Chapter 3: Qubits and Quantum States

In Chapter 2, we built the mathematical language of complex numbers, vectors, basis states, normalization, and inner products. Now we use that language to describe the basic object of quantum computing: the qubit.

A classical computer stores information in bits. A bit has one of two possible values:

$$0 \quad \text{or} \quad 1.$$

A quantum computer stores information in qubits. A qubit is not simply “a bit that can be both 0 and 1” in a vague way. More precisely, a qubit is described by a normalized vector in a two-dimensional complex vector space (Nielsen and Chuang, 2010). That sentence contains several ideas, so we will unpack it carefully.

By the end of this chapter, you should understand:

- what a single-qubit state is,
- what the notation $|0\rangle$ and $|1\rangle$ means,
- what superposition means mathematically,
- how amplitudes are related to probabilities,
- how multi-qubit states are written,
- why tensor products are needed,
- and how a quantum state over n qubits contains amplitudes for 2^n classical basis states.

This chapter is one of the most important foundations for Grover’s algorithm. Grover’s algorithm does not search by checking one candidate at a time. It prepares a quantum state whose amplitudes are spread over many candidates, then repeatedly changes those amplitudes so that marked answers become more likely to appear when measured.

To understand that, we first need to understand what quantum states are.

3.1 From classical bits to qubits

A classical bit has two possible states:

$$0$$

and

$$1.$$

In quantum computing, we represent the two basic states of a qubit using ket notation:

$$|0\rangle \quad \text{and} \quad |1\rangle.$$

The symbol $|0\rangle$ is read “ket zero,” and $|1\rangle$ is read “ket one.” Ket notation is part of Dirac notation, the standard notation used for vectors in quantum mechanics and quantum computing (Nielsen and Chuang, 2010).

For a single qubit, we may identify these basis states with column vectors:

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \qquad |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

These are the quantum versions of the two classical bit values. If a qubit is in state $|0\rangle$, then measuring it in the standard computational basis gives outcome 0 with probability 1. If it is in state $|1\rangle$, measuring it gives outcome 1 with probability 1.

The phrase computational basis means the standard basis used to label ordinary classical bit strings in a quantum computer. For one qubit, the computational basis is

$$\{|0\rangle, |1\rangle\}.$$

For two qubits, it will be

$$\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}.$$

For n qubits, it will contain all 2^n binary strings of length n.

This basis is called “computational” because it is the basis in which classical information is usually encoded and read out.

3.2 A single-qubit quantum state

A general single-qubit state is a linear combination of |0⟩ and |1⟩:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle.$$

Here:

- |ψ⟩ is the name of the quantum state,
- α is the amplitude of |0⟩,
- β is the amplitude of |1⟩,
- α and β are complex numbers.

The Greek letter ψ, pronounced “psi,” is commonly used for quantum states.

Because |0⟩ and |1⟩ are basis vectors, the expression

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

means the same thing as the vector

$$|\psi\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix}.$$

For example,

$$|\psi\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$$

corresponds to

$$\begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}.$$

Another possible state is

$$|\phi\rangle = \frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle,$$

which corresponds to

$$\begin{bmatrix} 1/\sqrt{2} \\ -1/\sqrt{2} \end{bmatrix}.$$

These two states have the same probability of giving 0 or 1 if measured immediately in the computational basis. But they are not the same state, because their amplitudes have different signs. Later, this difference in sign will matter greatly. Grover's algorithm uses phase changes—especially sign flips—to make amplitudes interfere constructively or destructively.

3.3 Normalization: why amplitudes cannot be arbitrary

A quantum state must be normalized. For a single qubit,

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

is a valid quantum state only if

$$|\alpha|^2 + |\beta|^2 = 1.$$

Here |α| means the magnitude of the complex number α, and |α|² means its squared magnitude. The rule that probabilities come from squared magnitudes of amplitudes is the standard Born rule in quantum mechanics and quantum computation (Nielsen and Chuang, 2010).

Why must the squared magnitudes add to 1? Because they will become probabilities, and the total probability of all possible measurement outcomes must be 1.

For example,

$$|\psi\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$$

is normalized because

$$\left|\frac{1}{\sqrt{2}}\right|^2 + \left|\frac{1}{\sqrt{2}}\right|^2 = \frac{1}{2} + \frac{1}{2} = 1.$$

But

$$|v\rangle = |0\rangle + |1\rangle$$

is not normalized, because

$$|1|^2 + |1|^2 = 2.$$

The vector $|0\rangle + |1\rangle$ points in the same direction as the normalized state

$$\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle,$$

but it is too long to be a valid quantum state.

Normalization is not a minor technical detail. In quantum computing, valid states are unit vectors. Quantum gates will preserve this unit length, and measurement probabilities will always sum to 1.

3.4 Superposition

A superposition is a linear combination of basis states.

For a single qubit,

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

is a superposition of $|0\rangle$ and $|1\rangle$ when both amplitudes are nonzero.

For example,

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

is a superposition. This state is often called “plus.”

Similarly,

$$|-\rangle = \frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle$$

is also a superposition. This state is called “minus.”

It is tempting to say that a qubit in state $|+\rangle$ is “both 0 and 1 at the same time.” That phrase can be useful as a first intuition, but it can also mislead. A more precise statement is:

> A qubit in superposition has amplitudes assigned to multiple basis states.

If we measure $|+\rangle$ in the computational basis, we get 0 with probability

$$\left|\frac{1}{\sqrt{2}}\right|^2 = \frac{1}{2},$$

and 1 with probability

$$\left|\frac{1}{\sqrt{2}}\right|^2 = \frac{1}{2}.$$

But before measurement, the state is not merely an ordinary random bit. The amplitudes can have signs and complex phases, and those phases can affect later computation through interference. This is one of the essential differences between quantum probability amplitudes and ordinary classical probabilities (Mermin, 2007).

For Grover’s algorithm, this is crucial. We will begin with a state that is a superposition over all candidate answers. Then the algorithm will change the signs and sizes of amplitudes so that correct candidates become more likely to appear when measured.

3.5 Amplitudes and probabilities

Let

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle.$$

If we measure this qubit in the computational basis, the probability of observing 0 is

$$P(0) = |\alpha|^2,$$

and the probability of observing 1 is

$$P(1) = |\beta|^2.$$

This is how amplitudes encode probabilities.

For example, consider

$$|\psi\rangle = \frac{\sqrt{3}}{2}|0\rangle + \frac{1}{2}|1\rangle.$$

The probability of measuring 0 is

$$\left|\frac{\sqrt{3}}{2}\right|^2 = \frac{3}{4}.$$

The probability of measuring 1 is

$$\left|\frac{1}{2}\right|^2 = \frac{1}{4}.$$

So this state is more likely to produce 0 than 1.

Now consider a state with a negative amplitude:

$$|\phi\rangle = \frac{\sqrt{3}}{2}|0\rangle - \frac{1}{2}|1\rangle.$$

The measurement probabilities are still

$$P(0) = \frac{3}{4}, \qquad P(1) = \frac{1}{4}.$$

The negative sign does not affect these probabilities directly, because

$$\left|-\frac{1}{2}\right|^2 = \frac{1}{4}.$$

However, the negative sign can affect future computation. If later operations combine amplitudes, positive and negative contributions can cancel or reinforce. This is called interference. Grover's algorithm depends on interference: it repeatedly arranges for amplitudes of marked states to grow while amplitudes of unmarked states shrink.

Complex phases work similarly. For example,

$$|\chi\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{i}{\sqrt{2}}|1\rangle$$

is a valid state because

$$\left|\frac{1}{\sqrt{2}}\right|^2 + \left|\frac{i}{\sqrt{2}}\right|^2 = \frac{1}{2} + \frac{1}{2} = 1.$$

Since |i| = 1, the amplitude i/√(2) has squared magnitude 1/2. Measuring in the computational basis gives 0 and 1 with equal probability. But the relative phase between the two amplitudes may matter in later gates.

A phase is the angular part of a complex number. For example, 1, -1, i, and -i all have magnitude 1, but they point in different directions in the complex plane. Quantum algorithms often work by changing phases and then converting those phase differences into probability differences.

3.6 Global phase and relative phase

Two quantum states can look different algebraically but represent the same physical state if they differ only by a global phase. A global phase is a complex number of magnitude 1 multiplying the entire state.

For example,

$$|\psi\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$$

and

$$|\psi'\rangle = -\frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle$$

differ by an overall factor of -1:

$$|\psi'\rangle = -|\psi\rangle.$$

This global sign does not change measurement probabilities. More generally, multiplying a quantum state by $e^{(i\theta)}$, where θ is a real angle, does not change observable measurement probabilities for that isolated state (Nielsen and Chuang, 2010).

But a relative phase does matter. Compare

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

with

$$|-\rangle = \frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle.$$

These two states are not related by multiplying the whole state by a single phase. Only the amplitude of $|1\rangle$ has changed sign relative to the amplitude of $|0\rangle$.

This relative sign can be detected by applying suitable quantum gates before measurement. In Chapter 5, we will meet the Hadamard gate, which transforms these states in different ways:

$$H|+\rangle = |0\rangle, \qquad H|-\rangle = |1\rangle.$$

So although |+⟩ and |-⟩ have the same immediate measurement probabilities in the computational basis, they are computationally different.

This distinction between global phase and relative phase is one of the quiet but powerful ideas behind Grover’s algorithm. The Grover oracle will mark solutions by changing their phase relative to nonsolutions.

3.7 Two-qubit states

A single qubit has two computational basis states:

$$|0\rangle,\ |1\rangle.$$

Two qubits have four computational basis states:

$$|00\rangle,\ |01\rangle,\ |10\rangle,\ |11\rangle.$$

These correspond to the four possible two-bit strings:

$$00,\ 01,\ 10,\ 11.$$

A general two-qubit state is a normalized linear combination of these four basis states:

$$|\psi\rangle = \alpha_{00}|00\rangle + \alpha_{01}|01\rangle + \alpha_{10}|10\rangle + \alpha_{11}|11\rangle.$$

The amplitudes

$$\alpha_{00},\ \alpha_{01},\ \alpha_{10},\ \alpha_{11}$$

are complex numbers, and they must satisfy

$$|\alpha_{00}|^2 + |\alpha_{01}|^2 + |\alpha_{10}|^2 + |\alpha_{11}|^2 = 1.$$

If we measure both qubits in the computational basis, then the probability of outcome 00 is

$$|\alpha_{00}|^2,$$

the probability of outcome 01 is

$$|\alpha_{01}|^2,$$

and so on.

For example,

$$|\psi\rangle = \frac{1}{2}|00\rangle + \frac{1}{2}|01\rangle + \frac{1}{2}|10\rangle + \frac{1}{2}|11\rangle$$

is normalized because

$$4\left(\frac{1}{2}\right)^2 = 1.$$

Each basis state has probability

$$\left|\frac{1}{2}\right|^2 = \frac{1}{4}.$$

So measuring this state gives each of the four two-bit strings with equal probability.

This is already the beginning of the pattern Grover’s algorithm uses. For n qubits, we can create a uniform superposition over 2^n possible bit strings. If those bit strings represent candidate answers, then the quantum state has an amplitude assigned to every candidate.

3.8 Tensor products: how qubits combine

We now need to answer a basic question:

> If one qubit is a vector with two components, how do we describe two qubits together?

The answer is the tensor product.

A tensor product is a mathematical operation that combines vector spaces. In quantum computing, the state space of a combined system is the tensor product of the state spaces of its parts (Nielsen and Chuang, 2010; Mermin, 2007).

For a first example, suppose the first qubit is in state

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

and the second qubit is in state

$$|1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

The combined two-qubit state is written

$$|0\rangle \otimes |1\rangle.$$

This is usually abbreviated as

$$|01\rangle.$$

The symbol ⊗ is read “tensor.”

To compute the tensor product of two column vectors, multiply every entry of the first vector by the entire second vector:

$$\begin{bmatrix} a \\ b \end{bmatrix} \otimes \begin{bmatrix} c \\ d \end{bmatrix} = \begin{bmatrix} ac \\ ad \\ bc \\ bd \end{bmatrix}.$$

So

$$|0\rangle \otimes |1\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \otimes \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \cdot 0 \\ 1 \cdot 1 \\ 0 \cdot 0 \\ 0 \cdot 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}.$$

This vector corresponds to $|01\rangle$, assuming the basis order

$$|00\rangle,\ |01\rangle,\ |10\rangle,\ |11\rangle.$$

Similarly,

$$|1\rangle \otimes |0\rangle = |10\rangle = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}.$$

The order matters:

$$|01\rangle \neq |10\rangle.$$

The first label belongs to the first qubit, and the second label belongs to the second qubit.

3.9 Tensor products of superpositions

Tensor products become especially important when qubits are in superposition.

Suppose one qubit is in state

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

If we have two qubits, each in state $|+\rangle$, the combined state is

$$|+\rangle \otimes |+\rangle.$$

Using distributive multiplication, we get

$$\begin{aligned} |+\rangle \otimes |+\rangle &= \left(\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle\right) \otimes \left(\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle\right) \\ &= \frac{1}{2}|00\rangle + \frac{1}{2}|01\rangle + \frac{1}{2}|10\rangle + \frac{1}{2}|11\rangle. \end{aligned}$$

This is the equal superposition over all two-bit strings.

Every basis state has amplitude 1/2, so every basis state has probability

$$\left|\frac{1}{2}\right|^2 = \frac{1}{4}.$$

This example generalizes. If we prepare n qubits all in state $|+\rangle$, then the combined state is an equal superposition over all 2^n bit strings:

$$|+\rangle^{\otimes n} = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle.$$

This notation deserves careful explanation.

The expression

$$|+\rangle^{\otimes n}$$

means

$$\underbrace{|+\rangle \otimes |+\rangle \otimes \cdots \otimes |+\rangle}_{n \text{ times}}.$$

The set

$$\{0,1\}^n$$

means all binary strings of length n. For example,

$$\{0,1\}^3 = \{000, 001, 010, 011, 100, 101, 110, 111\}.$$

The sum

$$\sum_{x \in \{0,1\}^n} |x\rangle$$

means “add one basis state |x⟩ for every n-bit string x.”

The coefficient

$$\frac{1}{\sqrt{2^n}}$$

is required for normalization. There are 2^n basis states, and each has squared amplitude

$$\left|\frac{1}{\sqrt{2^n}}\right|^2 = \frac{1}{2^n}.$$

So the total probability is

$$2^n \cdot \frac{1}{2^n} = 1.$$

This state will become central in Grover’s algorithm. If there are $N = 2^n$ possible candidates, then the uniform superposition gives every candidate the same starting amplitude:

$$\frac{1}{\sqrt{N}}.$$

Grover’s algorithm begins fairly: every candidate starts equally likely. Then the oracle and diffusion operator reshape the amplitudes.

3.10 Three qubits and the growth of the state space

For three qubits, the computational basis contains $2^3 = 8$ states:

$$|000\rangle,\ |001\rangle,\ |010\rangle,\ |011\rangle,\ |100\rangle,\ |101\rangle,\ |110\rangle,\ |111\rangle.$$

A general three-qubit state has the form

$$|\psi\rangle = \alpha_{000}|000\rangle+\alpha_{001}|001\rangle+\alpha_{010}|010\rangle+\alpha_{011}|011\rangle+\alpha_{100}|100\rangle+\alpha_{101}|101\rangle+\alpha_{110}|110$$

The normalization condition is

$$\sum_{x\in\{0,1\}^3} |\alpha_x|^2 = 1.$$

If we measure in the computational basis, the probability of seeing a particular bit string x is

$$|\alpha_x|^2.$$

For n qubits, the pattern is:

$$|\psi\rangle = \sum_{x\in\{0,1\}^n} \alpha_x|x\rangle,$$

with normalization

$$\sum_{x\in\{0,1\}^n} |\alpha_x|^2 = 1.$$

This is one reason quantum computing is mathematically rich. A state of n qubits is described by 2^n complex amplitudes. However, this does not mean we can directly read out all 2^n amplitudes as classical information. Measurement produces outcomes probabilistically, and obtaining full information about an unknown quantum state is much more limited than simply printing its amplitudes (Nielsen and Chuang, 2010).

This point is important enough to state plainly:

> A quantum state may involve amplitudes for exponentially many basis states, but measurement does not give us all those amplitudes.

Grover’s algorithm succeeds not by reading all amplitudes, but by changing them so that a useful answer becomes likely to appear when measured.

3.11 Product states and entangled states

Some multi-qubit states can be built as tensor products of individual qubit states. These are called product states.

For example,

$$|01\rangle = |0\rangle \otimes |1\rangle$$

is a product state.

The two-qubit equal superposition

$$\frac{1}{2}|00\rangle + \frac{1}{2}|01\rangle + \frac{1}{2}|10\rangle + \frac{1}{2}|11\rangle$$

is also a product state, because it equals

$$|+\rangle \otimes |+\rangle.$$

But not every two-qubit state can be written as a product of two single-qubit states. A state that cannot be written as a product state is called entangled. Entanglement is one of the central features distinguishing multi-particle quantum systems from classical systems (Nielsen and Chuang, 2010; Mermin, 2007).

A famous example is the Bell state

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle.$$

This state has amplitude only on $|00\rangle$ and $|11\rangle$. If measured in the computational basis, the two possible outcomes are:

$$00 \quad \text{with probability } \frac{1}{2},$$

and

$$11 \quad \text{with probability } \frac{1}{2}.$$

The outcomes of the two qubits are perfectly correlated: they are always the same.

Why is this not just two separate qubits each randomly being 0 or 1? Because the joint quantum state cannot be factored into

$$(\alpha|0\rangle + \beta|1\rangle) \otimes (\gamma|0\rangle + \delta|1\rangle).$$

To see this, expand a general product state:

$$(\alpha|0\rangle + \beta|1\rangle) \otimes (\gamma|0\rangle + \delta|1\rangle) = \alpha\gamma|00\rangle + \alpha\delta|01\rangle + \beta\gamma|10\rangle + \beta\delta|11\rangle.$$

For this to equal

$$\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle,$$

we would need

$$\alpha\gamma = \frac{1}{\sqrt{2}}, \qquad \beta\delta = \frac{1}{\sqrt{2}},$$

but also

$$\alpha\delta = 0, \qquad \beta\gamma = 0.$$

The first two equations require α,γ,β,δ all to be nonzero. But the last two equations require at least one factor in each product to be zero. These requirements cannot all be true at once. Therefore the Bell state is not a product state.

Grover's algorithm does not require us to begin with entanglement in the simple uniform superposition. The initial state

$$|+\rangle^{\otimes n}$$

is a product state. But as oracles and gates act on multiple qubits, entanglement can appear depending on the structure of the computation. For now, the most important lesson is that multi-qubit states are not merely lists of independent qubits. Their amplitudes describe the joint system as a whole.

3.12 Basis states as candidate answers

Grover's algorithm searches over candidates. In quantum computing, we usually encode those candidates as computational basis states.

Suppose we want to search over

$$N = 8$$

possible candidates. Since

$$8 = 2^3,$$

we can use three qubits. The candidates can be labeled by three-bit strings:

$$000,\ 001,\ 010,\ 011,\ 100,\ 101,\ 110,\ 111.$$

The quantum basis states are

$$|000\rangle,\ |001\rangle,\ |010\rangle,\ |011\rangle,\ |100\rangle,\ |101\rangle,\ |110\rangle,\ |111\rangle.$$

If the marked solution is candidate 101, then the basis state $|101\rangle$ represents the correct answer.

A general three-qubit state assigns an amplitude to every candidate:

$$|\psi\rangle = \sum_{x\in\{0,1\}^3} \alpha_x |x\rangle.$$

For example, the uniform starting state is

$$|\psi_0\rangle = \frac{1}{\sqrt{8}}\left(|000\rangle + |001\rangle + |010\rangle + |011\rangle + |100\rangle + |101\rangle + |110\rangle + |111\rangle\right).$$

In this state, every candidate has amplitude

$$\frac{1}{\sqrt{8}}.$$

Therefore every candidate has measurement probability

$$\left|\frac{1}{\sqrt{8}}\right|^2 = \frac{1}{8}.$$

At the start, the marked state $|101\rangle$ is no more likely than any other state. Grover’s algorithm will change that.

The oracle will flip the phase of the marked state:

$$|101\rangle \mapsto -|101\rangle,$$

while leaving the unmarked states unchanged. Then the diffusion operator will transform amplitudes so that the marked state’s amplitude grows. We will build these operations carefully in later chapters.

For now, it is enough to understand that a search space of N = 2^n candidates can be represented by n qubits, with each candidate corresponding to one computational basis state.

3.13 When N is not a power of two

Quantum registers naturally represent 2^n basis states when they contain n qubits. But real search spaces do not always have size exactly equal to a power of two.

For example, suppose there are

$$N = 10$$

candidates. We need enough qubits to represent at least 10 possibilities. Since

$$2^3 = 8$$

is too small, but

$$2^4 = 16$$

is large enough, we can use 4 qubits.

The computational basis then contains 16 states:

$$|0000\rangle,\ |0001\rangle, \ldots, |1111\rangle.$$

We can assign the first 10 basis states to real candidates and treat the remaining 6 states as invalid or unused.

For example:

$$0000 \to \text{candidate } 0,$$

$$0001 \to \text{candidate } 1,$$

and so on up to

$$1001 \to \text{candidate } 9.$$

The states

$$1010,\ 1011,\ 1100,\ 1101,\ 1110,\ 1111$$

do not correspond to valid candidates.

In practical algorithm design, the oracle must handle these extra states correctly. One common approach is to mark only valid solutions and never mark invalid states. Another approach is to build a state preparation routine that creates a superposition only over valid candidates, though that can be more difficult.

For most of our first study of Grover’s algorithm, we will assume

$$N = 2^n$$

because it makes the notation cleaner. Later, when thinking about real applications, we will return to the issue of encoding candidate spaces carefully.

3.14 Quantum registers

A quantum register is a collection of qubits treated as one system.

For example, a register of 3 qubits has computational basis states

$$|000\rangle,\ |001\rangle, \ldots, |111\rangle.$$

A register of n qubits has computational basis states

$$|x\rangle$$

where

$$x \in \{0,1\}^n.$$

Sometimes we interpret the bit string x as a binary number. For instance,

$$|101\rangle$$

can be interpreted as the binary representation of the number 5, because

$$101_2 = 5_{10}.$$

So a three-qubit register can represent numbers from 0 to 7:

$$|000\rangle = |0\rangle, \quad |001\rangle = |1\rangle, \quad |010\rangle = |2\rangle, \quad \ldots, \quad |111\rangle = |7\rangle.$$

This shorthand is common. Depending on context, the symbol $|5\rangle$ might mean the computational basis state whose binary label is 101. The number of qubits must be clear from context.

In Grover's algorithm, the main register usually stores the candidate answer. If the search space has $N = 2^n$ items, then the main register has n qubits.

Later, when we build oracles, we may also use extra qubits called ancilla qubits. An ancilla qubit is an additional working qubit used during a computation. Ancillas help compute intermediate values, store temporary information, and make operations reversible. We will study them carefully in Chapter 6 and Chapter 16.

3.15 The state vector viewpoint

It is helpful to think of a quantum state as a long column vector whose entries are amplitudes.

For one qubit:

$$|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$$

corresponds to

$$\begin{bmatrix} \alpha_0 \\ \alpha_1 \end{bmatrix}.$$

For two qubits:

$$|\psi\rangle = \alpha_{00}|00\rangle + \alpha_{01}|01\rangle + \alpha_{10}|10\rangle + \alpha_{11}|11\rangle$$

corresponds to

$$\begin{bmatrix} \alpha_{00} \\ \alpha_{01} \\ \alpha_{10} \\ \alpha_{11} \end{bmatrix}.$$

For three qubits, the state vector has 8 entries. For n qubits, it has 2^n entries.

This vector viewpoint will be very useful when we trace Grover's algorithm step by step. The oracle changes signs of certain entries. The diffusion operator changes amplitudes by reflecting them about their average. The full algorithm is a repeated transformation of this amplitude vector.

For example, suppose a two-qubit system begins in the uniform state

$$|\psi\rangle = \frac{1}{2}|00\rangle + \frac{1}{2}|01\rangle + \frac{1}{2}|10\rangle + \frac{1}{2}|11\rangle.$$

As a vector, this is

$$\begin{bmatrix} 1/2 \\ 1/2 \\ 1/2 \\ 1/2 \end{bmatrix}.$$

If an oracle marks $|10\rangle$ by flipping its phase, the state becomes

$$\frac{1}{2}|00\rangle + \frac{1}{2}|01\rangle - \frac{1}{2}|10\rangle + \frac{1}{2}|11\rangle,$$

or

$$\begin{bmatrix} 1/2 \\ 1/2 \\ -1/2 \\ 1/2 \end{bmatrix}.$$

Notice that the probabilities have not changed yet. Each squared magnitude is still 1/4. But the relative phase has changed. The diffusion operator will use that phase difference to alter the amplitudes.

This is the first glimpse of Grover's mechanism.

3.16 What a quantum state is not

Before moving on, we should correct three common misunderstandings.

First, a qubit is not simply a hidden classical bit whose value we do not know. If that were all it was, then amplitudes with negative or complex phases would not matter. But phases do matter because they affect interference.

Second, a superposition is not the same as having many classical computers running in parallel and then reading all their answers. Although an n-qubit state has amplitudes for 2^n basis states, measurement does not reveal all those amplitudes. Quantum algorithms must carefully arrange interference so that useful outcomes become likely (Nielsen and Chuang, 2010).

Third, the amplitude of a basis state is not its probability. The probability is the squared magnitude of the amplitude. This distinction matters because amplitudes can be negative or complex, while probabilities are real and nonnegative.

For example, the amplitudes

$$\frac{1}{2} \quad \text{and} \quad -\frac{1}{2}$$

give the same probability:

$$\frac{1}{4}.$$

But they can behave differently when combined with other

Document information

Chapter 3: Qubits and Quantum States

Project	Grover's Algorithm from First Principles
Document	Document 1.7
Author	mujirin
Verifier	Not verified
Downloaded	July 04, 2026 18:33 KST
Status	Working
Document link	https://theorytrace.com/projects/grovers-algorithm-from-first-principles/documents/chapter-3-qubits-and-quantum-states/