Quantum Commitments

Quantum Commitments: schemes where honest parties use quantum capabilities to commit

Why study quantum commitments?

1) Quantum commitments can be based on potentially weaker assumptions than one-way functions (which are necessary and sufficient classically)

2) If we want to commit to a quantum state, quantum commitments are necessary.

3) Open-ended: Quantum commitments could have uniquely quantum security properties.
Ex: Can we construct uncloneable commitments?

We won't cover ③ today, but it might be a good research project.

Example: Removing Interaction From Classical Commitments

Consider a keyed classical commitment:

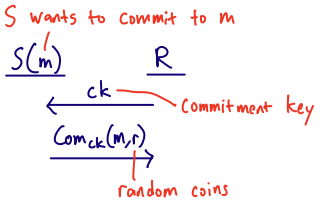

Recall that keys are sometimes necessary for binding to be possible, e.g., if Com is compressing or statistically hiding.

Claim: If S can send qubits, no interaction is needed.
High-level idea: S commits w.r.t. a uniform superposition of all ck,r and R verifies the superposition later.

Notation: C = commitment, D = decommitment

S(m) → C → R

Prepare the bipartite state

$$|\phi_m\rangle_{DC} := \sum_{ck,r} |ck,r\rangle_D |ck, \text{Com}_{ck}(m,r)\rangle_C$$

and send the C part.

To open, reveal m and send D. (m, D →) Check that the state is $|\phi_m\rangle$.

Hiding: Receiver sees $ck, \text{Com}_{ck}(m,r)$ for random ck,r

Binding (intuition): The original scheme is binding if ck is random. Since the receiver checks that the final state is

$$|\phi_m\rangle = \sum_{ck,r} |ck,r\rangle_D |ck, \text{Com}_{ck}(m,r)\rangle_C$$

the sender is forced to use a uniform superposition of ck.

In general, we can remove interaction from any interactive quantum/classical commitment by having the sender run the commit phase locally [BB21, Y22, GJMZ23, folklore]

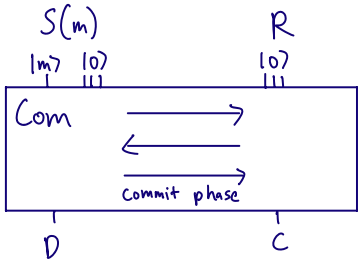

S decommits by sending D, which R can verify by running $\text{Com}^\dagger$ and measuring to see if all ancillas are 0, which checks that S followed the protocol.

So we can assume the following syntax:

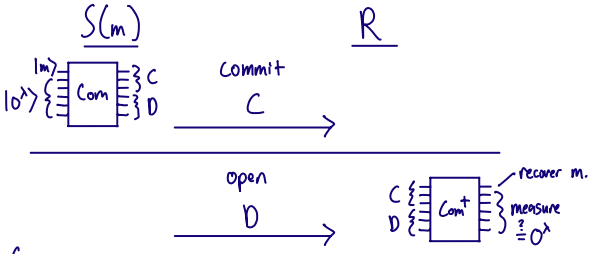

(Com is a public unitary specified by the scheme.)

Notation: ϕ_m denotes the state $\text{Com}(|m\rangle|0^\lambda\rangle)$.

Defining security can be subtle, but for now, we'll use the following definitions assuming $m \in \{0,1\}$.

- Hiding (computational): $(\phi_0)_C \approx (\phi_1)_C$ to any QPT adversary.
- Binding (statistical): $\text{TD}((\phi_0)_C, (\phi_1)_C) = 1 - \text{negl}(\lambda)$ (TD: trace distance)

Note: statistical hiding, i.e., $\text{TD}((\phi_0)_C, (\phi_1)_C) = \text{negl}$, is incompatible with statistical binding.

What Assumptions are Needed for Quantum Bit Commitment (QBC)?

Classically, one-way functions (OWFs) are necessary (and sufficient).

Intuition: If Com is comp. hiding + stat. binding, then $f(m,r) := \text{Com}(m,r)$ is a OWF (since inverting f breaks hiding).

But for QBCs, there's no obvious "hard-to-invert function" since the commitment is a quantum state.

This was recently formalized with an oracle separation.

[KQST23]: Classical oracle $\mathcal{O}$ s.t. relative to $\mathcal{O}$:
- P=NP (in particular, classical crypto doesn't exist)
- QBCs exist

/

Technically, [KQST23] showed that single-copy pseudorandom states exist, which imply QBCs.

We won't cover [KQST23] in this class.

Instead, we'll try to develop more intuition for what QBCs based on "non-OWF" assumptions might look like.

Detour: Haar-Random States and Unitaries

- Haar measure: Unique "translation-invariant" measure on unitary group. If U is a Haar-distributed random unitary, then for any fixed unitary A, the unitaries UA and AU are Haar-random.
- Haar-random state: $U|\psi\rangle$ for $U \leftarrow$ Haar and any fixed $|\psi\rangle$. Alternatively: random point on unit sphere in $\mathbb{C}^{2^n}$.
- Can sample a Haar-random state by sampling 2^n independent complex Gaussian amplitudes and normalizing.

Detour #2: The Schmidt Decomposition

What do bipartite pure states look like?

Ex: maximally entangled state on A, B

$$\sum_{i\in[N]} \frac{1}{\sqrt{N}} |i\rangle|i\rangle \quad (\dim(A)=\dim(B)=N)$$

In general, any bipartite state has a similar form:

Any $|\psi\rangle_{AB}$ can be written as $|\psi\rangle_{AB} = \sum_i \lambda_i |a_i\rangle_A |b_i\rangle_B$ where
- $\{|a_i\rangle\}_i$ and $\{|b_i\rangle\}$ are orthonormal states of A and B
- λ_i are non-negative reals satisfying $\sum_i \lambda_i^2 = 1$.

Proof: Write $|\psi\rangle = \sum_{jk} \alpha_{jk} |j\rangle_A |k\rangle_B$ + singular value decomposition on $\alpha = (\alpha_{jk})_{jk}$.

High-level idea: What if Com were a Haar-random unitary?

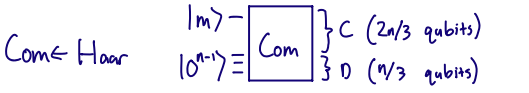

(This is not efficient, but we'll handle that later)

Assume all parties have oracle access to Com, Com†, which is enough to commit/verify.

Then $|\phi_0\rangle = \text{Com}|0^n\rangle$ and $|\phi_1\rangle = \text{Com}|1\rangle|0^{n-1}\rangle$ are two (nearly) independent Haar random states.

Consider the Schmidt decomposition of $|\phi_0\rangle$:

$$|\phi_0\rangle = \sum_{i\in[2^{n/3}]} \sqrt{p_i} \underbrace{|d_i\rangle_D}_{n/3 \text{ qubits}} \otimes \underbrace{|c_i\rangle_C}_{2n/3 \text{ qubits}}$$ where $\{|d_i\rangle\}_{i\in[2^{n/3}]}$, $\{|c_i\rangle\}_{i\in[2^{n/3}]}$ are sets of orthonormal vectors, and $\sum_i p_i = 1$.

For Haar-random $|\phi_0\rangle$: 1) Each $p_i \approx 1/2^{n/3}$
2) $\{|c_i\rangle\}$ is a set of $2^{n/3}$ Haar-random vectors in $2^{2n/3}$-dim space.

Similarly, $|\phi_1\rangle = \sum_{i\in[2^{n/3}]} \sqrt{p_i} \underbrace{|d_i'\rangle_D}_{n/3 \text{ qubits}} \otimes \underbrace{|c_i'\rangle_C}_{2n/3 \text{ qubits}}$

Let $\Pi_0 := \text{span}(|c_i\rangle)$ and $\Pi_1 := \text{span}(|c_i'\rangle)$.
Π_0, Π_1 are $2^{n/3}$-dim Haar-random subspaces in $2^{2n/3}$-dim space.

Key point: The receiver's view (either $(\phi_0)_C$ or $(\phi_1)_C$) is essentially a random state in Π_0 or Π_1.

Can verify that $TD((\phi_0)_C, (\phi_1)_C) = 1 - \text{negl}$. We won't prove this, but intuitively we'd expect Π_0, Π_1 to have tiny overlap.

Thus, binding is statistical and hiding must be computational.

Informal Claim: Breaking hiding seems <u>very</u> hard.
(i.e., distinguishing random states from Π_0 vs. Π_1)

Translation: computational hiding is a very weak assumption.

What can the adversary do?

It can try to learn Π_0, Π_1 using its oracle access to $Com, Com^\dagger$, but this takes $2^{\Omega(\lambda)}$ queries.

this isn't actually the set-up of [KQST23]

<u>Morally</u>, [KQST23] says that even if the adversary is also given an oracle that can answer arbitrary NP queries about the description of Com (i.e., the entries of the unitary), this task is still hard. (Note: such an oracle breaks any OWF)

Isn't this cheating? Haar random unitaries are exponentially complex, so this isn't a "realistic" scheme.

Fix: replace the Haar-random unitary with a $\text{poly}(\lambda)$-size random quantum circuit (based on discussions w/ Sam Gunn)

Example distributions:

- 2-qubit Haar-random gates, applied to random positions
- Random diagonal unitaries in Standard / Hadamard bases
- any fixed architecture + random (universal) gates
- many more...

Who picks the random circuit? We'll use the same trick from beginning of lecture:
The sender prepares a superposition over all circuits U of some fixed size, and the receiver checks it later.

classical description of circuit U

$$|\phi_m\rangle := \sum_U |U\rangle \underbrace{U|m,0^\lambda\rangle}|U\rangle$$

Send 2/3 of these qubits and $|U\rangle_C$

(Send the rest of $|\phi_m\rangle$ to open)

Receiver gets description of random U, and 2/3 of the qubits of $U(|m\rangle|0^\lambda\rangle)$.

Claim (won't prove): This is statistically binding.

Computational Assumption: Given the description of a random circuit U, it's computationally hard to distinguish

$$U|0^{\lambda+1}\rangle \text{ vs. } U|1,0^\lambda\rangle$$

given only 2/3 of the qubits.

[Gunn-Ma (unpublished), also mentioned in JLS18, CM22, BCQ22]

<u>Very Open</u>: Give evidence for/against hardness of this task.
So far, our evidence for hardness is that random circuits have been studied in many other fields, and seem to give rise to other hard computational tasks.
(quantum supremacy experiments, black hole radiation decoding)

Natural starting point: Consider the strongest assumption that we don't know how to break.

Strong Assumption: Same as above but U is a random t-design.