Lecture 7: The Quantum Fourier Transform and Its Applications

Phase Estimation, Order-Finding, and Shor's Factoring Algorithm

Moody T. Chu

North Carolina State University

Spring 2026

Lecture Overview

1. Review of Classical Discrete Fourier Transform (DFT)
2. Mathematical Definition and Circuit Architecture of the QFT
3. Quantum Phase Estimation (QPE) and Error Performance Bounds
4. The Quantum Order-Finding Problem and Continued Fractions
5. Shor's Factoring Algorithm and Its Reduction to Order-Finding

Goal: By the end of this lecture we will have built — entirely from first principles, and verified on a concrete numerical example — the quantum Fourier transformm that lies at the heart of quantum computation. Why that matters in practice, and how the subroutine is actually deployed against a live cryptosystem, is the story of Lecture 8.

Introduction to Phase & Periodicity

- The most spectacular discovery in quantum computing is that certain tasks, completely intractable classically, can be solved efficiently on a quantum architecture.
- **Primary Example:** Factoring an n-bit integer.
  - ◇ Classical (General Number Field Sieve): $\exp\left(\Theta\left(n^{1/3}\log^{2/3} n\right)\right)$.
    - ○ Θ stands for tight bound for both upper and lower operational limits.
  - ◇ Quantum (Shor's Algorithm): $O(n^2 \log n \log\log n)$.
- The *Quantum Fourier Transform (QFT)* is the mathematical engine behind factoring, phase estimation, and the general hidden subgroup problem.
- **Roadmap for this lecture:** build the QFT circuit (Sections 1–2), use it for phase estimation (Section 3), then apply phase estimation to order-finding and factoring (Sections 4–5).

Review of Classical DFT

- Given samples $\{f_0, f_1, \ldots, f_{N-1}\}$ of a signal at fixed intervals in the time domain,
  - The *Discrete Fourier Transform (DFT)* maps this signal into the frequency domain via

$$\mathfrak{F}_k := \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} e^{\imath \frac{2\pi}{N} jk} f_j, \quad k = 0, \ldots, N-1. \tag{1}$$

  - *Note: the conventional definition uses $e^{-\imath(\cdot)}$; what's shown above is technically the inverse transform. The sign is flipped here so that this formula aligns directly with the standard unitary QFT definition on the next slides.*
- This linear transformation can be expressed compactly as

$$\mathfrak{F} = W\mathbf{f}.$$

where W is an $N \times N$ unitary matrix.
  - Exploiting the cyclic structure of W leads to the Fast Fourier Transform (FFT), one of the most important and widely used algorithms in all of computing.

Algorithmic Complexity: FFT vs. QFT

- Fast Fourier Transform (FFT): Exploits the structural factorization of the unitary matrix W.
  - ◇ Algorithmic Cost: $O(N \log_2 N)$ arithmetic operations.
  - ◇ For a problem domain of size $N = 2^n$, this scales as $O(n2^n)$.
- Quantum Fourier Transform (QFT):
  - ◇ Operates directly on the computational basis amplitudes of a quantum state.
  - ◇ Algorithmic Cost: $O(n^2)$ elementary quantum gates — exponentially fewer gates than the FFT needs arithmetic operations, since $n = \log_2 N$.
- Crucial Caveat: Because the amplitudes are hidden within the wave function, we cannot read out the full DFT vector directly. QFT must act as a subroutine *inside* a larger algorithm.
  - ◇ QFT cannot simply be used as a faster drop-in replacement for FFT — the speedup is only accessible when the downstream algorithm needs information that survives measurement (such as a periodicity), not the raw output vector itself. We will see exactly how in Section 3.

QFT: Mathematical Definition

- Let $N = 2^n$, and let $|0\rangle, \ldots, |N-1\rangle$ be the orthonormal computational basis for an n-qubit Hilbert space.
- The QFT is defined by its action on a basis state $|j\rangle$:

$$\mathrm{QFT}\,|j\rangle = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} e^{\imath \frac{2\pi}{N} jk} \,|k\rangle\,. \tag{2}$$

- Given an arbitrary superposition input $\sum_{j=0}^{N-1} x_j \,|j\rangle$, linearity extends this to:

$$\mathrm{QFT}\left(\sum_{j=0}^{N-1} x_j \,|j\rangle\right) = \sum_{k=0}^{N-1} \underbrace{\left(\frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} e^{\imath \frac{2\pi}{N} jk} x_j\right)}_{\mathfrak{Q}_k} |k\rangle\,. \tag{3}$$

  - ◇ The coefficient $\mathfrak{Q}_k$ is exactly the classical DFT formula from the previous slide, applied to the amplitudes x_j — the QFT is, by design, in complete sync with the DFT.

Goal: An Efficient Circuit for the QFT

- The definition above is a $N \times N$ unitary matrix acting on n qubits — on its own, this gives no hint of an efficient circuit. The task for the rest of this section:
  - ◇ Show that an n-qubit QFT can be implemented out of 1-qubit and 2-qubit unitary gates.
  - ◇ Show that the number of gates required grows only *quadratically* in n, not exponentially.
- The key tool is **binary fraction notation**, which lets us rewrite the QFT's output as a clean product of single-qubit states.
  - ◇ If $j = j_1 2^{n-1} + j_2 2^{n-2} + \cdots + j_n 2^0$, write

$$j = j_1 j_2 \ldots j_n.$$

  - ◇ If $j = \frac{j_\ell}{2^1} + \frac{j_{\ell+1}}{2^2} + \cdots + \frac{j_m}{2^{m-\ell+1}}$, write

$$j = 0.j_\ell j_{\ell+1} \ldots j_m.$$

QFT: Product Representation Theorem

- **Theorem:** The output of the QFT can always be factored into an unentangled product state:

$$|j_1 \dots j_n\rangle \to \frac{(|0\rangle + e^{2\pi\imath 0.j_n}|1\rangle) \otimes \cdots \otimes (|0\rangle + e^{2\pi\imath 0.j_1 j_2 \dots j_n}|1\rangle)}{2^{n/2}} \tag{4}$$

- **A note on ordering:** the tensor factors above are written with the *shortest* binary fraction first ($0.j_n$) and the *longest* (full $0.j_1 \cdots j_n$) last. When we build the physical circuit a few slides from now, the wires will naturally produce these factors in the *opposite* order — wire 1 ends up holding the full fraction. This reversal is expected, and is exactly why the circuit ends with a layer of SWAP gates.

Derivation of the Product Form (Part I)

- If $k = \sum_{\ell=1}^{n} k_\ell 2^{n-\ell}$, then

$$\begin{cases} |k\rangle &= |k_1\rangle \otimes |k_2\rangle \otimes \cdots \otimes |k_n\rangle = \bigotimes_{\ell=1}^{n} |k_\ell\rangle\,, \\ \dfrac{k}{2^n} &= \sum_{\ell=1}^{n} k_\ell\, 2^{-\ell}. \end{cases}$$

- These two binary expansion rules let us rewrite the QFT definition, expanded fully on the next slide.

Derivation of the Product Form (Part I, cont.)

- Expand the definition of the transform using these binary expansion rules:

$$\begin{aligned}
\mathrm{QFT}\,|j\rangle &= \frac{1}{2^{n/2}} \sum_{k=0}^{2^n-1} e^{\imath 2\pi j \frac{k}{2^n}} \,|k\rangle \\
&= \frac{1}{2^{n/2}} \sum_{k_1=0}^{1} \cdots \sum_{k_n=0}^{1} e^{2\pi \imath j \left(\sum_{\ell=1}^{n} k_\ell 2^{-\ell}\right)} \,|k_1 \ldots k_n\rangle \\
&= \frac{1}{2^{n/2}} \sum_{k_1=0}^{1} \cdots \sum_{k_n=0}^{1} \bigotimes_{\ell=1}^{n} e^{2\pi \imath j k_\ell 2^{-\ell}} \,|k_\ell\rangle
\end{aligned}$$

  - The phase factors have been distributed to qubit states; the multi-index sum can now be pulled apart, qubit by qubit, on the next slide.

Derivation of the Product Form (Part II)

- Mathematically, a sum of tensor products of single-variable functions factors into a product of sums:

$$\sum_{k_1,\ldots,k_n} \bigotimes_{\ell=1}^{n} f(k_\ell) = \bigotimes_{\ell=1}^{n} \left(\sum_{k_\ell=0}^{1} f(k_\ell) \right).$$

- Applying this identity to the expression from Part I:

$$\begin{aligned} \mathrm{QFT}\,|j\rangle &= \frac{1}{2^{n/2}} \bigotimes_{\ell=1}^{n} \left[\sum_{k_\ell=0}^{1} e^{2\pi \imath j k_\ell 2^{-\ell}} |k_\ell\rangle \right] \\ &= \frac{1}{2^{n/2}} \bigotimes_{\ell=1}^{n} \left[|0\rangle + e^{2\pi \imath j 2^{-\ell}} |1\rangle \right]. \end{aligned}$$

- Each tensor factor is now a single qubit in a clean superposition — exactly the unentangled product form the theorem promised, though the phase $j2^{-\ell}$ still needs to be simplified.

Derivation of the Product Form (Part III)

- If $j = \sum_{s=1}^{n} j_s 2^{n-s}$, then

$$j2^{-\ell} = j_1 \dots j_{n-\ell} . j_{n-\ell+1} \dots j_n .$$

  - ◇ The integer portion $j_1 \dots j_{n-\ell}$ contributes an integer multiple of 2π to the phase $e^{2\pi\imath(\dots)}$, causing it to evaluate to 1.
  - ◇ Hence, $e^{2\pi\imath j2^{-\ell}} = e^{2\pi\imath 0.j_{n-\ell+1}\dots j_n}$ — only the fractional part of $j2^{-\ell}$ survives.

- Substituting back into Part II's result, the phase on the ℓ-th tensor factor simplifies completely:

$$\text{QFT}\,|j\rangle \quad = \quad \frac{1}{2^{n/2}} \bigotimes_{\ell=1}^{n} \left[|0\rangle + e^{2\pi\imath 0.j_{n-\ell+1}\dots j_n} |1\rangle \right].$$

- Checking the two endpoints confirms this matches the Product Representation Theorem: $\ell = 1$ gives $0.j_n$ (shortest), and $\ell = n$ gives $0.j_1 \cdots j_n$ (full fraction, longest). ■

The Controlled Phase Rotation Gate

- One fundamental component of the circuit is the single-qubit *conditional phase gate* $R_k = \begin{bmatrix} 1 & 0 \\ 0 & e^{2\pi\imath/2^k} \end{bmatrix}$, giving, in controlled configuration on target $|b\rangle$ from control $|a\rangle$:

$$\mathrm{CR}_k |a, b\rangle = e^{2\pi\imath \frac{a \cdot b}{2^k}} |a, b\rangle \tag{5}$$

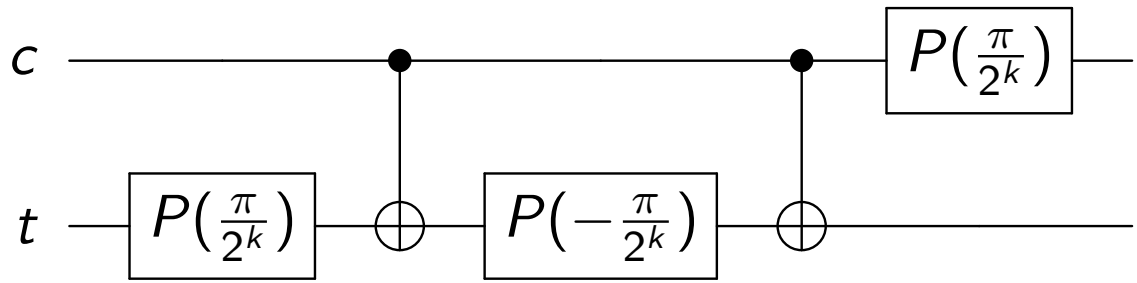

  - The phase-shift gate is $P(\phi) = \begin{bmatrix} 1 & 0 \\ 0 & e^{\imath\phi} \end{bmatrix}$. CR_k is symmetric in control/target, since $e^{2\pi\imath ab/2^k}$ is itself symmetric in a, b.

QFT : The 4-Qubit Circuit Architecture

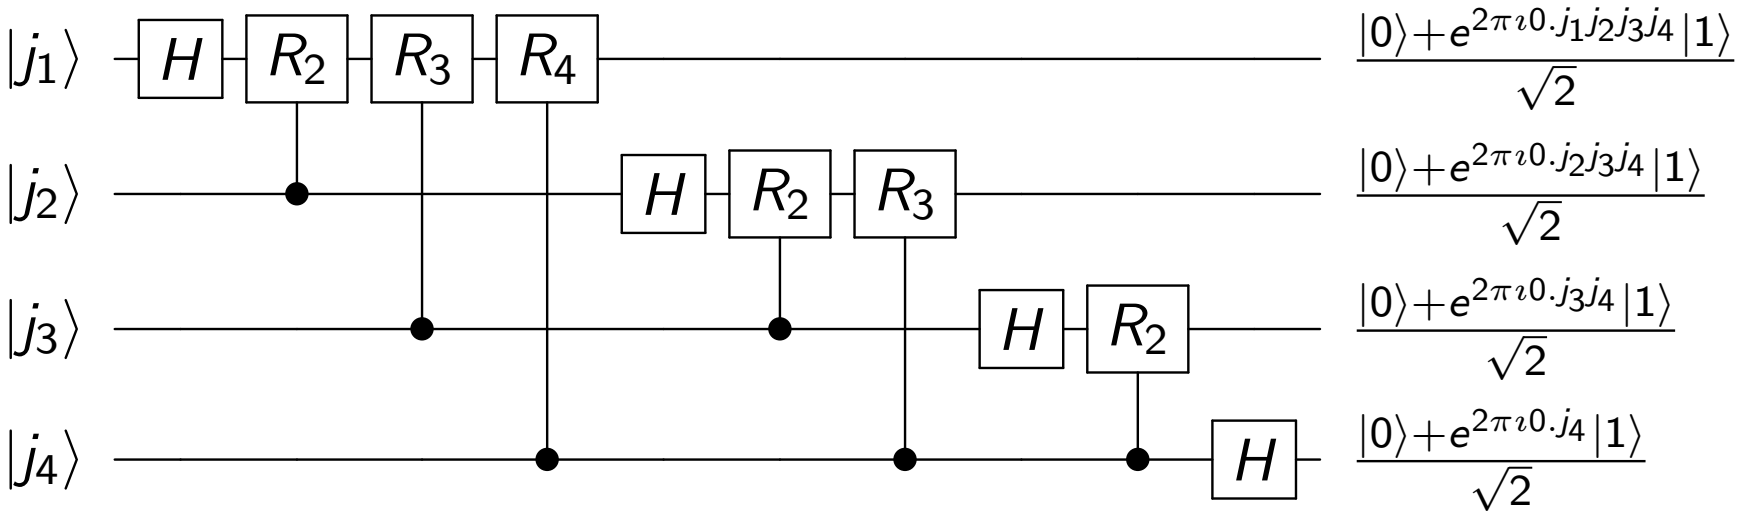

- Two key components, both verified on the previous slides:

$$H|j_s\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + (-1)^{j_s}|1\rangle\right) = \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath(0.j_s)}|1\rangle\right)$$

$$CR_k|j_c\rangle \otimes |1\rangle = e^{2\pi\imath\frac{j_c}{2^k}}|j_c\rangle \otimes |1\rangle = e^{2\pi\imath(0.\underbrace{0\ldots0}_{k-1}j_c)}|j_c\rangle \otimes |1\rangle$$

QFT Circuit: A Note on Wire Order

- Wire 1 (processed first) accumulates the *full* fraction $0.j_1j_2j_3j_4$, while wire 4 (processed last) ends with just $0.j_4$ — the reverse of the Product Representation Theorem's tensor order, exactly as flagged a few slides ago.
- This reversal is not a mistake in either statement: the theorem describes the state's abstract tensor structure, while the circuit describes the physical order qubits are processed in. They are reconciled by the SWAP layer at the very end of the circuit, which restores the standard $|j_1 \cdots j_n\rangle$-aligned wire order.

Step-by-Step State Evolution (Wires 1–3)

- **Step 1: Evolution of Wire 1**

$$|j_1\rangle \xrightarrow{H} \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath 0.j_1}|1\rangle\right) \xrightarrow{CR_2} \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath 0.j_1 j_2}|1\rangle\right)$$
$$\xrightarrow{CR_3} \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath 0.j_1 j_2 j_3}|1\rangle\right) \xrightarrow{CR_4} \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath 0.j_1 j_2 j_3 j_4}|1\rangle\right) = |\phi_1\rangle$$

- **Step 2: Evolution of Wire 2**

$$|j_2\rangle \xrightarrow{H} \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath 0.j_2}|1\rangle\right) \xrightarrow{CR_2} \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath 0.j_2 j_3}|1\rangle\right)$$
$$\xrightarrow{CR_3} \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath 0.j_2 j_3 j_4}|1\rangle\right) = |\phi_2\rangle$$

- **Step 3: Evolution of Wire 3**

$$|j_3\rangle \xrightarrow{H} \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath 0.j_3}|1\rangle\right) \xrightarrow{CR_2} \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath 0.j_3 j_4}|1\rangle\right) = |\phi_3\rangle$$

- Wire 4 needs only a single Hadamard: $|j_4\rangle \xrightarrow{H} \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi\imath 0.j_4}|1\rangle\right) = |\phi_4\rangle$.

QFT: The General Circuit Architecture

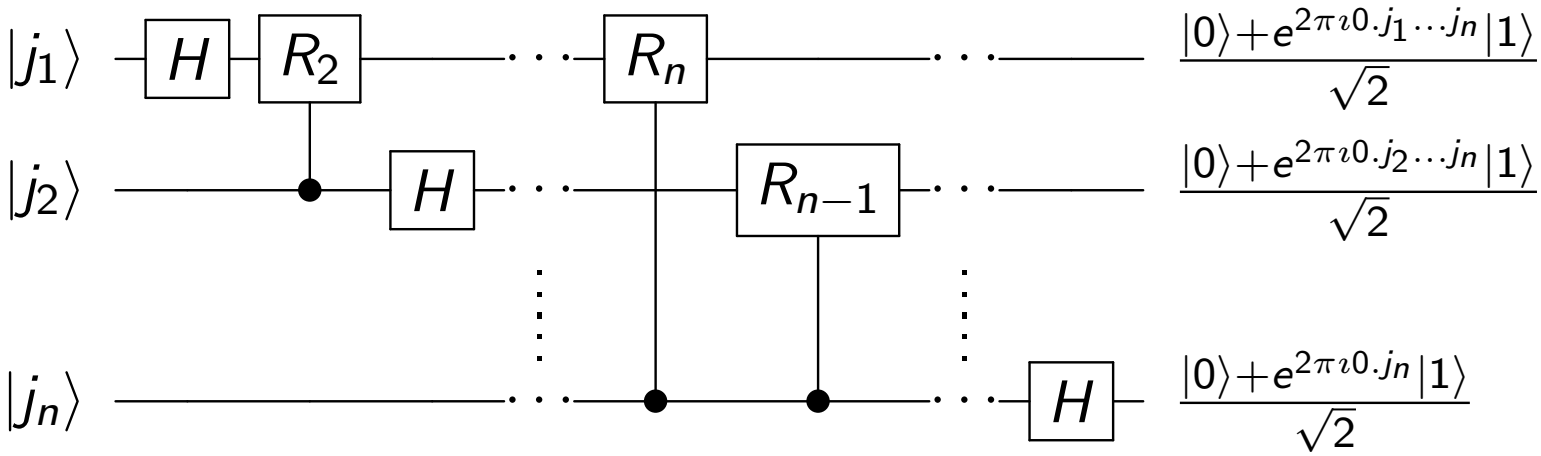

- Wire ℓ requires one Hadamard and $(n-\ell)$ controlled rotations $R_2, \ldots, R_{n-\ell+1}$ — exactly the pattern verified for $n = 4$ on the previous two slides, now generalized to arbitrary n.

QFT†: The 4-Qubit Circuit Architecture

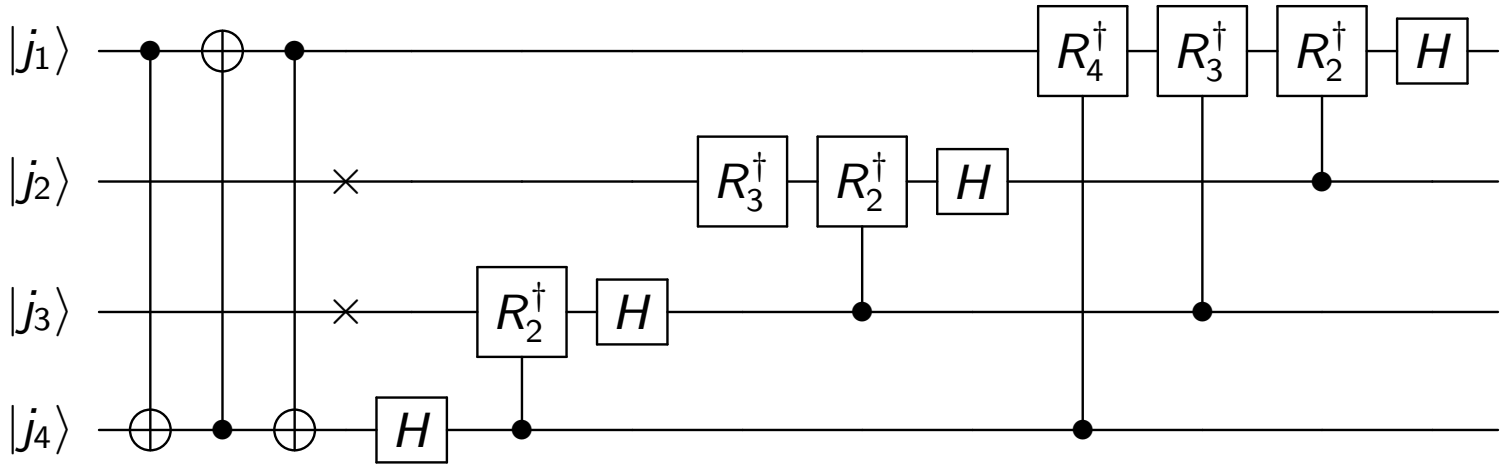

- Pre-emptive SWAP Layer: since the QFT ends with wire rearrangements, QFT† must immediately start by crossing them back (non-adjacent swaps via $CX_{1\to4} \cdot CX_{4\to1} \cdot CX_{1\to4}$; adjacent swaps directly).
- Bottom-to-Top Phase Unwinding: the rest of the pipeline flows upstream, undoing the forward QFT's gate sequence in reverse order, with every rotation conjugated: $R_k^\dagger = \begin{pmatrix} 1 & 0 \\ 0 & e^{-2\pi\imath/2^k} \end{pmatrix}$.

Quantitative Gate Complexity Analysis

- Let's count the exact number of operations required:
  - ◇ Qubit 1 requires 1 Hadamard and $n-1$ conditional rotations.
  - ◇ Qubit 2 requires 1 Hadamard and $n-2$ conditional rotations.
  - ◇ Qubit m requires 1 Hadamard and $n-m$ conditional rotations.
- Summing over all stages yields the total number of primary gates:

$$\sum_{m=1}^{n}(1+n-m)=\frac{n(n+1)}{2} \tag{6}$$

- To restore the standard $|j_1 \ldots j_n\rangle$-aligned wire order (undoing the reversal noted earlier), we append $\lfloor n/2 \rfloor$ SWAP gates. Total gate complexity remains $O(n^2)$.
  - ◇ SWAP $|ab\rangle = |ba\rangle$.
  - ◇ Without the SWAP layer, the gate cascade computes the Fourier coefficients with the qubits in bit-reversed order.

Quantum Phase Estimation (QPE)

- The Core Problem: Let U be a unitary operator. Suppose an eigenvalue is expressed as $e^{2\pi\imath\phi}$ for an unknown phase $\phi \in [0, 1)$.
  - ◇ Given an eigenstate $|u\rangle$ such that:

$$U|u\rangle = e^{2\pi\imath\phi}|u\rangle \tag{7}$$

  - ◇ Estimate the value of ϕ to t bits of precision.
- The algorithm requires two registers and the inverse QFT — this is precisely why we just spent an entire section building an efficient QFT circuit.

Why Is QPE Necessary? (I)

- Global phase invisibility:
  - ◇ By Born's Rule, the probability of observing any computational basis state $|j\rangle$ upon measurement of $U|u\rangle$ is:

$$P(|j\rangle) = |\langle j| U |u\rangle|^2 = |e^{2\pi\imath\phi} \langle j|u\rangle|^2 = |\langle j|u\rangle|^2$$

  - ◇ Because $|e^{2\pi\imath\phi}| = 1$, the phase ϕ completely vanishes from the measurement statistics if we just measure $U|u\rangle$ directly.
- Inaccessibility of amplitudes:
  - ◇ Information about a quantum state $|u\rangle = \sum_{j=0}^{2^m-1} c_j |j\rangle$ is encoded within its complex amplitudes.
  - ◇ Unlike classical data, individual amplitudes c_j cannot be read out directly.
  - ◇ Measuring the system causes the original state to collapse, losing information about all other amplitudes.

Why Is QPE Necessary? (II)

- Implicit operator representation:
  - ◇ In quantum algorithms, the operator U is specified as a circuit of local gates, not as explicit matrix entries.
  - ◇ We can physically evolve a state through it, but we cannot evaluate its matrix elements to compute direct entrywise ratios.
- **The upshot:** we need an indirect strategy that converts the invisible phase ϕ into something measurement *can* see — a computational basis state. That is exactly what the QPE circuit, built next, accomplishes.

QPE: Circuit Structural Layout

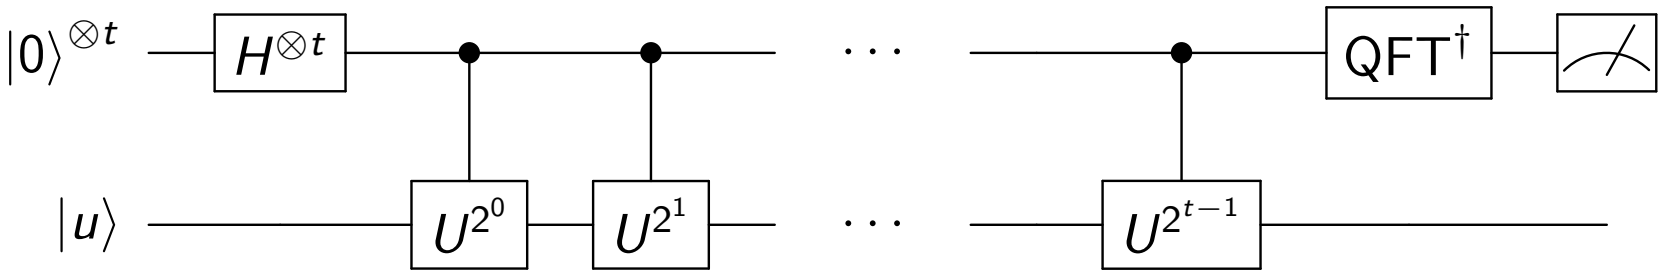

- Precision Register (t qubits): Determines the bits of accuracy for the phase approximation.
- Target Register ($|u\rangle$): Holds the eigenstate of the $2^m \times 2^m$ unitary operator U.
- The system exploits successive doubling of phases via controlled-U^{2^j} operations.

QPE: Multi-Wire Circuit Layout

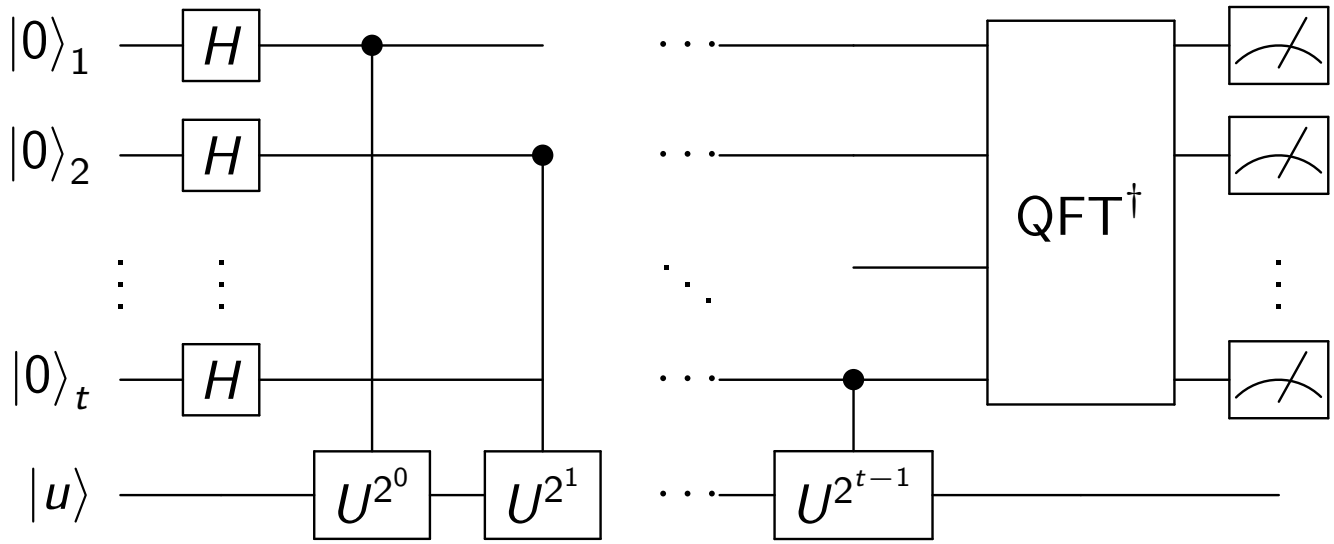

- Every single U^{2^j} operation is conditioned on its corresponding control qubit from the top register.
- *Parallel Phase Kickback:* If a specific control qubit is in the state $|1\rangle$, the unitary operator acts on the target register and kicks back its corresponding phase factor $e^{2\pi\imath\phi 2^j}$ onto that control qubit.

QPE: Mathematical Analysis, I

- Initialization:
  - Initial value: $|\Psi_0\rangle = |0\rangle^{\otimes t} |u\rangle$
  - Parallel Hadamards:

$$|\Psi_1\rangle = \frac{1}{2^{t/2}} \sum_{k=0}^{2^t-1} |k\rangle |u\rangle .$$

- Generate the controls:
  - Write $|k\rangle = |k_{t-1} \dots k_0\rangle$, where

$$k = k_{t-1}2^{t-1} + \cdots + k_1 2^1 + k_0 2^0.$$

  - Use the j-th qubit (k_j) of the estimation register as the control for the unitary operator U^{2^j}.
  - Simple math, but effective control:

$$U^k |u\rangle = U^{\sum_{j=0}^{t-1} k_j 2^j} |u\rangle = \left(U^{2^0}\right)^{k_0} \left(U^{2^1}\right)^{k_1} \dots \left(U^{2^{t-1}}\right)^{k_{t-1}} |u\rangle .$$

QPE: Mathematical Analysis, II

- Since $|u\rangle$ is an eigenstate,

$$U^k |u\rangle = e^{2\pi\imath\phi k} |u\rangle .$$

- Post U^{2^j} sequence (before QFT†):
  - ◇ When this circuit is applied to $|k\rangle$ in the summation, it produces $U^k |u\rangle$.
  - ◇ The phase $e^{2\pi\imath\phi k}$ is kicked back to the front of the state as a scalar.
  - ◇ Running through the entire summation:

$$|\Psi_2\rangle = \frac{1}{2^{t/2}} \sum_{k=0}^{2^t-1} e^{2\pi\imath\phi k} |k\rangle |u\rangle \tag{8}$$

QPE: Mathematical Analysis, III, (Ideal Case)

- Assume ϕ can be represented exactly using t bits,

$$\phi = 0.\phi_1\phi_2\ldots\phi_t.$$

  - Let $b = 2^t\phi$ (an integer!)
- With $N := 2^t$, rewrite $|\Psi_2\rangle$:

$$|\Psi_2\rangle = \frac{1}{2^{t/2}} \sum_{k=0}^{2^t-1} e^{\imath \frac{2\pi}{N} bk} |k\rangle |u\rangle . \tag{9}$$

  - The summation is precisely QFT($|b\rangle$), by the very definition from Section 2.
- Applying the inverse transformation ($\text{QFT}^\dagger$) yields:

$$|\Psi_3\rangle = (\text{QFT}^\dagger \otimes I) |\Psi_2\rangle = |b\rangle |u\rangle = |\phi_1\phi_2\ldots\phi_t\rangle |u\rangle . \tag{10}$$

- Measurement yields the exact phase with a probability of 1. The non-ideal case (real life) is analyzed next.

QPE: Mathematical Analysis, IV, (Non-Exact Phase)

- What does the circuit produce when ϕ does *not* have an exact t-bit representation?

$$
\begin{aligned}
\mathrm{QFT}^\dagger |\Psi_2\rangle &= \mathrm{QFT}^\dagger \left(\frac{1}{2^{t/2}} \sum_{k=0}^{2^t-1} e^{2\pi\imath\phi k} |k\rangle \right) \\
&= \frac{1}{2^{t/2}} \sum_{k=0}^{2^t-1} e^{2\pi\imath\phi k} \left(\mathrm{QFT}^\dagger |k\rangle \right) \\
&= \frac{1}{2^{t/2}} \sum_{k=0}^{2^t-1} e^{2\pi\imath\phi k} \left(\frac{1}{2^{t/2}} \sum_{m=0}^{2^t-1} e^{-2\pi\imath\frac{mk}{2^t}} |m\rangle \right) \\
&= \sum_{m=0}^{2^t-1} \left[\frac{1}{2^t} \sum_{k=0}^{2^t-1} e^{2\pi\imath\left(\phi-\frac{m}{2^t}\right)k} \right] |m\rangle
\end{aligned}
$$

QPE: Error Performance Bound

- If ϕ is not an exact t-bit fraction, let b be the nearest integer to $2^t\phi$, so the rounding/truncation error satisfies

$$0 \leq \delta \equiv \phi - \frac{b}{2^t} \leq \frac{1}{2^{t+1}}.$$

- Applying $\mathrm{QFT}^\dagger$, the amplitude for measuring the closest state $|b\rangle$ is:

$$c_b = \frac{1}{2^t}\sum_{k=0}^{2^t-1} e^{2\pi\imath\delta k} = \begin{cases} 1, & \text{if } \delta = 0, \\ \frac{1}{2^t}\frac{1-e^{2\pi\imath\delta 2^t}}{1-e^{2\pi\imath\delta}}, & \text{otherwise (geometric series).} \end{cases} \tag{11}$$

  - The probability of obtaining $|b\rangle$ is:

$$P(b) = |c_b|^2 = \frac{1}{2^{2t}}\frac{\sin^2(\pi\delta 2^t)}{\sin^2(\pi\delta)} \geq \frac{\sin^2(\theta)}{\theta^2} \geq \frac{4}{\pi^2} \approx 40.53\%$$

- This bound is tight: as $t \to \infty$ with δ at its worst-case value $1/2^{t+1}$, $P(b) \to 4/\pi^2$ exactly (verified numerically).

QPE: Interpreting Phase Estimation Output (I)

- After measuring the t-qubit estimation register to obtain an integer $m \in \{0, 1, \ldots, 2^t - 1\}$,
  - Single-shot estimation: Accept $\hat{\phi} = \frac{m}{2^t}$ directly. There is a $\geq 40.53\%$ baseline confidence that m is the absolute closest integer b to $2^t\phi$, with the error tightly bounded by $|\phi - \hat{\phi}| \leq \frac{1}{2^{t+1}}$.
  - Ensemble sampling (Spectral peak): Execute the circuit multiple times to construct an empirical distribution. Due to spectral leakage, the resulting histogram forms a sharp Dirichlet kernel; the mode (highest peak) identifies b.

QPE: Interpreting Phase Estimation Output (II)

- Algorithmic scaling (Width buffering): To obtain n bits of accuracy with a success probability bound of at least $1 - \epsilon$, the allocation of t qubits must satisfy:

$$t = n + \left\lceil \log\left(2 + \frac{1}{2\epsilon}\right) \right\rceil \tag{12}$$

Why this formula: the $\geq 4/\pi^2$ bound only guarantees a constant ($\approx 40\%$) success rate. Adding $\log(2 + 1/2\epsilon)$ extra qubits shrinks the probability of landing more than 1 bin away from the peak (a tail-sum over the $O(1/k)$-decaying neighboring amplitudes, analyzed when we revisit this calculation for order-finding) down to ϵ, boosting confidence to $1 - \epsilon$ at the cost of only a few extra qubits.

Preparing the Eigenstate: The Linear Algebra Paradox

- The foundational paradox: Finding an eigenvector of a $2^n \times 2^n$ unitary operator classically requires complexity $O(2^{3n})$.
  - ◇ Preparing $|u\rangle$ by classical means for QPE instantly destroys the exponential quantum speedup.
- The quantum bypass: Quantum systems do not compute vector coefficients in the classical sense.
  - ◇ Instead, operate on quantum registers, constructing desired states directly on hardware through three refined approaches:
    - ○ Algebraic superposition (Shor's method): Begin with a uniform superposition over all hidden eigenstates. This sum collapses to a simple computational basis state (e.g. $|1\rangle$), which is easy to prepare. *(This is exactly the strategy used in order-finding, Section 4.)*
    - ○ Quantum state filtering: Feed an imperfect "guess" state $|\chi\rangle$ into QPE. The phase measurement projects the system onto a pure eigenstate, effectively filtering out the desired component.
    - ○ Spectral Tracking (Adiabatic & VQE): Either adiabatically evolve a simple ground state through gradual parameter changes, or employ hybrid classical-quantum optimization to variationally adjust the state until it aligns with the target operator.

QPE: Applications

- Given a unitary operator U and an explicit eigenstate $|u\rangle$, QPE extracts the continuous phase ϕ from the eigenvalue $\lambda = e^{2\pi i\phi}$.
- Application: Many of the hardest problems in classical computer science (e.g., factoring, discrete logarithms, cryptography) are fundamentally concerned with uncovering a hidden, discrete period r, which can be addressed via QPE through the steps:
  1. Embed the classical periodic function into a quantum operator U.
  2. Prove that this operator is inherently unitary so it can execute on hardware.
  3. Construct structural eigenstates $|u_s\rangle$ whose eigenvalues $\lambda_s = e^{2\pi\imath\frac{s}{r}}$ encode the period r directly into their continuous phases.
- We now carry out exactly this program, for the single most consequential example: factoring.

The Order-Finding Problem

- **Definition:** Given two positive integers x and N,
  - $x < N$ and $\gcd(x, N) = 1$,
  - the *order* $ord_N(x)$ is defined as the smallest positive integer r such that:

$$x^r \equiv 1 \pmod{N} \tag{13}$$

- An example ($N = 15, x = 7$):
  - Orbits: $7^1 \equiv 7, \quad 7^2 \equiv 4, \quad 7^3 \equiv 13, \quad 7^4 \equiv 1 \pmod{15}$.
  - The sequence loops cyclically. $ord_{15}(7) = 4$.
- Computational Complexity:
  - Finding this order r is the absolute core computational bottleneck of classical integer factorization (we will see exactly why, in Section 5).
  - No efficient classical algorithm is known to find r for large numbers.

Step 1: Modular Multiplication Operator

- Given the target composite integer N, prepare a large enough state space:
  - Required bit width is $L := \lceil \log_2 N \rceil$.
  - An L-qubit register (Register 2) can store state vectors in the range $y \in [0, N-1]$.
- **Definition:** The modular multiplication operator U acting on this register:

$$U|y\rangle := |xy \pmod N\rangle, \quad 0 \leq y < N \tag{14}$$

- For this to be useful, U must be unitary — verified over the next three slides.

Step 2: U is unitary

- Claim: if $\gcd(x, N) = 1$, then the linear congruence mapping

$$y \mapsto xy \pmod{N}$$

  is a bijection over the set of integers coprime to N. Hence, it is a permutation.
  - ◇ The proof of bijectivity is the next two slides.
- A permutation of basis states maps orthonormal vectors to orthonormal vectors.
- Therefore, U represents a clean permutation matrix, guaranteeing that $U^\dagger U = I$. It is strictly unitary.

Justification of Bijectivity: Injectivity

- Group of units modulo N:

$$(\mathbb{Z}/N\mathbb{Z})^\times := \{y \in \mathbb{Z} \mid 0 \leq y < N \text{ and } \gcd(y, N) = 1\}.$$

  - ◇ This is a multiplicative group.
    - $\gcd(x, N) = 1, \gcd(y, N) = 1 \implies \gcd(xy, N) = 1.$
  - ◇ Its order is given by Euler's totient function, $\varphi(N)$.
- Given $\gcd(x, N) = 1$, consider the map $f : (\mathbb{Z}/N\mathbb{Z})^\times \to (\mathbb{Z}/N\mathbb{Z})^\times$ via

$$f(y) = xy \pmod{N}.$$

  - ◇ Taking the remainder modulo N does not alter common factors, i.e., $\gcd(xy \pmod{N}, N) = 1$.
  - ◇ Suppose $f(y_1) = f(y_2)$, i.e., $xy_1 \equiv xy_2 \pmod{N}$.
    - $x(y_1 - y_2) \equiv 0 \pmod{N}$ implies that $N \mid x(y_1 - y_2)$.
    - Since $\gcd(x, N) = 1$, it must be $N \mid (y_1 - y_2)$.
    - Because $y_1, y_2 \in [0, N-1]$, $0 \leq |y_1 - y_2| < N$.
    - The only multiple of N in this range is 0, forcing $y_1 = y_2$.

Justification of Bijectivity: Surjectivity

- Pigeonhole Principle: Any injective mapping from a finite set to itself is automatically surjective.
- Since $\gcd(x, N) = 1$,
  - ◇ Bézout's Identity guarantees the existence of a unique multiplicative inverse $x^{-1} \in (\mathbb{Z}/N\mathbb{Z})^{\times}$.
  - ◇ For any target element $z \in (\mathbb{Z}/N\mathbb{Z})^{\times}$, define $y = x^{-1}z \pmod{N}$.
  - ◇ Then $f(y) \equiv x(x^{-1}z) \equiv z \pmod{N}$, confirming every z has a preimage.
- Together with injectivity, this confirms f is a bijection — completing the proof that U is unitary.

Step 3: The Cyclic Orbit and Eigenvector

- From the cyclic orbit $\{1, x, \ldots, x^{r-1}\}$, observe

$$U \left| x^k \pmod N \right\rangle = \left| x^{k+1} \pmod N \right\rangle, \quad k = 0, 1, \ldots, r-1.$$

  - Boundary condition: $U \left| x^{r-1} \pmod N \right\rangle = |1\rangle$.

- To find the eigenvectors of a cyclic permutation matrix, define for each integer $0 \leq s < r$,

$$|u_s\rangle = \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} e^{-2\pi\imath \frac{sk}{r}} \left| x^k \pmod N \right\rangle \tag{15}$$

  - These are the $\mathrm{QFT}^\dagger$ of $|x^s \pmod N\rangle$.

- Because the orbital states $\left| x^k \pmod N \right\rangle$ are mutually orthogonal, $\langle u_s | u_{s'} \rangle = \delta_{ss'}$.

Step 3: Eigenvalue Verification

- $|u_s\rangle$ is an eigenstate of U:

$$\begin{aligned}
U|u_s\rangle &= \frac{1}{\sqrt{r}}\sum_{k=0}^{r-1} e^{-2\pi\imath\frac{sk}{r}} U\left|x^k \pmod N\right\rangle, \\
&= \frac{1}{\sqrt{r}}\sum_{k=0}^{r-1} e^{-2\pi\imath\frac{sk}{r}} \left|x^{k+1} \pmod N\right\rangle \\
&= \frac{1}{\sqrt{r}}\sum_{m=1}^{r} e^{-2\pi\imath\frac{s(m-1)}{r}} |x^m \pmod N\rangle \\
&= e^{2\pi\imath\frac{s}{r}} \underbrace{\left(\frac{1}{\sqrt{r}}\sum_{m=1}^{r} e^{-2\pi\imath\frac{sm}{r}} |x^m \pmod N\rangle\right)}_{|u_s\rangle}.
\end{aligned}$$

  - The $m = r$ term wraps around: $x^r \equiv x^0 = 1 \pmod N$, so it perfectly maps to the $m = 0$ term, confirming the sum in brackets equals $|u_s\rangle$ itself, with eigenvalue exactly $e^{2\pi\imath s/r}$.

The Paradox: The Catch-22 of Shor's Algorithm

- The phase is exactly
$$\phi = \frac{s}{r}.$$
  - ◇ If we feed $|u_s\rangle$ into QPE, it will output an estimate of $\frac{s}{r}$.
- The catch-22 paradox:
  - ◇ Standard QPE requires feeding an exact orbital eigenstate $|u_s\rangle$ into the target register.
  - ◇ Physically preparing $|u_s\rangle$ requires knowing the period r beforehand to calculate the phase coefficients $e^{-2\pi i sk/r}$.
  - ◇ The period r is the exact parameter the algorithm is trying to find.

The Superposition Trick

- Elegant solution: Consider summing *all* the eigenstates equally, rather than any single one:

$$\begin{aligned}\frac{1}{\sqrt{r}}\sum_{s=0}^{r-1}|u_s\rangle &= \frac{1}{r}\sum_{k=0}^{r-1}\left(\sum_{s=0}^{r-1}e^{-2\pi i\frac{sk}{r}}\right)\left|x^k \pmod N\right\rangle \\ &= \left|x^0 \pmod N\right\rangle = |1\rangle\end{aligned}$$

  - By the orthogonality of roots of unity (proven on the next slide), the inner sum vanishes everywhere due to destructive interference except when $k = 0$.

- **The payoff:** the state $|1\rangle$ requires *no* knowledge of r to prepare — it is just the standard computational basis state for the integer 1. Feeding $|1\rangle$ into the target register is therefore mathematically equivalent to feeding in an equal superposition of every eigenstate $|u_s\rangle$ simultaneously, sidestepping the paradox entirely.

Roots of Unity

- Define

$$S(k) := \sum_{s=0}^{r-1} e^{-2\pi\imath\frac{sk}{r}}, \quad \text{for } 0 \leq k < r \tag{16}$$

- Let

$$q_k = e^{-2\pi\imath\frac{k}{r}}, \quad k = 0, \ldots r-1.$$

Then $q_k^r = e^{-2\pi\imath k} = 1$ is a root of the equation $\omega^r = 1$.

  - ◇ Consider the factorization

$$\omega^r - 1 = (\omega - 1)(\omega^{r-1} + \ldots + \omega + 1) = 0.$$

  - ◇ Substituting $\omega = q_k$: if $k \neq 0$, then $q_k \neq 1$ but $q_k^r = 1$, so the first factor is nonzero, forcing the second factor (which equals $S(k)$) to vanish.

$$S(k) = \sum_{s=0}^{r-1} q_k^s = \begin{cases} r & \text{if } k = 0, \\ 0 & \text{otherwise.} \end{cases} \tag{17}$$

Circuit Mechanics: Feeding $|0\rangle^{\otimes t} \otimes |1\rangle$ into the QPE (I)

- Initialize Register 2 to the trivial computational state $|1\rangle$ — by the previous two slides, this is exactly the equal superposition $\frac{1}{\sqrt{r}} \sum_s |u_s\rangle$, so we get every eigenstate "for free," with no knowledge of r required. After the parallel Hadamards on Register 1:

$$|\psi_1\rangle = \frac{1}{2^{t/2}} \sum_{y=0}^{2^t-1} |y\rangle \otimes \left(\frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} |u_s\rangle \right) = \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} \left[\frac{1}{2^{t/2}} \sum_{y=0}^{2^t-1} |y\rangle \otimes |u_s\rangle \right].$$

- Applying the controlled-modular exponentiation operators U^y:

$$|\psi_2\rangle = \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} \left[\frac{1}{2^{t/2}} \sum_{y=0}^{2^t-1} e^{2\pi \imath \frac{s}{r} y} |y\rangle \right] \otimes |u_s\rangle$$

  - ◇ The phase kickback mechanism simultaneously writes the eigenvalues of all r components onto Register 1 in parallel.

Circuit Mechanics: Feeding $|0\rangle^{\otimes t} \otimes |1\rangle$ into the QPE (II)

- Applying $\mathrm{QFT}^{\dagger}$ to obtain

$$|\psi_3\rangle = \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} \left(\sum_{m=0}^{2^t-1} c_m(s) \, |m\rangle \right) \otimes |u_s\rangle = \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} \left| \widetilde{\phi_s} \right\rangle \otimes |u_s\rangle$$

- Each branch s of the superposition now carries its own "wave packet" $\left| \widetilde{\phi_s} \right\rangle$ in Register 1, encoding an approximation to the phase s/r. We analyze the shape of this wave packet over the next three slides.

Interpretation of $\left|\widetilde{\phi_s}\right\rangle$

- The control register state after applying $\mathrm{QFT}^\dagger$ is

$$\left|\widetilde{\phi_s}\right\rangle = \sum_{m=0}^{2^t-1} c_m(s)\,|m\rangle\,.$$

- This is exactly the same calculation as the QPE error-bound slide (Section 3), with $\phi = s/r$ in place of the general phase:

$$|c_m(s)| = \frac{1}{2^t}\frac{\left|\sin\left(\pi\left(2^t\frac{s}{r}-m\right)\right)\right|}{\left|\sin\left(\pi\left(\frac{s}{r}-\frac{m}{2^t}\right)\right)\right|}.$$

- Let b be the closest integer peak, with

$$-\frac{1}{2} < e := 2^t\frac{s}{r} - b < \frac{1}{2}.$$

The Local Tail

- If $m = b + k$ is a nearby state with an integer offset $k \ll 2^t$, then

$$|c_{b+k}(s)| = \frac{1}{2^t} \frac{|\sin(\pi(e-k))|}{\left|\sin\left(\frac{\pi(e-k)}{2^t}\right)\right|} \approx \frac{1}{2^t} \frac{|\sin(\pi e)|}{\frac{\pi|e-k|}{2^t}} = \frac{|\sin(\pi e)|}{\pi|k-e|}$$

- Amplitudes near the peak drop off at the scale $O(1/k)$.
  - ◇ No exponential decay.
  - ◇ This specific algebraic decay rate ensures that the probability mass leaks only into the immediate neighboring bins, which is exactly the leakage that the extra $\log(2 + 1/2\epsilon)$ qubits in the QPE scaling formula (Section 3) are designed to suppress.

Interpretation of $\left|\widetilde{\phi}_s\right\rangle$: Global Suppression

- If m is far away from b such that there is a fixed lower bound

$$\left|\frac{s}{r} - \frac{m}{2^t}\right| \geq \delta > 0,$$

  - ◇ the denominator of $c_m(s)$ is strictly bounded away from zero:

$$\left|\sin\left(\pi\left(\frac{s}{r} - \frac{m}{2^t}\right)\right)\right| \geq C_\delta > 0$$

  - ◇ The global suppression follows as

$$|c_m(s)| \leq \frac{1}{2^t \cdot C_\delta} = O(2^{-t})$$

- Together with the previous slide, the probability distribution is such that:
  - ◇ Nearby states leak power via an $O(1/k)$ tail.
  - ◇ Faraway states undergo total exponential suppression.

Interpretation of $\left|\widetilde{\phi_s}\right\rangle$: Final Verdict

- The state $\left|\widetilde{\phi_s}\right\rangle$ behaves as a localized wave packet.
  - ◇ It concentrates $\geq 40.5\%$ of its total probability mass onto the single nearest integer b (the same bound proven in Section 3).
  - ◇ It guarantees that classical sampling always extracts a sharp approximation of $\frac{s}{r}$.

$$\left|\widetilde{\phi_s}\right\rangle \approx \left|2^t \cdot \frac{s}{r}\right\rangle.$$

- Final measurement:
  - ◇ Each branch in ψ_3 has a uniform probability $P(s) = \frac{1}{r}$ to occur.
  - ◇ Associated with each orbital eigenstate $|u_s\rangle$ in Register 2 is a wavepacket $\left|\widetilde{\phi_s}\right\rangle$ in Register 1.
  - ◇ Every single valid phase value $\frac{s}{r}$ has an equal structural opportunity to be exposed by the measurement.
- The circuit outputs a single random fraction $\tilde{\phi} \approx \frac{s}{r}$, for a uniformly random, unknown s.
- It remains to extract the true denominator r by classical continued fractions, the topic of the next slides.

Order Extraction via Continued Fractions

- QPE acts on Register 1 to yield an observed binary fraction

$$\tilde{\phi} = \frac{m}{2^t} \approx \frac{s}{r}.$$

  How do we isolate the exact discrete denominator r?

- **Theorem (Legendre):** If a rational approximation $\frac{s}{r}$ satisfies:

$$\left|\tilde{\phi} - \frac{s}{r}\right| \leq \frac{1}{2r^2}$$

  then $\frac{s}{r}$ is guaranteed to be a unique, mathematically isolated convergent of the continued fraction expansion of $\tilde{\phi}$.

- Since r is not known, we use the problem parameter N (which *is* known) to set the register size t, ensuring Legendre's condition is automatically satisfied — worked out on the next slide.

Register Scaling Strategy (I)

- Modulus proxy: Since $r < N$, a conservative error floor based on N satisfies Legendre's condition:

$$\frac{1}{2N^2} < \frac{1}{2r^2}.$$

- QPE grid resolution: QPE with 2^t bins limits the maximum rounding distance to the nearest computational peak to:

$$\left|\tilde{\phi} - \frac{s}{r}\right| \leq \frac{1}{2^{t+1}}.$$

Register Scaling Strategy (II)

- Qubit allocation: Let $L := \lceil \log_2 N \rceil$ $(2^L \geq N)$ and select

$$t = 2L + 1.$$

Then the sampling error is compressed into:

$$\left| \tilde{\phi} - \frac{s}{r} \right| \leq \frac{1}{2^{t+1}} \leq \frac{1}{4N^2} < \frac{1}{2N^2} < \frac{1}{2r^2}$$

- Legendre's theorem now guarantees s/r is a genuine convergent. The classical continued fraction algorithm extracts r in polynomial time, needing only $O(L^3)$ operations.

The Continued Fraction Expansion (CFE): Architecture

- Every real number $\tilde{\phi} \in [0, 1)$ can be developed into a nested expansion generated by a discrete dynamical system:

$$\tilde{\phi} = [0; a_1, a_2, \ldots, a_n] = \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \ldots}}}$$

- Define

$$\frac{p_k}{q_k} = [0; a_1, a_2, \ldots, a_k] = \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{\ddots + \cfrac{1}{a_k}}}}.$$

- Need to calculate (a_k, p_k, q_k) for $k \geq 0$.
  - ◇ Textbooks often split the CFE algorithm into two parts (next slide).
  - ◇ In practice, both parts should be executed dynamically, interleaving the fraction updates at every single iteration to check the stopping criteria as early as possible.

CFE: Algorithm

Part 1. Computing a_k

Initialization: $\zeta_0 = \tilde{\phi} = \frac{m}{2^t}$.

Do:

$$a_k = \lfloor \zeta_k \rfloor$$

$$x_k = \zeta_k - a_k$$

$$\zeta_{k+1} = \frac{1}{x_k} \quad (\text{if } x_k \neq 0)$$

Part 2. Computing $\frac{p_k}{q_k}$

Initialization:

$$\begin{cases} p_{-2} = 0 \\ q_{-2} = 1 \end{cases} \quad \begin{cases} p_{-1} = 1 \\ q_{-1} = 0 \end{cases}$$

Do:

$$p_k = a_k p_{k-1} + p_{k-2}$$

$$q_k = a_k q_{k-1} + q_{k-2}$$

- Part 1 is essentially a discrete "Euclidean algorithm" for real numbers.
- The linear recurrence in Part 2 allows us to compute the fractions without re-evaluating the nested denominators, leading to $O(L^3)$ polynomial time complexity.

CFE: Stopping Criterion

- Criterion 1: $x_k = 0$ (Exact rational termination)
  - ◇ The input phase $\tilde{\phi} = \frac{m}{2^t}$ belongs to the field of rational numbers

$$\tilde{\phi} = \frac{m}{2^t} = \frac{p_k}{q_k}.$$

  - ◇ If $\gcd(s, r) = 1$, then $q_k = r$ exactly.
- Criterion 2: $q_k \geq N$ (Search space boundary overflow)
  - ◇ The algorithm has isolated a proper factor of the period.
  - ◇ Fixable, but more analysis is needed.
- Criterion 3: $x^{q_k} \equiv 1 \pmod{N}$ (Algebraic success)
  - ◇ If the test passes, then $r \mid q_k$.
  - ◇ Since $1 \leq q_k \leq r$, it must be $r = q_k$.
- In practice (and in the blind worked example two slides from now), Criterion 3 is the one actually used: it is the only criterion that can be checked without already knowing r.

Numerical Execution Trace

- Suppose the true hidden period is $r = 7$, and the phase-estimation circuit happens to sample the branch $s = 4$. So the exact phase is $\frac{s}{r} \approx 0.571428$.
- Using $t = 6$ in the QPE ($2^t = 64$), suppose QPE returns $m = 37$, yielding the empirical fraction $\tilde{\phi} = \frac{37}{64} = 0.578125$.
- Iterations through the CFE:

k	ζ_k	a_k	x_k	p_k	q_k	Convergent $\frac{p_k}{q_k}$
0	$37/64 \approx 0.5781$	0	$37/64$	0	1	$0/1$
1	$64/37 \approx 1.7297$	1	$27/37$	1	1	$1/1$
2	$37/27 \approx 1.3704$	1	$10/27$	1	2	$1/2$
3	$27/10 = 2.7000$	2	$7/10$	3	5	$3/5$
4	$10/7 \approx 1.4286$	**1**	$3/7$	**4**	**7**	$\mathbf{4/7}$

- At $k = 4$, the algorithm constructs $q_4 = 7$, exactly the true period r (verified: $\gcd(s, r) = \gcd(4, 7) = 1$, so Criterion 1 guarantees $q_4 = r$ exactly, matching Criterion 3's check $7^7 \equiv 1 \pmod{N}$ for the actual problem's N).

Why Does Finding the Order Allow Factorization?

- Suppose p and q are prime numbers and $N = pq$.
- If the order r of x is odd, choose another random x; otherwise, write

$$(x^{r/2} - 1)(x^{r/2} + 1) = x^r - 1 \equiv 0 \pmod{N}.$$

  - If $x^{r/2} + 1 \equiv 0 \pmod{N}$, choose a different x.
  - If $x^{r/2} + 1 \not\equiv 0 \pmod{N}$, then $d := \gcd(x^{r/2} - 1, N)$ must be either p or q.
    - $x^{r/2} - 1$ cannot be a multiple of N; otherwise, $x^{r/2} \equiv 1 \pmod{N}$, contradicting $ord_N(x) = r$ (since $r/2 < r$ would then be a smaller valid order).
    - Since $N = pq$ divides $(x^{r/2} - 1)(x^{r/2} + 1)$ but divides neither factor alone, N's two prime factors must split between the two factors — so $x^{r/2} - 1$ shares exactly one of p, q with N.

Generalizing Beyond $N = pq$

- This argument did not actually need $N = pq$ specifically.
  - For general composite N, the same $d = \gcd(x^{r/2} - 1, N)$ gives us a nontrivial factor of N.
  - Factorize d and N/d recursively to obtain all prime factors.

Complete Pipeline for Shor's Factoring Protocol

1. **Trivial Checks:** If N is even, return 2. If $N = a^b$ for integers $a \geq 2, b \geq 2$, return a (checked via classical root testing).
2. **Random Selection:** Choose a random integer x such that $1 < x < N$.
3. **Direct Hit:** Compute $\gcd(x, N)$ classically. If $\gcd(x, N) > 1$, return this value as a non-trivial factor.
4. **Quantum Subroutine:** Run the **Quantum Order-Finding Subroutine** using the initialization state $|1\rangle$ to discover the order r of $x \pmod N$.
5. **Condition Evaluation:** If r is even and $x^{r/2} \not\equiv -1 \pmod N$, compute and return the factors:

$$d = \gcd(x^{r/2} - 1, N)$$

6. **Fall Back:** If the conditions fail, loop back and select a new random x.

Blind Case Study: Factoring 15 from Ignorance

- To understand the true power of the algorithm, we must execute the steps acting **completely blind** to the hidden period r.
- The next few slides show, step by step, the execution of Shor's algorithm on $N = 15$ — the quantum machine never "knows" that $r = 4$ in advance; it is discovered purely from the measured output.

Step 1 to 3: Initialization & Setup

1. **Step 1: Parameter Selection**
   - Target Modulus: $N = 15$
   - Chosen Random Base: $x = 7$
   - Verification: $\gcd(7, 15) = 1$ (Valid base)
2. **Step 2: Register Allocation**
   - Register 1 (Precision): Requires $2^t \geq N^2 \implies 2^{11} = 2048$ ($t = 11$ qubits)
   - Register 2 (Target): Requires $\lceil \log_2(15) \rceil = 4$ qubits
3. **Step 3: Register State Initialization**

$$|\Psi_0\rangle = |0\rangle^{\otimes 11} |1\rangle$$

Step 4 to 6: The Quantum Pipeline

4. **Step 4: Create Superposition**

$$|\Psi_1\rangle = \left(H^{\otimes 11} \otimes I\right)|\Psi_0\rangle = \frac{1}{\sqrt{2048}} \sum_{k=0}^{2047} |k\rangle\,|1\rangle$$

5. **Step 5: Controlled-U Evaluation**
   - Operator: $U|y\rangle = |7y \pmod{15}\rangle$
   - State: $|\Psi_2\rangle = \frac{1}{\sqrt{2048}} \sum_{k=0}^{2047} |k\rangle \otimes U^k |1\rangle$
   - *Note: the underlying period r remains fully unobserved by the algorithm at this stage.*

6. **Step 6: Inverse QFT Transition**

$$|\Psi_3\rangle = \left(\mathrm{QFT}^\dagger \otimes I\right)|\Psi_2\rangle$$

Frequencies shift to localized spike states at multiples of $2^t/r$ (still unknown to us at this point).

Step 7 & 8: Hardware Output & CFE Start

7. **Step 7: Physical Measurement**
   - Hardware probes Register 1. State collapses into a peak index.
   - **Output Data: $\mathbf{k = 1536}$**. Quantum computer powers off.

8. **Step 8: CFE Variable Initialization**
   - Move to classical processor. Form input ratio:

$$\zeta_0 = \frac{k}{2^t} = \frac{1536}{2048} = 0.75$$

   - Initialize tracking convergence variables:

$$p_{-2} = 0,\ p_{-1} = 1,\ q_{-2} = 1,\ q_{-1} = 0$$

Step 9: Blind CFE Loop Iterations

9 Step 9: Execution of Recurrence Relations

- **Iteration 0:** $\zeta_0 = 0.75$

$$a_0 = \lfloor 0.75 \rfloor = 0, \quad f_0 = 0.75 - 0 = 0.75$$
$$p_0 = 0(1) + 0 = 0, \quad \mathbf{q_0 = 0(0) + 1 = 1}$$

Blind Order Check: $7^1 \equiv 7 \pmod{15} \neq 1$. (Fails). Advance: $\zeta_1 = \frac{1}{0.75} = \frac{4}{3}$.

- **Iteration 1:** $\zeta_1 = 4/3$

$$a_1 = \lfloor 4/3 \rfloor = 1, \quad f_1 = 4/3 - 1 = 1/3$$
$$p_1 = 1(0) + 1 = 1, \quad \mathbf{q_1 = 1(1) + 0 = 1}$$

Blind Order Check: $q_1 = 1$ has failed previously (same as q_0). Advance: $\zeta_2 = \frac{1}{1/3} = 3$.

Step 9 (Cont.): Period Discovery

- **Iteration 2:** $\zeta_2 = 3$

$$a_2 = \lfloor 3 \rfloor = 3, \quad f_2 = 3 - 3 = 0$$
$$p_2 = 3(1) + 0 = 3, \quad \mathbf{q_2 = 3(1) + 1 = 4}$$

- **Blind Order Check on Candidate** $r = q_2 = 4$**:**

$$7^4 = 2401$$

$$2401 = (160 \times 15) + 1 \implies \mathbf{7^4 \equiv 1} \pmod{\mathbf{15}}$$

- **Output Status:** The check passes (Criterion 3 from two slides ago). Remainder $f_2 = 0$ independently confirms this (Criterion 1). The loop terminates. The discovered hidden order is $\mathbf{r = 4}$.
- Convergent fraction $\frac{p_2}{q_2} = \frac{3}{4} \implies$ harmonic index $s = 3$.

Step 10: Classical Factor Splitting

10 Step 10: Greatest Common Divisor Calculations

- Parity Check: $r = 4$ is an even integer. (Passes).
- Non-Triviality Check:

$$x^{r/2} \pmod{N} \implies 7^{4/2} = 49 \equiv 4 \pmod{15}$$

  Since $4 \not\equiv -1 \pmod{15}$, criteria hold.
- Run Euclidean Splits:

$$\text{Factor } 1 = \gcd(4 - 1, 15) = \gcd(3, 15) = \mathbf{3}$$

$$\text{Factor } 2 = \gcd(4 + 1, 15) = \gcd(5, 15) = \mathbf{5}$$

Final Execution Output

The prime factors of 15 are isolated as **3** and **5**.

Summary of Lecture 7

- **The QFT** decomposes into an efficient $O(n^2)$-gate circuit via the product representation theorem, making it exponentially cheaper than the classical FFT on n qubits.
- **Quantum Phase Estimation** converts an eigenvalue phase ϕ — invisible to direct measurement — into a readable computational basis state, with success probability $\geq 4/\pi^2 \approx 40.53\%$ guaranteed even when ϕ has no exact finite binary representation.
- **Order-finding** encodes the unknown period r of $x^k \bmod N$ into exactly this kind of eigenvalue phase; the superposition trick (initialising with $|1\rangle$ rather than any individual eigenstate $|u_s\rangle$) sidesteps the catch-22 of needing r in advance.
- **Continued fractions** extract the exact integer r from QPE's output in classical polynomial time. Shor's algorithm chains all of this together into a complete factoring protocol, verified step-by-step on the blind example $N = 15$.

Looking Ahead: Lecture 8

- This lecture built the quantum subroutine. Lecture 8 asks: *what can an adversary actually do with it?*
- **The RSA cryptosystem** — its history (Diffie-Hellman 1976, MIT 1977, GCHQ's classified prehistory), the number theory that makes it work (Fermat's Little Theorem, Euler's totient), and exactly why its security collapses once factoring becomes easy.
- **Breaking RSA from the eavesdropper's perspective** — a more direct attack in which the adversary uses order-finding on the intercepted ciphertext itself to recover the plaintext, bypassing the need to factor N explicitly.
- **The Hidden Subgroup Problem** — the unifying abstraction that places Shor's algorithm, Simon's algorithm, and the discrete logarithm problem inside a single framework, and identifies the open non-Abelian frontier.