Universally-Optimal Distributed Algorithms for Known Topologies

Bernhard Haeupler[1], David Wajc[2], Goran Zuzic[3]

[1]ETH Zürich, Carnegie Mellon University [2]Stanford University ETH Zürich

Introduction

Distributed optimization includes some of the most fundamental and well-studied problems in distributed computing like the distributed minimum spanning tree (MST). Advancing our theoretical understanding of distributed optimization has led to a long line of work spanning over four decades that culminated in algorithms that are ***existentially optimal***, meaning that there exists some pathological worst-case topology on which no algorithm can do better.

Still, most networks of interest allow for exponentially faster algorithms. This motivates two questions:

- What network topology parameters determine the complexity of distributed optimization?
- Are there ***universally-optimal*** algorithms that are as fast as possible on ***every*** topology?

We resolve these 25-year-old open problems in the known-topology setting. Summary:

- The distributed runtime on any graph G is lower bounded (up to polylogs) by a combinatorial graph quantity called ***shortcut quality*** of G.
- Furthermore, at least in the known-topology setting, one can design algorithms that run time shortcut quality time, yielding ***universally-optimal algorithms***.
- We apply the same technique on multiple problems, including MST, $(1+\varepsilon)$-min cut, various approximate shortest paths problems, sub-graph connectivity, etc.

CONGEST model

We work in the classic CONGEST model. The upper bounds require an additional known-topology assumption.

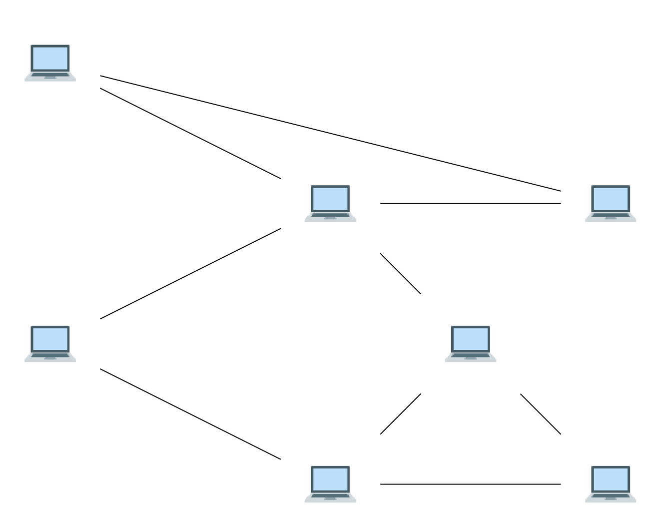

Figure 1. Your favorite large network G.

- Network topology is abstracted as an undirected graph.
- Nodes initially know only their neighborhood.
- Communication in synchronous rounds.
- Each round neighbors exchange $\tilde{O}(1)$-bit messages.
- Computation is free.
- Objective: minimize # rounds.
- (Known topology) Originally, each node knows the entire network (but not the problem input!).

Preliminary: Shortcut quality [Ghaffar, Haeupler, 2015]

In a graph $G = (V, E)$ we are given ***connected node-disjoint*** parts $\mathcal{P} = \{P_1, P_2, \ldots, P_k\}$, $P_i \subseteq V$. A shortcut of quality Q for $\mathcal{P}$ is:

1. (Shortcut edges) Each part P_i gets a set of edges $F_i \subseteq E$.
2. (Dilation) The diameter of $G[P_i] + G[F_i]$ is at most Q.
3. (Congestion) Each edge $e \in E$ is used by at most Q different F_i's.

Preliminary: algorithm design using shortcuts

Definition: $\mathrm{ShortcutQuality}(G)$ is the worst-case quality over all such $\mathcal{P}$.

Theorem [Ghaffari, Haeupler, 2015]: If any set of parts $\mathcal{P}$ admits a shortcut of quality Q (i.e., $\mathrm{ShortcutQuality}(G) \le Q$), and this shortcut can be efficiently constructed, then MST, approximate mincut, approximate shortest path, etc. can be solved in $\tilde{O}(Q)$ CONGEST rounds.

Our results

Theorem (Lower bound): for any graph G, any correct distributed algorithm for MST, SSSP, mincut, etc., requires $\tilde{\Omega}(\mathrm{ShortcutQuality} G)$ rounds.

Theorem (Upper bound): for any graph G, shortcuts of near-optimal quality $\tilde{O}(\mathrm{ShortcutQuality}(G))$ rounds in the known-topology setting. Therefore, MST (and other problems) can be solved in $\tilde{O}(\mathrm{ShortcutQuality}(G))$ rounds.

- Using a new type of an oblivious routing called "hop-constrained oblivious routing" [1].

Lower bound overview

1. We want to prove that $T_{MST}(G) \ge \mathrm{ShortcutQuality}(G)$ where T_{MST} denotes the optimal time to solve MST on G.
2. Suppose $T_{MST} \le Q$. Given parts $\mathcal{P}$, it is sufficient to construct shortcuts of quality $\tilde{O}(Q)$ on $\mathcal{P}$.
3. One can prove that it is sufficient to consider only paths $\mathcal{P}$, and not arbitrary subsets of nodes.
4. It is sufficient to show that the "distributed disjointness task" can be solved on arbitrary node-disjoint paths $\mathcal{P}$ in $\tilde{O}(Q)$ rounds. This would guarantee shortcuts of quality $\tilde{O}(Q)$ exist.
5. We find a "disjointness gadget" on a subset of the paths $\mathcal{P}' \subseteq \mathcal{P}$ of large size.
6. Using the disjointness gadget, we show that solving MST on $\mathcal{P}'$ (which can be done in Q rounda via the oracle) solves the distributed disjointness task. Therefore, shortcuts of quality $\tilde{O}(Q)$ exist.

Lower bound idea I

Definition: distributed disjointness task

- Alice and Bob have k bits $x \in \{0,1\}^k$ and $y \in \{0,1\}^k$ each.
- We are given k node-disjoint paths $P_1, \ldots, P_k$.
- Alice controls the heads of the paths $S = \{s_1, \ldots, s_k\}$; Bob controls the tails $T = \{t_1, \ldots, t_k\}$.
- What is the minimum amount of rounds until we know whether $\exists i \in \{1, \ldots, k\}$ such that $x_i = 1$ and $y_i = 1$.

Theorem: The distributed disjointness on each subset of paths $\{P_1, \ldots, P_k\}$ can be solved in at most Q rounds ***if and only if*** there exists shortcut of quality $\tilde{O}(Q)$ for $P_1, \ldots, P_k$.

- This follows from a network coding gap result [2].

Lower bound idea II

Definition: a disjointness gadget of a set of node-disjoint paths $P_1, \ldots, P_k$ is a connected subset of edges F that touches the heads/tails of each path, but does not otherwise intersect the interior.

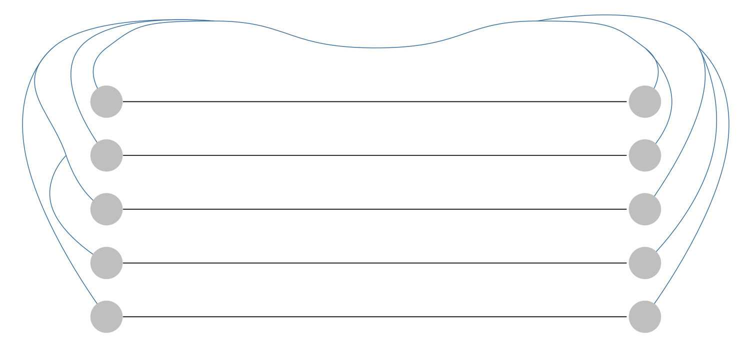

Figure 2. A disjointness gadget

Theorem (main technical contribution): Given any set of node-disjoint paths P, there exists a disjointness gadget on a subset $P' \subseteq P$ of size $|P'| \ge \frac{1}{\mathrm{poly}(\log n)} |P|$.

- The proof is combinatorial and fairly involved.

Relating MST and the distributed disjointness task

Observation: Let F be a disjointness gadget on node-disjoint paths $\mathcal{P}$. Using the MST oracle one can solve the distributed disjointness problem.

- Suppose Alice and Bob have inputs x and y, respectively.
- Assign the MST cost of 0 to (1) edges in F, (2) edges internal to paths $\mathcal{P}$ (not adjacent to the head or tail).
- In path P_i, the edge incident to the head is assigned a cost of x_i.
- In path P_i, the edge incident to the tail is assigned a cost of y_i.
- Easy to verify: there exists MST of cost 0 ***if and only if*** x and y are disjoint.

Conclusions and open questions

- Shortcut Quality of G is the graph parameter that characterizes distributed computing for MST, (approx) SSSP, (approx) mincut, connectivity verification, etc.
- We get algorithms that match this bound in the known topology setting, hence are universally optimal.
- New complexity class of inter-reducible problems in CONGEST.
- On a high-level: we make progress on many questions concerning joint optimization of congestion + dilation.

References

[1] Mohsen Ghaffari, Bernhard Haeupler, and Goran Zuzic.
Hop-constrained oblivious routing.
In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing (STOC)*. ACM, 2021.

[2] Bernhard Haeupler, David Wajc, and Goran Zuzic.
Network coding gaps for completion times of multiple unicasts.
In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*, 2020.