

A tight unconditional lower bound on distributed random walk computation

Danupon Nanongkai ^{*} Atish Das Sarma [†] Gopal Pandurangan [‡]

Abstract

We consider the problem of performing a random walk in a distributed network. Given bandwidth constraints, the goal of the problem is to minimize the number of rounds required to obtain a random walk sample. Das Sarma et al. [PODC'10] show that a random walk of length ℓ on a network of diameter D can be performed in $\tilde{O}(\sqrt{\ell D} + D)$ time. A major question left open is whether there exists a faster algorithm, especially whether the multiplication of $\sqrt{\ell}$ and \sqrt{D} is necessary.

In this paper, we show a tight unconditional lower bound on the time complexity of distributed random walk computation. Specifically, we show that for any n , D , and $D \leq \ell \leq (n/(D^3 \log n))^{1/4}$, performing a random walk of length $\Theta(\ell)$ on an n -node network of diameter D requires $\Omega(\sqrt{\ell D} + D)$ time. This bound is *unconditional*, i.e., it holds for any (possibly randomized) algorithm. To the best of our knowledge, this is the first lower bound that the diameter plays a role of multiplicative factor. Our bound shows that the algorithm of Das Sarma et al. is time optimal.

Our proof technique introduces a new connection between *bounded-round* communication complexity and distributed algorithm lower bounds with D as a trade-off parameter, strengthening the previous study by Das Sarma et al. [STOC'11]. In particular, we make use of the bounded-round communication complexity of the pointer chasing problem. Our technique can be of independent interest and may be useful in showing non-trivial lower bounds on the complexity of other fundamental distributed computing problems.

Keywords: Random walks, Distributed algorithms, Lower bounds, Communication complexity.

^{*}College of Computing, Georgia Institute of Technology, Atlanta, GA 30332, USA. E-mail: danupon@cc.gatech.edu

[†]Google Research, Google Inc., Mountain View, USA. E-mail: dassarma@google.com.

[‡]Division of Mathematical Sciences, Nanyang Technological University, Singapore 637371 and Department of Computer Science, Brown University, Providence, RI 02912, USA. E-mail: gopalpandurangan@gmail.com. Supported in part by NSF grant CCF-1023166 and by a grant from the United States-Israel Binational Science Foundation (BSF).

1 Introduction

The random walk plays a central role in computer science, spanning a wide range of areas in both theory and practice. The focus of this paper is on performing a random walk in distributed networks, in particular, decentralized algorithms for performing a random walk in arbitrary networks. The random walk is used as an integral subroutine in a wide variety of network applications ranging from token management [18, 4, 8], small-world routing [20], search [29, 1, 7, 17, 25], information propagation and gathering [5, 19], network topology construction [17, 22, 23], expander testing [12], constructing random spanning trees [6, 3, 2], distributed construction of expander networks [22], and peer-to-peer membership management [15, 30]. For more applications of random walks to distributed networks, see e.g. [11]. Motivated by the wide applicability of the random walk, [10, 11] consider the running time of performing a random walk on the synchronous distributed model. We now explain the model and problems before describing previous work and our results.

1.1 Distributed Computing Model

Consider a synchronous network of processors with unbounded computational power. The network is modeled by an undirected connected n -node multi-graph, where nodes model the processors and edges model the links between the processors. The processors (henceforth, nodes) communicate by exchanging messages via the links (henceforth, edges). The nodes 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 immediate neighborhood.

There are several measures to analyze the performance of algorithms on this model, 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 $O(\log n)$ bits can be sent through each edge in each direction. This is a standard model of distributed computation known as the *CONGEST* model [27] and has been attracting a lot of research attention during last two decades (e.g., see [27] and the references therein). We note that our result also holds on the *CONGEST*(B) model, where on each round at most B bits can be sent through each edge in each direction (see the remark after Theorem 1.1). We ignore this parameter to make the proofs and theorem statements simpler.

1.2 Problems

The basic problem is *computing a random walk where destination outputs source*, defined as follows. We are given a network $G = (V, E)$ and a source node $s \in V$. The goal is to devise a distributed algorithm such that, in the end, some node v outputs the ID of s , where v is a destination node picked according to the probability that it is the destination of a random walk of length ℓ starting at s . We assume the standard random walk where, in each step, an edge is taken from the current node v with probability proportional to $1/d(v)$ where $d(v)$ is the degree of v . Our goal is to output a true random sample from the ℓ -walk distribution starting from s .

For clarity, observe that the following naive algorithm solves the above problem in $O(\ell)$ rounds. The walk of length ℓ is performed by sending a token for ℓ steps, picking a random neighbor in each step. Then, the destination node v of this walk outputs the ID of s . The main objective of distributed random walk problem is to perform such sampling with significantly less number of rounds, i.e., in time that is sublinear in ℓ . On the other hand, we note that it can take too much time (as much as $\Theta(|E| + D)$ time) in the *CONGEST* model to collect all the topological information at the source node (and then computing the walk locally).

The following variations were also previously considered.

1. *Computing a random walk where source outputs destination:* The problem is almost the same as above except that, in the end, the source has to output the ID of the destination. This version is useful in nodes learning the topology of their surrounding networks and related applications such as a decentralized algorithm for estimating the mixing time [11].
2. *Computing a random walk where nodes know their positions:* Instead of outputting the ID of source or destination, we want each node to know its position(s) in the random walk. That is, if v_1, v_2, \dots, v_ℓ (where $v_1 = s$) is the result random walk starting at s , we want each node v_j in the walk to know the number j at the end of the process. This version is used to construct a random spanning tree in [11].

1.3 Previous work and our result

A key purpose of the random walk in many network applications is to perform node sampling. While the sampling requirements in different applications vary, whenever a true sample is required from a random walk of certain steps, typically all applications perform the walk naively — by simply passing a token from one node to its neighbor: thus to perform a random walk of length ℓ takes time linear in ℓ . Das Sarma et al. [10] showed this is not a time-optimal strategy and the running time can be made sublinear in ℓ , i.e., performing a random walk of length ℓ on an n -node network of diameter D can be done in $\tilde{O}(\ell^{2/3}D^{1/3})$ time where \tilde{O} hides polylog n . Subsequently, Das Sarma et al. [11] improved this bound to $\tilde{O}(\sqrt{\ell D} + D)$ which holds for all three versions of the problem.

There are two key motivations for obtaining sublinear time bounds. The first is that in many algorithmic applications, walks of length significantly greater than the network diameter are needed. For example, this is necessary in two applications presented in [11], namely distributed computation of a random spanning tree (RST) and computation of mixing time. More generally, many real-world communication networks (e.g., ad hoc networks and peer-to-peer networks) have relatively small diameter, and random walks of length at least the diameter are usually performed for many sampling applications, i.e., $\ell \gg D$.

The second motivation is understanding the time complexity of distributed random walks. Random walk is essentially a global problem which requires the algorithm to “traverse” the entire network. Classical “global” problems include the minimum spanning tree, shortest path etc. Network diameter is an inherent lower bound for such problems. Problems of this type raise the basic question whether n (or ℓ as the case here) time is essential or is the network diameter D , the inherent parameter. As pointed out in the seminal work of [16], in the latter case, it would be desirable to design algorithms that have a better complexity for graphs with low diameter. While both upper and lower bounds of time complexity of many “global” problems are known (see, e.g., [9]), the status of the random walk problem is still wide open.

A preliminary attempt to show a random walk lower bound is presented in [11]. They consider a restricted class of algorithms, where each message sent between nodes must be in the form of an interval of numbers. Moreover, a node is allowed to send a number or an interval containing it only after it receives such number. For this very restricted class of algorithms, a lower bound of $\Omega(\sqrt{\ell} + D)$ is shown [11] for the version where every node must know their positions in the end of the computation. While this lower bound shows a potential limitation of random walk algorithms, it has many weaknesses. First, it does not employ an information theoretic argument and thus does not hold for all types of algorithm. In stead, it assumes that the algorithms must send messages as intervals and thus holds only for a small class of algorithms, which does not even cover all deterministic algorithms. Second, the lower bound holds only for the version where nodes know their position(s), thereby leaving lower bounds for the other two random walk versions completely open. More importantly, there is still a gap of \sqrt{D} between lower and upper bounds, leaving a

question whether there is a faster algorithm.

Motivated by these applications, past results, and open problems, we consider the problem of finding lower bounds for random walk computation. In this work, we show an *unconditional* lower bound of $\Omega(\sqrt{\ell D} + D)$ for all three versions of the random walk computation problem. This means that the algorithm in [11] is optimal for all three variations. In particular, we show the following theorem.

Theorem 1.1. *For any n , D and ℓ such that $D \leq \ell \leq (n/(D^3 \log n))^{1/4}$, there exists a family of n -node networks of diameter D such that performing a random walk (any of the three versions) of length $\Theta(\ell)$ on these networks requires $\Omega(\sqrt{\ell D} + D)$ rounds.*

We note that our lower bound of $\Omega(\sqrt{\ell D} + D)$ also holds for the general $\mathcal{CONGEST}(B)$ model, where each edge has bandwidth B instead of $O(\log n)$, as long as $\ell \leq (n/(D^3 B))^{1/4}$. Moreover, one can also show a lower bound on simple graphs by subdividing edges in the network used in the proof and double the value of ℓ .

1.4 Techniques and proof overview

Our main approach relies on enhancing the connection between communication complexity and distributed algorithm lower bounds first studied in [9]. It has been shown in [9] that a fast distributed algorithm for computing a function can be converted into a two-party communication protocol with small message complexity to compute the same function. In other words, the communication complexity lower bounds implies the time complexity lower bounds of distributed algorithms. This result is then used to prove lower bounds on many *verification problems*. (In the verification problems, we are given H , a subgraph of the 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.) The lower bounds of verification problems are then used to prove lower bounds of approximation algorithms for many graph problems. Their work, however, does not make progress on achieving *any* unconditional lower bound on the random walk problem.

Further, while this approach has been successfully used to show lower bounds for several problems in terms of network size (i.e., n), it is not clear how to apply them to random walk computation. All the lower bounds previously shown are for *optimization* problems for well-defined metrics - for e.g. computing a minimum spanning tree. Random walk computation, on the other hand, is not deterministic; the input requires parameters such as the length of the walk ℓ and even if the source node and ℓ are fixed, the *solution* (i.e. the walk) is not uniquely determined. While even other problems, such as MST, can have multiple solutions, for optimization problems, *verification* is well-defined. It is not clear what it even means to verify whether a random walk is *correct*. For this reason, proving a lower bound of random walk computation through verification problems seems impossible.

Additionally, in terms of the theoretical bound we obtain, a key difficulty in our result is to introduce the graph parameter diameter (i.e., D) into the lower bound multiplied by ℓ . A crucial shortcoming in extending previous work in this regard is that the relationship between communication complexity and distributed computing shown in [9] does not depend on the network diameter D at all. In fact, such a relationship might not exist since the result in [9] is tight for some functions.

To overcome these obstacles, we consider a variation of communication complexity called *r-round two-party communication complexity*, which has been successfully used in, e.g., circuit complexity and data stream computation (see, e.g., [14, 26]). We obtain a new connection showing that a fast distributed algorithm for computing a function can be converted to a two-party communication protocol with a small message complexity *and number of rounds* to compute the same function. Moreover, the larger the network diameter is, the smaller the number of rounds will be. To obtain

this result one need to deal with a more involved proof; for example, the new proof does not seem to work for the networks previously considered [9, 13, 24, 28] and thus we need to introduce a novel network called $G(\Gamma, \kappa, \Lambda)$. This result and related definitions are stated and proved in Section 2.

A particular communication complexity result that we will use is that of Nisan and Wigderson [26] for the r -round pointer chasing problem. Using the connection established in Section 2, we derive a lower bound of any distributed algorithms for solving the pointer chasing problem on a distributed network. This result is in Section 3.

Finally, we prove Theorem 1.1 from the lower bound result in Section 3. The main idea, which was also used in [11], is to construct a network that has the same structure as $G(\Gamma, \kappa, \Lambda)$ (thus has the same diameter and number of nodes) but different edge capacities (depending on the input) so that a random walk follows a desired path (which is unknown) with high probability. This proof is in Section 4.

2 From bounded-round communication complexity to distributed algorithm lower bounds

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 function $f : \{0, 1\}^b \times \{0, 1\}^b \rightarrow \mathbb{R}$. At the end of the computation, the party receiving \bar{y} has to output the value of $f(\bar{x}, \bar{y})$. We consider two models of communication.

- *r -round direct communication:* This is a variant of the standard model in communication complexity (see [26] and references therein). Two parties can communicate via a bidirectional edge of unlimited bandwidth. We call the party receiving \bar{x} *Alice*, and the other party *Bob*. Two parties communicate in *rounds* where each round Alice sends a message (of any size) to Bob followed by Bob sending a message to Alice.
- *Communication through network $G(\Gamma, \kappa, \Lambda)$:* Two parties are distinct nodes in a distributed network $G(\Gamma, \kappa, \Lambda)$, for some integers Γ and Λ and real κ ; all networks in $G(\Gamma, \kappa, \Lambda)$ have $\Theta(\kappa\Gamma\Lambda^\kappa)$ nodes and a diameter of $\Theta(\kappa\Lambda)$. (This network is described below.) We denote the nodes receiving \bar{x} and \bar{y} by s and t , respectively.

We consider the *public coin randomized algorithms* under both models. In particular, we assume that all parties (Alice and Bob in the first model and all nodes in $G(\Gamma, \kappa, \Lambda)$ 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. In the first model, we focus on the message complexity, i.e., the total number of bits exchanged between Alice and Bob, denoted by $R_\epsilon^{r-cc-pub}(f)$. In the second model, we focus on the running time, denoted by $R_\epsilon^{G(\Gamma, \kappa, \Lambda), s, t}(f)$.

Before we describe $G(\Gamma, \kappa, \Lambda)$ in detail, we note the following characteristics which will be used in later sections. An essential part of $G(\Gamma, \kappa, \Lambda)$ consists of Γ *paths*, denoted by $\mathcal{P}^1, \dots, \mathcal{P}^\Gamma$ and nodes s and t (see Fig. 1). Every edge induced by this subgraph has infinitely many copies. The leftmost and rightmost nodes of each path are adjacent to s and t respectively. Ending nodes on the same side of the path (i.e., leftmost or rightmost nodes) are adjacent to each other. The following properties of $G(\Gamma, \kappa, \Lambda)$ follow from the construction of $G(\Gamma, \kappa, \Lambda)$ described in Section 2.1.

Lemma 2.1. *For any $\Gamma \geq 1$, $\kappa \geq 1$ and $\Lambda \geq 2$, network $G(\Gamma, \kappa, \Lambda)$ has $\Theta(\Gamma\kappa\Lambda^\kappa)$ nodes. Each of its path \mathcal{P}^i has $\Theta(\kappa\Lambda^\kappa)$ nodes. Its diameter is $\Theta(\kappa\Lambda)$.*




Figure 1: An example of $G(\Gamma, \kappa, \Lambda)$ where $\kappa = 2.5$ and $\Lambda = 2$. The dashed edges (in red) have one copy while other edges have infinitely many copies. Observe that $\phi_{10} = 6$ while $\phi'_{10} = 4$.

Proof. It follows from the construction of $G(\Gamma, \kappa, \Lambda)$ in Section 2.1 that the number of nodes in each path \mathcal{P}^i is $\sum_{j=-\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}}^{\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}} \phi'_j = \Theta(\kappa \Lambda^\kappa)$ (cf. Eq. (1)). Since there are Γ paths, the number of nodes in all paths is $\Theta(\Gamma \kappa \Lambda^\kappa)$. Each highway \mathcal{H}^i has $2\lceil \kappa \rceil \Lambda^i + 1$ nodes. Therefore, there are $\sum_{i=1}^{\lfloor \kappa \rfloor} (2\lceil \kappa \rceil \Lambda^i + 1)$ nodes in the highways. For $\Lambda \geq 2$, the last quantity is $\Theta(\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor})$. Hence, the total number of nodes is $\Theta(\Gamma \kappa \Lambda^\kappa)$.

To analyze the diameter of $G(\Gamma, \kappa, \Lambda)$, observe that each node on any path \mathcal{P}^i can reach a node in highway $\mathcal{H}^{\lfloor \kappa \rfloor}$ by traveling through $O(\kappa \Lambda)$ nodes in \mathcal{P}^i . Moreover, any node in highway \mathcal{H}^i can reach a node in highway \mathcal{H}^{i-1} by traveling through $O(\Lambda)$ nodes in \mathcal{H}^i . Finally, there are $O(\kappa \Lambda)$ nodes in \mathcal{H}^1 . Therefore, every node can reach any other node in $O(\kappa \Lambda)$ steps by traveling through \mathcal{H}^1 . Note that this upper bound is tight since the distance between s and t is $\Omega(\kappa \Lambda)$. \square

The rest of this section is devoted to prove the following theorem, which strengthens Theorem 3.1 in [9]. It says that if there is a fast ϵ -error algorithm for computing function f on any network $G(\Gamma, \kappa, \Lambda)$, then there is a fast bounded-round ϵ -error algorithm for Alice and Bob to compute f . Moreover, the number of rounds depends on the diameter of $G(\Gamma, \kappa, \Lambda)$ (which is $\Theta(\kappa \Lambda)$).

Theorem 2.2. *For any integers $\Gamma \geq 1$, $\Lambda \geq 2$, real $\kappa \geq 1$ and function $f : \{0, 1\}^b \times \{0, 1\}^b \rightarrow \mathbb{R}$, for any b , if $R_\epsilon^{G(\Gamma, \kappa, \Lambda), s, t}(f) \leq \kappa \Lambda^\kappa$ then f can be computed by a $\frac{8R_\epsilon^{G(\Gamma, \kappa, \Lambda), s, t}(f)}{\kappa \Lambda}$ -round direct communication protocol using at most $(2\kappa \log n)R_\epsilon^{G(\Gamma, \kappa, \Lambda), s, t}(f)$ communication bits in total. In other words,*

$$R_\epsilon^{\frac{8R_\epsilon^{G(\Gamma, \kappa, \Lambda), s, t}(f)}{\kappa \Lambda} - cc - pub}(f) \leq (2\kappa \log n)R_\epsilon^{G(\Gamma, \kappa, \Lambda), s, t}(f).$$

2.1 Description of $G(\Gamma, \kappa, \Lambda)$

$G(\Gamma, \kappa, \Lambda)$ has three parameters, a real $\kappa \geq 1$ and two integers $\Gamma \geq 1$ and $\Lambda \geq 2$. The two basic units in the construction of $G(\Gamma, \kappa, \Lambda)$ are *highways* and *paths* (see Fig. 1).

Highways. There are $\lfloor \kappa \rfloor$ highways, denoted by $\mathcal{H}^1, \mathcal{H}^2, \dots, \mathcal{H}^{\lfloor \kappa \rfloor}$. The highway \mathcal{H}^i is a path of $2\lceil \kappa \rceil \Lambda^i + 1$ nodes, i.e.,

$$V(\mathcal{H}^i) = \{h_0^i, h_{\pm \Lambda^{\lfloor \kappa \rfloor - i}}^i, h_{\pm 2\Lambda^{\lfloor \kappa \rfloor - i}}^i, h_{\pm 3\Lambda^{\lfloor \kappa \rfloor - i}}^i, \dots, h_{\pm \lceil \kappa \rceil \Lambda^i}^i\} \quad \text{and,}$$

$$E(\mathcal{H}^i) = \{(h_{-(j+1)\Lambda^{\lfloor \kappa \rfloor - i}}^i, h_{-j\Lambda^{\lfloor \kappa \rfloor - i}}^i), (h_{j\Lambda^{\lfloor \kappa \rfloor - i}}^i, h_{(j+1)\Lambda^{\lfloor \kappa \rfloor - i}}^i) \mid 0 \leq j < \lceil \kappa \rceil \Lambda^i\}.$$

We connect the highways by adding edges between nodes of the same subscripts, i.e., for any $0 < i \leq \lfloor \kappa \rfloor$ and $-\lceil \kappa \rceil \Lambda^i \leq j \leq \lceil \kappa \rceil \Lambda^i$, there is an edge between $h_{j\Lambda^{\lfloor \kappa \rfloor - i}}^i$ and $h_{j\Lambda^{\lfloor \kappa \rfloor - i}}^{i+1}$.

For any $-\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor} \leq j \leq \lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}$, define

$$\phi_j = \left\lfloor \frac{|j|}{\Lambda^{\lfloor \kappa \rfloor - 1}} \right\rfloor + 1.$$

Note that ϕ_j can be viewed as the number of nodes in \mathcal{H}_1 with subscripts between 0 and j , i.e.,

$$\phi_j = \begin{cases} |\{h_{j'}^1 \mid 0 \leq j' \leq j\}| & \text{if } j \geq 0 \\ |\{h_{j'}^1 \mid j \leq j' \leq 0\}| & \text{if } j < 0. \end{cases}$$

For any $j \geq 0$, we also define

$$\phi'_j = \phi'_{-j} = \min \left\{ \phi_j, \max(1, \lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor} - \sum_{j' > j} \phi_{j'}) \right\}.$$

We use ϕ'_j to specify the number of nodes in the paths (defined next), i.e., each path will have $\sum_{j=-\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}}^{\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}} \phi'_j$ nodes. Note that

$$\sum_{j=-\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}}^{\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}} \phi'_j = \Theta(\kappa \Lambda^{\lfloor \kappa \rfloor}). \quad (1)$$

Paths. There are Γ paths, denoted by $\mathcal{P}^1, \mathcal{P}^2, \dots, \mathcal{P}^\Gamma$. To construct each path, we first construct its subpaths as follows. For each node $h_j^{\lfloor \kappa \rfloor}$ in $\mathcal{H}^{\lfloor \kappa \rfloor}$ and any $0 < i \leq \Gamma$, we create a subpath of \mathcal{P}^i , denoted by \mathcal{P}_j^i , having ϕ'_j nodes. Denote nodes in \mathcal{P}_j^i in order by $v_{j,1}^i, v_{j,2}^i, \dots, v_{j,\phi'_j}^i$. We connect these paths together to form \mathcal{P}_j^i , i.e., for any $j \geq 0$, we create edges $(v_{j,\phi'_j}^i, v_{j+1,1}^i)$ and $(v_{-j,\phi'_{-j}}^i, v_{-(j+1),1}^i)$. Let

$$v_{-\infty}^i = v_{-\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}, \phi'_{-\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}}}^i \quad \text{and} \quad v_{\infty}^i = v_{\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}, \phi'_{\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}}}^i.$$

These two nodes can be thought of as the leftmost and rightmost nodes of path \mathcal{P}^i . We connect the paths together by adding edges between the leftmost (rightmost, respectively) nodes in the paths, i.e., for any i and i' , we add edges $(v_{-\infty}^i, v_{-\infty}^{i'})$ ($(v_{\infty}^i, v_{\infty}^{i'})$, respectively).

We connect the highways and paths by adding an edge from each node $h_j^{\lfloor \kappa \rfloor}$ to $v_{j,1}^i$. We also create nodes s and t and connect s (t , respectively) to all nodes $v_{-\infty}^i$ (v_{∞}^i , respectively). Finally, we make infinite copies of every edge except highway edges, i.e., those in $\cup_{i=1}^{\lfloor \kappa \rfloor} E(\mathcal{H}^i)$. Fig. 1 shows an example of $G(\Gamma, \kappa, \Lambda)$.

2.2 Terminologies

For any numbers i, j, i' , and j' , we say that $(i', j') \geq (i, j)$ if $i' > i$ or $(i' = i \text{ and } j' \geq j)$. For any $-\lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor} \leq i \leq \lceil \kappa \rceil \Lambda^{\lfloor \kappa \rfloor}$ and $1 \leq j \leq \phi'_i$, define the (i, j) -set as

$$S_{i,j} = \begin{cases} \{h_{i'}^x \mid 1 \leq x \leq \kappa, i' \leq i\} \cup \{v_{i',j'}^x \mid 1 \leq x \leq \Gamma, (i, j) \geq (i', j')\} \cup \{s\} & \text{if } i \geq 0 \\ \{h_{i'}^x \mid 1 \leq x \leq \kappa, i' \geq i\} \cup \{v_{i',j'}^x \mid 1 \leq x \leq \Gamma, (-i, j) \geq (-i', j')\} \cup \{r\} & \text{if } i < 0. \end{cases}$$

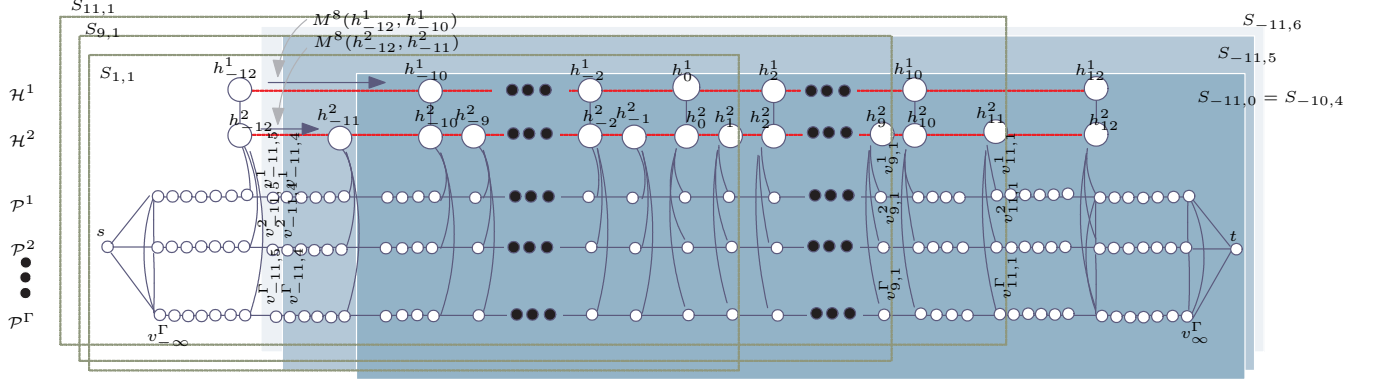


Figure 2: An example of round 11 in the proof of Theorem 2.2 (see detail in Example 2.5).

See Fig. 2 for an example. For convenience, for any $i > 0$, let $S_{i,0} = S_{i-1,\phi'_{i-1}}$ and $S_{-i,0} = S_{-(i-1),\phi'_{-(i-1)}}$, and, for any j , let $S_{\lceil\kappa\rceil\Lambda^{\lceil\kappa\rceil}+1,j} = S_{\lceil\kappa\rceil\Lambda^{\lceil\kappa\rceil},\phi'_{\lceil\kappa\rceil\Lambda^{\lceil\kappa\rceil}}}$ and $S_{-\lceil\kappa\rceil\Lambda^{\lceil\kappa\rceil}-1,j} = S_{-\lceil\kappa\rceil\Lambda^{\lceil\kappa\rceil},\phi'_{-\lceil\kappa\rceil\Lambda^{\lceil\kappa\rceil}}}$.

Let \mathcal{A} be any *deterministic* distributed algorithm run on $G(\Gamma, \kappa, \Lambda)$ for computing a function f . Fix any input strings \bar{x} and \bar{y} given to s and t 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 node v at the end of time τ during the execution $\varphi_{\mathcal{A}}(\bar{x}, \bar{y})$ by $\sigma_{\mathcal{A}}(v, \tau, \bar{x}, \bar{y})$. Let $\sigma_{\mathcal{A}}(v, 0, \bar{x}, \bar{y})$ be the state of the node v before the execution $\varphi_{\mathcal{A}}(\bar{x}, \bar{y})$ begins. Note that $\sigma_{\mathcal{A}}(v, 0, \bar{x}, \bar{y})$ is independent of the input if $v \notin \{s, t\}$, depends only on \bar{x} if $v = s$ and depends only on \bar{y} if $v = t$. Moreover, in two different executions $\varphi_{\mathcal{A}}(\bar{x}, \bar{y})$ and $\varphi_{\mathcal{A}}(\bar{x}', \bar{y}')$, a node reaches the same state at time τ (i.e., $\sigma_{\mathcal{A}}(v, \tau, \bar{x}, \bar{y}) = \sigma_{\mathcal{A}}(v, \tau, \bar{x}', \bar{y}')$) if and only if it receives the same sequence of messages on each of its incoming links.

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

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

is a vector of the states of the nodes of U at the end of time τ of the execution $\varphi_{\mathcal{A}}(\bar{x}, \bar{y})$. From now on, to simplify notations, when \mathcal{A} , \bar{x} and \bar{y} are clear from the context, we use $C_{i,j}^\tau$ to denote $C_{\mathcal{A}}(S_{i,j}, \tau, \bar{x}, \bar{y})$.

2.3 Proof of Theorem 2.2

Let $G = G(\Gamma, \kappa, \Lambda)$. Let f be the function in the theorem statement. Let \mathcal{A}_ϵ be any ϵ -error distributed algorithm for computing f on G . Fix a random string \bar{r} used by \mathcal{A}_ϵ (shared by all nodes in G) and consider the *deterministic* algorithm \mathcal{A} run on the input of \mathcal{A}_ϵ 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 $T_{\mathcal{A}} \leq \kappa\Lambda^\kappa$, 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 $(2\kappa \log n)T_{\mathcal{A}}$ communication bits in $8T_{\mathcal{A}}/(\kappa\Lambda)$ rounds, as follows. (We provide an example in the end of this section.)

Rounds, Phases, and Iterations. For convenience, we will name the rounds backward, i.e., Alice and Bob start at round $\lceil\kappa\rceil\Lambda^{\lceil\kappa\rceil}$ and proceed to round $\lceil\kappa\rceil\Lambda^{\lceil\kappa\rceil} - 1$, $\lceil\kappa\rceil\Lambda^{\lceil\kappa\rceil} - 2$, and so on. Each round is divided into two *phases*, i.e., when Alice sends messages and Bob sends messages (recall that Alice sends messages first in each iteration). Each phase of round r is divided into ϕ'_r *iterations*. Each iteration simulates one round of algorithm \mathcal{A} . We call the i^{th} iteration of round r when Alice (Bob, respectively) sends messages the *iteration* $I_{r,A,i}$ ($I_{r,B,i}$, respectively). Therefore,

in each round r we have the following order of iterations: $I_{r,A,1}, I_{r,A,2}, \dots, I_{r,A,\phi'_r}, I_{r,B,1}, \dots, I_{r,B,\phi'_r}$. For convenience, we refer to the time before communication begins as round $\lceil \kappa \rceil \Lambda^{\lceil \kappa \rceil} + 1$ and let $I_{r,A,0} = I_{r+1,A,\phi'_{r+1}}$ and $I_{r,B,0} = I_{r+1,B,\phi'_{r+1}}$.

Our goal is to simulate one round of algorithm \mathcal{A} per iteration. That is, after iteration $I_{r,B,i}$ finishes, we will finish the $(\sum_{r'=r+1}^{\lceil \kappa \rceil \Lambda^{\lceil \kappa \rceil}} \phi'_{r'} + i)^{th}$ round of algorithm \mathcal{A} . Specifically, we let

$$t_r = \sum_{r'=r+1}^{\lceil \kappa \rceil \Lambda^{\lceil \kappa \rceil}} \phi'_{r'}$$

and our goal is to construct a protocol with properties as in the following lemma.

Lemma 2.3. *There exists a protocol such that there are at most $\kappa \log n$ bits sent in each iteration and satisfies the following properties. For any $r \geq 0$ and $0 \leq i \leq \phi'_r$,*

1. *after $I_{r,A,i}$ finishes, Alice and Bob know $C_{r-i\Lambda^{\lceil \kappa \rceil}-1,1}^{t_r+i}$ and $C_{-r,\phi'_{r-i}}^{t_r+i}$, respectively, and*
2. *after $I_{r,B,i}$ finishes, Alice and Bob know $C_{r,\phi'_{r-i}}^{t_r+i}$ and $C_{-r+i\Lambda^{\lceil \kappa \rceil}-1,1}^{t_r+i}$, respectively.*

Proof. We first argue that the properties hold for iteration $I_{\lceil \kappa \rceil \Lambda^{\lceil \kappa \rceil} + 1, A, 0}$, i.e., before Alice and Bob starts communicating. After round $r = \lceil \kappa \rceil \Lambda^{\lceil \kappa \rceil}$ starts, Alice can compute $C_{r+1,0}^0 = C_{r+1,1}^0 = C_{r,\phi'_r}^0$ which contains the states of all nodes in $G(\Gamma, \kappa, \Lambda)$ except t . She can do this because every node except s and t has the same state regardless of the input and the state of s depends only on her input string \bar{x} . Similarly, Bob can compute $C_{-(r+1),0}^0 = C_{-(r+1),1}^0 = C_{r,\phi'_r}^0$ which depends only on his input \bar{y} .

Now we show that, if the lemma holds for any iteration $I_{r,A,i-1}$ then it also holds for iteration $I_{r,A,i}$ as well. Specifically, we show that if Alice and Bob know $C_{r-(i-1)\Lambda^{\lceil \kappa \rceil}-1,1}^{t_r+i-1}$ and $C_{-r,\phi'_{r-(i-1)}}^{t_r+i-1}$, respectively, then they will know $C_{r-i\Lambda^{\lceil \kappa \rceil}-1,1}^{t_r+i}$ and $C_{-r,\phi'_{r-i}}^{t_r+i}$, respectively, after Alice sends at most $\kappa \log n$ messages.

First we show that Alice can compute $C_{r-i\Lambda^{\lceil \kappa \rceil}-1,1}^{t_r+i}$ without receiving any message from Bob. Recall that Alice can compute $C_{r-i\Lambda^{\lceil \kappa \rceil}-1,1}^{t_r+i}$ if she knows

- $C_{r-i\Lambda^{\lceil \kappa \rceil}-1,1}^{t_r+i-1}$, and
- all messages sent to all nodes in $S_{r-i\Lambda^{\lceil \kappa \rceil}-1,1}$ at time $t_r + i$ of algorithm \mathcal{A} .

By assumption, Alice knows $C_{r-(i-1)\Lambda^{\lceil \kappa \rceil}-1,1}^{t_r+i-1}$ which implies that she knows $C_{r-i\Lambda^{\lceil \kappa \rceil}-1,1}^{t_r+i-1}$ since $S_{r-i\Lambda^{\lceil \kappa \rceil}-1,1} \subseteq S_{r-(i-1)\Lambda^{\lceil \kappa \rceil}-1,1}$. Moreover, observe that all neighbors of all nodes in $S_{r-i\Lambda^{\lceil \kappa \rceil}-1,1}$ are in $S_{r-(i-1)\Lambda^{\lceil \kappa \rceil}-1,1}$. Thus, Alice can compute all messages sent to all nodes in $S_{r-i\Lambda^{\lceil \kappa \rceil}-1,1}$ at time $t_r + i$ of algorithm \mathcal{A} . Therefore, Alice can compute $C_{r-i\Lambda^{\lceil \kappa \rceil}-1,1}^{t_r+i}$ without receiving any message from Bob.

Now we show that Bob can compute $C_{-r,\phi'_{r-i}}^{t_r+i}$ by receiving at most $\kappa \log n$ bits from Alice and use the knowledge of $C_{-r,\phi'_{r-i+1}}^{t_r+i-1}$. Note that Bob can compute $C_{-r,\phi'_{r-i}}^{t_r+i}$ if he knows

- $C_{-r,\phi'_{r-i}}^{t_r+i-1}$, and
- all messages sent to all nodes in $S_{-r,\phi'_{r-i}}$ at time $t_r + i$ of algorithm \mathcal{A} .

By assumption, Bob knows $C_{-r,\phi'_{r-i+1}}^{t_r+i-1}$ which implies that he knows $C_{-r,\phi'_{r-i}}^{t_r+i-1}$ since $S_{-r,\phi'_{r-i}} \subseteq S_{-r,\phi'_{r-i+1}}$. Moreover, observe that all neighbors of all nodes in $S_{-r,\phi'_{r-i}}$ are in $S_{-r,\phi'_{r-i+1}}$, except

$$h_{-(r+1)}^{[\kappa]}, h_{-([r/\Lambda]+1)}^{[\kappa]-1}, \dots, h_{-([r/\Lambda^i]+1)}^{[\kappa]-i}, \dots, h_{-([r/\Lambda^{[\kappa]}-1]+1)}^1.$$

In other words, Bob can compute all messages sent to all nodes in $S_{-r, \phi'_{-r-i}}$ at time $t_r + i$ except

$$M^{t_r+i}(h_{-(r+1)}^{[\kappa]}, h_{-r}^{[\kappa]}), \dots, M^{t_r+i}(h_{-([r/\Lambda^i]+1)}^{[\kappa]-i}, h_{-[r/\Lambda^i]}^{[\kappa]-i}), \dots, M^{t_r+i}(h_{-([r/\Lambda^{[\kappa]}-1]+1)}^1, h_{-([r/\Lambda^{[\kappa]}-1]}^1)$$

where $M^{t_r+i}(u, v)$ is the message sent from u to v at time $t_r + i$ of algorithm \mathcal{A} . Observe further that Alice can compute these messages because she knows $C_{r-(i-1)\Lambda^{[\kappa]}-1, 1}^{t_r+i-1}$ which contains the states of $h_{-(r+1)}^{[\kappa]}, \dots, h_{-([r/\Lambda^i]+1)}^{[\kappa]-i}, \dots, h_{-([r/\Lambda^{[\kappa]}-1]+1)}^1$ at time $t_r + i - 1$. (In particular, $C_{r-(i-1)\Lambda^{[\kappa]}-1, 1}^{t_r+i-1}$ is a superset of $C_{0,1}^{t_r+i-1}$ which contains the states of $h_{-(r+1)}^{[\kappa]}, \dots, h_{-([r/\Lambda^{[\kappa]}-1]+1)}^1$.) So, Alice can send these messages to Bob and Bob can compute $C_{-r, \phi'_{-r-i}}^{t_r+i}$ at the end of the iteration. Each of these messages contains at most $\log n$ bits since each of them corresponds to a message sent on one edge. Therefore, Alice sends at most $\kappa \log n$ bits to Bob in total. This shows the first property.

After Alice finishes sending messages, the two parties will switch their roles and a similar protocol can be used to show that the second property, i.e., if the lemma holds for any iteration $I_{r,B,i-1}$ then it also holds for iteration $I_{r,B,i}$ as well. That is, if Alice and Bob know $C_{r, \phi'_{r-(i-1)}}^{t_r+i-1}$ and $C_{-r+(i-1)\Lambda^{[\kappa]}-1, 1}^{t_r+i-1}$, respectively, then Bob can send $\kappa \log n$ bits to Alice so that they can compute $C_{r, \phi'_{r-i}}^{t_r+i}$ and $C_{-r+i\Lambda^{[\kappa]}-1, 1}^{t_r+i}$, respectively. \square

Let P be the protocol as in Lemma 2.3. Alice and Bob will run protocol P until round r' , where r' is the largest number such that $t_{r'} + \phi'_{r'} \geq T_{\mathcal{A}}$. Lemma 2.3 implies that after iteration $I_{r', B, T_{\mathcal{A}}-t_{r'}}$, Bob knows $C_{-r', \phi'_{-r'}-T_{\mathcal{A}}+t_{r'}}^{t_{r'}+T_{\mathcal{A}}-t_{r'}} = C_{-r', \phi'_{-r'}-T_{\mathcal{A}}+t_{r'}}^{T_{\mathcal{A}}}$ (note that $\phi'_{-r'} - T_{\mathcal{A}} + t_{r'} \geq 0$). In particular, Bob knows the state of node t at time $T_{\mathcal{A}}$, i.e., he knows $\sigma_{\mathcal{A}}(t, T_{\mathcal{A}}, \bar{x}, \bar{y})$. Thus, Bob can output the output of \mathcal{A} which is output from t .

Since \mathcal{A}_{ϵ} 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. The number of rounds is bounded as in the following claim.

Claim 2.4. *If algorithm \mathcal{A} finishes in time $T_{\mathcal{A}} \leq \lceil \kappa \rceil \Lambda^{\kappa}$ then $r' > \lceil \kappa \rceil \Lambda^{\kappa} - 8T_{\mathcal{A}} / (\lceil \kappa \rceil \Lambda)$. In other words, the number of rounds Alice and Bob need to simulate \mathcal{A} is $8T_{\mathcal{A}} / (\lceil \kappa \rceil \Lambda)$*

Proof. Let $R^* = 8T_{\mathcal{A}} / (\lceil \kappa \rceil \Lambda)$ and let $r^* = \Lambda^{[\kappa]} - R^* + 1$. Assume for the sake of contradiction that Alice and Bob need more than R^* rounds. This means that $r' < r^*$. Alice and Bob requiring more than R^* rounds implies that

$$\sum_{r=r^*}^{\lceil \kappa \rceil \Lambda^{[\kappa]}} \phi'_r = t_{r^*} + \phi'_{r^*} < T_{\mathcal{A}} \leq \lceil \kappa \rceil \Lambda^{\kappa}. \quad (2)$$

It follows that for any $r \geq r^*$,

$$\begin{aligned} \phi'_r &= \min \left(\phi(h_r^{k'}), \max(1, \lceil \kappa \rceil \Lambda^{\kappa} - \sum_{r' > r} \phi_{r'}) \right) && \text{(by definition of } \phi'_r) \\ &= \phi_r && \text{(because } \sum_{r' \geq r^*} \phi'_{r'} < \lceil \kappa \rceil \Lambda^{\kappa}) \\ &= \left\lfloor \frac{r}{\Lambda^{[\kappa]}-1} \right\rfloor + 1 && \text{(by definition of } \phi_r). \end{aligned}$$

Therefore, the total number of steps that can be simulated by Alice and Bob up to round r^* is

$$\begin{aligned}
\sum_{r=r^*}^{\lceil \kappa \rceil \Lambda^{\lceil \kappa \rceil}} \phi'_r &= \sum_{r=r^*}^{\lceil \kappa \rceil \Lambda^{\lceil \kappa \rceil}} \left(\left\lfloor \frac{r}{\Lambda^{\lceil \kappa \rceil - 1}} \right\rfloor + 1 \right) \\
&\geq \Lambda^{\lceil \kappa \rceil - 1} \sum_{i=1}^{\lfloor R^* / \Lambda^{\lceil \kappa \rceil - 1} \rfloor} (\lceil \kappa \rceil \Lambda - i) \\
&\geq \Lambda^{\lceil \kappa \rceil - 1} \cdot \frac{\lfloor R^* / \Lambda^{\lceil \kappa \rceil - 1} \rfloor (\lceil \kappa \rceil \Lambda - 1)}{2} \\
&\geq \frac{R^* \lceil \kappa \rceil \Lambda}{8} \\
&\geq T_{\mathcal{A}}
\end{aligned}$$

contradicting Eq. (2). □

Since there are at most $\kappa \log n$ bits sent in each iteration and Alice and Bob runs P for $T_{\mathcal{A}}$ iterations, the total number of bits exchanged is at most $(2\kappa \log n)T_{\mathcal{A}}$. This completes the proof of Theorem 2.2.

Example 2.5. Fig. 2 shows an example of the protocol we use above. Before iteration $I_{11,A,1}$ begins, Alice and Bob know $C_{11,1}^7$ and $C_{-11,5}^7$, respectively (since Alice and Bob already simulated \mathcal{A} for $\phi'_{12} = 7$ steps in round 12). Then, Alice computes and sends $M^8(h_{-12}^2, h_{-11}^2)$ and $M^8(h_{-12}^1, h_{-10}^1)$ to Bob. Alice and Bob then compute $C_{11,1}^8$ and $C_{-11,6}^8$, respectively, at the end of iteration $I_{11,A,1}$. After they repeat this process for five more times, i.e. Alice sends

$$\begin{aligned}
M^9(h_{-12}^2, h_{-11}^2), M^{10}(h_{-12}^2, h_{-11}^2), \dots, M^{13}(h_{-12}^2, h_{-11}^2), \quad \text{and} \\
M^9(h_{-12}^1, h_{-10}^1), M^{10}(h_{-12}^1, h_{-10}^1), \dots, M^{13}(h_{-12}^1, h_{-10}^1),
\end{aligned}$$

Bob will be able to compute $C_{-11,0}^{13} = C_{-10,4}^{13}$. Note that Alice is able to compute $C_{9,1}^8, C_{7,1}^9, \dots, C_{1,1}^{12}$ without receiving any messages from Bob so she can compute and send the previously mentioned messages to Bob.

3 The pointer chasing problem

In this section, we define the pointer chasing problem and prove its lower bound (Lemma 3.2) which will be used to prove Theorem 1.1 in the next section.

Informally, the r -round pointer chasing problem has parameters r and m and there are two players, which could be Alice and Bob or nodes s and t , who receive functions $f_A : [m] \rightarrow [m]$ and $f_B : [m] \rightarrow [m]$, respectively. The goal is to compute a function starting from 1 and alternatively applying f_A and f_B for r times each, i.e., compute $f_B(\dots f_A(f_B(f_A)))$ where f_A and f_B appear r times each. To be precise, let \mathcal{F}_m be the set of functions $f : [m] \rightarrow [m]$. For any $i \geq 0$ define $g^i : \mathcal{F}_m \times \mathcal{F}_m \rightarrow [m]$ inductively as

$$g^0(f_A, f_B) = 1 \quad \text{and} \quad g^i(f_A, f_B) \begin{cases} f_A(g^{i-1}(f_A, f_B)) & \text{if } i > 0 \text{ and } i \text{ is odd,} \\ f_B(g^{i-1}(f_A, f_B)) & \text{if } i > 0 \text{ and } i \text{ is even.} \end{cases}$$

Also define function $\text{PC}^{i,m}(f_A, f_B) = g^{2i}(f_A, f_B)$. The goal of the r -round pointer chasing problem is to compute $\text{PC}^{r,m}(f_A, f_B)$.

Observe that if Alice and Bob can communicate for r rounds then they can compute $\text{PC}^{r,m}$ naively by exchanging $O(r \log m)$ bits. Interestingly, Nisan and Wigderson [26] show that if Alice and Bob are allowed only $r - 1$ rounds then they essentially cannot do anything better than having Alice sent everything she knows to Bob.¹

Theorem 3.1. [26] $R_{1/3}^{(r-1)-cc-pub}(\text{PC}^{r,m}) = \Omega(m/r^2 - r \log m)$.

The pointer chasing problem on $G(\Gamma, \kappa, \Lambda)$. We now consider the pointer chasing problem on network $G(\Gamma, \kappa, \Lambda)$ where s and t receive f_A and f_B respectively. The following lemma follows from Theorem 2.2 and 3.1.

Lemma 3.2. For any $\kappa, \Gamma, \Lambda \geq 2, m \geq \kappa^2 \Lambda^{4\kappa} \log n, 16\Lambda^{\kappa-1} \geq r > 8\Lambda^{\kappa-1}, R_{1/3}^{G(\Gamma, \kappa, \Lambda), s, t}(\text{PC}^{r,m}) = \Omega(\kappa \Lambda^\kappa)$.

Proof. If $R_{1/3}^{G(\Gamma, \kappa, \Lambda), s, t}(\text{PC}^{r,m}) > \kappa \Lambda^\kappa$ then we are done so we assume that $R_{1/3}^{G(\Gamma, \kappa, \Lambda), s, t}(\text{PC}^{r,m}) \leq \kappa \Lambda^\kappa$. Thus,

$$R_{1/3}^{G(\Gamma, \kappa, \Lambda), s, t}(\text{PC}^{r,m}) \geq R_{1/3}^{\frac{8R_{1/3}^{G(\Gamma, \kappa, \Lambda), s, t}(\text{PC}^{r,m})}{\kappa \Lambda} - cc-pub}(\text{PC}^{r,m}) / (2\kappa \log n) \quad (3)$$

$$\geq R_{1/3}^{\frac{8\kappa \Lambda^\kappa}{\kappa \Lambda} - cc-pub}(\text{PC}^{r,m}) / (2\kappa \log n) \quad (4)$$

$$= \Omega((m(8\Lambda^{\kappa-1})^{-2} - 8\Lambda^{\kappa-1} \log m) / (\kappa \log n)) \quad (5)$$

$$= \Omega(\kappa \Lambda^\kappa) \quad (6)$$

where Eq. (3) is by Theorem 2.2 and the fact that $R_{1/3}^{G(\Gamma, \kappa, \Lambda), s, t}(\text{PC}^{r,m}) \leq \kappa \Lambda^\kappa$, Eq. (4) uses the fact that the communication does not increase when we allow more rounds and $R_{1/3}^{G(\Gamma, \kappa, \Lambda), s, t}(\text{PC}^{r,m}) \leq \kappa \Lambda^\kappa$, Eq. (5) follows from Theorem 3.1 with the fact that $16\Lambda^{\kappa-1} \geq r > 8\Lambda^{\kappa-1}$ and Eq. (6) is because $m \geq \kappa^2 \Lambda^{4\kappa} \log n$. \square

4 Proof of Theorem 1.1

An $\Omega(D)$ lower bound has already been shown (and is fairly straightforward) in [10]; so we focus on showing the $\Omega(\sqrt{\ell D})$ lower bound. Moreover, we will prove the theorem only for the version where destination outputs source. This is because we can convert algorithms for the other two version to solve this version by adding $O(D)$ rounds. To see this, observe that once the source outputs the ID of the destination, we can take additional $O(D)$ rounds to send the ID of the source to the destination. Similarly, if nodes know their positions, the node with position ℓ can output the source's ID by taking additional $O(D)$ rounds to request for the source's ID. Theorem 1.1, for the case where destination outputs source, follows from the following lemma.

Lemma 4.1. For any real $\kappa \geq 1$ and integers $\Lambda \geq 2$, and $\Gamma \geq 32\kappa^2 \Lambda^{6\kappa-1} \log n$, there exists a family of networks \mathcal{H} such that any network $H \in \mathcal{H}$ has $\Theta(\kappa \Gamma \Lambda^\kappa)$ nodes and diameter $D = \Theta(\kappa \Lambda)$, and any algorithm for computing the destination of a random walk of length $\ell = \Theta(\Lambda^{2\kappa-1})$ requires $\Omega(\sqrt{\ell D})$ time on some network $H \in \mathcal{H}$.

¹In fact this holds even when Alice and Bob are allowed r rounds but Alice cannot send a message in the first round.

Proof. We show how to compute $\text{PC}^{r,m}$ on $G = G(\Gamma, \kappa, \Lambda)$ by reducing the problem to the problem of sampling a random walk destination in some network H_{f_A, f_B} , obtained by restrict the number of copies of some edges in G , depending on input functions f_A and f_B . We let \mathcal{H} be the family of network H_{f_A, f_B} over all input functions. Note that for any input functions, an algorithm on H_{f_A, f_B} can be run on G with the same running time since every edge in G has more capacity than its counterpart in H_{f_A, f_B} .

Let $r = 16\Lambda^{\kappa-1}$ and $m = \kappa^2\Lambda^{5\kappa} \log n$. Note that $2rm \leq \Gamma$. For any $i \leq r$ and $j \leq m$, let

$$S^{i,j} = \mathcal{P}^{2(i-1)m+j} \quad \text{and} \quad T^{i,j} = \mathcal{P}^{2(i-1)m+m+j}.$$

That is, $S^{1,1} = \mathcal{P}^1, \dots, S^{1,m} = \mathcal{P}^m, T^{1,1} = \mathcal{P}^{m+1}, \dots, T^{1,m} = \mathcal{P}^{2m}, S^{2,1} = \mathcal{P}^{2m+1}, \dots, T^{r,m} = \mathcal{P}^{2rm}$. Let L be the number of nodes in each path. Note that $L = \Theta(\kappa\Lambda^\kappa)$ by Lemma 2.1. Denote the nodes in $S^{i,j}$ from *left to right* by $s_1^{i,j}, \dots, s_L^{i,j}$. (Thus, $s_1^{i,j} = v_{-\infty}^{2(i-1)m+j}$ and $s_L^{i,j} = v_\infty^{2(i-1)m+j}$.) Also denote the nodes in $T^{i,j}$ from *right to left* by $t_1^{i,j}, \dots, t_L^{i,j}$. (Thus, $t_1^{i,j} = v_\infty^{2(i-1)m+m+j}$ and $t_L^{i,j} = v_{-\infty}^{2(i-1)m+m+j}$.) Note that for any i and j , $s_1^{i,j}$ and $t_L^{i,j}$ are adjacent to s while $s_L^{i,j}$ and $t_1^{i,j}$ are adjacent to t .

Now we construct H_{f_A, f_B} . For simplicity, we fix input functions f_A and f_B and denote H_{f_A, f_B} simply by H . To get H we let every edge in G have one copy (thus with capacity $O(\log n)$), except the following edges. For any $i \leq r$, $j \leq m$, and $x < L$, we have $(6\Gamma\ell)^{2(i-1)L+x}$ copies of edges between nodes $s_x^{i,j}$ and $s_{x+1}^{i,j}$ and $(6\Gamma\ell)^{2(i-1)L+L+x}$ copies of edges between nodes $t_x^{i,j}$ and $t_{x+1}^{i,j}$. Note that these numbers of copies of edges are always the same, regardless of the input f_A and f_B .

Additionally, we have the following numbers of edges which depend on the input functions. First, s specifies the following number of edges between its neighbors. For any $i \leq r$, $j \leq m$, we have $(6\Gamma\ell)^{2(i-1)L+L}$ copies of edges between nodes $t_L^{i,j}$ and $s_1^{i, f_A(j)}$. These numbers of edges can be specified in one round since both $s_1^{i,j}$ and $t_L^{i, f_A(j)}$ are adjacent to s . Similarly, we have $(6\Gamma\ell)^{2(i-1)L+2L}$ copies of edges between nodes $t_1^{i,j}$ and $s_L^{i+1, f_B(j)}$ which can be done in one round since both nodes are adjacent to t . This completes the description of H .

Now we use any random walk algorithm to compute the destination of a walk of length $\ell = 2rL - 1 = \Theta(\Lambda^{2\kappa-1})$ on H by starting a random walk at $s_1^{1, f_A(1)}$. If the random walk destination is $t_L^{r,j}$ for some j , then node t outputs the number j ; otherwise, node t outputs an arbitrary number.

Now observe the following claim.

Claim 4.2. *Node t outputs $\text{PC}^{r,m}(f_A, f_B)$ with probability at least $2/3$.*

Proof. Let P^* be the path consisting of nodes $s_1^{1, f_A(1)}, \dots, s_L^{1, f_A(1)}, t_1^{1, f_B(f_A(1))}, \dots, t_L^{1, f_B(f_A(1))}, s_1^{1, f_A(f_B(f_A(1)))}, \dots, s_L^{i, g^{2^{i-1}}(f_A, f_B)}, t_1^{i, g^{2^i}(f_A, f_B)}, \dots, t_L^{r, g^{2^r}(f_A, f_B)}$. We claim that the random walk will follow path P^* with probability at least $2/3$. The node of distance $(2rL - 1)$ from $s_1^{1, f_A(1)}$ in this path is $t_L^{r, g^{2^r}(f_A, f_B)} = t_L^{r, \text{PC}^{r,m}(1)}$ and thus the algorithm described above will output $\text{PC}^{r,m}(1)$ with probability at least $2/3$.

To prove the above claim, consider any node u in path P^* . Let u' and u'' be the node before and after u in P^* , respectively. Let m' and m'' be the number of multiedges (u, u') and (u, u'') , respectively. Observe that $m'' \geq 6\Gamma\ell m'$. Moreover, observe that there are at most Γ edges between u and other nodes. Thus, if a random walk is at u , it will continue to u'' with probability at least $1 - \frac{1}{3\ell}$. By union bound, the probability that a random walk will follow P^* is at least $1 - \frac{1}{3}$, as claimed. \square

Thus, if there is any random walk algorithm with running time $O(T)$ on all networks in \mathcal{H} then we can use such algorithm to solve $\text{PC}^{r,m}$ (with error probability $1/3$) in time $O(T)$. Using the lower

bound of computing solving $PC^{r,m}$ in Lemma 3.2, the random walk computation also has a lower bound of $\Omega(\kappa\Lambda^\kappa) = \Omega(\sqrt{\ell D})$ as claimed. \square

To prove Theorem 1.1 with the given parameters n , D and ℓ , we simply set Λ and κ so that $\kappa\Lambda = D$ and $\Lambda^{2\kappa-1} = \Theta(\ell)$. This choice of Λ and κ exists since $\ell \geq D$. Setting Γ large enough so that $\Gamma \geq 32\kappa^2\Lambda^{6\kappa-1} \log n$ while $\Gamma = \Theta(n)$. (This choice of Γ exists since $\ell \leq (n/(D^3 \log n))^{1/4}$.) By applying the above lemma, Theorem 1.1 follows.

5 Conclusion

In this paper we prove a tight unconditional lower bound on the time complexity of distributed random walk computation, implying that the algorithm in [11] is time optimal. To the best of our knowledge, this is the first lower bound that the diameter plays a role of multiplicative factor. Our proof technique comes from strengthening the connection between communication complexity and distributed algorithm lower bounds initially studied in [9] by associating *rounds* in communication complexity to the distributed algorithm running time, with network diameter as a trade-off factor.

In light of the success in proving distributed algorithm lower bounds from communication complexity in this and the previous work [9], it is interesting to explore further applications of this technique. One interesting approach is to show a connection between distributed algorithm lower bounds and other models of communication complexity, such as multiparty and asymmetric communication complexity (see, e.g., [21]). Potential applications of this technique are distance-related problems such as shortest s - t path, single-source distance computation, and spanner construction. The lower bound of $\Omega(\sqrt{n})$ are shown in [9] for these types of problems and stronger lower bounds such as $\Omega(\sqrt{nD})$ or $\Omega(n)$ are still possible. Another problem left open is showing a lower bound of performing a long walk. For example, one can generate a *random spanning tree* by computing a walk of length equals the cover time (using the version where every node knows their positions). It is interesting to see if performing such a walk can be done faster. Additionally, the upper and lower bounds of the problem of generating a random spanning tree itself is very interesting since its current upper bound of $\tilde{O}(\sqrt{mD})$ [11] simply follows as an application of random walk computation [11] while no lower bound is known.

References

- [1] Lada A. Adamic, Rajan M. Lukose, Amit R. Puniyani, and Bernardo A. Huberman. Search in power-law networks. *Physical Review*, 64, 2001.
- [2] H. Baala, O. Flauzac, J. Gaber, M. Bui, and T. El-Ghazawi. A self-stabilizing distributed algorithm for spanning tree construction in wireless ad hoc networks. *J. Parallel Distrib. Comput.*, 63(1):97–104, 2003.
- [3] Judit Bar-Ilan and Dror Zernik. Random leaders and random spanning trees. In *Proceedings of the 3rd International Workshop on Distributed Algorithms*, pages 1–12, London, UK, 1989. Springer-Verlag.
- [4] Thibault Bernard, Alain Bui, and Olivier Flauzac. Random distributed self-stabilizing structures maintenance. In *ISSADS*, pages 231–240, 2004.
- [5] Ashwin R. Bharambe, Mukesh Agrawal, and Srinivasan Seshan. Mercury: supporting scalable multi-attribute range queries. In *SIGCOMM*, pages 353–366, 2004.
- [6] Andrei Z. Broder. Generating random spanning trees. In *FOCS*, pages 442–447, 1989.
- [7] Brian F. Cooper. Quickly routing searches without having to move content. In *IPTPS*, pages 163–172, 2005.
- [8] Don Coppersmith, Prasad Tetali, and Peter Winkler. Collisions among random walks on a graph. *SIAM J. Discret. Math.*, 6(3):363–374, 1993.
- [9] 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 *STOC*, 2011.
- [10] Atish Das Sarma, Danupon Nanongkai, and Gopal Pandurangan. Fast distributed random walks. In *PODC*, pages 161–170, 2009.
- [11] Atish Das Sarma, Danupon Nanongkai, Gopal Pandurangan, and Prasad Tetali. Efficient distributed random walks with applications. In *PODC*, pages 201–210, 2010.
- [12] Shlomi Dolev and Nir Tzachar. Spanders: distributed spanning expanders. In *SAC*, pages 1309–1314, 2010.
- [13] Michael Elkin. An Unconditional Lower Bound on the Time-Approximation Trade-off for the Distributed Minimum Spanning Tree Problem. *SIAM J. Comput.*, 36(2):433–456, 2006. Also in STOC’04.
- [14] Joan Feigenbaum, Sampath Kannan, Andrew McGregor, Siddharth Suri, and Jian Zhang 0004. Graph Distances in the Data-Stream Model. *SIAM J. Comput.*, 38(5):1709–1727, 2008. Also in SODA’05.
- [15] Ayalvadi J. Ganesh, Anne-Marie Kermarrec, and Laurent Massoulié. Peer-to-peer membership management for gossip-based protocols. *IEEE Trans. Comput.*, 52(2):139–149, 2003.
- [16] Juan A. Garay, Shay Kutten, and David Peleg. A Sublinear Time Distributed Algorithm for Minimum-Weight Spanning Trees. *SIAM J. Comput.*, 27(1):302–316, 1998. Also in FOCS ’93.

- [17] Christos Gkantsidis, Milena Mihail, and Amin Saberi. Hybrid search schemes for unstructured peer-to-peer networks. In *INFOCOM*, pages 1526–1537, 2005.
- [18] Amos Israeli and Marc Jalfon. Token management schemes and random walks yield self-stabilizing mutual exclusion. In *PODC*, pages 119–131, 1990.
- [19] David Kempe, Jon M. Kleinberg, and Alan J. Demers. Spatial gossip and resource location protocols. In *STOC*, pages 163–172, 2001.
- [20] Jon M. Kleinberg. The small-world phenomenon: an algorithmic perspective. In *STOC*, pages 163–170, 2000.
- [21] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, New York, NY, USA, 1997.
- [22] Ching Law and Kai-Yeung Siu. Distributed construction of random expander networks. In *INFOCOM*, 2003.
- [23] Dmitri Loguinov, Anuj Kumar, Vivek Rai, and Sai Ganesh. Graph-theoretic analysis of structured peer-to-peer systems: routing distances and fault resilience. In *SIGCOMM*, pages 395–406, 2003.
- [24] Zvi Lotker, Boaz Patt-Shamir, and David Peleg. Distributed MST for constant diameter graphs. *Distributed Computing*, 18(6):453–460, 2006. Also in PODC’01.
- [25] Qin Lv, Pei Cao, Edith Cohen, Kai Li, and Scott Shenker. Search and replication in unstructured peer-to-peer networks. In *ICS*, pages 84–95, 2002.
- [26] Noam Nisan and Avi Wigderson. Rounds in communication complexity revisited. *SIAM J. Comput.*, 22(1):211–219, 1993. Also in STOC’91.
- [27] David Peleg. *Distributed computing: a locality-sensitive approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [28] David Peleg and Vitaly Rubinfeld. 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.
- [29] Ming Zhong and Kai Shen. Random walk based node sampling in self-organizing networks. *Operating Systems Review*, 40(3):49–55, 2006.
- [30] Ming Zhong, Kai Shen, and Joel I. Seiferas. Non-uniform random membership management in peer-to-peer networks. In *INFOCOM*, pages 1151–1161, 2005.