

Sparsified Block Elimination for Directed Laplacians*

Richard Peng
University of Waterloo
Waterloo, Canada
y5peng@uwaterloo.ca

Zhuoqing Song
Fudan University
Shanghai, China
zqsong19@fudan.edu.cn

ABSTRACT

We show that the sparsified block elimination algorithm for solving undirected Laplacian linear systems from [Kyng-Lee-Peng-Sachdeva-Spielman STOC'16] directly works for directed Laplacians. Given access to a sparsification algorithm that, on graphs with n vertices and m edges, takes time $\mathcal{T}_S(m)$ to output a sparsifier with $N_S(n)$ edges, our algorithm solves a directed Eulerian system on n vertices and m edges to ϵ relative accuracy in time

$$O(\mathcal{T}_S(m) + N_S(n) \log n \log(n/\epsilon)) + \tilde{O}(\mathcal{T}_S(N_S(n)) \log n),$$

where the $\tilde{O}(\cdot)$ notation hides $\log \log(n)$ factors. By previous results, this implies improved runtimes for linear systems in strongly connected directed graphs, PageRank matrices, and asymmetric M-matrices. When combined with slower constructions of smaller Eulerian sparsifiers based on short cycle decompositions, it also gives a solver algorithm that, after pre-processing the matrix in $O(n^2 \log^{O(1)} n)$ time, takes $O(n \log^5 n \log(n/\epsilon))$ time per solve. At the core of our analyses are constructions of augmented matrices whose Schur complements encode error matrices.

CCS CONCEPTS

• Theory of computation → Data structures design and analysis; Graph algorithms analysis; Preconditioning; Random walks and Markov chains.

KEYWORDS

linear systems, matrices, Markov chains

ACM Reference Format:

Richard Peng and Zhuoqing Song. 2022. Sparsified Block Elimination for Directed Laplacians. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC '22)*, June 20–24, 2022, Rome, Italy. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3519935.3520053>

1 INTRODUCTION

The design of efficient solvers for systems of linear equations in graph Laplacian matrices and their extensions has been a highly

*A full version of this paper is available on arXiv [15]. Richard Peng is supported by National Science Foundation (NSF) under Grant No. 1846218. Part of this work was done at Georgia Tech. Zhuoqing Song is supported by Shanghai Science and Technology Program [Project No. 21JC1400600 and No. 20JC1412700].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '22, June 20–24, 2022, Rome, Italy

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9264-8/22/06...\$15.00

<https://doi.org/10.1145/3519935.3520053>

fruitful topic in algorithms. Laplacian matrices directly correspond to undirected graphs: off-diagonal entries are negations of edge weights, while the diagonal entries contain weighted degrees. Solvers for Laplacian matrices led to breakthroughs in fundamental problems in combinatorial optimization. Tools developed during such studies have in turn influenced data structures, randomized numerical linear algebra, scientific computing, and network science [17, 18].

An important direction in this Laplacian paradigm of designing graph algorithms is extending tools developed for undirected Laplacian matrices to directed graphs. Here a perspective from random walks and Markov chains leads to directed Laplacian matrices [8]. Such matrices have directed edge weights in off-diagonal entries, and weighted out-degrees on diagonals. In contrast to solving linear systems in undirected Laplacians, solving linear systems in directed Laplacians is significantly less well-understood. Almost-linear time [7] and nearly-linear time solvers [6] were developed very recently, and involve many more moving pieces.

In particular, the nearly-linear time algorithm from [6] combined block Gaussian elimination with single variables/vertex elimination, analyzed using matrix Martingales. In contrast, for undirected Laplacians, both block elimination [10] or matrix Martingales [11] can give different nearly-linear time solver algorithms, and there also exists more combinatorial approaches [9]. In this paper, we simplify this picture for directed Laplacian solvers by providing an analog of the sparsified Cholesky/multi-grid solver from [10]. This algorithm's running time is close to the limit of sparsification based algorithms: the running time of invoking a sparsification routine on its own output. Formally, we show:

THEOREM 1.1. *Given a strongly connected Eulerian Laplacian $L \in \mathbb{R}^{n \times n}$ and an error parameter $\epsilon \in (0, 1)$, we can process it in time $O(\mathcal{T}_S(m, n, 1)) + \tilde{O}(\mathcal{T}_S(N_S(n, 1), n, 1) \log n)$ so that, with high probability, given any vector $b \in \mathbb{R}^n$ with $b \perp \mathbf{1}$, we can compute a vector $x \in \mathbb{R}^n$ in time $O(N_S(n, 1) \log n \log(n/\epsilon))$ such that*

$$\|x - L^\dagger b\|_{U[L]} \leq \epsilon \|L^\dagger b\|_{U[L]},$$

where $U[L] = (L + L^\top)/2$.

This result improves the at least $\Omega(\log^5 n)$ factor overhead upon sparsification of the previous nearly-linear time directed Laplacian solver [6], and is analogous to the current best overheads for sparsification based solvers for undirected Laplacians [10]. From the existence of sparsifiers of size $O(n \log^4 n \epsilon^{-2})$ [5], we also obtain the existence of $O(n \log^5 n \log(n/\epsilon))$ time solver routines that require quadratic time preprocessing to compute. As with other improved solvers for directed Laplacians our improvements directly applies to applications of such solvers, including random walk related quantities [8], as well as PageRank / Perron-Frobenius vectors [1].

Our result complements recent developments of better sparsifiers of Eulerian Laplacians [5, 12, 13]. By analyzing a pseudocode that's entirely analogous to the undirected block-elimination algorithm from [10], we narrow the gap between Laplacian solvers for directed and undirected graphs. Our result also emphasizes the need for better directed sparsification routines. While there is a rich literature on undirected sparsification [3], the current best directed sparsification algorithms rely on expander decompositions, so have rather large logarithmic factor overheads. We discuss such bounds in detail in the full version.

Finally, our analysis of this more direct algorithm require better understanding the accumulation errors in Eulerian Laplacians and their partially eliminated states, known as Schur Complements. It was observed in [6] that these objects are significantly less robust than their undirected analogs. Our analysis of these objects rely on augmentations of matrices: constructing larger matrices whose Schur complements correspond to the final objects we wish to approximate, and bounding errors on these larger matrices instead. This approach has roots in symbolic computation, and can be viewed as a generalization of low-rank perturbation formulas such as Sherman-Morrison-Woodbury [19]. We believe both this algebraic technique, and the additional robustness properties of directed Schur Complements we show, may be of independent interest.

1.1 Related Works

Directed Laplacian matrices arise in problems related to directed random walks / non-reversible Markov chains, such as computations of stationary distributions, hitting times and escape probabilities. A formal treatment of applying an Eulerian solver to these problems can be found in [8] and [1]. Adaptations of Eulerian Laplacian solvers have also led to improved bounded-space algorithms for estimating random walk probabilities [2].

Our algorithm is most closely related to the previous nearly-linear time directed Laplacian solver [6]. That algorithm is motivated by single variable elimination and a matrix Martingale based analysis. However, it invokes both components of block elimination algorithms: finding strongly diagonally dominant subsets, and invoking sparsification as black-boxes. The runtime overhead of this routine over sparsification is at least $\log^5 n$: in [6], Lemma 5.1 gives that each phase (for a constant factor reduction) invokes sparsification $O(\log^2 n)$ times, and each call is ran with error at most $\frac{1}{O(\log^3 n)}$ (divided by $\log^2 n$ in Line 2 of Algorithm 2, and also by $\log n$ in Line 5 of Algorithm 3).

While our algorithms are directed analogs of the undirected block elimination routines from [10], our analyses rely on many structures developed in [6]. Specifically, our cumulative error during elimination steps is bounded via the matrix that's the sum of undirectifications of the intermediate directed matrices. On the other hand, we believe our algorithm is more natural: our sampling no longer needs to be locally unbiased, the per-step errors do not need to be decreased by polylog factors, and the algorithm is no longer divided into inner/outer phases. This more streamlined algorithm leads to our runtime improvements.

Our Schur Complement sparsification algorithm is based on the partial block elimination routine from [10], which is in turn based on a two-term decomposition formula for (pseudo-)inverses

from [16]. We remark that there is a good sparsification routine in the low space setting [2]. There is a subsequent algorithm that replaces this decomposition with directly powering via random walks [4] that's also applicable for sparsifying undirected Schur Complements. However, that algorithm relies on sparsifying 3-step random walk polynomials, which to our knowledge, is a subroutine that has not been studied in directed settings. As a result, we are unable to utilize this later development directly.

The existence of $O(n \log^4 n)$ sized sparsifiers in [5] relies on decomposing unit weighted graphs into short cycles and $O(n)$ extra edges. While this decomposition has a simple $O(m^2)$ time algorithm (peel off all vertices with degree < 3 , then return the lowest cross-edge in the BFS tree), the current fastest construction of it takes $m^{1+o(1)}$ time [5, 12, 13]. As a result, we need to instead invoke the more expensive, graph decomposition based, algorithms from [7] for sparsification. Also, we can only use the naive $O(m^2)$ construction of $O(\log n)$ -lengthed cycle decompositions (after an initial sparsification call to make $m = O(n \log^{O(1)} n)$) because the almost-linear time algorithm in [12] produces $O(\log^2 n)$ -lengthed cycles.

2 PRELIMINARY

2.1 Notations

The notation $\tilde{O}(\cdot)$ suppresses the $\text{polyloglog}(n)$ factors in this paper. We let $[n] = \{1, 2, \dots, n\}$.

Matrix: For matrix A , $\text{nnz}(A)$ denotes its number of nonzero entries. For subsets $T_1, T_2 \subseteq [n]$, $A_{T_1 T_2} \in \mathbb{R}^{|T_1| \times |T_2|}$ is the submatrix containing the entries with row indexes and column indexes in (T_1, T_2) ; and $A_{-T_1, -T_2}$ is the submatrix of A by removing the rows indexed by T_1 and columns indexed by T_2 . For vector $v \in \mathbb{R}^n$ and subset $C \subseteq [n]$, v_C is the subvector of v containing the entries indexed by C .

We use I_a , $0_{b \times c}$ to denote the identity matrix of size a and the b -by- c zero matrix, and we sometimes omit the subscripts when their sizes can be determined from the context. For any matrix $X \in \mathbb{R}^{a \times b}$ and set $T_1, T_2 \subseteq [n]$ with $|T_1| = a$, $|T_2| = b$, $\mathcal{P}(X, T_1, T_2, n)$ denotes an n -by- n matrix whose submatrix indexed by (T_1, T_2) equals X and all the other entries equal 0. In other words, $\mathcal{P}(X, T_1, T_2, n)$ can be regarded as replacing the submatrix indexed by (T_1, T_2) with X in the zero matrix $0_{n \times n}$.

For symmetric matrix $A \in \mathbb{R}^{n \times n}$, we use $\lambda_i(A)$ to denote its i -th smallest eigenvalue. For symmetric matrices $A, B \in \mathbb{R}^{n \times n}$, we use $A \succeq B$ ($A > B$) to indicate that for any $x \in \mathbb{R}^n$, $x^\top Ax \geq x^\top Bx$ ($x^\top Ax > x^\top Bx$). We define \leq , $<$ analogously. A square matrix $A \in \mathbb{R}^{n \times n}$ is positive semidefinite (PSD) iff A is symmetric and $A \succeq 0$; $A \in \mathbb{R}^{n \times n}$ is positive definite (PD) iff A is symmetric and $A > 0$.

For PSD matrix A , $A^{1/2}$ is its square root; A^\dagger denotes its Moore-Penrose pseudoinverse; $A^{\dagger/2}$ is the square root of its Moore-Penrose pseudoinverse.

Vector: 1_a , 0_b denote the a -dimensional all-ones vector and b -dimensional all-zeros vector; when their sizes can be determined from the context, we sometimes omit the subscripts.

For matrix $A \in \mathbb{R}^{n \times n}$, $\text{Diag}(A)$ is an n -by- n diagonal matrix with the same diagonal entries as A . For vector $x \in \mathbb{R}^n$, $\text{Diag}(x)$

denotes an n -by- n diagonal matrix with its i -th diagonal entry equalling x_i .

For any positive semidefinite matrix A , we define the vector norm $\|x\|_A = \sqrt{x^\top Ax}$. And $\|\cdot\|_p$ denotes the ℓ_p norm.

Matrix norm: $\|\cdot\|_p$ denotes the ℓ_p norm. For instance, for $A \in \mathbb{R}^{n \times n}$, $\|A\|_2 = \sqrt{\lambda_n(A^\top A)}$; $\|A\|_\infty = \max_{i \in [n]} \sum_{j=1}^n |A_{ij}|$. For matrix $B \in \mathbb{R}^{n \times n}$ and PSD matrix $A \in \mathbb{R}^{n \times n}$, we denote $\|B\|_{A \rightarrow A} = \sup_{\|x\|_A \neq 0} \frac{\|Bx\|_A}{\|x\|_A}$.

Schur complement: For $A \in \mathbb{R}^{n \times n}$ and F, C a partition of $[n]$ such that A_{FF} is nonsingular, the Schur complement of F in A is defined as $\text{Sc}(A, F) = A_{CC} - A_{CF}A_{FF}^{-1}A_{FC}$.

When we need to emphasize the support set of the entries that remain, we also denote $\text{Sc}(A, -C) = \text{Sc}(A, F)$.

2.2 (Directed) Laplacians, Symmetrizations

A matrix $L \in \mathbb{R}^{n \times n}$ is called a directed Laplacian iff $1^\top L = 0$ and all off-diagonal entries of L are non-positive, i.e., $L_{ii} = -\sum_{j:j \neq i} L_{ji}$ for all $i \in [n]$ and $L_{ij} \leq 0$ for all $i \neq j$. A (directed) Laplacian L can be associated with a (directed) graph $\mathcal{G}[L]$ whose adjacency matrix is $\hat{A} = \text{Diag}(L) - L^\top$. The in-degrees/out-degrees of L are defined as the in-degrees/out-degrees of $\mathcal{G}[L]$. For directed Laplacians, its out-degrees equal its diagonal entries. If $\mathcal{G}[L]$ is strongly connected, we say the (directed) Laplacian L is strongly connected.

In addition, if $L\mathbf{1} = \mathbf{0}^\top$, we call L an Eulerian Laplacian. These Laplacians have the property that in-degrees of vertices equal to out-degrees. The undirected Laplacian is a special case where $L = L^\top$. We often refer to these as symmetric Laplacians, or just Laplacians.

Symmetrization: For square matrix $A \in \mathbb{R}^{n \times n}$, we define its matrix symmetrization as $U[A] = \frac{A+A^\top}{2}$. For a directed Laplacian $L \in \mathbb{R}^{n \times n}$, we define its undirectification as

$$\mathcal{U}^G[L] = \frac{1}{2} (L + L^\top - \text{Diag}((L + L^\top)\mathbf{1})).$$

$\mathcal{U}^G[L]$ is called the undirectification because it is a symmetric Laplacian whose adjacency matrix is $U[\hat{A}]$, where $\hat{A} = \text{Diag}(L) - L^\top$ is the adjacency matrix of $\mathcal{G}[L]$. For an Eulerian Laplacian L , its matrix symmetrization coincides with its undirectification, i.e., $U[L] = \mathcal{U}^G[L]$. Eulerian Laplacians are critically important in solvers for directed Laplacians because they are the only setting in which the undirectification is positive semidefinite.

Row Column Diagonal Dominant (RCDD): A square matrix $A \in \mathbb{R}^{n \times n}$ is α -RCDD iff $\sum_{j \in [n] \setminus \{i\}} |A_{ij}| \leq \frac{1}{1+\alpha} A_{ii}$ and $\sum_{j \in [n] \setminus \{i\}} |A_{ji}| \leq \frac{1}{1+\alpha} A_{ii}$ for any $i \in [n]$. We also say A is RCDD if A is 0-RCDD.

2.3 Sparsification

All almost-linear time or faster solvers for directed Laplacians to date are built around sparsification: the approximation of graphs by ones with fewer edges. As it's difficult to even approximate reachability of directed graphs, [7] introduced the key idea of measuring approximations w.r.t. a symmetric PSD matrix. Such approximations are at the core of all subsequent algorithms, including ours.

Definition 2.1. (Asymmetrically bounded) Given a matrix $A \in \mathbb{R}^{n \times n}$ and a PSD matrix $U \in \mathbb{R}^{n \times n}$, A is *asymmetrically bounded*

by U iff $\ker(U) \subseteq \ker(A^\top) \cap \ker(A)$ and $\left\|U^{\dagger/2}AU^{\dagger/2}\right\|_2 \leq 1$. We denote it by $A \stackrel{\text{asym}}{\leq} U$.

By our definition, $A \stackrel{\text{asym}}{\leq} U$ is equivalent to $-A \stackrel{\text{asym}}{\leq} U$. The following lemma is changed slightly from Lemma B.2 of [7].

Fact 2.2. For any matrix $A \in \mathbb{R}^{n \times n}$ and PSD matrix $U \in \mathbb{R}^{n \times n}$, the following statements are equivalent:

- $A \stackrel{\text{asym}}{\leq} U$.
- $2x^\top Ay \leq x^\top Ux + y^\top Uy, \forall x, y \in \mathbb{R}^n$.

Definition 2.3. (Approximation of directed Laplacians via undirectification) Given matrix $A \in \mathbb{R}^{n \times n}$ and directed Laplacian $B \in \mathbb{R}^{n \times n}$, A is an ϵ -asymmetric approximation of B iff $A - B$ is asymmetrically bounded by $\epsilon \cdot \mathcal{U}^G[B]$.

In particular, for strongly connected Eulerian Laplacians A and B , A is an ϵ -asymmetric approximation of B iff

$$\left\|U[B]^{\dagger/2}(A - B)U[B]^{\dagger/2}\right\|_2 \leq \epsilon.$$

We will utilize sparsifiers for Eulerian Laplacians [5, 7], as well as implicit sparsifiers for products of directed adjacency matrices as black boxes throughout our presentations. The formal statements of these black boxes are below.

Theorem 2.4. (Directed Laplacian sparsification oracle) Given a directed Laplacian $L \in \mathbb{R}^{n \times n}$ with $\text{nnz}(L) = m$ and error parameter $\delta \in (0, 1)$, there is an oracle `ORASPARSELAPLACIAN` which runs in at most $\mathcal{T}_S(m, n, \delta)$ time, where $\mathcal{T}_S(m, n, \delta) = O((m \log^{O(1)} n + n \log^{O(1)} n)\delta^{-O(1)})$, to return with high probability a directed Laplacian \tilde{L} satisfying:

- (1) $\text{nnz}(\tilde{L}) \leq N_S(n, \delta)$ where $N_S(n, \delta) = O(n \log^{O(1)} n \delta^{-O(1)})$;
- (2) $\text{Diag}(\tilde{L}) = \text{Diag}(L)$;
- (3) $\tilde{L} - L \stackrel{\text{asym}}{\leq} \delta \cdot \mathcal{U}^G[L]$.

Remark 2.5. Conditions (2), (3) in Theorem 2.4 above are equivalent to \tilde{L} and L having the same in- and out-degrees, and

$$\left\|\mathcal{U}^G[L]^{\dagger/2}(\tilde{L} - L)\mathcal{U}^G[L]^{\dagger/2}\right\|_2 \leq \delta$$

respectively. By having the same in-degrees and out-degrees, we mean $\text{Diag}(\tilde{L}) = \text{Diag}(L)$ and $\tilde{L}\mathbf{1} = L\mathbf{1}$.

Lemma 2.6. (Lemma 3.18 of [7]) Let $x, y \in \mathbb{R}^n$ be nonnegative vectors with $\text{nnz}(x) + \text{nnz}(y) = m$ and let $\epsilon, p \in (0, 1)$. And we denote $G = (1^\top x) \text{Diag}(y) - xy^\top$. Then, there is a routine `SPARSEPRODUCT` which computes with probability at least $1 - p$ a nonnegative matrix A in $O\left(m\epsilon^{-2} \log \frac{m}{p}\right)$ time such that $\text{nnz}(A) = O\left(m\epsilon^{-2} \log \frac{m}{p}\right)$, $A - xy^\top \stackrel{\text{asym}}{\leq} \epsilon \cdot \mathcal{U}^G[G]$.

Given an Eulerian Laplacian $L \in \mathbb{R}^{n \times n}$ and a partition F, C of $[n]$, by invoking `ORASPARSELAPLACIAN` on subgraphs with edges inside (F, F) , (F, C) , (C, F) , (C, C) respectively, we can get a Laplacian sparsification procedure `SPARSEEULERIANFC` so that the sparsified Eulerian Laplacians returned by `SPARSEEULERIANFC` not only satisfy all the properties mentioned in Theorem 2.4, but also keep the in-degrees and out-degrees of the subgraph supported by (F, F) .

Analogously, a routine SPARSEPRODUCTFC can be constructed by applying SPARSEPRODUCT four time.

Explicit definitions of SPARSEEULERIANFC and SPARSEPRODUCTFC can be found in the full version.

2.4 Sufficiency of Solving Eulerian Systems to Constant Error

Previous works on solvers for directed Laplacians and their generalizations (to RCDD and M-matrices) established that it's sufficient to solve Eulerian systems to constant relative accuracy in their undirectification.

- The iterative refinement procedure shown in [7] shows that a constant accuracy solver can be amplified to one with ϵ relative accuracy in $O(\log(1/\epsilon))$ iterations.
- The stationary computation procedure in [8] showed that arbitrary strongly connected Laplacians with mixing time T_{mix} can be solved to 2-norm error ϵ by solving $O(\log(T_{mix}/\epsilon))$ systems in Eulerian Laplacians. This was subsequently simplified and extended to M-matrices and row-column-diagonally-dominant matrices in [1] (with an extra $\log n$ factor in running time). A purely random walk (instead of matrix perturbation) based outer loop is also given in the thesis of Peebles [14].

3 OVERVIEW

Our algorithm is based on sparse Gaussian elimination. Before we discuss the block version, it is useful to first describe how the single variable version works on Eulerian Laplacians.

Recall that Eulerian Laplacians store the (weighted) degrees on the diagonal, and the negation of the out edge weights from j in column j .

Suppose we eliminate vertex j . Then we need to add a rescaled version of row j to each row i where $L_{j,i}$ is non-zero. Accounting for $L_{jj} = d_j$, this weight for row i is given by $\frac{w_{j \rightarrow i}}{d_j}$, and the corresponding decrease in entry j, k is then

$$\frac{w_{j \rightarrow i} w_{k \rightarrow j}}{d_j}.$$

In other words, when eliminating vertex j , we add an edge $k \rightarrow i$ for each triple of vertices $k \rightarrow j \rightarrow i$, with weight given by above.

The effect of this elimination on the vector b can also be described through this ‘distribution’ of row j onto its out-neighbors. However, to start with, it's useful to focus on what happens to the matrix. The key observation for elimination based Laplacian solvers is that this new matrix remains a graph. In fact, it can be checked that this process exactly preserves the in and out degrees of all neighbors of i , so the graph also remains Eulerian.

However, without additional assumptions on the non-zero structures such as separators, directly performing the above process takes $O(n^3)$ time: the newly added entries quickly increases the density of the matrix until each row has $\Theta(n)$ entries. So the starting point of elimination based Laplacian solvers is to address the following two problems:

- (1) Keeping the results of elimination sparse.
- (2) Find vertices that are easy to eliminate.

3.1 Block Cholesky

One possible solution to the issue above is to directly sample the edges formed after eliminating each vertex. It leads to sparsified/incomplete Cholesky factorization based algorithms [6, 11], including the first nearly-linear time solver for directed Laplacians.

Our algorithm is based on eliminating blocks of vertices, and is closest to the algorithm from [10]. It aims to eliminate a block of $\Omega(n)$ vertices simultaneously. This subset, which we denote using F , is chosen to be almost independent. That is, F is picked so that each vertex in F has at least a constant portion of its out-degree going to $V \setminus F$, which we denote as C .

This property means that any random walk on F exits it with constant probability. This intuition, when viewed from iterative methods perspective, implies that power method allows rapid simulation of elimination onto $C = V \setminus F$. From a matrix perspective, it means these matrices are well-approximated by their diagonal. So subproblem $L_{FF}^{-1}b$ can be solved to high accuracy in $O(\log n)$ iterations via power method. We formalize the guarantees of such procedures, PR(\cdot) in Lemma 6.5 in Section 6.2.

Algorithm 1: PRECONDITION $\left(\left\{\tilde{S}^{(i)}, F_i\right\}_{i=1}^d, \mathbf{x}, N\right)$

Input: $(\alpha, \beta, \{\delta_i\}_{i=1}^d)$ -Schur complement chain $\left\{\left\{\tilde{S}^{(i)}\right\}_{i=1}^d, \{F_i\}_{i=1}^d\right\}$; vector $\mathbf{x} \in \mathbb{R}^n$; error parameter $\epsilon \in (0, 1)$

Output: vector $\mathbf{x} \in \mathbb{R}^n$

- 1 **for** $i = 1, \dots, d-1$ **do**
- 2 $\mathbf{x}_{F_i} \leftarrow \text{PR}\left(\tilde{S}_{F_i F_i}^{(i)}, \mathbf{x}_{F_i}, \text{Diag}\left(\tilde{S}^{(i)}\right)_{F_i F_i}^{-1}, \frac{1}{2}, N\right)$;
- 3 $\mathbf{x}_{C_i} \leftarrow \mathbf{x}_{C_i} - \tilde{S}_{C_i F_i}^{(i)} \mathbf{x}_{F_i}$
- 4 **end**
- 5 $\mathbf{x}_{F_d} \leftarrow \left(\tilde{S}^{(d)}\right)^\dagger \mathbf{x}_{F_d}$;
- 6 **for** $i = d-1, \dots, 1$ **do**
- 7 $\mathbf{x}_{F_i} \leftarrow \mathbf{x}_{F_i} - \text{PR}\left(\tilde{S}_{F_i F_i}^{(i)}, \tilde{S}_{F_i C_i}^{(i)} \mathbf{x}_{C_i}, \text{Diag}\left(\tilde{S}^{(i)}\right)_{F_i F_i}^{-1}, \frac{1}{2}, N\right)$
- 8 **end**
- 9 Let $\mathbf{x} \leftarrow \mathbf{x} - \frac{1^\top \mathbf{x}}{n} \cdot \mathbf{1}$;
- 10 Return \mathbf{x}

Compared to single-vertex elimination schemes, block elimination has the advantage of having less error accumulation. Single elimination can be viewed as eliminating 1 vertex per step, while we will show that the block method eliminates $\Omega(n)$ vertices in $O(\log \log n)$ steps. This smaller number of dependencies in turn provides us the ability to bound errors more directly.

Formally, given the partition $F, C \subseteq [n]$ with the permutation matrix P such that $PLP^\top = \begin{bmatrix} L_{FF} & L_{FC} \\ L_{CF} & L_{CC} \end{bmatrix}$, the block Cholesky factorization of $L \in \mathbb{R}^{n \times n}$ is given as

$$L = P^\top \begin{bmatrix} I & 0 \\ L_{CF} L_{FF}^{-1} & I \end{bmatrix} \cdot \begin{bmatrix} L_{FF} & 0 \\ 0 & \text{Sc}(L, F) \end{bmatrix} \cdot \begin{bmatrix} I & L_{FF}^{-1} L_{FC} \\ 0 & I \end{bmatrix} P.$$

Using the above factorization iteratively (with sparsification) generates a Schur complement chain $\{\tilde{S}^{(i)}, F_i\}_{i=1}^d$ (Definition 6.1), the solver algorithm loops through these and solves the subproblems via the projection / prolongation maps which are defined via the random walk on F_i respectively. Its pseudocode is given in Algorithm 1 for completeness.

We remark that the iteration numbers of the preconditioned Richardson iterations in Algorithm 1 can differ with each other. Here, we set a uniform N merely for simplicity. If we have unlimited (or quadratic) precomputation power, the method described above suffices to give us a fast solver. However, due to the exact Schur Complements being dense, the major remaining difficulty is to efficiently compute an approximate Schur complement.

3.2 Schur Complement Sparsification via Partial Block Elimination

The main bottleneck toward an efficient algorithm is the fast construction of approximate Schur complements. We will give such an algorithm whose running time is close to those of sparsification primitives via a partial block elimination process.

In simple terms, a step of this process squares the (F, F) block. Repeating this gives quadratic convergence. With α about $O(1)$, $O(\log \log n)$ iterations suffice, so the resulting error is easier to control than martingales.

LEMMA 3.1 ([16]). *For any diagonal matrix $D \in \mathbb{R}^{n \times n}$ and a matrix $A \in \mathbb{R}^{n \times n}$ with $D - A$ nonsingular, we have*

$$\begin{aligned} & (D - A)^{-1} \\ &= \frac{1}{2} D^{-1} + \frac{1}{2} (I + D^{-1} A) (D - A D^{-1} A)^{-1} (I + A D^{-1}). \end{aligned} \quad (1)$$

The main identity (1) gives rise to our definition for partial-block-eliminated Laplacian of L . Let $D = \text{Diag}(L)$ and $A = D - L$, let P be the permutation matrix such that $PLP^\top = \begin{bmatrix} L_{FF} & L_{FC} \\ L_{CF} & L_{CC} \end{bmatrix}$. The partial-block-eliminated operator $\Phi(L|D_{FF}, F)$ is defined as

$$\Phi(L|D_{FF}, F) = P^\top \begin{bmatrix} L_{FF} & -A_{FC} \\ -A_{CF} & 2L_{CC} \end{bmatrix} P - A_{:,F} D_{FF}^{-1} A_{F,:}.$$

We define the first exact partial-block-elimination of L by $L^{(1)} = \Phi(L|D_{FF}, F)$. Then, $\frac{1}{2}L^{(1)}$ is an Eulerian Laplacian which has the same Schur complement of F in L , i.e.,

$$\text{Sc}(L, F) = \frac{1}{2} \text{Sc}\left(L^{(1)}, F\right).$$

The 2-nd to the K -th partially-block-eliminated Laplacians are defining iteratively by $L^{(k+1)} = \Phi(L^{(k)}|D_{FF}, F)$. $L^{(k)}$ can also be regarded as a partially powered matrix of L , which uses the powering to obtain better spectral properties. Specifically, when we focus on the (F, F) block of $L^{(k)}$, it is easy to see that $\|D_{FF}^{-1} A_{FF}^{(k)}\|_\infty$ converges at a quadratic rate, where $A_{FF}^{(k)} = D_{FF} - L_{FF}^{(k)}$. Formal construction of the k -th partially-block-eliminated Laplacians and their properties can be found in the appendix of the full version.

To encounter the increasing density of $L^{(k)}$, sparsification black-boxes in Section 2.3 are naturally accompanied with the partial block elimination to yield a Schur complement sparsification method

(Algorithm 2). Our algorithm is essentially a directed variant of the one from [10]. A slight difference is that in the last step, we need to fix the degree discrepancies caused by approximating the strongly RCDD matrix by its diagonal.

The running time of Algorithm 2 is shown in Theorem 3.3. The k -th iterand $\tilde{L}^{(k)}$ of Algorithm 2 is termed the approximate k -th partially-block-eliminated Laplacian, while its exact version is just the (exact) k -th partially-block-eliminated Laplacian $L^{(k)}$ defined above. To show the running time of Algorithm 2, the most important thing is to provide relatively tight bounds for the difference $\tilde{L}^{(k)} - L^{(k)}$.

Algorithm 2: SPARSE SCHUR (L, F, δ)

Input: strongly connected Eulerian Laplacian $L \in \mathbb{R}^{n \times n}$;
partition F, C of $[n]$; error parameter $\delta \in (0, 1)$
Output: Sparse approximate Schur complement S

1 If $\text{nnz}(L) \geq O(N_S(n, \delta))$, call ORASPARSELAPLACIAN to
sparsify L with error parameter $O(\delta)$;
2 Find a permutation matrix P such that

$$PLP^\top = \begin{bmatrix} L_{FF} & L_{FC} \\ L_{CF} & L_{CC} \end{bmatrix};$$

3 Set $K \leftarrow O(\log \log \frac{n}{\delta})$, $\epsilon \leftarrow O\left(\frac{\delta}{K}\right)$, $\tilde{L}^{(0)} \leftarrow L$,
 $D \leftarrow \text{Diag}(\tilde{L}^{(0)})$, $\tilde{A}^{(0)} \leftarrow D - \tilde{L}^{(0)}$;

4 **for** $k = 1, \dots, K$ **do**

5 **for** $i \in F$ **do**

6 $\tilde{Y}^{(k,i)} \leftarrow$
 $\quad \text{SPARSEPRODUCTFC}\left(\tilde{A}_{:,i}^{(k-1)}, \left(\tilde{A}_{i,:}^{(k-1)}\right)^\top, \epsilon, F\right)$

7 **end**

8 Let $\tilde{Y}^{(k)} \leftarrow \sum_{i \in F} \frac{1}{D_{ii}} \tilde{Y}^{(k,i)}$,
 $\tilde{L}^{(k,0)} \leftarrow P^\top \begin{bmatrix} D_{FF} & -\tilde{A}_{FC}^{(k-1)} \\ -\tilde{A}_{CF}^{(k-1)} & 2\tilde{L}_{CC}^{(k-1)} \end{bmatrix} P - \tilde{Y}^{(k)}$;

9 $\tilde{L}^{(k)} \leftarrow \text{SPARSEEULERIANFC}\left(\tilde{L}^{(k,0)}, \epsilon, F\right)$ and
 $\tilde{A}^{(k)} \leftarrow P^\top \begin{bmatrix} D_{FF} & \text{Diag}(\tilde{L}_{CC}^{(k)}) \\ & \end{bmatrix} P - \tilde{L}^{(k)}$

10 **end**

11 **for** $i \in F$ **do**

12 $\tilde{X}^{(i)} \leftarrow \text{SPARSEPRODUCT}\left(\tilde{A}_{C,i}^{(K)}, \left(\tilde{A}_{i,C}^{(K)}\right)^\top, \epsilon\right)$

13 **end**

14 Let $\tilde{X} \leftarrow \sum_{i \in F} \frac{1}{D_{ii}} \tilde{X}^{(i)}$ and $\tilde{S}^{(0)} \leftarrow \frac{1}{2K} (\tilde{L}_{CC}^{(K)} - \tilde{X})$;

15 Compute a patching matrix $R \in \mathbb{R}^{|C| \times |C|}$ with

$$R_{2:|C|,1} = -\tilde{S}_{2:|C|,:}^{(0)} \mathbf{1}, R_{1,2:|C|} = -\mathbf{1}^\top \tilde{S}_{:,2:|C|}^{(0)},$$

$$R_{1,1} = -R_{1,2:|C|} \mathbf{1} - \mathbf{1}^\top R_{2:|C|,1} - \mathbf{1}^\top \tilde{S}^{(0)} \mathbf{1}, \text{ and } R_{ij} = 0 \text{ for } i \neq 1 \text{ and } j \neq 1 ;$$

16 Set $\hat{S} = \tilde{S}^{(0)} + R$;

17 Return $S = \text{ORASPARSELAPLACIAN}\left(\hat{S}, \delta/8\right)$

Remark 3.2. This permutation matrix P defined on line 2 of Algorithm 2 is only used to simplify the pseudocodes. We use the same D_{FF} in each iteration to simplify our analysis. It is possible to replace D_{FF} by $\text{Diag}(\tilde{L}^{(k-1)})_{FF}$ in iterations k and achieve similar running time.

THEOREM 3.3. (*Schur complement sparsification*) For a strongly connected Eulerian Laplacian $L \in \mathbb{R}^{n \times n}$, let F, C be a partition of $[n]$ such that L_{FF} is α -RCDD ($\alpha = O(1)$) and let $\delta \in (0, 1)$ be an error parameter, the subroutine SPARSE SCHUR (Algorithm 2) runs in time

$$O(\mathcal{T}_S(m, n, \delta)) + \tilde{O}\left(\mathcal{N}_S(n, \delta) \delta^{-2}, n, \delta\right) \log n$$

to return with high probability a strongly connected Eulerian Laplacian S satisfying $\text{nnz}(S) = O(\mathcal{N}_S(|C|, \delta))$ and

$$S - SC(L, F) \stackrel{\text{asym}}{\leq} \delta \cdot U[SC(L, F)].$$

Compared to the undirected analog from [10], powered directed matrices exhibit significantly more complicated spectral structures. To analyze them, we develop new interpretations of directed Schur complements based on matrix extensions.

3.3 Bounding Error Accumulations in Partially-Eliminated Laplacians by Augmented Matrices

When considering the approximate partial block elimination, in one update step, not only new sparsification errors are added into $\tilde{L}^{(k)}$, the errors accumulated from previous steps will multiply with each other and get possibly amplified. In addition, error accumulations in Schur complements of directed Laplacians are not as straightforward as their undirected counterparts. It's not the case that for two directed Eulerian Laplacians with the same undirectification, the undirectification of their Schur complements are the same. For instance, consider the undirected vs. directed cycle, eliminated till only two originally diametrically opposite vertices remain. The former has a Schur complement that has weight $2/n$, while the latter has a Schur complement that has weight 1.

By the definition of ϵ -asymmetric approximation, we need to essentially show the following inequality in order to obtain the approximations needed for a nearly-linear time algorithm:

$$\begin{aligned} \frac{1}{2^k} U[L^{(k)}] &= \frac{1}{2^k} U[\Phi(L^{(k-1)} | D_{FF}, F)] \\ &\leq O(1) \cdot U[L], \quad \forall 1 \leq k \leq K. \end{aligned}$$

Here significant difficulties arise due to the already complicated formula of $L^{(k)}$. So we instead express the exact and approximate partial block elimination as Schur complements of large augmented matrices introduced below.

In the remainder of Section 3 and the entire Section 4 and Section 5, unless otherwise specified, we assume $C = \{1, 2, \dots, |C|\}$ for simplicity. We define

$$F_a = \{b \in \mathbb{Z} : |C| + (a-1)|F| + 1 \leq b \leq |C| + a|F|\}.$$

Note that in our notation, $F = F_1$.

3.3.1 A Reformulation for Partial Block Elimination. For the exact and approximate k -th partially-block-eliminated matrices $L^{(k)}$, $\tilde{L}^{(k)}$, we define augmented matrices $\mathbf{M}^{(0,k)}, \tilde{\mathbf{M}}^{(0,k)}$ of size $2^k|F| + |C|$. We start with the construction of a desirable $\mathbf{M}^{(0,k)}$. To this end, we define a sequence of augmented matrices $\{\mathbf{M}^{(i,k)}\}_{i=0}^k$, where $\mathbf{M}^{(k,k)} = L^{(k)}$ and each $\mathbf{M}^{(i,k)}$ is a Schur complement of $\mathbf{M}^{(i-1,k)}$. Here we only give an informal explanation of how we construct $\mathbf{M}^{(i,k)}$. The formal definitions of these augmented matrices are given in Section 4. To begin with, for some fixed $k \in [K]$, we define $\mathbf{M}^{(k,k)} \stackrel{\text{def}}{=} L^{(k)}$.

Next, we take $\mathbf{M}^{(k-1,k)}$ and $\mathbf{M}^{(k-2,k)}$ as examples to show how we define such a sequence of matrices $\mathbf{M}^{(k-1,k)}, \dots, \mathbf{M}^{(0,k)}$.

Define $\mathbf{M}^{(k-1,k)}$ as follows

$$\mathbf{M}^{(k-1,k)} \stackrel{\text{def}}{=} \begin{bmatrix} 2L_{CC}^{(k-1)} & -A_{CF}^{(k-1)} & -A_{CF}^{(k-1)} \\ -A_{FC}^{(k-1)} & D_{FF} & -A_{FF}^{(k-1)} \\ -A_{FC}^{(k-1)} & -A_{FF}^{(k-1)} & D_{FF} \end{bmatrix}$$

Then, it follows by direct calculations that

$$\text{Sc}(\mathbf{M}^{(k-1,k)}, F_2) = L^{(k)}.$$

Next, we define $\mathbf{M}^{(k-2,k)}$ as follows

$$\mathbf{M}^{(k-2,k)} \stackrel{\text{def}}{=} \begin{bmatrix} 4L_{CC}^{(k-2)} & -A_{CF}^{(k-2)} & -A_{CF}^{(k-2)} & -A_{CF}^{(k-2)} & -A_{CF}^{(k-2)} \\ -A_{FC}^{(k-2)} & D_{FF} & & & \\ -A_{FC}^{(k-2)} & & D_{FF} & & -A_{FF}^{(k-2)} \\ -A_{FC}^{(k-2)} & & -A_{FF}^{(k-2)} & D_{FF} & \\ -A_{FC}^{(k-2)} & -A_{FF}^{(k-2)} & & & D_{FF} \end{bmatrix}$$

It follows by direct calculations that

$$\text{Sc}(\mathbf{M}^{(k-2,k)}, F_3 \cup F_4) = \mathbf{M}^{(k-1,k)}.$$

We will show $\frac{1}{2^k} U[L^{(k)}] \leq O(1) \cdot U[L]$ later by analyzing the properties of $\mathbf{M}^{(0,k)}$.

We believe this representation may be of independent interest. We also remark that these augmented matrices only arise during analysis, and are not used in the algorithms.

3.3.2 Bounding Error Accumulation in Algorithm 2. Next, we will mainly use Lemma 5.2 to bound the errors after taking Schur complements. However, in our analysis, iteratively applying Lemma 5.2 to bound $\frac{1}{2^k} (\tilde{L}^{(k)} - L^{(k)})$ will lead to more $\log n$ factors in the running time. To derive a tighter bound, we introduce another group of augmented matrices $\{\tilde{\mathbf{M}}^{(0,k)}\}$ which are defined by attaching sparsification errors to $\mathbf{M}^{(0,k)}$. $\tilde{\mathbf{M}}^{(0,k)}$ can help us disentangle the sparsification errors generated from different iterations and see how these errors accumulate as we do partial block eliminations more clearly. We use another group of augmented matrices $\{\mathbf{Q}^{(k)}\}$ to bound the difference between $\mathbf{M}^{(0,k)}$ and $\tilde{\mathbf{M}}^{(0,k)}$. The augmented matrix $\mathbf{Q}^{(k)}$ is defined as the sum of a group of “repetition matrix” (Section 4.2). $\mathbf{Q}^{(k)}$ adopts many properties similar to $\mathbf{M}^{(0,k)}$, so it is easy to analyze. Then, we can give tighter bound for $\tilde{L}^{(K)} - L^{(K)}$

using the robustness of Schur complements in this case with the properties of $Q^{(k)}$ (Section 5).

4 PARTIAL BLOCK ELIMINATION VIA AUGMENTED MATRICES

In this section, we introduce our augmented matrices based view of partial block elimination. As we will show later, after $O(\log \log n)$ steps of partial elimination, the (F, F) block of the approximate partially-block-eliminated Laplacian $\tilde{L}^{(k)}$ can be approximated by its diagonal “safely”. So, what remains is to bound the error accumulations in the difference $\frac{1}{2^k} \left(\text{Sc}(\tilde{L}^{(k)}, F) - \text{Sc}(L^{(k)}, F) \right)$, which we do by bounding differences in $\frac{1}{2^k} (L^{(k)} - \tilde{L}^{(k)})$.

4.1 A Reformulation of the Exact Partial Block Elimination

In this section, we provide a reformulation of the exact version of partial block elimination which is more friendly to error analysis. To be specific, our strategy is to construct a large matrix $\mathbf{M}^{(0,k)} \in \mathbb{R}^{(2^k|F|+|C|) \times (2^k|F|+|C|)}$ such that $L^{(k)}$ is a Schur complement of the large matrix $\mathbf{M}^{(0,k)}$. And there is a partition of $\mathbf{M}^{(0,k)}$ such that each block is a zero matrix or equals some submatrix of L . To construct $\mathbf{M}^{(0,k)}$, we will construct a sequence of augmented matrices $\{\mathbf{M}^{(i,k)}\}_{i=0}^k$ satisfying Lemma 4.1. Later, by analyzing the large matrix $\mathbf{M}^{(0,k)}$, we can derive tighter bounds for quantities related to $L^{(k)}$.

Now, we give a rigorous way to construct such a sequence of matrices $\{\mathbf{M}^{(i,k)}\}_{i=0}^k$ ($0 \leq i \leq k$).

We construct a sequence of bijections $\{\psi^{(i)}(\cdot)\}_{i=0}^K$ which indicate the “positions” of the blocks equalling $-A_{FF}^{(i)}$ in the large augmented matrix $\mathbf{M}^{(k-i,k)}$.

We start with $\psi^{(0)}(\cdot)$ and will define these $\psi^{(i)}$ iteratively. The mapping $\psi^{(0)}(\cdot)$ is defined as a trivial mapping from $\{1\}$ to $\{1\}$ with

$$\psi^{(0)}(1) = 1.$$

Then, assume we have defined $\psi^{(i-1)}(\cdot)$. Now, we define $\psi^{(i)}$ as follows:

$$\psi^{(i)}(a) = \begin{cases} a + 2^{i-1}, & a \in [2^{i-1}] \\ \psi^{(i-1)}(a - 2^{i-1}), & 2^{i-1} + 1 \leq a \leq 2^i \end{cases}$$

If $\psi^{(i-1)}(\cdot)$ is a bijection from $[2^{i-1}]$ to $[2^{i-1}]$, then $\psi^{(i)}(\cdot)|_{\{2^{i-1}+1, \dots, 2^i\}}$ is a bijection from $\{2^{i-1} + 1, \dots, 2^i\}$ to $[2^{i-1}]$. And by the definition, $\psi^{(i)}(\cdot)|_{[2^{i-1}]}$ is a bijection from $[2^{i-1}]$ to $\{2^{i-1} + 1, \dots, 2^i\}$. Then, $\psi^{(i)}$ is a bijection from $[2^i]$ to $[2^i]$.

It follows by induction that for any $k \in [K]$, $\psi^{(k)}(\cdot)$ is a bijection from $[2^k]$ to $[2^k]$. And by the fact that $2^{i-1} + 1 \leq \psi^{(i)}(a) \leq 2^i$ for $a \in [2^{i-1}]$ and $\psi^{(i)}(a) \in [2^{i-1}]$ for $2^{i-1} + 1 \leq a \leq 2^i$, we have the

following relation

$$\psi^{(j)}(a) \neq a, \forall 1 \leq j \leq K, a \in [2^j].$$

With the notations defined above, we define the matrix $\mathbf{M}^{(i,k)}$ as

$$\begin{aligned} \mathbf{M}^{(i,k)} = & 2^{k-i} \mathcal{P} \left(L_{CC}^{(i)}, C, C, 2^{k-i} |F| + |C| \right) \\ & + \sum_{a=1}^{2^{k-i}} \left(\mathcal{P} \left(D_{FF}, F_a, F_a, 2^{k-i} |F| + |C| \right) \right. \\ & + \mathcal{P} \left(-A_{FF}^{(i)}, F_a, F_{\psi^{(k-i)}(a)}, 2^{k-i} |F| + |C| \right) \\ & + \mathcal{P} \left(-A_{FC}^{(i)}, F_a, C, 2^{k-i} |F| + |C| \right) \\ & \left. + \mathcal{P} \left(-A_{CF}^{(i)}, C, F_a, 2^{k-i} |F| + |C| \right) \right). \end{aligned}$$

where the notation $\mathcal{P}(X, A, B, n)$ has been defined in Section 2, which means putting matrix X in the submatrix indexed by (A, B) in a zero matrix $0_{n \times n}$; $L^{(k)}$ is the exact k -th partially-block-eliminated Laplacian and formal definitions of $L^{(k)}$, $A^{(k)}$ are in the appendix of the full version.

We have the following properties of $\{\mathbf{M}^{(i,k)}\}$.

LEMMA 4.1. *For any $0 \leq k \leq K$, $0 \leq i \leq k$, $\mathbf{M}^{(i,k)}$ is an Eulerian Laplacian; $\mathbf{M}_{-[n],-[n]}^{(i,k)}$, $\mathbf{M}_{-C,-C}^{(i,k)}$ are α -RCDD; the Schur complement satisfies*

$$\text{Sc} \left(\mathbf{M}^{(i,k)}, \cup_{a=2^{k-i-1}+1}^{2^{k-i}} F_a \right) = \mathbf{M}^{(i+1,k)}.$$

Further,

$$\text{Sc} \left(\mathbf{M}^{(i,k)}, -[n] \right) = L^{(k)}.$$

In addition, for any $\mathbf{x} \in \mathbb{R}^n$, let

$$\widehat{\mathbf{x}} = (\underbrace{\mathbf{x}_C^\top \quad \mathbf{x}_F^\top \cdots \mathbf{x}_F^\top}_{2^k \text{ repetitions of } \mathbf{x}_F^\top})^\top,$$

then,

$$\widehat{\mathbf{x}}^\top \mathbf{M}^{(0,k)} \widehat{\mathbf{x}} = 2^k \mathbf{x}^\top L \mathbf{x}.$$

The following lemma answers a question in Section 3.3. That is, $\frac{1}{2^k} \mathbf{U} \left[L^{(k)} \right] \leq O(1) \mathbf{U} [L]$.

LEMMA 4.2. *For any $0 \leq k \leq K$,*

$$\frac{1}{2^k} \mathbf{U} \left[L^{(k)} \right] \leq \left(3 + \frac{2}{\alpha} \right) \mathbf{U} [L].$$

4.2 Bounding Error Accumulation Using Repetition Matrices

In this section, we define $\widetilde{\mathbf{M}}^{(0,k)}$ as an inexact version of $\mathbf{M}^{(0,k)}$ where sparsification errors accumulate. Then, we introduce a special type of augmented matrices, which we term *repetition matrices*, and bound the difference $\widetilde{\mathbf{M}}^{(0,k)} - \mathbf{M}^{(0,k)}$ in norms based on the matrices $\{Q^{(k)}\}$ which are defined as linear combinations of some *repetition matrices* of $\{\mathbf{M}^{(0,k)}\}$.

Before we define $\tilde{\mathbf{M}}^{(0,k)}$, we introduce our notations for the errors induced by sparsification:

$$\mathbf{E}^{(k)} = \tilde{\mathbf{L}}^{(k)} - \Phi(\tilde{\mathbf{L}}^{(k-1)} | D_{FF}, F),$$

$$E_X = \tilde{\mathbf{A}}_{CF}^{(K)} D_{FF}^{-1} \tilde{\mathbf{A}}_{FC}^{(K)} - \tilde{\mathbf{X}}.$$

We also denote in the rest of this paper,

$$\begin{aligned} \tilde{\mathbf{R}} = & \mathbf{R} + \frac{1}{2^K} \tilde{\mathbf{A}}_{CF}^{(K)} \left(D_{FF} - \tilde{\mathbf{A}}_{FF}^{(K)} \right)^{-1} \\ & - \frac{1}{2^K} \tilde{\mathbf{A}}_{CF}^{(K)} D_{FF}^{-1} \tilde{\mathbf{A}}_{FC}^{(K)}. \end{aligned}$$

Some elementary facts of the results of Algorithm 2 is given by the following lemma.

LEMMA 4.3. *With high probability, the following statements hold:*

- (1) $\{\tilde{\mathbf{L}}^{(k)}\}_{k=0}^K$ are Eulerian Laplacians;
- (2) $\{\tilde{\mathbf{A}}^{(k)}\}_{k=0}^K$ are nonnegative matrices satisfying

$$\left\| D_{FF}^{-1} \tilde{\mathbf{A}}_{FF}^{(k)} \right\|_{\infty} \leq \left(\frac{1}{1+\alpha} \right)^{2^k}, \quad \forall 0 \leq k \leq K;$$

(3) S, \tilde{S} are Eulerian Laplacians;

(4) The matrix $\tilde{\mathbf{R}}$ satisfies $\tilde{\mathbf{R}}\mathbf{1} = \tilde{\mathbf{R}}^T\mathbf{1} = \mathbf{0}$ and

$$\left\| \tilde{\mathbf{R}} \right\|_2 \leq \frac{n^2 \|D_{FF}\|_2}{2^{K-1}\alpha} \left(\frac{1}{1+\alpha} \right)^{2^K}.$$

Then, we can provide bounds for the one-step errors in the next lemma.

LEMMA 4.4. *The error matrices satisfies*

$$\mathbf{E}^{(k)} \stackrel{\text{asym}}{\leq} \epsilon_0 \mathbf{U} \left[\tilde{\mathbf{L}}^{(k-1)} \right], \quad (2)$$

$$E_X \stackrel{\text{asym}}{\leq} \epsilon \mathbf{U} \left[\text{Sc} \left(\tilde{\mathbf{L}}^{(K)}, F \right) \right], \quad (3)$$

where $\epsilon_0 = 2 \left(3 + \frac{2}{\alpha} \right) (2\epsilon + \epsilon^2)$.

In the remainder of this paper, we write $\epsilon_0 = 2 \left(3 + \frac{2}{\alpha} \right) (2\epsilon + \epsilon^2)$.

Recall that we define an augmented matrix $\mathbf{M}^{(0,k)}$ such that $\mathbf{L}^{(k)}$ is its Schur complement in Section 4.1. Now, we define $\tilde{\mathbf{M}}^{(0,k)}$ which is an inexact version of $\mathbf{M}^{(0,k)}$ to analyze the properties of $\tilde{\mathbf{L}}^{(k)}$. We first define

$$\begin{aligned} \mathcal{R} \left(k, a, E^{(i)} \right) = & \mathcal{P} \left(E_{FF}^{(i)}, F_a, F_{\psi^{(k-i)}(a)}, 2^{k-i} |F| + |C| \right) \\ & + \mathcal{P} \left(E_{FC}^{(i)}, F_a, C, 2^{k-i} |F| + |C| \right) \\ & + \mathcal{P} \left(E_{CF}^{(i)}, C, F_{\psi^{(k-i)}(a)}, 2^{k-i} |F| + |C| \right) \\ & + \mathcal{P} \left(E_{CC}^{(i)}, C, C, 2^{k-i} |F| + |C| \right). \end{aligned}$$

Then, we define the error matrices

$$\mathcal{E}^{(i,k)} = \sum_{a=1}^{2^{k-i}} \mathcal{R} \left(k, a, E^{(i)} \right)$$

and

$$\mathcal{E}^{(1:k,k)} = \sum_{i=1}^k \mathcal{E}^{(i,k)}.$$

The matrix $\tilde{\mathbf{M}}^{(0,k)}$ is defined as follows

$$\tilde{\mathbf{M}}^{(0,k)} = \mathbf{M}^{(0,k)} + \mathcal{E}^{(1:k,k)}.$$

LEMMA 4.5. *The Schur complement of $[2^k |F| + |C|] \setminus [n]$ in $\tilde{\mathbf{M}}^{(0,k)}$ satisfies:*

$$\text{Sc} \left(\tilde{\mathbf{M}}^{(0,k)}, -[n] \right) = \tilde{\mathbf{L}}^{(k)}.$$

To help bound $\mathcal{E}^{(1:k,k)}$, we define some special kinds of matrices termed “repetition matrices”. We will construct the matrices $\{Q^{(k)}\}$ as linear combinations of “repetition matrices”.

Definition 4.6. (“Repetition matrices”) We will use the following 3 kinds of “repetition matrices”: consider a matrix $A \in \mathbb{R}^{m \times m}$ and subset $C \subseteq [m]$ and $E = [m] \setminus C$,

- (1) the k -“repetition matrix” of A is defined as follows:

$$\text{Rep}(k, C, A) = \begin{bmatrix} kA_{CC} & A_{CE} & A_{CE} & \cdots & A_{CE} \\ A_{EC} & A_{EE} & \mathbf{0} & \cdots & \mathbf{0} \\ A_{EC} & \mathbf{0} & A_{EE} & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \mathbf{0} \\ A_{EC} & \mathbf{0} & \cdots & \mathbf{0} & A_{EE} \end{bmatrix},$$

where the repetition numbers of the blocks A_{CE}, A_{EC}, A_{EE} are k ;

- (2) $\text{Rep}^{+0}(k, C, A, N)$ is a larger matrix by appending all-zeros rows and columns to $\text{Rep}(k, C, A)$:

$$\text{Rep}^{+0}(k, C, A, N) = \begin{bmatrix} \text{Rep}(k, C, A) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{N \times N},$$

where $N \geq k |E| + |C|$ is used to indicate the size of the matrix $\text{Rep}^{+0}(k, C, A, N)$;

- (3) if F, F_+ is a partition of E , $\text{Rep}(k, C, F, A)$ is defined as a permutation of the k -“repetition matrix” of A , which has the following form:

$$\text{Rep}(k, C, F, A)$$

$$= \begin{bmatrix} kA_{CC} & A_{CF} & \cdots & A_{CF} & A_{CF^+} & \cdots & A_{CF^+} \\ A_{FC} & A_{FF} & & & & & \\ \vdots & & \ddots & & & & \\ A_{FC} & & & A_{FF} & & & \\ A_{F^+C} & & & & A_{F^+F^+} & & \\ \vdots & & & & & \ddots & \\ A_{F^+C} & & & & & & A_{F^+F^+} \end{bmatrix}.$$

Now, we define the matrices $\{Q^{(k)}\}_{0 \leq k \leq K}$ which are used to bound $\mathbf{M}^{(0,k)} - \tilde{\mathbf{M}}^{(0,k)}$ and then $\mathbf{L}^{(k)} - \tilde{\mathbf{L}}^{(k)}$. We define the matrices $\{Q^{(k)}\}_{0 \leq k \leq K}$ iteratively together with the error quantities $\{\gamma_k\}_{0 \leq k \leq K}$.

We start from $Q^{(0)} = U[L]$ and $\gamma_0 = 0$. If we have defined $\{Q^{(i)}\}_{0 \leq i < k}$, $\{\gamma_i\}_{0 \leq i < k}$, then $Q^{(k)} \in \mathbb{R}^{(2^k|F|+|C|) \times (2^k|F|+|C|)}$ and $\gamma_k \in \mathbb{R}_+$ are defined as follows:

$$\begin{aligned} Q^{(k)} &\stackrel{\text{def}}{=} \frac{k}{4k + \frac{2k}{\alpha} + \sum_{i=0}^{k-1} \gamma_i} U[\mathbf{M}^{(0,k)}] \\ &+ \frac{1}{4k + \frac{2k}{\alpha} + \sum_{i=0}^{k-1} \gamma_i} \sum_{i=0}^{k-1} \gamma_i \text{Rep}\left(2^{k-i}, C, F, Q^{(i)}\right) \\ &+ \frac{3 + \frac{2}{\alpha}}{4k + \frac{2k}{\alpha} + \sum_{i=0}^{k-1} \gamma_i} \sum_{i=0}^{k-1} \text{Rep}\left(2^{k-i}, C, F, U[\mathbf{M}^{(0,i)}]\right), \\ \gamma_k &\stackrel{\text{def}}{=} \left\| \text{Sc}\left(Q^{(k)}, -[n]\right)^{\dagger/2} \left(\tilde{L}^{(k)} - L^{(k)}\right) \text{Sc}\left(Q^{(k)}, -[n]\right)^{\dagger/2} \right\|_2. \end{aligned}$$

The following lemma shows some elementary properties of $Q^{(k)}$.

Fact 4.7. $Q^{(k)}$ is a Laplacian satisfying:

- (1) $U[\mathbf{M}^{(0,k)}] \leq \left(4 + \frac{2}{\alpha} + \frac{\sum_{i=0}^{k-1} \gamma_i}{k}\right) Q^{(k)}$;
- (2) $\text{Diag}(Q^{(k)}) = \text{Diag}(\mathbf{M}^{(0,k)})$;
- (3) $Q_{-C,-C}^{(k)}, Q_{-[n],-[n]}^{(k)}$ are α -RCDD;
- (4) $\left\| \left(Q^{(k)}\right)^{1/2} \text{Diag}(\mathbf{M}^{(0,k)})^{-1/2} \right\|_2^2 \leq 2$;

LEMMA 4.8. $\text{Sc}\left(Q^{(k)}, -C\right) \leq 2^k U[\text{Sc}(L, F)]$.

The following lemma shows that the sparsification errors attached to $\tilde{\mathbf{M}}^{(0,k)}$ can be bounded by $Q^{(k)}$.

LEMMA 4.9.

$$\mathcal{E}^{(1:k,k)} \stackrel{\text{asym}}{\leq} \epsilon_0 \left(4k + \frac{2k}{\alpha} + \sum_{i=0}^{k-1} \gamma_i\right) Q^{(k)}.$$

5 ROBUSTNESS OF SCHUR COMPLEMENTS AND FULL ERROR ANALYSIS

In this section, we show additional robustness properties of Schur complements suitable for analyzing errors on the augmented matrices. Specifically, we establish conditions on A, B, U where $A - B \stackrel{\text{asym}}{\leq} \epsilon \cdot U$, as well as the set to be eliminated, F , so that $\text{Sc}(A, F) - \text{Sc}(B, F) \stackrel{\text{asym}}{\leq} \delta \cdot \text{Sc}(U, F)$.

Using these properties, we bound the norms of errors in Schur complements of the $Q^{(k)}$ and γ_k . Such bounds allow us to complete the proof of Theorem 3.3.

The next lemma is used to prove Lemma 5.2 below.

LEMMA 5.1. Suppose that $L \in \mathbb{R}^{n \times n}$ is an Eulerian Laplacian, $D = \text{Diag}(L)$, W is PSD, $\left\|W^{1/2}D^{-1/2}\right\|_2 \leq a$, and the matrix $E \in \mathbb{R}^{n \times n}$ satisfies $E \stackrel{\text{asym}}{\leq} bW$ with $a^2b < 2$. Then the matrix $M = L + E$ satisfies:

$$MD^{-1}M^\top \leq \frac{1}{2 - a^2b} \left((4 + 2a^2b) U[L] + 2bW \right).$$

The following lemma shows the robustness of the Schur complements. It's used in the proof of Lemma 5.3 to bound γ_k .

LEMMA 5.2. Let $N \in \mathbb{R}^{n \times n}$ be an Eulerian Laplacian, let M be an n -by- n matrix, let $U \in \mathbb{R}^{n \times n}$ be PSD and F, C a partition of $[n]$. Suppose that U_{FF} is nonsingular, $U1 = 0$, N_{FF} is ρ -RCDD ($\rho > 0$), $U[N]_{FF} \geq \frac{1}{\mu} U_{FF}$, $U[N] \leq \beta U$, $\left\|U^{1/2} \text{Diag}(N)^{-1/2}\right\|_2 \leq a$, and the matrix $E = M - N$ satisfies $E \stackrel{\text{asym}}{\leq} b \cdot U$ with $b < \min\left\{\frac{2}{a^2}, \frac{1}{\mu}\right\}$.

Then, M_{FF}, N_{FF} are nonsingular and

$$\begin{aligned} &\text{Sc}(M, F) - \text{Sc}(N, F) \\ &\stackrel{\text{asym}}{\leq} b \left(1 + \frac{1}{\rho}\right) \frac{\mu(\beta(4 + 2a^2b) + 2b)}{(1 - \mu b)^2(2 - a^2b)} \cdot \text{Sc}(U, F). \end{aligned}$$

We can now obtain a relatively tight bound for $\text{Sc}(\tilde{L}^{(K)}, F) - \text{Sc}(L^{(K)}, F)$ by bounding γ_k iteratively.

LEMMA 5.3. For any $\delta_0 \in (0, 1)$, with a small $\epsilon = O\left(\frac{\delta_0}{K}\right)$ in Algorithm 2, the exact and approximate K -th partially-block-eliminated Laplacians $L^{(K)}, \tilde{L}^{(K)}$ satisfies

$$\begin{aligned} &\frac{1}{2^K} \left(\text{Sc}(\tilde{L}^{(K)}, F) - \text{Sc}(L^{(K)}, F) \right) \\ &\stackrel{\text{asym}}{\leq} O(\delta_0) \cdot U[\text{Sc}(L, F)]. \end{aligned} \tag{4}$$

Remark 5.4. Since $\tilde{S}^{(0)} = \frac{1}{2^K} (\tilde{L}_{CC}^{(K)} - \tilde{X})$ (in Algorithm 2), the $\frac{1}{2^K}$ factor on the LHS of (4) doesn't matter. The parameter choice $\epsilon = O\left(\frac{\delta_0}{K}\right)$ implies that when running Algorithm 2, the error accumulates linearly in k rather than exponentially.

6 A NEARLY-LINEAR TIME SOLVER

In this section, we complete the Sparsified Schur Complement based algorithm by invoking the nearly-linear time Schur complement sparsification procedure derived above in Sections 4 and 5. We first call this Schur complement sparsification procedure repeatedly to construct a sparse Schur complement chain in Section 6.1. Then, in Section 6.2, we show that this Schur complement chain gives a preconditioner PRECONDITION for the initial Eulerian Laplacian matrix. The full high accuracy solver then follows from invoking this preconditioner inside Richardson iteration.

6.1 Schur Complement Chains

We first define Schur complement chains over directed graphs, which is a variant of the Schur complement chain for undirected graphs in [10].

Definition 6.1. (Schur complement chain) Given a strongly connected Eulerian Laplacian $L \in \mathbb{R}^{n \times n}$, an $(\alpha, \beta, \{\delta_i\}_{i=1}^d)$ -Schur complement chain of L is a sequence of strongly connected Eulerian Laplacians and subsets $\left\{\tilde{S}^{(i)}\right\}_{i=1}^d, \{F_i\}_{i=1}^d$ satisfying

- (1) $\{F_i\}_{i=1}^d$ is a partition of $[n]$; each $\tilde{S}^{(i)}$ is supported on the submatrix indexed by (C_{i-1}, C_{i-1}) , where $C_i \stackrel{\text{def}}{=} [n] \setminus \left(\bigcup_{j=1}^i F_j\right)$ ($i = 0, 1, \dots, d-1$); $|C_i| \leq (1 - \beta)^i n$; $|F_d| = |C_{d-1}| = O(1)$;
- (2) For $1 \leq i \leq d-1$, $\tilde{S}_{F_i F_i}^{(i)}$ is α -RCDD;

Algorithm 3: Block Cholesky solver for directed Laplacians

Input: strongly connected Eulerian Laplacian $L \in \mathbb{R}^{n \times n}$;
query vectors $\{\mathbf{b}^{(q)}\}_{q=1}^Q \subseteq \mathbb{R}^n$ with each $\mathbf{b}^{(q)} \perp \mathbf{1}$;
error parameters $\{\epsilon_q\}_{q=1}^Q \subseteq (0, 1)$

Output: solutions $\{\mathbf{x}^{(q)}\}_{q=1}^Q \subseteq \mathbb{R}^n$

- 1 Call **SCHURCHAIN** ($L, 0.25, 0.1$) to compute a $\left\{0.25, 0.05, \left\{\frac{0.1}{i^2}\right\}_{i=1}^{O(\log n)}\right\}$ -Schur complement chain $\left\{\left\{\tilde{S}^{(i)}\right\}_{i=1}^d, \{F_i\}_{i=1}^d\right\}$ (Sections 4, 5, 6.1)
- 2 Generate the operator $Z(\mathbf{x}) =$
PRECONDITION $\left(\left\{\left\{\tilde{S}^{(i)}\right\}_{i=1}^d, \{F_i\}_{i=1}^d\right\}, \mathbf{x}, O(\log n)\right)$
(Section 6.2);
- 3 Using the preconditioned Richardson iteration with the preconditioner Z to solve the Laplacian systems: for each query vector $\mathbf{b}^{(q)}$, compute
 $\mathbf{x}^{(q)} \leftarrow \text{PR}\left(L, \mathbf{b}^{(q)}, Z(\cdot), 1, O(\log(n/\epsilon_q))\right)$

$$(3) \quad \tilde{S}^{(1)} - L \stackrel{\text{asym}}{\leq} \delta_1 \cdot U[L] \text{ and } \tilde{S}^{(i+1)} - \text{Sc}\left(\tilde{S}^{(i)}, F_i\right) \stackrel{\text{asym}}{\leq} \delta_{i+1} \cdot U\left[\text{Sc}\left(\tilde{S}^{(i)}, F_i\right)\right], \quad 1 \leq i \leq d-1;$$

$$(4) \quad U\left[\tilde{S}^{(1)}\right] \geq U[L] \text{ and } U\left[\tilde{S}^{(i+1)}\right] \geq U\left[\text{Sc}\left(\tilde{S}^{(i)}, F_i\right)\right], \quad 1 \leq i \leq d-1.$$

We also denote $F_0 = C_d = \emptyset, C_0 = [n]$ for notational simplicity.

Remark 6.2. Compared with the Schur complement chains for undirected graphs from [10], the only new condition is Condition (4). It guarantees the positive semi-definiteness of the symmetrization of the sparsified approximate Eulerian Laplacian \tilde{L} and the error-bounding matrix \tilde{B} defined in Section 6.2.

To construct a Schur complement chain, we first use the following lemma to find an α -RCDD subset F_1 , and then apply the Schur complement sparsification method SPARSE-SCHUR to compute $\tilde{S}^{(1)}$ which is an approximation for $\text{Sc}(L, F_1)$. Then, we repeat this process to get a desirable Schur complement chain.

LEMMA 6.3. (Theorem A.1 of [6]) Given an Eulerian Laplacian $L \in \mathbb{R}^{n \times n}$ with $\text{nnz}(L) = m$, the routine **FINDRCDBLOCK** outputs a subset $F \subseteq [n]$ such that $|F| \geq \frac{n}{16(1+\alpha)}$ and L_{FF} is α -RCDD in time $O\left(m \log \frac{1}{p}\right)$ with probability at least $1 - p$.

By Lemma 6.3, we can choose for instance $\alpha = 0.1$. So, we assume $\alpha = O(1)$, when analyzing the complexities below. Our method to construct a Schur complement chain is illustrated in Algorithm 4. Its running time is shown in Theorem 6.4.

THEOREM 6.4. Given a strongly connected Eulerian Laplacian $L \in \mathbb{R}^{n \times n}$ and parameters $\alpha = O(1), \delta \in (0, 1]$, the routine **SCHURCHAIN**

Algorithm 4: **SCHURCHAIN** (L, α, δ)

Input: strongly connected Eulerian Laplacian $L \in \mathbb{R}^{n \times n}$;
parameters $\alpha > 0, \delta \in (0, 1]$

Output: $\left(\alpha, \frac{1}{16(1+\alpha)}, \left\{\frac{\delta}{i^2}\right\}_{i=1}^d\right)$ -Schur complement chain $\left\{\left\{\tilde{S}^{(i)}\right\}_{i=1}^d, \{F_i\}_{i=1}^d\right\}$

- 1 Set $\delta'_i = \frac{\delta}{3i^2}$ for $i \geq 1$;
- 2 Compute $S^{(1)} \leftarrow \text{ORASPARSE-LAPLACIAN}(L, \delta'_1)$;
- 3 Let $\tilde{S}^{(1)} \leftarrow S^{(1)} + \frac{\delta'_1}{1-\delta'_1} U\left[S^{(1)}\right]$;
- 4 Set $i \leftarrow 0, C_0 = [n]$;
- 5 **while** $|C_i| > 100$ **do**
- 6 $i \leftarrow i + 1$;
- 7 $F_i \leftarrow \text{FINDRCDBLOCK}\left(\tilde{S}^{(i)}, \alpha\right)$;
- 8 $C_i \leftarrow C_{i-1} \setminus F_i$;
- 9 $S^{(i+1)} \leftarrow \text{SPARSE-SCHUR}\left(\tilde{S}^{(i)}, F_i, \delta_{i+1}\right)$;
- 10 $\tilde{S}^{(i+1)} \leftarrow S^{(i)} + \frac{\delta'_{i+1}}{1-\delta'_{i+1}} U\left[S^{(i)}\right]$
- 11 **end**
- 12 Return $\left\{\left\{\tilde{S}^{(i)}\right\}_{i=1}^d, \{F_i\}_{i=1}^d\right\}$

runs in time

$$O(\mathcal{T}_S(m, n, \delta)) + \tilde{O}\left(\mathcal{T}_S\left(\mathcal{N}_S(n, \delta) \delta^{-2}, n, \delta\right) \log n\right)$$

with high probability to return an $\left(\alpha, \frac{1}{16(1+\alpha)}, \left\{\frac{\delta}{i^2}\right\}_{i=1}^d\right)$ -Schur complement chain, where $d = O(\log n)$. In addition,

$$\sum_{i=1}^d \text{nnz}(\tilde{S}^{(i)}) = O(\mathcal{N}_S(n, \delta)).$$

6.2 Construction of the Preconditioner and the Solver

After constructing a desirable Schur complement chain, we use the Schur complement chain to construct a preconditioner and solve $\mathbf{L}\mathbf{x} = \mathbf{b}$ via the preconditioned Richardson iteration.

Consider a linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$, where \mathbf{b} is in the image space of \mathbf{A} . Given a preconditioner Z , the classical preconditioned Richardson iteration updates as follows:

$$\mathbf{x}^{(k+1)} \leftarrow \mathbf{x}^{(k)} + \eta Z(\mathbf{b} - \mathbf{A}\mathbf{x}^{(k)}).$$

We initialize $\mathbf{x}^{(0)} = \mathbf{0}$ for simplicity. This procedure is denoted by $\mathbf{x}^{(N)} = \text{PR}(\mathbf{A}, \mathbf{b}, Z, \eta, N)$.

We will use the following fundamental lemma to guarantee the convergence rate of the preconditioned Richardson iteration in our methods.

LEMMA 6.5. (Lemma 4.2 of [7]) Let $\mathbf{A}, \mathbf{Z}, \mathbf{U} \in \mathbb{R}^{n \times n}$, where \mathbf{U} is PSD and $\ker(\mathbf{U}) \subseteq \ker(\mathbf{Z}) = \ker(\mathbf{Z}^\top) = \ker(\mathbf{A}) = \ker(\mathbf{A}^\top)$. Let $\mathbf{b} \in \mathbb{R}^n$ be a vector inside the image space of \mathbf{A} . Denote the projection onto the image space of \mathbf{A} by $P_{\mathbf{A}}$. Denote $\mathbf{x}^{(N)} = \text{PR}(\mathbf{A}, \mathbf{b}, \mathbf{Z}, \eta, N)$.

Then, $\mathbf{x}^{(N)}$ satisfies

$$\left\| \mathbf{x}^{(N)} - \mathbf{A}^\dagger \mathbf{b} \right\|_U \leq \| \mathbf{P}_A - \eta \mathbf{Z} \mathbf{A} \|_{U \rightarrow U}^N \left\| \mathbf{A}^\dagger \mathbf{b} \right\|_U.$$

In addition, the preconditioned Richardson iteration is a linear operator with

$$\mathbf{x}^{(N)} = \eta \sum_{k=0}^{N-1} (\mathbf{P}_A - \eta \mathbf{Z} \mathbf{A})^k \mathbf{Z} \mathbf{b}.$$

Our construction for the preconditioner is illustrated in Algorithm 1. We analyze Algorithm 1 by representing the routine $\text{PR}(\cdot)$ as a linear operator which is equivalent to multiplying vector \mathbf{x} with the matrix $\Pi \tilde{\mathbf{Z}}$, where the definition of $\tilde{\mathbf{Z}} \in \mathbb{R}^{n \times n}$ is in the full version, and $\Pi = \mathbf{I} - \frac{\mathbf{1}\mathbf{1}^\top}{n}$ is the projection matrix onto the image space of \mathbf{L} . To analyze the quality of the preconditioner $\Pi \tilde{\mathbf{Z}}$, we need to provide bounds for $\Pi - \Pi \tilde{\mathbf{Z}} \mathbf{L}$. Define $\tilde{\mathbf{L}}$ as an approximation for \mathbf{L} with the errors induced by the Schur complement sparsification procedure

$$\tilde{\mathbf{L}} = \tilde{\mathbf{S}}^{(1)} + \sum_{i=1}^{d-1} \mathcal{P} \left(\tilde{\mathbf{S}}^{(i+1)} - \text{Sc} \left(\tilde{\mathbf{S}}^{(i)}, F_i \right), C_i, C_i, n \right).$$

Define an auxiliary matrix

$$\tilde{\mathbf{B}} = \delta_1 \mathbf{U} [\tilde{\mathbf{L}}] + \sum_{i=1}^{d-1} \delta_{i+1} \mathcal{P} \left(\mathbf{U} \left[\text{Sc} \left(\tilde{\mathbf{L}}, \cup_{j=1}^i F_j \right) \right], C_i, C_i, n \right).$$

We can prove that $O(1) \cdot \mathbf{U} [\mathbf{L}] \leq \mathbf{U} [\tilde{\mathbf{B}}] \leq O(\text{poly}(n)) \cdot \mathbf{U} [\mathbf{L}]$.

LEMMA 6.6. Given $\{\alpha, \beta, \{\delta_i\}_{i=1}^d\}$ -Schur complement chain with $d = O(\log n)$ and $\sum_{i=1}^d \delta_i \leq \frac{1}{4}$, by setting $N = O(\log n)$ in Algorithm 1, we have $\|\Pi - \Pi \tilde{\mathbf{Z}} \mathbf{L}\|_{\tilde{\mathbf{B}} \rightarrow \tilde{\mathbf{B}}} \leq \frac{1}{2}$.

Then, Theorem 1.1 follows by Lemma 6.4 and Lemma 6.6.

Using the smaller Eulerian Laplacian sparsifiers based on short cycle decompositions to sparsify the approximate Schur complements returned by Algorithm 2, we get the following solver which has quadratic processing time, but faster solve time.

COROLLARY 6.7. Given a strongly connected Eulerian Laplacian $\mathbf{L} \in \mathbb{R}^{n \times n}$, we can process it time $O(n^2 \log^{O(1)} n)$. Then, for each query vector $\mathbf{b} \in \mathbb{R}^n$ with $\mathbf{b} \perp \mathbf{1}$, we can compute a vector $\mathbf{x} \in \mathbb{R}^n$ with $\|\mathbf{x} - \mathbf{L}^\dagger \mathbf{b}\|_{U[\mathbf{L}]} \leq \epsilon \|\mathbf{L}^\dagger \mathbf{b}\|_{U[\mathbf{L}]}$ in time $O(n \log^5 n \log(n/\epsilon))$.

Remark 6.8. Combining Theorem 1.1 or Corollary 6.7 with Appendix D of [7] yields full solvers for strongly connected directed Laplacians.

REFERENCES

- [1] AmirMahdi Ahmadianejad, Arun Jambulapati, Amin Saberi, and Aaron Sidford. 2019. Perron-Frobenius Theory in Nearly Linear Time: Positive Eigenvectors, M-matrices, Graph Kernels, and Other Applications. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, January 6-9, 2019, Timothy M. Chan (Ed.)*. SIAM, San Diego, California, USA, 1387–1404. <https://doi.org/10.1137/1.9781611975482.85>
- [2] AmirMahdi Ahmadianejad, Jonathan A. Kelner, Jack Murtagh, John Peebles, Aaron Sidford, and Salil P. Vadhan. 2020. High-precision Estimation of Random Walks in Small Space. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, November 16-19, 2020, IEEE, Durham, NC, USA, 1295–1306*. <https://doi.org/10.1109/FOCS46700.2020.00123>
- [3] Joshua Batson, Daniel A. Spielman, Nikhil Srivastava, and Shang-Hua Teng. 2013. Spectral sparsification of graphs: theory and algorithms. *CACM* 56, 8 (Aug. 2013), 87–94. <https://doi.org/10.1145/2492007.2492029>
- [4] Dehua Cheng, Yu Cheng, Yan Liu, Richard Peng, and Shang-Hua Teng. 2015. Efficient Sampling for Gaussian Graphical Models via Spectral Sparsification. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015, July 3-6, 2015 (JMLR Workshop and Conference Proceedings, Vol. 40)*, Peter Grünwald, Elad Hazan, and Satyen Kale (Eds.). JMLR.org, Paris, France, 364–390. <http://proceedings.mlr.press/v40/Cheng15.html>
- [5] Timothy Chu, Yu Gao, Richard Peng, Sushant Sachdeva, Saurabh Sawlani, and Junxing Wang. 2018. Graph Sparsification, Spectral Sketches, and Faster Resistance Computation, via Short Cycle Decompositions. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, October 7-9, 2018, Mikkel Thorup (Ed.)*. IEEE Computer Society, Paris, France, 361–372. <https://doi.org/10.1109/FOCS.2018.00042>
- [6] Michael B Cohen, Jonathan Kelner, Rasmus Kyng, John Peebles, Richard Peng, Anup B Rao, and Aaron Sidford. 2018. Solving directed laplacian systems in nearly-linear time through sparse LU factorizations. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, October 7-9, 2018*. IEEE Computer Society, Paris, France, 898–909. <https://doi.org/10.1109/FOCS.2018.00089>
- [7] Michael B Cohen, Jonathan Kelner, John Peebles, Richard Peng, Anup B Rao, Aaron Sidford, and Adrian Vladu. 2017. Almost-linear-time algorithms for markov chains and new spectral primitives for directed graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, June 19-23, 2017, ACM, Montreal, QC, Canada, 410–419*. <https://doi.org/10.1145/3055399.3055463>
- [8] Michael B. Cohen, Jonathan A. Kelner, John Peebles, Richard Peng, Aaron Sidford, and Adrian Vladu. 2016. Faster Algorithms for Computing the Stationary Distribution, Simulating Random Walks, and More. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Irit Dinur (Ed.)*. IEEE Computer Society, Hyatt Regency, New Brunswick, New Jersey, USA, 583–592. <https://doi.org/10.1109/FOCS.2016.69>
- [9] Jonathan A. Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. 2013. A simple, combinatorial algorithm for solving SDD systems in nearly-linear time. In *Symposium on Theory of Computing Conference, STOC'13, June 1-4, 2013, Dan Boneh, Tim Roughgarden, and Joan Feigenbaum (Eds.)*. ACM, Palo Alto, CA, USA, 911–920. <https://doi.org/10.1145/2488608.2488724>
- [10] Rasmus Kyng, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Daniel A Spielman. 2016. Sparsified cholesky and multigrid solvers for connection laplacians. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, June 18-21, 2016*. ACM, Cambridge, MA, USA, 842–850. <https://doi.org/10.1145/2897518.2897640>
- [11] Rasmus Kyng and Sushant Sachdeva. 2016. Approximate Gaussian Elimination for Laplacians - Fast, Sparse, and Simple. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Irit Dinur (Ed.)*. IEEE Computer Society, Hyatt Regency, New Brunswick, New Jersey, USA, 573–582. <https://doi.org/10.1109/FOCS.2016.68>
- [12] Yang P. Liu, Sushant Sachdeva, and Zejun Yu. 2019. Short Cycles via Low-Diameter Decompositions. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, January 6-9, 2019, Timothy M. Chan (Ed.)*. SIAM, San Diego, California, USA, 2602–2615. <https://doi.org/10.1137/1.9781611975482.161>
- [13] Merav Parter and Eylon Yogev. 2019. Optimal Short Cycle Decomposition in Almost Linear Time. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019 (LIPIcs, Vol. 132)*, Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Patras, Greece, 89:1–89:14. <https://doi.org/10.4230/LIPIcs.ICALP.2019.89>
- [14] John Peebles. 2019. *Fast spectral primitives for directed graphs*. Ph. D. Dissertation. Massachusetts Institute of Technology. Available at: <https://dspace.mit.edu/handle/1721.1/124075>.
- [15] Richard Peng and Zhuoqing Song. 2021. Sparsified Block Elimination for Directed Laplacians. *arXiv preprint arXiv:2111.10257* (2021).
- [16] Richard Peng and Daniel A Spielman. 2014. An efficient parallel solver for SDD linear systems. In *Symposium on Theory of Computing, STOC 2014, May 31 - June 03, 2014*. ACM, New York, NY, USA, 333–342. <https://doi.org/10.1145/2591796.2591832>
- [17] Daniel A. Spielman. 2010. Algorithms, Graph Theory, and Linear Equations in Laplacian Matrices. In *Proceedings of the International Congress of Mathematicians 2010 (ICM 2010)*. World Scientific, Hyderabad, India, 2698–2722. https://doi.org/10.1142/9789814324359_0164
- [18] Shang-Hua Teng. 2010. The Laplacian Paradigm: Emerging Algorithms for Massive Graphs. In *Theory and Applications of Models of Computation, 7th Annual Conference, TAMC 2010, June 7-11, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6108)*, Jan Kratochvíl, Angsheng Li, Jiri Fiala, and Petr Kolman (Eds.). Springer, Prague, Czech Republic, 2–14. https://doi.org/10.1007/978-3-642-13562-0_2
- [19] Jan van den Brand. 2021. Unifying Matrix Data Structures: Simplifying and Speeding up Iterative Algorithms. In *4th Symposium on Simplicity in Algorithms, SOSA 2021, January 11-12, 2021, Hung Viet Le and Valerie King (Eds.)*. SIAM, Virtual Conference, 1–13. <https://doi.org/10.1137/1.9781611976496.1>