Filling Logarithmic Gaps in Distributed Complexity for Global Problems*

Hiroaki Ookawa and Taisuke Izumi

Nagoya Institute of Technology, Gokiso-cho, Showa-ku, Nagoya, Aichi, 466-8555, Japan
cht15031@nitech.jp, t-izumi@nitech.ac.jp

Abstract. Communication complexity theory is a powerful tool to show time complexity lower bounds of distributed algorithms for global problems such as minimum spanning tree (MST) and shortest path. While it often leads to nearly-tight lower bounds for many problems, polylogarithmic complexity gaps still lie between the currently best upper and lower bounds. In this paper, we propose a new approach for filling the gaps. Using this approach, we achieve tighter deterministic lower bounds for MST and shortest paths. Specifically, for those problems, we show the deterministic $\Omega(\sqrt{n})$-round lower bound for graphs of $O(n^{\epsilon})$ hop-count diameter, and the deterministic $\Omega(\sqrt{n/\log n})$ lower bound for graphs of $O(\log n)$ hop-count diameter. The main idea of our approach is to utilize the two-party communication complexity lower bound for a function we call *permutation identity*, which is newly introduced in this paper.

1 Introduction

In distributed computing theory, many graph problems are naturally treated as the problems in networks, where each vertex represents a computing entity (node) and each edge does a communication link between two nodes. The theory of *distributed graph algorithms* has been developed so far for the efficient in-network computation of graph problems. A crucial factor of distributed graph algorithms is *locality*. Local algorithms require each node to compute its output only by the interaction to the nodes within a bounded distance smaller than the diameter of the network. In other words, local algorithms must terminate within $o(D)$ rounds, where D is the hop-count diameter of the network. There are a number of problems allowing local solutions: Maximal matchings, colorings, independent sets, and so on. On the other hand, some of other graph problems (e.g., minimum spanning tree, shortest path, minimum cut, and so on) are known to have no local solution. They are called *global problems*. By the definition, the (worst-case) run of any algorithm for global problems inherently takes $\Omega(D)$ rounds.

For both local and global problems, the time complexity is one of the central measures to evaluate distributed algorithms. In this paper, we focus on the

* This work is supported in part by KAKENHI No. 25106507 and No. 25289114.

G.F. Italiano et al. (Eds.): SOFSEM 2015, LNCS 8939, pp. 377–388, 2015.

distributed complexity of two well-known global problems: Minimum spanning tree (MST) and shortest s-t path. As we stated above, these problems have trivial $\Omega(D)$-round lower bounds. If the communication bandwidth of each link is not bounded, every global problem has an optimal-time algorithm with $O(D)$ rounds: A process aggregates the whole information of the network, and computes the result locally. However the assumption of so rich bandwidth is far from real systems, and thus the challenge of global problems is to solve them in the environment with a limited bandwidth. Theoretically, such environments are called *CONGEST model*, where processes work under the round-based synchrony, and each link can transfer $O(\log n)$-bit messages per one round.

A seminal result about the lower bounds for global problems is the one by Das Sarma et al. [1], which exhibits that many problems, including MST and shortest s-t path, are more expensive tasks. Precisely, it shows that p $\Omega(\sqrt{n}/\log n + D)$-round lower bounds hold for many global problems even if D is small (i.e., $D = O(\log n)$). The core of this result is a general framework to obtain the lower bounds based on the reduction from two-party communication complexity by Yao [15]. Two-party communication complexity is a theory to reveal the amount of communication to compute a global function whose inputs are distributed among two players. The reduction framework in [1] induces the hardness of MST and shortest s-t path from the two-party communication complexity of *set-disjointness* function. While the framework is a powerful tool to bound the time complexity of global problems, all the bounds obtained by that approach have the form of $\Omega(f(n)/(m \log n))$, where $f(n)$ is the amount of information inherently exchanged among the network to solve the target problem, m is the number of the links where the information must be transferred, and $\log n$ factor is the bandwidth of each link (that is, $m \log n$ is the amount of information transmittable within a round). Unfortunately, that form does not strictly match the known corresponding upper bounds, which typically have the form of $O(f(n)\text{polylog}(n)/m)$. That is, for many global problems, the currently best bounds still have (poly)logarithmic gaps.

The objective of this paper is to close those gaps. For that goal, we introduce a new two-party function called *permutation identity*, whose deterministic communication complexity is slightly more expensive than set-disjointness, and show new reductions using it on the top of the framework by Das Sarma et al. [1]. Specifically, for MST and shortest s-t path, we show the deterministic $\Omega(\sqrt{n})$-round lower bound for graphs of $O(n^\epsilon)$ hop-count diameter, and the deterministic $\Omega(\sqrt{n/\log n})$ lower bound for graphs of $O(\log n)$ hop-count diameter. The comparison with the prior work are shown in Table 1. As far as we consider the complexity of *deterministic* and *exact* computation, our bounds beat the currently best ones.

2 Related Work

The paper by Das Sarma et al. [1] is the first one explicitly considering the distributed verification problem, which has given a general framework to lead lower