

Can Quantum Communication Speed Up Distributed Computation?

Michael Elkin^{*} Hartmut Klauck[†] Danupon Nanongkai[‡] Gopal Pandurangan[§]

Abstract

The focus of this paper is on *quantum distributed* computation, where we investigate whether quantum communication can help in *speeding up* distributed network algorithms. Our main result is that for certain fundamental network problems such as minimum spanning tree, minimum cut, and shortest paths, quantum communication *does not* help in substantially speeding up distributed algorithms for these problems compared to the classical setting.

In order to obtain this result, we extend the technique of Das Sarma et al. [SICOMP 2012] to obtain a uniform approach to prove non-trivial lower bounds for quantum distributed algorithms for several graph optimization (both exact and approximate versions) as well as verification problems, some of which are new even in the classical setting, e.g. tight randomized lower bounds for Hamiltonian cycle and spanning tree verification, answering an open problem of Das Sarma et al., and a lower bound in terms of the weight aspect ratio, matching the upper bounds of Elkin [STOC 2004]. Our approach introduces the *Server model* and *Quantum Simulation Theorem* which together provide a connection between distributed algorithms and communication complexity. The Server model is the standard two-party communication complexity model augmented with additional power; yet, most of the hardness in the two-party model is carried over to this new model. The Quantum Simulation Theorem carries this hardness further to quantum distributed computing. Our techniques, except the proof of the hardness in the Server model, require very little knowledge in quantum computing, and this can help overcoming a usual impediment in proving bounds on quantum distributed algorithms. In particular, if one can prove a lower bound for distributed algorithms for a certain problem using the technique of Das Sarma et al., it is likely that such lower bound can be extended to the quantum setting using tools provided in this paper and without the need of knowledge in quantum computing.

^{*}Department of Computer Science, Ben-Gurion University, Beer-Sheva, 84105, Israel. E-mail: elkinm@cs.bgu.ac.il.

[†]Division of Mathematical Sciences, Nanyang Technological University, Singapore 637371 & Centre for Quantum Technologies, National University of Singapore, Singapore 117543. E-mail: hklauck@gmail.com. Research at the Centre for Quantum Technologies is funded by the Singapore Ministry of Education and the National Research Foundation.

[‡]Faculty of Computer Science, University of Vienna, Austria. E-mail: danupon@gmail.com. Work partially done while at Division of Mathematical Sciences, Nanyang Technological University, Singapore 637371.

[§]Division of Mathematical Sciences, Nanyang Technological University, Singapore 637371 & Department of Computer Science, Brown University, Providence, RI 02912, USA. E-mail: gopalpandurangan@gmail.com. Supported in part by the following research grants: Nanyang Technological University grant M58110000, Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 2 grant MOE2010-T2-2-082, Singapore MOE AcRF Tier 1 grant MOE2012-T1-001-094, and a grant from the US-Israel Binational Science Foundation (BSF).

Contents

I	Overview	1
1	Introduction	1
2	The Setting	3
2.1	Quantum Distributed Computing Model	3
2.2	Distributed Graph Problems	4
3	Our Contributions	5
3.1	Lower Bound Techniques for Quantum Distributed Computing	6
3.2	Quantum Distributed Lower Bounds	9
3.3	Additional Results: Lower Bounds on Communication Complexity	12
4	Other Related Work	12
5	Conclusion and Open Problems	12
II	Proofs	14
6	Server Model Lower Bounds via Nonlocal Games (Lemma 3.2)	14
7	Server-model Lower Bounds for ham_n (Theorem 3.4)	16
8	The Quantum Simulation Theorem (Theorem 3.5)	18
9	Proof of main theorems (Theorem 3.6 & 3.8)	23
9.1	Proof of Theorem 3.6	23
9.2	Proof of Theorem 3.8	23
III	Appendix	25
A	Detailed Definitions	25
A.1	Quantum Distributed Network Models	25
A.2	Distributed Graph Verification Problems	28
A.3	Distributed Graph Optimization Problems	29
B	Detail of Section 6	31
B.1	Two-player XOR Games	31
B.2	From Nonlocal Games to Server-Model Lower Bounds	32
B.3	Lower Bound for $\text{IPmod}3_n$	33
C	Detail of Section 7	39

D	Detail of Section 8	40
D.1	Description of the network N	40
D.2	Simulation	41

Part I

Overview

1 Introduction

The power and limitations of distributed (network) computation have been studied extensively over the last three decades or so. In a distributed network, each individual node can communicate only with its neighboring nodes. Some distributed problems can be solved entirely via local communication, e.g., maximal independent set, maximal matching, coloring, dominating set, vertex cover, or approximations thereof. These are considered “local” problems, as they can be shown to be solved using *small* (i.e., *polylogarithmic*) communication (e.g., see [Lub86a, Pel00, Suoar]). For example, a maximal independent set can be computed in $O(\log n)$ time [Lub86b]. However, many important problems are “global” problems (which are the focus of this paper) from the distributed computation point of view. For example, to count the total number of nodes, to elect a leader, to compute a spanning tree (ST) or a minimum spanning tree (MST) or a shortest path tree (SPT), information necessarily must travel to the farthest nodes in a system. If exchanging a message over a single edge costs one time unit, one needs $\Omega(D)$ time units to compute the result, where D is the network diameter [Pel00]. If message size was unbounded, one can simply collect all the information in $O(D)$ time, and then compute the result. However, in many applications, there is *bandwidth restriction* on the size of the message (or the number of bits) that can be exchanged over a communication link in one time unit. This motivates studying global problems in the CONGEST model [Pel00], where each node can exchange at most B bits (typically B is small, say $O(\log n)$) among its neighbors in one time step. This is one of the central models in the study of distributed computation. The design of efficient algorithms for the CONGEST model, as well as establishing lower bounds on the time complexity of various fundamental distributed computing problems, has been the subject of an active area of research called (locality-sensitive) *distributed computing* for the last three decades (e.g., [Pel00, Elk04, DGP07, KKM⁺08, Suoar, DHK⁺12]). In particular, it is now established that $\tilde{\Omega}(D + \sqrt{n})$ ¹ is a fundamental lower bound on the running time of many important graph optimization (both exact and approximate versions) and verification problems such as MST, ST, shortest paths, minimum cut, ST verification etc [DHK⁺12].

The main focus of this paper is studying the power of distributed network computation in the *quantum setting*. More precisely, we consider the CONGEST model in the quantum setting, where nodes can use quantum processing, communicate over quantum links using quantum bits, and use exclusively quantum phenomena such as *entanglement* (e.g., see [DP08, BT08, GKM09]). A fundamental question that we would like to investigate is whether quantumness can help in speeding up distributed computation for graph optimization problems; in particular, whether the above mentioned lower bound of $\tilde{\Omega}(D + \sqrt{n})$ (that applies to many important problems in the classical setting) also applies to the quantum setting.

Lower bounds for local problems (where the running time is $O(\text{poly log } n)$) in the quantum setting usually follow directly from the same arguments as in the classical setting. This is because these lower bounds are proved using the “limited sight” argument: The nodes do not have time to get the information of the entire network. Since entanglement *cannot be used to replace communication* (by, e.g., Holevo’s theorem [Hol73] (also see [NC04, Nay99])), the same argument holds in

¹ $\tilde{\Omega}$ and \tilde{O} notations hide polylogarithmic factors.

the quantum setting with prior entanglement. This argument is captured by the notion of *physical locality* defined by Gavioille et al. [GKM09], where it is shown that for many *local* problems, quantumness does not give any significant speedup in time compared to the classical setting.

The above limited sight argument, however, does not seem to be extendible to *global* problems where the running time is usually $\Omega(D)$, since nodes have enough time to see the whole network in this case. In this setting, the argument developed in [DHK⁺12] (which follows the line of work in [PR00, LPSP06, Elk06, KKP11]) can be used to show tight lower bounds for many problems in the classical setting. However, this argument does not always hold in the quantum setting because it essentially relies on network “congestion”: Nodes cannot communicate fast enough (due to limited bandwidth) to get important information to solve the problem. However, we know that the quantum communication and entanglement can potentially *decrease the amount of communication* and thus there might be some problems that can be solved faster. One example that illustrates this point is the following *distributed verification of disjointness function* defined in [DHK⁺12, Section 2.3].

Example 1.1. Suppose we give b -bit string x and y to node u and v in the network, respectively, where $b = \sqrt{n}$. We want to check whether the inner product $\langle x, y \rangle$ is zero or not. This is called the *Set Disjointness* problem (Disj). It is easy to show that it is impossible to solve this problem in less than $D/2$ rounds since there will be no node having the information from both u and v if u and v are of distance D apart. (This is the very basic idea of the limited sight argument.) This argument holds for both classical and quantum setting and thus we have a lower bound of $\Omega(D)$ on both settings. [DHK⁺12, Lemma 4.1] shows that this lower bound can be significantly improved to $\tilde{\Omega}(b) = \tilde{\Omega}(\sqrt{n})$ in the classical setting, even when the network has diameter $O(\log n)$. This follows from the communication complexity of $\Omega(b)$ of Disj [BFS86, KS92, BYJKS04, Raz92] and the *Simulation Theorem* of [DHK⁺12]. This lower bound, however, does not hold in the quantum setting since we can simulate the known $O(\sqrt{b})$ -communication quantum protocol of [AA05] in $O(\sqrt{b}D) = O(n^{1/4}D)$ rounds. \square

Thus we have an example of a global problem that quantum communication gives an advantage over classical communication. This example also shows that the previous techniques and results from [DHK⁺12] does not apply to the quantum setting since [DHK⁺12] heavily relies on the hardness of the above distributed disjointness verification problem. A fundamental question is: “*Does this phenomenon occur for natural global distributed network problems?*”

Our paper answers the above question where we show that this phenomenon does not occur for many global graph problems. Our main result is that for fundamental global problems such as minimum spanning tree, minimum cut, and shortest paths, quantum communication *does not* help significantly in speeding up distributed algorithms for these problems compared to the classical setting. More precisely, we show that $\tilde{\Omega}(D + \sqrt{n})$ is a lower bound for these problems in the quantum setting as well. An $\tilde{O}(D + \sqrt{n})$ time algorithm for MST problem in the classical setting is well-known [KP98]. Recently, it has been shown that minimum cut also admits a distributed $(1 + \epsilon)$ -approximation algorithm in the same time in the classical setting [GK13, Su14, Nan14a, NS14]. Also, recently it has been shown that shortest paths admits an $\tilde{O}(D + \sqrt{n}D^{1/4})$ -time $(1 + \epsilon)$ -approximation and $\tilde{O}(\sqrt{n} + D)$ -time $O(\log n)$ -approximation distributed classical algorithms [LPS13, Nan14b]. Thus, our quantum lower bound shows that quantum communication does not speed up distributed algorithms for MST and minimum cut, while for shortest paths the speed up, if any, is bounded by $O(D^{1/4})$ (which is small for small diameter graphs).

In order to obtain our quantum lower bound results, we develop a uniform approach to prove non-trivial lower bounds for quantum distributed algorithms. This approach leads us to several

non-trivial quantum distributed lower bounds (which are the first-known quantum bounds for problems such as minimum spanning tree, shortest paths etc.), some of which are new even in the classical setting. Our approach introduces the *Server model* and *Quantum Simulation Theorem* which together provide a connection between distributed algorithms and communication complexity. The Server model is simply the standard two-party communication complexity model augmented with a powerful *Server* who can communicate for free but receives no input (cf. Def. 3.1). It is more powerful than the two-party model, yet captures most of the hardness obtained by the current quantum communication complexity techniques. The Quantum Simulation Theorem (cf. Theorem 3.5) is an extension of the Simulation Theorem of Das Sarma et al. [DHK⁺12] from the classical setting to the quantum one. It carries this hardness from the Server model further to quantum distributed computing. Most of our techniques require very little knowledge in quantum computing, and this can help overcoming a usual impediment in proving bounds on quantum distributed algorithms. In particular, if one can prove a lower bound for distributed algorithms in the classical setting using the technique of Das Sarma et al., then it is possible that one can also prove the same lower bound in the quantum setting in essentially the same way – the only change needed is that the proof has to start from problems that are hard on the server model that we provide in this paper.

2 The Setting

2.1 Quantum Distributed Computing Model

We study problems in a natural quantum version of the CONGEST(B) model [Pel00] (or, in short, the B -model), where each node can exchange at most B bits (typically B is small, say $O(\log n)$) among its neighbors in one time step. The main focus of the current work is to understand the time complexity of fundamental graph problems in the B -model in the *quantum setting*. We now explain the model. We refer the readers to Appendix A.1 for a more rigorous and formal definition of our model.

Consider a synchronous network of processors modeled by an undirected n -node graph, where nodes model the processors and edges model the links between the processors. The processors (henceforth, nodes) are assumed to have unique IDs. Each node has limited topological knowledge; in particular, it only knows the IDs of its neighbors and knows no other topological information (e.g., whether its neighbors are linked by an edge or not). The node may also accept some additional inputs as specified by the problem at hand.

The communication is synchronous, and occurs in discrete pulses, called *rounds*. All the nodes wake up simultaneously at the beginning of each round. In each round each node u is allowed to send an arbitrary message of B bits through each edge $e = (u, v)$ incident to u , and the message will arrive at v at the end of the current round. Nodes then perform an internal computation, which finishes instantly since nodes have infinite computation power. There are several measures to analyze the performance of distributed algorithms, a fundamental one being the *running time*, defined as the worst-case number of rounds of distributed communication.

In the quantum setting, a distributed network could be augmented with two additional resources: *quantum communication* and *shared entanglement* (see e.g., [DP08]). Quantum communication allows nodes to communicate with each other using *quantum bits (qubits)*; i.e., in each round at most B qubits can be sent through each edge in each direction. Shared entanglement allows nodes

to possess qubits that are entangled with qubits of other nodes². Quantum distributed networks can be categorized based on which resources are assumed (see, e.g., [GKM09]). In this paper, we are interested in the *most powerful model*, where both quantum communication and the *most general form of shared entanglement* are assumed: in a technical term, we allow nodes to share an arbitrary n -partite entangled state as long as it does not depend on the input (thus, does not reveal any input information). Throughout the paper, we simply refer to this model as quantum distributed network (or just distributed network, if the context is clear). All lower bounds we show in this paper hold in this model, and thus also imply lower bounds in weaker models.

2.2 Distributed Graph Problems

We focus on solving graph problems on distributed networks. We are interested in two types of graph problems: optimization and verification problems. In both types of problems, we are given a distributed network N modeled by a graph and some property \mathcal{P} such as “Hamiltonian cycle”, “spanning tree” or “connected component”.

In optimization problems, we are additionally given a (positive) weight function $w : E(N) \rightarrow \mathbb{R}_+$ where every node in the network knows weights of edges incident to it. Our goal is to find a subnetwork M of N of minimum weight that satisfies \mathcal{P} (e.g. minimum Hamiltonian cycle or MST) where every node knows which edges incident to it are in M in the end of computation. Algorithms can sometimes depend on the *weight aspect ratio* W defined as $W = \frac{\max_{e \in E(N)} w(e)}{\min_{e \in E(N)} w(e)}$.

In verification problems, we are additionally given a subnetwork M of N as the problem input (each node knows which edges incident to it are in M). We want to determine whether M has some property, e.g., M is a Hamiltonian cycle ($\text{Ham}(N)$), a spanning tree ($\text{ST}(N)$), or a connected component ($\text{Conn}(N)$), where every node knows the answer in the end of computation.

We use³ $Q_{\epsilon_0, \epsilon_1}^{*,N}(\text{Ham}(N))$ to refer to the quantum time complexity of Hamiltonian cycle verification problem on network N where for any 0-input M (i.e. M is not a Hamiltonian cycle), the algorithm has to output zero with probability at least $1 - \epsilon_0$ and for any 1-input M (i.e. M is a Hamiltonian cycle), the algorithm has to output one with probability at least $1 - \epsilon_1$. (We call this type of algorithm (ϵ_0, ϵ_1) -error.) When $\epsilon_0 = \epsilon_1 = \epsilon$, we simply write $Q_\epsilon^{*,N}(\text{Ham}(N))$. Define $Q_{\epsilon_0, \epsilon_1}^{*,N}(\text{ST}(N))$ and $Q_{\epsilon_0, \epsilon_1}^{*,N}(\text{Conn}(N))$ similarly.

We also study the *gap versions* of verification problems. For any integer $\delta \geq 0$, property \mathcal{P} and a subnetwork M of N , we say that M is δ -far⁴ from \mathcal{P} if we have to add at least δ edges from N and remove any number of edges in order to make M satisfy \mathcal{P} . We denote the problem of

²Roughly speaking, one can think of shared entanglement as a “quantum version” of shared randomness. For example, a well-known entangled state on two qubits is the *EPR* pair [EPR35, Bel64] which is a pair of qubits that, when measured, will either both be zero or both be one, with probability $1/2$ each. An EPR pair shared by two nodes can hence be used to, among other things, generate a shared random bit for the two nodes. Assuming entanglement implies shared randomness (even among all nodes), but also allows for other operations such as quantum teleportation [NC04], which replaces quantum communication by classical communication plus entanglement.

³We mention the reason behind our complexity notations. First, we use $*$ as in Q^* in order to emphasize that our lower bounds hold even when there is a shared entanglement, as usually done in the literature. Since we deal with different models in this paper, we put the model name after $*$. Thus, we have $Q^{*,N}$ for the case of distributed algorithm on a distributed network N , and $Q^{*,cc}$ and $Q^{*,sv}$ for the case of the standard communication complexity and the Server model (cf. Subsection 3.1), respectively.

⁴We note that the notion of δ -far should not be confused with the notion of ϵ -far usually used in property testing literature where we need to add and remove at least ϵ fraction of edges in order to achieve a desired property. The two notions are closely related. The notion that we chose makes it more convenient to reduce between problems on different models.

distinguishing between the case where the subnetwork M satisfies \mathcal{P} and is δ -far from satisfying \mathcal{P} the δ - \mathcal{P} problem (it is promised that the input is in one of these two cases). When we do not want to specify δ , we write $\text{Gap-}\mathcal{P}$.

Other graph problems that we are interested in are those in [DHK⁺12] and their gap versions. We provide definitions in Appendix A.1 for completeness.

3 Our Contributions

Our first contribution is lower bounds for various fundamental verification and optimization graph problems, some of which are new even in the classical setting and answers some previous open problems (e.g. [DHK⁺12]). We explain these lower bounds in detail in Section 3.2. The main implication of these lower bounds is that quantum communication does *not* help in substantially speeding up distributed algorithms for many of these problems compared to the classical setting. Notable examples are MST, minimum cut, s -source distance, shortest path tree, and shortest s - t paths. In Corollary 3.9, we show a lower bound of $\Omega(\sqrt{\frac{n}{B \log n}})$ for these problems which holds against any quantum distributed algorithm with any approximation guarantee. Due to the seminal paper of Kutten and Peleg [KP98], we know that MST can be computed exactly in $\tilde{O}(\sqrt{n} + D)$ time in the classical setting, and thus we cannot hope to use quantum communication to get a significant speed up for MST. Recently, Ghaffari and Kuhn [GK13] showed that minimum cut can be $(2 + \epsilon)$ -approximated in $\tilde{O}(\sqrt{n} + D)$ time in the classical setting, and Su [Su14] and Nanongkai [Nan14a] independently improved the approximation ratio to $(1 + \epsilon)$; this implies that, again, quantum communication does not help. More recently, Nanongkai [Nan14b] showed that s -source distance, shortest path tree, and shortest s - t paths, can be $(1 + o(1))$ -approximated in $\tilde{O}(\sqrt{n} D^{1/4} + D)$ time in the classical setting; thus, the speedup that quantum communication can provide for these problems, if any, is bounded by $O(D^{1/4})$. Moreover, if we allow higher approximation factor, the result of Lenzen and Patt-Shamir [LPS13] implies that we can $O(\log n)$ -approximate these problems in $\tilde{O}(\sqrt{n} + D)$ time; this upper bound together with our lower bound leaves no room for quantum algorithms to improve the time complexity. Besides the above lower bounds for optimization problems, we show the same lower bound of $\Omega(\sqrt{\frac{n}{B \log n}})$ for verification problems in Corollary 3.7. Das Sarma et al. [DHK⁺12] showed that these problems, except least-element list verification, can be solved in $\tilde{O}(\sqrt{n} + D)$ time in the classical setting; thus, once again, quantum communication does not help.

Our second contribution is the systematic way to prove lower bounds of quantum distributed algorithms. The high-level idea behind our lower bound proofs is by establishing a connection between quantum communication complexity and quantum distributed network computing. Our work is inspired by [DHK⁺12] (following a line of work in [PR00, LPSP06, Elk06, KKP11]) which shows lower bounds for many graph verification and optimization problems in the classical distributed computing model. The main technique used to show the classical lower bounds in [DHK⁺12] is the *Simulation Theorem* (Theorem 3.1 in [DHK⁺12]) which shows how one can use lower bounds in the standard two-party classical communication complexity model [KN97] to derive lower bounds in the “distributed” version of communication complexity. We provide techniques of the same flavor for proving quantum lower bounds. In particular, we develop the *Quantum Simulation Theorem*. However, due to some difficulties in handling quantum computation (especially the entanglement) we need to introduce one more concept: instead of applying the Quantum Simulation Theorem to

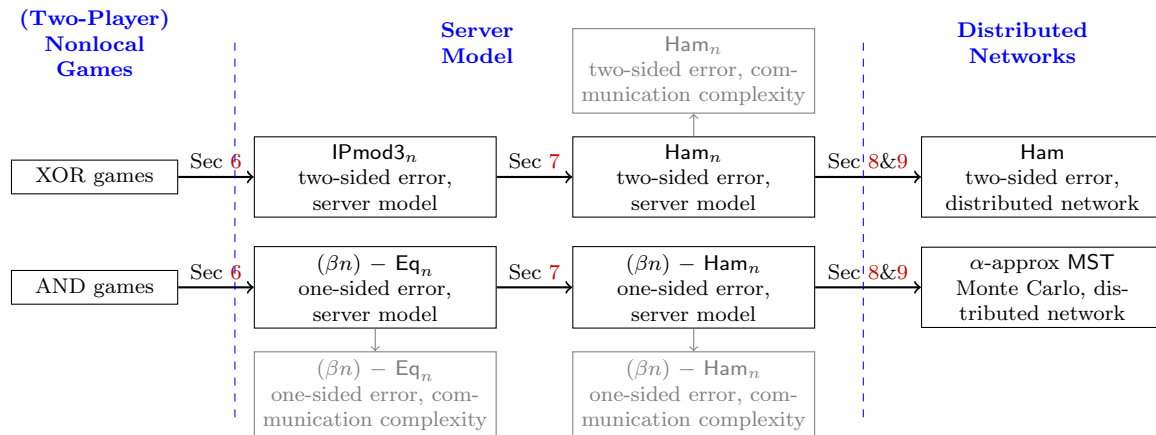


Figure 1: *Our proof structure. Lines in gray show the implications of our results in communication complexity.*

the standard two-party communication complexity model, we have to apply it to a slightly stronger model called *Server model*. We show that working with this stronger model does not make us lose much: several hard problems in two-party communication complexity remain hard in this model, so we can still prove hardness results using these problems. Quantum Simulation Theorem together with the Server model give us a tool to bring the hardness in the quantum two-party setting to the distributed setting. In Section 3.1, we give a more comprehensive overview of our techniques. Along the way, we also obtain new results in the standard communication complexity model, which we explain in Section 3.3.

3.1 Lower Bound Techniques for Quantum Distributed Computing

The high-level idea behind our lower bound proofs is establishing a connection between quantum communication complexity and quantum distributed network computing via a new communication model called the *Server model*, as shown in two middle columns of Fig. 1. This model is a generalization of the standard two-party communication complexity model in the sense that the Server model can simulate the two-party model; thus, lower bounds on this model imply lower bounds on the two-party network models. More importantly, we show that lower bounds on this model imply lower bounds on the quantum distributed model as well (cf. Section 8 & 9). This is depicted by the rightmost arrows in Fig. 1. In addition, we prove quantum lower bounds in the server model, some of which also imply *new* lower bounds in the two-party model for problems such as Hamiltonian cycle and spanning tree, even in the classical setting. This is done by showing that certain techniques based on *nonlocal games* can be extended to prove lower bounds on the Server model (cf. Section 6) as depicted by leftmost arrows in Fig. 1, and by reductions between problems in the Server models (cf. Section 7) as depicted by middle arrows in Fig. 1.

Definition 3.1 (Server Model). There are three players in the server model: Carol, David and the server. Carol and David receive the inputs x and y , respectively, and want to compute $f(x, y)$ for some function f . (Observe that the server receives no input.) Carol and David can talk to each other. Additionally, they can talk to the server. The catch here is that the server can send messages for *free*. Thus, the communication complexity in the server model counts only messages

sent by Carol and David.

We let $Q_{\epsilon_0, \epsilon_1}^{*,sv}(f)$ denote the communication complexity — in the quantum setting with entanglement — of computing function f where for any i -input (an input whose correct output is $i \in \{0, 1\}$) the algorithm must output i with probability at least $1 - \epsilon_i$. We will write $Q_{\epsilon}^{*,sv}(f)$ instead of $Q_{\epsilon, \epsilon}^{*,sv}(f)$. For the standard two-party communication complexity model [KN97], we use $Q_{\epsilon_0, \epsilon_1}^{*,cc}(f)$ to denote the communication complexity in the quantum setting with entanglement.

To the best of our knowledge, the Server model is different from other known models in communication complexity. Clearly, it is different from multi-party communication complexity since the server receives no input and can send information for *free*. Moreover, it is easy to see that the Server model, even without prior entanglement, is at least as strong as the standard quantum communication complexity model with shared entanglement, since the server can dispense any entangled state to Carol and David. Interestingly, it turns out that the Server model is *equivalent* to the standard two-party model in the classical communication setting, while it is not clear if this is the case in the quantum communication setting. This is the main reason that proving lower bounds in the quantum setting is more challenging in its classical counterpart.

To explain some issues in the quantum setting, let us sketch the proof of the fact that the two models are *equivalent* in the classical setting. Let us first consider the deterministic setting. The proof is by the following “simulation” argument. Alice will simulate Carol and the server. Bob will simulate David and the server. In each round, Alice will see all messages sent from the server to Carol and thus she can keep simulating Carol. However, she does not see the message sent from David to the server which she needs to simulate the server. So, she must get this message from Bob. Similarly, Bob will get from Alice the message sent from Carol to the server. These are the only messages we need in each round in order to be able to simulate the protocol. Observe that the total number of bits sent between Alice and Bob is exactly the number of bits sent by Carol and David to the server. Thus, the complexities of both models are exactly the same in the deterministic case. We can conclude the same thing for the public coin setting (where all parties share a random string) since Alice and Bob can use their shared coin to simulate the shared coin of Carol, David and the server.

The above argument, however, does not seem to work in the quantum setting. The main issue with a simulation along the lines of the one sketched above is that Alice and Bob cannot simulate a “copy” of the server each. For instance one could try to simulate the server’s state in a distributed way by maintaining the state that results by applying CNOT to every qubit of the server and a fresh qubit, and distribute these qubits to Alice and Bob. But then if the server sends a message to Carol, Bob would have to disentangle the corresponding qubits in his copy, which would require a message to Alice.

While we leave as an open question whether the two models are equivalent in the quantum setting, we prove that many lower bounds in the two-party model extend to the server model, via a technique called nonlocal games.

Lower Bound Techniques on the Server Model (Details in Section 6) We show that many hardness results in the two-party model (where there is no server) carry over to the Server model. This is the only part that the readers need some background in quantum computing. The main difficulty in showing this is that, the Server model, even *without* prior entanglement, is clearly at least *as strong as* the standard quantum communication complexity model (where there is no server) *with* shared entanglement, since the server can dispense any entangled state to Carol and

David. Thus, it is a challenging problem, which could be of an independent interest, whether *all* hard problems in the standard model remain hard in the server model.

While we do not fully answer the above problem, we identify a set of lower bound techniques in the standard quantum communication complexity model that can be carried over to the Server model, and use them to show that *many* problems remain hard. Specifically, we show that techniques based on the (two-player) *nonlocal* games (see, e.g., [LS09a, LZ10, KdW12]) can be extended to show lower bounds on the Server model.

Nonlocal games are games where two players, Alice and Bob, receive an input x and y from some distribution that is known to them and want to compute $f(x, y)$. Players cannot talk to each other; instead, they output one bit, say a and b , which are then combined to be an output. For example, in *XOR games* and *AND games*, these bits are combined as $a \oplus b$ and $a \wedge b$, respectively. The players' goal is to maximize the probability that the output is $f(x, y)$. We relate nonlocal games to the server model by showing that the XOR- and AND-game players can use an efficient server-model protocol to guarantee a good winning chance:

Lemma 3.2. (SERVER MODEL LOWER BOUNDS VIA NONLOCAL GAMES) *For any boolean function f and $\epsilon_0, \epsilon_1 \geq 0$, there is an (two-player nonlocal) XOR-game strategy \mathcal{A}' (respectively, AND-game strategy \mathcal{A}'') such that, for any input (x, y) , with probability $4^{-2Q_{\epsilon_0, \epsilon_1}^{*,sv}(f)}$, \mathcal{A}' (respectively, \mathcal{A}'') outputs $f(x, y)$ with probability at least $1 - \epsilon_{f(x, y)}$ (i.e. it outputs 1 with probability at least $1 - \epsilon_1$ and 0 with probability at least $1 - \epsilon_0$); otherwise (with probability $1 - 4^{-2Q_{\epsilon_0, \epsilon_1}^{*,sv}(f)}$), \mathcal{A}' outputs 0 and 1 with probability $1/2$ each (respectively, \mathcal{A}'' outputs 0 with probability 1).*

Roughly speaking, the above lemma compares two cases: in the “good” case \mathcal{A}' outputs the correct value of $f(x, y)$ with high probability (the probability controlled by ϵ_0 and ϵ_1) and in the “bad case” \mathcal{A}' simply outputs a random bit. It shows that if $Q_{\epsilon_0, \epsilon_1}^{*,sv}(f)$ is small, then the “good” case will happen with a non-negligible probability. In other words, the lemma says that if $Q_{\epsilon_0, \epsilon_1}^{*,sv}(f)$ is small, then the probability that the nonlocal game players win the game will be high.

This lemma gives us an access to several lower bound techniques via nonlocal games. For example, following the γ_2 -norm techniques in [LS09b, She11, LZ10] and the recent method of [KdW12], we show one- and two-sided error lower bounds for many problems on the server model (in particular, we can obtain lower bounds in general forms as in [Raz03, She11, LZ10]). These lower bounds match the two-party model lower bounds.

Graph Problems and Reductions between Server-Model Problems (Details in Section 7) To bring the hardness in the Server model to the distributed setting, we have to prepare hardness for the right problems in the Server model so that it is easy to translate to the distributed setting. In particular, the problems that we need are the following graph problems.

Definition 3.3 (Server-Model Graph Problems). Let G be a graph of n nodes⁵. We partition edges of G to $E_C(G)$ and $E_D(G)$, which are given to David and Carol, respectively. The two players have to determine whether G has some property, e.g., G is a Hamiltonian cycle (Ham_n)⁶, a spanning

⁵To avoid confusion, throughout the paper we use G to denote the input graph in the Server model and N and M to denote the distributed network and its subnetwork, respectively, unless specified otherwise. For any graph H , we use $V(H)$ and $E(H)$ to denote the set of nodes and edges in H , respectively.

⁶ Ham_n is used for the Hamiltonian cycle verification problem in the Server models, where n denotes the size of input graphs, and $\text{Ham}(N)$ is used for the Hamiltonian cycle verification problem on a distributed network N (defined in Section 2.2).

tree (ST_n), or is connected (Conn_n). For the purpose of this paper in proving lower bounds for distributed algorithms, we restrict the problem and assume that in the case of the Hamiltonian cycle problem $E_G(C)$ and $E_D(C)$ are both perfect matchings.

We also consider the gap version in the case of communication complexity. The notion of δ -far is slightly different from the distributed setting (cf. Section 2.2) in that we can add *any* edges to G instead of adding *only* edges in N to M . The *main challenge* in showing hardness results for these graph problems is that some of them, e.g. Hamiltonian cycle and spanning tree verification, are not known to be hard, even in the classical two-party model (they are left as open problems in [DHK⁺12]). To get through this, we derive several new reductions (using novel gadgets) to obtain this:

Theorem 3.4. (SERVER-MODEL LOWER BOUNDS FOR Ham_n) *There exist some constants $\epsilon, \beta > 0$ such that for any n , $Q_{\epsilon, \epsilon}^{*,sv}(\text{Ham}_n)$ and $Q_{0, \epsilon}^{*,sv}((\beta n)\text{-Ham}_n)$ are $\Omega(n)$.*

We prove Theorem 3.4 using elementary (but intricate) gadget-based reductions. Thus, no knowledge in quantum computing is required to understand this proof. Theorem 3.4 also leads to lower bounds that are new even in the classical two-party model. We discuss this in Section 3.3.

Quantum Simulation Theorem: From Server Model to Distributed Algorithms (Details in Section 8) To show the role of the Server model in proving distributed algorithm lower bounds, we prove a *quantum version* of the *Simulation Theorem* of [DHK⁺12] (cf. Section 8) which shows that the hardness of graph problems of our interest in the Server model implies the hardness of these problems in the quantum distributed setting (the theorem below holds for several graph problems but we state it only for the Hamiltonian Cycle verification problem since it is sufficient for our purpose):

Theorem 3.5 (Quantum Simulation Theorem). *For any $B, L, \Gamma \geq \log L, \beta \geq 0$ and $\epsilon_0, \epsilon_1 > 0$, there exists a B -model quantum network N of diameter $\Theta(\log L)$ and $\Theta(\Gamma L)$ nodes such that if $Q_{\epsilon_0, \epsilon_1}^{*,N}((\beta \Gamma)\text{-Ham}(N)) \leq \frac{L}{2} - 2$ then $Q_{\epsilon_0, \epsilon_1}^{*,sv}((\beta \Gamma)\text{-Ham}_\Gamma) = O((B \log L)Q_{\epsilon_0, \epsilon_1}^{*,N}((\beta \Gamma)\text{-Ham}(N)))$.*

In words, the above theorem states that if there is an (ϵ_0, ϵ_1) -error quantum distributed algorithm that solves the Hamiltonian cycle verification problem on N in at most $(L/2) - 2$ time, i.e. $Q_{\epsilon_0, \epsilon_1}^{*,N}(\text{Ham}(N)) \leq (L/2) - 2$, then the (ϵ_0, ϵ_1) -error communication complexity in the Server model of the Hamiltonian cycle problem on Γ -node graphs is $Q_{\epsilon_0, \epsilon_1}^{*,sv}(\text{Ham}_\Gamma) = O((B \log L)Q_{\epsilon_0, \epsilon_1}^{*,N}(\text{Ham}(N)))$. The same statement also holds for its gap version $((\beta \Gamma)\text{-Ham}(N))$. We note that the above theorem can be extended to a large class of graph problems. The proof of the above theorem does not need any knowledge in quantum computing to follow. In fact, it can be viewed as a simple modification of the Simulation Theorem in the classical setting [DHK⁺12]. The main difference, and the most difficult part to get our Quantum Simulation Theorem to work, is to realize that we must start from the Server model instead of the two-party model.

3.2 Quantum Distributed Lower Bounds

We present specific lower bounds for various fundamental verification and optimization graph problems. Some of these bounds are new even in the classical setting. To the best of our knowledge, our bounds are the first non-trivial lower bounds for fundamental global problems.

	Problems	Previous results	Our results
B -model distributed network	Ham, ST, MST verification	$\Omega(\sqrt{n/B \log n})$ deterministic, classical communication [DHK ⁺ 12, KKP11]	$\Omega(\sqrt{n/B \log n})$ two-sided error, quantum communication with entanglement
	Conn and other verification problems from [DHK ⁺ 12]	$\Omega(\sqrt{n/B \log n})$ two-sided error, classical communication [DHK ⁺ 12]	
	α -approx MST and other optimization problems from [DHK ⁺ 12]	$\Omega(\sqrt{n/B \log n})$ Monte Carlo, classical communication for $W = \Omega(\alpha n)$ [DHK ⁺ 12]	$\Omega(\min(\sqrt{n}, W/\alpha)/\sqrt{B \log n})$ Monte Carlo, quantum communication with entanglement
Communication complexity	Ham, ST, and other verification problems	$\Omega(n)$ one-sided error, classical communication [RS95]	$\Omega(n)$ two-sided error, quantum communication with entanglement
	Gap-Ham, Gap-ST, Gap-Conn, and other gap problems for $\Omega(n)$ gap	unknown	$\Omega(n)$ one-sided error, quantum communication with entanglement

Figure 2: Previous and our new lower bounds. We note that n is the number of nodes in the network in the case of distributed network and the number of nodes in the input graph in the case of communication complexity.

1. Verification problems We prove a *tight* two-sided error quantum lower bound of $\tilde{\Omega}(\sqrt{n})$ time, where n is the number of nodes in the distributed network and $\tilde{\Theta}(x)$ hides $\text{poly log } x$, for the *Hamiltonian cycle* and *spanning tree verification* problems. Our lower bound holds even in a network of small ($O(\log n)$) diameter.

Theorem 3.6 (Verification Lower Bounds). *For any B and large n , there exists $\epsilon > 0$ and a B -model n -node network N of diameter $\Theta(\log n)$ such that any (ϵ, ϵ) -error quantum algorithm with prior entanglement for Hamiltonian cycle and spanning tree verification on N requires $\Omega(\sqrt{\frac{n}{B \log n}})$ time. That is, $Q_{\epsilon, \epsilon}^{*, N}(\text{Ham}(N))$ and $Q_{\epsilon, \epsilon}^{*, N}(\text{ST}(N))$ are $\Omega(\sqrt{\frac{n}{B \log n}})$.*

Our bound implies a new bound on the classical setting which answers the open problem in [DHK⁺12], and is the *first randomized lower bound* for both graph problems, subsuming the deterministic lower bounds for Hamiltonian cycle verification [DHK⁺12], spanning tree verification [DHK⁺12] and minimum spanning tree verification [KKP11]. It is also shown in [DHK⁺12] that Ham can be reduced to several problems via deterministic classical-communication reductions. Since these reductions can be simulated by quantum protocols, we can use these reductions straightforwardly to show that all lower bounds in [DHK⁺12] hold even in the quantum setting.

Corollary 3.7. *The statement in Theorem 3.6 holds for the following verification problems: Connected component, spanning connected subgraph, cycle containment, e -cycle containment, bipartiteness, s - t connectivity, connectivity, cut, edge on all paths, s - t cut and least-element list. (See [DHK⁺12] and Appendix A.1 for definitions.)*

Fig. 2 compares our results with previous results for verification problems.

2. Optimization problems We show a *tight* $\tilde{\Omega}(\min(W/\alpha, \sqrt{n}))$ -time lower bound for any α -approximation quantum randomized (Monte Carlo and Las Vegas) distributed algorithm for the MST problem.




Figure 3: Previous and our new bounds (cf. Theorem 3.8 and Corollary 3.9) for approximating the MST problem in distributed networks when N and α are fixed. The dashed line (in blue) represents the deterministic upper bounds (Algorithms). The dotted line (in red) is the previous lower bound for randomized algorithms. The solid line (in black) represents the bounds shown in this paper. Note that the previous lower bounds hold only in the classical setting while the new lower bounds hold in the quantum setting even when entanglement is allowed.

Theorem 3.8 (Optimization Lower Bounds). *For any n , B , W and $\alpha < W$ there exists $\epsilon > 0$ and a B -model $\Theta(n)$ -node network N of diameter $\Theta(\log n)$ and weight aspect ratio W such that any ϵ -error α -approximation quantum algorithm with prior entanglement for computing the minimum spanning tree problem on N requires $\Omega(\frac{1}{\sqrt{B \log n}} \min(W/\alpha, \sqrt{n}))$ time.*

This result generalizes the bounds in [DHK⁺12] to the quantum setting. Moreover, this lower bound implies the same bound in the classical model, which improves [DHK⁺12] (see Fig. 3) and matches the deterministic upper bound of $O(\min(W/\alpha, \sqrt{n}))$ resulting from a combination of Elkin’s α -approximation $O(W/\alpha)$ -time deterministic algorithm [Elk06] and Peleg and Rubinovich’s $O(\sqrt{n})$ -time exact deterministic algorithm [GKP98, KP98] in the classical communication model. Thus this bound is tight up to a $\Theta(\sqrt{B \log n})$ factor. It is the first bound that is *tight for all values of the aspect ratio W* . Fig. 3 compares our lower bounds with previous bounds. By using the same reduction as in [DHK⁺12], our bound also implies that all lower bounds in [DHK⁺12] hold even in the quantum setting.

Corollary 3.9. *The statement in Theorem 3.8 also holds for the following problems: minimum spanning tree, shallow-light tree, s -source distance, shortest path tree, minimum routing cost spanning tree, minimum cut, minimum s - t cut, shortest s - t path and generalized Steiner forest. (See [DHK⁺12] and Appendix A.1 for definitions.)*

3.3 Additional Results: Lower Bounds on Communication Complexity

In proving the results in previous subsections, we prove several bounds on the Server model. Since the Server model is stronger than the standard communication complexity model (as discussed in Subsection 3.1), we obtain lower bounds in the communication complexity model as well. Some of these lower bounds are new even in the classical setting. In particular, our bounds in Theorem 3.4 lead to the following corollary. (Note that we use $Q_{\epsilon_0, \epsilon_1}^{*,cc}(\mathcal{P}_n)$ to denote the communication complexity of verifying property \mathcal{P} of n -node graphs on the standard quantum communication complexity model with entanglement.)

Corollary 3.10. *For any n and some constants $\epsilon, \beta > 0$, $Q_{\epsilon, \epsilon}^{*,cc}(\mathcal{P}_n) = \Omega(n)$, and $Q_{0, \epsilon}^{*,cc}((\beta n) - \mathcal{P}_n) \geq Q_{0, \epsilon}^{*,sv}((\beta n) - \mathcal{P}_n) = \Omega(n)$, where \mathcal{P}_n can be any of the following verification problems: Hamiltonian cycle, spanning tree, connectivity, s - t connectivity, and bipartiteness.*

To the best of our knowledge, the lower bounds for Hamiltonian cycle and spanning tree verification problems are the first two-sided error lower bounds for these problems, even in the classical two-party setting (only nondeterministic, thus one-sided error, lower bounds are previously known [RS95]). The bounds for Bipartiteness and s - t connectivity follow from a reduction from Inner Product given in [BFS86], and a lower bound for Connectivity was recently shown in [IKL⁺12]. We note that we prove the gap versions via a reduction from recent lower bounds in [KdW12] and observe new lower bounds for the gap versions of Set Disjointness and Equality.

4 Other Related Work

While our work focuses on solving graph problems in quantum distributed networks, there are several prior works focusing on other distributed computing problems (including communication complexity in the two-party or multiparty communication model) using quantum effects. We note that fundamental distributed computing problems such as leader election and byzantine agreement have been shown to solved better using quantum phenomena (see e.g., [DP08, TKM05, BOH05]). Entanglement has been used to reduce the amount of communication of a specific function of input data distributed among 3 parties [CB97] (see also the work of [BvDHT99, dW02, TS99] on multiparty quantum communication complexity).

There are several results showing that quantum communication complexity in the two-player model can be more efficient than classical randomized communication complexity (e.g. [BCW98, Raz99]). These results also easily extend to the so-called number-in-hand multiparty model (in which players have separate inputs). As of now no separation between quantum and randomized communication complexity is known in the number-on-the-forehead multiparty model, in which players' inputs overlap. Other papers concerning quantum distributed computing include [BR03, CKS10, KMT09, KMT10, PSK03, GBK⁺08].

5 Conclusion and Open Problems

In this paper, we derive several lower bounds for important network problems in a quantum distributed network. We show that quantumness does not really help in obtaining faster distributed algorithms for fundamental problems such as minimum spanning tree, minimum cut, and shortest

paths. Our approach gives a uniform way to prove lower bounds for various problems. Our technique closely follows the Simulation Theorem introduced by Das Sarma et al. [DHK⁺12], which shows how to use the two-party communication complexity to prove lower bounds for distributed algorithms. The main difference of our approach is the use of the Server model. We show that many problems that are hard in the quantum two-party communication setting (e.g. IPmod3) are also hard in the Server model, and show new reductions from these problems to graph verification problems of our interest. Some of these reductions give tighter lower bounds even in the classical setting.

Since the technique of Das Sarma et al. can be used to show lower bounds of many problems that are not covered in this paper (e.g. [FHW12, HW12, NDP11, LPS13, DMP13, Gha14, CHGK14]), it is interesting to see if these lower bounds remain valid in the quantum setting. Since most of these problems rely on a reduction from the set disjointness problem, the main challenge is to obtain new reductions that start from problems that are proved hard on the Server model such as IPmod3. One problem that seems to be harder than others is the random walk problem [NDP11, DNPT13] since the previous lower bound in the classical setting requires a *bounded-round* communication complexity [NDP11]. Proving lower bounds for the random walk problem thus requires proving a bounded-round communication complexity in the Server model as the first step. This requires different techniques since the nonlocal games used in this paper destroy the round structure of protocols.

It is also interesting to better understand the role of the Server model: Can we derive a quantum two-party version of the Simulation Theorem, thus eliminating the need of the Server model? Is the Server model *strictly* stronger than the two-party quantum communication complexity model? Also, it will be interesting to explore *upper* bounds in the quantum setting: Do quantum distributed algorithms help in solving other fundamental graph problems ?

Part II

Proofs

6 Server Model Lower Bounds via Nonlocal Games (Lemma 3.2)

In this section, we prove Lemma 3.2 which shows how to use nonlocal games to prove server model lower bounds. Then, we use it to show server-model lower bounds for two problems called *Inner Product mod 3* (denoted by $\text{IPmod}3_n$) and *Gap Equality* with parameter δ (denoted by $\delta\text{-Eq}_n$). These lower bounds will be used in the next section.

Our proof makes use of the relationship between the server model and *nonlocal games*. In such games, Alice and Bob receive input x and y from some distribution π that is known to the players. As usual they want to compute a boolean function $f(x, y)$ such as Equality or Inner Product mod 3. However, they cannot communicate to each other. Instead, each of them can send one bit, say a and b , to a referee. The referee then combines a and b using some function g to get an output of the game $g(a, b)$. The goal of the players is to come up with a strategy (which could depend on distribution π and function g) that maximizes the probability that $g(a, b) = f(x, y)$. We call this the *winning probability*. One can define different nonlocal games based on what function g the referee will use. Two games of our interest are *XOR*- and *AND*-games where g is *XOR* and *AND* functions, respectively.

Our proof follows the framework of proving two-party quantum communication complexity lower bounds via nonlocal games (see, e.g., [LS09a, LZ10, KdW12]). The key modification is the following lemma which shows that the XOR- and AND-game players can make use of an efficient server-model protocol to guarantee a good winning probability.

Lemma 3.2 (Restated). *For any boolean function f and $\epsilon_0, \epsilon_1 \geq 0$, there is an (two-player nonlocal) XOR-game strategy \mathcal{A}' (respectively, AND-game strategy \mathcal{A}'') such that, for any input (x, y) , with probability $4^{-2Q_{\epsilon_0, \epsilon_1}^{*, sv}(f)}$, \mathcal{A}' (respectively, \mathcal{A}'') outputs $f(x, y)$ with probability at least $1 - \epsilon_{f(x, y)}$ (i.e. it outputs 1 with probability at least $1 - \epsilon_1$ and 0 with probability at least $1 - \epsilon_0$); otherwise (with probability $1 - 4^{-2Q_{\epsilon_0, \epsilon_1}^{*, sv}(f)}$), \mathcal{A}' outputs 0 and 1 with probability 1/2 each (respectively, \mathcal{A}'' outputs 0 with probability 1).*

Proof. We prove the lemma in a similar way to the proof of Theorem 5.3 in [LS09a] (attributed to Buhrman). Consider any boolean function f . Let \mathcal{A} be any (ϵ_0, ϵ_1) -error server-model protocol for computing f with communication complexity T . We will construct (two-player) nonlocal XOR-games and AND-games strategies, denoted by \mathcal{A}' and \mathcal{A}'' , respectively, that *simulate* \mathcal{A} . First we simulate \mathcal{A} with an additional assumption that there is a “fake server” that sends messages to players (Alice and Bob) in the nonlocal games, but the two players in the games do not send any message to the fake server. Later we will eliminate this fake server. We will refer to parties in the server model as Carol, David, and the *real* server, while we call the nonlocal game players Alice, Bob, and the *fake* server.

Using teleportation (where we can replace a qubit by two classical bits when there is an entanglement; see, e.g., [NC04]), it can be assumed that Carol and David send $2T$ *classical* bits to the real server instead of sending T qubits (the server can set up the necessary entanglement for free). Assume that, on an input (x, y) , Carol and David send bits c_t and d_t in the t^{th} round, respectively. (We note one detail here that in reality c_t and d_t , for all t , are random variables. We will ignore this fact here to illustrate the main idea. More details are in Appendix B.)

Now, Alice, Bob and the fake server generate shared random strings $a_1 \dots a_t$ and $b_1 \dots b_t$ (this can be done since their states are entangled). These strings serve as a “guessed” communication sequence of \mathcal{A} . Alice, Bob and the fake protocol simulate Carol, David and the real protocol, respectively. However, in each round t , instead of sending bit c_t that Carol sends to the real server, Alice simply looks at a_t and continues playing if her guessed communication is the same as the real communication, i.e. $c_t = a_t$. Otherwise, she “aborts”: In the XOR-game protocol \mathcal{A}' she outputs 0 and 1 with probability $1/2$ each, and in the AND-game protocol \mathcal{A}'' she outputs 0. Bob does the same thing with d_t and b_t .

The fake server simply assumes it receives a_t and b_t and continues sending messages to Alice and Bob. Observe that the probability of never aborting is 4^{-T} (i.e., when the random strings $a_1 \dots a_T$ and $b_1 \dots b_T$ are the same as the communication sequences $c_1 \dots c_T$ and $d_1 \dots d_T$, respectively). If no one aborts, Alice will output Carol’s output while Bob will output 0 in the XOR-game protocol \mathcal{A}' and 1 in the AND-game protocol \mathcal{A}'' . If no one aborts, Alice, Bob and the fake server perfectly simulate \mathcal{A} and thus output $f(x, y)$ with probability at least $1 - \epsilon_{f(x, y)}$ in both protocols⁷. Otherwise (with probability at most $1 - 4^{-T}$) one or both players will abort and the output will be randomly 0 and 1 in \mathcal{A}' and 0 in \mathcal{A}'' . This is exactly what we claim in the theorem except that there is a fake server.

Now we eliminate the fake server. Notice that the fake server never receives anything from Alice and Bob. Hence we can assume that the fake server sends all his messages to Alice and Bob before the game starts (before the input is given), and those messages can be viewed as prior entanglement. We thus get standard XOR- and AND-game strategies without a fake server. \square

Now we define and prove lower bounds for $\text{IPmod}3_n$ and $\delta\text{-Eq}_n$. In both problems Carol and David are given n -bit strings x and y , respectively. In $\text{IPmod}3_n$, they have to output 1 if $(\sum_{i=1}^n x_i y_i) \bmod 3 = 0$ and 0 otherwise. In $\delta\text{-Eq}_n$, the players are *promised* that either $x = y$ or the hamming distance $\Delta(x, y) > \delta$ where $\Delta(x, y) = |\{i \mid x_i \neq y_i\}|$. They have to output 1 if and only if $x = y$. This theorem will be used in the next section.

Theorem 6.1. *For some $\beta, \epsilon > 0$ and any large n , $Q_{\epsilon, \epsilon}^{*, sv}(\text{IPmod}3_n)$ and $Q_{0, \epsilon}^{*, sv}((\beta n)\text{-Eq}_n)$ are $\Omega(n)$.*

Now we give a high-level idea of the proof of Theorem 6.1 (see Appendix B for detail).

To show that $Q_{\epsilon, \epsilon}^{*, sv}(\text{IPmod}3_n) = \Omega(n)$, we use an XOR-game strategy \mathcal{A}' and $\epsilon_0 = \epsilon_1 = \epsilon$ from Lemma 3.2. Using this we can extend the theorem of Linial and Shraibman [LS09b] from the two-party model to the server model and show that $Q_{\epsilon, \epsilon}^{*, sv}(f)$ is lower bounded by an *approximate* γ_2 norm: $Q_{\epsilon, \epsilon}^{*, sv}(f) = \Omega(\log \gamma_2^{2\epsilon}(A_f))$ for some matrix A_f defined by f . Using $f = \text{IPmod}3_n$, one can then extend the proof of Lee and Zhang [LZ10, Theorem 8] to lower bound $\log \gamma_2^{2\epsilon}(A_f)$ by an *approximate degree* $\deg_{2\epsilon}(f')$ of some function f' . Finally, one can follow the proof of Sherstov [She11] and Razborov [Raz03] to prove that $\deg_{2\epsilon}(f') = \Omega(n)$. Combining these three steps, we have

$$Q_{\epsilon, \epsilon}^{*, sv}(\text{IPmod}3_n) = \Omega(\log \gamma_2^{2\epsilon}(A_{\text{IPmod}3_n})) = \Omega(\deg_{2\epsilon}(f')) = \Omega(n).$$

We note that this technique actually extends all lower bounds we are aware of on the two-party model (e.g. those in [Raz03, She11, LZ10]) to the server model.

To prove that $Q_{0, \epsilon}^{*, sv}((\beta n)\text{-Eq}_n) = \Omega(n)$ for some $\beta, \epsilon > 0$, we use an AND-game strategy \mathcal{A}'' with $\epsilon_0 = 0$ and $\epsilon_1 = \epsilon = 1/2$ from Lemma 3.2. We adapt a recent result by Klauck and de Wolf

⁷That is, if $f(x, y) = 0$, they output 0 with probability at least $1 - \epsilon_0$ and, if $f(x, y) = 1$, they output 1 with probability at least $1 - \epsilon_1$

[KdW12], which shows that $Q_{0,1/2}^{*,cc}(f) \geq (\log \text{fool}^1(f))/4 - 1/2$. Here $\text{fool}^1(f)$ refers to the size of the 1-fooling set of f , which is defined to be a set $F = \{(x, y)\}$ of input pairs with the following properties.

- If $(x, y) \in F$ then $f(x, y) = 1$
- If $(x, y), (x', y') \in F$ then $f(x, y') = 0$ or $f(x', y) = 0$

We observe that the lower bound in [KdW12] actually applies to AND-games as follows. Suppose Alice and Bob receive inputs (x, y) , then perform local measurements on a shared entangled state, and output bits a, b . Then the probability that $a \wedge b = 1$ for a uniformly random $x, y \in F$ is at most $1/\text{fool}^1(f)$, if the probability that $a \wedge b = 1$ for (x, y) with $f(x, y) = 0$ is always 0.

Lemma 3.2 for the case of AND-games implies that there is an AND-game strategy \mathcal{A}'' such that if $f(x, y) = 0$ then \mathcal{A}'' always output 0 and if $f(x, y) = 1$ then \mathcal{A}'' outputs 1 with probability at least $(1 - \epsilon)4^{-2Q_{0,\epsilon}^{*,sv}(f)}$. This implies that $(1 - \epsilon)4^{-2Q_{0,\epsilon}^{*,sv}(f)} \leq 1/\text{fool}^1(f)$. In other words, if $\text{fool}^1(f) = 2^{\Omega(n)}$ then $Q_{0,1/2}^{*,sv}(f) = \Omega(n)$.

All that remains is to define a good fooling set for (βn) -Eq_n. Fix any $1/4 > \beta > 0$. The idea is to use a good error-correcting code to construct the fooling set. Recall that $\Delta(x, y)$ denote the Hamming distance between x and y . Let C be a set of n -bit strings such that the Hamming distance between any distinct $x, y \in C$ is at least $2\beta n$. Due to the Gilbert-Varshamov bound such codes C exist with $|C| \geq 2^{(1-H(2\beta))n} = 2^{\Omega(n)}$, where H denotes the binary entropy function. Hence we have $Q_{0,1/2}^{*,sv}((\beta n)\text{-Eq}_n) = \Omega(n)$.

7 Server-model Lower Bounds for Ham_n (Theorem 3.4)

In this section, we prove Theorem 3.4, which leads to new lower bounds for several graph problems as discussed in Section 3.3. The proof uses gadget-based reductions between problems on the Server model.

Theorem 3.4 (Restated). *For any n and some constants $\epsilon, \beta > 0$,*

$$Q_{\epsilon,\epsilon}^{*,sv}(\text{Ham}_n) = \Omega(n) \quad \text{and} \quad (1)$$

$$Q_{0,\epsilon}^{*,sv}((\beta n)\text{-Ham}_n) = \Omega(n). \quad (2)$$

We first sketch the lower bound proof of $Q_{\epsilon,\epsilon}^{*,sv}(\text{Ham}_n)$ and show later how to extend to the gap version. More detail can be found in Section C. We will show that for any $0 \leq \epsilon \leq 1$ and some constant c , $Q_{\epsilon,\epsilon}^{*,sv}(\text{IPmod}3_n) = O(Q_{\epsilon,\epsilon}^{*,sv}(\text{Ham}_{cn}))$. The theorem then immediately follows from the fact that $Q_{\epsilon,\epsilon}^{*,sv}(\text{IPmod}3_n) = \Omega(n)$ (cf. Theorem 6.1).

Let $x = x_1 \dots x_n$ and $y = y_1 \dots y_n$ be the input of IPmod₃_n. We construct a graph G which is an input of Ham_{cn} as follows. The graph G consists of n gadgets, denoted by G_1, \dots, G_n . For any $1 \leq i \leq n - 1$, gadgets G_i and G_{i+1} share exactly three nodes denoted by v_i^0, v_i^1, v_i^2 . Each gadget G_i is constructed based on the values of x_i and y_i as outlined in Fig. 4. The following observation can be checked by drawing G_i for all cases of x_i and y_i (as in Fig. 5).

Observation 7.1. *For any value of (x_i, y_i) , G_i consists of three paths where v_{i-1}^j is connected by a path to $v_i^{(j+x_i \cdot y_i) \bmod 3}$, for any $0 \leq j \leq 2$. Moreover, Alice's (respectively Bob's) edges, i.e. thin (red) lines (respectively thick (blue) lines) in Fig. 4, form a matching that covers all nodes except v_i^j (respectively v_{i-1}^j) for all $0 \leq j \leq 2$.*




Figure 4: The construction of gadget G_i . If $x_i = 0$ then Alice adds dashed thin edges (in red); otherwise she adds solid thin edges (in red). If $y_i = 0$ then Bob adds dashed thick edges (in blue); otherwise he adds solid thick edges (in blue).




Figure 5: Gadget G_i for different values of x_i and y_i . The main observation is that if $x_i \cdot y_i = 0$ then G_i consists of paths from v_{i-1}^j to v_i^j for all $0 \leq j \leq 2$. Otherwise, it consists of paths from v_{i-1}^j to $v_i^{(j+1) \bmod 3}$.




Figure 6: The graph G consists of gadgets G_1, \dots, G_n . The solid thick edges (in gray) linking between v_0^j and v_n^j , for $0 \leq j \leq 2$ represent the fact that $v_0^j = v_n^j$. Lines that appear in each gadget G_i depicts what we observe in Observation 7.1: solid thin lines (in red) represent paths that will appear in G_i if $x_i \cdot y_i = 0$, and dashed thick lines (in blue) represent paths that will appear in G_i if $x_i \cdot y_i = 1$.

Thus, when we put all gadgets together, graph G will consist of three paths connecting between nodes in $\{v_0^j\}_{0 \leq j \leq 2}$ on one side and nodes in $\{v_n^j\}_{0 \leq j \leq 2}$ on the other. How these paths look like depends on the structure of each gadget G_i which depends on the value of $x_i \cdot y_i$. The following lemma follows trivially from Observation 7.1.

Lemma 7.2. G consists of three paths P^0 , P^1 and P^2 where for any $0 \leq j \leq 2$, P^j has v_0^j as one end vertex and $v_n^{(j + \sum_{1 \leq i \leq n} x_i \cdot y_i) \bmod 3}$ as the other.

Now, we complete the description of G by letting $v_0^j = v_n^j$ for all $0 \leq j \leq 2$. It then follows that G is a Hamiltonian cycle if and only if $\sum_{1 \leq i \leq n} x_i \cdot y_i \bmod 3 \neq 0$ (see Fig. 6; also see Lemma C.3 and Fig. 12 in Section C). Thus we can check that $\sum_{1 \leq i \leq n} x_i \cdot y_i \bmod 3$ is zero or not by checking whether G is a Hamiltonian cycle or not. Theorem 3.4 now follows from Theorem 6.1.

To show a lower bound of $Q_{0,\epsilon}^{*,sv}((\beta n)\text{-Ham}_n)$, we reduce from $(\beta n)\text{-Eq}_n$ in a similar way using gadget G_i shown in Fig. 7. For any $1 \leq i \leq n-1$, gadget G_i and G_{i+1} share v_i^0 and v_i^1 , and we let $v_0^0 = v_0^1$ and $v_n^0 = v_n^1$. It is straightforward to show that if $x = y$, then G is a Hamiltonian cycle, and if $x_{i_j} \neq y_{i_j}$ for some $i_1 < i_2 < \dots < i_\delta$, then G consists of δ cycles where each cycle starts at gadget G_{i_j} and ends at gadget $G_{i_{j+1}}$. Note that our reduction gives a simplification of the rather complicated reduction in [DHK⁺12, Section 6].

8 The Quantum Simulation Theorem (Theorem 3.5)

In this section, we show that in the quantum setting, a server-model lower bound implies a B -model lower bound, as in Theorem 3.5.




Figure 7: Gadget G_i to reduce from (βn) -Eq $_n$ to (βn) -Ham $_n$.

Theorem 3.5 (Restated). *For any $B, L, \Gamma \geq \log L, \beta \geq 0$ and $\epsilon_0, \epsilon_1 > 0$, there exists a B -model quantum network N of diameter $\Theta(\log L)$ and $\Theta(\Gamma L)$ nodes such that if $Q_{\epsilon_0, \epsilon_1}^{*, N}((\beta \Gamma)\text{-Ham}(N)) \leq \frac{L}{2} - 2$ then $Q_{\epsilon_0, \epsilon_1}^{*, sv}((\beta \Gamma)\text{-Ham}_\Gamma) = O((B \log L) Q_{\epsilon_0, \epsilon_1}^{*, N}((\beta \Gamma)\text{-Ham}(N)))$.*

In words, the above theorem states that if there is an (ϵ_0, ϵ_1) -error quantum distributed algorithm that solves the Hamiltonian cycle verification problem on N in at most $(L/2) - 2$ time, i.e. $Q_{\epsilon_0, \epsilon_1}^{*, N}(\text{Ham}(N)) \leq (L/2) - 2$, then the (ϵ_0, ϵ_1) -error communication complexity in the server model of the Hamiltonian cycle problem on Γ -node graphs is $Q_{\epsilon_0, \epsilon_1}^{*, sv}(\text{Ham}_\Gamma) = O((B \log L) Q_{\epsilon_0, \epsilon_1}^{*, N}(\text{Ham}(N)))$. The same statement also holds for its gap version. We note that the above theorem can be extended to a large class of graph problems with some certain properties. We state it for only Ham for simplicity.

We give the proof idea here and provide full detail in Appendix D. Although we recommend the readers to read this before the full proof and believe that it is enough to reconstruct the full proof, this proof idea can be skipped without loss of continuity.

We note again that the main idea of this theorem essentially follows the ideas developed in a line of work in [PR00, Elk06, LPSP06, KKP11, DHK⁺12]. In particular, we construct a network following ideas in [PR00, Elk06, LPSP06, KKP11, DHK⁺12], and our argument is based on simulating the network by the three players of the server model. This idea follows one of many ideas implicit in the proof of the Simulation Theorem in [DHK⁺12] which shows how two players can simulate some class of networks. However, as we noted earlier, the previous proof does not work in the quantum setting, and it is still open whether the Simulation Theorem holds in the quantum setting. We instead use the server model. Another difference is that we prove the theorem for graph problems instead of problems on strings (such as Equality or Disjointness). This leads to some simplified reductions since reductions can be done easier in the communication complexity setting.

To explain the main idea, let us focus on the non-gap version of Hamiltonian cycle verification and consider a B -model network N' in Fig. 8 consisting of Γ paths, each of length L , where we have an edge between any pair of the leftmost (respectively, rightmost) nodes of paths. Now we will prove that if $Q_{\epsilon_0, \epsilon_1}^{*, N}(\text{Ham}(N)) \leq (L/2) - 2$ then $Q_{\epsilon_0, \epsilon_1}^{*, sv}(\text{Ham}_\Gamma) = 0$ (i.e. no communication is needed from Carol and David to the server!). Note that this statement is stronger than the theorem




Figure 8: The network N' used in the proof idea of Theorem 3.5 with sets S_C^t and S_D^t .

statement but it is not useful since N' has diameter $\Theta(L)$ which is too large. We will show how to modify N' to get the desired network N later.

Let paths in N' be P^1, \dots, P^Γ and nodes in path P^i be v_1^i, \dots, v_L^i . Let \mathcal{A} be an (ϵ_0, ϵ_1) -error quantum distributed algorithm that solves the Hamiltonian cycle verification problem on network N' ($\text{Ham}(N')$) in at most $(L/2) - 2$ time.

We show that Carol, David and the server can solve the Hamiltonian cycle problem on a Γ -node input graph without any communication, essentially by “simulating” \mathcal{A} on some input subnetwork M corresponding to the server-model input graph $G = (U, E_C \cup E_D)$ in the following sense. When receiving E_C and E_D , the three parties will construct a subnetwork M of N' (without communication) in such a way that M is a Hamiltonian cycle if and only if $G = (U, E_C \cup E_D)$ is. Next, they will simulate algorithm \mathcal{A} in such a way that, at any time t and for each node v_j^i in N' , there will be exactly one party among Carol, David and the server that knows *all information that v_j^i should know in order to run algorithm \mathcal{A}* , i.e., the state of v_j^i as well as the messages (each consisting of B quantum bits) sent to v_j^i from its neighbors at time t . The party that knows this information will pretend to be v_j^i and apply algorithm \mathcal{A} to get the state of v_j^i at time $t + 1$ as well as the messages that v_j^i will send to its neighbors at time $t + 1$. We say that this party *owns* v_j^i at time t . Details are as follows.

Initially at time $t = 0$, we let Carol own all leftmost nodes, and David own all rightmost nodes while the server own the rest, i.e. Carol, David and the server own the following sets of nodes respectively (see Fig. 8):

$$\begin{aligned} S_C^0 &= \{v_1^i \mid 1 \leq i \leq \Gamma\}, \\ S_D^0 &= \{v_L^i \mid 1 \leq i \leq \Gamma\}, \\ S_S^0 &= V(N') \setminus (S_C^0 \cup S_D^0). \end{aligned} \tag{3}$$

After Carol and David each receive a perfect matching, denoted by E_C and E_D respectively, on the node set $U = \{u_1, \dots, u_\Gamma\}$, they construct a subnetwork M of N' as follows. For any $i \neq j$,




Figure 9: The subnetwork M when the input perfect matchings are $E_C = \{(u_1, u_2), (u_3, u_4), \dots, (u_{\Gamma-1}, u_{\Gamma})\}$ and $E_D = \{(u_2, u_3), (u_4, u_5), \dots, (u_{\Gamma}, u_1)\}$ (M consists of all bold edges).

Carol marks $v_1^i v_1^j$ as participating in M if and only if $u_i u_j \in E_C$. Similarly, David marks $v_L^i v_L^j$ as participating in M if and only if $u_i u_j \in E_D$. The server marks all edges in all paths as participating in M . Fig. 9 shows an example. We note the following observation which relies on the fact that E_C and E_D are perfect matchings.

Observation 8.1. *The number of cycles in $G = (U, E_C \cup E_D)$ is the same as the number of cycles in M .*

Now the three parties start a simulation. Recall that at time $t = 0$ the three parties own nodes in the sets S_C^0 , S_D^0 and S_S^0 as in Eq.(3). Our goal is to simulate \mathcal{A} for one time step and make sure that Carol, David and the server own the following sets respectively (see Fig. 8):

$$\begin{aligned} S_C^1 &= \{v_1^i, v_2^i \mid 1 \leq i \leq \Gamma\}, \\ S_D^1 &= \{v_{L-1}^i, v_L^i \mid 1 \leq i \leq \Gamma\}, \\ S_S^1 &= V(N') \setminus (S_C^1 \cup S_D^1). \end{aligned} \tag{4}$$

To do this, the parties simulate \mathcal{A} on the nodes they own for one time step. This means that each of them will know the states and out-going messages at time $t = 1$ (i.e., after \mathcal{A} is executed once) of nodes they own. Observe that although Carol knows the state of v_1^i , for any i , at time $t = 1$, she is not able to simulate \mathcal{A} on v_1^i for one more step since she does not know the message sent from v_2^i to v_1^i at time $t = 1$. This information is known by the server who owns v_2^i at time $t = 0$. Thus, we let the server send this message to Carol. Additionally, for Carol to own node v_2^i at time $t = 1$, it suffices to let the server send the state of v_2^i and the message sent from v_3^i to v_2^i at time $t = 1$ (which are known by the server since it owns v_2^i and v_3^i at time $t = 0$). The messages sent from the server to David can be constructed similarly. It can be checked that after this communication the three parties own nodes as in Eq.(4) and thus they can simulate \mathcal{A} for one more step.

Using a similar argument as the above we can guarantee that at any time $t \leq (L/2) - 2$, Carol,

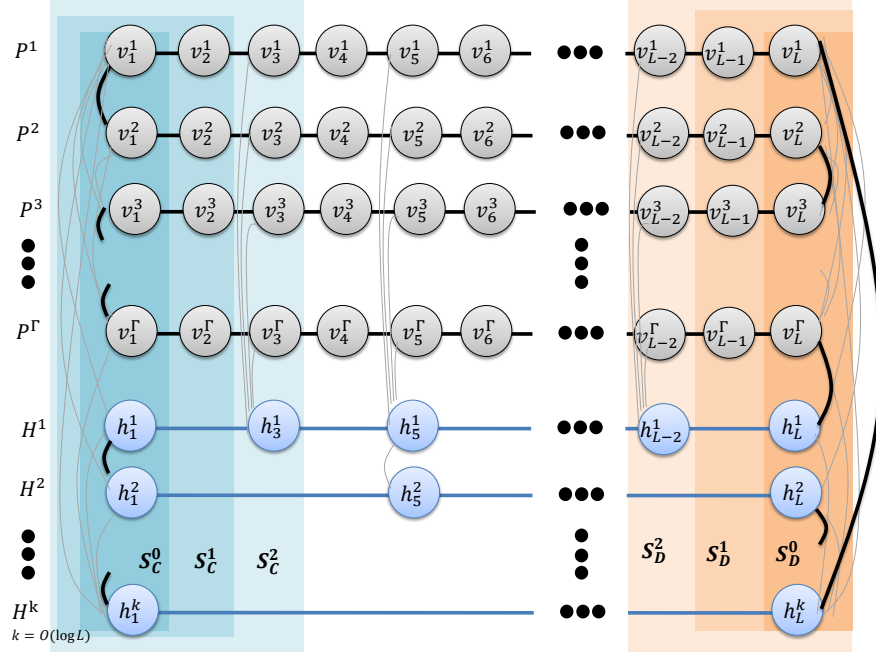


Figure 10: The network N consisting of network N' and some “highways” which are paths with nodes h_j^i (i.e., nodes in blue). Bold edges show an example of subnetwork M when the input perfect matchings are $E_C = \{(u_1, u_2), (u_3, u_4), \dots, (u_{\Gamma+k-1}, u_{\Gamma+k})\}$ and $E_D = \{(u_2, u_3), (u_4, u_5), \dots, (u_{\Gamma+k}, u_1)\}$. Pale edges are those in N but not in M .

David and the server own nodes in the following sets respectively:

$$\begin{aligned} S_C^t &= \{v_j^i \mid 1 \leq i \leq \Gamma, 1 \leq j \leq t+1\}, \\ S_D^t &= \{v_j^i \mid 1 \leq i \leq \Gamma, L-t \leq j \leq L\}, \\ S_S^t &= V(N') \setminus (S_C^t \cup S_D^t). \end{aligned}$$

Thus, if algorithm \mathcal{A} terminates in $(L/2) - 2$ steps then Carol, David and the server will know whether M is a Hamiltonian cycle or not with (ϵ_0, ϵ_1) -error by reading the output of nodes they own. By Observation 8.1, they will know whether $G = (U, E_C \cup E_D)$ is a Hamiltonian cycle or not with the same error bound.

Now we modify N' to get network N of small diameter. A simple idea to slightly reduce the diameter is to add a path having half the number of nodes of other paths and connect its nodes to every other node on the other paths (see path H^1 in Fig. 10). This path helps reducing the diameter from L to roughly $(L/2) - 2$ since any pair of nodes can connect in roughly $(L/2) - 2$ hops through this path. By adding about $O(\log L)$ such paths (with H^i having half the number of nodes of H^{i-1}) as in Fig. 10, we can reduce the diameter to $O(\log L)$. We call the new paths *highways*.

We can use almost the same argument as before to prove the theorem, by modifying sets S_C^t , S_D^t and S_S^t appropriately as in Fig. 10 and consider the input graph $G = (U, E_C \cup E_D)$ of $\Gamma + k$ nodes, where k is the number of highways. The exception is that now Carol and David have to speak a little. For example, observe that if the three parties want to own the states of S_C^1 , S_D^1 and S_S^1 at time $t = 1$, Carol has to send to the server the messages sent from node h_1^i to its right neighbor, for

all i . Since this message has size at most B , and the simulation is done for $Q_{\epsilon_0, \epsilon_1}^{*,N}(\text{Ham}(N))$ steps, Carol will send $O((B \log n) Q_{\epsilon_0, \epsilon_1}^{*,N}(\text{Ham}(N)))$ qubits to the server. David will have to send the same amount of information and thus the complexity in the server model is as claimed.

9 Proof of main theorems (Theorem 3.6 & 3.8)

9.1 Proof of Theorem 3.6

Theorem 3.6 (Restated). *For any B and large n , there exists $\epsilon > 0$ and a B -model n -node network N of diameter $\Theta(\log n)$ such that any (ϵ, ϵ) -error quantum algorithm with prior entanglement for Hamiltonian cycle and spanning tree verification on N requires $\Omega(\sqrt{\frac{n}{B \log n}})$ time. That is, $Q_{\epsilon, \epsilon}^{*,N}(\text{Ham}(N))$ and $Q_{\epsilon, \epsilon}^{*,N}(\text{ST}(N))$ are $\Omega(\sqrt{\frac{n}{B \log n}})$.*

We note from Theorem 3.4 that

$$Q_{\epsilon, \epsilon}^{*,sv}(\text{Ham}_\Gamma) > c'\Gamma \quad (5)$$

for some $\epsilon > 0$ and $c' > 0$. Let c be the constant in the big-Oh in Theorem 3.5. Let $L = \lfloor \frac{c'}{c} \sqrt{\frac{n}{B \log n}} \rfloor$ and $\Gamma = \lceil \sqrt{Bn \log n} \rceil$. Assume that

$$Q_{\epsilon, \epsilon}^{*,N}(\text{Ham}(N)) \leq L/2 \leq \frac{c'}{2c} \sqrt{\frac{n}{B \log n}}. \quad (6)$$

By Theorem 3.5, there is a network N of diameter $O(\log L) = O(\log n)$ and $\Theta(L\Gamma) = \Theta(n)$ nodes such that $Q_{\epsilon, \epsilon}^{*,sv}(\text{Ham}_\Gamma) \leq (cB \log L) Q_{\epsilon, \epsilon}^{*,N}(\text{Ham}(N)) \leq (cB \log L) \left(\frac{c'}{2c} \sqrt{\frac{n}{B \log n}} \right) \leq c' \sqrt{Bn \log n}$ where the second equality is by Eq. (6). This contradicts Eq.(5), thus proving that $Q_{\epsilon, \epsilon}^{*,N}(\text{Ham}(N)) > L/2 \geq \frac{c'}{4c} \sqrt{\frac{n}{B \log n}}$.

To show a lower bound of $Q_{\epsilon, \epsilon}^{*,N}(\text{ST}(N))$, let \mathcal{A} be an algorithm that solves spanning tree verification on N in $T_{\mathcal{A}}$ time. We can use \mathcal{A} to verify if a subnetwork M is a Hamiltonian cycle as follows. First, we check that all nodes have degree two in M (this can be done in $O(D)$ time). If not, M is not a Hamiltonian cycle. If it is, then M consists of cycles. Now we delete one edge e in M arbitrarily, and use \mathcal{A} to check if this subnetwork is a spanning tree. It is easy to see that this subnetwork is a spanning tree if and only if M is a Hamiltonian cycle. The running time of our algorithm is $T_{\mathcal{A}} + O(D)$. The lower bound of $Q_{\epsilon, \epsilon}^{*,N}(\text{Ham}(N))$ implies that $T_{\mathcal{A}} = \Omega(\sqrt{\frac{n}{B \log n}})$.

9.2 Proof of Theorem 3.8

Theorem 3.8 (Restated). *For any n , B , W and $\alpha < W$ there exists $\epsilon > 0$ and a B -model $\Theta(n)$ -node network N of diameter $\Theta(\log n)$ such that any ϵ -error α -approximation quantum algorithm with prior entanglement for computing the minimum spanning tree problem on N with weight function $w : E(N) \rightarrow \mathbb{R}_+$ such that $\frac{\max_{e \in E(N)} w(e)}{\min_{e \in E(N)} w(e)} \leq W$ requires $\Omega(\frac{1}{\sqrt{B \log n}} \min(W/\alpha, \sqrt{n}))$ time.*

We note from Theorem 3.4 that

$$Q_{0, \epsilon}^{*,sv}((\beta\Gamma)\text{-Ham}_\Gamma) > c'\Gamma \quad (7)$$

for some constant $\beta > 0$, $\epsilon > 0$ and $c' > 0$. Let c be the constant in the big-Oh in Theorem 3.5. Let $L = \lfloor \frac{c'}{c\sqrt{B\log n}} \min(\frac{W}{\alpha}, \sqrt{n}) \rfloor$ and $\Gamma = \lceil \sqrt{B\log n} \max(\frac{n\alpha}{W}, \sqrt{n}) \rceil$. We prove the following claim the same way we prove Theorem 3.6 in the previous section.

Claim 9.1. $Q_{0,\epsilon}^{*,N}((\beta\Gamma)\text{-Ham}) > \frac{L}{2} \geq \frac{c'}{4c} \min(W/\alpha, \sqrt{\frac{n}{B\log n}})$

Proof. Assume that

$$Q_{0,\epsilon}^{*,N}((\beta\Gamma)\text{-Ham}) \leq \frac{L}{2} \leq \frac{c'}{2c} \min(W/\alpha, \sqrt{\frac{n}{B\log n}}). \quad (8)$$

By Theorem 3.5, there is a network N of diameter $\Theta(\log L) = O(\log n)$ and $\Theta(L\Gamma) = \Theta(n)$ nodes such that

$$\begin{aligned} Q_{0,\epsilon}^{*,sv}((\beta\Gamma)\text{-Ham}_\Gamma) &\leq (cB\log L)Q_{0,\epsilon}^{*,N}((\beta\Gamma)\text{-Ham}) \\ &\leq (cB\log L)(L/2) \\ &\leq \frac{c'\sqrt{B\log n}}{2} \min(\frac{W}{\alpha}, \sqrt{n}) \\ &\leq \frac{c'\sqrt{B\log n}}{2} \max(\frac{n\alpha}{W}, \sqrt{n}) \\ &\leq c'\Gamma \end{aligned}$$

where the second equality is by Eq. (8) and the fourth inequality is because if $\frac{W}{\alpha} \leq \sqrt{n}$ then $\alpha \geq W/\sqrt{n}$ and thus $n\alpha/W \geq \sqrt{n} \geq W/\alpha$. This contradicts Eq.(7). \square

Now assume that there is an ϵ -error quantum distributed algorithm \mathcal{A} that finds an α -approximate MST in $T_{\mathcal{A}}$ time. We use \mathcal{A} to construct an $(0, \epsilon)$ -error algorithm that solves $(\beta\Gamma)\text{-Ham}(N)$ in $T_{\mathcal{A}} + O(D)$ time as follows. Let M be the input subnetwork. First we check if all nodes have degree exactly two in M . If not then M is not a Hamiltonian cycle and we are done. If it is then M consist of one cycle or more. It is left to check whether M is connected or not. To do this, we assign weight 1 to all edges in H and weight W to the rest edges. We use \mathcal{A} to compute an α -approximate MST T . Then we compute the weight of T in $O(D) = O(\log n)$ rounds. If T has weight at most $\alpha(n-1)$ then we say that H is connected; otherwise we say that it is $(\beta\Gamma)$ -far from being connected.

To show that this algorithm is $(0, \epsilon)$ -error, observe that, for any i , if H is i -far from being connected then the MST has weight at least $(n-1-i) + iW$ since the MST will contain at least i edges of weight W . If H is connected then the MST has weight exactly $n-1$ which means that T will have weight at most $\alpha(n-1)$ with probability at least $1-\epsilon$, and we will say that H is connected with probability at least $1-\epsilon$. Otherwise, if H is $(\beta\Gamma)$ -far from connected then T always have weight at least

$$(n-1-\beta\Gamma) + \beta\Gamma W \geq \beta\Gamma W \geq \beta(\sqrt{B\log n} \max(\frac{n\alpha}{W}, \sqrt{n}))W \geq \beta\sqrt{B\log n} \frac{n\alpha}{W} W \geq \alpha n > \alpha(n-1)$$

for large enough n (note that β is a constant), and we will always say that H is $(\beta\Gamma)$ -far from being connected. Thus algorithm is $(0, \epsilon)$ -error as claimed.

Part III

Appendix

A Detailed Definitions

A.1 Quantum Distributed Network Models

Informal descriptions

We first describe a *general* model which will later make it easier to define some specific models we are considering. We assume some familiarity with quantum computation (see, e.g., [NC04, Wat11] for excellent resources). A general distributed network N is modeled by a set of n processors, denoted by u_1, \dots, u_n , and a set of *bandwidth* parameters between each pair of processors, denoted by $B_{u_i u_j}$ for any $i \neq j$, which is used to bound the size of messages sent from u_i to u_j . Note that $B_{u_i u_j}$ could be zero or infinity. To simplify our formal definition, we let $B_{u_i u_i} = \infty$ for all i .

In the beginning of the computation, each processor u_i receives an input string x_i , each of size b . The processors want to cooperatively compute a global function $f(x_1, \dots, x_n)$. They can do this by communicating in *rounds*. In each rounds, processor u_i can send a message of $B_{u_i u_j}$ bits or qubits to processor u_j . (Note that u_i can send different messages to u_j and u_k for any $j \neq k$.) We assume that each processor has unbounded computational power. Thus, between each round of communication, processors can perform any computation (even solving an NP-complete problem!). The *time complexity* is the minimum number of rounds needed to compute the function f . We can categorize this model further based on the type of communication (classical or quantum) and computation (deterministic or randomized).

In this paper, we are interested in quantum communication when errors are allowed and nodes share entangled qubits. In particular, for any $\epsilon > 0$ and function f , we say that a quantum distributed algorithm \mathcal{A} is ϵ -*error* if for any input (x_1, \dots, x_n) , after \mathcal{A} is executed on this input any node u_i knows the value of $f(x_1, \dots, x_n)$ correctly with probability at least $1 - \epsilon$. We let $Q_\epsilon^{*,N}(N)$ denote the time complexity (number of rounds) of computing function f on network N with ϵ -error.

In the special case where f is a boolean function, for any $\epsilon_0, \epsilon_1 > 0$ we say that \mathcal{A} computes f with (ϵ_0, ϵ_1) -*error* if, after \mathcal{A} is executed on any input (x_1, \dots, x_n) , any node u_i knows the value of $f(x_1, \dots, x_n)$ correctly with probability at least $1 - \epsilon_0$ if $f(x_1, \dots, x_n) = 0$ and with probability at least $1 - \epsilon_1$ otherwise. We let $Q_{\epsilon_0, \epsilon_1}^{*,N}(N)$ denote the time complexity of computing boolean function f on network N with (ϵ_0, ϵ_1) -error.

Two main models of interest are the B -model (also known as $\mathcal{CONGEST}(B)$) and a new model we introduce in this paper called the *server model*. The B -model is modeled by an undirected n -node graph, where vertices model the processors and edges model the links between the processors. For any nodes (processors) u_i and u_j , $B_{u_i u_j} = B_{u_j u_i} = B$ if there is an edge $u_i u_j$ in the graph and $B_{u_i u_j} = B_{u_j u_i} = 0$ otherwise.

In the server model, there are three processors, denoted by *Carol*, *David* and the *server*. In each round, Carol and David can send one bit to each other and to the server while receiving an arbitrarily large message from the server, i.e. $B_{Carol, David} = B_{David, Carol} = B_{Carol, Server} = B_{David, Server} = 1$ and $B_{Server, Carol} = B_{Server, David} = \infty$.

We will also discuss the *two-party communication complexity model* which is simply the network

of two processors called Alice and Bob with bandwidth parameters $B_{Alice,Bob} = B_{Bob,Alice} = 1$. (Note that, this model is sometimes defined in such a way that only one of the processors can send a message in each round. The communication complexity in this setting might be different from ours, but only by a factor of two.)

When N is the server or two-party communication complexity model, we use $Q_\epsilon^{*,sv}(f)$ and $Q_\epsilon^{*,cc}(f)$ instead of $Q_\epsilon^{*,N}(f)$.

Formal definitions

Network States The *pure state* of a quantum network of n nodes with parameters $\{B_{u_i u_j}\}_{1 \leq i, j \leq n}$ is represented as a vector in a Hilbert space

$$\bigotimes_{1 \leq i, j \leq n} H_{u_i u_j} = H_{u_1 u_1} \otimes H_{u_1 u_2} \otimes \dots \otimes H_{u_1 u_n} \otimes H_{u_2 u_1} \otimes \dots \otimes H_{u_2 u_n} \otimes \dots \otimes H_{u_n u_n}$$

where \otimes is the tensor product. Here, $H_{u_i u_i}$, for any i , is a Hilbert space of arbitrary finite dimension representing the “workspace” of processor u_i . In particular, we let K be an arbitrarily large number (thus the complexity of the problem cannot depend on K) and $H_{u_i u_i}$ be a 2^K -dimensional Hilbert space. Additionally, $H_{u_i u_j}$, for any $i \neq j$, is a Hilbert space representing the $B_{u_i u_j}$ -qubit communication channel from u_i to u_j . Its dimension is $2^{B_{u_i u_j}}$ if $B_{u_i u_j}$ is finite and 2^K if $B_{u_i u_j} = \infty$.

The *mixed state* of a quantum network N is a probabilistic distribution over its pure states

$$\{(p_i, |\psi_i\rangle)\} \text{ with } p_i \geq 0 \text{ and } \sum_i p_i = 1.$$

We note that it is sometimes convenient to represent a mixed state by a *density matrix* $\rho = \sum_i p_i |\psi_i\rangle \langle \psi_i|$.

Initial state In the model *without* prior entanglement, the initial (pure) state of a quantum protocol on input (x_1, \dots, x_n) is the vector

$$|\psi_{x_1, \dots, x_n}^0\rangle = \bigotimes_{1 \leq i, j \leq n} |\psi_{x_1, \dots, x_n}^0(i, j)\rangle = |\psi_{x_1, \dots, x_n}^0(1, 1)\rangle |\psi_{x_1, \dots, x_n}^0(1, 2)\rangle \dots |\psi_{x_1, \dots, x_n}^0(n, n)\rangle$$

where $|\psi_{x_1, \dots, x_n}^0(i, j)\rangle$ for any $1 \leq i, j \leq n$ is a vector in $H_{u_i u_j}$ such that $|\psi_{x_1, \dots, x_n}^0(i, i)\rangle = |x_i, 0\rangle$ for any i and $|\psi_{x_1, \dots, x_n}^0(i, j)\rangle = |0\rangle$ for any $i \neq j$ (here, $|0\rangle$ represents an arbitrary unit vector independent of the input). Informally, this corresponds to the case where each processor u_i receives an input x_i and workspaces and communication channel are initially “clear”.

With prior entanglement, the initial (pure) state is a unit vector of the form

$$|\psi_{x_1, \dots, x_n}^0\rangle = \sum_w \left(\alpha_w \bigotimes_{1 \leq i, j \leq n} |\psi_{w, x_1, \dots, x_n}^0(i, j)\rangle \right) \quad (9)$$

where $|\psi_{w, x_1, \dots, x_n}^0(i, j)\rangle$ for any $1 \leq i, j \leq n$ is a vector in $H_{u_i u_j}$ such that $|\psi_{w, x_1, \dots, x_n}^0(i, i)\rangle = |x_i, w\rangle$ for any i and $|\psi_{w, x_1, \dots, x_n}^0(i, j)\rangle = |0\rangle$ for any $i \neq j$. Here, the coefficients α_w are arbitrary real numbers satisfying $\sum_w \alpha_w^2 = 1$ that is independent of the input (x_1, \dots, x_n) . Informally, this corresponds to the case where processors share entangled qubits in their workspaces.

Note that we can assume the global state of the network to be always a pure state, since any mixed state can be purified by adding qubits to the processor’s workspaces, and ignoring these in later computations.

Communication Protocol The communication protocol consists of rounds of *internal computation* and *communication*. In each internal computation of the t^{th} round, each processor u_i applies a *unitary transformation* to its incoming communication channels and its own memory, i.e. $H_{u_j u_i}$ for all j . That is, it applies a unitary transformation of the form

$$C_{t,u_i} \otimes \left(\bigotimes_{1 \leq j \leq n, k \neq i} I_{u_j u_k} \right) \quad (10)$$

which acts as an identity on $H_{u_j u_k}$ for all $1 \leq j \leq n$ and $k \neq i$. At the end of the internal computation, we require the communication channel to be clear, i.e. if we would measure any communication channel in the computational basis then we would get $|0\rangle$ with probability one. This can easily be achieved by swapping some fresh qubits from the private workspace into the communication channel. Note that the processors can apply the transformations corresponding to an internal computation simultaneously since they act on different parts of the network's state.

To define communication, let us divide the workspace $H_{u_i u_i}$ of processor u_i further to

$$H_{u_i u_i} = H_{u_i u_i,1} \otimes H_{u_i u_i,2} \otimes \dots \otimes H_{u_i u_i,n}$$

where $H_{u_i u_i,j}$ has the same dimension as $H_{u_i u_j}$. The space $H_{u_i u_i,j}$ can be thought of as a place where u_i prepares the messages it wants to send to u_j in each round, while $H_{u_i u_i,i}$ holds u_i 's remaining workspace. Now, for any $j \neq i$, u_i sends a message to u_j simply by swapping the qubits in $H_{u_i u_i,j}$ with those in $H_{u_i u_j}$. Note that u_i does not receive any information in this process since the communication channel $H_{u_i u_j}$ is clear after the internal computation. Also note that we can perform the swapping operations between any pair $i \neq j$ simultaneously since they act on different part of the network state. This completes one round of communication. We let

$$|\psi_{x_1, \dots, x_n}^t\rangle \quad (11)$$

denote the network state after t rounds of communication.

At the end of a T -round protocol, we compute the *output of processor u_i* as follows. We view part of $H_{u_i u_i}$ as an output space of u_i , i.e. $H_{u_i u_i} = H_{O_i} \otimes H_{W_i}$ for some H_{O_i} and H_{W_i} . We compute the output of u_i by measuring H_{O_i} in the computational basis. That is, if we let K' be the number of qubits in H_{O_i} and the network state after a T -round protocol be ψ_{x_1, \dots, x_n}^T then, for any $w \in \{0, 1\}^{K'}$,

$$Pr[\text{Processor } u_i \text{ outputs } w] = |\langle \psi_{x_1, \dots, x_n}^T | w \rangle|^2.$$

Fig. 11 depicts a quantum circuit corresponding to a communication protocol on three processors.

Error and Time Complexity For any $0 \leq \epsilon \leq 1$, we say that a quantum protocol \mathcal{A} on network N computes function f with ϵ -error if for any input (x_1, \dots, x_n) of f and any processor u_i , u_i outputs $f(x_1, \dots, x_n)$ with probability at least $1 - \epsilon$ after \mathcal{A} is executed. The ϵ -error time complexity of computing function f on network N , denoted by $Q_{\epsilon}^{*,N}(f)$, is the minimum T such that there exists a T -round quantum protocol on network N that computes function f with ϵ -error. We note that we allow the protocol to start with an entangled state. The $*$ in the notation follows the convention to contrast with the case that we do not allow prior entanglement (which is not considered in this

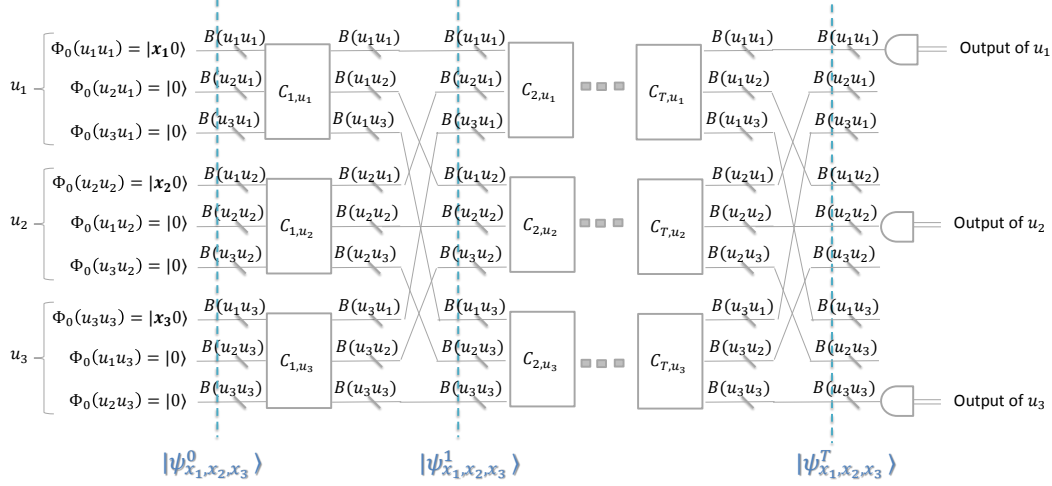


Figure 11: A circuit corresponding to T rounds of communication on general distributed network having 3 processors. The information flows from left to right and the line crossing each wire with a number $B_{u_i u_j}$ means that there are $B_{u_i u_j}$ qubits of information flowing through such wire. We note that the initial state in the picture is without entanglement.

paper). When N is the server model and two-party communication complexity model mentioned earlier, we use $Q_{\epsilon}^{*,sv}(f)$ and $Q_{\epsilon}^{*,cc}(f)$ respectively to denote the ϵ -error time complexity.

If f is a boolean function, we will sometimes distinguish between the error of outputting 0 and 1. For any $0 \leq \epsilon_0, \epsilon_1 \leq 1$ we say that \mathcal{A} computes f with (ϵ_0, ϵ_1) -error if for any input (x_1, \dots, x_n) of f and any processor u_i , if $f(x_1, \dots, x_n) = 0$ then u_i outputs 0 with probability at least $1 - \epsilon_0$ and otherwise u_i outputs 1 with probability at least $1 - \epsilon_1$. The time complexity, denoted by $Q_{\epsilon_0, \epsilon_1}^{*,N}(f)$ is defined in the same way as before. We will also use $Q_{\epsilon_0, \epsilon_1}^{*,sv}(f)$ and $Q_{\epsilon_0, \epsilon_1}^{*,cc}(f)$.

A.2 Distributed Graph Verification Problems

In the distributed network N , we describe its subgraph M as an input as follows. Each node u_i in N receives an n -bit binary string x_{u_i} as an input. We let $x_{u_i, u_1}, \dots, x_{u_i, u_n}$ be the bits of x_{u_i} . Each bit x_{u_i, u_j} indicates whether edge $u_i u_j$ participates in the subgraph M or not. The indicator variables must be consistent, i.e., for every edge $u_i u_j \in E(N)$, $x_{u_i, u_j} = x_{u_j, u_i}$ (this is easy to verify with a single round of communication) and if there is no edge between u_i and u_j in N then $x_{u_i, u_j} = x_{u_j, u_i} = 0$.

We define $M_{x_{u_1}, \dots, x_{u_n}}$, or simply M , to be subgraph of N having edges whose indicator variables are 1; that is,

$$E(M) = \{(u_i, u_j) \in E \mid \forall i \neq j, x_{u_i, u_j} = x_{u_j, u_i} = 1\}.$$

We list the following problems concerning the verification of properties of subnetwork M on distributed network N from [DHK⁺12].

- **connected spanning subgraph verification:** We want to verify whether M is connected and spans all nodes of N , i.e., every node in N is incident to some edge in M .

- **cycle containment verification:** We want to verify if M contains a cycle.
- **e -cycle containment verification:** Given an edge e in M (known to vertices adjacent to it), we want to verify if M contains a cycle containing e .
- **bipartiteness verification:** We want to verify whether M is bipartite.
- **s - t connectivity verification:** In addition to N and M , we are given two vertices s and t (s and t are known by every vertex). We would like to verify whether s and t are in the same connected component of M .
- **connectivity verification:** We want to verify whether M is connected.
- **cut verification:** We want to verify whether M is a cut of N , i.e., N is not connected when we remove edges in M .
- **edge on all paths verification:** Given two nodes u, v and an edge e . We want to verify whether e lies on all paths between u and v in M . In other words, e is a u - v cut in M .
- **s - t cut verification:** We want to verify whether M is an s - t cut, i.e., when we remove all edges $E(M)$ of M from N , we want to know whether s and t are in the same connected component or not.
- **least-element list verification [Coh97, KKM⁺08]:** The input of this problem is different from other problems and is as follows. Given a distinct rank (integer) $r(v)$ to each node v in the weighted graph N , for any nodes u and v , we say that v is the *least element* of u if v has the lowest rank among vertices of distance at most $d(u, v)$ from u . Here, $d(u, v)$ denotes the weighted distance between u and v . The *Least-Element List* (LE-list) of a node u is the set $\{\langle v, d(u, v) \rangle \mid v \text{ is the least element of } u\}$.
In the least-element list verification problem, each vertex knows its rank as an input, and some vertex u is given a set $S = \{\langle v_1, d(u, v_1) \rangle, \langle v_2, d(u, v_2) \rangle, \dots\}$ as an input. We want to verify whether S is the least-element list of u .
- **Hamiltonian cycle verification:** We would like to verify whether M is a Hamiltonian cycle of N , i.e., M is a simple cycle of length n .
- **spanning tree verification:** We would like to verify whether M is a tree spanning N .
- **simple path verification:** We would like to verify that M is a simple path, i.e., all nodes have degree either zero or two in M except two nodes that have degree one and there is no cycle in M .

A.3 Distributed Graph Optimization Problems

In the graph optimization problems \mathcal{P} on distributed networks, such as finding MST, we are given a positive weight $\omega(e)$ on each edge e of the network (each node knows the weights of all edges incident to it). Each pair of network and weight function (N, ω) comes with a nonempty set of *feasible solution* for problem \mathcal{P} ; e.g., for the case of finding MST, all spanning trees of N are feasible solutions. The goal of \mathcal{P} is to find a feasible solution that minimizes or maximizes the total weight.

We call such solution an *optimal solution*. For example, the spanning tree of minimum weight is the optimal solution for the MST problem. We let $W = \max_{e \in E(N)} \omega(e) / \min_{e \in E(N)} \omega(e)$.

For any $\alpha \geq 1$, an α -*approximate solution* of \mathcal{P} on weighted network (N, ω) is a feasible solution whose weight is not more than α (respectively, $1/\alpha$) times of the weight of the optimal solution of \mathcal{P} if \mathcal{P} is a minimization (respectively, maximization) problem. We say that an algorithm \mathcal{A} is an α -approximation algorithm for problem \mathcal{P} if it outputs an α -approximate solution for any weighted network (N, ω) . In case we allow errors, we say that an α -approximation T -time algorithm is ϵ -error if it outputs an answer that is not α -approximate with probability at most ϵ and always finishes in time T , regardless of the input.

Note the following optimization problems on distributed network N from [DHK⁺12].

- In the **minimum spanning tree** problem [Elk06, PR00], we want to compute the weight of the minimum spanning tree (i.e., the spanning tree of minimum weight). In the end of the process all nodes should know this weight.
- Consider a network with two cost functions associated to edges, weight and length, and a root node r . For any spanning tree T , the radius of T is the maximum length (defined by the length function) between r and any leaf node of T . Given a root node r and the desired radius ℓ , a **shallow-light tree** [Pel00] is the spanning tree whose radius is at most ℓ and the total weight is minimized (among trees of the desired radius).
- Given a node s , the **s -source distance** problem [Elk05] is to find the distance from s to every node. In the end of the process, every node knows its distance from s .
- In the **shortest path tree** problem [Elk06], we want to find the shortest path spanning tree rooted at some input node s , i.e., the shortest path from s to any node t must have the same weight as the unique path from s to t in the solution tree. In the end of the process, each node should know which edges incident to it are in the shortest path tree.
- The **minimum routing cost spanning tree** problem (see e.g., [KKM⁺08]) is defined as follows. We think of the weight of an edge as the cost of routing messages through this edge. The routing cost between any node u and v in a given spanning tree T , denoted by $c_T(u, v)$, is the distance between them in T . The routing cost of the tree T itself is the sum over all pairs of vertices of the routing cost for the pair in the tree, i.e., $\sum_{u, v \in V(N)} c_T(u, v)$. Our goal is to find a spanning tree with minimum routing cost.
- A set of edges E' is a **cut** of N if N is not connected when we delete E' . The **minimum cut** problem [Elk04] is to find a cut of minimum weight. A set of edges E' is an **s - t cut** if there is no path between s and t when we delete E' from N . The **minimum s - t cut** problem is to find an s - t cut of minimum weight.
- Given two nodes s and t , the **shortest s - t path** problem is to find the length of the shortest path between s and t .
- The **generalized Steiner forest** problem [KKM⁺08] is defined as follows. We are given k disjoint subsets of vertices V_1, \dots, V_k (each node knows which subset it is in). The goal is to find a minimum weight subgraph in which each pair of vertices belonging to the same subsets is connected. In the end of the process, each node knows which edges incident to it are in the solution.

B Detail of Section 6

B.1 Two-player XOR Games

We give a brief description of XOR games. AND game can be described similarly (their formal description is not needed in this paper). For a more detailed description as well as the more general case of nonlocal games see, e.g., [LS09a, Bri11] and references therein. An XOR game is played by three parties, Alice, Bob and a referee. The game is defined by \mathcal{X} and \mathcal{Y} which is the set of input to Alice and Bob, respectively, π , a joint probability distribution $\pi : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$, and a boolean function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$.

At the start of the game, the referee picks a pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ according to the probability distribution π and sends x to Alice and y to Bob. Alice and Bob then answer the referee with one-bit message a and b . The players win the game if the value $a \oplus b$ is equal to $f(x, y)$. In other words, Alice and Bob want the XOR of their answers to agree with f , explaining the name “XOR game.”

The goal of the players is to maximize the *bias* of the game, denoted by $\text{Bias}_\pi(f)$, which is the probability that Alice and Bob win minus the probability that they lose. In the classical setting, this is

$$\begin{aligned} \text{Bias}_\pi(f) &= \max_{\substack{a: \mathcal{X} \rightarrow \{-1, 1\}, \\ b: \mathcal{Y} \rightarrow \{-1, 1\}}} \sum_{(x, y) \in \mathcal{X} \times \mathcal{Y}} (-1)^{f(x, y)} \pi(x, y) (-1)^{a(x)} (-1)^{b(y)} \\ &= \max_{\substack{a \in \{-1, 1\}^{|\mathcal{X}|}, \\ b \in \{-1, 1\}^{|\mathcal{Y}|}}} \mathbb{E}_{(x, y) \sim \pi} [(-1)^{a(x)} (-1)^{b(y)} (-1)^{f(x, y)}]. \end{aligned}$$

In the quantum setting, Alice and Bob are allowed to play an *entangled strategy* where they may make use of an entangled state they share prior to receiving the input. That is, Alice and Bob start with some shared pure quantum state which is independent of the input and after they receive input (x, y) they make some projective measurements depending on (x, y) and return the result of their measurements to the referee. Formally, an XOR entangled strategy is described by a shared (pure) quantum state $|\psi\rangle \in \mathbb{C}^{d \times d}$ for some $d \geq 1$ and a choice of projective measurements $\{A_x^0, A_x^1\}$ and $\{B_y^0, B_y^1\}$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. When receiving input x and y , the probability that Alice and Bob output $(a, b) \in \{0, 1\}^2$ is $\langle \psi | A_x^a \otimes B_y^b | \psi \rangle$. Thus, the maximum correlation can be shown to be (see [Bri11] for details)

$$\text{Bias}_\pi(f) = \max \mathbb{E}_{(x, y) \sim \pi} [\langle \psi | (A_x^1 - A_x^0) \otimes (B_y^1 - B_y^0) | \psi \rangle (-1)^{f(x, y)}]$$

where the maximization is over pure states $|\psi\rangle$ and projective measurements $\{A_x^0, A_x^1\}$ and $\{B_y^0, B_y^1\}$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. In the rest of this paper, $\text{Bias}_\pi(f)$ always denotes the maximum correlation in the quantum setting. We let

$$Q^{*, \text{XOR}}(f) = \min_{\pi} \text{Bias}_\pi(f).$$

We note that while the players could start the game with a mixed state, it can be shown that pure entangled states suffice in order to maximize the winning probability (see, e.g., [Bri11]).

B.2 From Nonlocal Games to Server-Model Lower Bounds

Lemma B.1 (Lemma 3.2 restated). *For any boolean function f and $\epsilon_0, \epsilon_1 \geq 0$, there is an XOR-game strategy \mathcal{A}' and AND-game strategy \mathcal{A}'' such that, for any input (x, y) ,*

- *with probability $4^{-2Q_{\epsilon_0, \epsilon_1}^{*,sv}}(f)$, \mathcal{A}' and \mathcal{A}'' are able to simulate a protocol in the server model and hence output $f(x, y)$ with probability at least $1 - \epsilon_{f(x,y)}$;*
- *otherwise \mathcal{A}' outputs 0 and 1 with probability $1/2$ each, and \mathcal{A}'' outputs 0 with probability 1.*

Proof. We have sketched the proof in Section 6.1. We now provide more detail.

Let $c = Q_{\epsilon_0, \epsilon_1}^{*,sv}(f)$, i.e. Carol and David communicate with the server for c rounds where each of them sends one qubit to the server per round while the server sends them messages of arbitrary size. While Alice and Bob cannot run a protocol \mathcal{A} in the server model since they cannot communicate to each other, we show that they can obtain the output of \mathcal{A} with probability $\frac{1}{4^{2c}}$. To be precise, for any input (x, y) let $p_{x,y}$ and $q_{x,y}$ be the probability that $\mathcal{A}(x, y)$ is zero and one respectively. We will show that

(12)

Alice and Bob can obtain the final state of \mathcal{A} with probability 4^{-2c} and in that case output the correct answer with high probability. If they do not obtain that state one of them will output a random bit for XOR games and one of them will output 0 for AND games.

Hence the XOR game will accept with probability $\frac{1}{2}(1 - 4^{-2c}) + 4^{-2c}q_{x,y} = \frac{1}{2} + (q_{x,y} - \frac{1}{2})4^{-2c}$ and thus have a bias of at least $4^{-2c} \cdot \min\{1/2 - \epsilon_0, 1/2 - \epsilon_1\}$.

The AND game will accept 1-inputs with probability at least $q'_{x,y} \geq \frac{q_{x,y}}{4^{2c}}$. Furthermore if \mathcal{A} never accepts a 0-input, then neither will the AND game.

Let us first prove Statement (12) with an additional assumption that there is a “fake” server that Alice and Bob can receive a message from but cannot talk to (we will eliminate this fake server later). We will call this a fake server to distinguish it from the “real” server in the server model.

First let us note the Carol and David need not talk to each other, but can send their messages to the server who can pass them to the other player. Since the server can also set up entanglement between the three parties without cost, Carol, David and the server can use *teleportation* (see [NC04] for details) and we can assume that in protocol \mathcal{A} Carol and David send 2 classical bits per round to the server instead of one qubit. These two bits are also uniformly distributed, regardless of the state of the qubit.

Thus, for any input (x, y) , the messages sent by Carol and David in protocol \mathcal{A} will be $a, b \in \{0, 1\}^{2c}$ with some probability, say $p_{x,y,a,b}$. For simplicity, let us assume that each communication sequence (a, b) leads to a unique output of \mathcal{A} on input (x, y) (e.g., by requiring Carol and David to send their result to the server in the last round). Let $\mathcal{A}(x, y, a, b)$ be the output of the protocol \mathcal{A} on input (x, y) with communication sequence (a, b) . Then the probability that \mathcal{A} outputs zero and one is, respectively,

$$p_{x,y} = \sum_{(a,b): \mathcal{A}(x,y,a,b)=0} p_{x,y,a,b} \quad \text{and} \quad q_{x,y} = \sum_{(a,b): \mathcal{A}(x,y,a,b)=1} p_{x,y,a,b}.$$

The strategy of Alice and Bob who play the XOR and AND games is trying to “guess” this sequence.

In particular, Alice, Bob and the fake server will pretend to be Carol, David and the real server as follows. Before receiving the input, Alice, Bob and the fake server use their shared entanglement

to create two shared random strings of length $2c$, denoted by a' and b' , and start their initial entangled states with the same states of Carol, David and the server. In each round t of \mathcal{A} , Alice, Bob and the fake server will simulate Carol, David and the real server, respectively, as follows. Let $c_{t,1}$ and $c_{t,2}$ be two bits sent by Carol to the real server at round t . Alice will check whether the guessed communication sequence a' is correct by checking if $c_{t,1}$ and $c_{t,2}$ are the same as a'_{2t-1} and a'_{2t} which are the $(2t-1)^{th}$ and $(2t)^{th}$ bits of a' . If they are not the same then she will ‘abort’ which means that

- Alice will output 0 and 1 uniformly random if she is playing an XOR game, and
- Alice will output 0 if she is playing an AND game.

Similarly, Bob will check whether the guessed communication sequence b' is correct by checking b'_{2t-1} and b'_{2t} with two classical bits sent by David to the server. Moreover, the fake server will pretend that it receives a'_{2t-1} , a'_{2t} , b'_{2t-1} and b'_{2t} to execute \mathcal{A} and send huge quantum messages to Alice and Bob. Alice and Bob then execute \mathcal{A} using these messages. After $2c$ rounds (if no player aborts), the players output the following.

- In XOR games, Alice will send Carol’s output to the referee, and Bob will send 0 to the referee.
- In AND games, Alice will send Carol’s output to the referee, and Bob will send 1 to the referee.

Thus, if one or both players aborts then the output of an XOR game will be uniformly random in $\{0, 1\}$. For an AND game in case of a abort the players reject. Otherwise, the result of the XOR and AND games will be $\mathcal{A}(x, y, a, b)$. The probability that Alice and Bob do not abort, given that the communication sequence of \mathcal{A} on input (x, y) is a and b is $Pr[a' = a \wedge b' = b] = \frac{1}{4^{2c}}$.

This almost proves Statement (12) (thus the lemma) except that there is a fake server sending information to Alice and Bob in the XOR and AND game strategy. To remove the fake server, observe that we do not need an input in order to generate the messages the fake server sent to Alice and Bob. Thus, we change the strategy to the following. As previously done, before Alice, Bob and the fake server receive an input they generate shared random strings (a', b') and start with the initial states of Carol, David and the real server. In addition to this, the fake server use the string a' and b' to generate the messages sent by the real server to Carol and David. It then sends this information to Alice and Bob. We now remove the fake server completely and mark this point as a starting point of the XOR and AND games. After Alice and Bob receive input (x, y) , they simulate protocol \mathcal{A} as before. In each round, when they are supposed to receive messages from the fake server, they read messages that the fake server sent before the game starts. Since the fake server sends the same messages, regardless of when it sends, the result is the same as before. Thus, we achieve Statement (12) even when there is no fake server. This completes the proof of Lemma B.1. \square

B.3 Lower Bound for $\text{IPmod}3_n$

Using the above lemma, we prove the following lemma which extends the theorem of Linial and Shraibman [LS09b] from the two-party model to the server model. Our proof makes use of XOR games as in [LS09a] (attributed to Buhrman). For any boolean function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$,

let A_f be a $|\mathcal{X}|$ -by- $|\mathcal{Y}|$ matrix such that $A_f[x, y] = (-1)^{f(x, y)}$. Recall that for any matrix A , $\|A\|_1 = \sum_{i, j} |A_{i, j}|$.

Lemma B.2. *For boolean function f and $0 \leq \epsilon < 1/4$*

$$4^{2Q_\epsilon^{*,sv}(f)} \geq \max_M \frac{\langle A_f, M \rangle - 2\epsilon \|M\|_1}{\gamma_2^*(M)} = \gamma_2^{2\epsilon}(A_f).$$

Proof. We first prove the following claim.

Claim B.3. *For any boolean functions f, g on the same domain, probability distribution π and $0 \leq \epsilon \leq 1$,*

$$\text{Bias}_\pi(g) \geq \frac{\langle A_f, A_g \circ \pi \rangle - 2\epsilon}{4^{2Q_\epsilon^{*,sv}(f)}}.$$

Proof. First, suppose that when receive input (x, y) , Alice and Bob can somehow compute $f(x, y)$ and use this as an answer to the XOR game (e.g., Alice and Bob returns $f(x, y)$ and 1 to the referee respectively). What is the bias this strategy can achieve? Since the probability of winning is $\sum_{x, y: f(x, y)=g(x, y)} \pi(x, y)$, the bias is straightforwardly

$$\sum_{\substack{x, y \\ f(x, y)=g(x, y)}} \pi(x, y) - \sum_{\substack{(x, y) \\ f(x, y) \neq g(x, y)}} \pi(x, y) = \sum_{x, y} \pi(x, y) A_f[x, y] A_g[x, y] = \langle A_f, A_g \circ \pi \rangle$$

Let \mathcal{A} be an ϵ -error protocol for computing f in the server model and $\mathcal{A}(x, y)$ be the output of \mathcal{A} (which could be randomized) on input (x, y) . Now suppose that Alice and Bob use $\mathcal{A}(x, y)$ to play the XOR game. Then the winning probability will decrease by at most ϵ . Thus the bias is at least

$$\langle A_f, A_g \circ \pi \rangle - 2\epsilon. \tag{13}$$

Now suppose that Alice and Bob use protocol \mathcal{A}' from Lemma B.1 with $\epsilon_0 = \epsilon_1 = \epsilon$ to play the XOR game. With probability $1 - 4^{-2Q_\epsilon^{*,sv}(f)}$, \mathcal{A}' will output randomly; this means that the bias is 0. Otherwise, \mathcal{A}' will behave as an ϵ -error algorithm. Thus, we conclude from Eq.(13) that the bias is at least

$$4^{-2Q_\epsilon^{*,sv}(f)} (\langle A_f, A_g \circ \pi \rangle - 2\epsilon).$$

This completes the claim. □

Thus, for any π

$$4^{2Q_\epsilon^{*,sv}(f)} \geq \frac{\langle A_f, A_g \circ \pi \rangle - 2\epsilon}{\text{Bias}_\pi(g)}.$$

Note that $\text{Bias}_\pi(g) = \gamma_2^*(A_g \circ \pi)$ [Tsi87] (also see [LS09a, Theorem 5.2]). So,

$$4^{2Q_\epsilon^{*,sv}(f)} \geq \frac{\langle A_f, A_g \circ \pi \rangle - 2\epsilon}{\gamma_2^*(A_g \circ \pi)}.$$

Since this is true for any π and g ,

$$4^{2Q_\epsilon^{*,sv}(f)} \geq \max_{\pi, g} \frac{\langle A_f, A_g \circ \pi \rangle - 2\epsilon}{\gamma_2^*(A_g \circ \pi)} = \max_M \frac{\langle A_f, M \rangle - 2\epsilon \|M\|_1}{\gamma_2^*(M)}.$$

This proves the first inequality in Lemma B.2.

For the second inequality, we use Proposition 1 in [LZ10] (proved in [LS09a]) which states that for any norm Φ , matrix A and $0 \leq \alpha < 1$, the α -approximate norm is

$$\Phi^\alpha(A) = \max_W \frac{|\langle A, W \rangle| - \alpha \|W\|_1}{\Phi^*(W)}.$$

This means that $\gamma_2^{2\epsilon}(A_f) = \max_M \frac{|\langle A_f, M \rangle| - 2\epsilon \|M\|_1}{\gamma_2^*(M)}$ as claimed. \square

For finite sets X, Y , and E , a function $f : E^n \rightarrow \{0, 1\}$, and a function $g : X \times Y \rightarrow E$, the *block composition* of f and g is the function $f \circ g^n : X^n \times Y^n \rightarrow \{0, 1\}$ defined by $(f \circ g^n)(x, y) = f(g(x^1, y^1), \dots, g(x^n, y^n))$ where $(x^i, y^i) \in X \times Y$ for all $i = 1, \dots, n$. For any boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, let f' be such that, for all $x \in \{0, 1\}^n$, $f'(x) = -1$ if $f(x) = 0$ and $f'(x) = 1$ otherwise. The ϵ -approximate degree of f , denoted by $\deg_\epsilon(f)$ is the least degree of a real polynomial p such that $|f'(x) - p(x)| \leq \epsilon$ for all $x \in \{0, 1\}^n$. We say that g is *strongly balanced* if all rows and columns in the matrix A_g sum to zero. For any m -by- n matrix A , let $\text{size}(A) = m \times n$. We now prove a “server-model version” of Lee and Zhang’s theorem [LZ10, Theorem 8]. Our proof is essentially the same as their proof (also see [LS09a, Theorem 7.6]).

Lemma B.4. *For any finite sets X, Y , let $g : X \times Y \rightarrow \{0, 1\}$ be any strongly balanced function. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be an arbitrary function. Then*

$$Q_\epsilon^{*,sv}(f \circ g^n) \geq \deg_{4\epsilon}(f) \log_2 \left(\frac{\sqrt{|X||Y|}}{\|A_g\|} \right) - O(1)$$

for any $0 < \epsilon < 1/4$.

Proof. We simply follow the proof of Lee and Zhang [LZ10] and use Lemma B.2 instead of Linial-Shraibman’s theorem. First, we note the following inequality which follows from the definition of γ_2 : For any $\delta \geq 0$ and m -by- n matrix A ,

$$\gamma_2^\delta(A) = \min_{B: \|B-A\|_\infty \leq \delta} \gamma_2(B) \geq \min_{B: \|B-A\|_\infty \leq \delta} \frac{\|B\|_{tr}}{\sqrt{\text{size}(B)}} = \frac{\|A\|_{tr}^\delta}{\sqrt{\text{size}(A)}}$$

where the first and last equalities are by definition of the approximate norm (see, e.g., [LZ10, Definition 4]) and the inequality is by the definition of γ_2 norm (see, e.g., [LZ10, Definition 1]). Using $A = A_{f \circ g}$ which is an $|X|$ -by- $|Y|$ matrix, we have

$$\gamma_2^\delta(A_{f \circ g}) \geq \frac{\|A_{f \circ g}\|_{tr}^\delta}{\sqrt{\text{size}(A_{f \circ g})}}. \quad (14)$$

The following claim is shown in the proof of Theorem 8 in [LZ10].

Claim B.5 ([LZ10]).

$$\frac{\|A_{f \circ g}\|_{tr}^\delta}{\sqrt{\text{size}(A_{f \circ g})}} \geq \delta \left(\frac{\sqrt{|X||Y|}}{\|A_g\|} \right)^{\deg_{2\delta}(f)}. \quad (15)$$

Proof. We note the following lemma (noted as Lemma 1 in [LZ10]) which shows that there exists a *dual polynomial* of f which is a polynomial v which certifies that the approximate polynomial degree of f is at least a certain value.

Lemma B.6 ([She11, SZ09]). *For any $f : \{0, 1\}^n \rightarrow \{0, 1\}$, let f' be such that $f'(z) = (-1)^{f(z)}$ and $d = \deg_\delta(f)$. Then, there exists a function $v : \{0, 1\}^n \rightarrow \mathbb{R}$ such that*

1. $\langle v, \chi_T \rangle = 0$ for every character χ_T with $|T| < d$.
2. $\|v\|_1 = 1$.
3. $\langle v, f' \rangle \geq \delta$.

Let v be a dual polynomial of f as in the above lemma. We will use $B = (\frac{2^n}{\text{size}(A_g)})A_{v \circ g}$ as a “witness matrix”, i.e.,

$$B[x, y] = \frac{2^n}{\text{size}(A_g)^n} v(g(x_1, y_1), \dots, g(x_n, y_n)). \quad (16)$$

It follows that

$$\langle A_{f \circ g}, B \rangle = \frac{2^n}{\text{size}(A_g)^n} \langle M_{f \circ g}, A_{v \circ g} \rangle \quad (17)$$

$$= \frac{2^n}{\text{size}(A_g)^n} \sum_{x, y} f(g(x^1, y^1), \dots, g(x^n, y^n)) v(g(x^1, y^1), \dots, g(x^n, y^n)) \quad (18)$$

$$= \frac{2^n}{\text{size}(A_g)^n} \sum_{z \in \{0, 1\}^n} \left(f(z) v(z) \left(\sum_{\substack{x, y: \\ g(x^i, y^i) = z_i \\ \forall 1 \leq i \leq n}} 1 \right) \right) \quad (19)$$

$$= \frac{2^n}{\text{size}(A_g)^n} \sum_{z \in \{0, 1\}^n} \left(f(z) v(z) \prod_{i=1}^n \left(\sum_{\substack{x^i, y^i: \\ g(x^i, y^i) = z_i}} 1 \right) \right) \quad (20)$$

$$= \frac{2^n}{\text{size}(A_g)^n} \sum_{z \in \{0, 1\}^n} \left(f(z) v(z) \left(\sum_{\substack{x', y': \\ g(x', y') = z_i}} 1 \right)^n \right) \quad (21)$$

$$= \sum_{z \in \{0, 1\}^n} f(z) v(z) \quad (22)$$

$$= \langle f, v \rangle \quad (23)$$

$$\geq \delta \quad (24)$$

where Eq.(22) is because g is strongly balanced which implies that g is balanced, i.e. $g(x^i, y^i)$ is 0 (and 1) for half of its possible inputs (i.e. $\text{size}(A_g)/2$ entries of A_g are 1 (and -1)); thus,

$$\sum_{\substack{x', y': \\ g(x', y') = z_i}} 1 = \text{size}(A_g)/2.$$

A similar argument and the fact that $\|v\|_1 = 1$ can be used to show that

$$\|B\|_1 = 1. \quad (25)$$

Now we turn to evaluate the spectral norm $\|B\|$. As shown in [LZ10], the strongly balanced property of g implies that the matrices $\chi_T \circ g^n$ and $\chi_S \circ g^n$ are orthogonal for distinct sets $S, T \subseteq \{0, 1\}^n$. Note the following fact (Fact 1 in [LZ10]): For any matrices A' and B' of the same dimension, if $A'(B')^\dagger = (A')^\dagger B' = 0$ then $\|A + B\| = \max\{\|A\|, \|B\|\}$. Using this fact, we have

$$\|B\| = \frac{2^n}{\text{size}(A_g)^n} \left\| \sum_{T \subseteq [n]} \hat{v}_T A_{\chi_T \circ g^n} \right\| \quad (26)$$

$$= \frac{2^n}{\text{size}(A_g)^n} \max_T |\hat{v}_T| \|\hat{v}_T A_{\chi_T \circ g^n}\| \quad (\text{by the fact above}) \quad (27)$$

$$= \max_T 2^n |\hat{v}_T| \prod_i \frac{\|A_G^{T[i]}\|}{\text{size}(A_g)} \quad (28)$$

$$\leq \max_{T: \hat{v}^T \neq 0} \prod_i \frac{\|A_G^{T[i]}\|}{\text{size}(A_g)} \quad (29)$$

$$= \left(\frac{\|A_g\|}{\sqrt{\text{size}(A_g)}} \right)^d \left(\frac{1}{\text{size}(A_g)} \right)^{n/2} \quad (30)$$

where Eq.(29) is because $|\hat{v}_T| \leq 1/2^n$ as $\|v\|_1 = 1$ and Eq.(30) is because $\|J\| = \sqrt{\text{size}(A_g)}$.

We note that for any $0 \leq \epsilon < 1$, norm $\Phi: \mathbb{R}^n \rightarrow \mathbb{R}$ and vector $v \in \mathbb{R}^n$, the approximate norm is $\Phi^\epsilon(v) = \max_u \frac{|\langle v, u \rangle| - \epsilon \|u\|_1}{\Phi^*(u)}$ (see, e.g., [LS09a] and [LZ10, Proposition 1]). Note also that if Φ is the trace norm then its dual Φ^* is the spectral norm (this is noted in [LZ10]). Thus,

$$\|A_{f \circ g^n}\|_{tr}^{\delta/2} = \max_{B'} \frac{|\langle A_{f \circ g^n}, B' \rangle| - (\delta/2) \|B'\|_1}{\|B'\|} \quad (31)$$

$$\geq \frac{|\langle A_{f \circ g^n}, B \rangle| - \delta/2}{\|B\|} \quad (\text{by Eq.(25)}) \quad (32)$$

$$\geq \frac{\delta - \delta/2}{\|B\|} \quad (\text{by Eq.(24)}) \quad (33)$$

$$\geq (\delta/2) \left(\frac{\sqrt{\text{size}(A_g)}}{\|A_g\|} \right)^d (\text{size}(A_g))^{n/2} \quad (\text{by Eq.(30)}) \quad (34)$$

$$\geq (\delta/2) \left(\frac{\sqrt{\text{size}(A_g)}}{\|A_g\|} \right)^d \left(\sqrt{\text{size}(A_{f \circ g})} \right) \quad (35)$$

where the last inequality is because $\text{size}(A_{f \circ g}) = \text{size}(A_g)^n$. This completes the proof of the claim. \square

The lemma follows immediately from Eq.(14) and Eq.(15) by plugging in Lemma B.2:

$$4^{2Q_{\epsilon}^{*,sv}(f \circ g^n)} \geq \gamma_2^{2\epsilon}(A_{f \circ g^n}) \geq \frac{\|A_{f \circ g}\|_{tr}^{2\epsilon}}{\sqrt{\text{size}(A_{f \circ g})}} \geq (2\epsilon) \left(\frac{\sqrt{|X||Y|}}{\|A_g\|} \right)^{\deg_{4\epsilon}(f)}.$$

Lemma B.4 follows (the term 2ϵ will contribute to the term “ $-O(1)$ ”). \square

Now, we prove the lower bound for $\text{IPmod}3_n$. Our proof essentially follows Sherstov’s proof [She11] (also see [LS09a, Section 7.2.3]). We can assume w.l.o.g. that n is divisible by 4. Consider the *promise version* of $\text{IPmod}3_n$ where any n -bit string input $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ has the property that for any $0 \leq i \leq (n/4) - 1$,

$$\begin{aligned} x_{4i+1}x_{4i+2}x_{4i+3}x_{4i+4} &\in \{0011, 0101, 1100, 1010\} \quad \text{and} \\ y_{4i+1}y_{4i+2}y_{4i+3}y_{4i+4} &\in \{0001, 0010, 1000, 0100\}. \end{aligned}$$

Now we show that the claimed lower bound holds even in this case. This lower bound clearly implies the lower bound for the more general case of $\text{IPmod}3_n$ where no restriction is put on the input.

Observe that, for any $(x, y) \in \mathcal{X} \times \mathcal{Y}$, the function $\text{IPmod}3$ can be written as

$$f \circ g^{n/4}(x, y) = f(g(x_1 \dots x_4, y_1 \dots y_4), g(x_5 \dots x_8, y_5 \dots y_8), \dots, g(x_{n-3} \dots x_n, y_{n-3} \dots y_n))$$

where

$$g(x_{4i+1} \dots x_{4i+4}, y_{4i+1} \dots y_{4i+4}) = (x_{4i+1} \wedge y_{4i+1}) \vee (x_{4i+2} \wedge y_{4i+2}) \vee (x_{4i+3} \wedge y_{4i+3}) \vee (x_{4i+4} \wedge y_{4i+4})$$

for all $0 \leq i \leq (n/4) - 1$, and $f(z_1, \dots, z_{n/4}) = 1$ if $z_1 + \dots + z_{n/4}$ can be divided by 3 and 0 otherwise. Note that $\text{IPmod}3(x, y) = f \circ g^{n/4}(x, y)$ since the promise implies that $g(x_{4i+1} \dots x_{4i+4}, y_{4i+1} \dots y_{4i+4}) = 1$ if and only if $x_{4i+1}y_{4i+1} + \dots + x_{4i+4}y_{4i+4} = 1$. The matrix A_g is

$$A_g = \begin{matrix} & \begin{matrix} 0001 & 0010 & 1000 & 0100 \end{matrix} \\ \begin{matrix} 0011 \\ 0101 \\ 1100 \\ 1010 \end{matrix} & \begin{pmatrix} -1 & -1 & 1 & 1 \\ -1 & 1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix} \end{matrix}$$

which is clearly strongly balanced. It can be checked that this matrix has spectral norm $\|A_g\| = 2\sqrt{2}$ (see, e.g., [LS09a, Section 7.2.3]). Moreover, by Paturi [Pat92] (see also [dW10] and [She11, Theorem 2.6]), $\deg_{1/3}(f) = \Theta(n)$. Thus, Lemma B.4 implies that

$$\begin{aligned} Q_{1/12}^{*,sv}(f \circ g^n) &\geq \deg_{1/3}(f) \log_2 \left(\frac{\sqrt{4 \times 4}}{\|A_g\|} \right) - O(1) \\ &= \deg_{1/3}(f) \log_2 \sqrt{2} - O(1) \\ &= \Omega(n). \end{aligned}$$

We note that the same technique can be used to prove many bounds in the server model similar to bounds in [Raz03, She11, LZ10].




Figure 12: The resulted graph G in three situations depending on the value of $\sum_{1 \leq i \leq n} x_i \cdot y_i \bmod 3$. Dashed lines (in red) represent paths connecting v_0^0, \dots, v_0^2 and v_n^0, \dots, v_n^2 . Thick lines (in gray) show the fact that we identify nodes on two sides, i.e. $v_0^j = v_n^j$ for all $0 \leq j \leq 2$. Our main observation is that G is a Hamiltonian cycle if and only if $\sum_{1 \leq i \leq n} x_i \cdot y_i \bmod 3 \neq 0$ (cf. Lemma C.3).

C Detail of Section 7

First, let us recall that Alice and Bob construct a gadget G_i using x_i and y_i as shown in Fig. 4. Fig. 5 shows how G_i looks like for each possible value of x_i and y_i . It follows immediately that G_i always consist of three paths which connect v_{i-1}^j to $v_i^{(j+x \cdot y) \bmod 3}$, as in the following observation.

Observation C.1 (Observation 7.1 restated). *For any value of (x_i, y_i) , G_i consists of three paths where v_{i-1}^j is connected by a path to $v_i^{(j+x \cdot y) \bmod 3}$, for any $0 \leq j \leq 2$. Moreover, Alice's (respectively Bob's) edges, i.e. thin (red) lines (respectively thick (blue) lines) in Fig. 4, form a matching that covers all nodes except v_i^j (respectively v_{i-1}^j) for all $0 \leq j \leq 2$.*

Finally, we connect gadgets G_i and G_{i+1} together by identifying rightmost nodes of G_i with leftmost nodes of G_{i+1} , as shown in Fig. 6 (gray lines represent the fact that we identify rightmost nodes of G_n to leftmost nodes of G_1).

Lemma C.2 (Lemma 7.2 restated). *G consists of three paths P^0, P^1 and P^2 where for any $0 \leq j \leq 2$, P^j has v_0^j as one end vertex and $v_n^{(j+\sum_{1 \leq i \leq n} x_i \cdot y_i) \bmod 3}$ as the other.*

Proof. We will show that for any $2 \leq k \leq n$ and $0 \leq j \leq 2$, P^j has v_0^j as one end vertex and $v_k^{(j+\sum_{1 \leq i \leq k} x_i \cdot y_i) \bmod 3}$ as the other. We prove this by induction on k . Our claim clearly holds for $k = 2$ by Observation C.1. Now assume that this claim is true for any $2 \leq k \leq n - 1$, i.e., v_0^j is connected by a path to $v_k^{j'}$ where $j' = (j + \sum_{1 \leq i \leq k} x_i \cdot y_i) \bmod 3$. By Observation C.1, $v_k^{j'}$ is connected by a path to $v^{j''}$ where $j'' = (j' + x_{k+1} \cdot y_{k+1}) \bmod 3 = (j + \sum_{1 \leq i \leq k+1} x_i \cdot y_i) \bmod 3$ as claimed. \square

Lemma C.3. *Each player's edges form a perfect matching in G . Moreover, G is a Hamiltonian cycle if and only if $\sum_{1 \leq i \leq n} x_i \cdot y_i \bmod 3 \neq 0$.*

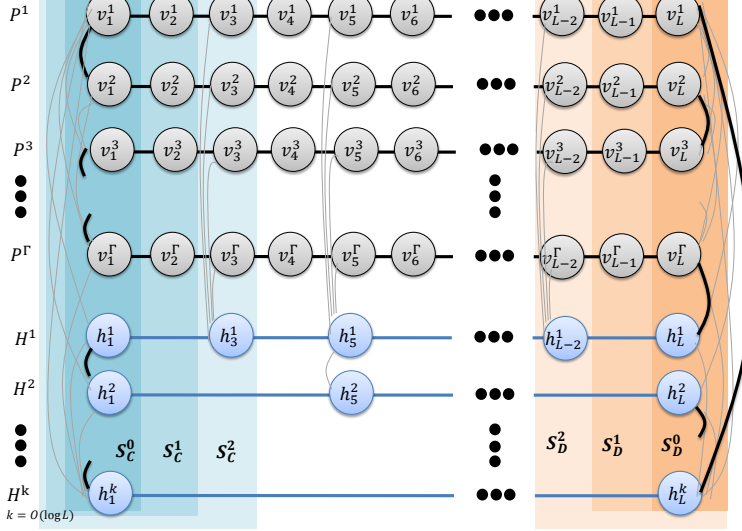


Figure 13: (Fig. 10 reproduced) The network N which consists of network N' and some “highways” which are paths with nodes h_j^i (i.e., nodes in blue). Bold edges show an example of subnetwork M when the input perfect matchings are $E_C = \{(u_1, u_2), (u_3, u_4), \dots, (u_{\Gamma+k-1}, u_{\Gamma+k})\}$ and $E_D = \{(u_2, u_3), (u_4, u_5), \dots, (u_{\Gamma+k}, u_1)\}$. Pale edges are those in N but not in M .

Proof. We consider three different values of $z = \sum_{1 \leq i \leq n} x_i \cdot y_i \pmod 3$ as shown in Fig. 12. If $z = 0$ then Lemma C.2 implies that v_0^j will be connected to v_n^j by a path, for all j . After we identify v_0^j with v_n^j we will have three distinct cycles, each containing a distinct $v_0^j = v_n^j$. If $z = 1$ then Lemma C.2 implies that v_0^j will be connected to $v_n^{(j+1) \pmod 3}$ by a path. After we identify v_0^j with v_n^j we will have one cycle that connects $v_0^0 = v_n^0$ to $v_n^1 = v_0^1$ then to $v_n^2 = v_0^2$. Similarly, if $z = 1$ then Lemma C.2 implies that v_0^j will be connected to $v_n^{(j+2) \pmod 3}$ by a path. After we identify v_0^j with v_n^j we will have one cycle that connects $v_0^0 = v_n^0$ to $v_n^2 = v_0^2$ then to $v_n^1 = v_0^1$. \square

D Detail of Section 8

Theorem D.1 (Theorem 3.5 restated). *For any $B, L, \Gamma \geq \log L, \beta \geq 0$ and $\epsilon_0, \epsilon_1 > 0$, there exists a B -model quantum network N of diameter $\Theta(\log L)$ and $\Theta(\Gamma L)$ nodes such that*

if $Q_{\epsilon_0, \epsilon_1}^{*, N}(\mathbf{P}(N)) \leq \frac{L}{2} - 2$ then $Q_{\epsilon_0, \epsilon_1}^{*, sv}(\mathbf{P}_\Gamma) = O((B \log L) Q_{\epsilon_0, \epsilon_1}^{*, N}(\mathbf{P}(N)))$

where \mathbf{P} can be replaced by \mathbf{Ham} and $(\beta\Gamma)$ -Conn.

D.1 Description of the network N

In this section we describe the network N as shown in Fig. 13. We assume that $L = 2^i + 1$ for some i . This can be assumed without changing the theorem statement by simply increasing L to the nearest number of this form.

The two basic units in the construction are *paths* and *highways*. There are Γ paths, denoted by $P^1, P^2, \dots, P^\Gamma$, each having L nodes, i.e., for $j = 1, 2, \dots, \Gamma$,

$$V(P^i) = \{v_1^i, \dots, v_L^i\} \quad \text{and} \quad E(P^i) = \{(v_j^i, v_{j+1}^i) \mid 1 \leq j \leq L-1\}.$$

We construct $k = \log_2(L-1)$ highways, denoted by H^1, \dots, H^k where H^i has the following nodes and edges.

$$V(H^i) = \{h_{1+j2^i}^i \mid 0 \leq j \leq \frac{L-1}{2^i}\} \quad \text{and} \\ E(H^i) = \{(h_{1+j2^i}^i, h_{1+(j+1)2^i}^i) \mid 0 \leq j \leq \frac{L-1}{2^i}\}.$$

For any node h_j^1 we add an edge (h_j^1, v_j^1) for any j . Moreover for any node h_j^i we add an edge (h_j^{i-1}, h_j^i) . Figure 13 depicts this network. We note the following simple observation.

Observation D.2. *The number of nodes in N is $n = \Theta(L\Gamma)$ and its diameter is $\Theta(\log L)$.*

D.2 Simulation

For any $0 \leq t \leq (L/2) - 2$, we partition $V(N)$ into three sets, denoted by S_C^t , S_D^t and S_S^t , as follows (also see Fig. 13).

$$S_C^t = \{v_j^i, h_j^i \mid 1 \leq i \leq \Gamma, 1 \leq j \leq t+1\}, \quad (36)$$

$$S_D^t = \{v_j^i, h_j^i \mid 1 \leq i \leq \Gamma, L-t \leq j \leq L\}, \quad (37)$$

$$S_S^t = V(N) \setminus (S_C^t \cup S_D^t). \quad (38)$$

Let \mathcal{A} be any quantum distributed algorithm on network N for computing a problem P (which is either **Ham** or $(\beta\Gamma) - \text{Conn}$). Let $T_{\mathcal{A}}$ be the worst case running time of algorithm \mathcal{A} (over all inputs). We note that $T_{\mathcal{A}} \leq (L/2) - 2$, as assumed in the theorem statement. We show that Carol, David and the server can solve problem P on $(\Gamma+k)$ -node input graph using small communication, essentially by “simulating” \mathcal{A} on some input subnetwork M corresponding to $G = (U, E_C \cup E_D)$ in the following sense. When receiving E_C and E_D , the three parties will construct a subnetwork M of N (without communication) in such a way that M is a 1-input of problem P (e.g., M is a Hamiltonian cycle) if and only if $G = (U, E_C \cup E_D)$ is. Next, they will simulate algorithm \mathcal{A} in such a way that, at any time t and for each node v_j^i in N , there will be exactly one party among Carol, David and the server that knows *all information that v_j^i should know in order to run algorithm \mathcal{A}* , i.e., the (quantum) state of v_j^i as well as the messages (each consisting of B quantum bits) sent to v_j^i from its neighbors at time t . The party that knows this information will pretend to be v_j^i and apply algorithm \mathcal{A} to get the state of v_j^i at time $t+1$ as well as the messages that v_j^i will send to its neighbors at time $t+1$. We say that this party *owns* v_j^i at time t . Details are as follows.

We will define a server-model protocol \mathcal{A}' that guarantees that, at any time t , Carol, David and the server will own nodes in sets S_C^t , S_D^t and S_S^t , respectively, at time t . That is, Carol’s workspace, denoted by $H_{C,C}$, contains all qubits in $H_{vv'}$, for any $v \in S_C^t$ and $v' \in V(N)$, resulting from t rounds of an execution of \mathcal{A} . Similarly, David’s (respectively the server’s) workspace, denoted by $H_{D,D}$

(respectively $H_{S,S}$), contains all qubits in $H_{vv'}$ for any $v \in S_D^t$ (respectively $v \in S_S^t$) and $v' \in V(N)$ resulting from t rounds of \mathcal{A} . In other words, if after t rounds of \mathcal{A} network N has state

$$|\psi_M^t\rangle = \sum_w \left(\alpha_w \bigotimes_{v,v' \in V(N)} |\psi_{w,M}^t(v,v')\rangle \right),$$

then we will make sure that the server model has state

$$|\Psi_G^t\rangle = \sum_w \left(\alpha_w \bigotimes_{i,j \in \{C,D,S\}} |\Psi_{w,G}^t(i,j)\rangle \right)$$

where $|\Psi_{w,G}^t(i,j)\rangle = |0\rangle$, for any $i,j \in \{C,D,S\}$ such that $i \neq j$, and for any $i \in \{C,D,S\}$

$$|\Psi_{w,G}^t(i,i)\rangle = \bigotimes_{v \in S_i^t, v' \in V(N)} |\psi_{w,M}^t(v',v)\rangle. \quad (39)$$

Let $\Gamma' = \Gamma + k$. Fix any Γ' -node input graph $G = (U, E_C \cup E_D)$ of problem P where E_C and E_D are edges given to Carol and David respectively. Let $U = \{u_1, \dots, u_{\Gamma'}\}$. For convenience, for any $1 \leq j \leq k$, let $v_1^{\Gamma+j} = h_1^j$ and $v_L^{\Gamma+j} = h_L^j$. We construct a subnetwork M of N as follows. For any $i \neq j$, we mark $v_1^i v_1^j$ as participating in M if and only if $u_i u_j \in E_C$. Note that this knowledge must be kept in qubits in $H_{v_1^i v_1^j}$ and $H_{v_L^i v_L^j}$ in network N as we require each node to know whether edges incident to it are in M or not. This means that this knowledge must be stored in $H_{C,C}$ since $v_1^i, v_1^j \in S_C^0$. This can be guaranteed without any communication since Carol knows E_C . Similarly, we mark $v_L^i v_L^j$ as participating in M if and only if $u_i u_j \in E_D$, and this information can be stored in $H_{D,D}$ without communication. Finally, we let all edges in all paths and highways be in M . This information is stored in $H_{S,S}$. An example of network M is shown in Fig. 13. To conclude, if the initial state of N with this subnetwork M is

$$|\psi_M^0\rangle = \sum_w \left(\alpha_w \bigotimes_{v,v' \in V(N)} |\psi_{w,M}^0(v,v')\rangle \right).$$

then the server model will start with state $|\Psi_G^0\rangle = \sum_w \left(\alpha_w \bigotimes_{i,j \in \{C,D,S\}} |\Psi_{w,G}^0(i,j)\rangle \right)$ where $|\Psi_{w,G}^0(i,j)\rangle = |0\rangle$, for any $i,j \in \{C,D,S\}$ such that $i \neq j$, and for any $i \in \{C,D,S\}$

$$|\Psi_{w,G}^0(i,i)\rangle = \bigotimes_{v \in S_i^0, v' \in V(N)} |\psi_{w,M}^0(v',v)\rangle.$$

Thus Eq.(39) holds for $t = 0$. We note the following simple observation.

Observation D.3. $G = (U, E_C \cup E_D)$ is a Hamiltonian cycle if and only if M is a Hamiltonian cycle. G is connected if and only if M is connected, and for any δ , G is δ -far from being connected if and only if M is δ -far from being connected.

Thus, Carol, David and the server can check whether G is a Hamiltonian cycle if they can check whether M is a Hamiltonian cycle. Similarly, they can check if G is connected or $(\beta\Gamma)$ -far from

being connected by checking M . So, if Eq.(39) can be maintained until \mathcal{A} terminates then we are done since each server-model player can pretend to be one of the nodes they own and measure the workspace of such node to get the property of M .

Now suppose that Carol, David and the server have maintained this guarantee until they have executed \mathcal{A} for $t-1$ steps, i.e., player i owns the nodes in S_i^{t-1} at time $t-1$. They maintain the guarantee at step t as follows. First, each player simulate the internal computation of \mathcal{A} on nodes they own. That is, for each node $v \in V(N)$, the player i such that $v \in S_i^{t-1}$ applies the transformation $C_{t,v}$ (cf. Section A.1) on qubits in workspace $\bigotimes_{v' \in V(N)} H_{v'v}$ which is maintained in $H_{i,i}$ at time $t-1$. This means that if after the internal computation N has state $|v_M^t\rangle = \sum_w \left(\alpha_w \bigotimes_{v,v' \in V(N)} |v_{w,M}^t(v,v')\rangle \right)$ then the server model will have state $|\Upsilon_G^t\rangle = \sum_w \left(\alpha_w \bigotimes_{i,j \in \{C,D,S\}} |\Upsilon_{w,G}^t(i,j)\rangle \right)$ where $|\Upsilon_w^t(i,j)\rangle = |0\rangle$, for any $i \neq j$, and $|\Upsilon_{w,G}^t(i,i)\rangle = \bigotimes_{v \in S_i^{t-1}, v' \in V(N)} |v_{w,M}^t(v',v)\rangle$ for any i . Note that the server model players can simulate the internal computation of \mathcal{A} without any communication since a player that owns node v has all information needed to simulate an internal computation of v (i.e., the state of v as well as all messages v received at time $t-1$).

At this point, for any $i \in \{C,D,S\}$, player i 's space contains the current state and out-going messages of every node $v \in S_i^{t-1}$. They will need to receive some information in order to guarantee that they own nodes in S_i^t . First, consider Carol. Let S'_C be the set of rightmost nodes in the set S_C^{t-1} , i.e. S'_C consists of v_{t+1}^i and h_j^i for all i and $j = \arg \max_j \{h_j^i \in S_C^{t-1}\}$.

Note that Carol already has the workspace and all incoming messages of nodes in $S_C^{t-1} \setminus S'_C$ at time t . This is because for any $v \in S_C^{t-1} \setminus S'_C$, Carol already has qubits in $H_{v'v}$ for all $v' \in V(N)$. For each $v \in S'_C$, Carol is missing the messages sent from v 's right neighbor; i.e., Carol does not have qubits in $H_{v_{t+2}^i v_{t+1}^i}$ and $H_{h_{j'}^i h_j^i}$ for all $i, j = \arg \max_j \{h_j^i \in S_C^{t-1}\}$ and $j' = \arg \min_{j'} \{h_{j'}^i \notin S_C^{t-1}\}$. Since $S'_C \subseteq S_C^t$, we need to make sure that Carol has all information of nodes in S'_C at time t .

For a non-highway node v_{t+1}^i , for all i , this can be done by letting the server who owns v_{t+2}^i send to Carol a message sent from v_{t+2}^i to v_{t+1}^i at time t , i.e., qubits in $H_{v_{t+2}^i v_{t+1}^i}$. For highway node h_j^i for all i and $j = \arg \max_j \{h_j^i \in S_C^{t-1}\}$, its right neighbor $h_{j'}^i$, where $j' = \arg \min_{j'} \{h_{j'}^i \notin S_C^{t-1}\}$, might be owned by David or the server. In any case, we let the owner of $h_{j'}^i$ send to Carol the message sent from $h_{j'}^i$ to h_j^i at time t , i.e., qubits in $H_{h_{j'}^i h_j^i}$. The cost of doing this is zero if $h_{j'}^i$ belongs to the server and at most B if $h_{j'}^i$ belongs to David since the message size is at most B . In any case, the total cost will be at most Bk since there are k highways. We can thus make sure that Carol gets the information of nodes in S_C^{t-1} at time t at the total cost of at most Bk .

In addition to this, Carol needs to get information of nodes in $S_C^t \setminus S_C^{t-1}$ at time t . This means that, for any $v \in S_C^t \setminus S_C^{t-1}$ she has to receive the qubits stored in $H_{v'v}$ for all $v' \in V(N)$. For any non-highway node $v_{t+2}^i \in S_C^t \setminus S_C^{t-1}$, it can be checked from the definition that v_{t+2}^i and all its neighbors are in $S_C^{t-1} \cup S_S^{t-1}$. So, we can make sure that Carol owns v_{t+2}^i by letting the server send to Carol the workspace of v_{t+2}^i and messages sent to v_{t+2}^i by its neighbors in S_S^{t-1} (i.e. qubits in $H_{v'v_{t+2}^i}$ for all $v' \in S_S^{t-1}$). This communication is again free. For a highway node h_j^i in $S_C^t \setminus S_C^{t-1}$, it can be checked from the definition that h_j^i as well as all its *non-highway* neighbors are in $S_C^{t-1} \cup S_S^{t-1}$. The only neighbor of h_j^i that might be in S_D^{t-1} is its right neighbor, say $h_{j'}^i$, in the highway. If $h_{j'}^i$ is in S_D^{t-1} then David has to send to Carol the message sent from $h_{j'}^i$ to h_j^i . This has cost at most B . So, Carol can obtain the workspace of h_j^i as well as all messages sent to h_j^i at the cost of B . Since there are k highway nodes in $S_C^t \setminus S_C^{t-1}$, the total cost for Carol to obtain information needed

to maintain nodes in $S_C^t \setminus S_C^{t-1}$ is Bk . We conclude that Carol can obtain all information needed to own nodes in S_C^t at time t at the cost of $2Bk$.

We can do the same thing to guarantee that David owns all nodes in S_D^t at time t at the cost of $2Bk$. Now we make sure that the server owns nodes in S_S^t . First, observe that the server already has the workspace of all nodes in S_S^t since $S_S^t \subseteq S_S^{t-1}$. Moreover, the server already has all messages sent to all non-highway nodes in S_S^t (i.e. v_j^i for all $t+2 \leq j \leq L-t-1$ and $1 \leq i \leq \Gamma$) since all of their neighbors are in S_S^{t-1} . Additionally, each leftmost highway node $h_j^i \in S_S^t$, for any i and $j = \arg \min_j \{h_j^i \in S_S^t\}$, has at most one neighbor in S_C^{t-1} (i.e., its right neighbor in the highway). Similarly, each rightmost highway node $h_j^i \in S_S^t$, for any i and $j' = \arg \max_{j'} \{h_{j'}^i \in S_S^t\}$, has at most one neighbor in S_D^{t-1} (i.e., its right neighbor in the highway). Thus, the server needs to obtain from Carol and David at most $2B$ qubits to maintain h_j^i and $h_{j'}^i$. Since there are k highways, the server needs at most $2kB$ qubits total from Carol and David. We thus conclude that the players can maintain Eq.(39) at the cost of $6kB = O(B \log L)$ qubits per round as desired.

As noted earlier, the server-model players will simulate \mathcal{A} until \mathcal{A} terminates. Then they can measure the workspace of nodes they own to check whether M is a 0- or 1-input of problem P . Observation D.3 implies that they can use this answer to answer whether G is a 0- or 1-input with the same error probability. Since each round of simulation requires a communication complexity of $O(B \log L)$ and the simulation is done for $T_A \leq Q_{\epsilon_0, \epsilon_1}^{*,N}(P(N))$ rounds, the total communication complexity is $O((B \log L)Q_{\epsilon_0, \epsilon_1}^{*,N}(P(N)))$ as claimed.

References

- [AA05] Scott Aaronson and Andris Ambainis. Quantum search of spatial regions. *Theory of Computing*, 1(1):47–79, 2005. Also in FOCS’03. 2
- [BCW98] Harry Buhrman, Richard Cleve, and Avi Wigderson. Quantum vs. classical communication and computation. In *STOC*, pages 63–68, 1998. 12
- [Bel64] John S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964. 4
- [BFS86] László Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *FOCS*, pages 337–347, 1986. 2, 12
- [BOH05] Michael Ben-Or and Avinatan Hassidim. Fast quantum byzantine agreement. In *STOC*, pages 481–485, 2005. 12
- [BR03] Harry Buhrman and Hein Röhrig. Distributed quantum computing. In *MFCS*, pages 1–20, 2003. 12
- [Bri11] Jop Briët. *Grothendieck Inequalities, Nonlocal Games and Optimization*. PhD thesis, Universiteit van Amsterdam, 2011. 31
- [BT08] A. Broadbent and A. Tapp. Can quantum mechanics help distributed computing? *SIGACT News*, 39(3):67–76, 2008. 1
- [BvDHT99] H. Buhrman, W. van Dam, P. Hoyer, and A. Tapp. Quantum multiparty communication complexity. *Physical Review A*, 60:2737–2741, 1999. 12
- [BYJKS04] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68(4):702–732, 2004. Also in FOCS’02. 2
- [CB97] Richard Cleve and Harry Buhrman. Substituting quantum entanglement for communication. *Physical Review A*, 56(2):1201–1204, 1997. 12

- [CHGK14] Keren Censor-Hillel, Mohsen Ghaffari, and Fabian Kuhn. Distributed connectivity decomposition. In *PODC*, 2014. [13](#)
- [CKS10] Bogdan S. Chlebus, Dariusz R. Kowalski, and Michal Strojnowski. Scalable quantum consensus for crash failures. In *DISC*, pages 236–250, 2010. [12](#)
- [Coh97] Edith Cohen. Size-Estimation Framework with Applications to Transitive Closure and Reachability. *J. Comput. Syst. Sci.*, 55(3):441–453, 1997. Also in FOCS’94. [29](#)
- [DGP07] Devdatt P. Dubhashi, Fabrizio Grandioni, and Alessandro Panconesi. Distributed Algorithms via LP Duality and Randomization. In *Handbook of Approximation Algorithms and Metaheuristics*. Chapman and Hall/CRC, 2007. [1](#)
- [DHK⁺12] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. Distributed verification and hardness of distributed approximation. *SIAM J. Comput.*, 41(5):1235–1265, 2012. [1](#), [2](#), [3](#), [5](#), [9](#), [10](#), [11](#), [13](#), [18](#), [19](#), [28](#), [30](#)
- [DMP13] Atish Das Sarma, Anisur Rahaman Molla, and Gopal Pandurangan. Distributed computation of sparse cuts. *CoRR*, abs/1310.5407, 2013. [13](#)
- [DNPT13] Atish Das Sarma, Danupon Nanongkai, Gopal Pandurangan, and Prasad Tetali. Distributed random walks. *J. ACM*, 60(1):2, 2013. [13](#)
- [DP08] Vasil S. Denchev and Gopal Pandurangan. Distributed quantum computing: a new frontier in distributed systems or science fiction? *SIGACT News*, 39(3):77–95, 2008. [1](#), [3](#), [12](#)
- [dW02] Ronald de Wolf. Quantum communication and complexity. *Theoretical Computer Science*, 287(1):337–352, 2002. [12](#)
- [dW10] Ronald de Wolf. A note on quantum algorithms and the minimal degree of ϵ -error polynomials for symmetric functions. *Quantum Information & Computation*, 8(10):943–950, 2010. [38](#)
- [Elk04] Michael Elkin. Distributed approximation: a survey. *SIGACT News*, 35(4):40–57, 2004. [1](#), [30](#)
- [Elk05] Michael Elkin. Computing almost shortest paths. *ACM Transactions on Algorithms*, 1(2):283–323, 2005. Also in PODC’01. [30](#)
- [Elk06] Michael Elkin. An Unconditional Lower Bound on the Time-Approximation Trade-off for the Distributed Minimum Spanning Tree Problem. *SIAM J. Comput.*, 36(2):433–456, 2006. Also in STOC’04. [2](#), [5](#), [11](#), [19](#), [30](#)
- [EPR35] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47(10):777–780, May 1935. [4](#)
- [FHW12] Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. Networks cannot compute their diameter in sublinear time. In *SODA*, pages 1150–1162, 2012. [13](#)
- [GBK⁺08] S. Gaertner, M. Bourennane, C. Kurtsiefer, A. Cabello, and H. Weinfurter. Experimental demonstration of a quantum protocol for byzantine agreement and liar detection. *PHYS.REV.LETT.*, 100:070504, 2008. [12](#)
- [Gha14] Mohsen Ghaffari. Near-optimal distributed approximation of minimum-weight connected dominating set. In *ICALP (2)*, 2014. [13](#)
- [GK13] Mohsen Ghaffari and Fabian Kuhn. Distributed minimum cut approximation. In *DISC*, pages 1–15, 2013. [2](#), [5](#)
- [GKM09] C. Gavaille, A. Kosowski, and M. Markiewicz. What can be observed locally? In *DISC*, pages 243–257, 2009. [1](#), [2](#), [4](#)

- [GKP98] J. Garay, S. Kutten, and D. Peleg. A sublinear time distributed algorithm for minimum-weight spanning trees. *SIAM J. on Computing*, 27:302–316, 1998. Also in FOCS’93. [11](#)
- [Hol73] A. S. Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii*, 9(3):3–11, 1973. English translation in *Problems of Information Transmission*, 9:177–183, 1973. [1](#)
- [HW12] Stephan Holzer and Roger Wattenhofer. Optimal distributed all pairs shortest paths and applications. In *PODC*, pages 355–364, 2012. [13](#)
- [IKL⁺12] Gábor Ivanyos, Hartmut Klauck, Troy Lee, Miklos Santha, and Ronald de Wolf. New bounds on the classical and quantum communication complexity of some graph properties. In *FSTTCS*, pages 148–159, 2012. [12](#)
- [KdW12] H. Klauck and R. de Wolf. Fooling one-sided quantum protocols. *Manuscript*, 2012. [8](#), [12](#), [14](#), [16](#)
- [KKM⁺08] Maleq Khan, Fabian Kuhn, Dahlia Malkhi, Gopal Pandurangan, and Kunal Talwar. Efficient distributed approximation algorithms via probabilistic tree embeddings. In *PODC*, pages 263–272, 2008. [1](#), [29](#), [30](#)
- [KKP11] Liah Kor, Amos Korman, and David Peleg. Tight bounds for distributed mst verification. In *STACS*, pages 69–80, 2011. [2](#), [5](#), [10](#), [19](#)
- [KMT09] H. Kobayashi, K. Matsumoto, and S. Tani. Ba: Exactly electing a unique leader is not harder than computing symmetric functions on anonymous quantum networks. In *PODC*, pages 334–335, 2009. [12](#)
- [KMT10] H. Kobayashi, K. Matsumoto, and S. Tani. Computing on anonymous quantum network. *CoRR*, abs/1001.5307, 2010. [12](#)
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, New York, NY, USA, 1997. [5](#), [7](#)
- [KP98] Shay Kutten and David Peleg. Fast Distributed Construction of Small k -Dominating Sets and Applications. *J. Algorithms*, 28(1):40–66, 1998. Also in PODC’95. [2](#), [5](#), [11](#)
- [KS92] Bala Kalyanasundaram and Georg Schnitger. The Probabilistic Communication Complexity of Set Intersection. *SIAM J. Discrete Math.*, 5(4):545–557, 1992. [2](#)
- [LPS13] Christoph Lenzen and Boaz Patt-Shamir. Fast routing table construction using small messages: extended abstract. In *STOC*, pages 381–390, 2013. [2](#), [5](#), [13](#)
- [LPSP06] Zvi Lotker, Boaz Patt-Shamir, and David Peleg. Distributed MST for constant diameter graphs. *Distributed Computing*, 18(6):453–460, 2006. Also in PODC’01. [2](#), [5](#), [19](#)
- [LS09a] T. Lee and A. Shraibman. Lower bounds in communication complexity. *Foundations and Trends in Theoretical Computer Science*, 3(4):263–398, 2009. [8](#), [14](#), [31](#), [33](#), [34](#), [35](#), [37](#), [38](#)
- [LS09b] N. Linial and A. Shraibman. Lower bounds in communication complexity based on factorization norms. *Random Struct. Algorithms*, 34(3):368–394, 2009. Also in STOC’07. [8](#), [15](#), [33](#)
- [Lub86a] M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1053, 1986. [1](#)
- [Lub86b] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1053, 1986. Also in STOC’85. [1](#)
- [LZ10] T. Lee and S. Zhang. Composition theorems in communication complexity. In *ICALP*

- (1), pages 475–489, 2010. 8, 14, 15, 35, 36, 37, 38
- [Nan14a] Danupon Nanongkai. Brief announcement: Almost-tight approximation distributed algorithm for minimum cut. In *PODC*, 2014. 2, 5
- [Nan14b] Danupon Nanongkai. Distributed Approximation Algorithms for Weighted Shortest Paths. In *STOC*, 2014. 2, 5
- [Nay99] A. Nayak. Optimal lower bounds for quantum automata and random access codes. In *FOCS*, pages 369–377, 1999. 1
- [NC04] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information (Cambridge Series on Information and the Natural Sciences)*. Cambridge University Press, 1 edition, January 2004. 1, 4, 14, 25, 32
- [NDP11] Danupon Nanongkai, Atish Das Sarma, and Gopal Pandurangan. A tight unconditional lower bound on distributed randomwalk computation. In *PODC*, pages 257–266, 2011. 13
- [NS14] Danupon Nanongkai and Hsin-Hao Su. Almost-tight distributed minimum cut algorithms. Manuscript, 2014. 2
- [Pat92] R. Paturi. On the degree of polynomials that approximate symmetric boolean functions (preliminary version). In *STOC*, pages 468–474, 1992. 38
- [Pel00] David Peleg. *Distributed computing: a locality-sensitive approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000. 1, 3, 30
- [PR00] David Peleg and Vitaly Rubinfeld. A Near-Tight Lower Bound on the Time Complexity of Distributed Minimum-Weight Spanning Tree Construction. *SIAM J. Comput.*, 30(5):1427–1442, 2000. Also in FOCS’99. 2, 5, 19, 30
- [PSK03] S. P. Pal, S. K. Singh, and S. Kumar. Multi-partite quantum entanglement versus randomization: Fair and unbiased leader election in networks, 2003. 12
- [Raz92] Alexander A. Razborov. On the Distributional Complexity of Disjointness. *Theor. Comput. Sci.*, 106(2):385–390, 1992. Also in ICALP’90. 2
- [Raz99] Ran Raz. Exponential separation of quantum and classical communication complexity. In *STOC*, pages 358–369, 1999. 12
- [Raz03] Alexander A. Razborov. Quantum communication complexity of symmetric predicates. *Izvestiya: Mathematics*, 67(1):145, 2003. 8, 15, 38
- [RS95] R. Raz and B. Spieker. On the “log rank”-conjecture in communication complexity. *Combinatorica*, 15(4):567–588, 1995. Also in FOCS’93. 10, 12
- [She11] Alexander A. Sherstov. The pattern matrix method. *SIAM J. Comput.*, 40(6):1969–2000, 2011. Also in STOC’08. 8, 15, 36, 38
- [Su14] Hsin-Hao Su. Brief announcement: A distributed minimum cut approximation scheme. In *SPAA*, 2014. 2, 5
- [Suoar] Jukka Suomela. Survey of local algorithms. *ACM Computing Surveys*, to appear. 1
- [SZ09] Y. Shi and Y. Zhu. Quantum communication complexity of block-composed functions. *Quantum Info. Comput.*, 9(5):444–460, May 2009. 36
- [TKM05] S. Tani, H. Kobayashi, and K. Matsumoto. Exact quantum algorithms for the leader election problem. In *STACS*, pages 581–592, 2005. 12
- [TS99] A. Ta-Shma. Classical versus quantum communication complexity. *SIGACT News*, 30(3):25–34, 1999. 12
- [Tsi87] B. Tsirelson. Quantum analogues of the bell inequalities: the case of two spatially separated domains. *Journal of Soviet Mathematics*, 36:557–570, 1987. 34

- [Wat11] J. Watrous. Guest column: an introduction to quantum information and quantum circuits 1. *SIGACT News*, 42(2):52–67, 2011. [25](#)