Distributed quantum inner product estimation

Yunchao Liu (UC Berkeley)
Joint work with Anurag Anshu (Harvard) and Zeph Landau (UC Berkeley)

arxiv: 2111.03273

Problem definition

Alice $\rho^{\otimes k}$ (unknown)

classical communication

Bob $\sigma^{\otimes k}$ (unknown)

$$\mathrm{Tr}(\rho\sigma) \approx ?$$

This talk: the optimal sample complexity of this task, i.e. smallest k such that the task is achievable

Some quick thoughts

- Q: What happens if allow quantum communication?
- A: $k = O(1/\varepsilon^2)$ suffices
  - Alice sends her copies to Bob
  - Bob performs the SWAP test

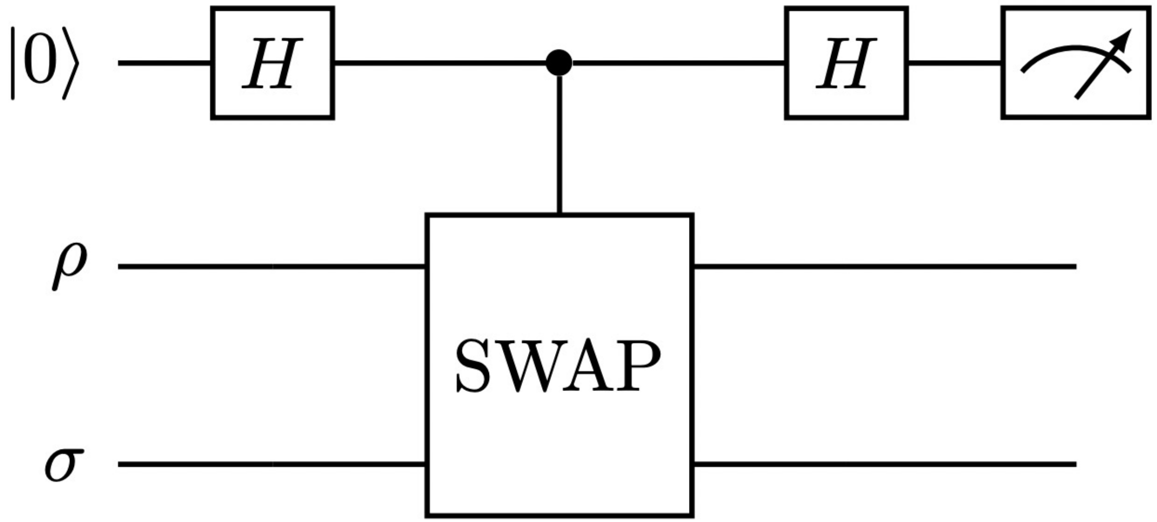

Some quick thoughts

- Q: Why do we care about $\mathrm{Tr}(\rho\sigma)$?
- A: $\mathrm{Tr}(\rho\sigma)$ itself doesn't have much operational meaning, but…
  - When one state is pure, $\mathrm{Tr}(\rho\sigma) = F(\rho,\sigma)$
  - $\mathrm{Tr}(\rho\sigma)$ is related to other (non-standard) distance metrics, such as
    - Hilbert-Schmidt distance $D_{HS}(\rho,\sigma) = \sqrt{\mathrm{Tr}((\rho-\sigma)^2)}$
    - "geometric mean" fidelity $F_{GM}(\rho,\sigma) = \frac{\mathrm{Tr}(\rho\sigma)}{\sqrt{\mathrm{Tr}(\rho^2)\mathrm{Tr}(\sigma^2)}}$
    - These distance metrics are determined by $\mathrm{Tr}(\rho\sigma)$, $\mathrm{Tr}(\rho^2)$, $\mathrm{Tr}(\sigma^2)$

Some quick thoughts

- Q: Why do we care about estimating $\mathrm{Tr}(\rho\sigma)$ in a distributed setting?
- A: Cross-platform verification [Elben et al'20]

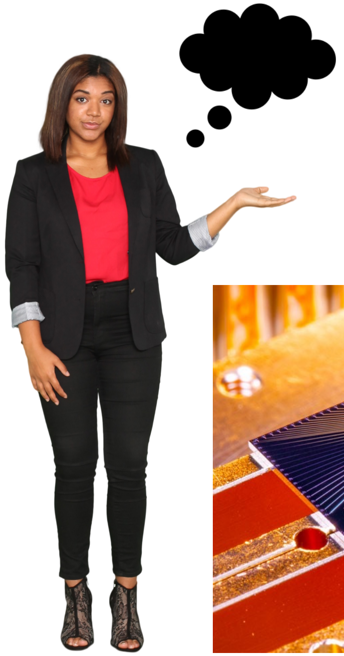

How do we compare our unknown quantum states
that live on different physical platforms?

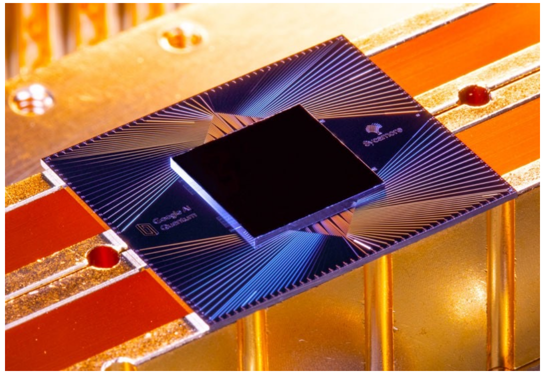

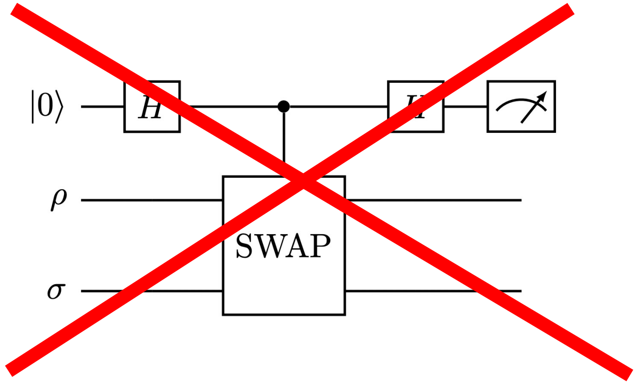

Problem definition

Alice $\rho^{\otimes k}$ — classical communication → Bob $\sigma^{\otimes k}$

$$\mathrm{Tr}(\rho\sigma) \approx ?$$

This talk: the optimal sample complexity of this task, i.e. smallest k such that the task is achievable

Probably depends on the model of measurement and communication?

Measurement models

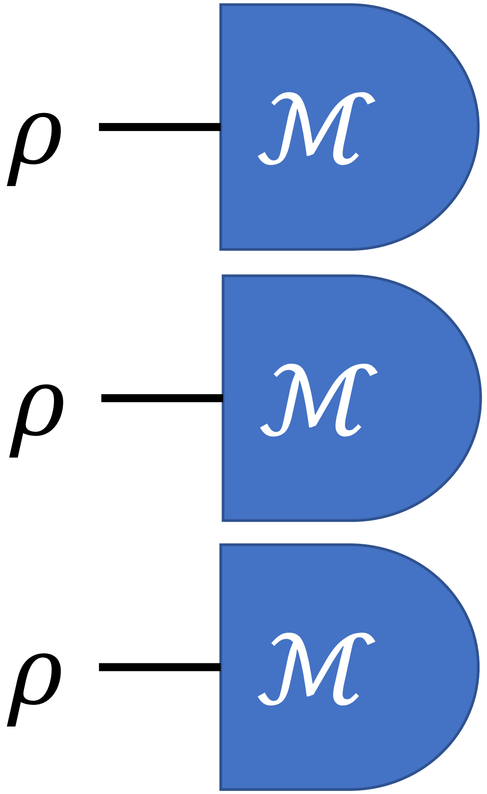

Single-copy measurements
Requires $\Theta(d^3)$ copies for tomography

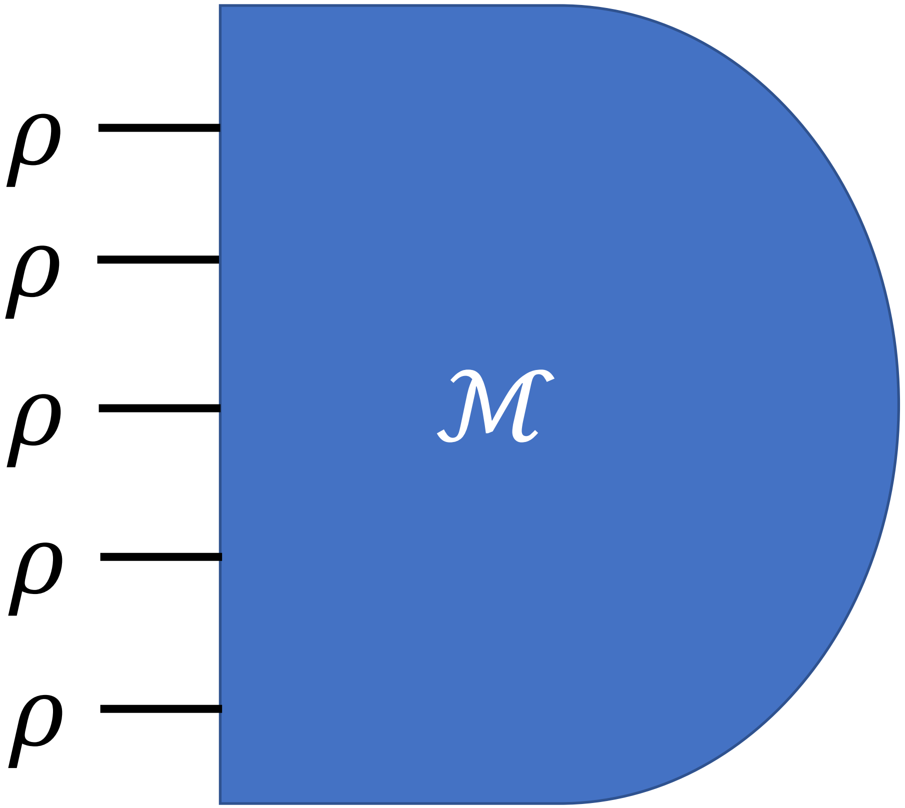

Multi-copy measurements
Requires $\Theta(d^2)$ copies for tomography

Communication models

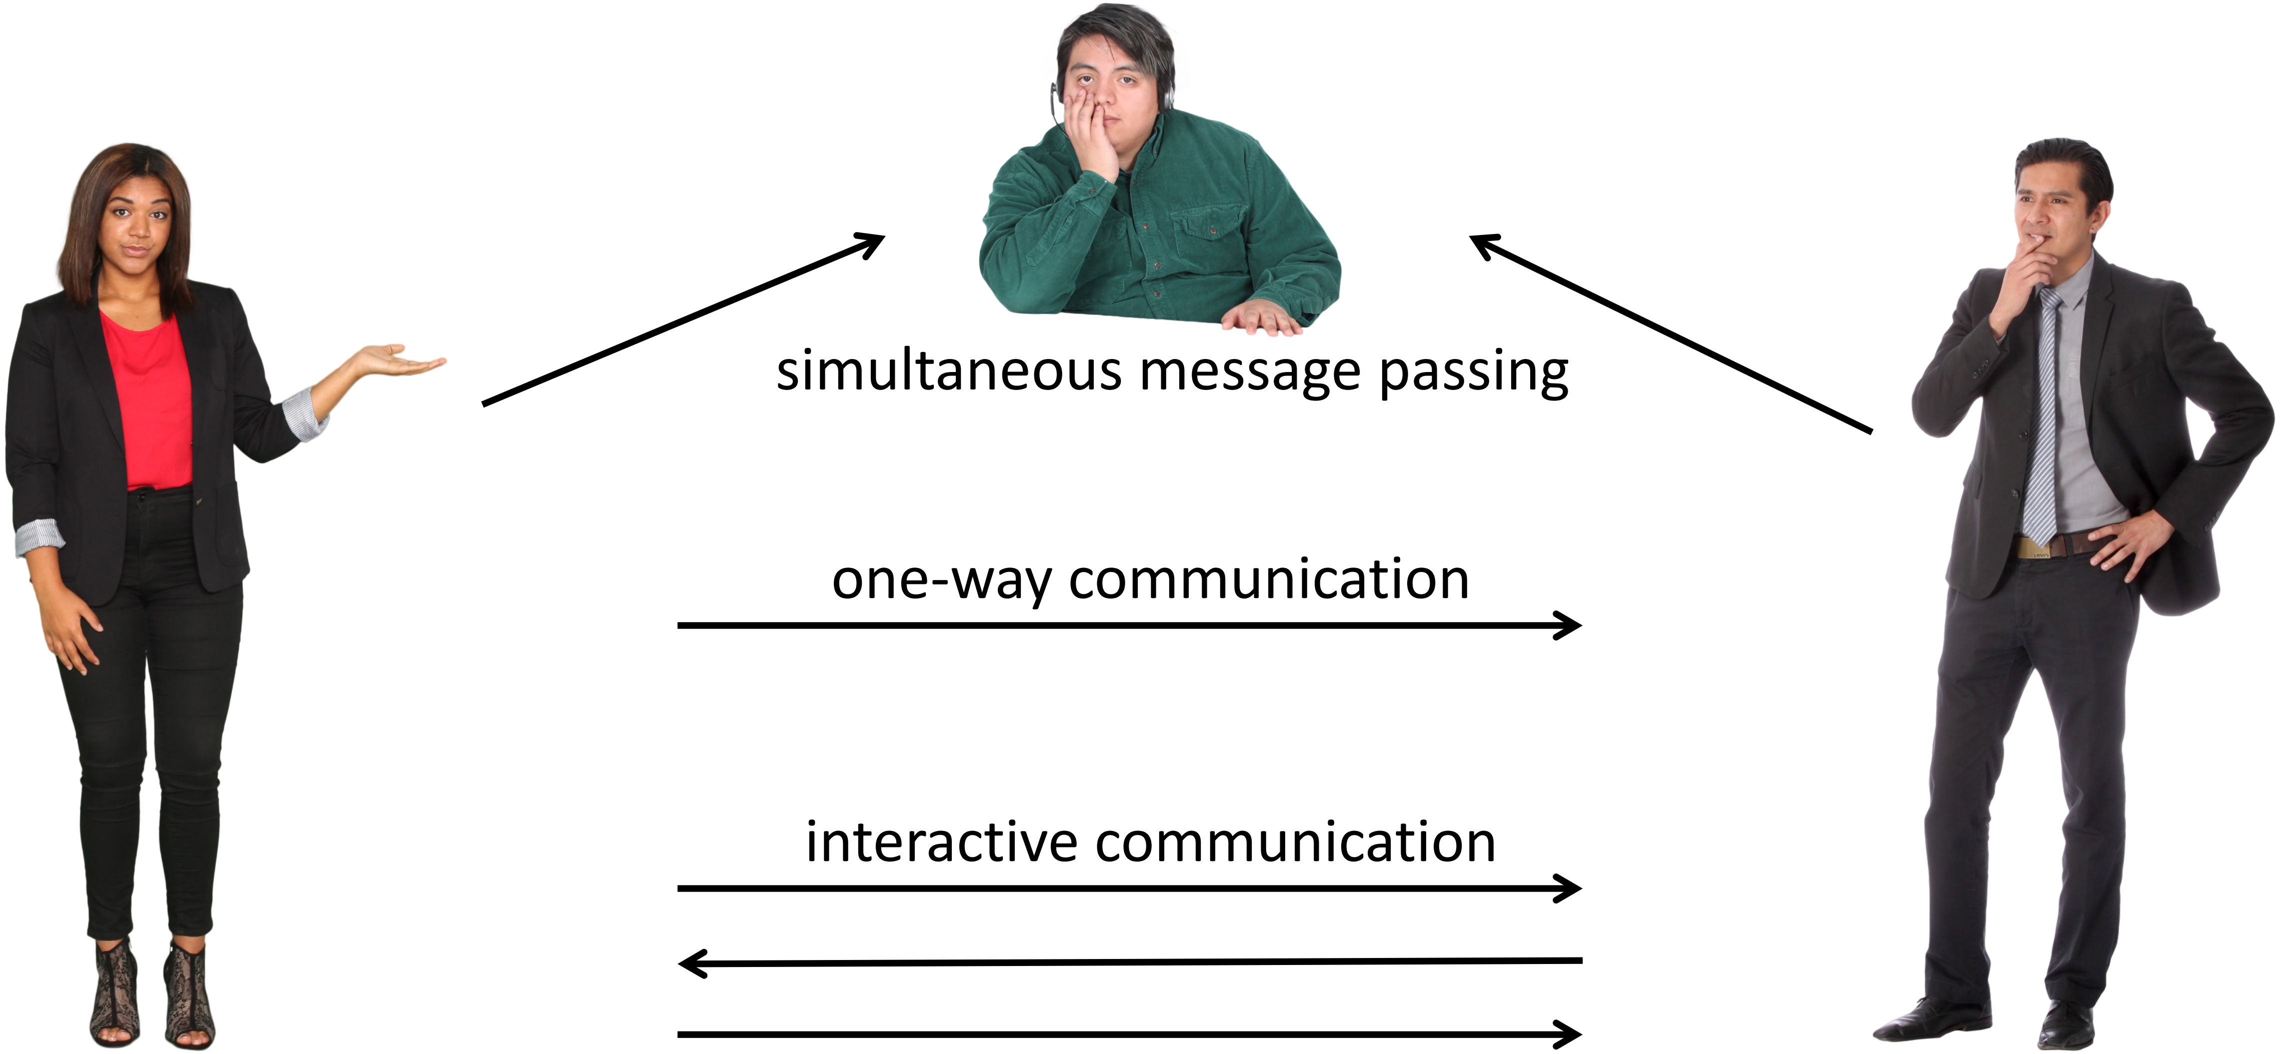

Result

- A priori the above $2\times 3 = 6$ models could lead to different sample complexity for the task, but we show this is not the case
- **Theorem.** The optimal sample complexity for distributed quantum inner product estimation is
  - $k = \Theta(\max\{\frac{1}{\varepsilon^2}, \frac{\sqrt{d}}{\varepsilon}\})$
- across all measurement and communication models

- When ε is constant, this gives $k = \Theta(2^{n/2})$ (n=#qubits)

Discussion

- **Theorem.** The optimal sample complexity for distributed quantum inner product estimation is $k = \Theta(\max\{\frac{1}{\varepsilon^2}, \frac{\sqrt{d}}{\varepsilon}\})$ across all measurement and communication models

- Regarding the cross-platform verification [Elben et al'20] task, we conclude that it requires less samples than tomography
- But still requires exponential samples (in #qubits), even with the most powerful measurements

Discussion

- **Theorem.** The optimal sample complexity for distributed quantum inner product estimation is $k = \Theta(\max\{\frac{1}{\varepsilon^2}, \frac{\sqrt{d}}{\varepsilon}\})$ across all measurement and communication models

- Shadow tomography [Aaronson'18]: linear functions of an unknown quantum state can be estimated sample-efficiently
- But our task is not sample-efficient… because the classical communication constraint seems to be a barrier for sample-efficiency

Discussion

- **Theorem.** The optimal sample complexity for distributed quantum inner product estimation is $k = \Theta(\max\{\frac{1}{\varepsilon^2}, \frac{\sqrt{d}}{\varepsilon}\})$ across all measurement and communication models

- Besides tomography, many examples are known which demonstrate large separation between single and multi-copy measurements for single-system property testing [BCL'20; ACQ'21; CCHL'21]
- But in our distributed setting, access to multi-copy measurements does not provide an advantage

Only need to prove two bounds

- Using single-copy measurements and simultaneous message passing, Alice and Bob can estimate inner product with $k = O(\max\{\frac{1}{\varepsilon^2}, \frac{\sqrt{d}}{\varepsilon}\})$ copies

- Even with multi-copy measurements and interactive communication, Alice and Bob require at least $k = \Omega(\max\{\frac{1}{\varepsilon^2}, \frac{\sqrt{d}}{\varepsilon}\})$ copies to estimate inner product

The upper bound

- Using single-copy measurements and simultaneous message passing, Alice and Bob can estimate inner product with $k = O(\max\{\frac{1}{\varepsilon^2}, \frac{\sqrt{d}}{\varepsilon}\})$ copies

- Idea: reduce quantum inner product to classical inner product using “correlated” classical shadows

Warm-up: how to estimate the inner product of two probability distributions?

- We can draw i.i.d. samples from two d-dim distributions p, q
- Want to estimate $f = \sum_{x=0}^{d-1} p_x \cdot q_x$

- Draw m samples $x_1, \dots, x_m \sim p$, $y_1, \dots, y_m \sim q$
- Collision estimator: output $\frac{1}{m^2} \sum_{j,k=1}^{m} 1[x_j = y_k]$

- Example: {101,111,010,101}, {110,000,101,111}
- Output=(1+1+0+1)/16=0.1875

Proof sketch

Shared randomness

1. Sample a random unitary U
2. Apply U to each copy of my state
3. Measure each copy in the computational basis, obtain bit strings $A = (a_1, \dots, a_k)$

1. Sample a random unitary U
2. Apply U to each copy of my state
3. Measure each copy in the computational basis, obtain bit strings $B = (b_1, \dots, b_k)$

Count #collisions between A and B

(Collision estimator)

Output a function of #collisions

Intuition

- To prove the sample complexity bound, we need to calculate the variance of the above estimator…
- Why is $O(\sqrt{d})$ the correct bound?
- **Intuition: birthday paradox:** expect to see collisions after drawing $k = O(\sqrt{d})$ samples from a d-dim uniform distribution
- Alice and Bob's measurement outcome distributions are close to uniform
  - When $k = o(\sqrt{d})$, never see any collision
  - When $k = O(\sqrt{d})$, see more collisions when inner product is large; fewer collisions when inner product is small

The lower bound

- Even with multi-copy measurements and interactive communication, Alice and Bob require at least $k = \Omega(\max\{\frac{1}{\varepsilon^2}, \frac{\sqrt{d}}{\varepsilon}\})$ copies to estimate inner product

Proof sketch: focus on a simpler problem

$|\phi\rangle^{\otimes k}$

$|\phi\rangle \sim \mathbb{C}^d$
(Haar measure)

$|\phi\rangle^{\otimes k}$

Which case are we in?

$|\phi\rangle^{\otimes k}$

$|\phi\rangle, |\psi\rangle \sim \mathbb{C}^d$
(Haar measure)

$|\psi\rangle^{\otimes k}$

The lower bound

- Even with multi-copy measurements and interactive communication, Alice and Bob require at least $k = \Omega(\sqrt{d})$ copies to decide
- Idea: symmetric subspace

Proof sketch

$|\phi\rangle^{\otimes k}$ $|\phi\rangle \sim \mathbb{C}^d$ $|\phi\rangle^{\otimes k}$

Which case are we in?

$|\phi\rangle^{\otimes k}$ $|\phi\rangle, |\psi\rangle \sim \mathbb{C}^d$ $|\psi\rangle^{\otimes k}$

Symmetric subspace

No matter which case, Alice (and Bob)'s state is of the form $|\phi\rangle^{\otimes k}$

Symmetric subspace:

$$\vee^k \mathbb{C}^d = \left\{ |\omega\rangle \in \left(\mathbb{C}^d\right)^{\otimes k} : P(\pi)|\omega\rangle = |\omega\rangle, \forall \pi \in S_k \right\}$$

$$\vee^k \mathbb{C}^d = \text{span}\left\{ |\phi\rangle^{\otimes k} : |\phi\rangle \in \mathbb{C}^d \right\}$$

POVM in the symmetric subspace: $\sum_i M_i = \Pi_{\text{sym}}$

"standard POVM" in the symmetric subspace:

$$\left\{ \binom{d+k-1}{k} |u\rangle\langle u|^{\otimes k} \mathrm{d}u \right\}$$

[Harrow'13] The Church of the Symmetric Subspace

Warm-up: “partial” tomography?

- Alice performs ”standard POVM” in the symmetric subspace, gets result $|u\rangle$
- Bob performs ”standard POVM” in the symmetric subspace, gets result $|v\rangle$
- They compute a function of $|u\rangle$ and $|v\rangle$ (can be implemented with simultaneous message passing)

- How many copies does this algorithm require? $k = O(\sqrt{d})$
- This gives evidence that Alice and Bob cannot do better than $O(\sqrt{d})$

Consider one-way protocol

Perform POVM $\{M_i\}$,
obtain result i

i →

Which case are we in?

Consider one-way protocol

- Case 1 (same state): Bob's state gets updated after seeing i
- $\rho = \frac{\binom{d+k-1}{k}}{\mathrm{Tr}(M_i \Pi_{\mathrm{sym}})} \mathbb{E}_{|\phi\rangle \sim \mathbb{C}^d} \mathrm{Tr}(M_i |\phi\rangle\langle\phi|^{\otimes k}) |\phi\rangle\langle\phi|^{\otimes k}$

- Case 2 (independent state): Bob's state is always the "maximally mixed state"
- $\sigma_{\mathrm{m}} = \frac{\Pi_{\mathrm{sym}}}{\binom{d+k-1}{k}}$

Which case are we in?

- Result: when $k = o(\sqrt{d})$, they are indistinguishable

Proof of indistinguishability

- $\rho = \frac{\binom{d+k-1}{k}}{\mathrm{Tr}(M_i \Pi_{\mathrm{sym}})} \mathbb{E}_{|\phi\rangle\sim\mathbb{C}^d} \mathrm{Tr}(M_i |\phi\rangle\langle\phi|^{\otimes k}) |\phi\rangle\langle\phi|^{\otimes k}$ is indistinguishable from $\sigma_{\mathrm{m}} = \frac{\Pi_{\mathrm{sym}}}{\binom{d+k-1}{k}}$ when $k = o(\sqrt{d})$

- Proof: think about the “measure-and-prepare” channel
- $\mathrm{MP}(\tau) = \binom{d+k-1}{k} \mathbb{E}_{|\phi\rangle\sim\mathbb{C}^d} \mathrm{Tr}(\tau \cdot |\phi\rangle\langle\phi|^{\otimes k}) |\phi\rangle\langle\phi|^{\otimes k}$
- Using Chiribella’s theorem [Chiribella’11], we show that the output of MP is indistinguishable from σ_{m} regardless of the input, when $k = o(\sqrt{d})$
- Can be generalized to a lower bound against arbitrary interactive communication

Discussion

- **Theorem.** The optimal sample complexity for distributed quantum inner product estimation is $k = \Theta(\max\{\frac{1}{\varepsilon^2}, \frac{\sqrt{d}}{\varepsilon}\})$ across all measurement and communication models

- What happens when allow a small amount (say $O(\log n)$ qubits) of quantum communication?
- Upper and lower bounds for other distributed quantum property estimation problems?