Amplitude Amplification

Generalise search by reflecting about a biased coherent preparation and a recognisable good subspace.

Undergraduate | 21 min lesson | Dr. Charith Jayasekara

By the end, you'll be able to

- Construct a preparation-state reflection using A and A†.
- Trace amplification for an initial success probability of one quarter.
- Identify inverse-access and recognition costs in an application.

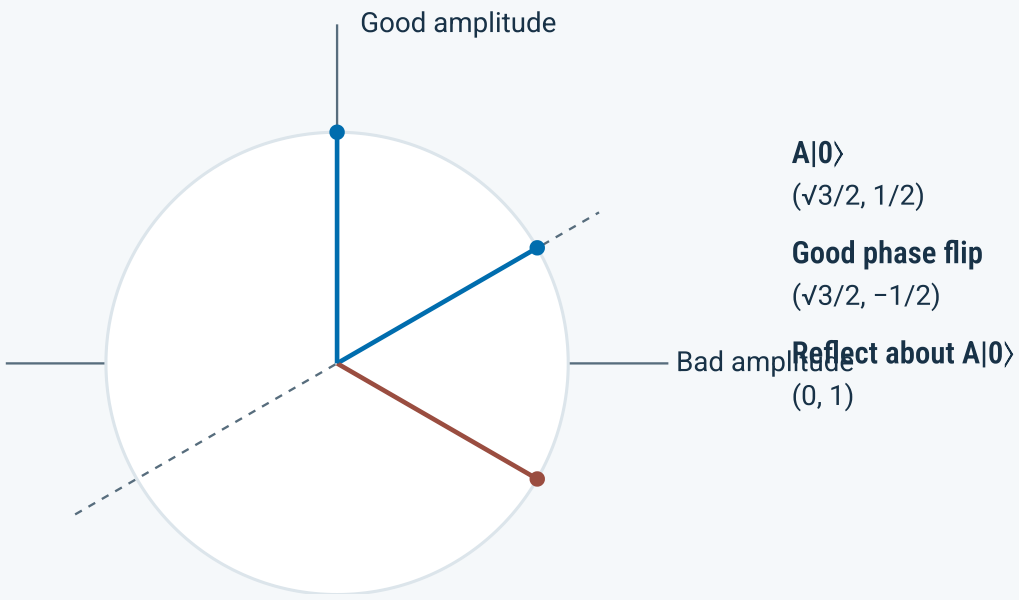
Improve an existing coherent success probability

Amplitude amplification generalises Grover search. The starting state need not be uniform over every candidate. A reversible procedure A prepares some state, and a predicate recognises good outcomes. If measuring that state succeeds with probability p, amplification can raise success using coherent calls to A, its inverse and suitable reflections.

Study Grover search first. The new idea is to reflect about whatever A prepares, rather than only the uniform state. The prerequisite that A can be run coherently and inverted is a real resource assumption. An arbitrary classical sampling service does not automatically supply it.

WORKED DIAGRAM · A07

Amplify the actual prepared state



Start at angle θ=π/6 with success 1/4. The good sign flip reflects across the bad axis, then reflection about the prepared state reaches |1⟩. This uses a biased preparation, not uniform diffusion.

State the input and output

You receive a unitary preparation A, a standard input zero state, and a phase test that identifies a good subspace. Write the prepared state as a good component plus a bad component. Their squared weights sum to one. The task is to produce a good classical outcome on measurement, with a specified success probability.

The procedure does not need to know every good label in advance. It needs an implementable coherent recognition test. It also does not reveal the entire prepared state. Like Grover, it returns a sampled candidate, and the application should verify or interpret that candidate according to its task description.

Separate the prepared state into normalised bad and good directions. The angle records the initial success probability.

$$|\psi\rangle = A|0\rangle = \sqrt{1-p}\,|B\rangle + \sqrt{p}\,|G\rangle = \cos\theta\,|B\rangle + \sin\theta\,|G\rangle, \qquad p = \sin^2\theta.$$

The states $|B\rangle$ and $|G\rangle$ are orthogonal and normalised. The basic nontrivial case has $0 < p < 1$ and $0 < \theta < \pi/2$.

Construct the two reflections

One reflection changes the sign of the good subspace. The other reflects about the state A applied to zero. To implement the latter, undo A, apply a reflection about zero, then reapply A. Up to a common sign convention, a full amplification step combines these reflections.

This construction explains why inverse access matters. If a preparation contains an uncontrolled measurement whose record is discarded, it may not admit the required simple inverse. Some randomised classical computations can be made reversible by retaining their random coins and workspace coherently, but that conversion has a cost and requires a plan to uncompute temporary information.

Reflecting about the actual preparation can be implemented by undoing and redoing it.

$$S_G = I - 2\Pi_G, \quad R_0 = 2|0\rangle\langle 0| - I, \quad R_\psi = AR_0A^\dagger = 2|\psi\rangle\langle\psi| - I, \quad Q = R_\psi S_G.$$

Π_G projects onto the good subspace. Reading right to left, R_ψ applies $A^\dagger$, then R_0, then A. This fixes the iterate's sign convention.

Checkpoint 2

PAUSE & THINK

Which operations implement reflection about A|0⟩, in execution order?

- ☐ A, then measure, then A
- ☐ H on every wire regardless of A
- ☐ A†, then R0, then A

Answer: A†, then R0, then A. Rψ=A R0 A†, with the rightmost operation executed first.

Worked one-qubit example

Let A prepare amplitudes square root of three over two for zero and one half for one. Define one as good. The initial success probability is one quarter, and the initial angle toward the good direction is 30 degrees. One amplification step should move that angle to 90 degrees, giving certain success in this special case.

The good-phase reflection changes the state to amplitudes square root of three over two and negative one half. The reflection about the original prepared state has matrix

$$R_\psi = \begin{pmatrix} 1/2 & \sqrt{3}/2 \\ \sqrt{3}/2 & -1/2 \end{pmatrix}.$$

Multiplying that matrix by the phase-flipped column gives zero in the first component and one in the second. The cancellation is exact. This example uses a biased preparation, so the uniform-state diffusion shortcut from basic Grover would be the wrong reflection. General amplification must reflect about the actual prepared state.

The matrix multiplication verifies the exact cancellation promised by the geometry.

$$\begin{pmatrix} 1/2 & \sqrt{3}/2 \\ \sqrt{3}/2 & -1/2 \end{pmatrix} \begin{pmatrix} \sqrt{3}/2 \\ -1/2 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \qquad p = \frac{1}{4}, \quad \theta = \frac{\pi}{6}.$$

The column on the left is the prepared state after the good sign flip. One amplification iterate produces $|1\rangle$ exactly for this chosen preparation.

Checkpoint 1

PAUSE & THINK

For p=1/4, what is the success angle after one basic amplification iterate?

- ☐ π/6
- ☐ π/2
- ☐ π

Answer: π/2. θ=π/6, so (2·1+1)θ=π/2. This particular preparation reaches success probability one.

Correctness through a two-dimensional plane

Normalise the good projection of A's output to define one direction and its bad projection to define another. The prepared state lies in their span. Let sine squared theta equal p. The good phase test and reflection about the preparation preserve this plane. Their composition rotates within it, giving success after k steps equal to sine squared of (2k plus one) theta.

For small p, reaching the first large-success region takes on the order of one over the square root of p amplification steps. Ordinary independent prepare-measure-check repetition needs on the order of one over p trials on average. The comparison assumes access to the coherent preparation, inverse and recognition reflection and counts those resources consistently.

Two reflections give a rotation in the bad/good plane.

$$Q^k|\psi\rangle = \cos((2k+1)\theta)|B\rangle + \sin((2k+1)\theta)|G\rangle, \qquad P_{\text{good}}(k) = \sin^2((2k+1)\theta).$$

This is the generalisation of the Grover formula to the state prepared by A.

Known and unknown success fractions

If p is known, choose an appropriate iteration count near a peak. Integer rounding and desired certainty matter; a basic fixed-phase iterate does not hit exact success for every p. Modified-phase methods can improve exactness or robustness under stated assumptions. If p is unknown, a single guessed iteration count may overshoot badly.

Randomised schedules, estimation-assisted approaches and fixed-point constructions address different versions of that problem. They have their own costs and guarantees. “Amplitude amplification” names a family of techniques, not a rule that repeatedly boosts success monotonically for any unknown input. State which variant is being used and what information about p it requires.

When it helps a larger algorithm

Suppose a coherent subroutine produces a useful answer rarely but can recognise success. Amplification can reduce the number of coherent attempts needed to obtain one good result. This pattern appears in search, state preparation and other probabilistic quantum subroutines. The improvement is relative to that subroutine’s initial success probability and access assumptions.

It does not eliminate expensive preparation. If A itself has a large data-loading or arithmetic cost, each amplification step pays for A and its inverse. If recognising a good result is as hard as solving the original problem, the phase test may be the bottleneck. An end-to-end analysis multiplies the iteration count by these costs and includes output verification.

Chapter summary

- Amplification works with a chosen coherent preparation, including biased ones.
- Reflect about A|0⟩ by applying A†, a zero-state reflection and A.
- The angle formula determines both the gain and the possibility of overshooting.

A common misconception

A collection of ordinary samples supplies everything amplitude amplification needs. The reflection construction requires coherent unitary preparation, its inverse and a coherent good-state test. Samples alone do not supply these operations.

Exercises

1. A coherent procedure succeeds with probability one in ten thousand. Compare the scaling of ordinary repetition and amplitude amplification, ignoring constants but counting a use of A as a resource.
2. A preparation service returns a random candidate but reveals no circuit and provides no inverse. Can you automatically apply the reflection construction above? Explain the missing resource.

Worked answers

1. Ordinary repetition has an average scale of ten thousand preparations. Amplification has a scale of one hundred coherent iterations, each including preparation/inverse and reflection work. These are scaling comparisons rather than exact runtime predictions.
2. No. The construction needs coherent unitary preparation and inverse access, plus the recognition reflection. A black-box classical sampler supplies only samples. Turning its implementation into a coherent reversible procedure may be possible in a particular setting, but must be explicitly designed and charged for.

Check your understanding

PAUSE & THINK

Why is a classical sampling service alone insufficient for this reflection construction?

- [ ] Its samples cannot be random
- [ ] The method also needs coherent preparation and inverse access
- [ ] Amplification requires a uniform initial state

Answer: The method also needs coherent preparation and inverse access. Reflecting about $A|0\rangle$ uses $A^\dagger$, a zero-state reflection and A. Ordinary samples do not automatically provide these coherent operations.

Sources

Brassard, Høyer, Mosca and Tapp develop amplitude amplification and estimation. The biased one-qubit calculation above makes the two reflections explicit for a small preparation.