CS 599 P1: Introduction to Quantum Computation Boston University, Fall 2025
Instructor: Alexander Poremba

LECTURE # 24: QUANTUM ADVANTAGE & CRYPTOGRAPHIC TESTS OF QUANTUMNESS

One of the major milestones of quantum computing is to demonstrate convincing *quantum advantage*: a task that can be performed efficiently by a quantum computer but is infeasible for any classical computer.

A canonical example is *Shor's algorithm*, which factors integers in polynomial time on a fault-tolerant quantum computer. Under widely believed assumptions in cryptography and complexity theory, integer factorization requires super-polynomial time on a classical computer. This constitutes a striking theoretical separation between quantum and classical computation.

However, this example immediately raises a practical concern. Shor's algorithm requires large, error-corrected quantum computers with thousands or millions of logical qubits. As of today, such machines do not exist. In fact, the largest numbers factored using Shor's algorithm on real hardware are

$$15 = 3 \times 5,$$

a result that is impressive scientifically, but hardly threatening to modern cryptography.

This motivates a broader and more urgent question:

> *Can we demonstrate quantum advantage using near-term quantum devices, without full error correction?*

To sharpen this question, it is useful to articulate what we want from such a demonstration. Ideally, a task exhibiting quantum advantage should satisfy the following three criteria:

1. **Feasible on near-term quantum hardware.** The task should be implementable using noisy, intermediate-scale quantum (NISQ) devices.
2. **Classically hard.** Simulating or reproducing the task should be infeasible for any efficient classical algorithm, under plausible complexity-theoretic assumptions.
3. **Easy to verify classically.** A classical observer should be able to efficiently check whether the quantum device performed the task correctly.

Achieving all three simultaneously is a major open problem in quantum computing. Much of modern research on quantum advantage and *proofs of quantumness* can be viewed as an attempt to navigate the tradeoffs between these requirements.

1 State of the Art

The current landscape of quantum advantage proposals can be summarized roughly as follows [AZ24]:

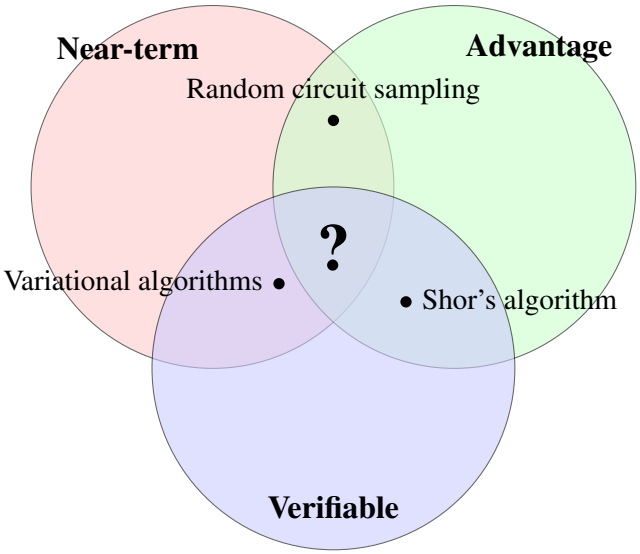

Figure 1: The state-of-the-art in quantum advantage protocols.

Random circuit sampling. Experiments by Google and others have implemented random quantum circuits on tens of qubits and sampled from the resulting output distributions. Under complexity-theoretic assumptions, reproducing these distributions appears to be classically intractable. This places random circuit sampling in the overlap between *near-term* and *advantage*. However, verification is subtle: checking correctness requires expensive brute-force classical computation.

Shor's algorithm. As discussed above, Shor's algorithm offers a clean and verifiable form of quantum advantage. Unfortunately, it lies far outside the near-term regime.

Variational algorithms. Variational quantum algorithms are widely implementable on NISQ devices and are easy to verify classically (one can simply evaluate the objective function). However, no convincing evidence yet shows that they provide provable quantum advantage.

The "holy grail" is the triple overlap: a protocol that is near-term feasible, provably classically hard, and efficiently verifiable. One promising route toward this goal uses *interaction*.

2 Interactive Quantum Advantage

Interaction is a powerful resource in theoretical computer science. Interactive proof systems allow a weak verifier to certify the outcome of a computation performed by a powerful prover. A famous example is the class IP, which is known to equal PSPACE—polynomial-space computation.

This suggests the following approach: instead of passively observing a quantum device, we interact with it using a carefully designed protocol.

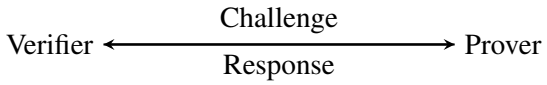

Output: YES / NO

At the end of the interaction, the verifier outputs either YES or NO. The goal is that:

- An honest *quantum* prover convinces the verifier with high probability.
- Any efficient *classical* prover fails with non-negligible probability.

A natural and crucial question is:

Can the verifier be entirely classical?

We begin with a warm-up example in which the verifier is still quantum.

2.1 Warm-up: A Qubit Test

The Qubit Test

1. The verifier samples random bits $x, \theta \in \{0,1\}$ and sends the state
$$H^\theta |x\rangle$$
to the prover.
2. After a short delay, the verifier sends θ to the prover.
3. The prover replies with a bit x'.
4. The verifier outputs YES if $x' = x$ and NO otherwise.

If the prover is quantum, they can simply store the qubit, measure it in the basis specified by θ, and recover x with certainty. By repeating this test many times, the verifier can become convinced that the prover can effectively handle measurements in incompatible bases, a distinctly quantum behavior.

However, this protocol still requires the verifier to prepare and send quantum states. All communication is classical *except* the first step. This leads to a tantalizing question:

> *Can we force the prover to generate the quantum state themselves, while the verifier retains "control" over it?*

Remarkably, the answer is yes—using cryptography.

3 Cryptographic Tests of Quantumness

We now describe a cryptographic protocol that enables a *classical verifier* to test quantumness via interaction [BCM+21, KMCVY22]. Such cyptographic tests of quantumness are inherently *verifiable* and may potentially even be more feasible than running Shor's algorithm on classically hard instances.

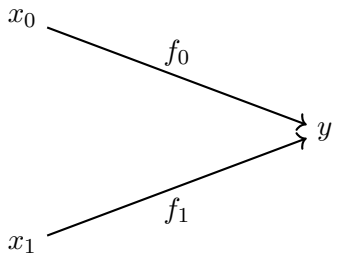

Figure 2: A *claw pair* consists of (x_0, x_1) such that $f_0(x_0) = f_1(x_1) = y$, for $y \in \{0,1\}^n$.

3.1 Trapdoor Claw-Free Functions

The key building block behind a cryptographic test of quantumness is a *trapdoor claw-free function* (TDF). A trapdoor claw-free function consists of a pair of efficiently computable functions

$$f_0, f_1 : \{0,1\}^n \to \{0,1\}^n$$

with the following properties:

1. **Injective with shared image.** For every image $y \in \{0,1\}$, there exist unique $x_0, x_1 \in \{0,1\}^n$ with

   $$f_0(x_0) = f_1(x_1) = y.$$

   The pair (x_0, x_1) is called a *claw* and is visualized in Figure 2.

2. **Claw-free.** Given public descriptions of the functions f_0 and f_1, no efficient (classical) algorithm can find a *claw pair* (x_0, x_1) such that $f_0(x_0) = f_1(x_1) = y$, for some image $y \in \{0,1\}^n$.

3. **Trapdoor.** There exists secret information $\mathtt{td}$ that allows one to efficiently recover the claw (x_0, x_1) from any image string y.

Such functions can be constructed from standard cryptographic assumptions, including Diffie–Hellman, Rabin's function, and (most importantly) even Learning with Errors (LWE)—a popular post-quantum secure cryptographic assumption which derives its hardness from worst-case lattice problems.

3.2 A Cryptographic Test of Quantumness

Cryptographic Test of Quantumness

1. The verifier samples (f_0, f_1, td) and sends (f_0, f_1) to the prover.
2. The prover prepares the state
$$\frac{1}{\sqrt{2^{n+1}}} \sum_{b\in\{0,1\}} \sum_{x\in\{0,1\}^n} |b, x\rangle\, |f_b(x)\rangle$$
and measures the second register, obtaining y. The state collapses to
$$\frac{1}{\sqrt{2}}\big(\, |0, x_0\rangle + |1, x_1\rangle\,\big).$$
The prover sends y to the verifier.
3. The verifier samples a random challenge bit $\mathsf{ch} \in \{0, 1\}$ and sends it to the prover.
4. If $\mathsf{ch} = 0$, the prover measures in the computational basis and sends (b, x). If $\mathsf{ch} = 1$, the prover measures in the Hadamard basis and sends (c, d).
5. The verifier checks:
   - If $\mathsf{ch} = 0$, that $f_b(x) = y$.
   - If $\mathsf{ch} = 1$, using td to compute (x_0, x_1) and verifying
   $$c = d \cdot (x_0 \oplus x_1) \pmod 2.$$

   The verifier outputs YES iff the check passes.

3.3 Why the (Honest) Quantum Prover Succeeds

The case when $\mathsf{ch} = 0$ is straightforward, so we focus on the case $\mathsf{ch} = 1$ instead.

Suppose the prover has reached the post-measurement state

$$|\psi\rangle = \frac{1}{\sqrt{2}}\big(|0, x_0\rangle + |1, x_1\rangle\big),$$

where $(x_0, x_1) \in \{0, 1\}^n \times \{0, 1\}^n$ are the unique preimages such that $f_0(x_0) = f_1(x_1) = y$.

The prover now measures this state in the Hadamard basis on all $n + 1$ qubits. Equivalently, the prover applies $H^{\otimes(n+1)}$ and measures in the computational basis. Recall that for any $z \in \{0, 1\}^m$,

$$H^{\otimes m}\, |z\rangle = \frac{1}{\sqrt{2^m}} \sum_{w\in\{0,1\}^m} (-1)^{w\cdot z}\, |w\rangle\,.$$

Applying $H^{\otimes(n+1)}$ to $|\psi\rangle$ yields

$$\frac{1}{\sqrt{2}}\left(H^{\otimes(n+1)}|0,x_0\rangle + H^{\otimes(n+1)}|1,x_1\rangle\right)$$
$$= \frac{1}{\sqrt{2}}\left(\frac{1}{\sqrt{2^{n+1}}}\sum_{c\in\{0,1\}}\sum_{d\in\{0,1\}^n}(-1)^{c\cdot 0+d\cdot x_0}|c,d\rangle + \frac{1}{\sqrt{2^{n+1}}}\sum_{c\in\{0,1\}}\sum_{d\in\{0,1\}^n}(-1)^{c\cdot 1+d\cdot x_1}|c,d\rangle\right)$$
$$= \frac{1}{\sqrt{2^{n+2}}}\sum_{c\in\{0,1\}}\sum_{d\in\{0,1\}^n}\left((-1)^{d\cdot x_0}+(-1)^{c+d\cdot x_1}\right)|c,d\rangle\,.$$

We now analyze when the amplitude of $|c,d\rangle$ is nonzero. Factoring out $(-1)^{d\cdot x_0}$, we get

$$(-1)^{d\cdot x_0}\left(1+(-1)^{c+d\cdot(x_0\oplus x_1)}\right).$$

The expression in the parentheses equals 2 if $c+d\cdot(x_0\oplus x_1)\equiv 0 \pmod 2$, and and equals 0 otherwise. Therefore, measurement outcomes (c,d) occur with nonzero probability if and only if

$$c = d\cdot(x_0\oplus x_1) \pmod 2.$$

Conditioned on this constraint, all valid outcomes (c,d) are equally likely.

Verifier Check. When $\mathsf{ch}=1$, the verifier uses the trapdoor td to compute the unique pair (x_0,x_1) satisfying $f_0(x_0)=f_1(x_1)=y$, and checks whether the prover's message (c,d) satisfies

$$c = d\cdot(x_0\oplus x_1) \pmod 2.$$

As shown above, an honest quantum prover always produces outcomes satisfying this condition. Hence the verifier accepts with probability 1 in the $\mathsf{ch}=1$ case.

Why this tests quantumness. An honest quantum prover can answer either challenge perfectly; though nver both challenges at once. This is because quantum measurements are destructive: the prover cannot simultaneously answer the pre-image test and the equation test simultaneously.

A classical prover, by contrast, can be rewound. One can show that any classical strategy that succeeds with high probability must necessarily be able to also answer *both* challenges. Using a careful extension of the protocol (which we do not cover in the lecture, see [KMCVY22]), one can turn this insight into fully-fledged reduction: the ability to answer both challenges yields a claw (x_0,x_1), thereby violating the claw-free property. Thus, under standard cryptographic assumptions, success in this protocol constitutes compelling evidence of genuine quantum behavior.

References

[AZ24] Scott Aaronson and Yuxuan Zhang. On verifiable quantum advantage with peaked circuit sampling, 2024. 1

[BCM+21] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device, 2021. 3

[KMCVY22] Gregory D. Kahanamoku-Meyer, Soonwon Choi, Umesh V. Vazirani, and Norman Y. Yao. Classically verifiable quantum advantage from a computational bell test. *Nature Physics*, 18(8):918–924, August 2022. 3, 6