Distributed Distance Approximation

Bertie Ancona
MIT, Cambridge, MA, USA
bancona@alum.mit.edu

Keren Censor-Hillel
Technion, Haifa, Israel
ckeren@cs.technion.ac.il

Mina Dalirrooyfard
MIT, Cambridge, MA, USA
minad@mit.edu

Yuval Efron
Technion, Haifa, Israel
efronyuv@gmail.com

Virginia Vassilevska Williams
MIT, Cambridge, MA, USA
virgi@mit.edu

Abstract

Diameter, radius and eccentricities are fundamental graph parameters, which are extensively studied in various computational settings. Typically, computing approximate answers can be much more efficient compared with computing exact solutions. In this paper, we give a near complete characterization of the trade-offs between approximation ratios and round complexity of distributed algorithms for approximating these parameters, with a focus on the weighted and directed variants.

Furthermore, we study *bi-chromatic* variants of these parameters defined on a graph whose vertices are colored either red or blue, and one focuses only on distances for pairs of vertices that are colored differently. Motivated by applications in computational geometry, bi-chromatic diameter, radius and eccentricities have been recently studied in the sequential setting [Backurs et al. STOC'18, Dalirrooyfard et al. ICALP'19]. We provide the first distributed upper and lower bounds for such problems.

Our technical contributions include introducing the notion of *approximate pseudo-center*, which extends the *pseudo-centers* of [Choudhary and Gold SODA'20], and presenting an efficient distributed algorithm for computing approximate pseudo-centers. On the lower bound side, our constructions introduce the usage of new functions into the framework of reductions from 2-party communication complexity to distributed algorithms.

2012 ACM Subject Classification Theory of computation

Keywords and phrases Distributed Computing, Distance Computation, Algorithms, Lower Bounds

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2020.30

Related Version A full version of this paper is available at `https://arxiv.org/abs/2011.05066`.

Funding This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme under grant agreement No 755839.

1 Introduction

The diameter and radius are central graph parameters, defined as the maximum and minimum eccentricities over all vertices, respectively, where the eccentricity of a vertex v is the maximum distance out of v. Computing the diameter and radius of a given graph are cornerstone

24th International Conference on Principles of Distributed Systems (OPODIS 2020).
Editors: Quentin Bramas, Rotem Oshman, and Paolo Romano; Article No. 30; pp. 30:1–30:17
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

problems with abundant applications. This is particularly the case in the context of distributed computing, where distances between nodes in a network (and in particular the graph diameter) directly influence the time it takes to communicate throughout the network.

We focus on computing the diameter, radius and eccentricities in the classic CONGEST model of distributed computation, in which n nodes of a synchronous network communicate by exchanging messages of $O(\log n)$ bits with their neighbors in the underlying network graph. In a seminal work, Frischknecht et al. [35] showed that the diameter is hard to compute in CONGEST, namely that $\tilde{\Omega}(n)$[1] rounds are required, even in undirected unweighted graphs. Abboud et al. [1] showed that the same holds for computing the radius. Both of these results are tight up to logarithmic factors due to algorithms that compute all pairs shortest paths (APSP) in a given unweighted, undirected graph in $O(n)$ rounds, see Lenzen and Peleg, and Peleg et al. [50, 53]. Recently, Bernstein and Nanongkai [14], presented an algorithm which computes exact APSP in a given weighted, directed graph in $\tilde{O}(n)$ rounds as well.

As computing the diameter and radius exactly in general graphs is hard, a natural relaxation is to settle for approximate computations. In an unweighted, undirected graph, a simple observation due to the triangle inequality is that computing a BFS tree from any node yields a 2-approximation to the diameter or radius, and a 3-approximation of all eccentricities.

Obtaining a more thorough understanding of the complexity landscape of computing approximations to these distance parameters has been an ongoing endeavour of the community. The current state of the art for diameter approximation is the algorithm by Holzer et al. [40] with round complexity of $O(\sqrt{n \log n} + D)$, that achieves a $\frac{3}{2}$-approximation of the diameter in a given unweighted, undirected graph (further discussion is deferred to Section 1.2).

However, many open cases have remained, and unveiling the full picture of the trade-offs between approximation ratios and round complexity for distance parameters in the CONGEST model has remained a central open problem. In this paper, we give a near-complete characterization of this trade-off for the problems of diameter, radius and eccentricities, focusing on the weighted and/or directed variants. For the problem of directed diameter, only the range $[\frac{3}{2}, 2]$ of approximation ratios remains open.

In some cases, originally motivated by computational geometry problems [4, 29, 46, 57], we are interested in a "bi-chromatic" definition of the parameters. In the bi-chromatic setting, the vertices are partitioned into two sets, S and $T = V \setminus S$, and the bi-chromatic eccentricity of a node $s \in S$ is the maximum distance from s to a node in T. The bi-chromatic diameter and radius are the maximum and minimum bi-chromatic eccentricities of nodes in S.

The bi-chromatic versions of diameter and radius have received much recent attention in the sequential setting [11, 24]. In this paper, we initiate the study of these problems in the CONGEST model, by providing upper and lower bounds for these problems. For example, we prove that a $\frac{5}{3}$-approximation to bi-chromatic diameter in an unweighted, undirected graph can be computed in $\tilde{O}(\sqrt{n} + D)$ rounds, and we prove this is tight in the sense that any improvement in the approximation ratio incurs a blowup in the round complexity to $\tilde{\Omega}(n)$.

A more comprehensive display of our results follows. Also, a comparison with previous work is depicted in Table 1 and Table 2 and is elaborated upon in Section 1.2.

1.1 Our contributions and techniques

As mentioned earlier, the *eccentricity* $ecc(v)$ of a vertex v is the distance $\max_{u \in V} d(v, u)$. The *diameter* D is the largest eccentricity in the graph, and the *radius* r is the smallest.

[1] Throughout the paper, $\tilde{O}$ and $\tilde{\Omega}$ are used to hide poly-logarithmic factors

Table 1 Upper bounds for the problems considered in this paper. A variant can be weighted, directed, both, or neither. Upper bounds hold for the listed variants and all subsets of those variants. Approximation factors are multiplicative but may omit additive error. The value k may be any integer greater or equal to 1. We denote the round complexity of the current best exact weighted SSSP algorithm by $T(SSSP)$, currently $\tilde{O}(\min\left\{\sqrt{nD}, \sqrt{n}D^{\frac{1}{4}} + n^{\frac{3}{5}}\right\} + D)$ by [34].

*for $k = 1$

Problem	Approx.	Variant	Upper Bound $\tilde{O}(\cdot)$	Reference
Diameter	Exact	wted dir	n	[14]
	$2 - \frac{1}{2^k}$		$n^{\frac{1}{k+1}} + D$	Theorem 9, [39]*
	2	wted dir	$T(SSSP)$	Corollary 4
	$2 + \epsilon$	wted	$\sqrt{n} + D$	[13]
		wted dir	$\sqrt{n}D^{1/4} + D$	Corollary 3
Radius	Exact	wted dir	n	[14]
	$2 - \frac{1}{2^k}$		$n^{\frac{1}{k+1}} + D$	Theorem 9
	2	wted dir	$T(SSSP)$	Corollary 4
	$2 + \epsilon$	wted	$\sqrt{n} + D$	Corollary 2
		wted dir	$\sqrt{n}D^{1/4} + D$	Corollary 3
Eccentricities	Exact	wted dir	n	[14]
	$3 - \frac{4}{2^k+1}$		$n^{\frac{1}{k+1}} + D$	Theorem 9
	2	wted dir	$T(SSSP)$	Corollary 4
	$2 + \epsilon$	wted	$\sqrt{n} + D$	Corollary 2
		wted dir	$\sqrt{n}D^{1/4} + D$	Corollary 3
Bi-chromatic Diameter	Exact	wted dir	n	[14]
	5/3		$\sqrt{n} + D$	Theorem 10
	2	wted	$T(SSSP)$	Theorem 11

Directed/weighted Radius and Eccentricities. We present a connection between the complexity of computing or approximating the Single Source Shortest Paths (SSSP) problem and the complexity of approximating radius, diameter and eccentricities. Formally, we prove the following theorem in Section 3.

▶ **Theorem 1.** *For any $\varepsilon \geq 0$, given a $(1 + \varepsilon)$-approximation algorithm $\mathcal{A}_\varepsilon$ for weighted and directed SSSP running in $T(n, \varepsilon, D)$ rounds, there exists an algorithm for $(2 + \varepsilon^3 + 3\varepsilon^2 + 4\varepsilon)$-approximate diameter, radius, and all eccentricities in $\tilde{O}(T(n, \varepsilon, D) + D)$ rounds on weighted, directed graphs.*

We now describe the challenges in proving the above and how we cope with them. A useful notion for distance parameters is the *center* of a graph, which is the vertex with the lowest eccentricity. Given the center c of a graph, we can easily approximate all eccentricities of a given graph by performing an SSSP algorithm rooted at c, and letting each node v estimate its eccentricity by outputting $d(v, c) + ecc(c)$. However, computing the center of a graph, or even its eccentricity (the radius), is a hard task that requires $\tilde{\Omega}(n)$ rounds [1].

For proving Theorem 1, we rely on an approach of Choudhary and Gold [22]. Here, one defines a notion of a *pseudo-center* and one then shows how to compute a pseudo-center of size $O(\log^2 n)$ sequentially in near-linear time. A pseudo-center C is a set of nodes, whose goal is to mimic the center of the graph, by promising that all eccentricities are at least the maximal distance between any node to the pseudo-center C. Using such a pseudo-center, one estimates the eccentricity of every node, similarly to the case of computing the actual center.

Table 2 Lower bounds for the problems considered in this paper. A variant can be weighted, directed, both, or neither. Lower bounds hold for the listed variants and all supersets of those variants. Approximation factors are multiplicative.

Problem	**Approx.**	**Variant**	**Lower Bound $\tilde{\Omega}(\cdot)$**	**Reference**
Diameter	$3/2-\varepsilon$		n	[1]
	$2-\varepsilon$	wted		[41]
	poly(n)	wted	$\sqrt{n}+D$	[49]
		dir		Theorem 8
Radius	$3/2-\varepsilon$		n	[1]
	$2-\varepsilon$	wted		Theorem 5
		dir		
	poly(n)	wted	$\sqrt{n}+D$	Corollary 6
		dir		
Eccentricities	$5/3-\varepsilon$		n	[1]
	$2-\varepsilon$	wted		[41]
		dir		Theorem 5
	poly(n)	wted	$\sqrt{n}+D$	Corollary 6
		dir		
Bi-chromatic Diameter	$5/3-\varepsilon$		n	Theorem 14
	$2-\varepsilon$	wted		[41]
		dir		Theorem 15
	poly(n)	wted	$\sqrt{n}+D$	Corollary 6
		dir		

The algorithm of [22] for computing a small pseudo-center can be viewed as a reduction to Single Source Shortest Paths (SSSP), which is very efficient in the sequential setting. However, the current state-of-the-art *distributed* complexity of computing exact SSSP is very costly, and hence we wish to avoid it. To overcome this, we introduce the notion of an *approximate pseudo-center*, which generalizes the notion of a pseudo-center. We prove that (i) an approximate pseudo-center of small size can be computed efficiently in a distributed manner (thus avoiding the complexities of exact SSSP), and (ii) an approximate pseudo-center is still sufficient for approximating the required distance parameters.

From Theorem 1, using the $(1+\varepsilon)$-approximate SSSP algorithms of [13,34], which run in $\tilde{O}((\sqrt{n}+D)/\varepsilon)$ rounds on weighted, undirected graphs and $\tilde{O}((\sqrt{n}D^{1/4}+D)/\varepsilon)$ rounds on weighted, directed graphs, respectively, we deduce the following corollaries:

▶ **Corollary 2.** *For any $\varepsilon = 1/\text{polylog}(n)$, there exists an algorithm for $(2+\varepsilon)$-approximate diameter, radius and all eccentricities running in $\tilde{O}(\sqrt{n}+D)$ rounds on nonnegative weighted graphs, with n nodes and hop-diameter D.*

▶ **Corollary 3.** *For any $\varepsilon = 1/\text{polylog}(n)$, there exists an algorithm for $(2+\varepsilon)$-approximate diameter, radius and all eccentricities running in $\tilde{O}(\sqrt{n}D^{1/4}+D)$ rounds on nonnegative weighted, directed graphs, with n nodes and hop-diameter D.*

Using the exact SSSP algorithm of Chechik and Mukhtar [21] we obtain the following.

▶ **Corollary 4.** *There exists an algorithm for 2-approximate radius, diameter and all eccentricities running in $\tilde{O}(\sqrt{n}D^{1/4}+D)$ rounds on nonnegative weighted, directed graphs, with n nodes and hop-diameter D.*

Regarding radius, the only previous result regarding the complexity of approximating the radius in the CONGEST model is due to [1], in which they showed that for any $\varepsilon > 0$, computing an $(3/2-\varepsilon)$-approximation to the radius in undirected, unweighted graphs requires

$\tilde{\Omega}(n)$ rounds. Abboud et al. [1] show that any algorithm computing an $(\frac{5}{3} - \varepsilon)$-approximation of all eccentricities requires $\tilde{\Omega}(n)$ rounds as well. Having a complete understanding of the relationship between approximation ratio and the round complexity of computing unweighted, undirected radius remains an intriguing open problem. As a step towards resolving this problem, we give a nearly full characterization of the approximation factor to round complexity mapping for radius in *weighted* or *directed* graphs in the CONGEST model.

In Section 4 we prove the following.

▶ **Theorem 5.** *Given any constant $\varepsilon > 0$, any algorithm (even randomized) computing an $(2 - \varepsilon)$-approximation to the weighted (directed) radius in a given weighted (directed) graph G requires $\tilde{\Omega}(n)$ rounds.*

A standard technique for proving lower bounds for the CONGEST model, is to reduce it from 2-party communication complexity. In the context of the distance parameters discussed in this work, this framework was used by [35] to show that any algorithm that distinguishes between networks with diameter 2 and 3 requires $\tilde{\Omega}(n)$ rounds. Later, [1] showed that this lower bound holds even when one considers sparse networks with only $O(n)$ edges (they also proved more results as discussed in the related work section).

Many of the papers that employ this framework, reduce from either the Set Disjointness function, the Equality function, or the Gap Disjointness function [10, 16, 23, 25]. In this work, we enhance this framework by showing lower bounds using reductions from other functions, which were not used previously to obtain lower bounds for the CONGEST model. Namely, in the proof of Theorem 5, we use the Tribes function, defined by Jayram et al. in [45], and the Hitting Set Existence (HSE) function, which is a communication complexity variant of a problem introduced by Abboud et al. in [3]. We elaborate upon this framework and the functions that we use in Section 2.

The following is a corollary of Theorem 7 and Theorem 8 which are stated below for the diameter, since any finite approximation to the radius, implies a finite approximation to the diameter, as $r \leq D \leq 2r$.

▶ **Corollary 6.** *Given any positive function $\alpha(n)$, any algorithm (even randomized) computing an $\alpha(n)$-approximation to the weighted (directed) radius in a given weighted (directed) graph G requires $\tilde{\Omega}(\sqrt{n} + D)$ rounds.*

Directed/Weighted Diameter. In previous work, Holzer and Pinsker [41] showed a lower bound of $\tilde{\Omega}(n)$ rounds for computing a $(2 - \varepsilon)$-approximation of the diameter of a given weighted graph. Shortly after, Becker et al. [13] designed an algorithm that computes a $(2 + o(1))$-approximation of weighted and directed diameter in $\tilde{O}(\sqrt{n}D^{1/4} + D)$ rounds. Such an algorithm makes one wonder, is there a smooth trade-off between the round complexity and the approximation ratio when going beyond a 2-approximation, for either the directed or weighted variants? In other words, can one further reduce the round complexity if we are willing to settle for a worse approximation ratio? For weighted diameter, this question was resolved by Lenzen et al. [49] in the negative, in the sense that the dependence on n in the algorithm of [13] is necessary (up to poly-logarithmic factors) for any approximation of the diameter in weighted or directed graphs. We give a proof of this result for completeness, and this allows us to more easily present a similar new result for the *bi-chromatic* diameter case. The bi-chromatic diameter is a variant of the diameter problem that is discussed later.

▶ **Theorem 7.** *Given any positive function $\alpha(n)$, any algorithm (even randomized) computing an $\alpha(n)$-approximation to the weighted diameter or bi-chromatic diameter in a given graph G requires $\tilde{\Omega}(\sqrt{n} + D)$ rounds.*

▶ **Theorem 8.** *Given any positive function $\alpha(n)$, any algorithm (even randomized) computing an $\alpha(n)$-approximation to the diameter in a given directed graph G requires $\tilde{\Omega}(\sqrt{n}+D)$ rounds.*

To prove these theorems we reduce from the problem of Spanning Connected Subgraph Verification (SCSV) to approximating these parameters. The SCSV problem is known to admit the above lower bound due to Das Sarma et al. [25]. The key challenge is to construct a reduction in a manner that can be efficiently simulated in CONGEST. The proofs of these theorems are given in full version of the paper.

Undirected and Unweighted Diameter, Radius and Eccentricities. Abbout et al. [1] show that for any $\varepsilon > 0$, any algorithm computing an $(\frac{3}{2} - \varepsilon)$-approximation of diameter or radius in unweighted undirected graphs has round complexity $\tilde{\Omega}(n)$. Furthermore, any algorithm computing an $(\frac{5}{3} - \varepsilon)$-approximation to all eccentricities has round complexity $\tilde{\Omega}(n)$. For upper bounds, the state of art for diameter approximation is an algorithm by Holzer et al. [40], computing a 3/2-approximation in $\tilde{O}(\sqrt{n \log n} + D)$ rounds. Fully understanding the mapping of approximation ratios in the range $[\frac{3}{2}, 2)$ for diameter and radius, and in the range $(\frac{5}{3}, 3)$ for all eccentricities, to their respective correct round complexity in the CONGEST model remains open. As a step towards resolving this open problem, in the full version of the paper, we present a simple distributed implementation of a sequential approximation algorithm of Cairo et al. [15] for diameter, radius and eccentricities with the following parameters.

▶ **Theorem 9.** *For any $k \in \mathbb{N}$, there exist algorithms that compute $(2 - \frac{1}{2^k})$-approximate diameter and radius and $(3 - \frac{4}{2^k+1})$-approximate eccentricities on unweighted, undirected graphs, that have running time of $\tilde{O}(n^{\frac{1}{k+1}} + D)$ rounds w.h.p.*

Bi-chromatic Diameter and Radius. To the best of our knowledge, no previous results regarding bi-chromatic distance parameters are known in distributed settings. Roughly speaking, these variants are defined using only distances between pairs of nodes in $S \times T$ where $S, T \subseteq V, T = V \setminus S$. D_{ST},R_{ST} respectively denote the ST-diameter $\max_{s \in S, t \in T} d(s,t)$ and the ST-radius $\min_{s \in S} \max_{t \in T} d(s,t)$ (also see Section 2.1). In the following, $T(SSSP)$ refers to the distributed complexity of exact weighted SSSP. The proofs of these theorems can be found in the full version of the paper

▶ **Theorem 10.** *There is an algorithm with complexity $\tilde{O}(\sqrt{n} + D)$ that given an undirected, unweighted graph $G = (V, E)$, and sets $S \subseteq V, T = V \backslash S$, w.h.p. computes a value D^*_{ST} such that $\frac{3D_{ST}}{5} - \frac{6}{5} \leq D^*_{ST} \leq D_{ST}$.*

▶ **Theorem 11.** *There is an algorithm with complexity $T(SSSP)$ that given an undirected graph $G = (V, E)$, and sets $S \subseteq V, T = V \backslash S$, computes a value D^* such that $\frac{D_{ST}}{2} - W/2 \leq D^* \leq D_{ST}$. Here W is the minimum edge weight in $S \times T$.*

We remark that using very similar algorithms to the ones of Theorem 10 and Theorem 11, one can obtain the following results, whose proofs we omit due to similarity to the main ideas in the proofs we provide for the above two theorems.

▶ Remark 12. There are algorithms with complexity $\tilde{O}(\sqrt{n} + D)$ that given an undirected, unweighted graph $G = (V, E)$, and sets $S, T \subseteq V$, compute w.h.p. the following.

1. A value R^*_{ST} such that $R_{ST} \leq R^*_{ST} \leq \frac{5R_{ST}}{3} + \frac{5}{3}$, in the case that $S = V \backslash T$.
2. A 2-approximation to all ST-eccentricities.
3. A 2-approximation to R_{ST}.

▶ Remark 13. There are algorithms with complexity $T(SSSP)$ that given an undirected graph $G = (V, E)$, and sets $S, T \subseteq V$, compute the following.

1. A value R^*_{ST} such that $R_{ST} \leq R^*_{ST} \leq 2R_{ST} + W$, in the case that $S = V \backslash T$. Here W is the minimum edge weight in $S \times T$.
2. A 3-approximation to all ST-eccentricities.
3. A 3-approximation to R_{ST}.

We complement these upper bounds with several lower bounds. We show that in the weighted case, one cannot hope to do better than a $\frac{5}{3}$-approximation for bi-chromatic diameter with $O(n^{1-\epsilon})$ rounds for some $\epsilon > 0$. Additionally, as a step towards realizing the complexity of finding a better than 2-approximation for directed diameter, we show that for *bi-chromatic* diameter, in which one is tasked with finding the largest distance between a pair of nodes in different sets of a given partition of the graph, finding such an approximation is a hard task. Formally, we prove the following theorems in the full version of the paper due to lack of space.

▶ **Theorem 14.** *For all constant $\epsilon > 0$, there is no $o(\frac{n}{\log^3 n})$ round algorithm for computing a $(\frac{5}{3} - \epsilon)$-approximation to the bi-chromatic diameter in an unweighted, undirected graph.*

▶ **Theorem 15.** *For all constant $\epsilon > 0$, there is no $o(\frac{n}{\log^2 n})$ round algorithm for computing a $(2 - \epsilon)$-approximation to the bi-chromatic diameter in a directed graph.*

Finally, we show that for both the directed and weighted cases, any approximation of the bi-chromatic diameter requires $\tilde{\Omega}(\sqrt{n} + D)$ rounds. The weighted case is proved as part of Theorem 7. In the full version of the paper we prove separately the directed case, which is stated formally as follows.

▶ **Theorem 16.** *Given any positive function $\alpha(n)$, any algorithm (even randomized) computing an $\alpha(n)$-approximation to the bi-chromatic diameter in a given directed graph G requires $\tilde{\Omega}(\sqrt{n} + D)$ rounds.*

1.2 Additional related work

The state of the art algorithm for 3/2-approximation of unweighted, undirected diameter [40] was preceded by a significant number of works. Notable examples are Holzer's and Wattenhofer's algorithm computing a 3/2-approximation of the diameter in undirected, unweighted graphs in $O(n^{3/4} + D)$ rounds [42], and the independent work of Peleg et al. [53], which achieves the same approximation in $O(D\sqrt{n}\log n)$ rounds. Later, Lenzen and Peleg [50] improved this upper bound to $O(\sqrt{n}\log n + D)$.

Approximations to more concrete variants of distance computations such as APSP and SSSP have been extensively studied in the CONGEST as well. Examples include the deterministic $(1 + o(1))$-approximation to APSP by Nanongkai [52], and the $(1 + \epsilon)$-approximation algorithm for SSSP of Becker et al. [13]. The near optimal algorithm of Bernstein and Nanongkai for APSP [14] was preceded by a series of papers that set to realize the complexity of APSP in CONGEST [5, 6, 8, 30, 43]. Given that [14] is a randomized Las Vegas algorithm, there remains a gap between between the best known deterministic and randomized algorithms for APSP, with the deterministic state of the art being $\tilde{O}(n^{4/3})$ [7]. For SSSP, the state of the art algorithm of [21] was also preceded by a series of improvements [13, 30, 31, 34, 37, 38, 48, 52] from the folklore $O(n)$ Bellman-Ford algorithm.

Approximations to distance computations have been studied in various distributed settings, such as the congested clique model. Starting from [18], which presented the first non trivial algorithms for both exact, and approximated APSP in the model. From there a series of works designed more and more efficient algorithms for approximating distances in the model [13, 17, 20, 26, 31, 32, 36], with the most recent work being the $poly(\log\log n)$ approximations for APSP and Multi Source Shortest Paths [27].

Conditional hardness results for these parameters are very well-studied in the sequential setting, within fine-grained complexity, under assumptions such as the Strong Exponential Time Hypothesis (SETH) [44]. For details, see e.g., the work of Backurs et al. [11] or the survey by Vassilevska Williams [56]. Returning to the CONGEST model, in some topologies such as planar graphs, work by Li and Parter [51] showed that the diameter of an unweighted, undirected graph can even be computed in a sublinear number of rounds.

The lower bound framework for reducing 2-party communication complexity to CONGEST was introduced by Peleg and Rubinovich in [54], in which they show that any algorithm solving the minimum spanning tree (MST) problem has round complexity $\tilde{\Omega}(\sqrt{n}+D)$. Since then, there has been a surge of lower bounds for the CONGEST model employing this framework; examples include [2, 10, 23, 25, 28, 33]. In an independent concurrent work, [9] show another angle of the landscape of the complexity of diameter approximation, proving that for any constant $\epsilon > 0$, any algorithm approximating the diameter of a given unweighted, undirected graph, within a factor of $(\frac{3}{5}+\epsilon)$, $(\frac{4}{7}+\epsilon)$, or $(\frac{6}{11}+\epsilon)$, must have a round complexity of at least $\tilde{\Omega}(n^{1/3})$, $\tilde{\Omega}(n^{1/4})$, or $\tilde{\Omega}(n^{1/6})$, respectively.

2 Preliminaries

2.1 The Model & Definitions

This paper considers the CONGEST model of computation. In this model, a synchronized network of n nodes is represented by an undirected, unweighted, simple graph $G = (V, E)$. In each round, each node can send a different message of $O(\log n)$ bits to each of its neighbors.

Next, we define the network parameters that we discuss in the paper.

▶ **Definition 17.** *Given a weighted, directed graph $G = (V, E)$, denote by $d(u, v)$ the weight of the lightest directed path starting at node u and ending at node v. If there is no such path, we define $d(u, v) = \infty$. Here, the weight of a path P is the sum of the weights of its edges. The eccentricity $ecc(u)$ of a node u is defined to be $\max_{v\in V} d(u, v)$. The radius r of G is defined to be $\min_{v\in V} ecc(v)$. The diameter D of G is defined to be $\max_{v\in V} ecc(v)$.*

The ST variants of these distance parameters are defined as follows.

▶ **Definition 18** (ST and bi-chromatic diameter, radius and eccentricities.)**.** *Given a weighted graph $G = (V, E)$, and two non empty subsets $S, T \subseteq V$, given $v \in S$, we define its ST-eccentricity by $ecc(v) = \max_{u\in T} d(v, u)$. We define the ST-diameter of G to be $D_{ST} = \max_{v\in S} ecc(v)$. The ST-radius of G is defined to be $R_{ST} = \min_{v\in S} ecc(v)$. When $S = V\backslash T$, the ST parameters are called bi-chromatic.*

2.2 The Communication Complexity Framework

The high level idea of applying the framework of reductions from 2-party communication complexity to obtain lower bounds in the CONGEST model is as follows. We pick some function $f : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}$, and then reduce any efficient communication protocol for it to an efficient CONGEST algorithm for the discussed problem. We start with our two players Alice (A) and Bob (B), each of them respectively receives a binary string of length k denoted by $x, y \in \{0,1\}^k$.

We construct a graph $G = (V, E)$ we call the *fixed graph construction*, and we partition the set of vertices V into the sets V_A, V_B. We call the cut induced by V_A, V_B the *communication cut*, and we denote the number of edges in this cut by $|cut|$.

Now, given x and the graph $G[V_A]$ (i.e., the subgraph of G induced by V_A), Alice modifies the graph $G[V_A]$ in any way that may depend only on x, and Bob does the same with y and $G[V_B]$. Denote the resulting graph by $G_{x,y}$, and denote its number of nodes by n.

The resulting graph $G_{x,y}$ should be constructed such that it has some property P (e.g. radius at least 3) iff $f(x, y) = 1$. Now, assuming there is an algorithm Alg in the CONGEST model that decides P in T rounds, Alice and Bob can simulate this algorithm on $G_{x,y}$, and the only communication required between them is for simulating messages that are sent on edges in the communication cut. Thus, Alice and bob can simulate $Alg(G_{x,y})$ while communicating $O(T \cdot |cut| \cdot \log n)$ bits of communication. Furthermore, by the property of $G_{x,y}$, deciding P on $G_{x,y}$ allows them to compute $f(x, y)$ with $O(T \cdot |cut| \cdot \log n)$ bits of communication. Therefore, a lower bound on the communication complexity of f, implies a lower bound on T, which is the round complexity of the distributed algorithm.

We next elaborate on the functions f that we use in our reductions.

▶ **Definition 19** (The Set Disjointness Problem (Disj) [55]). *Alice and Bob receive subsets $X, Y \subseteq [n]$, respectively, represented as binary vectors of length n. Their goal is to decide whether $X \cap Y = \emptyset$.*

It is known by [12, 47, 55] that the randomized communication complexity of Disj on inputs of size n is $\Omega(n)$.

▶ **Definition 20** (The Tribes (ListDISJ) Problem [45]). *Alice and Bob are given sets $A_i, B_i \in \{0, 1\}^N$ for each $i \in [N]$. They must output 1 if and only if there is some i such that A_i and B_i are disjoint, i.e. there is no j such that $A_{ij} = B_{ij} = 1$. We treat the inputs x and y as binary strings of length N^2, such that $x = A_1 \circ ... \circ A_N, y = B_1 \circ ... \circ B_N$. Here, $\circ$ refers to string concatenation.*

The Tribes function is defined in [45], where a lower bound of $\Omega(N^2)$ communication bits is proved, even for randomized protocols.

The full version of the paper contains discussion of additional functions which are employed to prove the results not present in this version.

3 Approximation Algorithms

3.1 Approximations for weighted directed variants

In this section, we prove our approximation algorithms, starting with the connection between the complexity of $SSSP$ and approximating distance parameters. Formally, we prove the following theorem, and then we deduce Corollaries 2, 3, and 4.

Theorem 1 *For any $\varepsilon \geq 0$, given a $(1 + \varepsilon)$-approximation algorithm $\mathcal{A}_\varepsilon$ for weighted and directed SSSP running in $T(n, \varepsilon, D)$ rounds, there exists an algorithm for $(2 + \varepsilon^3 + 3\varepsilon^2 + 4\varepsilon)$-approximate diameter, radius, and all eccentricities in $\tilde{O}(T(n, \varepsilon, D) + D)$ rounds on weighted, directed graphs.*

We briefly remind the reader of the discussion in the introduction regarding the theorem. In order to obtain fast algorithms and maintaining the quality of the approximation, we generalize the notion of *pseudo-center* defined by Choudhary and Gold [22] into *approximate pseudo-center*. We show how to compute such a set of small size, and we show that such a set suffices to obtain the approximations detailed in Theorem 1.

▶ **Definition 21.** *A α-approximate pseudo-center is a set C of nodes such that for all nodes $v \in V$, $ecc(v) \geq \max_{u \in V} \min_{c \in C} \{d(c,u)/\alpha\}$.*

We begin by showing that we can compute a small approximate pseudo-center efficiently.

▶ **Lemma 22.** *Given a $(1+\varepsilon)$-approximate, $T(n,\varepsilon,D)$-round SSSP algorithm $\mathcal{A}_\varepsilon$, there is a Las Vegas algorithm to compute a $(1+\varepsilon)^2$-approximate pseudo-center of size $O(\log^2(n))$ of a graph $G=(V,E)$ in $\tilde{O}(T(n,\varepsilon,D))$ rounds of communication, with high probability.*

Proof. Let the set C begin empty, and let W begin as the set V. Throughout the proof, running $\mathcal{A}_\varepsilon$ outward (inward) from a vertex $v \in V$ means computing the distances *from* v to the rest of the nodes (*to* v from the rest of the nodes). We repeat the following until W is empty:

- Assign each node in W to a set S independently with probability $\min\{1, 24\log(n)/|W|\}$. Resample if $|S| < 8\log n$ or $|S| > 36\log n$.
- Run $\mathcal{A}_\varepsilon$ outward from each node in S, and for all $u \in V$, compute estimated distances $d_{\mathcal{A}_\varepsilon}(S,u) = \min_{s\in S}\{d_{\mathcal{A}_\varepsilon}(s,u)\}$.
- Let a be the node with the largest estimated distance from S. Then, we broadcast $d_{\mathcal{A}_\varepsilon}(S,a)$ to all nodes in the graph using some BFS tree.
- Run $\mathcal{A}_\varepsilon$ inward from a, and remove all nodes u where $d_{\mathcal{A}_\varepsilon}(u,a) \geq d_{\mathcal{A}_\varepsilon}(S,a)$ from W.
- Add S to C.

First, we argue that C is a $(1+\varepsilon)^2$-approximate pseudo-center. We only remove a node u from W when $d_{\mathcal{A}_\varepsilon}(u,a) \geq d_{\mathcal{A}_\varepsilon}(S,a)$ for some sample S. Let a^* be the node that is truly farthest from S; then $d_{\mathcal{A}_\varepsilon}(S,a) \geq d(S,a^*)/(1+\varepsilon) \geq \max_{x\in V}\min_{c\in C}\{d(c,x)/(1+\varepsilon)\}$, because $S \subseteq C$. We also note that by similarly bounding the error of $\mathcal{A}_\varepsilon$, it holds that $d_{\mathcal{A}_\varepsilon}(u,a) \leq (1+\varepsilon)d(u,a) \leq (1+\varepsilon)ecc(u)$, so we may conclude that

$$(1+\varepsilon)ecc(u) \geq \max_{x\in V}\min_{c\in C}\{d(c,x)/(1+\varepsilon)\}.$$

In other words, $ecc(u) \geq \max_{x\in V}\min_{c\in C}\left\{d(c,x)/(1+\varepsilon)^2\right\}$, which meets the definition of a $(1+\varepsilon)^2$-pseudo-center.

Next, we argue that each iteration requires $\tilde{O}(T(n,\varepsilon,D))$ rounds. Using a Chernoff bound, it is simple to show that in each round, $8\log n \leq |S| \leq 36\log n$ with probability at least $1-1/n^4$, so we expect to resample a sub-constant number of times. We then run $\mathcal{A}_\varepsilon$ from each node in S and we run it again once to the node a, for a total of $O(\log n \cdot T(n,\varepsilon,D))$ rounds. The rest of each iteration involves a constant number of broadcasts that take $O(D)$ rounds in total.

Finally, we argue that with high probability, we only have $O(\log n)$ iterations in our algorithm. We do this by showing that in iteration i, the size of W reduces by at least half with high probability, i.e. $|W_i|/2 \geq |W_{i+1}|$. Consider the set $X \subseteq W_i$ of $|W_i|/2$ nodes with the smallest $d_{\mathcal{A}_\varepsilon}(u,a)$, $u \in W_i$. Note that S_i is a randomly sampled subset of W_i of size at least $8\log n$, and thus intersects X with probability at least $(1-1/n^5)$, as argued in Lemma 23 below [22] with no further assumptions.

All nodes in $W_i \backslash X$ are at least as far as any node in that intersection under $\mathcal{A}_\varepsilon$, by definition. This implies that for all $u \in W_i\backslash X$, $d_{\mathcal{A}_\varepsilon}(u,a) \geq d_{\mathcal{A}_\varepsilon}(S,a)$, which implies that all $|W_i|/2$ nodes of $W_i\backslash X$ will be removed from W_i in iteration i. ◀

▶ **Lemma 23** (Lemma 2.1 in [22])**.** *Let U be a universe set of size at most n, and let $S_1,\dots,S_n \subseteq U$ such that $|S_i| \geq L$ for each $i \in [n]$. Let c be some constant and $r = \frac{n(c+1)\ln n}{L}$. Let $S \subseteq U$ be a random subset of size r, then it holds that $S \cap S_i \neq \emptyset$ for all i with probability $1-n^{-c}$.*

Now that we showed how to compute an approximate pseudo-center, we show that it is sufficient for approximating the distance parameters as claimed.

▶ **Lemma 24.** *Given a $(1+\varepsilon)^2$-approximate pseudo-center C and a $(1+\varepsilon)$-approximate SSSP algorithm $\mathcal{A}_\varepsilon$ taking $T(n,\varepsilon,D)$ rounds, we may compute $(2+\varepsilon^3+3\varepsilon^2+4\varepsilon)$-approximate eccentricities for all nodes in $O(|C|\cdot T(n,\varepsilon,D)+D)$ rounds.*

Proof. First, we run $\mathcal{A}_\varepsilon$ to and from each node in C, so that each node $v \in V$ stores $d_{\mathcal{A}_\varepsilon}(c,v)$ and $d_{\mathcal{A}_\varepsilon}(v,c)$ for all $c \in C$. Each node u internally determines $\min_{c\in C}\{d_{\mathcal{A}_\varepsilon}(c,u)\}$. Then, using aggregation over a BFS tree, the nodes determine, and then broadcast the value $D_{\mathcal{A}_\varepsilon}(C) := \max_{u\in V}\min_{c\in C}\{d_{\mathcal{A}_\varepsilon}(c,u)\}$. Thus, the aggregation takes $O(D)$ rounds. Each node v approximates its eccentricity as $\max_{c\in C}\{d_{\mathcal{A}_\varepsilon}(v,c)\} + D_{\mathcal{A}_\varepsilon}(C)$.

First, note that this estimate is at least the true eccentricity of v, as each computed distance represents some path in the graph, and in this distance a path can go from v to any node in C and then any node in V.

We argue that this is a $(2+\varepsilon^3+3\varepsilon^2+4\varepsilon)$-approximation. The estimated distance $\max_{c\in C}\{d_{\mathcal{A}_\varepsilon}(v,c)\}$ is at most $(1+\varepsilon)\cdot ecc(v)$, because $\mathcal{A}_\varepsilon$ overestimates by at most a factor of $1+\varepsilon$. By our definition of $(1+\varepsilon)^2$-approximate pseudo-center, $D(C) \le (1+\varepsilon)^2 ecc(v)$. Our estimate $D_{\mathcal{A}_\varepsilon}(C)$ is at most $(1+\varepsilon)\cdot D(C)$, so $D_{\mathcal{A}_\varepsilon}(C) \le (1+\varepsilon)^3 ecc(v)$. Thus, $\max_{c\in C}\{d_{\mathcal{A}_\varepsilon}(v,c)\} + D_{\mathcal{A}_\varepsilon}(C) \le (1+\varepsilon+(1+\varepsilon)^3)\cdot ecc(v) = (2+\varepsilon^3+3\varepsilon^2+4\varepsilon)\cdot ecc(v)$.

We compute $\mathcal{A}_\varepsilon$ twice for each element of C, and broadcast a constant number of values to all nodes, so the total number of rounds is $O(|C|\cdot T(n,\varepsilon,D)+D)$. ◀

Proof of Theorem 1. Applying Lemma 22 and Lemma 24, given a $(1+\varepsilon)$-approximate algorithm $\mathcal{A}_\varepsilon$ for SSSP running in $T(n,\varepsilon,D)$ rounds, we may compute $(2+\varepsilon^3+3\varepsilon^2+4\varepsilon)$-approximations for all eccentricities in $O(\log^2(n)\cdot T(n,\varepsilon,D)+D)$ rounds. ◀

Using the $(1+\varepsilon)$-approximate SSSP algorithms of [13, 34], which run in $\tilde{O}((\sqrt{n}+D)/\varepsilon)$ rounds on weighted, undirected graphs and $\tilde{O}((\sqrt{n}D^{1/4}+D)/\varepsilon)$ rounds on weighted, directed graphs respectively, we achieve the following corollaries:

▶ **Corollary 2.** *For any $\varepsilon = 1/\text{polylog}(n)$, there exists an algorithm for $(2+\varepsilon)$-approximate diameter, radius and all eccentricities running in $\tilde{O}(\sqrt{n}+D)$ rounds on nonnegative weighted graphs, with n nodes and hop-diameter D.*

▶ **Corollary 3.** *For any $\varepsilon = 1/\text{polylog}(n)$, there exists an algorithm for $(2+\varepsilon)$-approximate diameter, radius and all eccentricities running in $\tilde{O}(\sqrt{n}D^{1/4}+D)$ rounds on nonnegative weighted, directed graphs, with n nodes and hop-diameter D.*

Using the exact SSSP algorithm of [21], which runs in $\tilde{O}(\sqrt{n}D^{1/4}+D)$ rounds, we obtain the following corollary.

▶ **Corollary 4.** *There exists an algorithm for 2-approximate radius, diameter and all eccentricities running in $\tilde{O}(\sqrt{n}D^{1/4}+D)$ rounds on nonnegative weighted, directed graphs, with n nodes and hop-diameter D.*

4 Hardness of Approximation

In this section, we prove the lower bound results of the paper. As stated, we use reductions from 2-party communication complexity. To formalize the reductions, we restate the following definition from Censor-Hillel et al. [19].

▶ **Definition 25** (Family of Lower Bound Graphs). *Given integers K and n, a Boolean function $f : \{0,1\}^K \times \{0,1\}^K \to \{0,1\}$ and some Boolean graph property or predicate denoted P, a set of graphs $\{G_{x,y} = (V, E_{x,y}) \mid x, y \in \{0,1\}^K\}$ is called a* family of lower bound graphs *with respect to f and P if the following hold:*

1. *The set of vertices V is the same for all the graphs in the family, and we denote by V_A, V_B a fixed partition of the vertices.*
2. *Given $x, y \in \{0,1\}^K$, the only part of the graph which is allowed to be dependent on x (by adding edges or weights, no adding vertices) is $G[V_A]$.*
3. *Given $x, y \in \{0,1\}^K$, the only part of the graph which is allowed to be dependent on y (by adding edges or weights, no adding vertices) is $G[V_B]$.*
4. *$G_{x,y}$ satisfies P if and only if $f(x, y) = 1$.*

The set of edges $E(V_A, V_B)$ is denoted by E_{cut}, and is the same for all graphs in the family.

We use the following theorem whose proof can be found in Censor-Hillel et al. [19], with $CC^R(f)$ denoting the randomized communication complexity of f.

▶ **Theorem 26.** *Fix a function $f : \{0,1\}^K \times \{0,1\}^K \to \{0,1\}$ and a predicate P. If there exists a family of lower bound graphs $\{G_{x,y}\}$ w.r.t f and P, then every randomized algorithm for deciding P takes $\Omega(CC^R(f)/(|E_{cut}| \log n))$ rounds.*

4.1 Lower bounds for radius

We start with proving our two lower bounds for weighted or directed radius approximations.

We divide the proof of Theorem 5 into two cases which we prove separately. We prove the weighted case here, and the proof of the directed case appears in the full version of the paper.

Theorem 5 [**Weighted case**] *For any $\varepsilon = 1/\text{poly}(n)$, $(2-\varepsilon)$-approximation of the radius of a weighted graph with n nodes requires $\Omega(n/\log n)$ rounds, even when the graph has constant hop-diameter.*

Proof. We reduce from the Tribes problem with vector sets A and B of size N. This construction is similar to that of [41, Theorem 7].

Figure 1 illustrates our family of lower bound graphs. We construct four cliques A^0, A^1, B^0, B^1 of size N, where the edges of the cliques have weight t, a value we will set later. Let K_i be the ith node in clique K. Add two nodes c_A and c_B.

Connect all nodes in A^0 to c_A with edges of weight t, and connect all nodes in B^0 to c_B with edges of weight t. Connect c_A and c_B with an edge of weight 1. For all $i \in [N]$ and $b \in \{0,1\}$, connect A_i^b and B_i^b with an edge of weight 1. Connect A_i^0 and A_j^1 with an edge of weight t if and only if $A_i[j] = 0$. Connect B_i^0 and B_j^1 with an edge of weight t if and only if $B_i[j] = 0$. Alice will simulate the nodes $A^0 \cup A^1 \cup \{c_A\}$, and Bob will simulate the nodes $B^0 \cup B^1 \cup \{c_B\}$.

First, we claim that if (A, B) is a "yes" instance of Tribes, then the radius is at most $t+2$. To show this, note that in this case, there must be some i such that the ith vectors of A and B are orthogonal. Consider the node A_i^0. It may reach in distance at most $t+1$ all nodes in $B^0 \cup A^0$, via a clique edge and an edge in the matching between A^0 and B^0. It may also reach $\{c_A, c_B\}$ in at most $t+1$. It may also reach all nodes in $A^1 \cup B^1$ in distance at most $t+2$, because for any j where $A_i[j] = 0$ or $B_i[j] = 0$, either A_i^0 may reach A_j^1 in distance t or B_i^0 may reach B_j^1 in distance t. Since A_i and B_i are orthogonal, this is true for all j. Thus the eccentricity of A_i^0 is at most $t+2$, which upper-bounds the radius.

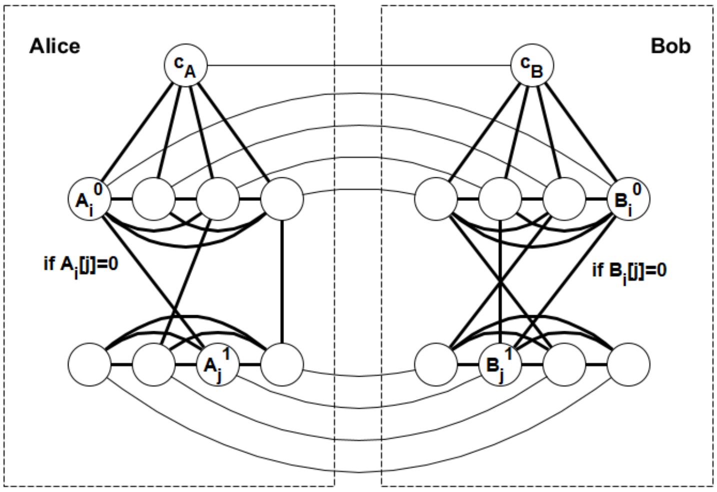

Figure 1 Sketch of Theorem 5, weighted case construction. Bold lines represent edges of weight t.

Second, we claim that if (A, B) is a "no" instance of Tribes, then the radius is at least $2t$. To see this, first note that c_A and c_B have eccentricity at least $2t$, because that is the shortest possible distance between them and $B^1 \cup A^1$. By the same argument, the eccentricity of all nodes in $A^1 \cup B^1$ is also at least $2t$. For all i, A_i and B_i are not orthogonal, which means that for all i there is some j such that neither A_i^0 nor B_i^0 has an edge to B_j^1 or A_j^1. Clearly any other path from B_i^0 or A_i^0 to B_j^1 or A_j^1 is at least of length $2t$, via a clique edge of weight t. Thus the eccentricities of all nodes are at least $2t$, so the radius is at least $2t$.

We set $t = \lceil \frac{4}{\varepsilon} \rceil$ so that a $(2 - \varepsilon)$-approximate radius algorithm needs to distinguish between $t + 2$ and $2t$. The constructed graph $G_{A,B}$ has $n = O(N)$ nodes with a cut of size $O(n)$, which by Theorem 26 and the lower bound of $\Omega(N^2)$ for the communication complexity of Tribes, implies that the radius algorithm requires $\Omega(n/\log n)$ rounds. ◀

References

1 Amir Abboud, Keren Censor-Hillel, and Seri Khoury. Near-linear lower bounds for distributed distance computations, even in sparse networks. In *Distributed Computing - 30th International Symposium, DISC 2016, Paris, France, September 27-29, 2016. Proceedings*, pages 29–42, 2016. doi:10.1007/978-3-662-53426-7_3.

2 Amir Abboud, Keren Censor-Hillel, Seri Khoury, and Ami Paz. Smaller cuts, higher lower bounds. *CoRR*, abs/1901.01630, 2019. arXiv:1901.01630.

3 Amir Abboud, Virginia Vassilevska Williams, and Joshua Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 377–391, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics. URL: http://dl.acm.org/citation.cfm?id=2884435.2884463.

4 Pankaj K. Agarwal, Herbert Edelsbrunner, and Otfried Schwarzkopf. Euclidean minimum spanning trees and bichromatic closest pairs. *Discret. Comput. Geom.*, 6:407–422, 1991. doi:10.1007/BF02574698.

5 Udit Agarwal and Vijaya Ramachandran. New and simplified distributed algorithms for weighted all pairs shortest paths. *CoRR*, abs/1810.08544, 2018. arXiv:1810.08544.

6 Udit Agarwal and Vijaya Ramachandran. Distributed weighted all pairs shortest paths through pipelining. In *2019 IEEE International Parallel and Distributed Processing Symposium,*

IPDPS 2019, Rio de Janeiro, Brazil, May 20-24, 2019, pages 23–32. IEEE, 2019. doi:10.1109/IPDPS.2019.00014.

7 Udit Agarwal and Vijaya Ramachandran. Faster deterministic all pairs shortest paths in congest model. In Christian Scheideler and Michael Spear, editors, *SPAA '20: 32nd ACM Symposium on Parallelism in Algorithms and Architectures, Virtual Event, USA, July 15-17, 2020*, pages 11–21. ACM, 2020. doi:10.1145/3350755.3400256.

8 Udit Agarwal, Vijaya Ramachandran, Valerie King, and Matteo Pontecorvi. A deterministic distributed algorithm for exact weighted all-pairs shortest paths in õ(n 3/2) rounds. In Calvin Newport and Idit Keidar, editors, *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, PODC 2018, Egham, United Kingdom, July 23-27, 2018*, pages 199–205. ACM, 2018. doi:10.1145/3212734.3212773.

9 Anonymous authors. Improved hardness of approximation of diameter in the congest model. In *submission*, 2020.

10 Nir Bachrach, Keren Censor-Hillel, Michal Dory, Yuval Efron, Dean Leitersdorf, and Ami Paz. Hardness of distributed optimization. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 238–247, 2019. doi:10.1145/3293611.3331597.

11 Arturs Backurs, Liam Roditty, Gilad Segal, Virginia Vassilevska Williams, and Nicole Wein. Towards tight approximation bounds for graph diameter and eccentricities. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 267–280, New York, NY, USA, 2018. ACM. doi:10.1145/3188745.3188950.

12 Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68(4):702–732, 2004. doi:10.1016/j.jcss.2003.11.006.

13 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 2017)*, volume 91 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:16, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.DISC.2017.7.

14 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 2019, pages 334–342, New York, NY, USA, 2019. ACM. doi:10.1145/3313276.3316326.

15 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 '16, pages 363–376, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics. URL: http://dl.acm.org/citation.cfm?id=2884435.2884462.

16 Keren Censor-Hillel and Michal Dory. Distributed spanner approximation. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing, PODC 2018, Egham, United Kingdom, July 23-27, 2018*, pages 139–148, 2018. doi:10.1145/3212734.3212758.

17 Keren Censor-Hillel, Michal Dory, Janne H. Korhonen, and Dean Leitersdorf. Fast approximate shortest paths in the congested clique. In Peter Robinson and Faith Ellen, editors, *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 74–83. ACM, 2019. doi:10.1145/3293611.3331633.

18 Keren Censor-Hillel, Petteri Kaski, Janne H. Korhonen, Christoph Lenzen, Ami Paz, and Jukka Suomela. Algebraic methods in the congested clique. *Distributed Comput.*, 32(6):461–478, 2019. doi:10.1007/s00446-016-0270-2.

19 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, October 16-20, 2017, Vienna, Austria, pages 10:1–10:16, 2017. doi:10.4230/LIPIcs.DISC.2017.10.

20 Shiri Chechik and Doron Mukhtar. Reachability and shortest paths in the broadcast CONGEST model. In Jukka Suomela, editor, *33rd International Symposium on Distributed Computing, DISC 2019, October 14-18, 2019, Budapest, Hungary*, volume 146 of *LIPIcs*, pages 11:1–11:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.DISC.2019.11.

21 Shiri Chechik and Doron Mukhtar. Single-source shortest paths in the CONGEST model with improved bound. In Yuval Emek and Christian Cachin, editors, *PODC '20: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, August 3-7, 2020*, pages 464–473. ACM, 2020. doi:10.1145/3382734.3405729.

22 Keerti Choudhary and Omer Gold. Extremal distances in directed graphs: Tight spanners and near-optimal approximation algorithms. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 495–514. SIAM, 2020. doi:10.1137/1.9781611975994.30.

23 Artur Czumaj and Christian Konrad. Detecting cliques in CONGEST networks. In *32nd International Symposium on Distributed Computing, DISC 2018, New Orleans, LA, USA, October 15-19, 2018*, pages 16:1–16:15, 2018. doi:10.4230/LIPIcs.DISC.2018.16.

24 Mina Dalirrooyfard, Virginia Vassilevska Williams, Nikhil Vyas, and Nicole Wein. Tight approximation algorithms for bichromatic graph diameter and related problems. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPIcs*, pages 47:1–47:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ICALP.2019.47.

25 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. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 363–372, New York, NY, USA, 2011. ACM. doi:10.1145/1993636.1993686.

26 Michael Dinitz and Yasamin Nazari. Massively parallel approximate distance sketches. In Pascal Felber, Roy Friedman, Seth Gilbert, and Avery Miller, editors, *23rd International Conference on Principles of Distributed Systems, OPODIS 2019, December 17-19, 2019, Neuchâtel, Switzerland*, volume 153 of *LIPIcs*, pages 35:1–35:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.OPODIS.2019.35.

27 Michal Dory and Merav Parter. Exponentially faster shortest paths in the congested clique. *CoRR*, abs/2003.03058, 2020. arXiv:2003.03058.

28 Andrew Drucker, Fabian Kuhn, and Rotem Oshman. On the power of the congested clique model. In *ACM Symposium on Principles of Distributed Computing, PODC '14, Paris, France, July 15-18, 2014*, pages 367–376, 2014. doi:10.1145/2611462.2611493.

29 Adrian Dumitrescu and Sumanta Guha. Extreme distances in multicolored point sets. *J. Graph Algorithms Appl.*, 8:27–38, 2004. doi:10.7155/jgaa.00080.

30 Michael Elkin. Distributed exact shortest paths in sublinear time. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 757–770. ACM, 2017. doi:10.1145/3055399.3055452.

31 Michael Elkin and Ofer Neiman. Hopsets with constant hopbound, and applications to approximate shortest paths. *SIAM J. Comput.*, 48(4):1436–1480, 2019. doi:10.1137/18M1166791.

32 Michael Elkin and Ofer Neiman. Linear-size hopsets with small hopbound, and constant-hopbound hopsets in RNC. In Christian Scheideler and Petra Berenbrink, editors, *The 31st ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA 2019, Phoenix, AZ, USA, June 22-24, 2019*, pages 333–341. ACM, 2019. doi:10.1145/3323165.3323177.

33 Orr Fischer, Tzlil Gonen, Fabian Kuhn, and Rotem Oshman. Possibilities and impossibilities for distributed subgraph detection. In *Proceedings of the 30th on Symposium on Parallelism in*

Algorithms and Architectures, SPAA 2018, Vienna, Austria, July 16-18, 2018, pages 153–162, 2018. doi:10.1145/3210377.3210401.

34 Sebastian Forster and Danupon Nanongkai. A faster distributed single-source shortest paths algorithm. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 686–697. IEEE Computer Society, 2018. doi:10.1109/FOCS.2018.00071.

35 Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 1150–1162, 2012. doi:10.1137/1.9781611973099.91.

36 François Le Gall. Further algebraic algorithms in the congested clique model and applications to graph-theoretic problems. In Cyril Gavoille and David Ilcinkas, editors, *Distributed Computing - 30th International Symposium, DISC 2016, Paris, France, September 27-29, 2016. Proceedings*, volume 9888 of *Lecture Notes in Computer Science*, pages 57–70. Springer, 2016. doi:10.1007/978-3-662-53426-7_5.

37 Mohsen Ghaffari and Jason Li. Improved distributed algorithms for exact shortest paths. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 431–444. ACM, 2018. doi:10.1145/3188745.3188948.

38 Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. A deterministic almost-tight distributed algorithm for approximating single-source shortest paths. In Daniel Wichs and Yishay Mansour, editors, *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 489–498. ACM, 2016. doi:10.1145/2897518.2897638.

39 Stephan Holzer, David Peleg, Liam Roditty, and Roger Wattenhofer. Brief announcement: Distributed 3/2-approximation of the diameter. In *Distributed Computing - 28th International Symposium*, DISC 2014, pages 562–564, 2014.

40 Stephan Holzer, David Peleg, Liam Roditty, and Roger Wattenhofer. Distributed 3/2-approximation of the diameter. In *Distributed Computing - 28th International Symposium, DISC 2014, Austin, TX, USA, October 12-15, 2014. Proceedings*, pages 562–564, 2014. URL: http://link.springer.com/content/pdf/bbm%3A978-3-662-45174-8%2F1.pdf.

41 Stephan Holzer and Nathan Pinsker. Approximation of Distances and Shortest Paths in the Broadcast Congest Clique. In Emmanuelle Anceaume, Christian Cachin, and Maria Potop-Butucaru, editors, *19th International Conference on Principles of Distributed Systems (OPODIS 2015)*, volume 46 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1–16, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.OPODIS.2015.6.

42 Stephan Holzer and Roger Wattenhofer. Optimal distributed all pairs shortest paths and applications. In *ACM Symposium on Principles of Distributed Computing, PODC '12, Funchal, Madeira, Portugal, July 16-18, 2012*, pages 355–364, 2012. doi:10.1145/2332432.2332504.

43 Chien-Chung Huang, Danupon Nanongkai, and Thatchaphol Saranurak. Distributed exact weighted all-pairs shortest paths in õ($n^{5/4}$) rounds. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 168–179. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.24.

44 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.

45 T. S. Jayram, Ravi Kumar, and D. Sivakumar. Two applications of information complexity. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, June 9-11, 2003, San Diego, CA, USA*, pages 673–682, 2003. doi:10.1145/780542.780640.

46 Naoki Katoh and Kazuo Iwano. Finding k farthest pairs and k closest/farthest bichromatic pairs for points in the plane. *Int. J. Comput. Geometry Appl.*, 5:37–51, 1995. doi:10.1142/S0218195995000040.

47 Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, New York, NY, USA, 1997.

48 Christoph Lenzen and Boaz Patt-Shamir. Fast routing table construction using small messages: extended abstract. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 381–390. ACM, 2013. doi:10.1145/2488608.2488656.

49 Christoph Lenzen, Boaz Patt-Shamir, and David Peleg. Distributed distance computation and routing with small messages. *Distributed Comput.*, 32(2):133–157, 2019. doi:10.1007/s00446-018-0326-6.

50 Christoph Lenzen and David Peleg. Efficient distributed source detection with limited bandwidth. In *ACM Symposium on Principles of Distributed Computing, PODC '13, Montreal, QC, Canada, July 22-24, 2013*, pages 375–382, 2013. doi:10.1145/2484239.2484262.

51 Jason Li and Merav Parter. Planar diameter via metric compression. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 152–163, 2019. doi:10.1145/3313276.3316358.

52 Danupon Nanongkai. Distributed approximation algorithms for weighted shortest paths. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 565–573. ACM, 2014. doi:10.1145/2591796.2591850.

53 David Peleg, Liam Roditty, and Elad Tal. Distributed algorithms for network diameter and girth. In *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part II*, pages 660–672, 2012. doi:10.1007/978-3-642-31585-5_58.

54 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. doi:10.1137/S0097539700369740.

55 A.A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992. doi:10.1016/0304-3975(92)90260-M.

56 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the International Congress of Mathematicians (ICM 2018)*, pages 3447–3487, 2018. doi:10.1142/9789813272880_0188.

57 Andrew Chi-Chih Yao. On constructing minimum spanning trees in k-dimensional spaces and related problems. *SIAM J. Comput.*, 11(4):721–736, 1982. doi:10.1137/0211059.