Windowed Quantum Phase Estimation: Signal Processing Approach to a Quantum Algorithm

Xue Wen
Samsung Research China - Beijing
xue.wen@samsung.com

ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) | 979-8-3503-6874-1/25/$31.00 ©2025 IEEE | DOI: 10.1109/ICASSP49660.2025.10888027

***Abstract*— Quantum phase estimation (QPE) is a key ingredient to some of the most important results of quantum computing. In this paper we examine the QPE algorithm through the looking glass of signal processing. This perspective helps to reveal a relation between the distribution of quantum phase estimates and the shape of the discrete Fourier transform (DFT) of a pure complex sinusoid. Guided by spectral shaping theory, we design a class of quantum circuits that are the quantum equivalents to classical windowing operators, for a family of window functions that have compact support in the DFT. These circuits have complexity linear in the number of target qubits, strictly lower than QPE itself. Using these circuits we are able to reshape QPE's performance characteristics. In particular, the asymptotic marginal cost for 2x reduction in fail rate (2x improvement in reliability) is reduced from one qubit to 1/3 qubit or lower, which can help reduce the qubit budget for a preset accuracy-reliability target, especially in high-reliability scenarios. This makes a demonstration of the value of signal processing principles in designing quantum algorithms.**

Keywords—quantum algorithms, quantum phase estimation, windowed Fourier transforms

I. Introduction

Quantum computing [1] exploits quantum-mechanical effects of physical systems to build quantum circuits and algorithms that show substantial speedup against classical counterparts in suitable tasks. States of quantum circuits are processed by unitary operators. It is no surprise that a fast algorithm to the unitary eigenvalue problem [2][3], called quantum phase estimation (QPE), enjoys special importance among quantum algorithms. QPE is key to several notable quantum algorithms claiming exponential speedup over classical alternatives, including integer factoring and discrete logarithm (Shor [4]), solving linear equations (HHL [5]), fitting data by linear regression [7], quantum versions of PCA [8] and SVD [9], and quantum counting [10]. The performance characteristics of QPE has profound implications on quantum computing.

From the signal processing perspective, QPE can be understood as estimating the frequency of a pure complex sinusoid from its discrete Fourier transform (DFT). Performance of the algorithm is determined by the shape of that discrete spectrum. We are interested in looking up whether classical spectral shaping techniques can help reshape the performance characteristics of QPE. In particular, we seek a quantum algorithm that implements *windowing*, a familiar operator in spectral analysis, with exponential speedup against its classical counterpart. Since windowing sharpens spectral peak of a complex sinusoid towards the true frequency, we expect windowed QPE to yield better estimates.

This paper is arranged with signal processing audience in mind. Section 2 reviews the fundamentals of quantum computing, which we retell in subspace terminology that is already familiar in signal processing. Section 3 recounts the QPE algorithm, and explains the relation between QPE performance, spectral shape, and windowing. Section 4 presents our quantum windowing circuit, and numerically examines how windowing affects QPE in terms of tradeoffs between accuracy, reliability, and resource. We conclude with further discussions in section 5.

II. Quantum computing

Quantum computers hold information in quantum-mechanical devices called quantum bits, or "qubits". Like classical bits, a qubit has two possible *eigenstates* $|0\rangle$ and $|1\rangle$. A qubit differs from a classical bit in that the qubit can be in *superposition* state, which is a normalized linear combination of eigenstates (consider eigenstates as orthonormal column vectors):

$$|\psi\rangle = a_0|0\rangle + a_1|1\rangle,\ a_0, a_1 \in \mathbb{C}, |a_0|^2 + |a_1|^2 = 1 \qquad (1)$$

The state space of a qubit is (isomorphic to) $\mathbb{C}^2$, in which valid states are those on the unit sphere. This has two canonical subspaces which are spanned by $|0\rangle$ and $|1\rangle$, respectively. We denote these subspaces as $|\boldsymbol{j} = 0\rangle$ and $|\boldsymbol{j} = 1\rangle$. Here the bold $\boldsymbol{j}$ stands for the qubit. $|\boldsymbol{j} = 0\rangle$ is simply the set of states in which qubit $\boldsymbol{j}$ remains in eigenstate $|0\rangle$, up to a global phase factor $e^{i\theta}, \theta \in \mathbb{R}$. We say that $|\psi\rangle$ has amplitude a_j, substate $a_j|j\rangle$, and energy $|a_j|^2$ in subspace $|\boldsymbol{j} = j\rangle$. Total energy of a state in all subspaces must sum to 1.

Things get more interesting when there are more than one qubits. A register of m qubits has 2^m eigenstates, indexed by all 2^m combinations of individual qubits' 0's and 1's. We may denote an m-qubit eigenstate either in binary form as $|b_{m-1} \dots b_1 b_0\rangle$, where $b_k \in \{0,1\}$ comes from the k^{th} qubit, or in integer form as $|j\rangle, j = \sum 2^k b_k$. An arbitrary state has the form $|\psi\rangle = \sum_{j=0}^{2^m-1} a_j|j\rangle, \sum_j |a_j|^2 = 1$. The state space is $\mathbb{C}^{2^m}$, to which the eigenstates form an orthonormal basis. Valid states are again those on the unit sphere. Notice that the dimension of the state space (hence data capacity of the register) grows exponentially with the number of qubits.

Eigenstate-aligned subspaces can be defined for an arbitrary bipartition of the m qubits. Choose a subset of n qubits, we denote this subset by $\boldsymbol{j}^n$ and its $(m-n)$-qubit complement by $\bar{\boldsymbol{j}^n}$. This split defines 2^n subspaces $|\boldsymbol{j}^n = j\rangle, j = 0, \dots, 2^n - 1$. An orthonormal basis of subspace $|\boldsymbol{j}^n = j\rangle$ consists of all 2^{m-n} eigenstates whose $\boldsymbol{j}^n$ bits evaluate to j. In each $|\boldsymbol{j}^n = j\rangle$ the subset $\boldsymbol{j}^n$ remains in eigenstate $|j\rangle$, while the other qubits ($\bar{\boldsymbol{j}^n}$) can be in any superposition state. The projection of a state $|\psi\rangle$ in $|\boldsymbol{j}^n = j\rangle$ is its substate in this subspace, which we denote $|\psi_{\boldsymbol{j}^n=j}\rangle$. To give it a concrete signal processing example, suppose we have a 1-D time signal $x = (x_0, x_1, \dots)$ of length 2^{16} encoded in the state of a 16-qubit register, i.e. $|\psi\rangle = \sum_{j=0}^{2^{16}-1} x_j|j\rangle$. Eigenstates of the 16 qubits form a complete time index from 0 to $2^{16} - 1$. Choosing the 6 highest qubits as $\boldsymbol{j}^n$, we get $2^6 = 64$ subspaces $|\boldsymbol{j}^6 = j\rangle, j = 0, \dots, 63$, of dimension 1024 each. This has the interpretation of splitting the whole signal into 64 disjoint time frames of length 1024, the j^{th} frame being the sequence of amplitudes of substate $|\psi_{\boldsymbol{j}^6=j}\rangle$. Conversely, the 10 lower qubits $\bar{\boldsymbol{j}^6}$ index another set of 1024 *dual* subspaces of dimension 64. The energy of a substate is the square sum of its amplitudes, in perfect agreement with the common notion of a signal frame's energy.

Quantum states are processed by unitary operators called gates and circuits, which may be thought as quantum computing analogies to instructions and algorithms. A gate operating on a single qubit is just a 2×2 unitary operator $U: \mathbb{C}^2 \to \mathbb{C}^2$, which takes input $|\psi\rangle$ to

TABLE I. EXAMPLES OF SINGLE-QUBIT GATES

Name	NOT	Hadamard	$\boldsymbol{R_k}$*	Y-rotation
Diagram symbol				
Matrix form	$\begin{bmatrix}0 & 1\\ 1 & 0\end{bmatrix}$	$\frac{1}{\sqrt{2}}\begin{bmatrix}1 & 1\\ 1 & -1\end{bmatrix}$	$\begin{bmatrix}1 & 0\\ 0 & e^{i2\pi/2^k}\end{bmatrix}$	$\begin{bmatrix}\cos\frac{\theta}{2} & -\sin\frac{\theta}{2}\\ \sin\frac{\theta}{2} & \cos\frac{\theta}{2}\end{bmatrix}$

* R_k also appears in a more familiar variant as "z-rotation" gate.

output $U|\psi\rangle$. A few examples are given in Table I, in which the matrix forms are to be left-multiplied to states written as column vectors. The first rows/columns in the matrices align to eigenstate $|0\rangle$, hence the Hadamard gate takes $|0\rangle$ to $(|0\rangle + |1\rangle)/\sqrt{2}$ ("equal superposition"), $|1\rangle$ to $(|0\rangle - |1\rangle)/\sqrt{2}$, and the R_k gate leaves $|0\rangle$ unchanged but takes $|1\rangle$ to $e^{i2\pi/2^k}|1\rangle$. These gates can be chained up to form single-qubit circuits. Circuit diagrams are read from left to right, each horizontal wire representing one qubit throughout its lifespan. For example the circuit 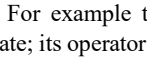 applies the R_k gate after X gate; its operator is $R_k X$, which is still unitary 2×2.

An intuitive way to explain the power of quantum computing is by considering what happens when a single-qubit gate is applied to one qubit, say $\boldsymbol{j}^1$, within an m-qubit register. Intuitively, this is a unitary operator on state space $\mathbb{C}^{2^m}$ and it shall leave any eigenstate of the other $m-1$ qubits unchanged. By the linearity of unitary operators, it is not hard to realize that the gate will remix the two substates $|\psi_{\boldsymbol{j}^1=0}\rangle$ and $|\psi_{\boldsymbol{j}^1=1}\rangle$, each of dimension 2^{m-1}, by

$$\left[|\psi_{\boldsymbol{j}^1=0}\rangle \quad |\psi_{\boldsymbol{j}^1=1}\rangle\right] \mapsto \left[|\psi_{\boldsymbol{j}^1=0}\rangle \quad |\psi_{\boldsymbol{j}^1=1}\rangle\right] U^\top, \qquad (2)$$

in which the substates are seen as column vectors, and $U \in \mathbb{C}^{2\times 2}$ is the gate's matrix. Equivalently, we may think the gate as being applied in all 2^{m-1} dual subspaces $|\boldsymbol{j}^{m-1} = k\rangle, k = 0, \dots, 2^{m-1}-1$, each of dimension 2. This is $O(2^m)$ parallelism by classical standards, brought about by one gate evolving an m-qubit register. In the 16-qubit example, one Hadamard gate on the highest qubit will compute the sum and difference between the first and second halves of the encoded signal. More generally, applying a circuit on a n-subset of m qubits "broadcasts" the n-qubit operator across all 2^{m-n} dual subspaces. For example, it is known that passing n qubits each through a Hadamard gate computes 2^n-point Hadamard transform in the state space. Applying this circuit with $n = 10$ on the 10 lowest qubits in our 16-qubit setup will return the Hadamard transforms of all 64 frames in substates $|\boldsymbol{j}^6 = j\rangle, j = 0, \dots, 63$.

Circuits consisting of single-qubit gates only are very limited as to what states are reachable from an initial state. To make full use of the state space the qubits must *interact* via multi-qubit gates. The only important multi-qubit gate is the controlled NOT (or CNOT) gate, in the sense that all unitary circuits can be equivalently built from single-qubit gates and CNOTs at reasonable overheads. CNOT is a gate operating on two qubits, called *control* ("c") and *target* ("t"), respectively. As shown in Fig. 1, CNOT takes eigenstate $|c, t\rangle$ to $|c, t \oplus c\rangle$, where $c, t \in \{0,1\}$, and $\oplus$ is the binary XOR operator. The interpretation is that CNOT does nothing in subspace $|\boldsymbol{c} = 0\rangle$, but applies NOT on qubit $\boldsymbol{t}$ in subspace $|\boldsymbol{c} = 1\rangle$. More generally, for any circuit U one can build its controlled version C-U $= \begin{bmatrix} I & \\ & U \end{bmatrix}$ at complexity linear in U, which takes an extra qubit $\boldsymbol{c}$ for input and applies U in subspace $|\boldsymbol{c} = 1\rangle$ only.

CNOT and NOT are examples of *permutation* circuits, which take each eigenstate to another eigenstate. Since unitary operators are invertible no two eigenstates are taken to the same target, so the outputs to the whole set of eigenstates must be its permutation. These circuits are said to simulate invertible classical algorithms, because they incur no new superposition. Conversely, all invertible classical algorithms can be quantumly simulated with linear overhead by a

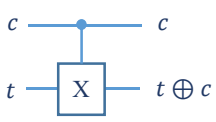

$$\text{C-NOT} = \begin{bmatrix} I & \\ & X \end{bmatrix} = \begin{bmatrix} 1 & 0 & & \\ 0 & 1 & & \\ & & 0 & 1 \\ & & 1 & 0 \end{bmatrix}$$

Fig. 1 Controlled NOT (CNOT) gate

permutation circuit. This allows using invertible classical subroutines in a quantum algorithm in the form of their quantum simulation. Applying a permutation circuit on a subset of qubits permutes their associated substates. Back to our 16-qubit example, a CNOT gate on the two highest qubits leaves the first half of the signal unchanged, but swaps the order of the third and fourth quarters; a classical increment operator on the top 6 qubits, i.e. $j \mapsto (j+1) \bmod 2^6$, circularly shifts the whole signal by 1024 points.

Quantum information in qubits can be converted to classical form by quantum *measurement*. Measuring a state $|\psi\rangle$ on a subset of n qubits, say $\boldsymbol{j}^n$, randomly reveals an eigenstate $j \in \{0, \dots, 2^n - 1\}$ as the classical outcome, with probability of seeing each j equal to the energy of substate $|\psi_{\boldsymbol{j}^n=j}\rangle$. Notice this is drawing a sample from an unstructured distribution of size 2^n in time $O(n)$. Depending on the outcome j, the system's state after measurement becomes its projection in subspace $|\boldsymbol{j}^n = j\rangle$, i.e. $|\psi_{\boldsymbol{j}^n=j}\rangle$, renormalized to unit energy. In our 16-qubit example, measuring the 6 higher qubits $\boldsymbol{j}^6$ randomly chooses a frame (prioritizing the louder ones), puts all energy in that frame, and makes all other frames zero. Besides retrieving the outcome, measurement is also used for its subspace projection capability to build non-unitary operators, as with the HHL algorithm [5]. In both cases the randomness involved may incur some chance of failure. Detectable failures result in wasted compute, which reduces an algorithm's *efficiency*; undetected failures lead to wrong results, which corrupts the algorithm's *reliability*. Measurement is drawn in circuit diagrams by the symbol , using a double wire to mark the output as a classical bit.

In summary, quantum computing offers data capacity, parallelism, and sampling speed that are exponential in the number of qubits. These are the main sources of quantum advantage against classical computing. Consequently, one need use sufficiently many qubits to achieve quantum advantage.

III. QUANTUM PHASE ESTIMATION AND ITS SIGNAL PROCESSING INTERPRETATION

Given unitary operator U and its eigenvector $|u\rangle$, QPE computes the corresponding eigenvalue $\lambda = e^{i\phi}$, or equivalently ϕ ("phase"), in classical form. The classical difficulty of the problem lies in the $O(2^m)$ size of the state space. The QPE algorithm proceeds in three steps: 1) use U and $|u\rangle$ to prepare a pure complex sinusoid $x_j = e^{i\phi j}, j = 0, \dots, 2^m - 1$, in an m-qubit target register; 2) quantumly compute the discrete Fourier transform (DFT) of this sinusoid in place; 3) measure the register to obtain outcome $\hat{k}$, and output estimate $\hat{\phi} = 2\pi\hat{k}/2^m$. Of the three steps, step 2 is well understood to be $O(m^2)$, thanks to a quantum Fourier transform circuit that does the job using m Hadamard and $m(m-1)/2$ controlled R_k gates, while step 3 is $O(m)$. Complexity of step 1 largely depends on the operator U. Several "best case" use cases [4][5] have bounded step 1 to $O(m^2)$, so that the overall complexity is comparable to quantum DFT. We are interested in applying windowing to the Fourier transform step without significant overhead. It needs be at most $O(m^2)$ to be applicable to general QPE use cases.

This windowing is motivated by observing steps 2, 3 of the QPE algorithm as estimating the frequency of sinusoid $x_j = e^{i\phi j}$ from its DFT, seeing j as "time". The DFT is computed in step 2 as the state of the m-qubit target register; the energy in the k^{th} DFT bin being that of eigenstate $|k\rangle$. By the rules of quantum measurement, step 3 draws a bin from the energy distribution over all bins, which is the

energy spectrum. Since we know from classical signal analysis that windowing concentrates spectral energy towards true frequency, it is natural to expect tighter estimates from windowed QPE.

More specifically, let $|x\rangle \in \mathbb{C}^{2^m}: x_j = e^{i\phi j}, \phi = 2\pi f$ be the sinusoid in the target register. We choose a window function $|w\rangle \in \mathbb{C}^{2^m}$, and feed their element-wise product $|wx\rangle$ to the quantum DFT. The energy spectrum from this DFT is given by

$$|\hat{x}_k^w|^2 = |\hat{w}(k/M - f)|^2, k = 0, \dots, M-1, \quad (3)$$

where $M = 2^m$, $\hat{w}(f), f \in \mathbb{R}$ is the discrete-time Fourier transform (DTFT) of $|w\rangle$, defined by $\hat{w}(f) \coloneqq \frac{1}{\sqrt{M}} \sum_{j=0}^{M-1} w_j e^{-i2\pi jf}$. Windowed QPE draws $\hat{k}$ from the distribution $P(k) = |\hat{x}_k^w|^2$. Typical window functions in spectral analysis have their energy concentrated about a main peak at some $\delta_{\text{pk}}^w \in (-1/2, 1/2]$ bins; real symmetric windows have $\delta_{\text{pk}}^w = 0$. The quantum phase estimate using a window of arbitrary δ_{pk}^w is given by $\hat{\phi} = 2\pi(\hat{k} - \delta_{\text{pk}}^w)/M$.

Fixing m, the performance of QPE is characterized by a tradeoff between accuracy and reliability. We say QPE meets performance target (n, ϵ) if it attains n-bit accuracy with probability at least $1 - \epsilon$ for all $f \in [0,1)$. Given f, the probability of failing to attain n-bit accuracy is explicitly

$$\epsilon_n(f) = \sum_{k=K_1^{(n)}}^{K_2^{(n)}} |\hat{w}(k/M - f)|^2, \quad (4)$$

$K_1^{(n)} = \lceil Mf + \delta_{\text{pk}}^w + 2^{m-n-1} \rceil$, $K_2^{(n)} = K_1^{(n)} + 2^m(1 - 2^{-n}) - 1$. Equation (4) sums over frequencies uniformly placed in the tails of $|\hat{w}|^2$ cut off at $\pm 2^{m-n-1}$ bins from the main peak. Standard QPE uses a rectangular window, which has notoriously long tails in the DTFT, and consequently high fail rates. It is possible to cut the fail rate by half, asymptotically, by adding one more qubit to the target register. We propose to reduce the fail rate by using windows with shorter tails. The goal is to do so at complexity no more than $O(m^2)$. We notice that the previous attempt at quantumly computing windowed Fourier transforms [12] did not meet this goal, as it had relied on naively applying a generic state preparation routine [13] at complexity at least $O(2^m)$.

IV. Quantum windowing circuit and its use in QPE

Fast quantum windowing can be achieved with a special class of window functions, the cosine family, whose DFTs have compact support about zero frequency. These have the explicit form

$$w_j \coloneqq \frac{1}{z_c\sqrt{M}} \sum_{r=R_-}^{R_+} c_r e^{i2\pi rj/M}, j = 0, \dots, M-1. \quad (5)$$

$R_- \le 0, R_+ \ge 0$ are small integers, $c_r \in \mathbb{C}$ is the r^{th} bin of the DFT, and $z_c = \sqrt{\sum_r |c_r|^2}$ a normalizer. These windows have "lobes" which are narrow intervals between regularly appearing zeros in the DTFT, with the main lobe striding $R_- - 1$ to $R_+ + 1$ bins. Hann, Hamming and Blackman windows are examples from this family. By polynomial factorization we can rewrite the above as

$$w_j = \frac{1}{z_d} \prod_{r=R_-}^{-1} \left(1 - d_r e^{-i2\pi j/M}\right) \prod_{r=1}^{R_+} \left(1 - d_r e^{i2\pi j/M}\right) \quad (6)$$

Fig. 2 gives the quantum circuit for computing DFT with this family of windows. The circuit uses $R \coloneqq R_+ - R_-$ auxiliary qubits $a_{R_-}, \dots, a_{-1}, a_1, \dots, a_{R_+}$, initialized to $|0\rangle$. The single-qubit $U(d)$ gates are parameterized by the values in (6). $FT^\dagger$ is the quantum DFT circuit. The $R_m^{\pm j}$ gates are linear phase shifters, mapping amplitudes of $|x\rangle$ by $x_j \mapsto x_j e^{\pm i2\pi j/M}$. Fig. 3 shows how the controlled version of R_m^j is built from m controlled R_k gates. We explain the circuit in Fig. 2 by showing that the subcircuit in dashed frame (bottom right) computes DFT of $|x\rangle$ with a window $w_j \propto 1 - d_{R_+} e^{i2\pi j/M}$, the last factor in (6); the rest is just nested application of the same construct. Starting with input $|x\rangle$ in the target register, the addition of auxiliary

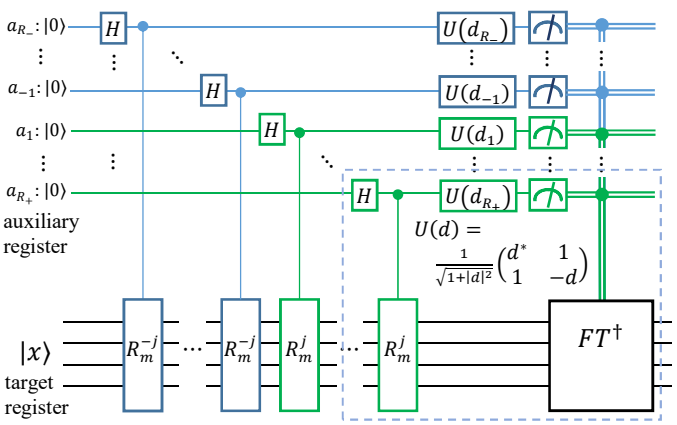

Fig. 2 Quantum circuit of windowed DFT

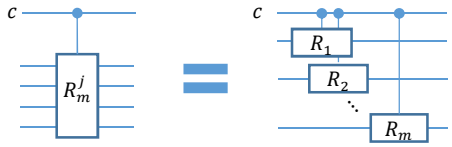

Fig. 3 Controlled linear phase shifter

qubit a_{R_+} in state $|0\rangle$ doubles the dimensionality of the state space, putting $|x\rangle$ in $|a_{R_+} = 0\rangle$, nothing in $|a_{R_+} = 1\rangle$. The Hadamard gate on a_{R_+} takes it to equal superposition of $|0\rangle$ and $|1\rangle$, effectually leaving identical copies of $|x\rangle$ in both subspaces. Controlled by a_{R_+}, the linear phase shifter maps $x_j \mapsto x_j e^{i2\pi j/M}$ in subspace $|a_{R_+} = 1\rangle$, while leaving $|x\rangle$ unchanged in $|a_{R_+} = 0\rangle$. The gate $U(d_{R_+})$ on a_{R_+} remixes the two substates, leaving correctly windowed signal $x_j - d_{R_+} x_j e^{i2\pi j/M}$ in $|a_{R_+} = 1\rangle$. Conditioned on measuring a_{R_+} to 1, this windowed substate renormalizes to become the full state in the target register. The quantum DFT then completes the job.

The windowing circuit (without DFT) uses R Hadamard gates, R $U(d)$ gates, and mR controlled R_k gates. This is strictly lower complexity than quantum DFT with $m(m-1)/2$ controlled R_k gates, for any m large enough to show quantum advantage. We note that for real-valued d the $U(d)$ gate reduces to rotation; for $d = 1$ it reduces to the Hadamard gate. Efficiency-wise the windowed DFT meets detectable failure if any of the R auxiliary qubits measure to 0. The conditioning on measurements is needed because windowing operators are generally non-unitary. We estimate the probability of measuring all auxiliary qubits to 1 to be between 0.16 and 0.5 for a typical set of windows ($P(\mathbf{1})$ column in Table II).

A. Numerical results

We study the effect of windowing on QPE by evaluating the performance profile expressed by (4) for several common cosine family windows listed in Table II. Direct evaluation of (4) becomes unwieldy for large m's due to exponential data size, so we seek to approximate (4) by numerical integration. To do so we factor out the faster oscillations from $|\hat{w}(k/M - f)|^2$:

$$|\hat{w}(k/M - f)|^2 = \frac{\sin^2 \pi Mf}{z_c^2 M^2} \cdot |W(k)|^2,$$

$$\text{where } W(k) \coloneqq \sum_r c_r e^{-i\pi \frac{r}{M}} \csc \pi \left(\frac{k-r}{M} - f\right). \quad (7)$$

Equation (7) is valid outside the main lobe. With $W(k)$ a smooth function and $\sin^2 \pi Mf / z_c^2 M^2$ independent of k, we are able to approximate (4) by integrating $|W|^2$ over $(K_1^{(n)} - 0.5, K_2^{(n)} + 0.5]$ using an off-the-shelf integrator. Calculations are carried out at 240-bit precision, which is sufficient for computing $\epsilon_n(f)$ with window types listed in Table II and m up to 64. We let $\epsilon_n = \sup_f \epsilon_n(f)$.

Table I Setup, attributes, and QPE results of evaluated window types

Window type	$c^{\lozenge\lozenge}$ ($\underline{c_0}$ underscored)	d ($d_{r<0}; d_{r>0}$)	Side lobe energy (%)	Spectral rolloff	δ_{pk} (bins)	$R^{\lozenge}$	$P(\mathbf{1})^{\lozenge}$	ϵ_{m-1} (%)	ϵ_{m-5} (%)	ϵ_n's rolloff$^{\lozenge\lozenge\lozenge}$
Rectangular	$\underline{1}$	-; -	9.7	6dB/oct $^{\square}$	0	0	-	0.19	0.013	2x/bit
Half sine	$\underline{1}, -1$	-; 1	0.5	12dB/oct	0.5	1	0.5	0.099	8.3E-6	8x/bit
Hann	$-1, \underline{2}, -1$	1; 1	0.05	18dB/oct	0	2	0.375	0.17	2.6E-8	32x/bit
Hamming	$-23, \underline{54}, -23$	0.559; 0.559	0.04	6dB/oct	0	2	0.341	0.13	2.0E-4	2x/bit
Hann$^{1.5}$	$-1, \underline{3}, -3, 1$	1; 1,1	0.007	24dB/oct	0.5	3	0.313	0.22	2.0E-10	128x/bit
Hann^2	$1, -4, \underline{6}, -4, 1$	1,1; 1,1	0.001	30dB/oct	0	4	0.273	0.26	2.7E-12	512x/bit
Blackman	$4, -25, \underline{42}, -25, 4$	0.25,1; 1,0.25	0.0002	18dB/oct	0	4	0.165	0.22	4.0E-9	32x/bit

$^{\lozenge}$ R: number of auxiliary qubits, equals the main lobe width minus 2. $P(\mathbf{1})$: probability of measuring auxiliary bits to all 1's.
$^{\lozenge\lozenge}$ Scaling all elements of c by a positive number results in the same window.
$^{\lozenge\lozenge\lozenge}$ Asymptotic scaling of fail rate for each additional bit of accuracy (roughly holds up to $m-5$).
$^{\square}$ "oct" stands for octave. Spectral roll-off of $3 \times p$ db/oct corresponds to $1/f^p$-like decay of energy spectrum.

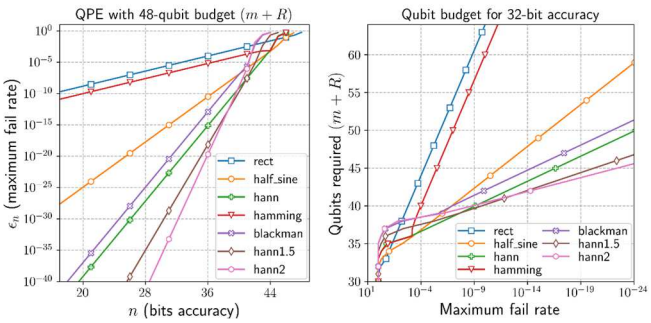

(a) (b)

Fig. 4 Numerical simulation of QPE: (a) performance characteristics at 48-qubit budget; (b) qubit budget for 32-bit accuracy at various fail rate bounds

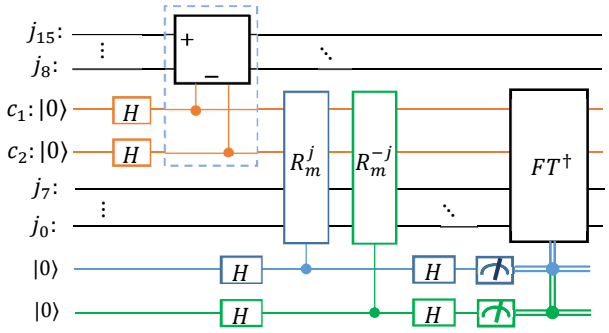

Fig. 5 Quantum STFT with Hann window and 75% frame overlap

The main results are shown in Fig. 4. Fig. 4(a) compares QPE performance profile (n, ϵ) at 48-qubit budget (including target and auxiliary qubits) across window types; Fig. 4 (b) compares resource-reliability curves, i.e. number of qubits needed to achieve specified fail rate, fixing 32-bit accuracy. The two plots closely mirror each other, which is the result of fixed 1:1 "exchange rate" between target qubits and bit precision of estimation. The curves consist of a head part where different windows show variable behaviours, and a tail part where all curves are asymptotically log-linear. The asymptotic slopes can be analytically derived from Fourier analysis theory, which we list in the last column of Table II. We also list the upper bounds of fail rates at $(m-1)$ and $(m-5)$ bits accuracy. Empirically we observe that the log-linear rule gives good approximation of ϵ_n for $n \leq m-5$, which can be useful for quick estimation of the performance profile, e.g. during dynamic planning of algorithms and resources.

An immediate observation from the results is that windowing changes QPE's performance characteristics. In terms of resource-reliability tradeoff, the asymptotic marginal cost for achieving 2x reliability (i.e. half fail rate), as visualized by the slopes of curves in Fig. 4 (b), is reduced from one qubit to 1/3, 1/5, 1/7 or 1/9 qubit, depending on the window type. Windowing does not affect the asymptotic cost of additional accuracy, which remains one qubit per bit accuracy. Fig. 4 also shows when windowing helps reduce overall qubit budget of QPE. Roughly speaking, rectangular (no windowing) has the lowest qubit budget for fail rates above 0.1; the half sine window has it for fail rates between 0.1 and 0.0003; then the Hann window takes over for fail rates down to 2×10^{-8}. Practical window selection will also consider factors like complexity and efficiency, but the big picture is clear: use a window if highly reliable estimation is desired.

V. Conclusion and Discussions

In this paper we have identified a link between the quantum phase estimation algorithm and classical frequency estimation from DFT, which immediately suggests using spectral shaping in QPE. This is accomplished by means of a quantum circuit that implements windowing in the quantum state space. By theoretical calculations we have confirmed that the method allows control of QPE's performance profile in terms of accuracy-reliability-resource tradeoffs, and can help reduce qubit budget in high-reliability use cases. We have also included a brief review of quantum computing basics in subspace terminology to help set up context for our discussion. For a more comprehensive introduction to the field in full quantum information vocabulary we suggest [1] for a good start.

Quantum computing and signal processing have largely gone separately ways in the past decades. Only recently did their paths cross, often in the form of using quantum circuits as alternative signal transformation and optimization resources [14][15][16], or that of applying signal processing methods in classical form to address various needs of quantum computing tasks [17][18][19], including alternative methods for the QPE task [20]. What we present in this paper is a demonstration of the interdisciplinarity working another way, in which signal processing theory directly motivates a new quantum circuit design. We expect more signal-inspired quantum algorithms to come up as quantum computing gains wider attention from signal processing communities.

Since windowing is a classic signal analysis tool by itself, an interesting question is whether the windowing circuit presented here will be useful for quantum signal processing tasks, such as quantum time-frequency analysis. For now the short answer is that we have yet to confirm a use case in which doing so brings concrete end-to-end advantage compared to classical computing. Were such a use case to arise someday, doing the windowing can be straightforward. For example, the circuit in Fig. 5 computes the short-time Fourier transform (STFT) of a 2^{16}-long input encoded in qubits $\boldsymbol{j}_{15} \ldots \boldsymbol{j}_0$, using 1024-point Hann windows with 75% overlap. The submodule in the dashed frame is a permutation circuit that simulates classical invertible subtraction $|j, c\rangle \mapsto |j - c, c\rangle$, where the deduction is modulo 2^8. There will be time alias between frames at the front and back because of the modular arithmetic, which can be eliminated by leaving the last 75% frame empty.

REFERENCES

[1] M. Nielsen and I. Chuang, Quantum Computation and Quantum Information. Cambridge: Cambridge University Press. 2000.

[2] A. Kitaev, "Quantum measurements and the Abelian stabilizer problem," Electronic Colloquium on Computational Complexity (ECCC), TR96-003, 1995.

[3] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, "Quantum algorithms revisited," Proc. R. Soc. London A, 454, 1998, pp 339–354.

[4] P.W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," SIAM Journal on Computing, 26 (5), 1997, pp 1484-1509.

[5] A.W. Harrow, A. Hassidim, and S. Lloyd, "Quantum algorithm for linear systems of equations," Physical Review Letters. 103 (15): 150502, 2009.

[6] N. Wiebe, D. Braun, and S. Lloyd, "Quantum algorithm for data fitting," Physical Review Letters, 109(5), 2012, p.050505.

[7] M. Schuld, I. Sinayskiy, and F. Petruccione, "Prediction by linear regression on a quantum computer," Physical Review A, 94(2), 2016, p.022342.

[8] S. Lloyd, M. Mohseni, and P. Rebentrost, "Quantum principal component analysis," Nature Physics, 10(9), 2014, pp.631-633.

[9] I. Kerenidis, and A. Prakash, "Quantum Recommendation Systems," 8th Innovations in Theoretical Computer Science Conference (ITCS), 2017.

[10] G. Brassard, P. Hoyer, and A. Tapp, "Quantum counting," 25th International Colloquium on Automata, Languages and Programming (ICALP'98), Aalborg, 1998.

[11] S. Mallat, A Wavelet Tour of Signal Processing (Chap.2). New York: Elsevier. 1999.

[12] H. Yin, D. Lu, and R. Zhang, "Quantum windowed Fourier transform and its application to quantum signal processing," International Journal of Theoretical Physics, 60, 2021, pp.3896-3918.

[13] P. Kaye, and M. Mosca, "Quantum networks for generating arbitrary quantum states," International Conference on Quantum Information, Optica Publishing Group, 2001, p. PB28.

[14] C. Yang, B. Li, Y. Zhang, et al., "A quantum kernel learning approach to acoustic modeling for spoken command recognition," ICASSP, 2023.

[15] A. Sharma, G. Uehara, V. Narayanaswamy, L. Miller, and A. Spanias, "Signal analysis-synthesis using the quantum Fourier transform," ICASSP, 2023.

[16] S. Dutta, A. Basarab, D. Kouamé, and B. Georgeot, "Quantum algorithm for signal denoising," IEEE Signal Processing Letters, December 2023.

[17] W. Shi, and R. Malaney, "Signal processing and quantum state tomography on noisy devices," ICASSP, 2023.

[18] J. Pérez-Guijarro, A. Pagès-Zamora, and J. R. Fonollosa, "Extension of Clifford data regression methods for quantum error mitigation," ICASSP, 2024.

[19] Y. Deville, and A. Deville, "Multi-stage tomography based on eigenanalysis for high-dimensional dense unitary quantum processes," arXiv 2407.13542.

[20] J.G. Smith, C.H. Barnes, and D.R. Arvidsson-Shukur, "Iterative quantum-phase-estimation protocol for shallow circuits," Physical Review A, 106(6), 062615, 2022.