Bounding the uncertainty of constrained adversaries

Frédéric Dupuis
Aarhus University

Joint work with
Omar Fawzi (ETH Zürich)
Stephanie Wehner (National University of Singapore)

arXiv:1305.1316

CEQIP 2014, 6 June 2014

Multiparty computation

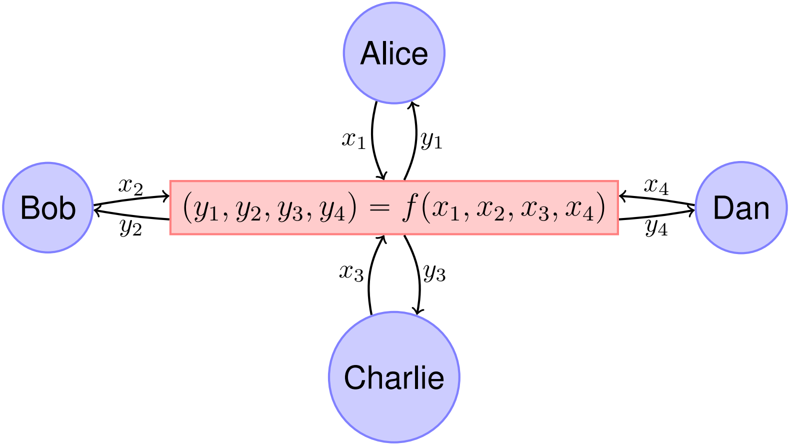

Examples: online voting, auctions, etc…

Multiparty computation

We want to implement this with no trusted third party:

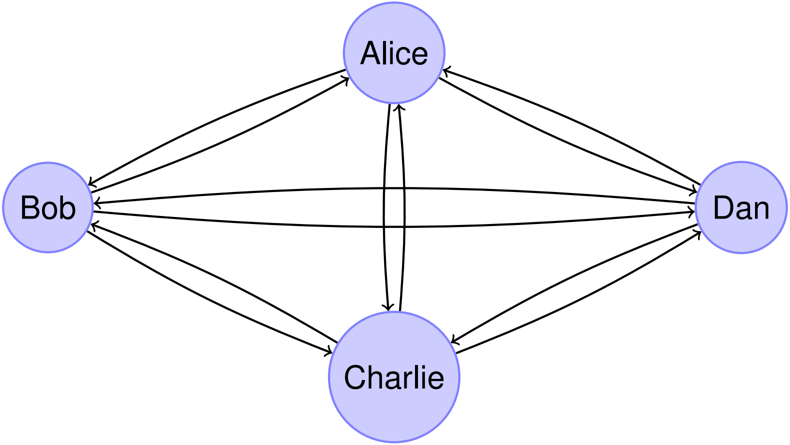

Oblivious transfer

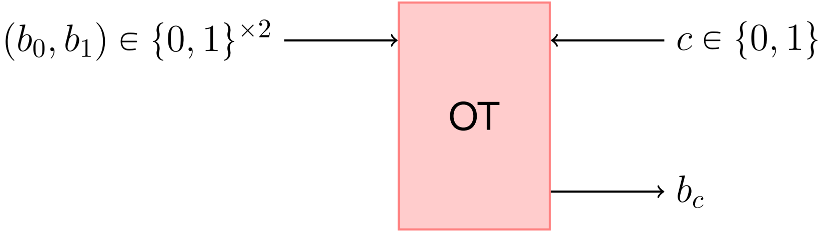

Bit commitment

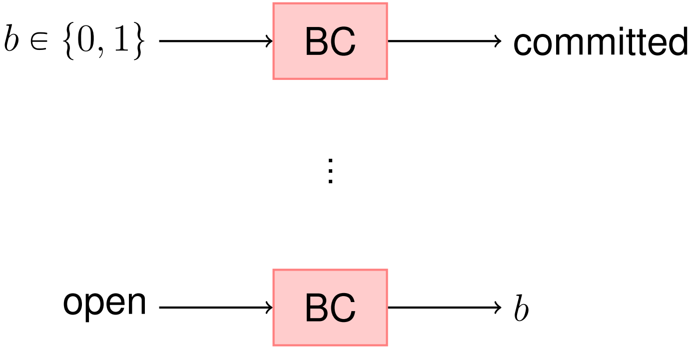

OT and BC

- Classically, BC is not enough for multiparty computation
- There exists a quantum protocol for OT using BC [Crépeau 1994]
- However: BC is impossible from scratch

Restricted adversaries

To make a BC protocol, we need to make assumptions:

- Computational assumptions: assume there is no efficient algorithm for solving certain problems
- Physical assumptions: assume an adversary is physically restricted in some way
  - Limited memory
  - Limited *quantum* memory
  - Noisy (quantum) memory
  - Noisy channel
  - Limited interaction between quantum systems
  - ...

Making use of the restrictions

- Goal of this talk: show how to make use of physical restrictions to construct protocols.
- Key idea: physical restriction $\Rightarrow$ bound on adversary's uncertainty about something

Measuring uncertainty

- How can we measure uncertainty?
- Entropy: $H(X)$, uncertainty about a random variable X:

$$H(X) = -\sum_x p_x \log p_x$$

- Why?
  - Compression: given n instances of X, we can compress it into $\approx nH(X)$ bits
  - Randomness extraction: given n instances of X, we can extract $\approx nH(X)$ bits of uniform randomness
  - What about just *one* instance of X? $H(X)$ is not good enough.

Measuring uncertainty

Why is $H(X)$ not good enough? Consider this distribution:

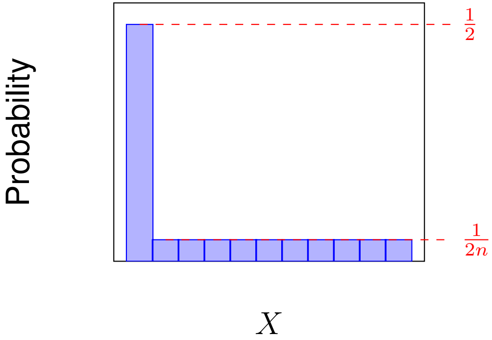

Can't really compress, can't extract more than 1 bit of randomness. But:

$$H(X) = -\frac{1}{2}\log(2) - \frac{1}{2}\log(n)$$

Measuring uncertainty

If we cannot use $H(X)$ to measure uncertainty, what should we use?

- Compression and randomness extraction require two different measures
- Compression: $H_{\max}(X)$ (won't talk about this)
- Randomness extraction: $H_{\min}(X)$

Min-entropy

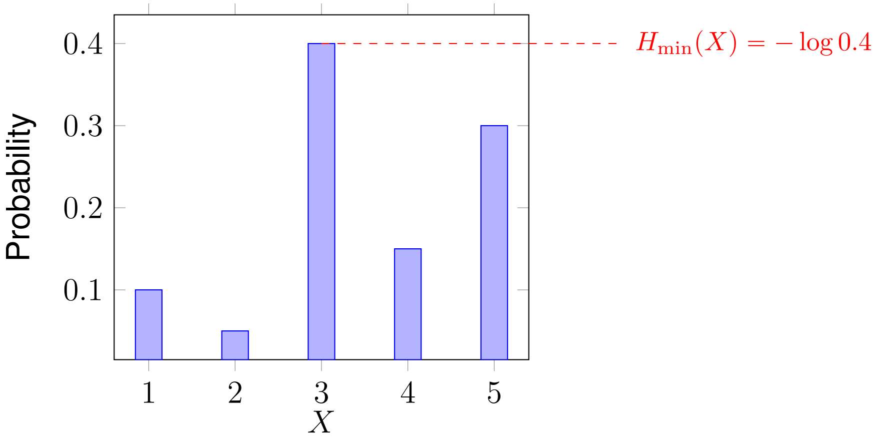

$H_{\min}(X) = -\log($probability of guessing $X)$.

Min-entropy: classical-quantum

What if we have quantum information about X?

- Alice has x with probability p_x
- Bob has ρ^x whenever Alice has x
- Represent this with the CQ state $\rho_{XB} := \sum_x p_x |x\rangle\langle x|_X \otimes \rho_B^x$.
- Bob tries to guess x by measuring his state

$$H_{\min}(X|B)_\rho := -\log(\text{probability of guessing } X).$$

Min-entropy: classical-quantum

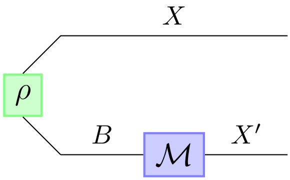

vs

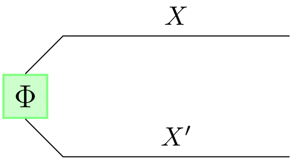

$$H_{\min}(A|B)_\rho := -\log d_X F(\Phi, \mathcal{M}(\rho))^2$$

Min-entropy: classical-quantum

Some properties of the min-entropy:

- Between 0 and $\log d$ (follows from the fact that the guessing probability must be between $1/d$ and 1)
- Can guess with probability 1: $H_{\min} = 0$
- Can't do better than $1/d$: $H_{\min} = \log d$

Min-entropy: fully quantum

What if X is now quantum as well?

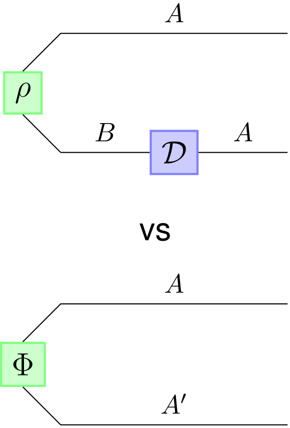

$$H_{\min}(A|B)_\rho := -\log d_A F(\Phi, \mathcal{D}(\rho))^2$$

Min-entropy: fully quantum

- In this case, the min-entropy can be *negative*!
- Example: maximally entangled state: $|\Phi\rangle = \sum_{x=1}^{d} |x\rangle_A \otimes |x\rangle_{A'}$ has a min-entropy of $H_{\min}(A|A')_\Phi = -\log d$.
- In general, $-\log d \leqslant H_{\min} \leqslant \log d$

Privacy amplification

- We have X, adversary has ρ_B^x, we somehow know that $H_{\min}(X|B) \geqslant k$.
- What can we do?
- We can extract $\approx k$ bits of uniform, independent randomness
- How? Apply a randomly chosen function $F(\cdot)$ to X

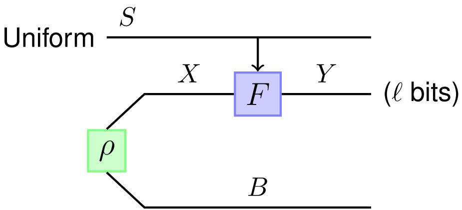

What we want at the output:

$$\mathrm{Unif}_{SY} \otimes \rho_B$$

Privacy amplification

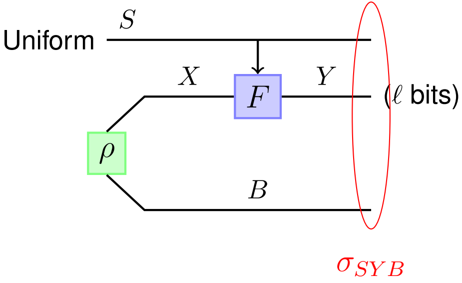

Theorem (Privacy amplification)

$$\left\|\sigma_{SYB} - \mathrm{Unif}_{SY} \otimes \rho_B\right\|_1 \leqslant \sqrt{2^{\ell - H_{\min}(X|B)_\rho}}$$

$\Rightarrow$ Just need ℓ to be a bit smaller than $H_{\min}(X|B)$.

Bounding the min-entropy

How can we get min-entropy bounds in protocols of interest?

- We want to be able to make statements such as $H_{\min}(A|E) \geqslant k$ where E is an adversary's information about some A of interest.
- Often, it is easy to make a statement about an intermediary step, but we want the bound to "survive" the rest of the protocol

Bounding the min-entropy

For example:

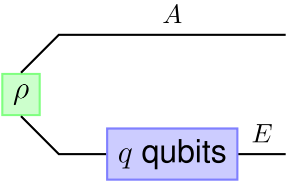

Very easy to bound the min-entropy:

$$H_{\min}(A|E) \geqslant -q$$

for any ρ.

Bounding the min-entropy

What if the honest parties then do something to A?

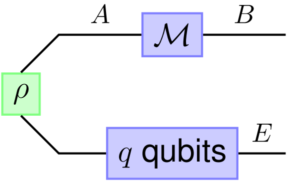

Some examples:

- Measure in random basis
- Sample random subsets of qubits
- Etc...

Bounding the min-entropy

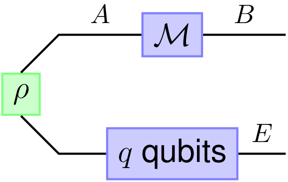

We want to be able to say

$$H_{\min}(B|E) \geqslant g(H_{\min}(A|E))$$

for an appropriate function g that will depend on $\mathcal{M}$.

A small caveat

- H_2 vs $H_{\min}$
- H_2 is "morally" equivalent to $H_{\min}$ (for example, privacy amplification still works with a bound on H_2 only)
- Can convert between the two:
  - For CQ states: $H_{\min}(X|B) \leqslant H_2(X|B) \leqslant 2H_{\min}(X|B)$
  - For general states: $H_{\min}(X|B) \leqslant H_2(X|B)$, and $H_2(X|B) + \log d \leqslant 2(H_{\min}(X|B) + \log d)$.

A general bound

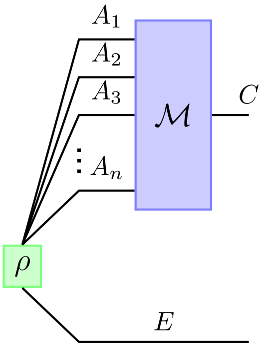

Theorem

$$\frac{1}{n}H_2(C|E) \gtrsim g\left(\frac{1}{n}H_2(A_1,\ldots,A_n|E)\right)$$

Sampling

Theorem

$$\frac{1}{k} H_2(A_S|ES) \gtrsim R\left(\frac{1}{n} H_2(A_1, \ldots, A_n|E)\right)$$

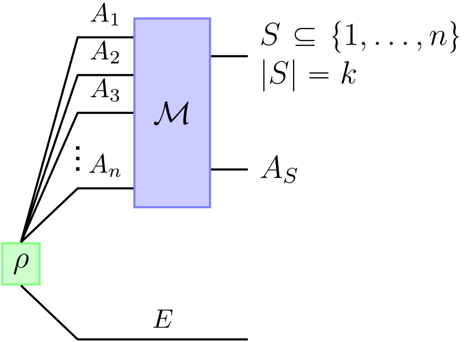

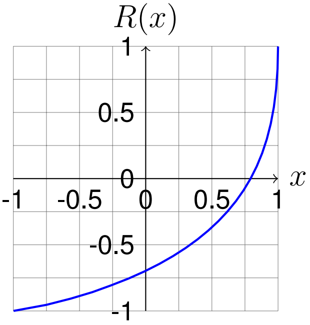

Sampling: the CQ case

Theorem

$$\frac{1}{k} H_2(X_S|ES)_\rho \gtrsim C\left(\frac{1}{n} H_2(X_1, \ldots, X_n|E)\right).$$

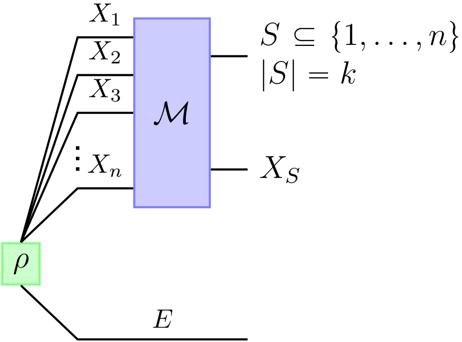

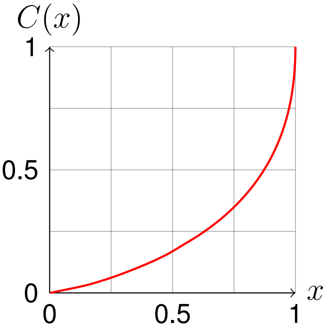

Sampling: CQ and fully quantum

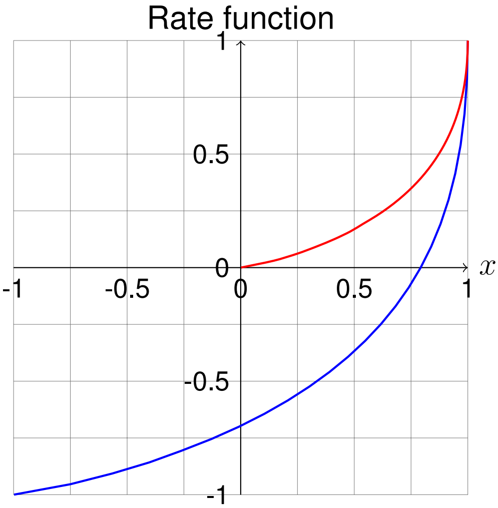

Measuring in a random BB84 basis

Theorem

$$\frac{1}{n} H_2(X^n | E\Theta^n)_\sigma \gtrsim \gamma\left(\frac{1}{n} H_2(A_1, \ldots, A_n | E)\right).$$

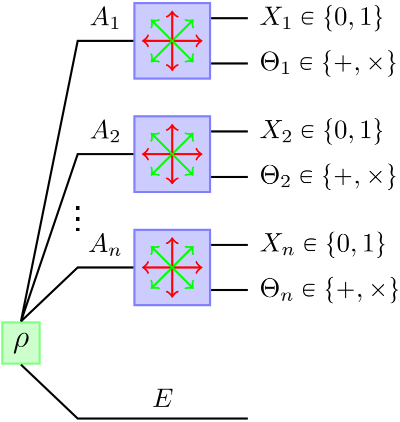

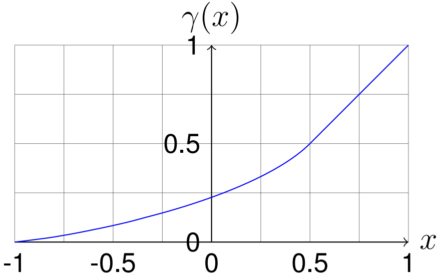

Bounded quantum storage model (BQSM)

At some point in the protocol, all parties are assumed to have at most q qubits of storage (but unlimited classical storage).

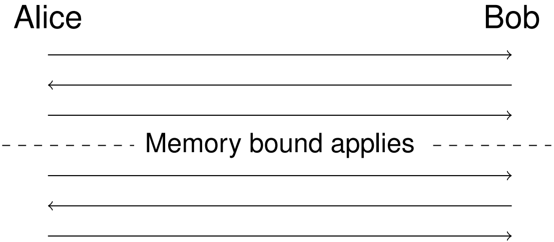

Weak string erasure

Bit commitment can in turn be reduced to *weak string erasure* [König, Wehner, Wullschleger 2012]:

Alice | Bob

$X^n \in_R \{0,1\}^n \longleftarrow$ WSE $\longrightarrow \mathcal{I} \subseteq_R [n], X_{\mathcal{I}}$

For security, we want:

- $\mathcal{I}$ is distributed uniformly over $[n]$ and is independent of anything Alice has.
- If Bob is dishonest, then $\frac{1}{n} H_{\min}(X^n|B)_\sigma \geqslant \lambda$, where $\sigma_{X^n B}$ is the state at the end of the protocol.

Weak string erasure

Given a protocol for weak string erasure with

$$\lambda \geqslant \Omega\left(\frac{\log n}{n}\right),$$

we can do bit commitment.

Protocol for weak string erasure

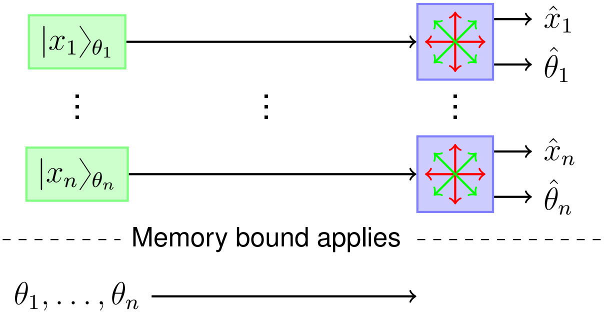

Output:

$x_1, \ldots, x_n$

Output:

$\mathcal{I} = \{i : \theta_i = \tilde{\theta}_i\}$

$\tilde{x}_{\mathcal{I}}$

Protocol for weak string erasure

Does this protocol satisfy the security definition?

- $\mathcal{I}$ uniform and independent. Yes: $\mathcal{I}$ only depends on the XOR of θ^n and $\tilde{\theta}^n \Rightarrow$ Alice has no control over it.
- We need that, for a dishonest Bob, $\frac{1}{n} H_{\min}(X^n|B)_\sigma \geqslant \lambda$.

We need our theorem to guarantee the second point.

Protocol for weak string erasure

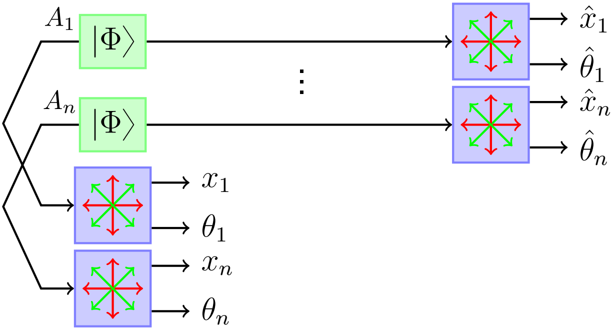

- - - - Memory bound: q qubits max - - - - - - - - - - -

$\theta_1, \ldots, \theta_n \longrightarrow$

Output:
$x_1, \ldots, x_n$

Output:
$\mathcal{I} = \{i : \theta_i = \tilde{\theta}_i\}$
$\tilde{x}_{\mathcal{I}}$

Protocol for WSE: dishonest Bob

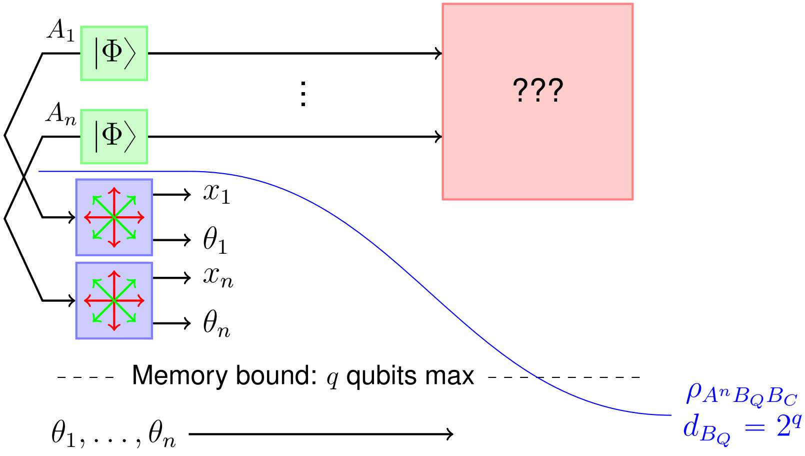

Output:
$x_1, \ldots, x_n$

Output:
$\mathcal{I} = \{i : \theta_i = \tilde{\theta}_i\}$
$\tilde{x}_{\mathcal{I}}$

Protocol for weak string erasure

Recall our theorem on measuring in random BB84 bases:

$$\frac{1}{n}H_2(X^n|B_Q B_C \Theta^n) \gtrsim \gamma\left(\frac{1}{n}H_2(A^n|B_Q B_C)\right)$$

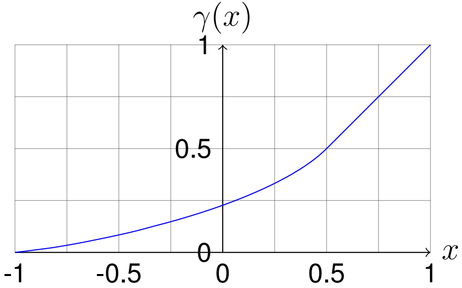

But we know that

$$H_2(A^n|B_Q B_C) \geqslant -q$$

because of the memory bound.

Protocol for weak string erasure

$$\frac{1}{n} H_2(A^n | B_Q B_C) \geqslant \frac{-q}{n}$$

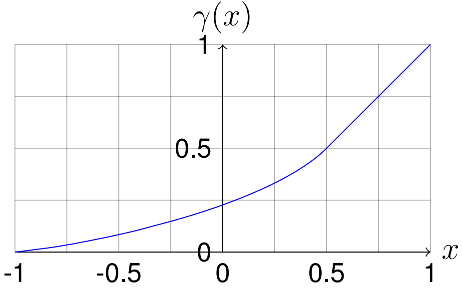

We get a nontrivial bound as soon as $q < n$!

Protocol for weak string erasure

- To get bit commitment, it enough for to require q to be at most
$$n - c\log^2 n - c\log n \log(1/\varepsilon).$$
- Since for $q = n$ we cannot have security, this is essentially optimal.
- Previous best: security for $q \approx 2n/3$.
- Also works for any other model in which we get a nontrivial bound on $H_2(A^n|B)_\rho$ (noisy memory model, etc).

Thank you

Thank you!

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

SOLIDUM PETIT IN PROFUNDIS · UNIVERSITAS ARHUSIENSIS ·

Centre for
Quantum
Technologies