

Distributed Laplacian Solving with Applications

Sebastian Forster, né Krinninger

University of Salzburg

SIROCCO 2022

Joint work with Gramoz Goranci, Yang P. Liu, Richard Peng, Xiaorui Sun, Tijn de Vos, and Mingquan Ye



This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 947702). Supported by the Austrian Science Fund (FWF): P 32863-N

Laplacian Paradigm

- Laplacian systems
- Spectral sparsifiers
- Electrical flow
- Effective resistance
- Expander decompositions
- Continuous optimization
- Interior-point methods
- Gradient descent
- Preconditioning
- ...



Laplacian Paradigm and Distributed Computing

Observation

Laplacian paradigm often yields inherently parallelizable algorithms

Laplacian Paradigm and Distributed Computing

Observation

Laplacian paradigm often yields inherently parallelizable algorithms

Basic operation:

- Vector \mathbf{j} : each node represents a coordinate
- Matrix \mathbf{L} : each edge represents a non-zero entry
- Matrix-vector multiplication $\mathbf{L}\mathbf{j}$: one round

Laplacian Paradigm and Distributed Computing

Observation

Laplacian paradigm often yields inherently parallelizable algorithms

Basic operation:

- Vector \mathbf{v} : each node represents a coordinate
- Matrix \mathbf{L} : each edge represents a non-zero entry
- Matrix-vector multiplication $\mathbf{L}\mathbf{v}$: one round

State of the art for (approximate) single-source shortest path, maximum flow, minimum-cost flow:

[Ghaffari, Karrenbauer, Kuhn, Lenzen, Patt-Shamir '15] [Becker, F, Karrenbauer, Lenzen '17] [Zuzic '21] [Anagnostides, Themis Gouleakis, Christoph Lenzen '21] [Zuzic, Goranci, Ye, Haeupler, Sun '22] [Rozhon, Grunau, Haeupler, Zuzic, Li '22]

Laplacian Systems

Goal

Solve linear system $\mathbf{L}x = \mathbf{b}$ / \mathbf{T} such that \mathbf{L} is a **Laplacian matrix**.

Laplacian Systems

Goal

Solve linear system $\mathbf{L}x = b$ / T such that \mathbf{L} is a **Laplacian matrix**.

Definition

The **Laplacian matrix** \mathbf{L} of graph $G = (V, E)$ is defined by

$$L_{ij} = \begin{cases} \sum_{k \sim i} 1 & \text{if } i = j, \\ -1 & \text{if } i \sim j, \\ 0 & \text{otherwise.} \end{cases}$$

Laplacian Systems

Goal

Solve linear system $\mathbf{L} \mathbf{j} = \mathbf{T}$ such that \mathbf{L} is a **Laplacian matrix**.

Definition

The **Laplacian matrix** \mathbf{L} of graph $G = (V, E)$ is defined by

$$L_{ij} = \begin{cases} \sum_{k \sim i} 1 & \text{if } i = j, \\ -1 & \text{if } i \sim j, \\ 0 & \text{otherwise.} \end{cases}$$

High-precision solver: Approximation of solution \mathbf{j} with \mathbf{j} s.t.

$$\mathbf{j} = \mathbf{L}^{-1} \mathbf{T} \approx \mathbf{L}^{-1} \mathbf{T}$$

Laplacian Systems

Goal

Solve linear system $\mathbf{A} \mathbf{x} = \mathbf{b}$ / \mathbf{T} such that \mathbf{A} is a **Laplacian matrix**.

Definition

The **Laplacian matrix** \mathbf{A} of graph $G = (V, E)$ is defined by

$$A_{ij} = \begin{cases} \deg(i) & \text{if } i = j, \\ -1 & \text{if } (i, j) \in E, \\ 0 & \text{otherwise.} \end{cases}$$


High-precision solver: Approximation of solution \mathbf{x} with $\mathbf{A} \mathbf{x} = \mathbf{b}$ s.t.

$$\|\mathbf{A} \mathbf{x} - \mathbf{b}\| \leq \epsilon \quad \mathbf{T} \mathbf{x} = \mathbf{z}$$

Prior work:

- $\mathcal{O}(n^3)$ sequential running time [Spielman, Teng '04]
- $\mathcal{O}(n^3)$ work, polylogarithmic depth [Peng, Spielman '14]

CONGEST Model



- Edges correspond to non-zero entries of matrix
- Each node holds one row/column of matrix
- Communication over edges in synchronous rounds
- Bandwidth γ per edge

Our Results for the CONGEST Model

Theorem ([F, Goranci, Liu, Peng, Sun, Ye])

\exists a randomized algorithm that, in the CONGEST model, computes a $(1 + \epsilon)$ -approximate maximum flow in $\tilde{O}(n^2 \log n / \epsilon)$ rounds with high probability.

Almost matches a $\Omega(n^2)$ lower bound

Our Results for the CONGEST Model

Theorem ([F, Goranci, Liu, Peng, Sun, Ye])

Let G be a graph with n nodes and m edges. Let Δ be the maximum degree of G . Let D be the diameter of G . Let $\epsilon \in (0, 1]$. Then there exists a randomized algorithm that runs in $\tilde{O}(n^2 \Delta^2 / \epsilon)$ time and uses $\tilde{O}(n^2 \Delta^2 / \epsilon)$ bits of communication, and achieves a $(1 - \epsilon)$ -approximation to the maximum flow value.

Almost matches a $\tilde{O}(n^2 \Delta^2 / \epsilon)$ flow bound

Implications

$\tilde{O}(n^2 \Delta^2 / \epsilon)$ ϵ -round algorithms in CONGEST model for the following problems:

- Maximum flow [Mądry '16]
- Unit capacity minimum cost flow [Cohen et al. '17]
- Negative weight shortest path [Cohen et al. '17]

First $\tilde{O}(n^2 \Delta^2 / \epsilon)$ -round algorithms for sparse, low-diameter graphs

Approximate Schur Complement

Definition (Schur complement)

For an $n \times n$ symmetric matrix A and a subset of $\mathcal{M} \subseteq \{1, \dots, n\}$, let P be a permutation matrix. Permute the rows/columns of A to write

$$P^T A P = \begin{bmatrix} A_{MM} & A_{M\bar{M}} \\ A_{\bar{M}M} & A_{\bar{M}\bar{M}} \end{bmatrix}$$

Then the *exact* Schur complement of A_{MM} in A is defined as

$$S = A_{\bar{M}\bar{M}} - A_{\bar{M}M} A_{MM}^{-1} A_{M\bar{M}}$$

Result of block Gaussian elimination

Approximate Schur Complement

Definition (Schur complement)

For an $n \times n$ symmetric matrix A and a subset of $V \subseteq \{1, \dots, n\}$, let P be a permutation that permutes the rows/columns of A to write


$$P^T A P = \begin{bmatrix} A_{VV} & A_{V\bar{V}} \\ A_{\bar{V}V} & A_{\bar{V}\bar{V}} \end{bmatrix}$$

Then the *exact* Schur complement of $A_{\bar{V}\bar{V}}$ onto V is defined as


$$S = A_{VV} - A_{V\bar{V}} A_{\bar{V}\bar{V}}^{-1} A_{\bar{V}V}$$

Result of block Gaussian elimination

Graphical interpretation:



Input graph



Schur complement

Approximate Schur Complement

Definition (Schur complement)

For an $n \times n$ symmetric matrix A and a subset of nodes $S \subseteq V$, let P be a permutation that permutes the rows/columns of A to write


$$P^T A P = \begin{bmatrix} A_{SS} & A_{S\bar{S}} \\ A_{\bar{S}S} & A_{\bar{S}\bar{S}} \end{bmatrix}$$

Then the L factor of the LU decomposition of $A_{\bar{S}\bar{S}}$ onto \bar{S} is defined as


$$L_{\bar{S}\bar{S}} = P_{\bar{S}\bar{S}}^T (L_{\bar{S}\bar{S}}^{\#})^{-1} P_{\bar{S}\bar{S}}$$

Result of block Gaussian elimination


Graphical interpretation:



Input graph



Schur complement



Sparsification

Minor Sparsifiers

Problem

Communication of edges “along” sparsifier edges may lead to too much congestion



Minor Sparsifiers

Problem

Communication of edges “along” sparsifier edges may lead to too much congestion

Solution

Vertex sparsifiers as **minors** of the communication graph





Minor Sparsifiers

Problem

Communication of edges “along” sparsifier edges may lead to too much congestion

Solution

Vertex sparsifiers as **minors** of the communication graph



Lemma

? Sfdj žhVfad_gf[b{Uf[a` [ha'h` Y_ [adebSde[[VdfS] Vē
/ - ǫ fcbg` Vež



Minor Sparsifiers

Problem

Communication of edges “along” sparsifier edges may lead to too much congestion

Solution

Vertex sparsifiers as **minors** of the communication graph



Lemma

\exists $S \subseteq E$ such that $G[S]$ is a (ϵ, ϵ) -sparsifier of G and $|S| \leq \epsilon |E|$

Key contribution: Parallel variant of [Li Schild '18]

Technical Details

- From [Kyng, Lee, Peng, Sachdeva, Spielman '16]: Repeated elimination of **almost independent sets** yields vertex sparsifier “chain” with recursion depth $\frac{1}{\epsilon} \log \frac{1}{\epsilon}$

Technical Details

- From [Kyng, Lee, Peng, Sachdeva, Spielman '16]: Repeated elimination of **almost independent sets** yields vertex sparsifier “chain” with recursion depth $\frac{1}{\epsilon} \log \frac{1}{\epsilon}$
 - Fast computation of inverse of submatrix of eliminated nodes using iterative method

Technical Details

- From [Kyng, Lee, Peng, Sachdeva, Spielman '16]: Repeated elimination of **almost independent sets** yields vertex sparsifier “chain” with recursion depth $\log n$
 - Fast computation of inverse of submatrix of eliminated nodes using iterative method
- In addition to Schur complement itself, we need to compute further information (linear operators)

Technical Details

- From [Kyng, Lee, Peng, Sachdeva, Spielman '16]: Repeated elimination of **almost independent sets** yields vertex sparsifier “chain” with recursion depth $\frac{1}{\epsilon} \log \frac{1}{\epsilon}$
 - Fast computation of inverse of submatrix of eliminated nodes using iterative method
- In addition to Schur complement itself, we need to compute further information (linear operators)
 - Obtain Schur complement from sampling random walks

Technical Details

- From [Kyng, Lee, Peng, Sachdeva, Spielman '16]: Repeated elimination of **almost independent sets** yields vertex sparsifier “chain” with recursion depth $\log n / \epsilon$
 - Fast computation of inverse of submatrix of eliminated nodes using iterative method
- In addition to Schur complement itself, we need to compute further information (linear operators)
 - Obtain Schur complement from sampling random walks
 - Algorithmically: estimate of $L_{\setminus V}$ in random walks

Technical Details

- From [Kyng, Lee, Peng, Sachdeva, Spielman '16]: Repeated elimination of **almost independent sets** yields vertex sparsifier “chain” with recursion depth $\frac{1}{\epsilon} \log \frac{1}{\epsilon}$
 - Fast computation of inverse of submatrix of eliminated nodes using iterative method
- In addition to Schur complement itself, we need to compute further information (linear operators)
 - Obtain Schur complement from sampling random walks
 - Algorithmically: estimate of $L_B^{-1} W_B^{-1} a$ in random walks
 - Computation introduces round-overhead of $\frac{1}{\epsilon} \log \frac{1}{\epsilon}$; can only work with $\frac{1}{\epsilon} \log \frac{1}{\epsilon}$ sophisticated recursion

Technical Details

- From [Kyng, Lee, Peng, Sachdeva, Spielman '16]: Repeated elimination of **almost independent sets** yields vertex sparsifier “chain” with recursion depth $\log n$
 - Fast computation of inverse of submatrix of eliminated nodes using iterative method
- In addition to Schur complement itself, we need to compute further information (linear operators)
 - Obtain Schur complement from sampling random walks
 - Algorithmically: estimate of $L_B^{-1} W_B^{-1} a$ in random walks
 - Computation introduces round-overhead of $\log n$; can only work with $\log n$ sophisticated recursion
- Minor sparsifiers: avoid sequential sampling of [Li Schild '18]

Technical Details

- From [Kyng, Lee, Peng, Sachdeva, Spielman '16]: Repeated elimination of **almost independent sets** yields vertex sparsifier “chain” with recursion depth $\log n$
 - Fast computation of inverse of submatrix of eliminated nodes using iterative method
- In addition to Schur complement itself, we need to compute further information (linear operators)
 - Obtain Schur complement from sampling random walks
 - Algorithmically: estimate of $L_B^{-1} W_B A^{-1}$ in random walks
 - Computation introduces round-overhead of $\log n$; can only work with $\log^2 n$ sophisticated recursion
- Minor sparsifiers: avoid sequential sampling of [Li Schild '18]
 - Identify “steady” edges that can be sampled independently
 - Requires recursive solution of linear system: edge reduction via ultra-sparsifiers

Technical Details

- From [Kyng, Lee, Peng, Sachdeva, Spielman '16]: Repeated elimination of **almost independent sets** yields vertex sparsifier “chain” with recursion depth $\log n$
 - Fast computation of inverse of submatrix of eliminated nodes using iterative method
- In addition to Schur complement itself, we need to compute further information (linear operators)
 - Obtain Schur complement from sampling random walks
 - Algorithmically: estimate of $L_B^{-1} W_B [a]$ in random walks
 - Computation introduces round-overhead of $\log n$; can only work with $\log^2 n$ sophisticated recursion
- Minor sparsifiers: avoid sequential sampling of [Li Schild '18]
 - Identify “steady” edges that can be sampled independently
 - Requires recursive solution of linear system: edge reduction via ultra-sparsifiers
 - Distortion of minor property in recursive calls

Implications

fi-round algorithms in CONGEST model for the following problems:

- Maximum flow [Mądry '16]
- Unit capacity minimum cost flow [Cohen et al. '17]
- Negative weight shortest path: [Cohen et al. '17]

What Next?

Implications

ϵ -round algorithms in CONGEST model for the following problems:

- Maximum flow [Madry '16]
- Unit capacity minimum cost flow [Cohen et al. '17]
- Negative weight shortest path: [Cohen et al. '17]

Question

Sublinear rounds in dense graphs?

What Next?

Implications

$\frac{1}{\epsilon}$ fi-round algorithms in CONGEST model for the following problems:

- Maximum flow [Mądry '16]
- Unit capacity minimum cost flow [Cohen et al. '17]
- Negative weight shortest path: [Cohen et al. '17]


Question

Sublinear #rounds in dense graphs?

Easier Question


Sublinear #rounds on the Broadcast Congested Clique?

Broadcast Congested Clique



- Nodes can communicate with all other nodes [Lotker et al. '05]
- Broadcast $fzWes_W$ message to all nodes [Drucker, Kuhn, Oshman '12]

Broadcast Congested Clique



- Nodes can communicate with all other nodes [Lotker et al. '05]
- Broadcast $fzWes_W$ message to all nodes [Drucker, Kuhn, Oshman '12]
- For many problems: only “trivialization” of CONGEST model upper bounds with $\quad / \#$ is known

Our Results for the BCC

Theorem ([F, de Vos '22])

$A \leq_{\text{BCC}} B$ iff \exists a TCF $f: A \rightarrow B$ such that f is a TCF reduction.

Our Results for the BCC

Theorem ([F, de Vos '22])

Let G be a graph with n nodes and m edges. Let $\epsilon \in (0, 1/2]$ and $\delta \in (0, 1]$. Then, there exists a randomized algorithm that, with probability at least $1 - \delta$, computes a $(1 \pm \epsilon)$ -approximation to the Broadcast Congested Clique (BCC) in $O\left(\frac{m}{\epsilon^2} \log \frac{1}{\delta}\right)$ rounds.

Other Results:

- On the Broadcast Congested Clique, a spectral sparsifier of quality $(1 \pm \epsilon)$ and size $O\left(\frac{m}{\epsilon^2}\right)$ can be computed in $O\left(\frac{m}{\epsilon^2}\right)$ rounds
- On the Broadcast Congested Clique, a Laplacian system can be solved up to high accuracy in $O\left(\frac{m}{\epsilon^2}\right)$ rounds
- On the Broadcast Congested Clique, certain Linear Programs can be solved in $O\left(\frac{m}{\epsilon^2}\right)$ rounds.

Main Idea and Challenges

Linear Programming

Minimize $\sum_j c_j x_j$ subject to $Ax = b, x \geq 0$

Linear Programming

Minimize $\sum_j c_j x_j$ subject to $Ax = b, x \geq 0$

Implementation of [Lee, Sidford '14]:

- Interior point method with $\tilde{O}(n^3)$ rankfliterations
- One linear system solve per iteration

Main Idea and Challenges

Linear Programming

Minimize $\sum_j c_j x_j$ subject to $Ax = b, x \geq 0$

Implementation of [Lee, Sidford '14]:

- Interior point method with \sqrt{n} rankfliterations
- One linear system solve per iteration

Minimum cost flow:

- Rank = #nodes
- Linear system has Laplacian matrix

Key Contribution

Iterative computation of spectral sparsifier [Koutis, Xu '16]:

- Compute a spanner
- Sample non-spanner edges with constant probability

Key Contribution

Iterative computation of spectral sparsifier [Koutis, Xu '16]:

- Compute a spanner
- Sample non-spanner edges with constant probability

Problem

On Broadcast Congested Clique, nodes cannot easily coordinate with neighbors on sampling incident edges

Key Contribution

Iterative computation of spectral sparsifier [Koutis, Xu '16]:

- Compute a spanner
- Sample non-spanner edges with constant probability

Problem

On Broadcast Congested Clique, nodes cannot easily coordinate with neighbors on sampling incident edges

Solution:

- Compute spanner on “probabilistic” graph
- Sample individual edges ad-hoc when needed
- Modification of spanner algorithm of [Baswana, Sen '07]

Optimization vs. Data Structures

[Lee, Sidford '14]:

#iterations: $\frac{1}{\epsilon} \log \frac{1}{\delta}$

Time per iteration: $\frac{1}{\epsilon} \log \frac{1}{\delta}$

[Chen et al. '22]:

#iterations: $\frac{1}{\epsilon} \log \frac{1}{\delta}$

Time per iteration: $\frac{1}{\epsilon} \log \frac{1}{\delta}$

Optimization vs. Data Structures

[Lee, Sidford '14]:

#iterations: $\sqrt{f_i}$

Time per iteration: $\sqrt{f_i}$

Iteration count carries over to
round complexity

[Chen et al. '22]:

#iterations: $\sqrt{L} / \sqrt{f_i}$

Time per iteration: $\sqrt{f_i}$

Running time improvement
does not improve round
complexity

Optimization vs. Data Structures

[Lee, Sidford '14]:

#iterations: $\tilde{O}(n^3)$

Time per iteration: $\tilde{O}(n^2)$

Iteration count carries over to
round complexity

[Chen et al. '22]:

#iterations: $\tilde{O}(n^2)$

Time per iteration: $\tilde{O}(n^3)$

Running time improvement
does not improve round
complexity

Question

Is $\tilde{O}(n^3)$ the right iteration count for min-cost flow LP?

Open Problem

Question

Is $\tilde{O}(n^3)$ the right round complexity for min-cost flow in the BCC?

Question

Is $\frac{1}{5}$ the right round complexity for min-cost flow in the BCC?

- Lower bounds in BCC at least not hopeless
[Frischknecht, Holzer, Wattenhofer '12] [Drucker, Kuhn, Oshman '14] [Censor-Hillel, Kaski, Korhonen, Lenzen, Paz, Suomela] [Holzer, Pinski '15] [Becker, Montealegre, Rapaport, Todinca '18]

Question

Is $\tilde{O}(n^2)$ the right round complexity for min-cost flow in the BCC?

- Lower bounds in BCC at least not hopeless
[Frischknecht, Holzer, Wattenhofer '12] [Drucker, Kuhn, Oshman '14] [Censor-Hillel, Kaski, Korhonen, Lenzen, Paz, Suomela] [Holzer, Pinski '15] [Becker, Montealegre, Rapaport, Todinca '18]
- Better upper bound already interesting for single-source reachability

Almost optimal Laplacian
solvers

Conclusion

Almost optimal Laplacian solvers

Broadcast Congested Clique is an interesting “burning glass”