# Quantum Algorithms for Discrete Log Require Precise Rotations

JIN-YI CAI, Computer Sciences, University of Wisconsin-Madison, Madison, United States
BEN YOUNG, Computer Sciences, University of Wisconsin-Madison, Madison, United States

Recently, Cai [3] showed that Shor's quantum factoring algorithm fails to factor large integers when algorithm's quantum Fourier transform (QFT) is corrupted by a vanishing level of random noise on the QFT's precise controlled rotation gates. We show that under the same error model, Shor's quantum discrete log algorithm, and its various modifications, fail to compute discrete logs modulo $P$ for a positive density of primes $P$ and a similarly vanishing level of noise. We also show that the same noise level causes Shor's algorithm to fail with probability $1 - o(1)$ to compute discrete logs modulo $P$ for randomly selected primes $P$.

CCS Concepts: • **Theory of computation → Quantum computation theory;**

Additional Key Words and Phrases: Discrete log, shor's algorithm, quantum fourier transform

**ACM Reference Format:**
Jin-Yi Cai and Ben Young. 2025. Quantum Algorithms for Discrete Log Require Precise Rotations. *ACM Trans. Quantum Comput.* 6, 3, Article 21 (June 2025), 18 pages. https://doi.org/10.1145/3736421

## 1 Introduction

*The discrete log problem and the QFT.* The **discrete log problem** (**DLP**) over $\mathbb{Z}_P^*$ (the multiplicative group of integers mod $P$) is defined as follows: given prime number $P$, nonzero natural number $g < P$ such that $g^0, g^1, g^2, \ldots, g^{P-2}$ generate all nonzero integers mod $P$, and an integer $y$ that is nonzero mod $P$, find the unique value $0 \leq d \leq P-2$ such that $g^d \equiv y \bmod P$. This $d$ is called the *discrete log value* of $y$ mod $P$ . The assumed hardness of this problem underlies the Diffie–Hellman key exchange [6], a widely-used cryptographic protocol. The importance of the DLP, and the problem of factoring integers, to modern cryptography (and the lack of any polynomial-time classical algorithm for these problems) make Shor's polynomial time quantum algorithms for these two problems [30, 31] two of the most famous results in quantum computing.

Shor's algorithms for the DLP and the factoring problem consist of classical pre- and post-processing and a **quantum Fourier transform** (**QFT**) exploiting the problems' underlying periodicity. The QFT is a central tool in quantum computing, forming the basis of a wide variety of quantum algorithms promising (if implemented exactly) an exponential speedup over their best known classical counterparts. The $n$-qubit QFT, or **quantum fast Fourier transform** (**QFFT**) [28]—the version used by Shor and analyzed in this work—is easily expressible as a quantum

Authors' Contact Information: Jin-Yi Cai, Computer Sciences, University of Wisconsin-Madison, Madison, Wisconsin, United States; email: jyc@cs.wisc.edu; Ben Young (Corresponding author), Computer Sciences, University of Wisconsin-Madison, Madison, Wisconsin, United States; e-mail: benyoung@cs.wisc.edu.


ACM 2643-6817/2025/06-ART21
https://doi.org/10.1145/3736421

circuit composed mostly of controlled-$R_k$ gates for $k = 2, \ldots, n$, where $R_k = \left[\begin{smallmatrix} 1 & 0 \\ 0 & e^{2\pi i/2^k} \end{smallmatrix}\right]$ is the single-qubit rotation about the $Z$-axis by angle $2\pi/2^k$.

*Noisy rotation gates.* The precision required to exactly implement the controlled-$R_k$ gates of the smallest angles $2\pi/2^k$ increases exponentially with the size of the input to the quantum algorithms presented in References [30, 31] using the QFT. We are far from the first to raise concern over this exponential dependency. In Reference [5], Coppersmith introduced an approximate version of the QFFT that simply omits the controlled-$R_k$ gates from the circuit for every $k \geq b$, where $b$ is a parameter much smaller than (but still increasing with) $n$. Coppersmith shows that Shor's factoring algorithm loses very little efficacy when its exact QFT is replaced by an approximate QFT, even for $b$ approximately on the order of $\log n$. Fowler and Hollenberg [12, 13] and Nam and Blümel [23] support this conclusion with numerical simulation for small $n$ and heuristic approximations for large $n$, showing that moderate sized values of $b$ may suffice to factor integers on the scale of those used in some current RSA schemes.

However, while these analyses offer *practical* support for the robustness of Shor's factoring algorithm on small inputs, they both suggest that, for fixed $b$, the success probability of Shor's algorithm decays exponentially in the number of bits of the integer to be factored. From a *theoretical* perspective, this suggests that, without arbitrarily precise quantum error correction, Shor's factoring algorithm will fail on sufficiently large inputs.

In Reference [24] (see also References [25, 26]), Nam and Blümel study the performance of the QFFT circuit when, instead of being removed entirely, the small controlled-$R_k$ gates are subject to four types of random noise. Again, their analysis suggests that the QFFT could remain effective on practical scales when subject to noise, but, for a fixed noise level, its performance decays exponentially with the size of the input.

In Reference [3], the first author gave, to our knowledge, the first rigorous proof under any error model that Shor's factoring algorithm fails when a vanishing level of noise is present on sufficiently large inputs. The particular error model applied in Reference [3] is as follows: for $k \geq b$, we replace each controlled-$R_k$ gate, which rotates about $Z$ by angle $2\pi/2^k$, with a noisy rotation about $Z$ by angle $2\pi(1 + \epsilon r)/2^k$, where $\epsilon$ is a global fixed noise level and $r$ is a Gaussian random variable. The noise variables $r$ on each gate are independent, and the same random perturbation on a gate applies to every state in a superposition to which the gate is applied. Under this error model, it is proved in [3] that, if $b + \log_2(1/\epsilon) < \frac{1}{3}\log_2(n) - c$ for some constant $c > 0$, then Shor's factoring algorithm fails to factor $n$-bit integers $pq$, both for random fixed-length primes $p, q$ with probability $1 - o(1)$, and for $p, q$ taken from a specifically defined set of primes of positive density.

*Noise model.* This work studies the effectiveness of Shor's discrete log algorithm under the error model of Cai [3] described in the previous paragraph. The origins of this error model in the context of the QFT go back to works of Coppersmith [5], Barenco, Ekert, Suominen, and Törmä [1], and Fowler and Hollenberg [12]. Coppersmith's banded QFFT, as discussed above, assumed exact implementations of all controlled-$R_k$ gates for $k < b$. Barenco et al. perform an empirical study of Coppersmith's banded QFFT when each remaining controlled-$R_k$ gate is not implemented exactly, but, similarly to our noise model, is subject to an additional rotation by a normally-distributed random angle. This noise models the effect of decoherence, where unwanted interactions with the surrounding environment induce random phase fluctuations. Since controlled-$R_k$ gates apply to two qubits and have more complex, time-consuming implementations in terms of potentially many basis gates, they are especially susceptible to decoherence. Fowler and Hollenberg analyze the performance of Shor's algorithm under the noise model of Barenco et al., and provide another justification for this model: any realization of the QFT that admits quantum error correction must

use only a finite set of gates with fault-tolerant implementations (e.g., the Hadamard and $\pi/8$ gates suffice to approximate any single-qubit rotation). Even with no noise present, arbitrary rotations cannot be constructed exactly from the finite fault-tolerant gate set; Fowler and Hollenberg use the noise model of Barenco et al. to simulate this discrepancy between exact and approximate rotations. Wei, Li, Hu, and Nori [32] study the effect of *dynamical phases*, coherent errors caused by operational delays between successive gates in a quantum circuit, on the performance of Shor's factoring algorithm. Their error model, like ours and those discussed above, considers noise affecting a qubit's phase, but does not apply noise to or between individual gates in the QFT. Nevertheless, the net effect is similar to (8) below: the expression for the probability of measuring a desired state is corrupted by a random phase angle added to each point in an otherwise carefully aligned sequence of points on the unit circle.

The error models of Barenco et al. and Fowler and Hollenberg are *absolute*, meaning the noise magnitude is independent of the angle of the controlled-rotation gate to which the noise is applied. Hence, for a fixed noise magnitude, sufficiently small angles will be completely overwhelmed by noise at the level of each individual gate. Nam and Blümel [24] introduced the *relative* error model, adopted in this work, in which the noise magnitude scales with the magnitude of the angle of rotation – that is, a noisy rotation applies an angle of $2\pi/2^k + 2\pi\epsilon r/2^k$ (cf. the description of our noise model above) instead of an angle with an expression of the form $2\pi/2^k + \epsilon r$. A negative result on the effectiveness of Shor's algorithm in the presence of noise is stronger in this relative error model, compared with the absolute model, because, in the relative model, individual gates retain some of their effectiveness even for arbitrarily small rotations. Our noise model is identical to the relative, uncorrelated (i.e., noise on distinct gates is modeled by distinct independent random variables) model of Nam and Blümel, but with the additional banding effect, in which noise only applies to gates $R_k$ with $k \geq b$. Again, our negative result is stronger under the banded error model, in which possibly less, but never more, noise is present (by setting $b = 0$, we recover Nam and Blümel's model). Nam and Blümel use relative, uncorrelated noise (and absolute and/or correlated noise) to model coherent errors in the QFT—-called *static defects*—arising from manufacturing errors such as improperly calibrated magnetic fields in an NMR quantum computer, or unwanted static electric fields in a trapped-ion quantum computer (see also Reference [25]).

*Other discrete log algorithms.* Many modifications to and extensions of Shor's original discrete log algorithm have been proposed. These modifications and extensions differ in classical pre- and post-processing, the algebraic structures over which the DLP is solved, and in certain details of the quantum part of the algorithm, but all apply some form of QFT for the purpose of period-finding. Indeed, since our error model and analysis are specific to the structure of the QFT, not the context in which the QFT is used, they apply to any quantum algorithm built on top of the QFT. The DLP over any cyclic group is an instance of the more general **Abelian Hidden Subgroup Problem** (**AHSP**), so efficient quantum algorithms for the AHSP also solve the DLP. In particular, the AHSP admits a QFT-based phase estimation algorithm [21], which was modified in Reference [18] to solve the DLP by computing the discrete log value one bit at a time. Quantum AHSP algorithms apply the QFT over general cyclic groups [20], which is not as easily implemented as the QFFT used by Shor, but has efficient quantum approximations [16, 19, 22]. However, at their cores, these approximations, like the QFFT, use circuits composed of controlled Z-rotation gates. Hence these approximate QFTs, and the algorithms relying on them, are similarly susceptible to noise.

The QFT has also been applied to solve the DLP over general groups (as long as the group operation can be computed efficiently) [2] and over other algebraic structures such as semigroups [4], and to give specialized algorithms for the DLP over elliptic curve groups [28], hyperelliptic curve groups [17], and for *short* discrete logs (in which the discrete log value is much smaller

than the group order) [9]. Additional algorithms have been proposed in which the size of the quantum circuit or number of quantum operations performed is reduced by a constant factor [7] or is asymptotically smaller [8] than in Shor's original algorithm, but more runs of the quantum circuit and more complex classical postprocessing are required.

Any quantum DLP algorithm using the QFFT can instead use the *semiclassical Fourier transform* [15] in which the QFFT's two-qubit controlled-$R_k$ gates are replaced by one-qubit classically-controlled $R_k$ gates; this can reduce the total number of qubits used by the algorithm [28]. See also [29]. We reiterate that all these extensions and modifications use a QFT circuit composed of (controlled) quantum $Z$-rotation gates (whether classically controlled or otherwise) to exploit the same periodicity of the DLP, hence are all susceptible to the same noise affecting Shor's original algorithm, and our analysis in this article applies.

*Noisy QFTs and the discrete log problem.* In this work, we show that there exists a constant $0 < c < 1$ such that, if the controlled rotation gates in the QFFT are subject to the error model from [3] with

$$b + \log_2(1/\epsilon) \leq \frac{1-c}{2}\log_2(n) - \Theta(1), \tag{1}$$

then Shor's algorithm [30, 31] for the DLP fails to find the discrete log modulo $P$, a prime of binary length $n$, of all but an exponentially small fraction of inputs $y \in \mathbb{Z}_P^*$, where $P$ is taken from a positive density of primes (Theorem 3.2), or $P$ is chosen uniformly at random with probability $1 - o(1)$ (Theorem 3.3). Although our main conclusions in Theorems 3.2 and 3.3 are similar to those of Reference [3], the proof is different at a technical level.

A key component of our proof is a technical lemma upper bounding the expected value of a sum of terms with factors of the form $e^{(2\pi i/a)\Sigma_k}$, where $\Sigma_k$ is, with high probability, a sum of $\Theta(n)$ (where $n$ is the binary length of the input) independent random noise variables $r$. Our model explicitly assumes each noise variable is normally distributed, and uses the fact that $\Sigma_k$ is then normally distributed. This is for the convenience of proof presentation (as is the case in Reference [3]). Our proof can be easily adapted if the noise is drawn from any distribution of bounded variance, by applying the central limit theorem.

Both analyses—for the DLP algorithm in this article and for Shor's factoring algorithm in Reference [3]—come down to arguing about the distribution of bits in the binary representations of certain integers to show that a sufficient number of random noise variables are included in the expression for the probability of measuring desired states. The theorems in Reference [3] use a bound on the power of 2 in the prime factorizations of certain integers appearing in the algorithm, but we employ a different technique featuring a counting argument instead.

The proofs of Theorems 3.2 and 3.3 show that when the algorithm measures the state after applying a noisy QFT, the probability that the measurement produces a state in some set of directly useful states is exponentially small. However, Shor notes in Reference [31] that an algorithm could still feasibly extract the discrete log value from states in a slightly larger set. Hence, we extend the specified sets of states in the proofs of Theorems 3.2, 3.3 to naturally enlarged sets of states that are "polynomially close" to states from which Shor's DLP algorithm explicity specifies it can extract the solution. This makes our results more robust, since they also preclude the success of any more flexible or slightly modified algorithms capable of extracting the solution from states that are close enough to the original algorithms' desired states. We give a natural definition of "polynomially close" in Section 4 and show that there is an exponentially small probability that the post-noisy-QFT measurement produces a state polynomially close to any state Shor defines in Reference [31] to be useful.

For fixed $b$ and $\epsilon$, once $n$ exceeds the product of some exponential expression in $b$ and polynomial expression in $1/\epsilon$, (1) is satisfied, hence Shor's algorithm for DLP fails on sufficiently large $n$ when noise exceeds this level. Section 5.2 presents the results of numerical simulations estimating the exact noise threshold after which this failure occurs. To place this threshold in the context of current quantum computers, Section 5.1 presents the results of quantum hardware experiments studying the efficacy of precise rotations such as $R_k$. We emphasize that it is still plausible that quantum computers can be built that can efficiently solve the DLP modulo integers on the scale of those used in current cryptosystems. However, our proof shows that Shor's algorithm for the DLP must apply *arbitrarily precise* controlled rotations to handle *arbitrarily large* inputs. In particular, the algorithm will fail for sufficiently large inputs on quantum computers lacking arbitrarily precise quantum error correction.

## 2 Preliminaries

For integers $a < b$, let $[a, b] = \{a, a+1, \ldots, b\}$ and $[a, b) = [a, b] \setminus \{b\}$.

We will use the following technical lemma (essentially a restatement of [3, Lemma 2]) bounding the squared norm of sums of unit norm random variables. It will be used to upper bound the probability of a quantum algorithm measuring a desired state.

Lemma 2.1. *For $a \in \mathbb{R}^+$, let $\omega_a = e^{2\pi i/a}$. Let $\{r_i \mid i \in [n]\}$ be i.i.d. Gaussian random variables drawn from $N(0,1)$, and let $\{J_k \subset [n] \mid k \in [K]\}$ be a finite collection of sets. Assume all except at most a fraction $\zeta$ of pairwise symmetric differences $J_k \Delta J_{k'}$ have cardinality at least $a^2 t$ for $k \neq k'$. Let $\Sigma_k = \sum_{i \in J_k} r_i$ and $\varphi_k \in [0, 2\pi)$. Then*

$$\mathrm{E}\left[\left|\omega_a^{\varphi_1+\Sigma_1} + \omega_a^{\varphi_2+\Sigma_2} + \ldots + \omega_a^{\varphi_K+\Sigma_K}\right|^2\right] \leq K + 2\zeta\binom{K}{2} + 2\binom{K}{2}e^{-2\pi^2 t}.$$

## 3 The Discrete Log Algorithm with Noise

### 3.1 Shor's Quantum Discrete Log Algorithm and the Quantum Fourier Transform

The setup of the DLP is as follows: given prime $P$, base $g \in \mathbb{Z}_P^*$ of order $P-1$, and $y \in \mathbb{Z}_P^*$, find the $0 \leq d \leq P-2$ such that $g^d \equiv y \bmod P$. Since $g$ is a generator of $\mathbb{Z}_P^*$, there is a one-to-one correspondence between $d$ values and input $y$ values. Suppose $P$ is an $n$-bit integer, so $2^{n-1} \leq P < 2^n$. We encode integers $0 \leq x < 2^n$ as $n$-qubit quantum states $|x\rangle = \left|x^{[0]}x^{[1]}\ldots x^{[n-1]}\right\rangle$, where $x^{[j]}$ is the value of the $j$th bit in the $n$-bit binary representation of $x$. In this section, we give an overview of Shor's quantum algorithm [31] to find $d$. We begin by preparing the state

$$\frac{1}{P-1}\sum_{u=0}^{P-2}\sum_{k=0}^{P-2}|u\rangle|k\rangle\left|g^u y^{-k} \bmod P\right\rangle = \frac{1}{P-1}\sum_{u=0}^{P-2}\sum_{k=0}^{P-2}|u\rangle|k\rangle\left|g^{u-dk} \bmod P\right\rangle, \tag{2}$$

in three $n$-qubit registers. Now the algorithm applies $n$-qubit QFTs to the first and second registers. The $n$-qubit QFT $F_{2^n}$ sends $|x\rangle$ to

$$F_{2^n}|x\rangle = \frac{1}{2^{n/2}}\sum_{v=0}^{2^n-1}\exp\left(2\pi i\frac{xv}{2^n}\right)|v\rangle. \tag{3}$$

Hence the state becomes

$$\frac{1}{2^n(P-1)}\sum_{u,k=0}^{P-2}\sum_{v,w=0}^{2^n-1}\exp\left(2\pi i\frac{uv+kw}{2^n}\right)|v\rangle|w\rangle\left|g^{u-dk} \bmod P\right\rangle. \tag{4}$$

Now, we measure the three registers. Let $0 \leq u^* \leq P-2$. For each $0 \leq k \leq P-2$, $u = dk + u^* \bmod (P-1)$ is the unique integer in the range $[0, P-2]$ satisfying $u - dk \equiv u^* \bmod (P-1)$. So,

letting $u_k = dk + u^* \bmod (P-1)$, the probability of obtaining $|v\rangle\,|w\rangle\,\left|g^{u^*}\right\rangle$ upon measuring the three registers is

$$\frac{1}{2^{2n}(P-1)^2}\left|\sum_{k=0}^{P-2}\exp\left(2\pi i\frac{u_k v + kw}{2^n}\right)\right|^2. \tag{5}$$

Certain useful Fourier peaks have a high probability of being measured. Let $\{z\}_{2^n}$ be the residue of $z \bmod 2^n$ in the range $-2^{n-1} < \{z\}_{2^n} \le 2^{n-1}$. Shor [31] shows that the probability of measuring some $|v\rangle\,|w\rangle$ in the first two registers with $v$ and $w$ satisfying

$$\left|\left\{vd + w - \frac{d}{P-1}\{v(P-1)\}_{2^n}\right\}_{2^n}\right| \le \frac{1}{2}, \tag{6}$$

and

$$|\{v(P-1)\}_{2^n}| < \frac{2^n}{12}, \tag{7}$$

is at least a positive constant, and that we can extract the discrete log value $d$ from such a pair $(v, w)$ with high probability.

### 3.2 The Noisy Quantum Fourier Transform

The exact $n$-qubit QFT $F_{2^n}$ is implemented using a quantum circuit composed of Hadamard gates and controlled-$R_k$ gates for $2 \le k \le n$, where $R_k$ is the single-qubit rotation about $Z$ by angle $2\pi/2^k$:

$$R_k = \begin{bmatrix} 1 & 0 \\ 0 & e^{2\pi i/2^k} \end{bmatrix}.$$

See [27, Section 5.1] for an explicit description of the circuit implementing $F_{2^n}$. We consider the scenario where there is some $b < n$ such that every controlled-$R_k$ gate for $k \ge b$ is accompanied by a small relative additive error. More precisely, we replace each controlled-$R_k$ gate in the circuit implementing $F_{2^n}$ by a controlled-$\widetilde{R_k}$ gate, where

$$\widetilde{R_k} = \begin{bmatrix} 1 & 0 \\ 0 & e^{2\pi i(1+\epsilon r)/2^k} \end{bmatrix}$$

is a $Z$-rotation of angle $2\pi(1+\epsilon r)/2^k$, where $r$ is an independent Gaussian random variable drawn from $N(0,1)$ and $\epsilon$ is a global parameter controlling the magnitude of the noise. With the exact rotations $R_k$, the $F_{2^n}$ circuit directly realizes the transformation

$$\begin{aligned} |x\rangle \mapsto \frac{1}{2^{n/2}} &\left(|0\rangle + \exp\left(2\pi i 0.x^{[n-1]}x^{[n-2]}\ldots x^{[0]}\right)|1\rangle\right) \\ &\left(|0\rangle + \exp\left(2\pi i 0.x^{[n-2]}\ldots x^{[0]}\right)|1\rangle\right) \\ &\vdots \\ &\left(|0\rangle + \exp\left(2\pi i 0.x^{[0]}\right)|1\rangle\right). \end{aligned}$$

After using swaps to reverse the order of the qubits, one can check that this operation is equivalent to the original expression for $F_{2^n}$ in (3) (see [27, Section 5.1]). When each controlled-$R_k$ gate is replaced by a controlled-$\widetilde{R_k}$ gate for $k \ge b$, the noisy circuit implements a transformation we call

$\widetilde{F_{2^n}}$, where

$$\widetilde{F_{2^n}}\,|x\rangle = \left(|0\rangle + \exp\left(2\pi i\left(0.x^{[n-1]}x^{[n-2]}\ldots x^{[0]} + \frac{\epsilon}{2^b}\left[\frac{r_0^{(0)}x^{[n-b]}}{2^0} + \ldots + \frac{r_{n-b}^{(0)}x^{[0]}}{2^{n-b}}\right]\right)\right)|1\rangle\right)$$
$$\left(|0\rangle + \exp\left(2\pi i\left(0.x^{[n-2]}\ldots x^{[0]} + \frac{\epsilon}{2^b}\left[\frac{r_0^{(1)}x^{[n-b-1]}}{2^0} + \ldots + \frac{r_{n-b-1}^{(1)}x^{[0]}}{2^{n-b-1}}\right]\right)\right)|1\rangle\right)$$
$$\vdots$$
$$\left(|0\rangle + \exp\left(2\pi i\left(0.x^{[b-1]}\ldots x^{[0]} + \frac{\epsilon}{2^b}r_0^{(n-b)}x^{[0]}\right)\right)|1\rangle\right)$$
$$\left(|0\rangle + \exp\left(2\pi i 0.x^{[b-2]}\ldots x^{[0]}\right)|1\rangle\right)$$
$$\vdots$$
$$\left(|0\rangle + \exp\left(2\pi i 0.x^{[0]}\right)|1\rangle\right)$$

and $r_0^{(0)},\ldots,r_{n-b}^{(0)},r_0^{(1)},\ldots,r_{n-b-1}^{(1)},\ldots,r_0^{(n-b)}$ are i.i.d. random variables drawn from $N(0,1)$.

### 3.3 Analysis Over a Positive Density of Primes

In this section, we show that there are a positive density of primes $P$ for which Shor's discrete log algorithm, as described in Section 3.1, has an exponentially small probability of solving the DLP over $\mathbb{Z}_P^*$ when forced to use the noisy QFT $\widetilde{F_{2^n}}$ in place of the exact transform $F_{2^n}$.

We begin with a result from number theory. For integer $x$, let $\mathcal{P}^+(x)$ be the largest prime dividing $x$.

Theorem 3.1 (Fouvry [11]). *There exist constants $c > 0$ and $n_0 > 0$ such that for all $x > n_0$,*

$$|\{\text{prime } p < x \mid \mathcal{P}^+(p-1) > p^{2/3}\}| \geq c\frac{x}{\log x}.$$

Since the number of primes at most $x$ is asymptotically $\frac{x}{\log x}$, Fouvry's theorem states that the set of primes $p$ satisfying $\mathcal{P}^+(p-1) > p^{2/3}$ has positive density in the set of all primes. Throughout, we assume that there is a $1/2 < c_1 < 1$ such that $P-1$ has a prime factor $\mathcal{P}^+(P-1) > P^{c_1}$. By Theorem 3.1, there is a set of primes $P$ of a positive density with $c_1 = 2/3$. However, we carry out the proof with a generic value $c_1$.

Recall that applying $F_{2^n}$ to the first two registers took the state in (2) to the state in (4). Suppose we instead apply $\widetilde{F_{2^n}}$ to the first two registers of the state in (2). Each noisy QFT comes with its own set of independent r.v.s, labeled as $r_\cdot^{(\cdot)}$ and $\rho_\cdot^{(\cdot)}$, respectively. We obtain the state

$$\frac{1}{2^n(P-1)}\cdot$$
$$\sum_{u=0}^{P-2}\sum_{k=0}^{P-2}\left(|0\rangle + \exp\left(2\pi i\left(0.u^{[n-1]}u^{[n-2]}\ldots u^{[0]} + \frac{\epsilon}{2^b}\left[\frac{r_0^{(0)}u^{[n-b]}}{2^0} + \ldots + \frac{r_{n-b}^{(0)}u^{[0]}}{2^{n-b}}\right]\right)\right)|1\rangle\right)$$
$$\left(|0\rangle + \exp\left(2\pi i\left(0.u^{[n-2]}\ldots u^{[0]} + \frac{\epsilon}{2^b}\left[\frac{r_0^{(1)}u^{[n-b-1]}}{2^0} + \ldots + \frac{r_{n-b-1}^{(1)}u^{[0]}}{2^{n-b-1}}\right]\right)\right)|1\rangle\right)$$
$$\vdots$$

$$\left(|0\rangle + \exp\left(2\pi i\left(0.u^{[b-1]}\dots u^{[0]} + \frac{\epsilon}{2^b} r_0^{(n-b)} u^{[0]}\right)\right)|1\rangle\right)\dots\left(|0\rangle + \exp\left(2\pi i 0.u^{[0]}\right)|1\rangle\right)$$
$$\left(|0\rangle + \exp\left(2\pi i\left(0.k^{[n-1]}k^{[n-2]}\dots k^{[0]} + \frac{\epsilon}{2^b}\left[\frac{\rho_0^{(0)}k^{[n-b]}}{2^0} + \dots + \frac{\rho_{n-b}^{(0)}k^{[0]}}{2^{n-b}}\right]\right)\right)|1\rangle\right)$$
$$\left(|0\rangle + \exp\left(2\pi i\left(0.k^{[n-2]}\dots k^{[0]} + \frac{\epsilon}{2^b}\left[\frac{\rho_0^{(1)}k^{[n-b-1]}}{2^0} + \dots + \frac{\rho_{n-b-1}^{(1)}k^{[0]}}{2^{n-b-1}}\right]\right)\right)|1\rangle\right)$$
$$\vdots$$
$$\left(|0\rangle + \exp\left(2\pi i\left(0.k^{[b-1]}\dots k^{[0]} + \frac{\epsilon}{2^b}\rho_0^{(n-b)}k^{[0]}\right)\right)|1\rangle\right)\dots\left(|0\rangle + \exp\left(2\pi i 0.k^{[0]}\right)|1\rangle\right)$$
$$\left|g^{u-dk} \bmod P\right\rangle.$$

Now, we measure the three registers. Instead of the probability expression in (5), the probability of measuring $|v\rangle\,|w\rangle\left|g^{u^*}\right\rangle$ after the noisy transform is

$$p(v,w,g^{u^*}) = \frac{1}{2^{2n}(P-1)^2}\left|\sum_{k=0}^{P-2}\exp\left(2\pi i\left[\right.\right.\right.$$
$$\sum_{t=0}^{n-1} v^{[t]}\left(0.u_k^{[n-t-1]}\dots u_k^{[0]}\right) + \sum_{\tau=0}^{n-1} w^{[\tau]}\left(0.k^{[n-\tau-1]}\dots k^{[0]}\right)$$
$$+\frac{\epsilon}{2^b}\left\{v^{[0]}\left(\frac{r_0^{(0)}u_k^{[n-b]}}{2^0} + \dots + \frac{r_{n-b}^{(0)}u_k^{[0]}}{2^{n-b}}\right) + v^{[1]}\left(\frac{r_0^{(1)}u_k^{[n-b-1]}}{2^0} + \dots + \frac{r_{n-b-1}^{(1)}u_k^{[0]}}{2^{n-b-1}}\right) + \dots\right.$$
$$\left.+ v^{[n-b]}\frac{r_0^{(n-b)}u_k^{[0]}}{2^0}\right\}$$
$$+\frac{\epsilon}{2^b}\left\{w^{[0]}\left(\frac{\rho_0^{(0)}k^{[n-b]}}{2^0} + \dots + \frac{\rho_{n-b}^{(0)}k^{[0]}}{2^{n-b}}\right) + w^{[1]}\left(\frac{\rho_0^{(1)}k^{[n-b-1]}}{2^0} + \dots + \frac{\rho_{n-b-1}^{(1)}k^{[0]}}{2^{n-b-1}}\right) + \dots\right.$$
$$\left.\left.\left.+ w^{[n-b]}\frac{\rho_0^{(n-b)}k^{[0]}}{2^0}\right\}\right]\right)\right|^2. \tag{8}$$

We will show that the probability of measuring a state satisfying (6) and (7) (states from which Shor's algorithm extracts the discrete log value), which was at least a positive constant in the noise-free case, is exponentially small, provided $n$ is sufficiently large compared with $b$ and $1/\epsilon$. Let

$$G = \{(v,w) \mid 0 \le v, w < 2^n, v \text{ and } w \text{ satisfy (6) and (7)}\}$$

and let $\pi_1 : \mathbb{Z}\times\mathbb{Z} \to \mathbb{Z}$ be the projection onto the first coordinate. So for $X \subset \mathbb{Z}\times\mathbb{Z}$ we have

$$\pi_1(X) = \{a \mid (a,b) \in X\}.$$

As Shor notes, for any $v$, there is exactly one $0 \le w < 2^n$ satisfying (6), so

$$|G| = |\pi_1(G)| \le 2^n. \tag{9}$$

Shor also notes that $|\pi_1(G)| \ge 2^n/12$.

We will ignore the $\rho_{\cdot}^{(\cdot)}$ terms in (8) entirely. When we apply Lemma 2.1, these $\rho_{\cdot}^{(\cdot)}$ are incorporated into the terms $\varphi_k$ (here and below, reintroducing such terms only increases the noise the algorithm must overcome). In fact, for $k \in [0, P-2]$, we consider only the error terms

$$\begin{aligned}&\frac{\epsilon}{2^b}\left(v^{[0]}r_0^{(0)}u_k^{[n-b]} + v^{[1]}r_0^{(1)}u_k^{[n-b-1]} + \ldots + v^{[n-b]}r_0^{(n-b)}u_k^{[0]}\right)\\ &= \frac{\epsilon}{2^b}\sum_{j=0}^{n-b} v^{[j]}u_k^{[n-b-j]}r_0^{(j)} = \frac{\epsilon}{2^b}\sum_{j\in J_k} r_0^{(j)}\end{aligned} \tag{10}$$

where

$$J_k = \{0 \le j \le n-b \mid v^{[j]}u_k^{[n-b-j]} = 1\}.$$

For any $0 < \delta < 1/2$ and $\ell \ge 1$, we have [10, Lemma 16.19]

$$\sum_{i=0}^{\lfloor \delta \ell \rfloor}\binom{\ell}{i} \le 2^{H_2(\delta)\ell},$$

where $H_2$ is the binary entropy function $H_2(\delta) = -\delta\log_2(\delta) - (1-\delta)\log_2(1-\delta)$. Therefore the number of 0-1 sequences of length $\ell$ with at most $\delta\ell$ one bits is at most $2^{H_2(\delta)\ell}$. For $v \in \pi_1(G)$, consider the sequence of bits $(v^{[0]}, v^{[1]}, \ldots, v^{[n-b]})$ of length $\ell = n-b+1$. Fix some $0 < \delta_1 < 1/2$; there are no more than $2^{H_2(\delta_1)(n-b+1)}$ 0-1 sequences of length $n-b+1$ with fewer than $\delta_1(n-b) < \delta_1(n-b+1)$ one bits. Then, there are at most $2^{b-1}\cdot 2^{H_2(\delta_1)(n-b+1)}$ 0-1 sequences of length $n$ with fewer than $\delta_1(n-b)$ one bits in positions $0, 1, \ldots, n-b$. Let

$$S_v = \{s : 0 \le s \le n-b, v^{[s]} = 1\} \quad \text{and} \quad G' = \{(v,w) \in G \mid |S_v| \ge \delta_1(n-b)\}.$$

The above argument shows that $|\pi_1(G) \setminus \pi_1(G')| \le 2^{b-1}\cdot 2^{H_2(\delta_1)(n-b+1)}$. Therefore, since $|\pi_1(G)| \ge 2^n/12$, the proportion of $v \in \pi_1(G)$ that are not in $\pi_1(G')$ is

$$\frac{|\pi_1(G) \setminus \pi_1(G')|}{|\pi_1(G)|} \le O(2^{(1-H_2(\delta_1))b}\cdot 2^{-(1-H_2(\delta_1))n}) = n^{O(1)}\cdot 2^{-(1-H_2(\delta_1))n} \tag{11}$$

for $b = O(\log n)$, which is exponentially small in $n$, as $0 < H_2(\delta_1) < 1$. Thus, the proportion of $v \in \pi_1(G)$ that are in $\pi_1(G')$ is exponentially close to 1.

We next use Lemma 2.1 to upper bound the probability of measuring any fixed $v \in \pi_1(G')$. First, define

$$S'_v = \{n-b-j \mid j \in S_v\}.$$

Then, for $k, k' \in [0, P-2]$,

$$J_k \Delta J_{k'} = \{s \in S'_v \mid u_k^{[s]} \oplus u_{k'}^{[s]} = 1\}. \tag{12}$$

Fix $k' \in [0, P-2]$. To apply Lemma 2.1, we aim at showing that, for most $k \in [0, P-2]$, $|J_k \Delta J_{k'}|$ is linear in $n$. Recall that $u_k = u^* + dk \bmod (P-1)$, for $k \in \{0, \ldots, P-2\}$, and recall our assumption that $P-1$ has an exponentially large prime factor $Q = \mathcal{P}^+(P-1) > P^{c_1}$. Then all but at most an exponentially small fraction $1/Q$ of $d \in \{0, \ldots, P-2\}$ have (additive) order $(P-1)/\gcd(d, P-1) \ge Q$ in $\mathbb{Z}_{P-1}$. For such $P$ and $d$, there are at least $P^{c_1} = \Omega(2^{c_1 n})$ distinct values $u_k$. Since inputs $y$ are in one-to-one correspondence with values $d$, for a positive density of primes $P$, we have

$$\Pr[d \text{ has order at least } Q \text{ in } \mathbb{Z}_{P-1}] \ge 1 - \frac{1}{Q}, \tag{13}$$

a probability exponentially close to 1, where the probability is over uniformly sampled input $y \in \{0, \ldots, P-2\}$. Until further notice, we assume $d$ has order at least $Q$ in $\mathbb{Z}_{P-1}$, hence that there are at least $P^{c_1} = \Omega(2^{c_1 n})$ distinct values $u_k$.

In light of (12), define integer $u_k \oplus u_{k'}$ so that $(u_k \oplus u_{k'})^{[s]} = u_k^{[s]} \oplus u_{k'}^{[s]}$. Consider the sequence of bits of $u_k \oplus u_{k'}$ at bit positions corresponding to indices in $S'_v$. Again applying the entropy bound, the number of 0-1 sequences of length $\ell = |S_v|$ with fewer than $\delta_2|S_v|$ one bits, for any $0 < \delta_2 < 1/2$, is $O(2^{H_2(\delta_2)|S_v|})$, so the total number of 0-1 sequences of length $n$ with fewer than $\delta_2|S_v|$ one bits at positions indexed by $S'_v$ is $O(2^{n-(1-H_2(\delta_2))|S_v|})$. For distinct $u_{k_1}, u_{k_2}$, we have $u_{k_1} \oplus u_{k'} \neq u_{k_2} \oplus u_{k'}$. As $k$ ranges in $[0, P-2]$, $u_k$ cycles through all distinct values $(P-1)/\gcd(d, P-1)$ times, achieving each distinct value exactly $\gcd(d, P-1)$ times. Thus, the proportion of $k \in [0, P-2]$ for which $u_k \oplus u_{k'}$ (viewed as a bit sequence of length $n$) has fewer than $\delta_2|S_v| = \Omega(n)$ one bits among those bits indexed by $S'_v$ equals the proportion of the $\Omega(2^{c_1 n})$ distinct values $u_k$ for which $u_k \oplus u_{k'}$ satisfies this property. By the discussion earlier in this paragraph, this proportion is

$$O\left(2^{n-(1-H_2(\delta_2))|S_v|}/2^{c_1 n}\right) = O\left(2^{(1-c_1)n-(1-H_2(\delta_2))|S_v|}\right). \tag{14}$$

For $v \in \pi_1(G')$, we have $|S_v| \geq \delta_1(n-b)$, so the expression (14) is, up to constant factors, at most

$$\zeta := 2^{(1-c_1)n-(1-H_2(\delta_2))\delta_1(n-b)} = 2^{\delta_1(1-H_2(\delta_2))b} \cdot 2^{(1-c_1-(1-H_2(\delta_2))\delta_1)n}.$$

Since $1/2 < c_1 < 1$, we may choose $\delta_1$ such that $0 < 1-c_1 < \delta_1 < 1/2$, ensuring that $0 < 1-\frac{1-c_1}{\delta_1} < 1$, then choose $\delta_2 < 1/2$ satisfying $H_2(\delta_2) < 1 - \frac{1-c_1}{\delta_1}$ to obtain $(1 - c_1 - (1 - H_2(\delta_2))\delta_1) < 0$. Then, assuming $b = O(\log n)$, $\zeta$ is exponentially small. In particular, for the constant $c_1 = 2/3$ given by Theorem 3.1, we may choose $\delta_1 = 0.4$ and $\delta_2 = 1/64$ so that $H_2(\delta_2) < 1 - \frac{1/3}{0.4}$ and obtain $(1 - c_1 - (1 - H_2(\delta_2))\delta_1) < -0.0202 < -1/50$.

The above reasoning applies for any fixed $k'$, so by (12), we conclude that the proportion of pairs $(k, k')$ for which $|J_k \Delta J_{k'}| \geq \delta_2|S_v|$ is $1 - O(\zeta)$. With this bound on $|J_k \Delta J_{k'}|$, we aim at applying Lemma 2.1 with $a := \frac{2^b}{\epsilon}$ and $t := n^c$ for some $0 < c < 1$. For $v \in \pi_1(G')$ and $n > b$, we have $\delta_2|S_v| \geq \delta_2\delta_1(n-b) \geq c^* n$ for some constant $0 < c^* < 1$. So, choosing $n$ large enough to also satisfy

$$b + \log_2(1/\epsilon) \leq \frac{1-c}{2}\log_2 n - \frac{1}{2}\log_2(1/c^*), \tag{15}$$

we have $\delta_2|S_v| \geq c^* n \geq (\frac{2^b}{\epsilon})^2 n^c$, so $|J_k \Delta J_{k'}| \geq (\frac{2^b}{\epsilon})^2 n^c$ for all but a $O(\zeta)$ fraction of pairs $(k, k')$.

Now, for $v \in \pi_1(G')$, Lemma 2.1 asserts that the expectation over the random noise bits $r_v^{(\cdot)}$ of the squared norm of the sum of exponentials in (8) is at most

$$(P-1) + 2\zeta\binom{P-1}{2} + 2\binom{P-1}{2}e^{-2\pi^2 n^c} = O(\max\{\zeta, e^{-2\pi^2 n^c}\}(P-1)^2),$$

since $c < 1$. Thus, for $(v, w) \in G'$ and any $u^*$, the expectation over the random noise bits of the whole expression in (8) is $\mathrm{E}[p(v, w, g^{u^*})] = O(\max\{\zeta, e^{-2\pi^2 n^c}\}/2^{2n})$. For any $(v, w) \in G$, let

$$p(v, w) := \sum_{u^*=0}^{P-2} p(v, w, g^{u^*}), \tag{16}$$

be the probability of measuring $|v\rangle\,|w\rangle$ in the first two registers. For each of the $P-1$ possible states $|g^{u^*}\rangle$ in the third register, $|G'| = |\pi_1(G')| \leq 2^n$, so

$$\sum_{(v,w)\in G'} \mathrm{E}[p(v, w)] = O\left(\max\{\zeta, e^{-2\pi^2 n^c}\}\frac{(P-1)2^n}{2^{2n}}\right) = O(\max\{\zeta, e^{-2\pi^2 n^c}\}). \tag{17}$$

Finally, recall the bound in (11) on the proportion of $v \in \pi_1(G)$ not in $\pi_1(G')$. Since $\pi_1$ is injective on $G$ (see (9)), we have, with (11),

$$\frac{|G \setminus G'|}{|G|} = \frac{|\pi_1(G) \setminus \pi_1(G')|}{|\pi_1(G)|} \leq n^{O(1)} \cdot 2^{-(1-H_2(\delta_1))n}. \tag{18}$$

Now (18) and (9) give

$$|G \setminus G'| \leq n^{O(1)} \cdot 2^{H_2(\delta_1)n}.$$

For each $(v, w) \in G \setminus G'$, the largest possible value of the expression for $p(v, w, g^{u^*})$ in (8) is $\frac{1}{2^{2n}}$, so $p(v, w) \leq \frac{P-1}{2^{2n}} = \Theta\left(\frac{1}{2^n}\right)$. Thus

$$\sum_{(v,w)\in G\setminus G'} \mathrm{E}[p(v, w)] \leq n^{O(1)} \cdot 2^{-(1-H_2(\delta_1))n}. \tag{19}$$

With $\delta_1 = 0.4$ as above, this quantity is at most $n^{O(1)}2^{-n/35}$. Now, adding the quantities in (17) and (19) gives

$$\begin{aligned}\sum_{(v,w)\in G} \mathrm{E}[p(v, w)] &= \sum_{(v,w)\in G'} \mathrm{E}[p(v, w)] + \sum_{(v,w)\in G\setminus G'} \mathrm{E}[p(v, w)] \\ &= O(\max\{\zeta, e^{-2\pi^2 n^c}\}) + n^{O(1)}2^{-(1-H_2(\delta_1))n},\end{aligned}$$

which is exponentially small in $n$. The success of Shor's algorithm for the DLP [31] is based on the nontrivial probability that we measure some $|v\rangle\,|w\rangle$ with $(v, w) \in G$. However, we have shown that this approach succeeds with exponentially small probability, when noise is present at the specified level. With (13) and (15), we summarize this result in the following theorem.

Theorem 3.2. *There exists a constant $0 < c < 1$ such that, for a positive density of primes $P$, if each controlled-$R_k$-gate in the quantum Fourier transform circuit is replaced by a controlled-$\widetilde{R_k}$-gate for all $k \geq b$, where $b + \log_2(1/\epsilon) \leq \frac{1-c}{2}\log_2 n - \Theta(1)$ and $n$ is the binary length of $P$, then for all but an exponentially small fraction of inputs $y \in \mathbb{Z}_P^*$, for the discrete log problem with respect to any generator $g \in \mathbb{Z}_P^*$, Shor's algorithm has an exponentially small probability, over quantum measurement and random noise, to find the discrete log value $d$ satisfying $g^d = y \bmod P$.*

### 3.4 Analysis Over a Random Prime

In this section, we study the performance of the noisy discrete log algorithm for random prime $P$. We will prove a result similar to that of the previous section: for probability $1 - o(1)$ over random choice of $P$ and $y \in Z_P^*$, the algorithm has an exponentially small probability of success. We will follow the analysis in Section 4 of Reference [3] regarding the performance of Shor's algorithm to factor integers $N = pq$ for random primes $p$ and $q$. Consider primes $P$ with binary length $n$: $Y \leq P \leq X$, where $X = 2^n - 1$, $Y = 2^{n-1}$. Let $\omega_P(d)$ be the order of $d$ in $\mathbb{Z}_{P-1}$. The proof of Theorem 3.2 uses Theorem 3.1 to assume that $P - 1$ has a prime factor of size at least $P^{c_1}$ for some $1/2 < c_1 < 1$, from which it concludes that $P$ has the property that $\Pr[\omega_P(d) \geq P^{c_1}]$ is exponentially close to 1, where the probability is over uniformly random $d \in \{0, \ldots, P-2\}$. We show below that a random prime $P$ has this property with probability exponentially close to 1.

Since $\mathbb{Z}_P^* \cong \mathbb{Z}_{P-1}$, it follows from [3, Lemma 6] that there is a constant $C > 0$ such that, for any $B > 1$,

$$\Pr\left[\omega_P(d) < \frac{P-1}{B}\right] \leq C\left(\frac{n}{B\log B}\right)^{1/2},$$

and the probability is over the choice of $Y \leq P \leq X$ and $d \in Z_P^*$. Setting $B = (P-1)/P^{c_1} \approx P^{1-c_1}$, we have

$$\Pr\left[\omega_P(d) < P^{c_1}\right] \leq O\left(\frac{n}{(1-c_1)P^{1-c_1}\log P}\right)^{1/2} = O\left(2^{-(1-c_1)n/2}\right),$$

which is exponentially small in $n$, giving the desired property.

Recall that, for generator $g$, inputs $y \in \mathbb{Z}_P^*$ are in one-to-one correspondence with discrete log values $d \in \{0, \ldots, P-2\}$. Combining the results of this section with the proof of Theorem 3.2, we have the following theorem.

THEOREM 3.3. *There exists a constant $0 < c < 1$ such that, if each controlled-$R_k$-gate in the quantum Fourier transform circuit is replaced by a controlled-$\widetilde{R_k}$-gate for all $k \geq b$, where $b + \log_2(1/\epsilon) \leq \frac{1-c}{2}\log_2 n - \Theta(1)$, then with probability $1 - o(1)$ for a random prime $P$ chosen uniformly from all primes of binary length $n$ and a random $y$ chosen uniformly from $\mathbb{Z}_P^*$, for the discrete log problem with respect to any generator $g \in \mathbb{Z}_P^*$, Shor's algorithm has an exponentially small probability, over quantum measurement and random noise, to find the discrete log value $d$ satisfying $g^d = y \bmod P$.*

## 4 Polynomial Relaxation

Recall that in Section 3.3, as part of the proof of Theorem 3.2, we showed that the probability that the measured state $|v\rangle\,|w\rangle$ satisfies (6) and (7) is exponentially small. However, it is conceivable that an algorithm could recover the discrete log value $d$ from a measured pair $(v, w)$ "polynomially close" to satisfying (6) and (7). In this section, we show that, upon measuring, the probability that the measured state $|v\rangle\,|w\rangle$ is such a pair $(v, w)$ is exponentially small, for a natural definition of "polynomially close".

To give a natural definition of polynomially close, we first define the QFT over more general cyclic groups. For any integer $N$, let $\omega_N = e^{2\pi i/N}$ be the basic $N$-th root of unity. Define the QFT $F_N$ over the cyclic group $\mathbb{Z}_N$ by, for $x \in \mathbb{Z}_N$,

$$F_N\,|x\rangle = \frac{1}{\sqrt{N}} \sum_{y=0}^{N-1} \omega_N^{xy}\,|y\rangle\,.$$

The QFFT $F_{2^n}$ in (3) is the QFT over $\mathbb{Z}_{2^n}$. In (4), we apply $F_{2^n}$ to both the first and second registers, or equivalently apply $F_{2^n} \otimes F_{2^n}$, the QFT over the product group $\mathbb{Z}_{2^n} \oplus \mathbb{Z}_{2^n}$ (a group of order $2^{2n}$). In [30], Shor presents an algorithm for the DLP using a QFT over $\mathbb{Z}_{P-1} \oplus \mathbb{Z}_{P-1}$ for *smooth* (with no large prime factors) $P-1$. The QFT over $\mathbb{Z}_{P-1} \oplus \mathbb{Z}_{P-1}$ leads to an exact [1] algorithm for DLP [22], but is difficult to implement for general $P-1$. Hence, in Reference [31], which is the formulation in Section 3.1, Shor instead uses a QFT over $\mathbb{Z}_{2^n} \oplus \mathbb{Z}_{2^n}$, where $2^n \approx P-1$. This introduces the residue mod $2^n$ operations $\{\cdot\}_{2^n}$ in (6) and (7). If we instead applied the QFT over $\mathbb{Z}_{P-1} \oplus \mathbb{Z}_{P-1}$, these $\{\cdot\}_{2^n}$ operators become $\{\cdot\}_{P-1}$, which causes (7) to become trivial, and reduces (6) to $\{vd + w\}_{P-1} = 0$. In other words, (7) only captures the discrepancy between the QFT over $\mathbb{Z}_{P-1} \oplus \mathbb{Z}_{P-1}$, the true group underlying the DLP, and the QFT over $\mathbb{Z}_{2^n} \oplus \mathbb{Z}_{2^n}$, the approximation used by Shor's algorithm in Reference [31]. Furthermore, this algorithm extracts the discrete log value $d$ only from the relationship (6) between $d$ and the known values $v$ and $w$. So to relax the set of desired states $|v\rangle\,|w\rangle$, we replace (6) with

$$\left| \left\{ vd + w - \frac{d}{P-1}\{v(P-1)\}_{2^n} \right\}_{2^n} \right| < \gamma \tag{20}$$

for $\gamma = \text{poly}(n)$ and simply replace (7) with

$$|\{v(P-1)\}_{2^n}| < \frac{2^n}{C}, \tag{21}$$

for constant $C$.

[1] "exact" means the algorithm succeeds (returns the correct discrete log value) with probability 1. However, this exactness assumes exactly precise quantum gates and uses amplitude amplification, which fundamentally assumes that the algorithm has a high success probability before amplitude amplification takes place, and we show that, under our noise model, it does not.

We now trace through how these changes affect the analysis of Shor's algorithm. Regarding the change from (6) to (20), we replace $G$ with

$$G^\gamma = \{(v, w) \mid 0 \le v, w < 2^n, v \text{ and } w \text{ satisfy (20) and (21)}\},$$

and hence replace $\pi_1(G)$ with $\pi_1(G^\gamma) = \{v \mid (v, w) \in G^\gamma\}$, and replace $G'$ with an analogously defined $(G^\gamma)'$. Observe that, for fixed $v$, the number of $w$ satisfying (20) is exactly $2\gamma$ regardless of the choice of $v$ (assuming $\gamma$ is chosen so that $(P-1)\gamma \notin \mathbb{Z}$). Hence, where before we had $|G| = |\pi_1(G)|$, we now have $|G| = \gamma|\pi_1(G)|$. Furthermore, $\pi_1(G)$ is still defined only by (21), and the change from 12 in (7) to an arbitrary constant $C$ in (21) makes only superficial difference; before we had $|\pi_1(G)| \ge 2^n/12$, and now we have $|\pi_1(G)| \ge 2^n/C$. Thus the analysis until (16) is unchanged, up to replacing any constant 12 with $C$, as it is concerned only with $\pi_1(G)$ and $|\pi_1(G)|$. The first difference comes immediately before (17), where we first consider $|G'|$ in the inequality $|G'| = |\pi_1(G')| \le 2^n$. We replace this inequality with $|(G^\gamma)'| = \gamma|\pi_1((G^\gamma)')| \le \gamma 2^n$, which necessitates multiplying the RHS of (17) by $\gamma$. Then, since the number of $w$ satisfying (20) is still independent of $v$, we have

$$\frac{|G^\gamma \setminus (G^\gamma)'|}{|G^\gamma|} = \frac{|\pi_1(G^\gamma) \setminus \pi_1((G^\gamma)')|}{|\pi_1(G^\gamma)|}, \tag{22}$$

matching the equality in (18). Finally, we must also replace $|G| = |\pi_1(G)| \le 2^n$ before (19) with $|G^\gamma| = \gamma|\pi_1(G^\gamma)| \le \gamma 2^n$. Then, by (11) and (22), we have $|G \setminus G'| \le \gamma \cdot n^{O(1)} \cdot 2^{H_2(\delta_1)n}$, so, following Section 3.3, we obtain an analogue of (19) with the RHS multiplied by $\gamma$ as well.

To recap, we obtain analogues of (17) and (19) with $G^\gamma$ in place of $G$ and $(G^\gamma)'$ in place of $G'$, with an extra factor of $\gamma$ on the RHS of both. Combining these two bounds as in Section 3.3, we obtain

$$\sum_{(v,w)\in G^\gamma} \mathrm{E}[p(v, w)] \le \gamma \left[ O(\max\{\zeta, e^{-2\pi^2 n^c}\}) + n^{O(1)} 2^{-(1-H_2(\delta_1))n} \right],$$

which, since $\gamma = \mathrm{poly}(n)$, is still exponentially small.

## 5 Experimental Results

### 5.1 Quantum Hardware Experiments

The experiments were performed on version 1.20.21 of `ibm_kyiv`, an IBM Eagle r3 **quantum processor** (**QPU**) on the IBM Quantum Platform. Version 1.20.21 of `ibm_kyiv` has median error rates of $1.197 \cdot 10^{-2}$ and $2.424 \cdot 10^{-4}$ for its basis gates ECR and $\sqrt{X}$, respectively, and median readout error rate of $9.300 \cdot 10^{-3}$. All reported experimental percentages are out of 500 trials (500 shots of the respective circuit).

Figure 1(a) shows the results of the following experiment, a simple test of the effectiveness of the (supposedly precise, IBM Quantum Platform built-in) controlled-$R_k$ gates used in the QFT. First, we prepare the two-qubit state $|1\rangle|0\rangle$. Then we apply $H$ to the second qubit, and apply $2^{\ell-1}$ consecutive $R_\ell$ gates controlled by the first qubit and acting on the second qubit. Since the first qubit is always in state $|1\rangle$, each of these gates acts as $R_\ell$ on the second qubit, so, if no noise is present and the rotations are exact, these gates in aggregate effect a $Z = R_1$ transformation on the second qubit. Thus, applying another Hadamard to and then measuring the second qubit, we should always obtain state $|1\rangle$. In other words, writing $CR_\ell$ for the controlled-$R_\ell$ operator, we have an operator $X_\ell$ such that

$$X_\ell |1\rangle|0\rangle := (I \otimes H)(CR_\ell)^{2^{\ell-1}}(I \otimes H)|1\rangle|0\rangle = |1\rangle|1\rangle .$$

However, Figure 1(a) shows that the experimental probability of measuring $|1\rangle$ after applying the circuit implementing $X_\ell$ to state $|1\rangle|0\rangle$ (in black/solid) for all $2^\ell = 4, 8, 16, 32, 64, 128, 256$ decays

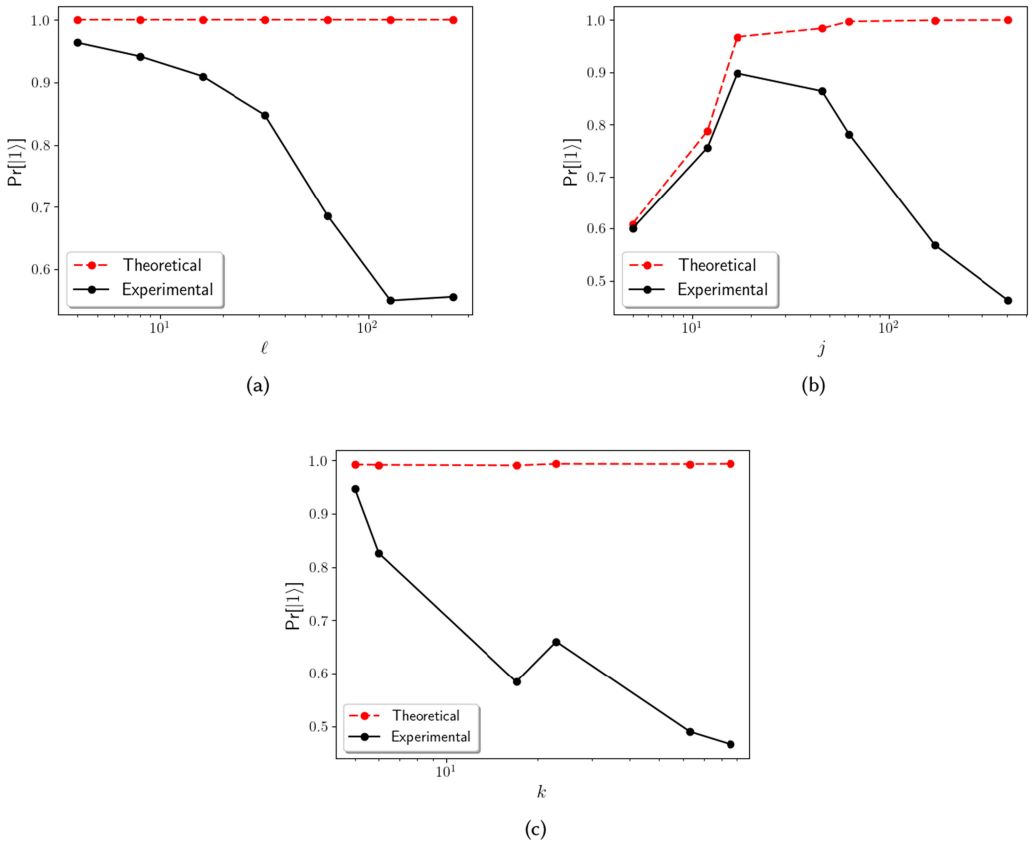


Fig. 1. Results of the three experiments discussed in Section 5.1, with the probability of measuring state $|1\rangle$ on the vertical axis and the respective experiment parameters on the horizontal axis.

to $\frac{1}{2}$ as $\ell$ increases. In the dashed line, the theoretical probability $\Pr[|1\rangle] = 1$ for all $\ell$ is shown. The experimental result confirms that as the angle $2\pi/2^\ell$ of rotation decreases, the precise rotations are overwhelmed by the noise and cannot accumulate to an overall rotation of angle $\pi$.

The above experiment uses the built-in quantum gates $CR_\ell$. We also try to perform an experiment using built-in quantum gates that are as basic as possible with the property that the subgroup generated by them includes arbitrarily small rotations. We report our result using the Hadamard gate and the $\pi/8$ gate $T$. More precisely, in order to express them in $SU(2)$, let $H = \frac{i}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$ and $T = \begin{bmatrix} e^{-i\pi/8} & 0 \\ 0 & e^{i\pi/8} \end{bmatrix}$ (up to an unimportant global phase, $H$ and $T$ are the familiar Hadamard and $\pi/8$ gates, respectively). Let $A = HT$. In the Bloch sphere representation, $A$ is a rotation around axis $(n_x, n_y, n_z)$ by angle $w$, that is,

$$A = \cos\left(\frac{w}{2}\right) I - i \cdot \sin\left(\frac{w}{2}\right)(n_x X + n_y Y + n_z Z),$$

where

$$\begin{aligned} w &\approx 2.593564246, \\ n_x = n_z &\approx -0.6785983445, \\ n_y &\approx 0.2810846377. \end{aligned}$$

Table 1. Angles of Rotation of $A^{j_n}$ for Continued Fraction Convergents $\frac{j_n}{d_n}$ of $\frac{2\pi}{w}$

| $j_n$ | $j_n w \bmod 2\pi$ |
|---|---|
| 5 | 0.4014506155 |
| 12 | −0.2931555843 |
| 17 | 0.1082950312 |
| 46 | −0.07656552182 |
| 63 | 0.03172950941 |
| 172 | −0.01310650300 |
| 407 | 0.005516503408 |

It can be shown that $w$ is not a rational multiple of $\pi$, and $2\pi/w \approx 2.42260638682$. We look for continued fraction expansions[2] of $2\pi/w$ and let $j_n/d_n$ be the $n$th continued fraction convergent of $2\pi/w$. By Dirichlet's Theorem, $|j_n/d_n - 2\pi/w| < 1/d_n^2$, and $d_n \to \infty$ as $2\pi/w$ is irrational. Hence $|j_n w - 2d_n\pi| < w/d_n \to 0$. $A^{j_n}$ is a rotation of the Bloch sphere about the same axis $(n_x, n_y, n_z)$ by angle $j_n w \bmod 2\pi \approx 0$, that is,

$$A^{j_n} = \cos\left(\frac{j_n w}{2}\right) I - i \sin\left(\frac{j_n w}{2}\right)(n_x X + n_y Y + n_z Z) \approx \cos(d_n\pi) I - i \sin(d_n\pi)(n_x X + n_y Y + n_z Z) = \pm I. \tag{23}$$

For increasing $n$, the rotation $A^{j_n}$ is an increasingly accurate approximation of the identity rotation, and thus (ignoring a phase factor) $A^{j_n} T \approx T$. Equivalently, $A^{j_n - 1} H \approx T^{-1}$. The angles of rotation of $A^{j_n}$ for the first seven values of $j_n$ are shown in Table 1.

Note that $T^4 = T^{-4} = Z$ is a rotation of $\pi$ about $(0, 0, 1)$ transforming $H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ to $H|1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. Thus

$$F_n := H(A^{j_n - 1} H)^4 H \approx X.$$

Figure 1(b) shows the theoretical probability $\Pr[|1\rangle]$ of measuring $|1\rangle$ after applying a circuit implementing $F_n$ to initial state $|0\rangle$ (dashed line) and the actual experimental outcomes (solid line). If we measure $F_n|0\rangle$ in the standard basis $\{|0\rangle, |1\rangle\}$ we should get back $|1\rangle$ with probability approaching 1 in theory. However, Figure 1(b) shows that in reality it does not approach 1. Rather, initially as $A^{j_n}$ becomes a better approximation of $\pm I$, the probability of measuring $|1\rangle$ increases according to theory. But as the (theoretical) angle of each rotation $A^{j_n}$ gets smaller, the actual percentage of measuring $|1\rangle$ decreases and getting a measurement of $|0\rangle$ or $|1\rangle$ becomes quite random, more in line with the general analysis of this article.

Figure 1(c) shows a variation of the previous experiment. Similarly to above, let $\frac{k_n}{e_n}$ be the $n$th continued fraction convergent of $\frac{\pi}{w}$. Then, as in (23), $A^{k_n}$ is a rotation around the axis through $(n_x, n_y, n_z)$ by angle $k_n w$ equivalent to $\epsilon_n$ or $\pi + \epsilon_n$ modulo $2\pi$ (depending on $e_n$ is even or odd, respectively), for small $\epsilon_n$. See Table 2.

Define angles $\phi$ and $\theta$ such that

$$(n_x, n_y, n_z) = (\cos(\phi)\sin(\theta), \sin(\phi)\sin(\theta), \cos(\theta)) =: B(\phi, \theta).$$

Then define a unitary operator $U$ that maps $|0\rangle$ (the point $(0, 0, 1)$ on the Bloch sphere) to $B(\phi, \theta)$, and maps $|1\rangle$ to the antipodal point on the sphere. Then, letting $R_Z(\theta) = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\theta} \end{bmatrix}$ be the rotation

[2]This idea is due to Kurt Girstmair [14] and we are grateful to Kurt for permission to use it for our experiment.

Table 2. Angles of Rotation of $A^{k_n}$ for Continued Fraction Convergents $\frac{k_n}{e_n}$ of $\frac{\pi}{w}$

| $k_n$ | $k_n w \bmod 2\pi$ | $\lvert\epsilon_n\rvert$ |
|---|---|---|
| 5 | 0.401450616 | 0.401450616 |
| 6 | 2.995014861 | 0.146577793 |
| 17 | 0.108295031 | 0.108295031 |
| 23 | 3.103309893 | 0.038282761 |
| 63 | 0.031729509 | 0.031729509 |
| 86 | 3.135039402 | 0.006553252 |

of the Bloch sphere around the $Z$-axis (the axis through the point $(0, 0, 1)$) by angle $\theta$, $A^{k_n}$ is equivalent to $R_Z(\epsilon_n)$ or $R_Z(\pi + \epsilon_n)$ under basis $U$. Let $t_n$ be the closest even integer to $\frac{\pi}{\epsilon_n}$. Then $A^{k_n t_n}$ is a rotation by angle close to $\pi \bmod 2\pi$, so

$$G_n := H^\dagger U^\dagger A^{k_n t_n} U H \approx H^\dagger Z H = X.$$

Hence $G_n |0\rangle \approx |1\rangle$. Figure 1(c) shows the theoretical and experimental probabilities of measuring $|1\rangle$ after applying a circuit implementing $G_n$ to state $|0\rangle$ for increasingly accurate convergent numerators $k_n$. Again, the experimental $\Pr(|1\rangle)$ decays to around $\frac{1}{2}$ as $n$ increases and noise overwhelms the precise $A^{k_n t_n}$ rotations, which no longer accumulate to an angle of $\pi \bmod 2\pi$.

### 5.2 Numerical Simulation

Figure 2 shows the results of some numerical experiments aimed at estimating the probability $\Pr[(v, w) \in G]$ that the state $|v\rangle |w\rangle$ measured by Shor's algorithm satisfies $(v, w) \in G$. This is the probability that the quantum algorithm likely reveals useful information for the discrete logarithm computation. For a given $n$, we sample a random $n$-bit prime $P$, a random $2 \le d \le P - 2$, a random $0 \le u^* \le P - 2$, and a random $(v, w) \in G$, then approximately compute the probability $p(v, w, g^{u^*})$ given in (8). We then use this to estimate $\Pr[(v, w) \in G]$, which is

$$\sum_{0 \le u^* \le P-2,\ (v,w) \in G} p(v, w, g^{u^*}) \approx (P - 1) \cdot \frac{2^n}{6} \cdot \mathbb{E}_{(v,w)\sim G,\, u^* \sim [0, P-2]}[p(v, w, g^{u^*})], \tag{24}$$

as $|G| \approx \frac{2^n}{6}$, and the value $p(v, w, g^{u^*})$ in (8) is approximately the same for all $0 \le u^* \le P - 2$. Each run of the computation of $p(v, w, g^{u^*})$ in (8), for a chosen $P, d, u^*, v, w$, is an exponential computation in $n$. To ease the computation slightly, we use only the noise random variables $r_0^{(\cdot)}, r_1^{(\cdot)}, \rho_0^{(\cdot)}, \rho_1^{(\cdot)}$, leading to a slight underestimate of the amount of noise present. (This is similar to what was done in the proof, see (10).) Each point in Figure 2 shows the average value of 1000 evaluations of $p(v, w, g^{u^*})$ split into 20 samples of $P$ and $d$, then 10 samples of $(v, w) \in G$ and $u^*$, and then 5 samples of the noise variables $r_0^{(\cdot)}, r_1^{(\cdot)}, \rho_0^{(\cdot)}, \rho_1^{(\cdot)}$, multiplied by $(P - 1)\frac{2^n}{6} \approx (P - 1)|G|$.

The points marked by a black x in Figure 2 show that $\Pr[(v, w) \in G] \approx 0.1$ for all $n$ in the noise-free case. We then consider noise level $\frac{\epsilon}{2^b} = n^{-\gamma}$ for $\gamma = \frac{2}{3}, \frac{1}{2}, \frac{1}{3}$. For $\gamma = \frac{2}{3}$ (the lowest level of noise among the three cases), $\Pr[(v, w) \in G]$ appears to stabilize around 0.02. For $\gamma = \frac{1}{2}$, $\Pr[(v, w) \in G]$ appears to decrease inverse polynomially in $n$, with best fit $2.961n^{-2.350}$. For $\gamma = \frac{1}{3}$, $\Pr[(v, w) \in G]$ decreases exponentially in $n$, roughly as $2^{-1.083n^{0.789}}$. In this case, with noise level still vanishing as a function in $n$, extrapolating to $n = 500$ gives $\Pr[(v, w) \in G]$ on the order of $10^{-44}$. Asymptotically, $\frac{\epsilon}{2^b} \approx n^{-1/3}$ is the noise level proved in the main theorems of this article, Theorems 3.2 and 3.3.

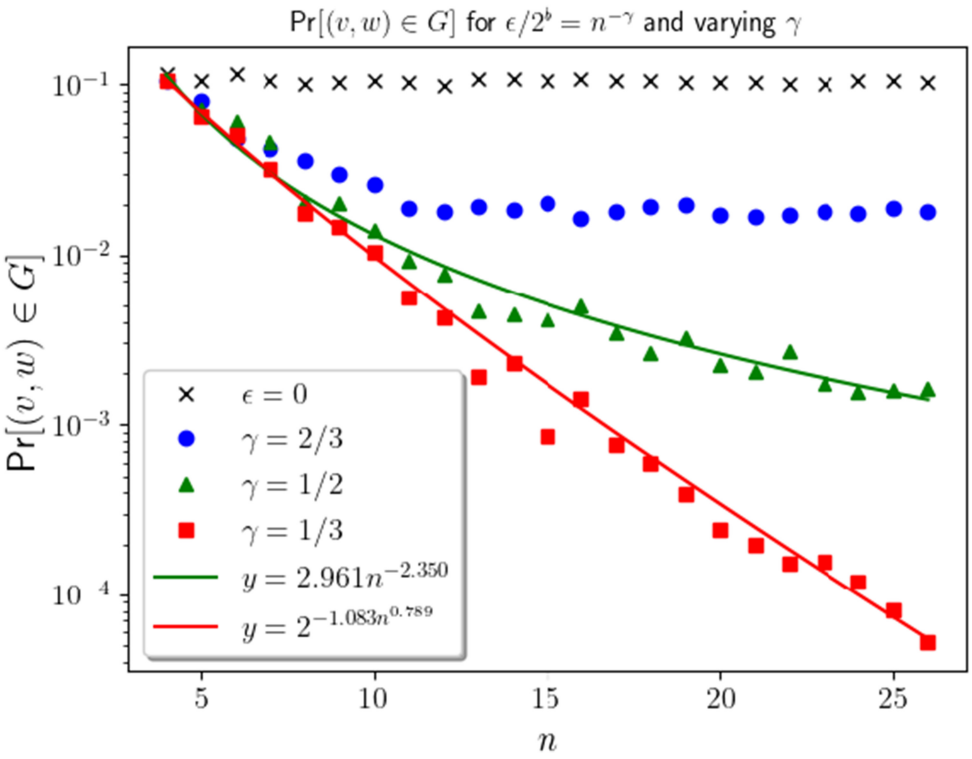


Fig. 2. Numerical simulation results: estimating the probability $\Pr[(v,w) \in G]$ for noise levels of $\frac{\epsilon}{2^b} = n^{-\gamma}$ by computing (8).

## Acknowledgements

The authors thank the anonymous referees for their detailed comments which helped us improve the content and presentation of the paper. The idea of using continued fraction convergents in our experiments is due to Kurt Girstmair. We thank him very much for his permission to let us use this idea in Section 5.1.

## References

[1] Adriano Barenco, Artur Ekert, Kalle-Antti Suominen, and Päivi Törmä. 1996. Approximate quantum Fourier transform and decoherence. *Phys. Rev. A* 54, 1 (Jul 1996), 139–146.

[2] Dan Boneh and Richard J. Lipton. 1995. Quantum cryptanalysis of hidden linear functions. In *Advances in Cryptology—CRYPT0' 95*, Don Coppersmith (Ed.). Springer, Berlin, 424–437.

[3] Jin-Yi Cai. 2024. Shor's algorithm does not factor large integers in the presence of noise. *Science China Information Sciences* 67, 7 (June 2024), 173501. DOI: https://doi.org/10.1007/s11432-023-3961-3

[4] Andrew M. Childs and Gábor Ivanyos. 2014. Quantum computation of discrete logarithms in semigroups. *Journal of Mathematical Cryptology* 8, 4 (2014), 405–416. DOI: https://doi.org/doi:10.1515/jmc-2013-0038

[5] D. Coppersmith. 2002. An approximate Fourier transform useful in quantum factoring. arXiv:quant-ph/0201067 [quant-ph] https://arxiv.org/abs/quant-ph/0201067

[6] W. Diffie and M. Hellman. 1976. New directions in cryptography. *IEEE Transactions on Information Theory* 22, 6 (1976), 644–654. DOI: https://doi.org/10.1109/TIT.1976.1055638

[7] Martin Ekerå. 2021. Quantum algorithms for computing general discrete logarithms and orders with tradeoffs. *Journal of Mathematical Cryptology* 15, 1 (2021), 359–407. DOI: https://doi.org/doi:10.1515/jmc-2020-0006

[8] Martin Ekerå and Joel Gärtner. 2024. Extending regev's factoring algorithm to compute discrete logarithms. In *Post-Quantum Cryptography - 15th International Workshop, PQCrypto 2024, Oxford, UK, June 12-14, 2024, Proceedings, Part*

*II (Lecture Notes in Computer Science, Vol. 14772)*, Markku-Juhani O. Saarinen and Daniel Smith-Tone (Eds.). Springer, 211–242. arXiv:2311.05545 [quant-ph]. DOI: 10.1007/978-3-031-62746-0_10

[9] Martin Ekerå and Johan Håstad. 2017. Quantum algorithms for computing short discrete logarithms and factoring RSA integers. In *Post-Quantum Cryptography*. Tanja Lange and Tsuyoshi Takagi (Eds.). Springer International Publishing, Cham, 347–363.

[10] Jörg Flum and Martin Grohe. 2006. *Parameterized Complexity Theory*. Springer, Berlin. DOI: https://doi.org/10.1007/3-540-29953-X

[11] Étienne Fouvry. 1985. Théorème de Brun-Titchmarsh; Application au théorème der Fermat. *Inventiones Mathematicae* 79 (1985), 383–408. Retrieved from http://eudml.org/doc/143202

[12] Austin G. Fowler and Lloyd C. L. Hollenberg. 2004. Scalability of Shor's algorithm with a limited set of rotation gates. *Physical Review A* 70, 3 (Sept. 2004), 032329. DOI: https://doi.org/10.1103/PhysRevA.70.032329 Publisher: American Physical Society.

[13] Austin G. Fowler and Lloyd C. L. Hollenberg. 2007. Erratum: Scalability of Shor's algorithm with a limited set of rotation gates [Phys. Rev. A **70** , 032329 (2004)]. *Physical Review A* 75, 2 (Feb. 2007), 029905. DOI: https://doi.org/10.1103/PhysRevA.75.029905

[14] Kurt Girstmair. 2024. Private communication.

[15] Robert B. Griffiths and Chi-Sheng Niu. 1996. Semiclassical fourier transform for quantum computation. *Phys. Rev. Lett.* 76, 17 (Apr 1996), 3228–3231. DOI: https://doi.org/10.1103/PhysRevLett.76.3228

[16] L. Hales and S. Hallgren. 2000. An improved quantum Fourier transform algorithm and applications. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*. IEEE, New York, NY, USA, 515–525. DOI: https://doi.org/10.1109/SFCS.2000.892139

[17] Yan Huang, Zhaofeng Su, Fangguo Zhang, Yong Ding, and Rong Cheng. 2020. Quantum algorithm for solving hyperelliptic curve discrete logarithm problem. *Quantum Information Processing* 19, 2 (2020), 62.

[18] Burton S. Kaliski. 2017. A quantum "Magic Box" for the discrete logarithm problem. *IACR Cryptol. ePrint Arch.* 2017 (2017), 745. Retrieved from https://api.semanticscholar.org/CorpusID:42139218

[19] Alexei Y. Kitaev. 1996. Quantum measurements and the Abelian Stabilizer Problem. Electron. *Colloquium Comput. Complex.* TR96-003 (1996). ECCC:TR96-003 https://eccc.weizmann.ac.il/eccc-reports/1996/TR96-003/index.html

[20] Chris Lomont. 2004. The Hidden Subgroup Problem—Review and Open Problems. (Nov. 2004). Retrieved from http://arxiv.org/abs/quant-ph/0411037 arXiv:quant-ph/0411037.

[21] Michele Mosca and Artur Ekert. 1999. The hidden subgroup problem and eigenvalue estimation on a quantum computer. In *Quantum Computing and Quantum Communications*. Colin P. Williams (Ed.). Springer, Berlin, 174–188.

[22] Michele Mosca and Christof Zalka. 2004. Exact quantum fourier transforms and discrete logarithm algorithms. *International Journal of Quantum Information* 02, 01 (March 2004), 91–100. DOI: https://doi.org/10.1142/S0219749904000109

[23] Y. S. Nam and R. Blümel. 2013. Scaling laws for Shor's algorithm with a banded quantum Fourier transform. *Physical Review A* 87, 3 (March 2013), 032333. DOI: https://doi.org/10.1103/PhysRevA.87.032333 Publisher: American Physical Society.

[24] Y. S. Nam and R. Blümel. 2014. Robustness of the quantum Fourier transform with respect to static gate defects. *Physical Review A* 89, 4 (April 2014), 042337. DOI: https://doi.org/10.1103/PhysRevA.89.042337. Publisher: American Physical Society.

[25] Y. S. Nam and R. Blümel. 2015. Performance scaling of the quantum Fourier transform with defective rotation gates. *Quantum Inf. Comput.* 15, 9&10 (2015), 721–736. DOI: 10.26421/QIC15.9-10-1

[26] Y. S. Nam and R. Blümel. 2015. Structural stability of the quantum Fourier transform. *Quantum Information Processing* 14, 4 (April 2015), 1179–1192. DOI: https://doi.org/10.1007/s11128-015-0923-2

[27] Michael A. Nielsen and Isaac L. Chuang. 2010. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, Cambridge, UK. DOI: https://doi.org/10.1017/CBO9780511976667

[28] John Proos and Christof Zalka. 2003. Shor's discrete logarithm quantum algorithm for elliptic curves. *Quantum Info. Comput.* 3, 4 (Jul 2003), 317–344.

[29] Oded Regev. 2025. An efficient quantum factoring algorithm. *J. ACM* 72, 1 (Jan. 2025), 10:1–10:13. DOI: 10.1145/3708471

[30] P. W. Shor. 1994. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*. IEEE, New York, NY, USA, 124–134. DOI: https://doi.org/10.1109/SFCS.1994.365700

[31] Peter W. Shor. 1997. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* 26, 5 (1997), 1484–1509. DOI: https://doi.org/10.1137/S0097539795293172

[32] L. F. Wei, Xiao Li, Xuedong Hu, and Franco Nori. 2005. Effects of dynamical phases in Shor's factoring algorithm with operational delays. *Phys. Rev. A* 71, 2 (Feb 2005), 022317.

Received 28 February 2024; revised 15 January 2025; accepted 12 April 2025