AMIR ABBOUD, Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Israel*

KEREN CENSOR-HILLEL, Department of Computer Science, Technion, Israel

SERI KHOURY, University of California, Berkeley, USA

AMI PAZ, Faculty of Computer Science, University of Vienna, Austria

This paper proves strong lower bounds for distributed computing in the CONGEST model, by presenting the *bit-gadget*: a new technique for constructing graphs with small cuts.

The contribution of bit-gadgets is twofold. First, developing careful sparse graph constructions with small cuts extends known techniques to show a near-linear lower bound for computing the diameter, a result previously known only for dense graphs. Moreover, the sparseness of the construction plays a crucial role in applying it to approximations of various distance computation problems, drastically improving over what can be obtained when using dense graphs.

Second, small cuts are essential for proving super-linear lower bounds, none of which were known prior to this work. In fact, they allow us to show near-quadratic lower bounds for several problems, such as exact minimum vertex cover or maximum independent set, as well as for coloring a graph with its chromatic number. Such strong lower bounds are not limited to NP-hard problems, as given by two simple graph problems in P which are shown to require a quadratic and near-quadratic number of rounds. All of the above are optimal up to logarithmic factors. In addition, in this context, the complexity of the all-pairs-shortest-paths problem is discussed.

Finally, it is shown that graph constructions for CONGEST lower bounds translate to lower bounds for the semi-streaming model, despite being very different in its nature.

CCS Concepts: • Theory of computation \rightarrow Distributed computing models.

Additional Key Words and Phrases: lower bounds, distributed algorithms

ACM Reference Format:

Amir Abboud, Keren Censor-Hillel, Seri Khoury, and Ami Paz. 2021. Smaller Cuts, Higher Lower Bounds. *ACM Trans. Algor.* 1, 1, Article 1 (January 2021), 41 pages. https://doi.org/10.1145/3469834

This paper is based on two conference papers, from DISC'16 [1] and DISC'17 [32] and contains additional, new results.

* This work was partly done when Amir Abboud was at IBM Almaden Research Center and at Stanford University.

Authors' addresses: Amir Abboud, amir.abboud@weizmann.ac.il, Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Israel ; Keren Censor-Hillel, ckeren@cs.technion.ac.il, Department of Computer Science, Technion, Israel; Seri Khoury, serikhourry@gmail.com, University of California, Berkeley, USA; Ami Paz, ami.paz@univie. ac.at, Faculty of Computer Science, University of Vienna, Austria.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

1549-6325/2021/1-ART1 \$15.00

https://doi.org/10.1145/3469834

Contents

Abstract		1
Contents		2
1	Introduction	2
1.1	Contributions	4
1.2	Additional Related Work	7
2	Preliminaries	9
2.1	Computational Models	9
2.2	Graph Parameters	9
2.3	Communication Complexity	9
2.4	Lower Bound Graphs	10
3	The Bit-Gadget Construction	11
4	Near-Linear Lower Bounds for Sparse Graphs	12
4.1	Exact Diameter	12
4.2	$(3/2 - \varepsilon)$ -Approximation of the Diameter	14
4.3	Radius	16
5	Near-Quadratic Lower Bounds for General Graphs	17
5.1	Minimum Vertex Cover	17
5.2	Graph Coloring	20
6	Quadratic and Near-Quadratic Lower Bounds for Problems in P	25
6.1	Weighted Cycle Detection	25
6.2	Identical Subgraphs Detection	26
7	Weighted APSP	29
7.1	A Linear Lower Bound for Weighted APSP	30
7.2	The Alice-Bob Framework Cannot Give a Super-Linear Lower Bound for Weighted	
	APSP	30
8	Streaming Lower Bounds	33
9	Discussion	36
Ackr	Acknowledgments	
Refe	References	

1 INTRODUCTION

This paper studies inherent limitations of distributed graph algorithms with bounded bandwidth, and shows new and strong lower bounds for some classical graph problems. A fundamental computational model for distributed networks is the CONGEST model [92], where the network graph represents n nodes that communicate in synchronous rounds in which $O(\log n)$ -bit messages are exchanged among neighbors.

Many lower bounds for the CONGEST model rely on reductions from two-party communication problems (see, e.g., [11, 31, 43, 44, 53, 66, 88, 89, 94, 99]). In this setting, two players, Alice and Bob, are given inputs of *K* bits and need to compute a single output bit according to some predefined function of their inputs.

The standard framework for reducing a two-party communication problem of computing a function f to deciding a graph predicate P in the CONGEST model is as follows. Given an instance (x, y) of the two-party problem f, a graph is constructed such that the value of P on it can be used to determine the value of f on (x, y). Some of the graph edges are fixed, while the existence

of some other edges depends on the inputs of Alice and Bob. Then, given an algorithm *ALG* for solving *P* in the CONGEST model, the vertices of the graph are split into two sets, V_A and V_B , and Alice simulates *ALG* over V_A while Bob simulates *ALG* over V_B . The only communication required between Alice and Bob in order to carry out this simulation is the content of messages sent in each direction over the edges of the cut $C = E(V_A, V_B)$. Using this technique for a two-party problem *f* on *K* bits with communication complexity CC(f, K) and a graph with a cut *C*, it can be proven that the complexity of *ALG* is in $\Omega(CC(f, K)/|C| \log n)$.¹

Thus, the lower bound achieved using the reduction depends on two parameters of the graph construction: (i) the size of the input, K, and (ii) the size of the cut, |C|. All previous constructions are both *dense* and have *large cuts*, which causes them to suffer from two limitations.

The first limitation is that lower bounds for global approximation tasks, such as approximating the diameter of the graph, which are typically obtained through stretching edges in the construction into paths by adding new nodes, must pay a significant decrease in the size of the input compared to the number of nodes because of their density. Together with their large cuts, this causes such lower bounds to stay well below linear. For example, the graph construction for the lower bound for computing the diameter [53] has $K = \Theta(n^2)$ and $|C| = \Theta(n)$, which gives an almost linear lower bound of $\Omega(n/\log n)$ using the set-disjointness problem whose communication complexity is known to be $\Theta(K)$ [76]. However, because the construction is dense, although the resulting graph construction for computing a $(3/2 - \epsilon)$ -approximation of the diameter [53] has a smaller cut of $|C| = \Theta(\sqrt{n})$, this comes at the price of supporting a smaller input size, of $K = \Theta(n)$, which gives a lower bound that is roughly a square-root of n.

The second limitation is that large, say, linear cuts, can inherently give only linear lower bounds at best. However, tasks such as computing an exact minimum vertex cover seem to be much harder for the CONGEST model, despite the inability of previous constructions to prove this.

In this paper, we present the bit-gadget technique for constructing graphs with small cuts, using which we prove strong lower bounds for the CONGEST model. Bit-gadgets allow us to have logarithmic-size cut, and are inspired by constructions that are used for proving *conditional* lower bounds in the field of fine-grained complexity for the sequential setting [3, 4, 30, 35, 98]. This novel connection between the two fields relies on the following intuition, which was made more formal in a follow-up paper [28]: many reductions in fine-grained complexity reduce from a problem (the Orthogonal Vectors problem) that happens to be of high communication complexity and therefore also serve to prove that the communication complexity of the end problems is high, which can lead to lower bounds in the CONGEST model.

Our constructions allow bringing lower bounds for approximate diameter and radius up to a near-optimal near-linear complexity. Furthermore, they allow us to obtain the first near-quadratic lower bounds for natural graph problems, such as computing a minimum vertex cover or a coloring with a minimal number of colors. These are near-optimal since all of these problems admit simple O(m) solutions in the CONGEST model. Notably, these are the first super-linear lower bounds for this model.

In addition, this paper discusses the complexity of the weighted all-pairs-shortest-paths problem. This is one of the most-studied problems in the CONGEST model, yet its complexity remains elusive. This problem was known to have at least almost-linear complexity; we improve this by presenting a linear lower bound, and also prove that the Alice-Bob technique discussed cannot achieve a super-linear lower bound for the problem.

¹In this paper, and in many others, the nodes are partitioned into disjoint sets, and this partition remains fixed over time. We remark that in earlier work, the partition of the graph nodes between Alice and Bob is not fixed, and their nodes are not disjoint: there are many nodes they both simulate, and in each round each player simulates less nodes. This technique requires a more involved analysis, and we do not discuss it further.

Finally, we show that graph constructions for lower bounds for the CONGEST model can be used directly to obtain lower bounds for the streaming model of computation in a black-box manner, and so we prove the lower bounds for problems such as computing a minimum vertex cover or a coloring with minimal number of colors.

1.1 Contributions

1.1.1 Distance Computation. Frischknecht et al. [53] showed that the diameter is surprisingly hard to compute: $\tilde{\Omega}(n)$ rounds are needed even in networks with constant diameter². This lower bound is nearly tight, due to an O(n) upper bound [66, 79, 93]. Naturally, approximate solutions are a desired relaxation, and were indeed addressed in several cornerstone studies [53, 64, 66, 79, 93], bringing us even closer to a satisfactory understanding of the time complexity of computing the diameter in the CONGEST model. Here we answer several central questions that remained elusive.

Sparse Graphs. The graphs constructed by Frischknecht et al. [53] have $\Theta(n^2)$ edges and constant diameter, and require any distributed algorithm for computing their diameter to spend $\tilde{\Omega}(n)$ rounds. Almost all large networks of practical interest are very sparse [81], e.g., the Internet in 2012 had roughly 4 billion nodes and 128 billion edges [85]. The only known lower bound for computing the diameter of a sparse network is obtained by a simple modification to the construction of [53] which yields a much weaker bound of $\tilde{\Omega}(\sqrt{n})$. Our first result is to rule out the possibility that the $\tilde{\Omega}(n)$ bound can be beaten significantly in sparse networks.

THEOREM 1. Any algorithm for computing the exact diameter requires $\Omega\left(\frac{n}{\log^2 n}\right)$ rounds, even in networks of $\Theta(n \log n)$ edges.

We remark that, as in [53], our lower bound holds even for networks with constant diameter and even against randomized algorithms. Due to simple transformations, e.g., adding a dummy node on each edge, our lower bound for computing the diameter also holds for the more strict definition of sparse graphs as having O(n) edges, up to a loss of a log factor in the round complexity. Our lower bound against approximating the diameter, which we discuss next, holds for networks of O(n)edges as well. Subsequent works have shown that the $\tilde{\Omega}(n)$ lower bound holds in planar graphs with edge weights [2] but not without them [82].

Approximation Algorithms. An important question is whether one can bypass this near-linear barrier by settling for an approximation to the diameter. An α -approximation algorithm to the diameter returns a value \hat{D} such that $D \leq \hat{D} \leq \alpha \cdot D$, where D is the true diameter of the network. From [53] we know that $\widehat{\Omega}(\sqrt{n+D})$ rounds are needed, even for computing a $(3/2-\varepsilon)$ -approximation to the diameter, for any constant $\varepsilon > 0$, while from [64] we know that a 3/2-approximation can be computed in $O(\sqrt{n \log n} + D)$ rounds. This raises the question of whether there is a sharp threshold at a 3/2-approximation factor, or whether a $(3/2 - \varepsilon)$ -approximation can also be obtained in a sub-linear number of rounds.

Progress towards answering this question was made by Holzer and Wattenhofer [66] who showed that any algorithm that needs to decide whether the diameter is 2 or 3 has to spend $\tilde{\Omega}(n)$ rounds. However, as the authors point out, their lower bound is not robust and does not rule out the possibility of a $(3/2 - \varepsilon)$ -approximation when the diameter is larger than 2, or an algorithm that is allowed an additive +1 error in addition to a multiplicative $(3/2 - \varepsilon)$ error.

As mentioned earlier, perhaps the main difficulty in extending the lower bound constructions of Frischknecht et al. [53] and Holzer and Wattenhofer [66] in order to resolve these gaps was that

²The notations $\widetilde{\Omega}$ and \widetilde{O} hide factors that are polylogarithmic in *n*.

1:5

their original graphs are dense. A natural way to go from a lower bound construction for exact algorithms to a lower bound for approximations is to subdivide each edge into a path; however, in dense graphs this dramatically blows up the number of nodes, resulting in much weaker bounds. The sparseness of our new construction allows us to tighten the bounds and negatively resolve the above question: we show a $\tilde{\Omega}(n)$ lower bound for computing a $(3/2 - \varepsilon)$ -approximation to the diameter, even if a constant additive approximation factor is also allowed.

THEOREM 2. For any constant $0 < \varepsilon < 1/2$, any algorithm for computing a $(3/2 - \varepsilon)$ -approximation of the diameter requires $\Omega\left(\frac{n}{\log^3 n}\right)$ rounds, even in networks of $\Theta(n)$ edges.

Radius. In many scenarios we want one special node to be able to efficiently send information to all other nodes. In this case, we would like this node to be the one that is closest to every other node, i.e., the *center* of the graph. The *radius* of the graph is the largest distance from the center, and it captures the number of rounds needed for the center node to transfer a message to all another node in the network. While radius and diameter are closely related, the previous lower bounds for diameter do not transfer to radius and it was conceivable that the radius of the graph could be computed much faster. Obtaining a non-trivial lower bound for radius is stated as an open problem in [66]. Another advantage of our technique is that it extends to computing the radius, for which we show that the same strong near-linear barriers above hold.

THEOREM 3. Any algorithm for computing the radius requires $\Omega\left(\frac{n}{\log^2 n}\right)$ rounds, even in networks of $\Theta(n \log n)$ edges.

Our techniques can also be used for proving lower bounds for approximating the network's radius, computing its eccentricities, and for verifying that a given subgraph is a spanner, even on sparse networks with a constant degree. The interested reader can find the details in [1].

1.1.2 Near-Quadratic Lower Bounds. High lower bounds for the CONGEST model can be obtained rather artificially, by forcing large inputs and outputs that must be exchanged, e.g., by having large edge weights, or by requiring a node to output its *t*-neighborhood for some value of *t*. However, until this work no super-linear lower bound for a natural problem was known, let alone near-quadratic bound. We remedy this state of affairs by showing quadratic and near-quadratic lower bound for several natural *decision* problems on graphs, where each input can be represented by $O(\log n)$ bits, and each output value consists of a single bit, or $O(\log n)$ bits. Specifically, using the bit-gadget we obtain graph constructions with small cuts that lead to the following lower bounds.

THEOREM 4. Any algorithm for computing a minimum vertex cover of the network or deciding whether there is a vertex cover of a given size requires $\Omega(n^2/\log^2 n)$ rounds.

This directly applies also to computing an exact maximum independent set, as the latter is the complement of an exact minimum vertex cover. This lower bound is in stark contrast to the recent $O(\log \Delta/\log \log \Delta)$ -round algorithm of [13] for obtaining a $(2 + \epsilon)$ -approximation to the minimum vertex cover.

An additional lower bound that we obtain using the bit-gadget is for coloring, as follows.

THEOREM 5. Any algorithm for coloring a χ -colorable network in χ colors, or for deciding if it is *c*-colorable for a given *c*, requires $\Omega(n^2/\log^2 n)$ rounds.

We further show that certain approximations of χ are hard, and we believe that such a lower bound should hold even for looser approximations. All these lower bounds hold for randomized algorithms which succeed with high probability.³

³An event occurs with high probability (w.h.p) if it occurs with probability $\frac{1}{n^c}$, for some constant c > 0.

We then show that not only NP-hard problems are near-quadratically hard in the CONGEST model, by presenting two simple problems that admit polynomial-time sequential algorithms, but require quadratic or near-quadratic time in the CONGEST model. The weighted cycle detection problem requires $\Omega(n^2/\log n)$ rounds, even when using randomized algorithms. The *identical* subgraph detection problem requires $\Omega(n^2)$ rounds deterministically, while we present an O(D)round randomized algorithm for it, yielding the strongest possible separation between deterministic and randomized complexities for global problems in the CONGEST model. A slight variant of this problem gives even stronger separation, for general problems: we prove an $\Omega(n^2)$ rounds lower bound for it, and give a constant-time randomized algorithm.

1.1.3 All Pairs Shortest Paths. An intriguing question in the CONGEST model is the complexity of computing exact weighted all-pairs-shortest-paths (APSP). The complexity of unweighted APSP is known to be $\Theta(n/\log n)$ [53, 67], both for deterministic and randomized algorithms. Several recent works study the complexity of computing weighted APSP [6, 7, 24, 46, 68], and the most recent results are a randomized $\tilde{O}(n)$ -round algorithm [24], and a deterministic $\tilde{O}(n^{3/2})$ -round algorithm [7].

We provide an extremely simple linear lower bound of $\Omega(n)$ rounds for weighted APSP, extending a construction of Nanongkai [88], which separates its complexity from that of the unweighted case. Moreover, we formally prove that the commonly used framework of reducing a two-party communication complexity problem to a problem in the CONGEST model, using a fixed-partition simulation, cannot provide a super-linear lower bound for weighted APSP, regardless of the function and the graph construction used. We then extend this claim for *t*-party communication complexity with a shared blackboard. For the randomized case, this is not surprising in light of the recent randomized $\tilde{O}(n)$ algorithm [24]; however, it shows that closing the gap for the deterministic case might require a new technique, unless the true complexity will turn out to be O(n).

1.1.4 Streaming Algorithms. The semi-streaming model of computation [49] is an important model for processing massive graphs. Here, a single processing unit with a bounded amount of memory obtains information of the graph edges one-by-one and is required to process them and return an output based on the graph properties. Usually, the memory is assumed to be of $O(n \operatorname{poly} \log n)$ bits for an *n*-node graph, and the number of allowed passes over the edges is one, constant, or logarithmic in n.

In the standard model, called the *edge arrival* model, the order of the edges is adversarial. In the node arrival model, the adversary is restricted in that edges must arrive grouped by nodes-all the edges connecting a node to the previous nodes arrive together; in the adjacency streaming model, all the edges adjacent to a node arrive together, regardless of the previous nodes; specifically, each edge appears twice, once for each of its endpoints. For simplicity, we prove our bounds for the edge arrival model, but it is immediate to check that they also apply to the node arrival and the adjacency streaming models.

We prove that constructions for lower bounds for the CONGEST model translate directly to give lower bounds for the semi-streaming model, and, with the standard parameters, impossibility results. Specifically, for the problems for which we obtain near-quadratic lower bounds in the CONGEST model, we establish that the product of the memory size and the number of passes in the streaming model must be quadratic in *n*.

The mentioned lower bound applies to vertex cover, maximum independent set, coloring and other problem in the semi-streaming model. Our construction for maximum independent set also easily translates to give the same lower bound for the maximum clique problem. Some bounds close to ours are known in the literature, as discussed next, but our work has several advantages: we

1:6

give a unified framework, using simple lower bound graphs and simple communication complexity problems, and these bounds are robust to multiple-pass algorithms and to variants of the model.

For vertex cover, there is a known lower bound of $\Omega(k^2)$ for deciding the existence of a cover of size k in one pass, and an algorithm for the problem using $\tilde{O}(k^2)$ memory [37]. Our work matches the lower bound for one-pass, and extends it to multiple-pass algorithms. The maximum independent set and maximum clique problems were also previously studied [27, 40, 58]; the known lower bounds also apply to these problems with a gap promise, so our construction is weaker in that sense. On the other hand, we improve upon the best results for maximum independent set [58] by a poly-logarithmic factor and in the simplicity of our construction, and on the results for maximum clique [27] in that we handle multiple-pass algorithms.

Note that the bounds we present are not only for problems with a linear output size, such as finding a maximum independent set or a coloring, but even for decision and computation problems, e.g., computing the size of a maximum independent set or approximating the chromatic number. For these problems, it is not even trivial that a linear memory is necessary.

Roadmap. In the following section we describe additional related work about the problems discussed in this paper. Section 2 contains our preliminaries. In Section 3 we define the bit gadget and discuss some of its properties. Sections 4, 5, and 6 contains our near-linear lower bounds, near-quadratic lower bounds for NP hard problems, and near-quadratic lower bounds for problems in P, respectively. Section 7 contains our results for computing APSP. Finally, in Section 8, we show how our results imply new lower bounds for the streaming model.

1.2 Additional Related Work

Vertex Coloring, Minimum Vertex Cover, and Maximum Independent Set: One of the most central problems in graph theory is vertex coloring, which has been extensively studied in the context of distributed computing (see, e.g., [15–20, 33, 38, 39, 48, 51, 52, 61, 83, 86, 95, 100] and references therein). The special case of finding a (Δ + 1)-coloring, where Δ is the maximum degree of a node in the network, has been the focus of many of these studies, but is a *local* problem, and can be solved in a much less than linear number of rounds. Much less attention was given to the problem of distributively coloring a graph with the minimal number of colors possible: it was discussed for rings [83], for planar graphs [5, 36], and for general graphs with a number of colors depending on minimal number of colors possible [100].

Our paper suggests a reason for this state of affairs: coloring with a minimal number of colors requires studying almost all the graph edges, so it is very likely that no nontrivial algorithms for this problem exist.

Another classical problem in graph theory is finding a minimum vertex cover (MVC). In distributed computing, the time complexity of approximating MVC has been addressed in several cornerstone studies [9, 10, 13, 20, 54, 55, 60, 72–75, 90, 96].

Finding a minimum size vertex cover is equivalent to finding a maximum size independent set, as mentioned above, but this equivalence is not approximation preserving. Distributed approximation algorithms for maximum independent set were studied in [12, 25, 41, 80]. Finally, finding a maximum independent set and finding maximum clique are equivalent in the sequential setting, but not in distributed or streaming settings. Nevertheless, our lower bounds for maximum independent set in the streaming model do translate to lower bounds for maximum clique.

Distance Computation: Computing a 2-approximation of the diameter D can be easily done by constructing a single BFS tree, in O(D) rounds. Getting close to 3/2-approximation for the diameter can be done in a sublinear number of rounds: A value \hat{D} satisfying $\lfloor 2D/3 \rfloor \leq \hat{D} \leq D$ can be computed

in $O(n^{3/4} + D)$ rounds [66] or in (incomparable) $O(D\sqrt{n} \log n)$ rounds [93]. These bounds were later improved to $O(\sqrt{n} \log n + D)$ rounds [79], and finally reduced to $O(\sqrt{n} \log n + D)$ rounds [64]. Using communication complexity techniques similar to the ones discussed in the current paper, it was recently shown that an exact computation of a (3/2)-approximation of the diameter cannot be done in sublinear time as it requires $\tilde{\Omega}(n)$ rounds, and achieving an approximation factor smaller that 11/6 requires at least time polynomial in n – specifically, $\tilde{\Omega}(n^{1/6})$ rounds are necessary [56].

Additional distance computation problems have been widely studied in the CONGEST model for both weighted and unweighted networks [1, 22, 53, 62, 64–67, 78, 79, 88, 93]. One of the most fundamental problems of distance computation is computing all pairs shortest paths. For unweighted networks, an upper bound of $O(n/\log n)$ was recently shown [67], matching an earlier lower bound [53]. Moreover, the possibility of bypassing this near-linear barrier for any constant approximation factor was ruled out [88]. For weighted APSP, the possibility of extending the lower bound from the unweighted case [88] was previously observed by Lenzen (see the second remark after Corollary 8.4 in his lecture notes [77]); for completeness, we provide a formal proof of this claim. For randomized weighted APSP, an $\tilde{O}(n^{5/3})$ -round algorithm was shown [46], improved by an $\tilde{O}(n^{5/4})$ -round algorithm [68], and finally by an $\tilde{O}(n)$ -round algorithm [24]. For the deterministic case, an $\tilde{O}(n^{3/2})$ -round algorithm was recently presented [6, 7].

All the aforementioned lower bounds are proved using reductions to two-party communication complexity problems. The first lower bounds were proved using a simulation argument, where the nodes are partitioned between Alice and Bob, and the partition dynamically changes during the course of the simulation [99]. In contrast, recent work uses a simpler simulation argument with a *fixed* partition [66]. In Section 7.2, we show that the fixed partition approach cannot be used to prove a super-linear lower bound for weighted APSP; however, this does not rule out the possibility to prove such a bound using the dynamic partition approach (recent work shows a randomized algorithm for the problem [24], but the deterministic case is still open [6]).

Streaming Algorithms: Streaming algorithms [63, 87] are a way to process massive information streams that cannot fit into the memory of a single machine. In this paper we focus on streaming algorithms for graph problems, and mainly on the semi-streaming algorithms [49], where the memory is assume to be in $\Theta(n \operatorname{poly} \log n)$.

Some problems solvable in the semi-streaming model include deciding connectivity and bipartiteness, building a minimum spanning tree, finding a 2-approximate maximum cardinality matching (all discussed in [49]), $(1 + \epsilon)(\Delta + 1)$ -coloring [23], finding a $(2 + \epsilon)$ -approximate maximum weight matching [91], building cut sparsifiers [8], spectral sparsifiers [71], spanners [21, 45, 47, 50], and counting subgraphs such as triangles [14], unweighted cycles [84], full bipartite graphs[29], and small graph minors [26]. A $(1 + \epsilon)$ -approximation to the single-source-shortest-paths problem can be found when $O(1/\epsilon^2)$ passes are allowed [22].

On the lower bounds side, maximum matching cannot be approximated better than $e/(e-1) \approx$ 1.58 factor [69] in the semi-streaming model. Deciding (s, t)-connectivity requires $\Omega(n)$ bits of memory, and $\Omega(n/R)$ bits if R passes on the input are allowed, and so does computing the connected components, testing planarity and more (see [63]). Maximum cut approximation was studied both for upper and lower bounds [70]. The mentioned lower bounds are achieved using reductions to communication complexity problems, and we essentially follow their footsteps in term of techniques, while achieving new lower bounds for different problems.

The Caro-Wei bound is a degree-sequence based lower bound on the size of a maximum independent set in a graph. Finding an independent set matching this bound was studied in [58], and evaluating the value of the bound was recently studied in [40]; note that such a set might not be a maximum independent set. There is a variety of upper and lower bounds for the maximum independent set and maximum clique problems, under a gap assumption: either the graph contains a large independent set (clique), or only a very small one [27, 59].

2 PRELIMINARIES

2.1 Computational Models

The CONGEST model: In the CONGEST model [92], the nodes of an undirected connected graph G = (V, E) on |V| = n nodes communicate over the graph edges in synchronous rounds. In each round, each node can send a message of $O(\log n)$ bits to each of its neighbors. The complexity measure of a distributed algorithm in this model is the number of rounds it needs in order to complete. A weighted graph G = (V, E, w) is a graph augmented with an edge weight function $w : E \rightarrow \{1, ..., W\}$. We assume that the maximum edge weight W is polynomial in n, and thus an edge weight, or a sum of O(n) edge weights, can be sent in a single message.

Each node is assumed to have a unique id in $\{1, ..., n\}$. At the beginning of an execution of an algorithm, each node knows its own id and, if the graph is weighted, also the weights of the edges adjacent to it. If the algorithm computes a graph parameter, it terminates when all nodes know the value of this parameter. If it outputs a labeling (e.g., a coloring or an indication of membership in a set) then each node should know its label.

The semi-streaming *model*: In the semi-streaming model [49], a single computational unit executes a centralized algorithm in order to process a large graph. The graph nodes are given in advance to the algorithm, and the edges are read one-by-one (with their weights), in an adversarial order. The algorithm is allowed to keep only *M* bits of memory, where usually $M = O(n \text{ poly } \log n)$, and to make only *R* passes over the input edges, where usually R = O(1) or $R = O(\log n)$.

2.2 Graph Parameters

We are interested in several classical graph problems. The distance between two nodes u, v in the graph, denoted d(u, v), is the minimum number of hops in a path between them in an unweighted graph, or the minimum weight of a path between them in a weighted graph. The *diameter* D of the graph is the maximum distance between two nodes in it. The *eccentricity* of a node u is $e(u) = \max_v \{d(u, v)\}$, and the *radius* of the graph is min_u $\{e(u)\}$. For a given integer i, an *i*-path in G is a simple path of i hops.

A vertex cover of a graph is a set $U \subseteq V$ such that for each edge $e \in E$ we have $e \cap U \neq \emptyset$. A *minimum vertex cover* is a vertex cover of minimum cardinality. An independent set is a set $U \subseteq V$ for which $u, v \in U \implies (u, v) \notin E$, and a clique is a set $U \subseteq V$ for which $u, v \in U \implies (u, v) \notin E$, and a clique is a set $U \subseteq V$ for which $u, v \in U \implies (u, v) \in E$. A *maximum independent set* is an independent set of maximum cardinality, and a *maximum clique* is a clique of maximum cardinality. A (proper) *c*-coloring of a graph is a function $f : V \to \{1, \ldots, c\}$ such that $(u, v) \in E \implies f(u) \neq f(v)$. The *chromatic number* χ of a graph is the minimum *c* such that a *c*-coloring of the graph exists.

2.3 Communication Complexity

In the two-party communication setting [76, 101], two players, Alice and Bob, are given two input strings, $x, y \in \{0, 1\}^K$, respectively, and need to jointly compute a function $f : \{0, 1\}^K \times \{0, 1\}^K \rightarrow \{\text{TRUE, FALSE}\}$ on their inputs. The *communication complexity* of a protocol π for computing f, denoted $CC(\pi)$, is the maximal number of bits Alice and Bob exchange in π , taken over all values of the pair (x, y). The *deterministic communication complexity* of f, denoted CC(f), is the minimum over $CC(\pi)$, taken over all deterministic protocols π that compute f.

In a *randomized protocol* π , Alice and Bob may each use a random bit string. A randomized protocol π computes f if the probability, over all possible bit strings, that π outputs f(x, y) is at

least 2/3. The *randomized communication complexity* of f, $CC^{R}(f)$, is the minimum over $CC(\pi)$, taken over all randomized protocols π that compute f.

For a vector x, let x[i] be the *i*-th bit in the string x. In the set-disjointness problem (DISJ_K), the function f is $DISJ_K(x, y)$, whose value is FALSE if there is an index $i \in \{0, ..., K-1\}$ such that x[i] = y[i] = 1, and TRUE otherwise. We say that x and y are disjoint if $DISJ_K(x, y) = TRUE$, and not disjoint otherwise. In the *Equality* problem (EQ_K), the function f is $EQ_K(x, y)$, whose output is TRUE if x = y, and FALSE otherwise. When K is clear from the context, or is determined in a later stage, we omit it from the notation.

Both the deterministic and randomized communication complexities of the DISJ_K problem are known to be $\Omega(K)$ [76, Example 3.22]. The deterministic communication complexity of EQ_K is in $\Omega(K)$ [76, Example 1.21], while its randomized communication complexity is in $\Theta(\log K)$ [97] (see also [76, Example 3.9]).

Remark: For some of our constructions which use the DISJ function, we need to exclude the all-0 or all-1 input vectors, in order to guarantee that the graphs are connected, as otherwise proving impossibility is trivial. However, this restriction does not change the asymptotic bounds for DISJ, since computing this function while excluding, e.g. the all-1 input vector, can be reduced to computing this function for inputs that are shorter by one bit (by having the last bit fixed to 0 or to 1).

2.4 Lower Bound Graphs

To prove lower bounds on the number of rounds necessary in order to solve a distributed problem in the CONGEST model, we use reductions from two-party communication complexity problems. The reductions are defined as follows.

DEFINITION 1. (Family of Lower Bound Graphs)

Fix an integer K, a function $f : \{0, 1\}^K \times \{0, 1\}^K \rightarrow \{\text{TRUE, FALSE}\}$ and a graph predicate P. A family of graphs $\{G_{x,y} = (V, E_{x,y}) \mid x, y \in \{0, 1\}^K\}$ with a partition $V = V_A \cup V_B$ is said to be a family of lower bound graphs for the CONGEST model w.r.t. f and P if the following properties hold:

- (1) Only the existence or the weight of edges in $V_A \times V_A$ may depend on x;
- (2) Only the existence or the weight of edges in $V_B \times V_B$ may depend on y;
- (3) $G_{x,y}$ satisfies the predicate P iff f(x,y) = TRUE.

We use the following theorem, which is standard in the context of communication complexitybased lower bounds for the CONGEST model (see, e.g. [1, 43, 53, 65]). Its proof is by a standard simulation argument.

THEOREM 6. Fix a function $f : \{0, 1\}^K \times \{0, 1\}^K \to \{\text{TRUE, FALSE}\}$ and a predicate P. If there is a family $\{G_{x,y}\}$ of lower bound graphs for the CONGEST model w.r.t. f and P with $C = E(V_A, V_B)$ then any deterministic algorithm for deciding P in the CONGEST model requires $\Omega(CC(f)/(|C|\log n))$ rounds, and any randomized algorithm for deciding P in the CONGEST model requires $\Omega(CC^R(f)/(|C|\log n))$ rounds.

PROOF. Let *ALG* be a distributed algorithm in the CONGEST model that decides *P* in *T* rounds. Given inputs $x, y \in \{0, 1\}^K$ to Alice and Bob, respectively, Alice constructs the part of $G_{x,y}$ for the nodes in V_A and Bob does so for the nodes in V_B . This can be done by items 1 and 2 in Definition 1, and since V_A and V_B are disjoint. Alice and Bob simulate *ALG* by exchanging the messages that are sent during the algorithm between nodes of V_A and nodes of V_B in either direction, while the messages within each set of nodes are simulated locally by the corresponding player without any communication. Since item 3 in Definition 1 also holds, we have that Alice and Bob correctly output



Fig. 1. The bit-gadget construction

f(x, y) based on the output of *ALG*. For each edge in the cut, Alice and Bob exchange $O(\log n)$ bits per round. Since there are *T* rounds and |C| edges in the cut, the number of bits exchanged in this protocol for computing *f* is $O(T | C| \log n)$. The lower bounds for *T* now follows directly from the lower bounds for CC(f) and $CC^{R}(f)$, where in the randomized case we note that an algorithm that succeeds with high probability definitely succeeds with probability at least 2/3.

In what follows, for each decision problem addressed, we describe a fixed graph construction G = (V, E) with a partition $V = V_A \cup V_B$, which we then generalize to a family of graphs $\{G_{x,y} = (V, E_{x,y}) \mid x, y \in \{0, 1\}^K\}$. We then show that $\{G_{x,y}\}$ is a family lower bound graphs w.r.t. to some communication complexity problem f and the required predicate P. By Theorem 6 and the known lower bounds for the two-party communication problem f, we deduce a lower bound for any algorithm for deciding P in the CONGEST model.

We use *n* for the number of nodes, *K* for the size of the input strings, and a third parameter *k* as an auxiliary parameter, usually for the size of the node-set that touches the edges that depend on the input. The number of nodes *n* determines *k* and *K*, and we usually only show the asymptotic relation between the three parameters and leave the exact values implicit.

3 THE BIT-GADGET CONSTRUCTION

The main technical novelty in our lower bounds comes from the ability to encode large communication complexity problems in graphs with small cuts. To this end, we use the following construction (see Figure 1).

Fix an integer k which is a power of 2, and start with two sets of k nodes each, $A = \{a^i \mid i \in \{0, ..., k-1\}\}$ and $B = \{b^i \mid i \in \{0, ..., k-1\}\}$. For each set $S \in \{A, B\}$, add two corresponding sets of log k nodes each, denoted $F_S = \{f_S^h \mid h \in \{0, ..., \log k - 1\}\}$ and $T_S = \{t_S^h \mid h \in \{0, ..., \log k - 1\}\}$. The latter are called the *bit-nodes* and they constitute the *bit-gadget*. Connect the nodes of each set $S \in \{A, B\}$ to their corresponding bit-nodes according to their indices, as follows. Let s^i be a node in a set $S \in \{A, B\}$, i.e., $s \in \{a, b\}$ and $i \in \{0, ..., k - 1\}$, and let i_h denote the *h*-th bit in the binary representation of *i*. For such s^i , define $bin(s^i) = \{f_S^h \mid i_h = 0\} \cup \{t_S^h \mid i_h = 1\}$, and connect s^i by an edge to each of the nodes in $bin(s^i)$. Finally, connect the bit-nodes: for each $h \in \{0, ..., \log k - 1\}$ connect f_A^h to f_B^h . Set $V_A = A \cup F_A \cup T_A$ and $V_B = V \setminus V_A$.



Fig. 2. Diameter lower bound construction. The dashed edges represent edges whose existence depends on x or y.

For example, consider the node $a^3 \in A$. The binary representation of the index i = 3 is $3_h = 0$ for all $h \ge 2$, and $3_0 = 3_1 = 1$. Hence, $bin(a^3) = \{f_A^h \mid h \ge 2\} \cup \{t_A^0, t_A^1\}$, and these are the nodes a^3 is connected to.

In the next sections, we augment the above construction with fixed nodes and edges, and then add some more edges according to the input strings, in order to create a family of graphs with some desired properties. The next claim exemplifies one of the basic properties of this construction.

CLAIM 1. For every $i, j \in \{0, ..., k-1\}$, if $i \neq j$ then $d(a^i, b^j) \leq 3$.

PROOF. Since $i \neq j$, there is $h \in \{0, ..., \log k - 1\}$ such that $i_h \neq j_h$. If $i_h = 1$ and $j_h = 0$, then the 3-path (a^i, t^h_A, f^h_B, b^j) connects the desired nodes; otherwise, $i_h = 0$ and $j_h = 1$, and the 3-path (a^i, f^h_A, t^h_B, b^j) connects the nodes.

It is not hard to also show that $d(a^i, b^j) \ge 3$ and that $d(a^i, b^i) = 5$, and we indeed prove similar claims in Section 4. In Section 5.1 we discuss the size and structure of a minimum vertex cover for this gadget.

4 NEAR-LINEAR LOWER BOUNDS FOR SPARSE GRAPHS

In this section we present our near-linear lower bounds for sparse networks. Sections 4.1 and 4.2 contain our lower bounds for computing the exact or approximate diameter, and Section 4.3 contains our lower bound for computing the radius.

4.1 Exact Diameter

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

Theorem 1 Any algorithm for computing the exact diameter requires $\Omega\left(\frac{n}{\log^2 n}\right)$ rounds, even in networks of $\Theta(n \log n)$ edges.

In order to prove Theorem 1, we describe a family of lower bound graphs with respect to the set-disjointness function and the predicate P that says that the graph has diameter at least 5. We

start by describing the fixed graph construction and then define the family of lower bound graphs and analyze its relevant properties.

The fixed graph construction: Start with the graph G and partition (V_A, V_B) described in Section 3, and add two nodes c_A, \bar{c}_A to V_A , and another two nodes c_B, \bar{c}_B to V_B (see Figure 2). For each set $S \in \{A, B\}$, connect all nodes in S to the center c_S , all the bit-nodes $F_S \cup T_S$ to \bar{c}_S , and the two centers c_S, \bar{c}_S to each other. Finally, connect \bar{c}_A to \bar{c}_B .

Adding edges corresponding to the strings x and y: Given two binary strings $x, y \in \{0, 1\}^k$, augment the graph defined above with additional edges, which defines $G_{x,y}$. For each $i \in \{0, ..., k-1\}$, if x[i] = 0 then add an edge between the nodes a^i and \bar{c}_A , and if y[i] = 0 then add an edge between b^i and \bar{c}_B .

CLAIM 2. For every u, v such that $(u, v) \notin (A \times B) \cup (B \times A)$, it holds that $d(u, v) \leq 4$.

PROOF. Observe that every node in $V \setminus (A \cup B)$ is connected to \bar{c}_A or to \bar{c}_B , and these two nodes are neighbors. Thus, the distance between every two nodes in $V \setminus (A \cup B)$ is at most 3. The claim follows from this, and from the fact that any node in A and any node in B are connected to a node in $V \setminus (A \cup B)$.

The following lemma is the main ingredient in proving that $\{G_{x,y}\}$ is a family of lower bound graphs.

LEMMA 1. The diameter of $G_{x,y}$ is at least 5 if and only if x and y are not disjoint.

PROOF. Assume that the sets are disjoint, i.e., for every $i \in \{0, ..., k-1\}$ either x[i] = 0 or y[i] = 0. We show that for every $u, v \in V$ it holds that $d(u, v) \le 4$. Consider the following cases:

- (1) $(u, v) \notin (A \times B) \cup (B \times A)$: By Claim 2, $d(u, v) \le 4$.
- (2) $u = a^i \in A$ and $v = b^j \in B$ (or vice versa) for $i \neq j$: Claim 1 implies $d(u, v) \leq 3$.
- (3) $u = a^i, v = b^i$ (or vice versa) for some *i*: By the assumption, either x[i] = 0 or y[i] = 0, and assume the former without loss of generality, implying that a^i is connected by an edge to \bar{c}_A . Thus, the path $(a^i, \bar{c}_A, \bar{c}_B, c_B, b^i)$ exists in the graph, and $d(a^i, b^i) \le 4$.

For the other direction, assume that the two sets are not disjoint, i.e., there is some $i \in \{0, ..., k-1\}$ for which x[i] = y[i] = 1. In this case, a^i is not connected by an edge to \bar{c}_A and b^i is not connected by an edge to \bar{c}_B . Note that V_A, V_B are disjoint, a^i belongs to V_A and b^i to V_B , so any path between a^i and b^i must go through an edge connecting a node from V_A and a node from V_B . Recall that $bin(a^i) = \{f_A^h \mid i_h = 0\} \cup \{t_A^h \mid i_h = 1\}$, and $bin(b^i)$ is similarly defined, with B instead of A. Fix a shortest path from a^i to b^i , and consider the following cases, distinguished by the first edge in the path crossing from V_A to V_B :

- (1) The path uses the edge (\bar{c}_A, \bar{c}_B) : Since a^i is not connected by an edge to \bar{c}_A , and b^i is not connected to \bar{c}_B , we have $d(a^i, \bar{c}_A) \ge 2$ and $d(b^i, \bar{c}_B) \ge 2$, so the length of the path is at least 5.
- (2) The path uses an edge (f_A^h, t_B^h) with $f_A^h \notin bin(a^i)$, or an edge (t_A^h, f_B^h) with $t_A^h \notin bin(a^i)$: For the first case, note that a^i and t_A^h are not connected by an edge, and they do not even have a common neighbor. Thus, $d(a^i, t_A^h) \ge 3$, and $d(a^i, b^i) \ge 5$. The case of (t_A^h, f_B^h) with $t_A^h \notin bin(a^i)$ is analogous.
- (3) The path uses an edge (f^h_A, t^h_B) with f^h_A ∈ bin(aⁱ), or an edge (t^h_A, f^h_B) with t^h_A ∈ bin(aⁱ): The definitions of bin(aⁱ) and bin(bⁱ) immediately imply t^h_B ∉ bin(bⁱ) for the first case, or f^h_B ∉ bin(bⁱ) for the second. The rest of the argument is the same as the previous: d(t^h_B, bⁱ) ≥ 3 or d(f^h_B, bⁱ) ≥ 3, both implying d(aⁱ, bⁱ) ≥ 5.



Fig. 3. Diameter approximation lower bound construction. Bold (blue) edges represent q-paths.

Thus, any path between a^i and b^i must have length at least 5.

Having constructed a family of lower bound graphs, we are now ready to prove Theorem 1.

Proof of Theorem 1: To complete the proof of Theorem 1, note that $n \in \Theta(k)$, and thus $K = |x| = |y| = \Theta(n)$. Furthermore, the only edges in the cut $E(V_A, V_B)$ are the edges between nodes in $F_A \cup T_A$ and nodes in $F_B \cup T_B$, and the edge (\bar{c}_A, \bar{c}_B) . Thus, in total, there are $\Theta(\log n)$ edges in the cut $E(V_A, V_B)$. Since Lemma 1 shows that $\{G_{x,y}\}$ is a family of lower bound graphs, we can apply Theorem 6 and deduce that any algorithm in the CONGEST model for deciding whether a given graph has a diameter at least 5 requires at least $\Omega(k/\log^2 n) = \Omega(n/\log^2 n)$ rounds. Finally, observe that the number of edges in the construction is $O(n \log n)$.

4.2 $(3/2 - \varepsilon)$ -Approximation of the Diameter

In this section we show how to modify our sparse construction presented in the previous section in order to achieve a near-linear lower bound even for computing a $(3/2 - \varepsilon)$ -approximation of the diameter.

Theorem 2 For any constant $0 < \varepsilon < 1/2$, any algorithm for computing a $(3/2 - \varepsilon)$ -approximation of the diameter requires $\Omega\left(\frac{n}{\log^3 n}\right)$ rounds, even in networks of $\Theta(n)$ edges.

As in the proof of Theorem 1, we show that there is a family of lower bound graphs with respect to the set-disjointness function and the predicate *P* that says that the graph has a diameter of length at least \tilde{D} , where \tilde{D} is an integer that may depend on *n*. Note that, unlike the case of proving a lower bound for computing the exact diameter, here we need to construct a family of lower bound graphs for which even an algorithm that computes a $(3/2 - \epsilon)$ -approximation to the diameter can be used to determine whether *P* holds.

The fixed graph construction: Start with our graph construction from the previous section and stretch it by replacing some edges by paths of length *q*, an integer that is chosen later. Apply the following changes to the construction described in the previous section (see Figure 3):

(1) Replace all the edges inside V_A and all the edges inside V_B by paths of length q. The cut edges and the edges that depend on the inputs remain intact.

(2) Add two additional sets of nodes $\overline{A} = \{\overline{a}^i \mid i \in \{0, \dots, k-1\}\}, \overline{B} = \{\overline{b}^i \mid i \in \{0, \dots, k-1\}\},$ each of size *k*. For each $i \in \{0, \dots, k-1\}$, connect \overline{a}^i to a^i and \overline{b}^i to b^i , by paths of length *q*.

The partition of nodes into V_A and V_B is similar to the previous: V_A is composed of the nodes in V_A in the previous construction, the set \overline{A} , and the nodes in the paths between them. V_B is composed of the rest of the nodes.

Adding edges corresponding to the strings x and y: Given two binary strings $x, y \in \{0, 1\}^k$, define $G_{x,y}$ by adding edges to the graph in a way similar to the one described in the previous section. That is, if x[i] = 0 then add an edge between the nodes a^i and \bar{c}_A , and if y[i] = 0 then add an edge between b^i and \bar{c}_B .

In this construction, the nodes $A \cup B \cup \{\bar{c}_A, \bar{c}_B\}$ serve as *hubs*, in the sense that any node is at distance at most q from one of the hubs and the hubs are at distance at most 2q + 2 from one another, implying $D \le 4q + 2$. The only exception to this is if the input strings are not disjoint, in which case there are a^i, b^i with $d(a^i, b^i) \ge 4q + 1$, implying $D \ge d(\bar{a}^i, \bar{b}^i) \ge 6q + 1$. Let us formalize these arguments.

CLAIM 3. For every $u, v \in A \cup B \cup \{\bar{c}_A, \bar{c}_B\}$, if there is no index $i \in \{0, ..., k-1\}$ such that $u = a^i$ and $v = b^i$, then $d(u, v) \le 2q + 1$.

PROOF. For each *i* there is a path from a^i to \bar{c}_B of length 2q+1, which passes through c_A and \bar{c}_A , so $d(a^i, \bar{c}_B) \leq 2q+1$, and also $d(a^i, \bar{c}_A) \leq 2q+1$. Similarly, $d(b^i, \bar{c}_A) \leq 2q+1$, and also $d(b^i, \bar{c}_B) \leq 2q+1$. For every $i, j \in \{0, ..., k-1\}$, there is a 2*q*-path from a^i to a^j through c_A , so $d(a^i, a^j) \leq 2q$, and similarly $d(b^i, b^j) \leq 2q$ using c_B .

We are left with the case of a^i and b^j where $i \neq j$, which is a simple extension of Claim 1. In this case, there must be some h such that $i_h \neq j_h$, and assume without loss of generality that $i_h = 1$ and $j_h = 0$. Hence, a^i is connected to t_A^h by a q-path, and b^j is similarly connected to f_B^h . Since t_A^h and f_B^h are connected by an edge, we have $d(a^i, b^j) \leq 2q + 1$, as desired.

The next lemma proves that $\{G_{x,y}\}$ is a family of lower bound graphs.

LEMMA 2. If x and y are disjoint then the diameter of $G_{x,y}$ is at most 4q + 2, and otherwise it is at least 6q + 1.

PROOF. Assume that the sets are disjoint, so for every $i \in \{0, ..., k-1\}$ either x[i] = 0 or y[i] = 0. Hence, for every *i* there is a (2q + 2)-path from a^i to b^i , whether through $\bar{c}_A, \bar{c}_B, c_B$ or through $c_A, \bar{c}_A, \bar{c}_B$. Hence, $d(a^i, b^i) \leq 2q + 2$, and together with Claim 2, we can conclude that for every $u', v' \in (A \cup B \cup \{\bar{c}_A, \bar{c}_B\})$ we have $d(u, v) \leq 2q + 2$.

For each node $v \in V$, there exists a node $v' \in (A \cup B \cup \{\bar{c}_A, \bar{c}_B\})$ such that $d(v, v') \leq q$. Consider any two nodes $u, v \in V$, and the nodes $u', v' \in (A \cup B \cup \{\bar{c}_A, \bar{c}_B\})$ closest to them. By the triangle inequality, $d(u, v) \leq d(u, u') + d(u', v') + d(v', v) \leq q + 2q + 2 + q = 4q + 2$, as desired.

Assume that the two sets are not disjoint, i.e., there is some $i \in \{0, ..., k-1\}$ such that x[i] = y[i] = 1. Hence, a^i is not connected directly to \bar{c}_A and b^i is not connected directly to \bar{c}_B . We show that $d(\bar{a}^i, \bar{b}^i) \ge 6q + 1$.

First, note that any path connecting \bar{a}^i and \bar{b}^i must go through the *q*-paths connecting \bar{a}^i to a^i and \bar{b}^i to b^i , hence it suffices to prove $d(a^i, b^i) \ge 4q + 1$. Observe that no shortest path from a^i to b^i uses an edge (a^j, \bar{c}_A) for $j \ne i$: using such an edge requires at least a (2q)-path from a^i to a^j or to \bar{c}_A , plus the edge (a^j, \bar{c}_A) itself, while there is a path of length merely 2q from a^i directly to a^j or to \bar{c}_A . A similar claim holds for b^i and edges of the form (b^j, \bar{c}_B) . Thus, a shortest path connecting a^i and b^i only uses *q*-paths that replace edges of the graph from the previous section, and cut edges. The proof of Lemma 1 shows that any such path must go through at least 4 edges



Fig. 4. Radius lower bound

which are internal to V_A or V_B , and one cut edge. In the current construction, this translates into 4 *q*-paths and an edge, which implies $d(a^i, b^i) \ge 4q + 1$, as desired.

Using this family of lower bound graphs, we prove Theorem 2.

Proof of Theorem 2: To complete the proof of Theorem 2, consider the predicate *P* of the diameter of the graph being at least 6q + 1. In order to make $\{G_{x,y}\}$ a family of lower bound graphs for which even a $(3/2 - \epsilon)$ -approximation algorithm decides *P*, we choose a constant *q* such that $(\frac{3}{2} - \varepsilon) \cdot (4q + 2) < (6q + 1)$, which holds for any $q > \frac{1}{2\varepsilon} - \frac{1}{2}$. Observe that $k \in \Theta(n/\log n)$ for any constant ε , and thus $K = |x| = |y| \in \Theta(n/\log n)$. Furthermore, the number of edges in the cut is $|E(V_A, V_B)| \in \Theta(\log n)$. By applying Theorem 6 to the above construction, we deduce that any algorithm in the CONGEST model for computing a $(3/2 - \epsilon)$ -approximation for the diameter requires at least $\Omega(n/\log^3(n))$ rounds. Finally, observe that we replaced most of the edges from the previous construction by a path of a constant number of nodes, so the total number of edges is linear in the number nodes, as desired. П

4.3 Radius

In this section we extend our sparse construction and show that computing the radius requires a near-linear number of rounds in the CONGEST model, even on sparse graphs.

Theorem 3 Any algorithm for computing the radius requires $\Omega\left(\frac{n}{\log^2 n}\right)$ rounds, even in networks of $\Theta(n \log n)$ edges.

The fixed graph construction: The graph construction for the radius is very similar to the one described in Section 4.1, with the following changes (see also Figure 4).

- (1) For each *h* ∈ {0,..., log *k* − 1}, add the edge (*f*^{*h*}_{*A*}, *t*^{*h*}_{*A*}).
 (2) Add a 2-path (*w*⁰, *w*¹, *w*²), and connect *w*⁰ to all the nodes in *A*. Add the nodes *w*⁰, *w*¹, *w*² to V_A .

Adding edges corresponding to the strings x and y: Given two binary strings $x, y \in \{0, 1\}^k$, we define $G_{x,y}$ as follows. If x[i] = 1, we add an edge between a^i and \bar{c}_A , and if y[i] = 1 we add an

edge between b^i and \bar{c}_B . Note that unlike the reduction from the previous section, here we add an edge if the corresponding bit is 1 rather than 0. Recall that the eccentricity of a node u is defined by $e(u) = \max_v \{d(u, v)\}$.

CLAIM 4. The graph family $\{G_{x,y}\}$ defined above has the following properties:

- (1) For every node $u \in V \setminus A$, it holds that $e(u) \ge 4$.
- (2) For every $a^i \in A$ and $u \in V \setminus \{b^i, c_B\}$, it holds that $d(a^i, u) \leq 3$.

PROOF. For $u \notin A \cup \{w^0, w^1, w^2\}$, any path from u to w^2 must go through a node $a^i \in A$. Since $d(w^2, a^i) = 3$ for every such $a^i \in A$, we get $d(w^2, u) > 3$. For $u \in \{w^0, w^1, w^2\}$, note that no a^i has a common neighbor with c_B , so $d(a^i, c_B) \ge 3$ for all a^i , so $d(u, c_B) > 3$, proving 1.

For 2, if $u \in \{w^0, w^1, w^2\}$ then $d(a^i, u) \leq 3$ by construction. If $u \in \{c_A, \bar{c}_A, \bar{c}_B\}$ then the path $(a^0, c_A, \bar{c}_A, \bar{c}_B)$ proves the claim. Every $v \in F_A \cup T_A$ is in $bin(a^i)$ or a neighbor of a node in $bin(a^i)$, so $d(a^i, v) \leq 2$, and every $u \in F_B \cup T_B$ is a neighbor of some $v \in F_A \cup T_A$, implying $d(a^i, u) \leq 3$. Finally, if $u = b^i$ for $i \neq j$ then $d(a^i, u) \leq 3$ an argument that is analogous to Claim 1. \Box

LEMMA 3. The strings x and y are disjoint if and only if the radius of G is at least 4.

PROOF. If the strings are not disjoint, i.e., there exists an $i \in \{0, ..., k-1\}$ such that x[i] = y[i] = 1, then the edges $(a^i, \bar{c}_A), (b^i, \bar{c}_B)$ exist in the graph. The 3-path $(a^i, \bar{c}_A, \bar{c}_B, b^i)$ implies $d(a^i, b^i) \leq 3$, and the 3-path $(a^i, \bar{c}_A, \bar{c}_B, c_B)$ implies $d(a^i, c_B) \leq 3$. The rest of the proof is analogous to Claim 4(2): $d(a^i, u) \leq 3$ for every $u \in V$, and $e(a^i) \leq 3$ as desired.

If the radius is at most 3, then by Claim 4(1) there must be a node $a^i \in A$ with $e(a^i) \leq 3$, which implies $d(a^i, b^i) \leq 3$. Since the nodes of $bin(a^i)$ and $bin(b^i)$ are not neighbors, there must be a 3-path connecting a^i and b^i that goes through the cut edge (\bar{c}_A, \bar{c}_B) . Hence, the edges (a^i, \bar{c}_A) and (b^i, \bar{c}_B) exist in $G_{\{x,y\}}$, so x[i] = y[i] = 1 and the strings are not disjoint.

Proof of Theorem 3: Note that the number of edges on the cut is $|E(V_A, V_B)| = \Theta(\log n)$, and that $K = k = \Theta(n)$. By Lemma 3, $\{G_{x,y}\}$ is a family of lower bound graphs, so we can apply Theorem 6 to the above construction to deduce that any algorithm in the CONGEST model for computing the radius of a sparse network requires at least $\Omega(n/\log^2(n))$ rounds.

5 NEAR-QUADRATIC LOWER BOUNDS FOR GENERAL GRAPHS

In this section we present the first super-linear lower bounds for natural graph problems in the CONGEST model. Section 5.1 introduces a relatively simple lower bound for the minimum vertex cover and maximum independent set algorithms, and Section 5.2 presents lower bounds for χ -coloring algorithms.

5.1 Minimum Vertex Cover

The first near-quadratic lower bound we present is for computing a minimum vertex cover, as stated in the following theorem.

Theorem 4 Any algorithm for computing a minimum vertex cover of the network or deciding whether there is a vertex cover of a given size requires $\Omega(n^2/\log^2 n)$ rounds.

A set of nodes is a vertex cover if and only if its complement is an independent set, which implies that finding the minimum size of a vertex cover is equivalent to finding the maximum size of an independent set. Thus, the following theorem is a direct corollary of Theorem 4.

THEOREM 7. Any distributed algorithm for computing a maximum independent set or for deciding whether there is an independent set of a given size requires $\Omega(n^2/\log^2 n)$ rounds.



Fig. 5. Lower bound graph for minimum vertex cover. The dashed edges represent edges whose existence depends on x or y.

Observe that a lower bound for deciding whether there is a vertex cover of some given size or not implies a lower bound for computing a minimum vertex cover. This is because computing the size of a given subset of nodes can be easily done in O(D) rounds using standard tools. Therefore, to prove Theorem 4 it is sufficient to prove its second part. We do so by describing a family of lower bound graphs with respect to the set-disjointness function and the predicate P that says that the graph has a vertex cover of size M, where M = M(n) is chosen later. We begin with describing the fixed graph construction G = (V, E) and then define the family of lower bound graphs and analyze its relevant properties.

The fixed graph construction: Start with two copies of the gadget described in Section 3. That is, the fixed graph (Figure 5) consists of four sets of size $k: A_1 = \{a_1^i \mid i \in \{0, \dots, k-1\}\}, A_2 = \{a_2^i \mid i \in \{0, \dots, k-1\}\}, B_1 = \{b_1^i \mid i \in \{0, \dots, k-1\}\}$ and $B_2 = \{b_2^i \mid i \in \{0, \dots, k-1\}\}$. Each such set *S* is connected to 2 log *k* nodes: $F_S = \{f_S^h \mid h \in \{0, \dots, \log k - 1\}\}$ and $T_S = \{t_S^h \mid h \in \{0, \dots, \log k - 1\}\}$, where $F_{A_1}, T_{A_1}, F_{B_1}, T_{B_1}$ constitute one bit-gadget, and $F_{A_2}, T_{A_2}, F_{B_2}, T_{B_2}$ constitute another. Partition the nodes into $V_A = A_1 \cup A_2 \cup F_{A_1} \cup T_{A_2} \cup T_{A_2} \cup T_{A_2}$ and $V_B = V \setminus V_A$.

Let s_{ℓ}^i be a node in a set $S \in \{A_1, A_2, B_1, B_2\}$, i.e., $s \in \{a, b\}$, $\ell \in \{1, 2\}$ and $i \in \{0, \dots, k-1\}$. As before, connect s_{ℓ}^i to $bin(s_{\ell}^i) = \{f_S^h \mid i_h = 0\} \cup \{t_S^h \mid i_h = 1\}$. In addition, in each of the sets A_1, A_2, B_1, B_2 , connect all the nodes to one another, forming a clique. Also, connect the nodes of

the bit-gadgets to form 4-cycles: for each $h \in \{0, ..., \log k - 1\}$ and $\ell \in \{1, 2\}$, connect the 4-cycle $(f_{A_\ell}^h, t_{A_\ell}^h, f_{B_\ell}^h, t_{B_\ell}^h)$.

The following two claims address the basic properties of vertex covers of G.

CLAIM 5. Any vertex cover of G must contain at least k - 1 nodes from each of the cliques in A_1, A_2, B_1 and B_2 , and at least $4 \log k$ bit-nodes.

PROOF. In order to cover all the edges of each of the cliques in A_1, A_2, B_1 and B_2 , any vertex cover must contain at least k - 1 nodes of each clique. For each $h \in \{0, ..., \log k - 1\}$ and $\ell \in \{1, 2\}$, in order to cover the edges of the 4-cycle $(f_{A_\ell}^h, t_{A_\ell}^h, f_{B_\ell}^h, t_{B_\ell}^h)$, any vertex cover must contain at least two of the cycle nodes.

CLAIM 6. If $U \subseteq V$ is a vertex cover of G of size $4(k-1) + 4\log k$, then there are two indices $i, j \in \{0, ..., k-1\}$ such that all four nodes $a_1^i, a_2^j, b_1^i, b_2^j$ are not in U.

PROOF. By Claim 5, U must contain k - 1 nodes from each clique A_1, A_2, B_1 and B_2 , and $4 \log k$ bit-nodes, so in each clique there is a node that is not in U. Let $a_1^i, a_2^j, b_1^{i'}, b_2^{j'}$ be the nodes in A_1, A_2, B_1, B_2 which are not in U, respectively. To cover the edges connecting a_1^i to $bin(a_1^i), U$ must contain all the nodes of $bin(a_1^i)$, and similarly, U must contain all the nodes of $bin(a_1^i)$. If $i \neq i'$ then there is an index $h \in \{0, \ldots, \log k - 1\}$ such that $i_h \neq i'_h$, so one of the edges $(f_{A_1}^h, t_{B_1}^h)$ or $(t_{A_1}^h, f_{B_1}^h)$ is not covered by U. Thus, it must hold that i = i'. A similar argument shows j = j'.

Adding edges corresponding to the strings x and y: Given two binary strings $x, y \in \{0, 1\}^{k^2}$, assume they are indexed by pairs of the form $(i, j) \in \{0, ..., k-1\}^2$. To define $G_{x,y}$, augment the graph G defined above with edges as following. For each such pair (i, j), if x[i, j] = 0, then add an edge between a_1^i and a_2^j , and if y[i, j] = 0 then we add an edge between b_1^i and b_2^j . To prove that $\{G_{xy}\}$ is a family of lower bound graphs, it remains to prove the following lemma.

LEMMA 4. The graph $G_{x,y}$ has a vertex cover of cardinality $M = 4(k-1) + 4\log k$ iff x and y are not disjoint.

PROOF. For the first implication, assume that x and y are not disjoint, and let $i, j \in \{0, ..., k-1\}$ be such that x[i, j] = y[i, j] = 1. Note that in this case a_1^i is not connected to a_2^j , and b_1^i is not connected to b_2^j . Define a set $U \subseteq V$ as

$$U \coloneqq (A_1 \setminus \{a_1^i\}) \cup (A_2 \setminus \{a_2^j\}) \cup (B_1 \setminus \{b_1^i\}) \cup (B_2 \setminus \{b_2^j\})$$
$$\cup \operatorname{bin}(a_1^i) \cup \operatorname{bin}(a_2^j) \cup \operatorname{bin}(b_1^i) \cup \operatorname{bin}(b_2^j)$$

and show that *U* is a vertex cover of $G_{x,y}$, as follows.

First, *U* covers all the edges inside the cliques A_1, A_2, B_1 and B_2 , as it contains k - 1 nodes from each clique. These nodes also cover all the edges connecting nodes in A_1 to nodes in A_2 and all the edges connecting nodes in B_1 to nodes in B_2 , since the edges (a_1^i, a_2^j) and (b_1^i, b_2^j) do not exist in $G_{x,y}$. Furthermore, *U* covers any edge connecting nodes $(A_1 \setminus \{a_1^i\}) \cup (A_2 \setminus \{a_2^j\}) \cup (B_1 \setminus \{b_1^i\}) \cup (B_2 \setminus \{b_2^j\})$ to the bit-gadgets. For each node $s \in \{a_1^i, a_2^j, b_1^i, b_2^j\}$, the nodes bin(*s*) are in *U*, so *U* also covers the edges connecting *s* to the bit-gadget. Finally, *U* covers all the edges inside the bit-gadgets, as from each 4-cycle $(f_{A_\ell}^h, t_{A_\ell}^h, f_{B_\ell}^h, t_{B_\ell}^h)$ it contains two non-adjacent nodes: if $i_h = 0$ then $f_{A_1}^h, f_{B_1}^h \in U$ and otherwise $t_{A_1}^h, t_{B_1}^h \in U$, and if $j_h = 0$ then $f_{A_2}^h, f_{B_2}^h \in U$ and otherwise $t_{A_2}^h, t_{B_2}^h \in U$. Thus, *U* is a vertex cover of size $4(k - 1) + 4 \log k$, as claimed.

For the other implication, let $U \subseteq V$ be a vertex cover of $G_{x,y}$ of size $4(k-1) + 4 \log k$. Since all the edges of *G* are also edges of $G_{x,y}$, *U* is also a cover of *G*, so Claim 6 implies that there are



Fig. 6. Lower bound graph for 3-coloring. Heavy edges represent a set of edges connecting a node to a set of nodes. Edges connecting c_h^1 and c_h^2 to other nodes in V_B are omitted.

indices $i, j \in \{0, ..., k-1\}$ such that $a_1^i, a_2^j, b_1^i, b_2^j$ are not in *U*. Since *U* is a cover, the graph does not contain the edges (a_1^i, a_2^j) and (b_1^i, b_2^j) , so we conclude that x[i, j] = y[i, j] = 1, which implies that *x* and *y* are not disjoint.

Having constructed the family of lower bound graphs, we are now ready to prove Theorem 4.

Proof of Theorem 4: Note that $n \in \Theta(k)$, and thus $K = |x| = |y| = \Theta(n^2)$, and furthermore, the only edges in the cut $E(V_A, V_B)$ are the edges between nodes in $F_{A_1} \cup T_{A_1} \cup F_{A_2} \cup T_{A_2}$ and nodes in $F_{B_1} \cup T_{B_1} \cup F_{B_2} \cup T_{B_2}$, which are in total $\Theta(\log n)$ edges. Since Lemma 4 shows that $\{G_{x,y}\}$ is a family of lower bound graphs, we can apply Theorem 6 to deduce that any algorithm in the CONGEST model for deciding whether a given graph has a cover of cardinality $M = 4(k-1) + 4\log k$ requires at least $\Omega(K/\log^2(n)) = \Omega(n^2/\log^2(n))$ rounds.

5.2 Graph Coloring

In this section we consider the problems of coloring a graph with χ colors, computing χ and approximating it. We prove the following theorem.

Theorem 5 Any algorithm for coloring a χ -colorable network in χ colors, or for deciding if it is *c*-colorable for a given *c*, requires $\Omega(n^2/\log^2 n)$ rounds.

The fixed graph construction: Define G = (V, E) as follows (see Figure 6 for the general construction, and Figure 9 for an example with specific *k* and inputs).

ACM Trans. Algor., Vol. 1, No. 1, Article 1. Publication date: January 2021.



(a) A failed attempt to color A_1 without using c_0 . The path $\bar{a}_1^0, \ldots, \bar{a}_3^1$ (marked by a double line) must be colored alternately with c_0 and c_2 , and so the edge $(\bar{a}_{3}^{1}, c_{a}^{1})$ is violated.



(b) A coloring of A_1 with a_1^1 colored c_0

Fig. 7. Part of the 3-coloring lower bound graph for k = 4, which assures that at least one node in A_1 is colored by c_0 (Proof of Claim 7)

Start with the family of graphs defined in Section 5.1, omitting the clique edges inside the four sets A_1, A_2, B_1, B_2 . Add the following two gadgets to the graph.

- (1) Add three nodes c_a^0, c_a^1, c_a^2 connected as a triangle, another three nodes c_b^0, c_b^1, c_b^2 connected as a triangle, and edges connecting c_a^i to c_b^j for each $i \neq j \in \{0, 1, 2\}$.
- (2) For each set $S \in \{A_1, A_2, B_1, B_2\}$, add two sets of nodes, $\overline{S} = \{\overline{s}_{\ell}^i \mid s_{\ell}^i \in S\}$ and $\overline{\overline{S}} = \{\overline{s}_{\ell}^i \mid s_{\ell}^i \in S\}$. For each $\ell \in \{1, 2\}$ and $i \in \{0, \dots, k-1\}$ connect a path $(s_{\ell}^i, \bar{s}_{\ell}^i, \bar{s}_{\ell}^i)$, and for each $\ell \in \{1, 2\}$ and $i \in \{0, \ldots, n-2\}$, connect \bar{s}_{ℓ}^i to \bar{s}_{ℓ}^{i+1} .

In addition, connect the gadgets by the following edges:

- (a) $(c_a^1, f_{A_1}^h), (c_a^1, t_{A_1}^h), (c_b^1, f_{B_1}^h)$ and $(c_b^1, t_{B_1}^h)$, for each $h \in \{0, \dots, \log k 1\}$. (b) $(c_a^2, f_{A_2}^h), (c_a^2, t_{A_2}^h), (c_b^2, f_{B_2}^h)$ and $(c_b^2, t_{B_2}^h)$, for each $h \in \{0, \dots, \log k 1\}$. (c) (c_a^2, a_1^i) and (c_a^1, \bar{a}_1^i) , for each $i \in \{0, \dots, k 1\}; (c_a^2, \bar{a}_1^0)$ and (c_a^2, \bar{a}_1^{k-1}) .
- (d) (c_b^2, b_1^i) and $(c_b^1, \overline{b}_1^i)$, for each $i \in \{0, \dots, k-1\}$; $(c_b^2, \overline{b}_1^0)$ and $(c_b^2, \overline{b}_1^{k-1})$.
- (e) (c_a^1, a_2^i) and (c_a^2, \bar{a}_2^i) , for each $i \in \{0, \dots, k-1\}$; (c_a^1, \bar{a}_2^0) and (c_a^1, \bar{a}_2^{k-1}) .
- (f) (c_b^1, b_2^i) and $(c_b^2, \overline{b}_2^i)$, for each $i \in \{0, \dots, k-1\}$; $(c_b^1, \overline{b}_2^0)$ and $(c_b^1, \overline{b}_2^{k-1})$.

Assume there is a proper 3-coloring of G. Denote by c_0, c_1 and c_2 the colors of c_a^0, c_a^1 and c_a^2 respectively. By construction, these are also the colors of c_b^0, c_b^1 and c_b^2 , respectively. In Section 5.1 we present a specific vertex cover, and mention that its complement is an independent set. In the current section, this independent set is colored by c_0 , and the part of the graph that did not appear in the previous section is used in order to guarantee that coloring this independent set by c_0 is the only valid option. The following claims are thus very similar to those appearing in the previous section.

CLAIM 7. In each set $S \in \{A_1, A_2, B_1, B_2\}$, at least one node is colored by c_0 .

PROOF. We start by proving the claim for $S = A_1$ (see figure 7). Assume, towards a contradiction, that none of the nodes of A_1 is colored by c_0 . All these nodes are connected to c_a^2 , so they are not colored by c_2 either, i.e., they all must be colored c_1 . Hence, all the nodes \bar{A}_1 are colored by c_0 and c_2 . The nodes of \bar{A}_1 are connected to c_a^1 , so they are colored by c_0 and c_2 as well.



(a) A failed attempt to color only a_1^0 and b_1^1 by c_0 leads to the violation of the edge $(f_{A_1}^0, t_{B_1}^0)$



(b) A valid coloring with a_1^1 and b_1^1 colored c_0

Fig. 8. Part of the 3-coloring lower bound graph for k = 4, which assures that if a_1^i is colored c_0 then so does b_1^i (proof of Claim 8)

Hence, we have a path $(\bar{a}_1^0, \bar{a}_1^0, \bar{a}_1^1, \bar{a}_1^1, \dots, \bar{a}_1^{k-1}, \bar{a}_1^{k-1})$ with an even number of nodes, starting in \bar{a}_1^0 and ending in \bar{a}_1^{k-1} . The colors of this path must alternate between c_0 and c_2 , but both its endpoint are connected to c_a^2 , so they must both be colored c_0 , a contradiction.

A similar proof shows the claim for $S = B_1$. For $S \in \{A_2, B_2\}$, we use a similar argument but change the roles of c_1 and c_2 .

CLAIM 8. For all $i, j \in \{0, ..., k-1\}$, the node a_1^i is colored by c_0 iff b_1^i is colored by c_0 and the node a_2^j is colored by c_0 iff b_2^j is colored by c_0 .

PROOF. Assume a_1^i is colored by c_0 , so all of its adjacent nodes $bin(a_1^i)$ can only be colored by c_1 or c_2 (see Figure 8). As all of these nodes are connected to c_a^1 , they must be colored by c_2 . By Claim 7, there is a node in $b_1^{i'} \in B_1$ that is colored by c_0 , and by a similar argument the nodes $bin(b_1^{i'})$ must also be colored by c_2 .

If $i \neq i'$ then there is a bit *h* such that $i_h \neq i'_h$, and there must be a pair of neighboring nodes $(f^h_{A_1}, t^h_{B_1})$ or $(t^h_{A_1}, f^h_{B_1})$ which are both colored by c_2 , a contradiction. Thus, the only option is i = i'.

An analogous argument shows that if b_1^i is colored by c_0 , then so does a_1^i . For a_2^j and b_2^j , similar arguments apply, where c_1 plays the role of c_2 .



Fig. 9. A coloring of a lower bound graph (proof of Lemma 5). In this example k = 4, and the inputs are x[i, j] = 1 for j = 3 (and all i), y[i, j] = 1 for |i - j| even. Note that x[1, 3] = y[1, 3] = 1.

Adding edges corresponding to the strings x and y: To get $G_{x,y}$ from G, add edges exactly as in the previous section: if x[i, j] = 0 then add (a_1^i, a_2^j) , and if y[i, j] = 0 then add (b_1^i, b_2^j) . The following lemma proves that $\{G_{x,y}\}$ is a family of lower bound graphs.

LEMMA 5. The graph $G_{x,y}$ is 3-colorable iff x and y are not disjoint.

PROOF. Assume $G_{x,y}$ is 3-colorable, and denote the colors of c_a^0 , c_a^1 and c_a^2 by c_0 , c_1 and c_2 respectively, as before. By Claim 7, there are nodes $a_1^i \in A_1$ and $a_2^j \in A_2$ that are both colored by c_0 . Hence, the edge (a_1^i, a_2^j) does not exist in $G_{x,y}$, implying x[i, j] = 1. By Claim 8, the nodes b_1^i and b_2^j are also colored c_0 , so y[i, j] = 1 as well, implying that x and y are not disjoint.

For the other direction, assume x and y are not disjoint, i.e., there is an index $(i, j) \in \{0, ..., k-1\}^2$ such that x[i, j] = y[i, j] = 1. Consider the following coloring (see Figure 9).

- (1) Color c_a^i and c_b^i by c_i , for $i \in \{0, 1, 2\}$.
- (2) Color the nodes a_1^i, b_1^i, a_2^j and b_2^j by c_0 . Color the nodes $a_1^{i'}$ and $b_1^{i'}$, for $i' \neq i$, by c_1 , and the nodes $a_2^{j'}$ and $b_1^{j'}$, for $j' \neq j$, by c_2 . (3) Color the nodes of $bin(a_1^i)$ by c_2 , and similarly color the nodes of $bin(b_1^i)$ by c_2 . Color the
- rest of the nodes in this gadget, i.e., $bin(a_1^{k-i})$ and $bin(b_1^{k-i})$, by c_0 . Similarly, color $bin(a_2^j)$ and $bin(b_2^j)$ by c_0 and $bin(a_2^{k-j})$ and $bin(b_2^{k-j})$ by c_1 . (4) Finally, color the nodes of the forms \bar{s}_{ℓ}^i and \bar{s}_{ℓ}^i as follows.
- - (a) Color \bar{a}_1^i and \bar{b}_1^i by c_1 , all nodes $\bar{a}_1^{i'}$ and $\bar{b}_1^{i'}$ with i' < i by c_0 , and all nodes $\bar{a}_1^{i'}$ and $\bar{b}_1^{i'}$ with i' > i by c_2 .
 - (b) Similarly, color \bar{a}_2^i and \bar{b}_2^i by c_2 , all nodes $\bar{a}_2^{i'}$ and $\bar{b}_2^{i'}$ with i' < i by c_0 , and all nodes $\bar{a}_2^{i'}$ and $\bar{b}_{2}^{i'}$ with i' > i by c_1 .
 - (c) Color all nodes $\bar{a}_1^{i'}$ and $\bar{b}_1^{i'}$ with i' < i by c_2 , and all nodes $\bar{a}_1^{i'}$ and $\bar{b}_1^{i'}$ with $i' \ge i$ by c_0 .



Fig. 10. A schematic figure of the lower bound graph for *c*-coloring, c = 9 (proof of Claim 9). Bold lines between sets represent the existence of all the edges between the two sets. The squares represent the previous graph, except for the color nodes which are explicitly indicated.

(d) Similarly, color all nodes $\bar{a}_2^{i'}$ and $\bar{b}_2^{i'}$ with i' < i by c_1 , and all nodes $\bar{a}_2^{i'}$ and $\bar{b}_2^{i'}$ with $i' \ge i$ by c_0 .

Checking all edges gives that the above coloring is indeed a proper 3-coloring of $G_{x,y}$, which completes the proof.

Having constructed the family of lower bound graphs, we are now ready to prove Theorem 5. **Proof of Theorem 5:** The edges in the cut $E(V_A, V_B)$ are the 6 edges connecting $\{c_a^0, c_a^1, c_a^2\}$ and $\{c_b^0, c_b^1, c_b^2\}$, and 2 edges from every 4-cycle of the nodes of $F_{A_1} \cup T_{A_1} \cup F_{B_1} \cup T_{B_1}$ and $F_{A_2} \cup T_{A_2} \cup F_{B_2} \cup T_{B_2}$, for a total of $\Theta(\log n)$ edges. Note that $n \in \Theta(k)$ and $K = k^2 \in \Theta(n^2)$. Lemma 5 shows that $\{G_{x,y}\}$ is a family of lower bound graphs with respect to DISJ_K and the predicate $\chi > 3$, so by applying Theorem 6 on the above partition we deduce that any algorithm in the CONGEST model for deciding whether a given graph is 3-colorable requires $\Omega(n^2/\log^2 n)$ rounds.

Any algorithm that computes χ of the input graph, or produces a χ -coloring of it, may be used to deciding whether $\chi \leq 3$, in O(D) additional rounds. Thus, the lower bound applies to these problems as well.

A lower bound for *c*-coloring: Our construction and proof naturally extend to handle *c*-coloring, for any $c \ge 3$. We prove the next theorem.

CLAIM 9. Any algorithm that decides if $\chi(G) \leq c$, for an integer $3 \leq c < n$ that may depend on n, requires $\Omega((n-c)^2/(c\log n + \log^2 n))$ rounds.

The proof of this claim is an extension of the proof of Theorem 5. Start with the graph $G_{x,y}$ defined above, add new nodes denoted c_a^i , $i \in \{3, ..., c-1\}$, and connect them to all of V_A and new nodes denoted c_b^i , $i \in \{3, ..., c-1\}$, and connect them to all of V_B and also to c_a^0 , c_a^1 and c_a^2 (see Figure 10). The nodes c_a^i are added to V_a , and the rest are added to V_b , which increases the cut size by $\Theta(c)$ edges.

Assume the extended graph is colorable by *c* colors, and denote by c_i the color of the node c_a^i (these nodes are connected by a clique, so their colors must be distinct). The nodes c_b^i , $i \in \{2, ..., c-1\}$ form a clique, and they are all connected to the nodes c_a^0 , c_a^1 and c_a^2 , so they are colored by the

1:25

colors $\{c_3, \ldots, c_{c-1}\}$, in some arbitrary order. All the original nodes of V_A are connected to c_a^i , $i \in \{3, \ldots, c-1\}$, and all the original nodes of V_B are connected to c_b^i , $i \in \{3, \ldots, c-1\}$, so the original graph must be colored by 3 colors, which we know is possible iff x and y are not disjoint. Thus, the newly defined family $\{G_{x,y}\}$ is a family of lower bound graphs, and we can easily prove the claim.

Proof of Claim 9: To construct $G_{x,y}$, we added 2c - 6 nodes to the graph, so now $K = k^2 = \Theta((n-c)^2)$. Thus, the new graphs constitute a family of lower bound graphs with respect to DISJ_K and the predicate $\chi > c$, the communication complexity of DISJ_K is in $\Omega(K^2) = \Omega((n-c)^2)$, the cut size is $\Theta(c + \log n)$, and Theorem 6 completes the proof.

6 QUADRATIC AND NEAR-QUADRATIC LOWER BOUNDS FOR PROBLEMS IN P

In this section we support our claim that what makes problems hard for the CONGEST model is not necessarily them being NP-hard problems. First, we address a class of subgraph detection problems, which requires detecting cycles of length 8 and a given weight, and show a near-quadratic lower bound on the number of rounds required for solving it, although its sequential complexity is polynomial. Then, we define a problem which we call the *Identical Subgraphs Detection* problem, in which the goal is to decide whether two given subgraphs are identical. While this last problem is rather artificial, it allows us to obtain a strictly quadratic lower bound for the CONGEST model.

6.1 Weighted Cycle Detection

In this section we show a lower bound on the number of rounds needed in order to decide if the graph contains a simple cycle of length 8 and weight *W*, such that *W* is a polylog(n)-bit value given as an input. Note that this problem can be solved easily in polynomial time in the sequential setting by simply checking all the $\binom{n}{8} \cdot 7!$ potential cycles of length 8.

THEOREM 8. Any distributed algorithm that decides if a weighted graph contains a simple cycle of length 8 and a given weight requires $\Omega(n^2/\log n)$ rounds.

Similarly to the previous sections, to prove Theorem 8 we describe a family of lower bound graphs with respect to the set-disjointness function and the predicate P that says that the graph contains a simple cycle of length 8 and weight W.

The fixed graph construction: The fixed graph construction G = (V, E) is defined as follows (see Figure 11). The set of nodes contains four sets A_1, A_2, B_1 and B_2 , each of size $k \ge 3$. For each set $S \in \{A_1, A_2, B_1, B_2\}$ there is a node c_S , which is connected to each of the nodes in S by an edge of weight 0. In addition there is an edge between c_{A_1} and c_{B_1} of weight 0 and an edge between c_{A_2} and c_{B_2} of weight 0. We set $V_A = A_1 \cup A_2 \cup \{c_{A_1}, c_{A_2}\}$ and $V_B = V \setminus V_A$.

Adding edges corresponding to the strings x and y: Given two binary strings $x, y \in \{0, 1\}^{k^2}$, we augment the fixed graph G defined above with additional edges, which defines $G_{x,y}$. If x[i, j] = 1, then we add an edge of weight $k^3 + ki + j$ between the nodes a_1^i and a_2^j . If y[i, j] = 1, then we add an edge of weight $k^3 - (ki + j)$ between the nodes b_1^i and b_2^j . We denote by InputEdges the set of edges depending on the input, i.e. the edges in $(A_1 \times A_2) \cup (B_1 \times B_2)$.

CLAIM 10. Any simple cycle of weight $2k^3$ contains exactly two edges from InputEdges, one from $A_1 \times A_2$ and one from $B_1 \times B_2$.

PROOF. The weight of each edge in InputEdges is in $\{k^3 - k^2 + 1, \dots, k^3 + k^2 - 1\}$ (the extremes are $w(b_1^{k-1}, b_2^{k-1})$ and $w(a_1^{k-1}, a_2^{k-1})$, if those edges exist), and all other edges weigh 0. A cycle of



Fig. 11. Lower bound graph for weighted cycle detection

weight $2k^3$ must contain at least two edges from InputEdges since $k^3 + k^2 - 1 < 2k^3$, and cannot contain three or more of these since $2k^3 < 3(k^3 - k^2 + 1)$.

Since the edges of $A_1 \times A_2$ weigh at least k^3 , and all but one of them weighs strictly more than that, two of these weigh more than $2k^3$. Similarly, two edges of $B_1 \times B_2$ weigh less than $2k^3$.

To prove that $\{G_{x,y}\}$ is a family of lower bound graphs, we prove the following lemma.

LEMMA 6. The graph $G_{x,y}$ contains a simple cycle of length 8 and weight $W = 2k^3$ if and only if x and y are not disjoint.

PROOF. For the first direction, assume that *x* and *y* are not disjoint, and let $i, j \in \{0, ..., k-1\}$ be such that x[i, j] = y[i, j] = 1. The 8-cycle $(a_1^i, c_{A_1}, c_{B_1}, b_1^i, b_2^j, c_{B_2}, c_{A_2}, a_2^j)$ has weight $w(a_1^j, a_2^i) + w(b_1^i, b_2^j) = k^3 + ki + j + k^3 - ki - j = 2k^3$, as needed.

For the other direction, assume that the graph contains a simple cycle *C* of length 8 and weight $2k^3$. By Claim 10, *C* contains two edges of the form $(a_1^i, a_2^j) \in A_1 \times A_2$ and $(b_1^{i'}, b_2^{j'}) \in B_1 \times B_2$. Since all other edge weights in *C* are 0, we conclude $w(a_1^i, a_2^j) + w(b_1^{i'}, b_2^{j'}) = 2k^3$, i.e. $k^3 + ki + j + k^3 - (ki' + j') = 2k^3$, or k(i - i') = (j' - j). The fact that |j' - j| < k implies i = i' and j = j', completing the proof.

Having constructed a family of lower bound graphs, we are now ready to prove Theorem 8. **Proof of Theorem 8:** Note that $n \in \Theta(k)$, and thus $K = |x| = |y| = \Theta(n^2)$. Furthermore, the only edges in the cut $E(V_A, V_B)$ are the edges (c_{A_1}, c_{B_1}) and (c_{A_2}, c_{B_2}) . Since Lemma 6 shows that $\{G_{x,y}\}$ is a family of lower bound graphs, we apply Theorem 6 on the above partition to deduce that any algorithm in the CONGEST model for deciding whether a given graph contains a simple cycle of length 8 and weight $W = 2k^3$ requires at least $\Omega(K/\log n) = \Omega(n^2/\log n)$ rounds.

6.2 Identical Subgraphs Detection

In this section we show the strongest possible, quadratic lower bound, for a global decision problem which can be solved in polynomial time in the sequential setting.

ACM Trans. Algor., Vol. 1, No. 1, Article 1. Publication date: January 2021.



Fig. 12. Lower bound graph for the identical subgraphs detection problem

Consider the following graph problem.

DEFINITION 2. (The Identical Subgraphs Detection Problem)

Given a weighted graph G = (V, E, w), a partition $V = V_A \cup V_B$, $|V_A| = |V_B|$, and node labeling $V_A = \{a^0, ..., a^{k-1}\}$ and $V_B = \{b^0, ..., b^{k-1}\}$, the Identical Subgraphs Detection problem is to determine whether the subgraph induced by V_A is identical to the subgraph induced by V_B , in the sense that for each $i, j \in \{0, ..., k-1\}$ it holds that $(a^i, a^j) \in E$ if and only if $(b^i, b^j) \in E$ and $w(a^i, a^j) = w(b^i, b^j)$ if these edges exist.

The identical subgraphs detection problem can be solved easily in linear time in the sequential setting by a single pass over the set of edges. However, as we prove next, it requires a quadratic number of rounds in the CONGEST model, for any deterministic solution (note that this restriction did not apply in the previous sections). We emphasize that in the distributed setting, the input to each node in A or B in this problem includes its enumeration as a_i or b_i , and the weights of its edges. The outputs of all nodes should be TRUE if the subgraphs are identical, and FALSE otherwise.

THEOREM 9. Any deterministic algorithm for solving the identical subgraphs detection problem requires $\Omega(n^2)$ rounds.

To prove Theorem 9 we describe a family of lower bound graphs.

The fixed graph construction: The fixed graph G = (V, E) is composed of two disjoint cliques on sets of k-nodes each, denoted $V_A = \{a^0, ..., a^{k-1}\}$ and $V_B = \{b^0, ..., b^{k-1}\}$, and one extra edge (a^0, b^0) (see Figure 12).

Adding edge weights corresponding to the strings x and y: Given two binary strings x and y, each of $K = \binom{k}{2} \log n$ bits, augment the graph G with additional edge weights as follows, to define $G_{x,y}$. For simplicity, assume that x and y are vectors of $(\log n)$ -bit numbers, each having $\binom{k}{2}$ entries enumerated as x[i, j] and y[i, j], with $i < j, i, j \in \{0, \ldots, k-1\}$. For each such i and j set the weights $w(a^i, a^j) = x[i, j]$ and $w(b^i, b^j) = y[i, j]$, and set $w(a^0, b^0) = 0$. Note that $\{G_{x,y}\}$ is a family of lower bound graphs with respect to EQ_K and the predicate P that says that the subgraphs are identical in the aforementioned sense.

Proof of Theorem 9: Note that $n \in \Theta(k)$, and thus $K = |x| = |y| = \Theta(n^2 \log n)$. Furthermore, the only edge in the cut $E(V_A, V_B)$ is the edge (a^0, b^0) . Since $\{G_{x,y}\}$ is a family of lower bound graphs, we can apply Theorem 6 on the above partition to deduce that because of the lower bound for EQ_{*K*}, any deterministic algorithm in the CONGEST model for solving the identical subgraphs detection problem requires at least $\Omega(K/\log n) = \Omega(n^2)$ rounds.

In a deterministic distributed algorithm for the identical subgraphs detection problem running on our family of lower bound graphs, information about essentially all the edges and weights in the subgraphs induced on V_A or V_B needs to be sent across the edge (a^0, b^0) . This might raise the suspicion that this problem is reducible to learning the entire graph, making the lower bound trivial. To argue that this is far from being the case, we present a randomized algorithm that solves the identical subgraphs detection problem in O(D) rounds and succeeds w.h.p. This has the additional benefit of providing the strongest possible separation between deterministic and randomized complexities for global decision problems in the CONGEST model, as the former is $\Omega(n^2)$ and the latter is at most O(D).

THEOREM 10. There is a randomized algorithm that solves the identical subgraphs detection problem with probability at least $1 - O(1/n^2)$ in O(D) rounds.

PROOF. Our starting point is the following randomized algorithm for the EQ_K problem, presented, e.g., in [76, Exersise 3.6]. Alice chooses a prime number p among the first K^2 primes uniformly at random. She treats her input string x as a binary representation of an integer $\bar{x} = \sum_{\ell=0}^{K-1} 2^{\ell} x_{\ell}$, and sends p and $\bar{x} \pmod{p}$ to Bob. Bob similarly computes \bar{y} , compares $\bar{x} \mod p$ with $\bar{y} \mod p$, and returns TRUE if they are equal and false otherwise. The error probability of this protocol is at most 1/K.

We present a simple adaptation of this algorithm for the identical subgraph detection problem. Consider the following encoding of a weighted induced subgraph on V_A : for each pair i, j of indices, we have $\lceil \log W \rceil + 1$ bits, indicating the existence of the edge and its weight (recall that $W \in \text{poly } n$ is the upper bound on the edge weights). This weighted induced subgraph is thus represented by a $K \in O(n^2 \log n)$ bit-string, denoted $x = x_0, \ldots, x_{K-1}$, and each pair (i, j) has a set $S_{i,j}$ of indices representing the edge (a^i, a^j) and its weight. The bits $\{x_\ell \mid \ell \in s_{i,j}\}$ are known to both a^i and a^j , and in the algorithm we use the node with smaller index in order to encode these bits. Similarly, a $K \in O(n^2 \log n)$ bit-string, denoted $y = y_0, \ldots, y_{K-1}$ encodes a weighted induced subgraph on V_B .

The Algorithm. Given a graph with nodes enumerated as in Definition 2, the algorithm starts with an arbitrary node, say a^0 , and constructs a BFS tree from it, which completes in O(D) rounds. Then, a^0 chooses a prime number p among the first K^2 primes uniformly at random and sends p to all the nodes over the tree, in another O(D) rounds.

Each node a^i computes the sum $\sum_{j>i} \sum_{\ell \in S_{i,j}} x_\ell 2^\ell \mod p$, and the nodes then aggregate these local sums modulo p up the tree, until a^0 computes the sum $\bar{x} \mod p = \sum_{j \neq i} \sum_{\ell \in S_{i,j}} x_\ell 2^\ell \mod p$. A similar procedure is then invoked by a^0 (not by b^0) w.r.t. \bar{y} . Finally, a^0 compares $\bar{x} \mod p$ and $\bar{y} \mod p$, and downcasts over the BFS tree its output, which is TRUE if these values are equal and is FALSE otherwise.

If the subgraphs are identical, a^0 always returns TRUE, while otherwise their encoding differs in at least one bit, and as in the case of EQ_K, a^0 returns TRUE falsely with probability at most $1/K \in O(1/n^2)$.

A more artificial problem can give even a larger gap: given the labeled graph *G* defined in our lower bound construction with a weight function satisfying $w(a^0, a^j) = 0$ and $w(b^0, b^j) = 0$ for all *j*, the *balanced hourglass problem* is to determine whether $w(a^i, a^j) = w(b^i, b^j)$ for all $i, j \in \{0, ..., k-1\}$ (see Figure 13) The same lower and upper bounds hold: for the lower bound, the inputs are of $\binom{k-1}{2} \log n$ bits, which is asymptotically the same; for the upper bound, we get O(D) = O(1). Thus, this problem gives asymptotically the maximal possible gap for any decision in the CONGEST model.



Fig. 13. Lower bound graph for the balanced hourglass problem

7 WEIGHTED APSP

In this section we use the following, natural extension of Definition 1, in order to address more general 2-party functions, as well as distributed problems that are not decision problems. For a non-decision graph problem P, it is useful to think about P as a multi-output problem that takes as input a graph G = (V, E) and gives an output to each of the nodes in V. For example, in the all-pairs-shortest-paths problem, the output of each node is the list of its distances to all other nodes in the graph.

Consider a function $f : \{0, 1\}^{K_1} \times \{0, 1\}^{K_2} \rightarrow \{0, 1\}^{L_1} \times \{0, 1\}^{L_2}$, and denote $f = (f_1, f_2)$, where $f_1 : \{0, 1\}^{K_1} \times \{0, 1\}^{K_2} \rightarrow \{0, 1\}^{L_1}$ represents the left hand side of f, and $f_2 : \{0, 1\}^{K_1} \times \{0, 1\}^{K_2} \rightarrow \{0, 1\}^{L_2}$ represents the right hand side. We define a family of lower bound graphs in a similar way as Definition 1, except that instead of a predicate, we have a multi-output graph problem, and we replace item 3 in the definition with a generalized requirement that says that for $G_{x,y}$, the outputs of the nodes (given by the multi-output problem) in V_A uniquely determine $f_1(x, y)$, and the outputs of the nodes in V_B uniquely determine $f_2(x, y)$. Let \mathcal{G} be the set of all possible weighted graphs of size n and let \mathcal{D} be the output domain of the nodes (i.e., the set of all possible outputs) for a specific graph problem P. The new definition is as follows.

DEFINITION 3. (Family of Lower Bound Graphs)

Fix an integer K, a function $f : \{0, 1\}^{K_1} \times \{0, 1\}^{K_2} \to \{0, 1\}^{L_1} \times \{0, 1\}^{L_2}$ and a multi-output graph problem $P : \mathcal{G} \to \mathcal{D}^n$. A family of graphs $\{G_{x,y} = (V, E_{x,y}) \mid x, y \in \{0, 1\}^K\}$ with a partition $V = V_A \cup V_B$ is said to be a family of lower bound graphs for the CONGEST model w.r.t. f and P if the following properties hold:

- (1) Only the existence or the weight of edges in $V_A \times V_A$ may depend on x;
- (2) Only the existence or the weight of edges in $V_B \times V_B$ may depend on y;
- (3) $f(x,y) = (f_1(x,y), f_2(x,y))$, where $f_1(x,y)$ can be determined by the outputs of the nodes in V_A in $P(G_{x,y})$, and $f_2(x,y)$ can be determined by the outputs of V_B .

Next, we argue that theorem similar to Theorem 6 holds for non-decision graph problems.

THEOREM 11. Fix a function $f : \{0, 1\}^{K_1} \times \{0, 1\}^{K_2} \rightarrow \{0, 1\}^{L_1} \times \{0, 1\}^{L_2}$ and a graph problem *P*. If there is a family $\{G_{x,y}\}$ of lower bound graphs that satisfies Definition 3 with $C = E(V_A, V_B)$ then any deterministic algorithm for solving *P* requires $\Omega(CC(f)/|C| \log n)$ rounds, and any randomized algorithm for deciding *P* requires $\Omega(CC^R(f)/|C| \log n)$ rounds.

The proof is by the same simulation argument as used in the proof of of Theorem 6. Notice that the only difference between the theorems, apart from the sizes of the inputs and outputs of f, are

with respect to item 3 in the definition of a family of lower bound graphs. However, the essence of this condition remains the same and this is all that is required for the proof: the outputs of the nodes in V_A determine the output of Alice for f(x, y), and outputs of the nodes in V_B determine the output of Bob.

7.1 A Linear Lower Bound for Weighted APSP

Nanongkai [88] showed that any algorithm in the CONGEST model for computing a poly(n)-approximation for weighted all pairs shortest paths (APSP) requires at least $\Omega(n/\log n)$ rounds. A similar result was also proved independently by Lenzen and Patt-Shamir [79]. Nanongkai also gave a $\Omega(n/\log n)$ lower bound for exactly computing unweighted APSP, and Lenzen [77] observed that this can be extended into an $\Omega(n)$ lower bound for weighted APSP. Here, we give a full proof of it, for completeness and future reference.

In this section we show that a slight modification to this construction yields an $\Omega(n)$ lower bound for computing exact weighted APSP. As explained in the introduction, this gives a separation between the complexities of the weighted and unweighted versions of APSP. At a high level, while we use the same simple topology for our lower bound as in [88], the reason that we are able to shave off the extra logarithmic factor is because our construction uses $O(\log n)$ bits for encoding the weight of each edge out of many optional weights, while in [88] only a single bit is used per edge for encoding one of only two options for its weight.

THEOREM 12. Any algorithm for computing weighted all pairs shortest paths requires at least $\Omega(n)$ rounds.

The reduction is from the following, perhaps simplest, two-party communication problem. Alice has an input string *x* of size *K* and Bob needs to learn the string of Alice. In terms of the above definition, this problem is computing the function $f : \{0, 1\}^K \times \{0, 1\}^0 \rightarrow \{0, 1\}^0 \times \{0, 1\}^K$ defined by $f(x, \bot) = (\bot, x)$. Any algorithm (possibly randomized) for solving this problem requires at least $\Omega(K)$ bits of communication, by a trivial information theoretic argument.

Notice that the problem of having Bob learn Alice's input is not a binary function as addressed in Section 2. Similarly, computing weighted APSP is not a decision problem, but rather a problem whose solution assigns a value to each node (which is its vector of distances from all other nodes). We therefore use the extended Theorem 11 stated above.

The fixed graph construction: The fixed graph G = (V, E) is composed of n - 2 nodes, $A = \{a_0, ..., a_{n-3}\}$, all connected to an additional node a, which is connected to another node b. Set $V_A = A \cup \{a\}$ and $V_B = \{b\}$.

Adding edge weights corresponding to the string x: Given a binary string x of size $K = (n-2) \log n$, assume for simplicity that x is a vector of n-2 numbers, each represented by $O(\log n)$ bits. To define G_x , add to G the edge weights $w(a^i, a) = x[i]$ for $i \in \{0, ..., n-3\}$, and w(a, b) = 0. It is straightforward to see that G_x is a family of lower bound graphs for the function f.

Proof of Theorem 12: To prove Theorem 12, note that $K = |x| = \Theta(n \log n)$, and that the cut $E(V_A, V_B)$ has a single edge, (a, b). Since $\{G_x\}$ is a family of lower bound graphs with respect to f on K bits, Theorem 11 implies that any algorithm in the CONGEST model for computing weighted APSP requires $\Omega(K/\log n) = \Omega(n)$ rounds.

7.2 The Alice-Bob Framework Cannot Give a Super-Linear Lower Bound for Weighted APSP

In this section we argue that a reduction from any 2-party function with a fixed partition of the graph into Alice and Bob's sides is incapable of providing a super-linear lower bound for computing



(a) A graph split into G_A and G_B . The shaded nodes are C: lightshaded are C_A and dark-shaded are C_B





(c) Using G'_A and distances in G_B , Alice computes the distances from V_A to all the graph node

Fig. 14. Lemma 7 and its proof applied to a specific graph. All unmarked edges weight 1

weighted all pairs shortest paths in the CONGEST model. A more detailed inspection of our analysis shows a stronger claim: our claim also holds for algorithms for the CONGEST-BROADCAST model, where in each round each node must send the same $O(\log n)$ -bit message to all of its neighbors. The following theorem states our claim.

THEOREM 13. Let $f : \{0, 1\}^{K_1} \times \{0, 1\}^{K_2} \rightarrow \{0, 1\}^{L_1} \times \{0, 1\}^{L_2}$ be a function and let $G_{x,y}$ be a family of lower bound graphs w.r.t. f and the weighted APSP problem. When applying Theorem 11 to f and $G_{x,y}$, the lower bound obtained for the number of rounds for computing weighted APSP is at most linear in n.

Roughly speaking, we show that given an input graph G = (V, E) with a partition $V = V_A \dot{\cup} V_B$, such that the graph induced by the nodes in V_A is simulated by Alice and the graph induced by nodes in V_B is simulated by Bob, Alice and Bob can compute weighted all pairs shortest paths by communicating $O(n \log n)$ bits of information for each node touching the cut $C = E(V_A, V_B)$ induced by the partition. In this way, we show that any attempt to apply Theorem 11 cannot give a lower bound higher than $\Omega(n)$: we consider an arbitrary function f, and an arbitrary family of lower bound graphs with respect to a function f and the weighted APSP problem, defined according to the extended definition from the beginning of Section 7. We then prove that Alice and Bob can compute weighted APSP, which determines their output for f, by exchanging only $O(|V(C)|n \log n)$ bits, where V(C) is the set of nodes touching C. This implies that CC(f) is at most $O(|V(C)|n \log n)$. Thus, the lower bound obtained by Theorem 11 cannot be better than $\Omega(n)$, and hence no super-linear lower bound can be deduced using this framework.

Formally, given a graph $G = (V_A \cup V_B, E)$, let $G_A = (V_A, E_A)$ be the subgraph induced by the nodes in V_A and let $G_B = (V_B, E_B)$ be the subgraph induced by the nodes in V_B (see Figure 14(a)). Let $C = E(V_A, V_B)$, and let V(C) denote the nodes touching the cut C, with $C_A = V(C) \cap V_A$ and $C_B = V(C) \cap V_B$. For a graph H, denote the weighted distance between two nodes u, v by wd_H(u, v).

LEMMA 7. Let $G = (V_A \cup V_B, E, w)$ be a weighted graph. Suppose that G_A, C_B, C and the values of w on E_A and C are given as input to Alice, and that G_B, C_A, C and the values of w on E_B and C are given as input to Bob.

Then, Alice can compute the distances in G from all nodes in V_A to all nodes in V and Bob can compute the distances from all nodes in V_B to all the nodes in V, using $O(|V(C)| n \log n)$ bits of communication.

PROOF. We describe a protocol for the required computation. For each node $u \in C_B$, Bob sends to Alice the weighted distances in G_B from u to all nodes in V_B , that is, Bob sends {wd_{G_B}(u, v) | $u \in C_B, v \in V_B$ } (or ∞ for pairs of nodes not connected in G_B). Alice constructs a virtual graph

 $G'_A = (V'_A, E'_A, w'_A)$ (see Figure 14(b)) with the nodes $V'_A = V_A \cup C_B$ and edges $E'_A = E_A \cup C \cup (C_B \times C_B)$. The edge-weight function w'_A is defined by $w'_A(e) = w(e)$ for each $e \in E_A \cup C$, and by $w'_A(u, v) = w_{G_B}(u, v)$ for $u, v \in C_B$, as received from Bob. Alice then computes the set of all weighted distances in G'_A , that is, $\{wd_{G'_A}(u, v) \mid u, v \in V'_A\}$.

Alice assigns her output for the weighted distances in *G* as follows (see Figure 14(c)). For two nodes $u, v \in V_A \cup C_B$, Alice outputs their weighted distance in G'_A , $wd_{G'_A}(u, v)$. For a node $u \in V'_A$ and a node $v \in V_B \setminus C_B$, Alice outputs $\min\{wd_{G'_A}(u, x) + wd_{G_B}(x, v) \mid x \in C_B\}$, where $wd_{G'_A}$ is the distance in G'_A as computed by Alice, and wd_{G_B} is the distance in G_B that was sent by Bob.

For Bob to compute his required weighted distances, similar information is sent by Alice to Bob, that is, Alice sends to Bob the weighted distances in G_A from each node $u \in C_A$ to all nodes in V_A . Bob constructs the analogous graph G'_B and outputs his required distance. The next paragraph formalizes this for completeness, but may be skipped by a convinced reader.

Formally, Alice sends $\{\mathrm{wd}_{G_A}(u,v) \mid u \in C_A, v \in V_A\}$. Bob constructs $G'_B = (V'_B, E'_B, w'_B)$ with $V'_B = V_B \cup C_A$ and edges $E'_B = E_B \cup C \cup (C_A \times C_A)$. The edge-weight function w'_B is defined by $w'_B(e) = w(e)$ for each $e \in E_B \cup C$, and $w'_B(u,v)$ for $u, v \in C_A$ is defined to be the weighted distance between u and v in G_A , as received from Alice (or ∞ if they are not connected in G_A). Bob then computes the set of all weighted distances in G'_B , $\{\mathrm{wd}_{G'_B}(u,v) \mid u, v \in V'_B\}$. Bob assigns his output for the weighted distances in G as follows. For two nodes $u, v \in V_B \cup C_A$, Bob outputs their weighted distance in G'_B , $\mathrm{wd}_{G'_B}(u,v)$. For a node $u \in V'_B$ and a node $v \in V_A \setminus C_A$, Bob outputs $\min\{\mathrm{wd}_{G'_B}(u,x) + \mathrm{wd}_{G_A}(x,v) \mid x \in C_A\}$, where $\mathrm{wd}_{G'_B}$ is the distance in G'_B as computed by Bob, and wd_{G_A} is the distance in G_A that was sent by Alice.

Complexity. Bob sends to Alice the distances from all nodes in C_B to all node in V_B , which takes $O(|C_B| |V_B| \log n)$ bits, and similarly Alice sends $O(|C_A| |V_A| \log n)$ bits to Bob. Since $|V_A| \le n$, $|V_B| \le n$ and $|C_A| + |C_B| = |V(C)|$, we have $|C_B| |V_B| \log n + |C_A| |V_A| \log n \le (|C_A| + |C_B|) n \log n = |V(C)| n \log n$, and the players exchange a total of $O(|V(C)| n \log n)$ bits.

Correctness. By construction, for every edge $(u, v) \in C_B \times C_B$ in G'_A with weight $\operatorname{wd}_{G'_A}(u, v)$, there is a corresponding shortest path $P_{u,v}$ of the same weight in G_B . Hence, for any path $P' = (v_0, v_1, \ldots, v_k)$ in G'_A between $v_0, v_k \in V'_A$, there is a corresponding path P_{v_0,v_k} of the same weight in G, where P is obtained from P' by replacing every two consecutive nodes v_i, v_{i+1} in $P \cap C_B$ by the path $P_{v_i,v_{i+1}}$ in G_B . Thus, $\operatorname{wd}_{G'_A}(v_0, v_k) \geq \operatorname{wd}_G(v_0, v_k)$.

On the other hand, for any shortest path $P = (v_0, v_1, \ldots, v_k)$ in G connecting $v_0, v_k \in V'_A$, there is a corresponding path P' of the same weight in G'_A , where P' is obtained from P by replacing any sub-path (v_i, \ldots, v_j) of P contained in G_B and connecting $v_i, v_j \in C_B$ by the edge (v_i, v_j) in G'_A . Thus, wd_G $(v_0, v_k) \ge$ wd_{G'_A} (v_0, v_k) . Alice thus correctly computes the weighted distances between pairs of nodes in V'_A .

It remains to argue about the weighted distances that Alice computes to nodes in $V_B \setminus C_B$. Any shortest path *P* in *G* connecting a node $u \in V'_A$ and a node $v \in V_B \setminus C_B$ must cross at least one edge of *C* and thus must contain a node in C_B . Therefore, $wd_G(u, v) = min\{wd_G(u, x) + wd_G(x, v) \mid x \in C_B\}$. Recall that we have shown that $wd_{G'_A}(u, x) = wd_G(u, x)$ for any $u, x \in V'_A$. The subpath of *P* connecting *x* and *v* is a shortest path between these nodes, and is contained in G_B , so $wd_{G_B}(x, v) = wd_G(x, v)$. Hence, the distance $min\{wd_{G'_A}(u, x) + wd_{G_B}(x, v) \mid x \in C_B\}$ returned by Alice is indeed equal to $wd_G(u, v)$.

The outputs of Bob are correct by an analogous arguments, completing the proof.

Proof of Theorem 13: Let $f : \{0, 1\}^{K_1} \times \{0, 1\}^{K_2} \rightarrow \{0, 1\}^{L_1} \times \{0, 1\}^{L_2}$ be a function and let $G_{x,y}$ be a family of lower bound graphs w.r.t. f and the weighted APSP problem. By Lemma 7, Alice and Bob can compute the weighted distances for any graph in $G_{x,y}$ while exchanging at most

1:33

 $O(|V(C)|n \log n)$ bits, which is in $O(|C|n \log n)$ bits. Since $G_{x,y}$ is a family of lower bound graphs w.r.t. f and weighted APSP, item 3 in the definition of lower bound graphs implies that they can use the solution of the APSP problem to compute f without further communication, implying $CC(f) = O(|C| n \log n)$. Therefore, when applying Theorem 11 to f and $G_{x,y}$, the lower bound obtained for the number of rounds for computing weighted APSP is $\Omega(CC(f)/|C| \log n)$, which is no higher than a bound of $\Omega(n)$.

Extending to t players: We argue that generalizing the Alice-Bob framework to a shared-blackboard multi-party setting is still insufficient for providing a super-linear lower bound for weighted APSP. Suppose that we increase the number of players in the above framework to t players, P_0, \ldots, P_{t-1} , each simulating the nodes in a set V_i in a partition of V in a family of lower bound graphs w.r.t. a t-party function f and weighted APSP. That is, the outputs of nodes in V_i for an algorithm ALG for solving a problem P in the CONGEST model, uniquely determines the output of player P_i in the function f. The function f is of the form $f : \{0, 1\}^{K_0} \times \cdots \times \{0, 1\}^{K_{t-1}} \rightarrow \{0, 1\}^{L_0} \times \cdots \times \{0, 1\}^{L_{t-1}}$.

The communication complexity CC(f) is the total number of bits written on the shared blackboard by all players. Denote by *C* the set of cut edges, that is, the edge whose endpoints do not belong to the same set V_i . Then, if *ALG* is an *R*-round algorithm, we have that writing $O(R|C|\log n)$ bits on the shared blackboard suffice for computing f, and so $R = \Omega(CC(f)/|C|\log n)$.

Consider the weighted APSP problem. Let f be a t-party function and let $G_{x_0,...,x_{t-1}}$ be a family of lower bound graphs w.r.t. f and weighted APSP. The players first write all the edges in C on the shared blackboard, for a total of $O(|C| \log n)$ bits. Then, each player P_i writes the weighted distances from all nodes in V_i to all nodes in $V(C) \cap V_i$. This requires no more than $O(|V(C)|n \log n)$ bits.

It is easy to verify that every player P_i can now compute the weighted distances from all nodes in V_i to all nodes in V, in a manner that is similar to that of Lemma 7.

This gives an upper bound on CC(f), i.e. $CC(f) = O(|V(C)| n \log n)$. A lower bound obtained by a reduction from f is $\Omega(CC(f)/|C| \log n)$, which is at most $\Omega(|V(C)|n \log n/(|C| \log n))$. But $|V(C)| \le 2|C|$, so the lower bound cannot actually be larger than $\Omega(n)$, as claimed.

Remark 1: Notice that the *t*-party simulation of the algorithm for the CONGEST model does not require a shared blackboard and can be done in the peer-to-peer multiparty setting as well, since simulating the delivery of a message does not require the message to be known globally. This raises the question of why would one consider a reduction to the CONGEST model from the stronger shared-blackboard model to begin with. Notice that our argument for *t* players does not translate to the peer-to-peer multiparty setting, because it assumes that the edges of the cut *C* can be made global knowledge within writing $|C| \log n$ bits on the blackboard. However, what our extension above shows is that if there is a lower bound that is to be obtained using a reduction from peer-to-peer *t*-party computation, *it must use a function f that is strictly harder to compute in the peer-to-peer setting than in the shared-blackboard setting*.

Remark 2: We suspect that a similar argument can be applied for the framework of non-fixed Alice-Bob partitions (e.g., [99]), but this requires precisely defining this framework, which we do not address in this paper.

8 STREAMING LOWER BOUNDS

In this section we show how our super-linear lower bound constructions can be used to prove lower bounds in the streaming model in a straightforward manner. The crux of our proofs is that our constructions use instances for the disjointness problem that have very large inputs; the cut sizes do not play a role here. In particular, the larger the input is, the higher the space lower bound is. We start by defining a family of lower bound graphs for the semi-streaming model.

DEFINITION 4. (Family of Streaming Lower Bound Graphs)

Fix an integer K, a function $f : \{0, 1\}^K \times \{0, 1\}^K \rightarrow \{\text{TRUE, FALSE}\}$ and a graph predicate P. A family of graphs $\{G_{x,y} = (V, E_{x,y}) \mid x, y \in \{0, 1\}^K\}$ with a partition of the edges $E = E_A \cup E_B$ is said to be a family of streaming lower bound graphs w.r.t. f and P if the following properties hold:

- (1) Only the existence of edges in E_A may depend on x;
- (2) Only the existence of edges in E_B may depend on y;
- (3) $G_{x,y}$ satisfies the predicate P iff f(x,y) = TRUE.

This definition is a variant of Definition 1. Unsurprisingly, the existence of such a family implies a lower bound in a way similar to Theorem 6, as stated next. Note that for the semi-streaming model, the cut size does not play a role.

THEOREM 14. Fix a function $f : \{0, 1\}^K \times \{0, 1\}^K \to \{\text{TRUE, FALSE}\}$ and a predicate P. If there is a family $\{G_{x,y}\}$ of streaming lower bound graphs w.r.t. f and P then any semi-streaming algorithm for deciding P in R passes and M bits of memory requires $RM = \Omega(CC(f))$ rounds, and any randomized semi-streaming algorithm for deciding P in R passes and M bits of memory requires $RM = \Omega(CC^R(f))$ rounds.

PROOF. Let *ALG* be a semi streaming algorithm for deciding *P* in *R* passes and *M* bits of memory. Given inputs $x, y \in \{0, 1\}^K$ to Alice and Bob, respectively, Alice and Bob simulate the execution of the algorithms, as follows. To simulate a pass of the algorithm, Alice executes the algorithm using the edges of E_A as input—she can do so since these edges depend only on *x*. Alice then sends the current state of the memory to Bob, who continues the execution of *ALG* from the point where Alice stopped, using the edges of E_B as input—he can do so since these edges depend only on *y*. This concludes a simulation of a single pass; if there are more passes left, Bob sends the current memory state to Alice, who continues the execution from this state. Otherwise, Bob knows the value of *P*, and by Item 3 in Definition 4, he can compute f(x, y).

To simulate a pass, Alice sent to Bob at most M bits, and in all passes but the last, Bob sent Alice another M memory bits at the most. This sums to R(2M - 1) bits of communication, using which the players have computed f(x, y) correctly. The lower bounds follows directly from the lower bounds for CC(f) and $CC^{R}(f)$.

Any family of lower bound graphs for the CONGEST model is also a family of streaming lower bound graphs: use exactly the same graphs, and set E_A to be the edges in $V_A \times V_A$ and in C, and E_B be the edges in $V_B \times V_B$. This implies the following, simple corollary.

COROLLARY 1. Fix a function $f : \{0,1\}^K \times \{0,1\}^K \to \{\text{TRUE, FALSE}\}$ and a predicate P. If there is a family $\{G_{x,y}\}$ of lower bound graphs for the CONGEST model w.r.t. f and P then any semi-streaming algorithm for deciding P using R passes and M bits of memory requires $RM = \Omega(CC(f))$ rounds, and any randomized semi-streaming algorithm for deciding P using R passes and M bits of memory requires $RM = \Omega(CC^R(f))$ rounds.

With this corollary in hand, the lower bounds from Section 5 easily extend to a series of lower bounds in the semi-streaming model.

THEOREM 15. Any algorithm in the semi-streaming model for the following problems that uses R passes and M bits of memory requires $RM = \Omega(n^2)$.

(1) Computing a minimum vertex cover or deciding whether there is a minimum vertex cover of a given size.

- (2) Computing a maximum independent set or deciding whether there is an independent set of a given size.
- (3) Computing a maximum clique or deciding whether there is a clique of a given size.
- (4) Computing a coloring of a graph with a minimal number of colors or deciding whether there is a coloring with a given number of colors.
- (5) Deciding if a graph contains an 8-cycle of a given weight.
- (6) Deterministically deciding the identical subgraphs detection problem.

In addition, any such algorithm for deciding whether a graph is c-colorable, for an input parameter $3 \le c < n$ that may depend on n, requires $RM = \Omega((n-c)^2)$.

PROOF. Theorem 15 is obtained from Corollary 1 and the constructions from Sections 5 and 6.

Item 1 follows from the lower bound graph of Section 5.1, as proven in Lemma 4. Item 2 follows from the same lemma, as the complement of a minimum vertex cover is a maximum independent set.

For Item 3, consider the complement graph of the graph from Lemma 4: a graph with the same node set but *complement edges*. An independent set translates into a clique in the complement graph, each non-edge is either fixed, depends solely on Alice's input, or depends solely on Bob's input, and so the complement graph is a family of lower bound graph w.r.t. DISJ and maximum clique. Theorem 14 completes the proof.

Item 4 follows from the construction in Section 5.2, and specifically from Lemma 5. Item 5 follows from the construction in Section 6.1, and specifically from Lemma 6. Item 6 follows from the construction in Section 6.2, and specifically from the proof of Theorem 9. In all these cases, the communication complexity problem is $DISJ_K$, where $K \in \Theta(n^2)$.

The *c*-coloring lower bound follows the construction in the proof of Claim 9, where $K \in \Theta((n - c)^2)$.

The lower bound for finding a graph's chromatic number χ also applies for $(4/3-\varepsilon)$ -approximation of χ , in a weak sense: since χ is integral, any $(4/3 - \varepsilon)$ -approximation algorithm must return the exact solution in the case $\chi = 3$, so the impossibility for distinguishing the cases $\chi \leq 3$ and $\chi \geq 4$ result also rules out a $(4/3 - \varepsilon)$ -approximation algorithm. However, it does not rule out such an algorithm that works only for graph with $\chi \geq 5$, and may distinguish, e.g., $\chi \leq 6$ from $\chi \geq 8$.

We now show a stronger result for the semi-streaming model: it is hard to get $(4/3 - \varepsilon)$ -approximation of χ even when a small additive error is allowed, and even if the algorithm should work only for graphs with *high chromatic number*. To this end, we show a lower bound for distinguishing between $\chi \leq 3c$ and $\chi \geq 4c$, for any integer *c*.

THEOREM 16. For any positive integer c, any algorithm in the semi-streaming model that distinguishes a graph G with $\chi(G) \leq 3c$ from a graph with $\chi(G) \geq 4c$ and uses R passes and M bits of memory, requires $RM = \Omega((n/c)^2)$.

PROOF. We show a family of lower bound graphs with respect to the DISJ_K function, where $K \in \Theta((n/c)^2)$, and the predicate $\chi \ge 4c$ (TRUE) or $\chi \le 3c$ (FALSE). The predicate is not defined for other values of χ .

We create a graph $G_{x,y}^c$, composed of c copies of $G_{x,y}$, the original graph from the proof of Theorem 5, each on n/c nodes. The *i*-th copy is denoted $G_{x,y}(i)$, and its nodes are partitioned into $V_A(i)$ and $V_B(i)$. Naturally, let $V_A = \bigcup_i V_A(i)$ and $V_B = \bigcup_i V_B(i)$. For every $i \neq j \in [c]$, we connect $G_{x,y}(i)$ and $G_{x,y}(j)$ by a complete bipartite graph, i.e., we add an edge from each node in $G_{x,y}(i)$ to each node in $G_{x,y}(j)$.

This construction immediately guarantees that each copy is colored using unique colors, that are not used in any other copy. By Theorem 5, if x and y are disjoint then each copy requires at

least 4 colors, and $\chi(G_{x,y}^c) \ge 4c$; otherwise, each copy is 3-colorable, and $\chi(G_{x,y}^c) = 3c$. Therefore, $G_{x,y}^c$ is a family of lower bound graphs.

The boosting reduction in the proof of Theorem 16 does not apply to the CONGEST model: the cut size is $\Theta(n^2)$, which only yields a constant lower bound in this model. It is natural to try to modify it by only connecting the nodes on each side of the cut by full bipartite graphs, i.e., connect each node in $V_A(i)$ and each node $V_A(j)$ with $i \neq j$, and similarly for $V_B(i)$ and $V_B(j)$. However, this does not guarantee that each copy is colored with a unique set of colors⁴ because a fourth color from $V_A(i)$ might also be used in $V_B(j)$ for $i \neq j$. The loss from this boosting reduction in the CONGEST model prevents us from getting a 3c vs. 4c lower bound, unless we start with a lower bound of 3 vs. k with a sufficiently large k > 4. It may be possible to prove such a claim, as described below.

A remark on approximating χ with larger factors. A beautiful construction of Guruswami and Khanna [57] proves the NP-hardness of 3 vs. 5 coloring (in the sequential setting) by giving a reduction from MaxIS. Their construction can be combined with our lower bounds for MaxIS to get a lower bound for $\chi \leq 3$ vs. $\chi \geq 5$ in the streaming setting, but it does not directly work in the CONGEST model—some special nodes in their construction (the "ground vertices") must be connected to nearly all other nodes, causing the cut size to become $\Omega(n)$. This issue can be alleviated by using two sets of ground vertices, one for Alice and one for Bob, and by modifying their "edge gadgets" in a simple but non-trivial way. We have chosen to omit this more involved construction because we believe a much higher lower bound is probably attainable by a different approach. In particular, in the sequential setting it is known that 3 vs. k for all constant k is NP-hard under variants of the Unique Games Conjecture [42]; whether such hardness can be transferred into the conGEST and streaming models is an interesting open question.

9 DISCUSSION

We introduced the bit-gadget, a powerful tool for constructing graphs with small cuts. Using the bit-gadget, we were able to prove new lower bounds for the CONGEST model for fundamental graph problems, such as computing the exact or approximate diameter, radius, minimum vertex cover, and the chromatic number of a graph.

Our lower bound for computing the radius answers an open question that was raised by Holzer and Wattenhofer[66]. Notably, our lower bound for computing the diameter implies a large gap between the complexity of computing a (3/2)-approximation, which can be done in $\widetilde{O}(\sqrt{n})$ rounds[64], and the complexity of computing a $(3/2 - \epsilon)$ -approximation, which we show to require $\widetilde{\Omega}(n)$ rounds. As there are no known lower bounds for computing a (3/2)-approximation, an intriguing open question that immediately arises is the complexity of (3/2)-approximation.

Furthermore, our bit-gadget allows us to show the first super-linear lower bounds for the CONGEST model, raising a plethora of open questions. First, we showed for some specific problems, namely, computing a minimum vertex cover, a maximum independent set and a χ -coloring, that they are nearly as hard as possible for the CONGEST model. However, we know that approximate solutions for some of these problems can be obtained much faster, in a polylogarithmic number of rounds or even less. For vertex cover, our lower bound can easily be amplified to any constant additive approximation⁵, but this also leaves huge gaps in our understanding of the trade-offs between approximation factors and efficiency for this problem. Thus, a family of specific open questions

 $^{^4}$ We have included such a wrong claim in a previous version of this work, and thank an anonymous reviewer of ACM TALG for spotting this error.

⁵We thank David Wajc for pointing this out.

is to characterize the exact trade-offs between approximation factors and round complexities for various optimization problems.

Finally, we propose a more general open question which addresses a possible classification of complexities of global problems in the CONGEST model. Some such problems have complexities of $\Theta(D)$, such as constructing a BFS tree. Others have complexities of $\tilde{\Theta}(D + \sqrt{n})$, such as finding an MST. Some problems have near-linear complexities, such as unweighted APSP. And now we know about the family of hardest problems for the CONGEST model, whose complexities are near-quadratic. Do these complexities capture all possibilities, as far as natural global graph problems are concerned? Or are there such problems with a complexity of, say, $\Theta(n^{1+\delta})$, for some constant $0 < \delta < 1$? A similar question was recently addressed in [34] for LCL problems the LOCAL model, and we propose investigating the possibility that such a hierarchy exists for the CONGEST model for certain classes of problems.

ACKNOWLEDGMENTS

We are grateful to Yuval Emek for many discussions and fruitful comments about the connection between communication complexity and streaming lower bounds, and for pointing out that our technique can be useful for streaming models. We also thank Ohad Ben Baruch, Michael Elkin, Yuval Filmus, Christoph Lenzen, Merav Parter and David Wajc for useful discussions, and the reviewers of ACM TALG journal for their thorough reading and the helpful comments and references.

This research has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 755839. This research is also supported in part by the Israel Science Foundation (grant 1696/14). We acknowledge the Austrian Science Fund (FWF) and netIDEE SCIENCE project P 33775-N.

REFERENCES

- Amir Abboud, Keren Censor-Hillel, and Seri Khoury. Near-linear lower bounds for distributed distance computations, even in sparse networks. In 30th International Symposium on Distributed Computing, DISC, pages 29–42, 2016.
- [2] Amir Abboud, Vincent Cohen-Addad, and Philip N Klein. New hardness results for planar graph problems in p and an algorithm for sparsest cut. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC, pages 996–1009, 2020.
- [3] Amir Abboud, Fabrizio Grandoni, and Virginia Vassilevska Williams. Subcubic equivalences between graph centrality problems, APSP and diameter. In ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 1681–1697, 2015.
- [4] Amir Abboud, Virginia Vassilevska Williams, and Joshua R. Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 377–391, 2016.
- [5] Pierre Aboulker, Marthe Bonamy, Nicolas Bousquet, and Louis Esperet. Distributed coloring in sparse graphs with fewer colors. In ACM Symposium on Principles of Distributed Computing, PODC, pages 419–425, 2018.
- [6] Udit Agarwal and Vijaya Ramachandran. A faster deterministic distributed algorithm for weighted apsp through pipelining. *CoRR*, abs/1804.05441, 2018.
- [7] Udit Agarwal, Vijaya Ramachandran, Valerie King, and Matteo Pontecorvi. A deterministic distributed algorithm for exact weighted all-pairs shortest paths in $\tilde{O}(n^{3/2})$ rounds. In Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC, pages 199–205, 2018.
- [8] Kook Jin Ahn and Sudipto Guha. Graph sparsification in the semi-streaming model. In Automata, Languages and Programming, 36th International Colloquium, ICALP, pages 328–338, 2009.
- [9] Matti Åstrand, Patrik Floréen, Valentin Polishchuk, Joel Rybicki, Jukka Suomela, and Jara Uitto. A local 2approximation algorithm for the vertex cover problem. In Proceedings of the 23rd International Symposium on Distributed Computing, DISC, pages 191–205, 2009.
- [10] Matti Åstrand and Jukka Suomela. Fast distributed approximation algorithms for vertex cover and set cover in anonymous networks. In Proceedings of the 22nd Annual ACM Symposium on Parallelism in Algorithms and Architectures, SPAA, pages 294–302, 2010.
- [11] Nir Bachrach, Keren Censor-Hillel, Michal Dory, Yuval Efron, Dean Leitersdorf, and Ami Paz. Hardness of distributed optimization. In PODC, pages 238–247. ACM, 2019.

- [12] Reuven Bar-Yehuda, Keren Censor-Hillel, Mohsen Ghaffari, and Gregory Schwartzman. Distributed approximation of maximum independent set and maximum matching. In Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC, pages 165–174, 2017.
- [13] Reuven Bar-Yehuda, Keren Censor-Hillel, and Gregory Schwartzman. A distributed (2+ε)-approximation for vertex cover in o(logδ/ε log log δ) rounds. In Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC, pages 3–8, 2016.
- [14] Ziv Bar-Yossef, Ravi Kumar, and D. Sivakumar. Reductions in streaming algorithms, with an application to counting triangles in graphs. In Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 623–632, 2002.
- [15] Leonid Barenboim. On the locality of some NP-complete problems. In Proceedings of the 39th International Colloquium on Automata, Languages, and Programming, ICALP, pages 403–415, 2012.
- [16] Leonid Barenboim. Deterministic (Δ + 1)-coloring in sublinear (in Δ) time in static, dynamic, and faulty networks. *J. ACM*, 63(5):47:1–47:22, 2016.
- [17] Leonid Barenboim and Michael Elkin. Deterministic distributed vertex coloring in polylogarithmic time. J. ACM, 58(5):23:1–23:25, 2011.
- [18] Leonid Barenboim and Michael Elkin. Combinatorial algorithms for distributed graph coloring. Distributed Computing, 27(2):79–93, 2014.
- [19] Leonid Barenboim, Michael Elkin, and Fabian Kuhn. Distributed (delta+1)-coloring in linear (in delta) time. SIAM J. Comput., 43(1):72–95, 2014.
- [20] Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. The locality of distributed symmetry breaking. J. ACM, 63(3):20:1–20:45, 2016.
- [21] Surender Baswana. Streaming algorithm for graph spanners single pass and constant processing time per edge. Inf. Process. Lett., 106(3):110–114, 2008.
- [22] Ruben Becker, Andreas Karrenbauer, Sebastian Krinninger, and Christoph Lenzen. Near-optimal approximate shortest paths and transshipment in distributed and streaming models. In Andréa W. Richa, editor, 31st International Symposium on Distributed Computing, DISC, volume 91 of LIPIcs, pages 7:1–7:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.DISC.2017.7.
- [23] Suman Kalyan Bera and Prantar Ghosh. Coloring in graph streams. CoRR, abs/1807.07640, 2018.
- [24] Aaron Bernstein and Danupon Nanongkai. Distributed exact weighted all-pairs shortest paths in near-linear time. In Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC, pages 334–342, 2019. doi:10.1145/3313276.3316326.
- [25] Marijke H. L. Bodlaender, Magnús M. Halldórsson, Christian Konrad, and Fabian Kuhn. Brief announcement: Local independent set approximation. In Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC, pages 93–95, 2016.
- [26] Ilaria Bordino, Debora Donato, Aristides Gionis, and Stefano Leonardi. Mining large networks with subgraph counting. In Proceedings of the 8th IEEE International Conference on Data Mining ICDM, pages 737–742, 2008.
- [27] Vladimir Braverman, Zaoxing Liu, Tejasvam Singh, N. V. Vinodchandran, and Lin F. Yang. New bounds for the CLIQUE-GAP problem using graph decomposition theory. *Algorithmica*, 80(2):652–667, 2018.
- [28] Karl Bringmann and Sebastian Krinninger. A note on hardness of diameter approximation. Inf. Process. Lett., 133:10–15, 2018. doi:10.1016/j.ipl.2017.12.010.
- [29] Luciana S. Buriol, Gereon Frahling, Stefano Leonardi, and Christian Sohler. Estimating clustering indexes in data streams. In Algorithms - ESA 2007, 15th Annual European Symposium, pages 618–632, 2007.
- [30] Massimo Cairo, Roberto Grossi, and Romeo Rizzi. New bounds for approximating extremal distances in undirected graphs. In Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 363–376, 2016.
- [31] Keren Censor-Hillel, Telikepalli Kavitha, Ami Paz, and Amir Yehudayoff. Distributed construction of purely additive spanners. Distributed Computing, 31(3):223–240, 2018.
- [32] Keren Censor-Hillel, Seri Khoury, and Ami Paz. Quadratic and near-quadratic lower bounds for the CONGEST model. In 31st International Symposium on Distributed Computing, DISC 2017, pages 10:1–10:16, 2017.
- [33] Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. An exponential separation between randomized and deterministic complexity in the LOCAL model. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS*, pages 615–624, 2016.
- [34] Yi-Jun Chang and Seth Pettie. A time hierarchy theorem for the LOCAL model. CoRR, abs/1704.06297, 2017.
- [35] Shiri Chechik, Daniel H. Larkin, Liam Roditty, Grant Schoenebeck, Robert Endre Tarjan, and Virginia Vassilevska Williams. Better approximation algorithms for the graph diameter. In Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 1041–1052, 2014.

- [36] Shiri Chechik and Doron Mukhtar. Optimal distributed coloring algorithms for planar graphs in the LOCAL model. CoRR, abs/1804.00137, 2018.
- [37] Rajesh Hemant Chitnis, Graham Cormode, Mohammad Taghi Hajiaghayi, and Morteza Monemizadeh. Parameterized streaming: Maximal matching and vertex cover. In Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 1234–1251, 2015.
- [38] Kai-Min Chung, Seth Pettie, and Hsin-Hao Su. Distributed algorithms for the lovász local lemma and graph coloring. In Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC, pages 134–143, 2014.
- [39] Richard Cole and Uzi Vishkin. Deterministic coin tossing with applications to optimal parallel list ranking. Information and Control, 70(1):32–53, 1986.
- [40] Graham Cormode, Jacques Dark, and Christian Konrad. Approximating the caro-wei bound for independent sets in graph streams. In Combinatorial Optimization - 5th International Symposium, ISCO, pages 101–114, 2018.
- [41] Andrzej Czygrinow, Michal Hanckowiak, and Wojciech Wawrzyniak. Fast distributed approximations in planar graphs. In Proceedings of the 22nd International Symposium on Distributed Computing, DISC, pages 78–92, 2008.
- [42] Irit Dinur and Igor Shinkar. On the conditional hardness of coloring a 4-colorable graph with super-constant number of colors. In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, pages 138–151. Springer, 2010.
- [43] Andrew Drucker, Fabian Kuhn, and Rotem Oshman. On the power of the congested clique model. In Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing, PODC, pages 367–376, 2014.
- [44] Michael Elkin. An unconditional lower bound on the time-approximation trade-off for the distributed minimum spanning tree problem. *SIAM J. Comput.*, 36(2):433–456, 2006.
- [45] Michael Elkin. Streaming and fully dynamic centralized algorithms for constructing and maintaining sparse spanners. ACM Trans. Algorithms, 7(2):20:1–20:17, 2011.
- [46] Michael Elkin. Distributed exact shortest paths in sublinear time. In 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC, pages 757–770, 2017.
- [47] Michael Elkin and Jian Zhang. Efficient algorithms for constructing (1 + ε, β)-spanners in the distributed and streaming models. In Proceedings of the Twenty-Third Annual ACM Symposium on Principles of Distributed Computing, PODC, pages 160–168, 2004. doi:10.1145/1011767.1011791.
- [48] Yuval Emek, Christoph Pfister, Jochen Seidel, and Roger Wattenhofer. Anonymous networks: randomization = 2-hop coloring. In Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC, pages 96–105, 2014.
- [49] Joan Feigenbaum, Sampath Kannan, Andrew McGregor, Siddharth Suri, and Jian Zhang. On graph problems in a semi-streaming model. *Theor. Comput. Sci.*, 348(2-3):207–216, 2005.
- [50] Joan Feigenbaum, Sampath Kannan, Andrew McGregor, Siddharth Suri, and Jian Zhang. Graph distances in the data-stream model. SIAM J. Comput., 38(5):1709–1727, 2008. doi:10.1137/070683155.
- [51] Pierre Fraigniaud, Cyril Gavoille, David Ilcinkas, and Andrzej Pelc. Distributed computing with advice: information sensitivity of graph coloring. *Distributed Computing*, 21(6):395–403, 2009.
- [52] Pierre Fraigniaud, Marc Heinrich, and Adrian Kosowski. Local conflict coloring. In Proceedings of the IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS, pages 625–634, 2016.
- [53] Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1150–1162, 2012.
- [54] Fabrizio Grandoni, Jochen Könemann, and Alessandro Panconesi. Distributed weighted vertex cover via maximal matchings. ACM Trans. Algorithms, 5(1):6:1–6:12, 2008.
- [55] Fabrizio Grandoni, Jochen Könemann, Alessandro Panconesi, and Mauro Sozio. A primal-dual bicriteria distributed algorithm for capacitated vertex cover. SIAM J. Comput., 38(3):825–840, 2008.
- [56] Ofer Grossman, Seri Khoury, and Ami Paz. Improved hardness of approximation of diameter in the CONGEST model. In DISC, volume 179 of LIPIcs, pages 19:1–19:16, 2020.
- [57] Venkatesan Guruswami and Sanjeev Khanna. On the hardness of 4-coloring a 3-colorable graph. SIAM J. Discret. Math., 18(1):30–40, 2004.
- [58] Bjarni V. Halldórsson, Magnús M. Halldórsson, Elena Losievskaja, and Mario Szegedy. Streaming algorithms for independent sets in sparse hypergraphs. *Algorithmica*, 76(2):490–501, 2016.
- [59] Magnús M. Halldórsson, Xiaoming Sun, Mario Szegedy, and Chengu Wang. Streaming and communication complexity of clique approximation. In Automata, Languages, and Programming - 39th International Colloquium, ICALP, pages 449–460, 2012.
- [60] Michal Hanckowiak, Michal Karonski, and Alessandro Panconesi. On the distributed complexity of computing maximal matchings. SIAM J. Discrete Math., 15(1):41–57, 2001.
- [61] David G. Harris, Johannes Schneider, and Hsin-Hao Su. Distributed (Δ+1)-coloring in sublogarithmic rounds. In Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC, pages 465–478, 2016.

- [62] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. A deterministic almost-tight distributed algorithm for approximating single-source shortest paths. In Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC, pages 489–498, 2016.
- [63] Monika Rauch Henzinger, Prabhakar Raghavan, and Sridhar Rajagopalan. Computing on data streams. In External Memory Algorithms, pages 107–118, 1998.
- [64] Stephan Holzer, David Peleg, Liam Roditty, and Roger Wattenhofer. Distributed 3/2-approximation of the diameter. In Proceedings of the 28th International Symposium on Distributed Computing, DISC, pages 562–564, 2014.
- [65] Stephan Holzer and Nathan Pinsker. Approximation of distances and shortest paths in the broadcast congest clique. In Proceedings of the 19th International Conference on Principles of Distributed Systems, OPODIS, pages 6:1–6:16, 2015.
- [66] Stephan Holzer and Roger Wattenhofer. Optimal distributed all pairs shortest paths and applications. In Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC, pages 355–364, 2012.
- [67] Qiang-Sheng Hua, Haoqiang Fan, Lixiang Qian, Ming Ai, Yangyang Li, Xuanhua Shi, and Hai Jin. Brief announcement: A tight distributed algorithm for all pairs shortest paths and applications. In Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA, pages 439–441, 2016.
- [68] Chien-Chung Huang, Danupon Nanongkai, and Thatchaphol Saranurak. Distributed exact weighted all-pairs shortest paths in $\tilde{O}(n^{5/4})$ rounds. In 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS, pages 168–179, 2017.
- [69] Michael Kapralov. Better bounds for matchings in the streaming model. In ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 1679–1697, 2013.
- [70] Michael Kapralov, Sanjeev Khanna, and Madhu Sudan. Streaming lower bounds for approximating MAX-CUT. In ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 1263–1282, 2015.
- [71] Jonathan A. Kelner and Alex Levin. Spectral sparsification in the semi-streaming setting. Theory Comput. Syst., 53(2):243–262, 2013.
- [72] Samir Khuller, Uzi Vishkin, and Neal E. Young. A primal-dual parallel approximation technique applied to weighted set and vertex covers. J. Algorithms, 17(2):280–289, 1994.
- [73] Christos Koufogiannakis and Neal E. Young. Distributed and parallel algorithms for weighted vertex cover and other covering problems. In Proceedings of the 28th Annual ACM Symposium on Principles of Distributed Computing, PODC, pages 171–179, 2009.
- [74] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. The price of being near-sighted. In Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 980–989, 2006.
- [75] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Local computation: Lower and upper bounds. J. ACM, 63(2):17:1–17:44, 2016.
- [76] Eyal Kushilevitz and Noam Nisan. Communication Complexity. Cambridge University Press, New York, NY, USA, 1997.
- [77] Christoph Lenzen. Theory of distributed systems, lecture 8: Distance approximation and routing, 2019. URL: https://www.mpi-inf.mpg.de/fileadmin/inf/d1/teaching/winter19/tods/ToDS_08.pdf. Last visited 09/11/2020.
- [78] Christoph Lenzen and Boaz Patt-Shamir. Fast partial distance estimation and applications. In Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC, pages 153–162, 2015.
- [79] Christoph Lenzen and David Peleg. Efficient distributed source detection with limited bandwidth. In Proceedings of the ACM Symposium on Principles of Distributed Computing, PODC, pages 375–382, 2013.
- [80] Christoph Lenzen and Roger Wattenhofer. Leveraging linial's locality limit. In Proceedings of the 22nd International Symposium on Distributed Computing, DISC, pages 394–407, 2008.
- [81] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. http://snap.stanford.edu/ data, June 2014.
- [82] Jason Li and Merav Parter. Planar diameter via metric compression. In Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC, pages 152–163, 2019.
- [83] Nathan Linial. Locality in distributed graph algorithms. SIAM J. Comput., 21(1):193-201, 1992.
- [84] Madhusudan Manjunath, Kurt Mehlhorn, Konstantinos Panagiotou, and He Sun. Approximate counting of cycles in streams. In Algorithms - ESA 2011 - 19th Annual European Symposium, pages 677–688, 2011.
- [85] Robert Meusel, Sebastiano Vigna, Oliver Lehmberg, and Christian Bizer. The graph structure in the web analyzed on different aggregation levels. J. Web Science, 1(1):33–47, 2015.
- [86] Thomas Moscibroda and Roger Wattenhofer. Coloring unstructured radio networks. Distributed Computing, 21(4):271– 284, 2008.
- [87] S. Muthukrishnan. Data streams: Algorithms and applications. Foundations and Trends in Theoretical Computer Science, 1(2), 2005.
- [88] Danupon Nanongkai. Distributed approximation algorithms for weighted shortest paths. In Proceedings of the Symposium on Theory of Computing, STOC,, pages 565–573, 2014.

ACM Trans. Algor., Vol. 1, No. 1, Article 1. Publication date: January 2021.

- [89] Danupon Nanongkai, Atish Das Sarma, and Gopal Pandurangan. A tight unconditional lower bound on distributed randomwalk computation. In Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC, pages 257–266, 2011.
- [90] Alessandro Panconesi and Romeo Rizzi. Some simple distributed algorithms for sparse networks. Distributed Computing, 14(2):97–100, 2001.
- [91] Ami Paz and Gregory Schwartzman. A $(2 + \epsilon)$ -approximation for maximum weight matching in the semi-streaming model. In *ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 2153–2161, 2017.
- [92] David Peleg. Distributed Computing: A Locality-Sensitive Approach. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [93] David Peleg, Liam Roditty, and Elad Tal. Distributed algorithms for network diameter and girth. In Proceedings of the 39th International Colloquium on Automata, Languages, and Programming, ICALP, pages 660–672, 2012.
- [94] David Peleg and Vitaly 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.
- [95] Seth Pettie and Hsin-Hao Su. Distributed coloring algorithms for triangle-free graphs. Inf. Comput., 243:263–280, 2015.
- [96] Valentin Polishchuk and Jukka Suomela. A simple local 3-approximation algorithm for vertex cover. Inf. Process. Lett., 109(12):642–645, 2009.
- [97] Alexander A. Razborov. On the distributional complexity of disjointness. Theor. Comput. Sci., 106(2):385-390, 1992.
- [98] Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In Proceedings of the Symposium on Theory of Computing Conference, STOC, pages 515–524, 2013.
- [99] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. Distributed verification and hardness of distributed approximation. SIAM J. Comput., 41(5):1235–1265, 2012.
- [100] Johannes Schneider, Michael Elkin, and Roger Wattenhofer. Symmetry breaking depending on the chromatic number or the neighborhood growth. *Theor. Comput. Sci.*, 509:40–50, 2013. doi:10.1016/j.tcs.2012.09.004.
- [101] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In Proceedings of the 11h Annual ACM Symposium on Theory of Computing, STOC, pages 209–213, 1979.