

Quantum Distributed Algorithms for Detection of Cliques

Keren Censor-Hillel ✉ 

Technion, Haifa, Israel

Orr Fischer ✉

Tel-Aviv University, Israel

François Le Gall ✉

Nagoya University, Aichi, Japan

Dean Leitersdorf ✉

Technion, Haifa, Israel

Rotem Oshman ✉

Tel-Aviv University, Israel

Abstract

The possibilities offered by quantum computing have drawn attention in the distributed computing community recently, with several breakthrough results showing quantum distributed algorithms that run faster than the fastest known classical counterparts, and even separations between the two models. A prime example is the result by Izumi, Le Gall, and Magniez [STACS 2020], who showed that triangle detection by quantum distributed algorithms is easier than triangle listing, while an analogous result is not known in the classical case.

In this paper we present a framework for fast quantum distributed clique detection. This improves upon the state-of-the-art for the triangle case, and is also more general, applying to larger clique sizes.

Our main technical contribution is a new approach for detecting cliques by encapsulating this as a search task for nodes that can be added to smaller cliques. To extract the best complexities out of our approach, we develop a framework for *nested* distributed quantum searches, which employ checking procedures that are quantum themselves.

Moreover, we show a circuit-complexity barrier on proving a lower bound of the form $\Omega(n^{3/5+\epsilon})$ for K_p -detection for any $p \geq 4$, even in the classical (non-quantum) distributed CONGEST setting.

2012 ACM Subject Classification Networks → Network algorithms; Theory of computation → Distributed algorithms

Keywords and phrases distributed graph algorithms, quantum algorithms, cycles, cliques, Congested Clique, CONGEST

Digital Object Identifier 10.4230/LIPIcs.ITCS.2022.35

Funding This project was partially supported by the European Union’s Horizon 2020 Research and Innovation Programme under grant agreement no. 755839. François Le Gall was supported by JSPS KAKENHI grants Nos. JP19H04066, JP20H05966, JP20H00579, JP20H04139, JP21H04879 and MEXT Quantum Leap Flagship Program (MEXT Q-LEAP) grants No. JPMXS0118067394 and JPMXS0120319794.

1 Introduction

Quantum distributed computing. The power of quantum computing in the distributed setting has recently been the subject of intensive investigations [15, 21, 25, 26]. The main difference between classical and quantum distributed computing is that the quantum setting, quantum information, i.e., quantum bits (qubits), can be sent through the edges of the



© Keren Censor-Hillel, Orr Fischer, François Le Gall, Dean Leitersdorf, and Rotem Oshman; licensed under Creative Commons License CC-BY 4.0

13th Innovations in Theoretical Computer Science Conference (ITCS 2022).

Editor: Mark Braverman; Article No. 35; pp. 35:1–35:25

Leibniz International Proceedings in Informatics



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

network instead of classical information (i.e., bits). Le Gall and Magniez [25] and Izumi and Le Gall [21], in particular, have shown the superiority of quantum distributed computing over classical distributed computing for two fundamental models, the CONGEST model and the CONGESTED CLIQUE model.

The (classical) CONGEST model is one of the most studied models in classical distributed computing. In this model, n nodes communicate with each other over the network by exchanging messages of $O(\log n)$ bits in synchronous rounds. All links and nodes are reliable and suffer no faults. Each node has a distinct identifier, but the network graph is not initially known to the nodes. In the quantum version of this model (which we denote QUANTUM CONGEST), as defined in [15, 25], the only difference is that the nodes can exchange quantum information: each message exchanged consists of $O(\log n)$ quantum bits instead of $O(\log n)$ bits in the classical case. In particular, initially the nodes of the network do not share any entanglement. Achieving quantum speedups in this setting is especially challenging since, as shown by Elkin et al. [15], the ability to send quantum information is not helpful for many crucial components of distributed algorithms (e.g., routing or broadcast of information). Le Gall and Magniez [25] have nevertheless showed the superiority of quantum distributed computing in this model: they constructed a $\tilde{O}(\sqrt{n})$ -round quantum algorithm for the exact computation of the diameter of the network (for networks with small diameter), while it is known that any classical algorithm in the CONGEST model requires $\tilde{\Omega}(n)$ rounds, even for networks with constant diameter [17].

The (classical) CONGESTED CLIQUE model is similar to the (classical) CONGEST model, but it separates the input to the problem we are working on from the communication topology: the input is some graph $G = (V, E)$, but the communication topology allows all nodes to communicate directly with one another (i.e., a clique). This model is close in flavor to massively-parallel computing, but we focus only on rounds and communication, not memory. The QUANTUM CONGESTED CLIQUE model is defined as the quantum version of the CONGESTED CLIQUE model: the only difference is again that each exchanged message consists of $O(\log n)$ quantum bits instead of $O(\log n)$ bits. Izumi and Le Gall [21] showed that quantum distributed algorithms can be more powerful than classical distributed algorithms in the CONGESTED CLIQUE model as well: they constructed a quantum algorithm faster than the best known classical algorithms for the All-Pair Shortest Path problem.

Distributed subgraph detection. In the past few years there has been a surge of works on classical distributed algorithms investigating the complexity of the subgraph detection problem, which asks to detect the existence of a specified subgraph within the input graph (which coincides with the communication network in the CONGEST and QUANTUM CONGEST models). For small subgraphs, the subgraph-detection problem is extremely *local*, and yet it is challenging to solve in the CONGEST model, due to the restricted bandwidth. Most prior work has focused on detecting p -cliques (denoted below K_p) and ℓ -cycles (denoted below C_ℓ), for small values of p and ℓ . In particular, in the p -clique detection problem (also called the K_p -freeness problem), the goal is to decide if the input graph contains a p -clique or not. If it contains a p -clique then at least one node must output “yes”. Otherwise all nodes must output “no”. (A detailed review of classical algorithms for p -clique detection is given at the end of this section.)

Izumi et al. [22] recently showed that quantum algorithms can give an advantage for triangle detection (i.e., K_3 detection), by constructing a quantum algorithm with round complexity $\tilde{O}(n^{1/4})$ in the QUANTUM CONGEST model (the best known classical algorithm for triangle detection, by Chang and Saranurak [9], has complexity $\tilde{O}(n^{1/3})$). The key technique used in [22] is *distributed Grover search*, which was introduced in [25] and consists in a distributed implementation of Grover’s celebrated quantum algorithm [20].

Our results. In this paper we further investigate the power of quantum distributed algorithms. Using a new approach based on *nested* distributed quantum searches, we obtain quantum algorithms for p -clique detection, in both the QUANTUM CONGEST model and the QUANTUM CONGESTED CLIQUE model, that outperform their classical counterparts.

We first consider clique detection in the QUANTUM CONGESTED CLIQUE model. Our results are summarized in the following theorem (the upper bounds we obtain for K_p -detection with $p \geq 5$ are actually even stronger – see Corollary 22 in Section 5).

► **Theorem 1.** *There exists a quantum algorithm that solves p -clique detection with success probability at least $1 - 1/\text{poly}(n)$ in the QUANTUM CONGESTED CLIQUE model with complexity $\tilde{O}(n^{1/5})$ for $p = 3$ and $\tilde{O}(n^{1-2/(p-1)})$ for $p \geq 4$.*

For all values $p \geq 3$, the quantum algorithms we obtain by Theorem 1 are faster than all known classical and quantum algorithms for p -clique detection.

We then investigate clique detection in the QUANTUM CONGEST model.

► **Theorem 2.** *There exists a quantum algorithm that solves p -clique detection with success probability at least $1 - 1/\text{poly}(n)$ in the QUANTUM CONGEST model with complexity $\tilde{O}(n^{1/5})$ for $p = 3$, and $\tilde{O}(n^{1-2/(p-1)})$ for $p \geq 7$.*

For all $p = 3$ and $p \geq 7$, the quantum algorithms we obtain by Theorem 2 are faster than all known algorithms for p -clique detection in the CONGEST or QUANTUM CONGEST model. For $4 \leq p \leq 6$, our approach currently does not lead to an improvement over the classical algorithms from [4], which are respectively $\tilde{O}(\sqrt{n})$, $\tilde{O}(n^{3/5})$ and $\tilde{O}(n^{2/3})$. The reason is, informally, that in the QUANTUM CONGEST model we (and all prior work on listing cliques) decompose the graph into well-connected clusters, and simulate the QUANTUM CONGESTED CLIQUE on each cluster; for very small cliques, the effort required to collect the edges needed at each cluster overwhelms any time savings we currently gain from the quantum search.

Finally, we consider lower bounds. While a tight $\Omega(\sqrt{n})$ lower bound for K_4 -detection is known in the (classical) CONGEST model [10], tight lower bounds are not known for larger cliques, and there is no known non-trivial lower bound for K_p -detection in the quantum setting. We show a barrier for proving an $\Omega(n^{3/5+\epsilon})$ lower bound for p -clique detection for any $\epsilon > 0$. Namely, such a bound would imply breakthrough results in the field of circuit complexity, which are far beyond the current state-of-the-art. We actually show this barrier for the (classical) CONGEST model, but since any lower bound in the QUANTUM CONGEST model also holds the CONGEST model, this barrier holds for the QUANTUM CONGEST model as well.¹

► **Theorem 3.** *For any constant integer $p \geq 4$ and any constant $\epsilon > 0$, proving a lower bound of $\Omega(n^{3/5+\epsilon})$ on p -clique detection in the CONGEST model would imply new lower bounds on high-depth circuits with constant fan-in and fan-out gates.*

Previously, several such barriers were known: in [13] it is shown that there is an absolute constant c such that proving a lower bound of the form $\Omega(n^{1-c})$ for C_{2k} -detection would imply breaking a circuit complexity barrier. In [14] such a barrier is shown for proving an $\Omega(n^\epsilon)$ lower bound on triangle detection, and in [5] a barrier for proving a lower bound of $\Omega(n^{1/2+\epsilon})$ for C_6 -detection is given.

¹ Note that the statement of Theorem 3 is actually interesting only for $p \geq 6$, since for $p = 4, 5$ we know that there exist algorithms beating this barrier: as already mentioned, algorithms with complexity $\tilde{O}(\sqrt{n})$ for $p = 4$ and $\tilde{O}(n^{3/5})$ for $p = 5$ are given in [4].

■ **Table 1** Our upper bounds for subgraph detection, and the corresponding known results in the classical setting. Here n denotes the number of nodes in the network. Note that in QUANTUM CONGESTED CLIQUE, the algorithms we obtain for K_p -detection with $p \geq 6$ are even faster than shown here, see Corollary 22.

Subgraph	CONGESTED CLIQUE		QUANTUM CONGESTED CLIQUE	
K_3	$\tilde{O}(n^{1/3})$	[9]	$\tilde{O}(n^{1/4})$	[22]
			$\tilde{O}(n^{1/5})$	Theorem 1
K_p ($p \geq 4$)	$\tilde{O}(n^{1-2/p})$	[4]	$\tilde{O}(n^{1-2/(p-1)})$	Theorem 1

Subgraph	CONGEST		QUANTUM CONGEST	
K_3	$\tilde{O}(n^{1/3})$	[9]	$\tilde{O}(n^{1/4})$	[22]
			$\tilde{O}(n^{1/5})$	Theorem 2
K_p ($p \geq 7$)	$\tilde{O}(n^{1-2/p})$	[4]	$\tilde{O}(n^{1-2/(p-1)})$	Theorem 2
	$\Omega(\sqrt{n})$	[10]		

We refer to Table 1 for a summary of our results and a detailed comparison with prior works.

Overview of our main technique. Our key approach is to encapsulate K_p -detection as a search task, and use a distributed implementation [25] of Grover search [20] to solve the task. Grover’s algorithm consists of alternating between quantum operations called *Grover diffusion operations*, and checking operations, also called *checking queries*. The total number of operations is $O(\sqrt{|X|})$, where X is the search domain. In the distributed implementation developed in [25], one specific node of the network (the *leader*) executes each Grover diffusion operation locally, but implements each query in a distributed way using a distributed checking procedure. All prior works using this framework ([21, 25, 22]) considered the setting where the checking procedure is a classical procedure. In this work, we consider checking procedures that themselves also apply distributed Grover searches; we develop a framework to describe such *nested* distributed quantum searches (see Lemma 5 in Section 2). Our framework can actually be applied in a completely “black-box” way to design quantum distributed algorithms (i.e., no knowledge of quantum computation is needed to apply this framework).

The main challenge in applying quantum search to the clique-detection problem is that the search-space is very large: we must search over the $\Theta(n^p)$ possibilities, and a naïve approach would require $\Theta(n^{p/2})$ quantum queries, which is extremely inefficient. Instead, we show that one can carefully split the search into nested stages, so that each stage adds a single node to the clique we are trying to find. Crucially, nesting the stages of the search allows us to re-use information computed in one stage for all the search queries in the next stage: in each stage, we have already found some ℓ -cliques, where $\ell < p$, and we want to add one more node to the cliques, to obtain $(\ell + 1)$ -cliques (until in the final stage we obtain p -cliques). To this end, the nodes collect some edges, which allow them to detect some $(\ell + 1)$ -cliques, and then use a nested search to try to complete the $(\ell + 1)$ -cliques into p -cliques.

Perhaps surprisingly, it turns out that in many cases it is not worthwhile to “start the search from scratch”: instead of using quantum search to detect p -cliques “from scratch”, it is more efficient to first *classically* list all q -cliques for some $q < p$, and then use quantum search to find an extension of some q -clique into a p -clique. This echoes the theme of

re-using information throughout the stages of the search: we precompute some information classically, which will be used by all stages of the search. For example, we show that to solve triangle-detection, we can improve on the algorithm from [22] by first classically listing *edges*, so that every node of the congested clique learns some set of edges that it will be responsible for trying to complete into a triangle, and then using quantum search to find a node that forms a triangle with some edge. This reduces the running time from $\tilde{O}(n^{1/4})$ rounds in [22] to $\tilde{O}(n^{1/5})$ rounds in our new algorithm.

More generally, we can solve the K_p -detection problem by first classically listing all instances of K_{p-1} in the graph, and then performing distributed Grover search over the nodes, to check if some $(p-1)$ -clique can be extended into a p -clique. Since there are n nodes to check, the Grover search will require \sqrt{n} quantum queries, and each query can be checked in $O(1)$ rounds (in the congested clique, it is possible to learn all neighbors of a given node in a single round). Thus, the overall running time we obtain will be $\tilde{O}(L_{p-1} + \sqrt{n})$, where L_{p-1} is the time required to list all $(p-1)$ -cliques. For example, 4-cliques can be listed in $L_4 = O(\sqrt{n})$ rounds [12], and this approach allows us to solve the K_5 -detection problem in roughly the same time complexity, $\tilde{O}(\sqrt{n})$. However, the cost L_{p-1} grows with p , so sometimes it is better to start from a smaller clique, K_q for $q < p-1$, and extend by more than a single node. This leads to our general nested-search-based approach, which starts by listing all copies of K_q for some $q < p$, and then uses nested quantum search to check if some q -clique can be extended into a p -clique.

Review of prior works on classical algorithms for clique detection. In the CONGEST model, the first sublinear algorithm for p -clique detection was obtained for $p = 3$ (i.e., triangle detection) by Izumi and Le Gall [23]. The complexity of triangle detection was then improved to $\tilde{O}(\sqrt{n})$ by Chang et al. [8], where n denotes the number of nodes, and then further to $\tilde{O}(n^{1/3})$ by Chang and Saranurak [9]. For p -cliques with $p \geq 4$, the first sublinear detection algorithm was constructed by Eden et al. [13]. These results were improved to $\tilde{O}(n^{p/(p+2)})$ rounds for all $p \geq 4$ by Censor-Hillel et al. [6], and very recently, $\tilde{O}(n^{1-2/p})$ rounds for all $p \geq 4$ by Censor-Hillel et al. [4]. Czumaj and Konrad [10] have shown the lower bound $\Omega(\sqrt{n})$ for p -clique detection for $p \geq 4$, which matches the upper bound from [4] for $p = 4$. Proving lower bounds for triangle detection, on the other hand, appears extremely challenging: it is known that for any $\varepsilon > 0$, showing a lower bound of $\Omega(n^\varepsilon)$ on triangle detection implies strong circuit complexity lower bounds [5] (see nevertheless [1] for a weaker, but still non-trivial, lower bound for triangle detection).²

In the powerful CONGESTED CLIQUE model, the best known upper bounds on the round complexity of p -clique detection is $O(n^{0.158})$ for $p = 3$, which is obtained by the algebraic approach based on matrix multiplication developed by Censor-Hillel et al. [7], and $O(n^{1-2/p})$ for any constant $p \geq 4$ [12].

Further related works on quantum distributed computing. There exist a few works investigating the power of quantum distributed computing in other models or settings (see also [2] and [11] for surveys). In the LOCAL model, separations between the computational powers of the classical and quantum algorithms have been also obtained [18, 26]. Over anonymous networks, zero-error quantum algorithms have been constructed for leader election [30]. Finally, quantum algorithms for byzantine agreements have also been investigated [3].

² Note that the algorithms from [8, 6, 9, 4] actually solve the listing version of the problem (which asks to list all p -cliques of the graph) as well. For the listing version, lower bounds matching the upper bounds from [9, 4] for all values of $p \geq 3$ are known [23, 28, 16].

Organization of the paper. The core conceptual message of the paper is contained in the first 10 pages: we describe our main technique in Section 2 and then, in Section 3, explain how to use this technique to construct fast quantum algorithms for clique detection in the QUANTUM CONGESTED CLIQUE model. Further sections then show how to apply the technique to construct fast algorithms in the QUANTUM CONGEST model (Section 4, which proves Theorem 2) and construct even faster algorithms in the QUANTUM CONGESTED CLIQUE model (Section 5, which proves Theorem 1). An overview of Theorem 3 is given in Section 6.

2 Nested Distributed Quantum Searches

In this section we present our main technique: nested distributed quantum searches. This is a generalization of a technique (called below *distributed Grover search*) used in prior quantum distributed works [21, 25, 22].

We note that implementing distributed quantum searches in a nested way is already allowed (but not used) in the framework introduced in [25]. Our main contribution in the current section is developing this approach into a full framework and describing its concrete implementation in the distributed setting.

Standard Grover search. We begin by informally describing the most standard framework for Grover search – as a technique to solve a search problem with black-box access.

Consider the following: given black-box access to a function $f: X \rightarrow \{0, 1\}$, for an arbitrary X , find an $x \in X$ such that $f(x) = 1$, if such an element exists. Grover’s quantum algorithm [20] solves this problem with high probability using $O(\sqrt{|X|})$ calls to the black box. Grover’s algorithm consists of $O(\sqrt{|X|})$ steps, where each step executes one quantum operation called the *Grover diffusion operation*, which does not use the black-box, and an operation called the *checking procedure*, which uses one call to the black-box.

Distributed Grover search. Let us present the basic quantum distributed search framework (distributed Grover search) introduced in [25]. In this distributed implementation, one specific node, called *the leader*, runs each Grover diffusion locally, but the checking procedure is implemented via a distributed algorithm.

Consider again a function $f: X \rightarrow \{0, 1\}$, for an arbitrary X , and the following search problem: one specified node (the leader) should find an element $x \in X$ such that $f(x) = 1$, or, if no such element exists the leader should output “not found”. Assume there exists a distributed algorithm \mathcal{A} , called the *checking procedure*, in which the leader is given $x \in X$ as input, and the leader returns $f(x)$ as output. The checking procedure \mathcal{A} is often described as a classical algorithm, but it can also be a quantum distributed algorithm.³

Let r be the round complexity of \mathcal{A} . The framework introduced in [25] shows that there is a quantum distributed algorithm that runs in $\tilde{O}(\sqrt{|X|} \cdot r)$ rounds and enables the leader to solve the above search problem with probability at least $1 - 1/\text{poly}(n)$. While the original statement in [25] was for the QUANTUM CONGEST model, as explained in [21], the same holds for the QUANTUM CONGESTED CLIQUE model.

³ As explained in [25], a classical procedure can easily be converted using standard techniques into a quantum procedure able to deal with superpositions of inputs.

► **Lemma 4** ([25]). *There is a quantum algorithm that runs in $\tilde{O}(\sqrt{|X|} \cdot r)$ rounds and enables the leader to solve the above search problem with probability at least $1 - 1/\text{poly}(n)$. This statement holds in both the QUANTUM CONGEST and the QUANTUM CONGESTED CLIQUE models.*

Nested distributed quantum searches. All prior works using quantum distributed search ([21, 25, 22]) used a classical checking procedure \mathcal{A} . The framework of [25] nevertheless allows quantum checking procedures. In particular, a distributed Grover search can be used as the checking procedure. We now present our framework for nested distributed quantum searches, consisting of k nested levels, where at each level:

1. The nodes run a distributed *setup* step for the current level, collecting information and preparing for the next search levels. The setup procedure in our results is classical, but in general it can be quantum.
2. We execute the next level of the search. Crucially, the information prepared during the setup will be re-used to evaluate all the nested queries in the next level (and subsequent levels).

Formally, let $f: X_1 \times \dots \times X_k \rightarrow \{0, 1\}$ be a function, for a constant $k \geq 2$ and sets X_1, \dots, X_k . The goal is finding $(x_1, \dots, x_k) \in X_1 \times \dots \times X_k$ where $f(x_1, \dots, x_k) = 1$, if such exists. For $\ell \in \{1, \dots, k-1\}$, $u \in V$, let $\text{setup}_\ell^u: X_1 \times \dots \times X_\ell \rightarrow \{0, 1\}^*$ be a function describing the setup data of u for the $(\ell+1)$ -th search. Let $\mathcal{S}_1, \dots, \mathcal{S}_{k-1}$ and \mathcal{C} be distributed algorithms with the following specifications.

- Algorithm \mathcal{S}_1 . Input: the leader is given $x_1 \in X_1$. Output: each node $u \in V$ outputs $\text{setup}_1^u(x_1)$.
- Algorithm \mathcal{S}_ℓ for any $\ell \in \{2, \dots, k\}$. Input: the leader is given $(x_1, \dots, x_\ell) \in X_1 \times \dots \times X_\ell$ and each node $u \in V$ is given $\text{setup}_{\ell-1}^u(x_1, \dots, x_{\ell-1})$. Output: each node $u \in V$ outputs $\text{setup}_\ell^u(x_1, \dots, x_\ell)$.
- Algorithm \mathcal{C} . Input: the leader is given $(x_1, \dots, x_k) \in X_1 \times \dots \times X_k$ and each node $u \in V$ is given $\text{setup}_{k-1}^u(x_1, \dots, x_{k-1})$. Output: the leader outputs $f(x_1, \dots, x_k)$.

Let s_1, \dots, s_{k-1} and c denote the round complexities of $\mathcal{S}_1, \dots, \mathcal{S}_{k-1}$ and \mathcal{C} , respectively. Applying Lemma 4 leads to the following result.

► **Lemma 5.** *There is a quantum algorithm that runs in*

$$\tilde{O}\left(\sqrt{|X_1|}\left(s_1 + \sqrt{|X_2|}\left(s_2 + \sqrt{|X_3|}\left(s_3 + \dots + \sqrt{|X_{k-1}|}\left(s_{k-1} + \sqrt{|X_k|}(s_k + c)\right)\right)\right)\right)\right)$$

rounds and enables the leader to output x_1, \dots, x_k such that $f(x_1, \dots, x_k) = 1$, or output that there are no such x_1, \dots, x_k , with probability at least $1 - 1/\text{poly}(n)$. This statement holds in both the QUANTUM CONGEST and the QUANTUM CONGESTED CLIQUE models.

3 The Power of Nested Quantum Search: Clique-Detection from Listing in the Quantum Congested Clique Model

In this section we describe our approach for taking an algorithm for K_p -listing in the CONGESTED CLIQUE or QUANTUM CONGESTED CLIQUE model, and extending it to K_{p+t} -detection (for some $t > 0$) using quantum search. We give two variations of the approach: the first uses the K_p -listing algorithm *as a black box*, so that any such algorithm can be used (for example, algorithms that perform better on certain classes of input graphs, etc.).

This also forms the basis of our QUANTUM CONGEST algorithms in Section 4. The second approach yields faster results, but it “opens the black box”, relying on the properties of the K_p -listing algorithm from [12]. Since it is more complicated, the second approach is described in Section 5.

To exploit the large bandwidth of the CONGESTED CLIQUE, we use Lenzen’s routing scheme for solving the *information distribution task*: for some $s \geq 1$, each $v \in V$ has at most $s \cdot n$ messages m_1^v, \dots, m_r^v , each of $O(\log n)$ bits, and each with a destination $\text{dest}(m_i^v) \in V$. Each $v \in V$ is the destination of at most $s \cdot n$ messages ($|\{m_i^u \mid u \in V \wedge i \in [s \cdot n] \wedge \text{dest}(m_i^u) = v\}| \leq s \cdot n$), and we wish to deliver each message m_i^v to its destination $\text{dest}(m_i^v)$.

► **Lemma 6** (Lenzen’s Routing Scheme[27]). *The information distribution task with parameter s can be solved in $O(s)$ rounds in CONGESTED CLIQUE.*

3.1 Warmup: Detecting Triangles in $\tilde{O}(n^{1/5})$ Rounds

We describe a simple triangle detection algorithm demonstrating the basic idea of our approach, and improving upon the state-of-the-art algorithm from [22].

In the algorithm, we partition the search-space V^3 into n shards, one per node, and each $v \in V$ checks if there is a triplet (u_1, u_2, u_3) in its shard that is a triangle in G . Each shard has the form $A_i \times A_j \times Q_k$, where $A_1, \dots, A_{n^{2/5}}$ partitions V into sets of $n^{3/5}$ nodes, and $Q_1, \dots, Q_{n^{1/5}}$ partitions V into sets of size $n^{4/5}$.⁴Note that the total number of shards is indeed $n^{2/5} \cdot n^{2/5} \cdot n^{1/5} = n$.

To check if its shard $A_i \times A_j \times Q_k$ contains a triangle, node v learns the edges $E(A_i, A_j) = E \cap (A_i \times A_j)$, and then, using a distributed quantum search, checks if some $w \in Q_k$ forms a triangle with some $\{u_1, u_2\} \in E(A_i, A_j)$. The search is not performed directly over Q_k : instead, we partition Q_k into *batches*, Q_k^1, \dots, Q_k^b , and search for a batch Q_k^ℓ with a node forming a triangle. Processing the nodes in batches allows us to fully utilize the bandwidth in the CONGESTED CLIQUE. However, we must balance the *size* of batches, which determines the time to check if a batch has a node completing a triangle, against the *number* of batches, which determines the number of quantum queries we will need to perform.

Detailed description of the algorithm. Consider a node v , and let $A_i \times A_j \times Q_k$ be the shard assigned to node v . The algorithm has two steps:

1. Node v learns $E(A_i, A_j)$, using Lenzen’s routing scheme.
2. Node v partitions Q_k into $n^{2/5}$ batches, $Q_k = \{Q_k^1, \dots, Q_k^{n^{2/5}}\}$, each containing $n^{2/5}$ nodes (since $|Q_k| = n^{4/5}$). We use a quantum search over $\ell \in [n^{2/5}]$ to check whether there exists a Q_k^ℓ containing a node $w \in Q_k^\ell$ that forms a triangle together with two nodes $u_1 \in A_i, u_2 \in A_j$.

Formally, we instantiate Lemma 4 with the search-space $X = [n^{2/5}]$ (i.e., the batch indices). The checking procedure \mathcal{A} checks an index $\ell \in [n^{2/5}]$ by routing $E(A_i \cup A_j, Q_k^\ell)$ to v (in parallel at all nodes). Then, v locally checks whether there is a $(u_1, u_2, w) \in A_i \times A_j \times Q_k^\ell$ such that $\{u_1, u_2\} \in E(A_i, A_j)$, $\{u_1, w\} \in E(A_i, Q_k^\ell)$, and $\{u_2, w\} \in E(A_j, Q_k^\ell)$; it sends ‘1’ to the leader if it found such a triplet, and ‘0’ otherwise.

⁴ To simplify the presentation, here and everywhere in the paper, when partitioning V into n^δ subsets, for a $\delta \in (0, 1)$, we assume n^δ is an integer and divides n . If this is not the case, one can replace n^δ by $\lceil n^\delta \rceil$, without affecting the asymptotic complexity.

Complexity. Step 1 requires $O(|A_i| \cdot |A_j|/n) = O(n^{2 \cdot 3/5 - 1}) = O(n^{1/5})$ rounds, using Lenzen’s routing scheme (Lemma 6). In Step 2, checking a particular batch Q_k^ℓ requires node v to learn $E(A_i \cup A_j, Q_k^\ell)$. As $|E(A_i \cup A_j, Q_k^\ell)| = O(n^{3/5 + 2/5}) = O(n)$, this can be done in $O(1)$ rounds using Lemma 6. By Lemma 4, since the search space is $[n^{2/5}]$, Step 2 takes $\tilde{O}(n^{1/5})$ rounds. In total, the algorithm takes $\tilde{O}(n^{1/5})$ rounds.

3.2 Extending K_p Listing to K_{p+t} Detection

Our triangle detection algorithm has the following structure: we view a triangle as an edge $\{u_1, u_2\}$, plus a node w connected to u_1, u_2 . We *classically* route information between nodes, so they can list the edges $\{u_1, u_2\}$ in the sets they are responsible for ($A_i \times A_j$). Then, we use *quantum search* to check if there is a node w forming a triangle with a listed edge.

We extend this idea to cliques of arbitrary sizes: given $q > 2$, take p, t where $p + t = q$. We view a q -clique as a p -clique $\{v_1, \dots, v_p\}$, plus a t -clique $\{u_1, \dots, u_t\}$ where u_1, \dots, u_t are all connected to v_1, \dots, v_p . We classically list all p -cliques, and then use quantum search to check for a t -clique forming a $q = (p + t)$ -clique with a listed p -clique.

We present two variants of this approach. The first takes a K_p -listing algorithm as a black box, making no assumptions about which p -cliques are found by which nodes. The second improves on the first by “opening the black box” and using properties of the K_p -listing algorithm of [12]: knowing which p -cliques are listed by each node reduces the amount of information we route during the quantum search, as some edges are not relevant to some nodes. We present the first variant here, and the second is given in Section 5. Note that our triangle detection algorithm is an instance of the second variant, since we exploit out knowledge of A_i, A_j to determine which edges ($E(A_i \cup A_j, Q_k^\ell)$) are learned by a given node as it evaluates batch Q_k^ℓ .

How should we explore the search-space $V^{(t)}$ of possible t -cliques that may extend a given p -clique to a $(p + t)$ -clique? One possibility is to partition it into batches, and search over them, as we did for triangles. However, the large search-space makes this inefficient: every node must learn the edges between every pair of nodes in the current batch, and since we can use at most n^2 batches⁵ to cover $V^{(t)}$, very soon we reach a situation where every node needs to learn *all* the edges. Instead, we use a *nested search*, building the t -clique node-by-node. The search is structured so that edges learned at a given level are re-used to evaluate *many* nested queries on following levels. See Fig. 1 for an example partitioning of the search-space.


The initial state. Let $V^{(p)}$ denote all subsets of V with cardinality p . When we begin, we assume copies of K_p have already been found: each node $u \in V$ has a subset $S_u^p \subseteq V^{(p)}$ of p -cliques it found. Let $S^p = \bigcup_{u \in V} S_u^p$ be all copies of K_p found by the nodes. We assume S^p is the set of all p -cliques in G .

We say that an algorithm A *extends* from K_p to K_{p+t} if, given sets $\{S_u^p\}_{u \in V}$, w.h.p., algorithm A outputs ‘1’ at all nodes iff G contains a $(p + t)$ -clique $\{v_1, \dots, v_{p+t}\} \in V^{(p+t)}$ such that $(v_1, \dots, v_p) \in S^p$.

► **Theorem 7.** *For every $p \geq 2, t \geq 1$, there is an algorithm that extends from K_p to K_{p+t} in $\tilde{O}(n^{1-1/2^t})$ rounds in QUANTUM CONGESTED CLIQUE.*

Proof. Fix in advance t partitions of V , where the ℓ -th partition divides V into $n^{1/2^{t-\ell}}$ sets, $V_1^\ell, \dots, V_{n^{1/2^{t-\ell}}}^\ell$, each of size $n^{1-1/2^{t-\ell}}$. Note that the partitions become increasingly finer, until at level $\ell = t$ we have n sets comprising a single node each. Our goal is to execute a

⁵ Otherwise we will need more than $\sqrt{n^2} = n$ quantum queries.



■ **Figure 1** Extending from K_p -listing to K_{p+2} -detection. We search $V \times V$ to check for a pair of nodes that can be added to an existing p -clique to form a $(p+2)$ -clique (in the figure, $|V| = 9$). A non-nested search requires $\sqrt{n} \cdot n = n$ queries. In contrast, in a nested search, we split the first dimension into \sqrt{n} batches, and search over them (requiring $n^{1/4}$ outer queries); to implement each outer query, all nodes send the edges corresponding to the current batch (requiring \sqrt{n} rounds), and then explore the second dimension ($|V| = n$) using \sqrt{n} inner quantum queries. The time of the entire search is $\tilde{O}(n^{1/4}(\sqrt{n} + \sqrt{n})) = \tilde{O}(n^{3/4})$.

nested quantum search to check if there is a t -tuple of indices (x_1, \dots, x_t) , such that there exist $v_1 \in V_{x_1}^1, \dots, v_t \in V_{x_t}^t$ forming a $(p+t)$ -clique together with some previously-listed p -clique $\{w_1, \dots, w_p\} \in S^p$.

We instantiate Lemma 5, executing a nested search with t levels over the domain $X_1 \times \dots \times X_t$, where $X_\ell = [n^{1/2^{t-\ell}}]$ for each $\ell = 1, \dots, t$. We search for an element satisfying the function

$$f(x_1, \dots, x_t) = 1 \Leftrightarrow \exists v_1 \in V_{x_1}^1 \dots \exists v_t \in V_{x_t}^t \exists \{w_1, \dots, w_p\} \in S^p : \{w_1, \dots, w_p, v_1, \dots, v_t\} \text{ is a } (p+t)\text{-clique in } G.$$

At level $\ell \leq t$ of the search, the setup we prepare takes the form of sets $\{S_u^{p+\ell}\}_{u \in V}$ at the nodes, where $S_u^{p+\ell} \subseteq V^{(p+\ell)}$ is a set of $(p+\ell)$ -cliques in G that node u has learned about. When initiating the search (" $\ell = 0$ "), we are given $\{S_u^p\}_{u \in V}$. The algorithm \mathcal{S}_ℓ that prepares the setup for level $\ell \geq 1$ is as follows:

- Every node $u \in V$ broadcasts the subset of $V_{x_\ell}^\ell$ that it is neighbors with:

$$N_{u, x_\ell} = \{v \in V_{x_\ell}^\ell : \{u, v\} \in E\}.$$

- Locally, each node $u \in V$ prepares $S_u^{p+\ell}$, by listing all the $(p+\ell)$ -cliques that it can form by taking a $(p+\ell-1)$ -clique from $S_u^{p+\ell-1}$ and appending to it a node from $V_{x_\ell}^\ell$:

$$S_u^{p+\ell} = \left\{ \{v_1, \dots, v_{p+\ell}\} \in V^{(p+\ell)} : \{v_1, \dots, v_{p+\ell-1}\} \in S_u^{p+\ell-1}, \right. \\
 \left. v_{p+\ell} \in V_{x_\ell}^\ell, \text{ and for each } i = 1, \dots, p+\ell-1 \text{ we have } v_{p+i} \in N_{v_i, x_\ell} \right\}.$$

The final algorithm \mathcal{C} that evaluates $f(x_1, \dots, x_t)$ simply has each node u inform the leader whether S_u^{p+t} is empty or not. If there is some node u with $S_u^{p+t} \neq \emptyset$, the leader outputs '1', and otherwise '0'.

Complexity. The size of the ℓ -th level partition is chosen to as to balance the setup cost against the time required for the remainder of the nested search: at level ℓ , the setup cost is $|V_{x_\ell}^\ell| = s_\ell = n^{1-1/2^{t-\ell}}$. Using Lemma 5, a backwards induction on ℓ shows the cost for levels $\ell + 1, \dots, t$ of the search is $\tilde{O}\left(n^{1-1/2^{t-\ell}}\right)$, matching the setup cost. The cost of the entire search (i.e., levels $\ell = 1, \dots, t$) is $\tilde{O}\left(n^{1-1/2^t}\right)$. ◀

By combining the classical clique-listing algorithm of [12] with Theorem 7, we obtain an algorithm for detection of K_p which improves on the state-of-the-art classical algorithm for $p \geq 5$.

► **Theorem 8.** *Given a graph $G = (V, E)$ and a clique size $p \geq 5$, it is possible to detect whether there exists an instance of K_p in G within $\tilde{O}(\min_{t \in \mathbb{N}} \max\{n^{1-\frac{2}{p-t}}, n^{1-\frac{1}{2^t}}\})$ rounds of the QUANTUM CONGESTED CLIQUE model.*

For instance, for $p = 5$, by taking $t = 1$, we get K_5 -detection in $\tilde{O}(n^{1/2})$ rounds of the QUANTUM CONGESTED CLIQUE model, improving on the classical runtime of $\tilde{O}(n^{3/5})$. We note that $\tilde{O}(n^{1/2})$ is the time required to classically list K_4 , so the quantum-search-based extension from 4-cliques to 5-cliques is “for free”.

4 Detection from Listing in the Quantum Congest Model

This section is devoted to the proof of Theorem 2.

► **Theorem 2.** *There exists a quantum algorithm that solves p -clique detection with success probability at least $1 - 1/\text{poly}(n)$ in the QUANTUM CONGEST model with complexity $\tilde{O}(n^{1/5})$ for $p = 3$, and $\tilde{O}(n^{1-2/(p-1)})$ for $p \geq 7$.*

We begin with the proof for triangle detection in QUANTUM CONGEST, and then proceed with K_p detection for $p \geq 7$.

At a very high level, our algorithms in this section use the framework of decomposing the graph into clusters of high conductance and working within each cluster in order to find the required subgraph, and then recursing over the edges remaining outside of clusters. To work within a cluster, throughout this section we will use the following expander decomposition and routing theorems.

Preliminaries. We begin by defining the notions of mixing time and conductance, which are used in the context of the expander decomposition. We note that we do not use these definitions directly, but rather, we use previously-proven lemmas that use these properties to obtain efficient routing or simulation procedures on such graphs.

The *conductance* of a graph G is $\Phi(G) = \min_{S \subseteq V} \frac{|E(S, V \setminus S)|}{\min(|E(S, V)|, |E(V, V \setminus S)|)}$, is the worst-case ratio between the number of edges crossing a cut in the graph, and the number of edges contained on either side of the cut.

The conductance is related to the *mixing time* of the graph, which, informally, is the number of steps required for a random walk starting from any vertex u to become close to its stationary distribution, where the probability of being at any given vertex v is roughly $\text{degree}(v)/2m$. The mixing time of a graph G is denoted $\tau_{\text{mix}}(G)$, and it is related to the conductance as follows: $\Theta\left(\frac{1}{\Phi(G)}\right) \leq \tau_{\text{mix}}(G) \leq \Theta\left(\frac{\log n}{\Phi^2(G)}\right)$ (see Corollary 2.3 in [24]).

► **Lemma 9** ([9, 4]). *Let $\delta > 0$ such that $m = n^{1+\delta}$. For any $\epsilon \in (0, 1)$, and constant $\gamma \in (0, 1)$, there is a constant $a_\gamma > 0$ dependent only on γ , such that a decomposition can be constructed in $\tilde{O}(n^\gamma)$ rounds, with high probability, in which the edges of the graph are partitioned into two sets, E_m, E_r , that satisfy the following conditions:*

1. Each connected component (cluster) C of E_m has conductance $\Phi(C) \geq (\epsilon/\log n)^{\alpha\gamma}$, and has average degree at least ϵn^δ .
2. For any cluster C and node $v \in V_C$, $\deg_{V_C}(v) \geq (\epsilon/\log n)^{\alpha\gamma} \deg_{V \setminus V_C}(v)$.
3. $E_r \leq \epsilon m$.

► **Theorem 10** (Expander routing [9, 19]). *Suppose $\tau_{\text{mix}}(G) = \text{polylog}(n)$ and let $0 \leq \gamma \leq 1$ be a constant. There is an $O(n^\gamma)$ -round algorithm that pre-processes the graph such that for any subsequent routing task where each node v is a source and a destination of at most $L \cdot \deg(v)$ messages of $O(\log n)$ bits, all messages can be delivered in $L \cdot \log^\alpha(n)$ rounds of the CONGEST model, with high probability, where α is a constant that depends only on γ .*

4.1 $\tilde{O}(n^{1/5})$ -Round Triangle Detection in the Quantum Congest Model

In this section we show an $\tilde{O}(n^{1/5})$ -round triangle detection algorithm in QUANTUM CONGEST. That is, we prove the following theorem, which is part of Theorem 2.

► **Theorem 11.** *There exists an algorithm that solves triangle detection with success probability at least $1 - 1/\text{poly}(n)$ in the QUANTUM CONGEST model with complexity $\tilde{O}(n^{1/5})$.*

As observed in [8, 9] (see also Theorem 4 in [22]), it is sufficient to solve this problem in high conductance graphs (with some additional edges incident to this graph) in order to obtain an algorithm for general graphs. Our QUANTUM CONGEST algorithm shares many similarities with the $\tilde{O}(n^{1/5})$ triangle detection quantum algorithm in the QUANTUM CONGESTED CLIQUE shown in Section 3, but requires a more “sparsity aware” approach for it to work in this more restricted model.

Specifically, in the FINDTRIANGLEINSUBNETWORK problem, as defined in [22], the input network $G' = (V', E_{\text{in}} \cup E_{\text{out}})$ is a connected network such that the mixing time of the graph $G_{\text{in}} = (V(E_{\text{in}}), E_{\text{in}})$ is at most $O(\text{polylog}(n))$ (e.g., its conductance is at least $\Omega(1/\text{polylog}(n))$), and each edge in E_{out} is incident to a node in $V(E_{\text{in}})$ so that $\deg_{G_{\text{in}}}(v) \geq \deg_{G' - G_{\text{in}}}(v)$ for every node $v \in G_{\text{in}}$, and the goal of the network is to determine whether G' is triangle-free.

► **Theorem 12** ([8, 9], Theorem 4 in [22]). *Assume that there exists an r -round distributed QUANTUM CONGEST algorithm A that solves the FINDTRIANGLEINSUBNETWORK problem with probability at least $1 - 1/n^3$ and uses only the edges in $E_{\text{in}} \cup E_{\text{out}}$ for communication. Then there exists an $O(r \log n + n^{0.1})$ -round QUANTUM CONGEST algorithm that solves the triangle finding problem over the whole graph $G = (V, E)$ with probability at least $1 - 1/\text{poly}(n)$.*

In the remainder of this subsection we describe an algorithm in QUANTUM CONGEST for the problem FINDTRIANGLEINSUBNETWORK, which assumes the input network has the restrictions mentioned above. We then plug this algorithm into Theorem 12 and obtain Theorem 11 as claimed. Specifically, we prove the following.

► **Theorem 13.** *There exists an algorithm that solves the FINDTRIANGLEINSUBNETWORK problem with success probability at least $1 - 1/n^3$ in the QUANTUM CONGEST model with complexity $\tilde{O}(n^{1/5})$.*

Denote by $m_{\text{in}} = |E_{\text{in}}|$ and by $n_{\text{in}} = |V(E_{\text{in}})|$ the number of edges and vertices inside the high conductance graph respectively, and by $\bar{m} = |E_{\text{in}} \cup E_{\text{out}}|$, $\bar{n} = |V'|$ the number of edges and vertices in G' respectively. Denote $\delta > 0$ such that $\bar{m} = \bar{n}^{1+\delta}$, that is, $O(\bar{n}^\delta)$ is the average degree in G' .

Computation units. Since $\deg_{G_{\text{in}}}(v) \geq \deg_{G' - G_{\text{in}}}(v)$ for every node $v \in G_{\text{in}}$, there is a constant $c > 1$ such that $\bar{m} \leq cm_{\text{in}}$. We define a *computation unit* to be a paired set (v, \tilde{E}) where $v \in V(E_{\text{in}})$, $\tilde{E} \subseteq E_{\text{in}}$, and such that $|\tilde{E}| = \lfloor \bar{n}^\delta / (2c) \rfloor$ and all edges in \tilde{E} are incident to v . We say that the node v is the *core* of the computation unit. Notice that a node can be a core of multiple computation units. A set of computation units is called *pairwise-disjoint* if for any two computation units $(v_1, E_1), (v_2, E_2)$ in the set, either the cores are different, i.e., $v_1 \neq v_2$, or the edges E_1 and E_2 are disjoint.

The reason we define computation units is as follows. In a nutshell, our algorithm will split the edges of the input into sets, such that each set defines a triangle finding *task* of checking whether any three edges in the set form a triangle and such that all possible triplets of edges are checked. To do so, we need to assign all such tasks to the nodes. However, a node with a smaller degree can receive less information to check compared with a node with a higher degree. This is where the computation units come into play: Each core of a computation unit will use the edges of its unit to send and receive information, so that it can process edges associated with a single such task. Now, the bandwidth of a node with a high degree will be exploited by having this node be a core of more computation units, and hence it will be responsible for solving more such tasks, compared with a node of a smaller degree.

First, we show that we have sufficient computation units. Consider a node v with internal degree at least $\bar{n}^\delta / 2c$ and split its edges into disjoint sets of size $\lfloor \bar{n}^\delta / 2c \rfloor$, and a remainder set of edges of size at most $\lfloor \bar{n}^\delta / 2c \rfloor$. Mark each such set of size $\lfloor \bar{n}^\delta / 2c \rfloor$ as a computation unit with the node v as its core.

► **Lemma 14.** *In the process above, the set of marked computation units is pairwise-independent and is of size at least $\bar{n}/2$.*

Proof. We note that the set is indeed pairwise independent, as for each node v the edges of the computation units with node v are edge-disjoint.

Recall that $\bar{m} \leq cm_{\text{in}}$ for the constant $c > 1$. Assume by contradiction that there are less than $\bar{n}/2$ computation units marked. Consider a graph \tilde{G}_{in} that is defined as a subgraph of G_{in} after removing all edges participating in any marked computation unit. By the assumption on the number of computation units and by the bound on the number of edges in a computation unit, there are at most $(\bar{n}/2) \cdot \lfloor \bar{n}^\delta / 2c \rfloor = \bar{n}^{1+\delta} / 4c$ edges removed. This implies that the number of edges in \tilde{G}_{in} is at least $m_{\text{in}} - \bar{n}^{1+\delta} / 4c \geq c\bar{m} - \bar{n}^{1+\delta} / 4c = \bar{n}^{1+\delta}(c - 1/4c) \geq \bar{n}^{1+\delta} / 2c$, where the last inequality holds for $c \geq \sqrt{3}/4$, which we have since $c > 1$.

This implies that the average degree in \tilde{G}_{in} is at least $\bar{n}^{1+\delta} / 2cn_{\text{in}} \geq \bar{n}^{1+\delta} / 2c\bar{n} \geq \bar{n}^\delta / 2c$, which implies that there is at least one node v with degree $\bar{n}^\delta / 2c$ in \tilde{G}_{in} , which contradicts the process above, as v could have marked an additional computation unit. ◀

The above process of constructing computation units can easily be computed by each node separately. Thus, in order to prove Theorem 13, we need a way to make the computation units globally known, in the sense that each has a unique identifier that is known to all nodes. To this end, we will use the following technical claim.

▷ **Claim 15** ([8] variant of Lemma 4.1). Let $G = (V, E)$ be a graph with $\text{polylog}(n)$ diameter. Assume every node has some integer value $f(v)$ which fits in a single $O(\log n)$ bit message. It is possible to give every node v in the network a new unique identifier $i_v \in [n]$, such that there is a globally known function $\tilde{f} : [n] \rightarrow \mathbb{N}$ that any node u can locally compute any of its values, and such that $\tilde{f}(i_v)$ is a 2-approximation to $f(v)$ for all $v \in V$. This algorithm requires $\text{polylog}(n)$ rounds of the CONGEST model.

We will also need the ability to have a leader node that coordinates quantum searches. To implement a leader, the diameter of the network naturally gets into the round complexity. Luckily, in [13], it was shown that for any H -freeness problem (i.e., H -detection), we may assume without loss of generality that the network has small diameter, stated as follows.

► **Lemma 16** ([13]). *Consider an H -freeness problem, where $|H| = k$. Let \mathcal{A} be a protocol that solves \mathcal{P} in time $T(n, D)$ with error probability $\rho = o(\frac{1}{n \log n})$. There is an algorithm \mathcal{A}' that solves \mathcal{P} with round complexity $\tilde{O}(T(n, O(k \log n)) + k \log^2 n)$ and error probability at most $c\rho n \log n + \frac{1}{\text{poly}(n)}$, for some constant c .*

Using the above *diameter reduction* technique, we assume that the diameter of the network is $O(\text{polylog}(n))$. Moreover, we assume that the network computes a leader node v^* and a BFS tree rooted at v^* , which can be done by the network in $O(\text{polylog}(n))$ rounds.

The last ingredient that we need is the notion of k -wise independent hash functions. For two integers a, b we say that a function $f : [a] \rightarrow [b]$ is a k -wise independent hash function if for any k distinct elements $x_1, \dots, x_k \in [a]$ and any k elements $y_1, \dots, y_k \in [b]$, it holds that $\Pr(f(x_1) = y_1 \wedge \dots \wedge f(x_k) = y_k) = 1/b^k$.

We will use the following tail bound for bounding the size of a given bin in a k -wise independent hash function.

► **Lemma 17** ([29], Theorem 5(II)(b)). *If X is the sum of k -wise independent random variables, each of which is confined to the interval $[0, 1]$ and has $\mu = E[X]$, then for $\alpha \geq 1$ and assuming $k \leq \lceil \alpha \mu e^{-1/3} \rceil$,*

$$\Pr(|X - \mu| \geq \alpha \mu) \leq e^{-\lfloor k/2 \rfloor}.$$

We are now ready to prove Theorem 13.

Proof of Theorem 13. First, each node splits its edges into computation units as describe above. This is done locally without communication. By Lemma 14, there are at least $\bar{n}/2$ computation units.

Recall that by Lemma 16, we can assume that the diameter of G' is $O(\text{polylog } n)$. The nodes run the procedure of Claim 15 to give to each node v a new unique ID i_v in $[\bar{n}]$, such that there is a globally known function $\tilde{f} : [n] \rightarrow \mathbb{R}$ that any node u can locally compute any of its values, and such that for any $v \in V$, $\tilde{f}(i_v)$ is a 2-approximation to the number of computation units of v . Using this knowledge, the network gives a unique ID in $[\bar{n}/4]$ to a set of $\lfloor \bar{n}/4 \rfloor$ marked computation units, such that all nodes know for each such computation unit its node ID (this can be done locally since the function $\tilde{f}(i_v)$ is globally known). Note that we have $\lfloor \bar{n}/2 \rfloor$ marked computation units, but we can only promise global knowledge of a 2-approximation on their number. This stage completes in $O(\text{polylog } n)$ rounds.

Following this, using standard computation of a maximum value, the nodes of the network find the node v^* in G_{in} with the highest ID and mark it as the leader node. The nodes then construct a BFS tree of G' from v^* and, using the BFS tree, v^* propagates $O(\text{polylog } n)$ unused random bits to the rest of the network. Since the diameter of the network is at most $O(\text{polylog } n)$, the above needs at most $O(\text{polylog } n)$ rounds.

Using the the propagated random bits, the nodes choose at random two hash functions, $h_A : [\bar{n}] \rightarrow [\bar{n}^{2/5}]$ and $h_S : [\bar{n}] \rightarrow [\bar{n}^{1/5}]$, uniformly and independently from a $\Theta(\log n)$ -wise independent hash function family. Let $A = \{A_1, \dots, A_{\bar{n}^{2/5}}\}$ be the sets such that $A_i = \{v \mid h_A(v) = i\}$, and let $S = \{S_1, \dots, S_{\bar{n}^{1/5}}\}$ be the sets such that $S_i = \{v \mid h_S(v) = i\}$.

Now, let P be a partition of the computation units of G_{in} into $\bar{n}^{4/5}$ sets of size at least $\bar{n}^{1/5}/4$ each. We arbitrarily index the sets in P by $P_{i,j}$ for $1 \leq i, j \leq \bar{n}^{2/5}$, and associate each set $P_{i,j}$ in P with a pair of sets (A_i, A_j) . In addition, for each $1 \leq i, j \leq \bar{n}^{2/5}$ and

each computation unit $C_{i,j,\ell}$ in $P_{i,j}$ (for ℓ in a range from 1 to a value that is at least $\bar{n}^{1/5}/4$), we associate $C_{i,j,\ell}$ with at most 4 sets in S (this is possible as there are at least $\bar{n}^{1/5}/4$ computation units in each $P_{i,j}$ and $\bar{n}^{1/5}$ sets in S). We denote the sets in S that are associated with $C_{i,j,\ell}$ by $S_{i,j,\ell,1}, \dots, S_{i,j,\ell,4}$ (the indices i, j can be omitted by being using the same 4 sets depending only on ℓ and being the same for every i, j).

The algorithm proceeds in two phases: For the first phase, consider a set of computation units $P_{i,j}$ in P , and define $E_1(i, j) = E(A_i, A_j)$. In the first phase, for each $1 \leq i, j \leq \bar{n}^{2/5}$, each core v of a computation unit in part $P_{i,j}$ of P learns all edges in $E_1(i, j)$. As h_A is a pairwise-independent hash function, for an edge $\{u, w\}$, we have $\Pr(h_A(u) = i \wedge h_A(w) = j) = (1/\bar{n}^{2/5})^2 = 1/\bar{n}^{4/5}$. Therefore, the number of edges that a core v is required to learn in expectation is $\mathbf{E}(|E_1(i, j)|) = \bar{m}/\bar{n}^{4/5} = \bar{n}^{\delta+1/5}$.

By the tail bound of Lemma 17, since we use an $\Omega(\log n)$ -wise independent hash function family (this is a logarithm of the total number of nodes n), the number of these edges is at most $\tilde{O}(\bar{n}^{\delta+1/5})$ w.h.p. (in n). Therefore, the core of a computation unit may learn these edges in $\tilde{O}(\bar{n}^{1/5})$ rounds using the routing scheme of Theorem 10 with a sufficiently small γ (since we have a degree of $O(\bar{n}^\delta)$ in G_{in} for the core using edges of this computation unit alone).

We now have that for every i, j , there are at least $\bar{n}^{1/5}/4$ computation units (those in $P_{i,j}$), whose cores know all edges in $E(A_i, A_j)$. Each such core is associated with 4 sets $S_{i,j,k,1}, \dots, S_{i,j,k,4}$, and what we would like to do in the second phase is for each core to check all edges from one of its S sets to nodes in A_i, A_j and detect a triangle. To leverage the power of the distributed Grover search, we split the task of each core into subtasks on which we can apply Lemma 4.

Formally, for the second phase, we define for each S_i an arbitrarily split of S_i into $\bar{n}^{2/5}$ batches, each of size $\tilde{O}(\bar{n}^{2/5})$ nodes, and denote these batches by $S_i^{(1)}, \dots, S_i^{(\bar{n}^{2/5})}$. We define the following protocols $\mathcal{A}_1, \dots, \mathcal{A}_{\bar{n}^{2/5}}$. The protocol \mathcal{A}_r is defined as follows: For each k such that $1 \leq k \leq \bar{n}^{1/5}$ and for each $1 \leq r \leq \bar{n}^{2/5}$, let $E_2(i, j, k, r) = \{(u, w) \mid u \in S_k^{(r)} \wedge w \in A_i \cup A_j\}$. In the protocol, each core v of a computation unit $C_{i,j,\ell}$ in $P_{i,j}$ learns all edges of $E_2(i, j, k, r)$ for at most 4 values of k which are assigned to it. Since a set $S_k^{(r)}$ has $\tilde{O}(\bar{n}^{2/5})$ nodes, the number of edges to be collected between nodes of $S_k^{(r)}$ and $A_i \cup A_j$ is in expectation $\mathbf{E}(|E_2(i, j, k, r)|) = \bar{n}^{1+\delta} \cdot (\bar{n}^{2/5}/\bar{n}) \cdot (\bar{n}^{3/5}/\bar{n}) = \tilde{O}(\bar{n}^\delta)$. By the tail bound of Lemma 17, the number of these edges is at most $\tilde{O}(\bar{n}^\delta)$ w.h.p. (in n), therefore the computation unit may learn these edges in $\tilde{O}(1)$ rounds using the routing scheme of Theorem 10. A computation unit in $P_{i,j}$ rejects if there is a triangle contained in $E_1(i, j) \cup E_2(i, j, k, r)$ for any of the 4 values of k which are assigned to it. Using the BFS tree, the network determines whether there was a computation unit that rejected, rejects if so, and otherwise accepts (notice that this propagation would not be needed in a non-quantum algorithm, as it is sufficient that one node rejects, but here we need the leader v^* to know this information).

The way these protocols are executed is as follows. In the second phase, the leader v^* performs the quantum procedure described in Lemma 4 on $\mathcal{A}_1, \dots, \mathcal{A}_{\bar{n}^{2/5}}$. Therefore, by Lemma 4, the quantum protocol rejects if and only if the network rejects in at least one of these procedures, and terminates after $\tilde{O}(\sqrt{\bar{n}^{2/5}}) = \tilde{O}(\bar{n}^{1/5}) = \tilde{O}(n^{1/5})$ rounds.

Correctness. If the graph is triangle-free then the network accepts, as a core v of a computation unit only rejects if it detects a triangle in the edges $E_1(i, j) \cup E_2^r(k)$ for some appropriate parameters i, j, r, k . Otherwise, let $\{v_1, v_2, v_3\}$ be vertices of a triangle in G' . Let $i, j \in [\bar{n}^{2/5}]$ such that $v_1 \in A_i, v_2 \in A_j$ and let $P_{i,j}$ be the set of computation units to which the pair A_i, A_j is mapped. Let $k \in [\bar{n}^{1/5}]$ and $r \in [\bar{n}^{2/5}]$ such that $v_3 \in S_k^{(r)}$ and let

$w \in P_{i,j}$ be the core of a computation unit in $P_{i,j}$ to which S_k is mapped. We note that the triangle is contained in $E_1(i, j) \cup E_2(i, j, k, r)$, and therefore the network rejects in the protocol \mathcal{A}_r , and hence our algorithm rejects as well. \blacktriangleleft

4.2 Quantum Detection of K_p for $p \geq 7$

We now show our method for enhancing the K_p -listing algorithm of [4] with a quantum procedure that allows us to detect a K_{p+1} instance. To ease the notation, we will be using $p + 1$ for the size of the clique that we are detecting, and hence note that the following statement is shifted, i.e., holds for $p \geq 6$.

► Theorem 18. *There exists a quantum algorithm that solves $(p + 1)$ -clique detection with success probability at least $1 - 1/\text{poly}(n)$ in the QUANTUM CONGEST model with complexity $\tilde{O}(n^{1-2/p})$, for $p \geq 6$.*

In a nutshell, the algorithm uses a conductance decomposition, and works on clusters in parallel. There are three ways in which the K_p -listing algorithm of [4] may find instances of K_p . The first two ways involve having a node learn its induced 2-hop neighborhood, and in these cases, this clearly gives detection of a K_{p+1} instance if such an instance exists, without further effort on our part. The third case is where we diverge from the algorithm of [4] by incorporating Grover searches inside the clusters, using Lemma 4.

During the algorithm, some parts use a simple listing of the edges in a node's induced 2-hop neighborhood. For completeness, we state here a formal claim and proof of how this is done.

▷ Claim 19. Given a graph $G = (V, E)$, and some value α , every node v such that $\deg(v) \leq \alpha$ can learn its induced 2-hop neighborhood in at most $O(\alpha)$ rounds of the CONGEST model.

Proof of Claim 19. Let v be such a node. Node v iterates over its at most α neighbors. When it considers neighbor i , it sends a message to all its neighbors asking if they neighbor i , to which they each respond whether or not they have an edge to node i . At the end of the iterations, after $O(\alpha)$ rounds, node v knows all of its induced 2-hop neighborhood.

Notice that it does not matter if v has a neighbor u which also tries to perform this search in parallel to v , in the odd rounds v and u just sends queries to each other across their shared edge, and in the even rounds they each respond to one another. \blacktriangleleft

We are now ready to prove Theorem 18.

Proof of Theorem 18. We begin by briefly explaining the K_p listing algorithm of [4]. The algorithm of [4] works in iterations, which are composed of stages. In every stage, some K_p in the graph may be listed. We split the listed K_p instances according to the stage in the algorithm in which they are listed. For each stage, we show how to detect in a quantum manner if there is an instance of K_p that is listed in that stage which can be extended to (i.e., is a part of) an instance of K_{p+1} . As such, if there is any instance of K_{p+1} in the graph, then we will certainly detect this as some instance of K_p could be extended to it.

We now show the K_p listing algorithm, and interject at appropriate places in order to perform quantum searches. Each iteration of the algorithm of [4] consists of the following steps.

1. **Exhaustive search detection stage.** Every node v with degree $O(n^{1/2})$ learns its induced 2-hop neighborhood in $O(n^{1/2})$ rounds using Claim 19, lists any cliques which it sees involving itself, and removes itself, along with its incident edges from the graph.

Detection of K_{p+1} : In particular, because the entire induced 2-hop neighborhood is learned, any instance of K_p that is listed that can be extended to an instance of K_{p+1} can be immediately listed in this way even without further communication, which proves our claim for this stage.

2. **Graph decomposition and cluster exhaustive search detection stage.** An expander decomposition is computed according to Lemma 9, with $\gamma = 0.1$ (any small constant would do here), while the above ensures that the average degree in the graph (and thus in the clusters, due to Lemma 9) is at least $\Omega(n^{1/2})$. In clusters with $O(n^{1-2/p})$ nodes, each node learns its induced 2-hop neighborhood in G using Claim 19 and lists all the instances of K_p which it is a part of, and finally removes itself and its incident edges from the graph. Because Lemma 9 promises that the number of edges that leave C is at most $\tilde{O}(|E_C|)$, where E_C is the set of edges inside C , then this completes in $\tilde{O}(n^{1-2/p})$ rounds.

Detection of K_{p+1} : As in the previous stage, because of the exhaustive search nature of this stage, we again obtain that any instance of K_p that is listed that can be extended to an instance of K_{p+1} can be immediately listed in this way even without further communication, which proves our claim for this stage.

3. **Cluster listing.** Within every remaining cluster C , some nodes are designated as *good* and the rest as *bad*. For a good node v , the nodes of C learn amongst themselves (collectively, not necessarily by a single node, and in particular not necessarily by v itself) all the edges in the induced 2-hop neighborhood of v . That is, the nodes of C request from the nodes neighboring the cluster to send in edges from outside C to nodes inside C , such that all the induced 2-hop neighborhood of v is known to the nodes in C . Now, the nodes in C ensure that every instance of K_p involving v and at least one other node in C is becomes known to some node in C – that is, these K_p instances are listed by C . The good nodes are then removed from the graph along with all of their edges.

This concludes the description of the CONGEST algorithm in [4]. The exact definition of good nodes is not required for our purpose. The proof of [4] shows that each iteration completes within $\tilde{O}(n^{1-2/p})$ rounds, and that after $\text{poly log}(n)$ iterations the remaining graph is empty and hence $\tilde{O}(n^{1-2/p})$ rounds are sufficient for the entire algorithm.

The cluster quantum detection stage. We now provide the quantum procedure that we run after the above **Cluster listing** stage, in order to detect an instance of K_{p+1} that contains an instance of K_p that was listed during this stage. First, let the C -degree of a node be its number of neighbors in C , let μ_C be the average C -degree of nodes in C , and let $H(C)$ be the high C -degree cluster nodes, namely, those of C -degree $\Omega(\mu_C)$. Notice that due to the invocation above of Lemma 9, it holds that $\mu_C = \Omega(n^{1/2})$. A more precise description of the algorithm in [4] is that the induced 2-hop neighborhood of each good node v becomes known to the nodes in $H(C)$ rather than to any node in C . Second, denote by $N^+(C)$ the set of nodes which have a neighbor in C (this includes all of C as C is connected, by definition) and by $N(C) = N^+(C) - C$ the neighbors of C outside of C . In the above **Cluster listing** stage we piggyback the degree (not C -degree) of every node in $N(C)$ onto the messages containing its edges, so that the degrees of all nodes in $N(C)$ are also known to nodes of $H(C)$.

The quantum process starts with the following. The nodes in C elect some arbitrary leader $v_C \in H(C)$. This can be done in $\text{poly log}(n)$ rounds because this is a bound on the diameter of the cluster, by Lemma 9. The nodes of $H(C)$ now broadcast within $H(C)$ the degrees of all the nodes in $N(C)$. This is done in two steps. In the first step, each node in $H(C)$ sends v_C the degrees of nodes in $N(C)$ which it knows about. As every node sends

and receives at most $|N(C)| = O(n)$ pieces of information, and every node in $H(C)$ has C -degree at least $\Omega(\mu_C) = \Omega(n^{1/2})$, this takes at most $\tilde{O}(n^{1/2})$ rounds, using the routing algorithm of Theorem 10 with a sufficiently small γ . In the second step, v_C makes this information known to all nodes in $H(C)$ using a simple doubling procedure: in each phase of this procedure, each informed node shares the information with a unique uninformed node. Each phase completes in $\tilde{O}(n^{1/2})$ rounds by the same argument using the routing algorithm of Theorem 10, and the number of phases is logarithmic. In a similar fashion, the nodes in $H(C)$ learn all the degrees of the nodes in C .

Knowing the degrees of all nodes in $N^+(C)$, the nodes $H(C)$ bucket $N^+(C)$ by degrees. That is, they compute $N^+(C) = N_1^+(C), \dots, N_{\log n}^+(C)$, such that the degree of any $v \in N_i^+(C)$ is in $[2^{i-1}, \dots, 2^i]$. Our goal is to iterate over the $\log n$ buckets, whereby in each bucket we check whether a node v in $N_i^+(C)$ can be used to extend some K_p instance, which is already listed by $H(C)$, into an instance of K_{p+1} . This is done as follows.

Fix an i between 1 and $\log n$. We perform a Grover search over $v \in N_i^+(C)$ using v_C as the leader of the search: In each query (i.e., over such a v), we broadcast within $H(C)$ all the neighbors of v which are known to $H(C)$, in order to try to extend any K_p listed by $H(C)$ to a K_{p+1} involving v . Let κ_p be such a K_p . As it is listed by $H(C)$, it must involve at least one good node in C , denoted $g \in \kappa_p$. Further, recall that every edge in the induced 2-hop neighborhood of g is known to $H(C)$. Thus, if v that can be used to extend κ_p to an instance of K_{p+1} , all the edges between v and κ_p are known to the nodes of $H(C)$. Therefore, if we manage to broadcast in $H(C)$ all the edges incident to v that are known to $H(C)$, then if there is a way to use v to extend a K_p listed by $H(C)$ to a K_{p+1} , we will certainly find it. Similarly to the above analysis, using the routing algorithm of Theorem 10 and the doubling procedure, we broadcast within $H(C)$ the edges incident to v which are known to $H(C)$. This completes in $O(2^i/\mu_C + 1)$ rounds, as the degree of v is at most 2^i .

In order to conclude the proof, we need to bound the size of each $N_i^+(C)$. In order to obtain the bound on the size of $N_i^+(C)$, we compute an upper bound on the number of *all* edges incident to nodes in $N^+(C)$ and use the bound on degrees of nodes in the bucket. To this end, we wish to show that $\sum_{v \in N^+(C)} \deg(v) = \tilde{O}(n \cdot \mu_C)$. To do so, we split the edges into three categories: E_1 – edges with both endpoints in C ; E_2 – edges with one endpoint in C ; E_3 – edges with both endpoints not in C . As μ_C is the average C -degree of the nodes in C , and n is the number of nodes in the entire graph, implying $|C| \leq n$, it holds that $E_1 = O(n \cdot \mu_C)$. It is ensured in Lemma 9 that for every $v \in C$, it holds that $\deg(v) = \tilde{O}(\deg_C(v))$, where $\deg_C(v)$ is the C -degree of v , implying $|E_2| = \tilde{O}(|E_1|) = \tilde{O}(n \cdot \mu_C)$. Further, it is ensured in Lemma 9 that $|E_3|/n = \tilde{O}(|E_1|/|C|) = \tilde{O}(\mu_C)$, implying that $|E_3| = \tilde{O}(n \cdot \mu_C)$. Thus, for every $N_i^+(C)$, due to the degrees in the bucket and due to the pigeonhole principle, it must be that $|N_i^+(C)| = \tilde{O}(n \cdot \mu_C / 2^i)$. Further, as n is the number of nodes in the graph, it also trivially holds that $|N_i^+(C)| \leq n$. All in all, we get that $|N_i^+(C)| = \tilde{O}(\min\{n, n \cdot \mu_C / 2^i\})$.

Finally, we can analyze the round complexity of the algorithm. The K_p listing algorithm is shown in [4] to take $\tilde{O}(n^{1-2/p})$ rounds. For each i , the checking procedure in our Grover search takes $\tilde{O}(2^i/\mu_C + 1)$ rounds. For every i such that $2^i/\mu_C \leq 1$, the checking procedure takes $O(1)$, and $|N_i^+(C)| = \tilde{O}(n)$, implying that the Grover search takes a total of $\tilde{O}(n^{1/2})$ rounds, using Lemma 4. For every i such that $2^i/\mu_C > 1$, the checking procedure takes $\tilde{O}(2^i/\mu_C + 1) = \tilde{O}(2^i/\mu_C)$, and $|N_i^+(C)| = \tilde{O}(n \cdot \mu_C / 2^i)$, implying that the Grover search takes a total of $\tilde{O}(\sqrt{n \cdot \mu_C / 2^i} \cdot (2^i/\mu_C)) = \tilde{O}(\sqrt{n \cdot (2^i/\mu_C)})$ rounds, using Lemma 4. As $2^i = O(n)$ and $\mu_C = \Omega(n^{1/2})$ (due to the guarantees of the invocation of Lemma 9, as stated above), this takes at most $\tilde{O}(n^{3/4})$ rounds.

We thus have that our quantum algorithm for detecting an instance of K_{p+1} for $p \geq 5$ completes in $\tilde{O}(n^{3/4} + n^{1-2/p})$ rounds, with probability at least $1 - 1/\text{poly}(n)$.

We note that we can decrease the $\tilde{O}(n^{3/4})$ part of the complexity, by slightly changing the listing algorithm of [4]. That is, the algorithm stated there performs K_p listing in $\tilde{O}(n^{1-2/p})$ rounds, while ensuring $\mu_C = \Omega(n^{1/2})$. It is implied in the proofs in [4] that, for any $1/2 \leq \delta < 1$, one can pay an additional $\tilde{O}(n^\delta)$ rounds (by performing an exhaustive search, as done in Claim 19) in order to ensure $\mu_C = \Omega(n^\delta)$. Further, notice that our Grover searches take a total of $\tilde{O}(n^{1/2} + n/\mu_C^{1/2}) = \tilde{O}(n^{1/2} + n^{1-\delta/2})$ rounds.

All in all, for any $1/2 \leq \delta < 1$, our algorithm requires $\tilde{O}(n^{1-2/p} + n^\delta + n^{1/2} + n^{1-\delta/2}) = \tilde{O}(n^{1-2/p} + n^\delta + n^{1-\delta/2})$ rounds. One can set $\delta = 2/3$, giving a final algorithm running in $\tilde{O}(n^{1-2/p} + n^{2/3} + n^{1-1/3}) = \tilde{O}(n^{1-2/p} + n^{2/3}) = \tilde{O}(n^{1-2/p})$ rounds, where the last transition is since $p \geq 6$. ◀

5 Faster Clique Detection in the Quantum Congested Clique Model

We present here an approach that improves upon our approach in Section 3 by taking into consideration *which* clique nodes know of which copies of K_p : instead of treating the K_p -listing algorithm as a black box, we explicitly use the K_p -listing algorithm of [12], so that we know which copies of K_p will be listed by each clique node, and what other information that node already has. We then search for $(p+t)$ -cliques by having each clique node learn only the edges that it needs to check if the p -cliques it has listed can be extended. (In the previous approach, edges could be learned by nodes that had no use for them, since they were not adjacent to any p -clique that the node had listed.)

Throughout this section we use the following notation: we let $\times(S_1, \dots, S_k)$ denote the Cartesian product $S_1 \times \dots \times S_k$. Also, we denote by $E(S_1, \dots, S_k)$ the edges $E \cap \bigcup_{i \neq j} (S_i \times S_j)$ that cross between any two sets S_i, S_j .

For presenting our algorithm for a general p , we first shortly review the listing algorithm of Dolev et al. [12]. Fix an arbitrary partition $S_1, \dots, S_{n^{1/p}}$ of the nodes of V , such that $|S_i| = n^{1-1/p}$ for each i , and a 1:1 mapping $g: V \rightarrow \{1, \dots, n^{1/p}\}^p$ assigning to each node $v \in V$ a p -tuple $g(v) = (i_1, \dots, i_p) \in \{1, \dots, n^{1/p}\}^p$. We assume for simplicity that $n^{1/p}$ is an integer; otherwise, we can adjust the set size to an integer, without changing the asymptotic complexity of the algorithm.

For each node $v \in V$ and index $j \in \{1, \dots, p\}$, let $T_j^v = S_{g(v)_j}$. (That is, if $g(v) = (i_1, \dots, i_p)$, then $T_1^v = S_{i_1}, \dots, T_p^v = S_{i_p}$.) Also, let $T^v := \bigcup_{j=1}^p T_j^v$. Node v is responsible for listing all p -cliques $(u_1, \dots, u_p) \in \times(T_1^v, \dots, T_p^v)$. To do so, node v needs to learn all edges in $E(T_1^v, \dots, T_p^v)$; there are at most $p(n^{1-1/p})^2 = O(n^{2-2/p})$ such edges, which, using Lenzen's routing scheme, can be collected in $O(n^{2-2/p-1}) = O(n^{1-2/p})$ rounds. After collecting the edges in $E(T_1^v, \dots, T_p^v)$, node v locally enumerates all p -cliques it sees.

5.1 K_{p+1} -Detection from K_p -Listing

We start by showing how to use the K_p listing algorithm for obtaining K_{p+1} -detection, stated as follows.

► **Theorem 20.** *For any $p \geq 3$, the K_{p+1} -detection problem in the QUANTUM CONGESTED CLIQUE can be solved in $\tilde{O}(n^{(1-1/p)/2} + n^{1-2/p})$ rounds.*

Proof. First, we let each node $v \in V$ collect all edges in $E(T^v \times T^v)$, and list all cliques in $\times(T_1^v, \dots, T_p^v)$, as described above.

Our goal now is for each node $v \in V$ to check whether there is a node $u \in V$ that extends a p -clique found by node v in the previous step into a $(p+1)$ -clique. In other words, node v searches for $(p+1)$ -cliques in $\times(T_1^v, \dots, T_p^v, V)$.

To speed up the search, we partition the nodes of V into $n^{1-1/p}$ batches, $Q_1, \dots, Q_{n^{1-1/p}}$, each of size $n^{1/p}$. We then use a Grover search (Lemma 4) to find an index i such that $\times(T_1^v, \dots, T_p^v, Q_i)$ contains a $(p+1)$ -clique (or determine that there is no such i).

We denote by A_i a query for determining whether $\times(T_1^v, \dots, T_p^v, Q_i)$ contains a $(p+1)$ -clique, and we now describe how each query is implemented. In A_i , node v needs to collect all edges in $E(T, Q_i)$; recall that $T = \bigcup_{j=1}^p T_j^v$, and its size is $|T| \leq p \cdot n^{1-1/p}$. Since $|Q_i| = n^{1/p}$, the number of edges v needs to learn about is $n^{1-1/p} \cdot n^{1/p} = n$, and using Lenzen's routing scheme, this can be done in a single round. After learning the relevant edges, node v locally searches for a $(p+1)$ -clique contained in the edges it has learned; it outputs “yes” if and only if it finds one.

The running time of the resulting algorithm is as follows. Executing the K_p -listing algorithm from [12] requires $O(n^{1-2/p})$ rounds. The Grover search requires $\sqrt{n^{1-1/p}}$ queries, each requiring a single round. Using Lemma 4, the overall running time is therefore $\tilde{O}(n^{(1-p)/2})$ for the quantum search, in addition to the $O(n^{1-2/p})$ rounds required for K_p -listing. Specifically, for K_4 -detection ($p = 3$), the total running time we obtain is $\tilde{O}(n^{1-2/3} + n^{1/2-1/6}) = \tilde{O}(n^{1/3})$. \blacktriangleleft

Theorem 20 gives, combined with the algorithm of Section 3.1, the statement of Theorem 1 in the introduction.

5.2 K_{p+t} -Detection from K_p -Listing

Now, we prove the following theorem, which shows how to obtain K_{p+t} -detection from K_p -listing.

► **Theorem 21.** *Given a graph $G = (V, E)$ and values p, t such that $t \leq 1 + \log(p-1)$, it is possible to detect if G contains an instance of K_{p+t} in $\tilde{O}(n^{(1-1/p)(1-1/2^t)} + n^{1-2/p})$ rounds of communication in the QUANTUM CONGESTED CLIQUE, w.h.p.*

Optimizing our results over all choices of p, t yields the following:

► **Corollary 22.** *The time required to solve K_{p+t} -detection in QUANTUM CONGESTED CLIQUE is*

$$\tilde{O}\left(\min_{p, t: t \leq 1 + \log(p-1)} \left(n^{(1-1/p)(1-1/2^t)} + n^{1-2/p}\right)\right).$$

Proof of Theorem 21.

High-level overview. We again first let each node $v \in V$ collect all edges in $E(T^v \times T^v)$, and list all cliques in $\times(T_1^v, \dots, T_p^v)$. Now we want to solve $(p+t)$ -clique detection, having already listed all p -cliques using the algorithm of [12].

To do this, we generalize the +1 extension from the previous section using a *recursive* search procedure, $\text{check}_k(P_1, \dots, P_k, F)$ where $k \leq p+t$. The procedure is given k sets of nodes, $P_1, \dots, P_k \subseteq V$, and the set F of all edges in $E(P_1, \dots, P_k)$; it recursively checks whether there is a k -clique $(u_1, \dots, u_k) \in \times(P_1, \dots, P_k)$ that can be extended into a $(p+t)$ -clique. The recursion begins at $k = p$, with the sets $P_1 = T_1^v, \dots, P_p = T_p^v$, and the edges that node v collected in the pre-processing stage; the recursion terminates at $k = p+t$, where node v enumerates all $(p+t)$ -tuples $(u_1, \dots, u_{p+t}) \in P_1 \times \dots \times P_{p+t}$ and checks using the edges of F whether one of them forms a $(p+t)$ -clique. Next, we explain how each internal level of the recursion is implemented. Fix t predetermined partitions of the nodes of V , where the i -th partition is given by $X_i^1, \dots, X_i^{n^{r_i}}$ for a parameter $r_i \in (0, 1)$, and all subsets in each partition

have the same size: $|X_i^1| = \dots = |X_i^{n^{r_i}}| = n^{1-r_i}$. (We assume again for simplicity that n^{r_i} is an integer.) For $i \in \{0, \dots, t-1\}$, when $\text{check}_{p+i}(P_1, \dots, P_{p+i}, F)$ is called, node v uses a Grover search to check if there is some set X_{i+1}^j such that $\text{check}_{p+i}(P_1, \dots, P_{p+i}, X_{i+1}^j, F')$ returns true, where $F' = E(P_1, \dots, P_{p+i}, X_{i+1}^j)$. Each query A_j in the Grover search takes the index $j \in [n^{r_i}]$ as input, learns all the edges in $F' \setminus F$ (the edges of F are already known), and calls $\text{check}_{p+i+1}(P_1, \dots, P_{p+i}, X_{i+1}^j, F')$.

The formal algorithm. Fix $p \geq 2$ and $t \geq 1$ satisfying the constraint. Also, for each $i = 1, \dots, t$, fix a partition $X_i^1, \dots, X_i^{n^{r_i}}$ of V , where

$$r_i = (1 - 1/p)/2^{t-i},$$

and $|X_i^1| = \dots = |X_i^{n^{r_i}}| = n^{1-r_i}$. The partitions are arbitrary, but fixed in advance, so no communication is necessary to compute them. Note that $r_i \geq 1/p$ for each $i = 1, \dots, t$.

Our algorithm is a depth- t nested quantum search, over the search space $[n^{r_1}] \times \dots \times [n^{r_t}]$, with the goal function $f : [n^{r_1}] \times \dots \times [n^{r_t}] \rightarrow \{0, 1\}$ such that $f(i_1, \dots, i_t) = 1$ iff there exists a p -clique $v_1, \dots, v_p \in V^p$, and there exist $u_1 \in X_1^{i_1}, \dots, u_t \in X_t^{i_t}$, such that $v_1, \dots, v_p, u_1, \dots, u_t$ is a $(p+t)$ -clique.

Before the quantum search begins, the nodes list all p -cliques in the graph; each node v learns all edges in $E(T_1^v, \dots, T_p^v)$ as described in the beginning of the section.

The setup algorithms $\mathcal{S}_1, \dots, \mathcal{S}_{t-1}$ are as follows: in \mathcal{S}_i , the leader disseminates the current partial search query $(j_1, \dots, j_i) \in [n^{r_1}] \times \dots \times [n^{r_i}]$. The nodes use Lenzen's routing scheme (Lemma 6) so that each $v \in V$ learns the edges $E(T_1^v \cup \dots \cup T_p^v \cup X_1^{j_1} \cup \dots \cup X_{i-1}^{j_{i-1}}, X_i^{j_i})$. The running time s_i of \mathcal{S}_i is $O(n^{1-1/p-r_i})$: we have $|T_1^v \cup \dots \cup T_p^v| = O(n^{1-1/p})$ (treating p as a constant), $|X_1^{j_1} \cup \dots \cup X_{i-1}^{j_{i-1}}| = O(n^{1-1/p})$ (by our assumption that $r_i \geq 1/p$ and treating t as a constant), and $|X_i^{j_i}| = O(n^{1-r_i})$. By Lemma 6, the information can be routed in $O(n^{1-1/p+1-r_i-1}) = O(n^{1-1/p-r_i})$ rounds.

The final classical evaluation procedure, \mathcal{C} , is as follows: the leader disseminates the query $(j_1, \dots, j_t) \in [n^{r_1}] \times \dots \times [n^{r_t}]$ to all nodes, and the nodes use Lenzen's routing scheme so that each $v \in V$ learns the edges $E(T_1^v \cup \dots \cup T_p^v \cup X_1^{j_1} \cup \dots \cup X_{t-1}^{j_{t-1}}, X_t^{j_t})$. Together with the edges in $\bigcup_{i < t} \text{setup}_i^t(j_1, \dots, j_i)$, each node v now knows $E(T_1^v, \dots, T_p^v, X_1^{j_1}, \dots, X_t^{j_t})$. Node v now checks whether it sees a $(p+t)$ -clique, and informs the leader. The running time of \mathcal{C} is $O(n^{1-1/p-r_t})$ (as above).

By Lemma 5, the overall running time of the quantum search is given by

$$\tilde{O}\left(n^{r_1/2} \left(n^{1-1/p-r_1} + n^{r_2/2} \cdot \left(n^{1-1/p-r_2} + \dots + n^{r_{t-1}/2} \left(n^{1-1/p-r_{t-1}} + n^{r_t/2} \cdot n^{1-1/p-r_t} \right) \right) \right) \right).$$

Denote by $g(i)$ the running time of the innermost $t-i+1$ levels of the search, starting from level i up to level t :

$$g(i) = \tilde{O}\left(n^{r_i/2} \left(n^{1-1/p-r_i} + \dots + n^{r_{t-1}/2} \left(n^{1-1/p-r_{t-1}} + n^{r_t/2} \cdot n^{1-1/p-r_t} \right) \right) \right).$$

We claim, by backwards induction on i , that

$$g(i) = \tilde{O}(n^{(1-1/p)(1-1/2^{t-i+1})}). \quad (1)$$

For the base case, $i = t$, we have

$$g(i) = \tilde{O}\left(n^{r_t/2} \cdot n^{1-1/p-r_t}\right) = \tilde{O}\left(n^{1-1/p-r_t/2}\right) = \tilde{O}\left(n^{1-1/p-(1-1/p)/2}\right) = \tilde{O}\left(n^{(1-1/p)/2}\right),$$

and indeed this matches (1). For the induction step, suppose the claim holds for i , and consider $g(i-1)$:

$$\begin{aligned} g(i-1) &= \tilde{O}\left(n^{r_{i-1}/2}\left(n^{1-1/p-r_{i-1}} + g(i)\right)\right) \\ &= \tilde{O}\left(n^{(1-1/p)/2^{t-i+1}/2}\left(n^{1-1/p-(1-1/p)/2^{t-i+1}} + n^{(1-1/p)(1-1/2^{t-i+1})}\right)\right) \\ &= \tilde{O}\left(n^{(1-1/p)/2^{t-i+2}} \cdot 2n^{(1-1/p)(1-1/2^{t-i+1})}\right) \\ &= \tilde{O}\left(n^{(1-1/p)(1-1/2^{t-i+2})}\right), \end{aligned}$$

again matching (1).

We thus obtain that the running time of the entire quantum search is

$$g(1) = \tilde{O}\left(n^{(1-1/p)(1-1/2^t)}\right),$$

and together with the complexity of listing the initial p -clique, we get the claimed complexity. \blacktriangleleft

A note on the restriction on p, t . Recall that we assumed throughout that $1 - r_i \leq 1 - 1/p$ for each $i = 1, \dots, t$, so that the sets of nodes we handle at each step never exceed the size of the sets in the p -clique-listing pre-processing step ($n^{1-1/p}$). In other words, we must have $r_i \geq \frac{1}{p}$. The value of r_i decreases with i , so it suffices to require that $r_1 \geq 1/p$. We now see that not every choice of p, t respects this condition: since we set $r_1 = \frac{1}{2^{t-1}}\left(1 - \frac{1}{p}\right)$, we require that p, t satisfy $t \leq 1 + \log(p-1)$.

For example, suppose we want to start by listing all triangles ($p=3$), and extend to K_5 (i.e., $t=2$). This is possible, since $2 \leq 1 + \log(3-1)$. However, extending from triangles to K_6 using the approach described here is not possible, because if we take $p=3$ and $t=3$, we get $3 \not\leq 1 + \log(3-1)$.

Although we think of p, t as constants, we observe that when p is large and $t \approx \log(p)$, the overall running time of the quantum part of our scheme is $\approx n^{1-2/p}$. Thus, the cost of extending from p -cliques to $(p+t)$ -cliques roughly matches the cost of the classical p -clique listing, which is $n^{1-2/p}$, and we get a quantum algorithm for K_{p+t} -detection that roughly matches the cost of classical K_p -listing.

6 Circuit-Complexity Barrier to Proving an $\Omega(n^{3/5+\alpha})$ -Round Lower Bound for Clique Detection in the Congest Model

In this section, we show that for any $\alpha > 0$, a lower bound of the form $\Omega(n^{3/5+\alpha})$ for K_p -detection in (non-quantum) CONGEST would imply strong circuit complexity results, far beyond the current state-of-the-art. As mentioned in the introduction, this barrier also applies to the QUANTUM CONGEST model.

Given a constant integer $\alpha > 0$, let \mathcal{F}_α be the family of Boolean circuits F where (a) F has $M \log M$ input wires for some integer M , which we interpret as encoding a graph \bar{G} on $O(M)$ edges, (b) All gates in F have constant fan-in and fan-out, (c) F has depth at most $R = M^\alpha$, and (d) F has at most $M^{1+\alpha}$ wires in total.

► **Theorem 23.** *If K_p -detection has a round complexity of $\Omega(n^{3/5+\alpha})$ for some constant $\alpha > 0$, then there is no circuit family contained in \mathcal{F}_α that solves K_p -detection.*

In other words, a lower bound of the form $\Omega(n^{3/5+\alpha})$ for any constant $\alpha > 0$ rules out the existence of a circuit with polynomial depth and super-linear size (in terms of wires) for an explicit problem, K_p -detection. Such a lower bound would be a major breakthrough in circuit complexity (see [13] for discussion). This gives the statement of Theorem 3 in the introduction.

The proof of Theorem 23 is essentially a reduction from K_p -detection in general graphs to K_p -detection in high-conductance graphs. It was shown in [13] that in high-conductance graphs, we can efficiently *simulate* a circuit from the family \mathcal{F}_α of certain size, so given such a circuit for K_p -detection, we can use it to solve the distributed K_p -detection problem in high-conductance graphs. To reduce from general graphs to high-conductance graphs, we use a similar approach to the clique detection algorithm of [4]: we first compute an expander-decomposition procedure from to partition the graph into high-conductance clusters with few inter-cluster edges, and have the nodes of the clusters learn the relevant sets of edges in the clusters' neighborhood, so that we either find a K_p -copy or can remove most of edges of the graph without removing a K_p -copy, and recurse on the remaining edges. We defer this result to the full version of the paper for lack of space.

References

- 1 Amir Abboud, Keren Censor-Hillel, Seri Khoury, and Christoph Lenzen. Fooling views: a new lower bound technique for distributed computations under congestion. *Distributed Computing*, 33(6):545–559, 2020.
- 2 Heger Arfaoui and Pierre Fraigniaud. What can be computed without communications? *SIGACT News*, 45(3):82–104, 2014.
- 3 Michael Ben-Or and Avinatan Hassidim. Fast quantum byzantine agreement. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, pages 481–485, 2005.
- 4 Keren Censor-Hillel, Yi-Jun Chang, François Le Gall, and Dean Leitersdorf. Tight distributed listing of cliques. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2021)*, pages 2878–2891, 2021.
- 5 Keren Censor-Hillel, Orr Fischer, Tzlil Gonen, François Le Gall, Dean Leitersdorf, and Rotem Oshman. Fast distributed algorithms for girth, cycles and small subgraphs. In *Proceedings of the 34th International Symposium on Distributed Computing (DISC 2020)*, pages 33:1–33:17, 2020.
- 6 Keren Censor-Hillel, François Le Gall, and Dean Leitersdorf. On distributed listing of cliques. In *Proceedings of the 39th ACM Symposium on Principles of Distributed Computing (PODC 2020)*, pages 474–482, 2020.
- 7 Keren Censor-Hillel, Petteri Kaski, Janne H. Korhonen, Christoph Lenzen, Ami Paz, and Jukka Suomela. Algebraic methods in the congested clique. *Distributed Computing*, 32(6):461–478, 2019.
- 8 Yi-Jun Chang, Seth Pettie, and Hengjie Zhang. Distributed triangle detection via expander decomposition. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2019)*, pages 821–840, 2019.
- 9 Yi-Jun Chang and Thatchaphol Saranurak. Improved distributed expander decomposition and nearly optimal triangle enumeration. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing (PODC 2019)*, pages 66–73, 2019.
- 10 Artur Czumaj and Christian Konrad. Detecting cliques in CONGEST networks. *Distributed Computing*, 33(6):533–543, 2020.
- 11 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.

- 12 Danny Dolev, Christoph Lenzen, and Shir Peled. “tri, tri again”: Finding triangles and small subgraphs in a distributed setting. In *Proceedings of the 26th International Symposium on Distributed Computing (DISC 2012)*, pages 195–209, 2012.
- 13 Talya Eden, Nimrod Fiat, Orr Fischer, Fabian Kuhn, and Rotem Oshman. Sublinear-time distributed algorithms for detecting small cliques and even cycles. In *Proceedings of the 33rd International Symposium on Distributed Computing (DISC 2019)*, pages 15:1–15:16, 2019.
- 14 Talya Eden, Nimrod Fiat, Orr Fischer, Fabian Kuhn, and Rotem Oshman. Sublinear-time distributed algorithms for detecting small cliques and even cycles (journal version). To appear in *Distributed Computing*, 2021.
- 15 Michael Elkin, Hartmut Klauck, Danupon Nanongkai, and Gopal Pandurangan. Can quantum communication speed up distributed computation? In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing (PODC 2014)*, pages 166–175, 2014.
- 16 Orr Fischer, Tzlil Gonen, Fabian Kuhn, and Rotem Oshman. Possibilities and impossibilities for distributed subgraph detection. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures (SPAA 2018)*, pages 153–162, 2018.
- 17 Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)*, pages 1150–1162, 2012.
- 18 Cyril Gavoille, Adrian Kosowski, and Marcin Markiewicz. What can be observed locally? In *Proceedings of the 23rd International Symposium on Distributed Computing (DISC 2009)*, pages 243–257, 2009.
- 19 Mohsen Ghaffari, Fabian Kuhn, and Hsin-Hao Su. Distributed MST and routing in almost mixing time. In *Proceedings of the 2017 ACM Symposium on Principles of Distributed Computing (PODC 2017)*, pages 131–140, 2017.
- 20 Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing (STOC 1996)*, pages 212–219, 1996.
- 21 Taisuke Izumi and François Le Gall. Quantum distributed algorithm for the all-pairs shortest path problem in the CONGEST-CLIQUE model. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing (PODC 2019)*, pages 84–93, 2019.
- 22 Taisuke Izumi, François Le Gall, and Frédéric Magniez. Quantum distributed algorithm for triangle finding in the CONGEST model. In *Proceedings of the 37th International Symposium on Theoretical Aspects of Computer Science (STACS 2020)*, pages 23:1–23:13, 2020.
- 23 Taisuke Izumi and François Le Gall. Triangle finding and listing in CONGEST networks. In *Proceedings of the 2017 ACM Symposium on Principles of Distributed Computing (PODC 2017)*, pages 381–389, 2017.
- 24 Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM journal on computing*, 18(6):1149–1178, 1989.
- 25 François Le Gall and Frédéric Magniez. Sublinear-time quantum computation of the diameter in CONGEST networks. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing (PODC 2018)*, pages 337–346, 2018.
- 26 François Le Gall, Harumichi Nishimura, and Ansis Rosmanis. Quantum advantage for the LOCAL model in distributed computing. In *Proceedings of the International Symposium on Theoretical Aspects of Computer Science (STACS 2019)*, pages 49:1–49:14, 2019.
- 27 Christoph Lenzen. Optimal deterministic routing and sorting on the congested clique. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing (PODC 2013)*, pages 42–50, 2013.
- 28 Gopal Pandurangan, Peter Robinson, and Michele Scquizzato. On the distributed complexity of large-scale graph computations. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures (SPAA 2018)*, pages 405–414, 2018.

- 29 Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff-Hoeffding bounds for applications with limited independence. In *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 1993)*, pages 331–340, 1993.
- 30 Seiichiro Tani, Hirota Kobayashi, and Keiji Matsumoto. Exact quantum algorithms for the leader election problem. *ACM Transactions on Computation Theory*, 4(1):1:1–1:24, 2012.