational efficiency	109
5.5	Riemann and the future of factorization	110
	References	112

Notation

A brief introduction to some of the notation used. Most of the notation listed will be formally introduced, but this can be used as a reference guide.

$R[\alpha]$	-	A ring extended by an element α
$F(\alpha)$	-	A field extended by an element α
\mathcal{O}_K	-	The ring of algebraic integers in a field extension K
$f.f.(R)$	-	The field of fractions generated by a ring R
$\text{irr}_F(a)$	-	The minimal polynomial of a over F
$[a] \in B$	-	The coset of a in structure B
$\langle a \rangle \subset B$	-	The ideal generated by a in structure B
$\mathcal{L}(s, \chi)$	-	Dirichlet L-function with character χ and exponent s
$a \sim B$	-	Draw a according to distribution B
$\pi(X)$	-	The number of primes below X

In this paper we use Bachmann–Landau Big-O notation extended by Knuth Big- Ω notation, such that for functions $f(x)$ and $g(x)$ and $N \in \mathbb{N}$:

$f(x) = \mathbf{o}(g(x))$	-	$\forall k \in \mathbb{R}_{>0} \exists N$ such that $\forall x \geq N : f(x) \leq k \cdot g(x)$
$f(x) = \mathbf{O}(g(x))$	-	$\exists k \in \mathbb{R}_{>0} \exists N$ such that $\forall x \geq N : f(x) \leq k \cdot g(x)$
$f(x) = \mathbf{\omega}(g(x))$	-	$\forall k \in \mathbb{R}_{>0} \exists N$ such that $\forall x \geq N : f(x) \geq k \cdot g(x)$
$f(x) = \mathbf{\Omega}(g(x))$	-	$\exists k \in \mathbb{R}_{>0} \exists N$ such that $\forall x \geq N : f(x) \geq k \cdot g(x)$

1 Introduction

In 1988 Pollard introduced a brand new factorization algorithm: The Number Field Sieve (NFS). This saw a first implementation in 1994 by Lenstra et al. and was used to factor the ninth Fermat number, F_9 .

The grandeur of this algorithm was in the conjectured complexity of

$$L_n\left(\frac{1}{3}, \sqrt[3]{\frac{64}{9}} + \mathbf{o}(1)\right),$$

where $n \in \mathbb{N}$ and for $a, b, x \in \mathbb{R}$:

$$L_x(a, b) = \exp(b((\log n)^a (\log \log n)^{1-a})).$$

This was a major improvement to the best algorithms at the time, such as the quadratic sieve, which were of complexity $L_n(\frac{1}{2}, b)$ for some constant b .

Despite abundant heuristics indicating the bound and copious research poured into a proof it was very difficult to show the complexity of the two major parts of the sieve:

1. Using the concept of smoothness to find a factor base.
2. Reducing this factor base to a factorization of a given $n \in \mathbb{N}$.

In chapter 3 we will be looking at the General Number Field Sieve (GNFS) how it was explicitly described in [17] and how the sieve is structured.

It wasn't until 2017 that Lee and Venkatesan, [14]. used a probabilistic and combinatorial approach to alter the algorithm in such a way that they managed to prove the following, ([14], theorem 2.1 & 2.3):

Theorem 1.1. There is a randomised variant of the Number Field Sieve which for each n finds congruences of squares $x^2 \equiv y^2 \pmod{n}$ in expected time:

$$L_n\left(\frac{1}{3}, \sqrt[3]{\frac{64}{9}} + \mathfrak{o}(1)\right)$$

Remark. Note that no claims are made regarding the factorization of n . This theorem only proves that there exists an algorithm for which the sieving process up until finding a congruence of squares, i.e. a pair (x, y) for which $x^2 \equiv y^2 \pmod{n}$, with a non-zero probability in the complexity bound given in the theorem above.

This version of the algorithm will be extensively studied in chapter 5 and finishes by proving the theorem above.

To do this we need to use a strong basis in analytic number theory, which we will introduce in chapter 4 along with a review of some probability theory. In chapter 2 we will recall the necessary concepts from algebraic number theory and Galois theory, but for those familiar with these subjects this chapter can be skipped.

2 Algebraic fundamentals

2.1 Field extensions

The reason the General Number Field Sieve is so effective is that it uses strong algebraic constructions to get a factorization, which allows for far fewer computations to be made to get an effective result. The discussion starts by considering these constructions and concepts, such as field extensions, in all generality, but will swiftly collapse down to the specific cases needed for this sieve. We begin by introducing some notation that we will use throughout.

Definition 2.1.

- Let K and L be fields such that $K \subset L$, then L is called a field extension of K and is denoted L/K .
- $K(\alpha)$ is the smallest field extension of K that contains α . If $\alpha \in K$ then $K(\alpha) = K$.
- $K[x]$ is the polynomial extension of K , such that

$$K[x] = \left\{ \sum_{i=0}^{\infty} a_i x^i \mid \forall i \in \mathbb{Z}^+ : a_i \in K \right\}.$$

Remark. To prevent confusion \subset and \supset are proper sub- and super-constructions, while \subseteq and \supseteq can have equality.

A field extension L/K can be seen as a K -vector space and when considered as such it becomes a natural next step to think about the dimension of L as a K -vectorspace. This is called the degree and is denoted $[L : K] = \dim_K(L)$. If there are multiple extensions, $L \supseteq M \supseteq K$, called a tower of extensions, called a tower of extensions then the degrees of these extensions are very well behaved.

Theorem 2.2. Tower Law

Let $K_n \supseteq K_{n-1} \supseteq \dots \supseteq K_0$ be fields. Then

$$[K_n : K] = [K_n : K_{n-1}] \dots [K_1 : K_0].$$

If $[L : K] < \infty$ the extension is called finite, else it is infinite. To use the sieve it is important to understand when such an extension is finite and when it isn't. Let $f(x)$ be an irreducible monic monovariate polynomial with coefficients in K of degree $\deg(f) = d$. As $f(x)$ is irreducible it has no roots in K , but it might have roots over other fields that contain K .

Definition 2.3. An irreducible polynomial f over a field K is said to be separable over K if it has no multiple zeros in a splitting field. That is there exists a field $L \supset K$ such that

$$f(x) = (x - \alpha_1)(x - \alpha_2) \dots (x - \alpha_d)$$

in L , where $f(\alpha_i) = 0$.

This definition is not extremely helpful and in practice it will be more useful to consider [[4], proposition 9.14] instead:

Proposition 2.4. If K is a subfield of \mathbb{C} then every irreducible polynomial over K is separable.

By the fundamental theorem of algebra it is known that a polynomial $f(x)$ with $\deg(f) = d$ has exactly d roots in \mathbb{C} . These roots are not exclusive, and a root α is the root of many polynomials. However there is a unique monic polynomial of lowest degree that is of specific importance:

Definition 2.5. Let $f \in K[x]$ be a polynomial with coefficients in a field K for which α is a root. Then f is the minimal polynomial of α in K if

- f is monic.
- f is irreducible.
- $\forall g \in K[x]$ such that $g(\alpha) = 0$, f divides g .

This polynomial is unique up to units.

Throughout this text we adopt the convention to denote the minimal polynomial for $\alpha \in K$ over K as $\text{irr}_K(\alpha)$. The minimal polynomial gives rise to the following equivalencies, which are vital in the background throughout.

Theorem 2.6. Let $f \in K[x]$ with $\deg(f) = d$ be a monic irreducible monovariate polynomial with root α . Let $\langle f(x) \rangle$ denote the ideal generated by $f(x)$. Then the following are equivalent:

1. $f(x)$ is the minimal polynomial for α in $K[x]$; $f(x) = \text{irr}_K(\alpha)$.
2. $K(\alpha) \cong K[x]/\langle f(x) \rangle$.
3. $K(\alpha)$ is a field.
4. $[K(\alpha) : K] = d$ and $\{1, \alpha, \alpha^2, \dots, \alpha^{d-1}\}$ is a basis for $K(\alpha)$ over K .

The above theorem reveals a condition for which the extensions $K(\alpha)$ is a finite degree extension of K . This condition is called algebraicity and for an extension L/K an element $\alpha \in L$ is called algebraic exactly if it is a root for a polynomial $f(x) \in K[x]$.

Remark. There are two special cases to consider:

- A complex number, $\alpha \in \mathbb{C}$, that is the root of a polynomial with coefficients in \mathbb{Z} is called an algebraic integer.
- A complex number, $\alpha \in \mathbb{C}$, that is the root of a polynomial with coefficients in \mathbb{Q} is called an algebraic number

It turns out that the structure of the algebraic numbers is very simple to describe.

Theorem 2.7. Every algebraic number is of the form $\frac{\alpha}{b}$ where α is algebraic over \mathbb{Z} and $b \in \mathbb{Z} \setminus \{0\}$.

As just discussed: If $\text{irr}_K(\alpha)$ has degree d , then it is only natural to ask about the other roots of $\text{irr}_K(\alpha)$. These are called the conjugates of α and have some interesting properties.

Definition 2.8. Let $\alpha \in \mathbb{C}$ be algebraic over $K \subseteq \mathbb{C}$. The conjugates of α , $\alpha = \alpha_1, \dots, \alpha_d$ are the roots of $\text{irr}_K(\alpha)$.

As $K \subset \mathbb{C}$ proposition 2.4 shows that any irreducible polynomial over K is separable, hence each α_i is distinct:

$$\text{irr}_K(\alpha) = (x - \alpha_1)(x - \alpha_2) \dots (x - \alpha_d).$$

Using the conjugates [[1], theorem 5.6.1] proves that we only need to consider extensions by a singular algebraic number:

Theorem 2.9. Let $K \subseteq \mathbb{C}$ and let $\alpha, \beta \in \mathbb{C}$ be algebraic over K . Then there exists $\gamma \in \mathbb{C}$ algebraic over K such that

$$K(\alpha, \beta) = K(\gamma).$$

This can easily be extended to any finite extension $K(\alpha_1, \dots, \alpha_n)$ by using the fact that $K(\alpha_1, \alpha_2) = (K(\alpha_1))(\alpha_2)$.

2.2 Algebraic number fields and Dedekind domains

To work effectively with the general number field sieve there are some restrictions we place on our extensions. The first restriction placed upon the extensions is the restriction that they be finite. Hence, from now on, every extension $L/K \subseteq \mathbb{C}$ is assumed finite. Special focus will be on algebraic number fields:

Definition 2.10. Let $K \subseteq \mathbb{C}$, then K is an algebraic number field if $K = \mathbb{Q}(\alpha_1, \dots, \alpha_n)$ such that for all $i \in \mathbb{N}$, α_i is an algebraic number.

By theorem 2.9 we can equate $K = \mathbb{Q}(\alpha)$ for some algebraic number $\alpha \in \mathbb{C}$. Moreover we know that any algebraic number can be written as $\alpha = \frac{\beta}{b}$ where β is an algebraic integer and $b \in \mathbb{Z} \setminus \{0\}$, but $\mathbb{Q}(\frac{\beta}{b}) \cong \mathbb{Q}(\beta)$ as $\frac{1}{b} \in \mathbb{Q}$ so it is safe to assume that α is in fact an algebraic integer.

Definition 2.11. Let A be the set of all algebraic integers and let K be an algebraic number field. Then the set $\mathcal{O}_K = A \cap K$ is called the ring of algebraic integers over K .

As the name suggests, \mathcal{O}_K is a ring. In fact, by [[1], theorem 6.1.4 and 6.1.6],

Theorem 2.12. Let K be an algebraic number field, then \mathcal{O}_K is integrally closed. i.e.

1. \mathcal{O}_K is an integral domain.

2. $\forall \alpha: \alpha$ algebraic over $\mathcal{O}_K \Rightarrow \alpha \in \mathcal{O}_K$.

In generality \mathcal{O}_K is surprisingly well behaved allowing for some interesting properties, including that it's field of fractions is K :

$$f.f.(\mathcal{O}_K) = \left\{ \frac{\alpha}{\beta} \mid \alpha, \beta \in \mathcal{O}_K \right\} = K.$$

Another of the properties which is important is [[1], theorem 6.5.3]:

Theorem 2.13. Let K be an algebraic number field. Then \mathcal{O}_K is a Noetherian domain.

Remark. A terribly inconvenient, yet common, convention is that the ideals $\mathfrak{I} \subseteq \mathcal{O}_K$ are just called ideals of K . Despite the confusion this might causes it will not pose a large problem, and will be clear from context, as in practice K will be a number field, hence only has trivial ideals.

It is nearly by definition that a maximal ideal of an integral domain is a prime ideal, but in general the converse is not true. However, in the case of \mathcal{O}_K for an algebraic number field K this converse is true.

Theorem 2.14. Let $K = \mathbb{Q}(\alpha)$ be an algebraic number field. Let $\mathfrak{P} \subseteq \mathcal{O}_K$ be a prime ideal of \mathcal{O}_K , then \mathfrak{P} is a maximal ideal of \mathcal{O}_K .

Now we have shown that \mathcal{O}_K is integrally closed, is a Noetherian domain, and each prime ideal of \mathcal{O}_K is a maximal ideal. These three define an important class of domains, which are very useful for the General Number Field Sieve.

Definition 2.15. An integral domain D that satisfies the following properties:

- D is a Noetherian domain.
- D is integrally closed.
- Each prime ideal of D is a maximal ideal.

is called a Dedekind domain.

Dedekind domains admit a very interesting property, as they generalize the idea of unique factorization into primes which we know by the fundamental theorem of arithmetic holds for all $z \in \mathbb{Z}$ up to units, ± 1 . In Dedekind domains this holds as well, however we must use the equivalent construction for number fields.

Theorem 2.16. Let D be a Dedekind domain. Any non-trivial integral ideal is a product of prime ideals and this factorization is unique up to order, i.e. for $\mathfrak{I} \subset D$ an ideal with factorization

$$\mathfrak{I} = \prod_{i \in I} \mathfrak{P}_i^{e_i} = \prod_{j \in J} \mathfrak{Q}_j^{f_j},$$

with prime ideals $\mathfrak{P}_i, \mathfrak{Q}_j$ then there exists a permutation $\sigma : I \rightarrow J$ such that $\mathfrak{P}_i = \mathfrak{Q}_j$ and $e_i = f_j$.

To prove this we need some auxiliary measures:

Proposition 2.17. Let D be a Dedekind domain. Then every non-zero ideal $\mathfrak{I} \in \mathcal{O}_K$ contains a product of one or more prime ideals.

Definition 2.18. Let D be an integral domain and let K be the quotient field of D . For each prime ideal \mathfrak{P} of D we define the set $\bar{\mathfrak{P}}$ by

$$\bar{\mathfrak{P}} = \{\alpha \in K : \alpha\mathfrak{P} \subseteq D\}.$$

This set is an ideal and is called a fractional ideal of D .

Proposition 2.19. Let D be a Dedekind domain. Let \mathfrak{P} be a prime ideal of D . Then

$$\mathfrak{P}\bar{\mathfrak{P}} = D.$$

With these facts recorded we proceed with the proof of theorem 2.16:

Proof. This proof will come in two parts: Existence and Uniqueness. Let D be a Dedekind domain such that there exists a non-trivial ideal of D that is not a product of prime ideals.

Existence: As D is a Dedekind domain, it is Noetherian, and so by the maximality principle there is a non-trivial ideal \mathfrak{A} of D that is maximal with respect to the property of not being a product of prime ideals. Then by proposition 2.17 there exists prime ideals $\mathfrak{P}_1, \dots, \mathfrak{P}_k$ of D such that

$$\mathfrak{P}_1 \dots \mathfrak{P}_k \subseteq \mathfrak{A}.$$

Now let k be the smallest positive integer for which such a product exists. If $k = 1$ then $\mathfrak{P}_1 \subseteq \mathfrak{A} \subset D$ is a prime ideal, hence it is maximal. So $\mathfrak{A} = \mathfrak{P}_1$. This contradicts our assumption that \mathfrak{A} was not a product of prime ideals. Hence $k \geq 2$. By proposition 2.19 we have that $\bar{\mathfrak{P}}_1 \mathfrak{P}_1 = D$ and so

$$\bar{\mathfrak{P}}_1 \mathfrak{P}_1 \dots \mathfrak{P}_k = D \mathfrak{P}_2 \dots \mathfrak{P}_k.$$

Hence we have that $\bar{\mathfrak{P}}_2 \mathfrak{A} \supseteq \mathfrak{P}_2 \dots \mathfrak{P}_k$. Moreover by proposition 2.19 we also have that $\mathfrak{A} \subseteq \bar{\mathfrak{P}}_1 \mathfrak{A}$.

Now assume that $\mathfrak{A} = \bar{\mathfrak{P}}_1 \mathfrak{A}$, then

$$\mathfrak{A} \supseteq \mathfrak{P}_2 \dots \mathfrak{P}_k.$$

Which contradicts the minimality of k . Now assume $\mathfrak{A} \subset \bar{\mathfrak{P}}_1 \mathfrak{A}$, then as the latter is an ideal of D , by the maximality property of \mathfrak{A} , we get

$$\bar{\mathfrak{P}}_1 \mathfrak{A} = \mathfrak{Q}_2 \dots \mathfrak{Q}_l,$$

for some prime ideals \mathfrak{Q}_j , $j \in \{2, \dots, l\}$. But then

$$\mathfrak{A} = \mathfrak{A}D = \mathfrak{A}\bar{\mathfrak{P}}_1 \mathfrak{P}_1 = \mathfrak{P}_1 \mathfrak{Q}_2 \dots \mathfrak{Q}_l,$$

which contradicts the way \mathfrak{A} was chosen. Hence every ideal of D is a product of prime ideals.

Uniqueness: Suppose now that there exists a maximal ideal \mathfrak{A} of D , such that

$$\mathfrak{A} = \mathfrak{P}_1 \dots \mathfrak{P}_k = \mathfrak{Q}_1 \dots \mathfrak{Q}_l.$$

By the maximality principle this is a choice we can make. Then as $\mathfrak{P}_1 \dots \mathfrak{P}_k \subseteq \mathfrak{Q}_1$ and \mathfrak{Q}_1 is a prime ideal we have that $\mathfrak{P}_i = \mathfrak{Q}_1$ for some $i \in \{1, \dots, k\}$. Relabelling \mathfrak{Q}_i as \mathfrak{Q}_1 we get $\mathfrak{Q}_1 = \mathfrak{P}_1$ as \mathfrak{P}_1 is a prime, hence maximal, ideal of D . Thus

$$\bar{\mathfrak{Q}}_1 \mathfrak{Q}_1 \dots \mathfrak{Q}_l = \mathfrak{A} \mathfrak{Q}_1 = \mathfrak{A} \bar{\mathfrak{P}}_1 = \bar{\mathfrak{P}}_1 \mathfrak{P}_1 \dots \mathfrak{P}_k = \mathfrak{P}_2 \dots \mathfrak{P}_k.$$

Now assume that $\mathfrak{A} = \mathfrak{A} \bar{\mathfrak{P}}_1$, then $\mathfrak{A} \bar{\mathfrak{P}}_1 \mathfrak{P}_1 = \mathfrak{A} \mathfrak{P}_1$ and so $\mathfrak{A} = \mathfrak{A} \mathfrak{P}_1$. Define the fractional ideal of \mathfrak{A} as

$$\bar{\mathfrak{A}} = \bar{\mathfrak{P}}_1 \dots \bar{\mathfrak{P}}_k.$$

Then

$$\mathfrak{A} \bar{\mathfrak{A}} = \mathfrak{P}_1 \dots \mathfrak{P}_k \bar{\mathfrak{P}}_1 \dots \bar{\mathfrak{P}}_k.$$

and so

$$D = \bar{\mathfrak{A}} \mathfrak{A} = \mathfrak{A} \bar{\mathfrak{P}}_1 \bar{\mathfrak{A}} = \mathfrak{P}_1.$$

but this contradicts the primality of \mathfrak{P}_1 , so our equality assumption was faulty. Hence $\mathfrak{A} \subset \mathfrak{A} \bar{\mathfrak{P}}_1$. Since $\mathfrak{A} \bar{\mathfrak{P}}_1$ is a maximal ideal of D , $\mathfrak{A} \bar{\mathfrak{P}}_1$ has exactly one factorization as a product of prime ideals. Thus we deduce that $k - 1 = l - 1$ and so $k = l$. Relabeling this gives that for all $i \in \{1, \dots, k\}$

$$\mathfrak{P}_i = \mathfrak{Q}_i,$$

proving that the decomposition into prime ideals is unique. □

2.3 Galois theory and the Frobenius element

In this section we give a very concise introduction to a few ideas from Galois Theory. If you have a basic understanding of finite Galois theory then there is no loss in simply skipping this section.

Definition 2.20. Let L/K be number fields. Then L is said to be normal over K if every irreducible polynomial $f \in K[x]$ that has a zero in L splits over L .

Normality is easily checked under certain conditions:

Definition 2.21. Let L/K be number fields. Then L is a splitting field for K if for the polynomial $f \in K[x]$:

- f splits over L .
- There is no M , $L/M/K$, such that $M \neq L$ and f splits over M .

Theorem 2.22. A field extension L/K is normal and finite if and only if L is a splitting field for some polynomial $f \in K[x]$.

We can also speak of separable extensions

Definition 2.23. Let K be a number field. Then a polynomial $f \in K[x]$ is said to be separable if it splits into linear terms in K .

Definition 2.24. Let L/K be number fields. Then L is said to be a separable extension of K if for every $\alpha \in L$, $\text{irr}_K(\alpha)$ is separable over K .

Remark. A finite, normal, and separable extension of a number field is also called a Galois extension.

For separable extensions we have an equivalence allowing us to reduce to one irreducible polynomial.:

Proposition 2.25. For an extensions L/K the following are equivalent:

1. L is a Galois extension of K .
2. L is the splitting field of a separable polynomial $f \in K[x]$.

With this in place we will now start talking about automorphisms. First we define the Galois group:

Definition 2.26. Let L/K be number fields. Then the Galois group of L over K is defined as the group of K -automorphisms, i.e. the group of automorphisms $\pi : L \rightarrow L$ that is fixed for all elements in K . The group operation is composition of maps and the group is denoted $\text{Gal}(L/K)$.

The existence of such a group is trivial as the identity function on L is a K -automorphism, so will always lie in the Galois group. We make a distinction between real and complex K -automorphisms: A K -automorphism $\sigma \in \text{Gal}(L/K)$ is real if $\sigma : L \rightarrow \mathbb{R}$ and complex if $\sigma : L \rightarrow \mathbb{C}$.

Definition 2.27. Let L/K be number fields with Galois group $\text{Gal}(L/K)$ such that $|\text{Gal}(L/K)| = n$. Then the signature of $\text{Gal}(L/K)$ is $\text{sign}(\text{Gal}(L/K)) = (r, s)$, where r is the number of real K -automorphisms and s is half the number of complex K -automorphisms in $\text{Gal}(L/K)$ such that $n = r + 2s$.

There is one element of the Galois group that is particularly interesting, which is the Frobenius element. The idea of this comes from the Frobenius automorphism, which we recall for good measure.

Definition 2.28. If K is a field of characteristic $p > 0$, the map $\phi : K \rightarrow K$ defined by $k \mapsto k^p$ is called the Frobenius monomorphism of K . When K is finite, ϕ is called the Frobenius automorphism.

To be able to define the Frobenius elements over extensions we have a little setting up to do. For this let L/K be a Galois extension with Galois group $\text{Gal}(L/K)$. Let $\mathfrak{P} \subseteq \mathcal{O}_L$ lying over

$\mathfrak{p} \subseteq \mathcal{O}_K$. Then we can define the decomposition group $D(\mathfrak{P} | \mathfrak{p})$ as the set of automorphisms in $\text{Gal}(L/K)$ which fix \mathfrak{P} :

$$D(\mathfrak{P} | \mathfrak{p}) = \{\sigma \in \text{Gal}(L/K) \mid \sigma(\mathfrak{P}) = \mathfrak{P}\}.$$

Definition 2.29. Let L/K be a Galois extension with Galois group $\text{Gal}(L/K)$, let $\mathfrak{P} \subseteq \mathcal{O}_L$ be a prime ideal unramified in L/K lying over $\mathfrak{p} \in \mathcal{O}_K$. Then the Frobenius element, $\text{Frob}_{\mathfrak{P}}$, is the element of $D(\mathfrak{P} | \mathfrak{p})$ that acts as the Frobenius automorphism on the residue field. I.e. the unique $\text{Frob}_{\mathfrak{P}} \in D(\mathfrak{P} | \mathfrak{p})$ such that

- $\text{Frob}_{\mathfrak{P}} \in D(\mathfrak{P} | \mathfrak{p})$.
- For all $\alpha \in \mathcal{O}_L$, $\text{Frob}_{\mathfrak{P}}(\alpha) \equiv \alpha^q \pmod{\mathfrak{P}}$, where $q = |\mathcal{O}_K/\mathfrak{p}|$.

This Frobenius element for \mathfrak{P} is uniquely determined, but is not unrelated. For example, if \mathfrak{P} and \mathfrak{P}' both lie over $\mathfrak{p} \subseteq \mathcal{O}_K$ then the Frobenius elements are conjugate. This means that each prime ideal in \mathcal{O}_K gives rise to a conjugacy class of Frobenius elements in $\text{Gal}(L/K)$.

Proposition 2.30. Let L/K be a Galois extension with Galois group $\text{Gal}(L/K)$, let $\mathfrak{P}, \mathfrak{P}' \subseteq \mathcal{O}_L$ be prime ideals unramified in L/K , both lying over $\mathfrak{p} \in \mathcal{O}_K$. Then there exists $\sigma \in \text{Gal}(L/K)$ such that

$$\text{Frob}_{\mathfrak{P}} = \sigma \text{Frob}'_{\mathfrak{P}'} \sigma^{-1}.$$

Proof. Let $\alpha \in \mathcal{O}_L$. Then

$$\sigma \text{Frob}_{\mathfrak{P}} \sigma^{-1}(\alpha) = \sigma \left((\sigma^{-1}(\alpha))^q + a \right),$$

for some $a \in \mathfrak{P}$, and

$$\sigma \left((\sigma^{-1}(\alpha))^q + a \right) = \alpha^q + \sigma(a) \equiv \alpha^q \pmod{\sigma\mathfrak{P}}.$$

□

This is all we need for now, but in chapter 5 we will be expand our theory a little more.

2.4 Norms, discriminants, and ramification

2.4.1 Norms and Traces

Let $K = \mathbb{Q}(\alpha)$ be a number field with Galois group $\text{Gal}(K/\mathbb{Q})$. Let \mathcal{O}_K be the ring of algebraic integers of K . It will prove useful in what is to come that we can relate the elements of a number field to rational number in \mathbb{Q} . There are two natural maps to consider:

Definition 2.31. Let L/K be an extension of number fields with $[L : K] = n$. Let $\sigma_1, \dots, \sigma_n$ be the K -automorphisms of L in \mathbb{C} . For any $\alpha \in L$ we define the following maps:

1. The trace of an element α :

$$\text{Tr}_{L/K}(\alpha) : L \rightarrow \mathbb{C},$$

$$\alpha \mapsto \sum_{i=1}^n \sigma_i(\alpha).$$

2. The field norm of an element α :

$$N_{L/K} : L \rightarrow \mathbb{C},$$

$$\alpha \mapsto \prod_{i=1}^n \sigma_i(\alpha)^{[L:K(\alpha)]}.$$

Remark. If it is clear from context what field we are taking the norm over we will often just write $N_{L/K} = N$.

We will mostly be concerned with Galois extensions which makes the field norm far simpler. For this consider L/K with Galois group $\text{Gal}(L/K)$, then for any $\alpha \in L$ the norm becomes

$$N_{L/K}(\alpha) = \prod_{\sigma \in \text{Gal}(L/K)} \sigma(\alpha).$$

These two maps have a few important properties we will make use of throughout.

Proposition 2.32.

1. $N_{L/K}$ and $\text{Tr}_{L/K}$ are homomorphisms from $L \rightarrow K$.
2. If $\alpha \in \mathcal{O}_L$ then $N_{L/K}(\alpha) \in \mathcal{O}_K$ and $\text{Tr}_{L/K}(\alpha) \in \mathcal{O}_K$.
3. Let $\alpha, \beta \in \mathcal{O}_L$ such that $\alpha \mid \beta$, then $N_{L/K}(\alpha) \mid N_{L/K}(\beta)$.
4. If $k \in K$, then $N_{L/K}(k) = k^n$ and $\text{Tr}_{L/K}(k) = nk$, where $[L : K] = n$.

Remark. Note that (2) implies that for all $\alpha \in \mathcal{O}_L$: $\alpha \in \mathcal{O}_L^* \Leftrightarrow N_{L/\mathbb{Q}}(\alpha) = \pm 1$ as $\mathcal{O}_{\mathbb{Q}} = \mathbb{Z}$ and $\mathbb{Z}^* = \pm 1$.

Despite the usefulness of the trace and field norm maps they are rather limited in that they are only defined for elements in the number field. We can also define the norm of an ideal.¹

Definition 2.33. Let K be a number field of finite degree. Let \mathcal{O}_K be the ring of algebraic integers of K . Let $\mathfrak{J} \subset \mathcal{O}_K$ be a non-trivial ideal of \mathcal{O}_K . Then

$$N(\mathfrak{J}) = |\mathcal{O}_K/\mathfrak{J}|.$$

However these definitions are not completely independent from each other, as they are related for principal ideals. In fact, for principal ideals the field and ideal norm overlap.

Theorem 2.34. Let K be a number field of finite degree. Let $\mathfrak{J} = \langle \alpha \rangle \subset \mathcal{O}_K$. Then

$$|\mathcal{O}_K/\mathfrak{J}| = N(\mathfrak{J}) = N(\langle \alpha \rangle) = |N(\alpha)|.$$

It will not come as a surprise that the ideal norm is also a homomorphism, such that for every two non-zero integral ideals $\mathfrak{A}, \mathfrak{B} \in \mathcal{O}_K$, $N(\mathfrak{A}\mathfrak{B}) = N(\mathfrak{A})N(\mathfrak{B})$. This is absolutely not trivial, but intuitive enough for us to understand to leave the proof to [[1], theorem 9.3.2]. Now we have a definition for the ideal norm and an explicit, and reasonably simple, way to compute this for principal ideals. Up until now we do not have a general way to compute the ideal norm,

¹There are multiple ways to define the norm of an ideal. The definition we adopt has the benefit of being completely algebraic, which is why we use it.

but we can get closer by recalling theorem 2.16 as the ideal norm is a homomorphism, so we just have to consider how to compute the ideal norm for prime ideals.

Theorem 2.35. Let K be an algebraic number field. Let $\mathfrak{P} \subset \mathcal{O}_K$. Then there exists a unique prime $p \in \mathbb{Z}$, such that

$$\mathfrak{P} \mid \langle p \rangle.$$

As this $p \in \mathbb{Z}$ is unique we say \mathfrak{P} lies over p , or equivalently p lies below \mathfrak{P} . This has an immediate effect on the ideal norm, as can be seen in the following theorem:

Theorem 2.36. Let K be an algebraic number field with $[K : \mathbb{Q}] = n$. Let $\mathfrak{P} \subset \mathcal{O}_K$ be a prime ideal lying over $p \in \mathbb{Z}$. Then

$$N(\mathfrak{P}) = p^f,$$

for some integer $f \in \{1, \dots, n\}$.

Proof. As \mathfrak{P} lies above p we have that $\mathfrak{P} \mid \langle p \rangle$. Hence $\langle p \rangle = \mathfrak{P}\mathfrak{Q}$ for some integral ideal $\mathfrak{Q} \subset \mathcal{O}_K$. As the norm is multiplicative we have that

$$N(\langle p \rangle) = N(\mathfrak{P})N(\mathfrak{Q}).$$

As the K -conjugates of p comprise p n times we have

$$N(p) = p^n,$$

and so

$$N(\langle p \rangle) = |N(p)| = N(\mathfrak{P})N(\mathfrak{Q}) = p^n.$$

□

The f in the theorem is called the inertial degree and is used to define one of the most interesting objects in number theory, the ideal class group, which will define momentarily. First we need some preliminary definitions and properties.

Definition 2.37. Let K be an algebraic number field with $[K : \mathbb{Q}] = n$. Let $p \in \mathbb{Z}$ be prime such that

$$\langle p \rangle = \prod_{i=1}^g \mathfrak{P}_i^{e_i}.$$

Then g is called the decomposition number, e_i is the ramification index for \mathfrak{P}_i , and the following properties hold:

1. The inertia degree, f_i , can be computed as $f_i = [\mathcal{O}_K/\mathfrak{P}_i : \mathbb{Z}/\langle p \rangle]$.
2. p is said to be ramified in K if, for some i , $e_i > 1$.
3. p is said to be unramified in K if for all i : $e_i = 1$.
4. The following equation holds:

$$\sum_{i=1}^g e_i f_i = n.$$

5. p is said to be (totally) split if for all i : $e_i = f_i = 1$, equivalently $g = n$.

Remark. This can also be defined generally for extensions L/K , and follows naturally from the above definition. For our purposes the above will suffice however.

2.4.2 Discriminants and Minkowski's bound

We now define the concept of the (absolute) discriminant of a number field and immediately apply it to a theorem by Dedekind:

Definition 2.38. Let K be a number field of finite degree with ring of algebraic integers \mathcal{O}_K . Let $\{a_1, \dots, a_n\}$ be an integral basis for \mathcal{O}_K and let $\{\sigma_1, \dots, \sigma_n\}$ be the set of K -automorphisms into \mathbb{C} . Then the (absolute) discriminant of K is

$$\Delta_K = \det \begin{pmatrix} \sigma_1(a_1) & \dots & \sigma_1(a_n) \\ \vdots & \ddots & \vdots \\ \sigma_n(a_1) & \dots & \sigma_n(a_n) \end{pmatrix}^2.$$

Theorem 2.39. Let K be an algebraic number field. Then the prime $p \in \mathbb{Z}$ ramifies in K if and only if $p \mid \Delta_K$.

Now we define the ideal class group:

Definition 2.40. Let K be an algebraic number field and let $\mathfrak{A}, \mathfrak{B} \subset \mathcal{O}_K$ be non-zero ideals. Then we say that $\mathfrak{A} \sim \mathfrak{B}$, that is they are equivalent, if there exists $\alpha, \beta \in \mathcal{O}_K \setminus \{0\}$ such that $\langle \beta \rangle \mathfrak{A} = \langle \alpha \rangle \mathfrak{B}$. The equivalence classes under \sim are called ideal classes and the set of ideal classes of K , denoted Cl_K , is called the ideal class group. Moreover, the cardinality of the ideal class group, $|\text{Cl}_K| = h_K$, is called the class number of K .

It is well known that h_K is finite and that Cl_K is an abelian group for K an algebraic number field. More can be said, especially about the generation of Cl_K , and we will record these in quick succession.

Theorem 2.41. Let $\mathfrak{J} \subseteq \mathcal{O}_K$ be a non-zero ideal. Then there exists a non-zero $\alpha \in \mathfrak{J}$ with

$$|N_{K/\mathbb{Q}}(\alpha)| \leq c_K N(\mathfrak{J}),$$

for

$$c_K = \left(\frac{4}{\pi}\right)^{r_2} \frac{n!}{n^n} \sqrt{|\Delta_K|},$$

where $2r_2 = |\{\sigma_i \mid \sigma_i \text{ a } \mathbb{Q}\text{-automorphism, } \sigma_i(K) \not\subseteq \mathbb{R}\}|$.

Remark. c_K is called the Minkowski bound.

Theorem 2.42. Any ideal class $c \in \text{Cl}_K$ contains an ideal \mathfrak{J} such that $N(\mathfrak{J}) \leq c_K$.

Theorem 2.43. Let K be a number field. Then Cl_K is generated by the classes of the prime ideals $[\mathfrak{P}]$, satisfying $N(\mathfrak{P}) \leq c_K$. In particular, any such \mathfrak{P} must lie above a prime $p \leq c_K$.

2.4.3 Relativity and Transitivity

In the last section we defined the absolute discriminant. It will not surprise the reader that there then also is a relative discriminant. To do this we need to consider the following:

Definition 2.44. Let L/K be number field with rings of integers \mathcal{O}_L and \mathcal{O}_K respectively. The fractional ideal ,

$$\mathfrak{C}_{\mathcal{O}_L|\mathcal{O}_K} = \{x \in L \mid \text{Tr}(x\mathcal{O}_L) \subset \mathcal{O}_K\},$$

is called Dedekind's complementary module. It's inverse,

$$\mathfrak{D}_{\mathcal{O}_L|\mathcal{O}_K} = \mathfrak{C}_{\mathcal{O}_L|\mathcal{O}_K}^{-1},$$

is called the different of $\mathcal{O}_L/\mathcal{O}_K$.

Remark. With the usual abuse of notation we will write $\mathfrak{D}_{L/K}$ for the different $\mathfrak{D}_{\mathcal{O}_L|\mathcal{O}_K}$

By definition Dedekind's complementary module contains \mathcal{O}_L and so $\mathfrak{D}_{L/K}$ is in fact an integral ideal. The different is well-behaved in towers of number fields:

Proposition 2.45. For a tower of fields $L/M/K/\mathbb{Q}$: $\mathfrak{D}_{L/K} = \mathfrak{D}_{L/M}\mathfrak{D}_{M/K}$.

This different ideal, or simply 'different', is key in the definition of the relative discriminant;

Definition 2.46. Let L/K be number fields. The relative discriminant $\Delta_{L/K}$ is the ideal of \mathcal{O}_K defined by the relative norm of the different:

$$\Delta_{L/K} = N_{L/K}(\mathfrak{D}_{L/K}).$$

This allows us to define a transitivity condition for the discriminant

Lemma 2.47. Let $L/M/K$ be a tower of number fields. Then

$$\Delta_{L/K} = \Delta_{M/K}^{[L:M]} N_{M/K}(\Delta_{L/M}).$$

Proof. Applying to $\mathfrak{D}_{L/K} = \mathfrak{D}_{L/M}\mathfrak{D}_{M/K}$ the norm $N_{L/K} = N_{M/K} \circ N_{L/M}$ we obtain

$$\Delta_{L/K} = N_{M/K}(\Delta_{L/M})N_{M/K}(\mathfrak{D}_{M/K}^{[L:M]}) = N_{M/K}(\Delta_{L/M})\Delta_{M/K}^{[L:M]}.$$

□

2.5 Dirichlet's Unit Theorem

In this last preliminary section we want to get an idea of how \mathcal{O}_K looks. Specifically, we want to get an idea of the units in \mathcal{O}_K . Note that the units of a ring form a group called the unit group and is denoted \mathcal{O}_K^\times .

Definition 2.48. Let $n \in \mathbb{N}$ and $\zeta_n = e^{2\pi i/n}$. Then ζ_n is a n th root of unity and the number field $\mathbb{Q}(\zeta_n)$ is called the n th cyclotomic field.

These cyclotomic fields are in many ways the easiest number fields to work with because they behave especially nicely when n is an odd prime. To emphasize this let $p \in \mathbb{Z}$ be an odd prime for the remainder of the section.

Theorem 2.49. Let $K = \mathbb{Q}(\zeta_p)$, then $\mathcal{O}_K = \mathbb{Z}[\zeta_p]$.

With that little note we return to what is at hand. Let $\mu(K)$ be the group of units contained in a number field K . Note that $\mu(K)$ is a cyclic group of order n under multiplication. The following proposition is then a tautology:

Proposition 2.50. Let K be a number field. Then $\mu(K) \subset \mathcal{O}_K^\times$.

The question now becomes: What are the other elements in \mathcal{O}_K^\times ? Recall from 2.3 that the signature of a field, $\text{sign}(K) = (r, s)$, where r is the number of real \mathbb{Q} -embeddings and $2s$ the number of complex \mathbb{Q} -embeddings.

Theorem 2.51. Let K be a number field. There exists a map $l : K \rightarrow \mathbb{R}^{r+s-1}$ such that $\ker(l)$ is finite and cyclic, and $\ker(l) \cong \mu(K)$. Moreover

$$l(\mathcal{O}_K^\times) \cong \mathbb{Z}^{r+s-1}.$$

This immediately leads us to the concluding theorem:

Theorem 2.52. Dirichlet's Unit Theorem Let K be a number field and let $\text{sign}(K) = (r, s)$. Then

$$\mathcal{O}_K^\times \cong \mu(K) \times \mathbb{Z}^{r+s-1}.$$

Proof. From theorem 2.51 we have that there exists a map l such that $\ker(l) \cong \mu(K)$ hence by the isomorphism theorems

$$K/\mu(K) = K/\ker(l) \cong \text{im}(l) = \mathbb{Z}^{r+s-1},$$

and as $\mu(K)$ is an abelian group and \mathbb{Z}^{r+s-1} is free we get that $K \cong \mu(K) \times \mathbb{Z}^{r+s-1}$. □

3 The General Number Field Sieve

For the entirety of this section let $n \in \mathbb{N}$ be a given integer that we wish to factor. Without loss of generality assume that this n is odd, else we could simply consider $m \in \mathbb{N}$ where $n = 2^s m$, and that n is composite, else a primality test, which can be done in polynomial time, such as the AKS Primality Test [18], will show that n does not need factoring. The General Number Field Sieve (GNFS) then uses a congruence of squares modulo n to perform integer factorization. Assume that (x, y) is a pair that admits a congruence of squares modulo n so that

$$x^2 \equiv y^2 \pmod{n},$$

where $x \neq \pm y$. Then we can use (x, y) to find a non-trivial factorization by using the well-established fact that $x^2 - y^2 = (x + y)(x - y)$. We then compute $\gcd(x + y, n) = d_1$ and $\gcd(x - y, n) = d_2$, in the hope that this gives a non-trivial factorization, that is $d_1, d_2 \neq 1$ or n . Assume that we are able to produce random integers x and y such that the above congruence holds. Then for the hardest case, $n = pq$ semiprime, Table 1 shows an exhaustive truth table:

Table 1: Factorization cases for $n = pq$

$p \mid x + y$	$p \mid x - y$	$q \mid x + y$	$q \mid x - y$	$\gcd(x + y, n)$	$\gcd(x - y, n)$	Factor?
T	T	T	T	n	n	F
T	T	T	F	n	p	T
T	T	F	T	p	n	T
T	F	T	T	n	q	T
T	F	T	F	n	1	F
T	F	F	T	p	q	T
F	T	T	T	q	n	T
F	T	T	F	q	p	T
F	T	F	T	1	n	F

It can be seen that we only have failure if p and q have the same conditions on the factorization of $x + y$ and $x - y$. This is because that means that $x \equiv \pm y \pmod{n}$. If that is not the case then we always find a non-trivial factorization. How the NFS operates to find such congruent pairs (x, y) is the goal of this section.

3.1 Introducing the algorithm

The Number Field Sieve is so named because of a sieving process happening both over \mathbb{Z} and a ring $\mathbb{Z}[\alpha]$ which are contained in the field \mathbb{Q} and the number field $\mathbb{Q}(\alpha)$ respectively. In this section we will give a description on how this is done, while we go into detail, introducing formal definitions and rigor, in Section 3.3. However, the following commuting diagram captures the whole process:

$$\begin{array}{ccc}
 & \mathbb{Z}[x] & \\
 x \mapsto m \swarrow & & \searrow \text{mod } f \\
 \mathbb{Z} & & \mathbb{Z}[x]/(f) \cong \mathbb{Z}[\alpha] \\
 \text{mod } n \searrow & & \swarrow x \mapsto m \text{ mod } n \\
 & \mathbb{Z}/n\mathbb{Z} &
 \end{array}$$

Before we can discuss how the NFS obtains a congruence we introduce the concept of smooth numbers.

Definition 3.1. Let $z \in \mathbb{Z}$ and let $B \in \mathbb{N}$. Then z is said to be B -smooth if in the prime factorization of z ,

$$z = \pm \prod_{i=1}^n p_i^{e_i},$$

it holds that $p_i \leq B$ for all $i = 1, \dots, n$.

It is clear that this definition is insufficient for $\mathcal{O}_{\mathbb{Q}(\alpha)}$ as this ring does not solely consist of integers, however that is easily amended.

For the sake of exposition we will consider $\mathcal{O}_{\mathbb{Q}(\alpha)} = \mathbb{Z}[\alpha]$. This is rarely the case for number fields in general as this only happens for monogenic fields, but in section 3.3 we will see that this assumption can be dealt with algebraically so we only need to concern ourselves with this case. Let $a - b\alpha \in \mathbb{Z}[\alpha]$ and let $\langle a - b\alpha \rangle$ be the ideal generated by it. As $\mathbb{Q}(\alpha)$ is an algebraic number field we get by that $\mathbb{Z}[\alpha]$ is a Dedekind domain, hence any ideal splits uniquely into a product of prime ideals:

$$\langle a - b\alpha \rangle = \prod_{i=1}^n \mathfrak{P}_i^{\eta_i},$$

Moreover by theorem 2.36 we know that any prime ideal of $\mathbb{Z}[\alpha]$ evaluates to a prime power under the norm map. This allows us to define an analogy to smoothness over \mathbb{Z} :

Definition 3.2. Let $a - b\alpha \in \mathbb{Z}[\alpha]$ and let $B \in \mathbb{N}$. Consider $\langle a - b\alpha \rangle \subseteq \mathbb{Z}[\alpha]$ with factorization into prime ideals

$$\langle a - b\alpha \rangle = \prod_{i=1}^n \mathfrak{P}_i^{\eta_i}.$$

Then $a - b\alpha$ is said to be B -smooth if for all $N(\mathfrak{P}_i) = p_i^{f_i}$ it holds that

$$p_i \leq B,$$

where $N : \mathbb{Q}(\alpha) \rightarrow \mathbb{Q}$ is the field norm.

Using these definitions we can, for some fixed m to be defined later, define sets $S_{\mathbb{Z}}$ and $S_{\mathbb{Z}[\alpha]}$ with smoothness bounds B and B' respectively, as follows:

$$S_{\mathbb{Z}} = \{(a, b) \in \mathbb{Z}^2 \mid \gcd(a, b) = 1, a - bm \text{ is } B\text{-smooth as in Definition 3.1}\}$$

$$S_{\mathbb{Z}[\alpha]} = \{(a, b) \in \mathbb{Z}^2 \mid \gcd(a, b) = 1, a - b\alpha \text{ is } B'\text{-smooth as in Definition 3.2}\}$$

The bulk of the work done in the GNFS is done in determining these sets and we will go in much more detail in the next section. If we succeed in finding a set $S \subseteq S_{\mathbb{Z}} \cap S_{\mathbb{Z}[\alpha]}$, such that $|S| \geq \pi(B) + \pi(B') + 1$, then the GNFS uses the remainder of its computation time to find a set $T \subset S$ such that

$$\prod_{(a,b) \in T} (a - bm) \text{ is a square in } \mathbb{Z}, \tag{1}$$

$$\prod_{(a,b) \in T} (a - b\alpha) \text{ is a square in } \mathbb{Z}[\alpha]. \tag{2}$$

It is not a certainty that such a set T fulfilling these equations can be found. If it is not then

the GNFS can be restarted with larger smoothness bounds B and B' . For now assume we have found such a set T . From here completion of the algorithm is nearly trivial:

First we observe that $f'(m)^2 \prod_{(a,b) \in T} (a-bm)$ must give a square in \mathbb{Z} , z^2 say, and $f'(\alpha)^2 \prod_{(a,b) \in T} (a-b\alpha)$ must give a square in $\mathbb{Z}[\alpha]$, ξ^2 say, for which there exist elementary methods to find the square roots z and ξ .

Remark. We have multiplied equation (1) and (2) by the square of the derivative of a polynomial at m and α respectively. This is to ensure that $z \in \mathbb{Z}$ and $\xi \in \mathbb{Z}[\alpha]$. A more thorough discussion of this follows in 3.3.

From here we compute $\gcd(z - N(\xi), n)$ and $\gcd(z + N(\xi), n)$. If this is non-trivial, then the result is a prime factor of n . Else we have to conclude failure and start again with different parameters.

This describes the algorithm in grand lines. Now it is time some mathematical rigor is introduced.

3.2 The algorithm Explained

3.2.1 Choosing a polynomial

The first step of the NFS is that we must choose a $d \in \mathbb{N}, d \neq 1$ and define $m = \lfloor n^{\frac{1}{d}} \rfloor$ s.t. $\gcd(m, n) = 1$. This choice of m gives rise to a polynomial $f \in \mathbb{Z}[x]$ such that $n \mid f(m)$. Over the years there have been many attempts at making the choice of polynomial as effective as possible, but the simplest and effectively cheapest² way to do this is to use the “base- m ” method. For each integer $m \in \mathbb{Z}$ we can write n as a linear combinations of powers of m such that

$$n = \sum_{i=0}^d c_i m^i.$$

²When we say “cheapest” we mean computationally most effective.

Then we can define a polynomial $f \in \mathbb{Z}[x]$ such that

$$f(x) = \sum_{i=0}^d c_i x^i,$$

where the $c_i \in \{0, \dots, m-1\}$ are the same as the base- m expansion of n . This guarantees that $n \mid f(m)$, in fact $f(m) = n$, and $\deg(f) = d$. By [[17], prop. 3.2] we have that the leading coefficient $c_d = 1$, such that $f \in \mathbb{Z}[x]$ is in fact monic.

Remark. Note that as $|c_i| \in \{0, \dots, m-1\}$ hence $c_i < m < n^{\frac{1}{d}}$ we can see that the discriminant of f satisfies that

$$|\Delta(f)| < d^{2d} n^{2-\frac{3}{d}}.$$

We may also assume that $f \in \mathbb{Z}[x]$ is irreducible. To see this assume, to the contrary, that $f \in \mathbb{Z}[x]$ is reducible. This means that there exist non-trivial polynomials $g, h \in \mathbb{Z}[x]$ such that

$$f(x) = g(x)h(x).$$

So a simple computation gives a non-trivial factorization, as $g(x)$ and $h(x)$ are assumed to be non-trivial. As we may now assume $f(x)$ to be monic and irreducible we have shown the following:

Proposition 3.3. Let $f(x)$ be a monovariate, monic, irreducible polynomial, with $\deg(f) = d$, in $\mathbb{Z}[x]$ with roots $\alpha_1, \dots, \alpha_d$, not necessarily in \mathbb{Z} . Then for all $i = 1, \dots, d$,

$$f(x) = \text{irr}_{\mathbb{Z}}(\alpha_i).$$

Note that by assumption $\deg(f) = d > 1$. As f is an irreducible polynomial over $\mathbb{Z}[x]$ we have that none of the roots of f lie in \mathbb{Z} , for else f would split into linear factors over \mathbb{Z} . This root, α say, lies in \mathbb{C} and as $f(\alpha) = 0$ this means that α is an algebraic integer. Moreover, by theorem 2.6 we have that $\mathbb{Q}(\alpha)$ is a field and has basis $\{1, \alpha, \dots, \alpha^{d-1}\}$ over \mathbb{Q} .

3.2.2 Sieving over \mathbb{Z}

Now that we have defined a polynomial $f \in \mathbb{Z}[x]$ and we have defined a field extension $\mathbb{Q}(\alpha)$ we can start sieving for pairs (a, b) such that equation (1) and (2) hold. For the sake of exposition we will focus on the mathematical aspects and therefore consider the pairs (a, b) over \mathbb{Z} and $\mathbb{Z}[\alpha]$ separately.

Definition 3.4. Let $u \in \mathbb{Z}$, then the set of integer pairs defined by

$$U = \{(a, b) \in \mathbb{Z} \mid |a| \leq u, 0 < b < u, \gcd(a, b) = 1\}$$

is called the universe of sieving.

Now let u be a large integer dependent on the to-factor n and define

$$\mathcal{B} = \{p \in \mathbb{Z} \mid p \text{ prime}, p \leq B\} \cup \{\pm 1\}$$

as the rational factor base of the sieve³, fully dependent on the chosen smoothness bound B .

Then there is a standard procedure to work through to fill up the set $S_{\mathbb{Z}}$:

³The rational factor base, normally, only contains the primes $p \leq B$. In our case, however, appending ± 1 improves the effectivity as we can now deal with $a - bm < 0$.

Algorithm 1 Procedure to populate $S_{\mathbb{Z}}$

Input: Universe U of (a, b) pairs, smoothness bound B

Output: Set $S_{\mathbb{Z}} = \{(a, b) \in U \mid a - bm \text{ is } B\text{-smooth, } \gcd(a, b) = 1\}$

Initialize array comprised of $a - mb$ for all $(a, b) \in U$

for each element in the array **do**

 compute the set, $P_{(a,b)}$, of primes p , $p \leq B$, such that $a \equiv bm \pmod{p}$

for each $p \in P_{(a,b)}$ **do** divide $a - bm$ by the maximal power of p , such that p does not divide the quotient, and replace $a - bm$ by this quotient

end for

if the the element in the array is ± 1 **and** $\gcd(a, b) = 1$ **then**

 add (a, b) to $S_{\mathbb{Z}}$

end if

end for

We will accept that this algorithm terminates by choosing B large enough and returns the set $S_{\mathbb{Z}}$ such that $|S_{\mathbb{Z}}| \geq \pi(B) + 1$. As we have recorded the vectors $\{e_p(a - bm)\}_{p \leq B}$ of p -exponents for all $p \leq B$ prime for each $a - bm$ we now define $M \in \mathcal{M}_{|S_{\mathbb{Z}}| \times \pi(B)}(\mathbb{Z})$ given by: $\forall (a_i, b_i) \in S_{\mathbb{Z}}, \forall p_j \in B$:

$$M = \begin{pmatrix} \text{sign}(a_1 - b_1 m) & e_{p_1}(a_1 - b_1 m) & \cdots & e_{p_{\pi(B)}}(a_1 - b_1 m) \\ \cdots & \cdots & \cdots & \cdots \\ \text{sign}(a_{|S_{\mathbb{Z}}|} - b_{|S_{\mathbb{Z}}|} m) & e_{p_1}(a_{|S_{\mathbb{Z}}|} - b_{|S_{\mathbb{Z}}|} m) & \cdots & e_{p_{\pi(B)}}(a_{|S_{\mathbb{Z}}|} - b_{|S_{\mathbb{Z}}|} m) \end{pmatrix},$$

where the sign-bit is defined as:

$$\text{sign}(a_i - b_i m) = \begin{cases} 1 & a_i - b_i m < 0 \\ 0 & \text{otherwise} \end{cases}.$$

To achieve the situation in which equation (1) holds we now want to find an independent subset of $S_{\mathbb{Z}}$. For this we first consider a pair $a - bm \in S_{\mathbb{Z}}$, then it is easily observed that for this to be

a square in \mathbb{Z} we must have that

$$x^2 = a - bm = \prod_{p \in \mathcal{B}} p^{e(p)} = \left(\prod_{p \in \mathcal{B}} p^{\frac{e(p)}{2}} \right)^2.$$

Hence

$$x = \pm \prod_{p \in \mathcal{B}} p^{e(p)}.$$

So it suffices for us to look for the independent subset of $S_{\mathbb{Z}}$ modulo \mathbb{F}_2 , the finite field of characteristic 2. Hence we define

$$M_2 = M \pmod{2} = (m_{i,j} \pmod{2})_{i \in \{1, \dots, |S_{\mathbb{Z}}|\}, j \in \{1, \dots, \pi(B)\}}.$$

Let $f(a_i + b_i m)$ be the i th row of M_2 . As $|S_{\mathbb{Z}}| > \pi(B)$ we have that there is a linearly dependent subset in M_2 and hence there is a non-trivial solution $T_{\mathbb{Z}} \subset S_{\mathbb{Z}}$ to the linear equation:

$$\sum_{(a,b) \in T_{\mathbb{Z}}} f(a - bm) = 0.$$

With this set we obtain (1):

$$\prod_{(a,b) \in T_{\mathbb{Z}}} (a - bm) \text{ is a square in } \mathbb{Z}.$$

3.2.3 Sieving over $\mathbb{Z}[\alpha]$

To sieve over $\mathbb{Z}[\alpha]$ we will attempt to have a similar construction as the process used for \mathbb{Z} , but to do this we will first have to deal with a few of the obstructions that arise. First recall from theorem 2.36 that for a prime ideal \mathfrak{P} we have that

$$N(\mathfrak{P}) = p^f,$$

where $p \in \mathbb{Z}$ is prime and $f \in \mathbb{N}$. Then p is called the prime lying below \mathfrak{P} and f is called the inertia degree of \mathfrak{P} . Moreover a prime ideal for which $f = 1$ is called a first degree prime. For a first degree prime ideal, \mathfrak{P} , we have that the index $[\mathbb{Z}[\alpha] : \mathfrak{P}] = p$, hence gives rise to the isomorphism

$$\mathbb{Z}[\alpha]/\mathfrak{P} \cong \mathbb{Z}/p\mathbb{Z},$$

hence $\mathbb{Z}[\alpha]/\mathfrak{P}$ is a field. This gives us a direct link between $\mathbb{Z}[\alpha]$ and $\mathbb{Z}/p\mathbb{Z}$:

Theorem 3.5. Let $f(x)$ be a monic irreducible polynomial with coefficients in \mathbb{Z} . Let α be a root of $f(x)$, then there is a bijective correspondence between the set \mathcal{P}_1 of first degree prime ideals and the set

$$\{(p, m) : p \text{ prime}, m \in \mathbb{Z}/p\mathbb{Z}, f(m) \equiv 0 \pmod{p}\}.$$

Proof. Let \mathfrak{p} be a first degree prime ideal of $\mathbb{Z}[\alpha]$. then $[\mathbb{Z}[\alpha] : \mathfrak{p}] = p$ for some prime integer p so that $\mathbb{Z}[\alpha]/\mathfrak{p} \cong \mathbb{Z}/p\mathbb{Z}$. There is a canonical ring epimorphism $\theta : \mathbb{Z}[\alpha] \rightarrow \mathbb{Z}[\alpha]/\mathfrak{p}$ such that $\ker(\theta) = \mathfrak{p}$, hence for $z \in \mathfrak{p} : p \mid \theta(z)$, moreover for $n \in \mathbb{Z}$ such that $p \mid n$ there is a $z \in \mathfrak{p}$ so $\theta(z) = n$. As θ is a homomorphism $\theta(1) = 1$ and so $\forall n \in \mathbb{Z} : \theta(n) \equiv n \pmod{p}$.

Now let $m = \theta(\alpha) \in \mathbb{Z}/p\mathbb{Z}$. If $f(x) = \sum_{i=0}^d a_i x^i$ with $a_d = 1$ and $a_i \in \mathbb{Z}$ then $\theta(f(\alpha)) \equiv 0 \pmod{p}$ as $f(\alpha) = 0$, and so

$$0 \equiv \theta(f(\alpha)) \equiv \theta\left(\sum_{i=1}^d a_i x^i\right) \equiv \sum_{i=1}^d a_i \theta(x)^i \equiv \sum_{i=1}^d a_i m^i \equiv f(m) \pmod{p},$$

hence $f(m) \equiv 0 \pmod{p}$ and \mathfrak{p} determines the unique pair (p, m) .

Conversely, let p be a prime integer and $m \in \mathbb{Z}/p\mathbb{Z}$ with $f(m) \equiv 0 \pmod{p}$. Then there is a natural ring epimorphism that maps polynomials in α to polynomials in r . In particular $\theta(a) \equiv a \pmod{p}$ for all $a \in \mathbb{Z}$ and $\theta(\alpha) \equiv m \pmod{p}$. Let $\mathfrak{p} = \ker(\theta)$ so that \mathfrak{p} is an ideal of $\mathbb{Z}[\alpha]$. Since θ is surjective that means that $\mathbb{Z}[\alpha]/\mathfrak{p} \cong \mathbb{Z}/p\mathbb{Z}$ and so $[\mathbb{Z}[\alpha] : \mathfrak{p}] = p$. This implication is

unique, hence we have the two unique implications,

$$(p, m) \leftrightarrow \mathfrak{p} \leftrightarrow (p, m),$$

proving the theorem. □

Hence finding these first degree prime ideals is equivalent to finding roots \pmod{p} for the minimal polynomial $\text{irr}_K(\alpha)$. Finding roots of polynomials over finite fields has a well documented background [5], for example Berlekamp's algorithm or Cantor-Zassenhaus. We may therefore assume that finding first degree prime ideals is easy⁴ and therefore assume we can find sufficient first degree prime ideals.

Now that we have restricted our prime ideals we may generalize the smoothness test for $\mathbb{Z}[\alpha]$. Buhler et. al suggested the following algorithm in [17]: For a smoothness bound $B' \in \mathbb{Z}$, for now further undefined, $S_{\mathbb{Z}[\alpha]}$ as defined above, and the polynomial f with root $\alpha = t$.

⁴And with easy we mean effective, and in that not adding to our overall complexity, as the size of $\deg(f)$ is far smaller than the size of n .

Algorithm 2 Procedure to populate $S_{\mathbb{Z}[\alpha]}$

Input: Universe U , polynomial f , smoothness bound B' **Output:** Set $S_{\mathbb{Z}[\alpha]} = \{(a, b) \mid a - b\alpha \text{ is } B'\text{-smooth, } \gcd(a, b) = 1\}$ **for** each prime $p \leq B'$ **do****for** each $(a, b) \in U$ such that $b \not\equiv 0 \pmod{p}$ **do** initialize an array populated by $N(a - b\alpha)$ **end for****for** each $r \in \mathbb{Z}/p\mathbb{Z}$ **do** compute the set $R(p) = \{r \mid f(r) \equiv 0 \pmod{p}\}$ **end for****for** each $r \in R(p)$ **do** **if** $a \equiv br \pmod{p}$ **then** retrieve $N(a - b\alpha)$ divide $N(a - b\alpha)$ by the maximal power of p , such that p does not divide the quotient, and replace $N(a - b\alpha)$ by this quotient. **end if****end for****end for****for** each element in the array corresponding to the pair (a, b) **do** **if** the element is ± 1 **then** add the pair (a, b) to $S_{\mathbb{Z}[\alpha]}$ **end if****end for**

Remark. It is clear that if $b \equiv 0 \pmod{p}$ then there are no integers with $(a, b) \in U$ and $N(\langle a - b\alpha \rangle) = N(a - b\alpha) \equiv 0 \pmod{p}$.

What is left to us is to find the square in $\mathbb{Z}[\alpha]$ that we wish to use. However following the

same procedure as the rational sieve would leave us, not with (2), but with

$$N \left(\prod_{(a,b) \in S_{\mathbb{Z}[\alpha]}} (a - b\alpha) \right) \text{ is a square in } \mathbb{Z}.$$

This is clearly a necessary condition for (2), but it is not sufficient. To combat this obstruction we recall the pairs $(p, R(p))$ as defined in algorithm 3.2.3 and define the following:

Proposition 3.6. Let $a, b \in \mathbb{Z}$, $\gcd(a, b) = 1$, and $p \in \mathbb{Z}$ prime. Let $r \in R(p)$. Then the function

$$e_{(p,r)}(a - b\alpha) = \begin{cases} \text{ord}_p(N(a - b\alpha)) & a - br \equiv 0 \pmod{p} \\ 0 & \text{otherwise} \end{cases}.$$

, where for $z \in \mathbb{Z} : \text{ord}_p(z) = f$ if $p^f \parallel z$, is well defined. Moreover

$$N(a - b\alpha) = \prod_{(p,r)} (p^{e_{(p,r)}(a - b\alpha)}).$$

To see how this can be used to produce the square in $\mathbb{Z}[\alpha]$ we need the last theorem of this section

Theorem 3.7. Let S' be a finite set of coprime integer pairs (a, b) fulfilling equation (2). Then for each prime number p and each $r \in R(p)$ we have

$$\sum_{(a,b) \in S_{\mathbb{Z}[\alpha]}} e_{(p,r)}(a - b\alpha) \equiv 0 \pmod{2}$$

Proof. For each prime ideal $\mathfrak{P} \in \mathbb{Z}[\alpha]$ define the group homomorphism $\varphi_{\mathfrak{P}} : \mathbb{Q}[\alpha]^* \rightarrow \mathbb{Z}$ such that

1. $\varphi_{\mathfrak{P}}(\beta) \geq 0$ for all $\beta \in \mathbb{Z}[\alpha], \beta \neq 0$
2. if $\beta \in \mathbb{Z}[\alpha], \beta \neq 0$, then $\varphi_{\mathfrak{P}}(\beta) > 0$ if and only if $\beta \in \mathfrak{P}$

3. for each $\beta \in \mathbb{Q}[\alpha]^*$ one has $\varphi_{\mathfrak{P}}(\beta) = 0$ for all but finitely many \mathfrak{P} , and

$$|N(\beta)| = \prod_{\mathfrak{P}} N(\mathfrak{P})^{\varphi_{\mathfrak{P}}(\beta)}.$$

Then $\varphi_{\mathfrak{P}}$ is a \mathfrak{P} -adic valuation from a multiplicative group of units to an additive group of integers, hence, for $\gcd(a, b) = 1$, if \mathfrak{P} is not a first degree prime then $\varphi_{\mathfrak{P}}(a - b\alpha) = 0$ and if \mathfrak{P} corresponds to the pair (p, m) as defined in theorem 3.6 then $\varphi_{\mathfrak{P}}(a - b\alpha) = e_{p,r}(a - b\alpha)$.

Now let $\mathfrak{P}_i \in \mathbb{Z}[\alpha]$ be a first degree prime ideal and let

$$\prod_{(a,b) \in S_{\mathbb{Z}[\alpha]}} (a - b\alpha) = \xi^2.$$

Then, as $\varphi_{\mathfrak{P}}$ is a homomorphism, we get that

$$\begin{aligned} \sum_{(a,b) \in S_{\mathbb{Z}[\alpha]}} e_{(p,r)}(a - b\alpha) &= \sum_{(a,b) \in S_{\mathbb{Z}[\alpha]}} \varphi_{\mathfrak{P}_i}(a - b\alpha) = \varphi_{\mathfrak{P}_i} \left(\prod_{(a,b) \in S_{\mathbb{Z}[\alpha]}} (a - b\alpha) \right) \\ &= \varphi_{\mathfrak{P}_i}(\xi^2) = 2\varphi_{\mathfrak{P}_i}(\xi) \equiv 0 \pmod{2}. \end{aligned}$$

□

If our assumptions hold that means we can now use this outcome to do a similar process as we did for the rational factor base and find the final set $S_{\mathbb{Z}[\alpha]}$ and use this to explicitly define the set $S = S_{\mathbb{Z}} \cap S_{\mathbb{Z}[\alpha]}$. However these have not been inconsequential and will all be explained in the following section. If we, however, suspend our disbelief for a moment longer we can finish the algorithm.

Assume that we have found $(z^2, \xi^2) \in \mathbb{Z} \times \mathbb{Z}[\alpha]$ such that they fulfil equations (1) and (2). Then all that rests us to do is finding the square roots. First consider the rational case, so z^2

and \mathbb{Z} . Then, by the sieving process, we know the prime factorization of z^2 :

$$z^2 = \prod_{p_i \in \mathbb{Z}} p_i^{e_i},$$

but then it is trivial to find z :

$$z = \prod_{p_i \in \mathbb{Z}} p_i^{\frac{e_i}{2}}.$$

In the algebraic case, i.e. ξ^2 and $\mathbb{Z}[\alpha]$, there is a bigger challenge, as there is no simple way to consider factorize ξ^2 . One naive approach would be to compute the root of $x^2 - \xi^2$, but this is usually not efficiently achievable. We instead take a theoretical approach, with an eye on computational efficiency, and let q be an odd prime. If $f \pmod q$ is irreducible in $\mathbb{F}_q[x]$, then $\mathbb{Z}[\alpha]/q\mathbb{Z}[\alpha] \cong \mathbb{F}_q[x]/(f \pmod q)$. It can be shown that with significant probability there exists an odd q such that $f \pmod q$ is irreducible. To see this consider

Theorem 3.8. Let $f \in \mathbb{Z}[x]$ be an irreducible polynomial of degree d , $d > 1$. Then the density, inside the set of all prime numbers, of the set of prime numbers q for which $f \pmod q$ factors in $\mathbb{F}_q[x]$ into distinct irreducible non-linear factors exists and is at least $\frac{1}{d}$

To prove this we need the following proposition:

Proposition 3.9. Let G be a finite group that acts transitively on a finite set Ω , with $\#\Omega = d > 1$. Then there are at least $\frac{\#G}{d}$ elements of G that act without fixed points on Ω .

Now we prove the theorem.

Proof. Let $\Gamma = \text{Gal}(f/\mathbb{Q})$, viewed as a permutation group of the set $A = \{\alpha_1, \dots, \alpha_d\}$ with roots of f . For each prime number q that does not divide the discriminant of f , there is a Frobenius element $\sigma_q \in \Gamma$ with the property that the degrees of the irreducible factors of $f \pmod q$ are the same as the lengths of the cycles of the permutation σ_q . Hence, we are interested in those q for which σ_q acts without fixed points on A . Then by the Chebotarev Density Theorem, 4.31, every subset $S \subset G$ that is closed under conjugation, the set of prime numbers σ_q belongs to has density $\frac{\#S}{\#G}$. The theorem follows from proposition 3.9. \square

So except for the extremely small probability that we can not choose any f for a specific n this means we can assume that there exists a q so $f \pmod q$ is irreducible in $\mathbb{F}_q[x]$. The following theorem completes the square root finding process.

Theorem 3.10. Let $\mathfrak{q} = q\mathbb{Z}[\alpha]$ be an ideal of $\mathbb{Z}[\alpha]/q\mathbb{Z}[\alpha]$. Then there exists a $\delta \in \mathbb{Z}[\alpha]$ such that

$$\delta^2 \xi^2 \equiv 1 \pmod{\mathfrak{q}}.$$

Proof. As $\mathbb{Z}[\alpha]/q\mathbb{Z}[\alpha] \cong \mathbb{F}_q[x]/f \pmod q$ we know that $|\mathbb{Z}[\alpha]/q\mathbb{Z}[\alpha]| = q^d$, where $d = \deg(f)$.

Now consider

$$\mathfrak{J} = q\mathbb{Z}[\alpha] = \left\{ \sum_{i=1}^{d-1} a_i \alpha^i : q \mid a_i \right\},$$

which is a degree d prime ideal in $\mathbb{Z}[\alpha]$. As $f \pmod q$ is assumed irreducible it follows that $f'(\alpha) \notin \mathfrak{J}$ and for each $(a, b) \in S_{\mathbb{Z}[\alpha]}$ we have that $a - b\alpha \notin \mathfrak{J}$ since $\gcd(a, b) = 1$. Therefore $\xi^2 = f'(\alpha)^2 \prod_{(a,b) \in S_{\mathbb{Z}[\alpha]}} (a - b\alpha) \notin \mathfrak{J}$.

Using Berlekamp's algorithm, [5], we find an element $\delta \pmod{\mathfrak{J}}$ such that $\delta^2 \xi^2 \equiv 1 \pmod{\mathfrak{J}}$, completing the proof. \square

Hence there exists an element δ that is the inverse of a square modulo \mathfrak{q} . Now we can apply Newton–Rhapson iteration ([5],[17]) to find approximations such that

$$\delta_j \equiv \frac{\delta_{j-1}(3 - \delta_{j-1}^2 \xi^2)}{2} \pmod{(q\mathbb{Z}[\alpha])^2}.$$

In a finite number of steps we will find a γ such that

$$\gamma \equiv \delta_j \xi \pmod{(q\mathbb{Z}[\alpha])^2},$$

where $\gamma = \xi$ in $\mathbb{Z}[\alpha]$. To complete the algorithm we compute $\gcd(z - N(\xi), n)$ and $\gcd(z + N(\xi), n)$ to find a factorization as described in section 3.1.

3.3 Dealing with obstructions

So far the technique we have been trying to describe can be captured in the following two bi-implications:

$$\text{Equation (1)} \Leftrightarrow \prod_{(a,b) \in T} (a - bm) \text{ has non-negative even exponents at all primes } p \leq B, \quad (3)$$

$$\text{Equation (2)} \Leftrightarrow \prod_{(a,b) \in T} (a - b\alpha) \text{ has even exponents at all prime ideals } \mathfrak{P} \in \mathbb{Z}[\alpha]. \quad (4)$$

It is clear to see that the first of these bi-implications holds, as we can find the root by simply dividing all exponents by 2. The second of these is not completely clear however. One of the first assumptions we made was that $\mathcal{O}_{\mathbb{Q}(\alpha)} = \mathbb{Z}[\alpha]$, but this is rarely the case. In fact this is something that was believed true until in the 19'th century and even featured in (eventually disproven) proposed proofs of Fermat's Last Theorem [6]. This assumption leads to four obstructions that have to be mitigated. For this let $\omega = \prod_{(a,b) \in T} (a - b\alpha)$.

1. $\mathbb{Z}[\alpha]$ is not necessarily $\mathcal{O}_{\mathbb{Q}(\alpha)}$. At best we know that $\mathbb{Z}[\alpha] \subseteq \mathcal{O}_{\mathbb{Q}(\alpha)}$, hence we can not assume that $\mathbb{Z}[\alpha]$ is a Dedekind domain so $\omega \mathcal{O}_{\mathbb{Q}(\alpha)}$ might not be the square of an ideal in $\mathbb{Z}[\alpha]$.
2. If $\omega \mathcal{O}_{\mathbb{Q}(\alpha)}$ is the square of some ideal \mathfrak{I} , it is not certain that \mathfrak{I} is a principal ideal.
3. If $\omega \mathcal{O}_{\mathbb{Q}(\alpha)}$ is the square of some principal ideal $\gamma \mathcal{O}_{\mathbb{Q}(\alpha)}$, it is not certain that $\omega = \gamma^2$ as ω agrees with γ^2 only up to units of $\mathcal{O}_{\mathbb{Q}(\alpha)}$.
4. And even if $\omega = \gamma^2$, then we are not assured that $\gamma \in \mathbb{Z}[\alpha]$

Clearly these are obstructions that break the number field sieve. To deal with obstruction 4 recall the standard fact that for any $\theta \in \mathcal{O}_{\mathbb{Q}(\alpha)}$ and $f(x) = \text{irr}_{\mathbb{Q}}(\alpha)$ we have that

$$f'(\alpha) \cdot \theta \in \mathbb{Z}[\alpha].$$

So we simply multiply equation (2) by $f'(\alpha)^2$ and equation (1) by $f'(m)^2$. The only condition we must apply is that $\gcd(f'(m), n) = 1$ or else the resulting element given by equation (1) may not be invertible, however this is simply checked and if this is not the case then we have found a factorization of n so we may simply assume that $f'(m)$ and n are coprime.

Now we attempt to subvert obstruction 1, which consequently allows us to negate both obstructions 2 and 3 as well. To do this we will consider the following chain:

$$V \supset V_1 \supset V_2 \supset V_3 = V \cap (\mathbb{Q}(\alpha)^\times)^2,$$

where V is the group generated by $\mathbb{Q}(\alpha)^\times$ with even exponents, i.e. $e_{(p,r)}(v) \equiv 0 \pmod{2}$ for all $v \in V$. This is a group with the following subgroups:

$$V_1 = \{v \in V \mid v\mathcal{O}_{\mathbb{Q}(\alpha)} = \mathfrak{I}^2 \text{ for some } \mathfrak{I} \subset \mathcal{O}_{\mathbb{Q}(\alpha)}\},$$

$$V_2 = \{v \in V_1 \mid v\mathcal{O}_{\mathbb{Q}(\alpha)} = \langle \vartheta \rangle^2 \text{ for some } \langle \vartheta \rangle \subset \mathcal{O}_{\mathbb{Q}(\alpha)}\},$$

$$V_3 = V \cap (\mathbb{Q}(\alpha)^\times)^2.$$

We attempt to bound the index $[V : V_3]$. Considered as a \mathbb{Z} -module $\mathcal{O}_{\mathbb{Q}(\alpha)}$ is free of rank $d = \deg(f)$ and by definition $\mathbb{Z}[\alpha] \subset \mathcal{O}_{\mathbb{Q}(\alpha)}$. It is a well-known identity in algebraic number theory that $\Delta(f) = [\mathcal{O}_{\mathbb{Q}(\alpha)} : \mathbb{Z}[\alpha]]^2 \cdot \Delta$. and as $[\mathbb{Q}(\alpha) : \mathbb{Q}] > 1$ it automatically follows that $\Delta > 0$ hence $[\mathcal{O}_{\mathbb{Q}(\alpha)} : \mathbb{Z}[\alpha]]$ is bounded by $\sqrt{\Delta(f)}$.

Remark. Recall that Δ is the discriminant of the number field as in Definition 2.38

As we have factorization into prime ideals in $\mathcal{O}_{\mathbb{Q}(\alpha)}$ there is a bijection between the prime ideals \mathfrak{Q} of $\mathcal{O}_{\mathbb{Q}(\alpha)}$ coprime to $[\mathcal{O}_{\mathbb{Q}(\alpha)} : \mathbb{Z}[\alpha]]$ and the ideals $\mathcal{I} \subset \mathbb{Z}[\alpha]$ coprime to this index. In fact if we only consider the prime ideals $\mathfrak{P} \subset \mathbb{Z}[\alpha]$ we get an isomorphism of local rings:

$$\mathbb{Z}[\alpha]_{\mathfrak{P}} \rightarrow (\mathcal{O}_{\mathbb{Q}(\alpha)})_{\mathfrak{Q}},$$

$$\mathfrak{Q} \mapsto \mathfrak{P} = \mathfrak{Q} \cap \mathbb{Z}[\alpha].$$

This allows us to generalize the definition of the map $e_{(p,r)}$ of proposition 3.6.

Proposition 3.11. Let $\mathfrak{P} \in \mathbb{Z}[\alpha]$ be a prime ideal of $\mathbb{Z}[\alpha]$. Then

$$e_{\mathfrak{P}}(a - b\alpha) = \sum_{\mathfrak{Q} \supset \mathfrak{P}} f(\mathfrak{Q}/\mathfrak{P}) e_{(q,r)}(a - b\alpha),$$

where \mathfrak{Q} lies over a prime $q \in \mathbb{Z}$.

In the case that \mathfrak{P} does not divide the index $[\mathcal{O}_{\mathbb{Q}(\alpha)} : \mathbb{Z}[\alpha]]$, $e_{\mathfrak{P}}(a - b\alpha) = e_{(q,r)}(a - b\alpha)$ as there will only be one \mathfrak{Q} such that $\mathfrak{P} \subset \mathfrak{Q}$ and $f(\mathfrak{Q}/\mathfrak{P}) = 1$.

Now consider the following map

$$\begin{aligned} V/V_1 &\rightarrow \bigoplus_{\mathfrak{Q} | [\mathcal{O}_{\mathbb{Q}(\alpha)} : \mathbb{Z}[\alpha]]} \mathbb{Z}/2\mathbb{Z}, \\ x &\mapsto (e_{\mathfrak{Q}}(x) \pmod{2})_{\mathfrak{Q}}. \end{aligned}$$

This is an injective homomorphism, hence V/V_1 is a \mathbb{F}_2 -vectorspace of dimension bounded by $|\{\mathfrak{Q} \mid \mathfrak{Q} \mid [\mathcal{O}_{\mathbb{Q}(\alpha)} : \mathbb{Z}[\alpha]]\}|$. As the number of rational primes dividing the index is no more than $\frac{1}{2}|\Delta(f)|$ and for each of these primes there are at most $\deg(f)$ prime ideals $\mathfrak{Q} \subset \mathcal{O}_{\mathbb{Q}(\alpha)}$ that divide it, we obtain

$$\dim_{\mathbb{F}_2}(V/V_1) \leq \frac{d}{2} \log \Delta(f).$$

This inequality is the first step to resolving the first obstruction. Further we will bound the dimensions $\dim_{\mathbb{F}_2}(V_1/V_2)$ and finally $\dim_{\mathbb{F}_2}(V_2/V_3)$ to obtain our final result.

As the class group, $\text{Cl}_{\mathbb{Q}(\alpha)}$, is a finite abelian group and that its order h is, [13], bounded by

$$h < \sqrt{|\Delta(f)|} \frac{d - 1 + \log|\Delta(f)|^{d-1}}{(d-1)!}.$$

Define the map $\kappa : V_1 \rightarrow \text{Cl}_{\mathbb{Q}(\alpha)}$ by $x \mapsto \mathfrak{J}$ where $x\mathcal{O}_{\mathbb{Q}(\alpha)} = \mathfrak{J}^2$. By definition $\ker(\kappa) = V_2$, hence

$$\dim_{\mathbb{F}_2}(V_1/V_2) \leq \frac{\log h}{\log 2}.$$

Not that V_2 consists of elements that are squares in $\mathbb{Q}(\alpha)^\times$ up to units in $\mathcal{O}_{\mathbb{Q}(\alpha)}$, hence by [[13], 8.3]

$$|V_2/V_3| \leq \left| \mathcal{O}_{\mathbb{Q}(\alpha)}^* / (\mathcal{O}_{\mathbb{Q}(\alpha)}^*)^2 \right| \leq d.$$

This leads to the following theorem:

Theorem 3.12. Let V be as above and suppose that $n > d^{2d^2} > 1$. Then the subgroup $V_3 = V \cap (\mathbb{Q}(\alpha)^\times)^2$ of squares in V satisfies

$$\dim_{\mathbb{F}_2}(V/V_3) \leq \log(n)^{\frac{3}{2}}$$

This shows that the general number field sieve algorithm is certain to generate an element in V , however we are not sure yet that this is also a square, hence is in V_3 . To obtain this we introduce the concept of quadratic characters. Let us start by proposing the idea over \mathbb{Z} : Assume $x, y \in \mathbb{Z}$ then x is a quadratic residue modulo a prime p if

$$x \equiv y^2 \pmod{p}.$$

This is easily tested by the Legendre symbol:

$$\left(\frac{x}{p} \right) = x^{\frac{p-1}{2}} \equiv (y^2)^{\frac{p-1}{2}} = y^{p-1} \equiv 1 \pmod{p}.$$

Now assume that x is a square, then it is also a square modulo p for every prime p such that $p \nmid x$. This means that for all $p \in \mathbb{Z}$, p prime:

$$\left(\frac{x}{p} \right) = 1$$

So if there is at least one p for which $\left(\frac{x}{p}\right) = -1$ then x is not a square modulo p and by extension not a square element. This generalizes well to $\mathbb{Q}(\alpha)$.

Let \mathfrak{P} be a first degree prime ideal of $\mathbb{Z}[\alpha]$ lying over p and let $(p, R(p))$ be the set as defined in algorithm 3.2.3. As there exists a ring homomorphism: $\pi : \mathbb{Z}[\alpha] \rightarrow \mathbb{Z}/p\mathbb{Z}$ with $\ker(\pi) = \mathfrak{P}$ by definition this gives rise to a Legendre symbol:

$$\left(\frac{\cdot}{\mathfrak{P}}\right) : \mathbb{Z}[\alpha] \rightarrow \mathbb{Z}/p\mathbb{Z} \rightarrow \{\pm 1, 0\},$$

such that for a non-square $x \in \mathbb{Z}[\alpha]$ we have $\left(\frac{x}{\mathfrak{P}}\right) = -1$ with probability $\frac{1}{2}$. Restricting to Legendre symbols coming from \mathfrak{P} over a prime $p \in \mathbb{Z}$ such that $p > B'$ we avoid the character value 0, and as a consequence of Chebotarev's Density Theorem we have that the Legendre symbols coming from \mathfrak{P} are equidistributed over the space of homomorphisms from V/V_3 to $\{\pm 1\}$, denoted $\text{Hom}(V/V_3, \{\pm 1\})$.

Now let $\chi_{\mathfrak{P}} = \left(\frac{\cdot}{\mathfrak{P}}\right)$ and consider the following lemma.

Lemma 3.13. Let k, r be non-negative integers, and let E be a k -dimensional \mathbb{F}_2 vector space. Then the probability that $k + r$ elements that are uniformly at random drawn from E form a spanning set for E is at least $1 - 2^{-r}$

Then the equidistribution over $\text{Hom}(V/V_3, \{\pm 1\})$ with the obtained bound on the dimension of V/V_3 as a \mathbb{F}_2 vector space from theorem 3.12 makes it overwhelmingly likely⁵ that a set of $B'' = \lfloor \frac{3(\log n)}{\log 2} \rfloor$ quadratic characters span the homomorphism space. If they do, then the converse to the final theorem of this section shows that if $\beta \in \mathbb{Z}[\alpha] \setminus \{0\}$ satisfies $\chi_{\mathfrak{Q}}(\beta) = 1$ for all first degree primes \mathfrak{Q} with $2\beta \notin \mathfrak{Q}$, then β is in fact a square in $\mathbb{Z}[\alpha]$. Furthermore, there can be a finite number of exceptions which is explored greater in [17].

Theorem 3.14. Let S be a finite set of integer pairs (a, b) such that $\gcd(a, b) = 1$ fulfilling that

⁵See [17] for details.

for some $\gamma \in \mathbb{Q}[\alpha]$:

$$\omega = \gamma^2.$$

Moreover let \mathfrak{q} be a first degree prime ideal corresponding to the pair (s, q) that does not divide $\langle a - b\alpha \rangle$ for any pair (a, b) and for which $f'(s) \not\equiv 0 \pmod{q}$. Then we have

$$\prod_{(a,b) \in S} \left(\frac{a - bs}{q} \right) = 1.$$

Proof. Let $\varphi : \mathbb{Z}[\alpha] \rightarrow \mathbb{Z}/q\mathbb{Z}$ be the ring homomorphism given by $\varphi(\alpha) = r \pmod{q}$ and let $\ker(\varphi) = \mathfrak{Q}$. Then \mathfrak{Q} is a first degree prime ideal corresponding to the pair (q, r) . Let $\chi_{\mathfrak{Q}} : \mathbb{Z}\alpha \rightarrow \{\pm 1\}$ and let $\psi_{\mathfrak{Q}} : \mathbb{Z}[\alpha]/\mathfrak{Q} \rightarrow \mathbb{Z}/q\mathbb{Z} \setminus \{0\}$, such that $\chi_{\mathfrak{Q}} = \text{Leg}_{\mathfrak{Q}} \circ \psi_{\mathfrak{Q}}$ where $\text{Leg}_{\mathfrak{Q}} : \mathbb{Z}/q\mathbb{Z} \rightarrow \{\pm 1\}$ is the Legendre symbol. Then

$$\chi_{\mathfrak{Q}}(a - b\alpha) = \left(\frac{a - br}{q} \right).$$

Letting, as discussed,

$$\xi^2 = f'(\alpha)^2 \prod_{(a,b) \in S} (a - b\alpha),$$

for some $\xi \in \mathbb{Z}[\alpha]$. By the hypothesis the factors on the left are not in \mathfrak{Q} , so we have that $\xi \notin \mathfrak{Q}$.

However

$$\chi_{\mathfrak{Q}}(\xi^2) = 1,$$

$$\chi_{\mathfrak{Q}}(f'(\alpha)^2) = 1,$$

so

$$\chi_{\mathfrak{q}} \left(\prod_{(a,b) \in S} \left(\frac{a - br}{q} \right) \right) = 1.$$

This completes the proof. □

So now we can be almost certain that we can find a square in $\mathbb{Z}[\alpha]$ and therefore have mitigated all obstructions. The only question left is: What is the complexity?

3.4 Computational efficiency

Let us record the number field sieve algorithm step-by-step:

Algorithm 3 General Number Field Sieve

Input: $n \in \mathbb{N}$, n odd composite, degree $d > 1$, universe U , smoothness bounds B and B'

Output: A factor of n or \neg for failure.

Choose $m = \lfloor n^{\frac{1}{d}} \rfloor$.

Define $f(x) \in \mathbb{Z}[x]$ as the polynomial with coefficient corresponding to the base- m expansion of n .

for each $(a, b) \in U$ **do**

 Run algorithm 3.2.2 to find $S_{\mathbb{Z}}$

 Run algorithm 3.2.3 to find $S_{\mathbb{Z}[\alpha]}$

 Use the techniques from section 3.3 to find quadratic character base S_Q such that $|S_Q| = B''$.

end for

if $|S_{\mathbb{Z}} \cap S_{\mathbb{Z}[\alpha]}| > |B| + |B'| + 1$ **then**

 Use a reduction algorithm over \mathbb{F}_2 to find an dependent subset $S \subset S_{\mathbb{Z}} \cap S_{\mathbb{Z}[\alpha]}$

 Compute $z^2 = \prod_{(a,b) \in S} (a - bm)$ and use its known prime factorization to find z

 Compute $\zeta^2 = \prod_{(a,b) \in S} (a - b\alpha)$ and use Newton–Rhapson iteration to find ζ

if $z \in \mathbb{Z}$ **and** $\zeta \in \mathbb{Z}[\alpha]$ **then**

return $\gcd(z - \zeta, n)$, $\gcd(z + \zeta, n)$

else find a new dependent subset S . If all such sets produce no result **return** \neg

end if

else return \neg

end if

It is generally considered difficult to do a complexity analysis on the number field sieve as it is stated in all its generality. Already when it was analysed in [17] there was a conjectured

complexity of

$$L_n \left(\frac{1}{3}, \sqrt[3]{\frac{64}{9}} + \mathbf{o}(1) \right),$$

and over the 30 years that followed it showed that the conjectured complexity held up heuristically, however there are some things that we are able to say. Following [13] we are able to carefully choose f and thereby $\deg(f)$ to optimize the smoothness bound $B^* = \max B, B'$ and parameter u for the universe of sieving U . The “basic” cost of the algorithm is computed to be $u^{2+o(1)} + y^{2+o(1)}$.

First we will address step 6. As the binary matrix that is formed will be sparse, i.e. there will be significantly more zeroes than ones, it allows us to use fast algorithms to find dependencies. Two of these extremely fast algorithms are Block–Lanczos ([19],[20]) and Block–Wiedemann [21] which both have similar complexity. The size of these matrices are at most $B^* \times B^*$, and by choosing $\log B^* \approx \log u$ we balance the contributions of U and the matrix reduction.

Looking at the sieving procedure we can see that a pair (a, b) is B^* -smooth if $(a - bm)N(a - b\alpha)$ is B^* -smooth, and thus, using that the size of m is approximately $\sqrt[d]{n}$, we may bound this product by

$$u \sqrt[d]{n} (d + 1) u^d \sqrt[d]{n} \approx n^{\frac{2}{d}} u^{d+1},$$

by choosing d optimally, namely

$$d \approx \left(\frac{3 \log n}{\log \log n} \right)^{\frac{1}{3}}.$$

The probability that a number x , of the right size, is smooth can be approximated by r^{-r} where $r = \frac{\log x}{\log u}$. In order to maximize this probability we choose

$$r = \frac{\log x}{\log u} \approx \frac{d'^2}{d'} + d' + 1,$$

for $d' = \sqrt{\frac{2 \log n}{\log u}}$.

This means we get dependency in the matrix if we have at least $y \approx u^2 r^{-r}$ (a, b) pairs. As we have assumed $\log y \approx \log u$ taking logarithms we get $\log u \approx r \log r$, equivalently $r \approx \frac{\log u}{\log \log u}$. Hence, by taking $\frac{2}{3}$ powers:

$$\log u (\log \log u)^{\frac{1}{3}} \approx 2(\log n)^{\frac{1}{3}}.$$

As we have, for $a, s, t \in \mathbb{R}$, that $s = t(\log t)^a$ goes to $t = (1 + o(1))s(\log s)^{-a}$ as t goes to infinity, we get that

$$\log y \approx \log u \approx 2(\log n)^{\frac{1}{3}} \left(\frac{1}{3} \log \log n \right)^{\frac{2}{3}} = \left(\frac{8}{9} \right)^{\frac{1}{3}} (\log n)^{\frac{1}{3}} (\log \log n)^{\frac{2}{3}}.$$

The final GCD computation, which even with the most simple implementation: Euclid's algorithm, has a complexity at most $O(\log n)$ and as such can be completely disregarded in the complexity analysis.

Hence, with this choice of basic parameters u, y and $d \approx d' \approx \left(\frac{3 \log n}{\log \log n} \right)^{\frac{1}{3}}$ we get the conjectured runtime of $L_n \left(\frac{1}{3}, \sqrt[3]{\frac{64}{9}} + \mathbf{o}(1) \right)$, however this is fully heuristic and very little strong can be proven rigorously for this particular form of the number field sieve even with modern techniques.

4 Preparing for randomness

4.1 Zeta function and the Prime Number Theorem

Before now we have been able to suffice with elementary and algebraic number theory and some Galois theory. To introduce randomness we will be drawing strongly on analytic number theory. It was Riemann who really gave analytic number theory a kickstart by introducing the Riemann zeta function.

Remark. Unless otherwise noted $p, q \in \mathbb{Z}$ are prime numbers.

Definition 4.1. Let $s \in \mathbb{C}$ and $n \in \mathbb{N}$, then the Riemann zeta function is defined as

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}.$$

This converges for all s such that $\operatorname{Re}(s) > 1$, has an analytic continuation to \mathbb{C} except for a simple pole at $s = 1$, and admits a functional equation

$$\pi^{-\frac{s}{2}} \Gamma\left(\frac{s}{2}\right) \zeta(s) = \pi^{-\frac{1-s}{2}} \Gamma\left(\frac{1-s}{2}\right) \zeta(1-s),$$

where $\Gamma(s) = \int_0^{\infty} e^{-t} t^{s-1} dt$. Moreover for $\operatorname{Re}(s) > 1$, $\zeta(s)$ admits an Euler product:

$$\zeta(s) = \prod_p \left(\frac{1}{1 - p^{-s}} \right).$$

Remark. The proofs for all the properties listed in the definition can be found in any undergraduate analytic number theory book. We suggest [8] or [10].

One of the first results of analytic number theory is the prime number theorem.

Theorem 4.2. Let $\pi(x) = \sum_{p \leq x} 1$ be the prime counting function. Then there exists a constant $a > 0$ such that

$$\pi(x) = \operatorname{Li}(x) + \mathbf{O}\left(\frac{x}{\log x} \exp\left(-a\sqrt{\log x}\right)\right) = \frac{x}{\log x} (1 + \mathbf{o}(1)),$$

where

$$\text{Li}(x) = \int_2^x \frac{1}{\log t} dt = \frac{x}{\log x} \left(1 + \mathbf{O}\left(\frac{1}{\log x}\right) \right)$$

is the logarithmic integral.

One of the most influential open questions in mathematics is related to the zeta function: The Riemann Hypothesis.

Proposition 4.3. [Riemann Hypothesis] Let $\zeta(s)$ be the Riemann zeta function. For $\text{Re}(s) > 0$, if $\zeta(s) = 0$ then $\text{Re}(s) = \frac{1}{2}$.

This hypothesis has been so extensively tested that it is often accepted as true in the mathematical community. Doing so there is a whole section of number theory that is considered ‘conditional’, which will become unconditional the moment that the Riemann Hypothesis is proven. There is a very interesting result by Hardy, proven in 1914, which we will record for good measure, [theorem 14.8, [8]].

Theorem 4.4. [Hardy’s Theorem] There exist infinitely many $t \in \mathbb{R}$ such that

$$\zeta\left(\frac{1}{2} + it\right) = 0.$$

However for the Riemann Hypothesis to be true this is not enough, as we want all the zeroes to be on the critical line of $\Re(s) = \frac{1}{2}$. This is still one of the major research areas in analytic number theory, but there are some results. The first such was by Selberg in 1942, who showed that there is some non-zero fraction of zeroes on the critical strip. This was later improved by Levinson who showed that at least 34.7% of the zeroes lie on this line in 1974-1975, and this was improved to 40% by Conrey in 1989. Since then there have been, to my knowledge, no significant improvements leaving this topic wide open for future research.

4.2 Generalizing the Riemann zeta function

4.2.1 Dirichlet characters

The Riemann zeta function is not the only function of its kind. The idea of the Riemann zeta function was generalized in two directions each with their own implications. Dirichlet introduced a generalization through the use of multiplicative character functions and Dirichlet series to define a whole family of meromorphic analytic arithmetic functions. Dedekind on the other hand chose the algebraic route and defined the Dedekind zeta functions, which are defined over number fields. For the former we shall introduce Dirichlet characters, which extends the idea behind quadratic characters.

Definition 4.5. Let G be a group. A character on G is a group homomorphism $\chi : G \rightarrow \mathbb{C}^\times$ and the set of characters of G is denoted \hat{G} .

Generally characters possess a few properties:

Proposition 4.6. Let χ be a G -character, then:

1. As χ is a homomorphism: $\forall g, g' \in G : \chi(gg') = \chi(g)\chi(g')$
2. $\forall G, \forall \chi : \chi(e_G) = 1$
3. $\forall G \exists \chi_0$ s.t. $\forall g \in G : \chi_0(g) = 1$. This is called the trivial or principal character.

These properties reveal that \hat{G} might have a group structure, and this is true for the finite case:

Lemma 4.7. If G is a finite group then so is \hat{G} .

There is one more property to consider for general characters before we can introduce Dirichlet characters and that is orthogonality:

Definition 4.8. Let G be a finite group with set (and as G is finite: group) of characters \hat{G} . Then G is said to have the orthogonality property if the following two equations hold:

$$\sum_{g \in G} \chi(g) = \begin{cases} \#G, & \chi = \chi_0 \\ 0, & \chi \neq \chi_0 \end{cases}$$

$$\sum_{\chi \in \hat{G}} \chi(g) = \begin{cases} \#\hat{G} & , g = e_G \\ 0 & , g \neq e_G \end{cases}$$

Remark. For the sake of clarity: Note that the first sum requires a fixed character and sums over all elements, while the second requires a fixed element and sums over all characters of G .

It is an elementary result from representation theory that the orthogonality property holds for all finite groups, and in particular finite cyclic groups, which is of importance as we shall now see:

Definition 4.9.

1. Let $q \in \mathbb{N}$. Then a Dirichlet character mod q is a character of the form $\chi : (\mathbb{Z}/q\mathbb{Z})^\times \rightarrow \mathbb{C}$.
2. Let χ be a Dirichlet character mod q , we extend χ to an arithmetic function $\chi : \mathbb{Z} \rightarrow \mathbb{C}$ as

$$\chi(n) = \begin{cases} \chi(n \pmod{q}) & , \gcd(n, q) = 1 \\ 0 & , \gcd(n, q) > 1 \end{cases}.$$

Moreover a Dirichlet character mod q has the orthogonality property:

$$\sum_{\substack{n=1 \\ \gcd(n,q)=1}}^d \chi(n) = \begin{cases} \phi(q) & , \chi = \chi_0 \\ 0 & , \chi \neq \chi_0 \end{cases},$$

and

$$\sum_{\chi \bmod q} \chi(n) = \begin{cases} \phi(q) & , n \equiv 1 \pmod{q} \\ 0 & , otherwise \end{cases} .$$

As the Dirichlet characters are defined over $\mathbb{Z}/q\mathbb{Z}$, and we will often consider q a prime, we have a vested interest in how characters of cyclic groups behave. It turns out that these are rather well-behaved:

Theorem 4.10. Assume that G is a finite cyclic group such that $|G| = n$ and $\langle a \rangle = G$. Then:

1. \hat{G} has exactly n elements:

$$\chi_k(a^m) = \zeta_n^{km}, k \in \{1, \dots, n\}.$$

where $\zeta_n = e^{\frac{2\pi i}{n}}$ is the n 'th root of unity.

2. G has the orthogonality property (either as defined in Definition 4.8 or as a Dirichlet character
3. \hat{G} is cyclic with generator χ_1 , hence $G \cong \hat{G}$.

Proof. Let $\chi \in \hat{G}$. Then $\chi(a) = \zeta_n^k$ for some $k \in \{1, \dots, n\}$. Hence

$$\chi(a^m) = \chi(a)^m = \zeta_n^{km},$$

proving part (1). By (1) \hat{G} is cyclic and generated by χ_1 , so $G \cong \hat{G}$ as required. To see that G has orthogonality of characters it is simple to see that χ follows the definition of orthogonality. To illustrate we show it according to Definition 4.8. Note that

$$\sum_{g \in G} \chi_0(g) = |G|$$

is trivial. Hence assume that $k \in \{1, \dots, n-1\}$. Then

$$\sum_{g \in G} \chi(g) = \sum_{m=0}^{n-1} \chi_k(a^m) = \sum_{m=0}^{n-1} \zeta_n^{km} = \frac{1 - \zeta_n^{kn}}{1 - \zeta_n^k} = 0.$$

By the duality of this identity the identity of \hat{G} holds too, hence we have shown (2). \square

If q is not prime we will still be dealing with a special group, as $(\mathbb{Z}/q\mathbb{Z})^\times$ will still be abelian, hence can be written as a direct product of cyclic groups. Therefore we can use theorem 4.10 to say the following:

Lemma 4.11. Let G_1, G_2 be finite cyclic groups and let $G \cong G_1 \times G_2$. Let $\chi_i \in \hat{G}_i$ for $i = 1, 2$ and define $\chi : G \rightarrow \mathbb{C}^\times$ via $\chi(g_1, g_2) = \chi_1(g_1)\chi_2(g_2)$. This is a character. Conversely, if $\chi \in \hat{G}$ then there exists a unique choice of $\chi_1 \in \hat{G}_1$ and $\chi_2 \in \hat{G}_2$ such that $\chi(g) = \chi_1(g_1)\chi_2(g_2)$. Furthermore $\hat{G} \cong \hat{G}_1 \times \hat{G}_2$ and G has orthogonality of characters and if G is finite abelian then $G \cong \hat{G}$.

It is clear that if q is an odd prime, then we can define

$$\chi(n) = \left(\frac{n}{p} \right),$$

where the right hand side is the Legendre symbol. This is a character, and it is easy to show that this is, in fact, a Dirichlet character. It is therefore that Dirichlet characters are a generalization of quadratic characters, which will be extremely useful when we randomize the number field sieve. Another very useful property is that, by definition, Dirichlet characters are (quasi-)periodic.

Definition 4.12. Let χ be a Dirichlet character mod q . We say that $d \in \mathbb{N}$ is a quasi-period if for all $m \equiv n \pmod{d}$ such that $\gcd(mn, q) = 1$: $\chi(m) = \chi(n)$. Moreover the least such d is called the conductor.

It is obvious from the definition that if there are two quasi-periods d_1, d_2 then there is a third: $\gcd(d_1, d_2)$. As q is always a quasi-period we have that the conductor of a Dirichlet character mod q is at most q . We can however say something much stronger about the conductor of q :

Proposition 4.13. Let χ be a Dirichlet character \pmod{q} with conductor d . Then $d \mid q$.

Proof. Let $g = \gcd(d, q)$. Suppose $m \equiv n \pmod{g}$ and $\gcd(mn, q) = 1$. By Euclid's algorithm there exists $x, y \in \mathbb{Z}$ such that $m - n = dx + qy$, thus

$$\chi(m) = \chi(m - qy) = \chi(dx + n) = \chi(n).$$

Hence g is a quasiperiod of χ . As d is the conductor that means $d \mid g$ and since $g = \gcd(d, q)$ we have that $d \mid q$. □

As per usual we have now found a way to consider characters “prime”:

Definition 4.14. A Dirichlet character modulo q is called primitive if it has conductor q .

Remark. Convention dictates that the trivial character χ_0 is not primitive.

This idea of primitive characters is especially useful when you are working with induced characters. Taking inspiration from the notation as it is used in representation theory consider $d \mid q$ and consider χ^* , a character \pmod{d} , and define

$$\chi(n) = \begin{cases} \chi^*(n) & , \gcd(n, q) = 1 \\ 0 & , otherwise \end{cases}.$$

Then $\chi(n)$ is a multiplicative character and has period q , so it is a Dirichlet character \pmod{q} . This is called the Dirichlet character \pmod{q} induced by a Dirichlet character \pmod{d} and we will denote it $\text{Ind}_d^q(\chi)$. We conclude with the following theorem about induced Dirichlet characters:

Theorem 4.15. Let χ be a Dirichlet character \pmod{q} with conductor d . Then there exists a unique character $\chi^* \pmod{d}$ such that $\chi(n) = \text{Ind}_d^q(\chi^*(n))$.

4.2.2 Dirichlet \mathcal{L} -functions

The definition of Dirichlet characters leads directly to the definition of Dirichlet \mathcal{L} -functions:

Definition 4.16. Let $q \in \mathbb{N}$ and let χ be a Dirichlet character modulo q . Then the Dirichlet \mathcal{L} -function associated to χ , denoted $\mathcal{L}(s, \chi)$, is defined as

$$\mathcal{L}(s, \chi) = \sum_{n=1}^{\infty} \frac{\chi(n)}{n^s},$$

for $\Re(s) > 1$.

As χ is a completely multiplicative character function we can immediately conclude that $\mathcal{L}(s, \chi)$ is absolutely convergent for $\Re(s) > 1$ and admits an Euler product in this region:

$$\mathcal{L}(s, \chi) = \prod_p \left(1 - \frac{\chi(p)}{p^s} \right)^{-1},$$

where p is prime. The behaviour of the trivial character and the non-trivial characters is quite different. In the case of the trivial character we can rewrite the Euler product:

$$\mathcal{L}(s, \chi_0) = \prod_p (1 - p^{-s}) \zeta(s).$$

This means $\mathcal{L}(s, \chi_0)$ behaves like $\zeta(s)$ and allows for a meromorphic continuation with a single simple pole at $s = 1$. When χ is not trivial the behaviour changes radically.

Theorem 4.17. Let $q \in \mathbb{N}$ and let χ be a Dirichlet character modulo q such that $\chi \neq \chi_0$. Then $\mathcal{L}(s, \chi)$ converges and is holomorphic for $\Re(s) > 0$.

Remark. $\Re(s) = 0$ is the abscissa of convergence for $\mathcal{L}(s, \chi)$, but despite the convergence we can not extend the definition of the Euler product of $\mathcal{L}(s, \chi)$ to $\Re(s) < 1$.

Knowing that we have convergence for $\Re(s) = 1$ in this case one might be interested in what these values actually are. This has been a problem for analytic number theorists for quite some time, but we can say the following:

Theorem 4.18. Let $Q \in \mathbb{N}$ and let χ be a Dirichlet character mod q . Then

$$\prod_{\chi \neq \chi_0} \mathcal{L}(1, \chi) \neq 0.$$

This is absolutely not trivial and a particularly concise proof is presented in [[8], theorem 4.9]. The first of two jewels in this is the following theorem by Dirichlet.

Theorem 4.19. Let $q \in \mathbb{N}$ and let $a \in \mathbb{Z}$ such that $\gcd(a, q) = 1$. Then there are infinitely many primes $p \equiv a \pmod{q}$

Remark. Despite not denoting the proof explicitly one the important step is that we show that

$$\sum_{p \equiv a \pmod{q}} \frac{1}{p}$$

diverges. The theorem follows immediately.

It might not surprise that with the generalization of the Riemann zeta function there follows a generalization of the Riemann Hypothesis:

Proposition 4.20. [Generalized Riemann Hypothesis] Let $\mathcal{L}(s, \chi)$ be a Dirichlet \mathcal{L} -function with χ a primitive character mod $q, q > 1$. Then for $\operatorname{Re}(s) \in (0, 1)$, if $\mathcal{L}(s, \chi) = 0$ then $\operatorname{Re}(s) = \frac{1}{2}$.

We will come back to Dirichlet \mathcal{L} -functions when we adapt the prime number theorem to arithmetic progressions in Section 4.3.

4.2.3 Dedekind zeta functions and Siegel zeroes

Now we consider the other generalization of the zeta function, though it can in this case very accurately be called an extension.

Definition 4.21. Let K be an algebraic number field and let $\mathfrak{J} \subset K$ be an ideal. Then the

Dedekind zeta function over K is defined as

$$\zeta_K(s) = \sum_{\mathfrak{J} \subseteq \mathcal{O}_K} \frac{1}{(N_K(\mathfrak{J}))^s}.$$

The Dedekind zeta function is for $\Re(s) > 1$ and even admits an Euler product for that region. To see this let $\mathfrak{P} \subseteq K$ be a prime ideal, then:

$$\zeta_K(s) = \prod_{\mathfrak{P} \subseteq \mathcal{O}_K} (1 - N_K(\mathfrak{P})^{-s}).$$

It was proven by Hecke that the Dedekind zeta function admits an analytic continuation to the whole complex plane and also admits a functional equation, so it is clear that in many ways the Dedekind zeta function behaves like the Riemann zeta function. This correlation is so close that it culminates in the following theorem:

Theorem 4.22. Landau, 1903 - Prime Ideal Theorem

Let K be an algebraic number field and let $\mathfrak{P} \subseteq \mathcal{O}_K$ be a prime ideal and define $\pi_K(x) = |\{\mathfrak{P} \in \mathcal{O}_K \mid N(\mathfrak{P}) \leq x\}|$. Then for $x \rightarrow \infty$,

$$\pi_K(x) \asymp \frac{x}{\log x}.$$

For our purposes it is sufficient to realize that the Dedekind zeta function over K behaves a lot like the Riemann zeta function, and that the same language can be used to talk about the Dedekind zeta function as what we use for the Riemann zeta function. In fact, there is also a Riemann Hypothesis for the Dedekind zeta functions:

Proposition 4.23. [Extended Riemann Hypothesis] Let K be a number field. For $\Re(s) \in (0, 1)$, if $\zeta_K(s) = 0$, then $\Re(s) = \frac{1}{2}$.

Remark. As the statements are equivalent for the Dirichlet \mathcal{L} -functions and the Dedekind zeta functions it is not uncommon to see the Extended Riemann Hypothesis be called the Generalized

Riemann Hypothesis for number fields. We will in this paper refer to both as the Generalized Riemann Hypothesis, GRH, and let context infer which version we are referring to.

4.2.4 Siegel zeroes

One of the main problems that we encounter when we consider both Dirichlet \mathcal{L} -functions and Dedekind zeta functions is that we can not generalize the zero free region. Especially, for the Riemann zeta function we were able to extend the zero-free region just past the $\Re(s) = 1$ line, and this may not be true for Dedekind zeta functions and Dirichlet \mathcal{L} -functions. In fact:

Theorem 4.24. There is a constant $c > 0$ such that if $\mathcal{L}(\sigma + it, \chi) = 0$ for some primitive complex Dirichlet character $\chi \pmod{q}$ then

$$\sigma < 1 - \frac{c}{\log q(|t| + 2)}.$$

If χ is a real primitive character then this holds for all zeroes of $\mathcal{L}(s, \chi)$ with at most one exception. The exceptional zero, if it exists is real and simple.

Later we will see a result by Landau, [5.29](#), which develops on this idea. It is clear to see that this theorem is formulated for Dirichlet \mathcal{L} -functions, but a similar result exists for Dedekind zeta function. We therefore introduce the following concept:

Definition 4.25. Let $\mathcal{L}(s, \chi)$ be a Dirichlet \mathcal{L} -function, then the hypothetical value $1 - \nu$ for $\nu \in (0, \frac{1}{2})$, such that

$$\mathcal{L}(1 - \nu, \chi) = 0,$$

for a Dirichlet character modulo q , χ , is called a Siegel-zero of $\mathcal{L}(s, \chi)$.

Equivalently, for $\zeta_K(s)$, if there exists a $\nu \in (0, \frac{1}{2})$ such that

$$\zeta_K(1 - \nu) = 0,$$

for a number field K , is called a Siegel-zero of $\zeta_K(s)$.

The culminating theorem from Siegel on ‘his’ zeroes was

Theorem 4.26. Siegel’s theorem For any $\epsilon > 0$ there exists a positive number $c(\epsilon)$ such that, if χ is a real primitive character \pmod{q} , then

$$L(1, \chi) > \frac{c(\epsilon)}{q^{-\epsilon}}.$$

Even in the definition we elude to these zeroes as hypothetical, as the assumption of the (Generalized) Riemann Hypothesis has as an immediate consequence that such zeroes do not exist. For the sake of this paper [14] does not assume the Riemann Hypothesis in any way, hence we must always concern ourselves with possible Siegel zeroes.

4.3 Arithmetic progressions

Now we turn to the next topic, which is a main factor in the Randomized Number Field Sieve: Smooth numbers on arithmetic progressions. To be able to do this we first need to know a little bit about prime numbers on arithmetic progressions and their distribution. First we note that it would be beneficial for us to have a prime counting function for arithmetic progressions. This turns out to be a simple modification:

Definition 4.27. For $\gcd(a, q) = 1$ the prime counting function for arithmetic progressions is defined as

$$\pi_{(a,q)}(x) = \pi(x, a, q) = \sum_{\substack{p \leq x \\ p \equiv a \pmod{q}}} 1.$$

The first question that comes to mind is one that was answered by Dirichlet: Are there infinitely many primes of that form? This is the case, as we stated in theorem 4.19:

Theorem 4.28. Let $x \in \mathbb{R}$ and $a, q \in \mathbb{Z}$ such that $\gcd(a, q) = 1$. Then there are infinitely many primes of the form $p \equiv a \pmod{q}$.

This immediately shows that just like for $\pi(x)$ we have that $\pi(x, a, q)$ diverges. Therefore, we infer, there must be something interesting to say about the asymptotic behaviour of the prime counting function. Not only is there something interesting to say, there is a whole lot to say as well:

Theorem 4.29. Let $x \in \mathbb{R}$, $a, q \in \mathbb{Z}$ such that $\gcd(a, q) = 1$, and let $\phi(x)$ be Euler's totient function. Then

$$\pi(x, a, q) \sim \frac{x}{\log x} \frac{1}{\phi(q)}.$$

This immediately shows that the density is independent of a as long as $\gcd(a, q) = 1$.

4.3.1 Extensions and Refinements

Before we start to talk about smooth numbers we want to consider if anything better than Dirichlet's function can be said, and we can if we allow for an ineffective bound. For this we define

$$\psi(x, a, q) = \sum_{\substack{n \leq x \\ n \equiv a \pmod{q}}} \Lambda(n),$$

where $\Lambda(n)$ is the von Mangoldt function:

$$\Lambda(n) = \begin{cases} \log p & \exists k \in \mathbb{N} : n = p^k \\ 0 & \text{otherwise} \end{cases}.$$

Theorem 4.30. Siegel-Walfisz theorem Let $\psi(x, a, q)$ and $\Lambda(n)$ be as defined above. Given $q \in \mathbb{N}$ there exists a positive constant $c(q)$ such that

$$\psi(x, a, q) = \frac{x}{\phi(q)} + \mathbf{O}\left(x \exp\left(-c(q)\sqrt{\log x}\right)\right).$$

The reason this is ineffective is because it is based on Siegel's theorem, [4.26](#), and this theorem gives no way to compute $c(n)$. Another way to go about this is looking at average cases. This probabilistic approach will be explored later and will lead to effective constant and a far stronger

bound.

It may not come as a surprise that the idea of prime numbers along arithmetic progressions can be expanded to number fields, as we have done with the Prime Ideal Theorem. For this we first consider what an ‘arithmetic progression’ is on a number field. This does not come intuitively, but the answer lies in Galois Theory. Let L/K be a Galois extension with Galois group $G = \text{Gal}(L/K)$. As L/K is a Galois extension the Frobenius element, $\text{Frob}_{\mathfrak{P}}$, defines a conjugacy class

$$C = \left\{ \text{Frob}_{\mathfrak{Q}} \mid \mathfrak{Q} \subset L \text{ s.t. } \mathfrak{Q} \text{ is a prime ideal and } \mathfrak{Q} \mid \mathfrak{P} \right\}.$$

It can be shown that this abides the rules of modular arithmetic and therefore can be used as an extension of the idea of arithmetic progressions. Using this we get the following result:

Theorem 4.31. Chebotarev Density Theorem Let L/K be Galois and let $\mathfrak{P} \subseteq K$ be a prime ideal. Moreover let $C \subseteq G$ be the conjugacy class defined above. Then

$$\{\mathfrak{P} \mid \mathfrak{P} \nmid \Delta_{L/K}, \text{Frob}_{\mathfrak{P}} \in C\}$$

has density $\frac{|C|}{|G|}$.

This has long been the best result available, but over the years it has been strengthened both under the GRH assumption and without. Later we will see a version which will be of particular interest to us as it is both unconditional and has effectively defined constants.

4.3.2 Smooth numbers on Arithmetic Progressions

Lastly, before we move onto a different topic, we discuss how smooth numbers behave on arithmetic progressions. First we need some definitions:

Definition 4.32. For $x, y, a, q \in \mathbb{N}$ and a Dirichlet character χ , we define:

$$\Psi(x, y) = |\{z \in \mathbb{N} \mid z < x, z \text{ is } y\text{-smooth}\}|,$$

$$\Psi_r(x, y) = |\{z \in \mathbb{N} \mid z < x, z \text{ is } y\text{-smooth, } \gcd(z, r) = 1\}|,$$

$$\Psi(x, y, a, r) = |\{z \in \mathbb{N} \mid z < x, z \text{ is } y\text{-smooth, } z \equiv a \pmod{r}\}|,$$

$$\Psi(x, y, \chi) = \sum_{z < x} \mathbf{1}_{\{z' \mid z' \text{ is } y\text{-smooth}\}}(z) \chi(z), \text{ and}$$

$$\rho(x, y) = \Psi(x, y)x^{-1}.$$

For many the ‘Holy Grail’ in this work is showing the following, [29]:

Conjecture 4.33. Let A be a given positive real number. Let y and q be large with $q \leq y^A$.

Then as $\frac{\log x}{\log y} \rightarrow \infty$ we have

$$\Psi(x, y, a, q) \sim \frac{1}{\phi(q)} \Psi_q(x, y).$$

Soundararajan proved in [29] that the above holds for $A \geq 4\sqrt{e} - \epsilon$ given a certain bound on y . The latter restriction was later removed by Harper, [30]. We will see more of Harper’s work later as we use these results.

We now introduce some bounds based on Hildebrand and Tenenbaum’s work, which will be presented proof-less here. For those especially interested in this we suggest [25],[26],[22]. We will avidly be working with these sets later on, but to do so we will first need a few results starting with [[14], fact 3.11]:

Proposition 4.34. Let $\epsilon > 0$ be arbitrary and let $3 \leq u \leq (1 - \epsilon) \frac{\log x}{\log \log x}$. Then:

$$\Psi\left(x, x^{\frac{1}{u}}\right) = x \exp\left(-u \left(\log u + \log \log u - 1 + \frac{\log \log u - 1}{\log u} + \mathbf{O}_\epsilon\left(\frac{\log \log^2 u}{\log^2 u}\right)\right)\right).$$

A very rough result that we use is a direct result of this fact

Corollary 4.35. Fix $0 < a < b \leq 1$. Then uniformly in $c, d > 0$:

$$\rho(L_x(b, d), L_x(a, c)) = L_x \left(b - a, \frac{d(b-a)}{c} \right)^{-1+\mathbf{o}(1)}.$$

Proof. Let

$$u = \frac{\log L_x(b, d)}{\log L_x(a, c)} = \frac{d}{c} \left(\frac{\log x}{\log \log x} \right)^{b-a}.$$

Then $u \rightarrow \infty$ and $u = \log \left(\frac{\log x}{\log \log x} \right)$. Hence

$$\begin{aligned} \rho(L_x(b, d), L_x(a, c)) &= \exp(-(1 + \mathbf{o}(1))u \log u) \\ &= \exp \left(-(1 + \mathbf{o}(1)) \frac{d(b-a)}{c} \log^{b-a}(x) (\log \log x)^{b-a} \right) \\ &= L_x \left(b - a, \frac{d(b-a)}{c} \right)^{-1+\mathbf{o}(1)}. \end{aligned}$$

□

If we are substantially more careful however there are tighter results, whose proofs go beyond the scope of this paper, by Hildebrand and Tenenbaum. These results allow short intervals to be sieved for smooth numbers effectively:

Theorem 4.36. Fix any $\epsilon > 0$. For any $x \geq 3$, $\log x \geq \log y \geq (\log \log x)^{\frac{5}{3}+\epsilon}$, $z \leq y^{\frac{5}{12}}$, the following estimate holds uniformly:

$$\Psi(x(1+z^{-1}), y) - \Psi(x, y) = \frac{\Psi(x, y)}{z} \left(1 + \mathbf{O} \left(\frac{\log(u+1)}{\log y} \right) \right).$$

Theorem 4.37. For any x, y , we set $u = \frac{\log x}{\log y}$. There exists a saddle-point, $\alpha = \alpha(x, y)$, such that for any $1 \leq c \leq y$:

$$\Psi(cx, y) = \Psi(x, y) c^{\alpha(x, y)} \left(1 + \mathbf{O} \left(\frac{1}{u} + \frac{\log y}{y} \right) \right), \text{ with}$$

$$\alpha(x, y) = \frac{\log\left(\frac{y}{\log x} + 1\right)}{\log y} \left(1 + \mathbf{O}_c\left(\frac{\log \log(y+1)}{\log y}\right)\right).$$

Theorem 4.38. Let $c > 0$ be a constant. Let $n \in N$ such that $\omega(n)$ is the number of prime factors of n (without multiplicity). Let n be a y -smooth number with $2 \leq y \leq x$ such that $\omega(n) \leq y^{c(\log(1+u))^{-1}}$. Then:

$$\Psi_n(x, y) = \frac{\phi(n)}{n} \Psi(x, y) \left(1 + \mathbf{O}\left(\frac{\log(1+u) \log(1+\omega(n))}{\log y}\right)\right).$$

This final theorem induces the following corollary which will prove crucial to us:

Corollary 4.39. Take $c > 0$ an arbitrary constant, and retain ω as in 4.38. Let $2 \leq y \leq x$ and with $\omega(n) \leq y^{c(\log(1+u))^{-1}}$. Then:

$$\Psi_n(x, y) = \frac{\phi(n)}{n} \Psi(x, y) \left(1 + \mathbf{O}_c\left(\frac{\log(1+u) \log(1+\omega(n))}{\log(y)}\right)\right) \left(1 + \mathbf{O}\left(\frac{\omega(n)}{y}\right)\right).$$

Proof. Let $n = sr$ for s y -smooth and r with no prime factor less than y . Then $\Psi_s(x, y) = \Psi_r(x, y)$, $\phi(n) = \phi(r)\phi(s)$ and, for p prime, $\phi(r)r^{-1} = \prod_{p|r} (1 - p^{-1}) = 1 + \mathbf{O}(\omega(n)y^{-1})$, which implies the given bound. \square

4.4 Probability measures and moments

The Randomized Number Field Sieve is a probabilistic algorithm and makes extensive use of the discrete uniform distribution. Define for $x \in [a, b]$ and for $y \in S$, where S is a finite set, the discrete uniform distribution as

$$\mathbb{P}(x) = \frac{1}{b - a + 1}, \text{ or}$$

$$\mathbb{P}(y) = \frac{1}{|S|}$$

This is in many ways one of the simplest distributions to work with and overall we will only need a limited amount of probability theory. There is however a need to understand how a distribution can be used to define a measure. For this recall the following definition:

Definition 4.40. A set function $\mu : \mathbb{F} \rightarrow \mathbb{R}$, for a field \mathbb{F} , is a probability measure if it satisfies these conditions:

1. $0 \leq \mu(A) \leq 1$ for $A \subseteq \mathbb{F}$.
2. $\mu(\emptyset) = 0$, $\mu(\mathbb{F}) = 1$.
3. $\mu(A_1, A_2, \dots, A_n)$, for a disjoint sequence of \mathbb{F} -sets, such that $\bigcup_{i=1}^{\infty} A_i \in \mathbb{F}$, then

$$\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu(A_i).$$

Remark. Note that if μ is a probability measure then the support of μ is any set $A \subset \mathbb{F}$ for which $\mu(A) = 1$. It is clear that this always exists by the second condition.

We can confirm that for any finite set S we can define a uniform measure. To see this let μ be the uniform distribution of some finite set $S = \{s_1, \dots, s_n\}$ such that $|S| = n$. Then for any $V \subset S$ with order $|V|$ we have that

$$\mu(V) = \frac{|V|}{|S|}.$$

Especially we have that

$$\mu(\emptyset) = 0, \mu(S) = 1,$$

and

$$0 \leq \mu(V) \leq 1.$$

Lastly for any collection of disjoint subsets $V_1, \dots, V_m \subset S$:

$$\mu\left(\bigcup_{i=1}^{\infty} V_i\right) = \mu\left(\bigcup_{i=1}^m V_i\right) = \frac{\sum_{i=1}^m |V_i|}{|S|} = \sum_{i=1}^m \frac{|V_i|}{|S|} = \sum_{i=1}^m \mu(V_i).$$

Similarly we can confirm that for any continuous interval $[a, b]$ we can define a uniform measure.

Definition 4.41. For any finite set S , we denote the uniform measure over S by

$$\mathcal{U}(S).$$

One of the things we are going to use our uniform measure on are the zero-centered half-open integer intervals of length L , which we will denote as follows:

$$\mathbb{I}(L) = \left[-\frac{1}{2}L, \frac{1}{2}L\right) \cap \mathbb{Z}.$$

We note here, for good measure, that this is in fact a finite set. In the section above we already discussed the Dirichlet convolution of arithmetic functions, but there is also a convolution of measures, which is normally defined as an integral, but we are interested in the discrete uniform measures as defined above. This allows us to restate the convolution of measures as follows:

Definition 4.42. For any two measures μ, ν over an additive group G we define the convolution of measures as

$$(\mu \star \nu)(x) = \sum_{y \in G} \mu(y)\nu(x - y).$$

Remark. Note that it is convention to denote the convolution of measures by $*$, but to be distinct we will denote the convolution of measures by \star and the Dirichlet convolution by $*$.

One of the big concepts in the RNFS is the avoidance of the General Riemann Hypothesis, and for this we need to consider the moments of a probability distribution. We will recognize the first two moments, mean and variance, of the discrete uniform distribution:

Definition 4.43. Let U be a discrete uniform distribution with support S , then the first moment, or mean, is defined as

$$\mathbb{E}(U) = \sum_{s \in S} \mathbb{P}_U(s).$$

Remark. By convention we denote subscripted to the probability, \mathbb{P} , the fixed variables. In this case we fix the distribution, but this is usually dropped when the distribution is clear from context.

Definition 4.44. Let U be a discrete uniform distribution with support S , then the second moment, or variance, of $x \in S$ is defined as

$$\text{Var}(x) = \mathbb{E}(x - \mathbb{E}(x))^2.$$

The idea of the RNFS is to remove the dependence of the analysis of number field sieve-type algorithms on the second moment, for which certain bounds on the complexity exist. To do this we consider a probabilistic technique which Lee and Venkatesan, in [14], have dubbed *stochastic deepening*. The core idea is as follows:

Lemma 4.45. Let x be a random variable with $\mathbb{E}(x) = \mu$. Given there exists a $K \geq 1$ such that $0 \leq x \leq K\mu$ uniformly, then there exists $i \in \{0, \dots, \lceil \log_2 K \rceil\}$ such that:

$$\mathbb{P}\left(x \geq \frac{2^i \mu}{1 + \lceil \log_2 K \rceil}\right) \geq \frac{1}{2^{i+1}}.$$

This lemma states that for non-negative random variables which are not too erratic, there is a substantial set where the value is large and whose contribution to the mean is large. This allows us to provide a search algorithm whose run times are shown to be near optimal without establishing variance bounds.

The explicit statement of such a search algorithm is not too important for us, as we are analysing the theoretical complexity, so we will not define such an algorithm explicitly. It will however become important in the proof of one of the key theorems, 5.2.

5 The Randomized Number Field Sieve

5.1 Introducing the algorithm

Recalling the GNFS algorithm, algorithm 3, we will focus in this section on giving a description of the differences between the sieve from chapter 3 and a randomized version with provable complexity. It is well known that the Number Field Sieve's runtime is dominated by two problems:

1. Finding a set S of sufficient (a, b) -pairs such that $a - bm$ is B -smooth and $a - b\alpha$ is B' -smooth.
2. Collapsing these (a, b) -pairs to find a subset $T \subseteq S$ such that a congruence of squares modulo the to-factor integer n arises.

In this version of the Number Field Sieve these two steps are randomized cleverly to produce a provable complexity. To do this there is one more change that needs to be made, and that is the choice of polynomial. It will be shown that by randomizing the choice of polynomial a rigorous bound can be obtained for the search. To start we consider the constants $\delta, \kappa, \sigma, \beta, \beta'$ under the conditions that

$$\kappa > \delta^{-1}, 2\sigma > \max(\beta, \beta') + \frac{\delta^{-1}}{3\beta}(1 + \mathbf{o}(1)) + \frac{\sigma\delta + \kappa}{3\beta'}(1 + \mathbf{o}(1)), \quad (5)$$

$$\delta^{-1} < \frac{\sigma\delta + \kappa}{2}, \quad (6)$$

and fix the smoothness bounds

$$B = L_n\left(\frac{1}{3}, \beta\right), B' = L_n\left(\frac{1}{3}, \beta'\right). \quad (7)$$

The culminating theorem of the randomized number field sieve is as follows:

Theorem 5.1. Let $\delta, \kappa, \sigma, \beta, \beta'$ be constants as defined in (1) and (2). Then for any n , the randomized number field sieve runs in expected time

$$L_n\left(\frac{1}{3}, \sqrt[3]{\frac{64}{9}} + \mathfrak{o}(1)\right),$$

and produces a pair x, y with $x^2 \equiv y^2 \pmod{n}$.

It is important to stress that this is a NFS-style algorithm, so we can not be sure that the congruence found does not give a trivial factorization. The algorithm is started in the same way as the GNFS by finding a polynomial, however in this randomized version we will insist that the polynomial is a bivariate homogeneous monic irreducible polynomial of bounded degree, that is:

$$f(x, y) \in \mathbb{Z}[x, y] : f(x, y) = \hat{f}_{m,n}(x, y) + R(x, y),$$

where $\hat{f}_{m,n}(x, y)$ is the polynomial given by n and m , such that $\hat{f}_{m,n}(m, 1)$ is the base- m expansion of n and

$$R(x, y) = \sum_{i=0}^{d-1} c_i (x - ym)x^{d-i-1}y^i,$$

such that every c_i is determined uniformly at random and $\deg(f) = d = \delta \sqrt[3]{\frac{\log n}{\log \log n}}$, d odd. Moreover we bound the choice of m by $m^d \leq n < 2m^d$.

Remark. The harshness of these bounds are all necessary to show a provable complexity, so we will often refer back to this introduction as “the defined parameters” without restating the exact parameters again.

Remark. We will, like [14], freely switch between considering the polynomial $f(x, y)$ as a bivariate polynomial and it’s monovariate equivalent $f(x) = f(x, 1)$.

The changes made to the polynomial selection method and the consequences from using random variables lead us to conclude the following theorem:

Theorem 5.2. Take $\delta, \kappa, \sigma, \beta$, and β' with the defined parameters. For any n , the randomized number field sieve can almost surely find an irreducible polynomial f of degree d and height at

most $L_n(\frac{2}{3})$, with α a root of f , $n \mid f(m)$, and

$$L_n(\frac{1}{3}, \max(\beta, \beta') + \mathbf{o}(1))$$

distinct pairs $a < |b| \leq L_n(\frac{1}{3}, \sigma)$ such that $(a - bm)$ is B -smooth and $a - b\alpha$ is B' smooth, in expected time at most $L_n(\frac{1}{3}, \lambda)$ for any

$$\lambda > \max(\beta, \beta') + \frac{\delta^{-1}}{3\beta}(1 + \mathbf{o}(1)) + \frac{\sigma\delta + \kappa}{3\beta'}(1 + \mathbf{o}(1)).$$

In particular, the probability that the randomized number field sieve fails to produce such a set is bounded above by $L_n(\frac{2}{3}, \kappa - \delta^{-1})^{-1+\mathbf{o}(1)}$.

The key purpose for the randomization of the polynomial selection process has to do with the sieving process over the number field, as it will cause the norm of a pair, $N(a - b\alpha)$, to become a random variable in \mathbb{Z} . This allows the consideration of smoothness over \mathbb{Z} and $\mathbb{Z}[\alpha]$ to be completely independent. We will explain this idea in more detail when we look at the sieving process over the number field.

This also leads us to question what the quadratic characters will then look like, and this is the second part where the RNFS differs significantly from the GNFS. Where in the GNFS we defined the quadratic characters as the Legendre symbols $(\frac{x}{p})$ modulo a prime p , we now have to account for our randomization and instead will have to choose maps from $\mathbb{Z}[\alpha]$ into \mathbb{F}_{p^k} stochastically and close to uniformly across all $k \log(p) < L_n(\frac{1}{3})$. This exponentially increases the sizes of the fields, but we will see they are necessary for the unconditional equidistribution of the characters. Once the sieving process is finished we are, with the exception of a set of bad f , guaranteed a reduction to a congruence of squares:

Theorem 5.3. Let B, B' be with the defined parameters. Let f be irreducible of degree d and height at most $L_n(\frac{2}{3}, \kappa)$, and let α be a root of f . Then for all but a $L_n(\frac{2}{3}, \frac{\kappa - \delta^{-1}}{2}(1 + \mathbf{o}(1)))^{-1}$

fraction of the set of f , if we are given

$$L_n\left(\frac{1}{3}, \max(\beta, \beta')\right)\Omega(\log \log n)$$

pairs $a < b \leq L_n(\frac{1}{3})$ such that $a - mb$ is B -smooth and $a - b\alpha$ is B' -smooth, we can find a congruence of squares modulo n in expected time at most

$$L_n\left(\frac{1}{3}, 2 \max\left(\frac{2\delta}{3}, \beta, \beta'\right)\right)^{1+\mathbf{o}(1)}.$$

Let's dive into the details.

5.2 Randomizing the polynomial

As we mentioned the algorithm is a GNFS based algorithm with two major differences: How the polynomial is chosen and how the quadratic characters are chosen. To start we will look at how the polynomial is chosen, and consequences related to that, before proving theorem 5.2. To do this let $\beta, \beta', \kappa, \sigma, \delta, B$, and B' be as defined in (5), (6), and (7).

To begin we make a few restrictions:

Definition 5.4. Let \mathcal{X} be the set of tuples (f, m, n, a, b) such that the following conditions hold:

1. $m \in \mathbb{Z}$ and $m \in \left[2^{-\frac{1}{d}}L_n\left(\frac{2}{3}, \delta^{-1}\right), L_n\left(\frac{2}{3}, \delta^{-1}\right)\right]$.
2. $f \in \mathbb{Z}[x, y]$, $\deg(f) = d = \delta \sqrt[3]{\frac{\log n}{\log \log n}}$, $2 \nmid d$, with coefficients bounded by $L\left(\frac{2}{3}, \kappa\right)(1 + \mathbf{o}(1))$.

Moreover let the coefficients $\{c_i\}_{i \in I}$ be drawn uniformly at random such that

$$c_i \in \mathbb{I}\left(2L_n\left(\frac{2}{3}, \kappa - \delta^{-1}\right)\right).$$

3. $a, b \in \mathbb{Z}$, $0 \leq a < |b| \in \left[\frac{1}{2}L_n\left(\frac{1}{3}, \sigma\right)\right]$, with $a - bm$ B -smooth and $f(a, b)$ B' -smooth.

Moreover let $\mathcal{X}_{f, m, n} = \{(a, b) \in \mathbb{Z}^2 : (f, m, n, a, b) \in \mathcal{X}\}$. The first thing we note from this

definition and the definition of our parameters is that $f(x, y)$ is not generated uniformly at random as it has a component solely determined by m and n , $\hat{f}_{m,n}(x, y)$, which is a problem as we try to obtain that $f(a, b)$ is as likely to be B' -smooth as a uniformly random integer of the same size. To mitigate this we shall use the observation that c_i is much larger than m to prove that the randomized part of $f(x, y)$ dominates and that $f(a, b)$ therefore can be considered as a uniformly distributed along any arithmetic progression of common difference $(a - mb)$.

Once we have achieved this we shall show that for most B -smooth moduli $a - mb$, the B' -smooth numbers are approximately uniformly distributed modulo $(a - mb)$. It is only then that we can show that $f(a, b)$ is as likely to be B' -smooth as a random integer of the same size. Once we have shown this we shall venture to prove that there are sufficient (a, b) -pairs such that $a - mb$ is B -smooth and $f(a, b)$ is B' -smooth, which will lead to a proof of theorem 5.2. A key in this is the observation that $f(a, b)$ lies on the arithmetic progression given by

$$\left\{ \hat{f}_{m,n}(a, b) + (a - mb)z : |z| \leq dL_n \left(\frac{2}{3}, \kappa - \delta^{-1} \right) b^d \right\}.$$

5.2.1 Uniform distribution of the polynomial

Continuing our exposition using the definitions and parameters we have set up until now we recall the following:

$$\forall i \in \{1, \dots, d-1\} : (c_i) \sim \mu = \mathcal{U} \left(\mathbb{I} \left(2L_n \left(\frac{2}{3}, \kappa - \delta^{-1} \right) \right)^d \right)$$

and denote $\mathbf{c} = (c_i)$ to be the coefficient vector of f .

Note that for $f(a, b)$ to be B' -smooth it is necessary for $a - bm$ to be B -smooth, however this is not something we can simply assume. Therefore we consider the definition of $f(x, y)$ and note that

$$f(a, b) \equiv \hat{f}_{m,n}(a, b) \pmod{a - bm}.$$

As $\gcd(a, b)^d \mid f(a, b)\hat{f}_{m,n}(a, b)$ we have that $\gcd(a, b)^d(a - bm) \mid R(a, b)$. Now let a and b be uniform in their ranges, then we can give an explicit description of the probability that $a - mb$ is B -smooth:

Proposition 5.5. Fix b in its interval. If a, m are uniformly random, then:

$$\mathbb{P}_{a,m}(a - bm \text{ is } B\text{-smooth}) = L_n \left(\frac{1}{3}, \frac{\delta^{-1}}{3\beta} (1 + \mathbf{o}(1)) \right)^{-1}.$$

Proof. We fix b . Note that a is uniformly random on an interval of length b , and m is uniformly random over an interval of length comparable to its largest value. In particular:

$$\begin{aligned} a - bm &\sim \mathcal{U} \left[-bL_n \left(\frac{2}{3}, \delta^{-1} \right), -b \left(2^{-\frac{1}{2d}} L_n \left(\frac{2}{3}, \delta^{-1} \right) + 1 \right) \right] \\ &= \mathcal{U} \left[-x (1 + z^{-1}), -x \right], \end{aligned}$$

for $x = L_n \left(\frac{2}{3}, \delta^{-1} (1 + \mathbf{o}(1)) \right)$ and $z \approx 2^{\frac{1}{2d}} - 1 = \mathbf{O}(d^{-1})$. Note that $d = \log \left(B^{\frac{5}{12}} \right)$, and that $\log \log B = \mathbf{O}(\log \log x)$. Hence from theorem 4.36 the number of smooth values of $a - mb$ is:

$$\frac{\Psi(x, B)}{z} \left(1 + \mathbf{O} \left(\frac{\log(u+1)}{\log B} \right) \right).$$

Since the range of values is of length $\frac{x}{z}$,

$$\mathbb{P}_{a,m}(a - bm \text{ is } B\text{-smooth}) = \rho(x, B) \left(1 + \mathbf{O} \left(\frac{\log(u+1)}{\log B} \right) \right).$$

Recall that $B = L_n \left(\frac{1}{3}, \beta \right)$ and $x = L_n \left(\frac{2}{3}, \delta^{-1} \right)^{-1+\log(1)}$. Furthermore, note that $\log u < \log \log n = \mathbf{o}(\log B)$. Hence by corollary 4.35

$$\rho \left(L_n \left(\frac{2}{3}, \delta^{-1} \right)^{1+\mathbf{o}(1)}, B \right) = L_n \left(\frac{1}{3}, \frac{\delta^{-1}}{3\beta} \right)^{-1+\mathbf{o}(1)}.$$

Absorbing the multiplicative $1 + \mathbf{o}(1)$ into the $\mathbf{o}(1)$ error term in the exponent to obtain

$$\mathbb{P}_{a,m}(a - bm \text{ is } B\text{-smooth}) = L_n \left(\frac{1}{3}, \frac{\delta^{-1}}{3\beta} (1 + \mathbf{o}(1)) \right)^{-1}.$$

□

Now we can fix the residue of $f(a, b) \pmod{a - mb}$ and consider only how the randomness of our coefficients c_i affect the polynomial. Therefore we will also want to be explicit about the coefficient vector, $c = (c_i)_{i \in \{1, \dots, d-1\}}$, for $f(x, y)$, hence we will, when when we need to be explicit, denote $f(x, y) = f_c(x, y)$.

Now we are set up to show that $f(a, b)$ can be considered as uniformly distributed along progressions of common difference $(a - mb)$.

Lemma 5.6. Let $a < b$, with $\gcd(a, b) = 1$, define $\varphi = \varphi_{a,b} : \mathbb{Z}^d \rightarrow \mathbb{Z}$ by the following

$$\varphi((v_0, \dots, v_{d-1})) = \sum_{i=0}^{d-1} v_i a^{d-i-1} b^i.$$

There exists a set $S \subseteq \mathbb{I}(4L_n(\frac{1}{3}, \sigma))^d$ such that φ bijects S and $\mathbb{I}(b^{d-1})$.

Proof. For each $i \geq 0$, we claim that for any $|t| \leq b^i + a^{i+1}$ there exists a representation

$$t = a^i x_0 + a^{i-1} b x_1 + \dots + b^i x_i,$$

with $|x_0|, \dots, |x_i| \leq a + b$. We proceed by induction on the number of terms. If $i = 0$ then the conclusion is trivial as $|t| \leq b^0 + a^1 = a + 1$ and so

$$t = a^0 x_0 = x_0,$$

for $|x_0| \leq a + b$, and as $a + 1 \leq a + b$ the result follows. Now let $i > 0$, then we may choose

$|y| < a$ such that

$$|t - ya^i| \leq b^i.$$

Letting $z \in \{0, \dots, b-1\}$ such that $za^i \equiv t - ya^i \pmod{b}$ we can set $x_0 = y + z$ and as $|y| < a$ and $|z| < b$ it follows that $|x_0| \leq a + b$. By construction $b \mid t - x_0a^i$ and

$$\left| \frac{t - x_0a^i}{b} \right| = \left| \frac{(t - ya^i) - za^i}{b} \right| \leq b^{i-1} + a^i.$$

Hence by the induction hypotheses such a representation exists for every t with $|t| \leq b^i + a^{i+1}$. Now we show the existence of S directly. Let $t \in \mathbb{I}(b^{d-1})$, then for any such t : $|t| \leq b^{d-1}$ and so the conditions of the above hold with $i = d-1$. Hence there exists a sequence $(x_i)_{i \in \{0, \dots, d-1\}}$ such that

$$\sum_{i=0}^{d-1} x_i a^{d-i-1} b^i = t,$$

with $|x_i| < a + b$ for all i . Hence the vector given by the sequence, $\mathbf{x}_t = (x_i)$, fulfills that $\varphi(\mathbf{x}_t) = t$. Moreover, $\mathbf{x}_t \in \mathbb{I}(2(a+b))^d$ and by Defn 2.4.3 this means $\mathbf{x}_t \in \mathbb{I}(4L_n(\frac{1}{3}, \sigma))^d$. Hence, taking $S = \{\mathbf{x}_t : t \in \mathbb{I}(b^{d-1})\}$ gives the construction of a single vector \mathbf{x}_t , hence φ is bijective on S . \square

Remark. Note that we have presupposed that $\gcd(a, b) = 1$, but it is a simple exercise to show that this is trivial. By the definition of $(a, b) \in \mathcal{X}_{f, m, n}$ we know that $a - mb$ is B -smooth, hence $\gcd(a, b)$ is B -smooth. Since we will only be interested in looking at B -smoothness to show the uniform randomness of $f(x, y)$ we can assume without loss of generality that $\gcd(a, b) = 1$ as division of two B -smooth numbers preserves B -smoothness.

We can now reformulate $f(x, y)$ using the map φ . For this consider m, n fixed and observe that:

$$f(a, b) = f_c(a, b) = f_{c'}(a, b) + (a - mb)\varphi_{(a, b)}(c - c').$$

This motivates the following:

Definition 5.7. Let S be the set given by the lemma 5.6. For any $l \leq b^{d-1}$, we define a set S_l

and a measure ν_l as follows:

$$S_l = \{v \in S \mid \varphi(v) = \mathbb{I}(l)\},$$

$$\nu_l = \mathcal{U}(S_l).$$

Particularly, ν_l gives a uniformly random element of S whose image under φ is in $\mathbb{I}(l)$.

It is clear to see that if $v \sim \nu_l$, then we can write

$$f_v(a, b) \sim f_{\mathbf{0}}(a, b) + (a - mb)\mathcal{U}(\mathbb{I}(l)).$$

This allows the consideration that for the measures ν_l with support S_l give us additive alterations to c which change $f(a, b)$ additively by $a - mb$ times a uniformly random value. This means that if f is indeed uniformly random, then we can replace the randomness of c over cosets of S_l with randomness of $f_c(a, b)$ over arithmetic progressions.

Definition 5.8. For $\bar{\mu} : X \rightarrow \mathbb{R}^+$ a measure and $F : X \rightarrow Y$ a function, we define the measure $F^{\bar{\mu}} : Y \rightarrow \mathbb{R}^+$ given by

$$F^{\bar{\mu}}(y) = \bar{\mu}(\{F^{-1}(y)\}) = \sum_{x:F(x)=y} \bar{\mu}(x).$$

Remark. It will be left to the reader to prove that this is in fact a measure. Moreover note that if F is bijective then $F^{\bar{\mu}}(y) = \bar{\mu}(x)$, for some unique $x \in X$.

Now consider $c \sim \nu_l$ such that we can write

$$f_c(a, b) \sim f_{\mathbf{0}}(a, b) + (a - mb)\mathcal{U}(\mathbb{I}(l)).$$

As μ is a uniform distribution on a cube of side $2L_n \left(\frac{2}{3}, \kappa - \delta^{-1}\right)$ and we have a bijective function $\varphi : S \rightarrow \mathbb{I}(b^{d-1})$ we can now show that φ^μ is actually close to a convolution of uniform distributions on intervals. In fact, by convolving with several distributions, ν_{l_i} , we can treat each coefficient in f as if it were independent and uniformly random. This is made rigorous in the following proposition:

Proposition 5.9. Fix a, b . There is a distribution ϑ such that ϑ is the convolution of uniform distribution on intervals of lengths $L_n \left(\frac{2}{3}, \kappa - \delta^{-1}\right) a^i b^{d-1-i}$ for $i \in \{0, \dots, d-1\}$ with

$$\|\varphi^\mu - \vartheta\|_1 = \mathbf{O} \left(L_n \left(\frac{2}{3}, (\kappa - \delta^{-1})(1 + \mathbf{o}(1)) \right)^{-1} \right),$$

and

$$|\mathbb{E}(\vartheta)| \leq \sum_{i=0}^{d-1} a^i b^{d-i-1}.$$

Proof. Recall that the convolution of distributions, given by \star , is defined by Definition 4.42 and define

$$\nu = \mu \star \left[\star_{i=0}^{d-1} \nu_{a^i b^{d-i-1}} \right].$$

From lemma 5.6, the support of $\nu_{a^i b^{d-i-1}}$ is contained in a cube of side $4L_n \left(\frac{1}{3}, \sigma\right)$. Hence the support of P of $\star_{i=0}^{d-1} \nu_{a^i b^{d-i-1}}$ is contained in a cube of side $4dL_n \left(\frac{1}{3}, \sigma\right)$. When $\|x\|_\infty < L_n \left(\frac{2}{3}, \kappa - \delta^{-1}\right) - 4dL_n \left(\frac{1}{3}, \sigma\right)$ and $p \in P$:

$$\mu(x-p) = \mu(x) = |\text{sup}(\mu)|^{-1},$$

so $\nu(x)$ is a convex combination of values in $\{\mu(x-p) : p \in P\} = \{\mu(x)\}$. Hence $\nu(x) = \mu(x)$ on the l_∞ ball of radius $L_n \left(\frac{2}{3}, \mathbf{o}(\kappa - \delta^{-1}) - 4dL_n \left(\frac{1}{3}, \sigma\right)\right)$. Then since $L_n \left(\frac{1}{3}, \sigma\right)$ is $L_n \left(\frac{2}{3}, \mathbf{o}(\kappa - \delta^{-1})\right)$:

$$\begin{aligned} \mathbb{P}_{x \sim \mu}(\nu(x) - \mu(x)) &\geq \left(1 - 4dL_n \left(\frac{2}{3}, \kappa - \delta^{-1} \right)^{-1+\mathbf{o}(1)} \right)^d \\ &\geq 1 - 4d^2 L_n \left(\frac{2}{3}, \kappa - \delta^{-1} \right)^{-1+\mathbf{o}(1)} = 1 - L_n \left(\frac{2}{3}, \kappa - \delta^{-1} \right)^{-1+\mathbf{o}(1)}. \end{aligned}$$

In particular, we have a bound on the l_1 distance between μ and ν :

$$\begin{aligned} \|\mu - \nu\|_1 &= \sum_{x \in \mathbb{Z}^d} |\nu(x) - \mu(x)| \\ &\leq \mathbb{P}_{x \sim \mu}(\nu(x) \neq \mu(x)) (\|\mu\|_\infty + \|\nu\|_\infty) \end{aligned}$$

$$= \mathbf{O} \left(L_n \left(\frac{2}{3}, (\kappa - \delta^{-1}) (1 + \mathbf{o}(1)) \right)^{-1} \right).$$

Now for fixed a, b we apply the map φ to μ and ν to obtain

$$\|\varphi^\mu - \varphi^\nu\|_1 \leq \|\mu - \nu\|_1 \leq \mathbf{O} \left(L_n \left(\frac{2}{3}, (\kappa - \delta^{-1}) (1 + \mathbf{o}(1)) \right)^{-1} \right),$$

and so the l_1 difference of the distributions $\varphi(\mu)$ and $\varphi(\nu)$ on \mathbb{Z} is small. From definition 5.7 $\varphi^\nu = \mathcal{U}(\mathbb{I}(l))$. Since applying our map φ to a measure commutes with convolution of measures:

$$\varphi^\nu = \varphi^\mu \star [\star_{i=0}^{d-1} \mathcal{U}(\mathbb{I}(a^i b^{d-i-1}))].$$

Since $c_i \sim \mathcal{U}(\mathbb{I}(L_n(\frac{2}{3}, \kappa - \delta^{-1})))$ are independent random variables,

$$c_i a^i b^{d-i-1} + \mathcal{U}(\mathbb{I}(a^i b^{d-i-1}))$$

is uniformly distributed along an interval of length $L_n(\frac{2}{3}, \kappa - \delta^{-1}) a^i b^{d-i-1}$. Note that $\mathbf{c} \sim \mu$. Hence there is a constant C such that for all $x \in \mathbb{Z}$:

$$\varphi^\nu(x) \sim \star_{i=0}^{d-1} \left[\mathcal{U} \left(\mathbb{I} \left(L_n \left(\frac{2}{3}, \kappa - \delta^{-1} \right) a^i b^{d-i-1} \right) \right) \right] (x - C).$$

The shift C accounts for the difference in expectation caused by the fact that the intervals associated with the independent random variables are not centred. However the centre of each of these intervals has modulus at most $\frac{1}{2} a^i b^{d-i-1} + \frac{1}{2}$, and so $|C| \leq \sum_{i=0}^{d-1} a^i b^{d-i-1}$. We take $\vartheta = \varphi^\nu$ to complete the proof. \square

Hence, this convolution allows us to replace $R(a, b)$ by $a - mb$ times some convolution of uniform measures on intervals. This proves that $f(a, b)$ is close to uniformly distributed along arithmetic progressions of common difference $(a - mb)$, so it rests to show that for most B -smooth moduli $a - mb$, the B' -smooth numbers are almost uniformly distributed modulo $a - mb$. To do this we introduce an estimation of smoothness, which is called goodness:

Definition 5.10. Fix a smoothness bound $B = L_n(\frac{1}{3})$ and $F = L_n(\frac{2}{3})$. Define $\epsilon(F, B, r, a) = \mathfrak{o}_n(1)$. We then say that a modulus r is B -good for F if uniformly over all $(a, r) = 1$:

$$\Psi(F, B, r, a) = \left(\frac{\Psi_r(F, B)}{\phi(r)} \right)^{1+\epsilon},$$

where ϕ is the Euler phi-function.

Remark. Recall from Definition 4.32

$$\Psi(x, y, r, a) = |\{z \in \mathbb{N} : z < x, z \text{ is } y\text{-smooth}, z \equiv a \pmod{r}\}|$$

and

$$\Psi_r(x, y) = |\{z \in \mathbb{N} : z < x, z \text{ is } y\text{-smooth}, \gcd(z, r) = 1\}|.$$

We will often suppress the ϵ as we are only interested in an error exponent up to $\mathfrak{o}(1)$. Moreover, we will say r is B -bad for F if it is not B -good for F and that r is B -good (dropping the F) if r is B -good for all $F \in [L_n(\frac{2}{3})\omega^{-1}, L_n(\frac{2}{3})]$ for $\omega = L_n(\frac{1}{3})$.

With this notion we can start constraining the behaviour of $f(a, b) \pmod{a - mb}$ for this we will need to characterise the moduli for which the smooth numbers are uniformly distributed across their residue classes. Once we have that it doesn't matter which residue class $f_c(a, b)$ lies in as it does not affect its probability of being smooth as c varies.

5.2.2 Uniform Distribution of Smooth Numbers

From Section 4.3 we have a result of [30] proving that, under some conditions, y -smooth numbers $\leq x$ are uniformly distributed over the $\phi(q)$ residue classes $a \pmod{q}$ with $\gcd(a, q) = 1$. However the conditions prevent us from concluding uniform distribution directly, hence if we can show that the distribution of B' -good numbers is close enough to uniform assuming $a - mb$ is B -smooth, to be considered as such then we are done. To do this we consider a Bombieri–Vinogradov style theorem proposed by Harper, [23]:

Theorem 5.11. Let c and K be fixed and effective constants. Then for any $\log^K F < B < F$ with $u = \frac{\log F}{\log B}$, and $Q \leq \sqrt{\Psi(F, B)}$:

$$\sum_{r \leq Q} \max_{(s,t)=1} \left| \Psi(F, B, r, s) - \frac{\Psi_r(F, B)}{\phi(r)} \right|$$

$$\ll \Psi(F, B) \left(e^{-\frac{cu}{\log^2(u)}} + B^{-c} \right) + Q \sqrt{\Psi(F, B)} \log^{7/2} F,$$

with an implied effective constant $C = C(c, k)$.

Remark. We remark that the definition of “effective” differs per application. For our purpose we may simply claim that $c < K$ are chosen fittingly for the to-factor n .

If we consider

$$Q_{max} = \max_{b,m} |a + bm| = L_n \left(\frac{2}{3}, \delta^{-1}(1 + \mathbf{o}(1)) \right),$$

we can reconsider our question to the equivalent notion of bounding the probability that a B -smooth modulus less than Q_{max} is B' -bad. This would immediately follow from showing that the number of B' -bad moduli below the Q -bound is much smaller than $\Psi(Q_{max}, B)$. We will, for our specific needs, bound the sum in theorem 5.11 as we know that the common difference in the arithmetic progression is known to be y -smooth. For this let $q = a - mb$, then:

Proposition 5.12. Let $\epsilon > 0$ be fixed. Then there exists effective constants K, c such that for any $\log^K x < y < x^{\frac{1}{\log \log x}}$ with $u = \frac{\log x}{\log y}$, $x^\epsilon \leq Q \leq \sqrt{\Psi(x, y)}$ and $\omega = \omega(1)$ with $\omega = y^{\mathbf{O}(1)}$:

$$\sum_{\substack{r \in [Q\omega^{-1}, Q] \\ r \text{ is } y\text{-smooth}}} \max_{\gcd(a,r)=1} \left| \Psi(x, y, r, a) - \frac{\Psi_r(x, y)}{\phi(r)} \right|$$

$$\ll \Psi(x, y) \rho(Q, y) \left(e^{-\frac{cu}{\log^2 u}} + y^{-c} \right) + Q \sqrt{\Psi(x, y)} \log^{7/2} x,$$

for some effective implied constant.

Remark. The proof for this is mostly an exercise in restating Harper’s [23] and is therefore omitted. The full proof can be found in [14].

Now we return to what we set out to prove: For most B -smooth moduli $a - mb$, the B' -smooth numbers are uniformly distributed. This follows immediately from the following two propositions:

Proposition 5.13. Fix a, b, m, n in their intervals and let f be uniformly random as before. Then:

$$\mathbb{P}_f (f(a, b) \text{ is } B'\text{-smooth} \mid (a - mb) \text{ is } B'\text{-good}) = L_n \left(\frac{1}{3}, \frac{\kappa + \sigma\delta}{3\beta'} (1 + \mathbf{o}(1)) \right)^{-1}.$$

Proof. Let $a - bm = r$. By proposition 5.9:

$$\begin{aligned} \mathbb{P}_{n,f} (f_{n,m}(a, b) \text{ is } B'\text{-smooth}) &= \mathbb{P}_{n,c} \left(\hat{f}_{m,n}(a, b) + r\varphi_{(a,b)}(c) \text{ is } B'\text{-smooth} \right) \\ &= \mathbb{P}_{n,\vartheta} \left(\hat{f}_{m,n}(a, b) + r\vartheta \text{ is } B'\text{-smooth} \right) + \mathbf{O} \left(L_n \left(\frac{2}{3}, (\kappa - \delta^{-1})(1 + \mathbf{o}(1)) \right)^{-1} \right). \end{aligned}$$

By the definition of the parameters we have $\kappa > \delta^{-1}$ and therefore $\kappa - \delta^{-1} > 0$. Now recall, from proposition 5.9, ϑ has $|\mathbb{E}(\vartheta)| \leq \sum_{i=0}^{d-1} a^i b^{d-i-1}$ and is sampled according to the convolution of uniform measures on intervals of length $L_n \left(\frac{2}{3}, \kappa - \delta^{-1} \right) a^i b^{d-i-1}$ for $i = 0, \dots, d-1$. Hence ϑ is unimodal with mode at some M satisfying

$$|M| \leq \sum_{i=0}^{d-1} a^i b^{d-i-1} < db^{d-i-1} = L_n \left(\frac{2}{3}, \sigma\delta(1 + \mathbf{o}(1)) \right).$$

Now define

$$F_{\max} = L_n \left(\frac{2}{3}, \kappa - \delta^{-1} \right) (a - mb) \sum_{i=0}^{d-1} a^i b^{d-i-1},$$

then the support of ϑ is contained in $[M - F_{\max}|r|^{-1}, M + F_{\max}|r|^{-1}]$. We choose an $\omega = L_n \left(\frac{2}{3}, \mathbf{o}(1) \right)$, such that $\omega \rightarrow \infty$, and set

$$Y = L_n \left(\frac{e}{3}, \kappa - \delta^{-1} \right) b^{d-1} \omega^{-1} = L_n \left(\frac{2}{3}, \kappa - \delta^{-1} + \sigma\delta - \mathbf{o}(1) \right).$$

Now we define a measure ϑ' to be

$$\vartheta'(x) = \begin{cases} \vartheta(\max(x, Y)) & x \geq 0 \\ \vartheta(\min(x, -Y)) & x < 0 \end{cases}.$$

Then

$$\|\vartheta' - \vartheta\|_1 \leq \mathbb{O}_{z \sim \vartheta}(|z| < Y) \leq 2Y \left(L_n \left(\frac{2}{3}, \kappa - \delta^{-1} \right) b^{d-1} \right)^{-1} = 2\omega^{-1},$$

from the definition of Y . Note that Y is much larger than M and so ϑ' is monotone decreasing away from 0; hence there are non-negative weights W_y for $y \in \mathbb{Z}$, with $W_y = 0$ for $|y| > F_{\max}|r|^{-1}$ such that:

$$\vartheta' = \sum_{y \geq Y} W_y \mathcal{U}([0, y]) + W_{-y} \mathcal{U}([-y, 0]),$$

and $\left| 1 - \sum_y W_y \right| \leq 2\omega^{-1}$. Hence we have:

$$\begin{aligned} \mathbb{P}_f(f_{m,n}(a, b) \text{ is } B'\text{-smooth}) &= \mathbf{O}(\omega)^{-1} + \\ &\sum_{y=Y}^{F_{\max}|r|^{-1}} W_y \mathbb{P} \left(\hat{f}_{m,n}(a, b) + r\mathcal{U}([0, y]) \text{ is } B'\text{-smooth} \right) \\ &+ W_{-y} \mathbb{P} \left(\hat{f}_{m,n}(a, b) + r\mathcal{U}([-y, 0]) \text{ is } B'\text{-smooth} \right). \end{aligned}$$

We note that $\mathbf{O}(\omega^{-1}) = L_n \left(\frac{2}{3}, \mathbf{o}(1) \right)^{-1}$ terms can be absorbed into our $\mathbf{o}(1)$ terms, and so it suffices to show that for any fixed, B' -good r and any $y \in [Y, F_{\max}|r|^{-1}]$:

$$\mathbb{P} \left(\hat{f}_{m,n}(a, b) + r\mathcal{U}([0, y]) \text{ is } B'\text{-smooth} \right) = L_n \left(\frac{1}{3}, \frac{\kappa + \sigma\delta}{3\beta'} \right)^{-1+\mathbf{o}(1)},$$

$$\mathbb{P} \left(\hat{f}_{m,n}(a, b) + r\mathcal{U}([-y, 0]) \text{ is } B'\text{-smooth} \right) = L_n \left(\frac{1}{3}, \frac{\kappa + \sigma\delta}{3\beta'} \right)^{-1+\mathbf{o}(1)}.$$

Since $\left| \hat{f}_{m,n}(a, b) \right| \leq \hat{F}_{\max} := Y L_n \left(\frac{2}{3} \right)^{-1}$, we can absorb the probability that the value on the left is negative or positive (respectively) in the above two equations. From the definition of B' -good

and corollary 4.39 , for any $x \in [r|Y - \hat{F}_{\max}, F_{\max} + \hat{F}_{\max}]$:

$$\Psi(x, B', r, s) = \frac{\Phi_r(x, B')}{\phi(r)} L_n \left(\frac{1}{3}, \mathbf{o}(1) \right) = \frac{\Psi(x, B')}{r} L_n \left(\frac{1}{3}, \mathbf{o}(1) \right),$$

and so to finish the estimate we observe that for any $x \in [r|Y - \hat{F}_{\max}, F_{\max} + \hat{F}_{\max}]$:

$$\rho(x, B') = \rho \left(L_n \left(\frac{2}{3}, \kappa + \sigma\delta \right), B' \right) = L_n \left(\frac{1}{3}, \frac{\kappa + \sigma\delta}{3\beta'} \right)^{-1+\mathbf{o}(1)}.$$

□

Now there is only one thing left to do: Show that $a - bm$ is B' -good significantly often when $a - bm$ is B -smooth, which we have assumed from the start and have a significant probability for from proposition 5.5.

Proposition 5.14. Fix any b . Then

$$\mathbb{P}_{a,m}(a - mb \text{ is } B'\text{-good} \mid a - mb \text{ is } B\text{-smooth}) = 1 - \mathbf{o}(1).$$

Proof. We begin by bounding the number of moduli which are F -bad for some

$$F \in \left[F_{\max} L_n \left(\frac{1}{3} \right)^{-1}, F_{\max} \right].$$

We fix $\omega = B'$ for concreteness. Observe that $\Psi(F, B') = F L_n \left(\frac{1}{3} \right)^{-1}$. Since $L_n \left(\frac{2}{3} \right) = \omega \left(L_n \left(\frac{1}{3} \right) \right)$.

$$Q \leq \sqrt{\Psi(F, B')} L_n \left(\frac{2}{3}, \frac{\epsilon}{4} \right)^{-1}.$$

Furthermore, for any K fixed, $\omega \left(\log^K F \right) = B' = \mathbf{o} \left(F^{\frac{1}{\log \log F}} \right)$. Hence we can apply proposition 5.12. Suppose that a modulus r is B -smooth and also B' -bad for F . Then for some residue a

with $\gcd(a, r) = 1$, the contribution to the LHS of proposition 5.12 for this r is at least

$$\frac{\Psi(F, B')}{\phi(r)}(1 + \mathbf{o}(1)) = \frac{\Psi(F, B')}{r} \geq \frac{\Psi(F, B')}{Q},$$

where for the first equality we use corollary 4.39, noting that $B \leq B'$ so r is B' -smooth, $u < \log \log n$ and the number of divisors of r is bounded by $\log r$ so that the multiplicative error is $1 + \mathbf{o}(1)$. Now

$$\begin{aligned} & \sum_{\substack{r \in [Q_{\max} \omega^{-1}, Q_{\max}] \\ r \text{ is } y\text{-smooth}}} \max_{\gcd(a, r)=1} \left| \Psi(F, B', r, a) - \frac{\Psi_r(F, B')}{\phi(r)} \right| \\ & \leq C \Psi(F, B') \rho(Q_{\max}, B') \left(e^{-\frac{cu'}{\log^2 u}} + B'^{-c} \right) + Q_{\max} \sqrt{\Psi(F, B')} \log^{7/2} F \\ & = CF \rho(F, B') \rho(Q_{\max}, B') \left(e^{-\frac{cu'}{\log^2 u}} + B'^{-c} \right) + Q_{\max} F^{\frac{1}{2}} \rho(Fc, B')^{\frac{1}{2}} \log^{7/2} F. \end{aligned}$$

First, we observe that F and Q_{\max} are $L_n(23)$, whilst B' is $L_n(\frac{1}{3})$. Hence both densities $\rho(Q_{\max}, B')$ and $\rho(F, B')$ are $L_n(\frac{1}{3})^{-1}$. From the definition of Q_{\max} and F_{\max} we have

$$Q_{\max} = L_n \left(\frac{2}{3}, \delta^{-1}(1 + \mathbf{o}(1)) \right), F_{\max} = L_n \left(\frac{2}{3}, (\kappa + \sigma\delta)(1 + \mathbf{o}(1)) \right),$$

and from the parameters defined in equation (2) we have $2\delta^{-1} < \kappa + \sigma\delta$. Hence $FQ_{\max}^{-2} = L_n(\frac{2}{3})$. Since up to order $L_n(\frac{2}{3})^{\mathbf{o}(1)}$, the first term is F and the second is $Q_{\max}F^{\frac{1}{2}}$, we deduce that the first term dominates the second. If r is B' -bad for F it contributes at least $\Psi(F, B')Q_{\max}^{-1}(1 + \mathbf{o}(1))$ to the sum on the left hand side. Hence the number of moduli which are in $[Q_{\max}\omega^{-1}, Q_{\max}]$, are B -smooth and B' -bad for F is at most

$$\begin{aligned} & \frac{Q_{\max}}{\Psi(F, B')} \Psi(F, B') \rho(Q_{\max}, B') (C + \mathbf{o}(1)) \left(e^{-\frac{cu'}{\log^2 u'}} + B'^{-c} \right), \\ & = \Psi(Q_{\max}, B') (C + \mathbf{o}(1)) \left(e^{-\frac{cu'}{\log^2 u'}} + B'^{-c} \right). \end{aligned}$$

If a modulus is B' -bad near F_{\max} , it must be B' -bad for some

$$F \in \left\{ F_{\max} L_n \left(\frac{1}{3} \right)^{-1}, F_{\max} \right\} \cup \left\{ 2^i : 2^i \in \left[F_{\max} L_n \left(\frac{1}{3} \right), F_{\max} \right] \right\},$$

which is a set of logarithmic size. We can absorb a logarithmic factor into the constants c, C , so the number of moduli which are in $[Q_{\max}\omega^{-1}, Q_{\max}]$, are B -smooth and B' -bad is at most:

$$\Psi(Q_{\max}, B') (C + \mathbf{o}(1)) \left(e^{-\frac{cu'}{\log^2 u'}} + B'^{-c} \right) = \mathbf{o}(\Psi(Q_{\max} B')).$$

Hence even assuming that every B -smooth number below $Q_{\max}\omega^{-1}$ is B' -bad gives:

$$\begin{aligned} & \mathbb{P}_{a,m}(a - mb \text{ is } B'\text{-good} \mid a - mb \text{ is } B\text{-smooth}) \\ & \geq 1 - \frac{\Psi(Q_{\max}\omega^{-1}, B') + \Psi(Q_{\max}, B')}{\Psi(Q_{\max}, B')} = 1 - \mathbf{o}(1). \end{aligned}$$

□

Hence the probability of a B -smooth modulus to be B' -good is nearing 1 and the probability for $f(a, b)$ to be B' -smooth given that $a - mb$ is a B -smooth modulus is comparable to a random integer of similar size to be B' -smooth under the same conditions. Hence we can now prove theorem 5.2.

5.2.3 Proof of theorem 5.2

Now that we have shown that the changes that randomization causes are controllable we are able to conclude that they are also sufficient to give a provable complexity. For this let us quickly reflect on our choice compared to the general number field sieve. Consider $(f, m, n, a, b) \in \mathcal{X}$ and let $\alpha \in \mathbb{C}$ be a root of the first coordinate of f , $f(\alpha, 1) = 0$. Then there is a ring homomorphism

$$\mathbb{Z}[\alpha] \rightarrow \mathbb{Z}/n\mathbb{Z}, (a - b\alpha) \mapsto (a - bm) \pmod{n}.$$

From the field norm we also have a multiplicative map from $\mathcal{O}_{\mathbb{Q}(\alpha)} \rightarrow \mathbb{Z}$, as we would expect from the general number field sieve. Moreover, denoting f_d for the leading coefficient of f , the field norm is guaranteed to be in $\frac{1}{f_d}\mathbb{Z}$ for any element of $\mathbb{Z} + \alpha\mathbb{Z}$. Hence the only thing that we can't be sure about is that f is irreducible as f is considered uniformly at random. If f was irreducible then we can apply the strategy from the general number field sieve to obtain a congruence of squares modulo n .

Lemma 5.15.

$$\mathbb{P}(f \text{ is reducible}) \leq L_N \left(\frac{2}{3}, \frac{\kappa - \delta^{-1} + \mathbf{o}(1)}{3} \right)^{-1}.$$

Proof. Fix m, n and let $H = 2L_N \left(\frac{2}{3}, \kappa - \delta^{-1} \right)$ be the range of each coefficient of the random part of our polynomial f . Note that if a polynomial over \mathbb{Z} is reducible if it is reducible modulo every prime. Hence if we bound the number of reducible polynomials modulo \mathbb{F}_p for each prime p and bound how often a polynomial is reducible modulo several primes p , we can get good bounds on the number of irreducible polynomials f . We count the reducible polynomials f with the Turán sieve, [11]. Let $\mathcal{A} = \{(c_{d-1}, \dots, c_0) \in \mathbb{Z}^d, |c_i| < H\}$ which we equate with the set of polynomials $f(x, y) = \hat{f}_{m,n}(x, y) + R(x, y)$ with f and $\hat{f}_{m,n}$ are both homogeneous of degree d with (c_i) the coefficients of R . For any prime r , let \mathcal{A}_r correspond to the subset of \mathcal{A} corresponding to irreducible polynomials $\pmod r$. Note that for any f to correspond to $g \pmod r$ we must have $(x - my) \mid \hat{f}_{m,n} - g \in \mathbb{F}_r[x, y]$ or equivalently $g(m, q) \equiv n \pmod r$. We do not insist that G is monic, although any irreducible g must be a scalar multiple of a monic irreducible. To estimate the number of irreducibles we accept the following fact, [[7], Chapter 2]:

For any $0 < i < r$, the number of monic irreducible g of degree D such that $g(m) \equiv i \pmod r$ is $\frac{r^{D-1}}{D(D-1)} + \mathbf{O}(r^{D/2})$.

Note that for any g over \mathbb{F}_r , such that $g(m) \equiv n$ with $r \ll \sqrt{H}$, there are

$$\left(\frac{H}{r} + \mathbf{O}(1) \right)^d = \left(\frac{H}{r} \right)^d + \mathbf{O} \left(\left(\frac{H}{r} \right)^{d-1} \right)$$

polynomials lying over g in \mathcal{A} , and none if $g(m) \not\equiv n \pmod{r}$. Hence by a union bound over the irreducibles \pmod{r} :

$$|\mathcal{A}_r| \leq \frac{H^d}{d(d-1)} + \mathbf{O}\left(\frac{H^d}{r^{\frac{d}{2}-1}}\right) + \mathbf{O}(H^{d-1}r),$$

$$|\mathcal{A}_r \cap \mathcal{A}_{r'}| \leq \frac{H^d}{d^2(d^2-1)} + \mathbf{O}\left(\frac{H^d}{r^{\frac{d}{2}-1}}\right) + \mathbf{O}\left(\frac{H^d}{r'^{\frac{d}{2}-1}}\right) + \mathbf{O}(H^{d-1}rr').$$

From the Turán sieve, considering all primes $r \leq p$ for any p much smaller than \sqrt{H} , the number of reducible polynomials f is much smaller than $H^d p^{-1} \log p + H^{d-1} p^2$, which for $p \sim H^{1/3} \log^{1/3} H$ is $H^{d-\frac{1}{3}} \log^{\frac{2}{3}} H$. Then the number of potential f for fixed n, m is H^d , and so the probability that f is reducible is at most

$$H^{-\frac{1}{3}} \log^{\frac{2}{3}} H = L_n \left(\frac{2}{3}, \frac{\kappa - \delta^{-1} + \mathbf{o}(1)}{3} \right)^{-1}.$$

□

Remark. If f is reducible we immediately assume that the algorithm fails, hence by sampling at most $L_n \left(\frac{1}{3}\right)$ polynomials the probability that any of them are reducible is $\mathbf{o}(1)$.

Now assume that f is irreducible. Then the following theorem gives us the last ingredient we need to prove theorem 5.2:

Theorem 5.16. With $\beta = \beta', \delta, \sigma, \kappa$ chosen subject to (5), (6), (7), and Definition 5.4:

$$\mathbb{E}_{m,f} (|\mathcal{X}_{f,m,n}|) \geq L_n \left(\frac{1}{3}, \tau \right),$$

with

$$\tau = 2\sigma - \frac{\delta^{-1}}{3\beta'}(1 + \mathbf{o}(1)) + \frac{\sigma\delta + \kappa}{3\beta}(1 + \mathbf{o}(1)).$$

Proof. As proposition 5.5 and proposition 5.12 randomise over a, m for any fixed b , and uniformly over n, f , for any b, n, f :

$$\mathbb{P}_{a,m} (a - bm \text{ is } B\text{-smooth and } B'\text{-good}) = L_n \left(\frac{1}{3}, \frac{\delta^{-1}}{3\beta}(1 + \mathbf{o}(1)) \right)^{-1}.$$

Since proposition 5.11 randomizes over f for any fixed a, b, m , we have for each fixed b that

$$\begin{aligned} & \mathbb{P}_{a,m,f}(a - bm \text{ is } B\text{-smooth and } B'\text{-good, } f(a, b) \text{ is } B'\text{-smooth}) \\ &= L_n \left(\frac{1}{3}, \frac{\delta^{-1}}{3\beta} + \frac{\kappa + \sigma\delta}{3\beta'} \right)^{-1+\mathbf{o}(1)}, \end{aligned}$$

as multiplicative factors of $1 + \mathbf{o}(1)$ may be absorbed into the $\mathbf{o}(1)$ in the exponent of the $L_n \left(\frac{1}{3}\right)$ terms. Summing over the $L_n \left(\frac{1}{3}, \sigma\right)^2$ choices for a fixed pair (a, b) :

$$\begin{aligned} \mathbb{E}_{m,f}(|\mathcal{X}_{f,m,n}|) &= \sum_{a,b} \mathbb{P}_{f,m,n}((f, n, m, a, b) \in \mathcal{X}) \\ &= L_n \left(\frac{1}{3}, \sigma \right) \sum_b \mathbb{P}_{f,m,n,a} \left(\begin{array}{l} (a - bm) \text{ is } B\text{-smooth} \\ \wedge f(a, b) \text{ is } B'\text{-smooth} \end{array} \right) \\ &\geq L_n \left(\frac{1}{3}, \sigma \right) \sum_b \mathbb{P}_{f,m,n,a} \left(\begin{array}{l} (a - bm) \text{ is } B\text{-smooth} \\ \wedge (a - bm) \text{ is } B'\text{-good} \\ \wedge f(a, b) \text{ is } B'\text{-smooth} \end{array} \right) \\ &\geq L_n \left(\frac{1}{3}, 2\sigma - \left(\frac{\delta^{-1}}{3\beta} \right) (1 + \mathbf{o}(1)) + \left(\frac{\sigma\delta + \kappa}{3\beta'} \right) (1 + \mathbf{o}(1)) \right). \end{aligned}$$

□

Proof of theorem 5.2. Let $\tau = 2\sigma - \frac{\delta^{-1}}{3\beta'} - \frac{\sigma\delta + \kappa}{3\beta}$, and note that:

$$\lambda \geq \max(\beta, \beta') + \frac{\delta^{-1}(1 + \mathbf{o}(1))}{3\beta} + \frac{(\sigma\delta + \kappa)(1 + \mathbf{o}(1))}{3\beta'} = \max(\beta, \beta') + 2\sigma - \tau + \mathbf{o}(1).$$

For any fixed pair (m, f) , we can use the hyper-elliptic curve method to examine any pair (a, b) for a suitable smoothness of $a - mb$ and $f(a, b)$ in $\max(B, B')^{\mathbf{o}(1)}$ time. Hence we can determine whether a pair (a, b) is in $\mathcal{X}_{f,m,n}$ in time $L_n \left(\frac{1}{3}, \mathbf{o}(1)\right)$.

By lemma 5.15 the probability that f is reducible is $L_n\left(\frac{2}{3}\right)$, and we have an unconditional uniform bound $|\mathcal{X}_{f,m,n}| \leq L_n\left(\frac{1}{3}, 2\sigma\right)$. Hence from theorem 5.16 we deduce that

$$\mathbb{E}_{f,m}(|\mathcal{X}_{f,m,n}| \mid f \text{ irreducible}) \geq L_n\left(\frac{1}{3}, \tau + \mathbf{o}(1)\right).$$

By using lemma 4.45 we can use stochastic deepening to consider $|\mathcal{X}_{f,m,n}|$ as a random variable of (f, m) , with $K \leq L_n\left(\frac{1}{3}, 2\sigma - \tau\right)$. Hence for some $j \leq 1 + \lceil \log_2 K \rceil = \mathbf{O}\left(\log^{\frac{1}{3}} n (\log \log n)^{\frac{2}{3}}\right)$, we have

$$\mathbb{P}_{f,m}\left(|\mathcal{X}_{f,m,n}| \geq 2^j L_n\left(\frac{1}{3}, \tau + \log(1)\right)\right) > 2^{-j},$$

absorbing logarithmic terms where appropriate.

Now consider the following: To find a collection of pairs the algorithm iterates through each $i \in \{0, \dots, 1 + \lceil \log_2 K \rceil\}$, and for each i generates 2^i pairs (f, m) , and for each pair (f, m) generates $2^{-i} L_n\left(\frac{1}{3}, \max(\beta, \beta') + 2\sigma - \tau + \mathbf{o}(1)\right)$ pairs (a, b) and tests for smoothness of $a - mb$ and $f(a, b)$. So if $|\mathcal{X}_{f,m,n}| > 2^i L_n\left(\frac{1}{3}, \tau + \mathbf{o}(1)\right)$ then the algorithm finds $L_n\left(\frac{1}{3}, \max(\beta, \beta') + \mathbf{o}(1)\right)$ pairs as required. Furthermore if

$$\mathbb{P}_{f,m}\left(|\mathcal{X}_{f,m,n}| \geq 2^i L_n\left(\frac{1}{3}, \tau + \log(1)\right)\right) > 2^{-i},$$

then with constant probability at least one of the pairs (m, f) satisfies this condition.

Since the time taken for a single i is $L_n\left(\frac{1}{3}, \max(\beta, \beta')\right) + 2\sigma = \tau + \mathbf{o}(1)$ we can absorb the logarithmic number of iterations into the $\mathbf{o}(1)$ term. Hence, iterating it at most a logarithmic number of time reduces the probability of failure to $L_n\left(\frac{2}{3}, \kappa - \delta^{-1}\right)$. Hence the expected time taken to complete the algorithm is

$$L_n\left(\frac{1}{3}, \max(\beta, \beta') + 2\sigma - \tau + \mathbf{o}(1)\right).$$

□

5.3 Algebraic obstructions Vol. 2: The congruence

Let S be the set of (a, b) pairs such that $a - mb$ is B -smooth and $f(a, b)$ is B' -smooth. The previous section provides a way to find a sufficient amount of (a, b) pairs in a sufficiently short time, which we now will assume to be fulfilled. Identically to the General Number Field Sieve algorithm we now start on the next procedure: sieving for the congruence. In Chapter 3 we were interested in finding a subset S_i such that

$$f'(m)^2 \prod_{(a,b) \in S_i} a - bm \text{ is a square in } \mathbb{Z},$$

$$f'(\alpha)^2 \prod_{(a,b) \in S_i} a - b\alpha \text{ is a square in } \mathbb{Z}[\alpha].$$

In the RNFS there will be a similar approach, but it can be recalled that we used particularly broad observations to avoid having to work with the ring of algebraic integers $\mathcal{O}_{\mathbb{Q}[\alpha]}$ for which the structure might be unknown or especially difficult. In this section we will not only tackle the procedure to find these squares, but in particular how we can avoid similar algebraic obstructions as we saw in section 3.3.

One of the first things to observe is that, given a set S such that for every (a, b) -pair $a - bm$ is B -smooth is that there is no reason for us not to use the standard method from the general number field sieve to find a square. Given an element $z \in \mathbb{Z}$ with prime factorization $\prod_{i=1}^d p_i^{\eta_i}$ it is easy to check if it is a square by ensuring that $\forall i \in \{1, \dots, d\} : 2 \mid \eta_i$. So given the factorizations of $a - mb$ for $B + 1$ (a, b) -pairs we can find a dependent subset in the same way we did with the general number field sieve in Step 3.

Once again it would be optimal if we could approach the $\mathbb{Z}[\alpha]$ problem in the same way, but as we saw in Chapter 3 this requires some justification. This justification is only made more

complex as we are now dealing with a uniformly random polynomial. We start in much the same fashion as we did with the general number field sieve by defining a set of pairs (p, r) such that $f(r, 1) \equiv 0 \pmod p$ for prime p and $r \in \{1, \dots, p-1\}$. Recall that the pairs (p, r) such that p is prime, $0 < r < p$ coprime to p , $p \mid f(r, 1)$ is in direct correspondence with the first degree prime ideals $\mathfrak{P} \subseteq \mathcal{O}_{\mathbb{Q}(\alpha)}$ such that $\mathfrak{P} \mid \langle p \rangle$, $N(\mathfrak{P}) = p$. Then from proposition 3.8 we get that the following function is well defined:

$$e_{(p,r)}(a - b\alpha) = \begin{cases} \text{ord}_r(N(a - b\alpha)) & a - br \equiv 0 \pmod p \\ 0 & \textit{otherwise} \end{cases}.$$

From the proof of theorem 3.7 we know that this map is well defined from $\mathbb{Q}(\alpha)^\times \rightarrow \mathbb{Z}$. Hence, as we assumed that we have found sufficient (a, b) -pairs we can use the same sieving techniques as in Section 3 to find a subset S such that $P(\alpha) = \prod_{(a,b) \in S} (a - b\alpha) \in \mathbb{Z}[\alpha]$. Moreover we can again be sure that $e_{r,s}(P(\alpha)) \equiv 0 \pmod 2$. Now it rests to us to show that we can, again by the use of characters, guarantee with high probability that the element found by $P(\alpha)$ is a square in $\mathbb{Z}[\alpha]$. To give a three line summary of our technique: we will apply the pigeonhole principle to a set H , to show that for a randomized field with a stochastic collection of characters with large conductor the number of ways in which an element might appear square and not be square is limited.

5.3.1 Quadratic characters

Now define the set we eluded to above:

$$H = \{z \in \mathbb{Q}(\alpha)^\times : \forall s < r, e_{(p,r)}(z) \equiv 0 \pmod 2\} / \{z^2 : z \in \mathbb{Q}(\alpha)^\times\}.$$

Because we attempt to use the pigeon hole principle we start by considering the size of this set, in fact a near-immediate result is the bounded dimension of H as a vector space:

Proposition 5.17. H is an \mathbb{F}_2 vector space of dimension at most

$$(\delta\kappa + \mathbf{o}(1)) \log_2(n) + \frac{\delta^2\kappa}{2 \log(2)} \frac{(\log(n))^{\frac{4}{3}}}{(\log \log(n))^{\frac{1}{3}}}.$$

The proof of this follows, bar a few modifications, from the argument in [[17], theorem 6.7]. The full proof can be found in [[14], p. 129-131], but we will accept this as a fact. To see that H is a vector space we only need to observe that $\mathbb{Q}(\alpha)^\times$ is commutative, hence every element of $h \in H$ can be represented as a coset $[h_0] = h_0 \cdot \{z^2 : z \in \mathbb{Q}(\alpha)^\times\}$. This means that for any $h \in H$, h^2 is the identity since it is equivalent to $h_0^2 \cdot \{z^2 : z \in (\mathbb{Q}(\alpha)^\times)\}$ and $h_0 \in \mathbb{Q}(\alpha)^\times$.

The next step is to find a significant amount of quadratic characters. For this let $\mathfrak{P} \subset \mathcal{O}_{\mathbb{Q}(\alpha)}$ be a prime ideal, as $N(\mathfrak{P}) = p^k$ and $\mathcal{O}_{\mathbb{Q}(\alpha)}$ is a Dedekind Domain we have that

$$\mathcal{O}_{\mathbb{Q}(\alpha)}/\mathfrak{P} = \mathbb{F}_{p^k}.$$

Hence we can identify \mathfrak{P} with a degree k monic irreducible polynomial $p_{\mathfrak{P}}$ over \mathbb{F}_r . We can say more: Given an irreducible polynomial g of degree k over \mathbb{F}_p , if $\gcd(f, g) = 1$ then the quotient $\mathbb{Q}(\alpha)/\langle g(\alpha) \rangle$ sends everything to 0, hence $\langle g(\alpha) \rangle$ is not prime. So we may assume $\gcd(f, g) > 1$, and since g is irreducible over \mathbb{F}_p that means that g is one of the irreducible factors of $f \pmod{p}$. In what follows we will regularly shift between following three equivalent concepts:

1. Prime ideal $\mathfrak{P} \subseteq \mathcal{O}_{\mathbb{Q}(\alpha)}$.
2. The irreducible polynomial divisor $p_{\mathfrak{P}}$ of $f(x, 1) \pmod{p}$.
3. Pair (p, r) , p prime and r s.t. $p_{\mathfrak{P}}(r, 1) \equiv 0 \pmod{p}$ for $p_{\mathfrak{P}}$ defined over \mathbb{F}_{p^k} .

The following definition extends the Legendre symbol.

Definition 5.18. Let K be a field and let $f(x) \in K[x]$ be a polynomial. Let $p(x) \in K[x]$ be an

irreducible polynomial, $p(x) \nmid f(x)$. Then the polynomial Legendre symbol is defined as follows:

$$\left(\frac{f(x)}{p(x)}\right)_L = \begin{cases} 1, & \exists g(x) \in K[x] \text{ s.t. } f(x) \equiv g(x)^2 \pmod{p(x)} \\ -1, & \nexists g(x) \in K[x] \text{ s.t. } f(x) \equiv g(x)^2 \pmod{p(x)} \end{cases}.$$

This gives a very natural extension to a Dirichlet character, $\chi_{\mathfrak{P}}$, defined by the following:

$$\chi_{\mathfrak{P}} : \mathcal{O}_{\mathbb{Q}(\alpha)} \rightarrow \mathbb{F}_{p^k}, \quad \langle a - b\alpha \rangle \mapsto \left(\frac{a - bx}{p_{\mathfrak{P}}}\right)_L \quad (8)$$

All that remains now is finding \mathfrak{P} , for which we will seek to factorize $f(x, 1) \pmod{p}$ and look at the irreducible divisors. Given a set \mathcal{F} of these characters $\chi_{\mathfrak{P}} = \chi_{(p,r)}$, we can define the map $\Psi_{\mathcal{F}} : H \rightarrow \mathbb{F}_2^{|\mathcal{F}|}$ ⁶ sending x to the tuple $(\chi_{(p,r)} \mid \chi_{(p,r)} \in \mathcal{F})$. We will rely on the random production of such a set \mathcal{F} to show that this almost surely makes $\ker(\Psi_{\mathcal{F}})$ small.

Lemma 5.19. There is a sampleable distribution Υ for pairs (p, r) such that $\chi_{(p,r)}$ is a character as in 8, such that for all but $\log \log n$ of the $h \in H$, considering map as a map from H to \mathbb{F}_2 :

$$\mathbb{P}_{\Upsilon} (\chi_{(p,r)} = -1) \geq \frac{1 + \mathbf{o}(1)}{2}.$$

Sampling according to Υ takes at most $L_n(\frac{1}{3}, c)$ time for a to-define c . Furthermore, each character, $\chi_{(p,r)}$, can be evaluated in time at most $L_n(\frac{1}{3}, \frac{c}{2})$.

5.3.2 Proof of lemma 5.19

To start assume that we have the finite-degree tower of number fields $L \supset K \supset \mathbb{Q}$, where L/K is a Galois extension with Galois group $\text{Gal}(L/K)$. Recall that Δ_L, Δ_K are the discriminants of L and K respectively, and let $d_{L/K} = [L : K]$, $d_L = [L : \mathbb{Q}]$ and $d_K = [K : \mathbb{Q}]$. Let $\mathfrak{P} \subseteq K$ be a prime ideal and let $Q = \{\mathfrak{Q} \subseteq L \mid \mathfrak{Q} \text{ lies over } \mathfrak{P}\}$, then from proposition 2.30 we know that the Frobenius elements of all $\mathfrak{Q} \in Q$ are conjugate, which allows us to give the following definition:

⁶Implicitly we are assuming $\mathbb{F}_2 \cong C_2$ by the map $b \mapsto (-1)^b$

Definition 5.20. Let $L \supset K \supset \mathbb{Q}$ be a finite degree tower of number fields with L/K Galois. Let $\mathfrak{P} \subset K$ be a prime ideal which is unramified in L and let $\text{Frob}_{\mathfrak{P}}$ be the Frobenius element of \mathfrak{P} in $\text{Gal}(L/K)$. Then the Artin symbol, $\left[\frac{L/K}{\mathfrak{P}} \right]$, is defined as the conjugacy class of the Frobenius automorphisms of L/K corresponding to the prime ideals $\mathfrak{Q} \subset L$ dividing \mathfrak{P} :

$$\left[\frac{L/K}{\mathfrak{P}} \right] = \{ \text{Frob}_{\mathfrak{Q}} \in \text{Gal}(L/K) \mid \mathfrak{Q} \mid \mathfrak{P} \}.$$

Remark. It is a common abuse of notation to write $\text{Frob}_{\mathfrak{P}}$ for any of the elements in $\left[\frac{L/K}{\mathfrak{P}} \right]$ and just consider the element well defined up to conjugacy.

This allows us to define the set

$$\pi_C(x) = \left| \left\{ \mathfrak{P} \mid \mathfrak{P} \subset K \text{ prime, } N_{K/\mathbb{Q}}(\mathfrak{P}) < x, \left[\frac{L/K}{\mathfrak{P}} \right] \in C \right\} \right|,$$

where $C \subseteq \text{Gal}(L/K)$ s.t. for all $g \in \text{Gal}(L/K)$: $gCg^{-1} = C$. Now let $K = \mathbb{Q}(\alpha)$ and let $h \in \mathcal{O}_K$ be an element of minimal norm representing a non-trivial element of H . Let $L = K(\sqrt{h})$, then by definition $[L : K] = 2$, $d_K = d$, $d_L = 2d$ and $\text{Gal}(L/K) = C_2$. Since $\text{Gal}(L/K)$ is abelian, we have that $\left[\frac{L/K}{\mathfrak{P}} \right]$ only contains one element. Hence the value $\left[\frac{L/K}{\mathfrak{P}} \right]$ corresponds exactly to the action of $\chi_{\mathfrak{P}}$ on h .

As h has minimal norm it fulfills the Minkowski bound, 2.41, hence:

$$N_{K/\mathbb{Q}}(h) \leq c_{K/\mathbb{Q}} = \sqrt{\Delta_K} \left(\frac{4}{n} \right)^{\frac{d}{2}} \frac{d!}{d^d} = n^{\delta\kappa(1+\mathfrak{o}(1))}.$$

By construction the different is generated by $2h$, and so is an integral ideal. As the relative discriminant is the norm of the different we get that:

$$\Delta_{L/\mathbb{Q}} \leq N_{K/\mathbb{Q}}(2h)\Delta_{K/\mathbb{Q}}^2 \leq n^{\delta\kappa(5+\mathfrak{o}(1))}.$$

We now need an improved version of Chebotarev's Density Theorem which we alluded to before.

Theorem 5.21. Unconditional Effective Chebotarev Density Th. Let $L/K/\mathbb{Q}$ be a tower of number fields with L/K Galois with $G = \text{Gal}(L/K)$. Let $C \subseteq G$ be a union of conjugacy classes such that $gCg^{-1} = C$ for all $g \in G$. Let $|\bar{C}|$ be the total number of conjugacy classes in G . Lastly let $1 - \nu$ be a Siegel zero of ζ_L if it exists and 0 otherwise. Then there exists $c > 0$ such that if $\log x \geq 10d_L \log^2(\Delta_L)$:

$$\left| \pi_{G'}(x) - \frac{|C|}{|G|} \text{Li}(x) \right| \leq \frac{|C|}{|G|} \text{Li}(x^{1-\nu}) + \mathbf{O} \left(x^{|\bar{C}|} \exp(-c_1 \sqrt{\frac{\log x}{d_L}}) \right).$$

Using the unconditional effective Chebotarev Density Theorem on the degree 2 extension L/K , then for \mathfrak{P} chosen uniformly at random with $N(\mathfrak{P}) \leq x$:

$$\left| \mathbb{P}(\chi_{\mathfrak{P}}(h) = 1) - \frac{1}{2} \right| < x^{-\nu(1+o(1))} + \mathbf{O} \left(2 \log x \exp \left(-c_1 \sqrt{\frac{\log x}{d_L}} \right) \right),$$

where $1 - \nu$ is a possible Siegel zero of the Dedekind zeta function over L , ζ_L . Ensuring that $\mathbb{P}(\chi_{\mathfrak{P}}(h) = 1) = \frac{1}{2} + \log(1)$ requires us to insist that $\log(x) = \omega(d_L(\log \log x)^2)$ and, if ζ_L has a Siegel zero, $\log(x) = \omega(\nu^{-1})$.

Note that, for each non-trivial coset, we can choose a representative $h \in \mathcal{O}_K$ of minimal norm and then run through the previous paragraph for each of these. Despite the element h being different every time we can get a similar result each time, dependent on h . Hence we make the following definition:

Definition 5.22. For a field K and h a minimal norm representative of an element of H , we define $L_h = K(\sqrt{h})$. For $\epsilon > 0$ we define the exceptional set:

$$E_{K,\epsilon} = \left\{ h \cdot \{z^2 \mid z \in K^\times\} \in H \text{ s.t. } \exists \nu : \zeta_{L_h}(1 - \nu) = 0, \nu^{-1} > L_n \left(\frac{1}{3}, \epsilon \right) \right\}.$$

It is clear that if $h \in K^\times$ and $k \in K^\times$ represent the same coset in H , then $L_h = L_k$. To justify this we only have to note that representing the same coset means that h and k differ by

a square:

$$h = kz^2,$$

hence $\sqrt{h} = \sqrt{kz^2} = \sqrt{k}z$, since $z \in K^\times$ we get that $\text{irr}_K(\sqrt{h}) = \text{irr}_K(\sqrt{k})$ and so $L_h = L_k$.

The exceptional set $E_{K,\epsilon}$ is the subset of H which cannot be reliably distinguished from 0 by characters induced by primes of size $e^{L_n(\frac{1}{3},\epsilon)}$. This means that if a Siegel zero of the form specified exists, then almost every prime ideal of this size induces a character which vanishes for some element of H . However we can limit the size of the exceptional set.

Lemma 5.23. Suppose that $K = \mathbb{Q}(\alpha)$ is a number field where α is a root of an irreducible $f = \hat{f} + (x - m)R$ where R is uniformly random. Then for $\epsilon = (\frac{1}{3} + \mathbf{o}(1)) \delta$,

$$\mathbb{P}_f \left(|E_{K,\epsilon}| > \frac{4}{4} \log \log n \right) \leq L_n \left(\frac{2}{3}, \frac{\kappa - \delta^{-1}}{3} (1 + \mathbf{o}(1)) \right)^{-1}.$$

For now we accept this lemma as fact, as this will allow us to prove lemma 5.19. The proof for this lemma is based on the sparseness of Siegel zeros for Dedekind zeta functions defined over L_h and will be shown in Section 5.3.3. Accepting the lemma, we know there exists an x such that $\log(x) = \omega(d_L(\log \log x)^2)$ and, if there is a Siegel zero for ζ_{L_h} , $\log(x) = \omega(\nu^{-1})$ for all but $\frac{4}{3} \log \log n$ of the $h \in H$. Moreover for all but at most $L_n \left(\frac{2}{3}\right)^{-1}$ of our maximum $L_n \left(\frac{1}{3}\right)$ polynomials f we have $\log x < L_n \left(\frac{1}{3}, \epsilon\right)$. Choosing to accept failure on all these polynomials still guarantees that the algorithm fails with probability $\mathbf{o}(1)$.

Now we turn to the sampleable distribution Υ . Recall that any prime for which $N(\mathfrak{P}) < x$ is guaranteed to divide a prime p with $p < x$. Moreover, if \mathfrak{P} is of degree k , then $p < \sqrt[k]{x}$. Equivalently, a k th degree prime ideal dividing p corresponds to a simple k 'th degree divisor of $f \pmod{p}$. The following algorithm samples Υ in such a way that it outputs ideals of which all but a small fraction are prime.

Algorithm 4 Sampling Υ for (prime) ideals.

Input: A polynomial f of degree d **Output:** A pair (p, r) Uniformly at random sample $k \in \mathbb{Z}/d\mathbb{Z}$ Uniformly at random sample $p \in \left(x^{(k+1)^{-1}}, x^{k^{-1}}\right)$ Test for primality of p with the Miller–Rabin test**if** r is prime **then**factor $f \bmod p$ **and** find the irreducible and unrepeated factors r_i of degree $< k$ **if** $|\{r_i\}| \geq j$ **then**Uniformly at random sample one r_i **and** define $r_i = r$ **else** Sample a new k and repeat the algorithm**end if****end if**

Remark. We will now make a few remarks regarding the runtime of the algorithm. First note that to obtain the required runtime of the sieve we need this algorithm to have runtime at most $\mathbf{O}(\log^4(n))$. Let's see how we obtain this:

- First note that the runtime-bound immediately means $\log x = L\left(\frac{1}{3}\right)$. This means we can not use deterministic primality tests, such as the AKS primality test we used for the GNFS. Especially we note that discarding a composite p using the Miller–Rabin primality test takes $\mathbf{O}(\log^3(x) \log \log x)$.
- The Miller–Rabin primality test discards composite p with probability $1 - \mathbf{O}(\log^{-1}(x))$ and so p is prime with probability $\mathbf{\Omega}((\log x)^{-1})$ and so any p produced here is prime with probability $1 - \mathbf{O}(d(\log x)^{-1})$.
- Factoring $f \bmod p$ is possible in time $\mathbf{O}((d \log x)^3)$ probabilistically.
- If p is composite, but slips through the cracks of the Miller–Rabin primality test, then the factorisation of $f \bmod p$ may fail. If it does not fail, then this p still induced a quadratic

character and therefore vanishes on the squares. Since we obtain at most one character from each p and are guaranteed to find a character if $k = d$ and p is prime, the fraction of characters which are not induced by primes is $\mathfrak{o}(1)$.

To finish the proof we need to show that this algorithm is sufficiently fast and that characters can be evaluated quickly. Moreover we, once again, have to show that these characters are sufficiently uniformly distributed that the bounds on $\mathbb{P}(\chi_{\mathfrak{P}} = 1)$ hold.

Proposition 5.24. The expected time taken to sample $(p, r) \sim \Upsilon$ as above is at most

$$L_n \left(\frac{1}{3}, (4 + \mathfrak{o}(1)) \epsilon \right).$$

Proof. We noted in our remark that each attempt from the start of the algorithm takes time $\mathfrak{O}((d \log x)^3)$ as the factorization of $f \pmod r$ is slowest. We are guaranteed to find a factor if our degree bound k is d , which happens with probability $\frac{1}{d}$, the integer r is prime with probability $\Omega(\frac{1}{\log x})$, and we successfully take an ideal in the final step with probability $\Omega(\frac{1}{d})$ if $k = d$. Hence the number of attempts needed to output a prime is bounded in expectation by $\mathfrak{O}(d^2 \log x)$. Hence the time taken to find an ideal is bounded in expectation by $\mathfrak{O}(d^5 \log^4(x)) = L_n(\frac{1}{3}, (4 + \mathfrak{o}(1))\epsilon)$. \square

Proposition 5.25. For any fixed h , $\mathbb{P}_{\Upsilon}(\chi_{(p,r)}(h) = -1) \geq \frac{1}{2d}(1 + \mathfrak{o}(1))$.

Proof. The distribution of primes \mathfrak{P} generated is uniform over $\mathfrak{P} \pmod \langle p \rangle$ for $p \in (x^{(k+1)^{-1}}, x^{k^{-1}})$ of degree at most k by the Prime Ideal Theorem, 4.22. This property also holds for a uniform distribution over primes of norm $\leq x$ by an adaptation of Dirichlet's theorem of primes in arithmetic progressions, 4.28. Thus the difference between Υ and a uniform distribution over primes of norm $\leq x$ is the distribution of the degree of these primes. The probability that Υ samples \mathfrak{P} with $N(\mathfrak{P}) \leq x$ and $\mathfrak{P} \mid \langle p \rangle$ for p is each of these intervals is $\frac{1}{d}$. Hence Υ pointwise dominates d^{-1} times the uniform distribution over all primes of norm below x . Therefore

$$\mathbb{P}_{\Upsilon}(\chi_{(p,r)}(h) = -1) \geq \frac{1}{d} \mathbb{P}_{N(\mathfrak{P}) \leq x}(\chi_{\mathfrak{P}}(h) = -1) = \frac{1}{2d}(1 + \mathfrak{o}(1)).$$

□

The final proposition immediately completes the proof of lemma 5.19, however it is quite technical as the logarithms we are interested in are quite large. We, therefore, have to track the arithmetic carefully.

Proposition 5.26. Evaluating the characters $\chi_{(p,r)}$ associated with the ideal \mathfrak{P} sampled as above on a term $a - b\alpha$ takes time at most $L_n \left(\frac{1}{3}, (2 + \mathbf{o}(1)) \epsilon\right)$.

Proof. Let $p = 2$, then $\chi_{(p,r)} = 1$. Hence we may assume $p > 2$. For any polynomial $P \in \mathbb{F}_p[x]$ let $|P| = p^{\deg(P)}$. Recall

$$\chi_{(p,r)}(a - b\alpha) = \left(\frac{a - bx}{r(x)} \right)_L.$$

Any constant c can reduce the computation to finding a Legendre symbol mod p :

$$\left(\frac{c}{P} \right)_L = c^{\frac{|P|-1}{2}} = \left(\frac{c}{p} \right)^{\frac{p^k-1}{r-1}} = \left(\frac{c}{p} \right)^k.$$

From [7] we call attention to the law of quadratic reciprocity for function fields by noting that for any two relatively prime monic irreducible polynomials over \mathbb{F}_p :

$$\left(\frac{P}{Q} \right) \left(\frac{Q}{P} \right) = (-1)^{\frac{|P|-1}{2} \frac{|Q|-1}{2}}.$$

Hence

$$\begin{aligned} \chi_{(p,r)}(a - b\alpha) &= \left(\frac{-b}{p} \right)^k \left(\frac{x - ab^{-1}}{r(x)} \right)_L = \left(\frac{-b}{p} \right)^k (-1)^{\frac{p-1}{2} \frac{p^{\deg(r)}-1}{2}} \left(\frac{r(x)}{x - ab^{-1}} \right) \\ &= \left(\frac{-b}{p} \right)^k (-1)^{\frac{p-1}{2} \frac{p^{\deg(r)}-1}{2}} \left(\frac{r(ab^{-1})}{x - ab^{-1}} \right) = \left(\frac{-b}{p} \right)^k (-1)^{\frac{p-1}{2} \frac{p^{\deg(r)}-1}{2}} \left(\frac{r(ab^{-1})}{p} \right). \end{aligned}$$

The parities of $\frac{p-1}{2}$ and $\frac{p^k-1}{2}$ are easily computed. Hence to compute $\chi_{(p,r)}(a - b\alpha)$ it suffices to compute $r(ab^{-1})$ and two legendre symbols modulo p in $\mathbf{O}(\log p)$ additions or subtractions of numbers of size at most p .

To compute $b^{-1} \pmod p$ requires the extended Euclidean algorithm to be run, which requires

$\mathbf{O}(\log p)$ additions of numbers of size at most p .

To compute $ab^{-1} \pmod p$ requires one multiplication.

To compute $s(ab^{-1}) \pmod p$ requires at most $\mathbf{O}(d)$ addition and multiplications modulo p .

Addition or subtraction of numbers of size p or modulo p takes $\mathbf{O}(\log p)$ steps, while multiplication modulo p takes $\mathbf{O}(\log^2(p))$ steps by iterative addition and doubling. Hence the computation time in total requires time

$$\mathbf{O}(d \log^2(p)) = \mathbf{O}(d^{-1} \log^2(x)) = L_n \left(\frac{1}{3}, (2 + \mathbf{o}(1)) \epsilon \right).$$

□

To complete the proof of the lemma we now simply take $c = (4 + \mathbf{o}(1)) \epsilon$.

5.3.3 Proof of lemma 5.23

As we remarked earlier this proof is dependent on the sparseness of Siegel zeroes over Dedekind zeta functions. To start this we need the following proposition, [[14], fact 6.22 and corollary 6.23], which we will state without proof.

Proposition 5.27. Let K be a number field and let \overline{K} be the normal closure of K . Let $c(K) = 4$ if K/\mathbb{Q} is normal and $c(K) = 4[\overline{K} : \mathbb{Q}]$ otherwise. Assume that there is a real Siegel zero of ζ_K , denoted $1 - \nu$, such that

$$1 - (c(K) \log |\Delta_K|)^{-1} \leq 1 - \nu \leq 1.$$

Then there is a quadratic field $F \subset K$ such that $\zeta_F(1 - \nu) = 0$.

Now we start the proof of lemma 5.23. To do this we may assume that $f(x) \in \mathbb{Q}(\alpha)[x]$ is an irreducible polynomial, as the probability that f is reducible can be absorbed in the error term by 5.15

Proposition 5.28. If $\overline{L}_h \supset L_h$ is the normal closure of $L_h = K(\sqrt{h})$, then $[\overline{L}_h : \mathbb{Q}] \leq 2^d d!$.

Proof. Let \overline{K} be the splitting field of K . By construction, \overline{K}/\mathbb{Q} is normal, especially that means it is Galois, and $[\overline{K} : \mathbb{Q}] \leq d!$. Let $\overline{G} = \text{Gal}(\overline{K}/\mathbb{Q})$. Given $h \in \mathcal{O}_K$, let \mathcal{O}_h be the orbit of h under the action of \overline{G} . Then $|\mathcal{O}_h| \leq d$. We adjoin square roots of each element of \mathcal{O}_h to \overline{K} to obtain a field \overline{L} . Then $[\overline{L} : \mathbb{Q}] \leq 2^d d!$.

Since degree 2 extensions are normal, the compositum of normal extensions is normal, and $\overline{L}/\overline{K}$ is a compositum of at most d degree 2 extensions, hence the extension $\overline{L}/\overline{K}$ and \overline{K}/\mathbb{Q} are normal. As any $\sigma \in \text{Aut}_{\mathbb{Q}}(\overline{K})$ acts on \mathcal{O}_h as a permutation we can extend σ to an element of $\text{Aut}_{\mathbb{Q}}(\overline{L})$. So \overline{L}/\mathbb{Q} is normal and $\overline{L}_h \subseteq \overline{L}$, concluding our hypothesis. \square

Hence from proposition 5.27, if $\nu^{-1} > 2^{d+2} d! \log \Delta_{L_h/\mathbb{Q}}$, then $1 - \nu$ must be a zero of some quadratic field $F_h = \mathbb{Q}(s_h) \subseteq L_h$. By assumption $2 \nmid [K : \mathbb{Q}]$, hence K has no quadratic subfields, which means that $F_h \cap K = \mathbb{Q}$ and so F_h is the only quadratic subfield of L_h . Moreover, L_h is the smallest field containing both F_h and K and since the $h \in H$ are not related by squares of K it holds that L_h doesn't contain any $h' \in H$ from another class. This assures that the L_h produced as h varies are all distinct fields and so all s_h must be distinct.

By the transitivity of the discriminant, lemma 2.47, we have, for the towers of number field $L_h/F_h/\mathbb{Q}$ and $L_h/K/\mathbb{Q}$ that

$$\Delta_{F_h/\mathbb{Q}}^d N_{F_h/\mathbb{Q}}(\Delta_{L_h/F_h}) = \Delta_{K/\mathbb{Q}}^2 N_{K/\mathbb{Q}}(\Delta_{L_h/K}).$$

As $\Delta_{L_h/K}$ is the relative discriminant, hence is an ideal, we can use the Minkowski Bound:

$$N_{K/\mathbb{Q}}(\Delta_{L_h/K}) \leq \sqrt{\Delta_{K/\mathbb{Q}}} \left(\frac{4}{\pi}\right)^d \left(\frac{d!}{d^d}\right) \leq \sqrt{\Delta_{K/\mathbb{Q}}}.$$

Since $\Delta_{K/\mathbb{Q}} \leq L_n \left(\frac{4}{3}, \frac{1}{2}\delta^2\kappa\right)$ and $\Delta_{L_h/\mathbb{Q}} \leq L_n \left(\frac{4}{3}, \frac{5}{4}\delta^2\kappa\right)$, we can conclude that

$$\Delta_{F_h/\mathbb{Q}} = \mathbf{O} \left(L_n \left(\frac{4}{3}, \frac{5}{4}\delta\kappa \right) \right).$$

By the contribution of the Euler product of the Dedekind zeta function:

$$\zeta_{F_h/\mathbb{Q}}(s) = \zeta(s) \mathcal{L}\left(s, \left(\frac{\Delta_{F_h/\mathbb{Q}}}{\cdot}\right)\right),$$

and by reciprocity $j \mapsto \left(\frac{\Delta_{F_h/\mathbb{Q}}}{j}\right)$ is a Dirichlet character modulo $\Delta_{F_h/\mathbb{Q}}$.

Now we invoke [[14], fact 6.24]:

Proposition 5.29. There is an effective constant $c > 0$ such that given two characters $\chi_r, \chi_{r'}$ of moduli r, r' respectively, with $\chi_r \chi_{r'}$ non-principal, then the product of Dirichlet \mathcal{L} -functions $\mathcal{L}(s, \chi_r) \mathcal{L}(s, \chi_{r'})$ has at most one real zero in $\left(1 - \frac{c}{\log rr'}, 1\right)$.

So if there are two characters modulo q and q' respectively there is at most one \mathcal{L} -function with a zero $1 - \nu$ and $\nu^{-1} > c \log qq'$ for some effective c . It immediately follows that there is at most one character with modulus in $[q, q^e]$ with a zero at $1 - \nu$ and $\nu^{-1} > (e + 1)c \log q$.

Since $\Delta_{F_h/\mathbb{Q}} < L_n \left(\frac{4}{3}, \left(\frac{5}{4} + \mathbf{o}(1)\right) \delta \kappa\right)$ and $\Delta_{F_h/\mathbb{Q}} \in \mathbb{Z}$ it follows that all possible discriminants can be covered by $\frac{4}{3} \log \log n$ ranges of form $[x, x^e]$. Hence there are at most $\frac{4}{3} \log \log n$ exceptional characters with exceptional zeroes such that

$$\nu^{-1} > (e + 1)c \log \left(L_n \left(\frac{4}{3}, \left(\frac{5}{4} + \mathbf{o}(1) \right) \delta \kappa \right) \right) = \mathbf{O} \left(\delta \kappa \log^{\frac{4}{3}}(n) \log \log(n)^{-\frac{1}{3}} \right),$$

as required. This bound is far weaker than the required bound of $\nu^{-1} > 2^{d+2} d! \log \Delta_{L_h/\mathbb{Q}}$, and so there are at most $\frac{4}{3} \log \log n$ extensions L_h/\mathbb{Q} with exceptional zeroes and $\nu^{-1} > 2^{d+2} d! \log \Delta_{L_h/\mathbb{Q}}$.

Finally:

$$2^{d+2} d! \log \Delta_{L_h/\mathbb{Q}} \leq d^{d(1+\mathbf{o}(1))} \log^{\mathbf{O}(1)}(n) = L_n \left(\frac{1}{3}, \frac{\delta}{3} (1 + \mathbf{o}(1)) \right).$$

This completes the proof of the lemma.

5.3.4 Proof of theorem 5.3

We have now removed all obstructions that were needed. It is therefore that we can now prove theorem 5.3. Recall that for a set of characters \mathcal{F} we define the map

$$\Psi_{\mathcal{F}} : H \rightarrow \mathbb{F}_2^{|\mathcal{F}|},$$

$$x \mapsto (\chi_{(p,r)}(x) : \chi_{(p,r)} \in \mathcal{F}).$$

With the work we have done up to now we are ready to produce our linear $\Psi_{\mathcal{F}}$ with small kernel, and this to produce a congruence of squares. Due to the size of some of the numbers involved, having $L_n(\frac{1}{3})$ bits, we track the arithmetic very closely. First, using lemma 5.19, we sample $4d \left(\delta\kappa \log n + \frac{\delta^2\kappa}{2\log 2} \frac{\log^{\frac{4}{3}}(n)}{\log \log n^{\frac{1}{3}}} \right)$ pairs (p_i, r_i) from Υ independently.

Note that this sample is of size $\mathbf{o}(\log^2(n)) = L_n(\frac{1}{3}, \mathbf{o}(1))$ and each individual sample takes at most $L_n(\frac{1}{3}, c)$, so we can produce the complete sample in $L_n(\frac{1}{3}, c + \mathbf{o}(1))$. After this process we have $M = 1 + B + dB' + 4d \left(\delta\kappa \log n + \frac{\delta^2\kappa}{2\log 2} \frac{\log^{\frac{4}{3}}(n)}{\log \log n^{\frac{1}{3}}} \right)$ pairs from Section 5.2.1 and the sample from Υ . Note that $M = L_n(\frac{1}{3}, \max(\beta, \beta'))^{1+\mathbf{o}(1)}$. For each of these, we need to evaluate each of our characters and as each character evaluates in $L_n(\frac{1}{3}, \frac{c}{2})$ we get that this process takes $L_n(\frac{1}{3}, \frac{c}{2} + \max(\beta, \beta'))^{1+\mathbf{o}(1)}$.

Fix some $h \in H \setminus \{1\}$ such that h is not in the exceptional set of size $\log \log n$. Each map $\chi_{(p_i, r_i)}$ is independent and induces a map in $\text{Hom}(H, \mathbb{F}_2)$ such that

$$\mathbb{P}(h \notin \ker(\chi_{(p_i, r_i)})) \geq \frac{1 + \mathbf{o}(1)}{2d}.$$

As a corollary it follows that:

$$\mathbb{P}(h \in \ker(\Psi_{\mathcal{F}})) \leq \left(1 - \frac{1 + \mathbf{o}(1)}{2d} \right)^{M - (1+B+dB')} \leq |H|^{-2+\mathbf{o}(1)}.$$

Hence by the union bound over the non-trivial elements of H the probability that any of these non-exceptional and non-zero elements is in the kernel is $\mathbf{o}(1)$. Hence with high probability the kernel of $\Psi_{\mathcal{F}}$ has size at most $\frac{4}{3} \log \log n$.

From here we proceed as normal. With the M pairs representing linear polynomials we use a fast kernel finding algorithm for sparse matrices, such as Block–Wiedemann, [21], to find a suitable subset S_i to construct a polynomial P_i in time

$$\mathbf{O}(M^2) = L_n \left(\frac{1}{3}, 2 \max(\beta, \beta')(1 + \mathbf{o}(1)) \right),$$

such that $P_i(m)$ is a square in \mathbb{Z} and $P_i(\alpha)$ is a square in $\mathcal{O}_{\mathbb{Q}(\alpha)}$ multiplied by one of at most $\frac{4}{3} \log \log n$ elements of H . Repeating the algorithm $l = \frac{4}{3} \log \log n$ times to generate polynomials P_1, \dots, P_l , we are able to guarantee that for some $i < j$, P_i and P_j lie over the same element h , hence $P_i P_j$ is a square in $\mathcal{O}_{\mathbb{Q}(\alpha)}$. In what follows we consider the $\binom{l}{2} \sim \frac{8}{9} (\log \log n)^2$ polynomials separately.

Now if $\gamma \in \mathcal{O}_{\mathbb{Q}(\alpha)}$ and $\gamma^2 \in \mathbb{Z}[\alpha]$, then $\gamma \cdot f'(\alpha) \in \mathbb{Z}[\alpha]$. Let $\mathcal{S} = S_i \Delta S_j$ and fix the polynomial

$$P = \left[\frac{\partial f}{\partial x}(x, 1) \right]^2 \prod_{(a,b) \in \mathcal{S}} (a - bx), \text{ and so } u^2 = \left[\frac{\partial f}{\partial x}(x, y) \right] (m, 1)^2 \prod_{a,b \in \mathcal{S}} (a - mb)$$

is a square in \mathbb{Z} . Hence u can be found by taking the product modulo n over all $p < B$ of p raised to half the total order of p in the terms $(a - mb)$ for $(a, b) \in \mathcal{S}$ and multiplying by $f'(m, 1)$.⁷ The reason for computing the square root in this fashion is to ensure we need only $M \log n$ additions and divisions and at most $M \log n$ modular multiplications to compute $u \pmod n$ from the exponents. This ensures polynomial running time.

⁷This technique differs minimally from the technique for \mathbb{Z} in the GNFS

Similarly, for at least one of the $\binom{l}{2}$ polynomials considered, there exists $v \in \mathbb{Z}[\alpha]$ such that

$$v^2 = \left[\frac{\partial f}{\partial x}(x, y) \right] (\alpha, 1)^2 \prod_{(a,b) \in \mathcal{S}} (a - b\alpha).$$

Using [24] we can compute square roots in the number field to find $v(m, 1) \pmod n$ in time $\mathbf{O}(M^2)$. We abuse notation and write $v(m)$ for the element of $\mathbb{Z}/n\mathbb{Z}$ obtained by substituting m for α . By the definition of f we have $f(m, 1) = n$ and so:

$$\begin{aligned} v(m)^2 \pmod{f(m, 1)} &= \left(\left[\frac{\partial f}{\partial x}(x, y) \right] (\alpha, 1)^2 \prod_{(a,b) \in \mathcal{S}} (a - b\alpha) \pmod{f(\alpha, 1)} \right) (m) \\ &= \left[\frac{\partial f}{\partial x}(x, y) \right] (m, 1)^2 \prod_{(a,b) \in \mathcal{S}} (a - mb) \pmod{(f(m, 1))} \\ &= u^2 \pmod n, \end{aligned}$$

and so we have constructed a congruence of squares in time

$$L_n \left(\frac{1}{3}, \max(2 \max(\beta, \beta'), \max(\beta, \beta') + \frac{c}{2}, c) \right)^{1+\mathbf{o}(1)}.$$

As $c \leq (\frac{4}{3} + \mathbf{o}(1)) \delta$, we can insist that we have at most $\frac{4}{3} \log \log n$ exceptional values of h , and our f lies off a set of probability at most $L_n \left(\frac{2}{3}, \frac{\kappa - \delta^{-1}}{3} (1 + \mathbf{o}(1)) \right)^{-1}$ the run time bound is as claimed, finishing the proof.

5.4 Computational efficiency

The computational efficiency boils down to proving theorem 5.1. For this we fix $n, \beta, \beta', \sigma, \delta$, and κ satisfying the conditions of equations 5 and 6. Then by theorem 5.3 we can extract a congruence of squares $\pmod n$ from $L_n \left(\frac{1}{3}, \max(\beta, \beta') + \mathbf{o}(1) \right)$ pairs $(a, b) \in \mathcal{X}_{f, m, n}$ for a fixed (m, f) in expected time

$$L_n \left(\frac{1}{3}, 2 \max \left(\frac{2\delta}{3}, \beta, \beta' \right) (1 + \mathbf{o}(1)) \right).$$

Theorem 5.2 tells us that a fixed (m, f) and this many pairs $(a, b) \in \mathcal{X}_{f,m,n}$ will be found in expected time

$$L_n \left(\frac{1}{3}, \max(\beta, \beta') + \frac{\delta^{-1}}{3\beta}(1 + \mathbf{o}(1)) + \frac{\sigma\delta + \kappa}{3\beta'}(1 + \mathbf{o}(1)) \right).$$

Hence we can run the RNFS to obtain a congruence of squares mod n with the expected time bounded by

$$L_n \left(\frac{1}{3}, \lambda(1 + \mathbf{o}(1)) \right), \lambda = \max \left(2 \max \left(\frac{2\delta}{3}, \beta, \beta' \right), \max(\beta, \beta') + \left(\frac{\delta^{-1}}{3\beta} + \frac{\kappa + \sigma\delta}{3\beta'} \right) \right).$$

Having chosen $\beta, \beta', \sigma, \delta, \kappa$ satisfying the conditions of 5 and 6 we can optimize the constants. Note that increasing the lesser of β and β' cannot increase λ or cause the conditions on the constants to be violated, so we can assume $\beta = \beta'$. We then compute:

$$2\sigma \geq \lambda \geq \min_{\beta, \delta} \left(\beta + \frac{2\delta^{-1} + \sigma\delta + \mathbf{o}(1)}{3\beta} \right) \geq \min_{\beta} \left(\beta + \frac{\sqrt{8\sigma} + \mathbf{o}(1)}{3\beta} \right) \geq 2\sqrt[4]{\frac{8\sigma}{9}} + \mathbf{o}(1).$$

Fix any $\epsilon > 0$ such that $\epsilon = \mathbf{o}(1)$. If we take $\beta = \beta' = \sigma = \frac{2\delta}{3} = \sqrt[3]{\frac{8}{9}} + \epsilon$ and $\kappa = \sqrt[3]{\frac{1}{3}} + \epsilon$ then the above are all equalities. Furthermore all the conditions are satisfied, giving $\lambda = 2\sqrt[3]{\frac{8}{9}} + \mathbf{o}(1)$.

This proves a runtime of

$$L_n \left(\frac{1}{3}, \sqrt[3]{\frac{64}{9}} + \mathbf{o}(1) \right).$$

5.5 Riemann and the future of factorization

As a final note we want to have a short discussion on some of the ways that the Generalized Riemann Hypothesis could impact the analysis performed on the Randomized number field sieve. As discussed: Until now we have not made any assumption regarding the GRH and only limited assumptions on the heuristics, but if we were to accept the GRH then much of our discussion becomes a lot simpler.

One of the main reasons it becomes simpler is because the Generalized Riemann Hypothesis

allows us to assume that there are no Siegel zeroes. This makes the discussion around algebraic obstructions a lot simpler. For example we obtain the following:

Proposition 5.30. Conditional on GRH, for $\epsilon = \log^{-\frac{1}{4}}(n) = \mathbf{o}(\delta)$,

$$\mathbb{P}_f (|E_{K,\epsilon}| > 0) = 0.$$

This follows automatically from our discussion of the size of this set and the fact there are no Siegel zeroes. This makes lemma 5.23 vacuous and allows us to sharpen the $L_n\left(\frac{1}{3}\right)$ bounds in lemma 5.19 polynomial in $\log n$.

The GRH also allows us to have far tighter effective bounds on the prime numbers in arithmetic progressions. Without going into details it would allow us to subvert the use of the adapted Bombieri–Vinogradov theorems and use an effective version of Chebotarev’s Density Theorem which is completely dependent on the GRH holding.

It is obvious that in a discussion like this it could easily be assumed that the GRH holds, but that would make the whole duscussion conditional and the beauty of this is that we make no grand assumptions whatsoever. To obtain a complexity that is equivalent to the heuristic complexity without relying on any conditions makes this the best possible outcome.

Does this mean we are done with the number field sieve after thirty years? No, definitely not. Over the years many different versions have found their way into mathematics and cryptography to solve other problems than factorization. For example the Tower Number Field Sieve which is an adapted algorithm to solve Discrete Logarithm Problems ([27], [28]) has been worked on for many years and it will be interesting to see if a similar randomization leads to a provable complexity for that as well.

References

Books

- [1] S. Alaca, K.S. Williams, *Introductory Algebraic Number Theory* Cambridge University Press, First edition, 2004.
- [2] R. Crandall, C. Pomerance, *Prime numbers, a computational perspective* Springer Verlag, First edition - corrected printing, 2002.
- [3] A.K. Lenstra, H.W. Lenstra Jr., *Development of the Number Field Sieve* Springer Verlag, 1993. Which is the source used for the papers [15], [16] cited below
- [4] I. Stewart, *Galois Theory*, CRC Press, Fourth edition, 2015
- [5] H. Cohen, *A Course in Computational Algebraic Number Theory* Springer Verlag, Fourth Edition, 2000
- [6] L. C. Washington, *Introduction to Cyclotomic Fields* Springer Verlag, Second edition, 1997
- [7] M. Rosen, *Number Theory in Function Fields*, Springer Verlag, 2002.
- [8] H.L. Montgomery, R.C. Vaughan, *Multiplicative Number Theory 1: Classical Theory*, Cambridge University Press, 2010.
- [9] J. Milne, *Class Field Theory*, Self-published, Version 4.02, viewed January-March 2020, www.jmilne.org/math/CourseNotes/
- [10] H. Davenport, *Multiplicative Number Theory*, Springer Verlag, Second edition, 1980
- [11] G. Greaves, *Sieve in Number Theory* Springer Verlag, First Edition, 2001

Papers

- [12] M.A. Morrison, J. Brillhart, *A Method of Factoring and the Factorization of F_7* Mathematics of Computation, American Mathematical Society, Volume 29, 1975.

- [13] P. Stevenhagen, *The number field sieve* Algorithmic Number Theory, MSRI Publications, Volume 44, 2008.
- [14] J.D. Lee, R. Venkatesan, *Rigorous analysis of a randomised number field sieve* Journal of Number Theory, Volume 187, 2018
- [15] J.M. Pollard, *Factoring with cubic integers* As printed in [3], Springer Verlag, 1993.
- [16] A.K. Lenstra, H.W. Lenstra Jr., M.S. Manasse, J.M. Pollard *The Number Field Sieve* As printed in [3], Springer Verlag, 1993
- [17] J.P. Buhler, H.W. Lenstra Jr., C. Pomerance *Factoring integers with the number field sieve* As printed in [3], Springer Verlag, 1993
- [18] A. Granville, *It is easy to determine whether a given integer is prime* Bull. Amer. Math. Soc., Volume 42, 2005
- [19] D. Coppersmith, *Solving linear equations over $GF(2)$: Block-Lanczos Algorithm* , Linear Algebra Applications, 1991.
- [20] P. Montgomery, *A Block-Lanczos Algorithm for Finding Dependencies over $GF(2)$* Adv. in Cryptology - Eurocrypt '95, 1995
- [21] D. Wiedemann *Solving sparse linear equations over finite fields* , IEEE Trans. Inform. Theory, Volume 32, 1986
- [22] A. Hildebrand, G. Tenenbaum, *Integers without large prime factors* , Jour. Th. des Nom. de Bordeaux, Volume 5, 1993
- [23] A.J. Harper, *Bombieri-Vinogradov and Barban-Davenport-Halberstam type theores for smooth numbers* , arXiv:1208.5992 [math.NT]
- [24] E. Thomé, *Square root algorithms for the number field sieve* , Proc. of the 4th Int. Workshop on Arith. in Finite Fields, Springer, 2012

- [25] A. Hildebrand, *On the number of positive integers $\leq x$ and free of prime factors $> y$* , J.Number Theory, 1986, Volume 22, p. 289–307
- [26] A. Hildebrand, G. Tenenbaum, *On integers free of large prime factors* , Trans. Amer. Math. Soc., 1986, Volume 296, p. 265–290
- [27] O. Schirokauer, *The impact of the number field sieve on the Discrete Logarithm Problem* , Algorithmic Number Theory, MSRI Publications, 2008, Volume 44
- [28] R. Barbulescu, et. al, *The Tower Number Field Sieve* , Advances in Cryptology, ASI-ACRYPT 2015, 2015, Volume 9453
- [29] K. Soundararajan, *The Distribution of Smooth Numbers in Arithmetic Progressions* CRM Proceedings and Lecture notes, American Mathematical Society, 2008, Volume 46, p. 115-128
- [30] A. Harper, *On a paper of K. Soundararajan on smooth numbers in arithmetic progressions* , Journal of Number Theory, Elsevier, 2012, Volume 132, p. 182-199.