Fooling Views: A New Lower Bound Technique for Distributed Computations under Congestion

Amir Abboud* Keren Censor-Hillel† Seri Khoury† Christoph Lenzen‡

June 28, 2021

Abstract

We introduce a novel lower bound technique for distributed graph algorithms under bandwidth limitations. We define the notion of *fooling views* and exemplify its strength by proving two new lower bounds for triangle membership in the CONGEST(B) model:

1. Any 1-round algorithm requires $B \geq c\Delta \log n$ for a constant $c > 0$.
2. If $B = 1$, even in constant-degree graphs any algorithm must take $\Omega(\log^* n)$ rounds.

The implication of the former is the first proven separation between the LOCAL and the CONGEST models for deterministic triangle membership. The latter result is the first non-trivial lower bound on the number of rounds required, even for *triangle detection*, under limited bandwidth. All previous known techniques are provably incapable of giving these bounds. We hope that our approach may pave the way for proving lower bounds for additional problems in various settings of distributed computing for which previous techniques do not suffice.

*IBM Almaden Research Center. amir.abboud@ibm.com.

†Technion, Department of Computer Science, {ckeren,serikhoury}@cs.technion.ac.il. Supported in part by the Israel Science Foundation (grant 1696/14).

‡MPI for Informatics, Saarland Informatics Campus, clenzen@mpi-inf.mpg.de.

1 Introduction

In a group of n players with names of $O(\log n)$ bits, how many days does it take for someone to find out if there are three players that are all friends with each other? All computation is free, and in the beginning, all players know their friends. Then, on each day, each player can send $B = O(\log n)$ bits privately to each of its friends.

This multi-party communication puzzle is known as the triangle detection problem in the popular CONGEST model of distributed computing. Its complexity is poorly understood. In the naive protocol, each player (node) tells all its friends about all its other friends (sends its neighborhood to every neighbor). This takes a single round in the LOCAL model and $O(n)$ rounds in CONGEST. A clever randomized protocol of Izumi and Le Gall [30] provides a solution with $O(n^{2/3}(\log n)^{2/3})$ rounds in CONGEST, and this is essentially all we know about the problem. Before our work, it could not be ruled out that the problem can be solved in $O(1)$ rounds, or even a single round, even with a bandwidth of $B = 1$!

Open Question 1. *What is the round complexity of triangle detection in the* CONGEST *model of distributed computing?*

Triangle detection is an extensively studied problem in most models of computation. In the centralized setting, the best known algorithm involves taking the cube of the adjacency matrix of the graph. It runs in $O(n^{2.3729})$ time and was found using a complex computer program [39,62]. If one wishes to avoid the impractical matrix multiplication, the problem can be solved in $O(n^3/\log^4 n)$ time [8, 19, 65]. Other works have designed algorithms for sparse graphs [6], for real-world graphs (e.g. [60]), for listing all the triangles [12, 29], for approximately counting their number [34], for weighted variants (e.g. [20]), and much more (an exhaustive list is infeasible). Moreover, conjectures about the time complexity of triangle detection and of its variants [1–3, 28, 64] are among the cornerstones of fine-grained complexity (see [63]). Other highly non-trivial algorithms were designed for it in settings such as: distributed models [15, 16, 21, 22, 25, 30, 32], quantum computing [11, 14, 38, 40–42, 47, 61], and Map-Reduce [4, 58]. It is truly remarkable that such a basic problem has lead to so much research.

From a technical perspective, Open Question 1 is one of the best illustrations of the lack and necessity of new techniques for proving lower bounds in distributed computing. In this work, we present a novel lower bound technique, providing a separation between the LOCAL and CONGEST(B) models for a problem for which previous techniques are *provably incapable* of doing so.[1] This is the first progress towards answering Open Question 1 from the lower bounds side.

1.1 Prior lower bound techniques and their limits

To date, there are essentially two techniques for deriving lower bounds for distributed graph algorithms. The first is the *indistinguishability* technique of Linial [45], which is the main source for lower bounds in the LOCAL model, where the message size is unrestricted. This technique argues that any r-round algorithm, regardless of message size, can be seen as a function that maps the r-hop neighborhood of a node to its output. Here, the topology of the graph is labeled by the unique $O(\log n)$-bit node identifiers and any other input provided to the nodes; for randomized algorithms, we simply give each node an infinite string of unbiased random bits as part of the input.

[1] CONGEST(B) stands for the synchronous model with a bandwidth of B bits. In particular, the standard LOCAL and CONGEST models correspond to CONGEST(∞) and CONGEST($\log n$), respectively.

This technique has resulted in a large number of *locality* lower bounds, e.g. [13, 36, 44, 45, 50]. For these problems, it is a long-standing open question whether higher lower bounds can be found in the CONGEST(B) model, see e.g. [52]. Note that part of its appeal is that one can entirely "forget" about the algorithm: for instance, a 2-round coloring algorithm is just interpreted as a function assigning a color to each possible 2-neighborhood, and a correct algorithm must assign distinct colors to any pair of neighborhoods that may belong to adjacent nodes in *any* feasible input graph; more generally, this gives rise to the so-called *r-neighborhood graph,* and showing that r rounds are insufficient for coloring with c colors equates to showing that the chromatic number of the r-neighborhood graph is larger than c.

Unfortunately, as this technique does not take bandwidth restrictions into account, it cannot show any separation between the LOCAL and CONGEST(B) models. Triangle detection is possibly one of the most extreme examples for this, as it can be solved in a single round in LOCAL but seems to require $n^{\Omega(1)}$ rounds in CONGEST($\log n$). Additional examples are symmetry breaking problems, such as Maximal Independent Set (MIS) and $(\Delta+1)$-Coloring, upon which we elaborate in Section 5.

The second tool available for generating distributed lower bounds is the *information bottleneck* technique, first introduced implicitly by Peleg and Rubinovich [56]. Here, the idea is to reduce a 2-party communication complexity problem (typically set disjointness) to a distributed problem, and argue that a fast distributed algorithm under limited bandwidth would imply a protocol exchanging few bits. This approach yields a large number of strong lower bounds for a wide range of *global* problems, in which there is no bound on the distance up to which a local change in the input may affect the output. Examples are, e.g., [17, 26, 43, 49, 59], but a complete list would justify an entire survey by its own.

Lower bounds based on information bottlenecks can also be proven for local problems. For instance, Drucker et al. show an $n^{\Omega(1)}$ lower bound for detecting k-cliques or k-cycles for any fixed $k > 3$ [22]. However, this technique is inherently incapable of proving lower bounds for many problems. In particular, it completely fails for the problem of triangle detection. As discussed by Drucker et al., this is because no matter how we divide the nodes of the graph among the two players, one of them will know about the triangle. One may, in principle, hope for lower bounds based on multi-party communication complexity or information complexity, but to date no such result is known.

Intuitively, a lower bound for triangle detection must combine the two techniques. We have to argue that when a small number of bits is sent, the nodes do not have enough information to distinguish between a distance-2 neighborhood in which there is a triangle and one in which there isn't. Interestingly, Drucker et al. prove that no such technique is possible without breakthroughs in circuit complexity in the related CONGESTED-CLIQUE model [22], where nodes can send messages to all other nodes, not only neighbors.[2] Indeed, it is still open whether triangle detection can be solved in $n^{o(1)}$ rounds even in this powerful model (it is likely that it is [22]), but now we know that the lack of super-constant lower bounds can be blamed on our inability to prove "computational hardness" results in CS: in a similar vein, we do not know whether 3SAT can be solved in polynomial or even linear time. In contrast, no such barrier is known in the standard CONGEST(B) model, where communication is limited to the input graph. This is an embarrassing situation, as we have a huge gap, but no well-known barrier to blame it on.

[2]This result applies to small output, e.g. the triangle detection problem, which is a decision problem. For *listing* all triangles in the graph, a tight bound of $\Omega(n^{1/3}/\log n)$ holds [21, 30].

With a clever usage of Rusza-Szemerdi graphs, Drucker et al. were also able to prove a strong $n^{1-o(1)}$ lower bound for triangle detection [22], under the restriction that each node sends the same message to all of its neighbors in each round (in a broadcast fashion). More specifically, they show that a deterministic protocol in the CLIQUE-BROADCAST model (and therefore also Congest-Broadcast) requires $\Omega(n/e^{\sqrt{\log n}})$ rounds, and that (essentially) the same lower bound holds for randomized protocols under the Strong Exponential Time Hypothesis, a popular conjecture about the time complexity of k-SAT. Unfortunately, even under such conjectures, we do not know how to get any non-trivial lower bound in the standard CONGEST(B) model.

Finally, it is worth pointing out a subtlety about the statement that 2-player communication complexity cannot provide *any* lower bound for triangle detection. This statement is fully accurate only under the assumption that nodes initially know the identifiers of their neighbors, as this renders it trivial to infer from the joint view of two neighbors whether they participate in a triangle. This assumption is known as KT_1, where KT_i means *Knowledge of Topology* up to distance i (excluding edges with both endpoints in distance i), as was first defined in [7]. The difference between KT_0, in which a node knows only its own identifier, and KT_1, in which a node knows also the identifiers of its neighbors, has been a focus of abundant studies, in particular concerning the *message complexity* of distributed algorithms (see, e.g., [7, 23, 33, 51, 54] and references therein).

Note that acquiring knowledge on the neighbors' identifiers requires no more than sending $O(\log n)$ bits over each edge, so the distinction between KT_0 and KT_1 is insubstantial for the round complexity in CONGEST(B) for $B \in \Omega(\log n)$; KT_1 is therefore the default assumption throughout wide parts of the literature. However, in KT_0 a lower bound of $\Omega(\frac{\log n}{B})$ on the round complexity of triangle detection follows from a simple counting argument. As, ultimately, the goal is to show lower bounds of $\omega(\log n)$, we consider KT_1 in this work.

1.2 Our contribution

In this paper, we introduce *fooling views*, a technique for proving lower bounds for distributed algorithms with congestion. We are able to show the first non-trivial round complexity lower bounds on triangle detection in KT_1, separating the LOCAL and CONGEST(B) models:

1. Triangle membership[3] in one round requires $B \geq c\Delta \log n$ for a constant $c > 0$ (Section 3).
2. If $B = 1$, triangle detection requires $\Omega(\log^* n)$ rounds, even if $\Delta = 2$, and even if the size of the network is constant and n is the size of the namespace (Section 4).

We stress that we do not view our main contribution as the bounds themselves: while the bandwidth lower bound for single-round algorithms is tight, it hardly comes as a surprise that such algorithms need to communicate the entire neighborhood. Additionally, we do not believe that with 1-bit messages, extremely fast triangle detection is possible. Rather, we present a novel *technique* that enables to separate the two models, which is infeasible with prior lower bound techniques. We hope this to be a crucial step towards resolving the large gap between lower and upper bounds, which in contrast to other models is not justified by, e.g., conditional hardness results.

The basic idea of fooling views is that they combine reasoning about locality with bandwidth restrictions. Framing this in terms of neighborhood graphs, this would mean to label neighborhoods by the information nodes have initially and the communication the algorithm performs over

[3] In the triangle membership problem, every node must indicate whether it participates in a triangle.

the edges incident to a node. However, this communication depends on the algorithm and the communication received in earlier rounds, enforcing more challenging inductive reasoning to prove multi-round lower bounds.

To capture the intuition for our technique for $B = 1$, think about a node that receives the same messages from its neighbors regardless of whether it participates in a triangle or not as a *fooled node*. Intuitively, in a triangle $\{u, w, v\}$ of a given network, if one of the nodes u, v and w is able to detect the triangle after t rounds of communication, then it may simply inform the other two nodes about the triangle during round number $t + 1$. Thus, it is crucial to maintain a perpetual state of confusion for all nodes involved in the triangle.

However, if the task is to detect whether a specific triple of IDs is connected by a triangle or not, then the nodes can solve this by simply exchanging only one bit of communication. Accordingly, our goal is to keep a large subset of the *namespace* fooled as long as possible. To this end, think of a triangle $\{u, v, w\}$ for which none of u, v and w is able to detect the triangle as a *fooled triangle*. Our main idea is to show that if there are many fooled triangles after t rounds, then there are many triangles among them that are fooled after $t + 1$ rounds as well. In order to express this intuition, one of our ingredients in the proof is the following extremal combinatorics result by Paul Erdös [24].

Theorem 1 ([24], Theorem 1)**.** *Any k-uniform hypergraph of n nodes which contains at least $n^{k-\ell^{1-k}}$ edges, must contain a complete k-partite k-uniform hypergraph such that each part of it is of size ℓ.*

Using this theorem of Erdös, we are able to show that if there are many fooled triangles after t rounds, then there is a set of nodes such that each triple in the set is a fooled triangle after $t + 1$ rounds. Blending counting and indistinguishability arguments with this theorem, we can derive our lower bound for multi-round algorithms.

Our $\Omega(\log^* n)$ bound serves as a proof of concept that our technique has the power to break through the bounds of previous techniques by demonstrating that this is indeed possible. Note that purely information-theoretic reasoning runs into the obstacle that in the KT_1 model, $\Theta(\log n)$ bits have already crossed each edge "before the algorithm starts." Accordingly, we argue that our approach represents a qualitative improvement over existing techniques. Proving lower bounds higher than $\log^* n$ requires new ideas, but we are hopeful that combining our technique with a more sophisticated analysis will lead to much higher lower bounds, both for triangle detection and for other non-global problems discussed above.

As an additional indication that the proposed technique is of wider applicability, we apply it to k-cycle detection for $k > 3$, for showing lower bounds on the bandwidth of optimal-round algorithms. However, here the information bottleneck technique is applicable again, and we do not obtain stronger bounds using our fooling views. The details of these constructions are given in Appendix A; the main body of the paper focuses on triangle detection. We conclude with open questions in Section 5.

1.3 Further related work

Edge-crossings. The basic topology components for our lower bound in Section 4 are triangles and 6-cycles. The main hardness that we show for a node in deciding whether it participates in a triangle or not comes from not knowing the neighbors of its neighbor. That is, a node is unable to distinguish between two triangles and one 6-cycle, because only difference between these two cases

is a single *edge crossing*, which are two node-disjoint edges for which we swap the endpoints from $\{w, x\}, \{w', x'\}$ to $\{w, x'\}, \{w', x\}$. Edge crossings have previously aided the construction of lower bounds, such as lower bounds for message complexity of broadcast [7] or of symmetry breaking [51], as well as lower bounds for proof-labeling schemes [10].

The Congest(1) model. While we view our results for the CONGEST(1) model more as a proof of concept for our technique rather than as a bound that attempts to capture the true complexity, this model has been attracting interest by itself in previous work. It has been shown in [48] that the cornerstone $O(\log n)$-round algorithms for maximal independent set [5,46] can be made to work with even with a bandwidth of $B = 1$, and this problem was studied also in, e.g., [9,35]. Using the standard framework of reduction from 2-party communication complexity, a $\Omega(\sqrt{n})$ lower bound for the number of rounds required for 4-cycle detection can be directly deduced from [22]. In fact, all lower bounds obtained using this framework are with respect to the bandwidth B, and therefore imply lower bounds also for the case of $B = 1$.

2 Model and Definitions

Our model is a network of n nodes, each having an ID in $[N]$, for some polynomial N in n. Each node starts with knowledge of its ID as well as the IDs of its neighbors. This is known as the KT_1 model, and differs from the KT_0 model in which each node starts only with knowledge of its own ID. The nodes communicate in synchronous rounds, in which each node can send a B-bit messages to each of its neighbors.

The model we consider is the CONGEST(B) model [55], where B is the bandwidth that is given for each message[4]. We will focus on the following problem.

Definition 1. *[**Triangle Membership**]. In the triangle membership problem, each node needs to detect whether it is a part of a triangle.*

We remark that our second lower bound also applies to the *triangle detection* problem, where it is sufficient that *some* node learns that there is a triangle, without being able to tell who participates. This is also the guarantee given by the sublinear-round algorithm of Izumi and Le Gall [30].

3 A Bandwidth Lower Bound for 1-round Triangle Membership

In this section we show the following theorem.

Theorem 2. *The triangle membership problem cannot be solved by a single-round algorithm in the* CONGEST(B) *model unless* $B \geq \log((\frac{n-2}{2(\Delta-1)})^{\Delta-1})$.

Observe that this implies an $\Omega(\Delta \log n)$ lower bound on B for $\Delta = O(n^{1-\epsilon})$, and an $\Omega(n)$ lower bound on B for $\Delta = \frac{n-2}{4} + 1$. This lower bound is significantly easier to obtain than the one in Section 4. This has the advantage of demonstrating the technique in a simpler context before delving into the proof of Theorem 3.

[4]For simplifying the exposition, we will assume that nodes send B bits in each round and cannot send less bits or remain silent. It is easy to verify that this only affects the constants in our asymptotic notation.

The main line of proof is to show that if the size of the messages is less than $\log(\frac{n-2}{2(\Delta-1)})^{\Delta-1}$, then there are three nodes $u, w, v \in V$, such that if the only neighbors of u are v and w, then u receives the same messages from v and w during the single communication round regardless of whether v and w are connected (see Figure 1). This is the standard notion of indistinguishability, given here under bandwidth restrictions for the first time.

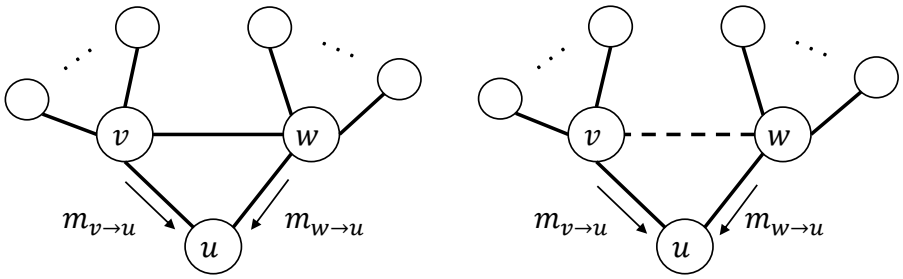

Figure 1: The node u receives the same messages regardless of whether v and w are connected.

For simplicity, we assume that each node has exactly Δ neighbors, except some special node u, which is fixed in the rest of this section, which has only two neighbors. Denote by $N(v)$ the set of neighbors of a node v, and let $m_{v\to u}(S)$ denote the message sent from v to u during the single round, given that $N(v) = S \cup \{u\}$.

The following two notions of fooling sets of nodes and fooling nodes are what allows us to capture indistinguishability in this setting. These are the specific shapes that our notion of *fooling views* takes for obtaining this result.

Definition 2. *Let $v \in V \setminus \{u\}$. A set of nodes S is called a* (v,u)-fooling set *if there is another set of nodes $S' \neq S$ such that $m_{v\to u}(S) = m_{v\to u}(S')$.*

A node $w \in V \setminus \{v,u\}$ is called a (v,u)-fooling node *if there are two sets of nodes $S \neq S'$, such that $w \in S$ and $w \notin S'$, and $m_{v\to u}(S) = m_{v\to u}(S')$.*

We denote by $F_{nodes}(v,u)$ the set of (v,u)-fooling nodes.

Our first step towards proving Theorem 2 is to show that for each $v \in V \setminus \{u\}$ there are many (v,u)-fooling nodes.

Lemma 1 (Many fooling nodes). *If $B < \log((\frac{n-2}{2(\Delta-1)})^{\Delta-1})$, then for each $v \in V \setminus \{u\}$ it holds that $|F_{nodes}(v,u)| \geq \frac{n-2}{2}$.*

Proof. Assume towards a contradiction that $|F_{nodes}(v,u)| < \frac{n-2}{2}$. Thus, there are $k' \geq \frac{n-2}{2}$ non (v,u)-fooling nodes, denote this set of non (v,u)-fooling nodes by $NF_{nodes}(v,u)$, and denote the family of sets of size $\Delta-1$ over nodes in $NF_{nodes}(v,u)$ by $NF_{sets}(v,u)$. It holds that

$$|NF_{sets}(v,u)| = \binom{|NF_{nodes}(v,u)|}{\Delta-1} \geq \binom{\frac{n-2}{2}}{\Delta-1} \geq \left(\frac{n-2}{2(\Delta-1)}\right)^{\Delta-1}$$

Observe that v must send a unique message to u on each of these sets, since otherwise, there are two sets $S_1 \neq S_2 \in NF_{sets}(v,u)$ such that $m_{v\to u}(S_1) = m_{v\to u}(S_2)$, and therefore, by Definition 2,

at least one node in $NF_{nodes}(v,u)$ is (v,u)-fooling. This implies that $B \geq \log((\frac{n-2}{2(\Delta-1)})^{\Delta-1})$, a contradiction. □

The strength of having many (v,u)-fooling nodes for every node v is that it implies that there is a node w^* that is a (v,u)-fooling node for many nodes v.

Lemma 2. *If $B < \log((\frac{n-2}{2(\Delta-1)})^{\Delta-1})$, then there is a node $w^* \in V$ and a set of nodes P_{w^*} of size $\frac{n-2}{2}$, such that for each $v \in P_{w^*}$, it holds that $w^* \in F_{nodes}(v,u)$.*

Proof. For each $w, v \in V$, let $X_{w,v}$ be an indicator defined as follows:

$$X_{w,v} = \begin{cases} 1 & \text{if } w \in F_{nodes}(v,u) \\ 0 & otherwise \end{cases}$$

Since

$$\sum_{w \in V \setminus \{u\}} \sum_{v \in V \setminus \{u\}} X_{w,v} = \sum_{v \in V \setminus \{u\}} \sum_{w \in V \setminus \{u\}} X_{w,v},$$

Lemma 1 gives that:

$$\sum_{w \in V \setminus \{u\}} \sum_{v \in V \setminus \{u\}} X_{w,v} \geq (n-1) \cdot \frac{n-2}{2}$$

Therefore, there must be a node w^* and a set of nodes P_{w^*} of size $\frac{n-2}{2}$, such that for each $v \in P_{w^*}$, it holds that $X_{w,v} = 1$. [5] □

Lemma 2 is what allows us to prove that u cannot solve triangle membership, as follows.

Proof of Theorem 2: Let w^* be the node provided by Lemma 2. It holds that $\binom{|P_{w^*}|}{\Delta-1} \geq (\frac{n-2}{2(\Delta-1)})^{\Delta-1}$. If w^* sends to u less than $\log(\frac{n-2}{2(\Delta-1)})^{\Delta-1}$ bits, then there are two sets of nodes $\{v_1, ..., v_{\Delta-1}\} \neq \{v'_1, ..., v'_{\Delta-1}\} \subseteq P_{w^*}$, such that $m_{w^* \to u}(\{v_1, ..., v_{\Delta-1}\}) = m_{w^* \to u}(\{v'_1, ..., v'_{\Delta-1}\})$. Let v^* be a node such that $v^* \in \{v_1, ..., v_{\Delta-1}\}$ and $v^* \notin \{v'_1, ..., v'_{\Delta-1}\}$. Since $v^* \in P_{w^*}$, by Definition 2, it holds that there are two sets of nodes $S \neq S'$, such that $w^* \in S$ and $w^* \notin S'$, and $m_{v^* \to u}(S) = m_{v^* \to u}(S')$.

To summarize, if the size of the messages is less than $\log(\frac{n-2}{2(\Delta-1)})^{\Delta-1}$, then there are two nodes $w, v \in V \setminus \{u\}$, such that if u is connected to both v and w, then u receives the same messages from v and w during the single communication round regardless of whether v and w are connected (see Figure 1). □

4 A Round Lower Bound for Triangle Detection with $B = 1$

Our main goal in this section is to prove the following theorem.

Theorem 3. *Any algorithm for triangle membership in the* CONGEST(1) *model requires $\Omega(\log^* n)$ rounds, even for graphs of maximum degree 2.*

[5] In fact, one can show that there are many such nodes as w^*, but we do not use this in this section.

The hard instances for this setting are as simple as tripartite graphs with degree 2, as follows. Let $\mathcal{G} = \{(V = A_1 \dot\cup A_2 \dot\cup A_3, E) \mid E \subseteq (A_1 \times A_2) \cup (A_1 \times A_3) \cup (A_2 \times A_3)\}$ be the family of tripartite graphs such that for each $G \in \mathcal{G}$, it holds that $|A_1| = |A_2| = |A_3| = n/3$, and for every $i \in \{1, 2, 3\}$, each node in A_i has exactly one neighbor in each of A_j and A_k, where $j \neq k \neq i \in \{1, 2, 3\}$.

In a nutshell, we show that there must be 6 nodes that cannot detect whether they constitute a 6-cycle or two triangles. Remarkably, this simple intuitive task requires us to develop novel machinery and make use of known results in extremal graph theory. Our approach has the compelling feature that – up to one round – in this setting triangle detection is equivalent to triangle membership, as a node detecting a triangle can infer that it is part of a triangle itself and just needs to inform its neighbors. Thus, we obtain the same lower bound for the, in general, possibly easier problem of triangle detection. Moreover, the lower bound depends only on the set of *possible* IDs, not the actual number of nodes in the graph.

Theorem 4. *Any algorithm for triangle detection in the* CONGEST(1) *model requires* $\Omega(\log^* N)$ *rounds, even for graphs of maximum degree 2 with only 6 nodes, given a namespace of size* N*.*

The main challenge in proving a lower bound for more than one round is that for $t > 1$ the communication between the nodes in round t does not depend only on their IDs and their sets of neighbors (as is the case for the first round), but rather it depends also on their *views* after $t - 1$ rounds, i.e., on the set of messages that the nodes receive during each of the first $t-1$ rounds. The view of a node after $t - 1$ rounds may depend on all the nodes in its $(t - 1)$-hop neighborhood.

Nevertheless, since the family of graphs $\mathcal{G}$ is defined such that, for each $G = (A_1 \dot\cup A_2 \dot\cup A_3, E) \in \mathcal{G}$, it holds that each node in each part of G has exactly one neighbor in each of the other two parts, any cycle in any $G \in \mathcal{G}$ must be a connected component. Therefore, if there is a triple $(u, w, x) \in A_1 \times A_2 \times A_3$ that is connected by a triangle (u, w, x) in G, then the communication between the nodes u, w and x during any round in G depends only on the nodes u, w and x. Similarly, if there is a 6-cycle $(u_1, w_1, x_1, u_2, w_2, x_2)$ in G, then the communication between the nodes $u_1, w_1, x_1, u_2, w_2, x_2$ during any round in G depends only on the nodes $u_1, w_1, x_1, u_2, w_2, x_2$.

Our main line of proof is to show that for any algorithm for triangle membership with $B = 1$, there exist $u_1, w_1, x_1, u_2, w_2, x_2 \in A_1 \dot\cup A_2 \dot\cup A_3$ such that u_1 receives the same messages during the first $\Omega(\log^* n)$ rounds in two different scenarios, the first of which is a scenario in which u_1 participates in a triangle (u_1, x_1, w_1), and the second is a scenario in which u_1 participates in a 6-cycle $(u_1, x_1, w_2, u_2, x_2, w_1)$. See Figure 2 for an illustration.

Given three nodes $u, w, x \in A_1 \dot\cup A_2 \dot\cup A_3$, we denote by $m^t_{w \to u}((u, w, x))$ the message sent from w to u during round t, given that u, w and x are connected by a triangle. Similarly, given six nodes $u_1, w_1, x_1, u_2, w_2, x_2 \in A_1 \dot\cup A_2 \dot\cup A_3$, we denote by $m^t_{w_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1))$ the message sent from w_1 to u_1 during round t, given that the nodes $u_1, w_1, x_2, u_2, w_2, x_1$ are connected by a 6-cycle $(u_1, x_1, w_2, u_2, x_2, w_1)$.

The structure of the proof for the lower bound on the number of rounds is as follows. In Section 4.1 we present the notion of *fooling sets of triangles* which shares a similar spirit to Definition 2 and is the basis of our fooling views technique in this section. Next, in section 4.2 we present the connection to 6-cycles and deduce our main result.

4.1 Fooling sets of triangles

Definition 3. *Fix a pair of nodes* $(u, w) \in (A_i \times A_j) \cap E$*, for* $i \neq j$*, and let* $X = \{x_1, ..., x_{|X|}\} \subseteq A_k$ *for* $k \neq i, j$*. A set of triangles* $\{(u, w, x) \mid x \in X\}$ *is called a* $(w, u)^t$-fooling set of triangles

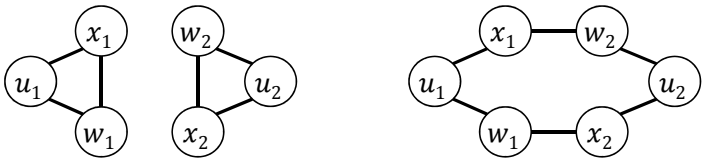

Figure 2: The structures we use are the triangle (u_1, x_1, w_1) and the 6-cycle $(u_1, x_1, w_2, u_2, x_2, w_1)$.

if w sends the same message to u during each of the first t rounds, in each of the triangles in $\{(u, w, x) \mid x \in X\}$. Formally, for each $1 \leq i \leq t$,

$$m^i_{w\to u}((u, w, x_1)) = m^i_{w\to u}((u, w, x_2)) = ... = m^i_{w\to u}((u, w, x_{|X|})).$$

We extend the notion of a fooling set of triangles to the notion of a *fooling rectangle of triangles*:

Definition 4. *Fix a node $u \in A_i$. A rectangle $W \times X \subseteq A_j \times A_k$, for $j \neq k \neq i$, is called a* u^t-fooling rectangle of triangles *if it satisfies the following properties.*

1. *For each $w \in W$, it holds that $\{(u, w, x) \mid x \in X\}$ is a $(w, u)^t$-fooling set of triangles.*
2. *For each $x \in X$, it holds that $\{(u, w, x) \mid w \in W\}$ is an $(x, u)^t$-fooling set of triangles.*

Finally, we extend the notion of a fooling rectangle of triangles to a *fooling cube of triangles*:

Definition 5. *A cube $U \times W \times X \subseteq A_1 \times A_2 \times A_3$ is called a* t-fooling cube of triangles *if:*

1. *For each $u \in U$, it holds that $W \times X$ is a u^t-fooling rectangle of triangles.*
2. *For each $w \in W$, it holds that $U \times X$ is a w^t-fooling rectangle of triangles.*
3. *For each $x \in X$, it holds that $U \times W$ is a x^t-fooling rectangle of triangles.*

The following observation follows immediately from Definition 5.

Observation 1. *$A_1 \times A_2 \times A_3$ is a 0-fooling cube of triangles.*

Our next step towards proving Theorem 3 is to show that any t-fooling cube of triangles of size s contains a $(t+1)$-fooling cube of triangles of size $\Omega(\sqrt{\log\log\log s})$.

Lemma 3. *If there is a t-fooling cube of triangles $U \times W \times X \subseteq A_1 \times A_2 \times A_3$ of size $|U| = |W| = |X| = s$, then there is a constant β and a $(t+1)$-fooling cube of triangles $U' \times W' \times X' \subseteq U \times W \times X$ of size $|U'| = |W'| = |X'| = \beta\sqrt{\log\log\log s}$.*

The key combinatorial ingredient for proving Lemma 3 is the following corollary of Theorem 1.

Corollary 1 (of Theorem 1)**.** *For every constant α, for sufficiently large s, any $s \times s \times s$ boolean cube that contains at least $e^{-1/\alpha^2}s^3$ entries that are 1, contains a 1-monochromatic subcube, such that each side of the subcube is of size $\alpha\sqrt{\log s}$.*

To prove Lemma 3, we first prove a weaker claim, which guarantees the existence of a sufficiently large cube of triangles satisfying only the first property in Definition 5. That is, we prove the existence of a cube $U' \times W' \times X' \subseteq U \times W \times X$ such that for each $u \in U'$, it holds that $W' \times X'$ is a u^{t+1}-fooling rectangle of triangles and $|U'| = |W'| = |X'| = \Theta(\sqrt{\log n})$. Then, applying this claim to the 3 different sides of the cube $U \times W \times X$ gives Lemma 3.

Claim 1. *If there is a t-fooling cube of triangles $U \times W \times X \subseteq A_1 \times A_2 \times A_3$ of size $|U| = |W| = |X| = s$, then there is a constant γ and a cube $U' \times W' \times X' \subseteq U \times W \times X$ of size $|U'| = |W'| = |X'| = \gamma\sqrt{\log s}$, such that for each $u \in U'$, it holds that $W' \times X'$ is a u^{t+1} fooling rectangle of triangles.*

Proof. Fix $u \in U$, and consider the messages that u receive from its neighbors in $W \cup X$ during round $t+1$. Since $B = 1$, for each $w \in W$, it holds that there exist $s/2$ nodes $x_1, ..., x_{s/2} \in X$ such that $m^{t+1}_{w\to u}((u,w,x_1)) = m^{t+1}_{w\to u}((u,w,x_2)) = ... = m^{t+1}_{w\to u}((u,w,x_{s/2}))$. Denote this specific message by $\tilde{m}^{t+1}_{w\to u}$. Therefore, defining an indicator variable $Y^{t+1}_{u,w,x}$ for each $(w,x) \in W \times X$ as follows

$$Y^{t+1}_{u,w,x} = \begin{cases} 1 & \text{if } m^{t+1}_{w\to u}((u,w,x)) = \tilde{m}^{t+1}_{w\to u} \\ 0 & otherwise \end{cases}$$

gives that

$$\sum_{w\in W}\sum_{x\in X} Y^{t+1}_{u,w,x} = s^2/2,$$

which implies

$$\sum_{x\in X}\sum_{w\in W} Y^{t+1}_{u,w,x} = \sum_{w\in W}\sum_{x\in X} Y^{t+1}_{u,w,x} = s^2/2.$$

It follows that there are at least $s/4$ nodes $x_1, ..., x_{s/4}$ in X, such that each $x_i \in \{x_1, ..., x_{s/4}\}$ has a set of nodes $P(x_i) \subseteq W$ of size $s/4$, such that for each $w \in P(x_i)$, it holds that $Y^{t+1}_{u,w,x_i} = 1$. Furthermore, for each x_i, it holds that there are $|P(x_i)|/2$ nodes $w^{x_i}_1, ..., w^{x_i}_{|P(x_i)|/2} \in P(x_i)$ such that $m^{t+1}_{x_i\to u}((u, w^{x_i}_1, x_i)) = ... = m^{t+1}_{x_i\to u}((u, w^{x_i}_{|P(x_i)|/2}, x_i))$. Denote this specific message by $\tilde{m}^1_{x_i\to u}$.

For each $(u,w,x) \in U \times W \times X$, let $Z^{t+1}_{u,w,x}$ be an indicator defined as follows:

$$Z^{t+1}_{u,w,x} = \begin{cases} 1 & \text{if } (m^{t+1}_{w\to u}((u,w,x)) = \tilde{m}^{t+1}_{w\to u}) \wedge (m^{t+1}_{x\to u}((u,w,x)) = \tilde{m}^{t+1}_{x\to u}) \\ 0 & otherwise \end{cases}$$

Notice that $Z^{t+1}_{u,w,x}$ is a boolean cube of size $s \times s \times s$, and that the above argument implies that it contains at least $s^3/32$ entries that are 1.

Therefore, by Corollary 1, the cube $Z^{t+1}_{u,w,x}$ contains a 1-monochromatic subcube such that each side of the subcube is of size $\sqrt{\log s/\log 32}$. Denote this subcube by $U' \times W' \times X'$, and the claim follows. □

Proof of Lemma 3: To finish the proof of Lemma 3, we apply Claim 1 on each of the three sides of the t-fooling cube of triangles $U \times W \times X$, and we deduce that there is a constant β and a $(t+1)$-fooling cube of triangles $U' \times W' \times X' \subseteq U \times W \times X$ of size $|U'| = |W'| = |X'| = \beta\sqrt{\log\log\log s}$. □

What we have so far, by Observation 1 and Lemma 3, is that there are sufficiently many triangles in which the nodes receive the same messages from their neighbors. What remains is to actually capture the two different scenarios of a node participating in a triangle or not, while having the node keep receiving the same messages in both. For this we next discuss 6-cycles.

4.2 Triangles and 6-Cycles

Having defined fooling structures of triangles in the previous section, our final step towards proving Theorem 3 is presenting a connection between fooling structures of triangles and 6-cycles. We start with the following definitions.

Definition 6. *Let $(u_1, w_1, x_1), (u_2, w_2, x_2) \in A_i \times A_j \times A_k$ be two disjoint triples, for $i \neq j \neq k \in \{1, 2, 3\}$. The triangle (u_1, w_1, x_1) and the 6-cycle $(u_1, w_1, x_2, u_2, w_2, x_1)$ are called $(3 \leftrightarrow 6, u_1)^t$-*fooling*, if u_1 receives the same messages from w_1 and x_1 during each of the first t rounds, in the triangle (u_1, w_1, x_1) and in the 6-cycle $(u_1, w_1, x_2, u_2, w_2, x_1)$. Formally, for each $1 \leq i \leq t$,*

$$m^i_{w_1 \to u_1}((u_1, w_1, x_1)) = m^i_{w_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1))$$
$$\wedge m^i_{x_1 \to u_1}((u_1, w_1, x_1)) = m^i_{x_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1)).$$

To capture the connection between triangles and 6-cycles, we extend the notion of a t-fooling cube of triangles into the notion of a $(3 \leftrightarrow 6)^t$-fooling cube.

Definition 7. *A cube $U \times W \times X \subseteq A_1 \times A_2 \times A_3$ is called a $(3 \leftrightarrow 6)^t$-*fooling cube*, if:*

1. *The cube $U \times W \times X$ is a t-fooling cube of triangles (see Definition 5).*
2. *For each pair of disjoint triples $(u_1, w_1, x_1), (u_2, w_2, x_2) \subseteq U \times W \times X$ it holds that:*
   (a) *The triangle (u_1, w_1, x_1) and the 6-cycle $(u_1, w_1, x_2, u_2, w_2, x_1)$ are $(3 \leftrightarrow 6, u_1)^t$-fooling (see Definition 6).*
   (b) *The triangle (u_1, w_1, x_1) and the 6-cycle $(u_1, w_1, x_1, u_2, w_2, x_2)$ are $(3 \leftrightarrow 6, w_1)^t$ fooling.*
   (c) *The triangle (u_1, w_1, x_1) and the 6-cycle $(u_1, w_2, x_2, u_2, w_1, x_1)$ are $(3 \leftrightarrow 6, x_1)^t$-fooling.*

Our goal next is to show that there is a $(3 \leftrightarrow 6)^{\Omega(\log^* n)}$-fooling cube $U \times W \times X$ of size $|U| = |W| = |X| \geq 2$. The following observation follows immediately from definition 7.

Observation 2. *$A_1 \times A_2 \times A_3$ is a $(3 \leftrightarrow 6)^0$-fooling cube of size $n/3$.*

In general, in life, it is good to know that there is a $(3 \leftrightarrow 6)^0$-fooling cube somewhere in the wild. However, as we are interested in a lower bound on the number of rounds, we need to show that the amazing $(3 \leftrightarrow 6)^0$-fooling cube contains *sufficiently large* sets of fooling cubes *during many rounds*. For this, we need to prove the following lemma.

Lemma 4. *If there is a $(3 \leftrightarrow 6)^t$-fooling cube $U \times W \times X \subseteq A_1 \times A_2 \times A_3$ of size $|U| = |W| = |X| = s$, then there is a $(3 \leftrightarrow 6)^{t+1}$-fooling cube $U' \times W' \times X' \subseteq U \times W \times X$ of size $|U'| = |W'| = |X'| = \Omega(\sqrt{\log \log \log s})$.*

Proof. By the first property of Definition 7, the cube $U \times W \times X$ is a t-fooling cube of triangles. Therefore, by Lemma 3, there is a $(t+1)$-fooling cube of triangles $U' \times W' \times X' \subseteq U \times W \times X$ of size $|U'| = |W'| = |X'| = \Omega(\sqrt{\log\log\log s})$. We show that $U' \times W' \times X'$ is also a $(3 \leftrightarrow 6)^{t+1}$-fooling cube. That is, we show that each pair of disjoint triples $(u_1, w_1, x_1), (u_2, w_2, x_2) \in U' \times W' \times X'$ satisfies Properties 2a, 2b and 2c of Definition 7.

We start with Property 2a. That is, we first prove that for each pair of disjoint triples $(u_1, w_1, x_1), (u_2, w_2, x_2) \in U' \times W' \times X'$ it holds that the triangle (u_1, w_1, x_1) and the 6-cycle $(u_1, w_1, x_2, u_2, w_2, x_1)$ are $(3 \leftrightarrow 6, u_1)^{t+1}$-fooling. Observe that since $U' \times W' \times X'$ is a $(3 \leftrightarrow 6)^t$-fooling cube, by Property 2a of Definition 7, it holds that for each $1 \leq i \leq t$,

$$m^i_{w_1 \to u_1}((u_1, w_1, x_1)) = m^i_{w_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1))$$
$$\wedge m^i_{x_1 \to u_1}((u_1, w_1, x_1)) = m^i_{x_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1)).$$

Therefore, in order to show that the triangle (u_1, w_1, x_1) and the 6-cycle $(u_1, w_1, x_2, u_2, w_2, x_1)$ are $(3 \leftrightarrow 6, u_1)^{t+1}$-fooling, it remains to show that the above holds also for $t+1$. That is, we need to show that

$$m^{t+1}_{w_1 \to u_1}((u_1, w_1, x_1)) = m^{t+1}_{w_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1))$$
$$\wedge m^{t+1}_{x_1 \to u_1}((u_1, w_1, x_1)) = m^{t+1}_{x_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1)).$$

Observe that since $U' \times W' \times X'$ is a $(t+1)$-fooling cube of triangles, it holds that

$$m^{t+1}_{w_1 \to u_1}((u_1, w_1, x_1)) = m^{t+1}_{w_1 \to u_1}((u_1, w_1, x_2)) \tag{1}$$
$$\wedge m^{t+1}_{x_1 \to u_1}((u_1, w_1, x_1)) = m^{t+1}_{x_1 \to u_1}((u_1, w_2, x_1)). \tag{2}$$

Furthermore, Since $U' \times W' \times X'$ is also a $(3 \leftrightarrow 6)^t$-fooling cube, by Property 2b of Definition 7, for the two triples $(u_1, w_1, x_2), (u_2, w_2, x_1)$, it holds that for each $1 \leq i \leq t$,

$$m^i_{u_1 \to w_1}((u_1, w_1, x_2)) = m^i_{u_1 \to w_1}((u_1, w_1, x_2, u_2, w_2, x_1))$$
$$\wedge m^i_{x_2 \to w_1}((u_1, w_1, x_2)) = m^i_{x_2 \to w_1}((u_1, w_1, x_2, u_2, w_2, x_1)),$$

which means that w_1 has the same view after t rounds both in the case that it participates in a triangle $((u_1, w_1, x_2))$, and in the case that it participates in a 6-cycle $((u_1, w_1, x_2, u_2, w_2, x_1))$. This, in turn, implies that it sends the same message to u_1 during round $t+1$, in these two scenarios, that is:

$$m^{t+1}_{w_1 \to u_1}((u_1, w_1, x_2)) = m^{t+1}_{w_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1)).$$

Combining this with Equation (1), gives that

$$m^{t+1}_{w_1 \to u_1}((u_1, w_1, x_1)) = m^{t+1}_{w_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1)).$$

Similarly, by Property 2c of Definition 7, for the two triples $(u_1, w_2, x_1), (u_2, w_1, x_2)$, it holds that for each $1 \leq i \leq t$,

$$m^i_{u_1 \to x_1}((u_1, w_2, x_1)) = m^i_{u_1 \to x_1}((u_1, w_1, x_2, u_2, w_2, x_1))$$
$$\wedge m^i_{w_2 \to x_1}((u_1, w_2, x_1)) = m^i_{w_2 \to x_1}((u_1, w_1, x_2, u_2, w_2, x_1)),$$

which implies that

$$m^{t+1}_{x_1 \to u_1}((u_1, w_2, x_1)) = m^{t+1}_{x_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1)).$$

Combining this with Equation (2), gives that

$$m^{t+1}_{x_1 \to u_1}((u_1, w_1, x_1)) = m^{t+1}_{x_1 \to u_1}((u_1, w_1, x_2, u_2, w_2, x_1)),$$

which completes the proof that for each pair of disjoint triples $(u_1, w_1, x_1), (u_2, w_2, x_2) \in U' \times W' \times X'$ it holds that the triangle (u_1, w_1, x_1) and the 6-cycle $(u_1, w_1, x_2, u_2, w_2, x_1)$ are $(3 \leftrightarrow 6, u_1)^{t+1}$-fooling, i.e., the cube $U' \times W' \times X'$ satisfies Property 2a of Definition 7. By symmetric arguments, $U' \times W' \times X'$ also satisfies Properties 2b and 2c of Definition 7. □

Proof of Theorem 3: Observe that it is sufficient to prove that there is a $(3 \leftrightarrow 6)^{\Omega(\log^* n)}$-fooling cube $U \times W \times X \subseteq A_1 \times A_2 \times A_3$ of size $|U| = |W| = |X| \geq 2$. By Observation 2, $A_1 \times A_2 \times A_3$ is a $(3 \leftrightarrow 6)^0$-fooling cube of size $n/3$. Furthermore, by Lemma 4, for any $t \geq 0$, if there is a $(3 \leftrightarrow 6)^t$-fooling cube $U^t \times W^t \times X^t$ of size $|U^t| = |W^t| = |X^t| = s$, then there is a $(3 \leftrightarrow 6)^{t+1}$-fooling cube $U^{t+1} \times W^{t+1} \times X^{t+1} \subseteq U^t \times W^t \times X^t$ of size $|U^{t+1}| = |W^{t+1}| = |X^{t+1}| = \Omega(\sqrt{\log \log \log s})$. Therefore, applying Lemma 4 repeatedly $\Omega(\log^* n)$ times implies that there is a $(3 \leftrightarrow 6)^{\Omega(\log^* n)}$-fooling cube $U \times W \times X \subseteq A_1 \times A_2 \times A_3$ of size $|U| = |W| = |X| \geq 2$. □

5 Discussion and Open Questions

Being a first-step type of contribution, this work would not be complete without pointing out many additional open questions for which our fooling views technique is a possible candidate as the road for making progress.

First, we raise the question of whether our specification of the triangle membership problem has an inherently different complexity than that of triangle detection. The latter is the standard way of phrasing decision problems in distributed computing (*everyone* outputs NO for a no instance, *someone* outputs YES for a yes instance). We believe that the bound we give for single-round algorithms in Section 3 should hold also for triangle detection, but this seems to require a deeper technical analysis. Our lower bound for the number of rounds in CONGEST(1) does hold also for triangle detection as explained in Section 4. Note that the sublinear algorithm of [30] solves triangle detection but it is not clear how to make it solve triangle membership. A closely-related problem is that of triangle listing, where all triangles need to be output. The work of [30] also gives the first sublinear algorithm for triangle listing, completing in $O(n^{3/4} \log n)$ rounds in the CONGEST model, as well as an $\Omega(n^{1/3}/\log n)$ lower bound (see also [53]).

Open Question 2. *Do the triangle membership, triangle detection, and triangle listing problems have different complexities in the* CONGEST *model?*

Our results in this paper are for deterministic algorithms, but we do not see a technical obstacle in making them work for randomized algorithms as well.

Open Question 3. *Is the deterministic complexity of triangle membership/detection strictly larger than that of its randomized complexity in the* CONGEST *model?*

Section 3 gives a tight bound on the bandwidth required for an optimal-round algorithm. In Appendix A we address the bandwidth of optimal-round algorithms for k-cycles. This question can be asked about further problems, even those for which we do not know yet what the exact round complexity is, such as various symmetry breaking problems.

Open Question 4. *What is the* bandwidth complexity *of optimal-round distributed algorithms for various problems?*

Whether gaps between the LOCAL and CONGEST models occur for symmetry breaking problems is a central open question, which we hope our technique can shed light upon. Prime examples are MIS and $(\Delta + 1)$-Coloring, but it is not known whether there is indeed a gap. The reason that reductions from 2-party communication problems are provably incapable of proving lower bounds for these problems is that any partial solution that is obtained by a greedy algorithm is extendable into a valid solution for the entire graph, which means that one player can solve the problem for its set of nodes and deliver only the state of nodes on the boundary to the other player, for completing the task. Another way to see why arguing about communication only will not suffice here is to notice that simulating the sequential greedy algorithm requires in fact very little communication in total, despite taking many rounds.

Open Question 5. *For various symmetry breaking problems, is the complexity in the* CONGEST *model strictly higher than its counterpart in the* LOCAL *model?*

Acknowledgements: We are grateful to Michal Dory, Eyal Kushilevitz, and Merav Parter for stimulating discussions. We are also grateful to Ivan Rapaport, Eric Remila, and Nicolas Schabanel, a point that was made by them helped in significantly simplifying Section 3.

References

[1] A. Abboud, A. Backurs, and V. V. Williams. If the current clique algorithms are optimal, so is Valiant's parser. In *FOCS*, pages 98–117, 2015.

[2] A. Abboud and V. Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *FOCS*, pages 434–443, 2014.

[3] A. Abboud, V. Vassilevska Williams, and H. Yu. Matching triangles and basing hardness on an extremely popular conjecture. In *STOC*, pages 41–50, 2015.

[4] F. N. Afrati, D. Fotakis, and J. D. Ullman. Enumerating subgraph instances using map-reduce. In *Data Engineering (ICDE), 2013 IEEE 29th International Conference on*, pages 62–73. IEEE, 2013.

[5] N. Alon, L. Babai, and A. Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms*, 7(4):567–583, 1986.

[6] N. Alon, R. Yuster, and U. Zwick. Color-coding. *Journal of the ACM (JACM)*, 42(4):844–856, 1995.

[7] B. Awerbuch, O. Goldreich, D. Peleg, and R. Vainish. A trade-off between information and communication in broadcast protocols. *J. ACM*, 37(2):238–256, 1990.

[8] N. Bansal and R. Williams. Regularity lemmas and combinatorial algorithms. *Theory of Computing*, 8(1):69–94, 2012.

[9] A. Bar-Noy, J. Naor, and M. Naor. One-bit algorithms. *Distributed Computing*, 4:3–8, 1990.

[10] M. Baruch, P. Fraigniaud, and B. Patt-Shamir. Randomized proof-labeling schemes. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 315–324, 2015.

[11] A. Belovs. Span programs for functions with constant-sized 1-certificates. In *STOC*, pages 77–84. ACM, 2012.

[12] A. Björklund, R. Pagh, V. V. Williams, and U. Zwick. Listing triangles. In *International Colloquium on Automata, Languages, and Programming*, pages 223–234. Springer, 2014.

[13] S. Brandt, O. Fischer, J. Hirvonen, B. Keller, T. Lempiäinen, J. Rybicki, J. Suomela, and J. Uitto. A lower bound for the distributed lovász local lemma. In *STOC*, pages 479–488, 2016.

[14] H. Buhrman, C. Durr, M. Heiligman, P. Hoyer, F. Magniez, M. Santha, and R. De Wolf. Quantum algorithms for element distinctness. In *CCC*, pages 131–137. IEEE, 2001.

[15] K. Censor-Hillel, E. Fischer, G. Schwartzman, and Y. Vasudev. Fast distributed algorithms for testing graph properties. In *DISC*, pages 43–56. Springer, 2016.

[16] K. Censor-Hillel, P. Kaski, J. H. Korhonen, C. Lenzen, A. Paz, and J. Suomela. Algebraic methods in the congested clique. In *PODC*, pages 143–152. ACM, 2015.

[17] K. Censor-Hillel, T. Kavitha, A. Paz, and A. Yehudayoff. Distributed construction of purely additive spanners. In *Proceedings of the 30th International Symposium on Distributed Computing (DISC), Paris, France, September 27-29, 2016*, pages 129–142, 2016.

[18] K. Censor-Hillel, S. Khoury, and A. Paz. Quadratic and near-quadratic lower bounds for the CONGEST model. *DISC 2017. Also in CoRR*, abs/1705.05646, 2017.

[19] T. M. Chan. Speeding up the four russians algorithm by about one more logarithmic factor. In *SODA*, pages 212–217, 2015.

[20] A. Czumaj and A. Lingas. Finding a heaviest vertex-weighted triangle is not harder than matrix multiplication. *SIAM Journal on Computing*, 39(2):431–444, 2009.

[21] D. Dolev, C. Lenzen, and S. Peled. " tri, tri again": Finding triangles and small subgraphs in a distributed setting. In *DISC*, pages 195–209. Springer, 2012.

[22] A. Drucker, F. Kuhn, and R. Oshman. On the power of the congested clique model. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 367–376, 2014.

[23] M. Elkin. A simple deterministic distributed MST algorithm, with near-optimal time and message complexities. In *PODC*, pages 157–163. ACM, 2017.

[24] P. Erdös. On extremal problems of graphs and generalized graphs. *Israel Journal of Mathematics*, 2(3):183–190, Sep 1964.

[25] P. Fraigniaud, I. Rapaport, V. Salo, and I. Todinca. Distributed testing of excluded subgraphs. In *DISC*, pages 342–356. Springer, 2016.

[26] S. Frischknecht, S. Holzer, and R. Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1150–1162, 2012.

[27] J. Håstad and A. Wigderson. The randomized communication complexity of set disjointness. *Theory of Computing*, 3(1):211–219, 2007.

[28] M. Henzinger, S. Krinninger, D. Nanongkai, and T. Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *STOC*, pages 21–30. ACM, 2015.

[29] A. Itai and M. Rodeh. Finding a minimum circuit in a graph. *SIAM Journal on Computing*, 7(4):413–423, 1978.

[30] T. Izumi and F. L. Gall. Triangle finding and listing in CONGEST networks. In *Proceedings of the ACM Symposium on Principles of Distributed Computing, (PODC)*, pages 381–389, 2017.

[31] S. Jukna. *Extremal Combinatorics - With Applications in Computer Science*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2001.

[32] J. Kari, M. Matamala, I. Rapaport, and V. Salo. Solving the induced subgraph problem in the randomized multiparty simultaneous messages model. In *SIROCCO*, pages 370–384. Springer, 2015.

[33] V. King, S. Kutten, and M. Thorup. Construction and impromptu repair of an MST in a distributed network with o(m) communication. In *PODC*, pages 71–80. ACM, 2015.

[34] M. N. Kolountzakis, G. L. Miller, R. Peng, and C. E. Tsourakakis. Efficient triangle counting in large graphs via degree-based vertex partitioning. *Internet Mathematics*, 8(1-2):161–185, 2012.

[35] K. Kothapalli, C. Scheideler, M. Onus, and C. Schindelhauer. Distributed coloring in $\tilde{O}(\sqrt{\log n})$ bit rounds. In *Proceedings of the 20th International Parallel and Distributed Processing Symposium (IPDPS)*, 2006.

[36] F. Kuhn, T. Moscibroda, and R. Wattenhofer. Local computation: Lower and upper bounds. *J. ACM*, 63(2):17, 2016.

[37] E. Kushilevitz and N. Nisan. Communication complexity. In *Cambridge University Press*, 1997.

[38] F. Le Gall. Improved quantum algorithm for triangle finding via combinatorial arguments. In *FOCS*, pages 216–225. IEEE, 2014.

[39] F. Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th international symposium on symbolic and algebraic computation*, pages 296–303. ACM, 2014.

[40] F. Le Gall and S. Nakajima. Multiparty quantum communication complexity of triangle finding. In *TQC, to appear*, 2017.

[41] F. Le Gall and S. Nakajima. Quantum algorithm for triangle finding in sparse graphs. *Algorithmica*, 79(3):941–959, 2017.

[42] T. Lee, F. Magniez, and M. Santha. Improved quantum query algorithms for triangle finding and associativity testing. In *SODA*, pages 1486–1502. SIAM, 2013.

[43] C. Lenzen and B. Patt-Shamir. Improved distributed steiner forest construction. In *PODC*, pages 262–271, 2014.

[44] N. Linial. Distributive graph algorithms global solutions from local data. In *Proceedings of the 28th Annual Symposium on Foundations of Computer Science*, FOCS, pages 331–335, 1987.

[45] N. Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.

[46] M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1053, 1986.

[47] F. Magniez, M. Santha, and M. Szegedy. Quantum algorithms for the triangle problem. *SIAM Journal on Computing*, 37(2):413–424, 2007.

[48] Y. Métivier, J. M. Robson, N. Saheb-Djahromi, and A. Zemmari. An optimal bit complexity randomized distributed MIS algorithm. *Distributed Computing*, 23(5-6):331–340, 2011.

[49] D. Nanongkai, A. D. Sarma, and G. Pandurangan. A tight unconditional lower bound on distributed randomwalk computation. In *PODC*, pages 257–266. ACM, 2011.

[50] M. Naor. A lower bound on probabilistic algorithms for distributive ring coloring. *SIAM J. Discrete Math.*, 4(3):409–412, 1991.

[51] S. Pai, G. Pandurangan, S. V. Pemmaraju, T. Riaz, and P. Robinson. Symmetry breaking in the congest model: Time- and message-efficient algorithms for ruling sets. *In DISC 2017. Also in CoRR*, abs/1705.07861, 2017.

[52] S. Pai, G. Pandurangan, S. V. Pemmaraju, T. Riaz, and P. Robinson. Symmetry breaking in the congest model: Time-and message-efficient algorithms for ruling sets. In *PODC, to appear*, 2017.

[53] G. Pandurangan, P. Robinson, and M. Scquizzato. Tight bounds for distributed graph computations. *CoRR*, abs/1602.08481, 2016.

[54] G. Pandurangan, P. Robinson, and M. Scquizzato. A time- and message-optimal distributed algorithm for minimum spanning trees. In *STOC*, pages 743–756. ACM, 2017.

[55] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. Society for Industrial and Applied Mathematics, 2000.

[56] D. Peleg and V. Rubinovich. A near-tight lower bound on the time complexity of distributed minimum-weight spanning tree construction. *SIAM J. Comput.*, 30(5):1427–1442, 2000.

[57] A. A. Razborov. On the distributional complexity of disjointness. *Theor. Comput. Sci.*, 106(2):385–390, 1992.

[58] A. D. Sarma, F. N. Afrati, S. Salihoglu, and J. D. Ullman. Upper and lower bounds on the cost of a map-reduce computation. In *VLDB*, volume 6, pages 277–288. VLDB Endowment, 2013.

[59] A. D. Sarma, S. Holzer, L. Kor, A. Korman, D. Nanongkai, G. Pandurangan, D. Peleg, and R. Wattenhofer. Distributed verification and hardness of distributed approximation. *SIAM J. Comput.*, 41(5):1235–1265, 2012.

[60] T. Schank and D. Wagner. Finding, counting and listing all triangles in large graphs, an experimental study. In *WEA*, pages 606–609. Springer, 2005.

[61] M. Szegedy. On the quantum query complexity of detecting triangles in graphs. *arXiv preprint quant-ph/0310107*, 2003.

[62] V. V. Williams. Multiplying matrices faster than coppersmith-winograd. In *STOC*, pages 887–898. ACM, 2012.

[63] V. V. Williams. Hardness of easy problems: basing hardness on popular conjectures such as the strong exponential time hypothesis (invited talk). In *LIPIcs-Leibniz International Proceedings in Informatics*, volume 43, 2015.

[64] V. V. Williams and R. Williams. Subcubic equivalences between path, matrix and triangle problems. In *FOCS*, pages 645–654, 2010.

[65] H. Yu. An improved combinatorial algorithm for boolean matrix multiplication. In *ICALP*, pages 1094–1105. Springer, 2015.

A k-Cycle Membership for $k \geq 4$

An immediate question is whether we can get lower bounds on the bandwidth for additional round-optimal algorithms. We show here how to generalize our lower bound technique to apply for detecting membership in larger cycles. However, curiously, as we show in Section A.1 for the sake of comparison, for cycles that are larger than 3, the standard approach of reductions from 2-party communication complexity problems allows for stronger lower bounds.

An optimal-round algorithm for solving k-cycle membership completes within exactly $\lfloor (k-1)/2 \rfloor$ rounds, by a simple (standard) indistinguishability argument (this is even with unlimited bandwidth). This is because after t rounds of communication a node may have information only about nodes at distance at most $t+1$ from it. Therefore, after $\lfloor (k-1)/2 \rfloor - 1$ rounds of communication, a node cannot distinguish whether it participates in a k-cycle or not.

To see how we generalize our technique, observe first that the lower bound given in Theorem 2 holds even when a specific node u is given as input to all the nodes, and only u needs to solve the triangle membership problem. This is helpful in extending our lower bound to the case of k-cycles when $k \geq 4$, and hence we define this problem formally.

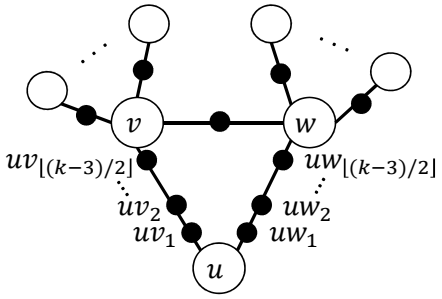

Figure 3: Constructing G' out of G.

Definition 8. *(Fixed-Node Triangle Membership). In the Fixed-Node Triangle Membership problem, all nodes are given the identity of a specific node u, and node u needs to detect whether it is a part of a triangle.*

The proof of Theorem 2 actually proves the following theorem.

Theorem 5. *The fixed-node triangle membership problem cannot be solved by a single-round algorithm in the* CONGEST(B) *model unless* $B \geq \log((\frac{n-2}{2(\Delta-1)})^{\Delta-1})$.

We formally extend the membership problem to larger cycles, as follows.

Definition 9. *(k-Cycle Membership). In the k-Cycle Membership problem, each node needs to detect whether it is a part of a k-cycle.*

Now, Theorem 5 can be used to prove the following.

Theorem 6. *Let $k \in O(n^{1-\epsilon})$, for some constant $0 < \epsilon < 1$. The k-cycle membership problem cannot be solved by an optimal $\lfloor(k-1)/2\rfloor$-round algorithm in the* CONGEST(B) *model unless $B \geq c\log((\frac{n-2}{2(\Delta-1)})^{\Delta-1})$, for some constant $c \geq 0$.*

We prove Theorem 6 by showing a reduction from the fixed-node triangle membership problem, given in Definition 8. That is, we show that an algorithm for solving the k-cycle membership problem can be used to solve the fixed-node triangle membership problem.

Proof of Theorem 6: First, we show how to construct an appropriate instance for the k-cycle membership problem, given an instance for the fixed-node triangle membership problem. We start with describing the construction for an odd value of k, and then show how to tweak it to handle even values of k as well.

Let $(G = (V, E), u)$ be an instance of the fixed-node triangle membership problem where $|V| = \widetilde{n}$, and let $E(u)$ be the set of edges incident to the node u. For an odd k, we define an instance of the k-cycle membership problem, $G' = (V', E')$, where $|V'| = n$, as follows. We replace each edge $(u, v) \in E(u)$ with a path P_{uv} of length $(k-3)/2+1$, going through $(k-3)/2$ new nodes, denoted $uv_1, \ldots, uv_{(k-3)/2}$, containing the edges $(u, uv_1), (uv_1, uv_2), \ldots, (uv_{(k-3)/2-1}, uv_{(k-3)/2}), (uv_{(k-3)/2}, v)$. For example, for $k = 7$, we replace each edge (u, v) by a path of three edges $(u, uv_1), (uv_1, uv_2), (uv_2, v)$, going through two new intermediate nodes uv_1 and uv_2 (see Figure 3). Formally, the new graph

$G' = (V', E')$ is defined as:

$$V' = V \cup \{uv_i \mid (u,v) \in E(u), 1 \leq i \leq (k-3)/2\}$$
$$E' = (E \setminus E(u)) \cup \{(u, uv_1)), (uv_1, uv_2), \ldots, (uv_{(k-3)/2-1}, uv_{(k-3)/2}), (uv_{(k-3)/2}, v) \mid (u,v) \in E(u)\}$$

The following observation follows directly from the construction above.

Observation 3. *For an odd k, the node u participates in a k-cycle in G' if and only if it participates in a triangle in G.*

Since an algorithm may use the IDs of the nodes, we need to also assign unique IDs to nodes in G'. We can do this in any consistent arbitrary manner, say, by assigning $ID_{G'}(x) = ID_G(x)$ if $x \in V$, and $ID(uv_i) = ID(v) \circ bin(i)$ for nodes in $V' \setminus V$, where $bin(i)$ is the binary representation of i. Notice that $\Delta(G') = \Delta(G)$.

Next, assume towards a contradiction, that there is an algorithm A' that solves the k-cycle membership problem in an optimal number of $(k-1)/2$ rounds in the CONGEST(B) model with $B < \log((\frac{\widetilde{n}-2}{2e(\Delta-1)})^{\Delta-1})$. We show an algorithm A that solves the fixed-node triangle membership problem on $(G = (V, E), u)$ in a single round of the CONGEST(B) model with $B < \log((\frac{\widetilde{n}-2}{2e(\Delta-1)})^{\Delta-1})$. As this contradicts Theorem 5, and $n = k\widetilde{n}$, this completes the proof for all values of k in $O(n^{1-\epsilon})$ for any constant $0 \leq \epsilon \leq 1$.

We construct A such that nodes in G simulate the nodes in G' running algorithm A', as follows. In A, each node in V, sends to its neighbors in G the message it sends in the first round of A' on G'. If $v \in V$ is a neighbor of u in G, then it sends to u the message it sends to $uv_{(k-3)/2}$ in A'. Then, for each of its neighbors v, the node u has all the messages that the node $uv_{(k-3)/2}$ receives in the first round of A'. Now, by backwards induction, for each i, $1 \leq i \leq (k-3)/2$, by local simulation at the node u, it knows the view of the node uv_i at the end of the first $(k-3)/2 - i + 1$ rounds of A'. This implies that u knows the message sent to it by uv_1 in round $(k-3)/2 + 1 = (k-1)/2$ of A'. Since in $(k-1)/2$ rounds of A' the node u knows whether it is in a k-cycle in G', by Observation 3, it thus knows in a single round whether it is in a triangle in G.

To handle even value of k, we construct G' in a similar manner of replacing each edge $(u,v) \in E(u)$ by a path of $(k-4)/2$ nodes. In addition, each edge $(v,w) \in E \setminus E(u)$, is replaced by a path of length two which consists of an additional node vw. Similarly to Observation 3, the node u participates in a k-cycle in G' if and only if it participates in a triangle in G. As in the case for odd k, given an algorithm A' for k-cycle membership in G', the simulation of $(k-4)/2 + 1 = (k-2)/2 = \lfloor (k-1)/2 \rfloor$ rounds of it in G requires only a single round of communication. Therefore, the reduction carries over for even values of k as well. Observe that here $n = k\widetilde{n} + \widetilde{n}^2$, therefore, as in the case of odd k, we achieve the same asymptotic lower bound as in the triangle membership problem, for any value of k in $O(n^{1-\epsilon})$. □

A.1 k-Cycle membership for $k \geq 4$ using communication complexity

Here we show a lower bound on the bandwidth needed for any optimal-round algorithm for solving k-cycle membership for $k \geq 4$ by using the standard framework of reduction from a 2-party communication complexity problem. For simplicity, we will show this for even values of k, but a similar construction works for odd values as well. For $k > 4$, the bound is larger compared with our proof of Section A.

Theorem 7. *The k-cycle membership problem cannot be solved by a deterministic optimal-round algorithm in the* CONGEST(B) *model unless $B \geq \Omega(\Delta^{\frac{k-4}{2}+1}\log(n))$, and it cannot be solved by a randomized optimal-round algorithm, which succeeds with high probability*[6]*, in the* CONGEST(B) *model unless $B \geq \Omega(\Delta^{\frac{k-4}{2}+1})$, for any integers n, Δ, k such that*

$$(\Delta-1)^{\frac{k-4}{2}+1} = O(n^{1-\epsilon}),$$

for some constant $0 < \epsilon < 1$.

In order to prove Theorem 7, we show a reduction from the 2-party communication complexity problem $DISJ_{K,S}$.

A 2-party communication complexity problem [37] consists of a function $f : \{0,1\}^K \times \{0,1\}^K \to \{\text{TRUE}, \text{FALSE}\}$, and two strings, $x, y \in \{0,1\}^K$, that are given as inputs for two players, Alice and Bob, respectively. The players exchange bits of communication in order to compute $f(x,y)$, according to a protocol π. The *communication complexity* $CC(\pi)$ of a protocol π for computing f is the maximal number of bits, taken over all input pairs (x,y), exchanged between Alice and Bob. The *communication complexity* $CC(f)$ of f is the minimum, taken over all protocols π that compute f, of $CC(\pi)$.

In the *S-Disjointness* problem ($DISJ_{K,S}$), each of the players Alice and Bob receives a K-bits input string containing exactly S ones, and the function f is $DISJ_{K,S}(x,y)$, whose output is FALSE if there is an index $i \in \{0,...,K-1\}$ such that $x_i = y_i = 1$, and TRUE otherwise. Observe that for $S > \frac{K}{2}$, the function is constant. For $S \leq \frac{K}{2}$ The deterministic communication complexity of $DISJ_{K,S}$ is known to be $\Omega(\log\binom{K}{S})$ [31,37][7], while its randomized communication complexity is known to be $\Omega(S)$ [27,57]. For the reduction, we adapt the formalization of *Family of Lower Bound Graphs* given in [18] to our setting.

Definition 10 (**Definition 1 (simplified) of [18]: Family of Lower Bound Graphs**). *Fix an integer K, a function $f : \{0,1\}^K \times \{0,1\}^K \to \{\text{TRUE}, \text{FALSE}\}$. The family of graphs $\{G_{x,y} = (V, E_{x,y}) \mid x, y \in \{0,1\}^K\}$, is said to be a family of* lower bound graphs w.r.t. f and k-cycle membership *if the following properties hold:*

(1) The set of nodes V is the same for all graphs, and we denote by $V = \{u\} \dot\cup V_A \dot\cup V_B \dot\cup \widetilde{V}$ a fixed partition of it;

(2) Only the existence of edges in $V_A \times \widetilde{V}$ may depend on x;

(3) Only the existence of edges in $V_B \times \widetilde{V}$ may depend on y;

(4) The node u participates in a k-cycle in $G_{x,y}$ iff $f(x,y) =$ FALSE.

Observe that given a family of lower bound graphs $\{G_{x,y} = (V, E_{x,y}) \mid x, y \in \{0,1\}^K\}$ w.r.t. to $DISJ_{K,S}$ and k-cycle membership, if Alice and Bob can simulate an algorithm for k-cycle membership on u, then by checking the output of u at the end of the algorithm they can solve $DISJ_{K,S}(x,y)$.

[6] We say that an event occurs with high probability if it occurs with probability $1 - \frac{1}{n^c}$, for some constant $c \geq 1$.

[7] It can be proved by the rank method for proving lower bounds for deterministic protocols in communication complexity. To read more about the rank method, see for example [37] section 1.4. For the proof of the lower bound on the rank of S-Disjointness, see [31], page 175.

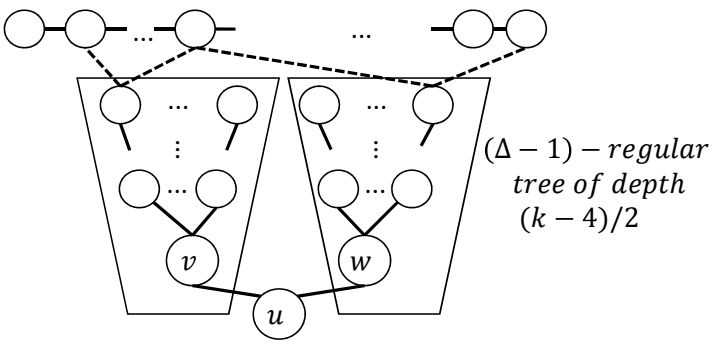

Figure 4: The reduction from Communication Complexity to k-cycle membership.

The proof of Theorem 7 is organized as follows. First, we construct a family of lower bound graphs, and next, we show that given an algorithm ALG for k-cycle membership with messages of size B, Alice and Bob can simulate ALG on $G_{x,y}$ by exchanging only $O(B)$ bits.

We now construct the following family of lower bound graphs, by describing a fixed graph construction $G = (V, E)$, which we then generalize to a family of graphs $\{G_{x,y} = (V, E_{x,y}) \mid x, y \in \{0,1\}^K\}$, which we show to be a family lower bound graphs w.r.t. to $DISJ_{K,S}$ and k-cycle membership.

The fixed graph construction: The fixed graph construction (Figure 4) consists of a tree T and a path P of size $n - |T|$. The tree T is a tree in which the root node of T, denoted by u, is connected to two nodes w and v, such that each of w and v is a root of a $(\Delta-1)$-regular tree of depth $\frac{k-4}{2}$. Denote by $leaves(w)$ and $leaves(v)$ the set of leaves of the tree rooted at w and the set of leaves of the tree rooted at v, respectively. We define V_A and V_B to be $leaves(w)$ and $leaves(v)$ respectively.

Adding edges corresponding to the inputs: Each of the players Alice and Bob receives as the input a set of nodes in P of size $(\Delta-1)^{\frac{k-4}{2}+1}$. Alice connects the nodes in her input to the $(\Delta-1)^{\frac{k-4}{2}}$ leaves of the tree rooted at v, such that each leaf is connected to $\Delta-1$ nodes in P. Similarly, Bob connects the nodes in his input to the $(\Delta-1)^{\frac{k-4}{2}}$ leaves of the tree rooted at w, such that each leaf in connected to $\Delta-1$ nodes in P. The following observation follows directly from the construction.

Observation 4. *The node u participates in a cycle of length k in $G_{x,y}$ if and only if the two sets of Alice and Bob are not disjoint.*

Therefore, given an algorithm ALG for k-cycle membership, if Alice and Bob can simulate ALG on u then they can solve S-Disjointness, where $S = (\Delta-1)^{\frac{k-4}{2}+1}$ and the size of the input stings is

$$K = n - |T| \geq n - \left(1 + 2\sum_{i=0}^{\frac{k-4}{2}} (\Delta-1)^i\right) \geq n - \left(1 + 2\frac{(\Delta-1)^{\frac{k-4}{2}+1} - 1}{\Delta-2}\right)$$

Observe that for

$$(\Delta-1)^{\frac{k-4}{2}+1} = O(n^{1-\epsilon})$$

for some constant $0 < \epsilon < 1$, it holds that $K = \Theta(n)$, and $\log \binom{K}{S} = \Omega(\Delta^{\frac{k-4}{2}+1} \log(n))$. It remains to show that given an algorithm for k-cycle membership with messages of size B, Alice and Bob can simulate ALG on u by exchanging only $O(B)$ bits.

Proof of Theorem 7: Let ALG be a $(k/2-1)$-round algorithm for solving the k-cycle membership problem. Let $\{m^1_{v\to u}, ..., m^{k/2-1}_{v\to u}\}$ and $\{m^1_{w\to u}, ..., m^{k/2-1}_{w\to u}\}$ be the two sets of messages sent from v and w to u during the $k/2-1$ rounds, where, e.g., $m^i_{v\to u}$ is the message sent from v to u in round i.

The crucial observation is that Alice and Bob can simulate the nodes v, w, and u during the first $k/2-2$ rounds without any communication, because the $k/2-2$ neighborhoods of these nodes are fixed. Therefore, in order for the players to compute the output of u after $k/2-1$ rounds, it suffices for Alice to send to Bob the message $m^{k/2-1}_{v\to u}$, and for Bob to send to Alice the message $m^{k/2-1}_{w\to u}$.

By Observation 4, this implies that $2B$ bits suffice for the players to correctly compute $DISJ_{K,S}$. Therefore, by the lower bounds on $DISJ_{K,S}$, any deterministic optimal-round algorithm for k-cycle membership requires messages of size $\Omega(\log \binom{K}{S}) = \Omega(\Delta^{\frac{k-4}{2}+1} \log(n))$, and any randomized optimal-round algorithm which succeeds with high probability requires $\Omega(S) = \Omega(\Delta^{\frac{k-4}{2}+1})$, for any integers n, Δ, k such that

$$(\Delta-1)^{\frac{k-4}{2}+1} = O(n^{1-\epsilon})$$

for some constant $0 < \epsilon < 1$.

□

We mention that one can use the set-disjointness function instead of $DISJ_{K,S}$ to obtain a lower bound of $(\Delta-1)^{\frac{k-4}{2}}$, by simply connecting each leaf to a single node on the path, based on the inputs. However, this gives a rather strong bound on Δ with respect to n because then the number of leaves must be also linear in n. Such a bound for Δ does not occur in our given construction.