Quantum Fourier Transform

Understand the QFT as a coherent basis change that encodes labels in phase patterns rather than printed arrays.

Undergraduate 22 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Calculate the four-point QFT of a basis state.
- Verify numeric output ordering and the role of final swaps.
- Explain why QFT gate efficiency is not automatic classical-array readout efficiency.

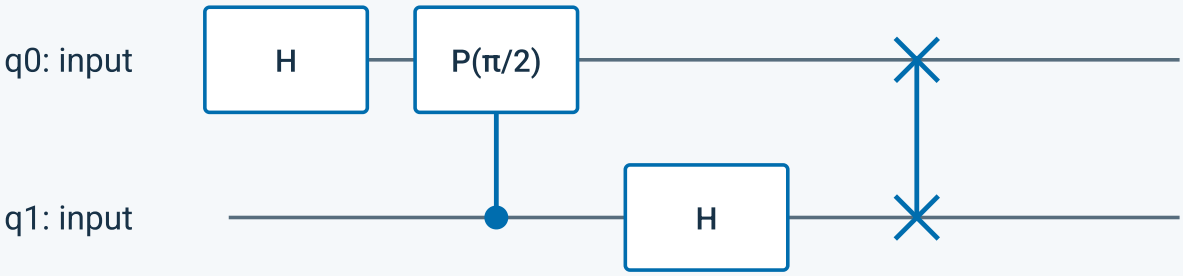
Turn a phase pattern into a basis label

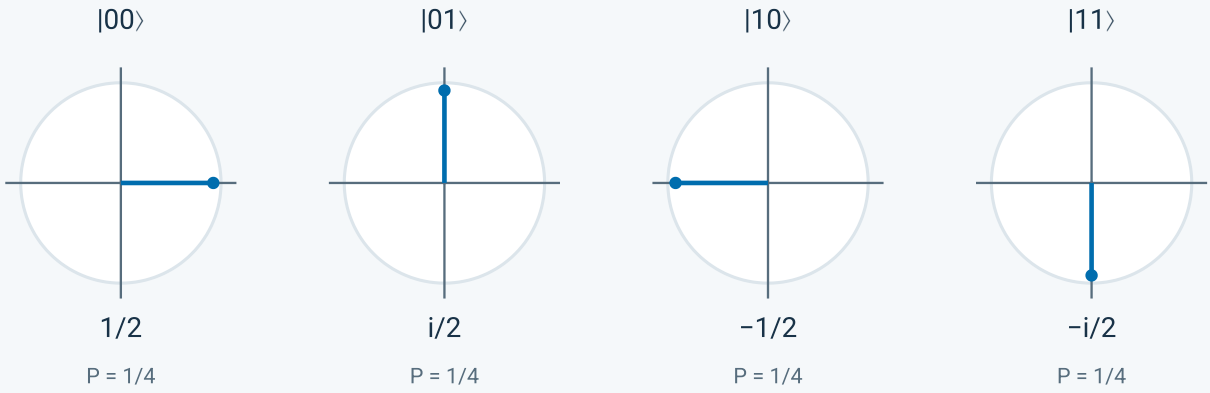
The quantum Fourier transform, or QFT, is a unitary change of basis. It maps basis labels to states with equal magnitudes and regularly spaced phases. Its inverse turns those structured phase patterns back into labels. That operation is central to phase estimation and period finding because those tasks naturally create phase patterns related to an unknown answer.

Review complex numbers, tensor products and phase. You can understand the four-dimensional example with the four complex directions one, i, negative one and negative i before considering a large circuit.

WORKED DIAGRAM · A11

The four-point QFT: gates and phases





With q0 leftmost and the positive-exponent convention, H on q0, controlled phase π/2, H on q1 and a final SWAP implement F4. Input 01 gives phases 1, i, −1, −i divided by two.

Definition and output setup

For N basis labels, this book uses the positive-exponent convention:

$$F_N|x\rangle = \frac{1}{\sqrt{N}} \sum_{y=0}^{N-1} e^{2\pi i x y/N} |y\rangle.$$

The inverse has the negative exponent. The operation transforms a quantum state. It does not receive an ordinary array and return a printed array of all Fourier coefficients. Preparing an arbitrary input state and reading a large classical output can be expensive. Those costs must be included before comparing a QFT-based application with a classical fast Fourier transform.

For qubit registers, N is usually a power of two. Basis labels are numeric bitstrings with q0 leftmost here. A circuit implementation may reverse wire order unless final swaps or an explicit output relabelling are included. Sign and ordering conventions should be written down before numerical comparisons.

Work the four-point transform

For x equal to one and N four, the output amplitudes are one half, i over two, negative one half and negative i over two. Every basis probability is one quarter. For x equal to zero, all four amplitudes are positive one half, also giving uniform probabilities. Their distinction is entirely in phase.

Applying the inverse QFT to either output recovers its original basis label exactly in the ideal model. A direct Z measurement immediately after QFT would not distinguish those two inputs, because their magnitudes agree. The transform is useful as part of a coherent computation that makes use of phase, not as a stand-alone measurement that reveals every input.

For an arbitrary four-entry column, calculate each transformed entry by multiplying one matrix row by the input. The first output amplitude is half the sum of the inputs. The next rows use the phase patterns above and their powers. This matrix calculation is a small, independent reference test for a gate implementation.

The four-point matrix puts the phase pattern directly beside its numeric example.

$$F_4 = \frac{1}{2}\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & i & -1 & -i \\ 1 & -1 & 1 & -1 \\ 1 & -i & -1 & i \end{pmatrix}, \qquad F_4|1\rangle = \frac{1}{2}\begin{pmatrix} 1 \\ i \\ -1 \\ -i \end{pmatrix}.$$

Rows and columns use numeric order $00, 01, 10, 11$. Every entry of this output column has squared magnitude $1/4$, while its phases differ.

Checkpoint 1

PAUSE & THINK

How do F4|0⟩ and F4|1⟩ differ?

- [ ] Their immediate Z probabilities differ
- [ ] Their relative phases differ
- [ ] Only one is normalised

Answer: Their relative phases differ. The states have equal amplitude magnitudes and different relative phases. An inverse QFT distinguishes the coherent patterns.

Why it is unitary

Take two different input labels. Their Fourier columns have an inner product equal to an average of roots of unity. Those phases are evenly spaced around the circle and sum to zero. A column's inner product with itself is one. The columns are therefore orthonormal, which makes the transform unitary and its conjugate transpose its inverse.

The same cancellation argument explains how the inverse recognises a matching phase pattern. At the correct label, contributions align. At other labels, they circle around and cancel. This is a many-component extension of H distinguishing plus and minus. H is the two-point Fourier transform under this convention.

The roots of unity sum to zero unless the two input labels match.

$$\langle x'|F_N^\dagger F_N|x\rangle = \frac{1}{N}\sum_{y=0}^{N-1} e^{2\pi i(x-x')y/N} = \delta_{x,x'}.$$

The symbol $\delta_{x,x'}$ equals one for equal labels and zero otherwise. This is the orthonormal-column test for unitarity.

A two-qubit circuit

For two qubits in the book's ordering, apply H to q0, a controlled phase of pi over two between q1 and q0, then H to q1, followed by a SWAP of the two output wires. This implements the positive-exponent four-point transform. Controlled phase is symmetric as a diagonal two-qubit matrix, but the wire roles in a larger synthesis still need consistent notation.

Test input 01. Before the final SWAP, the amplitudes have the phase pattern in reversed bit order. After the SWAP, the numeric-order column is one, i, negative one, negative i divided by two, matching the definition. Omitting the SWAP can be valid if the surrounding computation explicitly accounts for the reversal; silently comparing it with numeric-order output is not valid.

For n qubits, the standard exact circuit uses n H gates, a quadratic number of controlled phases and output swaps. Very small-angle rotations can be omitted in an approximate transform with a controlled error bound. A fault-tolerant implementation must also synthesise rotations to sufficient precision, so an abstract gate count is not a full physical cost.

Checkpoint 2

PAUSE & THINK

Why include final swaps in the numeric-order QFT circuit?

- [ ] To correct output bit reversal
- [ ] To measure all phases
- [ ] To change the transform from unitary to stochastic

Answer: To correct output bit reversal. The swaps align physical wire order with the stated numeric bitstring convention.

Periodic structure and a small example

Consider the two-qubit state consisting of 00 and 10 with equal amplitudes. Its labels differ by two. The four-point QFT gives equal support on 00 and 10 again, reflecting a reciprocal spacing relation. More general periodic states produce peaks related to the inverse period. Finite register length and periods that do not divide the grid introduce broadened peaks and require careful reconstruction.

This is why period-finding algorithms do not simply measure a state containing repeated function values and read the period directly. They use a Fourier transform to make the repetition influence a distribution over frequency-like labels, then process those labels classically. QFT supplies the basis change; the rest of the algorithm supplies the useful structure and interpretation.

A two-point periodic state shows cancellation and reinforcement across several Fourier rows.

$$F_4 \frac{|0\rangle + |2\rangle}{\sqrt{2}} = \sum_{y=0}^{3} \frac{1 + (-1)^y}{2\sqrt{2}} |y\rangle = \frac{|0\rangle + |2\rangle}{\sqrt{2}}.$$

Odd output labels cancel; even labels survive. In this small example the support is unchanged, while the general relation is reciprocal spacing.

Check a transform pair without a circuit library

The equal four-entry state is the QFT of basis zero. Applying the inverse transform should therefore concentrate every amplitude at zero. In this particular case, the forward transform does so too: each nonzero-frequency row contains a complete cycle of roots of unity and sums to zero. This is a valuable hand test because it checks cancellation across every output row, not only the magnitudes of one matrix column.

The transform applied twice has a simple interpretation: it reverses the numeric label modulo N. For four points, input one becomes three after two forward transforms, while zero and two return to themselves. Four forward transforms give the identity. These identities follow from Fourier-column orthogonality and provide independent implementation checks. They do not mean two forward transforms are always the inverse of one, or that every wire-reversal convention can be ignored. Test numeric labels one and three explicitly, where a mistaken sign or output reversal becomes visible rather than hidden by a symmetric input.

Predict and test

This lab uses three qubits and eight labels, extending the four-point calculation above. For input one, predict equal magnitudes and a phase step of pi over four. Select the inverse option to check recovery of the original basis label.

L07 · INTERACTIVE LAB

Runs on your device

Quantum Fourier transform

The positive-exponent QFT gives a basis state equal output magnitudes with a structured phase pattern. Numeric output ordering is used, without an implicit bit reversal.

MODEL · Ideal pure-state vector, complex double precision

ORDER · q0 is the leftmost bit

Initial setup · reference state

Exact probability



Statevector before readout · ideal simulation

Basis	Amplitude a + bi	Magnitude	Phase (rad)	Probability
|000⟩	0.353553 + 0i	0.353553	0.000000	12.5%
|001⟩	0.25 + 0.25i	0.353553	0.785398	12.5%
|010⟩	0 + 0.353553i	0.353553	1.570796	12.5%
|011⟩	−0.25 + 0.25i	0.353553	2.356194	12.5%
|100⟩	−0.353553 + 0i	0.353553	3.141593	12.5%
|101⟩	−0.25 − 0.25i	0.353553	-2.356194	12.5%
|110⟩	0 − 0.353553i	0.353553	-1.570796	12.5%
|111⟩	0.25 − 0.25i	0.353553	-0.785398	12.5%

Signs and complex phases are part of the state. Phase is undefined for a zero amplitude. A shared global phase changes no probabilities; relative phases can affect a later operation. Values are rounded to six decimal places.

Prediction to try: Predict the output magnitudes and the phase step between neighbouring indices. What should an inverse QFT recover?

Challenge: Compare basis inputs 0 and 1. Their QFT probabilities match; their phases differ. Verify the inverse for both.

Open the live experiment studio

Chapter summary

- The QFT is a coherent basis change with an explicit sign and label convention.
- Phase patterns can differ while all immediate Z probabilities agree.
- A useful algorithm supplies structure before the transform and interprets the later measurement.

A common misconception

A QFT measurement prints every Fourier coefficient. The output is a quantum state; one measurement gives one label. Preparing and reading large classical arrays have additional costs.

Code practice

Download the tested Qiskit notebook and follow the pinned environment instructions. Inspect the circuit, predict a result and run its independent assertions. The modular-action example is a small-instance demonstration; the notebook states its precision and post-processing limits.

Exercises

1. Write the four-point QFT output for input two. What are its Z probabilities?
2. Two QFT outputs have identical uniform histograms. Does that establish that the outputs are the same state? Give the zero-input and one-input counterexample.

Worked answers

1. The phases are one, negative one, one, negative one, all divided by two. Every squared magnitude is one quarter. The phase step is pi.
2. No. Zero maps to four positive halves, while one maps to one, i, negative one and negative i divided by two. Their phases differ and the inverse transform distinguishes them perfectly. A probability histogram discards the structure that the QFT encodes.

Check your understanding

PAUSE & THINK

Can a direct Z histogram distinguish the QFT of basis zero from the QFT of basis one?

- [ ] Yes, because their probabilities differ
- [ ] No; their magnitudes agree while their phases differ
- [ ] Only if zero has no amplitude

Answer: No; their magnitudes agree while their phases differ. Both have uniform Z probabilities. Their phase patterns differ and an inverse QFT distinguishes them perfectly.

Sources

IBM's phase-estimation procedure and QFT derivation defines the transform and its circuit construction. The official QFT teaching module provides further guided practice.