



RESEARCH ARTICLE

Unconditional correctness of recent quantum algorithms for factoring and computing discrete logarithms

Cédric Pilatte 

University of Oxford, United Kingdom; E-mail: cedric.pilatte@maths.ox.ac.uk.

Received: 8 May 2024; **Revised:** 8 October 2025; **Accepted:** 8 October 2025

2020 Mathematics Subject Classification: *Primary* – 11Y16; *Secondary* – 68Q12, 11M26, 11H06, 11L40

Abstract

In 1994, Shor introduced his famous quantum algorithm to factor integers and compute discrete logarithms in polynomial time. In 2023, Regev proposed a multidimensional version of Shor’s algorithm that requires far fewer quantum gates. His algorithm relies on a number-theoretic conjecture on the elements in $(\mathbb{Z}/N\mathbb{Z})^\times$ that can be written as short products of very small prime numbers. We prove a version of this conjecture using tools from analytic number theory such as zero-density estimates. As a result, we obtain an unconditional proof of correctness of this improved quantum algorithm and of subsequent variants.

Contents

1	Introduction	2
1.1	Context and quantum computing results	2
1.2	An overview of Regev’s algorithm	3
1.3	The number-theoretic conjecture behind Regev’s algorithm	4
1.4	Subgroup obstructions	5
1.5	Structure of the paper	6
1.6	Notation	7
2	An illustrative special case	7
3	Proof of the existence of a short lattice basis	10
3.1	Restricting to a convenient subgroup of $(\mathbb{Z}/N\mathbb{Z})^\times$	10
3.2	Lattice point counting via characters	12
3.3	Bounding the contribution of large-order characters	14
3.4	Geometry of numbers	18
3.5	Short basis vectors	19
4	Applications to quantum computing	21
4.1	Preparatory lemmas	21
4.2	The discrete logarithm problem	22
4.3	Factoring integers	23
A	Bounds for character sums over primes	24
	References	25

1. Introduction

1.1. Context and quantum computing results

Public key cryptography has become a crucial element of our global digital communication infrastructure. Notable examples include the Diffie-Hellman key exchange [6] and the RSA (Rivest-Shamir-Adleman) cryptosystem [19], which rely on the difficulty of finding discrete logarithms and factoring large numbers, respectively.

However, in 1994, Peter Shor [20] developed an algorithm capable of efficiently solving these problems using a quantum computer.

Theorem (Shor¹). *There is a quantum circuit having $O(n^2 \log n)$ quantum gates and $O(n \log n)$ qubits with the following property. There is a classical randomised polynomial-time algorithm that solves the factoring problem*

Input : *a composite integer $N \leq 2^n$*

Output : *a nontrivial divisor of N*

using $O(1)$ calls to this quantum circuit, and succeeds with probability $\Theta(1)$.

Shor's original article [20] also includes a similar algorithm to solve the discrete logarithm problem. This advancement prompted the development of postquantum cryptography, such as lattice-based cryptography, to ensure the security of our communications in the face of potential quantum computing breakthroughs [4].

Recently, Regev [18] devised a multidimensional variant of Shor's factoring algorithm that reduces the size of the quantum circuit, that is, the number of quantum gates, to $O(n^{3/2} \log n)$. A distinctive feature of Regev's algorithm is that the quantum circuit must be called $O(\sqrt{n})$ times, rather than a constant number of times for Shor's algorithm. This is not considered to be a serious drawback as the various complexity parameters of the quantum circuit are far more relevant metrics in quantum computing.

However, this remarkable work of Regev initially came with two main limitations.

First, the original version of Regev's algorithm requires $O(n^{3/2})$ qubits – many more than Shor's algorithm. The reason for this ultimately lies in the difficulty of performing classical computations on quantum computers in a space-efficient manner, due to the need for any quantum computation to be *reversible*. In a subsequent paper, Ragavan and Vaikuntanathan [17] improved Regev's algorithm to use only $O(n \log n)$ qubits, while maintaining the circuit size at $O(n^{3/2} \log n)$ quantum gates. This work, inspired by ideas of Kalinski [13], essentially matches the space cost of Shor's algorithm.

The second issue, which we address in this paper, is that Regev's algorithm [18] does not have theoretical guarantees, unlike Shor's algorithm. The correctness of Regev's algorithm is based on an *ad hoc* number-theoretic conjecture which we describe below. This unproven assumption cannot be avoided as it lies at the core of Regev's improvement on the circuit size. The variant of Ragavan and Vaikuntanathan [17] also crucially relies on this conjecture.

In this paper, we prove a version of Regev's conjecture. This allows us to unconditionally prove the correctness of (slightly modified versions of) the algorithms of Regev [18] and Ragavan and Vaikuntanathan [17]. More precisely, we obtain the following algorithmic result.

Theorem 1.1. *There is a quantum circuit having $O(n^{3/2} \log^3 n)$ quantum gates and $O(n \log^3 n)$ qubits with the following property. There is a classical randomised polynomial-time algorithm that solves the factoring problem*

¹The version of Shor's algorithm stated here incorporates some small improvements from [21] (bounded number of calls) and [10] (fast integer multiplication). The number of qubits in Shor's algorithm can be brought down to $O(n)$, though this typically requires a larger number of quantum gates (see, e.g., [8, 12]).

Input : a composite integer $N \leq 2^n$

Output : a nontrivial divisor of N

using $O(\sqrt{n})$ calls to this quantum circuit, and succeeds with probability $\Theta(1)$.

This unconditionally establishes the results in [17, 18] up to logarithmic factors.

Regev's algorithm was adapted by Ekerå and Gärtner [7] to the discrete logarithm problem. Their paper [7] also uses the space-saving arithmetic of Ragavan and Vaikuntanathan [17] to obtain a quantum circuit with $O(n^{3/2} \log n)$ gates and $O(n \log n)$ qubits for computing discrete logarithms.

Once more, the correctness of the algorithm by Ekerå and Gärtner [7] relies on an unproven hypothesis, which can be viewed as a stronger form of Regev's conjecture. Our methods also apply to this stronger statement (again, with minor technical adjustments). Thus, we get an analogue of Theorem 1.1 for the discrete logarithm problem, that is, an unconditional proof of the results in [7] up to logarithmic factors.

Theorem 1.2. *There is a quantum circuit having $O(n^{3/2} \log^3 n)$ quantum gates and $O(n \log^3 n)$ qubits with the following property. There is a classical randomised polynomial-time algorithm that solves the discrete logarithm problem*

Input : an integer $N \leq 2^n$ and elements $g, y \in (\mathbb{Z}/N\mathbb{Z})^\times$ such that $y \in \langle g \rangle$

Output : an integer x such that $g^x \equiv y \pmod{N}$

using $O(\sqrt{n})$ calls to this quantum circuit, and succeeds with probability $\Theta(1)$.

Moreover, the slight technical modifications that we need to introduce to make all these quantum algorithms unconditional are compatible with the error-correction results of [7, 17]. Finally, our results extend to further variants of Regev's algorithm, such as computing discrete logarithms or multiplicative orders modulo N for several elements simultaneously (see [7]).

Remark 1.3. By being slightly more careful in the space usage of our quantum algorithms, it is possible to reduce the number of qubits to $O(n \log^2 n)$ for Theorems 1.1 and 1.2. We will not show this here in order to keep the quantum computing prerequisites to a minimum.

1.2. An overview of Regev's algorithm

Let us start by recalling the basic idea behind Shor's factoring algorithm [20].

Let N be the integer to be factored, which can be assumed to be odd and with at least two distinct prime factors. Shor proved that there is an efficient quantum algorithm to find the multiplicative order of any invertible element modulo N . With this in hand, factoring N is straightforward. The first step is to pick an integer $1 < a < N$ uniformly at random. If $\gcd(a, N) > 1$, the factoring task is trivial. Otherwise, with probability at least $1/2$, the order of a modulo N , denoted r , is even and satisfies $a^{r/2} \not\equiv \pm 1$ (i.e., $a^{r/2}$ is a nontrivial square root of 1 modulo N). Whenever this is the case, we get a nontrivial divisor of N , namely $\gcd(N, a^{r/2} - 1)$.

The reason that Shor's algorithm uses $O(n^2 \log n)$ quantum gates, where $n := \lceil \log_2 N \rceil$, comes from the part of the quantum circuit that performs modular exponentiation. Consider the task of computing a power $a^M \pmod{N}$ on a classical computer, where $1 < a < N - 1$ and $M \leq N$. This computation can be performed efficiently using the well-known square-and-multiply method, which involves $O(\log M)$ multiplications of two integers modulo N . Since two n -bit integers can be multiplied in time $O(n \log n)$ by [10], the complexity of this modular exponentiation problem is thus $O(n^2 \log n)$.² For Shor's algorithm, a similar modular exponentiation needs to be performed quantumly, which requires $O(n^2 \log n)$ quantum gates.

²Note that, even if a is a small integer, say $a = 2$, this procedure still takes time $O(n^2 \log n)$, because most multiplications will involve n -bit integers.

Regev's improvement of Shor's algorithm is made possible by combining two key ideas.

The first idea in Regev's algorithm is to work in *higher-dimensional* space and replace the random parameter a by several integers $b_1, \dots, b_d$ (chosen in a specific way, as we explain below). Eventually, the optimal choice of dimension turns out to be $d \asymp \sqrt{\log N}$. Similarly to Shor's algorithm, the problem of factoring N easily reduces to the task of finding a vector $(e_1, \dots, e_d) \in \mathbb{Z}^d$ such that $\prod_{i=1}^d b_i^{e_i}$ is a nontrivial square root of 1 modulo N .

Using additional tools such as the LLL lattice reduction algorithm, Regev [18] generalised Shor's quantum algorithm to efficiently find all such vectors $(e_1, \dots, e_d)$ in a ball of radius $N^{O(1/d)}$ centred at the origin.

Regev's algorithm only succeeds if there indeed exists a vector $v = (e_1, \dots, e_d) \in \mathbb{Z}^d$ such that

- (i) $\|v\|_2 \leq N^{O(1/d)}$, and
- (ii) $\prod_{i=1}^d b_i^{e_i}$ is a nontrivial square root of 1 modulo N .

If the parameters $b_1, \dots, b_d$ are 'sufficiently multiplicatively independent' modulo N , one might heuristically expect the existence of a vector $v = (e_1, \dots, e_d)$ satisfying Item i and Item ii. This is, roughly speaking, what Regev needs to assume. To state his conjecture properly, it remains to specify how the parameters $b_1, \dots, b_d$ are chosen. This is a crucial point, as the idea of working in dimension d does not, on its own, offer any advantage over Shor's algorithm.

The second key insight of Regev is to choose $b_1, \dots, b_d$ to be *very small* compared to N . We mentioned that the most costly part of Shor's quantum circuit lies in the modular exponentiation step. Similarly, the number of gates of Regev's circuit is dominated by the cost of computing an expression of the form

$$\prod_{i=1}^d b_i^{M_i} \pmod{N} \quad (1)$$

in the quantum setting, where the exponents $M_1, \dots, M_d$ are $\leq N^{O(1/d)}$. Regev observed that (1) can be computed classically in time $O(n^{3/2} \log n)$ if $|b_i| \leq (\log N)^{O(1)}$ for all i (for $d \asymp \sqrt{n} \asymp \sqrt{\log N}$). This can be achieved by cleverly ordering the intermediate multiplications so that most of them involve integers with much fewer than n bits. This classical procedure can then be turned into a quantum version that uses $O(n^{3/2} \log n)$ gates.

Shor's algorithm corresponds to the $d = 1$ case of Regev's algorithm where the parameter b_1 is an element of $(\mathbb{Z}/N\mathbb{Z})^\times$ chosen uniformly at random. In this case, simple considerations from elementary number theory immediately imply the existence of an integer e_1 satisfying Item i and Item ii with probability $\gg 1$. This would hold more generally for $d \geq 1$ if $b_1, \dots, b_d$ were independent, uniformly distributed random elements of $(\mathbb{Z}/N\mathbb{Z})^\times$. However, for Regev's algorithm, the b_i 's are far from uniformly distributed in $(\mathbb{Z}/N\mathbb{Z})^\times$ as they are constrained to lie in a very small subset of $(\mathbb{Z}/N\mathbb{Z})^\times$.

1.3. The number-theoretic conjecture behind Regev's algorithm

Regev's algorithm [18] and its space-efficient variant [17] crucially rely on the possibility of choosing very small integers $b_1, \dots, b_d$ such that $\prod_{i=1}^d b_i^{e_i}$ is a nontrivial square root of 1 modulo N for some $e_1, \dots, e_d$ with $|e_i| \leq N^{O(1/d)}$. We remind the reader that $n \asymp \log N$ and $d \asymp \sqrt{n}$.

The least restrictive bound³ on the b_i 's that still allows for a quantum circuit with $\tilde{O}(n^{3/2})$ gates is to have $|b_i| \leq \exp(\tilde{O}(d))$, where the $\tilde{O}(\cdot)$ notation possibly hides a factor $(\log n)^{O(1)}$. Moreover, it is natural to set the b_i 's to be prime numbers (as in [7, 17, 18]) in order to avoid obvious multiplicative relations between them.

³The bound stated in Regev's paper [18] is $|b_i| \leq (\log N)^{O(1)}$, but this condition can be relaxed somewhat, as observed by Ragavan (private communication).

In this paper, we choose $b_1, \dots, b_d$ to be independent random prime numbers less than $d^{10^3 d}$. With these parameters, we can now state a version of Regev's conjecture⁴ that follows from our results.

Corollary 1.4. *Let $N > 2$ be an integer. Let $d := \lceil \sqrt{\log N} \rceil$ and $X := d^{10^3 d}$.*

Let $\mathbf{b}_1, \dots, \mathbf{b}_d$ be i.i.d. random variables, each uniformly distributed in the set of primes $\leq X$ not dividing N .

Then, with high probability, every $x \in \langle \mathbf{b}_1, \dots, \mathbf{b}_d \rangle$ can be expressed as

$$x \equiv \prod_{i=1}^d \mathbf{b}_i^{e_i} \pmod{N}$$

for some integers e_i with $|e_i| \leq e^{O(d)}$ for all $1 \leq i \leq d$.

The analogue of Regev's algorithm for the discrete logarithm problem, proposed by Ekerå and Gärtner [7], relies on a stronger version of Regev's conjecture. It concerns the geometry of a certain lattice $\mathcal{L}$ which encodes all multiplicative dependencies between the $\mathbf{b}_i$'s (modulo N). We prove the following version of it.⁵

Corollary 1.5. *Let $N > 2$ be an integer. Let $d := \lceil \sqrt{\log N} \rceil$ and $X := d^{10^3 d}$.*

Let $\mathbf{b}_1, \dots, \mathbf{b}_d$ be i.i.d. random variables, each uniformly distributed in the set of primes $\leq X$ not dividing N .

Let $\mathcal{L}$ be the random lattice defined by

$$\mathcal{L} := \left\{ (e_1, \dots, e_d) \in \mathbb{Z}^d : \prod_{i=1}^d \mathbf{b}_i^{e_i} \equiv 1 \pmod{N} \right\}. \quad (2)$$

Then, with high probability, this lattice $\mathcal{L}$ has a basis consisting of vectors of Euclidean norm $\leq e^{O(d)}$.

The main technical result of this paper is Theorem 3.18, of which Corollary 1.5 is a special case. In turn, Corollary 1.4 is a direct consequence of Corollary 1.5.

1.4. Subgroup obstructions

Let N, d and $\mathbf{b}_1, \dots, \mathbf{b}_d$ be as in Corollary 1.4. This corollary shows that, with high probability, every element x in the subgroup generated by $\mathbf{b}_1, \dots, \mathbf{b}_d$ can be written as a *short* product of these $\mathbf{b}_i$ modulo N . To obtain the full strength of Regev's assumption [17, Conjecture 1], it would be necessary to show that this subgroup $\langle \mathbf{b}_1, \dots, \mathbf{b}_d \rangle$ contains a nontrivial square root of 1 modulo N .

Unfortunately, it is not possible to prove this in full generality without considerable advances on the well-known *least quadratic nonresidue problem* in number theory. For any prime number p , write $n(p)$ for the smallest positive integer a which is a quadratic nonresidue (i.e., not a square) modulo p . The best known asymptotic upper bound for $n(p)$ is Burgess's classical result [5] that

$$n(p) \ll_{\varepsilon} p^{\frac{1}{4\sqrt{e}} + \varepsilon}.$$

For simplicity, suppose that the integer N to be factored is a product of two equally sized primes $p_1, p_2 \equiv 3 \pmod{4}$. Recall that, for Regev's algorithm, the parameters $\mathbf{b}_i$ are chosen to be less than $\exp(\tilde{O}(d))$. With current techniques, we cannot rule out the possibility that both $n(p_1)$ and $n(p_2)$ are larger than this threshold. If this is the case, all $\mathbf{b}_i$ will be quadratic residues modulo both p_1 and p_2 , which implies that $\langle \mathbf{b}_1, \dots, \mathbf{b}_d \rangle$ is contained in the subgroup of squares modulo N . In particular, since $p_1, p_2 \equiv 3 \pmod{4}$, the subgroup $\langle \mathbf{b}_1, \dots, \mathbf{b}_d \rangle$ does not contain any nontrivial square roots of 1 modulo N .

⁴Compare with [18, Theorem 1.1] or [17, Conjecture 1].

⁵Compare with [7, Assumption 1], or Conjecture E.1 in the full version of [17].

This issue can be approached from several perspectives.

1. Assuming the Generalised Riemann Hypothesis, Ankeny proved the much stronger bound $n(p) \ll (\log p)^2$ [2]. This suggests that quadratic residues may no longer be an obstruction in this case, and indeed, under GRH, it is straightforward to show that $\langle \mathbf{b}_1, \dots, \mathbf{b}_d \rangle$ contains a nontrivial square root of 1 modulo N with high probability. Assuming GRH would also considerably simplify the proof of Theorem 3.18 and remove a logarithmic factor for the gate and qubit costs in Theorems 1.1 and 1.2. However, in this paper we seek fully unconditional results.
2. It is possible to prove the following partial result unconditionally: for *almost all* odd N with at least two distinct prime factors, the subgroup $\langle \mathbf{b}_1, \dots, \mathbf{b}_d \rangle$ contains a nontrivial square root of 1 modulo N with high probability (with $d = d(N)$ and $\mathbf{b}_i = \mathbf{b}_i(N)$ as in Corollary 1.4). By Corollary 1.4, this implies that Regev's original algorithm (with these parameters $\mathbf{b}_i$) finds a nontrivial divisor of N with high probability, for *almost all* N . More precisely, the set $E \subset \mathbb{N}$ of exceptional values of N for which the above does not hold can be shown to satisfy

$$|E \cap [1, x]| \leq \exp((\log x)^{1/2+o(1)})$$

for $x \geq 1$. This partial result can be proved using character sums, the Landau-Page theorem [15, Corollary 11.10] and a grand zero-density estimate [11, Theorem 1 (1.8)], by means of a case analysis depending on the prime factorisation of N .

3. In the present work, we follow a different approach to obtain a completely unconditional result that applies to *all* N , by slightly modifying the algorithm itself. While the key to the efficiency of Regev's algorithm is that the $\mathbf{b}_i$'s are small, one can tolerate a bounded number of large $\mathbf{b}_i$'s.⁶ We use this extra flexibility to overcome the subgroup obstructions. The simplest way to proceed is to allow for one of the parameters, say $\mathbf{b}_1$, to be uniformly distributed in $(\mathbb{Z}/N\mathbb{Z})^\times$. As with Shor's algorithm, this ensures that the subgroup $\langle \mathbf{b}_1 \rangle$, and hence also $\langle \mathbf{b}_1, \dots, \mathbf{b}_d \rangle$, contains a nontrivial square root of 1 with probability $\gg 1$.

Remark 1.6. In his paper [18], Regev proposed to select the parameters $b_1, \dots, b_d$ deterministically, for example by setting b_i to be the i -th smallest prime for $1 \leq i \leq d$. While such a deterministic choice of parameters is likely to work for *almost all* N , this seems to be very difficult to prove.

Remark 1.7. In this paper, the parameters $\mathbf{b}_i$ are chosen to be random primes (not dividing N) under a certain threshold X . Another natural choice would have been to let $\mathbf{b}_1, \dots, \mathbf{b}_d$ be i.i.d. random variables uniformly distributed in the set of all integers $\leq X$ coprime to N . The proofs in this paper would carry over to this framework if an analogue of Proposition 3.9 for short character sums $\sum_{n \leq x} \chi(n)$ were available. Assuming the Generalised Riemann Hypothesis, the work of Granville and Soundararajan [9, Theorem 2] gives a bound of the required strength. However, this approach does not seem sufficient to obtain an unconditional result (the bounds in [9] become too weak when $L(s, \chi)$ has zeroes close to the line $\operatorname{Re} s = 1$).

1.5. Structure of the paper

We begin in Section 2 with a proof of a simple special case of our main theorem. This is intended to illustrate our overall strategy and motivate the need for the more technical approach required in the general case. The purpose of this section is to build intuition; it is not formally required for the main argument.

Our main technical result is Theorem 3.18, which shows that lattices $\mathcal{L}$ similar to that in Corollary 1.5 have a basis of short vectors with high probability. Using simple geometry of numbers (see Section 3.4), we reduce this problem to estimating the number of lattice points in balls of growing radii. Unfortunately, we are unable to obtain a suitable lattice point count for $\mathcal{L}$ directly. We resolve this by considering a

⁶In fact, this observation was already needed to adapt Regev's algorithm to the discrete logarithm problem [7].

different lattice $\mathcal{L}_M$ from the start of the argument (using the lemmas in Section 3.1). In Section 3.2, we expand the lattice point count for $\mathcal{L}_M$ in terms of Dirichlet characters modulo N . This produces a main term, which can be estimated precisely, and an error term. The heart of the proof lies in using a zero-density estimate for Dirichlet characters modulo N to bound this error term unconditionally. Finally, we prove our quantum algorithmic applications (Theorems 1.1 and 1.2) in Section 4.

1.6. Notation

We write $f \ll g$ or $f = O(g)$ if $|f| \leq Cg$ for some absolute constant $C > 0$. If instead, C depends on a parameter θ , we write $f \ll_\theta g$ or $f = O_\theta(g)$. The notation $f \asymp g$ or $f = \Theta(g)$ means that $f \ll g$ and $g \ll f$.

A *character* of a finite abelian group G is a homomorphism $\chi : G \rightarrow \mathbb{C}^\times$. The *order* of χ , denoted by $\text{ord}(\chi)$, is the least positive integer n such that χ^n is the trivial character 1. The group of all characters of G is denoted by $\hat{G}$. If $g_1, \dots, g_k \in G$, we write $\langle g_1, \dots, g_k \rangle$ for the subgroup of G generated by these elements.

A nontrivial square root of 1 modulo N is an element $a \in (\mathbb{Z}/N\mathbb{Z})^\times$ such that $a^2 \equiv 1 \pmod{N}$ and $a \not\equiv \pm 1 \pmod{N}$. Euler's totient function and the prime counting function are denoted by φ and π , respectively. We use the notation $\|\cdot\|_2$ for the Euclidean norm. We write $\log$ for the natural logarithm and $\log_2$ for the logarithm in base 2.

Random variables are typically written in bold font, such as $\mathbf{b}_1, \dots, \mathbf{b}_d$. A statement will be said to occur *with high probability* if it holds with probability tending to 1 as $N \rightarrow \infty$.

We refer the reader to [16] for a detailed textbook on quantum computing, and [4, Chapter 2] for a brief overview.

2. An illustrative special case

In this section, we prove a weak version of our main theorem in a simplified setting. This special case can be used to obtain an unconditional factoring algorithm for integers N of a special form: RSA moduli that are products of two ‘safe primes’.⁷

Definition 2.1 (RSA moduli with safe primes). Let $\mathcal{M}_{\text{RSA}}$ be the set of integers N of the form $N = P \cdot Q$, where $P, Q \geq N^{1/4}$ are distinct primes such that $(P-1)/2$ and $(Q-1)/2$ are also prime.

For $N \in \mathcal{M}_{\text{RSA}}$, the multiplicative group $(\mathbb{Z}/N\mathbb{Z})^\times$ has a very convenient structure, allowing us to give a relatively short proof of the following proposition.

Proposition 2.2. Let $N \in \mathcal{M}_{\text{RSA}}$ be a sufficiently large integer. Let $d := \lceil \sqrt{\log N} \rceil$ and $X := d^d$.

Let $\mathbf{b}_1, \dots, \mathbf{b}_d$ be i.i.d. random variables, each uniformly distributed in the set of primes less than X . Let $\mathbf{b}_0$ be a random variable uniformly distributed in $(\mathbb{Z}/N\mathbb{Z})^\times$, independent from the $\mathbf{b}_i$'s.

Let $x \in (\mathbb{Z}/N\mathbb{Z})^\times$. Then, with probability $\gg 1$, there are integers $0 \leq e_0, \dots, e_d \leq e^{O(d)}$ such that

$$\prod_{i=0}^d \mathbf{b}_i^{e_i} \equiv x \pmod{N}.$$

Applying Proposition 2.2 (with x being a nontrivial square root of 1 modulo N) shows that Regev's factoring algorithm, with random parameters $\mathbf{b}_0, \dots, \mathbf{b}_d$ as in that proposition⁸, admits an unconditional proof of correctness for integers $N \in \mathcal{M}_{\text{RSA}}$.

It should be noted that Theorem 3.18 not only removes the restriction $N \in \mathcal{M}_{\text{RSA}}$, it also provides a stronger geometric conclusion required for applications to the discrete logarithm problem.

⁷Note that this restrictive assumption is generally not satisfied by RSA moduli used in practice.

⁸Again, the presence of one potentially large parameter $\mathbf{b}_0$ circumvents the subgroup obstructions discussed in the introduction, while not significantly affecting the efficiency of the quantum algorithm.

Notation 2.3. For the rest of this section, we fix $N, d, X, \mathbf{b}_0, \dots, \mathbf{b}_d$ and x as in the statement of Proposition 2.2. In particular, $N \in \mathcal{M}_{\text{RSA}}$ is assumed to be sufficiently large. Moreover, there are primes P, Q, P' and Q' such that $N = P \cdot Q$, $P = 2P' + 1$ and $Q = 2Q' + 1$. Hence,

$$G := (\mathbb{Z}/N\mathbb{Z})^\times \cong (\mathbb{Z}/2\mathbb{Z})^2 \times \mathbb{Z}/P'\mathbb{Z} \times \mathbb{Z}/Q'\mathbb{Z} \cong \widehat{G}. \quad (3)$$

To prove Proposition 2.2, we require the following lemma from analytic number theory, that ultimately relies on a zero-density estimate for Dirichlet L -functions.

Lemma 2.4. *For every $j \geq 0$, the number of characters $\psi \in \widehat{G}$ such that*

$$\left| \mathbb{E}_{p \leq X} \psi(p) \right| \geq e^{-j}$$

is at most $e^{O(d)} N^{O(j/\log X)}$. Here, $\mathbb{E}_{p \leq X}$ denotes the average over primes $p \leq X$.

We omit the proof of Lemma 2.4, as it is somewhat tangential to the aim of this explanatory section.⁹

Proof of Proposition 2.2. Let $H = \lfloor e^{Cd} \rfloor$, where C is a large absolute constant to be chosen later.

We will show that, with probability $\gg 1$, there exist integers $0 \leq h_1, \dots, h_d < H$ such that

$$x \equiv \mathbf{b}_0 \mathbf{b}_1^{h_1} \cdots \mathbf{b}_d^{h_d} \pmod{N},$$

which clearly implies the desired conclusion. By orthogonality of characters modulo N , we have

$$\sum_{0 \leq h_1, \dots, h_d < H} \mathbf{1}_{\mathbf{b}_0 \mathbf{b}_1^{h_1} \cdots \mathbf{b}_d^{h_d} \equiv x \pmod{N}} = \frac{1}{\varphi(N)} \sum_{\chi \in \widehat{G}} \chi(\mathbf{b}_0 x^{-1}) \prod_{i=1}^d \sum_{0 \leq h < H} \chi^h(\mathbf{b}_i). \quad (4)$$

Let χ_0 be the trivial character and χ_1, χ_2, χ_3 be the three characters of order 2 in $\widehat{G}$. Each of these four characters χ_j satisfies $\chi_j(\mathbf{b}_i) \in \{\pm 1\}$ for all i , so the term $\sum_{0 \leq h < H} \chi_j^h(\mathbf{b}_i)$ is either equal to H or lies in $\{0, 1\}$. Hence, the contribution of these four characters to (4) is

$$\frac{H^d}{\varphi(N)} + \frac{1}{\varphi(N)} \sum_{j=1}^3 \chi_j(\mathbf{b}_0 x^{-1}) C_j(\mathbf{b}_1, \dots, \mathbf{b}_d). \quad (5)$$

for some quantities $C_j(\mathbf{b}_1, \dots, \mathbf{b}_d) \geq 0$. Since $\mathbf{b}_0$ is uniformly distributed in G , by (3) we have

$$\mathbb{P} \left(\chi_1(\mathbf{b}_0 x^{-1}) = \chi_2(\mathbf{b}_0 x^{-1}) = \chi_3(\mathbf{b}_0 x^{-1}) = 1 \right) = \frac{1}{4},$$

as this is the probability that $\mathbf{b}_0 x^{-1}$ lies in the subgroup of squares modulo N , which has index 4 in G . Therefore, with probability $\geq 1/4$, the expression (5) is bounded below by $H^d/\varphi(N)$.

As a result, Proposition 2.2 reduces to showing that the contribution of the remaining characters $\chi \in \widehat{G} \setminus \{\chi_0, \chi_1, \chi_2, \chi_3\}$ to (4) is, with high probability, at most $H^d/(2\varphi(N))$ in absolute value. By Chebyshev's inequality, it suffices to prove that

$$\mathbb{E}[|\mathbf{z}|^2] = o(H^{2d}), \quad (6)$$

where $\mathbf{z}$ is the random variable defined by

$$\mathbf{z} := \sum_{\substack{\chi \in \widehat{G} \\ \text{ord}(\chi) > 2}} \left| \prod_{i=1}^d \sum_{0 \leq h < H} \chi^h(\mathbf{b}_i) \right|.$$

⁹This lemma is a consequence of Propositions 3.9 and 3.10; see Section 3.3 and Appendix A for further details.

By the triangle inequality for the L^2 -norm and the definition of the random variables $\mathbf{b}_i$, we have

$$\mathbb{E}[|\mathbf{z}|^2]^{1/2} \leq \sum_{\substack{\chi \in \widehat{G} \\ \text{ord}(\chi) > 2}} \left(\mathbb{E} \left[\left| \prod_{i=1}^d \sum_{0 \leq h < H} \chi^h(\mathbf{b}_i) \right|^2 \right] \right)^{1/2} = \sum_{\substack{\chi \in \widehat{G} \\ \text{ord}(\chi) > 2}} \left(\mathbb{E}_{p \leq X} \left| \sum_{0 \leq h < H} \chi^h(p) \right|^2 \right)^{d/2}.$$

Expanding the square and exchanging the order of summation, we obtain

$$\mathbb{E}_{p \leq X} \left| \sum_{0 \leq h < H} \chi^h(p) \right|^2 \leq \sum_{0 \leq h_1, h_2 < H} \left| \mathbb{E}_{p \leq X} \chi^{h_1 - h_2}(p) \right| \leq H \sum_{|h| < H} \left| \mathbb{E}_{p \leq X} \chi^h(p) \right|.$$

Therefore,

$$\mathbb{E}[|\mathbf{z}|^2]^{1/2} \leq H^{d/2} \sum_{\substack{\chi \in \widehat{G} \\ \text{ord}(\chi) > 2}} \left(\sum_{|h| < H} \left| \mathbb{E}_{p \leq X} \chi^h(p) \right| \right)^{d/2}. \quad (7)$$

To estimate the right-hand side of (7), we partition the set of characters as follows. For $j \geq 0$, let

$$E_j := \left\{ \chi \in \widehat{G} \setminus \{\chi_0, \chi_1, \chi_2, \chi_3\} : \max_{0 < |h| < H} \left| \mathbb{E}_{p \leq X} \chi^h(p) \right| \in (e^{-j-1}, e^{-j}] \right\}.$$

Notice that, by (3) and Definition 2.1, every character χ of order greater than 2 has order at least $\min(P', Q') \gg N^{1/4}$. Hence, provided N is sufficiently large, for any χ of order greater than 2, the characters χ^h for $0 < |h| < H$ are all distinct. In addition, for any fixed $\psi \in \widehat{G} \setminus \{\chi_0, \chi_1, \chi_2, \chi_3\}$ and $0 < |h| < H$, there are at most four characters χ such that $\chi^h = \psi$. By Lemma 2.4, these observations imply that

$$|E_j| \ll H e^{O(d)} N^{O(j/\log X)} \ll H e^{O(d)} e^{O(jd/\log d)} \quad (8)$$

for all $j \geq 0$.¹⁰

Let $j \geq 0$ and let $\chi \in E_j$. By Lemma 2.4, there are at most $e^{O(d)}$ integers $0 \leq |h| < H$ such that $|\mathbb{E}_{p \leq X} \chi^h(p)| > d^{-2}$. By definition of E_j , all other integers $0 < |h| < H$ satisfy

$$\left| \mathbb{E}_{p \leq X} \chi^h(p) \right| \leq \min(d^{-2}, e^{-j}).$$

Hence,

$$\left(\sum_{|h| < H} \left| \mathbb{E}_{p \leq X} \chi^h(p) \right| \right)^{d/2} \leq (e^{O(d)} + 2H \min(d^{-2}, e^{-j}))^{d/2} \leq e^{O(d^2)} + e^{O(d)} H^{d/2} \min(d^{-d}, e^{-jd/2}).$$

Summing this over all $\chi \in E_j$ and $j \geq 0$, we can bound (7) by

$$\mathbb{E}[|\mathbf{z}|^2]^{1/2} \leq e^{O(d^2)} \varphi(N) H^{d/2} + e^{O(d)} H^d \sum_{j \geq 0} |E_j| \min(d^{-d}, e^{-jd/2}).$$

Note that $\varphi(N) \leq N \leq e^{d^2}$. Bounding $|E_j|$ using (8) and summing the resulting geometric series (using that N is sufficiently large), we obtain

¹⁰It is here that we crucially rely on the particular structure of $(\mathbb{Z}/N\mathbb{Z})^\times$ for $N \in \mathcal{M}_{\text{RSA}}$. These arguments can be extended somewhat, but additional ideas are required to handle the general case (the most difficult case being when N has many prime factors).

$$\mathbb{E}[|z|^2]^{1/2} \leq (e^{O(d^2)} H^{-d/2} + e^{O(d)} H d^{-d}) H^d.$$

Recalling that $H = \lfloor e^{Cd} \rfloor$ we see that, if C is a sufficiently large absolute constant, the right-hand side is $o(H^d)$ as $N \rightarrow \infty$. This establishes (6) and concludes the proof of Proposition 2.2. $\square$

Remark 2.5. For Theorem 3.18, we obtain (with high probability) an asymptotic formula for quantities similar to (4), which is then combined with techniques from geometry of numbers to prove the existence of a short lattice basis.

This requires a more careful analysis of the small-order characters, which may be more numerous than in the special case treated above. Moreover, the analysis of large-order characters is much more delicate, as an exceptional character ψ (in the sense of Lemma 2.4) may have more representations of the form $\psi = \chi^h$ (for some $\chi \in \widehat{G}$ and $|h| < H$) than the previous argument can handle. This issue is resolved by performing Fourier analysis on a well-chosen subgroup of $(\mathbb{Z}/N\mathbb{Z})^\times$, introduced in Section 3.1.

3. Proof of the existence of a short lattice basis

3.1. Restricting to a convenient subgroup of $(\mathbb{Z}/N\mathbb{Z})^\times$

The goal of this preliminary section is to define a subgroup of the character group of $(\mathbb{Z}/N\mathbb{Z})^\times$ that does not have too many elements of small order. This will be crucial for Section 3.3, to reduce the influence of potential counterexamples to the Generalised Riemann Hypothesis. For example, we need to avoid having many characters $\chi_1, \dots, \chi_k$ modulo N such that $\chi_i^2 = \psi$ for all i , where ψ is an exceptional character (in the sense that $L(s, \psi)$ has a zero very close to the line $\operatorname{Re} s = 1$). The definition of the suitable subgroup depends on the precise structure of $(\mathbb{Z}/N\mathbb{Z})^\times$.

Notation 3.1. If G is a finite abelian group (written multiplicatively) and $M \geq 1$, we write $G^M := \{\chi^M : \chi \in G\}$ for the subgroup of M th powers in G .¹¹

Lemma 3.2. *Let G be a finite abelian group (written multiplicatively) and $M \geq 1$. There is an isomorphism*

$$\iota : \widehat{G^M} \xrightarrow{\sim} \widehat{G}^M$$

such that, for any $\chi \in \widehat{G^M}$, $\iota(\chi) = \widetilde{\chi}^M$ where $\widetilde{\chi} \in \widehat{G}$ is an arbitrary extension of χ to G .

Proof. Pontryagin duality for finite abelian groups is an equivalence of categories, and therefore an exact functor. In simple terms, this implies that injections turn into surjections when switching to the dual, and vice versa. Hence, the M th power map $G \twoheadrightarrow G^M$, $g \mapsto g^M$, induces an injection $\iota : \widehat{G^M} \hookrightarrow \widehat{G}$ defined by

$$\iota(\chi)(g) := \chi(g^M)$$

for every $\chi \in \widehat{G^M}$ and $g \in G$. Moreover, the restriction map $\widehat{G} \rightarrow \widehat{G^M}$ is surjective, using exactness of Pontryagin duality again. Thus, every $\chi \in \widehat{G^M}$ can be extended to a character $\widetilde{\chi}$ on G , and

$$\chi(g^M) = \widetilde{\chi}(g^M) = \widetilde{\chi}^M(g).$$

This implies that the image of ι is contained in $\widehat{G}^M$. Since $|\widehat{G^M}| = |G^M| = |\widehat{G^M}|$, the lemma follows. $\square$

¹¹In particular, G^M is not the M -fold direct product of G .

Definition 3.3. Let $N \geq 1$. For every $h \geq 1$ we define

$$K(h) := \frac{|(\mathbb{Z}/N\mathbb{Z})^\times|}{|((\mathbb{Z}/N\mathbb{Z})^\times)^h|}.$$

In other words, $K(h)$ is the size of the kernel of the h th power map in $(\mathbb{Z}/N\mathbb{Z})^\times$.

Lemma 3.4. Let $N > 2$ be an integer and $d := \lceil \sqrt{\log N} \rceil$.

For every prime p , define m_p to be the largest non-negative integer such that

$$K(p^{m_p}) \geq p^{dm_p/10}.$$

Let $M_* = M_*(N) := \prod_p p^{m_p}$. Then the following holds.

1. We have $M_* \leq e^{10d}$.
2. For every $h \geq 1$, we have $K(M_*h)/K(M_*) \leq h^{d/10}$.

Proof. 1. For every prime p , since $K(p^{m_p})$ is a power of p dividing $|(\mathbb{Z}/N\mathbb{Z})^\times| = \varphi(N)$, we have

$$\prod_p K(p^{m_p}) \mid \varphi(N).$$

By definition of m_p , we know that $K(p^{m_p}) \geq p^{dm_p/10}$ for every p , and thus

$$M_*^{d/10} = \prod_p p^{dm_p/10} \leq \prod_p K(p^{m_p}) \leq \varphi(N) \leq N.$$

Recalling that $d = \lceil \sqrt{\log N} \rceil$, it follows that $M_* \leq e^{10d}$.

2. In any finite abelian group (written multiplicatively) and for any coprime integers a and b , there is an isomorphism

$$\ker(g \mapsto g^a) \times \ker(g \mapsto g^b) \cong \ker(g \mapsto g^{ab}).$$

Thus, if $h = \prod_p p^{e_p}$ is the prime factorisation of h , we see that $K(M_*) = \prod_p K(p^{m_p})$ and $K(M_*h) = \prod_p K(p^{m_p+e_p})$, where all products are finite. Whenever $e_p \geq 1$, we have

$$K(p^{m_p+e_p}) < p^{d(m_p+e_p)/10}, \quad K(p^{m_p}) \geq p^{dm_p/10}$$

by definition of m_p . Thus $K(p^{m_p+e_p})/K(p^{m_p}) \leq p^{de_p/10}$, and this last inequality also holds when $e_p = 0$. We conclude that

$$\frac{K(M_*h)}{K(M_*)} = \prod_p \frac{K(p^{m_p+e_p})}{K(p^{m_p})} \leq \prod_p p^{de_p/10} = h^{d/10}.$$

as claimed. □

Lemma 3.5. Let $N > 2$, $d := \lceil \sqrt{\log N} \rceil$, $G := (\mathbb{Z}/N\mathbb{Z})^\times$ and let $M_* \geq 1$ be the integer defined in the statement of Lemma 3.4.

Let $h \geq 1$. Let $\chi_1, \dots, \chi_n \in \widehat{G^{M_*}}$ be distinct characters such that

$$\chi_1^h = \chi_2^h = \dots = \chi_n^h,$$

Then $n \leq h^{d/10}$.

Proof. Let $f_h : G^{M_*} \rightarrow G^{M_*}$ be the h th power map, $f(\chi) := \chi^h$. By assumption, the n distinct characters $\chi_1\chi_n^{-1}, \chi_2\chi_n^{-1}, \dots, \chi_n\chi_n^{-1}$ lie in the kernel of f . Hence,

$$n \leq |\ker(f)| = \frac{|G^{M_*}|}{|(G^{M_*})^h|} = \frac{|((\mathbb{Z}/N\mathbb{Z})^\times)^{M_*}|}{|((\mathbb{Z}/N\mathbb{Z})^\times)^{M_*h}|} = \frac{K(M_*h)}{K(M_*)},$$

which is $\leq h^{d/10}$ by Lemma 3.4. $\square$

3.2. Lattice point counting via characters

To prove Theorem 3.18, we will need to study the following quantities.

Definition 3.6. Let $N \geq 1$, $G = (\mathbb{Z}/N\mathbb{Z})^\times$ and $\chi \in \widehat{G}$. For all $d \geq 1$ and $b_1, \dots, b_d \in G$, we define the quantity

$$F_{\chi, H}(b_1, \dots, b_d) := \prod_{i=1}^d \sum_{|h| \leq H} \chi^h(b_i).$$

The following lemma relates these expressions to a lattice point counting problem.

Lemma 3.7. Let $N, M, d, H \geq 1$ be integers. Let $b_1, \dots, b_d \in G := (\mathbb{Z}/N\mathbb{Z})^\times$. The number of vectors $(e_1, \dots, e_d) \in \mathbb{Z}^d \cap [-H, H]^d$ such that

$$\prod_{i=1}^d b_i^{Me_i} \equiv 1 \pmod{N}$$

is equal to

$$\frac{1}{|\widehat{G}^M|} \sum_{\chi \in \widehat{G}^M} F_{\chi, H}(b_1, \dots, b_d). \quad (9)$$

Proof. By orthogonality of characters in the abelian group G^M , we have

$$\mathbf{1}_{\prod_i b_i^{Me_i} \equiv 1 \pmod{N}} = \frac{1}{|\widehat{G}^M|} \sum_{\psi \in \widehat{G}^M} \psi \left(\prod_{i=1}^d b_i^{Me_i} \right) = \frac{1}{|\widehat{G}^M|} \sum_{\psi \in \widehat{G}^M} \prod_{i=1}^d \psi(b_i^{Me_i}).$$

Observe that $\psi(b_i^{Me_i}) = \iota(\psi)(b_i^{e_i})$, where ι is the isomorphism $\widehat{G}^M \xrightarrow{\sim} \widehat{G}^M$ defined in Lemma 3.2. We may thus rewrite this equality as

$$\mathbf{1}_{\prod_i b_i^{Me_i} \equiv 1 \pmod{N}} = \frac{1}{|\widehat{G}^M|} \sum_{\chi \in \widehat{G}^M} \prod_{i=1}^d \chi(b_i^{e_i}).$$

Therefore, the number of vectors $(e_1, \dots, e_d) \in \mathbb{Z}^d \cap [-H, H]^d$ such that $\prod_{i=1}^d b_i^{Me_i} \equiv 1 \pmod{N}$ is

$$\sum_{\substack{(e_1, \dots, e_d) \in \mathbb{Z}^d \\ \max_i |e_i| \leq H}} \mathbf{1}_{\prod_i b_i^{Me_i} \equiv 1 \pmod{N}} = \frac{1}{|\widehat{G}^M|} \sum_{\chi \in \widehat{G}^M} \sum_{\substack{(e_1, \dots, e_d) \in \mathbb{Z}^d \\ \max_i |e_i| \leq H}} \prod_{i=1}^d \chi^{e_i}(b_i) = \frac{1}{|\widehat{G}^M|} \sum_{\chi \in \widehat{G}^M} F_{\chi, H}(b_1, \dots, b_d)$$

as claimed. $\square$

It is fairly straightforward to estimate $F_{\chi,H}(b_1, \dots, b_d)$ when χ is a character of small order in $\widehat{G}$. The contribution of these small-order characters gives the expected main term for (9).

Lemma 3.8 (Main term). *Let $N, M, d \geq 1$ be integers. Let $b_1, \dots, b_d \in G := (\mathbb{Z}/N\mathbb{Z})^\times$. Uniformly for all integers $H \geq e^{31d}$, we have*

$$\frac{1}{|\widehat{G}^M|} \sum_{\chi \in \widehat{G}^M} F_{\chi,H}(\underline{b}) = \left(1 + O\left(\frac{e^{31d}}{H}\right)\right) \frac{(2H+1)^d}{|\langle b_1^M, \dots, b_d^M \rangle|} + \frac{1}{|\widehat{G}^M|} \sum_{\substack{\chi \in \widehat{G}^M \\ \text{ord}(\chi) \geq e^{10d}}} |F_{\chi,H}(\underline{b})|, \quad (10)$$

where $F_{\chi,H}(\underline{b})$ is short for $F_{\chi,H}(b_1, \dots, b_d)$.

Proof. Define

$$\langle b_1, \dots, b_d \rangle^\perp := \{\chi \in \widehat{G} : \forall i \in [d], \chi(b_i) = 1\}.$$

Observe that

$$|\langle b_1, \dots, b_d \rangle^\perp \cap \widehat{G}^M| = |\{\chi \in \widehat{G}^M : \forall i \in [d], \chi(b_i^M) = 1\}| = \frac{|G^M|}{|\langle b_1^M, \dots, b_d^M \rangle|} \quad (11)$$

where the first equality follows from Lemma 3.2 and the second from the canonical isomorphism $H_1^\perp \cong \widehat{G_1}/H_1$ for any subgroup H_1 of a finite abelian group G_1 .

Clearly, if $\chi \in \langle b_1, \dots, b_d \rangle^\perp \cap \widehat{G}^M$, we have $F_{\chi,H}(\underline{b}) = (2H+1)^d$. Hence, by (11), those characters contribute

$$\frac{1}{|\widehat{G}^M|} \sum_{\chi \in \langle b_1, \dots, b_d \rangle^\perp \cap \widehat{G}^M} F_{\chi,H}(\underline{b}) = \frac{(2H+1)^d}{|\langle b_1^M, \dots, b_d^M \rangle|}$$

to the sum (10).

For any character $\chi \in \widehat{G}^M \setminus \langle b_1, \dots, b_d \rangle^\perp$ of order $\leq e^{10d}$, and any $i \in [d]$ such that $\chi(b_i) \neq 1$, we can bound

$$\left| \sum_{|h| \leq H} \chi(b_i)^h \right| \leq \text{ord}(\chi) \leq e^{10d}$$

by the geometric series formula. Thus, defining

$$I_\chi := \{i \in [d] : \chi(b_i) \neq 1\},$$

we have

$$|F_{\chi,H}(\underline{b})| \leq e^{10d|I_\chi|} (2H+1)^{d-|I_\chi|}.$$

Consequently,

$$\sum_{\substack{\chi \in \widehat{G}^M \setminus \langle b_1, \dots, b_d \rangle^\perp \\ \text{ord}(\chi) \leq e^{10d}}} |F_{\chi,H}(\underline{b})| \ll \sum_{\substack{I \subset [d] \\ I \neq \emptyset}} e^{10d|I|} (2H+1)^{d-|I|} \sum_{\substack{\chi \in \widehat{G}^M \\ \text{ord}(\chi) \leq e^{10d} \\ I_\chi = I}} 1. \quad (12)$$

Fix some nonempty set $I \subset [d]$ and complex numbers $(z_i)_{i \in I}$. Let $C_{I, (z_i)}$ be the set of all characters $\chi \in \widehat{G}^M$ such that

$$\chi(b_i) = \begin{cases} z_i & \text{if } i \in I \\ 1 & \text{if } i \in [d] \setminus I. \end{cases}$$

Note that $C_{I, (z_i)}$ is either the empty set, or a coset of $\langle b_1, \dots, b_d \rangle^\perp \cap \widehat{G}^M$ in $\widehat{G}^M$. Moreover, if χ has order at most e^{10d} , this set $C_{I, (z_i)}$ can only be nonempty if all z_i are roots of unity of order $\leq e^{10d}$, and there are $\leq e^{20d}$ such roots of unity. We conclude that

$$\sum_{\substack{\chi \in \widehat{G}^M \\ \text{ord}(\chi) \leq e^{10d} \\ I_\chi = I}} 1 \leq \sum_{(z_i)_{i \in I}} |C_{I, (z_i)}| \leq e^{20d|I|} |\langle b_1, \dots, b_d \rangle^\perp \cap \widehat{G}^M|.$$

Thus, we can bound the right-hand side of (12) by

$$\ll (2H+1)^d |\langle b_1, \dots, b_d \rangle^\perp \cap \widehat{G}^M| \sum_{\substack{I \subset [d] \\ I \neq \emptyset}} (e^{30d} H^{-1})^{|I|} \ll (2H+1)^d |\langle b_1, \dots, b_d \rangle^\perp \cap \widehat{G}^M| 2^d e^{30d} H^{-1}$$

where we used that $H \geq e^{30d}$ in the last step. By (11), this means that

$$\frac{1}{|\widehat{G}^M|} \sum_{\substack{\chi \in \widehat{G}^M \setminus \langle b_1, \dots, b_d \rangle^\perp \\ \text{ord}(\chi) \leq e^{10d}}} |F_{\chi, H}(\underline{b})| \ll e^{31d} H^{-1} \frac{(2H+1)^d}{|\langle b_1^M, \dots, b_d^M \rangle|},$$

which completes the proof. $\square$

3.3. Bounding the contribution of large-order characters

In this section, we bound the error term coming from large-order characters, on average over primes $b_1, \dots, b_d$ in a short interval.

We will need the following ingredients from classical analytic number theory.

Proposition 3.9 (Character sums over primes). *Let $1/2 \leq \alpha \leq 1$. Let $q, x \geq 2$. Let χ be a nonprincipal character modulo q whose Dirichlet L -function $L(s, \chi)$ has no zero in the rectangle*

$$\{s \in \mathbb{C} : \alpha < \text{Re } s \leq 1, |\text{Im } s| \leq x^{1-\alpha}\}.$$

Then

$$\frac{1}{\pi(x)} \sum_{p \leq x} \chi(p) \ll x^{-(1-\alpha)} \log^3(qx).$$

Proposition 3.9 is a standard consequence of the explicit formula for $L(s, \chi)$ and is proved in Appendix A. The logarithmic factors can be somewhat improved, but such refinements are irrelevant here.

The Generalised Riemann Hypothesis is the claim that, for every Dirichlet character χ , $L(s, \chi)$ has no zero ρ with $\frac{1}{2} < \text{Re } \rho < 1$. Assuming GRH, Proposition 3.9 thus implies almost square-root cancellation for character sums over primes. Proposition 3.10 will serve as an unconditional substitute for GRH.

Proposition 3.10 (Zero-density estimate for a fixed modulus). *Uniformly for $\frac{4}{5} \leq \alpha \leq 1$ and $q, T \geq 1$, we have*

$$\sum_{\chi \pmod{q}} \mathcal{N}(\alpha, T, \chi) \ll_{\varepsilon} (qT)^{(2+\varepsilon)(1-\alpha)}.$$

where $\mathcal{N}(\alpha, T, \chi)$ denotes the number of zeros (with multiplicity) of $L(s, \chi)$ in the rectangle

$$\{s \in \mathbb{C} : \alpha < \operatorname{Re} s \leq 1, |\operatorname{Im} s| \leq T\}.$$

Proof. This is [11, Theorem 1 (1.7)]. □

Our goal is to control the total contribution of all large-order characters in (10). We will do so in Proposition 3.14 (restricting to a suitable subgroup of $\widehat{G}$). We first prove the following L^2 bound for a single large-order character.

Lemma 3.11. *Let $N > 2$ be an integer and let $G = (\mathbb{Z}/N\mathbb{Z})^\times$. Let $d = \lceil \sqrt{\log N} \rceil$ and $X = d^{10^3 d}$.*

Let $\mathbf{b}_1, \dots, \mathbf{b}_d$ be i.i.d. random variables, each uniformly distributed in the set of primes less than X not dividing N . Let $\chi \in \widehat{G}$ be a character of order $\geq e^{10d}$. Then, for every $H \geq e^{10d}$,

$$\mathbb{E} \left[|F_{\chi, H}(\mathbf{b}_1, \dots, \mathbf{b}_d)|^2 \right] \ll d^{-10d} H^{2d}.$$

Proof. By Proposition 3.10 (the zero-density estimate), the number of Dirichlet characters modulo N whose associated L -function has a zero in the region

$$\left\{ s \in \mathbb{C} : 1 - \frac{1}{10d} < \operatorname{Re} s \leq 1, |\operatorname{Im} s| \leq X \right\} \quad (13)$$

is bounded by

$$\ll (NX)^{(2+1)/(10d)} \ll e^d. \quad (14)$$

If a nonprincipal character ψ modulo N has no zero in the region (13), then by Proposition 3.9 we have

$$\mathbb{E} [\psi(\mathbf{b}_1)] \ll X^{-1/(10d)} (\log(NX))^3 \ll d^{-100} d^6 \ll d^{-94}. \quad (15)$$

We now expand $\mathbb{E} \left[|F_{\chi, H}(\mathbf{b}_1, \dots, \mathbf{b}_d)|^2 \right]$ as

$$\begin{aligned} \mathbb{E} \left[\left| \prod_{i=1}^d \sum_{|h| \leq H} \chi^h(\mathbf{b}_i) \right|^2 \right] &= \prod_{i=1}^d \mathbb{E} \left[\left| \sum_{|h| \leq H} \chi^h(\mathbf{b}_i) \right|^2 \right] \\ &= \left(\sum_{|h_1| \leq H} \sum_{|h_2| \leq H} \mathbb{E} [\chi^{h_1-h_2}(\mathbf{b}_1)] \right)^d \\ &\leq (2H+1)^d \left(\sum_{|h| \leq 2H} |\mathbb{E} [\chi^h(\mathbf{b}_1)]| \right)^d. \end{aligned}$$

We can use (15) to bound the term $\mathbb{E} [\chi^h(\mathbf{b}_1)]$, unless χ^h is principal or $L(s, \chi^h)$ has a zero in the region (13). By (14), the number of values of h with $|h| \leq 2H$ for which one of these two situations occurs is

$$\ll \left(1 + \frac{H}{\operatorname{ord}(\chi)} \right) e^d \ll e^{-10d} e^d H \ll e^{-d} H.$$

By (15), we deduce that, for all but $O(e^{-d}H)$ values of $|h| \leq 2H$, the bound $\mathbb{E} [\chi^h(\mathbf{b}_1)] \ll d^{-94}$ holds. Hence

$$\mathbb{E} [|F_{\chi,H}(\mathbf{b}_1, \dots, \mathbf{b}_d)|^2] \leq (2H+1)^d \left(O(e^{-d}H) + O(d^{-94}H) \right)^d \ll e^{O(d)} d^{-94d} H^{2d} \ll d^{-10d} H^{2d},$$

as desired. $\square$

Lemma 3.11 applies to all characters χ of order at least e^{10d} , but gives a relatively weak upper bound. In Lemma 3.13, we will prove that a stronger bound can be obtained if a small set of exceptions is allowed. We begin by describing the exceptional characters in the following lemma.

Lemma 3.12. *There is some absolute constant $t_0 \geq 1$ such that the following holds.*

Let N, G, d, X and $\mathbf{b}_i$ be as in Lemma 3.11. Let $H \geq e^{10d}$.

Let $\chi \in \widehat{G}$ be a character of order $\geq e^{10d}$ such that

$$\mathbb{E} [|F_{\chi,H}(\mathbf{b}_1, \dots, \mathbf{b}_d)|^2] > e^{-2td} H^{2d}$$

for some t with $t_0 \leq t \leq 2d$.

Then there are integers $-2H \leq h_1, h_2 \leq 2H$ with $0 < h_2 - h_1 \leq e^{3t}$ such that both $L(s, \chi^{h_1})$ and $L(s, \chi^{h_2})$ have a zero in the region

$$\left\{ s \in \mathbb{C} : 1 - \frac{t}{100d} < \operatorname{Re} s \leq 1, |\operatorname{Im} s| \leq X \right\}. \quad (16)$$

Proof. As in the proof of Lemma 3.11, we have

$$\mathbb{E} [|F_{\chi,H}(\mathbf{b}_1, \dots, \mathbf{b}_d)|^2] = \mathbb{E} \left[\left| \prod_{i=1}^d \sum_{|h| \leq H} \chi^h(\mathbf{b}_i) \right|^2 \right] \leq (2H+1)^d \left(\sum_{|h| \leq 2H} |\mathbb{E} [\chi^h(\mathbf{b}_1)]| \right)^d. \quad (17)$$

Let I_1 be the set of all $-2H \leq h \leq 2H$ such that χ^h is principal. Since χ has order $\geq e^{10d}$, we have

$$|I_1| \ll 1 + e^{-10d} H \ll e^{-10d} H.$$

Let I_2 be the set of all $-2H \leq h \leq 2H$ such that $L(s, \chi^h)$ has a zero in the region (16). By contradiction, suppose that the conclusion of Lemma 3.12 does not hold. Then any sub-interval of $[-2H, 2H]$ of length e^{3t} contains at most one element of I_2 . This implies that

$$|I_2| \ll 1 + e^{-3t} H \ll e^{-3t} H.$$

Moreover, for any integer $h \in [-2H, 2H] \setminus (I_1 \cup I_2)$, the character χ^h is nonprincipal and has no zero in the region (16), which by Proposition 3.9 implies that

$$\mathbb{E} [\chi^h(\mathbf{b}_1)] \ll X^{-t/(100d)} (\log(NX))^3 \ll d^{-10t} d^6 \ll e^{-3t}.$$

Therefore, we can bound

$$\sum_{|h| \leq 2H} |\mathbb{E} [\chi^h(\mathbf{b}_1)]| \ll |I_1| + |I_2| + e^{-3t} H \ll e^{-10d} H + e^{-3t} H \ll e^{-3t} H.$$

Hence, by (17) we get

$$\mathbb{E} [|F_{\chi,H}(\mathbf{b}_1, \dots, \mathbf{b}_d)|^2] \leq e^{O(d)} e^{-3td} H^{2d},$$

which contradicts the assumption in the statement if t_0 is chosen to be sufficiently large. $\square$

We wish to use a zero-density estimate again to show that there are few characters satisfying the conclusion of Lemma 3.12. For this step to work, we need to restrict to the subgroup $\widehat{G}^{M_*}$ of the full character group $\widehat{G}$, where M_* is the integer defined in Lemma 3.4.

Lemma 3.13. *Let $t_0 \geq 1$ be the constant from Lemma 3.12. Let N, G, d, X and $\mathbf{b}_i$ be as in Lemma 3.11. Let $H \geq e^{10d}$ and let $M_* \geq 1$ be the integer defined in Lemma 3.4.*

For every $t \geq t_0$,

$$\left| \left\{ \chi \in \widehat{G}^{M_*} : \text{ord}(\chi) \geq e^{10d}, \mathbb{E} [|F_{\chi,H}(\mathbf{b}_1, \dots, \mathbf{b}_d)|^2] > e^{-2td} H^{2d} \right\} \right| \ll e^{td/2}. \quad (18)$$

Proof. If $t \geq 2d$, the bound (18) is trivially satisfied as the right-hand side is $\gg N$. We may thus assume that $t_0 \leq t \leq 2d$, in which case Lemma 3.12 applies.

Let E_t be the set of all characters ψ modulo N such that $L(s, \psi)$ has a zero in the rectangle defined in (16). By Proposition 3.10, this set E_t has size

$$|E_t| \ll (NX)^{(2+1/2)t/(100d)} \ll e^{td/20}. \quad (19)$$

For every $\chi \in \widehat{G}$ of order $\geq e^{10d}$, if $\mathbb{E} [|F_{\chi,H}(\mathbf{b}_1, \dots, \mathbf{b}_d)|^2] > e^{-2td} H^{2d}$ then by Lemma 3.12 we can write

$$\chi^h = \psi_1 \overline{\psi_2}$$

for some integer $0 < h < e^{3t}$ and some characters $\psi_1, \psi_2 \in E_t$. Hence, the left-hand side of (18) is

$$\leq \sum_{\psi_1, \psi_2 \in E_t} \sum_{0 < h < e^{3t}} \left| \left\{ \chi \in \widehat{G}^{M_*} : \chi^h = \psi_1 \overline{\psi_2} \right\} \right|. \quad (20)$$

Since $\widehat{G}^{M_*}$ and $\widehat{G}^{M_*}$ are isomorphic, Lemma 3.5 implies that

$$\left| \left\{ \chi \in \widehat{G}^{M_*} : \chi^h = \psi_1 \overline{\psi_2} \right\} \right| \leq h^{d/10} \leq e^{3td/10}. \quad (21)$$

Inserting Eqs. (19) and (21) into (20), we conclude that the left-hand side of (18) is

$$\leq |E_t|^2 e^{3t} e^{3td/10} \leq e^{3t} e^{4td/10} \ll e^{td/2}$$

as claimed. $\square$

Combining the previous lemmas, we obtain a suitable bound for the sum of second moments of $F_{\chi,H}(\mathbf{b}_1, \dots, \mathbf{b}_d)$ over all large-order characters in $\widehat{G}^{M_*}$.

Proposition 3.14 (Large-order characters). *Let N, G, d, X and $\mathbf{b}_i$ be as in Lemma 3.11. Let $M_* \geq 1$ be the integer defined in Lemma 3.4. For $H \geq e^{10d}$, we have*

$$\sum_{\substack{\chi \in \widehat{G}^{M_*} \\ \text{ord}(\chi) \geq e^{10d}}} \mathbb{E} [|F_{\chi,H}(\mathbf{b}_1, \dots, \mathbf{b}_d)|^2]^{1/2} \ll d^{-2d} H^d.$$

Proof. By Lemma 3.11, we know that $\mathbb{E} [|F_{\chi,H}(\mathbf{b})|^2]^{1/2} \leq C d^{-5d} H^d$ for all characters $\chi \in \widehat{G}$ of order $\geq e^{10d}$, where $C > 0$ is an absolute constant (as before, $\mathbf{b}$ stands for $\mathbf{b}_1, \dots, \mathbf{b}_d$). Therefore,

$$\sum_{\substack{\chi \in \widehat{G}^{M_*} \\ \text{ord}(\chi) \geq e^{10d}}} \mathbb{E} [|F_{\chi,H}(\mathbf{b})|^2]^{1/2} \leq \sum_{m=m_0}^{+\infty} \sum_{\substack{\chi \in \widehat{G}^{M_*} \\ \text{ord}(\chi) \geq e^{10d}}} e^{-m+1} H^d \mathbf{1}_{e^{-m} < \mathbb{E} [|F_{\chi,H}(\mathbf{b})|^2]^{1/2} H^{-d} \leq e^{-m+1}}(\chi)$$

where $m_0 := \lfloor 5d \log d - \log C + 1 \rfloor$. We may assume that $m_0 \geq \max(t_0 d, 4d \log d)$, since otherwise $d \ll 1$ and Proposition 3.14 is trivially satisfied (given that $|\widehat{G}^{M_*}| \leq N \ll 1$ when $d \ll 1$).

Applying Lemma 3.13 with $t := m/d$, we obtain that for every $m \geq t_0 d$,

$$\left| \left\{ \chi \in \widehat{G}^{M_*} : \text{ord}(\chi) \geq e^{10d}, \mathbb{E} [|F_{\chi, H}(\mathbf{b})|^2]^{1/2} > e^{-m} H^d \right\} \right| \ll e^{m/2}.$$

Therefore,

$$\sum_{\substack{\chi \in \widehat{G}^{M_*} \\ \text{ord}(\chi) \geq e^{10d}}} \mathbb{E} [|F_{\chi, H}(\mathbf{b})|^2]^{1/2} \ll \sum_{m=m_0}^{+\infty} e^{m/2} e^{-m+1} H^d \ll e^{-m_0/2} H^d$$

which is $O(d^{-2d} H^d)$ as $m_0 \geq 4d \log d$. □

3.4. Geometry of numbers

In this section, we show how to pass from good estimates on the number of lattice points in certain regions to the existence of a short basis for the lattice.

Lemma 3.15. *Let $L \geq 1$ be an integer. Cover the cube $[-L, L]^d$ by $(2L)^d$ cubes of side length 1 in the obvious way. Label these unit cubes $C_1, \dots, C_{(2L)^d}$ (in any order). Let $V \subset \mathbb{R}^d$ be a hyperplane through the origin. Then the number of unit cubes C_i intersecting V is at most $(d+1)(2L)^{d-1}$.*

Proof. Let $e_1, \dots, e_d$ be the standard basis for $\mathbb{R}^d$. Let $v \in \mathbb{R}^d$ be a unit vector orthogonal to V . Without loss of generality, since $\|v\|_2 = 1$, we may assume that $\langle v, e_1 \rangle \geq 1/\sqrt{d}$.

Suppose that V intersects two unit cubes C_i and C_j where $C_i = C_j + k e_1$ for some integer $k \geq 0$. Then, there is some $p \in C_i$ and some $w \in \mathbb{R}^d$ with $\|w\|_\infty \leq 1$ such that $p \in V$ and $p + k e_1 + w \in V$. Therefore, $k e_1 + w \in V$ and thus $\langle k e_1 + w, v \rangle = 0$. Noting that

$$\langle k e_1 + w, v \rangle \geq k \langle e_1, v \rangle - \|w\|_2 \|v\|_2 \geq \frac{k}{\sqrt{d}} - \sqrt{d},$$

we deduce that $k \leq d$.

We have thus proved that, for any cube C_i , there are at most $d+1$ cubes of the form $C_i + k e_1$ for some $k \in \mathbb{Z}$ which intersect V . The lemma follows. □

The next lemma allows us to convert information about the number of lattice points in cubes into the existence of short linearly independent lattice vectors.

Lemma 3.16. *Let $d \geq 1$. Let $\Lambda \subset \mathbb{R}^d$ be a full-rank lattice. Let $2 \leq H_0 < H_1$ be real numbers such that H_1/H_0 is an integer. Suppose that, for $i \in \{0, 1\}$,*

$$|\Lambda \cap [-H_i, H_i]^d| = \theta_i \frac{(2H_i + 1)^d}{\text{vol}(\mathbb{R}^d/\Lambda)} \quad (22)$$

where $\theta_0, \theta_1 > 0$ satisfy $\frac{H_1}{H_0} > \frac{\theta_0}{\theta_1} d \left(\frac{5}{2}\right)^d$. Then $\Lambda \cap [-H_1, H_1]^d$ contains d linearly independent vectors.

Proof. By contradiction, suppose that there exists a linear hyperplane V containing all the points $v \in \Lambda \cap [-H_1, H_1]^d$.

Let $L = H_1/H_0$. The large cube $[-H_1, H_1]^d$ can be covered by $(2L)^d$ axis-parallel cubes of side length H_0 in the natural way. By Lemma 3.15, we can bound

$$|\Lambda \cap [-H_1, H_1]^d| \leq (d+1)(2L)^{d-1} \sup_C |\Lambda \cap C|,$$

where the supremum runs over all (not necessarily centred) axis-parallel cubes $C \subset \mathbb{R}^d$ of side length H_0 . If C is such a cube and $v \in \Lambda \cap C$, then $C \subset v + [-H_0, H_0]^d$, which implies that

$$|\Lambda \cap C| \leq |\Lambda \cap (v + [-H_0, H_0]^d)| = |\Lambda \cap [-H_0, H_0]^d|.$$

Thus,

$$|\Lambda \cap [-H_1, H_1]^d| \leq (d+1)(2L)^{d-1} |\Lambda \cap [-H_0, H_0]^d|.$$

Plugging in our lattice point estimate (22), we get

$$\theta_1(2H_1 + 1)^d \leq (d+1)(2L)^{d-1} \theta_0(2H_0 + 1)^d.$$

Using $2H_1 + 1 \geq 2LH_0$ and $d+1 \leq 2d$, this implies that $L \leq \frac{\theta_0}{\theta_1} d(2 + \frac{1}{H_0})^d$, contradicting the inequality in the statement of the lemma. $\square$

The linearly independent vectors given by Lemma 3.16 can be upgraded to a genuine basis for Λ by standard geometry of numbers, namely Mahler's theorem. We state this fact in a slightly more general situation.

Lemma 3.17. *Let $d \geq 1$. Let $\Lambda_1, \Lambda_2 \subset \mathbb{R}^d$ be full-rank lattices such that $M\Lambda_1 \subset \Lambda_2$ for some integer $M \geq 1$. Suppose that Λ_1 contains d linearly independent vectors in $[-H, H]^d$ for some $H > 0$. Then, Λ_2 admits a basis where each basis vector has Euclidean norm $\leq d^{3/2}MH$.*

Proof. By assumption, there are linearly independent vectors $v_1, \dots, v_d \in \Lambda_1 \cap [-H, H]^d$. Then $Mv_1, \dots, Mv_d$ are linearly independent vectors of Λ_2 such that $\max_{i \in [d]} \|Mv_i\|_2 \leq \sqrt{d}MH$.

Let $B \subset \mathbb{R}^d$ be the unit ball for the Euclidean norm. Let $0 < \lambda_1 \leq \dots \leq \lambda_d$ be the successive minima¹² of B with respect to Λ_2 . Since $Mv_1, \dots, Mv_d \in \Lambda_2$ are linearly independent, we deduce from the above that $\lambda_d \leq \sqrt{d}MH$.

By Mahler's theorem (see [22, Theorem 3.34]), we conclude that Λ_2 admits a basis of vectors of Euclidean norm at most $d\lambda_d \leq d^{3/2}MH$, which is what we needed to show. $\square$

3.5. Short basis vectors

We can now prove our main technical result, which may be of independent interest.

Theorem 3.18. *Let $N > 2$ be an integer. Let $d := \lceil \sqrt{\log N} \rceil$ and $X := d^{10^3 d}$.*

Let $\mathbf{b}_1, \dots, \mathbf{b}_d$ be i.i.d. random variables, each uniformly distributed in the set of primes $\leq X$ not dividing N . Let $r \geq 0$ and let $\mathbf{x}_1, \dots, \mathbf{x}_r$ be arbitrary¹³ random variables taking values in $(\mathbb{Z}/N\mathbb{Z})^\times$.

Then, with probability $1 + O(d^{-d})$, the lattice

$$\mathcal{L} := \left\{ (e_1, \dots, e_d, f_1, \dots, f_r) \in \mathbb{Z}^{d+r} : \prod_{i=1}^d \mathbf{b}_i^{e_i} \prod_{i=1}^r \mathbf{x}_i^{f_i} \equiv 1 \pmod{N} \right\}$$

has a basis consisting of vectors of Euclidean norm $\ll e^{42(d+r)}$.

Remark 3.19. The constant 42 in the exponent is by no means the limit our techniques. For example, using smooth cutoffs in Definition 3.6 would significantly reduce the error term in Lemma 3.8 and hence lower this constant. We have not performed such optimisations to keep the paper as simple as possible.

Proof of Theorem 3.18. Let $M := M_*(N)$ be the integer defined in Lemma 3.4. We introduce the auxiliary lattice

¹²See [22, Definition 3.29] for the definition of successive minima.

¹³In particular, the $\mathbf{x}_i$ are not assumed to be independent or identically distributed.

$$\mathcal{L}_M := \left\{ (e_1, \dots, e_d, f_1, \dots, f_r) \in \mathbb{Z}^{d+r} : \prod_{i=1}^d \mathbf{b}_i^{Me_i} \prod_{i=1}^r \mathbf{x}_i^{Mf_i} \equiv 1 \pmod{N} \right\}.$$

By Lemmas 3.7 and 3.8, we have the lattice point estimate

$$|\mathcal{L}_M \cap [-H, H]^{d+r}| = \frac{(1 + O(e^{31(d+r)H^{-1}}))(2H+1)^{d+r}}{|\langle \mathbf{b}_1^M, \dots, \mathbf{b}_d^M, \mathbf{x}_1^M, \dots, \mathbf{x}_r^M \rangle|} + \frac{1}{|\widehat{G}^M|} \sum_{\substack{\chi \in \widehat{G}^M \\ \text{ord}(\chi) \geq e^{10d}}} |F_{\chi, H}(\underline{\mathbf{b}}, \underline{\mathbf{x}})| \quad (23)$$

for every integer $H \geq e^{31(d+r)}$, where

$$F_{\chi, H}(\underline{\mathbf{b}}, \underline{\mathbf{x}}) := F_{\chi, H}(\mathbf{b}_1, \dots, \mathbf{b}_d, \mathbf{x}_1, \dots, \mathbf{x}_r) = \left(\prod_{i=1}^d \sum_{|h| \leq H} \chi^h(\mathbf{b}_i) \right) \left(\prod_{i=1}^r \sum_{|h| \leq H} \chi^h(\mathbf{x}_i) \right).$$

Since the $\mathbf{x}_i$ have unknown distributions, we will use the trivial bound

$$|F_{\chi, H}(\underline{\mathbf{b}}, \underline{\mathbf{x}})| \leq (2H+1)^r |F_{\chi, H}(\mathbf{b}_1, \dots, \mathbf{b}_d)|.$$

Applying Proposition 3.14, we deduce that, for all $H \geq e^{10d}$,

$$\sum_{\substack{\chi \in \widehat{G}^{M*} \\ \text{ord}(\chi) \geq e^{10d}}} \mathbb{E} [|F_{\chi, H}(\underline{\mathbf{b}}, \underline{\mathbf{x}})|^2]^{1/2} \ll d^{-2d} (2H+1)^{d+r}.$$

By the L^2 triangle inequality and Chebyshev's inequality, this implies that, for fixed $H \geq e^{10d}$,

$$\mathbb{P} \left(\sum_{\substack{\chi \in \widehat{G}^M \\ \text{ord}(\chi) \geq e^{10d}}} |F_{\chi, H}(\underline{\mathbf{b}}, \underline{\mathbf{x}})| > d^{-d} (2H+1)^{d+r} \right) \ll d^{-2d}. \quad (24)$$

Let $H_0 := d \lceil e^{31(d+r)} \rceil$ and $H_1 := \lceil d(d+r)(5/2)^{d+r} \rceil H_0$. We apply Eqs. (23) and (24) twice, once with $H = H_0$ and once with $H = H_1$ to obtain the following: with probability $1 + O(d^{-d})$, the two estimates

$$|\mathcal{L}_M \cap [-H_0, H_0]^{d+r}| = (1 + O(1/d)) \frac{(2H_0+1)^{d+r}}{|\langle \mathbf{b}_1^M, \dots, \mathbf{b}_d^M, \mathbf{x}_1^M, \dots, \mathbf{x}_r^M \rangle|}$$

and

$$|\mathcal{L}_M \cap [-H_1, H_1]^{d+r}| = (1 + O(1/d)) \frac{(2H_1+1)^{d+r}}{|\langle \mathbf{b}_1^M, \dots, \mathbf{b}_d^M, \mathbf{x}_1^M, \dots, \mathbf{x}_r^M \rangle|}$$

simultaneously hold.

We now apply Lemma 3.16. By our choice of H_0 and H_1 , the inequality in the statement is satisfied provided that d is sufficiently large. We can assume that d is large enough as, for $d \ll 1$, we have $N \ll 1$ and Theorem 3.18 is trivially true. Thus, Lemma 3.16 implies that $\mathcal{L}_M \cap [-H_1, H_1]^{d+r}$ contains $d+r$ linearly independent vectors, with probability $1 + O(d^{-d})$. By Lemma 3.17, since $M\mathcal{L} \subset \mathcal{L}_M$, we conclude that, with probability $1 + O(d^{-d})$, $\mathcal{L}$ admits a basis of vectors of Euclidean norm at most

$$\ll d^{3/2} M H_1 \ll d^{3/2} e^{10d} d^2 (d+r) (5/2)^{d+r} e^{31(d+r)} \ll e^{42(d+r)}$$

using the bound $M \leq e^{10d}$ given by Lemma 3.4. This completes the proof. $\square$

4. Applications to quantum computing

In this section, we prove the correctness of efficient quantum algorithms for factoring and for the discrete logarithm problem by applying our version of Regev's number-theoretic conjecture, Theorem 3.18.

4.1. Preparatory lemmas

Lemma 4.1. *Let $N, d, m \geq 2$ be integers. There is a classical algorithm that, given integers $0 \leq a_1, \dots, a_d \leq 2^m$ and exponents $t_1, \dots, t_d \in \{0, 1\}$, computes the product*

$$\prod_{i=1}^d a_i^{t_i} \pmod{N}$$

in time $O(md(\log d) \log(md))$.

Proof. This is similar to [18, p5] or [17, Lemma 12] but for general values of m . Without loss of generality, assume that d is a power of 2, say $d = 2^l$. We proceed to compute this product in a binary tree fashion. Let $T(k)$ be the complexity of multiplying any k of these numbers $a_i^{t_i}$ modulo N . Note that $T(2k) \leq 2T(k) + O(M(mk))$ where $M(x)$ is the time needed to multiply two integers having at most x bits. By the work of Harvey and van der Hoeven [10], it is known that $M(x) = O(x \log x)$, which leads to the bound

$$T(d) \ll \sum_{j=0}^l 2^j M(md/2^{j+1}) \ll \sum_{j=0}^l md \log(md/2^{j+1}) \ll md(\log d) \log(md).$$

as claimed. $\square$

We can turn this into a quantum circuit with the following well-known fact.

Lemma 4.2. *Any classical circuit can be ‘compiled’ into a reversible quantum circuit that carries out the same computations, with the number of gates and qubits used being proportional to the size of the classical circuit.*

Proof. This well-known fact is explained in Section A.1 of the full version of [17]. $\square$

Lemma 4.3. *Let N be a sufficiently large integer. Let $d := \lceil \sqrt{\log N} \rceil$ and $X = d^{10^3 d}$. Let $k = d^4$.*

Let $\mathbf{n}_1, \dots, \mathbf{n}_k$ be i.i.d. random variables uniformly distributed in $\{1, \dots, X\}$. Then, the probability that at least d of these $\mathbf{n}_i$ are prime numbers not dividing N is $1 + O(1/N)$.

Proof. Since N has $\ll \log N$ prime factors, and since the number of primes $\leq X$ is $\gg X/\log X$, we have

$$\mathbb{P}(\mathbf{n}_1 \text{ is prime and } \mathbf{n}_1 \nmid N) \geq \frac{c}{\log X}$$

for some absolute constant $c > 0$. Thus, if E_i is the event that $\mathbf{n}_i$ is not prime or divides N , then by the union bound and independence,

$$\mathbb{P}\left(\bigcup_{\substack{I \subset [k] \\ |I| > k-d}} \bigcap_{i \in I} E_i\right) \leq \sum_{l < d} \binom{k}{l} \left(1 - \frac{c}{\log X}\right)^{k-l} \leq dk^d e^{-c(k-d)/\log X} \ll e^{-d^2}$$

as needed. $\square$

4.2. The discrete logarithm problem

We now prove Theorem 1.2. For convenience, we restate it here.

Theorem 1.2. *There is a quantum circuit having $O(n^{3/2} \log^3 n)$ quantum gates and $O(n \log^3 n)$ qubits with the following property. There is a classical randomised polynomial-time algorithm that solves the discrete logarithm problem*

Input : *an integer $N \leq 2^n$ and elements $g, y \in (\mathbb{Z}/N\mathbb{Z})^\times$ such that $y \in \langle g \rangle$*

Output : *an integer x such that $g^x \equiv y \pmod{N}$*

using $O(\sqrt{n})$ calls to this quantum circuit, and succeeds with probability $\Theta(1)$.

Proof of Theorem 1.2. Suppose we are given an integer $N > 2$ and elements $g, y \in (\mathbb{Z}/N\mathbb{Z})^\times$ such that y lies in the subgroup generated by g in $(\mathbb{Z}/N\mathbb{Z})^\times$. Let $d := \lceil \sqrt{\log N} \rceil$ and $X := d^{10^3 d}$. Let $\mathbf{b}_1, \dots, \mathbf{b}_d$ be i.i.d. random variables, each uniformly distributed in the set of primes $\leq X$ not dividing N .

Consider the random lattice

$$\mathcal{L}_{\mathbf{b}, g, y} := \left\{ (e_1, \dots, e_d, f_1, f_2) \in \mathbb{Z}^{d+2} : \left(\prod_{i=1}^d \mathbf{b}_i^{e_i} \right) g^{f_1} y^{f_2} \equiv 1 \pmod{N} \right\}.$$

By Theorem 3.18, with probability $1 + O(d^{-d})$, this lattice has a basis consisting of vectors of Euclidean norm $\ll e^{42d}$.

Computing the discrete logarithm of y with respect to the base g reduces (with a polynomial-time classical algorithm) to computing a short basis for $\mathcal{L}_{\mathbf{b}, g, y}$. To see why this is the case, suppose that we managed to compute a basis $v_1, \dots, v_{d+2}$ for $\mathcal{L}_{\mathbf{b}, g, y}$ with $\max_i \|v_i\|_2 \ll e^{42d}$. It is then easy to find a vector in $\mathcal{L}_{\mathbf{b}, g, y}$ of the form $(0, \dots, 0, x, 1)$ for some integer x (note that such a vector exists since we assume that $y \in \langle g \rangle$). Indeed, this amounts to solving a linear system with integer coefficients. The complexity of solving an integer linear system is polynomial in the dimensions of the matrix and the number of bits of the coefficients, which are both $O(d)$ since $\max_i \|v_i\|_2 \ll e^{42d}$ (see [3]). This yields an integer x such that $g^x \equiv y \pmod{N}$.

The algorithm for solving the discrete logarithm problem thus proceeds as follows.

1. Generate d primes $b_1, \dots, b_d$ independently and uniformly at random in the set of primes $\leq X$ not dividing N . To do this, it suffices to generate d^4 independent random integers $\leq X$. By Lemma 4.3, the required conditions will be satisfied for at least d of those with probability $1 + O(1/N)$. This first step takes polynomial time on a classical computer using the AKS primality test [1].
2. Use the procedure described by Ekerå and Gärtner [7] to obtain a basis for the lattice $\mathcal{L}_{\mathbf{b}, g, y}$. This involves making $O(d)$ calls to a quantum circuit, followed by polynomial-time classical postprocessing. This step is now guaranteed to succeed with probability $\Theta(1)$, using the fact that $\mathcal{L}_{\mathbf{b}, g, y}$ has a basis of vectors of Euclidean norm $\ll e^{42d}$ with probability $1 + O(d^{-d})$ (Theorem 3.18).
3. Compute the discrete logarithm of y in classical polynomial time using this basis, as explained above.

It remains to analyse the gate and space costs of the quantum part of this algorithm.

The only modification needed to the analysis of the quantum circuit in [7] comes from the fact that $b_1, \dots, b_d$ are not quite as small as in [7]. In [7] (and more generally in [17, 18]), the primes b_i are assumed to have $O(\log d)$ bits. In the present situation, we instead have the bound $b_i \leq X = d^{10^3 d}$, that is, each b_i has $O(d \log d)$ bits.

The only place in [7] where the assumption on the size of the b_i 's is used is in [7, Lemma 3]. In turn, the only point in the proof of [7, Lemma 3] where this assumption is needed is when [17, Lemma 5] is invoked (as a black box). The proof of [17, Lemma 5] is given in [17, Section 5], and the size assumption on the b_i 's only comes up in [17, Lemma 12].

The result [17, Lemma 12] essentially¹⁴ states that there is a quantum circuit using $O(d \log^3 d)$ gates and $O(d \log^3 d)$ qubits to perform the computation of $\prod_{i=1}^d a_i^{t_i}$ where $t_i \in \{0, 1\}$ and a_i are integers on $O(\log d)$ bits. This is a special case of Lemma 4.1 (used together with Lemma 4.2) applied with $m = O(\log d)$. In our case, we just apply Lemma 4.1 with $m = O(d \log d)$, together with Lemma 4.2 to convert the classical circuit into a quantum one. This yields a quantum circuit having $O(d^2 \log^3 d)$ gates and $O(d^2 \log^3 d)$ qubits to compute $\prod_{i=1}^d b_i^{t_i}$.

The number of gates of the quantum circuit [17, Lemma 5] is

$$O(d(n \log n + d \log^3 d)) = O(n^{3/2} \log n)$$

(because [17, Lemma 12] is used $O(d)$ times, see [17, Algorithm 5.2] and the surrounding explanations). Note that this is the same as for Regev's algorithm (see [18, p5]). In our case, the number of gates is

$$O(d(n \log n + d^2 \log^3 d)) = O(n^{3/2} \log^3 n).$$

The space optimisations of Ragavan and Vaikuntanathan keep the total number of qubits for their circuit [17, Lemma 5] under

$$O(n \log n + d \log^3 d) = O(n \log n),$$

due to the way the qubits are used and restored in the main loop of [17, Algorithm 5.2]. In our situation, the number of qubits is

$$O(n \log n + d^2 \log^3 d) = O(n \log^3 n).$$

In summary, our final quantum circuit has $O(n^{3/2} \log^3 n)$ gates and $O(n \log^3 n)$ qubits. As noted in [7], the fact that the two elements g and y are not small (g and y can be as large as N , as opposed to the b_i 's) does not affect these complexity bounds. $\square$

Remark 4.4. In the proof of Theorem 1.2, we used the naive Lemma 4.2 to convert a classical circuit into a quantum one. As mentioned in Remark 1.3, it is possible to save a factor of $\log n$ in the space cost for Theorems 1.1 and 1.2 by reusing certain qubits when performing the quantum computation corresponding to Lemma 4.1.

4.3. Factoring integers

In this section, we prove Theorem 1.1.

Lemma 4.5. *Let $N > 1$ be an odd integer with at least two distinct prime factors. Let $\mathbf{x}$ be a random variable, uniformly distributed in $(\mathbb{Z}/N\mathbb{Z})^\times$. Then $\langle \mathbf{x} \rangle$ contains a nontrivial square root of 1 modulo N with probability $\geq 1/2$.*

Proof. This elementary number-theoretic fact is well-known – it was already needed for Shor's algorithm [20]. See [16, Appendix A4.3] for a detailed proof. $\square$

Theorem 1.1. *There is a quantum circuit having $O(n^{3/2} \log^3 n)$ quantum gates and $O(n \log^3 n)$ qubits with the following property. There is a classical randomised polynomial-time algorithm that solves the factoring problem*

Input : a composite integer $N \leq 2^n$

Output : a nontrivial divisor of N

using $O(\sqrt{n})$ calls to this quantum circuit, and succeeds with probability $\Theta(1)$.

¹⁴There are extra details pertaining to the precise use of qubits (e.g., restoring the ancilla qubits to $|0\rangle$), but these will be the same in our setup.

Proof of Theorem 1.1. We can assume that N is odd and not a perfect prime power, as otherwise it is easy to factor N in polynomial time with a classical computer (see [16, Exercise 5.17]).

Let $d := \lceil \sqrt{\log N} \rceil$ and $X := d^{10^3 d}$. The probabilistic algorithm to find a nontrivial divisor of N goes as follows.

1. Generate d primes $b_1, \dots, b_d$ independently and uniformly at random in the set of primes $\leq X$ not dividing N . Sample another integer x uniformly chosen in $(\mathbb{Z}/N\mathbb{Z})^\times$. As in the proof of Theorem 1.2, this can be done classically in polynomial time with probability $1 + O(1/N)$.
2. Use the algorithm of Ekerå and Gärtner [7] to obtain a basis for the lattice

$$\mathcal{L}_{b,x} := \left\{ (e_1, \dots, e_d, f) \in \mathbb{Z}^{d+1} : \left(\prod_{i=1}^d b_i^{e_i} \right) x^f \equiv 1 \pmod{N} \right\}.$$

This involves making $O(d)$ calls to a quantum circuit, followed by polynomial-time classical post-processing. By Theorem 3.18, $\mathcal{L}_{b,x}$ has a basis of vectors of Euclidean norm $\ll e^{42d}$ with probability $1 + O(d^{-d})$, which means that this step is guaranteed to succeed with probability $\Theta(1)$.

3. Using the short basis for $\mathcal{L}_{b,x}$ computed in the previous step, find the vector of the form $(0, \dots, 0, r)$ in $\mathcal{L}_{b,x}$ with $r \geq 1$ as small as possible. This involves solving a linear system with integer coefficients, which can be done efficiently as in the proof of Theorem 1.2. This integer r is the order of x in $(\mathbb{Z}/N\mathbb{Z})^\times$. By Lemma 4.5, with probability $\geq 1/2$, this order r will be even and the element $x^{r/2}$ will be a nontrivial square root of 1 modulo N . Hence, N divides the product $(x^{r/2} - 1)(x^{r/2} + 1)$ but neither term individually, which implies that $\gcd(N, x^{r/2} - \text{mod } N)$ is a nontrivial divisor of N .

The analysis of this algorithm is identical to the corresponding part of the proof of Theorem 1.2. $\square$

A. Bounds for character sums over primes

In this appendix, we prove Proposition 3.9. We start by stating the truncated explicit formula for $L(s, \chi)$.

Lemma A.1. *Let $q \geq 2$. Let χ be a nonprincipal Dirichlet character modulo q . For $x \geq T \geq 2$, we have*

$$\sum_{n \leq x} \chi(n) \Lambda(n) = - \sum_{\substack{\rho = \beta + i\gamma \\ |\gamma| \leq T}} \frac{x^\rho - 1}{\rho} + O\left(\frac{x \log^2(xq)}{T}\right)$$

where the sum runs over all nontrivial zeros ρ of $L(s, \chi)$ with multiplicity.

Proof. This is [14, Theorem 11.3]. $\square$

We will also need the following standard bound for the number of zeros of $L(s, \chi)$ in the critical strip at some height t .

Lemma A.2. *Let $q \geq 1$ and let χ be a Dirichlet character modulo q . Let $t \in \mathbb{R}$. The number of zeros $\rho = \beta + i\gamma$ of $L(s, \chi)$ in the rectangle $0 \leq \beta \leq 1$, $t \leq \gamma \leq t + 1$ is $\ll \log(q(|t| + 2))$, where zeros are counted with multiplicity.*

Proof. This is [15, Theorem 10.17]. $\square$

Proof of Proposition 3.9. By Lemmas A.1 and A.2, choosing $T = x^{1-\alpha} - 1$, we have

$$\sum_{n \leq x} \chi(n) \Lambda(n) \ll \log(qx) \sum_{\substack{t \in \mathbb{Z} \\ |t| \leq x^{1-\alpha} - 1}} \max_{\substack{\rho = \beta + i\gamma \\ |\gamma - t| \leq 1/2}} \left| \frac{x^\rho - 1}{\rho} \right| + x^\alpha \log^2(qx),$$

where the maximum is over all zeros of $L(s, \chi)$ in the specified region. Since $|x^\rho| \leq x^\alpha$ for all zeros $\rho = \beta + i\gamma$ with imaginary part $|\gamma| \leq x^{1-\alpha}$ by our zero-free rectangle assumption, we can bound

$$\left| \frac{x^\rho - 1}{\rho} \right| \ll \frac{x^\alpha}{1 + |\gamma|}$$

(using for example that $\frac{x^\rho - 1}{\rho} = \int_1^x t^{\rho-1} dt$ for $|\gamma| < 1$). Thus, we obtain

$$\sum_{n \leq x} \chi(n) \Lambda(n) \ll \log(qx) x^\alpha \sum_{|t| \leq x} \frac{1}{1 + |t|} + x^\alpha \log^2(qx) \ll x^\alpha \log^2(qx).$$

Discarding perfect prime powers, which contribute $O(x^{1/2} \log x)$, and using partial summation, we get

$$\sum_{p \leq x} \chi(p) \ll \frac{x^\alpha \log^2(qx)}{(1 - \alpha) \log x}.$$

We may assume that $1 - \alpha \geq \frac{1}{\log x}$, as otherwise Proposition 3.9 is trivial. If this is the case, we conclude that

$$\frac{1}{\pi(x)} \sum_{p \leq x} \chi(p) \ll \frac{x^\alpha \log^2(qx)}{\pi(x)} \ll x^{-(1-\alpha)} \log^3(qx)$$

as claimed. □

Acknowledgments. I am indebted to my advisors, Ben Green and James Maynard, for their invaluable advice and support. I also wish to thank Seyoon Ragavan for his kind explanations, and in particular for showing me that the upper bound for the parameters in Regev’s algorithm could be relaxed considerably.

Competing interest. The author has no competing interest to declare.

Financial support. The author is supported by the Mathematical Institute of the University of Oxford, by a Saven European Scholarship and by a Jane Street Graduate Research Fellowship.

References

- [1] M. Agrawal, N. Kayal and N. Saxena, ‘Primes is in P’, *Ann. Math.* **160** (2004), 781–793.
- [2] N. C. Ankeny, ‘The least quadratic non residue’, *Ann. Math.* **55**(1) (1952), 65–72.
- [3] E. H. Bareiss, ‘Sylvester’s identity and multistep integer-preserving Gaussian elimination’, *Math. Comput.* **22**(103) (1968), 565–578.
- [4] D. J. Bernstein, J. Buchmann and E. Dahmen, *Post-quantum Cryptography* (Springer, Berlin, 2009).
- [5] D. A. Burgess, ‘The distribution of quadratic residues and non-residues’, *Mathematika* **4**(2) (1957), 106–112.
- [6] W. Diffie and M. E. Hellman, ‘New directions in cryptography’, *IEEE Trans. Inf. Theory* **22**(6) (1976), 644–654.
- [7] M. Ekerå and J. Gärtner, *Extending Regev’s Factoring Algorithm to Compute Discrete Logarithms*, *Post-quantum cryptography*. Part II, Lecture Notes in Computer Science, vol. 14772 (Springer, Cham, 2024), 211–242.
- [8] C. Gidney, ‘Asymptotically efficient quantum Karatsuba multiplication’, Preprint, 2019, [arXiv:1904.07356](https://arxiv.org/abs/1904.07356).
- [9] A. Granville and K. Soundararajan, ‘Large character sums’, *J. Am. Math. Soc.* **14**(2) (2001), 365–397.
- [10] D. Harvey and J. van der Hoeven, ‘Integer multiplication in time $O(n \log n)$ ’, *Ann. Math.* **193**(2) (2021), 563–617.
- [11] M. Jutila, ‘On Linnik’s constant’, *Math. Scand.* **41**(1) (1977), 45–62.
- [12] G. D. Kahanamoku-Meyer and N. Y. Yao, ‘Fast quantum integer multiplication with zero ancillas’, Preprint, 2024, [arXiv:2403.18006](https://arxiv.org/abs/2403.18006).
- [13] B. S. Kaliski Jr, ‘Targeted Fibonacci exponentiation’, Preprint, 2017, [arXiv:1711.02491](https://arxiv.org/abs/1711.02491).
- [14] D. Koukoulopoulos, *The Distribution of Prime Numbers*, vol. 203 (American Mathematical Society, Providence, RI, 2019).
- [15] H. L. Montgomery and R. C. Vaughan, *Multiplicative Number Theory I: Classical Theory*, Cambridge Studies in Advanced Mathematics (Cambridge University Press, Cambridge, 2006).
- [16] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2010).

- [17] S. Ragavan and V. Vaikuntanathan, *Space-efficient and Noise-robust Quantum Factoring*, Advances in Cryptology – CRYPTO 2024, Part VI, Lecture Notes in Computer Science, vol. 14925 (Springer, Cham, 2024), 107–140.
- [18] O. Regev, ‘An efficient quantum factoring algorithm’, *J. ACM* **72**(1) (2025), Art. 10, 1–13.
- [19] R. L. Rivest, A. Shamir and L. Adleman, ‘A method for obtaining digital signatures and public-key cryptosystems’, *Commun. ACM* **21**(2) (1978), 120–126.
- [20] Peter W. Shor, Algorithms for quantum computation: discrete logarithms and factoring, *Proceedings 35th Annual Symposium on Foundations of Computer Science*, IEEE, 1994, pp. 124–134.
- [21] P. W. Shor, ‘Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer’, *SIAM J. Comput.* **26**(5) (1997), 1484–1509.
- [22] T. Tao and Van H. Vu, *Additive Combinatorics*, vol. 105 (Cambridge University Press, Cambridge, 2006).