On Solving Asymmetric Diagonally Dominant Linear Systems in Sublinear Time

Tsz Chiu Kwok[1] Zhewei Wei[2] **Mingji Yang**[2]

[1]Shanghai University of Finance and Economics

[2]Renmin University of China

ITCS 2026

Background: Nearly-Linear-Time Laplacian/SDD Solvers

- solving system of linear equations is a fundamental problem
- solvers for special classes of systems have been extensively studied

- breakthrough: nearly-linear-time solvers for Laplacian / symmetric diagonally dominant (SDD) systems [Spielman-Teng '04]
- $\Longrightarrow$ Laplacian Paradigm

Background: Laplacian/SDD Systems

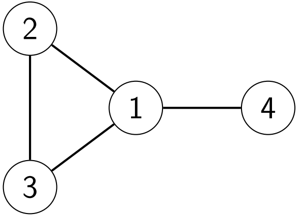

- graph Laplacian $\mathbf{L}_G = \mathbf{D}_G - \mathbf{A}_G$
- $\mathbf{D}_G$: diagonal degree matrix, $\mathbf{A}_G$: adjacency matrix
- an example **Laplacian system**:

$$\begin{pmatrix} 3 & -1 & -1 & -1 \\ -1 & 2 & -1 & 0 \\ -1 & -1 & 2 & 0 \\ -1 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \boldsymbol{x}(1) \\ \boldsymbol{x}(2) \\ \boldsymbol{x}(3) \\ \boldsymbol{x}(4) \end{pmatrix} = \begin{pmatrix} \boldsymbol{b}(1) \\ \boldsymbol{b}(2) \\ \boldsymbol{b}(3) \\ \boldsymbol{b}(4) \end{pmatrix}$$

- an example **symmetric diagonally dominant (SDD) system**:

$$\begin{pmatrix} 3 & 0 & 1 & -0.5 \\ 0 & 2 & -1 & 0 \\ 1 & -1 & 2 & 0 \\ -0.5 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \boldsymbol{x}(1) \\ \boldsymbol{x}(2) \\ \boldsymbol{x}(3) \\ \boldsymbol{x}(4) \end{pmatrix} = \begin{pmatrix} \boldsymbol{b}(1) \\ \boldsymbol{b}(2) \\ \boldsymbol{b}(3) \\ \boldsymbol{b}(4) \end{pmatrix}$$

Background: Nearly-Linear-Time RDD/CDD Solvers

- generalization to nearly-linear-time row/column diagonally dominant (RDD/CDD) solvers [CKPPSV '16] [CKKPPRS '18]
- an example RDD system:

$$\begin{pmatrix} 4 & -2 & 2 \\ 1 & 3 & -1 \\ -1 & 0 & 2 \end{pmatrix} \begin{pmatrix} \boldsymbol{x}(1) \\ \boldsymbol{x}(2) \\ \boldsymbol{x}(3) \end{pmatrix} = \begin{pmatrix} \boldsymbol{b}(1) \\ \boldsymbol{b}(2) \\ \boldsymbol{b}(3) \end{pmatrix}$$

- CDD systems are defined in the natural column-wise way

Sublinear-Time Solvers

- [Andoni-Krauthgamer-Pogrow, ITCS '19]: algorithm for **solving a single entry** of "well-conditioned" SDD systems in sublinear time
- a natural question: can we solve RDD/CDD systems in sublinear time?
- our answer: yes for "well-structured" RDD/CDD systems
  - provided that we properly characterize what it means by "well-structured"

Problem Formulation

- for an RDD/CDD system $\mathbf{M}\boldsymbol{x} = \boldsymbol{b}$ where $\mathbf{M} \in \mathbb{R}^{n \times n}$ and $\boldsymbol{b} \in \mathrm{range}(\mathbf{M})$
- given standard oracle access to $\mathbf{M}$, $\boldsymbol{b}$, and $\boldsymbol{t}$
- **our goal**: compute an approximation of $\boldsymbol{t}^\top \boldsymbol{x}^*$
- where $\boldsymbol{x}^*$ is a particular solution to the system determined by $\mathbf{M}$ and $\boldsymbol{b}$

- examples: single-pair *Personalized PageRank* and *effective resistance* on graphs

Main Contributions

1. **generalization** of the formulation and results in [AKP19] to RDD/CDD systems, via generalizing spectral gap to a novel concept called maximum p-norm gap
2. **more complexity upper bounds** by adapting techniques for local graph algorithms: random-walk sampling, local push, and bidirectional method

a **general and unified framework** for understanding local solvers and graph algorithms

Contents

Power Series Expansion of $\boldsymbol{x}^*$

- we decompose $\mathbf{M}$ as $\mathbf{M} = \mathbf{D_M} - \mathbf{A_M}$, with $\mathbf{D_M}$ being the diagonal part
- **for certain types of** $\mathbf{M}$, we have

$$\mathbf{M}\boldsymbol{x} = \boldsymbol{b} \iff \left(\mathbf{I} - \mathbf{D}_\mathbf{M}^{-1}\mathbf{A_M}\right)\boldsymbol{x} = \mathbf{D}_\mathbf{M}^{-1}\boldsymbol{b} \iff \boldsymbol{x} = \sum_{\ell=0}^{\infty}\left(\mathbf{D}_\mathbf{M}^{-1}\mathbf{A_M}\right)^{\ell}\mathbf{D}_\mathbf{M}^{-1}\boldsymbol{b}$$

- to guarantee convergence for general RDD/CDD $\mathbf{M}$, we choose

$$\boldsymbol{x}^* := \frac{1}{2}\sum_{\ell=0}^{\infty}\left(\frac{1}{2}\left(\mathbf{I} + \mathbf{D}_\mathbf{M}^{-1}\mathbf{A_M}\right)\right)^{\ell}\mathbf{D}_\mathbf{M}^{-1}\boldsymbol{b}$$

Theorem (Property of $\boldsymbol{x}^*$)

1. $\boldsymbol{x}^*$ is well-defined and satisfies $\mathbf{M}\boldsymbol{x}^* = \boldsymbol{b}$;
2. if $\mathbf{M}$ is SDD, then $\boldsymbol{x}^* = \mathbf{D}_\mathbf{M}^{-1/2}\left(\mathbf{D}_\mathbf{M}^{-1/2}\mathbf{M}\mathbf{D}_\mathbf{M}^{-1/2}\right)^{+}\mathbf{D}_\mathbf{M}^{-1/2}\boldsymbol{b}$, matching [AKP19].

Truncated Power Series

$$\boldsymbol{x}^* := \frac{1}{2}\sum_{\ell=0}^{\infty}\left(\frac{1}{2}\left(\mathbf{I}+\mathbf{D}_{\mathbf{M}}^{-1}\mathbf{A}_{\mathbf{M}}\right)\right)^{\ell}\mathbf{D}_{\mathbf{M}}^{-1}\boldsymbol{b}$$

- for a truncation parameter L (to be determined later), we approximate $\boldsymbol{x}^*$ by

$$\boldsymbol{x}_L^* := \frac{1}{2}\sum_{\ell=0}^{L-1}\left(\frac{1}{2}\left(\mathbf{I}+\mathbf{D}_{\mathbf{M}}^{-1}\mathbf{A}_{\mathbf{M}}\right)\right)^{\ell}\mathbf{D}_{\mathbf{M}}^{-1}\boldsymbol{b}$$

- we need to upper bound $\left|\boldsymbol{t}^\top\boldsymbol{x}_L^* - \boldsymbol{t}^\top\boldsymbol{x}^*\right| \le \varepsilon$ in terms of some quantity of $\mathbf{M}$

p-Norm Gaps: Intuition

$$\boldsymbol{x}^* := \frac{1}{2} \sum_{\ell=0}^{\infty} \left(\frac{1}{2} \left(\mathbf{I} + \mathbf{D}_{\mathbf{M}}^{-1} \mathbf{A}_{\mathbf{M}} \right) \right)^{\ell} \mathbf{D}_{\mathbf{M}}^{-1} \boldsymbol{b}$$

- intuitively, stronger diagonal dominance of $\mathbf{M}$ implies faster convergence
- if $\mathbf{M}$ is RDD, consider $1 - \left\| \mathbf{D}_{\mathbf{M}}^{-1} \mathbf{A}_{\mathbf{M}} \right\|_{\infty} = \min_{j \in [n]} \left\{ \frac{\mathbf{M}(j,j) - \sum_{k \neq j} |\mathbf{M}(j,k)|}{\mathbf{M}(j,j)} \right\} \geq 0$
- if $\mathbf{M}$ is CDD, the quantity is $1 - \left\| \mathbf{A}_{\mathbf{M}} \mathbf{D}_{\mathbf{M}}^{-1} \right\|_1 \geq 0$
- this hints us to consider the quantity $1 - \left\| \mathbf{D}_{\mathbf{M}}^{-1/q} \mathbf{A}_{\mathbf{M}} \mathbf{D}_{\mathbf{M}}^{-1/p} \right\|_p$
  - $p \in [1, \infty]$, $1/p + 1/q = 1$
- however, this quantity can be zero, making it useless

p-Norm Gaps

- we define the **p-norm gap** of $\mathbf{M}$ as

$$\gamma_p(\mathbf{M}) := 1 - \left\| \frac{1}{2}\left(\mathbf{I} + \mathbf{D}_{\mathbf{M}}^{-1/q}\mathbf{A}_{\mathbf{M}}\mathbf{D}_{\mathbf{M}}^{-1/p}\right)\Big|_{\mathrm{range}\left(\mathbf{I}-\mathbf{D}_{\mathbf{M}}^{-1/q}\mathbf{A}_{\mathbf{M}}\mathbf{D}_{\mathbf{M}}^{-1/p}\right)} \right\|_p,$$

  where $p \in [1,\infty]$, $1/p + 1/q = 1$
- **maximum p-norm gap**: $\gamma_{\max}(\mathbf{M}) := \max_{p\in[1,\infty]} \gamma_p(\mathbf{M})$

Theorem (Maximum p-Norm Gap)

1. If $\mathbf{M}$ is RDD/CDD, then $0 < \gamma_{\max}(\mathbf{M}) \leq 1$;
2. if $\mathbf{M}$ is SDD, then $\gamma_{\max}(\mathbf{M}) = \gamma_2(\mathbf{M}) = \gamma(\mathbf{M})$, the spectral gap of $\mathbf{M}$.

Truncation Error in Terms of the Maximum p-Norm Gap

Theorem (Truncation Error Bound)

Suppose $0 < \gamma \leq \gamma_{\max}(\mathbf{M})$. To ensure that $\left|\boldsymbol{t}^\top \boldsymbol{x}_L^* - \boldsymbol{t}^\top \boldsymbol{x}^*\right| \leq \varepsilon$, it suffices to set

$$L := \widetilde{\Theta}\left(\frac{1}{\gamma}\right).$$

- our formulation reduces to that in [AKP19] when $\mathbf{M}$ is SDD

Contents

Interpreting PageRank Computation

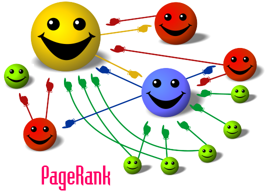

- **Personalized PageRank vector**: $\boldsymbol{\pi}_{G,\alpha,\boldsymbol{s}} = \alpha \boldsymbol{s} + (1-\alpha)\mathbf{A}_G^\top \mathbf{D}_G^{-1} \boldsymbol{\pi}_{G,\alpha,\boldsymbol{s}}$
  - $\alpha \in (0,1)$ is the decay factor
  - $\boldsymbol{s}$ is the source probability distribution
  - $\mathbf{D}_G$ is the diagonal outdegree matrix

equivalently, $$\left(\mathbf{D}_G - (1-\alpha)\mathbf{A}_G^\top\right)\left(\mathbf{D}_G^{-1}\boldsymbol{\pi}_{G,\alpha,\boldsymbol{s}}\right) = \alpha \boldsymbol{s}$$

- $\mathbf{D}_G - (1-\alpha)\mathbf{A}_G^\top$ is invertible CDD and is SDD for undirected G
- $\frac{1}{2}\alpha$ lower bounds the maximum p-norm gap of $\mathbf{D}_G - (1-\alpha)\mathbf{A}_G^\top$

Interpreting Effective Resistance Computation

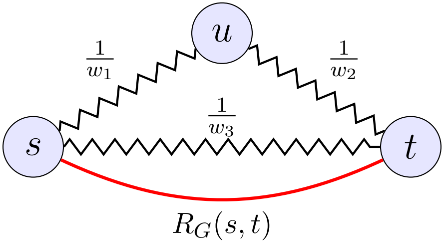

- consider connected undirected graph G
- **effective resistance** $R_G(s,t) = (\boldsymbol{e}_s - \boldsymbol{e}_t)^\top \mathbf{L}_G^+ (\boldsymbol{e}_s - \boldsymbol{e}_t)$
- when $\mathbf{M} = \mathbf{L}_G, \boldsymbol{b} = \boldsymbol{t} = \boldsymbol{e}_s - \boldsymbol{e}_t$ in our framework, $\boldsymbol{t}^\top \boldsymbol{x}^* = R_G(s,t)$

- $\mathbf{L}_G$ is SDD
- $\gamma_{\max}(\mathbf{L}_G) = \gamma(\mathbf{L}_G)$, the spectral gap

Contents

Random-Walk Sampling (1/2)

- we aim to estimate:

$$\boldsymbol{t}^\top \boldsymbol{x}_L^* = \frac{1}{2}\boldsymbol{t}^\top \sum_{\ell=0}^{L-1} \left(\frac{1}{2} \left(\mathbf{I} + \mathbf{D}_{\mathbf{M}}^{-1} \mathbf{A}_{\mathbf{M}} \right) \right)^\ell \mathbf{D}_{\mathbf{M}}^{-1} \boldsymbol{b}$$

- when $\mathbf{M}$ is RDD, $\frac{1}{2}\left(\mathbf{I} + \mathbf{D}_{\mathbf{M}}^{-1} |\mathbf{A}_{\mathbf{M}}|\right)$ is row substochastic
- we can estimate the quantity by **sampling random walks** of length $\ell \in [0, L-1]$
- the random walk starts from distribution $|\boldsymbol{t}|/\|\boldsymbol{t}\|_1$ and transitions via $\frac{1}{2}\left(\mathbf{I} + \mathbf{D}_{\mathbf{M}}^{-1} |\mathbf{A}_{\mathbf{M}}|\right)$

Random-Walk Sampling (2/2)

using the Hoeffding bound, we prove:

Theorem

- Suppose $\mathbf{M}$ is RDD and we are given $0 < \gamma \leq \gamma_{\max}(\mathbf{M})$.
- Suppose we can simulate one step of the random walk in $O(1)$ time.
- Then there exists a randomized algorithm that computes an estimate $\hat{x}$ such that $\Pr\left\{\left|\hat{x} - \boldsymbol{t}^\top \boldsymbol{x}^*\right| \leq \varepsilon \left\|\mathbf{D}_{\mathbf{M}}^{-1}\boldsymbol{b}\right\|_\infty\right\} \geq \frac{3}{4}$ in time

$$\widetilde{O}\left(\|\boldsymbol{t}\|_1^2 \gamma^{-3} \varepsilon^{-2}\right).$$

- this generalizes the algorithmic result in [AKP19] for SDD systems
- [AKP19] proves a complexity lower bound of $\widetilde{\Omega}\left(1/\gamma_{\max}(\mathbf{M})^2\right)$
- we prove a complexity lower bound of $\Omega(1/\varepsilon)$

Counterparts for CDD Systems

- we have

$$\boldsymbol{t}^\top \boldsymbol{x}_L^* = \frac{1}{2}\boldsymbol{t}^\top \sum_{\ell=0}^{L-1} \left(\frac{1}{2}\left(\mathbf{I} + \mathbf{D}_{\mathbf{M}}^{-1}\mathbf{A}_{\mathbf{M}}\right)\right)^{\ell} \mathbf{D}_{\mathbf{M}}^{-1}\boldsymbol{b}$$
$$= \frac{1}{2}\boldsymbol{b}^\top \sum_{\ell=0}^{L-1} \left(\frac{1}{2}\left(\mathbf{I} + \mathbf{D}_{\mathbf{M}}^{-1}\mathbf{A}_{\mathbf{M}}^\top\right)\right)^{\ell} \mathbf{D}_{\mathbf{M}}^{-1}\boldsymbol{t}$$

- if $\mathbf{M}$ is CDD, $\frac{1}{2}\left(\mathbf{I} + \mathbf{D}_{\mathbf{M}}^{-1}\left|\mathbf{A}_{\mathbf{M}}^\top\right|\right)$ is row substochastic
- results for RDD systems can be adapted to CDD ones by **swapping the roles of b and t**

Local Push: Forward Push for PageRank

- push operation [Andersen-Chung-Lang '06] for computing single-source PageRank:

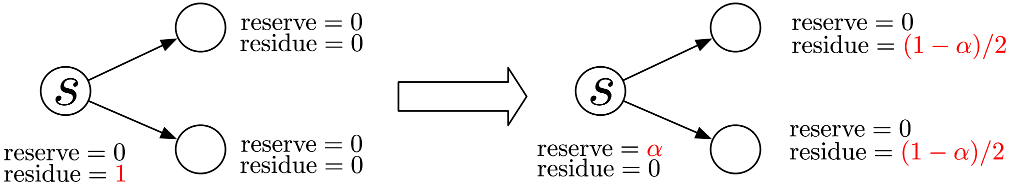

- iteratively perform push operations until all residues are small
- reserves serve as the estimate for $\boldsymbol{\pi}_{G,\alpha,\boldsymbol{e}_s}$
- residues capture the approximation error through an invariant equation

The Algebraic Push Primitive

$$\boldsymbol{x}_L^* = \frac{1}{2}\sum_{\ell=0}^{L-1}\left(\frac{1}{2}\left(\mathbf{I}+\mathbf{D}_{\mathbf{M}}^{-1}\mathbf{A}_{\mathbf{M}}\right)\right)^{\ell}\mathbf{D}_{\mathbf{M}}^{-1}\boldsymbol{b}$$

- `Push` initializes residues as $\mathbf{D}_{\mathbf{M}}^{-1}\boldsymbol{b}$
- propagates residues via $\frac{1}{2}\left(\mathbf{I}+\mathbf{D}_{\mathbf{M}}^{-1}\mathbf{A}_{\mathbf{M}}\right)$

- if $\mathbf{M}$ is RDD, `Push` admits closed-form **accuracy guarantee**
- if $\mathbf{M}$ is CDD, `Push` admits closed-form **complexity bound**

- for RCDD systems, `Push` admits both accuracy and complexity bounds

The Bidirectional Method (1/2)

- `BiPPR` [Lofgren-Banerjee-Goel '16] for estimating $\boldsymbol{\pi}_{G,\alpha,\boldsymbol{e}_s}(t)$

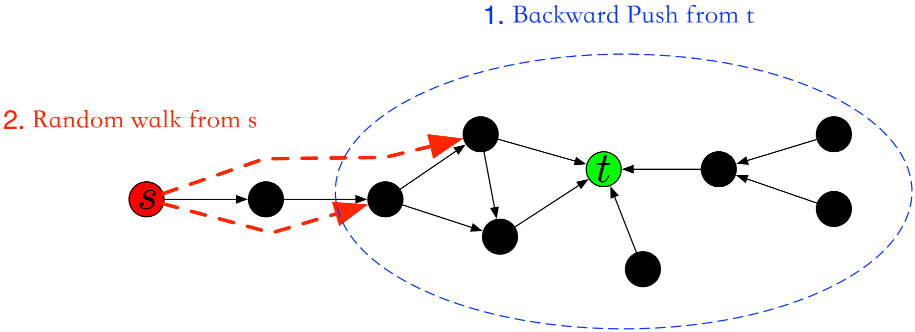

The Bidirectional Method (2/2)

- we aim to estimate:

$$\boldsymbol{t}^\top \boldsymbol{x}_L^* = \frac{1}{2}\boldsymbol{t}^\top \sum_{\ell=0}^{L-1} \left(\frac{1}{2}\left(\mathbf{I} + \mathbf{D}_{\mathbf{M}}^{-1}\mathbf{A}_{\mathbf{M}}\right)\right)^{\ell} \mathbf{D}_{\mathbf{M}}^{-1}\boldsymbol{b}$$

- the invariant property of `Push` implies that

$$\begin{aligned} &\boldsymbol{t}^\top \boldsymbol{x}_L^* - \frac{1}{2}\boldsymbol{t}^\top \left(\sum_{\ell=0}^{L-1} \boldsymbol{p}^{(\ell)} + \boldsymbol{r}^{(L-1)}\right) \\ &= \frac{1}{2}\boldsymbol{t}^\top \sum_{\ell=0}^{L-1} \left(\frac{1}{2}\left(\mathbf{I} + \mathbf{D}_{\mathbf{M}}^{-1}\mathbf{A}_{\mathbf{M}}\right)\right)^{\ell} \left(\sum_{\ell'=0}^{\min(L-\ell-1, L-2)} \boldsymbol{r}^{(\ell')}\right) \end{aligned}$$

- when $\mathbf{M}$ is RDD, we can sample random walks to estimate this difference
- method: `Push` from $\mathbf{D}_{\mathbf{M}}^{-1}\boldsymbol{b}$ + random-walk sampling from $\boldsymbol{t}$ + parameter balancing

Contents

Conclusions and Open Problems

Conclusions

- "well-structured" RDD/CDD systems can be solved in sublinear time
- this general framework unifies sublinear SDD solvers and local graph algorithms for PageRank / effective resistance estimation

Open Problems

- bridge the gaps between upper and lower bounds
- relate the p-norm gaps to combinatorial properties
- find more applications of sublinear RDD/CDD solvers

- Thank you!