

Even-Cycle Detection in the Randomized and Quantum CONGEST Model*

Pierre Fraigniaud¹, Maël Luce¹, Frédéric Magniez¹, and Ioan Todinca²

¹ Université Paris Cité, CNRS, IRIF, Paris, France

² Université d'Orléans, INSA-Centre Val de Loire, LIFO, Orléans, France

Abstract

We show that, for every $k \geq 2$, C_{2k} -freeness can be decided in $O(n^{1-1/k})$ rounds in the CONGEST model by a randomized Monte-Carlo distributed algorithm with one-sided error probability $1/3$. This matches the best round-complexities of previously known algorithms for $k \in \{2, 3, 4, 5\}$ by Drucker et al. [PODC'14] and Censor-Hillel et al. [DISC'20], but improves the complexities of the known algorithms for $k > 5$ by Eden et al. [DISC'19], which were essentially of the form $\tilde{O}(n^{1-2/k^2})$. Our algorithm uses colored BFS-explorations with threshold, but with an original *global* approach that enables to overcome a recent impossibility result by Fraigniaud et al. [SIROCCO'23] about using colored BFS-exploration with *local* threshold for detecting cycles.

We also show how to quantize our algorithm for achieving a round-complexity $\tilde{O}(n^{1/2-1/2k})$ in the quantum setting for deciding C_{2k} freeness. Furthermore, this allows us to improve the known quantum complexities of the simpler problem of detecting cycles of length *at most* $2k$ by van Apeldoorn and de Vos [PODC'22]. Our quantization is in two steps. First, the congestion of our randomized algorithm is reduced, to the cost of reducing its success probability too. Second, the success probability is boosted using a new quantum framework derived from sequential algorithms, namely Monte-Carlo quantum amplification.

1 Introduction

For every fixed graph H , H -freeness is the problem consisting in deciding whether any given graph G contains H as a subgraph. In the distributed setting, in which each vertex of G is a computing element, the decision rule is specified as: G is H -free if and only if all nodes of G accept. That is, if H is a subgraph of G , then at least one node of G must reject, otherwise all nodes must accept.

H -freeness has been extensively studied in various distributed models, including the standard CONGEST model (see the recent survey [5, 6]). Recall that, in this latter model, a network is modeled as a simple connected n -vertex graph G , where the computing nodes occupy the vertices of G , and they exchange messages along the edges of G . All nodes start computing at the same time, and they perform in lockstep, as a sequence of synchronous rounds. At each round, every

*Research supported in part by the European QuantERA project QOPT (ERA-NET Cofund 2022-25), the French ANR projects DUCAT (ANR-20-CE48-0006) and QUDATA (ANR-18-CE47-0010), the French PEPR integrated project EPiQ (ANR-22-PETQ-0007), and the QuanTEdu-France project (22-CMAS-0001).

node sends a (possibly different) message to each of its neighbors in G , receives the message sent by its neighbors, and performs some local computation. The main constraint imposed on computation by the CONGEST model is on the size of the messages exchanged during each round, which must not exceed $O(\log n)$ bits in n -node networks [32].

The specific case of *cycle-detection*, i.e., deciding C_k -freeness for a fixed $k \geq 3$, has been the focus of several contributions to the CONGEST model (cf. Table 1) — unless specified otherwise, all the algorithms mentioned there are in the randomized setting, i.e., they are Monte-Carlo algorithms with 1-sided error probability $1/3$. More generally, for $\varepsilon \in (0, 1)$, a randomized distributed algorithm \mathcal{A} solves C_k -freeness with one-sided error probability ε if, for every input graph G ,

- If G contains a cycle C_k as a subgraph, then, with probability at least $1 - \varepsilon$, at least one node of G outputs reject;
- Otherwise, all nodes of G output accept with probability 1.

For odd cycles of length at least 5, the problem is essentially solved. Indeed, for every $k \geq 3$, C_k -freeness can be decided in $O(n)$ rounds by a *deterministic* algorithm [30], and it was shown that, for every odd $k \geq 5$, deciding C_k -freeness requires $\tilde{\Omega}(n)$ rounds, even for randomized algorithms [15]. The case of triangles, i.e., deciding C_3 -freeness, is however still open. The best known upper bound on the round-complexity of deciding C_3 -freeness is $\tilde{O}(n^{1/3})$ [11]. However, it was shown that any polynomial lower bound for this problem, i.e., any bound of the form $\Omega(n^\alpha)$ with $\alpha > 0$ constant, would imply major breakthroughs in circuit complexity [16].

Reference	Cycle	Complexity	Framework
[11]	C_3	$\tilde{O}(n^{1/3})$	rand.
[15, 30]	$C_{2k+1}, k \geq 2$	$\tilde{\Theta}(n)$	O det. / $\tilde{\Omega}$ rand.
[15]	C_4	$\tilde{\Theta}(\sqrt{n})$	O det. / $\tilde{\Omega}$ rand.
[30]	$C_{2k}, k \geq 2$	$\tilde{\Omega}(\sqrt{n})$	rand.
[10]	$C_{2k}, k \in \{2, 3, 4, 5\}$	$O(n^{1-1/k})$	rand.
[16]	$C_{2k}, k \geq 6$ even	$\tilde{O}(n^{1-2/(k^2-2k+4)})$	rand.
[16]	$C_{2k}, k \geq 7$ odd	$\tilde{O}(n^{1-2/(k^2-k+2)})$	rand.
[10]	$\{C_\ell \mid 3 \leq \ell \leq 2k\}, k \geq 2$	$\tilde{O}(n^{1-1/k})$	rand.
this paper	$C_{2k}, k \geq 2$	$O(n^{1-1/k})$	rand.
[8]	C_3	$\tilde{O}(n^{1/5})$	quant.
[9]	C_4	$\tilde{O}(n^{1/4})$	quant.
[33]	$\{C_\ell \mid 3 \leq \ell \leq 2k\}, k \geq 2$	$\tilde{O}(n^{1/2-1/4k+2})$	quant.
this paper	$C_{2k}, k \geq 2$	$\tilde{O}(n^{1/2-1/2k})$	quant.
this paper	$C_{2k}, k \geq 2$	$\tilde{\Omega}(n^{1/4})$	quant.
this paper	$C_{2k+1}, k \geq 2$	$\tilde{\Theta}(\sqrt{n})$	quant.
this paper	$\{C_\ell \mid 3 \leq \ell \leq 2k\}, k \geq 2$	$\tilde{O}(n^{1/2-1/2k})$	quant.

Table 1: Summary of the results about deciding C_k -freeness in CONGEST.

The landscape of deciding the presence of an even-size cycle is more contrasted. It was first shown that C_4 -freeness can be decided in $O(\sqrt{n})$ rounds, and that this complexity is essentially optimal thanks to a lower bound of $\tilde{\Omega}(\sqrt{n})$ rounds [15]. It was then shown that the lower bound of $\tilde{\Omega}(\sqrt{n})$ rounds applies to deciding C_{2k} -freeness too, for every $k \geq 2$ [30]. However, any lower bound of the form $\Omega(n^{1/2+\alpha})$ with $\alpha > 0$ for deciding C_{2k} -freeness for some $k \geq 3$ would imply

new lower bounds in circuit complexity, which are considered hard to obtain (see [10]). On the positive side, it was first proved that, for every $k \geq 2$, C_{2k} -freeness can be decided in $O(n^{1-1/k(k-1)})$ rounds [22]. This was later improved in [16], where it is proved that C_{2k} -freeness can be decided in $O(n^{1-2/(k^2-2k+4)})$ for even $k \geq 4$, and in $O(n^{1-2/(k^2-k+2)})$ for odd $k \geq 3$. However, better results were obtained for small cycles, with algorithms for deciding C_{2k} -freeness performing in $O(n^{1-1/k})$ rounds, for $k \in \{3, 4, 5\}$ [10]. These algorithms were obtained by analyzing the congestion caused by colored BFS-explorations, and by showing that if the congestion exceeds a certain threshold then there must exist a $2k$ -cycle in the network. Unfortunately, it was later shown (see [23]) that this technique does not extend to larger cycles, of length $2k$ for $k \geq 6$.

A problem related to deciding C_k -freeness is deciding whether there is a cycle of length *at most* $2k$. Surprisingly, this problem is actually simpler, because deciding C_k -freeness can take benefit from the fact that there are no cycles of length less than k . This fact was exploited in [10] to design an algorithm deciding $\{C_\ell \mid 3 \leq \ell \leq 2k\}$ -freeness in $O(n^{1-1/k})$ rounds, for any $k \geq 2$.

Moreover, the CONGEST model can take benefits of quantum effects, as far as detecting subgraphs is concerned. Indeed, it was proved that deciding C_3 -freeness (respectively, C_4 -freeness) can be decided in $\tilde{O}(n^{1/5})$ rounds [8] (resp., $\tilde{O}(n^{1/4})$ rounds [9]) in the quantum CONGEST model. Similarly, $\{C_\ell \mid 3 \leq \ell \leq 2k\}$ -freeness can be decided in $\tilde{O}(n^{1/2-1/4k+2})$ rounds [33] in the quantum CONGEST model. However, since any lower bound on the round-complexity of a problem in quantum CONGEST is also a lower bound on the round-complexity of the same problem in the classical version of CONGEST, the hardness of designing lower bounds for triangle-freeness and for C_{2k} -freeness for $k \geq 3$ in the classical setting also holds in the quantum setting.

1.1 Our Results

Our results are summarized in Table 1. We considerably simplify the landscape of results about deciding C_k -freeness, and we extend it to the quantum CONGEST framework.

Our first main contribution shows that, for every $k \geq 2$, the complexity of deciding C_{2k} -freeness is $O(n^{1-1/k})$, hence extending the results of Drucker et al. [15] and Censor-Hillel et al. [10] to all $k \geq 6$, and improving the complexity of the algorithms by Eden et al. [16]. Specifically, we show the following.

Theorem 1. *For every integer $k \geq 2$, and every real $\varepsilon > 0$, there is an algorithm that solves C_{2k} -freeness with one-sided error probability ε in $O(\log^2(1/\varepsilon) \cdot 2^{3k} k^{2k+3} \cdot n^{1-1/k})$ rounds in the CONGEST model.*

Our second main contribution is related to analysing the speedup that can be obtained by allowing nodes to handle entangled quantum resources, yet exchanging at most $O(\log n)$ qubits at each round. We show that quantum resources enable to obtain a quadratic speedup for deciding cycle-freeness. Specifically, we show the following.

Theorem 2. *In the quantum CONGEST model, the following holds:*

- *For every integer $k \geq 2$, there is a quantum algorithm deciding C_{2k} -freeness with one-sided error probability $1/\text{poly}(n)$ in $k^{O(k)} \cdot \text{polylog}(n) \cdot n^{1/2-1/2k}$ rounds, and any algorithm deciding C_{2k} -freeness with one-sided error probability at most $1/3$ performs in at least $\tilde{\Omega}(n^{1/4})$ rounds.*
- *For every integer $k \geq 1$, there is a quantum algorithm deciding C_{2k+1} -freeness with one-sided error probability $1/\text{poly}(n)$ in $\tilde{O}(\sqrt{n})$ rounds, and, for every integer $k \geq 2$, any quantum*

algorithm deciding C_{2k+1} -freeness with one-sided error probability at most $1/3$ performs in at least $\tilde{\Omega}(\sqrt{n})$ rounds.

Therefore, our quantum algorithm for deciding C_4 -freeness performing in $\tilde{O}(n^{1/4})$ rounds is, up to logarithmic multiplicative factors, optimal. For $k \geq 3$, we face the same difficulty as in the classical framework for designing lower bounds of deciding C_{2k} -freeness. For odd cycles, we show that quantum algorithms enable quadratic speedup too, and this is essentially optimal, up to logarithmic factors. That is, the complexity of deciding C_{2k+1} -freeness in the quantum CONGEST model is $\tilde{\Theta}(\sqrt{n})$ rounds for all $k \geq 2$. As for the (classical) CONGEST model, the case of C_3 -freeness remains open in absence of any non trivial lower bound. However, we point out that our new quantum framework enables to improve the quantum complexity of detecting cycles of length at most $2k$ from [33]. For every $k \geq 2$, our algorithm for $\{C_\ell \mid 3 \leq \ell \leq 2k\}$ -freeness performs in $\tilde{O}(n^{1/2-1/2k})$ instead of $\tilde{O}(n^{1/2-1/4k+2})$.

1.1.1 Our Technique for CONGEST

As in previous papers about deciding C_{2k} -freeness [10, 17], we separate the case of detecting *light* cycles from the case of detecting *heavy* cycles, where the former are $2k$ -cycles containing solely nodes the degrees of which do not exceed a specific bound d_{max} , and the latter are $2k$ -cycles that are not light. We use the same technique as [10] for detecting light cycles, with the same bound $d_{max} = n^{1/k}$ on the degrees. Our main contribution is a new technique for detecting heavy cycles, i.e., cycles containing at least one node of degree larger than $n^{1/k}$.

A central technique for detecting $2k$ -cycles is *color coding* [1], which is implemented in the distributed setting as the so-called *colored BFS-exploration* [21, 24]. Roughly, it consists in each node v picking a color $c(v) \in \{0, \dots, 2k - 1\}$ uniformly at random (u.a.r.), and the goal is to check whether there is a cycle $C = (u_0, \dots, u_{2k-1})$ with $c(u_i) = i$ for every $i \in \{0, \dots, 2k - 1\}$. For this purpose, some nodes colored 0 launch a BFS-exploration, sending their identifiers to their neighbors colored 1 and $2k - 1$. Then, at every round $i \in \{1, \dots, k - 1\}$, nodes colored i forward all the identifiers received from neighbors colored $i - 1$ to their neighbors colored $i + 1$, while nodes colored $2k - i$ forward all the identifiers received from neighbors colored $2k - (i - 1)$ to their neighbors colored $2k - (i + 1)$. If a node colored k receives the same identifier from a neighbor colored $k - 1$ and $k + 1$, a $2k$ -cycle has been detected. By repeating about $(2k)^{2k}$ times the random choice of the colors, the probability that an existing $2k$ -cycle is well colored is at least $2/3$, which guarantees the desired success probability. Using this approach, the issue is to control congestion, that is, to control the number of distinct identifiers that a same node has to forward during the BFS-exploration.

For controlling the congestion in colored BFS-exploration, an elegant approach has been presented in [10], that we refer to as *local threshold*. In essence, it consists in selecting a single source node s u.a.r., which triggers all its neighbors colored 0, asking them to launch a colored BFS-exploration. A key point is that, as shown in [10], for $k \in \{2, 3, 4, 5\}$, there exists a (constant) threshold $\tau_k \geq 1$ such that, a constant fraction of sources s are either in a $2k$ -cycle or will not cause any node to receive more than τ_k identifiers. Therefore, each attempt to find a $2k$ -cycle triggered by a selected source s takes a constant number of rounds (namely, at most $k \cdot \tau_k$ rounds). Now, the probability that the selected source s has a neighbor in a cycle is at least $1/n^{1/k}$, and thus it is sufficient to repeat $O(n^{1-1/k})$ times the random choice of s for finding a $2k$ -cycle with constant probability (if it exists). Unfortunately, the local threshold technique was proved to suffer from

some limitation [23], namely it is not extendable to the detection of larger $2k$ -cycles, i.e., to $k > 5$.

To overcome the inherent limitation of the local threshold technique, we adopt a *global threshold* approach. Specifically, instead of repeating $O(n^{1-1/k})$ times the choice of a single source s that triggers the colored BFS-exploration, we directly select a set S of $O(n^{1-1/k})$ random sources, and we only need to repeat $O(1)$ times this choice. Now, we show that there exists a (global) threshold $\tau_k(n) = O(n^{1-1/k})$ such that, for each choice of S , if a node has to forward more than $\tau_k(n)$ identifiers at some round, then there must exist a $2k$ -cycle. Establishing this result is the key point in our algorithm. It requires to analyze in detail the successive rounds of the BFS-exploration which may eventually cause a node v to receive a set I_v of more than $\tau_k(n)$ identifiers, and to use this analysis for constructing a $2k$ -cycle. This cycle is itself the union of a path P alternating between nodes in S and nodes whose identifiers are in I_v , and two vertex-disjoint paths connecting the two end points of P to v .

1.1.2 Our Technique for Quantum CONGEST

One of the main tools to speed-up distributed algorithm in the quantum framework is the *distributed Grover search* [26]. Indeed, (sequential) Grover search enables solving problems such as minimizing a function, or searching for a preimage, in a time-complexity that is often quadratically faster than the classical (i.e., non quantum) time-complexity. Transferred to the quantum CONGEST model, Grover search was first used to design sublinear algorithms for computing the diameter of a graph [26]. A nested version of Grover search was also used recently for detecting cliques [8] and a parallel one for cycles [33].

We take one step further in the use of Grover search, by defining an encapsulated quantum framework for distributed computing that we call *distributed quantum Monte-Carlo amplification* (Theorem 3). We show that, given any distributed (quantum, or randomized) Monte-Carlo algorithm with small one-sided *success* probability ε and round-complexity R , there exists a quantum algorithm with constant (e.g., $2/3$) one-sided *success* probability, whose round-complexity is roughly $\sqrt{1/\varepsilon} \cdot R$. Observe that boosting the success probability to a constant would have required $1/\varepsilon$ iterations in a non-quantum setting. Our boosting technique is quadratically faster. There is however a cost to pay. Indeed, this is ignoring an additional term $D/\sqrt{\varepsilon}$, where D is the diameter of the graph, that appears in the overall quantum complexity. This may not be an issue for problems whose round-complexity inherently depends on the diameter (e.g., computing an MST), but this is an issue for “local problems” such as H -freeness. Nevertheless, we can eliminate the diameter dependence by employing the standard *diameter reduction* technique [17].

In order to apply our amplification technique, we first *decrease* the success probability of our classical algorithm. This allows us to decrease its congestion too. Specifically, our classical algorithm has constant success probability, and congestion $O(n^{1-1/k})$. Indeed, in colored BFS-exploration performed by the classical algorithm, every node participating to the exploration forwards information that comes from at most $\tau_k(n)$ nodes, where $\tau_k(n) = O(n^{1-1/k})$ is the global threshold. This yields a congestion at most $\tau_k(n)$, and thus a round complexity at most $\tau_k(n)$. By not activating systematically the nodes of the tested set in the colored BFS-exploration, but by activating each of them independently with probability $\varepsilon = O(1/\tau_k(n))$, we show that the congestion decreases to $O(\varepsilon \tau_k(n)) = O(1)$, and the success probability drops down to ε . At this point, we can apply our Monte-Carlo amplification technique to get one sided-error probability $1/\text{poly}(n)$ with round-complexity $\tilde{O}(D \cdot \sqrt{\tau_k(n)})$ rounds, where D is the diameter of the graph. Finally, the diameter dependence is eventually removed, up to polylog factors, by using the technique in [17].

1.2 Related Work

There is a vast literature on distributed algorithms for deciding whether the input graph G includes a fixed given graph H as a subgraph, or as an induced subgraph. We refer to the recent survey [5, 6] for a detailed description of the recent progress in the matter, including the typical techniques. We just clarify here a few points that sometimes cause confusion. First, deciding H -freeness is also referred to as *detecting H* . Second, subgraph detection has a more demanding variant, called subgraph *listing*. In the latter, each occurrence of H must be reported by at least one node. Subgraph detection and subgraph listing each have their *local* variants. Local detection imposes that each node outputs accept or reject based on whether it is a part of a copy of H or not, and local listing imposes that each node outputs the list of occurrences of the subgraph H it belongs to. In this paper, we focused on deciding C_k -freeness, i.e., cycle detection. Besides cycles, two families of graphs have been extensively studied in the framework of distributed H -freeness, trees [30] and cliques [7, 13, 16]. Subgraph detection as well as subgraph listing have been studied in other distributed computing models, such as CONGESTED CLIQUE (see, e.g., [14]). The *testing* variant of the problem has also been considered, in which the goal is to decide whether the input graph is H -free or contains *many* copies of H (measured, e.g., by the number of edges of G that must be deleted to obtain a graph that is H -free). In this latter framework, it is possible to design decision algorithms performing in $O(1)$ rounds in CONGEST (see, e.g., [21, 25]).

In the quantum setting, it was shown in [18] that the quantum CONGEST model is not more powerful than the classical CONGEST model for many important graph-theoretical problems. Nonetheless, it was later shown in [26] that computing the diameter of the network can be solved more efficiently in the quantum setting. Since then, other quantum speed-ups have been discovered, including subgraph detection [8, 33]. We also mention that a speed-up similar than ours has already been established for detecting C_4 , with round complexity $\tilde{O}(n^{1/4})$ in an unpublished work [9]. This latter contribution is directly using a distributed Grover search [26], and not our Monte-Carlo amplification. The implementation of Grover in [9] is thus a bit more involved. In particular, a leader node is used for deciding which nodes are activated or not, based on a token sampling in $[1, \sqrt{n}]$. While the token sampling is decentralized, the activation of the nodes for a given token is centralized before applying the distributed Grover search, leading to a multiplicative term equal to the diameter D . Finally, we note that in the CONGESTED CLIQUE model as well, quantum algorithms faster than the best known classical algorithms have been discovered, starting for the All-Pair Shortest Path problem [29], and recently for clique-detection [8]. In the LOCAL model, which is another fundamental model in distributed computing, separations between the computational powers of the classical and quantum versions of the model have been reported for some problems [27, 31], but other papers have also reported that quantum effects do not bring significant benefits for other problems (e.g., approximate graph coloring [12]).

2 Cycle Detection in the Classical CONGEST Model

This section is entirely dedicated to proving Theorem 1. For this purpose, we first describe the algorithm claimed to exist in the statement of the theorem (Algorithm 1), then we analyze it.

2.1 Algorithm description

In Algorithm 1, $k \geq 2$ and $\epsilon > 0$ are fixed parameters, and $G = (V, E)$ is the input graph. The only prior knowledge given to each node $v \in V$, apart from k , ϵ , and the identifier $\text{id}(v)$, is the size $n = |V|$ of the input graph. Algorithm 1 is using a variant of *color-BFS with threshold* [10], displayed as Procedure color-BFS in Algorithm 1, detailed next.

2.1.1 Procedure Color-BFS with Threshold

The syntax of a call to color-BFS with threshold is

$$\text{color-BFS}(k, H, c, X, \tau),$$

where k is the fixed parameter, H is a subgraph of the input graph G , $c : V(H) \rightarrow \{0, \dots, 2k - 1\}$ is a (non-necessarily proper) coloring of the vertices of H , $X \subseteq V(H)$ is a set of vertices, and $\tau \geq 0$ is an integer called *threshold*. In all calls to color-BFS, graph H will be an induced subgraph of G , and in particular every node of G will locally know whether it belongs to H or not. Similarly, every node will know whether it belongs to set X or not.

In our variant of color-BFS, only the nodes $x \in X$ initiate the search for a $2k$ -cycle, and the search is performed in H only, i.e., not necessarily in the whole graph G . Specifically, every node $x \in X$ colored 0, i.e., every node $x \in X$ for which $c(x) = 0$, launches the search by sending its identifier $\text{id}(x)$ to all its neighbors in H (cf. Instruction 15).

Then, as specified in the for-loop at Instruction 16, for every $i = 1, \dots, k - 1$, every node $v \in V(H)$ colored i having received a collection of identifiers $I_v \subseteq \{\text{id}(x) \mid x \in X\}$ from its neighbors in H colored $i - 1$ aims at forwarding I_v to all its neighbors in H colored $i + 1$. However, it does forward I_v to its neighbors only if I_v is not too large, namely only if the number $|I_v|$ of identifiers in I_v does not exceed the threshold τ . Instead, if v has collected too many identifiers, i.e., if $|I_v| > \tau$, then v simply discards I_v , and forwards nothing.

Similarly, for every $i = 2k - 1, \dots, k + 1$, every node $v \in V(H)$ colored i having received a collection I_v of identifiers from its neighbors in H colored $i + 1 \bmod 2k$ forwards I_v to all its neighbors in H colored $i - 1$ whenever $|I_v| \leq \tau$, and discards I_v otherwise.

Eventually, at Instruction 25, if a node $v \in V(H)$ colored k receives a same identifier $\text{id}(x)$ of some node $x \in X$ from a neighbor in H colored $k - 1$, and from a neighbor in H colored $k + 1$, then v rejects. Observe that when v rejects, then by construction there is a $2k$ -cycle containing v and x . Indeed $\text{id}(x)$ traveled to v along two paths $(x = u_0, u_1, \dots, u_k = v)$ and $(x = u_0, u_{2k-1}, \dots, u_k = v)$ of length k , with different internal vertices.

2.1.2 The Cycle Detection Algorithm

Our cycle detection algorithm is detailed in Algorithm 1. It essentially consists of three calls to $\text{color-BFS}(k, H, c, X, \tau)$, for three different graphs H , three different sets X , and one threshold τ . The first set is the set U of *light* nodes in G (see Instruction 1), i.e., the set of nodes

$$U = \{u \in V \mid \deg(u) \leq n^{1/k}\}.$$

The second set is denoted by S . It is constructed randomly at Instructions 3-4, by having each node independently deciding whether to enter set S with probability $p = \Theta(1/n^{1/k})$. In other words, if

Algorithm 1 Deciding C_{2k} -freeness in $G = (V, E)$ with one sided-error ε

```

1:  $U \leftarrow \{u \in V \mid \deg(u) \leq n^{1/k}\};$  ▷  $U$  is the set of “light” nodes
2:  $\hat{\varepsilon} \leftarrow \ln(3/\varepsilon); \quad p \leftarrow \hat{\varepsilon} \cdot 2k^2/n^{1/k};$  ▷ Setting of the selection probability  $p$ 
3: every node  $u \in V$  selects itself with probability  $p$ ;
4:  $S \leftarrow \{\text{selected nodes}\};$  ▷ In expectation,  $|S| = np$ 
5:  $W \leftarrow \{u \in V \setminus S \mid |N_G(u) \cap S| \geq k^2\};$  ▷  $u \in W$  iff  $u$  has at least  $k^2$  selected neighbors
6:  $K \leftarrow \hat{\varepsilon} \cdot (2k)^{2k}; \quad \tau \leftarrow k2^k \cdot np;$  ▷ Setting of #repetitions  $K$ , and threshold  $\tau$ 
7: for  $r = 1$  to  $K$  do ▷ A sequence of  $K$  search phases
8:   every node  $u \in V$  picks a color  $c(u) \in \{0, \dots, 2k - 1\}$  u.a.r.;
9:   color-BFS( $k, G[U], c, U, \tau$ ); ▷ Search for  $2k$ -cycles with light nodes only
10:  color-BFS( $k, G, c, S, \tau$ ); ▷ Search for  $2k$ -cycles with at least one selected node
11:  color-BFS( $k, G[V \setminus S], c, W, \tau$ ); ▷ Search for other  $2k$ -cycles
12: end for
13: every node that has not output reject at a previous round outputs accept.

14: procedure color-BFS( $k, H, c, X, \tau$ )
15:   every node  $x \in X$  with  $c(x) = 0$  sends its ID to its neighbors in  $H$ 
16:   for  $i = 1$  to  $k - 1$  do
17:     for every node  $v \in V(H)$  with  $c(v) = i$  (resp.,  $c(v) = 2k - i$ ) do
18:        $I_v \leftarrow \{\text{IDs received from neighbors in } H \text{ colored } i - 1 \text{ (resp. } 2k - i - 1)\}$ 
19:       if  $|I_v| \leq \tau$  then
20:          $v$  forwards  $I_v$  to its neighbors in  $H$  colored  $i + 1$  (resp.  $2k - i - 1$ )
21:       end if
22:     end for
23:   end for
24:   for every node  $v \in V(H)$  colored  $k$  do
25:     if  $v$  receives a same ID from two neighbors respectively colored  $k - 1$  and  $k + 1$  then
26:        $v$  outputs reject
27:     end if
28:   end for
29: end procedure

```

B_p denotes the random variable with Bernoulli distribution of parameter p ,

$$S = \{u \in V \mid B_p(u) = 1\}.$$

Note that the expected size of S is $\Theta(n^{1-1/k})$. As we shall show later, the probability $p = \Theta(1/n^{1/k})$ can be set so that, for every node u of degree at least $n^{1/k}$, the probability that u has at least a constant number $|N_G(u) \cap S|$ of neighbors in S is constant. This leads us to setting our third set, W , as the set of nodes that have at least k^2 selected neighbors (see Instruction 5), that is,

$$W = \{u \in V \setminus S \mid |N_G(u) \cap S| \geq k^2\}.$$

Once these three sets U , S , and W have been constructed, they remain fixed for the rest of the algorithm. Finally, a threshold

$$\tau = \Theta(n^{1-1/k})$$

is set at Instruction 6. This threshold will be used in all color-BFS calls in the for-loop at Instruction 7.

The for-loop at Instruction 7 is performed a constant number K of times. It will be shown later that choosing K as a sufficiently large constant will be sufficient to guarantee that, if there is a $2k$ -cycle C in G , then, with constant probability, its nodes will be colored consecutively from 0 to $2k - 1$ in a run of the loop. Indeed, at each iteration of the for-loop, every node $u \in V$ picks a color $c(u) \in \{0, \dots, 2k - 1\}$ uniformly at random, at Instruction 8. Provided with this coloring c , the algorithm proceeds as a sequence of three different color-BFS calls.

The first color-BFS (see Instruction 9) aims at detecting the presence of a $2k$ -cycle containing light nodes only. Therefore, X is set to U , the graph H is merely the subgraph $G[U]$ of G induced by U , and the threshold is τ . Then, since $G[U]$ has maximum degree $n^{1/k}$, as we shall see in Lemma 1, the threshold cannot be exceeded. As a consequence, thanks to Fact 4, if the coloring c does color the nodes of a $2k$ -cycle consecutively from 0 to $2k - 1$, then this cycle will be detected.

The second color-BFS (see Instruction 10) aims at detecting the presence of a $2k$ -cycle containing at least one selected node, i.e., at least one node in S . Therefore, X is set to S , and the graph H is the whole graph G . The threshold is again set to τ , which exceeds the expected size of S . Therefore, thanks to Fact 4, if there is a $2k$ -cycle including a node in S , this cycle will be detected.

The third color-BFS (see Instruction 11) is addressing the “general case”, that is the detection of a cycle containing at least one heavy node (i.e., with at least one node not in U), and not containing any selected nodes (i.e., with no nodes in S). The search is therefore performed in the subgraph $G[V \setminus S]$ of G . The crucial point in the third color-BFS is the choice of the set X initializing the search. We set it to the aforementioned set $W = \{u \in V \setminus S \mid |N_G(u) \cap S| \geq k^2\}$. That is, the search is activated only by nodes neighboring sufficiently many, i.e., at least k^2 , selected nodes. As for the first two color-BFS calls, the threshold is set to τ . This may appear counter intuitive as W may be larger than $O(n^{1-1/k})$. Nevertheless, we shall show that if a node $v \in G[V \setminus S]$ has to forward more than $O(n^{1-1/k})$ identifiers of nodes in W , then there must exist a $2k$ -cycle in G passing through S , which would have been detected at the previous step.

Finally, each node having performed K iterations of the for-loop, without rejecting during any execution of the $3K$ color-BFS procedures, accepts. This completes the description of the algorithm.

2.2 Analysis of Algorithm 1

Let us start the analysis of Algorithm 1 by a collection of basic technical facts.

2.2.1 Preliminary Results

We first show that when a $2k$ -cycle exists, by setting the constant K large enough, there will be a run of the loop at Instruction 7 of Algorithm 1 for which the nodes of the cycle are colored consecutively from 0 to $2k - 1$ by the coloring c .

Fact 1. *Let $\alpha > 0$, and $K \geq \alpha(2k)^{2k}$. Let $(u_0, u_1, \dots, u_{2k-1})$ be any sequence of $2k$ nodes of G . Then, with probability at least $1 - e^{-\alpha}$, there is at least one iteration of the loop at Instruction 7 in Algorithm 1 such that $c(u_i) = i$, for $i = 0, \dots, 2k - 1$.*

Proof. For each iteration, the probability that this event occurs is $(1/2k)^{2k}$, due to the independent choices of the colors of the nodes. The choices of coloring being also independent at each iteration

of the loop, the probability that the event never occurs for any of the K iterations is

$$\left(1 - \frac{1}{(2k)^{2k}}\right)^K \leq \exp\left(-\frac{K}{(2k)^{2k}}\right) \leq e^{-\alpha},$$

as claimed. \square

We carry on with two standard claims stating that, with high probability, the set S is not too large, and vertices with large degree (in G) have sufficiently many neighbors in S .

Fact 2. *Let $\alpha > 0$. If every node selects itself with probability $p = \alpha/n^{1/k}$ at Instruction 3 of Algorithm 1, then*

$$\Pr[|S| \leq 2\alpha n^{1-1/k}] \geq 1 - e^{-\alpha/3}.$$

Proof. The cardinality of S is a random variable that follows a binomial distribution of parameters $(n, \alpha/n^{1/k})$. We use the following Chernoff Bound: for any $\delta \geq 0$,

$$\Pr[|S| \geq (1 + \delta)\mathbb{E}[|S|]] \leq \exp\left(-\frac{\delta^2\mathbb{E}[|S|]}{2 + \delta}\right).$$

With $\delta = 1$ and $\mathbb{E}[|S|] = \alpha n^{1-1/k}$, we conclude:

$$\Pr[|S| \geq 2\alpha n^{1-1/k}] \leq \exp\left(-\frac{\alpha n^{1-1/k}}{3}\right) \leq e^{-\alpha/3},$$

as claimed. \square

Fact 3. *Let $\alpha > 0$, and let $v \in V$ such that $\deg(v) > n^{1/k}$. If every node selects itself with probability $p = \alpha/n^{1/k}$ at Instruction 3 of Algorithm 1, then*

$$\Pr[|N_G(v) \cap S| \geq \alpha/2] \geq 1 - e^{-\alpha/8}.$$

Proof. The degree $|N_G(v) \cap S|$ of v in S follows a binomial law of parameters $(\deg(v), p)$. We use the following Chernoff Bound: for any $0 < \delta < 1$,

$$\Pr[|N_G(v) \cap S| \leq (1 - \delta)\mathbb{E}[|N_G(v) \cap S|]] \leq \exp\left(-\frac{\delta^2\mathbb{E}[|N_G(v) \cap S|]}{2}\right).$$

For $\delta = 1/2$, we have $\mathbb{E}[|N_G(v) \cap S|] = p \deg(v) \geq \alpha$, from which it follows that

$$\Pr[|N_G(v) \cap S| \leq \alpha/2] \leq \Pr[|N_G(v) \cap S| \leq p \deg(v)/2] \leq \exp\left(-\frac{p \deg(v)}{8}\right) \leq e^{-\alpha/8},$$

as desired. \square

Let us now introduce some notations used for analysing Procedure `color-BFS`(k, H, c, X, τ) in Algorithm 1. Let

$$X_0 = \{x \in X \mid c(x) = 0\}$$

be the set of nodes that send their identifiers at Instruction 15 of `color-BFS`(k, H, c, X, τ). We define, for any node v in H colored i or $2k - i$ with $i \in \{1, \dots, k - 1\}$, the subset $X_0(v) \subseteq X_0$ of nodes connected to v by a “well colored” path in H , as follows.

- If $c(v) \in \{1, \dots, k-1\}$ then

$$X_0(v) = \{x \in X_0 \mid \text{there exists a path } (x, v_1, \dots, v_{c(v)-1}, v) \text{ in } H \text{ such that,} \\ \text{for every } j \in \{1, \dots, c(v)-1\}, c(v_j) = j\}. \quad (1)$$

- If $c(v) \in \{2k-1, \dots, k+1\}$ then

$$X_0(v) = \{x \in X_0 \mid \text{there exists a path } (x, v_{2k-1}, \dots, v_{c(v)+1}, v) \text{ in } H \text{ such that} \\ \text{for every } j \in \{c(v)-1, \dots, 2k-1\}, c(v_j) = j\}. \quad (2)$$

With this construction in hand, we can state a first general fact for any instance of `color-BFS`(k, H, c, X, τ) from Algorithm 1.

Fact 4. *Procedure `color-BFS`(k, H, c, X, τ) satisfies the following two properties.*

- *Let $(v_1, v_2, \dots, v_{k-1})$ be a path in H such that $c(v_i) = i$ for all $i \in \{1, \dots, k-1\}$. If $|X_0(v_{k-1})| \leq \tau$, then, for every $i \in \{1, \dots, k-1\}$, $|I_{v_i}| \leq \tau$ and $I_{v_i} = X_0(v_i)$.*
- *Let $(v_{2k-1}, v_{2k-2}, \dots, v_{k+1})$ be a path in H such that $c(v_i) = i$ for all $i \in \{k+1, \dots, 2k-1\}$. If $|X_0(v_{k+1})| \leq \tau$, then, for every $i \in \{k+1, \dots, 2k-1\}$, $|I_{v_i}| \leq \tau$ and $I_{v_i} = X_0(v_i)$.*

Proof. We prove the result for nodes with colors between 1 and $k-1$ only, as the result for nodes with colors between $2k-1$ and $k+1$ holds by the same arguments. From Instructions 15 and 18, we get that for every node v in H , every identifier in I_v is the one of some node in $X_0(v)$. All along a path (v_1, \dots, v_{k-1}) , if $|X_0(v_i)| \leq \tau$ for every $i \in \{1, \dots, k-1\}$, then the condition of Instruction 19 is always satisfied for I_{v_i} , and the identifiers of all the nodes $X_0(v_i)$ are in I_{v_i} . The result follows after noticing that $X_0(v_i) \subseteq X_0(v_{k-1})$ for every $i \in \{1, \dots, k-1\}$. \square

2.2.2 Analysis of the first two `color-BFS`

The next lemma states that, if G contains a $2k$ -cycle composed of consecutively colored nodes which are all light, that is, with small degree, then the first call of `color-BFS` leads at least one node to reject.

Lemma 1. *Suppose that G contains a $2k$ -cycle $C = (u_0, \dots, u_{2k-1})$ in $G[U]$, and that c is a coloring for which $c(u_i) = i$ for every $i \in \{0, \dots, 2k-1\}$. Then u_k rejects in `color-BFS`($k, G[U], c, U, \tau$).*

Proof. Every node in U has degree at most $n^{1/k}$. By a direct induction on $i = 1, \dots, k-1$, we get that $|U_0(u_i)| \leq n^{i/k} \leq \tau$ for every $i \in \{0, \dots, k-1\}$, and the same holds for u_{2k-i} . Fact 4 implies that $|I_{u_i}| \leq \tau$ and $|I_{u_{2k-i}}| \leq \tau$ for every $i \in \{1, \dots, k-1\}$. Therefore the identifier of u_0 is forwarded along the two paths $(u_0, u_1, \dots, u_{k-1}, u_k)$ and $(u_0, u_{2k-1}, \dots, u_{k+1}, u_k)$, and u_k rejects. \square

The next lemma establishes that if S is of size at most τ , and if G contains a well colored cycle for which the node colored 0 is in S , then the second call of `color-BFS` leads the node colored k in the cycle to reject.

Lemma 2. *Suppose that G contains a $2k$ -cycle $C = (u_0, \dots, u_{2k-1})$ with $u_0 \in S$, and let c be a coloring such that $c(u_i) = i$ for every $i \in \{0, \dots, 2k-1\}$. If $|S| \leq \tau$ then u_k rejects in `color-BFS`(k, G, c, S, τ).*

Proof. As $S_0(v) \subseteq S$ for every node $v \in G$, $|S_0(v)| \leq |S|$, and thus in particular for every node $v \in \{u_1, \dots, u_{k-1}\} \cup \{u_{2k-1}, \dots, u_{k+1}\}$. Fact 4 then implies that $|I_{u_i}| \leq \tau$ and $|I_{u_{2k-i}}| \leq \tau$ for every $i \in \{1, \dots, k-1\}$. Therefore, the identifier of node u_0 is forwarded along the two branches of the cycle, and u_k rejects. \square

2.2.3 Analysis of the third color-BFS

The purpose of this section is to prove the following lemma which, combined with Lemmas 1 and 2, ensures the correctness of our algorithm. Informally, Lemma 3 below states that if G contains no cycle passing through S , but contains a well colored cycle passing through some u_0 having sufficiently many neighbours in S , then the third call to color-BFS rejects. This lemma is crucial as it does not only provide the key of the proof of Theorem 1, but it is also central in the design of our algorithm for quantum CONGEST.

Lemma 3. *Suppose that G contains a $2k$ -cycle $C = (u_0, \dots, u_{2k-1})$ with $c(u_i) = i$ for every $i \in \{0, \dots, 2k-1\}$, which is not intersecting S , but such that u_0 has at least k^2 neighbors in S . If $|S| \leq \tau/(k2^{k-1})$ then u_k rejects in color-BFS($k, G[V \setminus S], c, W, \tau$).*

The proof of Lemma 3 is based on the following combinatorial property. If $C \cap S = \emptyset$ for every $2k$ -cycle C in G , then, for every $v \in V \setminus S$, the set $W_0(v)$ defined below is small, which implies that I_v is also small in color-BFS($k, G[V \setminus S], c, W, \tau$). More specifically, the following holds.

Lemma 4 (Density lemma). *Let $S, W_0, V_1, \dots, V_{k-1}$ be non-empty disjoint subsets of vertices of a graph G , and let us assume that every vertex $w_0 \in W_0$ has at least k^2 neighbors in S . For every $i \in \{1, \dots, k-1\}$, and every $v \in V_i$, let us define*

$$W_0(v) = \{w \in W_0 \mid G \text{ contains a path } (w, v_1, \dots, v_i = v) \text{ s.t., for every } j \in \{1, \dots, i-1\}, v_j \in V_j\}.$$

If $|W_0(v)| > 2^{i-1}(k-1)|S|$ for some $i \in \{1, \dots, k-1\}$ and some $v \in V_i$, then G contains a $2k$ -cycle intersecting S .

Before proving Lemma 4, let us observe that it directly implies Lemma 3.

Proof of Lemma 3. Let us apply Lemma 4 twice, once with

$$V_i = \{v \in V \setminus S \mid c(v) = i\}$$

for every $i \in \{1, \dots, k-1\}$, and once with

$$V_i = \{v \in V \setminus S \mid c(v) = 2k - i\}.$$

We get that, for color-BFS($k, G[V \setminus S], c, W, \tau$), if $C \cap S = \emptyset$ for every $2k$ -cycle C in G then, for every node $v \in V \setminus S$ colored i or $2k - i$ with $i \in \{1, \dots, k-1\}$, we have $|W_0(v)| \leq 2^{i-1}(k-1)|S|$. Using the fact that $|S| \leq \tau/(k2^{k-1})$, we derive that $|W_0(v)| \leq \tau$. It infers by Fact 4 that $I_{u_i} = W_0(u_i)$ for all $i \in \{1, \dots, k-1\} \cup \{2k-1, \dots, k+1\}$. In particular the identifier of u_0 is forwarded along the two branches of the cycle, and u_k rejects. \square

The rest of this subsection is devoted to the proof of the Density Lemma (Lemma 4).

As a warm-up, let us prove the Density Lemma for $i = 1$, in order to provide the reader with an intuition of the proof before getting into the (considerably more complicated) general case. The case $i = 1$ will also illustrate why the lemma is called ‘‘Density lemma’’.

Warm Up: The Case $i = 1$. Let $v \in V_1$, and let us consider the bipartite graph $H(v)$ with vertex partition S and $W_0(v)$. (Recall that $W_0(v) \subseteq W_0$.) Suppose that $|W_0(v)| > (k-1)|S|$. Let us show that one can construct a path P , of length $2(k-1)$, with both end points in $W_0(v)$. This path together with v forms a $2k$ -cycle. Path P exists merely because the graph $H(v)$ is dense. More precisely, $H(v)$ has degeneracy¹ at least k . Indeed, since every $w \in W_0$ has at least k^2 neighbours in S , we get that

$$|E(H(v))| \geq k^2|W_0(v)| \geq 2k|W_0(v)| \geq k(|W_0(v)| + |S|) = k|V(H(v))|.$$

Consequently, $H(v)$ contains a non-empty subgraph of minimum degree at least k , obtained by repeatedly removing all vertices of degree less than k . In a bipartite graph of minimum degree k , one can greedily construct a path of $2k$ vertices, starting from any vertex, and extending it $2k-1$ times with a new vertex that has not been used before. The extension is possible as long as the path has no more than $k-1$ vertices in each of the partitions. In particular, one can construct a path with $2k-1$ edges with one extremity in S and one in $W_0(v)$. Therefore there exists a path P with $2(k-1)$ edges, starting and ending on vertices of $W_0(v)$, thus achieving the proof for $i = 1$. \square

We aim at extending the proof above to arbitrary values of i . Let i be the smallest index for which the condition of the lemma, i.e., $|W_0(v)| > 2^{i-1}(k-1)|S|$ for some $v \in V_i$, holds. A naive approach consists in considering again the bipartite graph $H(v)$ with edge set $E(S, W_0(v))$. As in the case $i = 1$, one may argue that $H(v)$ is dense enough to contain a path P of length $2(k-i)$, starting and ending on some vertices of $W_0(v)$. Both endpoints of path P have some path to v , say P' and P'' respectively, of length i . If paths P' and P'' were disjoint (except for v), then they would form, together with path P , a $2k$ -cycle. Unfortunately, with a greedy construction of path P as in the case $i = 1$, there are no reasons why paths P' and P'' should be disjoint, hence the need to refine the analysis.

Intuition of the proof. (Also see Fig. 1.) Let us construct a sequence of nested subgraphs of $H(v)$, denoted by

$$\text{IN}(v, 0), \text{IN}(v, 1), \dots, \text{IN}(v, 2q),$$

where $q = \lfloor \frac{k-1}{2} \rfloor$. The construction starts from $\text{IN}(v, 2q)$, down to $\text{IN}(v, 0)$, with $\text{IN}(v, \gamma-1) \subseteq \text{IN}(v, \gamma)$ for every $\gamma \in \{1, \dots, 2q\}$. We prove (cf. Lemma 6) that, if $\text{IN}(v, 0)$ is not empty, then G contains a $2k$ -cycle passing through S . Then we conclude by showing (cf. Lemma 7) that if the graph $\text{IN}(v_j, 0)$ is empty for every $j \in \{1, \dots, i\}$ and every $v_j \in V_j$, then $|W_0(v)| \leq 2^{i-1}(k-1)|S|$.

In order to construct the decreasing sequence of graphs $\text{IN}(v, \gamma)$, $\gamma = 2q, \dots, 0$, each vertex v “selects” a subset of edges of $H(v)$, denoted by $\text{OUT}(v)$. One can think of these edge-sets as edges selected by vertex v for its neighbors of color $c(v) + 1$, and vertices of color $c(v) + 1$ will only be allowed to choose among edges selected for them by their neighbors of color $c(v)$. More specifically, if v is colored 0, then $\text{OUT}(v)$ is simply the set of edges incident to v connecting v to a vertex of S . Then, for any vertex v of color $i \geq 1$, v constructs the sequence $\text{IN}(v, \gamma)$ starting from the union of $\text{OUT}(v')$ for all neighbors v' of v with color $i-1$. The sequence $\text{IN}(v, \gamma)$ is constructed by repeatedly removing edges incident to vertices of “small” degree, and $\text{OUT}(v)$ corresponds to some of the edges removed during the process. This construction will ensure that if, for some vertex v of color i , $\text{IN}(v, 0) \neq \emptyset$, then G contains a $2k$ -cycle (Lemma 6), and the reason for this is that one can construct the three following paths (see Fig. 1):

¹A graph is d -degenerate if each of its subgraphs has a vertex of degree at most d ; The degeneracy of a graph is the smallest value d for which it is d -degenerate.

1. a path P of length $2(k - i) - 1$ in $\text{IN}(v, 2q)$, starting from some $w \in W_0(v)$, and ending on some $s \in S$ of high degree;
2. a path $P' = (v'_0, v'_1, \dots, v'_i - 1, v'_i)$ of length i from the endpoint $w = v'_0$ of P to $v = v'_i$, such that, for every $j < i$, v'_j is colored j , and the edge incident to w in P is contained in the graph $\text{OUT}(v'_j)$;
3. a path $P'' = (s, v''_0, \dots, v''_i)$ of length $i + 1$ from the other endpoint s of P to $v = v''_i$, where each v''_j is colored j , such that P'' intersects P only in s , and intersects P' only in v . For constructing this third path P'' , we heavily rely on the maximum degrees of the graphs $\text{OUT}(v_j)$ (which are upper bounded by a function of j), and on the fact that the edge of P'' incident to s can be chosen to avoid all sets $\text{OUT}(v_j)$, for all $j < i$. This will ensure that P'' and P' are disjoint.

The paths P , P' , and P'' together form the cycle of length $2k$. Provided with this rough intuition, let us proceed with the formal construction of the graphs IN and OUT .

Sparsification. Let $E(S, W_0)$ denote the set of edges of G having one endpoint in S , and the other in W_0 . Let $\text{OUT}(v), \text{IN}(v) \subseteq E(S, W_0)$ be defined inductively on $i = 1, \dots, k - 1$ for any node $v \in V_i$. $\text{OUT}(v)$ and $\text{IN}(v)$ are constructed as sets of edges, and we slightly abuse notation by also denoting $\text{OUT}(v)$ and $\text{IN}(v)$ as the *graphs* induced by these edge sets.

For every $w \in W_0$, $\text{OUT}(w)$ is defined as the set of edges between w and S in $G = (V, E)$, that is, for every $w \in W_0$,

$$\text{OUT}(w) = E(\{w\}, S) = \{\{w, s\} \in E \mid s \in S\}. \quad (3)$$

Let us set $V_0 = W_0$. Then, for every $i \in \{1, \dots, k - 1\}$, and every $v \in V_i$, we inductively define

$$\text{IN}(v) = \bigcup_{\{v, v'\} \in E : v' \in V_{i-1}} \text{OUT}(v'). \quad (4)$$

Let $q = \lfloor \frac{k-i}{2} \rfloor$. We now inductively define an intermediate increasing sequence of subsets

$$\text{IN}(v, 0) \subseteq \text{IN}(v, 1) \subseteq \dots \subseteq \text{IN}(v, 2q - 1) \subseteq \text{IN}(v, 2q) \subseteq \text{IN}(v)$$

as follows:

1. Initialization:

$$\text{IN}(v, 2q) = \{\{s, v\} \in \text{IN}(v) \mid (s \in S) \wedge (\text{deg}_{\text{IN}(v)}(s) > 2^{i-1}(k - 1))\}. \quad (5)$$

2. Induction from 2γ to $2\gamma - 1$:

$$\text{IN}(v, 2\gamma - 1) = \{\{w, s\} \in \text{IN}(v, 2\gamma) \mid (w \in W_0) \wedge (\text{deg}_{\text{IN}(v, 2\gamma)}(w) > 2\gamma)\}. \quad (6)$$

3. Induction from $2\gamma - 1$ to $2\gamma - 2$:

$$\text{IN}(v, 2\gamma - 2) = \{\{s, w\} \in \text{IN}(v, 2\gamma - 1) \mid (s \in S) \wedge (\text{deg}_{\text{IN}(v, 2\gamma - 1)}(s) > 2\gamma - 1)\}. \quad (7)$$

Finally, we define $\text{OUT}(v)$ as:

$$\begin{aligned} \text{OUT}(v) = & \left\{ \{s, v\} \in \text{IN}(v) \mid (s \in S) \wedge (\text{deg}_{\text{IN}(v)}(s) \leq 2^{i-1}(k - 1)) \right\} \\ & \bigcup_{\gamma=1}^q \left\{ \{s, v\} \in \text{IN}(v, 2\gamma - 1) \mid (s \in S) \wedge (\text{deg}_{\text{IN}(v, 2\gamma - 1)}(s) \leq 2\gamma - 1) \right\}. \end{aligned} \quad (8)$$

This construction is central because if $\text{IN}(v, 0) \neq \emptyset$ then there is a $2k$ -cycle in the graph, passing through S (Lemma 6). Let us first state two preliminary results on $\text{IN}(v)$ and $\text{OUT}(v)$.

Fact 5. *Let $i \in \{1, \dots, k-1\}$, and let $v \in V_i$. For any $s \in S \cap \text{IN}(v)$, we have*

$$\deg_{\text{OUT}(v)}(s) \leq 2^{i-1}(k-1).$$

Proof. By construction thanks to Eq. (8), and by the fact that $2q-1 \leq k-2 \leq 2^{i-1}(k-1)$. \square

Lemma 5. *Let $i \in \{1, \dots, k-1\}$, and let $v \in V_i$. For every edge $\{s, w\} \in \text{IN}(v)$, with $s \in S$ and $w \in W_0$, there is a path $(w, v_1, \dots, v_{i-1}, v)$ such that, for every $j \in \{1, \dots, i-1\}$, $v_j \in V_j$ and $\{s, w\} \in \text{OUT}(v_j)$. In particular, $\text{IN}(v) \subseteq E(S, W_0(v))$.*

Proof. The second claim, $\text{IN}(v) \subseteq E(S, W_0(v))$, is a direct consequence of the first. The first statement is proven by induction on $i \geq 1$. For $i = 1$, i.e., for $v \in V \setminus S$ with $c(v) = 1$, and for $\{s, w\} \in \text{IN}(v)$, it follows from Eq. (4) that there exists $v' \in V \setminus S$ such that $\{v, v'\} \in E$, $c(v') = 0$, and $\{s, w\} \in \text{OUT}(v')$. By construction, $\text{OUT}(v') = E(\{v'\}, S)$ when $c(v') = 0$ (see Eq. (3)), therefore $w = v'$, and the required path is just the edge $\{w, v\}$.

Assume now that the lemma holds for color $i \geq 1$. Fix $v \in V \setminus S$ such that $c(v) = i+1$. Again, by Eq. (4), if $\{s, w\} \in \text{IN}(v)$, then there exists $v' \in V \setminus S$ such that $\{v, v'\} \in E$, $c(v') = i$, and $\{s, w\} \in \text{OUT}(v')$. By construction (see Eq. (8)), $\text{OUT}(v') \subseteq \text{IN}(v')$. As a consequence, we can apply the induction hypothesis on $\{s, w\} \in \text{IN}(v')$, which provides us with a path $(w, v_1, \dots, v_{i-1}, v')$. We further extend this path using the edge $\{v, v'\} \in G[V \setminus S]$, which concludes the induction. \square

Cycle construction. The following lemma implies that if there exists $v \in V \setminus S$ such that $\text{IN}(v, 0) \neq \emptyset$, then there is a $2k$ -cycle intersecting S . Such a cycle is exhibited in Figure 1 for $k = 5$ and $i = 2$. Associated with Lemma 7 that provides the existence of such a v , it will prove Lemma 4.

Lemma 6. *Assume that $\text{IN}(v, 0) \neq \emptyset$ for some $v \in V \setminus S$ with $c(v) = i \in \{1, \dots, k-1\}$. Then there is a $2k$ -cycle C in G such that $C \cap S \neq \emptyset$.*

Proof. The existence of a $2k$ -cycle in G going through S is shown by constructing three simple paths P (Claim 1), P' and P'' (Claim 2), the union of which is a cycle. In Figure 1, path $P = (w, s_3, w_2, s_1, w'_2, s)$, path $P' = (w, v'_1, v)$ and path $P'' = (s, w'', v''_1, v)$.

Claim 1. *There is a simple path P with $2(k-i)$ nodes in $W_0 \cup S$, alternating between nodes in W_0 and nodes in S , the edges of which are all in $\text{IN}(v, 2q)$.*

Proof of claim. The construction is done by induction on γ . A path

$$P_\gamma = (s_{2\gamma+1}, \dots, w_2, s_1, w'_2, \dots, s'_{2\gamma+1}) \subseteq \text{IN}(v, 2\gamma)$$

is aimed at being extended to a path $P_{\gamma+1} \subseteq \text{IN}(v, 2\gamma+2)$. The extension uses the fact that both extremities of P_γ have some incident edges in $\text{IN}(v, 2\gamma)$.

For the base case $\gamma = 0$, let $s_1 = s'_1 \in S$ be any node with an incident edge in $\text{IN}(v, 0) \neq \emptyset$. Observe that, in this base case, the path is trivial, formed by a unique vertex of graph $\text{IN}(v, 0)$.

For the induction step, we start from a path $P_\gamma = (s_{2\gamma+1}, \dots, s_1, \dots, s'_{2\gamma+1})$ in $\text{IN}(v, 2\gamma)$, from S to S . Note that this path is also a path in $\text{IN}(v, 2\gamma+2) \supseteq \text{IN}(v, 2\gamma)$. Moreover, since $s_{2\gamma+1}$ has




Figure 1: The case of a 10-cycle (i.e., $k = 5$). In the figure, $\text{IN}(v, 0) \neq \emptyset$ as $v \in V_2$. Here we have $q = 1$, and thus the considered graphs for the proof are $\text{IN}(v, 0) \subseteq \text{IN}(v, 1) \subseteq \text{IN}(v, 2) \subseteq \text{IN}(v)$. Regarding the proof of Claim 1, we have $\deg_{\text{IN}(v, 1)}(s_1) > 1$, and thus there exist vertices w_2 and w'_2 in $\text{IN}(v, 1)$. Since $\deg_{\text{IN}(v, 2)}(w_2) > 2$, and since $\deg_{\text{IN}(v, 2)}(w'_2) > 2$, there are two vertices s_3 and s'_3 in $\text{IN}(v, 2)$. And since $\deg_{\text{IN}(v)}(s_3) > 8$, there exists a vertex w in $\text{IN}(v)$. Regarding the proof of Claim 2, we have $\deg_{\text{IN}(v)}(s) > 8$ and $\deg_{\text{OUT}(v')} (s'_3) \leq 4$. Therefore, there exists a vertex w'' in $\text{IN}(v)[s] \setminus (\{w, w_2, w'_2\} \cup \text{OUT}(v'))$.

some incident edge in $\text{IN}(v, 2\gamma)$, and since $\text{IN}(v, 2\gamma) \subseteq \text{IN}(v, 2\gamma + 1)$, $s_{2\gamma+1}$ must have large degree in $\text{IN}(v, 2\gamma + 1)$. Specifically, thanks to Eq. (7), we have

$$\deg_{\text{IN}(v, 2\gamma+1)}(s_{2\gamma+1}) = \deg_{\text{IN}(v, 2\gamma)}(s_{2\gamma}) > 2\gamma + 1.$$

Thus, there is a vertex $w_{2\gamma+2} \in W_0$ such that

$$w_{2\gamma+2} \notin \{w_2, w'_2, \dots, w_{2\gamma}, w'_{2\gamma}\} \text{ and } \{w_{2\gamma+2}, s_{2\gamma+1}\} \in \text{IN}(v, 2\gamma + 1).$$

Similarly, we can find $w'_{2\gamma+2}$ adjacent to $s'_{2\gamma+1}$ in $\text{IN}(v, 2\gamma + 1)$ (see edges (s_1, w_2) and (s'_1, w'_2) in Figure 1). The degrees of $w_{2\gamma+2}$ and $w'_{2\gamma+2}$ in $\text{IN}(v, 2\gamma + 2)$ are equal to their degrees in $\text{IN}(v, 2\gamma + 1)$, and thus they are both at least $2\gamma + 2$ (by Eq (6)). This establishes the existence of two new vertices $s_{2\gamma+3}$ and $s'_{2\gamma+3}$ (see edges (w_2, s_3) and (w'_2, s'_3) in Figure 1). The extended path is then

$$P_{\gamma+1} = (s_{2\gamma+3}, w_{2\gamma+2}, s_{2\gamma+1}, \dots, s_1, \dots, s'_{2\gamma+1}, w'_{2\gamma+2}, s'_{2\gamma+3}),$$

which concludes the induction step. The path P_q has $4q + 1$ vertices, and has its end points in S (path $P_1 = (s_3, w_2, s_1, w'_2, s'_3)$ in Figure 1).

If $k - i$ is even, then $4q + 1 = 2(k - i) + 1$. So the path P is merely P_q without one of its two end points.

If $k - i$ is odd (as in Figure 1), then the path P_q has $2(k - i) - 1$ vertices, and we need one more step to reach the desired length $2(k - i)$. Since $\{s_{2q+1}, w_{2q}\} \in \text{IN}(v, 2q)$, by using Eq. (5), we get

$$\deg_{\text{IN}(v)}(s_{2q+1}) = \deg_{\text{IN}(v, 2q)}(s_{2q+1}) > 2^{i-1}(k - 1) > 2q = k - i - 1.$$

Therefore we can select $w \notin \{w_2, w'_2, \dots, w_{2q}, w'_{2q}\}$ such that $\{s_{2q+1}, w\} \in \text{IN}(v, 2q)$ (see edge (s_3, w) in Figure 1). Then P is P_q augmented with w connected to its end point s_{2q+1} . \diamond

Let us now connect the two extremities of P to node v by using two well-colored node-disjoint paths.

Claim 2. *There are two simple paths:*

1. $P' = (w, v'_1, \dots, v'_{i-1}, v)$, with $i + 1$ nodes from the extremity $w \in W_0$ of P to v , such that, for every $j \in \{1, \dots, i - 1\}$, $v'_j \notin P$, and $c(v'_j) = j$.
2. $P'' = (s, w'', v''_1, \dots, v''_{i-1}, v)$, with $i + 2$ nodes from the extremity $s \in S$ of P to v , where $w'' \in \text{IN}(v) \setminus (P \cup P')$, and, for every $j \in \{1, \dots, i - 1\}$, $v''_j \notin P \cup P'$, and $c(v''_j) = j$.

Proof of claim. We first construct the path P' . Since the edge in P incident to w belongs to $\text{IN}(v, 2q) \subseteq \text{IN}(v)$, Lemma 5 directly gives us a path $(w, v'_1, \dots, v'_{i-1}, v)$ in $G[V \setminus S]$ such that $v'_j \in V_j$ for all $j \in \{1, \dots, i - 1\}$ ($P' = (w, v'_1, v)$ in Figure 1). This path P' intersects P only in w because the nodes of P are contained in $W_0 \cup S$, and the only intersection of P' with this set is w .

Let us now construct the second path P'' . Let s be the extremity of P in S . Let us denote by $\text{IN}(v)[s]$ the set of edges incident to s in $\text{IN}(v)$. On the one hand, by Claim 1, the edge of P incident to s is in $\text{IN}(v, 2q)$. By Eq. (5), the degree of s in $\text{IN}(v, 2q)$ is larger than $2^{i-1}(k - 1)$. Since $\text{IN}(v, 2q) \subseteq \text{IN}(v)$, we have

$$|\text{IN}(v)[s]| = \deg_{\text{IN}(v)}(s) > 2^{i-1}(k - 1).$$

On the other hand,

$$\left| \text{IN}(v)[s] \cap \left(P \cup \left(\bigcup_{j=1}^{i-1} \text{OUT}(v'_j) \right) \right) \right| \leq |W_0 \cap P| + \sum_{j=1}^{i-1} \deg_{\text{OUT}(v'_j)}(s).$$

Moreover, thanks to Fact 5, $\deg_{\text{OUT}(v'_j)}(s) \leq 2^{j-1}(k - 1)$ for every $j \in \{1, \dots, i - 1\}$. As a consequence,

$$\left| \text{IN}(v)[s] \cap \left(P \cup \left(\bigcup_{j=1}^{i-1} \text{OUT}(v'_j) \right) \right) \right| \leq k - i + \sum_{j=1}^{i-1} 2^{j-1}(k - 1) \leq 2^{i-1}(k - 1).$$

It follows that

$$\left| \text{IN}(v)[s] \cap \left(P \cup \left(\bigcup_{j=1}^{i-1} \text{OUT}(v'_j) \right) \right) \right| < |\text{IN}(v)[s]|.$$

Therefore, there exists an edge $e = \{s, w''\} \in \text{IN}(v)$ which is neither in P , nor in any $\text{OUT}(v'_j)$ for every $j \in \{1, \dots, i - 1\}$ (see Figure 1). Finally, by Lemma 5, there exists a path $(w'', v''_1, \dots, v''_{i-1}, v)$ in $G[V \setminus S]$ such that, for every $j \in \{1, \dots, i - 1\}$, $c(v''_j) = j$, and $e \in \text{OUT}(v''_j)$. Since $e \notin \text{OUT}(v'_j)$ for any $j \in \{1, \dots, i - 1\}$, and since nodes in P are either in S , or colored 0, none of the nodes v''_1, \dots, v''_{i-1} are in $P \cup P'$, which completes the proof of the claim. \diamond

By construction, the paths P, P' , and P'' from Claims 1-2 satisfy that $P \cup P' \cup P''$ is a cycle. By denoting $|Q|$ the number of vertices of a path Q , the length of $P \cup P' \cup P''$ is

$$(|P| - 1) + (|P'| - 1) + (|P''| - 1) = (2(k - i) - 1) + ((i + 1) - 1) + ((i + 2) - 1) = 2k.$$

which completes the proof of Lemma 6. \square

Congestion. To complete the proof of Lemma 4, we establish a lemma that allows us to bound the size of $W_0(v)$ as a function of the size of $\text{OUT}(v)$ whenever all sets $\text{IN}(v)$ are empty. Its contrapositive provides the hypothesis needed to apply Lemma 6, and prove Lemma 4.

Lemma 7. *Let $i \in \{1, \dots, k-1\}$. Assume that, for every $j \in \{1, \dots, i\}$, every $v_j \in V_j$ satisfies $\text{IN}(v, 0) = \emptyset$. Then, for every $v \in V_i$, $|W_0(v)| \leq 2^{i-1}(k-1)|S|$.*

Proof. We first show that, for every $w \in W_0(v)$, there is an edge incident to w in $\text{OUT}(v)$. In other words, we show that w is of degree at least 1 in $\text{OUT}(v)$. We first prove the more general inequality

$$\deg_{\text{OUT}(v)}(w) \geq k^2 - 2iq,$$

by induction on the index $i \in \{0, \dots, k-1\}$ of set V_i containing node v . In fact, this inequality also holds for the case $v \in W_0$, which constitutes the base case “ $i = 0$ ” of our induction, i.e., $W_0(v) = \{v\}$. In this case, the only scenario is $w = v$. Recall that $\text{OUT}(w) = E(\{w\}, S)$ (see Eq. (3)). Since $w \in W_0$ has at least k^2 neighbors in S , we have $\deg_{\text{OUT}(w)}(w) \geq k^2$, as desired.

Suppose now that the inequality holds for vertices colored $i-1$, and let $v \in V_i$ and $w \in W_0(v)$. By definition of $W_0(v)$, there exists a neighbor v' of v colored $i-1$ such that $w \in W_0(v')$. By induction,

$$\deg_{\text{OUT}(v')}(w) \geq k^2 - 2(i-1)q.$$

As $\text{OUT}(v') \subseteq \text{IN}(v)$, it also holds that

$$\deg_{\text{IN}(v)}(w) \geq k^2 - 2(i-1)q.$$

We are going to estimate the maximal number of edges incident to w that can be removed, i.e., that appear in $\text{IN}(v)$ but not in $\text{OUT}(v)$. By construction of $\text{OUT}(v)$ and $\text{IN}(v, 0)$, every edge $\{s, w\} \in \text{IN}(v)$ satisfies exactly one of the three following statements:

1. $\{s, w\} \in \text{IN}(v, 0)$;
2. There exists $\gamma \in \{1, \dots, q\}$ such that $\{s, w\} \in \text{IN}(v, 2\gamma) \setminus \text{IN}(v, 2\gamma - 1)$;
3. $\{s, w\} \in \text{OUT}(v)$.

By our assumption, $\text{IN}(v, 0) = \emptyset$, excluding case 1. Let us now consider case 2, if it occurs. Let $\gamma \in \{1, \dots, q\}$ be the largest integer such that $\text{IN}(v, 2\gamma) \setminus \text{IN}(v, 2\gamma - 1)$ contains some edge incident to w . Fix any $\{s, w\} \in \text{IN}(v, 2\gamma) \setminus \text{IN}(v, 2\gamma - 1)$. Since $\{s, w\} \notin \text{IN}(v, 2\gamma - 1)$, we have $\deg_{\text{IN}(v, 2\gamma)}(w) \leq 2\gamma$ (cf. Eq. 6), and none of the edges incident to w in $\text{IN}(v, 2\gamma)$ are kept in $\text{IN}(v, 2\gamma - 1)$. In other words, vertex w does not appear in graph $\text{IN}(v, 2\gamma - 1)$, and therefore it does not appear in any of the graphs $\text{IN}(v, \alpha)$, for any $\alpha \leq 2\gamma - 1$. That is, all edges incident to w that are in case 2 have been suppressed simultaneously, they all belong to $\text{IN}(v, 2\gamma) \setminus \text{IN}(v, 2\gamma - 1)$. Therefore, we conclude the induction step as follows:

$$\begin{aligned} \deg_{\text{OUT}(v)}(w) &= \deg_{\text{IN}(v)}(w) - \deg_{\text{IN}(v, 2\gamma)}(w) \\ &\geq [k^2 - 2(i-1)q] - 2\gamma \\ &\geq k^2 - 2iq. \end{aligned}$$

Using this inequality we deduce that

$$\deg_{\text{OUT}(v)}(w) \geq k^2 - 2iq \geq k^2 - i(k-i) \geq k^2 - (k-1)^2 > 0.$$

Thus every vertex $w \in W_0$ has at least one incident edge in $\text{OUT}(v)$, which implies that $|W_0(v)| \leq |\text{OUT}(v)|$. By Fact 5, in graph $\text{OUT}(v)$, each vertex contained in S has degree at most $2^{k-2}(k-1)$, hence $|\text{OUT}(v)| \leq 2^{k-2}(k-1)|S|$, which proves our lemma. \square

Proof of Lemma 4. Applying first the contrapositive of Lemma 7, Lemma 4 directly follows from Lemma 6. \square

Proof of Theorem 1. Let us first analyze the round complexity of Algorithm 1, and then prove that this algorithm accepts and rejects with the specified success probabilities.

Complexity. By construction, the threshold τ ensures that the complexities of the three calls $\text{color-BFS}(k, G[U], c, U, \tau)$, $\text{color-BFS}(k, G, c, S, \tau)$, and $\text{color-BFS}(k, G[V \setminus S], c, W, \tau)$ are each of at most $k\tau$ rounds. The for-loop at Instruction 7 is performed K times. So, overall, the total number of rounds performed by Algorithm 1 is $Kk\tau = O(\log^2(1/\varepsilon) \cdot 2^{3k} k^{2k+3} n^{1-1/k})$ rounds.

Acceptance without error. First note that a node u may reject in Algorithm 1 only while performing one of the calls to color-BFS . The rejection by u is thus caused by some identifier $\text{id}(u_0)$ that has traversed two paths, on the one hand a path u_0, u_1, \dots, u_k , and on the other hand a path $u_0, u_{2k-1}, u_{2k-2}, \dots, u_k$, where u_i is colored i , for every $i \in \{0, \dots, 2k-1\}$, and has reached $u = u_k$ from both neighbors u_{k-1} and u_{k+1} . These paths form together a $2k$ -cycle $u_0, \dots, u_{k-1}, u_k, u_{k+1}, \dots, u_{2k-1}$. Therefore, any node that rejects does so rightfully. In other words, if G contains no $2k$ -cycles, then the probability that all nodes accept is 1, as desired.

Rejection probability. We now prove that when there is a $2k$ -cycle C in G , Algorithm 1 rejects with probability $1 - \varepsilon$. We do the analysis by considering three cases, not necessarily disjoint but covering every possible scenarios. Each of them will reject with probability at least $1 - \varepsilon$, leading to the claimed global rejection. The probability events will be analyzed using Facts 1, 2, 3. For any considered $2k$ -cycle $C = (u_0, \dots, u_{2k-1})$, we will assume that $c(u_i) = i$ for every $i \in \{0, \dots, 2k-1\}$. This is indeed the case, with probability at least $1 - \varepsilon/3$, for at least one coloring of the loop (Fact 1 with $\alpha = \hat{\varepsilon}$). Moreover, we will also assume that $|S| \leq 4k^2 \hat{\varepsilon} n^{1-1/k}$, since this occurs with probability at least $1 - \varepsilon/3$ (Fact 2 with $\alpha = 2k^2 \hat{\varepsilon}$, using that $k \geq 2$). Lastly, when considering $u_0 \in C$ with $\deg(u_0) \geq n^{1/k}$, we will assume that $|N_G(u_0) \cap S| \geq k^2$, since this occurs with probability at least $1 - \varepsilon/3$ (Fact 3 with $\alpha = 2k^2 \hat{\varepsilon}$).

Case 1: $C \subseteq G[U]$. Let $C = (u_0, \dots, u_{2k-1})$ with $u_i \in U$ for every $i \in \{0, \dots, 2k-1\}$. By Lemma 1, when a coloring of the for-loop satisfies $c(u_i) = i$ for every $i \in \{0, \dots, 2k-1\}$, node u_k rejects in $\text{color-BFS}(k, G[U], c, U, \tau)$. Thus Algorithm 1 rejects with probability at least $1 - \varepsilon/3$.

Case 2: $C \cap S \neq \emptyset$. Let $C = (u_0, \dots, u_{2k-1})$ with $u_0 \in S$. If a coloring of the for-loop satisfies $c(u_i) = i$ for every $i \in \{0, \dots, 2k-1\}$, and if $|S| \leq 4k^2 \hat{\varepsilon} n^{1-1/k}$, then, by Lemma 2, u_k rejects in $\text{color-BFS}(k, G, c, S, \tau)$. Thus Algorithm 1 rejects with probability at least $1 - 2\varepsilon/3$, using union bound.

Case 3: $C \cap S = \emptyset$ and $C \cap (V \setminus U) \neq \emptyset$. Let $C = (u_0, \dots, u_{2k-1})$ with $u_0 \in V \setminus (S \cup U)$, in particular $\deg(u_0) \geq n^{1/k}$. Using union bound, with probability at least $1 - \varepsilon$, we get (1) a coloring of the for-loop satisfying $c(u_i) = i$ for every $i \in \{0, \dots, 2k-1\}$, (2) $|S| \leq$

$4k^2\hat{\varepsilon}n^{1-1/k}$, and (3) $|N_G(u_0) \cap S| \geq k^2$. Thanks to Lemma 3, since $u_0 \in W_0$, u_k rejects in color-BFS($k, G[V \setminus S], c, W, \tau$).

The analysis of these three cases completes the proof. \square

3 Quantum complexity

The objective of this section is to prove that by using an adaptation of Grover search to the distributed setting, Algorithm 1 can be sped up to $\tilde{O}(n^{1/2-1/2k})$ rounds in a quantum setting, as claimed in Theorem 2. The whole section is devoted to the proof of Theorem 2.

3.1 Preliminaries

3.1.1 Quantum amplification

We first review the framework for distributed quantum search [26] in the simple case of a search problem (instead of a general optimization problem), which can be seen as a distributed implementation of Grover's search [28], and its generalizations such as amplitude amplification [3]. Then we show how this can be used as in the sequential setting (see [2, Theorem 2.2]) for boosting the success probability of any one-sided error distributed algorithm. Those procedures require a centralized control, thus a specific node v_{lead} will have to play that role in distributed environments.

Let X be a finite set, and let $f: X \rightarrow \{0, 1\}$. Let v_{lead} be any fixed vertex of a network (e.g., an elected leader), whose purpose is to find $x \in X$ such that $f(x) = 1$ (assuming such elements exist). Vertex v_{lead} has two (randomized or quantum) distributed procedures at its disposal:

- **Setup:** Sample $x \in X$ (or create a superposition of elements $x \in X$) such that $f(x) = 1$ with probability p_{found} (when measuring x in the quantum setting). We let T_{Setup} denote the round-complexity of **Setup**.
- **Checking:** Compute $f(x)$, given input $x \in X$. We let T_{Checking} denote the round-complexity of **Checking**.

Assume one wants to amplify the success probability of **Setup** as follows: (1) vertex v_{lead} has to sample $x \in X$ in the support of **Setup**, and (2) x should satisfy $f(x) = 1$ with high probability whenever $p_{\text{found}} \geq \varepsilon$. A randomized strategy consists in executing **Setup** $\Theta(1/\varepsilon)$ times, and returning any of the sampled values of x satisfying $f(x) = 1$, using **Checking**. The success probability can be made arbitrarily high assuming that $p_{\text{found}} \geq \varepsilon$. This procedure has round complexity $O((T_{\text{Setup}} + T_{\text{Checking}})/\varepsilon)$. But quantumly we can do quadratically better in ε .

Lemma 8 (Distributed quantum search [26, Theorem 7]). *Let $f, v_{\text{lead}}, \text{Setup}, \text{Checking}, T_{\text{Setup}}, T_{\text{Checking}}$ and p_{found} be defined as above. For any $\delta > 0$, there is a quantum distributed algorithm with round complexity*

$$O\left(\log(1/\delta) \cdot \frac{1}{\sqrt{\varepsilon}} (T_{\text{Setup}} + T_{\text{Checking}})\right),$$

*such that (1) node v_{lead} returns x in the support of **Setup**, and (2) $f(x) = 1$ with probability at least $1 - \delta$ whenever $p_{\text{found}} \geq \varepsilon$.*

Remark. Note that, actually, the statement above slightly differs from Theorem 7 in [26] in several aspects. First, it is restricted to the simplest case of search problems (instead of optimization ones). Second, there is no mention of any initialization procedure that we let outside our framework for the sake of simplifying the presentation. Third, procedures **Setup** and **Checking** may be deterministic or randomized. Indeed, in the distributed setting, they can be converted into quantum procedures using standard techniques similar to those used for the sequential setting (see [26]). Last, point (1) was not explicit in [26], but this is a well-known property of Grover’s search and amplitude amplification.

As an application, and similarly to the sequential case (see for instance [2, Theorem 2.2]), Lemma 8 can be used to amplify the *success* probability of any one-sided error Monte-Carlo distributed algorithm. This amplification result is of independent interest, and may be used for other purposes, beyond the design of H -freeness decision algorithms. We call it *Distributed quantum Monte-Carlo amplification*.

Theorem 3 (Distributed quantum Monte-Carlo amplification). *Let \mathcal{P} be a Boolean predicate on graphs. Assume that there exists a (randomized or quantum) distributed algorithm \mathcal{A} that decides \mathcal{P} with one-sided success probability ε on input graph G , i.e.,*

- *If G satisfies \mathcal{P} , then, with probability 1, \mathcal{A} accepts at all nodes;*
- *If G does not satisfy \mathcal{P} , then, with probability at least ε , \mathcal{A} rejects in at least one node.*

Assume further that \mathcal{A} has round-complexity $T(n, D)$ for n -node graphs of diameter at most D . Then, for any $\delta > 0$, there exists a quantum distributed algorithm \mathcal{B} that decides \mathcal{P} with one-sided error probability δ , and round-complexity $\text{polylog}(1/\delta) \cdot \frac{1}{\sqrt{\varepsilon}}(D + T(n, D))$.

Proof. Let \mathcal{A} be the given distributed algorithm. First we recast the problem in the framework of Lemma 8. Let $X = \{\text{accept}, \text{reject}\}$ and $f: X \rightarrow \{0, 1\}$ be such that $f(\text{accept}) = 0$ and $f(\text{reject}) = 1$. Now define **Setup** as an algorithm which (1) selects a leader node v_{lead} , (2) runs \mathcal{A} , (3) broadcasts the existence of a rejecting node to v_{lead} , (4) outputs *accept* when all nodes accepts, and *reject* otherwise. Thus $T_{\text{Setup}} = T(n, D) + O(D)$. The procedure **Checking** is trivial since v_{lead} simply transforms *accept* in 0 and *reject* in 1, which requires no rounds, thus $T_{\text{Checking}} = 0$. Let us call \mathcal{B} the resulting algorithm after applying Lemma 8 with these procedures, and with the same parameter ε . By construction, \mathcal{B} has the claimed round-complexity. For the analysis of correctness, we consider two cases.

- Assume that \mathcal{A} accepts with probability 1, that is, \mathcal{A} samples *reject* with probability 0. Then \mathcal{B} also samples *reject* with probability 0 since it samples in the support of \mathcal{A} .
- Assume now that \mathcal{A} rejects with probability at least ε . Then \mathcal{A} samples *reject* with the same probability, that is $p_{\text{found}} \geq \varepsilon$. Thus \mathcal{B} will sample *reject* with probability at least $1 - \delta$.

We conclude that \mathcal{B} is an algorithm that solves \mathcal{P} with one-sided error probability δ . □

3.1.2 Diameter reduction

As demonstrated in [17], one can assume that the network has small diameter when looking for a forbidden connected subgraph. As observed in [8, 33], this assumption is valid for both quantum and randomized algorithms in the CONGEST model.

Lemma 9 ([17, Theorem 15]). *Let H be any fixed, connected k -vertex graph. Let \mathcal{A} be a randomized (resp., quantum) algorithm that decides H -freeness, with round-complexity $T(n, D)$ in n -node graphs of diameter at most D , and error probability $\rho(n) = o(1/(n \log n))$. Then there is a randomized (resp., quantum) algorithm \mathcal{A}' that decides H -freeness with round-complexity $\text{polylog}(n)(T(n, O(k \log n)) + k)$, and error probability at most $(c\rho(n) \cdot n \log n + 1/\text{poly}(n))$, for some constant $c > 0$.*

This result is based on a powerful technique of network decomposition from [19], which leads to the following (classical) preprocessing step, where nodes get one or more colors.

Lemma 10 ([17, Theorem 17]). *Let $G = (V, E)$ be an n -node graph and let $k \geq 2$ be an integer. There is a randomized algorithm with round complexity $k \text{polylog}(n)$ and error probability at most $1/\text{poly}(n)$ that constructs a set of clusters of diameter $O(k \log n)$ such that (1) each node is in at least one cluster, (2) the clusters are colored with $\gamma = O(\log n)$ colors, and (3) clusters of the same color are at distance at least k from each other in G .*

We now briefly review the reduction in Lemma 9 to convince the reader that this reduction applies to quantum protocols too. The following can be viewed as a sketch of proof of Lemma 9.

Construction in the proof of Lemma 9. We define \mathcal{A}' as follows. First \mathcal{A}' computes the network decomposition of Lemma 10 with parameter $2k + 1$. Therefore, the clusters of a same color are at pairwise distance at least $2k + 1$. Define $G(i, k)$ as the graph induced by all vertices of color $i \in [\gamma]$, and their k -neighborhood. Observe that each connected component of $G(i, k)$ is of diameter $O(k \log n)$. Indeed when we enlarge the clusters of color i as in $G(i, k)$ by adding their k -neighborhood, the subgraphs of $G(i, k)$ resulting from different clusters are disjoint. Moreover, they are not connected in $G(i, k)$. Therefore each connected component of $G(i, k)$ is obtained by the enlargement of one cluster, with its k -neighborhood. Thus the diameter of the component is at most the diameter of the cluster, plus $2k$.

Then, sequentially, for each color $i \in [\gamma]$, \mathcal{A}' runs \mathcal{A} in parallel on each connected component of $G(i, k)$. For the analysis, we first assume that the network decomposition has succeeded, and satisfies the conclusions of Lemma 10. Since connected components of $G(i, k)$ have diameter $O(k \log n)$, and since we used $O(\log n)$ different colors, the round-complexity of \mathcal{A}' is as claimed. The correctness of \mathcal{A}' is a direct consequence of the following observation. G contains a copy of H as a subgraph if and only if there exists a color $i \in [c]$ such that $G(i, k)$ contains a copy of H as a subgraph. For the overall error probability, simply take the union bound of all the potential failure events.

3.2 Quantum algorithm

3.2.1 Congestion reduction

Our classical algorithm has a constant success probability (cf. Theorem 1), and its three calls to the subroutine `color-BFS` have complexities upper bounded by $O(n^{1-1/k})$. In order to speed up this algorithm, we first reduce its success probability to $\varepsilon = \Theta(1/n^{1-1/k})$ such that the `color-BFS` subroutines have $O(1)$ round-complexities. Classically, one would have to repeat this process $1/\varepsilon$ times in order to get constant success probability, whereas $1/\sqrt{\varepsilon}$ quantum rounds suffices thanks to Theorem 3.

Recall that the complexity $O(n^{1-1/k})$ of `color-BFS` is a consequence of the setting of the threshold $\tau = O(n^{1-1/k})$. Indeed, every node only needs to send information that comes from at most τ nodes x_0 , leading to a congestion of τ , and thus a round complexity of τ . By activating each node x_0 with probability $\varepsilon = O(1/\tau)$ uniformly at random, one can expect to reduce the congestion to $O(\varepsilon\tau) = O(1)$, and to reduce the success probability to ε . Then, we apply Theorem 3 to get constant success probability within $O(D \times \sqrt{\tau})$ rounds, where D is the graph diameter. Note that, later on, we will eliminate the diameter dependence by employing the standard diameter reduction offered by Lemma 9. This technique has already been used for detecting C_4 with round complexity $\tilde{O}(n^{1/4})$ in an unpublished work [9], using directly Lemma 8, and not our new Theorem 3. The implementation is then a bit more involved. Each node initially samples a private token in $[1/\varepsilon]$, activated nodes are then decided by the leader choosing a token and broadcasting it. While the token sampling is decentralized, the activation of nodes of a given token must be centralized before applying Lemma 8, which yields also a multiplicative term D . However, for our purpose, we used a novel, and more general randomized distributed algorithm (Theorem 1), together with a more encapsulated quantum framework (Theorem 3) that eases its application.

3.2.2 Application to `color-BFS`

To speedup Algorithm 1, we replace `color-BFS`(k, H, c, X, τ) with a randomized protocol called `randomized-color-BFS`(k, H, c, X, τ), described in Algorithm 2. This randomized protocol has smaller round complexity, and thus has smaller success probability. Compared to `color-BFS`(k, H, c, X, τ), the modifications are the following:

- A node colored 0 does not systematically initiate a search, but does so with probability $1/\tau$ (cf. Instruction 1).
- Instead of using the threshold τ , the randomized algorithm uses a much smaller (constant) threshold, merely equal to 4 (cf. Instruction 5).

Algorithm 2 `randomized-color-BFS`(k, H, c, X, τ)

```

1: every node  $x \in X$  with  $c(x) = 0$  sends its ID to its neighbors in  $H$  with probability  $1/\tau$ 
2: for  $i = 1$  to  $k - 1$  do
3:   for every node  $v \in V(H)$  with  $c(v) = i$  (resp.,  $c(v) = 2k - i$ ) do
4:      $I_v \leftarrow \{\text{IDs received from neighbors in } H \text{ colored } i - 1 \text{ (resp. } 2k - i - 1)\}$ 
5:     if  $|I_v| \leq 4$  then
6:        $v$  forwards  $I_v$  to its neighbors in  $H$  colored  $i + 1$  (resp.  $2k - i - 1$ )
7:     end if
8:   end for
9: end for
10: for every node  $v \in V(H)$  colored  $k$  do
11:   if  $v$  receives a same ID from two neighbors respectively colored  $k - 1$  and  $k + 1$  then
12:      $v$  outputs reject
13:   end if
14: end for

```

The protocol `randomized-color-BFS` has round-complexity $O(1)$ since no node ever forwards more than 4 identifiers. Let us analyse its success probability. In order to take into account the fact that

the first instruction of randomized-color-BFS(k, H, c, X, τ) is randomized, we define the random set

$$X'_0 = \{x \in X_0 \mid x \text{ sends } \text{id}(x) \text{ at Instruction 1 of Algorithm 2}\}.$$

For $X'_0(v)$, we proceed accordingly to the definition of $X_0(v)$ in Eq. (1) and (2) for every $v \in H$ colored i , or $2k - i$, for every $i \in \{1, \dots, k - 1\}$. More formally,

$$X'_0(v) = \{x \in X_0(v) \cap X'_0\}.$$

Then $X'_0(v)$ satisfies a fact similar to Fact 4 for $X_0(v)$.

Fact 6.

- *Let $(v_1, v_2, \dots, v_{k-1})$ be a path in H such that $c(v_i) = i$ for every $i \in \{1, \dots, k - 1\}$. If $|X'_0(v_{k-1})| \leq 4$ then, for every $i \in \{1, \dots, k - 1\}$, $|I_{v_i}| \leq 4$ and $I_{v_i} = X'_0(v_i)$.*
- *Let $(v_{2k-1}, v_{2k-2}, \dots, v_{k+1})$ be a path in H such that $c(v_i) = i$ for every $i \in \{k+1, \dots, 2k-1\}$. If $|X'_0(v_{k+1})| \leq 4$ then, for every $i \in \{k+1, \dots, 2k-1\}$, $|I_{v_i}| \leq 4$ and $I_{v_i} = X'_0(v_i)$.*

Proof. We prove the result only for nodes colored between 1 and $k - 1$ as, by symmetry, the same holds for nodes colored between $2k - 1$ and $k + 1$. It follows from Instructions 1 and 4 that, for every $v \in H$, each identifier in I_v corresponds to a node in $X'_0(v)$. Now, along a path (v_1, \dots, v_{k-1}) , as long as $|X'_0(v_i)| \leq 4$ for every $i \in \{1, \dots, k - 1\}$, the condition of Instruction 5 is always satisfied for I_{v_i} , and every node in $X'_0(v_i)$ belongs to I_{v_i} . The result follows thanks to the fact that $X'_0(v_i) \subseteq X'_0(v_{k-1})$ for every $i \in \{1, \dots, k - 1\}$. \square

We can now state a first statement on the success probability of randomized-color-BFS, which is not yet suitable for applying Theorem 3. Note that the two cases considered in the lemma below do not cover all scenarios.

Lemma 11. *Protocol randomized-color-BFS(k, H, c, X, τ) satisfies the following.*

- *If there are no $2k$ -cycles in H , then, with probability 1, all nodes accept.*
- *If there is a $2k$ -cycle $C = (u_0, \dots, u_{2k-1})$ in H , with $u_0 \in X$, $c(u_i) = i$ for every $i \in \{0, \dots, 2k - 1\}$, $|X_0(u_{k-1})| \leq \tau$, and $|X_0(u_{k+1})| \leq \tau$, then, with probability at least $1/(2\tau)$, at least one node rejects.*

Proof. The first claim is straightforward. Let us prove the second claim. As each element of X_0 belongs to X'_0 with probability $1/\tau$, we get $u_0 \in X'_0$ with probability $1/\tau$. Assume that $u_0 \in X'_0$. Then $u_0 \in I_{u_1} \cap I_{u_{2k-1}}$. Conditioned to $u_0 \in X'_0$, the expected sizes of $X'_0(u_{k-1}) \setminus \{u_0\}$ and $X'_0(u_{k+1}) \setminus \{u_0\}$ are at most 1. Thus, by Markov's inequality, each of them is at least 4 with probability at most $1/4$. Therefore by the union bound,

$$\Pr \left[(|X'_0(u_{k-1}) \setminus \{u_0\}| \leq 3) \wedge (|X'_0(u_{k+1}) \setminus \{u_0\}| \leq 3) \right] \geq \frac{1}{2}.$$

It follows that, with probability at least $1/2\tau$, we have $u_0 \in I_{u_1} \cap I_{u_{2k-1}}$, $|X'_0(u_{k-1})| \leq 4$, and $|X'_0(u_{k+1})| \leq 4$. Using Fact 6, we derive that, for every $i \in \{1, \dots, k - 1\}$, $|I_{u_i}| \leq 4$ and $|I_{u_{2k-i}}| \leq 4$, meaning that u_k rejects after receiving u_0 's identifier along path (u_1, \dots, u_{k-1}) , and along path $(u_{2k-1}, \dots, u_{k+1})$. \square

3.2.3 Application to cycle detection

We now have all the ingredients to prove the upper bound in Theorem 2 regarding deciding C_{2k} -freeness in quantum CONGEST.

Lemma 12. *For every $k \geq 2$, there is a randomized distributed algorithm \mathcal{A} solving C_{2k} -freeness with one-sided success probability $1/(3\tau)$, and running in $k^{O(k)}$ rounds.*

Proof. Let \mathcal{A} be Algorithm 1 where color-BFS is replaced by randomized-color-BFS. To analyze \mathcal{A} we refer to the analysis of Algorithm 1 with $\epsilon = 1/3$ in Theorem 1, and only highlight the differences.

- Complexity: By construction, the new threshold 4 in randomized-color-BFS, instead of τ in color-BFS, ensures that the overall complexity is $4kK$, that is $O(k(2k)^{2k}) = k^{O(k)}$ rounds.
- Acceptation probability: As a node can only reject in randomized-color-BFS(k, H, c, X, τ), then by Lemma 11, Algorithm \mathcal{A} always accepts when there are no $2k$ -cycle in G .
- Rejection probability: We prove that if there is $2k$ -cycle C in G then Algorithm \mathcal{A} rejects with probability $\Omega(1/\tau)$. Observe that the proofs of Lemmas 1, 2 and 3 consist in upper bounding by τ the size of the sets $X_0(u_{k-1})$ and $X_0(u_{k+1})$, where $C = (u_0, \dots, u_{2k-1})$ is the considered cycle. In particular, whenever those lemmas are applied, we can use Lemma 11 instead, on the same considered cycle, in order to lower bound the rejection probability by $1/(2\tau)$. With this modification in mind, the rest of the proof of Theorem 1 applies, leading to a rejection probability of at least $(1 - \epsilon)/(2\tau) = 1/(3\tau)$.

This completes the proof. □

The upper bound for even cycles in Theorem 2 is now proved in the following lemma.

Lemma 13. *There is a quantum distributed algorithm \mathcal{A} solving C_{2k} -freeness with one sided error probability $1/\text{poly}(n)$, and running in $k^{O(k)} \cdot \text{polylog}(n) \cdot n^{1/2-1/2k}$ rounds.*

Proof. First, by Lemma 12, we get a randomized algorithm with one-sided success probability $1/3\tau$, and round-complexity $k^{O(k)}$. Then, we apply Theorem 3 to amplify this algorithm to get a one-sided error probability $1/\text{poly}(n)$, in $\text{polylog}(n) \cdot \sqrt{\tau} \cdot (k^{O(k)} + D)$ rounds. Finally, we get rid of the diameter factor by applying Lemma 9, which yields one-sided error probability $1/\text{poly}(n)$, and round-complexity

$$\text{polylog}(n)[(k^{O(k)} + O(k \log n))\sqrt{\tau} + 2k] = k^{O(k)}\text{polylog}(n) \cdot \sqrt{\tau} = k^{O(k)}\text{polylog}(n) \cdot n^{\frac{1}{2} - \frac{1}{2k}},$$

which completes the proof. □

3.3 Quantum Lower bounds

In this section, we prove the lower bounds stated in Theorem 2. Let us first consider even length cycles.

3.3.1 Even-Length Cycles

We show that, for any $k \geq 2$, deciding C_{2k} -freeness requires $\Omega(n^{1/4}/\log n)$ rounds in the quantum CONGEST model. The proof relies on a reduction from the Set-Disjointness problem in the two-party quantum communication framework, to the C_{2k} -freeness problem in the quantum CONGEST

model. The reduction is the same as in the classical case. For $k = 2$, it is based on the construction of [15], and for $k \geq 3$, on the one of [30]. Recall that, in the Set-Disjointness problem, each of two players, typically referred to as Alice and Bob, gets a subset of a universe of size N , say $[N]$, and they must decide whether there is an element in common in their subsets, while exchanging as few qubits (in the quantum setting) as possible. The total number of exchanged qubits is called the quantum communication complexity. The number of rounds of interactions between the two players is the round complexity. Note that the input of Alice (resp., Bob) can be represented as a binary string $x \in \{0, 1\}^N$ (resp., $y \in \{0, 1\}^N$), where $x_i = 1$ (resp., $y_i = 1$) whenever element i is in the input set of Alice (resp., Bob).

The reduction in [15] and [30] are based on picking a gadget graph G with N edges, e_1, e_2, \dots, e_N , and constructing a graph H composed on two subgraphs G_A and G_B of G connected by a perfect matching. These subgraphs are obtained from the inputs x and y of Alice and Bob, as follows. For every $i \in [N]$, Alice (resp., Bob) keeps edge e_i if $x_i = 1$ (resp., $y_i = 1$), and discards it otherwise. The gadget graphs in [15] and [30] are different, but the construction is the same once the gadget graph is fixed. The gadget graph in [15] has $N = \Theta(n^{3/2})$ edges, while the gadget graph in [30] has $N = \Theta(n)$ edges. Let us denote by H' the graph resulting from the former, and by H'' the graph resulting from the latter.

- It was shown in [15] that if there is a $T(n)$ -round CONGEST algorithm deciding C_4 -freeness in the n -node graphs H' then there is a $T(n)$ -round two-party communication protocol for Set-Disjointness of size $N = \Theta(n^{3/2})$ with communication complexity $O(T(n) \cdot n \cdot \log n)$ bits.
- Similarly, it was shown in [30] that if there is a $T(n)$ -round CONGEST algorithm deciding C_{2k} -freeness for $k \geq 3$ in the n -node graphs H'' then there is a $T(n)$ -round two-party communication protocol for Set-Disjointness of size $N = \Theta(n)$ with communication complexity $O(T(n) \cdot \sqrt{n} \cdot \log n)$ bits.

By exactly the same arguments, the same two properties hold for quantum CONGEST algorithms. Now, it is known [4] that, in the quantum two-party communication model, for every $r \geq 1$, any r -round communication protocol solving Set Disjointness for sets of size N has communication complexity $\Omega(r + \frac{N}{r})$ qubits. As a consequence, we get the following.

- For $k = 2$, any quantum algorithm solving C_4 -freeness in $T(n)$ rounds must satisfy

$$T(n) \cdot n \cdot \log n = \Omega(N/T(n)),$$

with $N = \Theta(n^{3/2})$, and therefore $T(n) = \Omega(n^{1/4}/\sqrt{\log n})$.

- For $k \geq 3$, any quantum algorithm solving C_{2k} -freeness in $T(n)$ rounds must satisfy

$$T(n) \cdot \sqrt{n} \cdot \log n = \Omega(N/T(n)),$$

with $N = \Theta(n)$. Therefore $T(n) = \Omega(n^{1/4}/\sqrt{\log n})$, as claimed.

3.3.2 Odd-Length Cycles

Let us now move on with establishing the lower bound stated in Theorem 2 for cycles of odd lengths. That is, we show that, for any $k \geq 2$, the round complexity of C_{2k+1} -freeness is $\tilde{\Theta}(\sqrt{n})$ in quantum CONGEST. Once again, we use reduction from set disjointness. A gadget graph H'' with $\Theta(n^2)$ edges has been constructed in [15], for which it was proved that if there is a $T(n)$ -round CONGEST

algorithm deciding C_{2k+1} -freeness in the n -node graphs H'' then there is a $T(n)$ -round two-party communication protocol for Set-Disjointness of size $N = \Theta(n^2)$ with communication complexity $O(T(n) \cdot n \cdot \log n)$ bits. The same holds for quantum CONGEST. Therefore, using again the lower bound for set disjointness in [4], we get that any quantum algorithm solving C_{2k+1} -freeness in $T(n)$ rounds must satisfy

$$T(n) \cdot n \cdot \log n = \Omega(N/T(n)),$$

with $N = \Omega(n^2)$. Therefore $T(n) = \Omega(\sqrt{n}/\log n)$, as claimed.

3.4 Quantum Upper Bound for Deciding Odd Cycles

We finally show that the lower bound for odd cycles C_{2k+1} with $k \geq 2$ established in the previous section is tight, which completes the proof of Theorem 2. Specifically, we show that there is a simple one-sided error randomized algorithm with success probability $\Omega(1/n)$ for deciding C_{2k+1} -freeness. As for the even case, the algorithm consists of looking for a well colored cycle, using a repetition of $K = O(1)$ random coloring c of the vertices of the graph, but with colors taken in $\{0, \dots, 2k\}$ (instead of $\{0, \dots, 2k-1\}$). For every coloring c , we apply a procedure similar to randomized-color-BFS($k, G, c, V, 4$) described in Algorithm 2. Indeed, for odd cycles the procedure only differs by the fact that we look for a cycle (u_0, \dots, u_{2k}) instead of (u_0, \dots, u_{2k-1}) . That is, a node colored k checks reception of a same identifier transmitted along a path colored $0, 1, \dots, k-1, k$ (of length k), and a path colored $0, 2k, \dots, k+1, k$ (of length $k+1$). For any node u , and for any coloring c , we have $|V_0(u)| \leq |V| \leq n$. Therefore, by using the same arguments as in Lemma 11, we get an algorithm with one-sided success probability $\Omega(1/n)$, and constant round-complexity. Indeed, whenever a $(2k+1)$ -cycle C is well colored, the node u_0 with color 0 in C has probability $1/n$ of sending its identifier. Moreover, if this happens then, with constant probability, no node of the cycle C will receive more than a constant number of identifiers. One can apply our quantum boosting technique (cf. Theorem 3), combined with the diameter reduction technique (cf. Lemma 9). This results in an algorithm with round-complexity $\tilde{O}(\sqrt{n})$, and error probability $1 - 1/\text{poly}(n)$.

3.5 Quantum Upper bound for Deciding Cycles of Bounded Length

For every $k \geq 2$, let $F_{2k} = \{C_\ell \mid 3 \leq \ell \leq 2k\}$. In this section, we show how to use our quantum algorithm deciding C_{2k} -freeness for deciding F_{2k} -freeness, in $\tilde{O}(n^{1/2-1/2k})$ rounds. In a nutshell, the (quantum) algorithm in [33] uses a different bound d_{max} compared to [10], and quantize only the search for heavy cycles. Instead, we keep the same bound $d_{max} = n^{1/k}$, but we quantize the search of both light and heavy cycles. Our quantum algorithm is rejecting with probability $1 - 1/\text{poly}(n)$ if there is a cycle of length $\ell \in \{3, \dots, 2k\}$, and is accepting otherwise.

Our quantum algorithm results from quantizing the classical algorithm for F_{2k} -freeness in [10], in the same way we quantized our classical algorithm for C_{2k} -freeness (see Algorithm 1). More specifically, we sequentially check the existence of cycles for pairs of lengths, by deciding $C_{2\ell-1}$ - and $C_{2\ell}$ -freeness conjointly, for every $\ell \in \{2, \dots, k\}$. For every ℓ , the decision algorithm works under the assumption that there are no cycles of length at most $2(\ell-1)$, as if there were such cycles, they would have been detected when testing a smaller pair of length. For a fixed ℓ , the algorithm is quasi-identical to Algorithm 1 with just four differences, listed below (we refer to the instructions in Algorithm 1).

- Instruction 5: We set W as the set of *all* neighbors of the set S , with no restrictions on the degrees of the nodes.
- Instruction 6: the threshold is now set to $\tau = 2np$.
- Instructions 10 and 11 are merged into a single $\text{color-BFS}(m, G, c, W, \tau)$.

In addition, in any color-BFS aiming at detecting $(2\ell - 1)$ -cycles, nodes colored $\ell + 1$ also forwards the received identifiers to neighbors colored $\ell - 1$, which reject if one of those identifiers is equal to one received from neighbors colored $\ell - 2$.

The detection of light cycles by $\text{color-BFS}(m, G[U], c, U, \tau)$ performs in $O(n^{1-1/\ell})$ rounds. On the other hand, regarding the simplified detection of heavy cycles using $\text{color-BFS}(m, G, c, W, \tau)$, if a node $v \in V$ receives more than $|S|$ identifiers of nodes in W , then two of them must be neighbors of a same node $s \in S$. The two sequentially colored paths from s to v induced by the forwarding of those two identifiers then form a cycle of length $\leq 2\ell$.

4 Conclusion

Thanks to our work, which complements previous work on the matter, the complexity landscape of deciding C_k -freeness in CONGEST is roughly as follows. In the classical setting, for $k > 3$ odd, deciding C_k -freeness takes $\tilde{\Theta}(n)$ rounds, and, for $k \geq 4$ even, deciding C_k -freeness takes $\tilde{O}(n^{1-2/k})$ rounds. In the quantum setting, for $k > 3$ odd, deciding C_k -freeness takes $\tilde{\Theta}(\sqrt{n})$ rounds, and, for $k \geq 4$ even, deciding C_k -freeness takes $\tilde{O}(n^{1/2-1/k})$ rounds. That is, quantum effects allowed us to design quantum algorithms with quadratic speedup compared to the best classical algorithms. Interestingly, in order to get quantum speed-up for even cycles, we first establish a trade-off between congestion, and round complexity. Such trade-offs might be exhibited for other distributed tasks, which would automatically lead to a quantum speed-up for these tasks as well.

We have already mentioned the difficulty of designing lower bounds for triangle-freeness, as well as for C_{2k} -freeness for $k \geq 3$. Yet, it seems that for C_{2k} -freeness and $k \geq 3$, the design of classical algorithms with complexity $O(n^{1-1/k-\alpha})$, or quantum algorithms with complexity $O(n^{1/2-1/2k-\alpha})$, with $\alpha > 0$, would require entirely new techniques.

Finally, it is worth mentioning that, as far as randomized algorithms for cycle detection are concerned, the randomized color-coding phases can often be replaced by deterministic protocols based on [20] (see, e.g., [25, 30] for its application). However, as for most distributed algorithms detecting subgraphs, our algorithm needs to pick a set S of vertices at random. For 4-cycles, randomization is not necessary, but we do not know whether randomization is necessary or not for detecting larger cycles of even length in a sublinear number of rounds.

References

- [1] Noga Alon, Raphael Yuster, and Uri Zwick. Color coding. In *Encyclopedia of Algorithms*, pages 335–338. 2016.
- [2] Andis Ambainis. Quantum search algorithms. *SIGACT News*, 35(2):22–35, 2004.
- [3] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum counting. In *25th Int. Coll. on Automata, Languages and Programming (ICALP)*, volume 1443 of *LNCS*, pages 820–831. Springer, 1998.

- [4] Mark Braverman, Ankit Garg, Young Kun Ko, Jieming Mao, and Dave Touchette. Near-optimal bounds on the bounded-round quantum communication complexity of disjointness. *SIAM Journal on Computing*, 47(6):2277–2314, 2018.
- [5] Keren Censor-Hillel. Distributed subgraph finding: Progress and challenges. In *48th Int. Coll. on Automata, Languages, and Programming (ICALP)*, volume 198 of *LIPICs*, pages 3:1–3:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [6] Keren Censor-Hillel. Distributed subgraph finding: Progress and challenges, 2022.
- [7] Keren Censor-Hillel, Yi-Jun Chang, François Le Gall, and Dean Leitersdorf. Tight distributed listing of cliques. In *32nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2878–2891, 2021.
- [8] Keren Censor-Hillel, Orr Fischer, François Le Gall, Dean Leitersdorf, and Rotem Oshman. Quantum distributed algorithms for detection of cliques. In *13th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 215 of *LIPICs*, pages 35:1–35:25. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [9] Keren Censor-Hillel, Orr Fischer, François Le Gall, Dean Leitersdorf, and Rotem Oshman. Private communication, 2023.
- [10] 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 *34th International Symposium on Distributed Computing (DISC)*, volume 179 of *LIPICs*, pages 33:1–33:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [11] Yi-Jun Chang and Thatchaphol Saranurak. Improved distributed expander decomposition and nearly optimal triangle enumeration. In *38th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 66–73. ACM, 2019.
- [12] Xavier Coiteux-Roy, Francesco d’Amore, Rishikesh Gajjala, Fabian Kuhn, François Le Gall, Henrik Lievonen, Augusto Modanese, Marc-Olivier Renou, Gustav Schmid, and Jukka Suomela. No distributed quantum advantage for approximate graph coloring, 2024.
- [13] Artur Czumaj and Christian Konrad. Detecting cliques in CONGEST networks. In *32nd International Symposium on Distributed Computing (DISC)*, volume 121 of *LIPICs*, pages 16:1–16:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [14] Danny Dolev, Christoph Lenzen, and Shir Peled. ”tri, tri again”: Finding triangles and small subgraphs in a distributed setting. In *26th Int. Symp. on Distributed Computing (DISC)*, volume 7611 of *LNCS*, pages 195–209. Springer, 2012.
- [15] Andrew Drucker, Fabian Kuhn, and Rotem Oshman. On the power of the congested clique model. In *33rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 367–376, 2014.
- [16] Talya Eden, Nimrod Fiat, Orr Fischer, Fabian Kuhn, and Rotem Oshman. Sublinear-time distributed algorithms for detecting small cliques and even cycles. In *33rd International Symposium on Distributed Computing (DISC)*, volume 146 of *LIPICs*, pages 15:1–15:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

- [17] Talya Eden, Nimrod Fiat, Orr Fischer, Fabian Kuhn, and Rotem Oshman. Sublinear-time distributed algorithms for detecting small cliques and even cycles. *Distributed Computing*, 35(3):207–234, June 2022.
- [18] Michael Elkin, Hartmut Klauck, Danupon Nanongkai, and Gopal Pandurangan. Can quantum communication speed up distributed computation? In *35th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 166–175, 2014.
- [19] Michael Elkin and Ofer Neiman. Distributed strong diameter network decomposition. *Theoretical Computer Science*, 922:150–157, 2022.
- [20] Paul Erdős, András Hajnal, and J. W. Moon. A problem in graph theory. *The American Mathematical Monthly*, 71(10):1107–1110, 1964.
- [21] Guy Even, Orr Fischer, Pierre Fraigniaud, Tzlil Gonen, Reut Levi, Moti Medina, Pedro Montealegre, Dennis Olivetti, Rotem Oshman, Ivan Rapaport, and Ioan Todinca. Three notes on distributed property testing. In *31st Int. Symp. on Distributed Computing (DISC)*, volume 91 of *LIPICs*, pages 15:1–15:30. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- [22] Orr Fischer, Tzlil Gonen, Fabian Kuhn, and Rotem Oshman. Possibilities and impossibilities for distributed subgraph detection. In *30th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 153–162, 2018.
- [23] Pierre Fraigniaud, Maël Luce, and Ioan Todinca. On the power of threshold-based algorithms for detecting cycles in the CONGEST model. In *30th Int. Coll. on Structural Information and Communication Complexity (SIROCCO)*, volume 13892 of *LNCS*, pages 459–481. Springer, 2023.
- [24] Pierre Fraigniaud, Pedro Montealegre, Dennis Olivetti, Ivan Rapaport, and Ioan Todinca. Distributed subgraph detection, 2017.
- [25] Pierre Fraigniaud and Dennis Olivetti. Distributed detection of cycles. *ACM Trans. Parallel Comput.*, 6(3):12:1–12:20, 2019.
- [26] François Le Gall and Frédéric Magniez. Sublinear-time quantum computation of the diameter in CONGEST networks. In *39th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 337–346, 2018.
- [27] Cyril Gavoille, Adrian Kosowski, and Marcin Markiewicz. What can be observed locally? In *Distributed Computing, 23rd International Symposium (DISC)*, pages 243–257, 2009.
- [28] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *28th ACM Symposium on the Theory of Computing (STOC)*, pages 212–219, 1996.
- [29] Taisuke Izumi and François Le Gall. Quantum distributed algorithm for the All-Pairs Shortest Path problem in the CONGEST-CLIQUE model. In *40th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 84–93, 2019.
- [30] Janne H. Korhonen and Joel Rybicki. Deterministic subgraph detection in broadcast CONGEST. In *21st International Conference on Principles of Distributed Systems (OPODIS)*, volume 95 of *LIPICs*, pages 4:1–4:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

- [31] François Le Gall, Harumichi Nishimura, and Ansis Rosmanis. Quantum advantage for the LOCAL model in distributed computing. In *36 Int. Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 49:1–49:14, 2019.
- [32] David Peleg. *Distributed computing: a locality-sensitive approach*. SIAM, 2000.
- [33] Joran van Apeldoorn and Tijn de Vos. A framework for distributed quantum queries in the CONGEST model. In *41st ACM Symposium on Principles of Distributed Computing (PODC)*, pages 109–119, 2022.