Last time: PQ soundness of Blum's Hamiltonian Cycle protocol assuming commitments are collapse-binding

Today: 1) Do collapse-binding commitments exist?
2) Is Blum ZK against quantum attacks?

Collapsing Hash Functions

Want: compressing $H: \{0,1\}^n \to \{0,1\}^{n/2}$ where $H(x)$ is collapse-binding commitment to x.

(Previously we defined C.B. for 1-bit messages, but the def easily extends to longer messages.

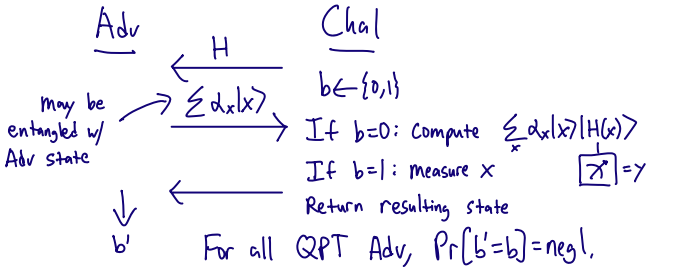

For all QPT Adv, $\Pr[b'=b] = \text{negl}$.

Today: Lossy Functions $\Rightarrow$ Collapsing hash functions

Homework: Collapsing hash $\Rightarrow$ stat. hiding, collapse-binding commitment.

Cryptographic Lossy Functions

Two function families F_{inj}, F_{lossy} where

- $f \leftarrow F_{inj}$ is injective,
- $f \leftarrow F_{lossy}$ has small image (e.g., $|\{f(x)\}_{x \in \{0,1\}^n}| = 2^{n/10}$)
- $F_{inj} \approx_c F_{lossy}$

(For both F_{inj} / F_{lossy}, $f: \{0,1\}^n \to \{0,1\}^{n+k}$)

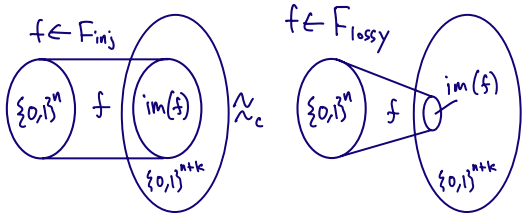

How is this related to collapsing?

Trivial: $f \leftarrow F_{inj}$ is collapsing.

Consequence: $f \leftarrow F_{lossy}$ is collapsing.

Pf: If A breaks collapsing of $f \leftarrow F_{lossy}$, then Reduction^A can distinguish $f \leftarrow F_{inj}$ vs. $f \leftarrow F_{lossy}$ by testing whether A breaks collapsing of f.

We want a shrinking function $H: \{0,1\}^n \to \{0,1\}^{n/2}$.

($f \leftarrow F_{lossy}$ is info-theoretically compressing, but output is long)

Idea: Compose with a function $h: \{0,1\}^k \to \{0,1\}^{n/2}$ that is injective on image(f) for lossy f.

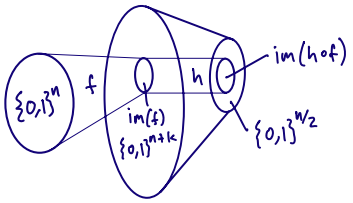

$h \circ f$ is collapsing because if adv outputs $\sum \alpha_x |x\rangle$,

Measuring x $\approx_c$ Measuring f(x) $\approx_s$ Measuring h(f(x))

(since f is collapsing) (since h is injective)

Observe: Pairwise independent h suffices!

For any $z_1 \neq z_2$ in im(f),

$$\Pr_h[h(z_1) = h(z_2)] = \frac{1}{2^{n/2}}$$

If $|\text{image}(f)| \leq 2^{n/10}$, we have $2^{n/5}$ pairs (z_1, z_2),

so by union bound:

$$\Pr_h[h \text{ is injective on im}(f)] \leq \frac{2^{n/5}}{2^{n/2}} = \text{negl}.$$

Summary: If $f: \{0,1\}^n \to \{0,1\}^k$ is a lossy function and $h: \{0,1\}^k \to \{0,1\}^{n/2}$ is pairwise indep, then $h \circ f: \{0,1\}^n \to \{0,1\}^{n/2}$ is collapsing.

What do lossy functions look like?

Learning w/ Errors (LWE) gives a simple construction

$f(x) = \lfloor A \cdot x \bmod p \rceil$ (LWE $\Rightarrow$ hard to find x)

where A is a random square matrix mod p,

$x \in \{0,1\}^n$, $\lfloor \cdot \rceil$ is rounding (drop lower order bits)

Under suitable parameters:

Injective: A is random

Lossy: A is low rank + small noise $\approx_c$ by LWE

Homework: prove this works

Open: Collapsing from other crypto assumptions (see [Zhandry22] for recent progress)

New topic: PQ Zero Knowledge of Blum

(moving on from collapsing, which is only relevant for soundness)

Recall Blum's protocol for Hamiltonicity

$P(G,\sigma)$		$V(G)$
• Permute vertices w/ random π.	many commitments $\downarrow$ c $\longrightarrow$	
• Commit to each bit of $adj(\pi(G))$ and π.	$\longleftarrow$ b	$b \leftarrow \{0,1\}$
• If $b=0$, open everything • If $b=1$, open π, σ	$\longrightarrow$ z	Accept if openings are valid and: • $b=0$: graph is $\pi(G)$ • $b=1$: valid cycle

Last time: If commitments are collapse-binding, and QPT $\tilde{P}$ convinces V w/ prob $\frac{1}{2}+\varepsilon$, then we can extract an H-cycle from $\tilde{P}$ w/ prob poly(ε).

Defining Zero Knowledge

[GMR85]: Protocol is ZK if view of any malicious $\tilde{V}$ interacting with $P(w)$ can be efficiently simulated without w.

(can be computational / statistical)

Classical ZK of Blum

(1) Show Blum is "honest verifier" ZK.
(2) Extend to full ZK.

Honest Verifier ZK (HVZK)

Protocol is HVZK if honest verifier can be simulated.
(much weaker property than full ZK)

HVZK Sim: sample b first, then pick c,z.

HV Sim

1) $b \leftarrow \{0,1\}$
2) Generate c,z:
   - If $b=0$, generate c,z using random permutation of G.
   - If $b=1$, generate c,z using random cycle graph.
3) Output (c,b,z).

By commitment hiding, $(c,b,z) \leftarrow$ HVSim looks like honest protocol execution.

Full ZK: Simulating a Malicious $\tilde{V}$

Malicious $\tilde{V}$ can pick $\tilde{b}$ adaptively based on c.

Is HVSim $\rightarrow (c,b,z)$ still useful?

Could $\tilde{V}(c)$ always output $\tilde{b}=1-b$?

No! Otherwise $\tilde{V}$ breaks hiding. So we have:

$$\Pr\left[\tilde{b}=b \;\middle|\; \substack{(c,b,z)\leftarrow HVSim \\ \tilde{b}\leftarrow \tilde{V}(c)}\right] \approx \frac{1}{2}.$$

Natural idea: repeat until $\tilde{b}=b$.

Simulator $\hat{V}$

$(c,b,z) \leftarrow$ HVSim

$c \longrightarrow$ $\tilde{V}$ $\longleftarrow \tilde{b}$ (rewind if $\tilde{b} \neq b$)

If $b=\tilde{b}$: $z \longrightarrow$ (Give up after λ rewinds)

This takes ≈ 2 tries in expectation, takes $> \lambda$ tries (fails) w/ prob $\approx \frac{1}{2^\lambda}$

Post-Quantum ZK of Blum [Watrous 09]

Goal: Simulate view of malicious quantum $\tilde{V}$ (may have initial state $|\psi\rangle$)

Even if $\tilde{V}$ is quantum:

PQ-hiding of commitments $\Rightarrow \Pr[\tilde{b}=b] \approx \frac{1}{2}$ ($\pm$ negl)

Problem: if $\tilde{b} \neq b$, how do we rewind?

Step 0: Define unitary U corresponding to
1) $(c,b,z) \leftarrow$ HVSim
2) Run $V(|\psi\rangle, c) \rightarrow \tilde{b}$

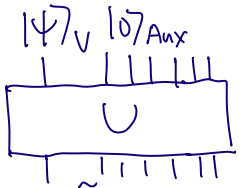

Aux contains:
- $(c,b,z) \leftarrow$ HVSim
- Workspace of HVSim

$\Pr[\tilde{b}=b] \approx \frac{1}{2}$.

Goal: Amplify this probability.

Formally: Define projector $\Pi_{eq} = \sum_b |b,b\rangle\langle b,b| \otimes I$

Want state $\Pi_{eq} U(|\psi\rangle|0\rangle_{Aux})$ (normalized)

= state conditioned on "guessing" $\tilde{b} = b$.

Plan

1. Solve simplified "2-dim" version of problem
2. Recall Jordan's Lemma
3. Full ZK sim via Jordan + crypto

Abstract task: We have efficient $\{\Pi, I-\Pi\}$.
Given $|v\rangle$ s.t. $\|\Pi|v\rangle\|^2 = p$
output $\Pi|v\rangle/\sqrt{p}$ with prob ≈ 1

(ZK Simulation: $\Pi = U^\dagger \Pi_{eq} U$, $|v\rangle = |\psi\rangle|0\rangle_{Aux}$, $p \approx 1/2$)

Easy case (similar to amplitude amplification):
Assume we can measure $\{|v\rangle\langle v|, I - |v\rangle\langle v|\}$

(this won't work for ZK)

Solution: 1) Start with $|v\rangle$

2) Alternate $|v\rangle\langle v|$ and Π binary projective measurements until a Π outcome occurs.

- Why does this work?
- How long does it take?

Let $|w\rangle = \Pi|v\rangle/\sqrt{p}$.

Observe: State never leaves 2-D subspace

$$|v\rangle = \sqrt{p}|w\rangle + \sqrt{1-p}|w^\perp\rangle$$
$$|v^\perp\rangle = \sqrt{1-p}|w\rangle - \sqrt{p}|w^\perp\rangle$$

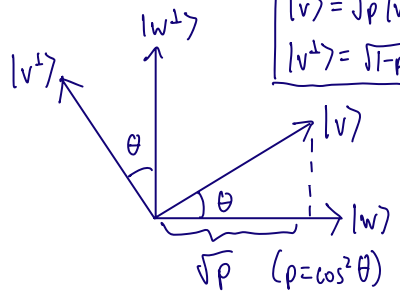

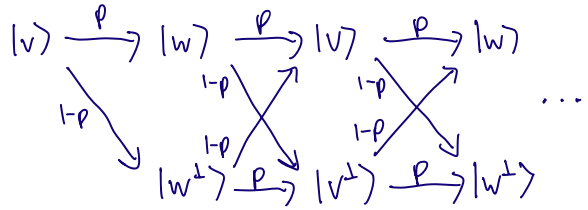

Claim: $\Pr[\text{reach } |w\rangle \text{ in } \le \frac{\lambda}{p} \text{ steps}] \approx 1$

But what if we can't measure $|v\rangle\langle v|$?

We'll replace this with some other projector (to be determined later)

Detour: Jordan's Lemma

Intuition: A pair of 1-dim subspaces (lines thru 0) intersect at an angle.

Jordan's Lemma generalizes this to higher dimensions.

Lemma:

For any pair of projectors Π_A, Π_B acting on $\mathcal{H}$, we can decompose $\mathcal{H} = \bigoplus_j S_j$ where

- each S_j is 1- or 2-dimensional
- each S_j is invariant under Π_A, Π_B.

Each 2D S_j looks like

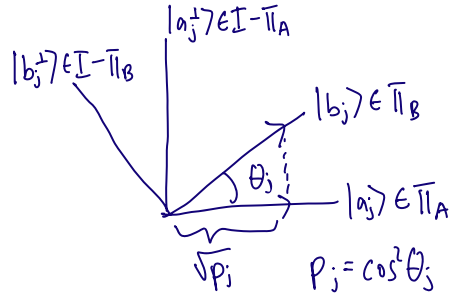

How can we pick these subspaces?

For any evector $|\psi\rangle$ of $\Pi_A \Pi_B \Pi_A$, can pick a subspace $\text{Span}(|\psi\rangle, |\phi\rangle)$ where

$$\Pi_B |\psi\rangle = \sqrt{p}\,|\phi\rangle$$

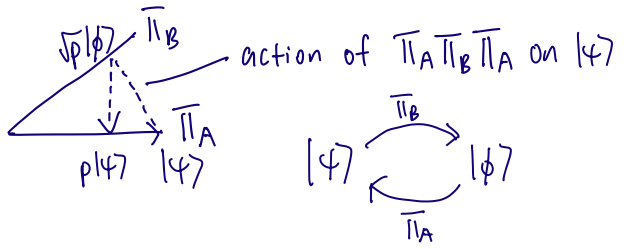

Aside: When one of Π_A, Π_B is rank 1, only one non-trivial ($p_j \neq 0$) subspace S_j,

(Since $|v\rangle$ is the only non-trivial evector of $|v\rangle\langle v| \Pi_B |v\rangle\langle v| = \langle v|\Pi_B|v\rangle \cdot |v\rangle\langle v|$)

In other words, mapping $|v\rangle$ to $\Pi|v\rangle/\sqrt{p}$ given $\{|v\rangle\langle v|, I - |v\rangle\langle v|\}$ is easy because $|v\rangle$ is an evector of $|v\rangle\langle v| \cdot \Pi \cdot |v\rangle\langle v|$.

> Observation: Suffices to implement any binary projective measurement $\{\Pi_A, I - \Pi_A\}$ where $|v\rangle$ is an evector of $\Pi_A \Pi_B \Pi_A$.

In this case, we can map $|v\rangle \to \Pi_B|v\rangle/\sqrt{p}$ by alternating Π_A, Π_B measurements.

Back to ZK Simulation

Recall: Malicious verifier has state $|\psi\rangle$, we want to output $\Pi_{eq} U(|\psi\rangle|0\rangle)$.

Idea: Pick $\Pi_A := I \otimes |0\rangle\langle 0|_{Aux}$.

Claim: If commitments are perfectly hiding, $|\psi\rangle|0\rangle_{Aux}$ is an evector of $\Pi_A \Pi_B \Pi_A$ w/ evalue $1/2$.

Pf: Let $|v_j\rangle|0\rangle$ be an evector of $\Pi_A \Pi_B \Pi_A$.

$$\langle v_j|\langle 0|\Pi_A \Pi_B \Pi_A |v_j\rangle|0\rangle = \left\| \Pi_{eq} U |v_j\rangle|0\rangle \right\|^2 = \Pr\left[\begin{array}{l}\tilde{V}(|v_j\rangle) \text{ outputs } b' = b \\ \text{given } c \text{ from } (c,b,z) \leftarrow \text{HVSim}\end{array}\right] = \frac{1}{2}$$

Since c is independent of b by hiding

Since all eigenvalues are $1/2$, any state $|\psi\rangle|0\rangle$ is an evector w/ evalue $1/2$.

So I can pick the Jordan decomposition to guarantee $|\psi\rangle|0\rangle$ is in a 2-D subspace.

Then alternating projectors works!

Recap: 1) Goal: create the state

$$\Pi_{eq} U(|\psi\rangle|0\rangle_{Aux})$$

corresponding to a successful simulation

2) Start w/ $|\psi\rangle|0\rangle$, alternate $\{U^\dagger \Pi_{eq} U, I - U^\dagger \Pi_{eq} U\}$ and $\{I \otimes |0\rangle\langle 0|_{Aux}, I - I \otimes |0\rangle\langle 0|_{Aux}\}$ measurements until a $U^\dagger \Pi_{eq} U$ outcome occurs;

(cut off after $2\lambda - 1$ measurements)

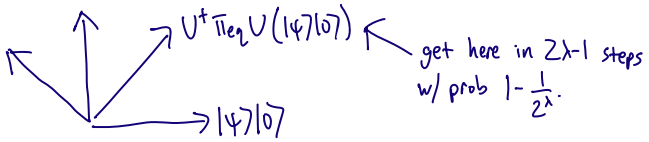

3) Apply U to get $\Pi_{eq} U(|\psi\rangle|0\rangle)$

Extension to stat/computational hiding

Intuition: $|\psi\rangle|0\rangle$ should "behave like" a $\frac{1}{2}$-evector of $\Pi_A \Pi_B \Pi_A$ when we alternate Π_A, Π_B measurements. If not, this should contradict stat/comp. hiding.

Claim: All non-zero eigenvalues of $\Pi_A \Pi_B \Pi_A$ lie in $\left[\frac{1}{2} - \text{negl}, \frac{1}{2} + \text{negl}\right]$.

Pf: Suppose not. Then there exists $|\psi\rangle|0\rangle \in \text{image}(\Pi_A)$ s.t. $\left| \|\Pi_B|\psi\rangle|0\rangle\|^2 - \frac{1}{2} \right| > \frac{1}{\text{poly}}$.

Then $\tilde{V}(|\psi\rangle)$ can guess b w/ $\frac{1}{\text{poly}}$ advantage, which breaks hiding.

We'll show that the same simulator as in the perfect hiding case still works to get negl-accuracy simulation.

High level idea: Sim is close to Sim' where Sim' behaves like the simulator in the perfect case.

Define "close" projector Π_B' s.t. $|\psi\rangle|0\rangle$ is a $\frac{1}{2}$-evector of $\Pi_A \Pi_B' \Pi_A$.

Construction: For each (Π_A, Π_B)-subspace S_j, define $|b_j'\rangle := \frac{1}{\sqrt{2}}(|a_j\rangle + |a_j^\perp\rangle)$.

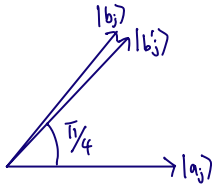

Let $\Pi_B' := \sum_j |b_j'\rangle\langle b_j'|$.

Note: $\langle a_j|b_j\rangle \approx \langle a_j|b_j^\perp\rangle = \frac{1}{\sqrt{2}}$ and $\langle b_j|b_j'\rangle \approx 1$

Claim: $\|\Pi_B - \Pi_B'\|_{op} = \text{negl}(\lambda)$.

(Note: Π_B' does not correspond to an efficient measurement.)

By construction, $|\psi\rangle|0\rangle$ is a $\frac{1}{2}$-evector of $\Pi_A \Pi_B' \Pi_A$.

Let $A := \{\Pi_A, I - \Pi_A\}$, where Π_A corresponds to "accept." (same for B, B')

Sim: Alternate $B, A, B, A, \ldots$ until B accepts

Sim': Alternate $B', A, B', A, \ldots$ until B' accepts (give up after $2\lambda - 1$ steps)

Let $p := \|\Pi_B|\psi\rangle|0\rangle\|^2$, $p' := \|\Pi_{B'}|\psi\rangle|0\rangle\|^2$. Note $p \approx p' = \frac{1}{2}$

Want: Sim maps $|\psi\rangle|0\rangle$ to $\approx \frac{\Pi_B|\psi\rangle|0\rangle}{\sqrt{p}}$. (this implies valid ZK sim)

(where $\approx$ means negl close)

Claim 1: Sim' output $\approx_{\text{negl}} \Pi_{B'}|\psi\rangle|0\rangle / \sqrt{p'}$

Claim 2: $\frac{\Pi_{B'}|\psi\rangle|0\rangle}{\sqrt{p'}} \approx_{\text{negl}} \frac{\Pi_B|\psi\rangle|0\rangle}{\sqrt{p}}$

Claim 3: Sim output $\approx$ Sim' output.

Pf of 1:

$|\psi\rangle|0\rangle$ is $\frac{1}{2}$ evector of $\Pi_A, \Pi_{B'}$, so only source of "error" is truncation at $2\lambda - 1$ steps.

$\Pr[\text{fail to get } \Pi_{B'} \text{ output in } 2\lambda - 1 \text{ steps}] = \frac{1}{2^\lambda}$.

Pf of 2:

Suffices to show inner product is ≈ 1. Write $|\psi\rangle|0\rangle = \sum_j \alpha_j |a_j\rangle$.

$\Pi_B|\psi\rangle|0\rangle = \sum_j \alpha_j \langle b_j|a_j\rangle |b_j\rangle \approx \sum_j \alpha_j \frac{1}{\sqrt{2}}|b_j\rangle$

$\Pi_{B'}|\psi\rangle|0\rangle = \sum_j \alpha_j \langle b_j^\perp|a_j\rangle |b_j^\perp\rangle = \sum_j \alpha_j \cdot \frac{1}{\sqrt{2}}|b_j^\perp\rangle$

$\Rightarrow \frac{\langle\psi|\langle 0|\Pi_{B'}\Pi_B|\psi\rangle|0\rangle}{\sqrt{p'}\sqrt{p}} \approx 1$

Pf of 3: Define $U_B = X_W \otimes \Pi_B + I_W \otimes (I - \Pi_B)$.

(i.e., write measurement outcome onto W)

Sim and Sim' can be coherently implemented as a sequence of $O(\lambda)$ unitaries that only differ by $U_B / U_{B'}$.

Claim: $\|U_B - U_{B'}\|_{op} = \text{negl}$.

($\Rightarrow$ for any $|\psi\rangle$, $U_B|\psi\rangle$ is negl-close to $U_{B'}|\psi\rangle$)

$\|U_B - U_{B'}\|_{op} \leq 2\|\Pi_B - \Pi_{B'}\|_{op} = \text{negl}$