Almost Universally Optimal Distributed Laplacian Solvers via Low-Congestion Shortcuts

Themis Gouleakis
National University of Singapore

August 26th, 2022

Joint work with Ioannis Anagnostides, Bernhard Haeupler, Christoph Lenzen, and Goran Zuzic

Graph Laplacians and Laplacian systems

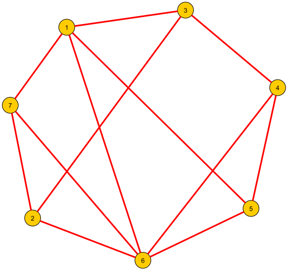

- Let $G = (V, E, \vec{w} > 0)$
- The *Laplacian* of G is:

$$\mathbf{L}(G)_{u,v} = \begin{cases} \sum_{\{u,z\} \in E} w(u,z) & \text{If } u = v, \\ -w(u,v) & \text{otherwise.} \end{cases}$$

- Laplacian system solving: Find **x** such that **L**(G)**x**=**b**.
  - Note that: **L**(G) is: (i) **symmetric**, (ii) **PSD** and (iii) **weakly diagonally dominant**

Prior work and applications

- **First** nearly-linear algorithm by [Spielman Teng 04]
- Applications in **graph algorithms**:
  - Max flow [Madry 13],[Madry 16]
  - Bipartite matching [Madry 13],[Madry 16], [Cohen Madry Sankowski Vladu 17]
  - min-cost flow [Cohen Madry Sankowski Vladu 17], [Axiotis Madry Vladu 20]
  - multicommodity flow [Koutis Miller Peng 10], [Kelner Lee Orecchia Sidford 14]
- **More applications:** Simulating electrical circuits, machine learning, image processing, network analysis, etc.

(SUPPORTED) CONGEST MODEL

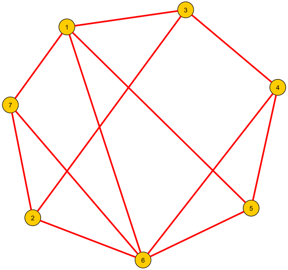

- Synchronous message passing on the edges of the *communication* graph $\bar{G} = (V, E)$.
- $O(\log n)$ message size/internal computation each round.
- complexity=$\sharp$ rounds.
- Supported CONGEST variant: Each node knows $\bar{G}$ in advance.

NODE CAPACITATED CLIQUE (NCC) MODEL

CONGESTED CLIQUE model [Lotker Pavlov Patt-Shamir Peleg 03]

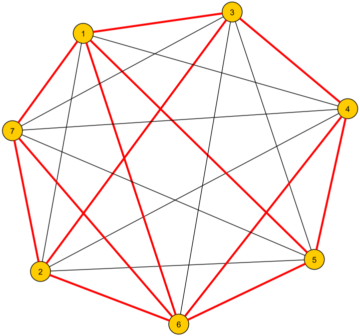

- *Complete* communication graph: $\bar{G} = K_{|V|}$.
- *Synchronous* messages of size $O(\log n)$ bits
- complexity=# rounds.
- NCC variant: Each node may *only* send or receive $O(\log n)$ messages in each round.

Hybrid Models

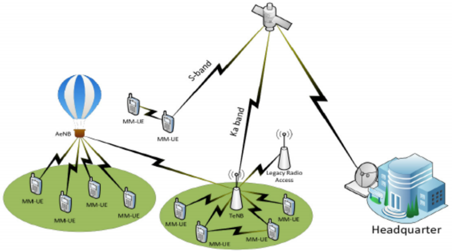

- The HYBRID model was recently introduced in [Augustine Hinnenthal Kuhn Scheideler Schneider 20]

Hybrid Models

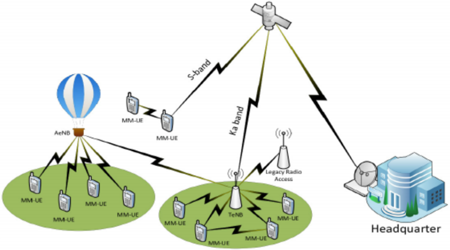

- The HYBRID model was recently introduced in [Augustine Hinnenthal Kuhn Scheideler Schneider 20]
- It enhances the standard LOCAL model with limited global power

HYBRID MODELS

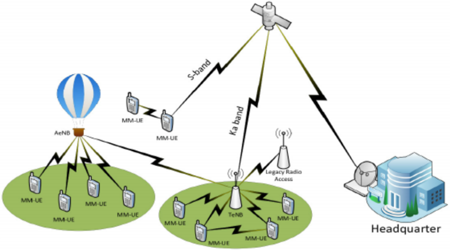

- The HYBRID model was recently introduced in [Augustine Hinnenthal Kuhn Scheideler Schneider 20]
- It enhances the standard LOCAL model with limited global power
- The global network is realized using the node-capacitated clique (NCC)

Hybrid Models

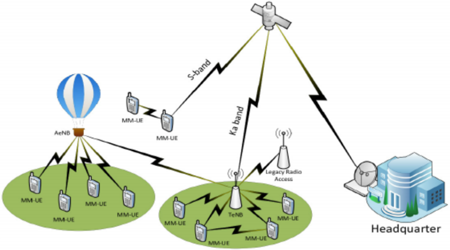

- The HYBRID model was recently introduced in [Augustine Hinnenthal Kuhn Scheideler Schneider 20]
- It enhances the standard LOCAL model with limited global power
- The global network is realized using the node-capacitated clique (NCC)
    - Can communicate with at most $O(\log n)$ other entities per round

Hybrid Models

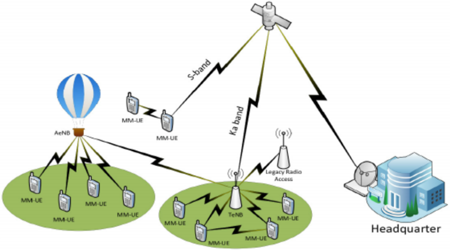

- The HYBRID model was recently introduced in [Augustine Hinnenthal Kuhn Scheideler Schneider 20]
- It enhances the standard LOCAL model with limited global power
- The global network is realized using the node-capacitated clique (NCC)
  - Can communicate with at most $O(\log n)$ other entities per round
- HYBRID leads to much faster algorithms that each model separately

Prior Work

CONGEST model:

- (almost) Existentially optimal Laplacian solver in $n^{o(1)}(\sqrt{n} + D)\log(1/\epsilon)$ rounds. [Forster Goranci Liu Peng Sun 20]
  - Worst- case $\tilde{\Omega}(\sqrt{n} + D)$ round lower bound
- Structure helps: e.g MST and Min-Cut on planar graphs in $\tilde{O}(D)$ rounds via *low-congestion shortcuts*. [Ghaffari Haeupler 16]

HYBRID models (LOCAL + NCC and CONGEST + NCC):

- Polynomial improvements for general graphs in APSP, SSSP [Kuhn and Schneider 20], [Censor-Hillel Leitersdorf Polosukhin 20][Censor-Hillel Leitersdorf Polosukhin 21] [Anagnostides G 21]
- Exponential improvements (i.e $\text{poly}\log n$ round algorithms) in MST, Min-cut, etc . [Feldmann Hinnenthal Scheideler 20], [Anagnostides G 21]

(Congested) Part-Wise Aggregation

Congested Part-Wise Aggregation Problem

- Let $G = (V(G), E(G))$
- Suppose $P_1, \ldots, P_k \subseteq V(G)$ such that $G[P_i]$ is *connected* and each $v \in V(G)$ is contained in at most ρ parts P_i.
- Goal: Each node $v \in V(G)$ learns $\bigoplus_{w \in P_i} \vec{x}(w)$ (e.g min,sum, AND...)

- Setting $\rho = 1$ defines the standard (non-congested) version.
- How fast can this problem be solved?

Low-congestion shortcuts

Low-congestion shortcuts [Ghaffari Haeupler 16]

- Let $G = (V(G), E(G))$
- Suppose $P_1 \uplus \cdots \uplus P_k \subseteq V(G)$ such that $G[P_i]$ is *connected*
- Let $H_1, \ldots, H_k$ such that
  - The diameter of $G[P_i] + H_i$ is at most d (**dilation**)
  - Every $e \in E(G)$ appears in at most c (**congestion**) subgraphs.
- $Q = c + d$ is called **shortcut quality**.

- Remark: Given a shortcut of quality Q, we can solve this problem in $\tilde{O}(Q)$ rounds.
- $SQ(G)$ is the minimum shortcut quality of the worst case valid partition of the vertices of G

Low-congestion shortcuts

Let Q be the number of rounds required to solve part-wise aggregation.

- For general graphs $Q = \Omega(\sqrt{n} + D)$, but for more structured families Q is substantially smaller (e.g $O(D)$ for planar graphs [Ghaffari Haeupler 16])
- For MST one can obtain $O(Q \log n)$ complexity
- For approximate Min-Cut $\widetilde{O}(Q)$

Laplacian solver via Part-wise aggregation

- Suppose that: solving ρ-congested part-wise aggregation problem takes $Q(\rho) = O(\rho^c \mathcal{Q})$ rounds for $c \geq 1$.
  - The constant $\mathcal{Q}$ will depend on the communication model

Theorem

Consider a weighted $\overline{n}$-node graph $\bar{G}$ for which the above holds for some $Q(\rho) = O(\rho^c \mathcal{Q})$, where c is a universal constant and $\mathcal{Q} = \mathcal{Q}(\bar{G})$ is some parameter. Then, we can solve any Laplacian system after $\overline{n}^{o(1)} \mathcal{Q} \log(1/\epsilon)$ rounds.

REDUCTION TO $\rho = 1$?

- Can we reduce solving the ρ-congested part-wise aggregation problem to solving 1-congested part-wise aggregation problems?

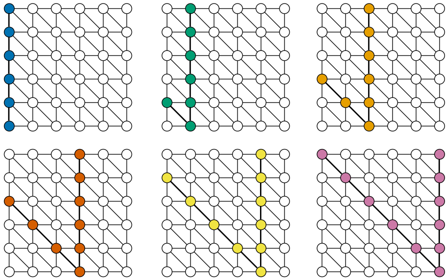

- We would need $\Omega(\sqrt{n})$ problems even for $\rho = 2$.

Layered graph

- We introduce the *layered graph* $\hat{G}_\rho$

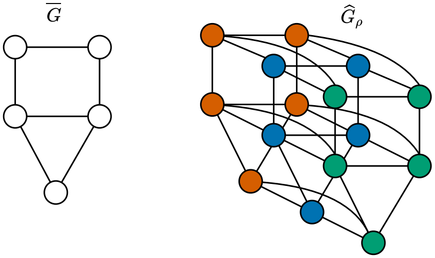

Lemma

Let $\bar{G}$ be an $\bar{n}$-node graph and let $\mathbb{Z}_{\geq 1} \ni \rho \leq \text{poly}(\bar{n})$. If any (1-congested) part-wise aggregation on $\hat{G}_{O(\rho)}$ takes τ CONGEST rounds on $\hat{G}_{O(\rho)}$, then there exists an $\tilde{O}(\rho \cdot \tau)$-round CONGEST algorithm on $\bar{G}$ that solves any ρ-congested part-wise aggregation instance on $\bar{G}$.

PROOF SKETCH

Special case: Each P_i in the part-wise aggregation problem is a path.

- Let $\mathcal{P} = \{P_1, P_2, \ldots, P_k\}$, where $P_i \subseteq V(\bar{G})$ is a simple path and E_i: edges in P_i.
- Goal: Find $\mathcal{P}' = \{P'_1, P'_2, \ldots, P'_k\}$, where $P'_i \subseteq V(\hat{G}_{O(\rho)})$ such that: solving 1-congested PWA on $\mathcal{P}' \Rightarrow \rho$-congested PWA on $\mathcal{P}$.
- Consider subgraph $G^* := (V(\bar{G}), \biguplus_{i=1}^{k} E_i)$ with $\Delta(G') = 2\rho$.
- Color the edges $\{u, v\} \in E(G')$ with colors $c(u, v) \in \{1, \ldots, O(\Delta)\}$ in $O(\log n)$ $CONGEST$ rounds [Johansson 99]
- Construct disjoint paths $P'_i \subseteq V(\hat{G}_{O(\rho)})$: For each $\{u, v\} \in E(G^*)$, add $u_{c(u,v)}, v_{c(u,v)}$ to P'_i.
- Simulate τ-round 1-congested PWA algorithm on $\hat{G}_{O(\rho)}$ ($\Leftrightarrow \rho-$ congested PWA on $\bar{G}$) using $\tilde{O}(\tau \cdot \rho)$ rounds on $\bar{G}$.

OUR RESULTS FOR $CONGEST$

Theorem

For any $\bar{n}$-node graph $\bar{G}$ and any $\mathbb{Z}_{\geq 1} \ni \rho \leq \text{poly}(\bar{n})$, we have that $SQ(\hat{G}_\rho) = \tilde{O}(SQ(\bar{G}))$.

Corollary

There exists a randomized distributed algorithm that, for any $\bar{n}$-node graph $\bar{G}$ and $\rho \in \mathbb{Z}_{\geq 1} \leq \text{poly}(\bar{n}gou)$, solves with high probability any ρ-congested part-wise aggregation instance on $\bar{G}$ with the following guarantees:

- In $CONGEST$, the algorithm terminates in at most $\rho \cdot \text{poly}\Big(SQ(\bar{G})\Big) \cdot \overline{n}^{o(1)}$ rounds.
- In Supported-$CONGEST$, the algorithm terminates in $\tilde{O}(\rho \cdot SQ(\bar{G}))$ rounds.

OUR RESULTS FOR $HYBRID$

Lemma [Augustine Ghaffari Gmyr Hinnenthal Scheideler Kuhn Li 19]

Let $\overline{G}$ be an $\overline{n}$-node communication network. Then, we can solve with high probability any ρ-congested part-wise aggregation problem on $\overline{G}$ after $O(\rho + \log \overline{n})$ rounds of NCC.

Theorem

Consider any n-node graph. There exists a distributed Laplacian solver in the HYBRID model with round complexity $n^{o(1)} \log(1/\epsilon)$, where $\epsilon > 0$ is the error of the solver.

Lower bounds-Universal optimality

Theorem

Consider a graph $\bar{G}$ with shortcut quality $SQ(\bar{G})$. Then, solving a Laplacian system on $\bar{G}$ with $\epsilon \leq \frac{1}{2}$ requires $\widetilde{\Omega}(SQ(\bar{G}))$ rounds even in the Supported-$CONGEST$ model.

- Combined with our $O(n^{o(1)}SQ(\bar{G}))$ upper, this implies universal optimality up to an $n^{o(1)}$ factor.
  - That is, optimal for any given graph G (in the worst case input).
- Improvement over the known existentially optimal $n^{o(1)}(\sqrt{n} + D)$ algorithm by [Forster Goranci Liu Peng Sun 20].

Summary- Open problems

- We provide an (almost) universally optimal distributed algorithm for Laplacian systems in $CONGEST$ and $HYBRID$.
- We use a *unified approach* via ρ-congested part-wise aggregation.
- Open problem: Efficient shortcut construction for general graphs in the CONGEST model?
- Can the techniques be applied to other graph problems?