Linear Combination of Unitaries

Construct a linear combination of unitaries and account for its success branch.

Research | 22 min lesson | Dr. Charith Jayasekara

By the end, you'll be able to

- Describe PREPARE, SELECT and inverse PREPARE.
- Calculate an ancilla-conditioned amplitude and probability.
- Include postselection or amplification in resource costs.

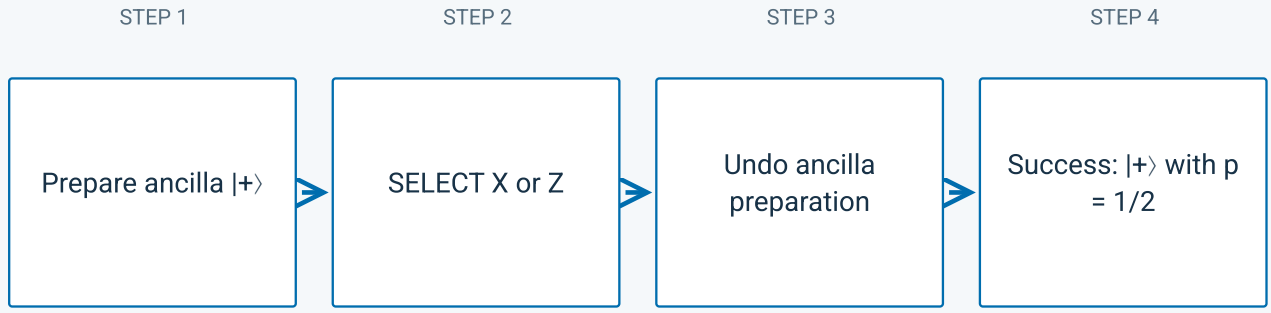
Add operations coherently with an ancilla

A linear combination of unitaries, or LCU, expresses an operator as a weighted sum of unitary operations. The sum need not itself be unitary, so it cannot generally be applied as an ordinary deterministic gate on the system alone. An ancilla records which operation is selected, and interference makes a desired combination appear in a particular ancilla branch.

Prerequisites: tensor products, controlled operations and amplitude normalisation. A primary simulation algorithm uses LCU to implement a truncated Taylor expansion of time evolution. Berry and colleagues.

WORKED DIAGRAM · A24

LCU of X and Z



For A = (X+Z)/2 acting on |0⟩, the success branch is (|0⟩+|1⟩)/2, whose squared norm is 1/2.

The access setup

Suppose an operator A is a sum of nonnegative coefficients times unitaries. Let alpha be the coefficient sum. A PREPARE routine creates an index superposition whose amplitudes are the square roots of the coefficients divided by alpha. A SELECT routine applies the corresponding unitary to the system conditioned on that index. Undoing PREPARE creates a zero-ancilla block equal to A divided by alpha.

Negative or complex coefficients can be absorbed into the selected unitaries' phases, leaving nonnegative weights. This is an algebraic convention, not an instruction to discard signs. The cost includes implementing the index preparation, the controlled selection and their inverses. A long list of cheap terms can still require a nontrivial coherent addressing circuit.

Worked example: average X and Z

Let A equal one half X plus one half Z. Use one ancilla prepared in plus. SELECT applies X when the ancilla is zero and Z when it is one. Apply H to the ancilla again. Starting with system zero, the joint state before the final H is a superposition of ancilla zero with system one and ancilla one with system zero.

After the final H, the zero-ancilla branch is

$$\frac{|1\rangle + |0\rangle}{2} = A|0\rangle.$$

Its squared norm is one half. Measuring the ancilla and obtaining zero therefore succeeds with probability one half, leaving the normalised system in plus. The other ancilla outcome is a different branch and must not be silently ignored when reporting cost.

In this example A equals H divided by square root of two. Directly implementing H would be simpler because we recognise that special structure. The LCU construction is valuable as a general technique, not as the best implementation of every small sum. The toy calculation makes the success amplitude and normalisation visible.

Why postselection needs accounting

If success probability is p, independent postselected attempts need roughly one over p repetitions on average. Amplitude amplification can improve the scaling under suitable coherent access to the preparation and its inverse. Those extra operations and reflections are part of the resource count. Saying "condition on ancilla zero" does not make that event deterministic or free.

Success probability can depend on the input state through the norm of A acting on that state. This is expected for a nonunitary target operation. A classical simulator that simply normalises A times the state has skipped the physical success mechanism. To model an actual protocol, keep both branches or explicitly charge the success probability.

Step-by-step construction

Choose a decomposition and calculate alpha. Build and test PREPARE on the index register. Implement SELECT with the declared ordering and phase conventions. Apply PREPARE, SELECT and inverse PREPARE to the combined input. Verify that the projected zero-index block matches A divided by alpha on every basis input. Then choose a postselection or coherent amplification strategy.

For the two-term example, testing both zero and one is enough to establish the system block by linearity. Include an imaginary superposition input as a numerical regression check. Comparing only success probabilities would miss a relative sign in the selected operations, so inspect the projected amplitudes in a small simulation.

Preparing, selecting and unpreparing the ancilla produces the weighted sum in its zero block.

$$A = \sum_j \alpha_j U_j, \quad \alpha_j \geq 0, \quad \alpha = \sum_j \alpha_j, \qquad (\langle 0| \otimes I)\,\mathrm{PREPARE}^\dagger \mathrm{SELECT}\,\mathrm{PREPARE}\,(|0\rangle \otimes I) = \frac{A}{\alpha}.$$

PREPARE creates amplitudes $\sqrt{\alpha_j/\alpha}$ on label j; SELECT applies U_j on that label. In the worked average of X and Z, the weights are each one half and $\alpha = 1$, so the success branch is $A|0\rangle$ with probability one half. Complex coefficient

phases can instead be included in the selected unitaries.

From sums to Hamiltonian simulation

The exponential of a Hamiltonian can be approximated by a polynomial or truncated series. Expanding products of its unitary terms creates a larger LCU. An ancilla-assisted implementation then realises the desired approximate operator, with an amplification construction arranged to recover the nearly unitary evolution. Segmentation and truncation choices control the error and cost.

The full simulation algorithm is more than postselecting a single arbitrary sum. Its analysis allocates error and uses the structure of a near-unitary target. Do not transfer a success claim from that specialised construction to an unrelated nonunitary matrix without checking the assumptions.

Compare two decompositions of the same operator

The worked average of X and Z has success probability one half on every normalised input because its square is one half times the identity. That is a special property of these anticommuting terms. For the average of I and Z, success depends strongly on the input: it is one on zero, zero on one and one half on plus.

This comparison shows why one successful test input cannot establish a universal success rate. Calculate the norm of the encoded operator acting on each relevant state or prove a state-independent relation. An algorithm may require a lower bound over an entire input family, not merely a convenient example.

A different decomposition can also change the coefficient sum without changing the target operator. Cancellations in a poor decomposition can make the normalisation unnecessarily large and reduce useful amplitude. Optimising the decomposition requires balancing that factor against the complexity of preparing coefficients and selecting terms coherently.

Complexity and limitations

A useful LCU decomposition balances the coefficient sum, number of terms, selection cost and preparation cost. A small alpha can improve amplitude scaling, while a complicated decomposition can make SELECT expensive. Term count alone is not a complete resource metric. The representation of the input operator is an essential assumption.

LCU is a building block for simulation and matrix-function algorithms. It does not mean arbitrary matrices become cheap gates. Large normalisation factors, low success amplitude, hard data access or expensive output extraction can dominate. Compare the full workflow with classical methods for the requested observable or sampling task.

Exercises and worked answers

Exercise 1. Replace A by the average of I and Z. What happens on zero and one?

<details>
<summary>Check the worked answer</summary>

A is the zero-state projector. On zero, success has probability one and the state remains zero. On one, the success branch has norm zero, so postselection on success never occurs. The input-dependent probability is part of the operation.
</details>

Exercise 2. Why cannot a simulator's automatic renormalisation replace the ancilla protocol?

Check the worked answer

Renormalisation describes the conditional state but omits the probability of obtaining that condition and the alternative outcomes. A physical implementation must account for those outcomes or use a justified coherent amplification method with its associated cost.

Check your understanding

PAUSE & THINK

What is the ancilla-zero success probability in the worked LCU example?

- [ ] 1
- [ ] 1/2
- [ ] 1/4

Answer: 1/2. The unnormalised system branch has two amplitudes 1/2, so its squared norm is 1/4 + 1/4 = 1/2.

Sources

- Berry and colleagues