An efficient quantum parallel repetition theorem and applications

John Bostanci
Columbia University

Luowen Qian
Boston University

Nicholas Spooner
University of Warwick & NYU

Henry Yuen
Columbia University

Short Plenary Talk at Quantum Information Processing (QIP) 2024

[arXiv:2311.10681](https://arxiv.org/abs/2311.10681)

Hardness Amplification

- Starting point: a weakly hard task (best success probability is say ¾)
- Goal: a “truly” hard task (cannot do much better than trivial)
- Examples: hard functions, distinguish two distributions, interactive arguments, MIP
- Applications: circuit lower bounds, PCP theorems, cryptography

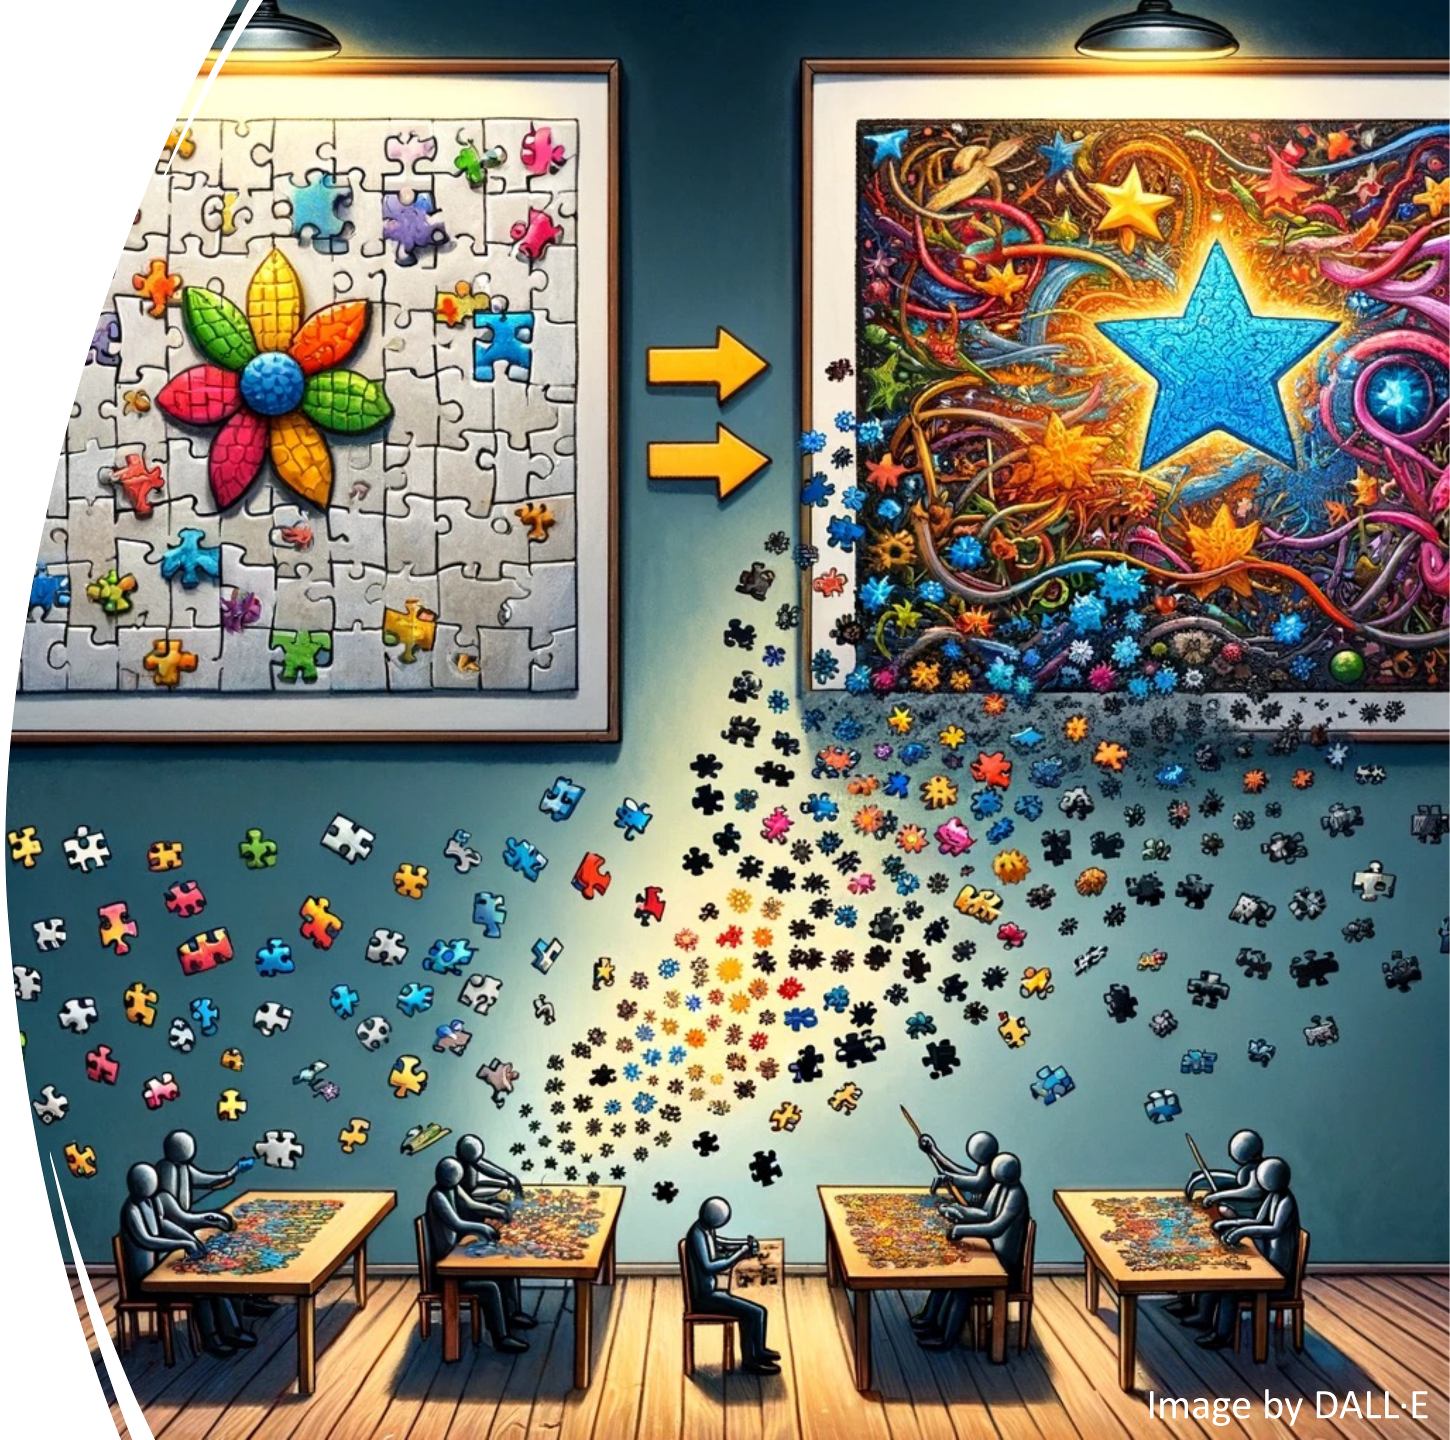

Hardness amplification of security game

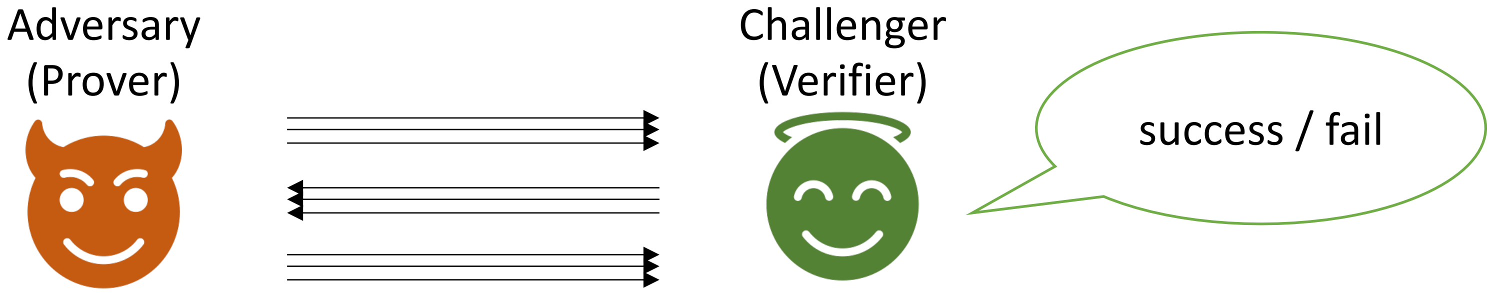

- Starting point: a weakly secure primitive against *efficient* adversaries
- Goal: construct a fully secure primitive
- Constraints: preserve desirable properties
  - Time efficiency
  - Round complexity
  - Zero knowledge
  - …

Simplest, natural & generic approach: parallel repetition (aka direct product)

Efficient parallel repetition: classical landscape

- k-fold parallel repetition of any 3-message protocol with computational security ε yields a *tight* security of $\varepsilon^k + \mathrm{negl}$ [Bellare, Impagliazzo, Naor'97; Canetti, Halevi, Steiner'05]
  - Parallel repetition probably does not work for 4-message protocols in general [BIN97; Pietrzak, Wikström'12]
  - Negligible loss is also probably inherent [Dodis, Jain, Moran, Wichs'12]
- Parallel repetition also works if the protocol is partially simulatable (3-message, public coin, random-terminating) […; Berman, Haitner, Tsfadia'20] or wrapped in FHE [Chung, Liu'10]
  - Possible to preprocess any r-round protocol incurring a multiplicative cost of order r or λ in efficiency (also tight for each approach)

(Post-)quantum security games

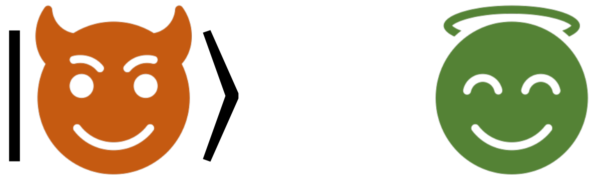

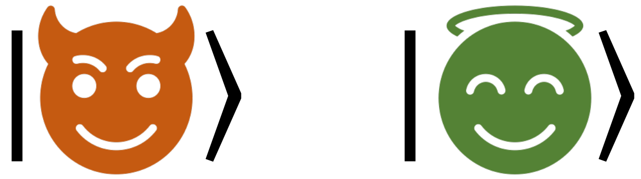

Post-quantum cryptography:
secure existing cryptography against quantum adversaries
(challenger is still classical)

Quantum cryptography:
go beyond existing cryptography through quantum information and quantum computing

Our quantum efficient parallel repetition

- k-fold parallel repetition of any 3-message quantum protocol with computational security ε yields a *tight* security of $\varepsilon^k + \mathrm{negl}$
  - Parallel repetition does not work for 4-message post-quantum protocols assuming post-quantum concurrent non-malleable commitments
  - Negligible loss is inherent even for post-quantum assuming exponentially hard post-quantum extended second-preimage resistant hash functions
- *Round collapse*: compile any protocol to a 3-message quantum protocol while preserving computational security
  - Same transformation as QIP [Kitaev, Watrous'00; Kempe, Kobayashi, Matsumoto, Vidick'07]
  - Multiplicative loss of $O(r^{3.322})$ for r-round (definitely not tight)

Uniformity of reduction

Core (classical/quantum) proof strategy: reduction

- "Good" k-fold adversary A
  $\Rightarrow$ construct "good" 1-fold adversary B

A has success probability δ^k ($\delta \gg \varepsilon$)

$\Rightarrow$ we want B to also succeed with probability $\approx \delta$

- A could be non-uniform!

$$\exists \alpha \colon \Pr_{A,C}[\langle A(\alpha), C\rangle = 1] = \delta^k$$

- Uniform reduction: B uses the same advice (constructive: desirable for win-win philosophy)
- Possible classically: $\Pr_{B,C}[\langle B(\alpha), C\rangle = 1] \approx \delta$

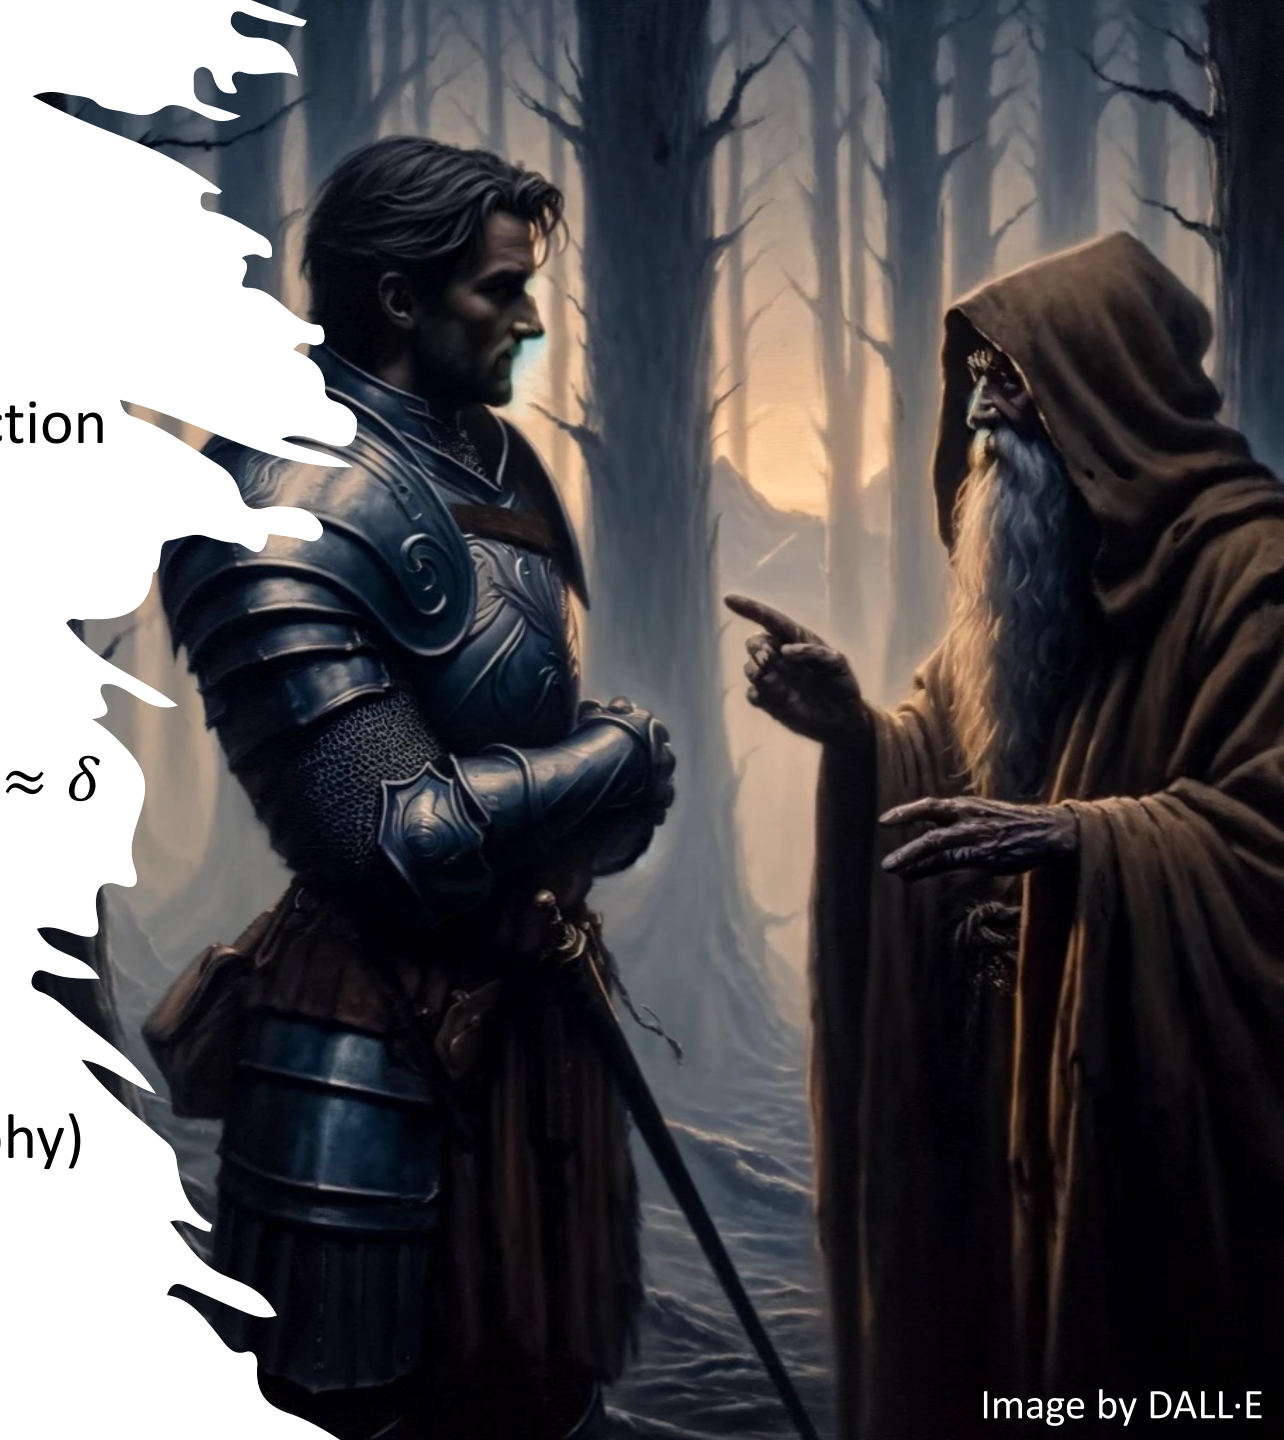

Trouble with randomized/quantum advice

$$\forall A, \alpha: \left(\Pr_{A,C}[\langle A(\alpha), C\rangle = 1] = \delta^k \ \Rightarrow \ \Pr_{B,C}[\langle B(\alpha), C\rangle = 1] \approx \delta \right)$$

- Q: What if advice is randomized? $\exists D: \Pr_{A,C,\alpha\sim D}[\langle A(\alpha), C\rangle = 1] = \delta^k$
- A: "Uniform reduction" is now impossible!
  - D samples a "trapdoor" with probability δ^k and samples "abort" otherwise
  - B must either work with a known "good sample" or try to find one by taking $\approx \delta^{-k}$ samples from D (assuming only black-box access to A and D)
  - Problematic as long as $\Pr[B] > \Pr[A]$
- Same issue with quantum advice via purification
  - We can work with either a single known "good eigenstate" or take $\approx \delta^{-k}$ copies (*best possible* uniform reduction)

Cryptographic applications

- Amplification of quantum primitives with a 2-message security game
  - Commitment schemes, EFI pairs… (posed in Yan'22 and Brakerski, Canetti, Q'23)
  - Quantum money schemes from weakly unforgeable ones (posed in Aaronson and Christiano'13)
  - Quantum lightning schemes (existential unforgeable quantum money)
- Any weakly-sound (quantum) honest-verifier zero-knowledge (QHVZK) argument ⇒ 3-message negligibly-sound QHVZK arguments
  - ➢ Preserves succinctness but not classical communication
- Amplification for any *post-quantum* 3-message argument
- …

Why is even post-quantum non-obvious?

- ➢ Classical reductions for parallel repetition must "rewind" many times, notoriously problematic for quantum adversaries
  - ➢ Rewinding: feed an adversary with one message, obtain some information, go back and feed a different message
  - ➢ Quantumly, obtaining some information is measuring, which disturbs the adversary's success probability; cloning internal states is also impossible
- ➢ Quantumly unrewindable (contrived) protocols exist:
  - Relative to a quantum oracle [Ambainis, Rosmanis, Unruh'14]
  - Assuming quantum hardness of Learning with Errors [Brakerski, Christiano, Mahadev, Vazirani, Vidick'21]
- ➢ Quantumly rewinding techniques developed for zero knowledge and succinct arguments do not immediately apply [Watrous'09; Unruh'12; Chia, Chung, Yamakawa'21; Chiesa, Ma, Spooner, Zhandry'22; Lombardi, Ma, Spooner'22]

Yao's XOR lemma = parallel repetition

- Predicate $f\colon \pm 1^n \to \pm 1$ is ε-hard to predict over D if any poly-time A,
$$\mathrm{E}_{x\sim D}[A(x)\cdot f(x)] \leq \varepsilon + \mathrm{negl}$$
- Yao's XOR lemma (1982): if f is ε-hard to predict over D, then $f^{\oplus k}(x_1,\ldots,x_k) \coloneqq \prod_i f(x_i)$ is ε^k-hard to predict over $D^{\otimes k}$ [Levin'87]

Equivalent to parallel repetition up to some loss:

- XOR lemma $\Rightarrow$ parallel repetition — intuitively easy [Viola, Widgerson'08]
- XOR lemma $\Leftarrow$ parallel repetition — Goldreich–Levin [Goldreich, Nisan, Widgerson'11]
- Extremely similar proof techniques

XOR lemma for quantum predicates

- Quantum predicates ρ_+ and ρ_- with disjoint support ($\rho_+\rho_- = 0$)
- ε-unpredictable if poly-time $A(\rho_+) - A(\rho_-) \leq 2\varepsilon + \mathrm{negl}$
- Our parallel repetition theorem + quantum commitment duality

[Hhan, Morimae, Yamakawa'23]

$\Rightarrow$ k-fold XOR of ρ_+, ρ_- is $\varepsilon^{k/2}$-unpredictable

(posed in Colisson'19 and Brakerski'23)

- Better loss than classical GNW11 proof (like how quantum Goldreich–Levin is also more efficient)
- Application: average-case hardness amplification for "quantum-input decision PSPACE"

Proof for baby case: 2-fold 2-message

- Start with classical *baby* case:
  2-fold 2-message tight parallel repetition with *non-uniform* reduction from Levin's isolation lemma (1987) and CHS05
- Adapt to post-quantum
- Adapt to fully quantum (handwavy)
- See paper:
  - Extension to many folds
  - Proof of best possible uniform reduction
  - Other applications and details

2-fold 2-message parallel repetition

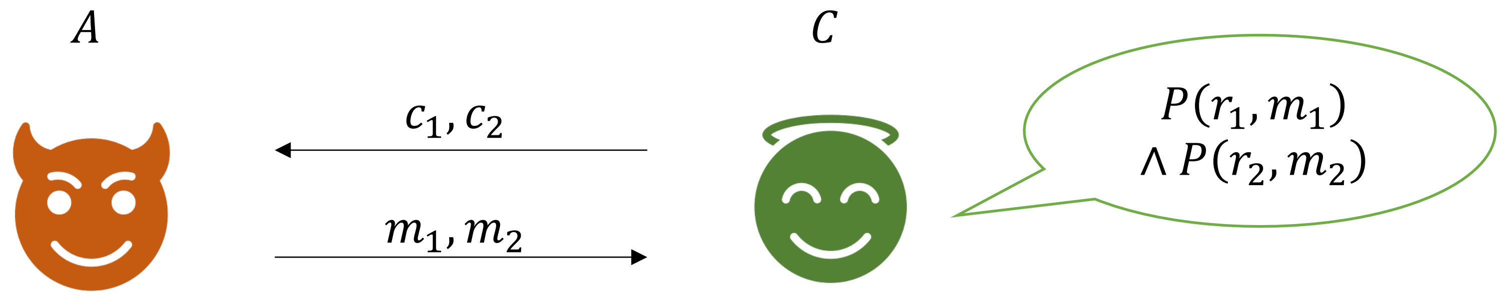

- Winning fold $\#i$ event $G_i \coloneqq P(r_i, m_i)$
- $\Pr[G_1 \wedge G_2] = \delta^2$
- We want to have tight bounds!
  $\Rightarrow$ Reduction should succeed with probability $\approx \delta$
- Hope: $\Pr[G_1] \geq \delta$ or $\Pr[G_2] \geq \delta \Rightarrow$ contradiction?

2-fold 2-message parallel repetition (careful)

- Winning fold #i event $G_i \coloneqq P(r_i, m_i)$
- $\delta^2 = \Pr[G_1 \wedge G_2]$

  $= \Pr[G_1] \cdot \Pr[G_2 | G_1]$ (applying Bayes' rule)

1. $\Pr[G_1] \geq \delta \Rightarrow$ contradiction: reduction honestly simulate fold #2
2. $\Pr[G_2 | G_1] \geq \delta \Rightarrow$?
   - Not meaningful!
   - Conditioning on G_1 may significantly change the distribution on r_2
   - No reduction ☹

2-fold 2-message parallel repetition (careful-er)

- Winning fold $\#i$ event $G_i \coloneqq P(r_i, m_i)$
- $\delta^2 = \Pr[G_1 \wedge G_2]$

$$= \Pr[G_1] \cdot \Pr[G_2 | G_1]$$

$$= \mathrm{E}_{r_2}[\Pr[G_1] \cdot \Pr[G_2 | G_1]]$$

1. $\exists r_2 \colon \Pr[G_1] \geq \delta \Rightarrow$ still contradiction!
   - Reduction hardwires that r_2 as advice (non-uniform)
2. $\forall r_2 \colon \Pr[G_1] \leq \delta \Rightarrow \mathrm{E}_{r_2}[\Pr[G_2 | G_1]] \geq \delta \Rightarrow$ also contradiction
   - Reduction: "rejection sample" m_2 until G_1
     (randomly sample r_1, run A, output m_2 if G_1, otherwise rewind to beginning)

Great, how about post-quantum?

1. $\exists r_2 \colon \Pr[G_1] \geq \delta \Rightarrow$ contradiction
   - Reduction hardwires that r_2 as advice (non-uniform)
   - Still works!
2. $\mathrm{E}_{r_2}[\Pr[G_2|G_1]] \geq \delta \Rightarrow$ contradiction?
   - Reduction: "rejection sample" m_2 until G_1
     (randomly sample r_1, run A, output m_2 if G_1, otherwise rewind to beginning)
   - Can reset to beginning if $|\alpha\rangle$ is clonable/classical,
     or if we have many copies of $|\alpha\rangle$ (ok but not ideal [Bitansky, Brakerski, Kalai'22])
   - Fails harder for 3-message ☹

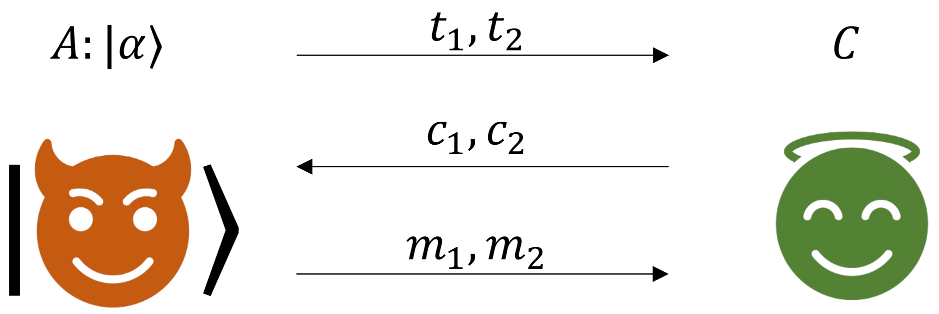

Fully quantum (very handwavy)

Back to 2-message…

1. $\exists |r_2\rangle : \Pr[G_1] \geq \delta \Rightarrow$ contradiction
   - Reduction hardwires that $|r_2\rangle$ as advice (non-uniform)
2. $\mathrm{E}_{|r_2\rangle}[\Pr[G_2|G_1]] \geq \delta \Rightarrow$ contradiction?
   - Reduction: "rejection sample" $|m_2\rangle$ until G_1?
   - Natural idea: alternating projectors from quantum rewinding
     [Watrous09, CCY21, CMSZ22, LMS22]
   - Issues: (1) measures singular value causing disturbance
     (2) possible unnecessary amplitude causing destructive interference
   - Solution: Quantum Singular Value Transform (QSVT)

Quantum Singular Value Transform (QSVT)

[Gilyén, Su, Low, Wiebe'19]

Unification of most quantum algorithms (except QFT and classical)

- Given a block encoding of $A = \sum_i \varsigma_i |w_i\rangle\langle v_i|$ and a low-degree odd polynomial p: $[-1,1] \to [-1,1]$, QSVT approximates $\sum_i p(\varsigma_i)|w_i\rangle\langle v_i|$
- **Uniform singular value amplification**: Given $PQ = \sum_i \varsigma_i |w_i\rangle\langle v_i|$, we can efficiently approximate the map $\sum_i \frac{\varsigma_i}{\gamma} |w_i\rangle\langle v_i|$ on all $\varsigma_i < \gamma$ given access to $\mathrm{C}_P\mathrm{NOT}, \mathrm{C}_Q\mathrm{NOT}$ gates
- We use uniform singular value amplification to do "coherent post-selection ≈ sampling quantum conditional distributions"

Conclusions

- We adapt recent quantum algorithmic and rewinding techniques to prove efficient 3-message parallel repetition theorem and XOR lemmas with best possible uniform reduction (see paper)
- We show how to quantumly efficiently round collapse other protocols to 3-message

Future work:

- Quantize other parallel repetition theorems (partially simulatable or FHE wrapped protocols)
- Investigate more rewinding reductions ☺

Thank you! Questions?