Security Reductions in A Quantum World

Mark Zhandry (Princeton & NTT Research)

Security Proofs

Crypto security "proof" **=** Computational Assumption Ꝑ **+** Reduction from Ꝑ

Should be well-studied and widely believed

Concrete assumptions: Hardness of Factoring, DLog, LWE

Generic assumptions: ∃ OWF, ∃ PKE

In other words, if you can break scheme, you can solve Ꝑ

Enter Quantum

Thm [Shor'94]: $\exists$ Quantum polynomial time (QPT) algorithms solving FACTORING, DLOG

↓

Post-Quantum Crypto = developing crypto secure against quantum attacks

Post-Quantum Security Proofs

Post-quantum security "proof" **=** *Post-quantum* Assumption Ꝑ **+** *Post-quantum* Reduction

Should be well-studied and widely believed
Concrete assumptions: (Quantum) hardness of LWE, …
Generic assumptions: ∃ (quantum immune) OWF, PKE

If you can break scheme *with a quantum computer*, then you can solve Ꝑ *with a quantum computer*

Main Takeaway

Post-quantum security "proof"	≠	*Post-quantum* Assumption P	+	**Classical Reduction**

BAD NEWS: Most crypto literature = classical reduction	**GOOD NEWS:** Most results translate to quantum trivially	**BUT:** ∃ notable exceptions

Even those working with post-quantum tools

Outline for Today

1st hour: 4 illustrative examples

- Increasing PRG stretch – black box reductions
- PRFs – interaction
- Coin tossing – rewinding
- Goldreich-Levin – running adversary many times

2nd hour: Begin seeing new post-quantum techniques

Example 1: PRG Length Extension

$x \in \{0,1\}^n$

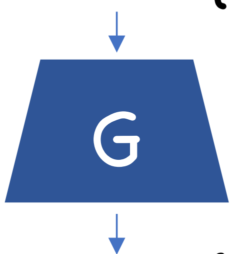

$y \in \{0,1\}^m$

(m>n)

Def: G is a secure pseudorandom generator (PRG) if, $\forall$ PPT A, $\exists$ negligible ε such that

$$| \Pr[A(y)=1] - \Pr[A(G(x))=1] | < \varepsilon$$

ε called "advantage" of A

Example 1: PRG Length Extension

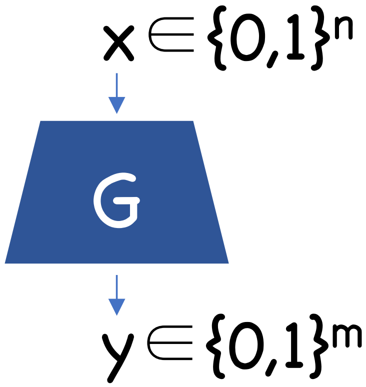

Suppose **m=n+1**. How to get larger stretch?

Solution: G_2 =

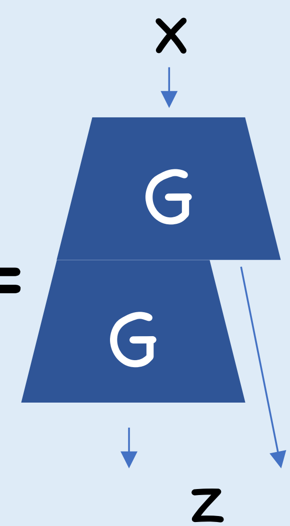

Thm: If G is secure, then so is G_2

Example 1: PRG Length Extension

Proof: Suppose G_2 insecure. Then $\exists$ PPT A, non-negl ε such that

$$| \Pr[A(y)=1] - \Pr[A(G_2(x))=1] | \geq \varepsilon$$

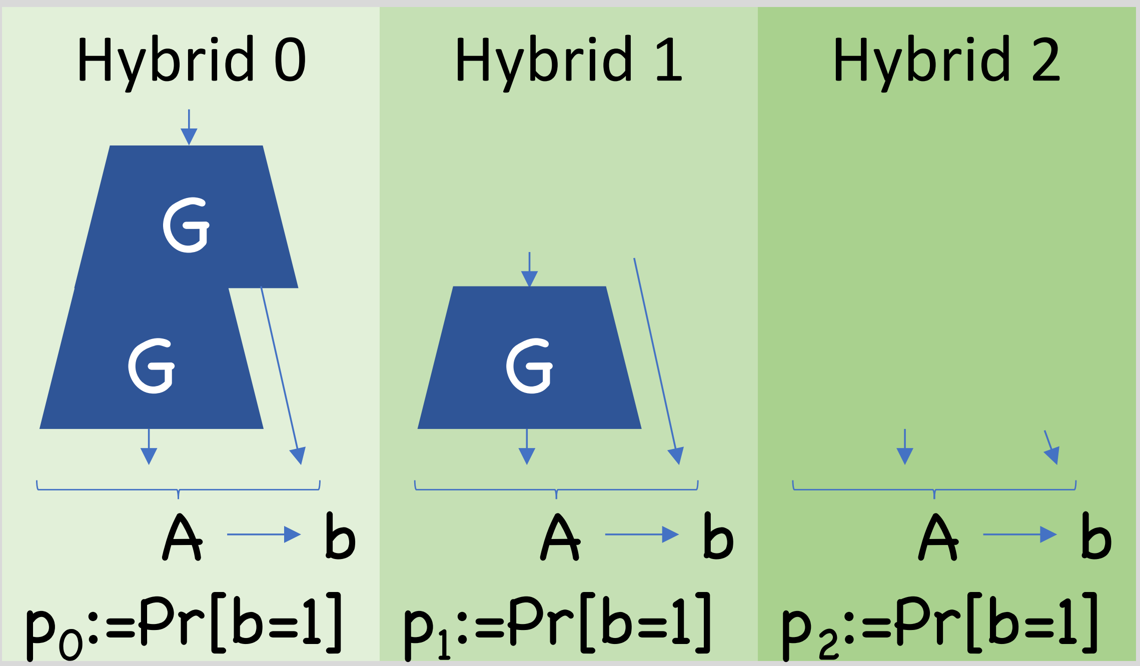

Example 1: PRG Length Extension

Proof: Suppose G_2 insecure. Then $\exists$ PPT A, non-negl ε such that

$$| p_2 - p_0 | \geq \varepsilon$$

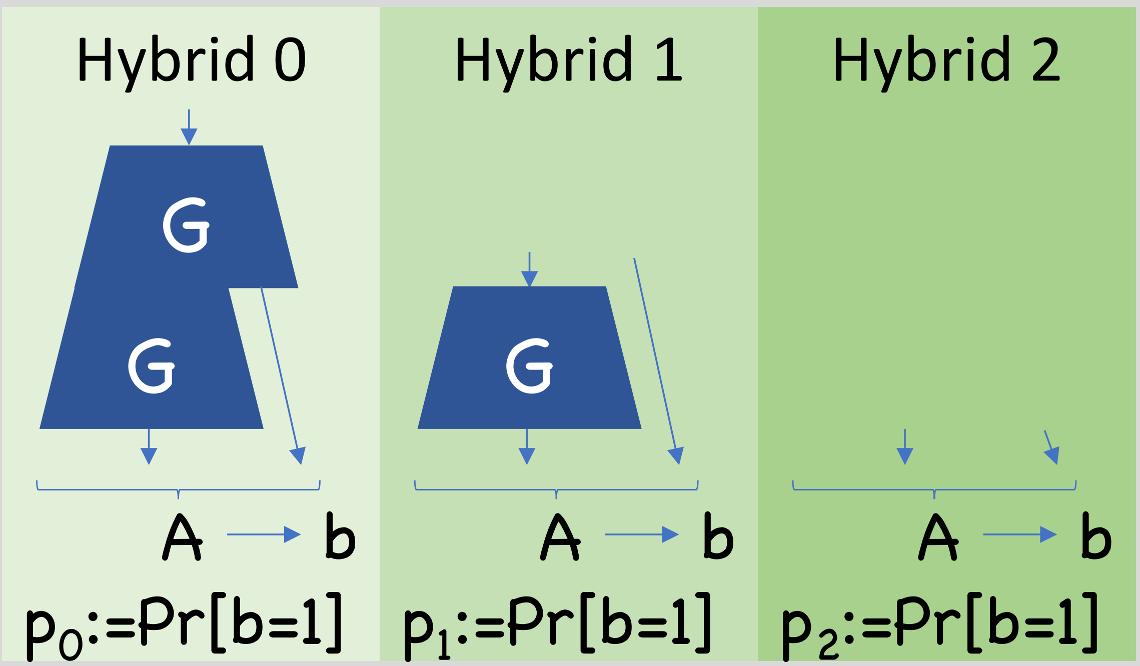

Either: $|p_1 - p_0| \geq \varepsilon/2$

$B(y_0,y_1) = A(G(y_0),y_1)$

Or: $|p_2 - p_1| \geq \varepsilon/2$

$B(y_0,y_1) = A(y_0,y_1,\$)$

In either case, B has advantage ε/2 against security of G

Example 1: PRG Length Extension

What about quantum?

Def: G is a **post-quantum** secure PRG if,
$\forall$ **Q**PT A, $\exists$ negligible ε such that
$$| \Pr[A(y)=1] - \Pr[A(G(x))=1] | < \varepsilon$$

Thm: If G is post-quantum secure, then so is G_2

Example 1: PRG Length Extension

Proof: Suppose G_2 **PQ** insecure. Then $\exists$ **Q**PT A, non-negl ε s.t.

$$| p_2 - p_0 | \geq \varepsilon$$

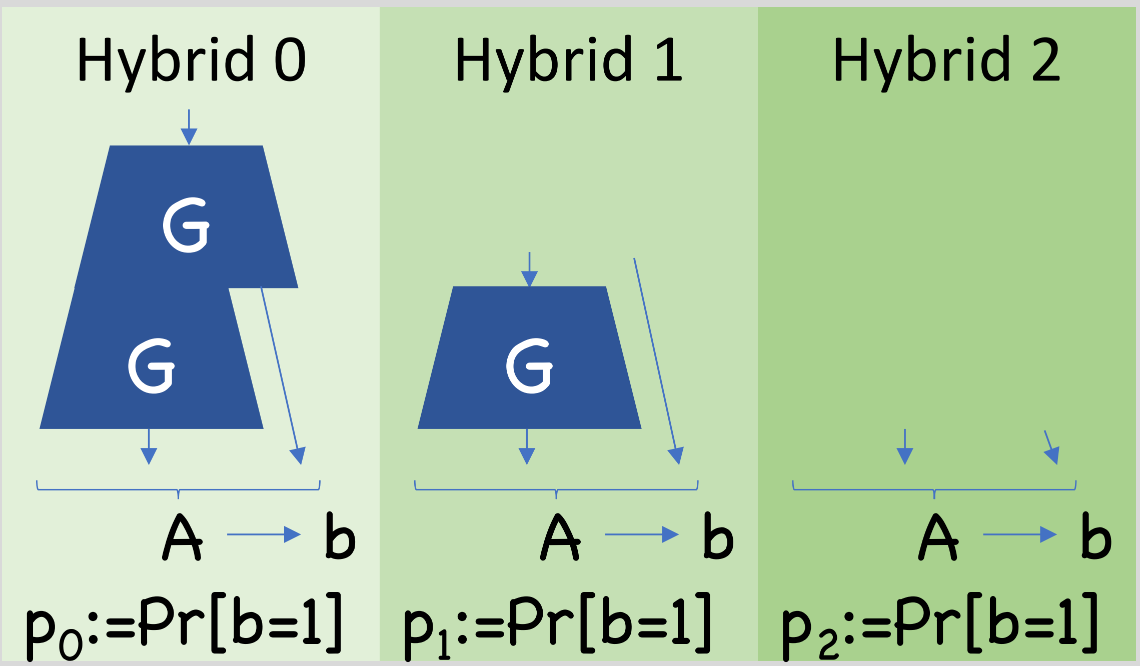

Either:
$|p_1 - p_0| \geq \varepsilon/2$

$B(y_0,y_1) = A(G(y_0),y_1)$

Or:
$|p_2 - p_1| \geq \varepsilon/2$

$B(y_0,y_1) = A(y_0,y_1,\$)$

In either case, B has advantage ε/2 against **PQ** security of G

Example 1: PRG Length Extension

Proof for G_2 doesn't care how A works internally, as long as it has non-negligible advantage

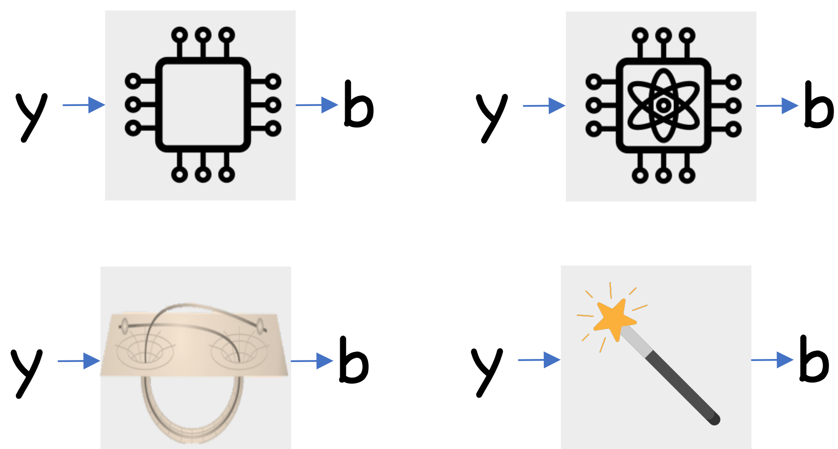

That is, proof treats A as "black box"

Example 1: PRG Length Extension

Key Takeaway: As long as reduction treats A as a *non-interactive single-run* black box, reduction likely works in quantum setting

Example 2: PRFs

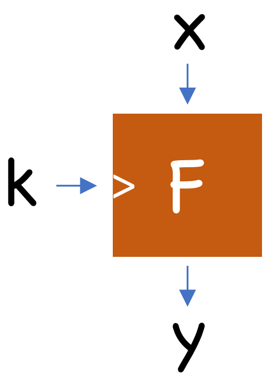

Def: F is a secure pseudorandom function (PRF) if, $\forall$ PPT A, $\exists$ negligible ε such that

$$| \Pr[A^{F(k,\cdot)}()=1] - \Pr[A^{R(\cdot)}()=1] | < \varepsilon$$

Notes:

- k random
- R uniformly random function
- $A^{O(\cdot)}$ means A makes queries on x, receives O(x)

Example 2: PRFs

What is a post-quantum PRF?

$A^{|O(\cdot)\rangle}$ means quantum queries:

$\Sigma \alpha_{x,y} |x,y\rangle$

↓

$\Sigma \alpha_{x,y} |x, y \oplus O(x)\rangle$

Def: F is a **PQ** secure PRF if, $\forall$ **Q**PT A, $\exists$ negligible ε such that

$$| \Pr[A^{F(k,\cdot)}()=1] - \Pr[A^{R(\cdot)}()=1] | < \varepsilon$$

Def: F is a **Fully Quantum** secure PRF if, $\forall$ **Q**PT A, $\exists$ negligible ε such that

$$| \Pr[A^{|F(k,\cdot)\rangle}()=1] - \Pr[A^{|R(\cdot)\rangle}()=1] | < \varepsilon$$

Example 2: PRFs

Is there a difference? YES!

Proof: Embed Simon's oracle/period finding

$$PRF'((k,z), x) = PRF(k, \{x, x \oplus z\})$$

Example 2: PRFs

Ok. Which definition do we want?

It depends

Example 2a: PRFs → CPA-secure encryption

$$\mathrm{Enc}(k,m) = \begin{array}{l} r \leftarrow \$ \\ c = (r, F(k,r) \oplus m) \end{array}$$

Encrypter (honest) chooses r → always classical

PQ security suffices

Example 2: PRFs

Ok. Which definition do we want?

It depends

Example 2b: PRFs → MAC

$$MAC(k,m) = F(k,m)$$

Security model lets attacker choose m, but signer (honest) actually computes MAC

Can attacker force signer to MAC superpositions?

Example 2: PRFs

Ok. Which definition do we want?

It depends

Example 2c: PRFs → Pseudorandom quantum states

[Ji-Liu-Song'18,Brakerski-Shmueli'19]

$$\sum_x (-1)^{F(k,x)} |x\rangle$$

Generation of state makes superposition query to F

Need full quantum security

Example 2: PRFs

So then, what does a classical proof give us?

Example 2: PRFs

PRG→PRF

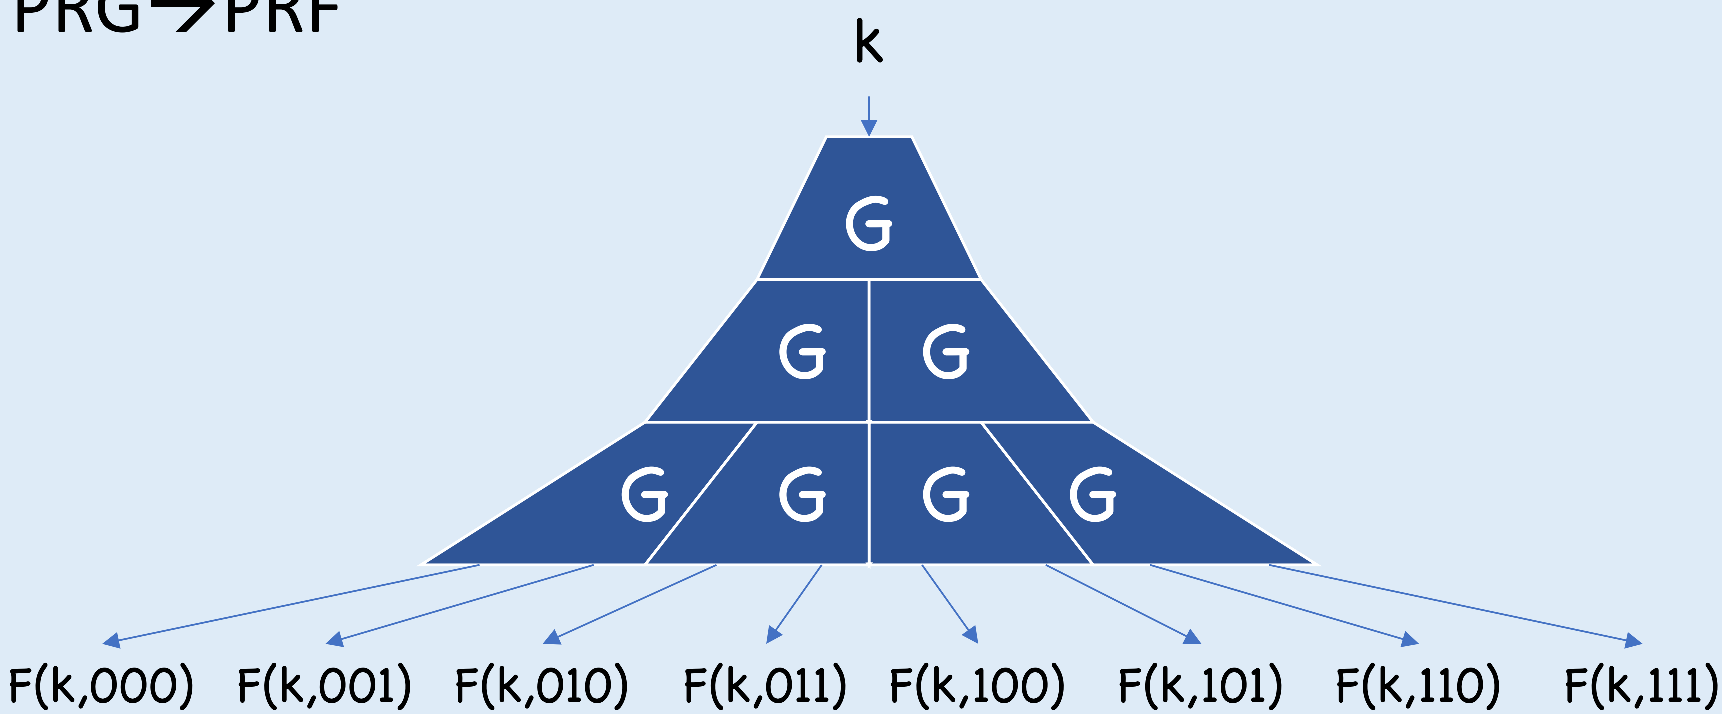

Example 2: PRFs

PRG→PRF

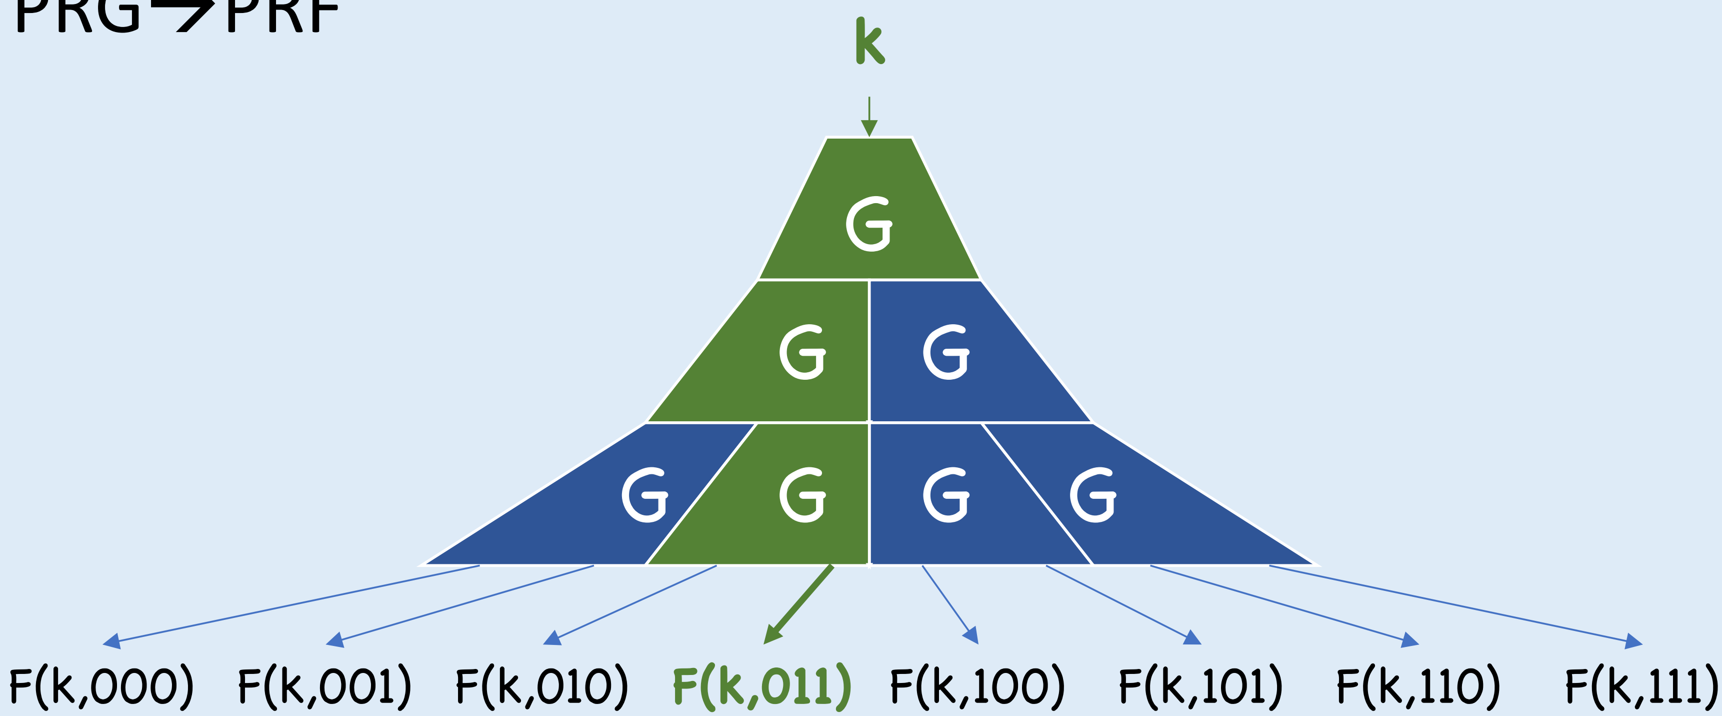

Example 2: PRFs

Classical proof, step 1: Hybrid

Hybrid 0 (F(k, ·)):

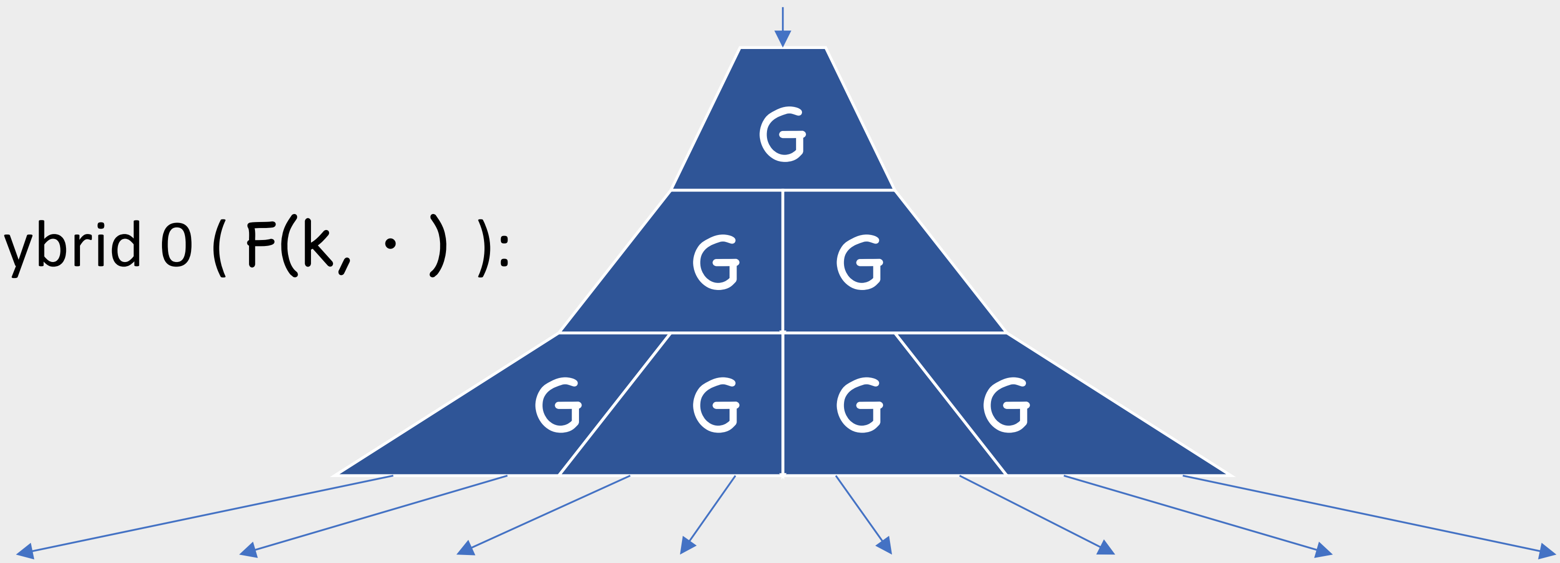

Example 2: PRFs

Classical proof, step 1: Hybrid

Hybrid 1:

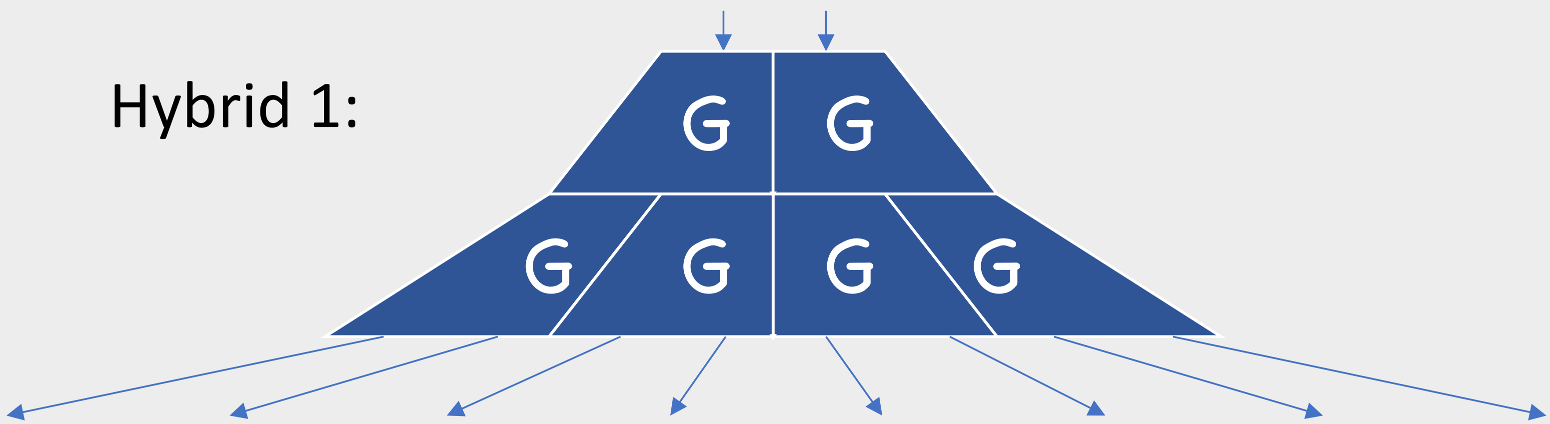

Example 2: PRFs

Classical proof, step 1: Hybrid

Hybrid 2:

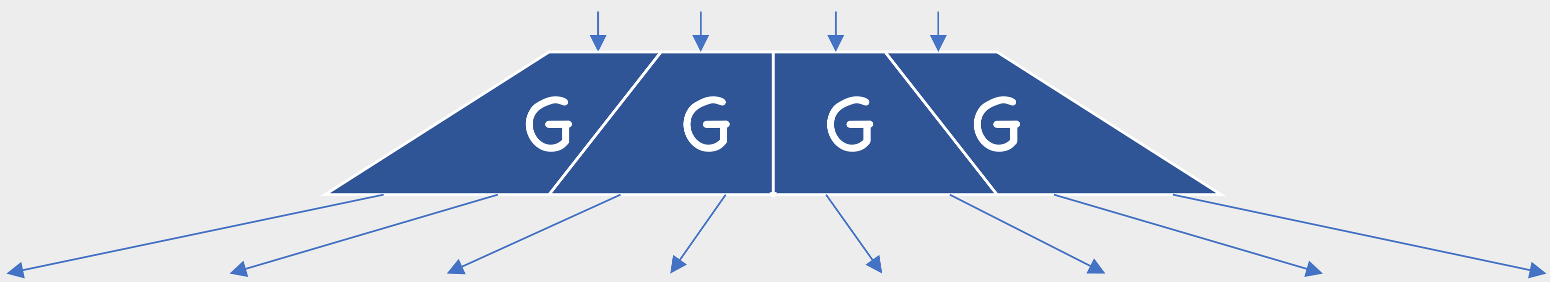

Example 2: PRFs

Classical proof, step 1: Hybrid

Hybrid n ($R(\cdot)$):

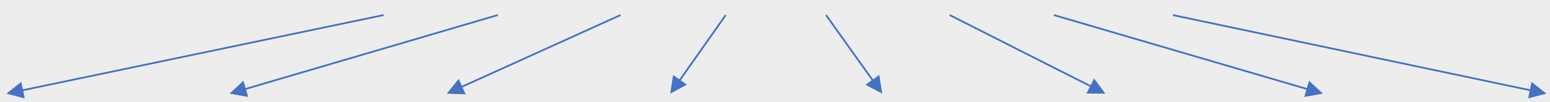

Example 2: PRFs

Classical proof, step 1: Hybrid

$$\exists i \text{ s.t. } | \Pr[A^{\text{Hybrid } i+1}() = 1] - \Pr[A^{\text{Hybrid } i}() = 1] | \geq \varepsilon/n$$

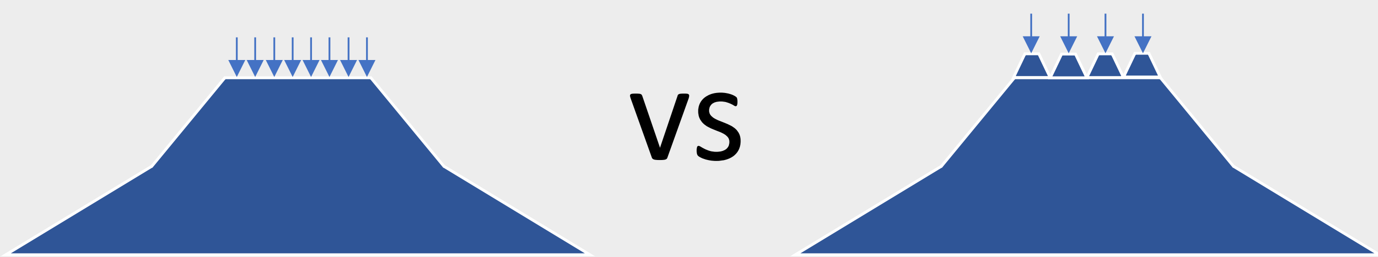

Step 1 makes sense if A classical, post-quantum, or fully quantum

Example 2: PRFs

Classical proof, step 2: Another hybrid

Hybrid i.0 = i:

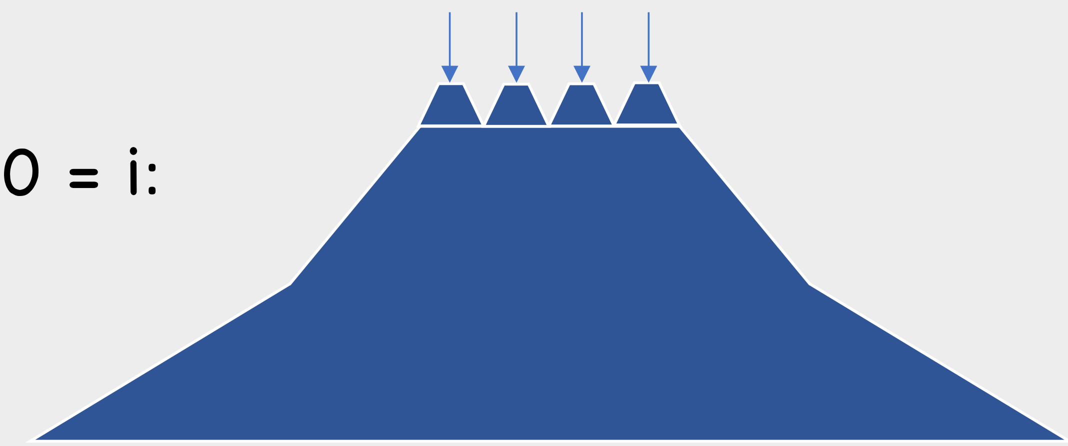

Example 2: PRFs

Classical proof, step 2: Another hybrid

Hybrid i.1:

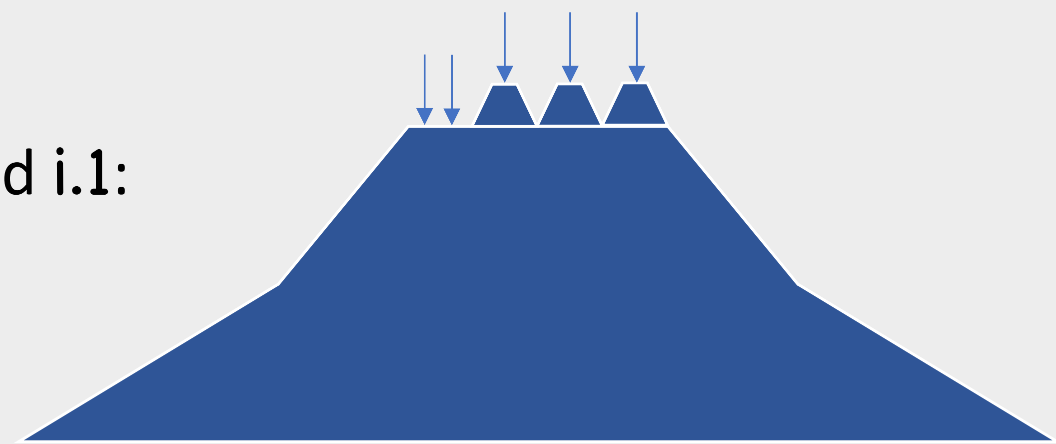

Example 2: PRFs

Classical proof, step 2: Another hybrid

Hybrid i.2:

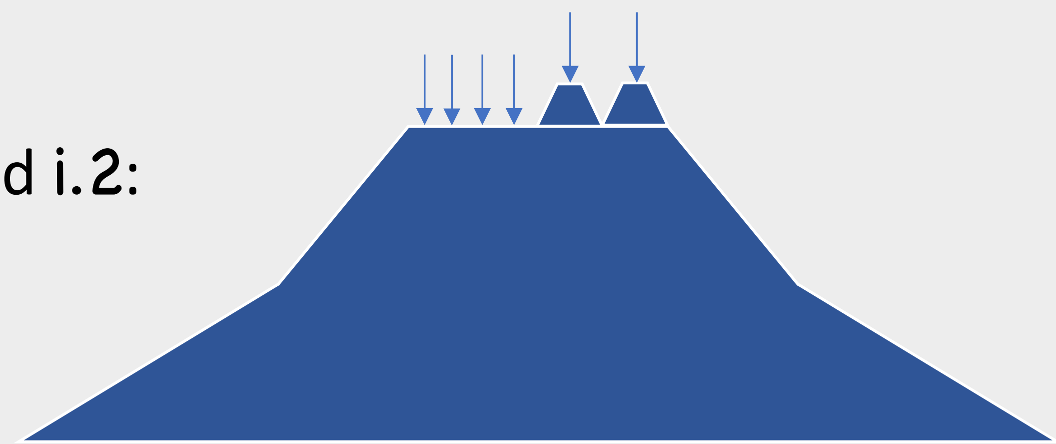

Example 2: PRFs

Classical proof, step 2: Another hybrid

Hybrid i.3:

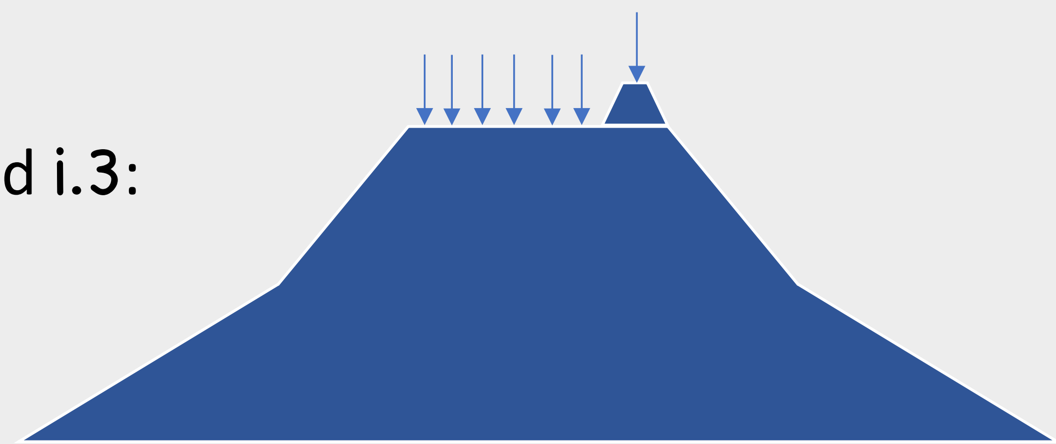

Example 2: PRFs

Classical proof, step 2: Another hybrid

Hybrid $i.2^i = i+1$:

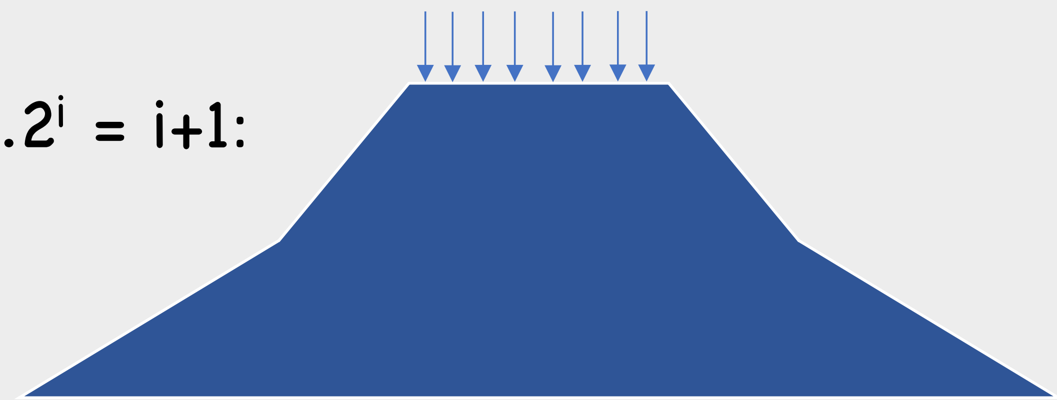

Problem: 2^i loss potentially exponential

Example 2: PRFs

Classical proof, step 2: Another hybrid

Solution: lazy/on-the-fly sampling

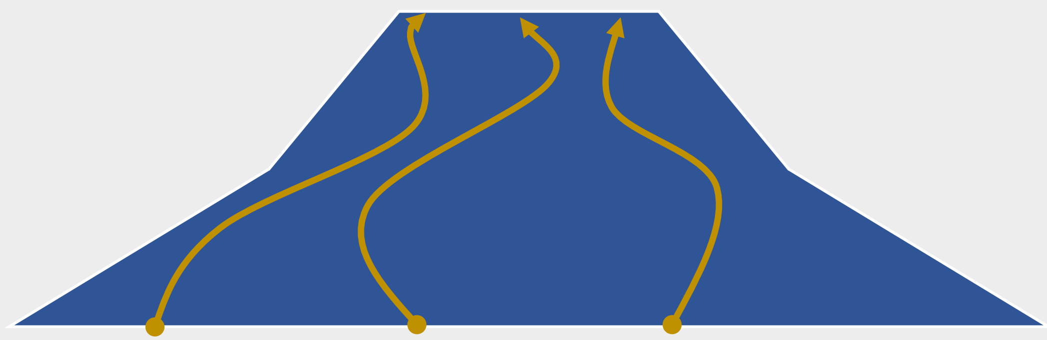

q queries → Only hybrid over **q** “active” positions

Example 2: PRFs

Proof doesn't care how A works internally, as long as it has non-negligible advantage

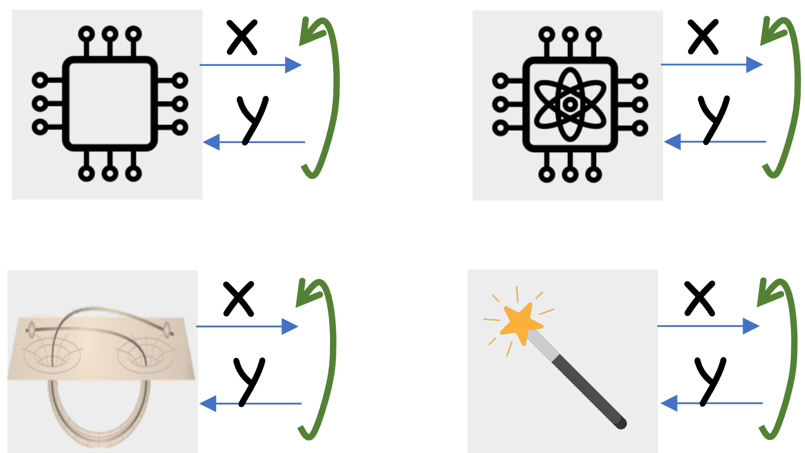

➔ Also post-quantum reduction

Example 2: PRFs

What about full quantum security?

Even single query touches **everything**

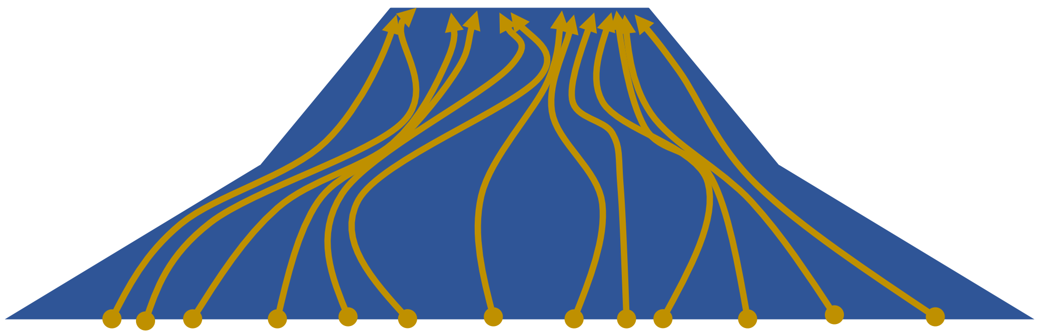

Lazy sampling?

Embedding challenges?

Example 2: PRFs

What about full quantum security?

Classical proof is black box, but requires classical queries

Can the proof be fixed for full quantum security?

Topic for 2nd hour…

Example 2: PRFs

Key Takeaway: As long as reduction treats $\mathcal{A}$ as a *single-run* black box (potentially w/ *classical* interaction), reduction likely works in quantum setting

! But if interaction is quantum, all bets are off

Example 3: Coin Tossing

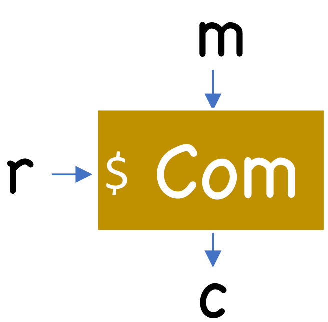

Def: *Com* is (computationally) binding if, $\forall$ PPT *A*, $\exists$ negligible ε such that

$$\Pr\left[\begin{array}{c} m_0 \neq m_1 \wedge \\ Com(m_0,r_0)=Com(m_1,r_1) \end{array} : (m_0,r_0,m_1,r_1) \leftarrow A()\right] < \varepsilon$$

Also want hiding, but we will ignore

Example 3: Coin Tossing

Simple protocol:

$b_A \leftarrow \{0,1\}$

$r \leftarrow \$$

$c = com(b_A,r)$ →

← b_B

b_A,r →

$b_B \leftarrow \{0,1\}$

Verify $c = com(b_A,r)$

pass: $b = b_A \oplus b_B$

fail: $b = \perp$

Example 3: Coin Tossing

Proof that Alice can't bias **b**:

Let $\mathcal{A}$ be supposed adversary

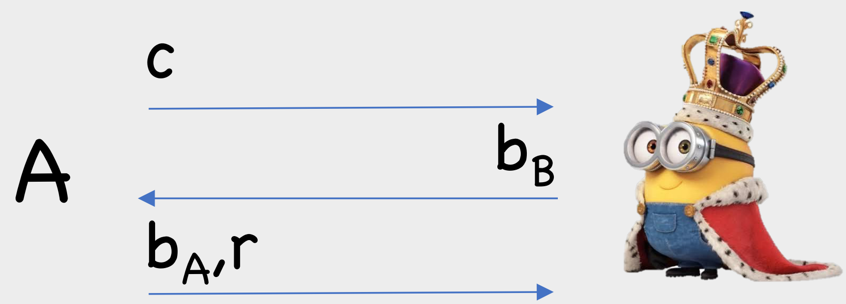

$\Pr[b=0] > \frac{1}{2}+\varepsilon$ ➡ For both $b_B=0$ and $b_B=1$, good chance $b_A=b_B$ and $\text{Com}(b_A,r)=c$

Example 3: Coin Tossing

Proof that Alice can't bias b:

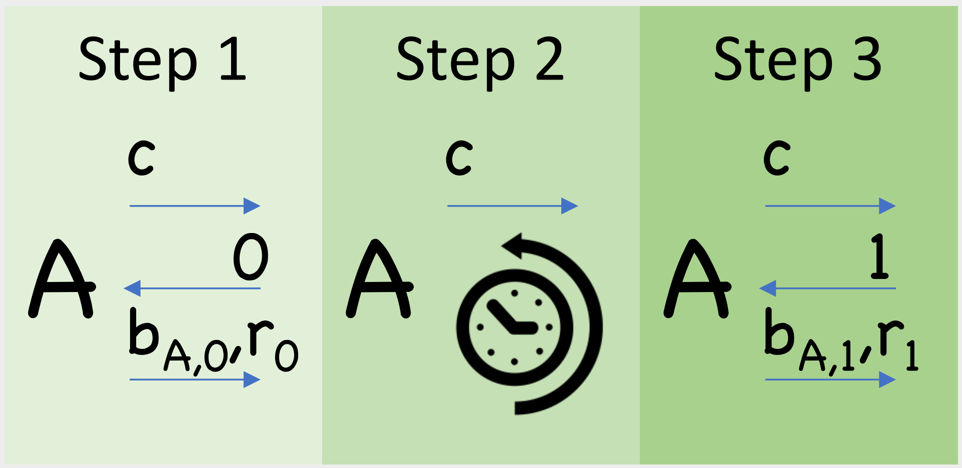

$$\Pr\left[\begin{array}{c} b_{A,0} = 0 \ \wedge \ b_{A,1} = 1 \ \wedge \\ Com(b_{A,0},r_0) = Com(b_{A,1},r_1) = c \end{array}\right] \geq poly(\varepsilon)$$

Example 3: Coin Tossing

What if A is quantum?

Def: *Com* is **post-quantum** (computationally) binding if, $\forall$ **Q**PT A, $\exists$ negligible ε such that

$$\Pr\left[\begin{array}{c} m_0 \neq m_1 \wedge \\ Com(m_0,r_0)=Com(m_1,r_1) \end{array} : (m_0,r_0,m_1,r_1) \leftarrow A()\right] < \varepsilon$$

Define coin-tossing goal similarly

Note: adversary's interaction unchanged (unlike Ex 2)

Example 3: Coin Tossing

Proof that **quantum** Alice can't bias **b**?

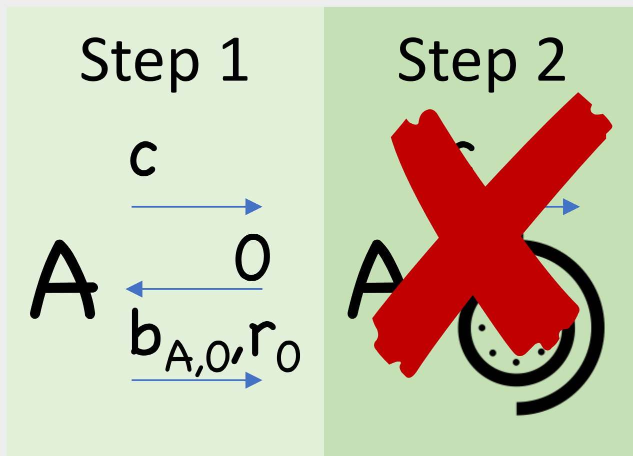

Measurement principle: extracting $b_{A,0}, r_0$ irreversibly altered A's state

Example 3: Coin Tossing

Thm (Ambainis-Rosmanis-Unruh'14,Unruh'16):
$\exists$ PQ binding *Com* s.t. Alice has a near-perfect strategy

I.e., quantumly, ability to produce either of two values isn't the same as ability to produce both simultaneously

Example + how to overcome topic for tomorrow

Example 3: Coin Tossing

Key Takeaway: As long as reduction treats A as a *single-run* black box (potentially w/ *classical* interaction), reduction likely works in quantum setting

! But if interaction is quantum, all bets are off

! But if rewinding A, all bets are off

Example 4: Goldreich-Levin

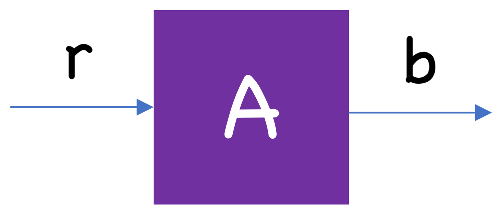

Stateless/rewindable

"GL assumption": A is PPT, $\exists x$: $\Pr[A(r) = \langle r,x \rangle] \geq \frac{1}{2} + \varepsilon$

Thm:

$\exists$ PPT

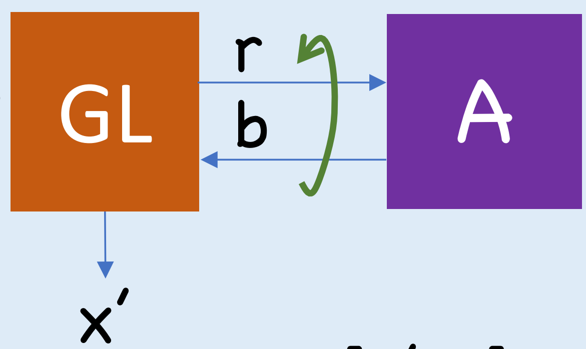

$\Pr[x'=x] \geq \delta = \text{poly}(\varepsilon)$

Example 4: Goldreich-Levin

What happens in quantum setting?

Proof of GL doesn't care how A works internally, as long as "GL Assumption" holds for **all** queries

A has classical description (even if quantum alg.) ✓

Good enough for most applications, e.g. OWF → PRG [HILL'99]

But what if A contains quantum state?

Example 4: Goldreich-Levin

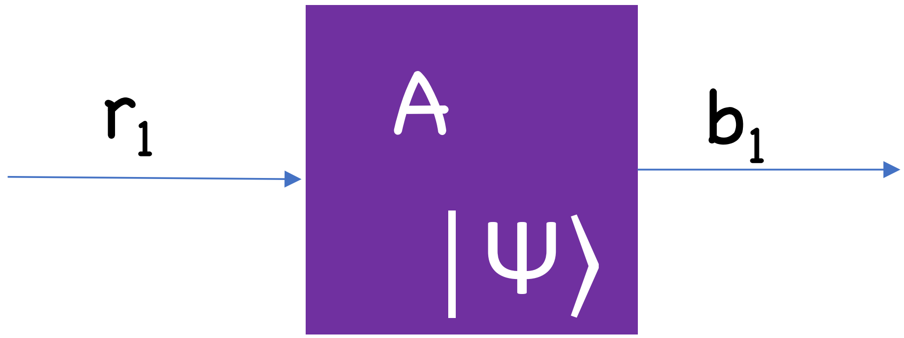

Measurement principle: extracting b_1 irreversibly altered $|\Psi\rangle$

GL assumption may not hold for 2nd query

Example 4: Goldreich-Levin

Thm (Adcock-Cleve'01): ∃ single-query quantum GL algorithm

Proof:

$\sum_{r \in \{0,1\}^n} |r\rangle$ → A → $\sum_{r \in \{0,1\}^n} (-1)^{A(r)} |r\rangle$ → QFT/ $H^{\otimes n}$ → 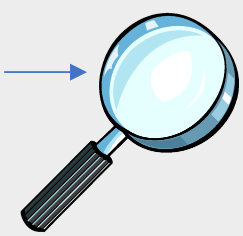

Results in tighter security reductions!

Example 4: Goldreich-Levin

Key Takeaway: As long as reduction treats $\mathcal{A}$ as a black box, potentially w/ *classical* interaction or w/ rewinding to *classical* value, reduction likely works in quantum setting

! But if interaction is quantum, all bets are off

! If rewinding to *quantum* state, all bets are off

Roadmap

New Quantum Attack Models

Quantum rewinding

Quantum Random Oracle Model