

Distributed Quantum Property Testing with Communication Constraints

Mina Doosti, Ryan Sweke, Chirag Wadhwa

1 Introduction

Recent years have seen a massive surge in the study of algorithms for testing and learning quantum states [MW16; AA23]. This line of work has resulted in tight characterizations of the resource requirements for a wide variety of both testing and learning problems, under many different assumptions on the algorithm, such as its ability to perform adaptive measurements and its ability to perform coherent measurements on multiple copies of the unknown quantum state. However, the vast majority of previous work in this area has been in the *centralized* setting, where the algorithm itself has access, in some way, to copies of the unknown quantum state.

In this work, we study the problem of learning and testing quantum states in the *distributed* setting, where the copies of the unknown quantum state(s) are distributed among multiple distributed nodes, all of which communicate with a central node running a learning/testing algorithm. This setting has a variety of motivations. Firstly, as quantum communication networks develop, this setting is natural and allows us to characterize the potential and limitations of inference over such networks. Additionally, this problem is the natural quantum analogue of classical statistical inference in the distributed setting, introduced in Ref. [ACT19], whose theoretical foundations have proven essential to the development of large-scale federated learning protocols, in which data is distributed among multiple parties.

In this distributed setting, we study quantum inference problems, where one wishes to learn or test some property of an unknown quantum state ρ from its copies. As illustrated in Figure 1, this setting involves m distributed nodes $\{N_i \mid i \in [m]\}$, each holding a single copy of ρ . Each such node communicates with a central node N_c , which should output a candidate solution to the problem. We formalize the resources available to each node in Definition 1.1 below, and illustrate specific instantiations in Figure 1.

Definition 1.1: ((n_c, n_q, R, E) -model for distributed quantum inference) Given m distributed nodes $N_1, \dots, N_m$ and a central node N_c , we define the (n_c, n_q, R, E) model as the model in which no communication is allowed between distributed nodes, only one-way communication is allowed from the distributed nodes to the central node, and:

1. The distributed nodes each possess a single copy of ρ ; the central node N_c has no such copies.
2. At most n_c classical bits and n_q qubits can be sent from any distributed node N_i to the central node N_c .
3. All nodes have access either to private sources of randomness or a shared public source of randomness, indicated by $R \in \{\text{public}, \text{private}\}$.
4. Each pair of neighboring nodes shares E Bell pairs. (assume the distributed nodes sit on the vertices of a graph).

We start by making the following observations. Firstly, when *unlimited quantum communication* is allowed between the distributed and central nodes, one immediately recovers the unconstrained centralized setting. Indeed, in this case, each distributed node can simply send their quantum state to the central node, which can then run any testing algorithm – even one requiring adaptive or multi-copy coherent measurements. Secondly, in the case where the communication channels between distributed nodes and the central node are purely classical, but *unlimited classical communication* is involved, then one recovers the centralized setting in which only *non-adaptive single-copy* measurements are possible. More specifically, in this case, each distributed node can perform the predetermined single-copy measurements required by the centralized testing algorithm, then send the classical measurement results to the central node for classical post-processing. Taken together, we have the following observation:




Figure 1: An illustration of the (n_c, n_q, R, E) model for distributed quantum inference, as per Definition 1.1. Each distributed node $\{N_i\}_{i \in [m]}$ holds a single copy of ρ and communicates with the central node N_c via a communication channel limited to n_c bits and n_q qubits. The central node should output a valid solution to the inference problem. We use $R \in \{\text{public}, \text{private}\}$ to distinguish between the public and private coin setting (illustrated in panels (a) and (b) respectively), and E to denote the number of Bell pairs shared between each neighbouring distributed node.

The problem of distributed quantum inference is primarily interesting under communication constraints!

In light of this, our attention here is focused precisely on inference in this setting of limited communication. Moreover, for concreteness, we focus primarily on the fundamental testing problem of *quantum state certification*, the quantum analogue of *distribution identity testing* [Can20].

Definition 1.2: (Quantum State Certification) Given $\epsilon, \delta \in (0, 1]$, a complete classical description of a quantum state σ , and multiple copies of an unknown quantum state ρ , an algorithm is said to succeed at (ϵ, δ) -certification of σ if it behaves as follows:

1. If $\rho = \sigma$, output “Accept” with probability at least $1 - \delta$.
2. If $\|\rho - \sigma\|_1 \geq \epsilon$, output “Reject” with probability at least $1 - \delta$.

When δ is unspecified, we mean the algorithm succeeds with probability at least $\frac{3}{4}$.

The copy complexity of state certification is well-understood in the centralized setting with entangled measurements [OW15; BOW19; OW26], as well as with single-copy measurements [BCL20; CLO22; CHL+22]. Thus, we are concerned with the following concrete question:

What is the complexity of distributed state certification with communication constraints?

For certification of n -qubit states, we are thus interested in regimes with $n_c, n_q < n$. We note that when $n_q = 0$ but $n_c < n$, the (n_c, n_q, R, E) model is equivalent to the centralized setting in which the algorithm can only make non-adaptive single-copy measurements with a further restriction: each POVM can have at most 2^{n_c} outcomes. The complexity of state certification under this restriction has recently been characterized by Liu and Acharya [LA24]; see Table 1 for a translation of their results to our setting. In this work, we will focus primarily on the complementary regime of $n_c = 0$ and $0 < n_q < n$, i.e., the setting where only limited quantum communication is possible.

2 Our results

	Distribution identity testing	Quantum state certification	
	$2^{n_c} \leq d, n_q = 0, E = 0$	$2^{n_c} \leq d, n_q = 0, E = 0$	$n_c = 0, 2^{n_q} \leq d, E = 0$
Public-coin	$\Theta\left(\frac{d}{\sqrt{2^{n_c}} \epsilon^2}\right)$ [ACT19; ACT20]	$\Theta\left(\frac{d^2}{\sqrt{2^{n_c}} \epsilon^2}\right)$ [LA24]	$\Theta\left(\frac{d^2}{2^{n_q} \epsilon^2}\right)^\dagger$ Theorems 2.1 and 2.2
Private-coin	$\Theta\left(\frac{d^{3/2}}{2^{n_c} \epsilon^2}\right)$ [ACT19; ACT20]	$\Theta\left(\frac{d^3}{2^{n_c} \epsilon^2}\right)$ [LA24]	$\Omega\left(\frac{d^3}{4^{n_q} \epsilon^2}\right)^\dagger$ Theorem 2.3

Table 1: Summary of our results and prior work in the distributed setting. $\dagger$ indicates lower bounds holding only against mixedness-preserving strategies.

We summarize our results and those of relevant prior work in Table 1. In the setting where the distributed nodes can only send *quantum* messages, with public coins, we obtain the following upper bound for distributed certification:

Theorem 2.1: *Let $1 \leq n_q \leq \lfloor \log_2(d) \rfloor$, and $\epsilon, \delta \in (0, 1]$. In the distributed $(0, n_q, \text{public}, 0)$ -setting, the complexity of (ϵ, δ) -certification of d -dimensional states is at most $\mathcal{O}\left(\frac{d^2}{2^{n_q} \epsilon^2}\right)$.*

Note that when $2^{n_q} \geq \Omega(d)$, we achieve an $\mathcal{O}(d/\epsilon^2)$ upper bound, which is tight even in the centralized setting. Further, comparing Theorem 2.1 to the $\mathcal{O}\left(\frac{d^2}{2^{n_c/2} \epsilon^2}\right)$ upper bound for state certification with public coins and limited *classical* communication, we see that quantum communication provides a quadratic advantage over classical communication for this problem. Moreover, we show that the above upper bound is tight under the assumption that the action of each distributed node is *mixedness preserving*, i.e., when node N_i receives the maximally mixed state $\rho = \frac{\mathbb{1}_d}{d}$, its message to N_c is the maximally mixed state of dimension $d_q \triangleq 2^{n_q}$, i.e., $\frac{\mathbb{1}_{d_q}}{d_q}$. In particular, we show:

Theorem 2.2: *Let $d \geq 6, 1 \leq n_q \leq \lfloor \log_2(d) \rfloor$ and $0 < \epsilon < c$, where $c < 1$ is an absolute constant. In the distributed $(0, n_q, \text{public}, 0)$ -setting, assuming the action of each distributed node is mixedness preserving, the complexity of ϵ -certification of d -dimensional states is at least $\Omega\left(\frac{d^2}{2^{n_q} \epsilon^2}\right)$.*

Note that this mixedness-preserving assumption is indeed satisfied by our algorithm of Theorem 2.1, implying that, in the public-coin setting, $\Theta\left(\frac{d^2}{2^{n_q} \epsilon^2}\right)$ is the optimal complexity for state certification with such strategies. Further, in the private-coin setting, we show a stronger lower bound, implying a strict separation between these settings.

Theorem 2.3: *Let $d \geq 6, 1 \leq n_q \leq \lfloor \log_2(d) \rfloor$ and $0 < \epsilon < c$, where $c < 1$ is an absolute constant. In the distributed $(0, n_q, \text{private}, 0)$ -setting, assuming the action of each distributed node is mixedness preserving, the complexity of ϵ -certification of d -dimensional states is at least $\Omega\left(\frac{d^3}{4^{n_q} \epsilon^2}\right)$.*

Given our lower bounds, one observes a strong dimension-dependent gap between the public- and private-coin settings (for constant n_q , this is of order d), highlighting that shared randomness is a powerful *resource* in distributed learning and testing, as also observed in the settings of classical communication (see Table 1). In the resource-driven mindset of our framework, this is particularly interesting because it shows that the *classical* resource of *shared randomness* can be leveraged to reduce the use of another, often more expensive, resource: quantum communication. Notably, although the shared random seed is independent of the hypothesis state, it still leads to a substantial reduction in the quantum communication required for distributed certification. Finally, while Theorems 2.2 and 2.3 only hold against mixedness-preserving strategies, we conjecture that these bounds are tight for *all* strategies.

We refer the interested reader to the full version of this paper [DSW26] for an overview of our techniques and complete proofs of our results, along with auxiliary results on testing with shared entanglement.

References

- [AA23] Anurag Anshu and Srinivasan Arunachalam. *A survey on the complexity of learning quantum states*. 2023. arXiv: [2305.20069](https://arxiv.org/abs/2305.20069) [quant-ph]. URL: <https://arxiv.org/abs/2305.20069> (page 1).
- [ACT19] Jayadev Acharya, Clément L Canonne, and Himanshu Tyagi. “Inference under Information Constraints: Lower Bounds from Chi-Square Contraction”. In: *Proceedings of the Thirty-Second Conference on Learning Theory*. Ed. by Alina Beygelzimer and Daniel Hsu. Vol. 99. Proceedings of Machine Learning Research. PMLR, June 2019, pp. 3–17. URL: <https://proceedings.mlr.press/v99/acharya19a.html> (pages 1, 3).
- [ACT20] Jayadev Acharya, Clément L. Canonne, and Himanshu Tyagi. “Inference Under Information Constraints II: Communication Constraints and Shared Randomness”. In: *IEEE Transactions on Information Theory* 66.12 (2020). Preliminary version in ICML 2019. Preprint available at arXiv:1905.08302, pp. 7856–7877. DOI: [10.1109/TIT.2020.3028439](https://doi.org/10.1109/TIT.2020.3028439) (page 3).
- [BCL20] Sebastien Bubeck, Sitan Chen, and Jerry Li. “Entanglement is necessary for optimal quantum property testing”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 692–703 (page 2).
- [BOW19] Costin Bădescu, Ryan O’Donnell, and John Wright. “Quantum state certification”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 2019, pp. 503–514 (page 2).
- [Can20] Clément L. Canonne. *A Survey on Distribution Testing: Your Data is Big. But is it Blue?* Graduate Surveys 9. Theory of Computing Library, 2020, pp. 1–100. DOI: [10.4086/toc.gs.2020.009](https://doi.org/10.4086/toc.gs.2020.009). URL: <http://www.theoryofcomputing.org/library.html> (page 2).
- [CHL+22] Sitan Chen, Brice Huang, Jerry Li, Allen Liu, and Mark Sellke. “Tight bounds for state tomography with incoherent measurements”. In: *arXiv preprint arXiv:2206.05265* 4 (2022) (page 2).
- [CLO22] Sitan Chen, Jerry Li, and Ryan O’Donnell. “Toward instance-optimal state certification with incoherent measurements”. In: *Conference on Learning Theory*. PMLR. 2022, pp. 2541–2596 (page 2).
- [DSW26] Mina Doosti, Ryan Sweke, and Chirag Wadhwa. “Distributed quantum property testing with communication constraints”. In: *arXiv preprint arXiv:2604.05962* (2026) (page 3).
- [LA24] Yuhan Liu and Jayadev Acharya. “Quantum state testing with restricted measurements”. In: *arXiv preprint arXiv:2408.17439* (2024) (pages 2, 3).
- [MW16] Ashley Montanaro and Ronald de Wolf. *A Survey of Quantum Property Testing*. Graduate Surveys 7. Theory of Computing Library, 2016, pp. 1–81. DOI: [10.4086/toc.gs.2016.007](https://doi.org/10.4086/toc.gs.2016.007). URL: <http://www.theoryofcomputing.org/library.html> (page 1).
- [OW15] Ryan O’Donnell and John Wright. “Quantum spectrum testing”. In: *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. 2015, pp. 529–538 (page 2).
- [OW26] Ryan O’Donnell and Chirag Wadhwa. “Instance-optimal quantum state certification with entangled measurements”. In: *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*. 2026, pp. 398–409 (page 2).