Distributed Verification and Hardness of Distributed Approximation *

Atish Das Sarma[†]
Google Research
Mountain View, USA.
dassarma@google.com

Stephan Holzer
ETH Zurich
Zurich, Switzerland
stholzer@tik.ee.ethz.ch

Liah Kor
Weizmann Institute of Science
Rehovot, Israel
liah.kor@weizmann.ac.il

Amos Korman[‡]
CNRS & LIAFA, Univ. Paris 7
Paris, France
amos.korman@liafa.jussieu.fr

Danupon Nanongkai
University of Vienna & Georgia
Institute of Technology
Austria & USA
danupon@gmail.com

Gopal Pandurangan[§]
Nanyang Technological
University & Brown University
Singapore & USA
gopalpandurangan@gmail.com

David Peleg
Weizmann Institute of Science
Rehovot, Israel
david.peleg@weizmann.ac.il

Roger Wattenhofer
ETH Zurich
Zurich, Switzerland
wattenhofer@tik.ee.ethz.ch

ABSTRACT

We study the *verification* problem in distributed networks, stated as follows. Let H be a subgraph of a network G where each vertex of G knows which edges incident on it are in H . We would like to verify whether H has some properties, e.g., if it is a tree or if it is connected (every node knows in the end of the process whether H has the specified property or not). We would like to perform this verification in a decentralized fashion via a distributed algorithm. The time complexity of verification is measured as the number of rounds of distributed communication.

In this paper we initiate a systematic study of distributed verification, and give almost tight lower bounds on the running time of distributed verification algorithms for many

fundamental problems such as connectivity, spanning connected subgraph, and $s - t$ cut verification. We then show applications of these results in deriving strong unconditional time lower bounds on the *hardness of distributed approximation* for many classical optimization problems including minimum spanning tree, shortest paths, and minimum cut. Many of these results are the first non-trivial lower bounds for both exact and approximate distributed computation and they resolve previous open questions. Moreover, our unconditional lower bound of approximating minimum spanning tree (MST) subsumes and improves upon the previous hardness of approximation bound of Elkin [STOC 2004] as well as the lower bound for (exact) MST computation of Peleg and Rubinfeld [FOCS 1999]. Our result implies that there can be no distributed approximation algorithm for MST that is significantly faster than the current exact algorithm, for *any* approximation factor.

Our lower bound proofs show an interesting connection between communication complexity and distributed computing which turns out to be useful in establishing the time complexity of exact and approximate distributed computation of many problems.

*A full version of this paper is available as [5] at <http://arxiv.org/abs/1011.3049>

[†]Part of the work done while at Georgia Institute of Technology.

[‡]Supported by the ANR projects ALADDIN and PROSE and by the INRIA project GANG. Also supported by a France-Israel cooperation grant ("Mutli-Computing" project) from the France Ministry of Science and Israel Ministry of Science.

[§]Supported in part by the following grants: Nanyang Technological University grant M58110000, US NSF grant CCF-1023166, and a grant from the US-Israeli Binational Science Foundation (BSF).

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'11, June 6–8, 2011, San Jose, California, USA.

Copyright 2011 ACM 978-1-4503-0691-1/11/06 ...\$10.00.

Categories and Subject Descriptors

C.2.4 [Computer Systems Organization]: Computer-Communication Networks—*Distributed Systems*; F.0 [Theory of Computation]: General; G.2.2 [Mathematics of Computing]: Discrete Mathematic—*Graph Theory*

General Terms

Algorithms, Theory

Keywords

Distributed Algorithms, Graph Algorithms, Lower Bound, Time Complexity, Communication Complexity, Minimum Spanning Tree, Shortest Path

1. INTRODUCTION

Large and complex networks, such as the human society, the Internet, or the brain, are being studied intensely by different branches of science. Each individual node in such a network can directly communicate only with its neighboring nodes. Despite being restricted to such *local* communication, the network itself should work towards a *global* goal, i.e., it should organize itself, or deliver a service.

In this work we investigate the possibilities and limitations of distributed/decentralized computation, i.e., to what degree local information is sufficient to solve global tasks. Many tasks can be solved entirely via local communication, for instance, how many friends of friends one has. Research in the last 30 years has shown that some classic combinatorial optimization problems such as matching, coloring, dominating set, or approximations thereof can be solved using small (i.e., polylogarithmic) local communication. For example, a maximal independent set can be computed in time $O(\log n)$ [22], but not in time $\Omega(\sqrt{\log n / \log \log n})$ [17] (n is the network size). This lower bound even holds if message sizes are unbounded.

However “many” important optimization problems are “global” problems from the distributed computation point of view. To count the total number of nodes, to determining the diameter of the system, or to compute a spanning tree, information necessarily must travel to the farthest nodes in a system. If exchanging a message over a single edge costs one time unit, one needs $\Omega(D)$ time units to compute the result, where D is the network diameter. If message size was unbounded, one can simply collect all the information in $O(D)$ time, and then compute the result. Hence, in order to arrive at a realistic problem, we need to introduce communication limits, i.e., each node can exchange messages with each of its neighbors in each step of a synchronous system, but each message can have at most B bits (typically B is small, say $O(\log n)$). However, to compute a spanning tree, even single-bit messages are enough, as one can simply breadth-first-search the graph in time $O(D)$ and this is optimal [24].

But, can we *verify* whether an existing spanning tree indeed is a correct spanning tree?! In this paper we show that this is not generally possible in $O(D)$ time – instead one needs $\Omega(\sqrt{n} + D)$ time. (Thus, in contrast to traditional non-distributed complexity, verification is harder than computation in the distributed world!). Our paper is more general, as we show interesting lower and upper bounds (these are almost tight) for a whole selection of verification problems. Furthermore, we show a key application of studying such verification problems to proving strong unconditional time lower bounds on exact and approximate distributed computation for many classical problems.

1.1 Technical Background and Previous Work

Distributed Computing. Consider a synchronous network of processors with unbounded computational power. The network is modeled by an undirected n -vertex graph, where vertices model the processors and edges model the links between the processors. The processors (henceforth, vertices) communicate by exchanging messages via the links (henceforth, edges). The vertices have limited global knowledge, in particular, each of them has its own local perspective of the network (a.k.a graph), which is confined to its im-

mediate neighborhood. The vertices may have to compute (cooperatively) some global function of the graph, such as a spanning tree (ST) or a minimum spanning tree (MST), via communicating with each other and running a distributed algorithm designed for the task at hand. There are several measures to analyze the performance of such algorithms, a fundamental one being the running time, defined as the worst-case number of *rounds* of distributed communication. This measure naturally gives rise to a complexity measure of problems, called the *time complexity*. On each round at most B bits can be sent through each edge in each direction, where B is the bandwidth parameter of the network. The design of efficient algorithms for this model (henceforth, the B model), as well as establishing lower bounds on the time complexity of various fundamental graph problems, has been the subject of an active area of research called (locality-sensitive) *distributed computing* (see [24] and references therein.)

Distributed Algorithms, Approximation, and Hardness. Much of the initial research focus in the area of distributed computing was on designing algorithms for solving problems exactly, e.g., distributed algorithms for ST, MST, and shortest paths are well-known [24, 23]. Over the last few years, there has been interest in designing distributed algorithms that provide approximate solutions to problems. This area is known as *distributed approximation*. One motivation for designing such algorithms is that they can run faster or have better communication complexity albeit at the cost of providing suboptimal solution. This can be especially appealing for resource-constrained and dynamic networks (such as sensor or peer-to-peer networks). For example, there is not much point in having an optimal algorithm in a dynamic network if it takes too much time, since the topology could have changed by that time. For this reason, in the distributed context, such algorithms are well-motivated even for network optimization problems that are not NP-hard, e.g., minimum spanning tree, shortest paths etc. There is a large body of work on distributed approximation algorithms for various classical graph optimization problems (e.g., see the surveys by Elkin [7] and Dubhashi et al. [6], and the work of [14] and the references therein).

While a lot of progress has been made in the design of distributed approximation algorithms, the same has not been the case with the theory of lower bounds on the approximability of distributed problems, i.e., *hardness* of distributed approximation. There are some inapproximability results that are based on lower bounds on the time complexity of the exact solution of certain problems and on integrality of the objective functions of these problems. For example, a fundamental result due to Linial [20] says that 3-coloring an n -vertex ring requires $\Omega(\log^* n)$ time. In particular, it implies that any $3/2$ -approximation protocol for the vertex-coloring problem requires $\Omega(\log^* n)$ time. On the other hand, one can state inapproximability results assuming that vertices are computationally limited; under this assumption, any NP-hardness inapproximability result immediately implies an analogous result in the distributed model. However, the above results are not interesting in the distributed setting, as they provide no new insights on the roles of locality and communication [10].

There are but a few significant results currently known on the hardness of distributed approximation. Perhaps the first important result was presented for the MST problem by

Elkin in [10]. Specifically, he showed strong *unconditional* lower bounds (i.e., ones that do not depend on complexity-theoretic assumptions) for distributed approximate MST (more on this result below). Later, Kuhn, Moscibroda, and Watzenhofer [17] showed lower bounds on time approximation trade-offs for several problems.

1.2 Distributed Verification

The above discussion summarized two major research aspects in distributed computing, namely studying distributed algorithms and lower bounds for (1) exact and (2) approximate solutions to various problems. The third aspect — that turns out to have remarkable applications to the first two — called *distributed verification*, is the main subject of the current paper. In distributed verification, we want to efficiently check whether a given subgraph of a network has a specified property via a distributed algorithm¹. Formally, given a graph $G = (V, E)$, a subgraph $H = (V, E')$ with $E' \subseteq E$, and a predicate Π , it is required to decide whether H satisfies Π (i.e., when the algorithm terminates, every node knows whether H satisfies Π). The predicate Π may specify statements such as “ H is connected” or “ H is a spanning tree” or “ H contains a cycle”. (Each vertex in G knows which of its incident edges (if any) belong to H .) The goal is to study bounds on the time complexity of distributed verification. The time complexity of the verification algorithm is measured with respect to parameters of G (in particular, its size n and diameter D), independently from H .

We note that verification is different from construction problems, which have been the traditional focus in distributed computing. Indeed, distributed algorithms for constructing spanning trees, shortest paths, and other problems have been well studied ([24, 23]). However, the corresponding verification problems have received much less attention. To the best of our knowledge, the only distributed verification problem that has received some attention is the MST (i.e., verifying if H is a MST); the recent work of Kor et al. [15] gives a $\Omega(\sqrt{n}/B + D)$ deterministic lower bound on distributed verification of MST, where D is the diameter of the network G . That paper also gives a matching upper bound (see also [16]). Note that distributed *construction* of MST has rather similar lower and upper bounds [25, 11]. Thus in the case of the MST problem, verification and construction have the same time complexity. We later show that the above result of Kor et al. is subsumed by the results of this paper, as we show that verifying *any* spanning tree takes so much time.

Motivations. The study of distributed verification has two main motivations. The first is understanding the complexity of verification versus construction. This is obviously a central question in the traditional RAM model, but here we want to focus on the same question in the distributed model. Unlike in the centralized setting, it turns out that verification is *not* in general easier than construction in the distributed setting! In fact, as was indicated earlier, distributively verifying a spanning tree turns out to be harder than constructing it in the worst case. Thus understanding the complexity of verification in the distributed model is also important. Second, from an algorithmic point of view, for some problems, understanding the verification problem

can help in solving the construction problem or showing the inherent limitations in obtaining an efficient algorithm. In addition to these, there is yet another motivation that emerges from this work: We show that distributed verification leads to showing *strong unconditional lower bounds on distributed computation (both exact and approximate)* for a variety of problems, many hitherto unknown. For example, we show that establishing a lower bound on the spanning connected subgraph verification problem leads to establishing lower bounds for the minimum spanning tree, shortest path tree, minimum cut etc. Hence, studying verification problems may lead to proving hardness of approximation as well as lower bounds for exact computation for new problems.

1.3 Our Contributions

In this paper, our main contributions are two fold. First, we initiate a systematic study of *distributed verification*, and give almost tight uniform lower bounds on the running time of distributed verification algorithms for many fundamental problems. Second, we make progress in establishing strong hardness results on the distributed approximation of many classical optimization problems. Our lower bounds also apply seamlessly to exact algorithms. We next state our main results (the precise theorem statements are in the respective sections as mentioned below).

1. Distributed Verification. We show a lower bound of $\Omega(\sqrt{n}/(B \log n) + D)$ for many verification problems in the B model, including *spanning connected subgraph*, *s-t connectivity*, *cycle-containment*, *bipartiteness*, *cut*, *least-element list*, and *s - t cut* (cf. Section 4). These bounds apply to *randomized* algorithms as well, and clearly hold also for asynchronous networks. Moreover, it is important to note that our lower bounds apply even to graphs of small diameter ($D = O(\log n)$). (Indeed, the problems studied in this paper are “global” problems, i.e., the network diameter of G imposes an inherent lower bound on the time complexity.)

Additionally, we show that another fundamental problem, namely, the spanning tree verification problem (i.e., verifying whether H is a spanning tree) has the same lower bound of $\Omega(\sqrt{n}/(B \log n) + D)$ (cf. Section 6). However, this bound applies to only deterministic algorithms. This result strengthens the lower bound result of MST verification by Kor et al. [15]. Moreover, we note the interesting fact that although finding a spanning tree (e.g., a breadth-first tree) can be done in $O(D)$ rounds [24], verifying if a given subgraph is a spanning tree requires $\tilde{\Omega}(\sqrt{n} + D)$ rounds! Thus the verification problem for spanning trees is harder than its construction in the distributed setting. This is in contrast to this well-studied problem in the centralized setting. Apart from the spanning tree problem, we also show deterministic lower bounds for other verification problems, including *Hamiltonian cycle* and *simple path*.

Our lower bounds are almost tight as we show that there exist algorithms that run in $O(\sqrt{n} \log^* n + D)$ rounds (assuming $B = O(\log n)$) for all the verification problems addressed here (cf. Full version).

2. Bounds on Hardness of Distributed Approximation. An important consequence of our verification lower bound is that it leads to lower bounds for exact and approximate distributed computation. We show the unconditional time lower bound of $\Omega(\sqrt{n}/(B \log n) + D)$ for approximat-

¹Such problems have been studied in the sequential setting, e.g., Tarjan[27] studied verification of MST.

ing many optimization problems, including MST, shortest $s-t$ path, shortest path tree, and minimum cut (Section 5). The important point to note is that the above lower bound applies for *any* approximation ratio $\alpha \geq 1$. Thus the same bound holds for exact algorithms also ($\alpha = 1$). All these hardness bounds hold for randomized algorithms. As in our verification lower bounds, these bounds apply even to graphs of small ($O(\log n)$) diameter. Figure 1 summarizes our lower bounds for various diameters.

Our results improve over previous ones (e.g., Elkin’s lower bound for approximate MST and shortest path tree [10]) and subsumes some well-established exact bounds (e.g., Peleg and Rubinfeld lower bound for MST [25]) as well as shows new strong bounds (both for exact and approximate computation) for many other problems (e.g., minimum cut), thus answering some questions that were open earlier (see the survey by Elkin [7]).

The new lower bound for approximating MST simplifies and improves upon the previous $\Omega(\sqrt{n/(\alpha B \log n)} + D)$ lower bound by Elkin [10], where α is the approximation factor. [10] showed a *tradeoff* between the running time and the approximation ratio of MST. Our result shows that approximating MST requires $\Omega(\sqrt{n/(B \log n)} + D)$ rounds, regardless of α . Thus our result shows that there is actually no trade-off, since there can be no distributed approximation algorithm for MST that is significantly faster than the current exact algorithm [19, 9], for any approximation factor $\alpha > 1$.

2. OVERVIEW OF TECHNICAL APPROACH

We prove our lower bounds by establishing an interesting connection between communication complexity and distributed computing. Our lower bound proofs consider the family of graphs evolved through a series of papers in the literature [10, 21, 25]. However, while previous results [10, 25] rely on counting the number of states to analyze the *mailing problem* (along with some sophisticated techniques for the variant, called *corrupted mail problem*, in the case of approximation algorithm lower bounds) and use Yao’s method [30] (with appropriate input distributions) to get lower bounds for randomized algorithms, our results are achieved using the following three steps of simple reductions, as follows.

(Section 3) First, we reduce the lower bounds of problems in the standard communication complexity model [18] to the lower bounds of the equivalent problems in the “distributed version” of communication complexity. Specifically, we relate the *communication* lower bound from the standard communication complexity model [18] to compute some appropriately chosen function f , to the distributed *time* complexity lower bound for computing the same function in a specially chosen graph G . In the standard model, Alice and Bob can communicate directly (via a bidirectional edge of bandwidth one). In the distributed model, we assume that Alice and Bob are some vertices of G and they together wish to compute the function f using the communication graph G . The choice of graph G is critical. We use a graph called $G(\Gamma, d, p)$ (parameterized by Γ , d and p) that was first used in [10]. We show a reduction from the standard model to the distributed model, the proof of which relies on certain observations similar to those used in previous results (e.g., [25]).

(Section 4) The connection established in the first step allows us to bypass the state counting argument and Yao’s

method, and reduces our task in proving lower bounds of verification problems to merely picking the right function f to reduce from. The function f that is useful in showing our randomized lower bounds is the *set disjointness function*, which is the quintessential problem in the world of communication complexity with applications to diverse areas and has been studied for decades (see a recent survey in [3]). Following the result well known in communication complexity [18], we show that the distributed version of this problem has an $\Omega(\sqrt{n/(B \log n)})$ lower bound on graphs of small diameter. We then reduce this problem to the verification problems using simple reductions similar to those used in data streams [12]. The set disjointness function yields randomized lower bounds and works for many problems (see Fig. 2), but it does not reduce to certain other problems such as spanning tree. To show lower bounds for these and a few other problems, we use a different function f called *equality*. However, this reduction yields only deterministic lower bounds for the corresponding verification problems.

(Section 5) Finally, we reduce the verification problem to hardness of distributed approximation for a variety of problems to show that the same lower bounds hold for approximation algorithms as well. For this, we use a reduction whose idea is similar to one used to prove hardness of approximating TSP (Traveling Salesman Problem) on general graphs (see, e.g., [28]): We convert a verification problem to an optimization problem by introducing edge weight in such a way that there is a large gap between the optimal values for the cases where H satisfies, or does not satisfy a certain property. This technique is surprisingly simple, yet yields strong unconditional hardness bounds — many hitherto unknown, left open (e.g., minimum cut [7] and some that improve over known ones (e.g., MST and shortest path tree) [10]. As mentioned earlier, our approach shows that approximating MST by *any* factor needs $\tilde{\Omega}(\sqrt{n})$ time, while the previous result due to Elkin gave a bound that depends on α (the approximation factor), i.e. $\tilde{\Omega}(\sqrt{n/\alpha})$, using more sophisticated techniques.

Fig. 2 summarizes these reductions that will be proved in this paper. Due to the space constraint, we focus on the proofs towards the lower bound of approximating minimum spanning tree in the main paper. Other results can be found in the Full version.

3. FROM COMMUNICATION COMPLEXITY TO DISTRIBUTED COMPUTING

Consider the following problem. There are two parties that have unbounded computational power. Each party receives a b -bit string, for some integer $b \geq 1$, denoted by $\bar{x}$ and $\bar{y}$ in $\{0, 1\}^b$. They both want to together compute $f(\bar{x}, \bar{y})$ for some boolean function $f : \{0, 1\}^b \times \{0, 1\}^b \rightarrow \{0, 1\}$. We consider two models of communication.

- *Direct communication*: This is the standard model in communication complexity. Two parties can communicate via a bidirectional edge of bandwidth one. We call the party receiving $\bar{x}$ *Alice*, and the other party *Bob*. At the end of the process, Bob will output $f(\bar{x}, \bar{y})$.
- *Communication through network $G(\Gamma, d, p)$* : Two parties are distinct vertices in a B model distributed network, called $G(\Gamma, d, p)$, for some parameters Γ , d , and p ; the network has $\Theta(\Gamma d^p)$ vertices and a diameter of

Diameter D	Previous lower bound for MST and shortest-path tree [10] (for exact algorithms, use $\alpha = 1$)	New lower bound for MST, shortest-path tree and all problems in Fig. 2.
$n^\delta, 0 < \delta < 1/2$	$\Omega(\sqrt{\frac{n}{\alpha B}})$	$\Omega(\sqrt{\frac{n}{B}})$
$\Theta(\log n)$	$\Omega(\sqrt{\frac{n}{\alpha B \log n}})$	$\Omega(\sqrt{\frac{n}{B \log n}})$
Constant ≥ 3	$\Omega((\frac{n}{\alpha B})^{\frac{1}{2} - \frac{1}{2D-2}})$	$\Omega((\frac{n}{B})^{\frac{1}{2} - \frac{1}{2D-2}})$
4	$\Omega((\frac{n}{\alpha B})^{1/3})$	$\Omega((\frac{n}{B})^{1/3})$
3	$\Omega((\frac{n}{\alpha B})^{1/4})$	$\Omega((\frac{n}{B})^{1/4})$

Figure 1: Lower bounds of randomized α -approximation algorithms on graphs of various diameters. Bounds in the first column are for the MST and shortest path tree problems [10] while those in the second column are for these problems and many problems listed in Fig. 2. We note that these bounds almost match the $O(\sqrt{n \log^* n} + D)$ upper bound for the MST problem [11, 19] and are independent of the approximation-factor α .

Figure 2: Problems and reductions between them to obtain randomized and deterministic lower bounds. For all problems, we obtain lower bounds as in Fig. 1

$\Theta(2p + 2)$. (This network was first defined in [10] and described below.) We denote the vertex receiving $\bar{x}$ by s and the vertex receiving $\bar{y}$ by r . At the end of the process, r will output $f(\bar{x}, \bar{y})$.

We consider time lower bounds for *public coin randomized algorithms* under both models. In particular, we assume that all parties (Alice and Bob in the first model and all vertices in $G(\Gamma, d, p)$ in the second model) share a random bit string of infinite length. For any $\epsilon \geq 0$, we say that a randomized algorithm $\mathcal{A}$ is ϵ -error if for any input, it outputs the correct answer with probability at least $1 - \epsilon$, where the probability is over all possible random bit strings. The running time of $\mathcal{A}$, denoted by $T_{\mathcal{A}}$, is the number of rounds in the worst case (over all inputs and random strings). Let $R_{\epsilon}^{cc-pub}(f)$ and $R_{\epsilon}^{G(\Gamma, d, p), s, r}(f)$ denote the best time complexity of ϵ -error algorithms in the models of direct communication and communication through graph $G(\Gamma, d, p)$, respectively. We are particularly interested in the case where we pick b to be Γ (for any choice of Γ). The rest of this section is devoted to showing that if there is a fast ϵ -error algorithm for computing f on $G(\Gamma, d, p)$, then there is a fast ϵ -error algorithm for Alice and Bob to compute f .

Theorem 3.1 *For any Γ , d , p , B , $\epsilon \geq 0$, and function*

$$f : \{0, 1\}^{\Gamma} \times \{0, 1\}^{\Gamma} \rightarrow \{0, 1\}, \text{ if } R_{\epsilon}^{G(\Gamma, d, p), s, r}(f) < \frac{d^p - 1}{2}$$

then

$$R_{\epsilon}^{cc-pub}(f) \leq 2dpBR_{\epsilon}^{G(\Gamma, d, p), s, r}(f).$$

We first describe the graph $G(\Gamma, d, p)$ with parameters Γ , d and p and distinct vertices s and r .

The graph $G(\Gamma, d, p)$ [10]: The two basic units in the construction are *paths* and a *tree*. There are Γ paths, denoted by $\mathcal{P}^1, \mathcal{P}^2, \dots, \mathcal{P}^{\Gamma}$, each having d^p vertices, i.e., for $\ell = 1, 2, \dots, \Gamma$,

$$V(\mathcal{P}^{\ell}) = \{v_0^{\ell}, \dots, v_{d^p-1}^{\ell}\} \text{ and}$$

$$E(\mathcal{P}^{\ell}) = \{(v_i^{\ell}, v_{i+1}^{\ell}) \mid 0 \leq i < d^p - 1\}.$$

There is a tree, denoted by $\mathcal{T}$ having depth p where each non-leaf vertex has d children (thus, there are d^p leaf vertices). We denote the vertices of $\mathcal{T}$ at level ℓ from left to right by $u_0^{\ell}, \dots, u_{d^{\ell}-1}^{\ell}$ (so, u_0^0 is the root of $\mathcal{T}$ and $u_0^p, \dots, u_{d^p-1}^p$ are the leaves of $\mathcal{T}$). For any ℓ and j , the leaf vertex u_j^p is connected to the corresponding path vertex v_j^{ℓ} by a *spoke* edge (u_j^p, v_j^{ℓ}) . Finally, we set the two special vertices (which will receive input strings $\bar{x}$ and $\bar{y}$) as $s = u_0^p$ and $r = u_{d^p-1}^p$. Fig. 3(a) depicts this graph. The number of vertices in $G(\Gamma, d, p)$ its diameter are analyzed in [10], as follows.

Figure 3: (a) The Graph $G(\Gamma, d, p)$ (here, $d = 2$). (b) Example of H for the spanning connected subgraph problem (marked with thick red edges) when $\bar{x} = 0 \dots 10$ and $\bar{y} = 1 \dots 00$.

Lemma 3.2 [10] *The number of vertices in $G(\Gamma, d, p)$ is $n = \Theta(\Gamma m)$ and its diameter is $2p + 2$.*

Terminologies: For $1 \leq i \leq \lfloor (d^p - 1)/2 \rfloor$, define the i -left and the i -right of the path $\mathcal{P}^\ell$ as

$$L_i(\mathcal{P}^\ell) = \{v_j^\ell \mid j \leq d^p - 1 - i\} \quad \text{and} \quad R_i(\mathcal{P}^\ell) = \{v_j^\ell \mid j \geq i\},$$

respectively. Define the i -left of the tree $\mathcal{T}$, denoted by $L_i(\mathcal{T})$, as the union of set $S = \{u_j^p \mid j \leq d^p - 1 - i\}$ and all ancestors of all vertices in S in $\mathcal{T}$. Similarly, the i -right $R_i(\mathcal{T})$ of the tree $\mathcal{T}$ is the union of set $S = \{u_j^p \mid j \geq i\}$ and all ancestors of all vertices in S . Now, the i -left and i -right sets of $G(\Gamma, d, p)$ are the union of those left and right sets,

$$L_i = \bigcup_{\ell} L_i(\mathcal{P}^\ell) \cup L_i(\mathcal{T}) \quad \text{and} \quad R_i = \bigcup_{\ell} R_i(\mathcal{P}^\ell) \cup R_i(\mathcal{T}).$$

For $i = 0$, the definition is slightly different; we set $L_0 = V \setminus \{r\}$ and $R_0 = V \setminus \{s\}$. See Fig. 3(a).

Let $\mathcal{A}$ be any deterministic distributed algorithm run on graph $G(\Gamma, d, p)$ for computing a function f . Fix any input strings $\bar{x}$ and $\bar{y}$ given to s and r respectively. Let $\varphi_{\mathcal{A}}(\bar{x}, \bar{y})$ denote the execution of $\mathcal{A}$ on $\bar{x}$ and $\bar{y}$. Denote the state of the vertex v at the end of round t during the execution $\varphi_{\mathcal{A}}(\bar{x}, \bar{y})$ by $\sigma_{\mathcal{A}}(v, t, \bar{x}, \bar{y})$. In two different executions $\varphi_{\mathcal{A}}(\bar{x}, \bar{y})$ and $\varphi_{\mathcal{A}}(\bar{x}', \bar{y}')$, a vertex reaches the same state at time t (i.e., $\sigma_{\mathcal{A}}(v, t, \bar{x}, \bar{y}) = \sigma_{\mathcal{A}}(v, t, \bar{x}', \bar{y}')$) if and only if it receives the same sequence of messages on each of its incoming links.

For a given set of vertices $U = \{v_1, \dots, v_\ell\} \subseteq V$, a configuration

$$C_{\mathcal{A}}(U, t, \bar{x}, \bar{y}) = \langle \sigma_{\mathcal{A}}(v_1, t, \bar{x}, \bar{y}), \dots, \sigma_{\mathcal{A}}(v_\ell, t, \bar{x}, \bar{y}) \rangle$$

is a vector of the states of the vertices of U at the end of round t of the execution $\varphi_{\mathcal{A}}(\bar{x}, \bar{y})$. We note the following crucial observation used in [25] and many later results.

Observation 3.3 *For any set $U \subseteq U' \subseteq V$, $C_{\mathcal{A}}(U, t, \bar{x}, \bar{y})$ can be uniquely determined by $C_{\mathcal{A}}(U', t-1, \bar{x}, \bar{y})$ and all messages sent to U from $V \setminus U'$ at time t .*

PROOF. Recall that the state of each vertex v in U can be uniquely determined by its state $\sigma_{\mathcal{A}}(v, t-1, \bar{x}, \bar{y})$ at time $t-1$ and the messages sent to it at time t . Moreover, the messages sent to v from vertices inside U' can be determined by $C_{\mathcal{A}}(U', t, \bar{x}, \bar{y})$. Thus if the messages sent from vertices in $V \setminus U'$ are given then we can determine all messages sent to U at time t and thus we can determine $C_{\mathcal{A}}(U, t, \bar{x}, \bar{y})$. $\square$

From now on, to simplify notation, when $\mathcal{A}$, $\bar{x}$ and $\bar{y}$ are clear from the context, we use C_{L_t} and C_{R_t} to denote $C_{\mathcal{A}}(L_t, t, \bar{x}, \bar{y})$ and $C_{\mathcal{A}}(R_t, t, \bar{x}, \bar{y})$, respectively. The lemma below states that C_{L_t} (C_{R_t} , respectively) can be determined by $C_{L_{t-1}}$ ($C_{R_{t-1}}$, respectively) and dp messages generated by some vertices in R_{t-1} (L_{t-1} respectively) at time t . It essentially follows from Observation 3.3 and an observation that there are at most d^p edges linking between vertices in $V \setminus R_{t-1}$ ($V \setminus L_{t-1}$ respectively) and vertices in R_t (L_t respectively).

Lemma 3.4 *Fix any deterministic algorithm $\mathcal{A}$ and input strings $\bar{x}$ and $\bar{y}$. For any $0 < t < (d^p - 1)/2$, there exist functions g_L and g_R , B -bit messages $M_1^{L_{t-1}}, \dots, M_{d^p}^{L_{t-1}}$ sent by some vertices in L_{t-1} at time t , and B -bit messages $M_1^{R_{t-1}}, \dots, M_{d^p}^{R_{t-1}}$ sent by some vertices in R_{t-1} at time t such that*

$$C_{L_t} = g_L(C_{L_{t-1}}, M_1^{R_{t-1}}, \dots, M_{d^p}^{R_{t-1}}) \quad (1)$$

$$C_{R_t} = g_R(C_{R_{t-1}}, M_1^{L_{t-1}}, \dots, M_{d^p}^{L_{t-1}}). \quad (2)$$

PROOF. We prove Eq. (2) only. (Eq. (1) is proved in exactly the same way.) Observe that all neighbors of all path vertices in R_t are in R_{t-1} . Similarly, all neighbors of all leaf vertices in $V(\mathcal{T}) \cap R_t$ are in R_{t-1} . Moreover, for any non-leaf tree vertex u_i^ℓ (for some ℓ and i), if u_i^ℓ is in R_t then its parent and vertices $u_{i+1}^\ell, u_{i+2}^\ell, \dots, u_{d^\ell-1}^\ell$ are in R_{t-1} . For any $\ell < p$ and t , let $u^\ell(R_t)$ denote the leftmost vertex that is at level ℓ of $\mathcal{T}$ and in R_t , i.e., $u^\ell(R_t) = u_i^\ell$ where i is such that $u_i^\ell \in R_t$ and $u_{i-1}^\ell \notin R_t$. (For example, in Fig. 3(a), $u^{p-1}(R_1) = u_0^{p-1}$ and

$u^{p-1}(R_2) = u_1^{p-1}$.) Finally, observe that for any i and ℓ , if u_i^ℓ is in R_t then all children of u_i^ℓ are in R_t (otherwise, all children of u_i^ℓ are not in R_t and so is u_i^ℓ , a contradiction). Thus, all edges linking between vertices in R_t and $V \setminus R_{t-1}$ are in the following form: $(u^\ell(R_t), u')$ for some ℓ and child u' of $u^\ell(R_t)$.

Setting $U' = R_{t-1}$ and $U = R_t$ in Observation 3.3, we have that C_{R_t} can be uniquely determined by $C_{R_{t-1}}$ and messages sent to $u^\ell(R_t)$ from its children in $V \setminus R_{t-1}$. Note that each of these messages contains at most B bits since they correspond to a message sent on an edge in one round.

Observe further that, for any $t < (d^p - 1)/2$, $V \setminus R_{t-1} \subseteq L_{t-1}$ since L_{t-1} and R_{t-1} share some path vertices. Moreover, each $u^\ell(R_t)$ has d children. Therefore, if we let $M_1^{L_{t-1}}, \dots, M_{dp}^{L_{t-1}}$ be the messages sent from children of $u^0(R_t)$, $u^1(R_t), \dots, u^{p-1}(R_t)$ in $V \setminus R_{t-1}$ to their parents (note that if there are less than dp such messages then we add some empty messages) then we can uniquely determine C_{R_t} by $C_{R_{t-1}}$ and $M_1^{L_{t-1}}, \dots, M_{dp}^{L_{t-1}}$. Eq. (2) thus follows. $\square$

Using the above lemma, we can now prove Theorem 3.1.

PROOF OF THEOREM 3.1. Let f be the function in the theorem statement. Let $\mathcal{A}_\epsilon$ be any ϵ -error distributed algorithm for computing f on graph $G(\Gamma, d, p)$. Fix a random string $\bar{r}$ used by $\mathcal{A}_\epsilon$ (shared by all vertices in $G(\Gamma, d, p)$) and consider the *deterministic* algorithm $\mathcal{A}$ run on the input of $\mathcal{A}_\epsilon$ and the fixed random string $\bar{r}$. Let $T_{\mathcal{A}}$ be the worst case running time of algorithm $\mathcal{A}$ (over all inputs). We only consider $T_{\mathcal{A}} < (d^p - 1)/2$, as assumed in the theorem statement. We show that Alice and Bob, when given $\bar{r}$ as the public random string, can simulate $\mathcal{A}$ using $2dpT_{\mathcal{A}}$ communication bits, as follows.

Alice and Bob make $T_{\mathcal{A}}$ iterations of communications. Initially, Alice computes C_{L_0} which depends only on $\bar{x}$. Bob also computes C_{R_0} which depends only on $\bar{y}$. In each iteration $t > 0$, we assume that Alice and Bob know $C_{L_{t-1}}$ and $C_{R_{t-1}}$, respectively, before the iteration starts. Then, Alice and Bob will exchange at most $2dpB$ bits so that Alice and Bob know C_{L_t} and C_{R_t} , respectively, at the end of the iteration.

To do this, Alice sends to Bob the messages $M_1^{L_{t-1}}, \dots, M_{dp}^{L_{t-1}}$ as in Lemma 3.4. Alice can generate these messages since she knows $C_{L_{t-1}}$ (by assumption). Then, Bob can compute C_{R_t} using Eq. (2) in Lemma 3.4. Similarly, Bob sends dp messages to Alice and Alice can compute C_{L_t} . They exchange at most $2dpB$ bits in total in each iteration since there are $2dp$ messages, each of B bits, exchanged.

After $T_{\mathcal{A}}$ iterations, Bob knows $C(R_{T_{\mathcal{A}}}, T_{\mathcal{A}}, \bar{x}, \bar{y})$. In particular, he knows the output of $\mathcal{A}$ (output by r) since he knows the state of r after $\mathcal{A}$ terminates. He thus outputs the output of r .

Since $\mathcal{A}_\epsilon$ is ϵ -error, the probability (over all possible shared random strings) that $\mathcal{A}$ outputs the correct value of $f(\bar{x}, \bar{y})$ is at least $1 - \epsilon$. Therefore, the communication protocol run by Alice and Bob is ϵ -error as well. Moreover, Alice and Bob communicates at most $2dpBT_{\mathcal{A}}$ bits. The theorem follows. $\square$

4. RANDOMIZED LOWER BOUNDS FOR DISTRIBUTED VERIFICATION

In this section, we present randomized lower bounds for many verification problems for graph of various diameters, as shown in Fig. 1.

The general theorem is below. For brevity, in the main section of the paper, we prove the theorem only for the spanning connected subgraph verification problem. This will be useful later in proving many hardness of approximation results. In this problem, we want to verify whether H is connected and spans all nodes of G , i.e., every node in G is incident to some edge in H . Definitions of other problems and proofs of their lower bounds are in Full version.

Theorem 4.1 *For any $p \geq 1$, $B \geq 1$, and $n \in \{2^{2p+1}pB, 3^{2p+1}pB, \dots\}$, there exists a constant $\epsilon > 0$ such that any ϵ -error distributed algorithm for any of the following problems requires $\Omega((n/(pB))^{\frac{1}{2} - \frac{1}{2(2p+1)}})$ time on some $\Theta(n)$ -vertex graph of diameter $2p+2$ in the B model: Spanning connected subgraph, connectivity, s - t connectivity, k -components, bipartiteness, cycle containment, e -cycle containment, cut, s - t cut, least-element list $[4, 14]$, and edge on all paths.*

In particular, for graphs with diameter $D = 4$, we get $\Omega((n/B)^{1/3})$ lower bound and for graphs with diameter $D = \log n$ we get $\Omega(\sqrt{n/(B \log n)})$. Similar analysis also leads to a $\Omega(\sqrt{n/B})$ lower bound for graphs of diameter n^δ for any $\delta > 0$, and $\Omega((n/B)^{1/4})$ lower bound for graphs of diameter 3 using the same analysis as in [10]. We note that the lower bound holds even in the public coin model where every vertex shares a random string. To prove the theorem, we need the lower bound for computing *set disjointness function*.

Definition 4.2 (Set Disjointness function) *Given two b -bit strings $\bar{x}$ and $\bar{y}$, the set disjointness function, denoted by $\text{disj}(\bar{x}, \bar{y})$, is defined to be 1 if the inner product $\langle \bar{x}, \bar{y} \rangle$ is 0 (i.e., $x_i = 0$ or $y_i = 0$ for every $1 \leq i \leq b$) and 0 otherwise. We refer to the problem of computing disj function on $G(\Gamma, d, p)$ on Γ -bit input strings given to s and r by $\text{DISJ}(G(\Gamma, d, p), s, r, \Gamma)$.*

The following lemma is a consequence of Theorem 3.1 and the communication complexity lower bound of computing disj .

Lemma 4.3 *For any Γ, d, p , there exists a constant $\epsilon > 0$ such that any ϵ -error algorithm solving $\text{DISJ}(G(\Gamma, d, p), s, r, \Gamma)$ requires $\Omega(\min(d^p, \frac{\Gamma}{dpB}))$ time.*

PROOF. If $R_\epsilon^{G(\Gamma, d, p), s, r}(\text{disj}) \geq (d^p - 1)/2$ then $R_\epsilon^{G(\Gamma, d, p), s, r}(\text{disj}) = \Omega(d^p)$ and we are done. Otherwise, Theorem 3.1 implies that $R_\epsilon^{cc-pub}(\text{disj}) \leq 2dpB \cdot R_\epsilon^{G(\Gamma, d, p), s, r}(\text{disj})$. Now we use the fact that $R_\epsilon^{cc-pub}(\text{disj}) = \Omega(\Gamma)$ for the function disj on Γ -bit inputs, for some $\epsilon > 0$ [1, 13, 2, 26] (also see [18, Example 3.22] and references therein). It follows that $R_\epsilon^{G(\Gamma, d, p), s, r}(\text{disj}) = \Omega(\Gamma/(dpB))$. $\square$

The lower bound of spanning connected subgraph verification essentially follows from the following lemma.

Lemma 4.4 *For any Γ , $d \geq 2$ and p , there exists a constant $\epsilon > 0$ such that any ϵ -error distributed algorithm for*

spanning connected subgraph verification on graph $G(\Gamma, d, p)$ can be used to solve the $\text{DISJ}(G(\Gamma, d, p), s, t, \Gamma)$ problem on $G(\Gamma, d, p)$ with the same time complexity.

PROOF. Consider an ϵ -error algorithm $\mathcal{A}$ for the spanning connected subgraph verification problem, and suppose that we are given an instance of the $\text{DISJ}(G(\Gamma, d, p), s, t, \Gamma)$ problem with input strings $\bar{x}$ and $\bar{y}$. We use $\mathcal{A}$ to solve this instance of set disjointness problem as follows.

First, we mark all path edges and tree edges as participating in H . All spoke edges are marked as not participating in subgraph H , except those incident to s and r for which we do the following: For each bit x_i , $1 \leq i \leq \Gamma$, vertex s indicates that the spoke edge (s, v_0^i) participates in H if and only if $x_i = 0$. Similarly, for each bit y_i , $1 \leq i \leq \Gamma$, vertex r indicates that the spoke edge (r, v_{dp-1}^i) participates in H if and only if $y_i = 0$. (See Fig. 3(b).)

Note that the participation of all edges, except those incident to s and r , is decided independently of the input. Moreover, one round is sufficient for s and r to inform their neighbors the participation of edges incident to them. Hence, one round is enough to construct H . Then, algorithm $\mathcal{A}$ is started.

Once algorithm $\mathcal{A}$ terminates, vertex r determines its output for the set disjointness problem by stating that both input strings are disjoint if and only if spanning connected subgraph verification algorithm verified that the given subgraph H is indeed a spanning connected subgraph.

Observe that H is a spanning connected subgraph if and only if for all $1 \leq i \leq \Gamma$ at least one of the edges (s, v_0^i) and (r, v_{dp-1}^i) is in H ; thus, by the construction of H , H is a spanning connected subgraph if and only if the input strings $\bar{x}, \bar{y}$ are disjoint, i.e., for every i either $x_i = 0$ or $y_i = 0$. Hence the resulting algorithm has correctly solved the given instance of the set disjointness problem. $\square$

Using Lemma 4.3, we obtain the following result.

Corollary 4.5 *For any Γ, d, p , there exists a constant $\epsilon > 0$ such that any ϵ -error algorithm for spanning connected subgraph verification problem requires $\Omega(\min(d^p, \frac{\Gamma}{dpB}))$ time on some $\Theta(\Gamma d^p)$ -vertex graph of diameter $2p + 2$.*

In particular, if we consider $\Gamma = d^{p+1}pB$ then

$$\Omega(\min(d^p, \Gamma/(dpB))) = \Omega(d^p).$$

Moreover, by Lemma 3.2, $G(d^{p+1}pB, d, p)$ has $n = \Theta(d^{2p+1}pB)$ vertices and thus the lower bound $\Omega(d^p)$ becomes

$$\Omega((n/(pB))^{\frac{1}{2} - \frac{1}{2(2p+1)}}).$$

Theorem 4.1 (for the case of spanning connected subgraph) follows.

5. HARDNESS OF DISTRIBUTED APPROXIMATION

In this section we show a time lower bound of $\Omega(\sqrt{n/(B \log n)})$ for approximation algorithms of many problems. For distributed approximation problems such as MST, we assume that a weight function $\omega : E \rightarrow \mathbb{R}^+$ associated with the graph assigns a nonnegative real weight $\omega(e)$ to each edge $e = (u, v) \in E$. Initially, the weight $\omega(e)$ is known only to the adjacent vertices, u and v . We assume that the edge

weights are bounded by a polynomial in n (the number of vertices). It is assumed that B is large enough to allow the transmission of any edge weight in a single message.

We show the hardness of distributed approximation for many problems, as in the theorem below. For brevity, we only prove the theorem for the minimum spanning tree problem here. Definitions and proofs of other problems can be found in Full version.

Theorem 5.1 *For any polynomial function $\alpha(n)$, numbers $p, B \geq 1$, and $n \in \{2^{2p+1}pB, 3^{2p+1}pB, \dots\}$, there exists a constant $\epsilon > 0$ such that any $\alpha(n)$ -approximation ϵ -error distributed algorithm for any of the following problems requires $\Omega((\frac{n}{pB})^{\frac{1}{2} - \frac{1}{2(2p+1)}})$ time on some $\Theta(n)$ -vertex graph of diameter $2p + 2$ in the B model: minimum spanning tree [10, 25], shortest s - t path, s -source distance [8], s -source shortest path tree [10], minimum cut [7], minimum s - t cut, maximum cut, minimum routing cost spanning tree [29], shallow-light tree [24], and generalized Steiner forest [14].*

Recall that in the minimum spanning tree problem, we are given a connected graph G and we want to compute the minimum spanning tree (i.e., the spanning tree of minimum weight). At the end of the process each vertex knows which edges incident to it are in the output tree.

Recall the following standard notions of an *approximation algorithm*. We say that a randomized algorithm $\mathcal{A}$ is α -approximation ϵ -error if, for any input instance $\mathcal{I}$, algorithm $\mathcal{A}$ outputs a solution that is at most α times the optimal solution of $\mathcal{I}$ with probability at least $1 - \epsilon$. Therefore, in the minimum spanning tree, an α -approximation ϵ -error algorithm should output a number that is at most α times the total weight of the minimum spanning tree, with probability at least $1 - \epsilon$.

PROOF OF THEOREM 5.1. (This proof is only for the case of minimum spanning tree.) Let $\mathcal{A}_\epsilon$ be an $\alpha(n)$ -approximation ϵ -error algorithm for the minimum spanning tree problem. We show that $\mathcal{A}_\epsilon$ can be used to solve the spanning connected subgraph verification problem using the same running time.

To do this, construct a weight function on edges in G , denoted by ω , by assigning weight 1 to all edges in H and $n\alpha(n)$ to all other edges. Note that constructing ω does not need any communication since each vertex knows which edges incident to it are in H . Now we find the weight W of the minimum spanning tree using $\mathcal{A}_\epsilon$ and announce that H is a spanning connected subgraph if and only if W is less than $n\alpha(n)$.

Now we show that the weighted graph (G, ω) has a spanning tree of weight less than $n\alpha(n)$ if and only if H is a spanning connected subgraph of G and thus the algorithm above is correct: Suppose that H is a spanning connected subgraph. Then, there is a spanning tree that is a subgraph of H and has weight $n - 1 < n\alpha(n)$. Thus the minimum spanning tree has weight less than $n\alpha(n)$. Conversely, suppose that H is not a spanning connected subgraph. Then, any spanning tree must contain an edge not in H . Therefore, any spanning tree has weight at least $n\alpha(n)$ as claimed. $\square$

Our MST lower bound here matches the lower bound of exact MST algorithms and improves the lower bound of

$\Omega(\sqrt{\frac{n}{\alpha B}})$ by Elkin [10]. Our lower bound for s -source distance complements the results in [8].

6. DETERMINISTIC LOWER BOUNDS

We show the following lower bound of deterministic algorithms for problems listed in the theorem below. We note that our lower bound of spanning tree verification simplifies and generalizes the lower bound of *minimum* spanning tree verification shown in [15]. Due to space constraints, problem definitions and proofs are placed in Full version.

Theorem 6.1 *For any $p, B \geq 1$, and $n \in \{2^{2p+1}pB, 3^{2p+1}pB, \dots\}$, any deterministic distributed algorithm for any of the following problems requires $\Omega((\frac{n}{pB})^{\frac{1}{2} - \frac{1}{2(2p+1)}})$ time on some $\Theta(n)$ -vertex graph of diameter $O(2p+2)$ in the B model: Hamiltonian cycle, spanning tree, and simple path verification.*

7. CONCLUSION

We initiate the systematic study of verification problems in the context of distributed network algorithms and present a uniform lower bound for several problems. We also show how these verification bounds can be used to obtain lower bounds on exact and approximation algorithms for many problems.

Several problems remain open. A general direction for extending all of this work is to study similar verification problems in special classes of graphs, e.g., a complete graph. A few specific open questions include proving better lower or upper bounds for the problems of shortest s - t path, single-source distance computation, shortest path tree, s - t cut, minimum cut. (Some of these problems were also asked in [7].) Also, showing randomized bounds for Hamiltonian path, spanning tree, and simple path verification remains open.

8. REFERENCES

- [1] L. Babai, P. Frankl, and J. Simon. Complexity classes in communication complexity theory (preliminary version). In *FOCS*, pages 337–347, 1986.
- [2] Z. Bar-Yossef, T. S. Jayram, R. Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68(4):702–732, 2004. Also in *FOCS'02*.
- [3] A. Chattopadhyay and T. Pitassi. The Story of Set Disjointness. *SIGACT News*, 41(3):59–85, 2010.
- [4] E. Cohen. Size-Estimation Framework with Applications to Transitive Closure and Reachability. *J. Comput. Syst. Sci.*, 55(3):441–453, 1997. Also in *FOCS'94*.
- [5] A. Das Sarma, S. Holzer, L. Kor, A. Korman, D. Nanongkai, G. Pandurangan, D. Peleg, and R. Wattenhofer. Distributed verification and hardness of distributed approximation. *CoRR*, abs/1011.3049, 2010.
- [6] D. P. Dubhashi, F. Grandioni, and A. Panconesi. Distributed Algorithms via LP Duality and Randomization. In *Handbook of Approximation Algorithms and Metaheuristics*. Chapman and Hall/CRC, 2007.
- [7] M. Elkin. Distributed approximation: a survey. *SIGACT News*, 35(4):40–57, 2004.
- [8] M. Elkin. Computing almost shortest paths. *ACM Transactions on Algorithms*, 1(2):283–323, 2005. Also in *PODC'01*.
- [9] M. Elkin. A faster distributed protocol for constructing a minimum spanning tree. *J. Comput. Syst. Sci.*, 72(8):1282–1308, 2006. Also in *SODA'04*.
- [10] M. 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. Also in *STOC'04*.
- [11] J. A. Garay, S. Kutten, and D. Peleg. A Sublinear Time Distributed Algorithm for Minimum-Weight Spanning Trees. *SIAM J. Comput.*, 27(1):302–316, 1998. Also in *FOCS '93*.
- [12] M. R. Henzinger, P. Raghavan, and S. Rajagopalan. Computing on data streams. In J. M. Abello and J. S. Vitter, editors, *External memory algorithms*, pages 107–118. American Mathematical Society, Boston, MA, USA, 1999.
- [13] B. Kalyanasundaram and G. Schnitger. The Probabilistic Communication Complexity of Set Intersection. *SIAM J. Discrete Math.*, 5(4):545–557, 1992.
- [14] M. Khan, F. Kuhn, D. Malkhi, G. Pandurangan, and K. Talwar. Efficient distributed approximation algorithms via probabilistic tree embeddings. In *PODC*, pages 263–272, 2008.
- [15] L. Kor, A. Korman, and D. Peleg. Tight bounds for distributed MST verification. *STACS*, 2011.
- [16] A. Korman and S. Kutten. Distributed verification of minimum spanning trees. *Distributed Computing*, 20(4):253–266, 2007. Also in *PODC'06*.
- [17] F. Kuhn, T. Moscibroda, and R. Wattenhofer. What cannot be computed locally! In *PODC*, pages 300–309, 2004.
- [18] E. Kushilevitz and N. Nisan. *Communication complexity*. Cambridge University Press, New York, NY, USA, 1997.
- [19] S. Kutten and D. Peleg. Fast Distributed Construction of Small k -Dominating Sets and Applications. *J. Algorithms*, 28(1):40–66, 1998. Also in *PODC'95*.
- [20] N. Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
- [21] Z. Lotker, B. Patt-Shamir, and D. Peleg. Distributed MST for constant diameter graphs. *Distributed Computing*, 18(6):453–460, 2006. Also in *PODC'01*.
- [22] M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1053, 1986. Also in *STOC'85*.
- [23] N. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
- [24] D. Peleg. *Distributed computing: a locality-sensitive approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [25] D. Peleg and V. Rubinovich. A Near-Tight Lower Bound on the Time Complexity of Distributed Minimum-Weight Spanning Tree Construction. *SIAM J. Comput.*, 30(5):1427–1442, 2000. Also in *FOCS'99*.
- [26] A. A. Razborov. On the Distributional Complexity of

- Disjointness. *Theor. Comput. Sci.*, 106(2):385–390, 1992. Also in ICALP’90.
- [27] R. E. Tarjan. Applications of path compression on balanced trees. *J. ACM*, 26(4):690–715, 1979.
- [28] V. V. Vazirani. *Approximation Algorithms*. Springer, July 2001.
- [29] B. Y. Wu, G. Lancia, V. Bafna, K.-M. Chao, R. Ravi, and C. Y. Tang. A polynomial-time approximation scheme for minimum routing cost spanning trees. *SIAM J. Comput.*, 29(3):761–778, 1999. Also in SODA’98.
- [30] A. C.-C. Yao. Probabilistic Computations: Toward a Unified Measure of Complexity. In *FOCS*, pages 222–227, 1977.