


Generalized Inner Product Estimation with Limited Quantum Communication

Srinivasan Arunachalam 

IBM Quantum, Almaden, CA, USA

Louis Schatzki  

Electrical and Computer Engineering, University of Illinois Urbana-Champaign, IL, USA

Abstract

In this work, we consider the fundamental task of distributed inner product estimation when allowed limited communication. Suppose Alice and Bob are given k copies of an unknown n -qubit quantum state $|\psi\rangle, |\phi\rangle$ respectively, are allowed to send q qubits to one another, and the task is to estimate $|\langle\psi|\phi\rangle|^2$ up to constant additive error. We show that $k = \Theta(\sqrt{2^{n-q}})$ copies are essentially necessary and sufficient for this task (extending the work of Anshu, Landau and Liu (STOC'22) who considered the case when $q = 0$). Additionally, we also consider the task when the goal of the players is to estimate $|\langle\psi|M|\phi\rangle|^2$, for arbitrary Hermitian M . For this task we show that certain norms on M determine the sample complexity of estimating $|\langle\psi|M|\phi\rangle|^2$ when using only classical communication.

2012 ACM Subject Classification Theory of computation \rightarrow Quantum information theory

Keywords and phrases Quantum property testing, Quantum Distributed Algorithms

Digital Object Identifier 10.4230/LIPIcs.STACS.2025.11

Related Version *Full Version:* <https://arxiv.org/pdf/2410.12684> [4]

Funding We acknowledge support from IBM through the Illinois-IBM Discovery Accelerator Institute.

1 Introduction

The construction of quantum networks is a major goal in quantum information science [22, 27, 24, 15]. These networks are envisioned to consist of disjoint quantum processing nodes with quantum communication interlinks, likely implemented through photonics. While the nodes may consist of the same type of qubit, this is not a requirement, and work has gone into designing heterogeneous networks where the nodes have different physical compositions [19]. With this in mind it is of interest to develop distributed protocols in a platform-independent framework.

In this paper we consider a natural *distributed* variant of the inner product estimation problem: suppose there are two spatially separated parties Alice and Bob who are each given copies of an unknown d -dimensional quantum state $|\psi\rangle$ and $|\phi\rangle$ respectively and their goal is to compute $|\langle\psi|M|\phi\rangle|^2$ for some Hermitian operator M . We will be interested in the difficulty of this task depending on how much entanglement Alice and Bob share. The core motivation for our setup is that distributing entanglement between nodes in a network is not a free resource and adds time and difficulty to a protocol. Thus, it is useful to minimize the amount of entanglement required. Similar concerns arise in various papers on distributed quantum computing [28, 18, 1].

We call this the *generalized distributed inner product estimation* problem (GIPE). The setting of $M = \mathbb{I}$ in GIPE (which we refer to as *distributed inner product estimation* DIPE) is one of the important steps in cross-platform verification proposed in [16] where the goal is to test if two quantum computers (say built on different platforms) are preparing the same quantum state; a fundamental problem for near-term devices. A natural constraint in these papers is how the quantum computers communicate with one another and in these works they



© Srinivasan Arunachalam and Louis Schatzki;
licensed under Creative Commons License CC-BY 4.0

42nd International Symposium on Theoretical Aspects of Computer Science (STACS 2025).

Editors: Olaf Beyersdorff, Michał Pilipczuk, Elaine Pimentel, and Nguyễn Kim Thăng;

Article No. 11; pp. 11:1–11:17



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



consider either only classical communication or limited quantum communication (the latter being the focus of our work). Apart from verification, on a theoretical level, understanding the complexity of a question, as fundamental as inner product estimation, in a distributed manner is a natural question.

The seminal work of Buhrman, Cleve, Watrous and Wolf [9] on *quantum fingerprinting* introduced the so-called *swap test*. This subroutine takes as input one copy of unknown quantum states $|\psi\rangle$ and $|\phi\rangle$ and outputs a bit whose bias equals $|\langle\psi|\phi\rangle|^2$; allowing us to estimate the *inner product* between two unknown quantum states without knowing anything about the states themselves. This subroutine has found applications in several areas in quantum computing from complexity theory, learning theory, entanglement theory, to optimization. Although so widely used, there still remain some open fundamental questions about estimating inner products, which is the topic of this work.

There are two techniques to solve GIPE under different settings: (i) If allowed quantum communication, then Alice can send a few copies of the d -dimensional quantum state over to Bob, then Bob could perform a variant of the swap test to estimate $|\langle\psi|M|\phi\rangle|^2$; (ii) if allowed only classical communication, i.e., Alice and Bob only send classical messages, then a recent work of Anshu et al. [3] showed that $\Theta(\sqrt{d})$ copies of $|\psi\rangle, |\phi\rangle$ are necessary and sufficient in order to estimate $|\langle\psi|\phi\rangle|^2$. This setting is often referred to as *local operations and classical communication* (LOCC) and has received a lot of attention in quantum information theory [13, 7, 17]. Although (ii) is surprising in that, the sample complexity is quadratically better than full-state tomography, the exponential (since we typically work with $d = 2^n$, where n is the number of qubits) nature of the complexity seems rather large in practice. Their work raises two natural questions, which will be the focus here:

1. With the advent of near-term quantum devices, sending the full-quantum state (as in the protocol (i) above) might be hard, but it is plausible that one could send a *few* qubits of quantum message. If this were possible, then what is the complexity of DIPE when allowed q -qubit messages?
2. The work of Anshu et al. [3] showed how to solve GIPE when $M = \mathbb{I}$, but other natural properties of quantum states are captured by letting M be an arbitrary operator, for example, predicting few-body properties, such as two-point correlation functions, entanglement entropy of small subsystems, expectation values of observables, projector onto a subspace or perhaps a Pauli string, then M may be different than \mathbb{I} . In this case, understanding the complexity of GIPE for arbitrary M is relevant.

These questions were explicitly raised by Anshu et al. [3] (and as open question 21 in [2]) and in this work we answer both.

1.1 Results

Our first result considers the setup where Alice and Bob may communicate a few qubits in an interactive protocol but not necessarily enough to teleport copies of their states. For simplicity assume that Alice and Bob each have n -qubit states and can transmit in total q qubits of communication. As we mentioned above, the sample complexity when $q = 0$ was shown to be $\sqrt{2^n}$ by [3] and when $q = n$, the sample complexity is $O(1)$. Is there some saving in sample complexity when allowed q qubits of communication? Our first result shows a smooth interpolation between both these settings.

► **Theorem 1 (Informal).** *Suppose Alice and Bob are given k copies of n -qubit states $|\psi\rangle$ and $|\phi\rangle$ respectively can communicate $\Theta(q)$ qubits along with performing LOCC. Then it is necessary and sufficient for them to obtain*

$$k = \Theta\left(2^{(n-q)/2}\right),$$

*copies of their states to estimate $|\langle\phi|\psi\rangle|^2$ up to constant accuracy with high probability.*¹

The upper bound $k = \mathcal{O}\left(2^{(n-q)/2}\right)$ shows that quantum communication may help. However, to reach sub-exponential scaling, a large amount of quantum communication is necessary i.e., unless $q = n - \text{polylog}(n)$, our lower bound is still exponential in n . Informally, entanglement helps, but not too much.

Our next result concerns the sample complexity, under LOCC, of GIPE, i.e., for arbitrary Hermitian operators M . A natural possibility when estimating $|\langle\psi|M|\phi\rangle|^2$ is that, perhaps Alice and Bob only need to confirm that their states have a large overlap in one of several smaller subspaces “within M ”. Then, it is conceivable that this task would be much easier than full inner product estimation (i.e., when $M = \mathbb{I}$). We show that this is indeed the case! Let M_ε be M restricted to subspaces with eigenvalue of magnitude at least $\varepsilon/2$. We prove the following near-optimal results for general GIPE.

► **Theorem 2 (Informal).** *For every M with $\|M\| \leq 1$ and $\varepsilon > 0$, to estimate $|\langle\phi|M|\psi\rangle|^2$ to error ε , it is sufficient to obtain*

$$k = \mathcal{O}(\max\{1/\varepsilon^2, \|M_\varepsilon\|_2/\varepsilon\})$$

copies of $|\psi\rangle, |\phi\rangle$ and

$$k = \Omega(\max\{1/\varepsilon^2, \|M_\varepsilon\|_2/\sqrt{\varepsilon}\})$$

copies are necessary.

Interestingly, this implies that M having both positive and negative eigenvalues may not render estimation harder than definite operators. Take Pauli strings as an example, an important set of operators in quantum information. These have eigenspaces of eigenvalues ± 1 each of dimension 2^{n-1} . Then, in evaluating $|\langle\phi|M|\psi\rangle|^2$ there may be cancellations between components of the states lying in eigenspaces of different signs. However, our results show that estimating such an M is no harder than estimating \mathbb{I} , which is positive definite and basis independent. While our upper and lower bounds do not exactly match, this is partially due to choice of presentation. The lower bound we prove is actually $\Omega(\|M_\varepsilon\|_1/\varepsilon\sqrt{d_\varepsilon})$, where d_ε is the dimension of the support of M_ε . Converting 1 to 2 norm yields the lower bound as presented above. However, in some cases this more precise bound is tight. Take, for example, M a projector onto a subspace or a Pauli string. Then, we obtain the lower bound $\Omega(\|M_\varepsilon\|_2/\varepsilon)$, *exactly* matching the upper bound.

Proof sketch. In [3], they prove their lower bound using a very interesting connection to quantum cloning: in particular, they show that if one could solve a decision-version of DIPE, then one could have cloned the states “well-enough”. One of the main challenges in proving our lower bound was that the cloning-based lower bounds do not work when

¹ We will state the constants in this theorem more clearly when proving this theorem.

allowed even few qubits of communication in the cloning channel (which is the setting in which we would like lower bounds). Instead, here we showed that one could use the notion of *robustness of entanglement*, a concept from entanglement theory and split a protocol with a quantum channel into a sum of LOCC protocols. Proving this is a small technical lemma which we combine with the lower bound of [3] to obtain our overall lower bound of $\sqrt{2^{n-q}}$. In order to prove our upper bound, we use ideas from the celebrated Johnson-Lindenstrauss lemma. Our main idea is to perform projections onto random subspaces, i.e., Alice and Bob pick random subspaces and project their states onto these subspaces, maintaining the inner product between their states with high probability. Further, this protocol is completely agnostic to the states they started with. Now holding smaller-dimensional projections of their states, they can use classical communication to detect when their subspace-projections have collisions, use their shared quantum channel and perform SWAP tests to estimate the inner product within the subspace. With some analysis, we are able to show that the overall sample complexity for this scales as $\sqrt{2^{n-q}}$ as well, matching our lower bound.

For estimating bilinear forms, our protocol is fairly similar to the one in [3] except that we need to deal with some technicalities due to the Hermitian operator M . The key observation we first make is that $|\langle \phi | M | \psi \rangle|^2$ and $|\langle \phi | M_\varepsilon | \psi \rangle|^2$ can only differ by at most $\varepsilon/2$. So it suffices for Alice and Bob to restrict themselves to the support of M_ε . Then, they can simultaneously measure the magnitude of the components of their states in this subspace as well as the weighted inner product of said components. The calculations here are similar to the one in [3] wherein one needs to compute the mean variance of the estimator, but a little bit more technical in order to bound all the terms in terms of $\|M_\varepsilon\|_2$. The lower bound follows from realizing that M_ε when restricted to its support, is not too far off from identity and then we can invoke the lower bounds for inner product estimation can be applied. We remark that the lower bound here is much more subtle than the one in [3] since one needs to project onto the eigenspaces of M carefully in order to get the optimal dependence on $\|M_\varepsilon\|_2$.

1.2 Related works and open questions

Like we mentioned earlier, our work answers two open questions raised by Anshu et al. [3]. There have been a few more papers since their work: Hinsche et al. [21] look at specific instantiations of DIPE (when the protocol used by Alice and Bob are Pauli sampling and the quantum states are structured) and prove some positive and negative results to this end. Chen et al. [10] extend the work of [3] by proving better bounds for non-trace-preserving quantum channels, Chen et al. [11] also looked at quantum property testing in the LOCC model.

Distinguishability under LOCC has a rich history in the quantum information literature. Bennett, et al. demonstrated that there exist product bases which are easily distinguished, but are indistinguishable under LOCC [7]. Enough entanglement renders this task easy as the two parties can simply teleport states. Subsequent works showed that much less entanglement may suffice for this task [14, 29]. Unlike our framework, these papers concerned themselves with measurements of a single copy and further identifying states from some known ensemble, rather than learning some property.

A natural open question. A tantalizing open question left open by [3] and also this work is the analogue of DIPE in the mixed case setting. Here the goal is as follows: Alice has copies of ρ , Bob has copies of σ and they engage in LOCC. They need to distinguish if $\rho = \sigma$ or $\|\rho - \sigma\|_{tr} \geq \varepsilon$, promised one of them is the case. Using the bounds obtained in [3], we have an upper bound of $O(d^2/\varepsilon^4 + d^{1.5}/\varepsilon^2)$. As for lower bounds, [5] showed that if Alice also had copies of σ (instead of Bob having them), then one can solve this property testing task using

$\mathcal{O}(d/\varepsilon^2)$ copies (this procedure requires a highly-entangling operation between copies of ρ and σ). Furthermore, if we restricted the measurements to be single-copy measurements and if we fix $\sigma = \mathbb{I}/d$, then $\Omega(d^{3/2})$ copies are necessary [12]. Despite several attempts in making progress on this question, we have not been able to obtain a lower bound better than $\Omega(d)$ and believe that the sample complexity of mixed state DIPE should be $\mathcal{O}(d)$. We leave this as an open question.

2 Preliminaries

2.1 Quantum States and Measurements

Pure quantum states are unit vectors in \mathbb{C}^d . A qubit is a state in \mathbb{C}^2 and n -qubits is a state in $(\mathbb{C}^2)^{\otimes n} \cong \mathbb{C}^{2^n}$. Mixed states, also known as density matrices, ρ , are positive semi-definite operators on \mathbb{C}^d with unit trace. Pure states correspond to rank 1 mixed states and we will routinely use the notation $|\psi\rangle$ to refer to the density matrix corresponding to a pure state $|\psi\rangle$. Measurements of quantum systems are described by positive operator valued measures (POVMs), ensembles of positive semi-definite operators $\{M_i\}_i$ such that $\sum_i M_i = \mathbb{I}$. The probability of observing an outcome i is given by $\text{Tr}[M_i\rho]$. By $\|M\|$ we denote the operator norm of an operator.

In quantum computation states are manipulated via unitary evolution. That is, $|\psi\rangle$ is mapped to $U|\psi\rangle$ for some unitary U . An important unitary we will make repeated usage of is the SWAP gate on a bipartite state. This has the action $\text{SWAP}|\psi\rangle \otimes |\phi\rangle = |\phi\rangle \otimes |\psi\rangle$.

Operators on two Hilbert spaces \mathcal{H}_A and \mathcal{H}_B lie in the tensor product space $\mathcal{B}(\mathcal{H}_A \otimes \mathcal{H}_B)$. A positive semi-definite operator M acting on these tensor product space is said to be separable if it admits a decomposition $M = \sum_i A_i \otimes B_i$, where each A_i and B_i are positive semi-definite. We will be interested in separable POVMs, where each aspect M_i is separable. These are measurements that cannot create entanglement between distributed parties. However, such measurements may still go beyond those achievable with classical communication. With this in mind, one can define local operations and classical communication (LOCC) [13]. This can roughly be defined as all protocols where Alice performs some measurement in her lab, communicates the result to Bob, who then performs a measurement in his lab and communicates the result to Alice and so on. However, the set of LOCC measurements is mathematically unwieldy and it is generally easier to prove results regarding separable measurements. An e -bit is a standard resource of entanglement in quantum information. This is a shared two qubit resource state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. With n e -bits and classical communication, two distributed parties can teleport an n qubit state [23].

2.2 Subroutines

We will make repeated reference to the SWAP test, which can be used to estimate the overlap between two states $|\psi\rangle$ and $|\phi\rangle$ [6, 9].

► **Definition 3** (SWAP test). *Given two mixed states ρ and σ and an ancilla qubit initialized in the state $|0\rangle_E$, the SWAP test performs the unitary $(H_E \otimes \mathbb{I}_{AB})(|0\rangle\langle 0| \otimes \mathbb{I}_{AB} + |1\rangle\langle 1| \otimes \text{SWAP}_{AB})(H_E \otimes \mathbb{I}_{AB})$ and measures the ancilla in the computation basis. The measurement probabilities are given by*

$$p(0) = \frac{1 + \text{Tr}[\rho\sigma]}{2}, \quad p(1) = \frac{1 - \text{Tr}[\rho\sigma]}{2}.$$

11:6 Generalized Inner Product Estimation with Limited Quantum Communication

Via standard amplification arguments, $\mathcal{O}(1/\varepsilon^2)$ trials are sufficient to estimate $\text{Tr}[\psi\phi]$ to error ε . Using block-encodings, this can be extended to estimating $|\langle\phi|M|\psi\rangle|^2$ for an arbitrary hermitian M such that $\|M\| \leq 1$. Again, standard amplification arguments show that $\mathcal{O}(1/\varepsilon^2)$ trials suffices to estimate $|\langle\phi|M|\psi\rangle|^2$. We will also use the standard POVM on the symmetric subspace. The symmetric group S_k has a natural action on the state space of $(\mathbb{C}^d)^{\otimes k}$ given by

$$\pi \bigotimes_{i=1}^k |\psi_i\rangle = \bigotimes_{i=1}^k |\psi_{\pi^{-1}(i)}\rangle \quad \forall \pi \in S_k . \quad (1)$$

► **Definition 4** (Symmetric Subspace). *The k -copy symmetric subspace of a vector space $V \cong \mathbb{C}^d$, is given by*

$$\vee^k \mathbb{C}^d = \left\{ |\psi\rangle \in (\mathbb{C}^d)^{\otimes k} \mid \pi |\psi\rangle = |\psi\rangle \forall \pi \in S_k \right\} .$$

The symmetric subspace has a natural spanning set given by $\left\{ |\phi\rangle^{\otimes k} \mid |\phi\rangle \in \mathbb{C}^d \right\}$ [20]. Due to $\vee^k \mathbb{C}^d$ being an irreducible representation of the unitary group (under the action $U^{\otimes k}$), we can define a uniform POVM on the symmetric subspace, which is known as the standard POVM.

► **Definition 5** (Standard POVM on $\vee^k \mathbb{C}^d$). *The standard POVM on $\vee^k \mathbb{C}^d$ is the continuous POVM with elements*

$$\left\{ \binom{d+k-1}{k} |\varphi\rangle\langle\varphi|^{\otimes k} d\varphi \mid |\varphi\rangle \in \mathbb{C}^d \right\} . \quad (2)$$

We require one last fact about the symmetric subspace:

► **Fact 6** (Haar moments). *If $|\psi\rangle$ is drawn from the Haar measure on \mathbb{C}^d , then*

$$\mathbb{E}_\psi[|\psi\rangle] = \frac{\mathbb{I}}{d} , \quad \mathbb{E}_\psi[|\psi\rangle\langle\psi|^{\otimes 2}] = \frac{1}{d(d+1)} (\mathbb{I} + \text{SWAP}) . \quad (3)$$

3 Estimating Inner Product

In this section we consider the following task: Alice and Bob each have copies of unknown d -dimensional states $|\phi\rangle$ and $|\psi\rangle$. They may use any amount of classical communication and some restricted quantum communication. Their goal is to estimate $|\langle\phi|\psi\rangle|^2$ using as few samples of $|\psi\rangle$ and $|\phi\rangle$ as possible. Here we will assume the desired measurement precision is some constant.

► **Theorem 7.** *Suppose Alice and Bob share a q -qubit entangled state, can perform measurements on k copies of their respective d -dimensional states $|\psi\rangle$ and $|\phi\rangle$ respectively and engage in unbounded classical communication. If they are able to estimate $|\langle\psi|\phi\rangle|^2$ to accuracy ε with high probability, then $\Omega(\max\{1/\varepsilon^2, \sqrt{d/q}/\varepsilon\})$ are necessary.*

► **Theorem 8.** *Suppose Alice and Bob share a $10q$ -qubit entangled state, can perform measurements on k copies of their respective d -dimensional states $|\psi\rangle$ and $|\phi\rangle$ respectively and engage in unbounded classical communication. It suffices to obtain $k = O(\sqrt{d/q}/\varepsilon^2)$ copies, in order to estimate $|\langle\psi|\phi\rangle|^2$ to accuracy ε with high probability.*

We remark that there is a constant-factor difference in the amount of entanglement bounds in the upper and lower bound. This largely seems to stem from the sample complexity still being greater than 1 even when $q = n$: say that Alice is able to transmit her state exactly and Bob does the swap test. Even then, in order to compute the inner product $|\langle\psi|\phi\rangle|^2$ to error 0.1, they need to repeat the swap test constantly many times. Our lower bound for this setting would imply a constant lower bound as well (matching the upper bound, but not “exactly” since it would be a constant-factor off). Regardless, our lower bound can be understood as saying that entanglement does not help for inner product estimation unless the shared entanglement dimension scales with d .

3.1 Lower bound in main theorem

To obtain a lower bound we consider the decision problem constructed in [3] (which is called the DIPE, *distributed inner product estimation* problem). Alice and Bob are promised to be in one of the following two scenarios. Their version restricts Alice and Bob to LOCC protocols. However, here we allow them to share some resource state σ as well.

- **Definition 9** (Distributed Inner Product Estimation, Decision Version). *Alice and Bob each have access to k copies of the states $|\phi\rangle$ and $|\psi\rangle$ in \mathbb{C}^d respectively and are asked to decide, using LOCC and perhaps a resource state σ , which of the following two scenarios they are in:*
- **YES:** $|\phi\rangle = |\psi\rangle$ and they are Haar random
 - **NO:** $|\phi\rangle, |\psi\rangle$ are Haar random.

Say that Alice and Bob are able to transmit a r dimensional quantum message. Because of quantum teleportation, this communication channel is equivalent, under LOCC, to sharing a r -dimensional maximally entangled state. Going forward, we let $\sigma = \frac{1}{r} \sum_{i,j=1}^r |ii\rangle\langle jj|$ be the state Alice and Bob share.

We also introduce a measure of entanglement called the *robustness of entanglement*.

- **Definition 10** (Robustness of entanglement [26]). *Any quantum state $\sigma \in \mathcal{B}(\mathcal{H}_A \otimes \mathcal{H}_B)$ can be decomposed as $\sigma = (1+s)\sigma^+ - s\sigma^-$, where σ^+ and σ^- are both separable states and $s \in \mathbb{R}_{\geq 0}$. The minimum value of s over all such decompositions is called the robustness of entanglement. We denote this minimum value by $E(\sigma)$.*

In particular, we will use the following lemma.

- **Lemma 11** ([26]). *If σ is a pure bipartite state, then the robustness of entanglement is $E(\sigma) = (\sum_i \lambda_i)^2 - 1$, where $\{\lambda_i\}_i$ are the Schmidt coefficients.*

Thus, the robustness of entanglement of σ is $E(\sigma) = r - 1$ (since all of its Schmidt coefficients are $1/\sqrt{r}$). In [3] they show the following result:

- **Theorem 12** ([3]). *Let M be a separable measurement, then*

$$|\mathbb{E}_{\phi,\psi} [M \cdot \text{Tr}[\phi^{\otimes k} \otimes (\phi^{\otimes k} - \psi^{\otimes k})]]| \leq e^{k^2/d} - 1. \quad (4)$$

With this we now prove our main theorem lower bound.

Proof of lower bound of Thm 7. Say that there exists a separable protocol \mathcal{A} that uses k copies of the states $|\phi\rangle$ and $|\psi\rangle$ and the resource state σ and returns an estimate of $|\langle\psi|\phi\rangle|^2$ to accuracy ε with probability at least $3/4$. Because of anti-concentration of the Haar measure, Alice and Bob will be able to use this protocol to solve problem 9 with high probability. Let $B([\phi], \varepsilon)$ be the ball of radius $\varepsilon \leq \pi/2$ around $[\phi]$ in $\mathbb{C}\mathbb{P}(d)$ with respect to the Fubini-Study

metric, which we will denote as $d([\phi, \psi]) = \arccos |\langle \psi | \phi \rangle|$. Then, it is known that, under the uniform distribution, the volume of $B([\phi], \varepsilon)$ for $\varepsilon \leq \pi/2$ is given by $\sin^{2d-2} \varepsilon$ [8]. If $|\langle \psi | \phi \rangle|^2 > \varepsilon$, then $d([\psi, \phi]) \leq \arccos \sqrt{\varepsilon}$ (and $\arccos \sqrt{\varepsilon} \leq \pi/2$ for $\varepsilon \in [0, 1]$). Thus, for constant ε it follows that $|\langle \phi | \psi \rangle|^2 \leq \varepsilon$ with exponentially (in d) high probability in case two of 9. Thus, if Alice and Bob use the protocol \mathcal{A} , they will solve problem 9 with probability at least $2/3$ as well. This implies that there is some separable POVM $\{M, \mathbb{I} - M\}$ such that

$$1/3 \leq |\mathbb{E}_{\phi, \psi} \text{Tr}[M(\phi^{\otimes k} \otimes (\phi^{\otimes k} - \psi^{\otimes k}) \otimes \sigma)]|. \quad (5)$$

We now show that the RHS of the inequality above can be upper bounded by $(1 + 2E(\sigma))(e^{k^2/d} - 1)$. To see this, split σ into $\sigma = (1 + E(\sigma))\sigma^+ - E(\sigma)\sigma^-$, where σ^+ and σ^- are both separable states. Then, it follows that

$$\begin{aligned} \text{Tr}[M(\phi^{\otimes k} \otimes (\phi^{\otimes k} - \psi^{\otimes k}) \otimes \sigma)] &= (1 + E(\sigma))\text{Tr}[M(\phi^{\otimes k} \otimes (\phi^{\otimes k} - \psi^{\otimes k}) \otimes \sigma^+)] \\ &\quad + E(\sigma)\text{Tr}[M(\phi^{\otimes k} \otimes (\psi^{\otimes k} - \phi^{\otimes k}) \otimes \sigma^-)]. \end{aligned} \quad (6)$$

Note that the sign of σ^- has been absorbed into the trace in the second term above. Now, because σ^\pm are separable, each term above could be replaced by a separable measurement on $\phi^{\otimes k} \otimes (\phi^{\otimes k} - \psi^{\otimes k})$ without any reference state. To see this, decompose σ^+ into $\sigma^+ = \sum_i p_i \rho_i^A \otimes \rho_i^B$. Using spectral decompositions $\rho_i^A = \sum_k \lambda_{i,k} |u_{i,k}\rangle\langle u_{i,k}|$ and $\rho_i^B = \sum_k \nu_{i,k} |v_{i,k}\rangle\langle v_{i,k}|$, we arrive at

$$\text{Tr}[M(\phi^{\otimes k} \otimes (\phi^{\otimes k} - \psi^{\otimes k}) \otimes \sigma^+)] \quad (7)$$

$$= \sum_i p_i \sum_{j,k} \lambda_{i,k} \nu_{i,j} \text{Tr}[M(\phi^{\otimes k} \otimes \psi^{\otimes k} \otimes |u_{i,j}\rangle\langle u_{i,j}| \otimes |v_{i,k}\rangle\langle v_{i,k}|)]. \quad (8)$$

As a separable measurement, we have $M = \sum_t A_t \otimes B_t$. Now define a new measurement

$$M' := \sum_t \sum_i p_i \left(\sum_j \lambda_{i,j} A_t^{i,j,j} \right) \otimes \left(\sum_k \nu_{i,k} B_t^{i,k,k} \right), \quad (9)$$

where $A_t = \sum_{j,j'} A_t^{i,j,j'} \otimes |u_{i,j}\rangle\langle u_{i,j'}|$ and $B_t = \sum_{k,k'} B_t^{i,k,k'} \otimes |v_{i,k}\rangle\langle v_{i,k'}|$. As a sum of positive operators, this is positive. As a convex combination of operators majorized by \mathbb{I} , this is also majorized by \mathbb{I} . It then follows that M' is a separable POVM. From Thm 12 it then follows that

$$\begin{aligned} |\mathbb{E}_{\phi, \psi} [\text{Tr}[M(\phi^{\otimes k} \otimes (\phi^{\otimes k} - \psi^{\otimes k}) \otimes \sigma^+)]]| &= |\mathbb{E}_{\phi, \psi} [\text{Tr}[M'(\phi^{\otimes k} \otimes (\phi^{\otimes k} - \psi^{\otimes k}) \otimes \sigma^+)]]| \\ &\leq e^{k^2/d} - 1. \end{aligned} \quad (10)$$

The same proof shows that

$$|\mathbb{E}_{\phi, \psi} [\text{Tr}[M(\phi^{\otimes k} \otimes (\psi^{\otimes k} - \phi^{\otimes k}) \otimes \sigma^-)]]| \leq e^{k^2/d} - 1. \quad (11)$$

Thus, we have that

$$|\mathbb{E}_{\phi, \psi} [\text{Tr}[M(\phi^{\otimes k} \otimes (\phi^{\otimes k} - \psi^{\otimes k}) \otimes \sigma)]]| \leq (1 + 2E(\sigma))(e^{k^2/d} - 1). \quad (12)$$

Putting the above along with our lower bound in Eq.(5) shows

$$1/3 \leq |\mathbb{E}_{\phi, \psi} \text{Tr}[M(\phi^{\otimes k} \otimes (\phi^{\otimes k} - \psi^{\otimes k}) \otimes \sigma)]|. \leq (1 + 2E(\sigma))(e^{k^2/d} - 1). \quad (13)$$

Rearranging, this implies that

$$k = \Omega\left(\sqrt{\frac{d}{\log E(\sigma)}}\right). \quad (14)$$

So, if Alice and Bob share q Bell pairs, then $E(\sigma) = 2^q - 1$ and we get the desired lower bound.

Using techniques similar to [3, Section 5.3], wherein they showed how to use standard concentration techniques to reduce the decision version of DIPE (for which we have a $\Omega(\sqrt{d/q})$ lower bound) to prove a lower bound for estimating the inner product between $|\phi\rangle$ and $|\psi\rangle$ with a factor of $1/\varepsilon$. In particular, they prove that if there exists a protocol that solves the ε -version of inner product estimation then it can be used to solve DIPE with high probability. Then, the lower bound on DIPE follows from their reduction. \blacktriangleleft

3.2 Upper bound in main theorem

We now give a protocol which achieves the promised sample complexity of Thm 8. At a high level, the protocol works by randomly projecting $|\phi\rangle$ and $|\psi\rangle$ into small-dimensional subspaces. Since such a projection maintains the inner product between two states with high probability, Alice can simply send Bob these smaller dimensional states and Bob can perform a swap test. The reader may notice that this assumes that Alice can transmit an arbitrary state in a q -dimensional subspace of a d -dimensional space. This, however, does not lead to any issues since Alice *knows* the subspace the state lies in. Then, any maximally entangled state on a q -dimensional subspace can be converted to a maximally entangled state on the desired subspace via a unitary transformation. From there, Alice and Bob simply perform a teleportation protocol.

► **Theorem 13.** *Let $q = \Omega(\log d)$. Using $k = \mathcal{O}(\sqrt{d/q})$ copies of $|\psi\rangle$ and $|\phi\rangle$, there is a protocol that uses $\Theta(1)$ r -dimensional quantum messages, single copy measurements, and returns $|\langle\psi|\phi\rangle|^2$ up to constant error with high probability*

Proof. Let $\psi = |\psi\rangle\langle\psi|$ and $\phi = |\phi\rangle\langle\phi|$. Fix some unitary U . For convenience of notation we will assume that r divides d (working with qubits, if Alice and Bob may communicate $\log(r)$ qubits then this holds). Divide the d -dimensional Hilbert space into d/r subspaces each of dimension r . Let this decomposition be given by a collection of orthogonal projectors $\{P_i\}_{i=1}^{d/r}$. Alice and Bob perform the protocol in Figure 1.

Our theorem will follow by repeating the protocol above constantly many times, each time by sending a r -dimensional quantum message. To see the correctness, we will require the following fact.

► **Fact 14** ([25, Fact 2]). *Let $1 \leq q \leq d$. Let $v \in \mathbb{C}^d$ be a unit vector and P_i be an arbitrary projector onto a r -dimensional subspace. Let U be drawn from the unitary Haar measure. Then,*

$$\Pr_U \left[\|P_i U v\|_2 \notin (1 \pm \Delta) \sqrt{\frac{r}{d}} \right] \leq 4 \exp \left\{ -\frac{\Delta^2 r}{16} \right\}, \quad (15)$$

By linearity this extends to arbitrary vectors in \mathbb{C}^d , as the following corollary demonstrates.

► **Corollary 15.** *Let $1 \leq r \leq d$. Let $u \in \mathbb{C}^d$ be any vector and P_i an arbitrary projector onto a r -dimensional subspace. Let U be drawn from the unitary Haar measure. Then,*

$$\Pr_U \left[\|P_i U u\|_2 \notin (1 \pm \Delta) \sqrt{\frac{r}{d}} \|u\|_2 \right] \leq 4 \exp \left\{ -\frac{\Delta^2 r}{16} \right\}, \quad (16)$$

11:10 Generalized Inner Product Estimation with Limited Quantum Communication

Input: Alice gets $|\psi\rangle^{\otimes k}$, Bob gets $|\phi\rangle^{\otimes k}$

Output: ε -approximation of $|\langle\psi|\phi\rangle|^2$

1. Using shared randomness, Alice and Bob sample a unitary U from Haar measure.
2. Alice and Bob apply the POVM $\{UP_iU^\dagger\}_{i=1}^{\lceil d/r \rceil}$ on each of the k copies and record which copies were projected onto which subspaces.
3. Using a two-way classical communication they determine which copies lie now in the same subspaces. Alice sends Bob a constant number of such projected states which can be paired with a post-projection state of Bob's. Say the number of such pairs is m .
4. Bob performs a SWAP test between the post-projected pairs of states lying in the same subspaces. We say a SWAP test succeeds if Bob measures $|0\rangle$ on the ancilla register. Let the number of successes be s .
5. Bob declares the inner product between $|\psi\rangle$ and $|\phi\rangle$ to be $2s/m - 1$.

■ **Figure 1** Protocol to estimate inner product using shared entanglement.

Proof. This follows readily from Fact 14 via linearity. Apply Fact 14 to $\frac{u}{\|u\|_2}$ while multiplying both $\|P_iUu/\|u\|_2\|_2$ and $(1 \pm \Delta)\sqrt{\frac{r}{d}}$ by $\|u\|_2$. ◀

For a fixed unitary U , let $|\psi'_i\rangle = \frac{P_iU|\psi\rangle}{|P_iU|\psi|}$ and $|\phi'_i\rangle = \frac{P_iU|\phi\rangle}{|P_iU|\phi|}$ be the normalized projections of $|\psi\rangle$ and $|\phi\rangle$ into the subspace given by P_i . We will now show that, with high probability over the choice of U , $|\langle\psi'_i|\phi'_i\rangle|^2 \approx |\langle\psi|\phi\rangle|^2$.

► **Lemma 16.** *Let $\Delta > 0$. Then,*

$$\Pr_U \left[\left| |\langle\phi'_i|\psi'_i\rangle|^2 - |\langle\psi|\phi\rangle|^2 \right| \leq 16\Delta \text{ for every } i \in [d/r] \right] \geq 1 - O(d/r \cdot \exp(-\Delta^2 r/16)). \quad (17)$$

Proof. Using Fact 14, with probability at least $1 - 16 \exp\{-\frac{\Delta^2 r}{16}\}$, the following four conditions hold

$$\|P_iU|\psi\rangle\|_2, \|P_iU|\phi\rangle\|_2 \in (1 \pm \Delta)\sqrt{\frac{r}{d}}, \quad (18)$$

$$\|P_iU(|\psi\rangle - |\phi\rangle)\|_2 \in (1 \pm \Delta)\sqrt{\frac{r}{d}}\|\psi - \phi\|_2, \quad (19)$$

$$\|P_iU(|\psi\rangle - i|\phi\rangle)\|_2 \in (1 \pm \Delta)\sqrt{\frac{r}{d}}\|\psi - i\phi\|_2. \quad (20)$$

We now follow steps similar to those of [25, Theorem 1] to obtain

$$|\langle\phi_i|\psi_i\rangle - \langle\psi|\phi\rangle| \leq 8\Delta. \quad (21)$$

Let $x, y \in \mathbb{C}$ be such that $|x - y| \leq \Delta$. Then,

$$|x|^2 - |y|^2 = (|x| - |y|)(|x| + |y|) \leq \Delta(|x| + |y|). \quad (22)$$

In our case, this implies that

$$|\langle\phi_i|\psi_i\rangle|^2 \in |\langle\psi|\phi\rangle|^2 \pm 16\Delta. \quad (23)$$

Taking a union bound over all d/r subspaces completes the proof. ◀

Now, after receiving m pairs of post-projected states, Bob then performs m many SWAP tests. A SWAP test on the pair $|\psi'_i\rangle$ and $|\phi'_i\rangle$ succeeds with probability $\frac{1}{2} + \frac{1}{2}|\langle\phi'_i|\psi'_i\rangle|^2 \in \frac{1}{2} + \frac{1}{2}|\langle\phi|\psi\rangle| \pm 8\Delta$. The expected value of the sample average s/m then satisfies

$$\left| \mathbb{E} \left[\frac{s}{m} \right] - \frac{1}{2} - \frac{1}{2}|\langle\psi|\phi\rangle|^2 \right| \leq 8\Delta, \quad (24)$$

which implies $2\mathbb{E} \left[\frac{s}{m} \right] - 1 \in |\langle\phi|\psi\rangle|^2 \pm 16\Delta$, and by a Hoeffding bound, we get

$$2 \left| \frac{s}{m} - \mathbb{E} \left[\frac{s}{m} \right] \right| \leq \delta, \quad (25)$$

with probability at least $1 - 2\exp(-2m\delta^2)$. In this case, the total error of the estimator is at most $16\Delta + \delta$. For constant error, $m = O(1)$ suffices.

It remains to determine k needs to be to obtain $m = O(1)$ projected pairs with high probability. To this end, let $E_{a,b}$ be an indicator variable for if Alice's a th projection falls into the same subspace as Bob's b th projection. The expected number of collisions is given by

$$\sum_{a,b} \mathbb{E}[E_{a,b}] = k^2 \sum_i \text{Tr}[P_i U \psi U^\dagger] \text{Tr}[P_i U \phi U^\dagger] \geq k^2 (1 - \Delta)^4 \frac{r}{d}, \quad (26)$$

where the inequality follows from $\|P_i U |\psi\rangle\|, \|P_i U |\phi\rangle\| \geq (1 - \Delta)\sqrt{\frac{r}{d}}$. By a Hoeffding bound, it suffices to take $k = \Omega(\sqrt{d/r})$ to obtain a collision with high probability. As we only require a constant number of pairs, this can simply be repeated a constant number of times to obtain a constant number of collisions with high probability. Thus, the protocol succeeds for $k = \Omega(\sqrt{d/r})$. It remains to pick Δ and δ such that the claims go through: we require that $\frac{d}{r} \exp(-\Delta^2 r/16) = O(1)$, which can be achieved by letting $\Delta = O(1)$ and $r = \Omega(\log d)$. ◀

4 Generalized Distributed Inner Product Estimation

Now we turn to bilinear forms $f : \mathbb{C}^d \otimes \mathbb{C}^d \rightarrow \mathbb{C}$. Any such form can be expressed as $f(u, v) = u^\dagger M v$ for a matrix M . Here we will assume that M is Hermitian, which implies that $f(u, v) = \overline{f(v, u)}$. Without loss of generality we will assume M to be normalized such that $\|M\| = 1$. Due to the unphysical nature of global phases, $f(|\psi\rangle, |\phi\rangle)$ is less meaningful than $|f(|\psi\rangle, |\phi\rangle)|^2$ and we concern ourselves entirely with this value instead. Of course, the special case of $M = \mathbb{I}$ corresponds to inner product estimation. Fixing such an M , let P_ε be a projector which annihilates all eigenspaces of M of eigenvalue less than $\varepsilon/2$. We define d_ε to be the dimension of the support of P_ε . Now define $M_\varepsilon = P_\varepsilon M P_\varepsilon$. The intuition behind this is that we can drop eigenspaces with small weights in estimating $|\langle\phi|M|\psi\rangle|^2$. The following lemma formalizes this.

► **Lemma 17.** *Let M be a Hermitian operator and $M_\varepsilon = P_\varepsilon M P_\varepsilon$, then*

$$\left| |\langle\phi|M|\psi\rangle|^2 - |\langle\phi|M_\varepsilon|\psi\rangle|^2 \right| \leq \varepsilon/2. \quad (27)$$

Proof.

$$|\text{Tr}[M\psi M\phi] - \text{Tr}[M_\varepsilon\psi M_\varepsilon\phi]| = |\text{Tr}[\psi(M\phi M - M_\varepsilon\phi M_\varepsilon)]| \quad (28)$$

$$= |\text{vec}(\psi)^\dagger (M \otimes M - M_\varepsilon \otimes M_\varepsilon)^T \text{vec}(\phi)| \quad (29)$$

$$\leq \|M \otimes M - M_\varepsilon \otimes M_\varepsilon\| \leq \varepsilon/2, \quad (30)$$

where the second line uses the linear map $\text{vec} : \mathcal{B}(\mathcal{H}) \rightarrow \mathcal{H} \otimes \mathcal{H}$ defined by $\text{vec}(|i\rangle\langle j|) = |i\rangle \otimes |j\rangle$ and the identity $\text{vec}(AXB) = (A \otimes B)^T \text{vec}(X)$. ◀

Input: An operator M . Alice gets $|\psi\rangle^{\otimes k}$, Bob gets copies of $|\phi\rangle^{\otimes k}$.

Output: An ε -approximation of $\langle\psi|M|\phi\rangle^2$.

1. Alice and Bob each perform the two-outcome measurement $\{P_\varepsilon, \mathbb{I} - P_\varepsilon\}$ on each of the k copies of their respective states and obtain s_a and s_b copies of the states projected into one of the two subspace. If $s_a = 0$ or $s_b = 0$ then Alice and Bob simply output 0.
2. Alice and Bob independently perform the standard POVM in the symmetric subspace on the support of M_ε , obtaining the classical outcomes $|u\rangle$ and $|v\rangle$.
3. Alice communicates $|u\rangle$ and s_a to Bob who outputs

$$w := \frac{(d_\varepsilon + s_a)(d_\varepsilon + s_b)}{k^2} |\langle u|M_\varepsilon|v\rangle|^2 - \frac{\text{Tr}[M^2]}{k^2}. \quad (31)$$

■ **Figure 2** Protocol to estimate $|\langle\phi|M_\varepsilon|\psi\rangle|^2$ using only LOCC.

Using this lemma, if Alice and Bob can estimate $|\langle\phi|M_\varepsilon|\psi\rangle|^2$ up to precision $\varepsilon/2$, that directly implies an ε -approximation to the quantity $|\langle\phi|M|\psi\rangle|^2$. For the remainder of this section, we will be primarily focused on upper and lower bounds for estimating $|\langle\phi|M_\varepsilon|\psi\rangle|^2$.

With this notation, in this section our main result will be the following theorem, proving close-to-optimal upper and lower bounds for estimating bilinear forms on $|\psi\rangle, |\phi\rangle$.

► **Theorem 18.** *For every M with $\|M\| \leq 1$, with only LOCC, it is sufficient to obtain*

$$k = O(\max\{1/\varepsilon^2, \|M_\varepsilon\|_2/\varepsilon\})$$

copies of $|\psi\rangle, |\phi\rangle$ to estimate $|\langle\phi|M|\psi\rangle|^2$ to error ε with high prob. and it necessary to obtain

$$k = \Omega(\max\{1/\varepsilon^2, \|M_\varepsilon\|_2/\sqrt{\varepsilon}\})$$

copies to produce an ε -approximation to $|\langle\phi|M|\psi\rangle|^2$.

4.1 Upper Bound

In Figure 2 we outline a protocol that estimates $|\langle\phi|M_\varepsilon|\psi\rangle|^2$. Recall that the standard POVM on the symmetric subspace, def 5, has continuous aspects $\left\{ \binom{d+k-1}{k} \varphi^{\otimes k} \mid |\varphi\rangle \in \mathbb{C}^d \right\}$. This can be extended to $\vee^k W$ for arbitrary subspaces $W \subseteq \mathbb{C}^d$:

$$\left\{ \binom{\dim W + k - 1}{k} \varphi^{\otimes k} \mid |\varphi\rangle \in W \right\}. \quad (32)$$

The second step of the protocol of Figure 2 has Alice and Bob implementing this POVM for $\text{Im}P_\varepsilon \subseteq \mathbb{C}^d$. Since they may not have k copies after the projection step, this is a POVM on $\vee^{s_a} \text{Im}P_\varepsilon$ and $\vee^{s_b} \text{Im}P_\varepsilon$, respectively. First note that if $P_\varepsilon|\psi\rangle = 0$ or $P_\varepsilon|\phi\rangle = 0$ then they always output 0, which must be a good estimate in this case. Thus, we assume that $|\psi_\varepsilon\rangle = P_\varepsilon|\psi\rangle / \|P_\varepsilon|\psi\rangle\|$ and $|\phi_\varepsilon\rangle = P_\varepsilon|\phi\rangle / \|P_\varepsilon|\phi\rangle\|$ both exist. The technical lemmas that we prove are the mean and variance of our estimators, whose proofs we defer to the full version [4].

► **Lemma 19.** *The expected value of the estimator w given in Figure 2 is*

$$\mathbb{E}[w] = |\langle\phi|M_\varepsilon|\psi\rangle|^2 + \frac{\text{Tr}[M_\varepsilon^2\phi_\varepsilon]}{k} \text{Tr}[P_\varepsilon\psi] + \frac{\text{Tr}[M_\varepsilon^2\psi_\varepsilon]}{k} \text{Tr}[P_\varepsilon\phi]. \quad (33)$$

► **Lemma 20.** *The variance of the estimator of Figure 2 is upper bounded by*

$$\text{Var}(w) = \mathcal{O}\left(\frac{1}{k} + \frac{\|M_\varepsilon\|_2^2}{k^2} + \frac{\|M_\varepsilon\|^4}{k^4}\right). \quad (34)$$

Our estimator is biased, but the bias is at most $2/k$. Taking

$$k = \Omega\left(\max\left\{\frac{1}{\varepsilon^2}, \frac{\|M_\varepsilon\|_2^2}{\varepsilon}\right\}\right),$$

ensures that both the variance and bias are small enough that Alice and Bob's estimate is within $\varepsilon/2$ of $|\langle\phi|M_\varepsilon|\psi\rangle|^2$ and thus within ε of $|\langle\phi|M|\psi\rangle|^2$ (with high probability).

4.2 Lower Bound

We now prove the claimed lower bounds in Theorem 18. The first, $\Omega(1/\varepsilon^2)$ follows from lemma 13 in [3]. We restate their argument here slightly adapted to our setup.

► **Lemma 21.** *Say there is an algorithm acting on $\rho^{\otimes k} \otimes \sigma^{\otimes k}$ that outputs an estimate of $\text{Tr}[M\rho M\sigma]$ to accuracy ε with high probability. Then, $k = \Omega(1/\varepsilon^2)$.*

Proof. Let $|0\rangle$ be such that $M|0\rangle = \pm 1$, which exists for since $\|M\| = 1$ and the unit sphere is compact. Let $|1\rangle$ be an eigenvector of M orthogonal to $|0\rangle$. Consider the states

$$|\psi_0\rangle = \sqrt{\frac{1}{2} - \varepsilon}|0\rangle + \sqrt{\frac{1}{2} + \varepsilon}|1\rangle, \quad |\psi_1\rangle = \sqrt{\frac{1}{2} + \varepsilon}|0\rangle + \sqrt{\frac{1}{2} - \varepsilon}|1\rangle. \quad (35)$$

Now, say that Alice is given k copies of $|\psi_0\rangle$ or $|\psi_1\rangle$ and Bob is given k copies of $|0\rangle$. We have that

$$|\langle\psi_0|M|0\rangle|^2 = \frac{1}{2} - \varepsilon, \quad |\langle\psi_1|M|0\rangle|^2 = \frac{1}{2} + \varepsilon. \quad (36)$$

Thus, if Alice and Bob can estimate $|\langle\psi|M|\phi\rangle|^2$ to accuracy ε , they can distinguish between these two states. However, the fidelity between these states is given by

$$F(\psi_0^{\otimes k}, \psi_1^{\otimes k}) = (1 - 4\varepsilon^2)^k \geq 1 - 4k\varepsilon^2, \quad (37)$$

which implies that $k = \Omega(1/\varepsilon^2)$. ◀

Before proving the second lower bound, we give some intuition. Say that $M \geq 0$. If $|\psi\rangle$ and $|\phi\rangle$ are drawn independently from the Haar measure, then

$$\mathbb{E}[\text{Tr}[M\psi M\phi]] = \text{Tr}[M^2]/d^2. \quad (38)$$

If they are identical, that is $\psi = \phi$ always, then instead the expected value is

$$\mathbb{E}[\text{Tr}[M\psi M\psi]] = \frac{1}{d(d-1)} (\text{Tr}[M^2] + \text{Tr}[M]^2). \quad (39)$$

The difference between these two is roughly $\text{Tr}[M^2]/d^2$. Thus, if Alice and Bob are able to estimate $|\langle\psi|M|\phi\rangle|^2$ to this order, then they can solve DIPE (that we defined as Problem 9). However, $\text{Tr}[M^2]/d^2$ may be quite small. Restricted to the support of M_ε , we instead have that $\varepsilon^2/4 \leq \text{Tr}[M_\varepsilon^2]/d_\varepsilon^2 \leq 1$. We will show that if Alice and Bob can estimate $|\langle\phi|M_\varepsilon|\psi\rangle|^2$, then they can solve the following decision problem, which is known to require $k = \Omega(\sqrt{\dim \mathcal{H}}/\varepsilon)$ samples [3] (when $\varepsilon \leq 0.01$).

11:14 Generalized Inner Product Estimation with Limited Quantum Communication

► **Definition 22** (Inner product estimation, ε decision version). *Let $\varepsilon > 0$. Alice and Bob are given k copies each of pure states $|\psi\rangle, |\phi\rangle \in \mathcal{H}$, for some finite dimensional Hilbert space \mathcal{H} . They are asked to decide, using LOCC, which of the following two scenarios they are in:*

■ **YES instance:**

$$|\psi\rangle = \sqrt{1 - \varepsilon}e^{i\theta} |0\rangle + \sqrt{\varepsilon} |\chi\rangle, \quad |\phi\rangle = \sqrt{1 - \varepsilon}e^{i\theta'} |0\rangle + \sqrt{\varepsilon} |\chi\rangle, \quad (40)$$

where $|\chi\rangle$ is drawn from the Haar measure on the orthogonal complement of $|0\rangle$ and θ and θ' are independently and uniformly drawn from $[0, 2\pi)$.

■ **NO instance:**

$$|\psi\rangle = \sqrt{1 - \varepsilon}e^{i\theta} |0\rangle + \sqrt{\varepsilon} |\chi\rangle, \quad |\phi\rangle = \sqrt{1 - \varepsilon}e^{i\theta'} |0\rangle + \sqrt{\varepsilon} |\varphi\rangle, \quad (41)$$

where $|\chi\rangle$ and $|\varphi\rangle$ are drawn independently from the Haar measure on the orthogonal complement of $|0\rangle$ and θ and θ' are independently and uniformly drawn from $[0, 2\pi)$.

We will require the two following technical lemmas in proving the lower bound. Here Lipschitz constants are taken to be with respect to the Euclidean norm $\|\cdot\|_2$.

► **Fact 23** (Levy's Lemma). *If $f : \mathbb{S}^d \rightarrow \mathbb{R}$ is λ -Lipschitz, then*

$$\Pr_u[|f(u) - \mathbb{E}[f(u)]| \geq t] \leq 2e^{-\frac{dt^2}{2\lambda^2}}, \quad (42)$$

where u is drawn from the Haar measure on \mathbb{S}^d .

► **Lemma 24.** *Let $c < 1$ be a sufficiently small constant. Say there is a protocol acting on $\rho^{\otimes k} \otimes \sigma^{\otimes k}$ via LOCC that outputs an estimate of $\text{Tr}[M\rho M\sigma]$ to accuracy $c\varepsilon$ with high probability. Then, $k = \Omega(\|M_\varepsilon\|_1/\varepsilon\sqrt{d_\varepsilon})$.*

Proof. Let $M|0\rangle = |0\rangle$ be stabilized by M . Such a vector exists by considering either M or $-M$, and we choose the appropriate sign without loss of generality. Let $W = \{|\psi\rangle \in \mathbb{C}^d \mid \langle 0|\psi\rangle = 0\} \cap \text{Im}P_\varepsilon$ and $\tilde{M}_\varepsilon = M_\varepsilon|_W$. Without loss of generality we assume that \tilde{M}_ε is a (semi)definite matrix. Indeed, let $\tilde{M}_\varepsilon = \tilde{M}_\varepsilon^+ - \tilde{M}_\varepsilon^-$ then $\|\tilde{M}_\varepsilon\|_2^2 = \|\tilde{M}_\varepsilon^+\|_2^2 + \|\tilde{M}_\varepsilon^-\|_2^2$. Of course, this means that $\|\tilde{M}_\varepsilon^+\|_2 \geq \|\tilde{M}_\varepsilon\|_2/\sqrt{2}$ or $\|\tilde{M}_\varepsilon^-\|_2 \geq \|\tilde{M}_\varepsilon\|_2/\sqrt{2}$. Thus, we restrict ourselves further and consider the subspace, again labeled by W , in the support of the operator with a larger norm.

We will now show that estimating $|\langle \phi|M|\psi\rangle|^2$ suffices to solve problem 22 with high probability where $\mathcal{H} = \mathbb{C}|0\rangle \oplus W$. Let $\tilde{d} := \dim W$. Set the precision parameter in problem 22 to be $\delta := \varepsilon\tilde{d}/200\text{Tr}[\tilde{M}_\varepsilon]$ and define $\vartheta := \theta - \theta'$. Since $\text{Tr}[\tilde{M}_\varepsilon] \geq \varepsilon\tilde{d}/2$, $\delta \leq 0.01$ as required. Let f be Alice and Bob's estimate of $|\langle \phi|M|\psi\rangle|^2$, and say that it is within $c \cdot \varepsilon$ of the true value with probability at least 0.99. Then, if f lies in the range $[(1 - \delta)^2 - 3c \cdot \varepsilon, (1 - \delta)^2 + c \cdot \varepsilon]$, they REJECT. Otherwise, they ACCEPT. We split the proof into the two cases.

YES instance: in this case

$$|\langle \phi|M|\psi\rangle|^2 = (1 - \delta)^2 + \delta^2\text{Tr}[M\chi]^2 + 2\delta(1 - \delta)\cos\vartheta\text{Tr}[M\chi]. \quad (43)$$

Thus,

$$\Pr[f \in [(1 - \delta)^2 - \frac{\delta}{8}, (1 - \delta)^2 + x]] \quad (44)$$

$$\leq \Pr[\delta^2\text{Tr}[M\chi]^2 + 2\delta(1 - \delta)\cos\vartheta\text{Tr}[M\chi] \in [-\frac{\delta}{8} - c\varepsilon, \frac{\delta}{8} + c\varepsilon]] + 0.01. \quad (45)$$

$\text{Tr}[M\chi]$ is 2-Lipschitz:

$$|\langle\psi|M|\psi\rangle - \langle\phi|M|\phi\rangle| = |(\langle\psi| - \langle\phi|)M|\psi\rangle - \langle\phi|M(|\phi\rangle - |\psi\rangle)| \quad (46)$$

$$\leq |(\langle\psi| - \langle\phi|)M|\psi\rangle| + |\langle\phi|M(|\phi\rangle - |\psi\rangle)| \quad (47)$$

$$\leq 2\|\psi\rangle - |\phi\rangle\|, \quad (48)$$

where the final inequality follows from $\|M\| = 1$ and Cauchy-Schwarz. Then, using fact 23, it holds that

$$\Pr\left[|\text{Tr}[M\chi] - \mathbb{E}_\chi[\text{Tr}[M\chi]]| \leq \frac{7}{\sqrt{\tilde{d}}}\right] > 0.99. \quad (49)$$

We now have that $\frac{7}{\sqrt{\tilde{d}}} \leq \frac{\text{Tr}[\tilde{M}_\varepsilon]}{4\tilde{d}}$, since $28/\sqrt{\tilde{d}} \leq \varepsilon$. But notice that if $\varepsilon = o(1/\sqrt{\tilde{d}})$, then the lower bound $k = \Omega(1/\varepsilon^2) = \Omega(d)$ dominates since $\|M_\varepsilon\|_2/\sqrt{\varepsilon} \leq \sqrt{\tilde{d}}/\sqrt{\varepsilon} = o(1/\varepsilon^2)$. The above bounds imply that we can take $\text{Tr}[M\chi]^2 \leq 25\text{Tr}[\tilde{M}_\varepsilon]^2/16\tilde{d}^2$. We are then interested in the probability

$$\Pr\left[2\delta(1-\delta)\frac{3\text{Tr}[\tilde{M}_\varepsilon]}{4\tilde{d}}\cos\vartheta \in \left[-\frac{\delta}{8} - c\varepsilon - \delta^2\frac{25\text{Tr}[\tilde{M}_\varepsilon]^2}{16\tilde{d}^2}, \frac{\delta}{8} + c\varepsilon - \delta^2\frac{\text{Tr}[\tilde{M}_\varepsilon]^2}{\tilde{d}^2}\right]\right]. \quad (50)$$

Now, by construction $\delta \leq 0.01$ and thus $2(1-\delta) \geq 99/50$. It follows that

$$2\delta(1-\delta)\frac{3\text{Tr}[\tilde{M}_\varepsilon]}{4\tilde{d}} \geq \frac{297}{40000}\varepsilon. \quad (51)$$

Then, the above probability is upper bounded by

$$\Pr\left[\cos\vartheta \in \left[-\frac{1}{2} - \frac{25\varepsilon}{4752}, \frac{1}{2} + \frac{\varepsilon}{297}\right]\right] \leq \Pr[\cos\vartheta \in [-0.51, 0.51]] \leq 0.34. \quad (52)$$

In total, they accept in the YES case with probability at least 0.66.

NO instance: in this case

$$|\langle\varphi|M|\chi\rangle|^2 = (1-\delta)^2 + \delta^2\text{Tr}[M\varphi M\chi] + 2\delta(1-\delta)\text{Re}(e^{i\vartheta}\langle\varphi|M|\chi\rangle). \quad (53)$$

By symmetry, $\mathbb{E}[\text{Re}(e^{i\vartheta}\langle\varphi|M|\chi\rangle)] = 0$. Now, fixing $|\varphi\rangle$, $|\langle\varphi|M|\chi\rangle|$ is 1-Lipschitz in $|\chi\rangle$. This implies that $|\langle\varphi|M|\chi\rangle| \leq 8/\sqrt{\tilde{d}}$ with probability at least 0.999. This then implies that $\text{Tr}[M\varphi M\chi] \leq 84/\tilde{d}$. Then, we have that

$$|\langle\varphi|M|\chi\rangle|^2 - (1-\delta)^2 \leq \frac{\delta}{\sqrt{\tilde{d}}}\left(\frac{84}{\sqrt{\tilde{d}}} + 16\right). \quad (54)$$

Further, $\delta/\sqrt{\tilde{d}} \leq 1/100\sqrt{\tilde{d}}$. We assume that $84/\sqrt{\tilde{d}} < 1$ (as otherwise the lower bound is constant anyways) and obtain

$$|\langle\varphi|M|\chi\rangle|^2 - (1-\delta)^2 \leq \frac{17}{100\sqrt{\tilde{d}}}. \quad (55)$$

We require that this is less than $c\varepsilon/3$. This requires that $\varepsilon = \Omega(1/\sqrt{\tilde{d}})$, but, as previously stated, we assume this as otherwise the other lower bound dominates. The total probability for this instance to be rejected is then at least 0.9.

Thus, Alice and Bob are able to solve problem 22 with high probability. Then, it must be that $k = \Omega(\sqrt{\tilde{d}}/\delta) = \Omega(\text{Tr}[\tilde{M}_\varepsilon]/\varepsilon\sqrt{\tilde{d}})$. ◀

► **Corollary 25.** *Let $c < 1$ be a sufficiently small constant. Say there is a protocol acting on $\rho^{\otimes k} \otimes \sigma^{\otimes k}$ via LOCC that outputs an estimate of $\text{Tr}[M\rho M\sigma]$ to accuracy $c\varepsilon$ with high probability. Then, $k = \Omega(\|M_\varepsilon\|_2/\sqrt{\varepsilon})$.*

Proof. Lemma 24 yields a lower bound of $k = \Omega(\text{Tr}[\tilde{M}_\varepsilon]/\varepsilon\sqrt{\tilde{d}})$. Using $\text{Tr}[\tilde{M}_\varepsilon] \geq \|\tilde{M}_\varepsilon\|_2^2$ and $\|\tilde{M}_\varepsilon\|_2 \geq \sqrt{\varepsilon\tilde{d}/2}$, we arrive at our claimed lower bound of $k = \Omega(\|M_\varepsilon\|_2/\sqrt{\varepsilon})$. ◀

References

- 1 Pablo Andres-Martinez, Tim Forrer, Daniel Mills, Jun-Yi Wu, Luciana Henaut, Kentaro Yamamoto, Mio Muraio, and Ross Duncan. Distributing circuits over heterogeneous, modular quantum computing network architectures. *Quantum Science and Technology*, 9(4):045021, 2024.
- 2 Anurag Anshu and Srinivasan Arunachalam. A survey on the complexity of learning quantum states. *Nature Reviews Physics*, 6(1):59–69, 2024.
- 3 Anurag Anshu, Zeph Landau, and Yunchao Liu. Distributed quantum inner product estimation. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 44–51, 2022. doi:10.1145/3519935.3519974.
- 4 Srinivasan Arunachalam and Louis Schatzki. Distributed inner product estimation with limited quantum communication. *arXiv preprint*, 2024. doi:10.48550/arXiv.2410.12684.
- 5 Costin Bădescu, Ryan O’Donnell, and John Wright. Quantum state certification. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 503–514, 2019. doi:10.1145/3313276.3316344.
- 6 Adriano Barenco, Andre Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello. Stabilization of quantum computations by symmetrization. *SIAM Journal on Computing*, 26(5):1541–1557, 1997. doi:10.1137/S0097539796302452.
- 7 Charles H Bennett, David P DiVincenzo, Christopher A Fuchs, Tal Mor, Eric Rains, Peter W Shor, John A Smolin, and William K Wootters. Quantum nonlocality without entanglement. *Physical Review A*, 59(2):1070, 1999.
- 8 Michael Brannan. Alice and bob meet banach: The interface of asymptotic geometric analysis and quantum information theory, 2021.
- 9 Harry Buhrman, Richard Cleve, John Watrous, and Ronald De Wolf. Quantum fingerprinting. *Physical review letters*, 87(16):167902, 2001.
- 10 Kean Chen, Qisheng Wang, Peixun Long, and Mingsheng Ying. Unitarity estimation for quantum channels. *IEEE Transactions on Information Theory*, 69(8):5116–5134, 2023. doi:10.1109/TIT.2023.3263645.
- 11 Kean Chen, Qisheng Wang, and Zhicheng Zhang. Local test for unitarily invariant properties of bipartite quantum states. *arXiv preprint*, 2024. doi:10.48550/arXiv.2404.04599.
- 12 Sitan Chen, Jerry Li, Brice Huang, and Allen Liu. Tight bounds for quantum state certification with incoherent measurements. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1205–1213. IEEE, 2022. doi:10.1109/FOCS54457.2022.00118.
- 13 Eric Chitambar, Debbie Leung, Laura Mančinska, Maris Ozols, and Andreas Winter. Everything you always wanted to know about locc (but were afraid to ask). *Communications in Mathematical Physics*, 328:303–326, 2014.
- 14 Scott M Cohen. Understanding entanglement as resource: Locally distinguishing unextendible product bases. *Physical Review A – Atomic, Molecular, and Optical Physics*, 77(1):012304, 2008.
- 15 Jacob P Covey, Harald Weinfurter, and Hannes Bernien. Quantum networks with neutral atom processing nodes. *npj Quantum Information*, 9(1):90, 2023.

- 16 Andreas Elben, Benoît Vermersch, Rick van Bijnen, Christian Kokail, Tiff Brydges, Christine Maier, Manoj K Joshi, Rainer Blatt, Christian F Roos, and Peter Zoller. Cross-platform verification of intermediate scale quantum devices. *Physical review letters*, 124(1):010504, 2020.
- 17 Heng Fan. Distinguishability and indistinguishability by local operations and classical communication. *Physical Review Letters*, 92(17):177905, 2004.
- 18 Ranjani G Sundaram, Himanshu Gupta, and CR Ramakrishnan. Efficient distribution of quantum circuits. In *35th International Symposium on Distributed Computing (DISC 2021)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.
- 19 Giovanni Guccione, Tom Darras, Hanna Le Jeannic, Varun B Verma, Sae Woo Nam, Adrien Cavallès, and Julien Laurat. Connecting heterogeneous quantum networks by hybrid entanglement swapping. *Science advances*, 6(22):eaba4508, 2020.
- 20 Aram W Harrow. The church of the symmetric subspace. *arXiv preprint*, 2013. [arXiv:1308.6595](https://arxiv.org/abs/1308.6595).
- 21 Marcel Hinsche, Marios Ioannou, Sofiene Jerbi, Lorenzo Leone, Jens Eisert, and Jose Carrasco. Efficient distributed inner product estimation via pauli sampling. *arXiv preprint*, 2024. [arXiv:2405.06544](https://arxiv.org/abs/2405.06544).
- 22 Christopher Monroe, Robert Raussendorf, Alex Ruthven, Kenneth R Brown, Peter Maunz, L-M Duan, and Jungsang Kim. Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Physical Review A*, 89(2):022317, 2014.
- 23 Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010.
- 24 A Pirker, J Wallnöfer, and W Dür. Modular architectures for quantum networks. *New Journal of Physics*, 20(5):053054, 2018.
- 25 Pranab Sen. A quantum johnson-lindenstrauss lemma via unitary t-designs. *arXiv preprint*, 2018. [arXiv:1807.08779](https://arxiv.org/abs/1807.08779).
- 26 Guifré Vidal and Rolf Tarrach. Robustness of entanglement. *Physical Review A*, 59(1):141, 1999.
- 27 Stephanie Wehner, David Elkouss, and Ronald Hanson. Quantum internet: A vision for the road ahead. *Science*, 362(6412):eaam9288, 2018.
- 28 Jun-Yi Wu, Kosuke Matsui, Tim Forrer, Akihito Soeda, Pablo Andrés-Martínez, Daniel Mills, Luciana Henaut, and Mio Muraō. Entanglement-efficient bipartite-distributed quantum computing. *Quantum*, 7:1196, 2023. [doi:10.22331/q-2023-12-05-1196](https://doi.org/10.22331/q-2023-12-05-1196).
- 29 Zhi-Chao Zhang, Fei Gao, Tian-Qing Cao, Su-Juan Qin, and Qiao-Yan Wen. Entanglement as a resource to distinguish orthogonal product states. *Scientific reports*, 6(1):30493, 2016.