Quantum Information and Computation, Vol. 15, No. 9&10 (2015) 0721–0736

PERFORMANCE SCALING OF THE QUANTUM FOURIER TRANSFORM WITH DEFECTIVE ROTATION GATES

YUN SEONG NAM [a]
Department of Physics, Wesleyan University
Middletown, Connecticut 06459-0155, USA

REINHOLD BlÜMEL
Department of Physics, Wesleyan University
Middletown, Connecticut 06459-0155, USA

Received June 26, 2014
Revised Juanuary 12, 2015

Addressing Landauer's question concerning the influence of static gate defects on quantum information processor performance, we investigate analytically and numerically the case of the quantum Fourier transform (QFT) with defective controlled rotation (CROT) gates. Two types of defects are studied, separately and in combination: systematic and random. Analytical scaling laws of QFT performance are derived with respect to the number of qubits n, the size δ of systematic defects, and the size ϵ of random defects. The analytical results are in excellent agreement with numerical simulations. In addition, we present an unexpected result: The performance of the defective QFT does not deteriorate with increasing n, but approaches a constant that scales in ϵ. We derive an analytical formula that accurately reproduces the ϵ scaling of the performance plateaus. Overall, we observe that the CROT gates may exhibit static and random defects of the order of 30% and larger, and still result in satisfactory QFT performance. Thus we answer Landauer's question in the case of the QFT: far from being lethal, the QFT can tolerate tremendous static gate defects and still perform its function. The extraordinary robustness of the QFT with respect to static gate defects displayed in our numerical and analytical calculations should be a welcome boon for laboratory and industrial realizations of quantum circuitry.

Keywords: Quantum Fourier Transform, Performance, Defects
Communicated by: R Cleve & G Milburn

1 Introduction

Unless its operation is trivial, any quantum information processor, if it is designed to communicate its output to the "classical" world of our everyday existence, is necessarily a probabilistic device. Therefore, how well it performs under realistic conditions, is primarily a function of the probability of success with which it returns a desired result. While ultimately the speed of execution might also play a role in characterizing the performance of a quantum information processor, for the purposes of this paper, it is sufficient to define performance as the success probability of a quantum information processor to deliver a desired result in a single run of the processor. Realistically, when operating a quantum information processor,

[a]E-mail address: ynam@wesleyan.edu

i.e., a physical realization of a quantum algorithm, there are many different kinds of noise and circuit defects, static and dynamic, that may negatively impact the performance of the processor. Well-known negative influences, e.g., are quantum decoherence [1] and unwanted coupling between qubits [2]. An early list of mechanisms and effects that may be fatal for quantum computing was compiled by Landauer [3]. Manufacturing flaws in the physical implementations of quantum gates is one of the adverse mechanisms stressed in [3] and forms the main topic of investigation of this paper. We call these manufacturing flaws static gate defects [4], and, from the outset, it is indeed not clear how sensitively the performance of a quantum information processor reacts to static imperfections of its gates. Taking the controlled rotation gate (CROT) as an example: how accurately does the rotation angle have to be implemented for the processor to work with satisfactory performance? How does the performance scale with the size of the gate errors and the number of qubits?

In this paper we answer these questions quantitatively for the case of the quantum Fourier transform (QFT) as a specific example of a quantum information processor with static defects in its CROT gates. Since the purpose of the QFT is to find the embedded periodicity of an input state, we define our performance measure as the success probability of finding the periodicity (see Sec. 3 for detail).

Typically, a QFT circuit, both the coherent two-qubit version [1] and the semiclassical version [5], contains controlled rotation (CROT) gates, rotating the phase angle of qubit number t (the target qubit) by an angle θ_j, if the control qubit c is in state $|1\rangle$. The rotation angle θ_j is given by [1]

$$\theta_j = \frac{\pi}{2^j}, \tag{1}$$

where $j = |t - c|$.

Previously [4, 6], in a first attempt at addressing Landauer's question [3] of what happens to the performance of quantum circuits if they are not precisely implemented, but exhibit static defects, we subjected the QFT to four different classes of static defects in the CROT gates, characterized by $\theta_j \to \theta_j + R$, where R is a random number. Four different cases were studied in detail [4, 6], characterized by whether R is correlated with j or not, or whether R is an absolute or a relative error.

In this paper we relax the condition of symmetric (on average) perturbation around the exact CROT rotation angle (1) and, in addition to random defects, introduce a systematic offset angle such that the class of static defects to be studied takes the form

$$\theta_j \to \tilde{\theta}_j = \frac{\delta + \epsilon r}{2^j}, \tag{2}$$

where δ is the offset, ϵ denotes the random defect strength, and r is a flat-distributed random variable taking values between -1 and 1, not correlated with j. This means that each CROT gate of type j in the circuit will have a different rotation angle $\tilde{\theta}_j$ that corresponds to a different realization of the random variable r. This is closest to a possible experimental realization of the QFT, where successive gates, even if they are of the same type (same j) will, in general, have different defects. The exact, defect-free case is included and corresponds to the case $\delta = \pi$ and $\epsilon = 0$.

The introduction of the static offset δ is not just an incremental change with respect to the case of purely random defects, since $\delta = 0$, $\epsilon \neq 0$, e.g., corresponds to truly random CROT

gates, and, for $0 < \delta \ll 1$, the effects of a small positive offset in angle, masked by large random defects, may be studied. The surprise here is that acceptable performance is already achieved in cases where δ is small but ϵ is *larger* than δ, seemingly overpowering the average drift ($\delta \neq 0$) toward the exact rotation angle π.

Apart from deriving performance scaling laws for static defects that include systematic errors, we also report a qualitatively new phenomenon. Although, with increasing number of qubits, an increasing number of defective CROT gates are added into the QFT circuit, the performance of the curcuit does not deteriorate, but approaches a plateau. This observation is counterintuitive and quite unexpected. It holds strictly in the ideal case where external influences on the QFT circuit, such as ambient noise, are neglected. The presence of ambient noise breaks the plateau behavior at a critical number of qubits n_c that depends on the ambient noise level. However, $n_c \to \infty$ for noise level $\to 0$.

We shall see that the QFT circuit exhibits astonishing resilience to static gate defects, even if the circuit consists of thousands of qubits, relevant for applications in quantum cryptography [1]. This may be of importance for all researchers and engineers who are interested in designing and building quantum information and computation hardware based on the gate model of quantum computing.

Our paper is structured in the following way. First, in Sec. 2, we briefly review some historical background pertaining to the performance of the QFT with respect to gate defects and justify our choice of defect model. We also clarify the fundamental difference between spontaneous, correctable errors introduced by decoherence and systematic, uncorrectable errors introduced by faulty hardware. In Sec. 3, we numerically study the performance of the QFT circuit equipped with defective CROT gates. We will see that acceptable performance is obtained for large ϵ even if δ is small. In Sec. 4, we investigate how the performance of the QFT depends on the number n of qubits. It is in this section that we encounter the plateauing phenomenon, i.e., the fact that the performance of the QFT stays high for increasing n even though the quantum state is processed by an increasing number of defective gates. Then, in Sec. 5, we derive an analytical scaling law for the performance of the QFT with defective CROT gates. We discuss our results in Sec. 6, and summarize and conclude our paper in Sec. 7.

2 Our defect model in the historical context

The problem of fragility of quantum computing has been pointed out by many physicists since the inception of quantum computing. In particular, Landauer [3] was among the first to compile a list of physical processes that may prove detrimental to quantum computing. Since then, much effort has been devoted to overcoming many of the obstacles pointed out by Landauer, culminating in novel paradigms such as quantum error correction and fault-tolerant quantum computing (see [7] and references therein). Furthermore, in the pursuit of a full-scale, physical quantum computer, useful for practical purposes, many numerical studies of the QFT have been performed throughout the past decades, for instance the investigation of approximation techniques [8, 9, 10], the effects of imprecise quantum operations [11, 12], and modeling the effects of decoherence [9, 12, 13].

The recent discovery of surface codes and their ingenious working principle further illustrate the relentless efforts toward making quantum computing a practical reality. By encoding

logical quantum information in a non-local degree of freedom, surface codes allow us to detect and localize local errors occurring on physical qubits that realize a logical qubit. This results in an astonishingly improved level of protection of quantum information, raising the threshold level, i.e., the maximal physical, unencoded error rate that may be tolerated to still result in a lower logical, encoded error rate, to 1% [14]. With the availability of software such as LIQUID [15] and AutoTune [16, 17], the research frontiers continue to determine experimental targets that quantum hardware should aim for, and anticipated further developments will be of interest to many quantum engineers and experimentalists.

Among the numerous investigations available in the literature, we find that our work is most closely related to the works of Coppersmith [8], Barenco *et al.* [9], and Fowler and Hollenberg [10]. Coppersmith, in 1994 [8], discovered that the QFT may be approximated with banding, i.e., removing all rotation gates $\theta_j = \pi/2^j$ with $j > b$, where b is the chosen bandwidth, without suffering from significant loss of performance for a sufficiently large b. Followed by Barenco *et al.* in 1996 [9], the QFT was then further investigated, this time subjected to both banding and decoherence, modeled by extra phase-angle rotations applied every time the phase rotation gate θ_j was applied. This, in fact, corresponds closely to one of the four types of defects studied in [4], where, in addition, an explicit analytical scaling result is reported. A more extensive numerical simulation was then performed by Fowler and Hollenberg in 2004 [10]. Not only was the QFT affected by banding and extra phase-angle rotations similar to the cases studied in [9], but was also the justification for the choice of the noisy rotation gate model provided, this time with fault-tolerant construction of the gates in mind.

Our work attempts to extend this line of research in the following way. First, quantum error correction will be a crucial component of an implementation of a quantum computer, and will require significant additional resources in terms of qubits and quantum gates. Each logical qubit is represented by a set of physical qubits and each logical quantum gate is replaced by a fault-tolerant operation on the physical qubits. For a given quantum error correction code, only a discrete set of quantum gates can be implemented fault-tolerantly. Therefore, each rotation gate must be decomposed into a sequence of gates drawn from the fault-tolerant basis. According to the Solovay-Kitaev theorem, if the basis is universal, then such a sequence can always be found [18]. Now, the obtained sequence, of course, results in an approximate rotation gate, provided that the sequence has to be finite. While recent developments in techniques and new methods of decomposition [19, 20, 21, 22, 23] allow us to straightforwardly obtain the sequence with an arbitrarily chosen accuracy (see for instance [24] for open-source software), the fact that the resulting rotation is approximate remains true. Thus, considering that resources required for quantum error correction are demanding, investigating how sensitively a quantum computer (in our case the QFT) reacts to deviations from the exact operation of its rotation gates becomes important, as low sensitivity can help ease the accuracy requirement, which implies a step forward into the direction of practical quantum computing.

While fault-tolerant implementation via sequence decomposition and surface codes allow us to fight many types of errors [14], including some of the systematic [25] and intrinsic [26, 27, 28] types, some errors will always remain in quantum gates and states. This is obvious, since, in general, an error correcting code never fully restores the corrupted state;

rather when compared to the unencoded counterparts, the encoded one performs better [7, 29]. Accordingly, in practice, there will always be some residual errors left in individual decomposed gates of a rotation gate, for instance, and they will (diffusively) accumulate through the pulse sequences. It is then interesting to see how the QFT circuit, used in many algorithms, performs under various types of errors, which forms the main topic of this paper. If the QFT is very tolerant to these errors, this in turn means we may relax required precision, saving on quantum resources. The exact description of these errors will of course be hardware dependent, and we will learn more about the expression of the errors as hardware advances. For now, without the availability of the hardware at present, we conclude that our model of the rotation gate in (2) closely models an actual experimental realization. The strengths δ and ϵ of systematic and random errors may then be adjusted to reflect the actual, physical realization architecture of the quantum gate. The more general case, modeled by a general 2×2 unitary matrix, will be discussed briefly in Sec. 6.

3 Defective Phase Rotation Gates

The exact n-qubit QFT of an integer input state $|a\rangle$ is defined as

$$|a'\rangle = \frac{1}{\sqrt{2^n}} \sum_{b=0}^{2^n-1} \Phi(a,b)|b\rangle, \tag{3}$$

where

$$\Phi(a,b) = \exp(2\pi i ab/2^n) \tag{4}$$

is the associated normalized phase between the input state $|a\rangle$ and the output state $|b\rangle$. Constructing (4) on a gate–by–gate level, we obtain

$$\Phi(a,b) = \prod_{m=0}^{n-1} e^{i\pi a_{[m]} b_{[n-m-1]}} \prod_{j=1}^{n-1} \prod_{l=0}^{n-j-1} e^{i a_{[l]} b_{[n-l-j-1]} \theta_j}, \tag{5}$$

where θ_j is defined in (1), and the subscript $[\ldots]$ of $a_{[m]}$ and $b_{[n-m-1]}$ denotes the mth and $(n-m-1)$th binary digit of a and b, respectively. We note that the m product in (5) arises from the Hadamard operations in the QFT, and the j and l products are due to the CROT gates in the QFT.

In this paper we study the QFT equipped with defective CROT gates, i.e., we replace the exact rotation angle θ_j in (5) with the modified angles $\tilde{\theta}_j$, defined in (2). Clearly, the introduction of defects into the CROT gates of the QFT reduces the ability of the QFT to detect the periodicity ω of an input state. We call the success probability with which ω can be determined in a single run of the QFT the performance of the QFT. For simplicity we compute the success probability as the probability $\tilde{P}$ of finding the periodicity ω of an input state with a single run of classical post processing (CPP_1) via the continued fraction method [1]. In order to form a first impression of how static defects affect the performance of the QFT, we choose $n = 10$ and $\omega = 30$ and plot in Fig. 1 the performance $\tilde{P}$ of the defective QFT, obtained via CPP_1, as a function of δ for three different values of ϵ.

Clearly, according to Fig. 1, the performance sacrifice is less than 50% when we shift δ by up to $\pm\pi/3$ away from the peak locations at $\pm\pi$. This is surprising since the result points to

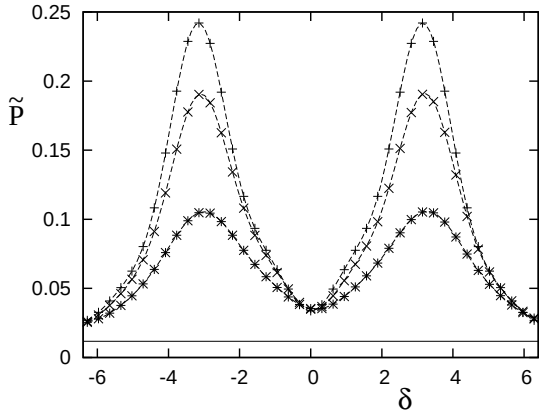

Fig. 1. Success probability $\tilde{P}$ of finding the periodicity $\omega = 30$ with a 10-qubit, defective QFT as a function of rotation-angle offset δ. Shown are three cases with random defect strengths $\epsilon = 0$ (pluses), $\epsilon = 0.5$ (crosses), and $\epsilon = 1$ (asterisks). Smooth, dashed, interpolation lines are drawn through the data points to guide the eye. The horizontal line corresponds to the pure-chance performance of the QFT.

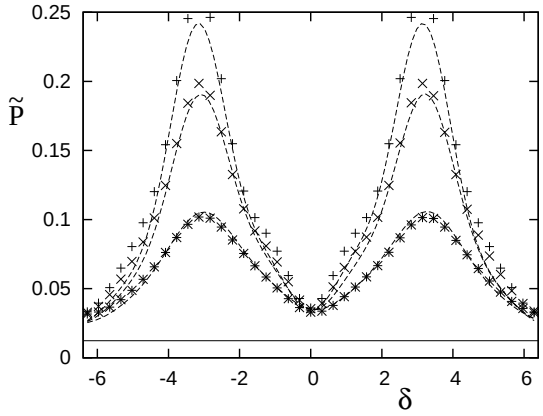

Fig. 2. Success probability of finding the periodicity $\omega = 30$ with a 14-qubit, defective QFT as a function of rotation-angle offset δ. Shown are three cases with random defect strengths $\epsilon = 0$ (pluses), $\epsilon = 0.5$ (crosses), and $\epsilon = 1$ (asterisks). The smooth, dashed lines are imported from Fig. 1 to show that the widths of the performance peaks centered around $\pm\pi$ are not altered. The horizontal line corresponds to the pure-chance performance of the QFT.

the fact that we can use rotation gates with large defects in actual hardware implementations of the QFT circuit, and are still able to find the periodicity ω with reasonable probability. In order to further confirm this observation, we show, in Fig. 2, the numerical results of the performance of the QFT with $n = 14$ and $\omega = 30$. Surprisingly, as demonstrated by the dashed lines in Fig. 2, imported from Fig. 1, the performance is nearly identical to the case with $n = 10$ (see Fig. 1). In particular, the widths of the performance peaks around $\delta = \pi$ (the δ corresponding to the exact QFT), are unaltered.

Another surprising observation, found in both Figs. 1 and 2, is the fact that even for $\delta = 0$, which corresponds to completely random CROT gates, we obtain a positive performance of the defective QFT of about 3%. This may be contrasted with the pure-chance performance of the QFT, which is obtained by assuming that each output state appears with equal probability. In this pure-chance case, the QFT is about 1% effective. Therefore, at $\delta = 0$, the 3% performance of the defective QFT is clearly significant compared to the pure-chance case. The 2% difference is explained by the action of the Hadamard gates.

4 Performance scaling: numerical results

The question of why the interpolation lines in Fig. 1 represent the data points in Fig. 2 so well is hard to answer analytically. Numerically, however, we find a new phenomenon that corroborates this observation. In Fig. 3, we plot the performance $\tilde{P}$ for a fixed periodicity $\omega = 30$ and fixed offset $\delta = \pi$ as a function of the number of qubits n of the defective QFT for four different defect strengths ϵ. We see that for all ϵ shown in Fig. 3, the performance $\tilde{P}$ quickly reaches a plateau for increasing n. This is unexpected since, naively, one might think that more qubits n lead to more defective CROT gates, which pick up more errors, and thus reduce the performance. Heuristically, we obtain the following functional form of the plateaus

$$\tilde{P}_{n\to\infty}(n,\omega,\epsilon) = \tilde{P}_{n\to\infty}(n,\omega,\epsilon=0)e^{-\gamma\epsilon^2}, \tag{6}$$

represented by the solid lines in Fig. 3. This behavior may be explained in the following way.

First, $\tilde{P}_{n\to\infty}(n,\omega,\epsilon=0)$ is expected to be

$$\tilde{P}_{n\to\infty}(n,\omega,\epsilon=0) = \frac{\nu(\omega)}{\omega}, \tag{7}$$

where ω is the periodicity of the input state and $\nu(\omega)$ is the co-prime counting function that returns the number of co-primes of ω that are larger than 0 and smaller than ω. This is so, since the CPP_1-based performance criterion essentially picks those Fourier peaks, say indexed with q, whose indices are relatively prime to ω. Since the index q runs from 0 to $\omega - 1$ inclusively, and the effective probability contribution that the qth peak and its vicinity represent is approximately $1/\omega$ for an exact QFT, we obtain (7). We note that for the current case of $\omega = 30$, $\nu(\omega = 30) = 8$, i.e., $\tilde{P}_{n\to\infty}(\omega = 30, \epsilon = 0) = 8/30$. The initial sub-plateau performance for small n (see Fig. 3) is also expected since in such cases the classical post-processing criterion does not include enough vicinity around the useful qth peak.

The ϵ-scaling in (6) is motivated by the ϵ-scaling results obtained in Ref. [4]. We have previously shown [4] that the effect of random defects with strength ϵ on the normalized performance $P = \tilde{P}(\epsilon)/\tilde{P}(\epsilon = 0)$, of a quantum computer equipped with a defective QFT is Gaussian, i.e., the normalized performance P of the quantum computer decreases according

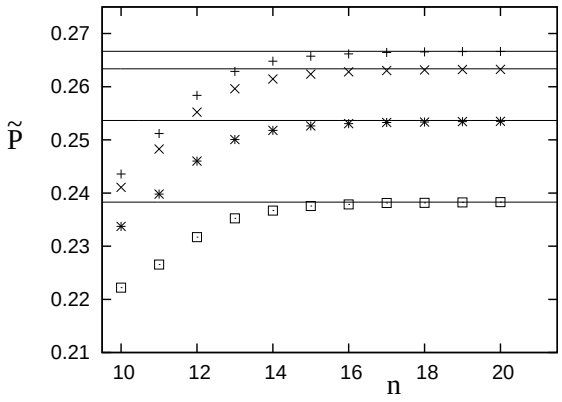

Fig. 3. Success probability of finding the periodicity $\omega = 30$ with an n-qubit QFT with fixed $\delta = \pi$ as a function of the number of qubits n. Shown are four cases with random defect strengths $\epsilon = 0$ (pluses), $\epsilon = 0.1$ (crosses), $\epsilon = 0.2$ (asterisks), and $\epsilon = 0.3$ (squares). The solid lines are the analytical prediction (6) with $\gamma = 1.25$ and $\tilde{P}_{n\to\infty}(\omega = 30, \epsilon = 0) = 8/30$.

to a Gaussian function in ϵ. The difference with respect to [4] is that, in the current paper, the success criterion incorporates the classical post-processing criterion CPP_1, a more advanced success criterion compared with the non-trivial nearest-peak criterion used in Ref. [4]. Therefore, while we expect that the Gaussian functional form in ϵ carries over, it is not justified to use the analytical result for the coefficient γ of ϵ^2 as computed in Ref. [4] in (6). As a consequence, we leave γ as a free fit parameter. For the specific case of $\omega = 30$, we find $\gamma = 1.25$.

5 Performance scaling: analytical results

Having identified the reasons for why the performances of the defective QFTs with $n = 10$ and $n = 14$ qubits are nearly identical (see Figs. 1 and 2), we now continue our investigation of performance scaling in n, δ, and ϵ, with an emphasis on analytical results. It is of utmost importance to assess the behavior of the width of the performance peaks in Figs. 1 and 2 as a function of n, since practical implementations of the QFT will always be realized with real-world CROT gates that have some non-zero defects. If the peaks would narrow quickly, say exponentially, as a function of n, we would have to impose impossibly small defect tolerances on the CROT gates of the QFT to be of any use in practical applications. Luckily, as we shall see, this is not the case. As already indicated by Figs. 1 and 2, the following analytical calculations will show that the width of the performance peaks stays constant as a function of n. In order to show this, and in order to get an analytical grip on the problem, we start by defining the normalized performance proxy

$$\hat{P}(s_0, \delta) = \left| \frac{1}{K(s_0)} \sum_{k=0}^{K(s_0)-1} \tilde{\Phi}(s_0 + k\omega, \hat{b}, \delta) \right|^2 , \tag{8}$$

where $K(s_0)$ is the number of elements of the equivalence class $[s_0]$ that comprises the integers

$s_0 + k\omega < 2^n$, where $k \geq 0$ is an integer, $\hat{b}$ is the proxy output state, and $\tilde{\Phi}$ denotes the amplitude defined in (5), where the exact phase angle θ_j defined in (1) is replaced by $\tilde{\theta_j}$ defined in (2). For the time being we will also set $\epsilon = 0$. According to (8), the performance proxy is the probability to obtain $\hat{b}$ as the result of a measurement of the output state of the defective QFT. The probability $\hat{P}(s_0, \delta)$ is a "proxy" measure because $|\hat{b}\rangle$ is assumed to represent all other states that may be obtained with high probability as the result of a measurement of the QFT output state, and it is assumed additionally that classical post processing of $\hat{b}$ reveals ω with certainty. Currently, both assumptions are necessary to make any analytical headway in the analysis of the performance scaling of the defective QFT for large n. Both assumptions are benign, since we checked previously [4, 6, 30] that (i) the probability of obtaining output states $|b\rangle$ in the vicinity of a main Fourier peak produced by the QFT, i.e., the set of states from which the proxy state $|\hat{b}\rangle$ is chosen, are correlated in the sense that they respond in unison to the introduction of static defects into the gates of the QFT, and (ii) the probability of obtaining ω via classical post-processing from any of the states $|b\rangle$ close to a main Fourier peak is indeed 1 or close to 1.

The central ingredient in the proxy measure $\hat{P}(s_0, \delta)$ defined in (8) are the amplitudes $\tilde{\Phi}(s_0 + k\omega, \hat{b}, \delta)$ and it is convenient to define the phases $\Delta(k, \delta)$ according to

$$\tilde{\Phi}(s_0 + k\omega, \hat{b}, \delta) \equiv e^{i\Delta(k,\delta)}. \tag{9}$$

Evaluating the amplitudes $\tilde{\Phi}$ for the special case $\delta = 0$, we notice that there are two equivalence classes, denoted by $[[1]]$ and $[[-1]]$, depending on whether $\tilde{\Phi}(s_0 + k\omega, \hat{b}, \delta = 0)$ evaluates to 1 or to -1 for given $s_0 + k\omega$ and $\hat{b}$. Since for $\delta = 0$ (and $\epsilon = 0$; see assumptions above) only the m product in (5) is active, whether $\tilde{\Phi}(s_0 + k\omega, \hat{b}, \delta = 0)$ is 1 or -1 depends entirely on the number of non-zero $(s_0 + k\omega)_{[m]}\hat{b}_{[n-m-1]}$ terms in the exponent of the m product in (5). We also note that the phases $\Delta(k, \delta = 0)$ take only two values. We denote them by $\Delta_1(k, \delta = 0) = 0$ for amplitudes in equivalence class $[[1]]$ and $\Delta_{-1}(k, \delta = 0) = \pi$ for amplitudes in equivalence class $[[-1]]$.

We now examine the amplitudes $\tilde{\Phi}(s_0 + k\omega, \hat{b}, \delta)$ in the other special case, i.e., $\delta = \pi$. Recalling that we currently have $\epsilon = 0$, $\delta = \pi$ is the case in which the QFT is exact. In this case, in order to produce the sharp peaks in b, characteristic for the exact QFT, all amplitudes $\tilde{\Phi}(s_0 + k\omega, \hat{b}, \delta)$, independent of k, have nearly the same value, which corresponds to a sharp clustering of the phases $\Delta(k, \delta = \pi)$ around a constant, denoted by Δ_π.

We now turn to the general, interpolating case, i.e., we examine the amplitudes $\tilde{\Phi}(s_0 + k\omega, \hat{b}, \delta)$ in the δ interval $0 < \delta < \pi$. Apparently, switching from $\delta = 0$ to $\delta = \pi$, the phases $\Delta_1(k, \delta = 0) = 0$ and $\Delta_{-1}(k, \delta = 0) = \pi$, corresponding to the equivalence classes $[[1]]$ and $[[-1]]$, have to coalesce into the single phase $\Delta_1(k, \pi) = \Delta_{-1}(k, \pi) = \Delta_\pi$ at $\delta = \pi$. We confirmed numerically that the phase difference between the amplitudes of the two equivalence classes $[[1]]$ and $[[-1]]$ is a decreasing function of δ, and moreover that it is linear according to

$$\Delta_{-1}(k, \delta) - \Delta_1(k, \delta) = \pi - \delta. \tag{10}$$

We check immediately that this formula reproduces the correct results in the two special cases $\delta = 0$ and $\delta = \pi$. Denoting by $\mu_1(s_0)$ and $\mu_{-1}(s_0)$ the number of elements of the two equivalence classes $[[1]]$ and $[[-1]]$, respectively, the performance proxy (8) may be evaluated

directly with the result

$$\hat{P}(s_0,\delta)=\frac{\mu_1(s_0)^2+\mu_{-1}(s_0)^2+2\mu_1(s_0)\mu_{-1}(s_0)\cos(\pi-\delta)}{K(s_0)^2}. \tag{11}$$

This formula shows that the performance of the QFT with respect to δ is independent of n. Hence, the analytical formula (11), which is n-independent, supports our numerical result (see Sec. 3) that the width of the peaks in Fig. 1 and 2 stays constant when n is increased from $n=10$ to $n=14$. The main power of (11), however, is that it is not only applicable in the low-n case (for instance $n=10$, $n=14$) but is valid in the large-n case (for instance $n\sim 1000$), not reachable by numerical calculations. We note that the worst case scenario [the minimum of $\hat{P}(s_0,\delta)$ under the constraint $\mu_1(s_0)+\mu_{-1}(s_0)=K(s_0)$] occurs for $\mu_1=\mu_{-1}$, in which case the performance proxy simplifies to

$$\hat{P}^{(\mathrm{W})}(\delta)=\frac{1+\cos(\pi-\delta)}{2}, \tag{12}$$

which, up to a normalization factor, captures the behavior of $\tilde{P}$ in Figs. 1 and 2 reasonably well.

Although the theory developed above assumes $\epsilon=0$, it also works for non-zero ϵ. We see this in the following way. Replacing θ_j in (5) with $\tilde{\theta}_j$ [see (2)], we may write approximately

$$\tilde{\Phi}(a,b)\approx\prod_{m=0}^{n-1}e^{i\pi a_{[m]}b_{[n-m-1]}}\prod_{j=1}^{n-1}\prod_{l=0}^{n-j-1}e^{ia_{[l]}b_{[n-l-j-1]}\delta/2^j}\prod_{j'=1}^{n-1}\prod_{l'=0}^{n-j-1}e^{ia_{[l']}b_{[n-l'-j'-1]}\epsilon r/2^{j'}}, \tag{13}$$

where we expanded the double product over j,l into a quadruple product over j,l,j',l', assuming that, because of the near statistical independence of $a_{[l]}b_{[n-l-j-1]}$ and $a_{[l']}b_{[n-l'-j'-1]}$, the off-diagonal elements average to a total factor 1, effectively reducing the quadruple product to the original double product on average. The form (13) of $\tilde{\Phi}(a,b)$ has the advantage that for $\delta=\pi$ the product in (13) for an input state with periodicity ω has already been evaluated in [4] and the resulting performance, consistent with the notation in Sec. 5, was denoted by $\tilde{P}(n,\omega,\epsilon)$. Thus, assuming that δ is close enough to π that the phase angles resulting from the m,j,l products for $\delta\neq\pi$ are still closely clustered (as is the case in the exact QFT), the double product over j',l' in (13) results in $\tilde{P}(n,\omega,\epsilon)$ and the normalized performance P of the QFT may be written in the form

$$P\approx\hat{P}(s_0,\delta)\frac{\tilde{P}(n,\omega,\epsilon)}{\tilde{P}(n,\omega,\epsilon=0)}. \tag{14}$$

In the limit $n\to\infty$, using (6), we obtain

$$P_{n\to\infty}\approx\hat{P}(s_0,\delta)e^{-\gamma\epsilon^2}, \tag{15}$$

which applies if the success criterion is based on CPP_1. Since $\hat{P}$ is normalized to 1 at $\delta=\pi$ and $\epsilon=0$, $P_{n\to\infty}$ may now be written approximately in the form

$$P_{n\to\infty}(\delta,\epsilon)\approx\left[\frac{1+\cos(\delta-\pi)}{2}\right]e^{-\gamma\epsilon^2}, \tag{16}$$

where we used $\hat{P}^{(\mathrm{W})}$ of (12). The analytical formula (16) qualitatively explains the double-hump structure in Figs. 1 and 2. It also explains why the widths of the peaks in Figs. 1 and 2 are constant, independent of ϵ.

Formula (13) predicts $P_{n\to\infty}(\delta,\epsilon)=0$ for $\delta=0$, independent of ϵ, while Figs. 1 and 2 show a positive performance of the defective QFT of about 3%. This positive offset, previously explained due to the action of the Hadamard gates, is still included in formula (11). It disappeared in (16) only because we used the worst-case performance estimate (12) as input for (16). Still, in terms of an estimate and in terms of a concise analytical formula, a 3% underestimate of the actual performance of the defective QFT by (16) is acceptable.

In general, (16) may be used as a convenient first estimate of the performance of the defective QFT in situations where classical simulations are too expensive or simply impossible to perform.

6 Discussion

Generally in quantum computing and quantum information processing, analytical results are necessary since the n-regions of practical interest, i.e., n of the order of 1000 [1, 31], are completely inaccessible to numerical simulation on classical, digital computers. Thus, only analytical techniques are powerful enough to explore these large-n regions. Following this line of reasoning, we derived the scaling law (16) of the performance of the QFT equipped with defective CROT gates as defined in (2). Since γ in (16) is a free parameter when the performance measure is based on CPP_1 instead of using the nearest non-trivial peak criterion as in Ref. [4], we need to better understand how γ behaves as a function of ω. To this end, we plot, in Fig. 4 (plus symbols), $\gamma(\omega)$ for all ω ranging from 3 to 30, excluding powers of 2, for which $\gamma=0$ (since the QFT performs perfectly in these cases, independently of ϵ [4, 6, 30]). The dotted line $\sim\ln(\omega)$ in Fig. 4 is drawn to guide the eye. It suggests that γ scales logarithmically in ω on average.

An important remark regarding γ is to be noted here. Although in Ref. [4] we derived analytically that γ is a constant, strictly speaking, even with the nearest-peak criterion, the coefficient γ carries an ω-dependence. In fact, we find that the dependence of γ on ω for the nearest-peak criterion is nearly identical to that found in Fig. 4, i.e., $\gamma\sim\ln(\omega)$. This discrepancy between the analytical, constant γ and the numerical, ω-scaling γ arises, because the analytical procedures used in Ref. [4], which yield the performance scaling formula that is ω independent, assume statistical independence wherever possible, i.e., the problem of solving for the performance scaling of the defective QFT is essentially reduced to how much drift in phase angle one obtains in (5) for a single integer input state toward a single integer output state. Since the multiplicity of the integer states for a periodic input state is then taken into account via an averaging process, together with the same procedure applied to the output state, the analytical performance scaling formulas in Ref. [4] are then ω independent. As is correctly pointed out in Ref. [4], the ω-dependence arises by taking the ω-induced residual coupling into account, effectively resulting in $\gamma\to\gamma(\omega)$. This modification is n-independent, as demonstrated in Fig. 8 of Ref. [4], which is consistent with our observations in Fig. 4 that γ, for a given ω, is (trivially) independent of n.

Previously, in Ref. [30], in connection with Shor's algorithm, we showed that the periodicities ω associated with the semiprimes N to be factored scale exponentially in n, where n was

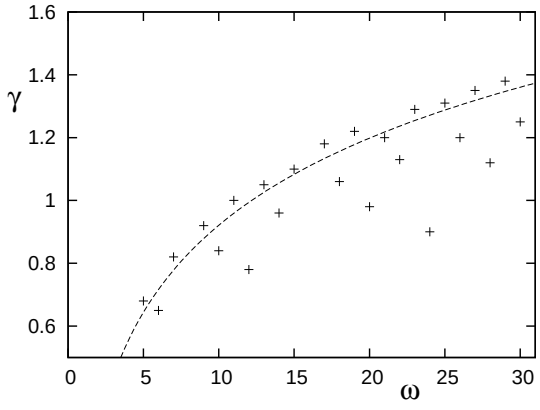

Fig. 4. Coefficient γ of ϵ^2 in the exponent of the performance scaling law (6) as a function of ω. The dotted line $\sim \ln(\omega)$ is drawn to guide the eye.

chosen such that it scales linearly in the bit-length of N. Therefore, used in connection with Shor factoring, γ scales linearly in n. This, however, should not distract from the fact that γ is primarily dependent on ω, and not on n. It becomes dependent on n only indirectly if, e.g. in connection with Shor's algorithm, the problem is posed in such a way that a one-to-one relationship exists between n and ω. In such cases, then, it does not matter whether we consider γ as a function of n or as a function of ω.

Because, as shown in Fig. 4, an approximate one-to-one correspondence exits between γ and ω, for *fixed* ω, we have γ independent of n. As a consequence, for fixed ω, the performance of the defective QFT does not deteriorate as a function of n. So far, we have been unable to uncover the underlying reason for this observation. Since the performance plateau does not show up using simpler performance criteria, such as the nearest-peak criterion [4, 6, 30], the performance plateau must have something to do with the use of our more realistic post-processing criterion CPP_1, which captures the actual performance characteristics of the defective QFT more accurately. This indicates that the reason for the plateau is number-theoretic in origin (CPP_1 involves a continued fractions analysis of output states $|b\rangle$ [1]) and certainly has something to do with the way defective CROT gates channel probability out of the main Fourier peaks into other useful states, bypassing states that are not useful for period-finding.

In an attempt to rule out any trivial explanations of the plateauing behavior shown in Fig. 3, for instance saturation of Fourier space with states useful for period-finding, we investigate the fraction ρ of useful states with respect to all possible output states. For instance, for the present example of $\omega = 30$, we find $\rho_{n=10} = 12/2^{10}$, $\rho_{n=15} = 404/2^{15}$, and $\rho_{n=20} = 12892/2^{20}$, leading to, approximately, $\rho \approx 1/80$, which is constant, independent of n. This means that for a totally randomized Fourier spectrum, we expect the success probability to be $\tilde{P} \approx 1/80$. This does not correspond to our observations, thus ruling out a simple, probabilistic explanation for the origin of the plateauing behavior. We note that the constancy of ρ is an interesting mathematical problem in itself, which might be accessible to a number-theoretic study of the CPP_1 criterion itself.

Since the total number of output states increases exponentially in n, and since for fixed ω the ratio of useful states to the total number of output states is a fixed constant ρ, the number of useful states for a fixed ω, then, increases exponentially in n as well. Hence, the average probability that a measurement of the output state of the QFT yields a state $|b\rangle$ useful for period-finding according to CPP_1, decreases exponentially in n according to $\tilde{P}_b \sim 1/(\rho 2^n)$. This means, that in the presence of some fixed background noise, not all states $|b\rangle$ that are theoretically useful for period finding will be useful in practice; in order to be useful in practice, their probability $\tilde{P}_b$ has to rise above some probability cut-off $\tilde{P}_{\mathrm{noise}}$, which depends on the background noise level. Therefore, we anticipate that despite the presence of exponentially many states that are theoretically useful, the probability cut-off may seriously limit the number of states that can be reliably used for period-finding, thus fundamentally modifying the plateau behavior of QFT performance.

In order to investigate this expectation more quantitatively, we plot, in Fig. 5, the performance $\tilde{P}^* = \sum_b \tilde{P}_b^*$ of the defective QFT with $\delta = \pi$ and $\epsilon = 0.3$ subjected to various cut-off levels $\tilde{P}_{\mathrm{noise}}$, normalized by $\tilde{P} = \sum_b \tilde{P}_b$, its noise-free counterpart, as a function of the number of qubits n. We define $\tilde{P}_b^* = \tilde{P}_b$ if $\tilde{P}_b > \tilde{P}_{\mathrm{noise}}$ and $\tilde{P}_b^* = 0$ if $\tilde{P}_b \leq \tilde{P}_{\mathrm{noise}}$. Thus, if a measurement of the output state of the QFT yields a theoretically useful state $|b\rangle$ with probability $\tilde{P}_b \leq \tilde{P}_{\mathrm{noise}}$, it is rejected from our performance measure. As might have been expected, for small n, where only relatively few useful states $|b\rangle$ with relatively large probabilities (probabilities have to sum to 1) exist, the background noise has little effect in suppressing these useful states and the plateau is relatively undisturbed. This is clearly demonstrated in Fig. 5, which shows a performance plateau for small n. However, Fig. 5 also shows that from some critical n_c on, when the probability level in the beneficial states, on average, starts to dive below the specified noise level, the plateau behavior breaks, the performance deteriorates, and continues to deteriorate for all $n > n_c$. The cut-off $\tilde{P}_{\mathrm{noise}}$ controls n_c, the onset of performance decay: the higher the noise cut-off $\tilde{P}_{\mathrm{noise}}$, the smaller n_c, i.e., the earlier the deterioration of the performance occurs in n. The plateau and its break at $n = n_c$ are predictions that can be checked experimentally.

The peeling off of useful output states $|b\rangle$ with rising noise cut-off $\tilde{P}_{\mathrm{noise}}$ provides us with a justification for choosing the nearest-peak criterion as a measure of period-finding success [4, 6, 30]: In the presence of noise, the peaks nearest to the theoretically expected main peak locations of the QFT output state, i.e., the highest peaks in the spectrum, are the ones that survive the longest as a function of n. Thus, the nearest-peak criterion may be used to measure QFT performance in situations with large background noise, while for low background noise levels, CPP_1 is more appropriate.

Now, an important question arises: How does the fraction ρ scale in ω? This needs to be investigated since, if ρ were a constant, say $\rho = 1/80$, as in the case $\omega = 30$, then, with probability 1 in 80 we would be able to obtain the periodicity of an input state simply with trial and error (picking random numbers we would be successful in 1 out of 80 trials on average), no matter how large ω. We might use this scheme to factor large integers, which cannot possibly be true. Therefore, we anticipate that ρ is a decreasing function of ω and the only remaining question is the functional form of this decrease.

In order to answer this question, Fig. 6 shows ρ as a function of ω. We see that ρ decreases approximately as $\rho \sim 1/\omega$ as ω increases. Therefore, as suspected, in the case of exponentially

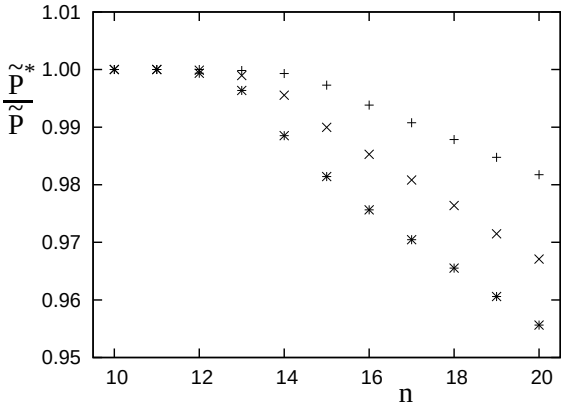

Fig. 5. Normalized success probability of finding the periodicity $\omega = 30$ as a function of the number of qubits n for fixed $\delta = \pi$ and $\epsilon = 0.3$, subjected to different noise levels $\tilde{P}_{\text{noise}}$. Shown are the data with $\tilde{P}_{\text{noise}} = 1 \times 10^{-5}$ (pluses), $\tilde{P}_{\text{noise}} = 3 \times 10^{-5}$ (crosses), and $\tilde{P}_{\text{noise}} = 7 \times 10^{-5}$ (asterisks). The transitions occur at $n_c = 14$, 13, and 12 for $\tilde{P}_{\text{noise}} = 1 \times 10^{-5}$, 3×10^{-5}, and 7×10^{-5}, respectively.

large ω, this functional behavior prohibits us from detecting the periodicity ω by simple trial and error. Moreover, since those ω that are of interest in cryptanalysis are exponentially large (see, e.g., [1, 30, 31]), we conclude, on the basis of Fig. 6, that ρ decays exponentially in n with exponentially increasing ω.

A note on our use of the CPP_1 criterion instead of the complete CPP criterion [1] is in order. We used CPP_1 for convenience only, expediting the (surprisingly time-consuming) classical post-processing procedure. If CPP_1 decides that an output state $|b\rangle$ is useful for period-finding, it is also useful according to CPP. Conversely, if CPP_1 rejects a state $|b\rangle$, CPP might find it useful. Thus, CPP_1 does not do anything wrong, it simply underestimates (slightly) the performance of the defective QFT. Thus, all performance estimates in this paper are lower bounds; the actual performance of the defective QFT is even better.

Most of our numerical simulations were performed with $n = 10$ or $n = 14$ qubits, and only the calculations in Fig. 3, 5, and 6 were performed with up to $n = 20$ qubits. This seems on the low side given the formidable computer power regularly available in modern research labs. Indeed, we had access to a state-of-the-art 300-core cluster computer on which all of our numerical computations were performed. However, as is generally known, the processing power of quantum computers is exponentially larger than the processing power of a classical digital computer, and while it is known that a digital computer can simulate all operations of a quantum computer, this simulation is exponentially inefficient. Therefore, at $n = 20$, even our 300-core cluster computer runs into a wall (the execution times become too large in real time) and if ensembles have to be simulated (in our case statistical ensembles of QFTs with systematic and random gate errors), the reasonable limit in classical computer resources that was available to us to be expended on this project dictated an n cut-off for our simulations at $n = 14$. We note that in Fig. 3, for instance, each point in the plot is the result of averaging over 10000 ensembles of random numbers. Clearly, larger n would be desirable, in particular

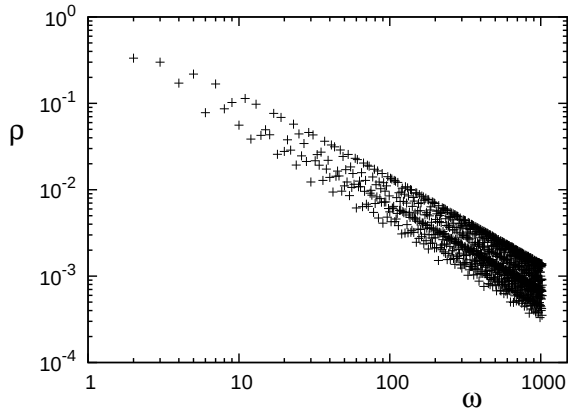

Fig. 6. Fraction ρ of states useful for period-finding to all possible 2^n output states as a function of the periodicity ω of the input state with ω ranging up to 2^{10}. The data shown are obtained from computing ρ with $n = 20$. Apparently, ρ scales like $\sim 1/\omega$.

for checking our analytical scaling laws. However, even with the help of supercomputers, because of the exponential explosion of required computational resources with n, the increase in accessible n values would only increase marginally. If not anything else, this situation illustrates two important points. (i) The awesome power of quantum computers (by dwarfing the classical computers used for their simulation) and (ii) the need for analytical formulas to access the large-n regimes forever inaccessible to classical computers.

While the current paper focuses on the model characterized by (2), we did run numerical simulations with a generic 2×2 unitary matrix, perturbed around the exact operation. Preliminary results show that population or amplitude errors in the rotation gate in fact result in noticeable deterioration of the performance of the QFT, apparently destroying the plateauing behavior shown in Fig. 3. This suggests that when it comes to manufacturing quantum hardware, one should try to minimize the amplitude errors as much as possible in order to fully take advantage of the performance plateau behavior associated with the phase-angle errors. The exact scaling relation of the combined errors are the subject of current investigations and their results will be reported elsewhere.

7 Summary and conclusions

In this paper, using the example of the QFT with defective CROT gates, we address Landauer's question [3] of whether static gate defects are fatal for quantum information processors. In general we find that the QFT is astonishingly resilient to the introduction of static defects, called "manufacturing flaws" by Landauer [3]. We investigate two types of static defects in the CROT rotation angles, systematic and random. Overall we find that even for defect strengths on the order of 30% and larger, acceptable QFT performance can be guaranteed, corroborating our findings in previous research [4, 6, 30] that already pointed to the extraordinary robustness of quantum information processors with respect to static gate defects. This means that static gate defects should not be an immediate source of worry for engineers and scientists working on designing and implementing these information processors as actual hardware com-

ponents. We also find that using a more advanced classical post-processing criterion (CPP_1, as described in Sec. 3) instead of the simpler nearest-peak criterion [4, 6, 30], the performance of the defective QFT reaches a performance plateau as a function of the number of qubits n, a prediction that may be checked experimentally. This result is significant both for the theoretical properties of the QFT and for practical applications, since it means that despite the proliferation of defective gates for increasing n, the performance stays the same. While this is a firm result, confirmed numerically by our simulations and supported by analytical arguments, the microscopic mechanism according to which this performance plateau arises currently remains an open problem awaiting further numerical and analytical work. While the fundamental origin of the performance plateaus is not clear at present, our analytical scaling laws reproduce the plateau levels with high accuracy.

References

1. M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
2. I. Garcia-Mata, K. M. Frahm, and D. L. Shepelyansky, Phys. Rev. A **75**, 052311 (2007).
3. R. Landauer, in *Quantum Computing and Communications*, edited by M. Brooks (Springer, London, 1999), p. 61.
4. Y. S. Nam and R. Blümel, Phys. Rev. A **89**, 042337 (2014).
5. R. B. Griffiths and C.-S. Niu, Phys. Rev. Lett. **76**, 3228 (1996).
6. Y. S. Nam and R. Blümel, Phys. Rev. A **88**, 062310 (2013).
7. F. Gaitan, *Quantum Error Correction and Fault Tolerant Quantum Computing* (CRC Press, Boca Raton, FL, 2008).
8. D. Coppersmith, arXiv:quant-ph/0201067.
9. A. Barenco, A. Ekert, K.-A. Suominen, and P. Törmä, Phys. Rev. A **54**, 139 (1996).
10. A. G. Fowler and L. C. L. Hollenberg, Phys. Rev. A **70**, 032329 (2004).
11. J. I. Cirac and P. Zoller, Phys. Rev. Lett. **74**, 4091 (1995).
12. J. Niwa, K. Matsumoto, and H. Imai, Phys. Rev. A **66**, 062317 (2002).
13. C. Miquel, J. P. Paz, and R. Perazzo, Phys. Rev. A **54**, 2605 (1996).
14. A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Phys. Rev. A **86**, 032324 (2012).
15. D. Wecker and K. M. Svore, arXiv:1402.4467v1 [quant-ph].
16. A. G. Fowler, A. C. Whiteside, A. L. McInnes, and A. Rabbani, Phys. Rev. X **2**, 041003 (2012).
17. A. G. Fowler, arXiv:1307.0689v1 [quant-ph].
18. A. Y. Kitaev, Russ. Math. Surveys **52**, 1191 (1997).
19. V. Kliuchnikov, D. Maslov, and M. Mosca, arXiv:1212.0822v2 [quant-ph].
20. N. J. Ross and P. Selinger, arXiv:1403.2975v1 [quant-ph].
21. P. Selinger, arXiv:1212.6253v2 [quant-ph].
22. A. Bocharov, M. Roetteler, and K. M. Svore, arXiv:1404.5320v2 [quant-ph].
23. A. Bocharov, M. Roetteler, and K. M. Svore, arXiv:1409.3552v2 [quant-ph].
24. P. Selinger, Newsynth: exact and approximate synthesis of quantum circuits, http://www.mathstat.dal.ca/∼selinger/newsynth/, 2013.
25. K. R. Brown, A. W. Harrow, and I. L. Chuang, Phys. Rev. A **70**, 052318 (2004).
26. A. G. Fowler and J. M. Martinis, Phys. Rev. A **89**, 032316 (2014).
27. Y. Tomita and K. M. Svore, arXiv:1404.3747v3 [quant-ph].
28. J. Ghosh and A. G. Fowler, arXiv:1406.2404v1 [quant-ph].
29. S. J. Devitt, W. J. Munro, and K. Nemoto, Rep. Prog. Phys. **76** 076001 (2013).
30. Y. S. Nam and R. Blümel, Phys. Rev. A **87**, 032333 (2013).
31. N. D. Mermin, *Quantum Computer Science* (Cambridge Univ. Press, Cambridge, 2007).