Lecture 4: Quantum Gates and Circuits

Reversibility, Logic, and Protocols

Moody T. Chu

North Carolina State University

Spring 2026

Lecture Overview

1. Fundamentals: Bits vs. Qubits
2. Thermodynamics of Computation: Landauer's Principle
3. Reversible Logic: XOR, Toffoli, and Ancilla Bits
4. Simulating Classical Logic: Reversible AND and NAND
5. Single-Qubit Operations and Pauli Matrices
6. The Hadamard Transformation and Initialization
7. Multi-Qubit Systems and Entanglement
8. Protocol 1: Superdense Coding
9. Protocol 2: Quantum Teleportation
10. Universality and the Solovay-Kitaev Theorem

The Unit of Information: Classical Bit

- A classical bit exists in a discrete state $c \in \{0, 1\}$.
  - ◇ Bits cannot exist in a state between 0 and 1.
- Physically represented by voltages, currents, or magnetization.
- Logic operations are functions $f : \{0, 1\}^n \to \{0, 1\}^m$.
- Most classical gates are **irreversible** (many-to-one).
  - ◇ Examples: AND, OR, XOR (when bits are discarded) — given only the output, the input cannot be recovered.

The Unit of Information: Qubit

- A qubit is a unit vector $|\psi\rangle$ in a 2D complex Hilbert space $\mathbb{C}^2$.
- General state: $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$ with $|\alpha|^2 + |\beta|^2 = 1$.
- Computation involves moving this vector via **unitary transformations** — and every unitary transformation is, by definition, **invertible** ($U^{-1} = U^\dagger$).
  - This is the central theme of today's lecture: quantum logic gates are always reversible, unlike most classical gates.

Irreversibility and Information Loss

- In an AND gate, if the output is 0, we cannot reconstruct the input: $(0,0)$, $(0,1)$, and $(1,0)$ all map to output 0.
- This loss of information is not just a mathematical curiosity — it has physical consequences.
- Erasing information reduces the number of possible states of the computer, which (as we will see) has a thermodynamic cost.

Landauer's Principle

- Erasing one bit of information increases the entropy of the environment by at least $k_B \ln 2$.
- This results in a minimum heat dissipation of $\Delta Q = k_B T \ln 2$ per bit erased.
- For modern processors, this heat generation is a fundamental scaling limit.
- **Implication:** if computation never erases information (i.e., is *reversible*), this thermodynamic floor can in principle be avoided.

The Case for Reversibility

- If computation is reversible, no information is erased.
- Theoretically, a perfectly reversible computer could operate with zero energy dissipation (Landauer's bound is not triggered).
- Quantum mechanics is *inherently* reversible: unitary evolution is always invertible (Lecture 3, Postulate 2).
- **Consequence:** any classical computation we want to run on a quantum computer must first be recast in a reversible form.

Implementing Classical Functions Reversibly

- Given any classical function $f : \{0,1\}^n \to \{0,1\}^m$ (possibly irreversible), we build a reversible unitary U_f by *keeping the input* and XOR-ing the output into a fresh register:

$$U_f : |x\rangle\,|y\rangle \;\longmapsto\; |x\rangle\,|y \oplus f(x)\rangle\,.$$

- By preserving the input register $|x\rangle$, the map becomes one-to-one (injective), hence invertible.
- **Verification that U_f is its own inverse** (when y starts at a fixed ancilla):

$$\begin{aligned} U_f\big(U_f\,|x\rangle\,|y\rangle\big) &= U_f\,|x\rangle\,|y \oplus f(x)\rangle \\ &= |x\rangle\,|(y \oplus f(x)) \oplus f(x)\rangle = |x\rangle\,|y\rangle\,, \end{aligned}$$

  using $a \oplus a = 0$ for any bit a.
- We now build U_f concretely for the basic logic gates XOR, AND, and NAND.

Reversible XOR (CNOT)

- Standard XOR:
  - A logical operation that outputs true (1) only when the inputs differ.
  - $(x, y) \to x \oplus y$ is irreversible (both inputs are lost; only the single output bit remains).
- To make it reversible, we keep the control bit and apply U_f with $f(x) = x$: $(x, y) \to (x, x \oplus y)$. This is the **CNOT** (controlled-NOT) gate.

x	y	x'	y'
0	0	0	0
0	1	0	1
1	0	1	1
1	1	1	0

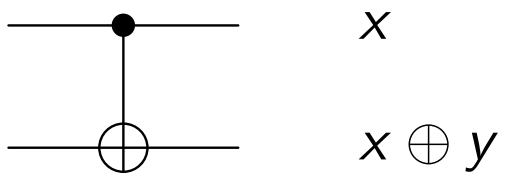

The CNOT Gate: Matrix Representation

- CNOT is an operator acting on two qubits.
  - ◇ It is a 4×4 unitary matrix acting on $\mathbb{C}^2 \otimes \mathbb{C}^2$, with rows/columns ordered $|00\rangle, |01\rangle, |10\rangle, |11\rangle$:

$$CNOT = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

- CNOT **cannot** be written as a tensor product of two 2×2 matrices.
  - ◇ It inherently connects/entangles the two subsystems (as we will see explicitly in Section 7).

The Toffoli Gate (CCNOT)

- A 3-bit reversible gate: $(x, y, z) \to (x, y, z \oplus xy)$.
- The target bit z is flipped if and only if both controls x and y are 1.

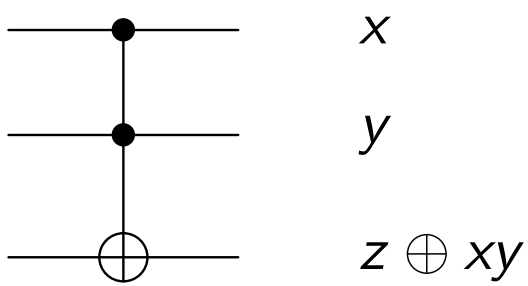

- The Toffoli gate is **universal** for classical reversible computation: any Boolean function can be implemented using only Toffoli gates and ancilla bits (proved in the next two slides).
- It can be constructed from **6 CNOT gates** together with single-qubit H and T, $T^\dagger$ gates — a standard textbook decomposition.

Toffoli Truth Table

Before		After	
Control	Target	Control	Target
$\|00\rangle$	$\|0\rangle$	$\|00\rangle$	$\|0\rangle$
$\|00\rangle$	$\|1\rangle$	$\|00\rangle$	$\|1\rangle$
$\|01\rangle$	$\|0\rangle$	$\|01\rangle$	$\|0\rangle$
$\|01\rangle$	$\|1\rangle$	$\|01\rangle$	$\|1\rangle$
$\|10\rangle$	$\|0\rangle$	$\|10\rangle$	$\|0\rangle$
$\|10\rangle$	$\|1\rangle$	$\|10\rangle$	$\|1\rangle$
$\|11\rangle$	$\|0\rangle$	$\|11\rangle$	$\|1\rangle$
$\|11\rangle$	$\|1\rangle$	$\|11\rangle$	$\|0\rangle$

- Every row is unchanged except the last two, where both controls are $|1\rangle$ and the target flips — exactly $z \oplus xy$.

Reversible AND Logic

- We can simulate an AND gate using a Toffoli gate by setting the third input (**ancilla**) to $|0\rangle$.
- Result: $z' = 0 \oplus xy = xy$ — exactly the AND function, with x, y preserved for reversibility.

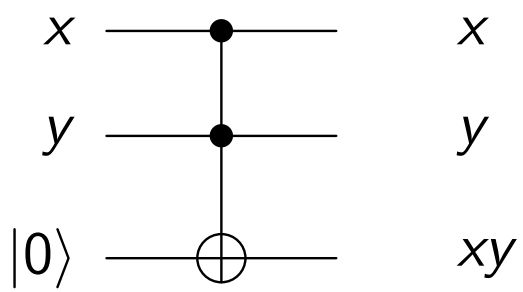

Reversible NAND Logic

- A NAND (Not-AND) operation is a fundamental digital logic gate that acts as the inverse of an AND gate.
- Setting the ancilla to $|1\rangle$ instead produces the NAND operation.
- Result: $z' = 1 \oplus xy = \text{NAND}(x, y)$.
- Since NAND is classically universal (any Boolean circuit can be built from NAND gates alone), **this proves the Toffoli gate is universal for classical logic** — completing the claim from two slides ago.

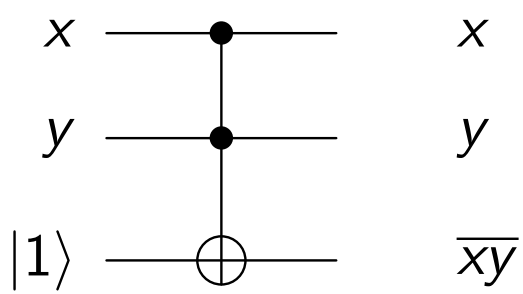

Single Qubit Gates: The Pauli Matrices

- The most common single-qubit gates are the Pauli matrices:

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

- Properties of Pauli matrices:
  - All are unitary, Hermitian, traceless, and mutually orthogonal (with respect to the trace inner product $\langle A, B \rangle = \mathrm{Tr}(A^\dagger B)$).
  - $\{I, X, Y, Z\}^{\otimes n}$ spans the space of $2^n \times 2^n$ Hermitian matrices.

Quantum Gates: Pauli X

- The quantum NOT gate: $X|0\rangle = |1\rangle$, $X|1\rangle = |0\rangle$.

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

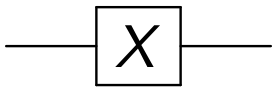

Quantum Gates: Pauli Z

- The phase flip gate: $Z|0\rangle = |0\rangle$, $Z|1\rangle = -|1\rangle$.

$$Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

- Unlike X, this leaves the measurement statistics in the computational basis unchanged — the effect of Z is only visible in a superposition.

The Hadamard Gate (H)

- Maps computational basis states to superposition states:

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

$$H|0\rangle = |+\rangle := \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \qquad H|1\rangle = |-\rangle := \frac{|0\rangle - |1\rangle}{\sqrt{2}}.$$

- H is its own inverse: $H^2 = I$ (direct computation, or note H is both unitary and Hermitian).

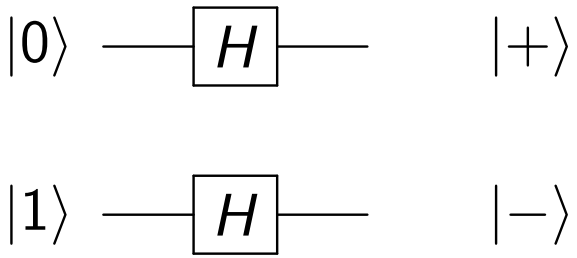

(The two wires above are two *independent* single-qubit examples, not an entangled 2-qubit circuit.)

The Effect of the Hadamard Gate on Multiple Qubits

- Applying H to each qubit of a register, independently:

$$H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle),$$

$$(H \otimes H)(|0\rangle \otimes |0\rangle) = H|0\rangle \otimes H|0\rangle = \frac{1}{2}\big(|0\rangle_2 + |1\rangle_2 + |2\rangle_2 + |3\rangle_2\big),$$

  where $|k\rangle_2$ denotes the 2-qubit basis state whose binary representation is k (e.g. $|2\rangle_2 = |10\rangle$).
- In general, for an n-qubit register:

$$H^{\otimes n}|0\rangle_n = \frac{1}{\sqrt{2^n}} \sum_{0 \le x < 2^n} |x\rangle_n .$$

- This is an **equally weighted superposition** of all possible n-qubit basis states — a good starting point for any quantum algorithm.
- **Example ($n = 100$):** applying H to each of 100 qubits initialized to $|0\rangle$ produces a single quantum state that simultaneously encodes all $2^{100} \approx 10^{30}$ basis states. This is the power of *quantum parallelism*.

Hadamard Initialization Formula

- A cornerstone of quantum algorithms is the preparation of a uniform superposition, generalized to act on *any* starting basis state $|z\rangle_n$ (not just $|0\rangle_n$):

$$H^{\otimes n}|z\rangle_n = \sum_{0\le x<2^n} \frac{(-1)^{x\cdot z}}{\sqrt{2^n}}|x\rangle_n .$$

- Here, for $|x\rangle_n = |x_{n-1}\cdots x_0\rangle$ and $|z\rangle_n = |z_{n-1}\cdots z_0\rangle$ in binary, the **bitwise dot product** is

$$x\cdot z := x_{n-1}z_{n-1} \oplus \cdots \oplus x_0 z_0 \pmod 2.$$

- Setting $z = 0$ recovers the previous slide's formula, since $x\cdot 0 = 0$ for all x.
- This formula is the engine behind algorithms like Deutsch–Jozsa and Simon's algorithm, covered later in the course.

Phase and T-Gates

- The Z gate is a special case of a more general family of **phase gates**.
- **Phase Gate (S):** adds a relative phase of i to the $|1\rangle$ component.

$$S = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}, \qquad S^2 = \begin{bmatrix} 1 & 0 \\ 0 & i^2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} = Z.$$

- $\pi/8$ **Gate (T):** adds a relative phase of $e^{i\pi/4}$.

$$T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}, \qquad T^2 = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/2} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix} = S.$$

- These provide fine-grained rotations around the Z-axis of the Bloch sphere, and (as we will see in Section 10) T is essential for universal quantum computation.

Bloch Sphere Rotations

- Any unitary operation on a single qubit can be viewed as a **rotation** of the Bloch sphere:

$$R_{\hat{n}}(\theta) = e^{-i\theta(\hat{n}\cdot\vec{\sigma})/2}, \qquad \vec{\sigma} = (X, Y, Z).$$

- Using $X^2 = Y^2 = Z^2 = I$, the exponential collapses to a closed form (analogous to Euler's formula):

$$\begin{aligned} R_x(\theta) &= \cos(\theta/2)I - i\sin(\theta/2)X, \\ R_y(\theta) &= \cos(\theta/2)I - i\sin(\theta/2)Y, \\ R_z(\theta) &= \cos(\theta/2)I - i\sin(\theta/2)Z. \end{aligned}$$

- **Question:** why $\theta/2$ in the operator, if the Bloch-sphere rotation angle is θ?

Geometry of the 1/2 Factor

- The Bloch vector $\vec{r} = (r_x, r_y, r_z)$ is defined by Pauli expectation values (Lecture 3): $r_i = \langle\psi| \sigma_i |\psi\rangle$.
- Under $U = e^{-i\theta X/2} = cI - isX$ (with $c = \cos\frac{\theta}{2}$, $s = \sin\frac{\theta}{2}$), the new coordinates are $r_i' = \langle\psi| U^\dagger \sigma_i U |\psi\rangle$.
- **Conjugating Y:** using $U^\dagger = cI + isX$ and $XY = iZ$, $YX = -iZ$, $XYX = -iZ$:

$$\begin{aligned} U^\dagger Y U &= (cI + isX)\, Y\, (cI - isX) \\ &= c^2 Y - ics(YX) + ics(XY) - i^2 s^2 (XYX) \\ &= c^2 Y - csZ - csZ + s^2 Y \\ &= (c^2 - s^2)\, Y - (2sc)\, Z = \cos(\theta)\, Y - \sin(\theta)\, Z, \end{aligned}$$

  using $c^2 - s^2 = \cos\theta$, $2sc = \sin\theta$.
- By an identical computation: $U^\dagger Z U = \sin(\theta)\, Y + \cos(\theta)\, Z$.

Bridging Operators and Coordinates

- The Bloch vector coordinates are the expectation values

$$r_y = \langle\psi| Y |\psi\rangle, \qquad r_z = \langle\psi| Z |\psi\rangle.$$

- When the state evolves to $|\psi'\rangle = U|\psi\rangle$ under $U = R_x(\theta)$, the new coordinate r'_y is:

$$r'_y = \langle\psi'| Y |\psi'\rangle = \langle U\psi| Y |U\psi\rangle = \langle\psi| U^\dagger Y U |\psi\rangle.$$

- Substituting the operator identity $U^\dagger Y U = \cos(\theta) Y - \sin(\theta) Z$ from the previous slide:

$$\begin{aligned} r'_y &= \langle\psi| \left(\cos(\theta) Y - \sin(\theta) Z\right) |\psi\rangle \\ &= \cos(\theta) \langle\psi| Y |\psi\rangle - \sin(\theta) \langle\psi| Z |\psi\rangle = \cos(\theta)\, r_y - \sin(\theta)\, r_z. \end{aligned}$$

- Similarly, using $U^\dagger Z U = \sin(\theta) Y + \cos(\theta) Z$:
$r'_z = \sin(\theta) r_y + \cos(\theta) r_z$.

The 1/2 Factor, Resolved

- Collecting the previous slide's two relations, the Bloch vector transforms by exactly a standard 2D rotation by angle θ in the YZ-plane:

$$\begin{pmatrix} r_y' \\ r_z' \end{pmatrix} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} r_y \\ r_z \end{pmatrix}.$$

- **This confirms:** $R_x(\theta)$ rotates the Bloch sphere by angle θ about the x-axis, even though the operator itself was defined using $\theta/2$.
- **Why the factor of 2?** This is a hallmark of the *spin-*$\frac{1}{2}$ (double-cover) structure of $SU(2)$: a 2π rotation of the operator ($\theta = 2\pi$) returns the state vector to $-|\psi\rangle$ (an overall sign), while the Bloch sphere itself — and hence every physical observable — returns to its starting point after only $\theta = 2\pi$ in real space, consistent with $\cos(2\pi) = 1, \sin(2\pi) = 0$ in the rotation matrix above.

Measurement and the Born Rule

- Measurement is an **irreversible** transition from quantum information to classical information (unlike every gate seen so far).
- For $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$, the probability of outcome 0 is $|\alpha|^2$ (Born rule, Lecture 3).
- The state collapses to $|0\rangle$ or $|1\rangle$ after measurement.

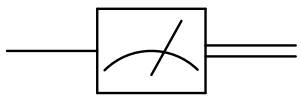

Bit

Multi-Qubit Systems

- Two qubits are described by states in $\mathbb{C}^2 \otimes \mathbb{C}^2 \cong \mathbb{C}^4$ (Lecture 3, Postulate 4).
- Basis vectors: $|00\rangle, |01\rangle, |10\rangle, |11\rangle$.
- General state: $\sum_{i,j \in \{0,1\}} \alpha_{ij} |ij\rangle$.
  - Normalization: $\sum_{i,j} |\alpha_{ij}|^2 = 1$.

Entanglement: The Bell States

- Entanglement occurs when the state cannot be factored into independent qubit states (Lecture 3).
- The most famous examples are the Bell states:
  - $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$
  - $|\Psi^+\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$
- **Question:** how do we *create* entanglement using the gates we have just introduced?

Bell State Circuit

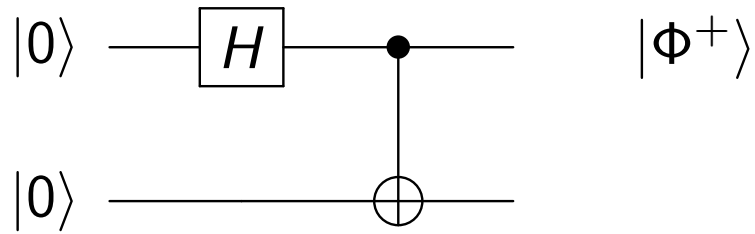

- Generating entanglement using a Hadamard followed by a CNOT:
  - Here is the math:

$$|00\rangle \xrightarrow{H\otimes I} \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}}\right) \otimes |0\rangle = \frac{|00\rangle + |10\rangle}{\sqrt{2}}$$
$$\xrightarrow{CNOT} \frac{|00\rangle + |11\rangle}{\sqrt{2}} = |\Phi^+\rangle$$

- **Exercise:** design a circuit that generates $|\Psi^+\rangle$. (Hint: start from $|01\rangle$ instead of $|00\rangle$.)

Superdense Coding: Concept

Alice and Bob share a Bell pair $|\Phi^+\rangle$. Alice can transmit **two classical bits** to Bob by sending only **one qubit** (her half of the pair).

- Alice performs a local Pauli operation based on her 2-bit message.
- She sends her qubit to Bob.
- Bob performs a joint ("Bell") measurement on both qubits to decode.

Superdense Coding: The Encoding

- Depending on the message $00, 01, 10, 11$, Alice applies one of the Pauli matrices $I_2, \sigma_x, i\sigma_y, \sigma_z$, respectively, to her (first) qubit of the shared pair $|\Phi^+\rangle$.

message	transformation U on $\|\Phi^+\rangle$	state sent
$0 = 00$	$I_2 \otimes I_2$	$\|\psi_0\rangle = \frac{1}{\sqrt{2}}(\|00\rangle + \|11\rangle)$
$1 = 01$	$\sigma_x \otimes I_2$	$\|\psi_1\rangle = \frac{1}{\sqrt{2}}(\|10\rangle + \|01\rangle)$
$2 = 10$	$i\sigma_y \otimes I_2$	$\|\psi_2\rangle = \frac{1}{\sqrt{2}}(\|01\rangle - \|10\rangle)$
$3 = 11$	$\sigma_z \otimes I_2$	$\|\psi_3\rangle = \frac{1}{\sqrt{2}}(\|00\rangle - \|11\rangle)$

- Alice sends only her single qubit over to Bob.
  - Each of the four resulting states $|\psi_0\rangle, \ldots, |\psi_3\rangle$ is, in fact, one of the four Bell states — a unique, perfectly distinguishable 2-qubit state.
  - But Bob does not yet *have* that 2-qubit state in hand: he only holds the two individual particles, which he must jointly measure to extract the encoded message.

Superdense Coding: The Decoding

- Bob applies the following circuit (CNOT, then Hadamard on the first qubit) to the joint state he received from Alice:

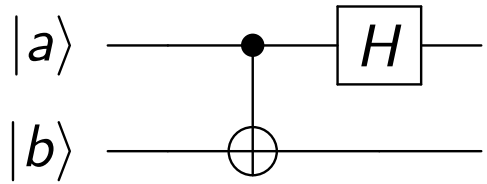

state received	after CNOT	after H
$\|\psi_0\rangle$	$\frac{1}{\sqrt{2}}(\|00\rangle + \|10\rangle)$	$\|00\rangle$
$\|\psi_1\rangle$	$\frac{1}{\sqrt{2}}(\|11\rangle + \|01\rangle)$	$\|01\rangle$
$\|\psi_2\rangle$	$\frac{1}{\sqrt{2}}(\|01\rangle - \|11\rangle)$	$\|11\rangle$
$\|\psi_3\rangle$	$\frac{1}{\sqrt{2}}(\|00\rangle - \|10\rangle)$	$\|10\rangle$

A Note on the Decoded Bit Order

- Comparing the two tables: message $00 \to$ decoded 00, and $01 \to 01$ — but message $10 \to$ decoded 11, and $11 \to$ decoded 10.
- This is **not** an error: the decoding circuit performs a *Bell-basis measurement*, and the natural output ordering of that measurement is

$$|\Phi^+\rangle \to |00\rangle\,, \quad |\Psi^+\rangle \to |01\rangle\,, \quad |\Psi^-\rangle \to |11\rangle\,, \quad |\Phi^-\rangle \to |10\rangle\,,$$

  which does *not* coincide with the decimal ordering $00, 01, 10, 11$ of the original messages.
- In practice this causes no difficulty: Alice and Bob simply agree in advance on the fixed lookup table above (exactly as given on the previous two slides) — Bob's measured bits always determine the message *uniquely*, just not in the same numerical order it was sent.

Superdense Coding: All in One Circuit

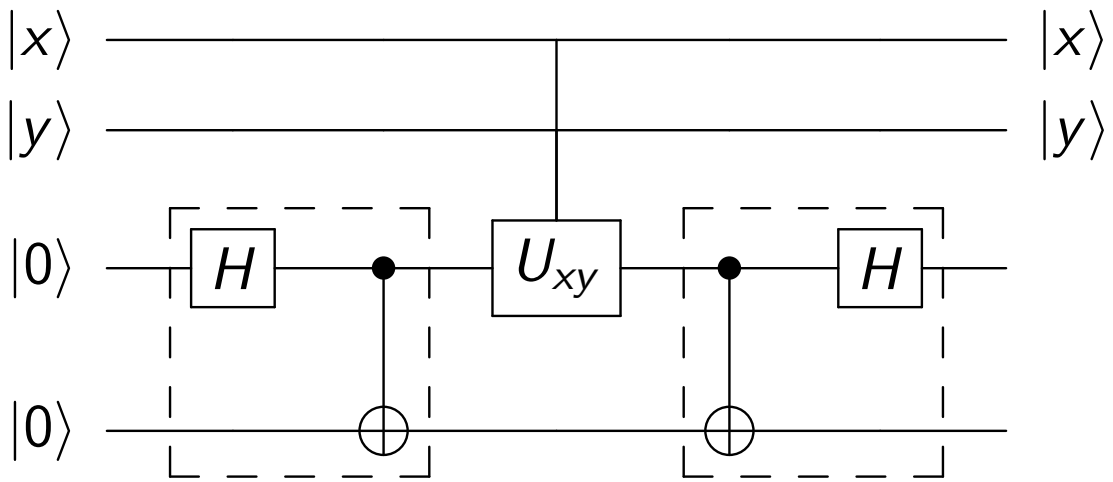

- The top two wires are Alice's classical message bits x, y — they pass through *untouched*; the thin vertical lines indicate that they determine which gate U_{xy} is applied below, not that they are entangled with it.
- $U_{xy} \in \{I_2,\ \sigma_x,\ i\sigma_y,\ \sigma_z\}$ is selected by (x, y) exactly as in the encoding table: $U_{00} = I$, $U_{01} = \sigma_x$, $U_{10} = i\sigma_y$, $U_{11} = \sigma_z$.
- The left dashed box is the **EPR gate** (H then CNOT), preparing $|\Phi^+\rangle$ on the bottom two wires; the right dashed box is Bob's EPR^{-1} (CNOT then H), exactly as verified on the decoding slide.

Why Is This Significant?

- Alice only needs to act on, and send, a *single* qubit of a shared pair to convey **two** bits of classical information.
  - ◇ "Super" refers to that jump from 1 bit to 2 bits, made possible by the properties of a shared entangled state.
- Bob can fully decode the message from the joint state:
  - ◇ The second decoded qubit alone distinguishes $\{00, 10\}$ from $\{01, 11\}$.
  - ◇ The first decoded qubit alone distinguishes $\{00, 01\}$ from $\{10, 11\}$.
- The protocol's power rests entirely on Alice and Bob *sharing* an entangled pair beforehand — without that shared entanglement, one qubit can carry at most one classical bit (Holevo's bound).
  - ◇ Exercise: try a few other Bell states as the shared resource. Does the ordering $\{I_2, \sigma_x, i\sigma_y, \sigma_z\}$ of the applied gates matter?

Quantum Teleportation: The Goal

Transfer an *unknown* qubit state $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$ from Alice to Bob, using a pre-shared entangled pair and two classical bits.

- **Note:** no information travels faster than light — Bob needs Alice's classical bits to complete the protocol.
- **No-Cloning:** Alice's original copy of $|\psi\rangle$ is destroyed during her measurement (Lecture 3); only one copy of $|\psi\rangle$ ever exists.

Quantum Teleportation: Setup

- Alice holds the unknown state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ and her half of a pre-shared pair $|\Phi^+\rangle$; Bob holds the other half. The 3-qubit starting state is:

$$|\psi\rangle \otimes |\Phi^+\rangle = \frac{1}{\sqrt{2}}\big(\alpha|000\rangle + \alpha|011\rangle + \beta|100\rangle + \beta|111\rangle\big).$$

- Alice applies CNOT (her $|\psi\rangle$ qubit controlling her EPR-half), then Hadamard to $|\psi\rangle$:

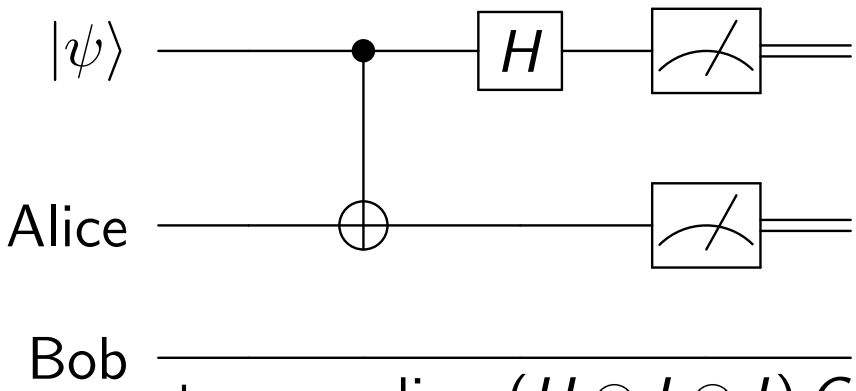

- Before measurement, expanding $(H \otimes I \otimes I)\, CNOT_{1\to 2}\left(|\psi\rangle\,|\Phi^+\rangle\right)$ gives:

$$\frac{1}{2}\big(\alpha|000\rangle + \beta|001\rangle + \beta|010\rangle + \alpha|011\rangle + \alpha|100\rangle - \beta|101\rangle - \beta|110\rangle + \alpha|111\rangle\big).$$

Quantum Teleportation: Alice's Measurement

- Grouping the 8 terms above by Alice's two qubits (the first two kets in each term), the state collapses upon measurement to one of:

$$|00\rangle \otimes (\alpha\,|0\rangle + \beta\,|1\rangle),$$
$$|01\rangle \otimes (\beta\,|0\rangle + \alpha\,|1\rangle),$$
$$|10\rangle \otimes (\alpha\,|0\rangle - \beta\,|1\rangle),$$
$$|11\rangle \otimes (-\beta\,|0\rangle + \alpha\,|1\rangle).$$

- Each of the four outcomes $|00\rangle, |01\rangle, |10\rangle, |11\rangle$ occurs with probability $\frac{1}{4}$. (Why?) *Each term in the expansion above has the same coefficient magnitude* $\frac{1}{2}$*, so each grouped outcome has total probability* $|\alpha|^2/4 + |\beta|^2/4 = 1/4$.
- Alice now holds only two *classical* bits — she no longer possesses any copy of $|\psi\rangle$ (consistent with the No-Cloning theorem).
- Bob's third qubit is left in one of the four states above, each a simple transformation of the original $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$.

Quantum Teleportation: Bob's Correction

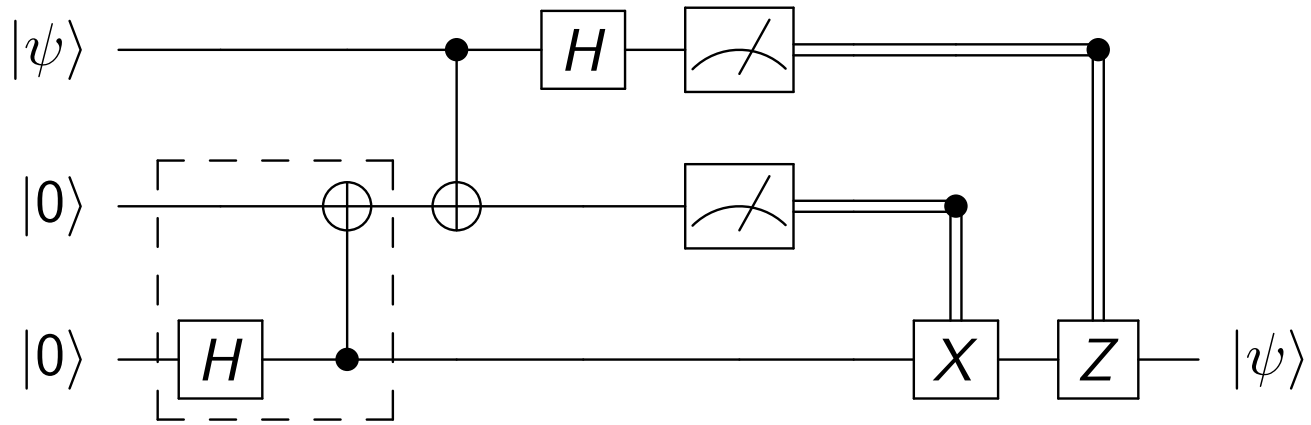

- The full circuit, left to right: Bob's EPR-half is prepared with H, then a CNOT (dashed box) creates the shared pair with Alice's EPR-half; Alice then runs her CNOT/H/measurement as before.
- Bob receives two classical bits (b_1, b_2) from Alice and applies the corresponding Pauli "fix":
  - If $b_2 = 1$ (from Alice's EPR-half qubit), apply X.
  - If $b_1 = 1$ (from the $|\psi\rangle$ qubit), apply Z.
- This restores Bob's qubit to exactly $\alpha\,|0\rangle + \beta\,|1\rangle$.
- It does not matter which party's qubit "controls" the entangling gate when the shared pair is first created — the resulting entanglement is symmetric.

Universal Gate Sets

A set of gates is **universal** if any unitary U can be approximated to error ϵ using only those gates.

- Standard set: $\{H, S, CNOT, T\}$.
- A gate U is a **Clifford gate** if it maps Pauli operators to other Pauli operators under conjugation ($UPU^\dagger$ is again a Pauli operator, for every Pauli P).
  - Any circuit consisting only of Clifford gates $\{H, S, CNOT\}$ can be simulated efficiently on a classical computer in polynomial time (Gottesman-Knill theorem).
  - Clifford gates alone are therefore *not* "computationally universal" — they offer no quantum advantage by themselves.
- The T-gate is **non-Clifford**.
  - It allows for non-trivial rotations not reachable by Clifford gates alone.
  - Adding T to the Clifford set provides the necessary complexity for universal quantum computation.

Solovay-Kitaev Theorem

- Guarantees that any universal gate set can approximate an arbitrary unitary *efficiently*.
- The required gate count grows only as $O\big(\text{polylog}(1/\epsilon)\big)$.
  - A function $f(x)$ is **polylogarithmic** if $f(x) = O\big((\log x)^c\big)$ for some constant $c > 0$.
  - Consequently, doubling the precision (halving ϵ) only requires a modest, polynomial increase in the gate-sequence length — not an exponential blow-up.
- This makes a small, fixed set of gates viable for implementing *any* quantum algorithm to arbitrary precision.

Solovay-Kitaev: Practical Implications

- **Hardware Independence**
  - ◇ We do not need a hardware-native gate for every possible rotation angle θ.
  - ◇ A small, discrete set (e.g., $\{H, T, \text{CNOT}\}$) is sufficient to simulate any $U \in SU(2^n)$.
- **Fault Tolerance & Error Correction**
  - ◇ It justifies focusing error-correction resources on a limited gate set.
  - ◇ Since non-Clifford gates (like T) are expensive to make fault-tolerant, Solovay-Kitaev helps minimize the required T-count.
- **Quantum Compilation**
  - ◇ Acts as the "compiler" that translates high-level abstract unitaries into machine-level instructions.
  - ◇ Modern compilers (Qiskit, Cirq) use optimized versions of this theorem to synthesize circuits.

Summary of Lecture 4

- **Reversibility:** classical gates are usually irreversible (information loss $\rightarrow$ heat, Landauer's principle); quantum gates are always reversible (unitary).
- **Reversible classical logic:** CNOT implements reversible XOR; Toffoli (with ancillas) implements universal classical logic reversibly.
- **Single-qubit gates:** Pauli $\{X, Y, Z\}$, Hadamard H, phase gates $\{S, T\}$; every single-qubit unitary is a Bloch-sphere rotation $R_{\hat{n}}(\theta)$.
- **Entanglement:** CNOT following H generates Bell states from product states.
- **Protocols:** superdense coding sends 2 classical bits per shared qubit; teleportation sends 1 qubit's worth of quantum information using a shared pair plus 2 classical bits.
- **Universality:** $\{H, S, CNOT, T\}$ is universal; Solovay-Kitaev guarantees efficient approximation using a small fixed gate set.

Looking Ahead: Lecture 5

Having established the circuit model, we now turn to what quantum circuits can *compute*. In Lecture 5 we will cover:

- **Computational models:** the Turing machine and the Church-Turing thesis.
- **Complexity classes:** P, NP, BQP, and where quantum computing fits.
- **The Deutsch-Jozsa algorithm:** the first concrete demonstration of quantum speedup.