CMSC5719 - Lecture 5 (Fall 2025)
An Introductory Lecture on Quantum Computing

Xiao Liang
https://xiao-liang.github.io

Department of Computer Science and Engineering

The Chinese University of Hong Kong

(September 29, 2025)

Chapter 1
Genesis: The (Simplified) Four Postulates

One aspect of teaching a CS-oriented Quantum Computing course that I find somewhat frustrating is the limited opportunity to discuss the *quantum mechanics* underpinning it. Instead, I have to adopt an axiomatic approach, asking students to accept *the four postulates* of quantum mechanics as fundamental truths upon which the entire course is built. This approach is dictated by two major reasons:

1. In a postgraduate quantum computing course, there are many topics to cover related to quantum computing itself, and the lecture time is finite. As a result, I cannot dedicate much time to exploring the quantum mechanics behind the subject in depth.

2. As human beings, we do not yet fully understand the "why" behind quantum mechanics. While quantum physicists have an extraordinary grasp of *what* quantum mechanics predicts and *how* to apply it, the reasons behind its postulates remain one of the great open questions in physics. This unanswered "why" points to deeper mysteries about the nature of reality and the underlying structure of the universe.

Given these constraints, the approach I usually take is to provide as many "evidences" or "rationales" as possible to validate the four postulates, in the hope of helping you "feel comfortable" with them. For now, this is the best I (and even modern physics) can do.

With this perspective in mind, I would like you to watch this video from 14:54 to 22:27. It is an interview with Richard Feynman, where he discusses the boundaries one must respect when probing the truth of events. I hope this will help you appreciate why we accept the four postulates of quantum mechanics as the "foundation" of quantum computing without delving further into why they must be true. Exploring such questions would take us into the cutting edge of quantum mechanics research — territory that remains unresolved even today.

By the way, the entire one-hour interview with Feynman is highly engaging and inspiring. If you have time, I strongly encourage you to watch the full video. Feynman's insights are valuable for anyone pursuing creativity, curiosity, or research in any field. No matter what career path or research direction you follow, there is much to learn from his words.

1.1 Postulate 1

A qubit is a **unit** vector on $\mathbb{C}^2$. That is, we can always express it as

$$\mathsf{state} = a_0 \cdot \mathbf{e}_0 + a_1 \cdot \mathbf{e}_1,$$

where $|a_0|^2 + |a_1|^2 = 1$.

We use Dirac notation: $|0\rangle = \mathbf{e}_0$, $|1\rangle = \mathbf{e}_1$.

this extends to n-dimensional vector space, simply let $|k\rangle = \mathbf{e}_k$ for $k \in \{0, 1, \ldots, n-1\}$.

$\langle a|$ is the conjugate transpose for $|a\rangle$. Then, $\langle a|a\rangle$ is simply the inner product.

Example 1.1.1. We will often see the following two qubits. They together are referred to as the *Hadamard basis* (of $\mathbb{C}^2$):

$$|+\rangle := \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \quad \text{and} \quad |-\rangle := \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle).$$

1.2 Postulate 2

Putting two quantum registers together, we will combine their own space. **And** any unit vector in this new space is a valid quantum state (not necessarily the separable ones.)

That is, you will have a new unit vector in $\mathbb{C}^4$. Essentially, you have a new space spanned by the 4 combination of the basis for each of the two.

- Mathematically, this is instantiated by "tensor product."
- (Not necessarily the separable ones:) In partucular, you could have $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$.
- Note that we can keep appending more qubits. When we have n qubits, we will be in space $\mathbb{C}^{2^n}$. That is, we represent exponential amount of information

1.3 Postulate 3

A quantum state can *only* evolve by unitary matrices. Notation-wise:

$$|\psi\rangle \to U\,|\psi\rangle\,.$$

Unitary matrices are square matrices over the complex field satisfying $UU^\dagger = I$, or equivalently $U = U^\dagger$ (which is also equivalent to $U^\dagger = U^{-1}$).

Some remarks:

1. This is a result of Shrödinger's equation

   $$i\hbar\frac{d}{dt}\,|\psi(t)\rangle = \hat{H}\,|\psi(t)\rangle\,.$$

   We don't provide further explanation. Again, we focus on what we can do obeying these postulates, instead of questions why these postulates must hold.

2. This postulate is consistent with Postulate 1 — unitary operations preserve the length of quantum states. Therefore, if you begin with a unit vector, it remains a unit vector after applying any sequence of unitary operations.

3. Recall that we use quantum state to represente data/information. They are the quantum analogue of classical numbers like $1, 2, 3, \ldots$. Classically, we have freedom to do addition, multiplication, division etc. But in the quantum setting, we can only do unitaries!

Example 1.3.1. Some examples of unitary matrices:

- The Pauli matrices[1] :

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

  Calculate the result of applying these matrices to $|0\rangle$ and $|1\rangle$.

- The Hadamard matrix $H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$. Verify by yourself that $H|0\rangle = |+\rangle$ and $H|1\rangle = |-\rangle$. (Think: what is the inverse of H?)

1.4 Postulate 4

You can measure a state, governed by the following rule: When you measure the state

$$|\psi\rangle = a_0|0\rangle + a_1|1\rangle + \ldots a_{2^n-1}|2^n - 1\rangle,$$

it "collapses" to $|k\rangle$ with probability $|a_k|^2$, for all $k \in \{0, 1, \ldots, 2^n - 1\}$.

- Note that this is consistent with the requirement that $\sum_i |a_i|^2 = 1$.
- Xiao: Show examples of measuring $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ on the whiteboard. Xiao!

[1] These matrices are named after the renowned physicist Wolfgang Ernst Pauli, recipient of the Nobel Prize in Physics in 1945.

Chapter 2

Classical Computation as a Special Case of Quantum Computation

Quantum mechanics is a more general theory than classical mechanics—there is no inconsistency between the two, and the former *incorporates* the latter. Therefore, any computation that can be performed by a classical computer must also be implementable on a quantum computer. However, this is not immediately evident from the four postulates of quantum computing that we described earlier.

- **Question:** How can we implement any classically-computable function on a quantum computer?

To answer this question, we first need to characterize "classical computation."

- **Boolean Circuits**: A model based on logic gates (AND, OR, NOT) applied to Boolean input bits.

Theorem 2.0.1 (Informal)**.** The set of gates $\{AND, NOT\}$ are *universal* in the sense that once you can implement them, you can compute any classically-computable fuction. ◊

2.1 Quantum Computing Simulates Classical Circuits

Due to Theorem 2.0.1, to show that quantum computing is a superset of classical computing, it suffices to show how to compute AND and NOT gates.

Simulating the NOT gate. We can do that for NOT using the so-called Pauli-X,[1]

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

where it is obvious that $X|b\rangle = |\neg b\rangle$ for all $b \in \{0, 1\}$.

Simulating the AND gate. In contrast, it is much harder to imagine a unitary implementing the AND gate:

- They have a fan-in of 2 and a fan-out of 1, which means they cannot be represented by unitary matrices. However, according to Postulate 2, only unitary operations are allowed in the quantum world.

To implement an AND gate in a quantum circuit, we must embed it in a larger *reversible unitary operation*. A common approach is to use the *Toffoli gate*, which is a 3-qubit unitary operation defined as:

[1] This is named after the renowned physicist Wolfgang Ernst Pauli, recipient of the Nobel Prize in Physics in 1945.

$$\text{Toffoli}\,|a,b,c\rangle = |a,b,c \oplus (a \cdot b)\rangle$$

for all $a, b, c \in \{0,1\}$.

How the Toffoli Gate Implements AND

- The Toffoli gate takes three qubits: two **control qubits** a and b, and one **target qubit** c.
- The target qubit c is flipped if and only if both a and b are **1**.
- If we initialize $c = 0$, then after applying the Toffoli gate, c will store the result of $a \cdot b$, effectively computing the AND function.

Importantly, the first two qubits (on the output wires) serves as "auxiliary" qubits in the procedure above. Very roughly, You can think of it this way: these bits are there to ensure that necessary information is preserved without loss, even though you might not be directly interested in that information (since your goal is only to compute the summation); They ensure that your computing is *Reversible*, which is a necessary condition for quantum computing (Postulate 2).

Matrix Representation of the Toffoli Gate The Toffoli gate is represented by the following 8×8 unitary matrix:

$$U_{\text{Toffoli}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

This matrix swaps the states $|110\rangle$ and $|111\rangle$, implementing the controlled transformation.

2.2 Quantizing any Classical Boolean Circuits

The above procedure we did for AND can indeed be generalized to arbitrary boolean functions. We state this as the following informal theorem without presenting its proof:

Theorem 2.2.1 (informal)**.** For any classical circuit computing a function $f : \{0,1\}^n \to \{0,1\}^m$, we can always build a quantum circuit U_f that agrees with f on every classical input, with the help of "ancilla" qubits:

$$U_f\,|x\rangle_{\mathsf{in}}\,|z\rangle_{\mathsf{out}}\,|0\ldots0\rangle_{\mathsf{ancilla}} = |x\rangle_{\mathsf{in}}\,|f(x) \oplus z\rangle_{\mathsf{out}}\,|\text{garbage}\rangle_{\mathsf{ancilla}}\,.$$

$\Diamond$

You don't need to focus on the "ancilla" qubits—they are present due to certain technical reasons that we won't cover in this introductory talk.

If you only expose the x and z registers, the computation can be written as

$$\mathcal{O}_f : |x,z\rangle \mapsto |x, f(x) \oplus z\rangle\,.$$

Such an access to f is known as the **Quantum Oracle Access to a classical function** f. Note that O_f may not be a unitary anymore.

2.3 Phase Oracle and Phase Kickback Trick

For any classical function $f : \{0,1\}^n \to \{0,1\}$ with a single-bit output, and any $x \in \{0,1\}^n$, it holds that

$$\begin{aligned}
\mathcal{O}_f |x\rangle |-\rangle &= \mathcal{O}_f |x\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \\
&= |x\rangle \frac{1}{\sqrt{2}}(|f(x)\rangle - |f(x) \oplus 1\rangle) \\
&= \begin{cases} |x\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) & f(x) = 0 \\ |x\rangle \frac{1}{\sqrt{2}}(|1\rangle - |0\rangle) & f(x) = 1 \end{cases} \\
&= \begin{cases} |x\rangle |-\rangle & f(x) = 0 \\ -|x\rangle |-\rangle & f(x) = 1 \end{cases} \\
&= (-1)^{f(x)} |x\rangle |-\rangle
\end{aligned}$$

If we hide the ancilla register that contains $|-\rangle$, we will have a oracle:

$$\mathsf{PhO}_f : |x\rangle \mapsto (-1)^{f(x)} |x\rangle .$$

This is referred to as the "Phase Oracle."
The above derivation of is referred to as the "Phase Kickback Trick" by some authors.

Chapter 3

Advantages of Quantum Computing over Classical Computing

3.1 The First Example: Deutsch Algorithm

First quantum algorithm. In the mid-1980s by David Deutsch [Deu85].

Problem Statement: Given oracle access to a boolean function $f : \{0,1\} \to \{0,1\}$, decide whether $f(0) = f(1)$ or $f(0) \neq f(1)$.

Xiao: Draw the truth table on the whiteboard. Xiao!

Quantum Supremacy: Classical algorithms need to make 2 queries to answer this qustion. But Deutsch uses only 1 quantum query.

Algorithm: Best illustrated using phase oracle (or phase kickback trick):

1. **Prepare uniform superposition:** Apply H to $|0\rangle$ to get $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.
2. **Query phase oracle PhO_f:**

$$\frac{(-1)^{f(0)}|0\rangle + (-1)^{f(1)}|1\rangle}{\sqrt{2}} = \begin{cases} (-1)^{f(0)}|+\rangle & f(0) = f(1) \\ (-1)^{f(0)}|-\rangle & f(0) \neq f(1) \end{cases}.$$

3. **Hadamard it back:** Apply Hamamard again. We get

$$\begin{cases} (-1)^{f(0)}|0\rangle & f(0) = f(1) \\ (-1)^{f(0)}|1\rangle & f(0) \neq f(1) \end{cases}$$

4. **Measure:** in computational basis.

Note that the last two steps together are typically called *measure in Hadamard basis*.

Its Circuit: Xiao: the following diagram is not accurate. It would have been better if I used a single input and output wire. The U_f can be encapsulated inside a phase oracle. Xiao!

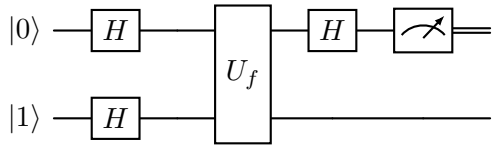

3.2 Example 2: Quantum Key Distribution

Classically, it is provably true that *key distribution* relies on strong cryptographic assumptions that public-key encryption exists. This relies on our faith in the existence of mathematically hard problems, e.g., Factoring, Diffie-Hellman, etc.

QKD is secure without any assumptions! We call it "information-theoretically" secure, in contrast to the "computationally secure" case of the classical setting.

To show QKD, we first need to talk about no-cloning theorem:

3.2.1 No Cloning Theorem

Theorem 3.2.1 (No Cloning Theorem)**.** For any Hilbert space $\mathcal{H}$, there is no unitary operator U on $\mathcal{H} \otimes \mathcal{H}$ such that for all normalized states $|\psi\rangle_A \in \mathcal{H}$ and $|e\rangle_B \in \mathcal{H}$,

$$U(|\psi\rangle_A|e\rangle_B) = |\psi\rangle_A|\psi\rangle_B.$$

◊

Proof. Assume for the sake of contradiction that we have a universal cloning machine U such that for an arbitrary state $|\psi\rangle$,

$$U(|\psi\rangle_A|e\rangle_B) = |\psi\rangle_A|\psi\rangle_B. \tag{3.1}$$

Now, choose $|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.

According to Equation (3.1), we have

$$U\left(\frac{1}{\sqrt{2}}(|0\rangle_A + |1\rangle_A)|e\rangle_B\right) = \frac{1}{\sqrt{2}}(|0\rangle_A + |1\rangle_A)\frac{1}{\sqrt{2}}(|0\rangle_B + |1\rangle_B). \tag{3.2}$$

On the other hand, by the linearity of U, we obtain

$$U\left(\frac{1}{\sqrt{2}}(|0\rangle_A + |1\rangle_A)|e\rangle_B\right) = \frac{1}{\sqrt{2}}U(|0\rangle_A|e\rangle_B) + \frac{1}{\sqrt{2}}U(|1\rangle_A|e\rangle_B) = \frac{1}{\sqrt{2}}(|0\rangle_A|0\rangle_B) + \frac{1}{\sqrt{2}}(|1\rangle_A|1\rangle_B). \tag{3.3}$$

However, Equation (3.2) is not equal to Equation (3.3), which leads to a contradiction.

■

3.2.2 The QKD Protocol

Classically, QKD requires hardness assumptions. The existence of it implies the existence of OWF, which at least implies $\mathsf{P} \neq \mathsf{NP}$.

Quantumly, QKD exists *unconditionally*. This is a result by Brassard and Bennett in year 1984 [BB84]. The protocol is called BB84 protocol.

3.2.3 The First version of BB84:

Prerequisite: Xiao: Explain the following part: Encoding classical bits in quantum manner. Xiao!

- Encoding in standard/computational basis.
- Encoding in Hadamard basis.

Figure 3.1: Commercially Available QKD Devices

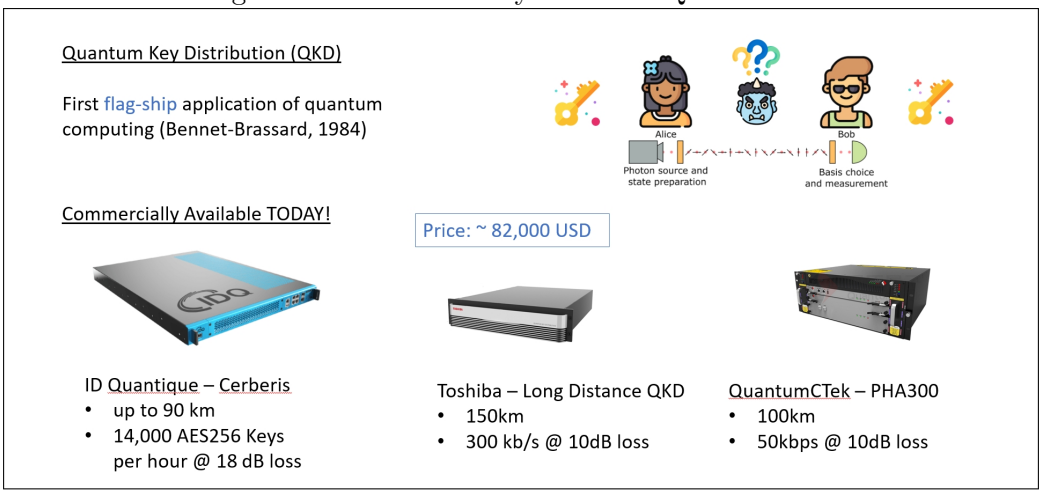

The First version of the Protocol:

1. Alice chooses uniformly at random a pair of strings $x, y \in \{0,1\}^n$.
2. Alice then generates an n-qubit state $|\psi\rangle$, where she uses the bits of y to determine which basis to encode her qubits in:
   - If $y_i = 0$, encode x_i in the standard basis $\{|0\rangle, |1\rangle\}$.
   - If $y_i = 1$, encode x_i in the Hadamard basis $\{|+\rangle, |-\rangle\}$.

   This gives you $|\psi_i\rangle$
3. **(Quantum Round.)** Alice sends the quantum state $|\psi\rangle \coloneqq |\psi_1\rangle \dots |\psi_n\rangle$ to Bob.
4. Bob picks a string $y' \in \{0,1\}^n$ uniformly at random. Bob uses the bits of y' to determine the basis in which to measure each qubit of the state received from Alice:
   - If $y'_i = 0$, he measures in the standard basis $\{|0\rangle, |1\rangle\}$.
   - If $y'_i = 1$, he measures in the Hadamard basis $\{|+\rangle, |-\rangle\}$.

   He then records the result of each measurement as a bit string x'
5. **(Classical Rounds.)** Alice and Box exchange the y and y' strings.

Xiao: Draw a picture of Alice and Bob. Be explicit that there should be a *quantum* channel for $|\psi\rangle$, and a *classical* channel for y and y'. Xiao!

Claim 3.2.1. Here are some important observations:

- **Correctness (for "matching" positions).** For all i, $x_i = x'_i$ if $y_i = y'_i$. Also, in expectation, $\frac{n}{2}$ positions would be "matching" (i.e., $y_i = y'_i$).
- **"Weak" Security:** If Eve (Eavesdropper) only corrupts the **Classical Channel**, then the protocol is secure.

However, Eve may want to steal information from the quantum channel as well. How do we deal with that? This brings us to the next section:

3.2.4 The Eventual BB84 Protocol

Idea 1: Let's ask Alice and Bob to check *all the "matching" positions*. Then, No-cloning theorem implies that Eve cannot steal **too many** positions from the quantum channel, otherwise he would be caught.

- As a simply calculation, if Eve steals a constant fraction, say $0.1n$, of qubits from $|\psi\rangle$, then he will be caught with probability $(\frac{1}{2})^{0.1n}$. This is exponentially small so that we don't need to worry about it.
- But what if Eve just steals only a few positions so that Alice and Bob didn't notice? This introduce errors to the "matching " position. Fortunately, this is a well-studied problem in information theory. We can solve this using *Error Correction Code*.
- We emphasize that this Error Correction Code only help Alice and Bob to correct any errors, so that the "matching" positions really match. However, note that Eve may have already learned a few information of the "matching" positions.

However, we cannot ask Alice and Bob to check all the positions — our eventual goal is to have them agreed on some secret string! This brings us to our second idea

Idea 2:

- Split the "matching" positions in half randomly. Check one half of it by revealing the positions to gain confidence that "only a few positions have been stolen by Eve."
- For the remaining "unrevealed" half, using ECC to make sure Alice and Bob agree on the same bits.
- Ideally, we would like to use the "unrevealed" half as the final secret string. However, as mentioned earlier, ECC only ensure correctness, but not privacy. Eve may have already stolen a few bits out of there.
- Fortunately, this is another well-studied topic in cryptography called "privacy amplification." We can use *randomness extractor* to "distill" the final secret string.

The Final Split:

- Start from $n = 16k$ qubits.
- In expectation, we will have $8k$ "matching" positions
- Reveal have of them for detecting if Eve steals too many bits
- There will be $4k$ unrevealed positions. Apply ECC code to it, we will get k correct, almost secret string
- apply randomness extractor to it, we will get, say, $0.1k$ correct, secret string.

If we set $k = 1000$, eventually, we will have 100-bit secret string. Starting from here, they can use symmetric key encryption to expand this to whatever length they want.

Chapter 4
Closing Remarks

Quantum computing is a vast and rapidly evolving field, full of exciting research opportunities. You can study how to physically implement quantum computers, explore potential real-world applications—such as drug discovery, protein simulation, and combinatorial optimization—or dive deep into its theoretical foundations. Quantum computing is a rare example of a computing model where the hardware, software, and theory are all still in their infancy.

As computer theorists, we typically don't focus on the hardware—that's best left to physicists. What we care about is the theory: designing efficient quantum algorithms, analyzing their complexity, and understanding the fundamental limits of quantum computation, just as we've done for classical computing.

4.0.1 Understanding Computational Power

In classical computation, we have developed a well-established framework to understand what can and cannot be computed efficiently. This is the domain of computational complexity theory. It helps us classify problems, understand their intrinsic difficulty, and guide algorithm design.

With quantum computing, however, we are dealing with a new species of computation. The boundaries of quantum computational power are still largely unknown. What problems can quantum computers solve efficiently? What problems remain hard for them? How does quantum computation compare to classical computation in terms of power and limitations? These are deep and fundamental questions that remain open.

4.0.2 Quantum Computing and Cryptography

Quantum algorithms have already demonstrated the ability to break classical cryptographic assumptions. For example, Shor's algorithm breaks RSA, a cornerstone of modern cryptography. This poses a major challenge: How do we design cryptographic systems that remain secure against quantum attacks? This has given rise to the field of post-quantum cryptography, and also motivates the study of quantum cryptographic protocols that are secure by the laws of quantum mechanics themselves.

4.0.3 Quantum-Exclusive Effects and New Possibilities

Quantum mechanics introduces unique phenomena—such as the no-cloning theorem—that have no classical analogue. These enable entirely new functionalities that are impossible in classical settings.

One well-known example is Quantum Key Distribution (QKD), which allows two parties to exchange a cryptographic key with information-theoretic security, guaranteed by the laws of physics.

There are other fascinating ideas too:

- Quantum money: a form of digital currency that is physically impossible to counterfeit, due to the no-cloning property of quantum states. (We refer Interested readers to Appendix A for more details.)

- Quantum anti-piracy: schemes that protect digital content from being copied, leveraging quantum principles.

These examples show how quantum mechanics can unlock new kinds of secure or otherwise impossible functionalities. As theorists, we are interested in asking: **What other useful tasks can be enabled by quantum-exclusive effects?** This is a rich and largely unexplored direction.

4.0.4 Final Thoughts

Quantum computing offers computer theorists a new frontier to explore—one that challenges our classical intuitions and opens the door to radically new paradigms. Whether it's understanding quantum complexity, designing new algorithms, or leveraging quantum mechanics for cryptographic or informational tasks, this field promises a wealth of deep, fascinating questions.

Thank you for joining this introduction. I hope this lecture has sparked your curiosity and inspired you to explore further!

Bibliography

[BB84] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing, Bangalore, Dec. 1984*, pages 175–179, 1984. 8

[Deu85] David Deutsch. Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985. 7

Appendix A
Wiesner's Private-Key Quantum Money

In the early 1970s, Stephen Wiesner proposed one of the first applications of quantum information: a scheme for creating unclonable quantum money. The idea was to leverage the **no-cloning theorem** and the uncertainty principle to construct banknotes that cannot be counterfeited, even by someone with unlimited computational power.

Motivation

Classically, secure money requires cryptographic assumptions (e.g., digital signatures or hard problems). Wiesner's scheme is information-theoretically secure based on the laws of quantum mechanics.

The Setup

A quantum banknote consists of:

- A classical serial number s, which is unique and publicly visible.
- A quantum state $|\psi_s\rangle$, which consists of n qubits, each randomly prepared in one of four possible states:

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

- The bank maintains a secret database mapping s to the preparation basis and bit for each qubit.

Each qubit is prepared in either the **computational basis** ($|0\rangle, |1\rangle$) or the **Hadamard basis** ($|+\rangle, |-\rangle$), chosen uniformly at random.

Issuing a Banknote

To issue a quantum banknote:

1. The bank generates a random serial number s.
2. For each of n positions:
   - Choose a random basis: either Z (computational) or X (Hadamard).
   - Choose a random bit $b \in \{0, 1\}$.
   - Prepare the corresponding quantum state: $|0\rangle$, $|1\rangle$, $|+\rangle$, or $|-\rangle$.
3. The banknote is the pair $(s, |\psi_s\rangle)$. The bank stores $(s, \text{basis}, \text{bit})$ in a secure database.

Verifying a Banknote

To verify $(s, |\psi_s\rangle)$:

1. The bank looks up s in its database and retrieves the correct basis and bit string.
2. The bank measures each qubit in the correct basis (either Z or X).
3. It compares the measurement outcomes to the expected bit string.
4. If all outcomes match, the banknote is accepted as valid.

Security Intuition

The security stems from two key quantum principles:

- **No-cloning theorem:** An unknown quantum state cannot be copied.
- **Measurement disturbance:** Measuring a qubit in the wrong basis irreversibly disturbs its state.

Thus, even if a counterfeiter obtains a valid quantum banknote, they cannot make a perfect copy. Any attempt to measure the qubits to learn the state will likely disturb it, making the counterfeit detectable.

Example

Suppose the bank creates a 4-qubit banknote with serial number $s = 1234$, and chooses:

Qubit	Basis	Bit
1	Z	0
2	X	1
3	Z	1
4	X	0

Then, the quantum state is:

$$|\psi_s\rangle = |0\rangle \otimes |-\rangle \otimes |1\rangle \otimes |+\rangle$$

Only the bank knows the basis choices. A counterfeiter trying to copy the state without this information will inevitably introduce errors.

Limitations and Later Developments

Wiesner's scheme is **private**, meaning only the bank can verify the money. Publicly verifiable quantum money remains an active area of research, with proposals using cryptographic assumptions or complexity-theoretic constructions.

Conclusion

Wiesner's quantum money is a beautiful and foundational idea in quantum information science. It was one of the first proposals to demonstrate a practical task that is impossible classically but achievable with quantum mechanics.

List of Corrections