Quantum Computing 101, Part 1 – Hello Quantum World

Aaron Turner, March 2021

Introduction

This is the first part of a blog series on **quantum computing**, broadly derived from CERN's Practical introduction to quantum computing video series, Michael Nielson's Quantum computing for the very curious / determined, and the following (widely regarded as definitive) references:

- [Hidary] *Quantum Computing: An Applied Approach*
- [Nielsen & Chuang] *Quantum Computing and Quantum Information* [a.k.a. "Mike & Ike"]
- [Yanofsky & Mannucci] *Quantum Computing for Computer Scientists*

My objective is to keep the mathematics to an absolute minimum (albeit not quite zero), in order to engender an **intuitive understanding**. You can think it as a quantum computing cheat sheet.

What is quantum computing?

Quantum computing may be defined as computation (in the form of **circuits** and **algorithms**) that (unlike classical computation) specifically exploits the properties of **quantum mechanics**:

- quantum **superposition** (defined in this Part)
- quantum **entanglement** (defined in Part 2)
- quantum **interference** (defined in Part 4)

What quantum computing does *not* promise

Turing machines, the **lambda calculus**, and **general recursive functions** are three different ways of describing [a subset of the class of all possible mathematical functions known as] the set of **effectively computable** functions, and all three have been formally proven to be equivalent.

It is widely believed by mathematicians and computer scientists that these equivalent definitions accurately characterise the notion of **computability**, i.e. what may and may not be computed in finite time by a (classical) computer. This is known as the **Church-Turing thesis**. Any function that cannot in principle be computed in finite time by a Turing machine is called **uncomputable**.

It is **not** the case that some functions that are uncomputable on a classical computer suddenly become computable on a quantum computer. The set of functions that may be computed in finite time on a quantum computer is exactly the same as the set of functions that may be computed in finite time on a classical computer. There are **no benefits** to be gained in terms of computability.

What quantum computing *may* promise

Although this is not yet something that can be formally proved, there is strong evidence to suggest that quantum computers may be **much faster** than classical computers for at least some tasks.

This is known as **quantum supremacy**.

At the time of writing, two groups have claimed to have demonstrated quantum supremacy:

- in October 2020, Google claimed to have achieved in 200 seconds a quantum computation that would take 10,000 years on a classical supercomputer (IBM later questioned this claim, saying that it would only take 2.5 days, but that's still a thousand times faster…)

- in December 2020, a group at the University of Science and Technology of China claimed to have achieved in 200 seconds a quantum computation that would take 2.5 billion years on a classical supercomputer (for a **quantum advantage**, or overall speedup, of 10^{14})

Physical vs logical quantum computing

Quantum algorithms execute on **quantum circuits** (of which more later!) In order to write, and execute, quantum algorithms, all you need to know is how the quantum circuits work **logically** – it's not necessary to know how the **implementations** of these circuits work, as (logically) they are all equivalent. The **physical technologies** being used to implement quantum circuits include:

- neutral atom qubits
- nitrogen-vacancy centre-in-diamond qubits
- nuclear magnetic resonance qubits
- photonic qubits
- silicon-based spin qubits
- superconducting qubits
- topological qubits
- trapped ion qubits

Running your own quantum algorithms

In order to **execute** a quantum algorithm, you need access to a either a **quantum computer** or a **simulation** of one. The IBM Quantum Experience provides access to both – and it's free! The Quirk quantum circuit simulator and Qiskit open source SDK are also excellent free resources.

Later, we will be executing the **Hello World** of quantum computing on a **real** quantum computer!

Bits and qubits

Computations are **initialised** to a particular **state**, **evolve** through a **finite sequence** of states (specified by an **algorithm**), and then **terminate**; the **final** state is the **result** of the computation.

The instantaneous state of a computation is recorded in a classical computer by a number of **bits** (each being either 0 or 1), and in a quantum computer by a number of **quantum bits**, or **qubits**.

Mathematically, the state of a single qubit is represented by 2-element column **vector**

$$\begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix}$$

where α_0 and α_1 are **complex numbers** (called **probability amplitudes**, or just **amplitudes** for short), and $|\alpha_0|^2 + |\alpha_1|^2 = 1$, where $|\alpha|$ is the **modulus** of complex number α.

The constraint $|\alpha_0|^2 + |\alpha_1|^2 = 1$ is called the **normalisation condition**. An arbitrary pair of complex numbers α_0 and α_1 that do not satisfy this constraint do not represent a valid qubit state!

Dirac notation and basis states

We will use Dirac's **ket notation** $|\psi\rangle$ to denote the state of an arbitrary qubit.

In particular, we will use $|0\rangle$ as shorthand for

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

and $|1\rangle$ as shorthand for

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

The states $|0\rangle$ and $|1\rangle$ are called **basis states**; they are the quantum analogues of 0 and 1.

Superposition

Now that we have the two basis states $|0\rangle$ and $|1\rangle$, we can rewrite the state of a generic qubit as

$$
\begin{aligned}
|\psi\rangle &= \alpha_0|0\rangle + \alpha_1|1\rangle \\
&= \alpha_0 \begin{pmatrix}1\\0\end{pmatrix} + \alpha_1 \begin{pmatrix}0\\1\end{pmatrix} \\
&= \begin{pmatrix}\alpha_0\\0\end{pmatrix} + \begin{pmatrix}0\\\alpha_1\end{pmatrix} \\
&= \begin{pmatrix}\alpha_0\\\alpha_1\end{pmatrix}
\end{aligned}
$$

The way to interpret this is that the qubit state $|\psi\rangle$ is a **simultaneous mixture** of $|0\rangle$ and $|1\rangle$, i.e. it is effectively in **both** states $|0\rangle$ **and** $|1\rangle$ **at the same time** – this is known as a **superposition**.

Intuitively, as will become clear in the next section, the $\boldsymbol{\alpha_0}$ **amplitude** tells us how much the qubit $|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$ **favours** $|\mathbf{0}\rangle$, and the $\boldsymbol{\alpha_1}$ **amplitude** tells us how much it **favours** $|\mathbf{1}\rangle$.

Measurement

The **result** of a quantum computation is always a **classical value** (i.e. 0 or 1, **not** $|0\rangle$ or $|1\rangle$). The way we extract this result is to **measure** the value of a qubit. **Born's rule**, a law of quantum mechanics, implies that, when measured, a generic qubit $|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$ will yield (classical) 0 with **probability** $|\alpha_0|^2$ and (classical) 1 with probability $|\alpha_1|^2$. (Hence $|\alpha_0|^2 + |\alpha_1|^2$ **must** = 1.)

If a measurement yields 0, then the state of the qubit **immediately after being measured** will have **collapsed** to $|0\rangle$ (and any subsequent measurement will yield 0 with probability 1), otherwise it will have **collapsed** to $|1\rangle$ (and any subsequent measurement will yield 1 with probability 1). This is simply how the underlying quantum mechanics of the universe works! Note that **qubits cannot be copied** in order to circumvent this effect (the **no-cloning theorem**).

The following **quantum circuit** (hint: read these from **left to right**) initialises the quantum state (generally depicted as a single line) to $|0\rangle$, which it immediately measures via the **measurement gate** (depicted as a dial), which then yields a classical bit (generally depicted as a double line).

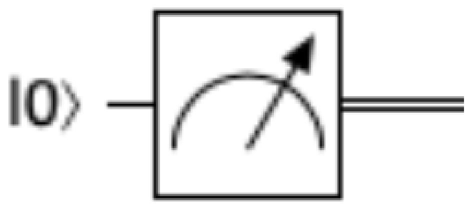

The Bloch sphere

The qubit state $|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$ may be visualised as **a point on the surface of a unit sphere**

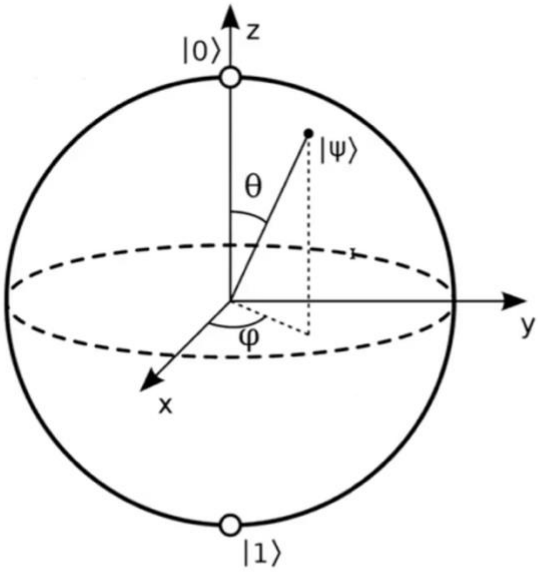

where:

- the basis state $|0\rangle$ is at the north pole of the sphere
- the basis state $|1\rangle$ is at the south pole of the sphere
- θ is the angle between the north-south axis and a line from the centre of the sphere to $|\psi\rangle$
- φ is the angle eastwards from an arbitrary point on the equator to the longitude of $|\psi\rangle$
- θ and φ ($0 \leq \theta \leq \pi$, $0 \leq \varphi < 2\pi$) may be calculated from α_0 and α_1

Note that **any basis states** such as $|0\rangle$ and $|1\rangle$ must be **orthogonal**, i.e. **opposite** each other.

Single-qubit operations (gates)

In order to manipulate the state of a qubit (in order that a quantum algorithm may evolve), we need some way to apply **operations** to qubit states. Because single-qubit states are 2-element **complex vectors**, we can represent a single-qubit operation (gate) U as a 2 × 2 **complex matrix**

$$U = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

We can then calculate the result of applying U to state $\alpha_0|0\rangle + \alpha_1|1\rangle$ via **matrix multiplication**

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} = \begin{pmatrix} a\alpha_0 + \mathrm{b}\alpha_1 \\ c\alpha_0 + \mathrm{d}\alpha_1 \end{pmatrix}$$

In order for the result on the RHS to be a **valid** state it must satisfy the normalisation condition $|a\alpha_0 + \mathrm{b}\alpha_1|^2 + |c\alpha_0 + \mathrm{d}\alpha_1|^2 = 1$. In order to **preserve normalisation**, a single-qubit operation

$$U = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

must be **unitary** (length preserving), meaning that U multiplied by its **conjugate transpose** $U^\dagger$ yields the identity matrix (here, the order of matrix multiplication doesn't matter). In other words

$$UU^\dagger = \begin{pmatrix} a & b \\ c & d \end{pmatrix}\begin{pmatrix} \bar{a} & \bar{c} \\ \bar{b} & \bar{d} \end{pmatrix} = U^\dagger U = \begin{pmatrix} \bar{a} & \bar{c} \\ \bar{b} & \bar{d} \end{pmatrix}\begin{pmatrix} a & b \\ c & d \end{pmatrix} = I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

where the **complex conjugate** $\bar{\alpha}$ of complex number $\alpha = x + iy$ is defined as $\bar{\alpha} = x - iy$.

The I (or Identity) gate

The I gate, defined as

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

leaves its input unchanged.

If $|\psi\rangle = \begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} = \alpha_0|0\rangle + \alpha_1|1\rangle$ then

$$I\,|\psi\rangle = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} = \begin{pmatrix} \alpha_0 \\ \alpha_1 \end{pmatrix} = |\psi\rangle$$

$$I\,|0\rangle = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} = |0\rangle$$

$$I\,|1\rangle = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} = |1\rangle$$

The X (or NOT) gate (a.k.a. Pauli gate σ_X)

The X gate, defined as

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

swaps the α_0 and α_1 amplitudes.

If $|\psi\rangle = \begin{pmatrix}\alpha_0\\ \alpha_1\end{pmatrix} = \alpha_0|0\rangle + \alpha_1|1\rangle$ then

$$X\,|\psi\rangle = \begin{pmatrix}0 & 1\\ 1 & 0\end{pmatrix}\begin{pmatrix}\alpha_0\\ \alpha_1\end{pmatrix} = \begin{pmatrix}\alpha_1\\ \alpha_0\end{pmatrix} = \alpha_1|0\rangle + \alpha_0|1\rangle$$

$$X\,|0\rangle = \begin{pmatrix}0 & 1\\ 1 & 0\end{pmatrix}\begin{pmatrix}1\\ 0\end{pmatrix} = \begin{pmatrix}0\\ 1\end{pmatrix} = |1\rangle$$

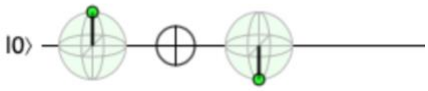

This quantum circuit starts by initialising the state to $|0\rangle$, as indicated by the first Bloch sphere, after which an X gate (denoted by the circle with the cross in it) is applied, resulting in the output indicated by the second Bloch sphere

$$X\,|1\rangle = \begin{pmatrix}0 & 1\\ 1 & 0\end{pmatrix}\begin{pmatrix}0\\ 1\end{pmatrix} = \begin{pmatrix}1\\ 0\end{pmatrix} = |0\rangle$$

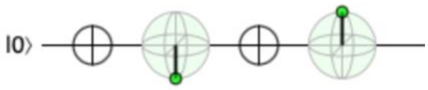

Here we use an X gate to flip the initial $|0\rangle$ into a $|1\rangle$, prior to the Bloch-X-Bloch sequence as above

The H (or Hadamard) gate

The H gate, defined as

$$H = \frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}$$

creates a **superposed** (i.e. **mixed** $\{|0\rangle, |1\rangle\}$) state from a **non-superposed** basis state.

If $|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$ then [note the two new definitions, $|+\rangle$ and $|-\rangle$]

$$H\,|\psi\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}\begin{pmatrix}\alpha_0\\ \alpha_1\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}\alpha_0 + \alpha_1\\ \alpha_0 - \alpha_1\end{pmatrix} = \frac{1}{\sqrt{2}}(\alpha_0 + \alpha_1)\,|0\rangle + \frac{1}{\sqrt{2}}(\alpha_0 - \alpha_1)\,|1\rangle$$

$$H\,|0\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}\begin{pmatrix}1\\ 0\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}1\\ 1\end{pmatrix} = \frac{\begin{pmatrix}1\\ 0\end{pmatrix}+\begin{pmatrix}0\\ 1\end{pmatrix}}{\sqrt{2}} = \frac{|0\rangle + |1\rangle}{\sqrt{2}} = |+\rangle$$

This quantum circuit starts by initialising the state to $|0\rangle$, as indicated by the first Bloch sphere, after which a Hadamard gate (denoted by the square with an H in it) is applied, resulting in the output indicated by the second Bloch sphere

$$H\,|1\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}\begin{pmatrix}0\\ 1\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}1\\ -1\end{pmatrix} = \frac{\begin{pmatrix}1\\ 0\end{pmatrix}-\begin{pmatrix}0\\ 1\end{pmatrix}}{\sqrt{2}} = \frac{|0\rangle - |1\rangle}{\sqrt{2}} = |-\rangle$$

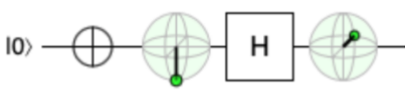

Here we use an X gate to flip the initial $|0\rangle$ into a $|1\rangle$, prior to the Bloch-H-Bloch sequence as above

Note that $|+\rangle$ and $|-\rangle$ are **orthogonal** (Bloch opposites), and therefore candidate **basis states**.

Hello World

The following circuit (the “Hello World” of quantum computing)

pushes the initial $|0\rangle$ through a Hadamard gate, thereby creating a superposition ($|+\rangle$) of the basis states $|0\rangle$ and $|1\rangle$, which it then measures. The probability of the measurement yielding 0 is $\left|\frac{1}{\sqrt{2}}\right|^2$, and the probability of the measurement yielding 1 is also $\left|\frac{1}{\sqrt{2}}\right|^2$, which evaluates to 0.5. In other words, this quantum circuit, executed repeatedly, generates a perfectly random bitstream.

In general (although not always), we don’t just execute a quantum algorithm once, as we would on a classical computer. Instead, we execute it many times in order to obtain a probabilistic result.

Executing this algorithm N = 1024 times on a **real** quantum computer returned the following:

What we effectively get is a histogram depicting the **probability distribution** over basis states represented by the superposed state. It’s not perfectly 50:50 between 0 and 1, for two reasons:

1. it is **literally** a random process, and will only **tend towards** the ideal result as N increases
2. **actual** quantum machines are **not perfect** – in particular, qubits are vulnerable to quantum errors due to decoherence (interactions with their environment) and other quantum noise.

Phase

The concept of **phase** is of fundamental importance in quantum computation.

A complex number α expressed in **rectangular form** (also called **Cartesian form**)

$$\alpha = x + iy$$

may be re-expressed (via basic **trigonometry**) in **polar form**

$$\alpha = r\cos\varphi + i\,r\sin\varphi = r(\cos\varphi + i\sin\varphi)$$

where

$$r = |\alpha| = \sqrt{x^2 + y^2}$$

φ is the positive angle from the real axis on the corresponding **Argand diagram**

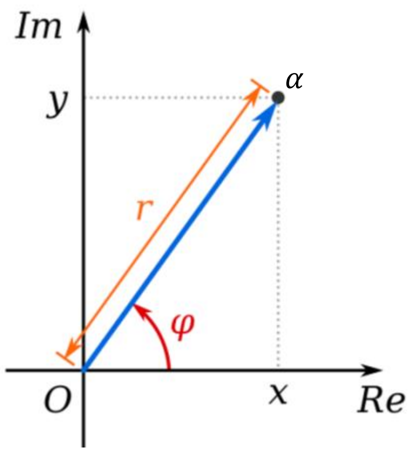

φ = **Arg(α)**, calculated as follows

$$Arg(\alpha) = \begin{cases} \tan^{-1}\left(\frac{y}{x}\right), & \text{when x > 0} \quad \text{[top-right and bottom-right quadrants]} \\ \frac{\pi}{2}, & \text{when x = 0, y > 0} \quad \text{[degenerate case, zero denominator]} \\ 0 \text{ [by convention]}, & \text{when x = 0, y = 0} \quad \text{[degenerate case, zero denominator]} \\ -\frac{\pi}{2}, & \text{when x = 0, y < 0} \quad \text{[degenerate case, zero denominator]} \\ \tan^{-1}\left(\frac{y}{x}\right) + \pi, & \text{when x < 0, y} \geq \text{0} \quad \text{[top-left quadrant]} \\ \tan^{-1}\left(\frac{y}{x}\right) - \pi, & \text{when x < 0, y < 0} \quad \text{[bottom-left quadrant]} \end{cases}$$

Calculated as above, φ is always in the **interval** $(-\pi, \pi]$. When (and only when!) φ (as calculated above) is negative, adding 2π will move it into the equivalent interval $[0, 2\pi)$, if this is preferred.

Thanks to **Euler's formula** ($e^{i\varphi} = \cos\varphi + i\sin\varphi$), α may be re-expressed in **exponential form**

$$\alpha = re^{i\varphi}$$

(commonly called a **phasor** in engineering), where

$e^{i\varphi}$ is called the **phase factor**

φ (always expressed in **radians**) is called the **phase** [note: $d° = \frac{d\pi}{180}$ rad].

Properties of $\bar{\alpha}$, $|\alpha|$, and $e^{i\pi}$

The following properties will prove useful when thinking about phase:

(1) for a complex number α in exponential form

$$\alpha = re^{i\varphi}$$

its complex conjugate $\bar{\alpha}$ is simply

$$\bar{\alpha} = re^{-i\varphi}$$

(2) complex conjugation distributes over multiplication

$$\overline{\alpha\beta} = \bar{\alpha}\bar{\beta}$$

(3) the modulus $|\alpha|$ and complex conjugate $\bar{\alpha}$ of complex number α are related as follows

$$|\alpha|^2 = \alpha\bar{\alpha}$$

(4) the special case of Euler's formula with $\varphi = \pi$ yields **Euler's identity**

$$e^{i\pi} = \cos\pi + i\sin\pi = -1$$

(5) the special case of Euler's formula with $\varphi = \frac{\pi}{2}$ yields

$$e^{i\frac{\pi}{2}} = \cos\frac{\pi}{2} + i\sin\frac{\pi}{2} = i$$

Global phase

Consider a generic quantum state

$|\psi_0\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle$

If we were to measure $|\psi_0\rangle$

we would obtain 0 with probability $|\alpha_0|^2$

we would obtain 1 with probability $|\alpha_1|^2$

Now imagine a second quantum state equal to $|\psi_0\rangle$ multiplied by a **global phase factor** $e^{i\varphi}$

$|\psi_1\rangle = e^{i\varphi}|\psi_0\rangle = e^{i\varphi}\alpha_0|0\rangle + e^{i\varphi}\alpha_1|1\rangle$

If we were to measure $|\psi_1\rangle$

we would obtain 0 with probability $\left|e^{i\varphi}\alpha_0\right|^2$

$$= \left(e^{i\varphi}\alpha_0\right)\overline{\left(e^{i\varphi}\alpha_0\right)}$$

$$= \left(e^{i\varphi}\alpha_0\right)\left(\overline{e^{i\varphi}}\,\overline{\alpha_0}\right)$$

$$= \left(e^{i\varphi}\alpha_0\right)\left(e^{-i\varphi}\overline{\alpha_0}\right)$$

$$= e^{(i\varphi-i\varphi)}\alpha_0\overline{\alpha_0}$$

$$= e^{0}|\alpha_0|^2$$

$$= |\alpha_0|^2$$

similarly, we would obtain 1 with probability $|\alpha_1|^2$

In other words, the measurement probabilities for $|\psi_0\rangle$ and $|\psi_1\rangle$ are **identical**. Which means that, despite the fact that these states differ **internally**, they are nevertheless **indistinguishable** from the perspective of an **external observer**. For this reason, we say that $|\psi_0\rangle$ and $|\psi_1\rangle$ are **equal up to the global phase factor $e^{i\varphi}$**. For all practical purposes, $|\psi_0\rangle$ and $|\psi_1\rangle$ are **equivalent**

$|\psi_0\rangle \equiv |\psi_1\rangle$.

In short, **global phase doesn't matter!**

Relative phase

We say that two **amplitudes** α and β differ by a **relative phase factor** $e^{i\varphi}$ if there exists some real number φ such that $\alpha = e^{i\varphi}\beta$. (Note that if φ = 0 then the relative phase factor is $e^{i0} = 1$.)

More generally, we say that two **quantum states**

$$|\psi_0\rangle = \alpha_{00}|0\rangle + \alpha_{01}|1\rangle$$

and

$$|\psi_1\rangle = \alpha_{10}|0\rangle + \alpha_{11}|1\rangle$$

differ by a relative phase if

(1) the $|0\rangle$ amplitudes α_{00} and α_{10} differ by some relative phase factor $e^{i\varphi_0}$

(2) the $|1\rangle$ amplitudes α_{01} and α_{11} differ by some relative phase factor $e^{i\varphi_1}$

(3) $e^{i\varphi_0} \neq e^{i\varphi_1}$

(Note that if $e^{i\varphi_0} = e^{i\varphi_1}$ then this would be equivalent to a **global** phase difference!)

For example, the quantum states

$$H\,|0\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}1\\1\end{pmatrix} = \frac{|0\rangle + |1\rangle}{\sqrt{2}} = |+\rangle$$

and

$$H\,|1\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}1\\-1\end{pmatrix} = \frac{|0\rangle - |1\rangle}{\sqrt{2}} = |-\rangle$$

differ by a relative phase because (1) the $|0\rangle$ amplitudes differ by a relative phase factor of $1 = e^{i0}$, (2) the $|1\rangle$ amplitudes differ by a relative phase factor of $-1 = e^{i\pi}$, (3) and $1 \neq -1$.

At first glance, it might appear that $|+\rangle$ and $|-\rangle$ are indistinguishable by an external observer, as was the case with global phase. However, when the phase difference between two states is **relative**, not **global**, it is always possible to distinguish between them simply by **measuring in a different basis** (in this case, by putting a Hadamard gate before the measurement dial). As a result, unlike global phase shifts, **relative phase shifts may be used to encode information**.

In short, **relative phase matters!**

