Quantum Counting

Estimate the number of marked labels by reading Grover's eigenphase and converting it into a count.

Undergraduate 21 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Relate the marked fraction to Grover eigenphases.
- Convert either conjugate phase branch into the same count.
- Explain finite-resolution and controlled-power costs.

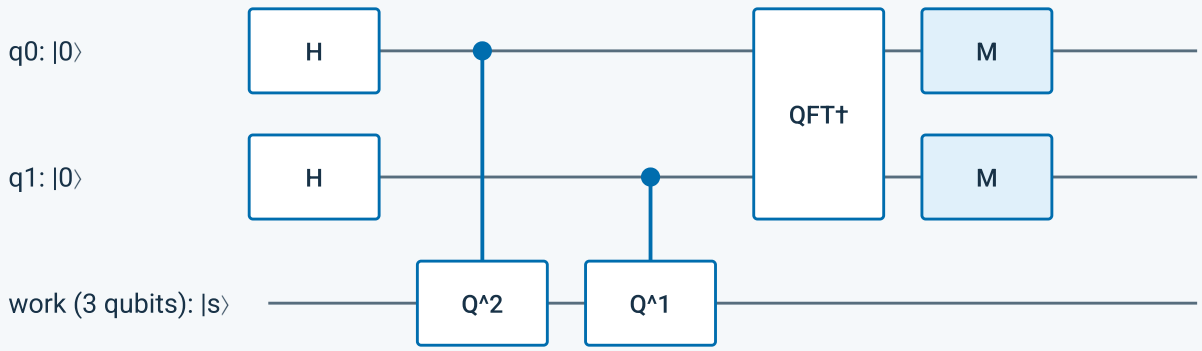
Count by learning a rotation angle

Quantum counting estimates how many labels satisfy a predicate without listing them all. Grover's iterate rotates by an angle determined by the marked fraction. Phase estimation can learn that rotation angle, and a classical calculation converts it into an estimate of the count. The algorithm combines Grover search with phase estimation.

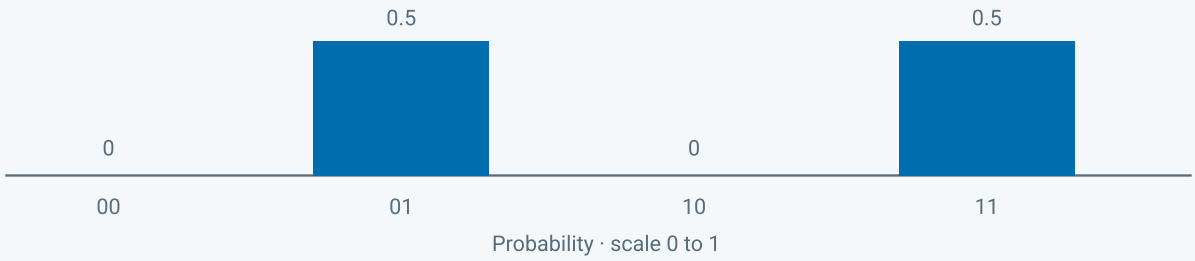
The intuitive bridge is simple: if many labels are marked, the prepared uniform state already has a large good component, and the Grover rotation is larger. If few are marked, the angle is smaller. Counting reads that dynamical parameter. It does not obtain a full table of predicate values or enumerate the marked records.

WORKED DIAGRAM · A08

Counting: two phases, one count



Counting-register output



Phase 1/4 → marked count 4

Phase 3/4 → marked count 4

For N=8, M=4, eigenphase labels 01 and 11 each have probability 1/2 and both convert to count four.

Setup and resources

There are N labels and M marked labels. You can prepare a uniform superposition and coherently evaluate the predicate. The desired output is an approximation to M, with an accuracy and failure probability that must be specified. The basic phase-estimation construction needs controlled powers of the Grover iterate, including the necessary preparation and reflection operations.

If the required error is expressed in the marked fraction, then multiplying by N converts it into absolute count error. Distinguishing a count of ten from eleven in a huge space is a more demanding task than estimating a large marked fraction to a few percentage points. Precision belongs in the problem statement.

The procedure

Build the Grover iterate for the predicate. Prepare the search register uniformly. Add a counting register, apply H to its qubits, and use them to control successive powers of the iterate. Apply the inverse Fourier transform to the counting register and measure it. Interpret the sampled eigenphase through the relation between the rotation angle and M over N.

The uniform search state is a superposition of two Grover eigenvectors. Consequently, ideal phase estimation can return either of two conjugate phases. Both encode the same marked fraction. The reconstruction rule should handle that symmetry rather than treating one branch as an error.

The marked fraction fixes the two eigenphases of the Grover iterate used in counting.

$$\sin^2\theta = \frac{M}{N}, \qquad Q|v_\pm\rangle = e^{\pm 2i\theta}|v_\pm\rangle, \qquad \phi_+ = \frac{\theta}{\pi}, \quad \phi_- = 1 - \frac{\theta}{\pi}.$$

M is the marked count and N the number of labels. The phase fractions $\phi_\pm$ are measured in turns. Use $Q = (2|s\rangle\langle s| - I)O_f$, as in A06; these formulas assume $0 < M < N$.

A measured phase estimate converts back to a count using the same rule for either branch.

$$\hat{\phi} = \frac{j}{2^t}, \qquad \widehat{M} = N\sin^2(\pi\hat{\phi}), \qquad \sin^2(\pi(1-\phi)) = \sin^2(\pi\phi).$$

Here t is the counting-register width and j its numeric measurement label. A finite-grid estimate $\widehat{M}$ can be nonintegral.

Work an exactly representable example

Take N equal to eight and M equal to four. The marked fraction is one half, so theta is pi over four. The Grover eigenvalues have phases plus or minus twice theta, corresponding to fractions one quarter and three quarters of a full turn. A two-qubit phase-estimation register can represent those fractions exactly.

The measured counting labels are therefore 01 or 11 in numeric two-bit order. Convert either phase using N times sine squared of pi times the phase fraction. For one quarter, sine squared of pi over four is one half, giving four. For three quarters, the sine squared is also one half. The two outcomes encode the same count.

This example shows why phase estimation on a non-eigenstate input can still be useful. It samples eigenphase branches weighted by the input's overlaps. The application does not need to know which branch occurred if both lead to the same desired classical quantity.

For eight labels with four marked, two counting bits are sufficient for both phase branches.

$$\theta = \frac{\pi}{4}, \quad \phi \in \left\{\frac{1}{4}, \frac{3}{4}\right\}, \quad j \in \{1, 3\} = \{01_2, 11_2\}, \quad \widehat{M} = 8 \cdot \frac{1}{2} = 4.$$

Each branch has probability $1/2$ for the uniform search-state input in the ideal example. Different measured labels can therefore give the same answer.

Checkpoint 1

PAUSE & THINK

With N=8 and two counting bits, label 11 is measured. What count estimate follows?

- [ ] 6
- [ ] 4
- [ ] 3

Answer: 4. Label 11 means phase fraction 3/4. Its sine squared is 1/2, so the estimated count is four.

Finite precision changes the answer

Now let M be two while keeping N eight. The marked fraction is one quarter, theta is pi over six, and the phase fractions are one sixth and five sixths. A power-of-two phase grid cannot represent these exactly. The phase-estimation output spreads across neighbouring grid values, and the converted count can be nonintegral.

Do not round every output immediately and assume the rounded integer is guaranteed correct. A rigorous procedure relates phase precision and failure probability to count error, possibly repeats the experiment and combines results. Counts near the boundaries have different sensitivity from counts near the middle because the sine-squared conversion is nonlinear.

A small phase error changes the count most strongly near the middle of its range.

$$\frac{dM}{d\phi} = \pi N \sin(2\pi\phi), \qquad |\widehat{M} - M| \leq \pi N \, |\widehat{\phi} - \phi|.$$

The inequality follows from the largest derivative magnitude and applies to a consistent unwrapped phase branch. It is a simple sensitivity bound, rather than a complete confidence guarantee for a QPE run.

Checkpoint 2

PAUSE & THINK

What happens when the true phase fraction is 1/6 and the phase grid has eight points?

- ☐ Several outcomes can occur even with perfect gates
- ☐ The circuit must always return an exact 1/6 label
- ☐ Every off-grid result proves hardware noise

Answer: Several outcomes can occur even with perfect gates. A finite binary grid cannot exactly encode 1/6. Interpret the resulting distribution with a precision and confidence analysis.

Why the estimate is related to the true count

The Grover state remains in the good/bad plane. Its rotation eigenvalues encode theta, where sine squared theta equals M over N. Phase estimation learns the eigenvalue phase. The classical sine-squared conversion therefore reconstructs the marked fraction. The proof has two parts: the Grover eigenphase relation and the phase-estimation accuracy guarantee.

If either part is approximated in implementation, its error contributes to the final count estimate. Imperfect controlled powers, approximate reflections and finite phase resolution are different errors. A resource estimate should choose their tolerances so the total output accuracy meets the task description.

Complexity and comparison

Ordinary random sampling estimates a marked fraction to additive error epsilon with a sample scale proportional to one over epsilon squared for constant confidence. Canonical quantum amplitude-estimation/counting methods can improve coherent query scaling to roughly one over epsilon under their stronger access model. The number of counting qubits grows logarithmically with the phase-grid size, but the total controlled-iterate work does not: implementing powers can require many base iterate calls.

If the goal is an exact count or a relative error for a very rare marked set, the relevant bounds need their additional dependence on M and N. A statement such as “counting takes logarithmic time because it uses a small counting register” confuses storage width with evolution cost. Include oracle construction and coherent depth before interpreting the query advantage as application runtime.

Chapter summary

- Counting estimates Grover’s angle instead of enumerating marked labels.
- Conjugate phase branches give the same count through $N \sin^2(\pi\varphi)$.
- Finite precision, coherent implementation error and confidence must all fit the desired count accuracy.

A common misconception

Two conjugate phase outcomes imply two different marked counts. The sine-squared conversion is symmetric under $\varphi \leftrightarrow 1-\varphi$. Both outcomes encode the same fraction.

Exercises

1. With N eight and an exactly measured phase one quarter, what count estimate follows? Does phase three quarters imply a different count?
2. Why does adding one counting qubit generally increase more than just memory use in the canonical controlled-powers circuit?

Worked answers

1. The estimate is eight times sine squared of pi over four, which is four. Three quarters gives the same sine squared and therefore the same estimate. The two eigenphase branches are symmetric.
2. The new qubit controls a higher power of the iterate. If powers are implemented by repeated base calls, the coherent work approximately doubles at the largest scale. Counting-register width is not the same as the number of oracle invocations or total circuit depth.

Check your understanding

PAUSE & THINK

For N=8, do measured phases 1/4 and 3/4 imply different counts?

- [ ] Yes: two and six
- [ ] No: both imply four
- [ ] Only 1/4 is a valid branch

Answer: No: both imply four. The count estimate is $N \sin^2(\pi\varphi)$. Both phase fractions have sine squared 1/2, giving four.

Sources

Brassard and colleagues connect amplitude estimation with approximate counting. IBM's phase-estimation procedure supplies the controlled-power and inverse-Fourier mechanism used in the canonical construction.