

Sparsified Cholesky and Multigrid Solvers for Connection Laplacians*

Rasmus Kyng[†]
Yale University
rasmus.kyng@yale.edu

Yin Tat Lee[‡]
M.I.T.
yintat@mit.edu

Richard Peng
Georgia Tech
rpeng@cc.gatech.edu

Sushant Sachdeva[§]
Yale University
sachdeva@cs.yale.edu

Daniel A. Spielman[¶]
Yale University
spielman@cs.yale.edu

December 8, 2015

Abstract

We introduce the *sparsified Cholesky* and *sparsified multigrid* algorithms for solving systems of linear equations. These algorithms accelerate Gaussian elimination by sparsifying the nonzero matrix entries created by the elimination process.

We use these new algorithms to derive the first nearly linear time algorithms for solving systems of equations in connection Laplacians—a generalization of Laplacian matrices that arise in many problems in image and signal processing.

We also prove that every connection Laplacian has a linear sized approximate inverse. This is an LU factorization with a linear number of nonzero entries that is a strong approximation of the original matrix. Using such a factorization one can solve systems of equations in a connection Laplacian in linear time. Such a factorization was unknown even for ordinary graph Laplacians.

*This paper incorporates and improves upon results previously announced by a subset of the authors in [LPS15].

[†]Supported by NSF grant CCF-1111257.

[‡]Supported in part by NSF awards 0843915 and 1111109. Part of this work was done while visiting the Simons Institute for the Theory of Computing, UC Berkeley.

[§]Supported by a Simons Investigator Award to Daniel A. Spielman.

[¶]Supported by AFOSR Award FA9550-12-1-0175, NSF grant CCF-1111257, a Simons Investigator Award to Daniel A. Spielman, and a MacArthur Fellowship.

1 Introduction

We introduce the *sparsified Cholesky* and *sparsified multigrid* algorithms for solving systems of linear equations. Two advantages of these algorithms over other recently introduced nearly-linear time algorithms for solving systems of equations in Laplacian matrices [Vai90, ST14, KMP10, KMP11, KOSZ13, CKM⁺14] are:

1. They give nearly-linear time algorithms for solving systems of equations in a much broader class of matrices—the connection Laplacians and Hermitian block diagonally dominant matrices. Connection Laplacians [SW12, BSS13] are a generalization of graph Laplacians that arise in many applications, including celebrated work on cryo-electron microscopy [SS11a, SS12, ZS14], phase retrieval [ABFM14, MTW14], and many image processing problems (e.g. [OSB15, ANKKS⁺12]). Previous algorithms for solving systems of equations in graph Laplacians cannot be extended to solve equations in connection Laplacians because the previous algorithms relied on some form of low stretch spanning trees, a concept that has no analog for the more general connection Laplacians.
2. They provide linear-sized approximate inverses of connection Laplacian matrices. That is, for every n -dimensional connection Laplacian \mathbf{M} , the sparsified Cholesky factorization algorithm produces a block-upper-triangular matrix \mathbf{U} and block diagonal \mathbf{D} with $O(n)$ nonzero entries so that $\mathbf{U}^T \mathbf{D} \mathbf{U}$ is a constant-factor approximation of \mathbf{M} . Such matrices \mathbf{U} and \mathbf{D} allow one to solve systems of equations in \mathbf{M} to accuracy ϵ in time $O(m \log \epsilon^{-1})$, where m is the number of nonzero entries of \mathbf{M} . Even for ordinary Laplacian matrices, the existence of such approximate inverses is entirely new. The one caveat of this result is that we do not yet know how to compute those approximate inverses in nearly linear time.

The sparsified Cholesky and sparsified multigrid algorithms work by sparsifying the matrices produced during Gaussian elimination. Recall that Cholesky factorization is the version of Gaussian elimination for symmetric matrices, and that the high running time of Gaussian elimination comes from *fill*—new nonzero matrix entries that are created by row operations. If Gaussian elimination never produced rows with a super-constant number of entries, then it would run in linear time. The sparsified Cholesky algorithm accelerates Gaussian elimination by sparsifying the rows that are produced by elimination, thereby guaranteeing that the elimination will be fast. Sparsified Cholesky is inspired by one of the major advances in algorithms for solving linear equations in Laplacian matrices—the Incomplete Cholesky factorization (ICC) [MV77]. However, ICC merely drops entries produced by elimination, whereas sparsification also carefully increases ones that remain. The difference is crucial, and is why ICC does not provide a nearly-linear time solver.

To control the error introduced by sparsification, we have to be careful not to do it too often. This means that our algorithm actually chooses a large set of rows and columns to eliminate at once, and then sparsifies the result. This is the basis of our first algorithms, which establish the existence of linear time solvers and linear sized approximate inverses, after precomputation. This precomputation is analogous to the procedure of computing a matrix inverse: the approximate inverse takes much less time to apply than to compute.

To produce entire algorithms that run in nearly linear time (both to compute and apply the approximate inverse) requires a little more work. To avoid the work of computing the matrix obtained by eliminating the large set of rows and columns, we design a fast algorithm for approximating it quickly. This resulting algorithm produces a solver routine that does a little more work at each level, and so resembles a multigrid V-cycle [TOS00]. We call the resulting algorithm the *sparsified multigrid*. We note that Krishnan, Fattal, and Szeliski [KFS13] present experimental results from the use of a sparsification heuristic in a multigrid algorithm for solving problems in computer vision.

Our new algorithms are most closely related to the Laplacian solver recently introduced by Peng and Spielman [PS14]: unlike the other Laplacian solvers, they rely only on sparsification and do not directly rely on “support theory” preconditioners or any form of low stretch spanning trees. This is why our algorithms can solve a much broader family of linear systems. To sparsify without using graph theoretic algorithms, we employ a recently developed algorithm of Cohen *et. al.* [CLM⁺14] that allows us to sparsify a matrix by solving systems of equations in a subsampled matrix.

1.1 Connection Laplacians and Block DD Matrices

In this section, we define block diagonally dominant (bDD) matrices—the most general family of matrices such that the associated systems of linear equations can be solved by our algorithms. We begin by defining our motivating case: the connection Laplacians

Connection Laplacians may be thought of as a generalization of graph Laplacians where every vertex is associated with a vector, instead of a real number, and every edge is associated with a unitary matrix. Like graph Laplacians, they describe a natural quadratic form. Let $\mathbf{M}_{[i,j]} \in \mathbb{C}^{r \times r}$ be the unitary matrix associated with edge (i, j) , and let $w_{i,j}$ be the (nonnegative) weight of edge (i, j) . We require that $\mathbf{M}_{[i,j]} = \mathbf{M}_{[j,i]}^*$, where $*$ denotes the conjugate transpose. The quadratic form associated with this connection Laplacian is a function of vectors $\mathbf{v}^{(i)} \in \mathbb{C}^r$, one for each vertex i , that equals

$$\sum_{(i,j) \in E} w_{i,j} \|\mathbf{v}^{(i)} - \mathbf{M}_{[i,j]} \mathbf{v}^{(j)}\|^2.$$

The matrix corresponding to this quadratic form is a block matrix with blocks $\mathbf{M}_{[i,j]}$. Most applications of the connection Laplacian require one to either solve systems of linear equations in this matrix, or to compute approximations of its smallest eigenvalues and eigenvectors. By applying the inverse power method (or inverse Lanczos), we can perform these eigenvector calculations by solving a logarithmic number of linear systems in the matrix (see [ST14, Section 7]).

The matrices obtained from the connection Laplacian are a special case of block diagonally dominant (bDD) matrices, which we now define. Throughout this paper, we consider block matrices having entries in $\mathbb{C}^{r \times r}$, where $r > 0$ is a fixed integer. We say that $\mathbf{M} \in (\mathbb{C}^{r \times r})^{m \times n}$ has m block-rows, and n block-columns. For $i \in [m], j \in [n]$, we let $\mathbf{M}_{[i,j]} \in \mathbb{C}^{r \times r}$ denote the i, j -block in \mathbf{M} , and $\mathbf{M}_{[j]}$ denote the j^{th} block-column, i.e., $\mathbf{M}_{[j]} = [(\mathbf{M}_{[1,j]})^*, (\mathbf{M}_{[2,j]})^*, \dots, (\mathbf{M}_{[m,j]})^*]^*$. For sets $F \subseteq [m], C \subseteq [n]$, we let $\mathbf{M}_{[F,C]} \in (\mathbb{C}^{r \times r})^{|F| \times |C|}$ denote the block-submatrix with blocks $\mathbf{M}_{[i,j]}$ for $i \in F, j \in C$. \mathbf{M} is block-diagonal, if $\mathbf{M}_{[i,j]} = 0$ for $i \neq j$. *We emphasize that all computations are done over \mathbb{C} , not over a matrix group.*

To define bSDD matrices, let $\|\mathbf{A}\|$ denote the operator norm¹ of a matrix \mathbf{A} . For Hermitian matrices \mathbf{A}, \mathbf{B} , write $\mathbf{A} \succeq \mathbf{B}$ iff $\mathbf{A} - \mathbf{B}$ is positive semidefinite.

Definition 1.1. A Hermitian block-matrix $\mathbf{M} \in (\mathbb{C}^{r \times r})^{n \times n}$ is block diagonally dominant (or bDD) if

$$\text{for all } i \in [n], \quad \mathbf{M}_{[i,i]} \succeq \mathbb{I}_r \cdot \sum_{j:j \neq i} \|\mathbf{M}_{[i,j]}\|,$$

where $\mathbb{I}_r \in \mathbb{C}^{r \times r}$ denotes the identity matrix.

Equivalently, a Hermitian \mathbf{M} is a bDD matrix iff it can be written as $\mathbf{D} - \mathbf{A}$ where \mathbf{D} is block-diagonal, and $\mathbf{D}_{[i,i]} \succeq \mathbb{I}_r \sum_j \|\mathbf{A}_{[i,j]}\|$ (see Lemma B.1).

Throughout the paper we treat r as a constant. The hidden dependence on r , of the running times of the algorithms we present, is polynomial. The major results of this paper are the following.

¹Recall that the operator norm is the largest singular value of \mathbf{A} and square root of the largest eigenvalue of $\mathbf{A}^* \mathbf{A}$.

Theorem 1.2 (Sparsified Multigrid). *There is an algorithm that, when given a bDD matrix \mathbf{M} with n block rows and m nonzero blocks, produces a solver for \mathbf{M} in $O(m \log n + n \log^{2+o(1)} n)$ work and $O(n^{o(1)})$ depth so that the solver finds ϵ -approximate solutions to systems of equations in \mathbf{M} in $O((m + n \log^{1+o(1)} n) \log(1/\epsilon))$ work and $O(\log^2 n \log \log n \log(1/\epsilon))$ depth.*

Previously, the existence of nearly-linear time solvers was even unknown for the 1-dimensional case ($r = 1$) when the off-diagonal entries were allowed to be complex numbers.

We can use the above algorithm to find approximations of the smallest eigenvalues and eigenvectors of such matrices at an additional logarithmic cost.

Theorem 1.3 (Sparsified Cholesky). *For every bDD matrix \mathbf{M} with n block-rows there exists a diagonal matrix \mathbf{D} and an upper triangular matrix \mathbf{U} with $O(n)$ nonzero blocks so that*

$$\mathbf{U}^T \mathbf{D} \mathbf{U} \approx_{3/4} \mathbf{M}.$$

Moreover, linear equations in \mathbf{U} , \mathbf{U}^T , and \mathbf{D} can be solved with linear work in depth $O(\log^2 n)$, and these matrices can be computed in polynomial time.

These matrices allow one to solve systems of linear equations in \mathbf{M} to ϵ -accuracy in parallel time $O(\log^2 n \log^{-1} \epsilon)$ and work $O(m \log^{-1} \epsilon)$. Results of this form were previously unknown even for graph Laplacians.

In the next two sections we explain the ideas used to prove these theorems. Proofs may be found in the appendices that follow.

2 Background

We require some standard facts about the \preceq order on matrices.

Fact 2.1. *For \mathbf{A} and \mathbf{B} positive definite, $\mathbf{A} \succ \mathbf{B}$ if and only if $\mathbf{B}^{-1} \succ \mathbf{A}^{-1}$.*

Fact 2.2. *If $\mathbf{A} \succ \mathbf{B}$ and \mathbf{C} is any matrix of compatible dimension, then $\mathbf{C} \mathbf{A} \mathbf{C}^T \succ \mathbf{C} \mathbf{B} \mathbf{C}^T$.*

We say that \mathbf{A} is an ϵ -approximation of \mathbf{B} , written $\mathbf{A} \approx_\epsilon \mathbf{B}$, if $e^\epsilon \mathbf{B} \succ \mathbf{A} \succ e^{-\epsilon} \mathbf{B}$. This relation is symmetric. We say that $\tilde{\mathbf{x}}$ is an ϵ -approximate solution to the system $\mathbf{A} \mathbf{x} = \mathbf{b}$ if $\|\tilde{\mathbf{x}} - \mathbf{A}^{-1} \mathbf{b}\|_{\mathbf{A}} \leq \epsilon \|\mathbf{x}\|_{\mathbf{A}}$, where $\|\mathbf{x}\|_{\mathbf{A}} = (\mathbf{x}^T \mathbf{A} \mathbf{x})^{1/2}$. If $\epsilon < 1/2$, $\mathbf{A} \approx_\epsilon \mathbf{B}$ and $\mathbf{B} \tilde{\mathbf{x}} = \mathbf{b}$, then $\tilde{\mathbf{x}}$ is a 2ϵ approximate solution to $\mathbf{A} \mathbf{x} = \mathbf{b}$.

Fact 2.3. *If $\mathbf{A} \approx_\epsilon \mathbf{B}$ and $\mathbf{B} \approx_\delta \mathbf{C}$, then $\mathbf{A} \approx_{\epsilon+\delta} \mathbf{C}$.*

We now address one technicality of dealing with bDD matrices: it is not immediate whether or not a bDD matrix is singular. Moreover, if it is singular, the structure of its null space is not immediately clear either. Throughout the rest of this paper, we will consider bDD matrices to which a small multiple of the identity have been added. These matrices will be nonsingular. To reduce the problem of solving equations in a general bDD matrix \mathbf{M} to that of solving equations in a nonsingular matrix, we require an estimate of the smallest nonzero eigenvalue of \mathbf{M} .

Claim 2.4. *Suppose that all nonzero eigenvalues of \mathbf{M} are at least μ and $\mathbf{Z} \approx_\epsilon (\mathbf{M} + \epsilon \mu \mathbb{I})^{-1}$ for some $0 < \epsilon < 1/2$. Given any \mathbf{b} such that $\mathbf{M} \mathbf{x} = \mathbf{b}$ for some \mathbf{x} , we have*

$$\|\mathbf{x} - \mathbf{Z} \mathbf{b}\|_{\mathbf{M}}^2 \leq 6\epsilon \|\mathbf{x}\|_{\mathbf{M}}^2.$$

Hence, we can solve systems in \mathbf{M} by approximately solving systems in $\mathbf{M} + \epsilon\mu\mathbf{I}$. Any lower bound μ on the smallest nonzero eigenvalue of \mathbf{M} will suffice. It only impacts the running times of our algorithms in the numerical precision with which one must carry out the computations.

The above claim allows us to solve systems in \mathbf{M} that have a solution. If \mathbf{M} is singular and we want to apply its pseudoinverse (that is, to find the closest solution in the range of \mathbf{M}), we can do so by pre and post multiplying by \mathbf{M} to project onto its range. The resulting algorithm, which is implicit in the following claim, requires applying a solver for \mathbf{M} three times. It also takes $O(\log \kappa)$ times as long to run, where κ is the finite condition number² of \mathbf{M} .

Claim 2.5. *Let κ be an upper bound on the finite condition number of \mathbf{M} . Given an error parameter $0 < \epsilon < 1$, let $\delta = \epsilon/56\kappa^3$. If all nonzero eigenvalues of \mathbf{M} are at least μ , and if $\mathbf{Z} \approx_\delta (\mathbf{M} + \epsilon\mu\mathbf{I})^{-1}$, then*

$$\mathbf{M}\mathbf{Z}^3\mathbf{M} \approx_{4\epsilon} \mathbf{M}^+.$$

3 Overview of the algorithms

We index the block rows and columns of a block matrix by a set of vertices (or indices) V . When we perform an elimination, we eliminate a set $F \subset V$, and let $C = V - F$. Here, F stands for “fine” and C stands for “coarse”. In contrast with standard multigrid methods, we will have $|F| \leq |C|$.

To describe block-Cholesky factorization, we write the matrix \mathbf{M} with the rows and columns in F first:

$$\mathbf{M} = \begin{bmatrix} \mathbf{M}_{[F,F]} & \mathbf{M}_{[F,C]} \\ \mathbf{M}_{[C,F]} & \mathbf{M}_{[C,C]} \end{bmatrix}.$$

Cholesky factorization writes the inverse of this matrix as

$$\mathbf{M}^{-1} = \begin{bmatrix} \mathbf{I} & -\mathbf{M}_{[F,F]}^{-1}\mathbf{M}_{[F,C]} \\ 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{M}_{[F,F]}^{-1} & 0 \\ 0 & S_c(\mathbf{M}, F)^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & 0 \\ -\mathbf{M}_{[C,F]}\mathbf{M}_{[F,F]}^{-1} & \mathbf{I} \end{bmatrix}, \quad (1)$$

where

$$S_c(\mathbf{M}, F) \stackrel{\text{def}}{=} \mathbf{M}_{[C,C]} - \mathbf{M}_{[C,F]}\mathbf{M}_{[F,F]}^{-1}\mathbf{M}_{[F,C]}$$

is the Schur complement of \mathbf{M} with respect to F .

Our algorithms rely on two fundamental facts about *bDD* matrices: that the Schur complement of a *bDD* matrix is a *bDD* matrix (Lemma B.3) and that one can sparsify *bDD* matrices. The following theorem is implicit in [dCSHS11].

Theorem 3.1. *For every $\epsilon \leq 1$, every *bDD* matrix \mathbf{M} with n r -dimensional block rows and columns can be ϵ -approximated by a *bDD* matrix having at most $10nr/\epsilon^2$ nonzero blocks.*

We use the identity (1) to reduce the problem of solving a system of equations in \mathbf{M} to that of solving equations in its Schur complement. The easiest part of this is multiplication by $\mathbf{M}_{[C,F]}$: the time is proportional to the number of nonzero entries in the submatrix, and we can sparsify \mathbf{M} to guarantee that this is small.

The costlier part of the reduction is the application of the inverse of $\mathbf{M}_{[F,F]}$ three times. This would be fast if $\mathbf{M}_{[F,F]}$ were block-diagonal, which corresponds to F being an independent set. We cannot find a sufficiently large independent set F , but we can find a large set that is almost independent. This results in a matrix $\mathbf{M}_{[F,F]}$ that is well approximated by its diagonal, and thus linear equations in this matrix can be quickly solved to high accuracy by a few Jacobi iterations (see Theorem 3.7 and Section E).

²The finite condition number is the ratio of the largest singular value to the smallest *nonzero* singular value.

We can prove the existence of linear sized approximate inverses by explicitly computing $Sc(\mathbf{M}, F)$, sparsifying it, and then recursively applying the algorithm just described. To make this algorithm efficient, we must compute a sparse approximation to $Sc(\mathbf{M}, F)$ without constructing $Sc(\mathbf{M}, F)$. This is the problem of *spectral vertex sparsification*, and we provide a fast algorithm for this task in Sections 3.4 and G.

3.1 Schur Complement Chains

We encode this recursive algorithm by a *Schur complement chain* (SCC). An SCC defines a linear operator that can be used to approximately solve equations in an initial matrix $\mathbf{M}^{(0)}$. If the *bDD* matrix \mathbf{M} is sparse, then it is the same as $\mathbf{M}^{(0)}$; if not, then $\mathbf{M}^{(0)}$ is a sparse approximation to \mathbf{M} . Let F_1 be the first set of vertices eliminated, $\mathbf{M}^{(1)}$ a sparse approximation to the Schur complement of $\mathbf{M}^{(0)}$ with respect to F_1 , and $\mathbf{Z}^{(1)}$ an operators that approximates the inverse of $\mathbf{M}_{[F_1, F_1]}^{(0)}$.

Definition 3.2 (Schur Complement Chain). An ϵ -Schur complement chain (ϵ -SCC) for a matrix \mathbf{M}^0 indexed by vertex set V is a sequence of operators and subsets,

$$\left((\mathbf{M}^{(1)}, \mathbf{Z}^{(1)}), \dots, (\mathbf{M}^{(d)}, \mathbf{Z}^{(d)}); F_1, \dots, F_d \right),$$

so that for $C_0 = V$ and $C_{i+1} = C_i \setminus F_{i+1}$, $|C_d| \leq 1000$ and for $1 \leq i \leq d$,

$$\mathbf{M}^{(i)} \approx_{\epsilon_i} Sc\left(\mathbf{M}^{(i-1)}, F_i\right) \quad \text{and} \quad 0 \preceq (\mathbf{Z}^{(i)})^{-1} - \mathbf{M}_{[F_i, F_i]}^{(i-1)} \preceq \epsilon_i \cdot Sc\left(\mathbf{M}^{(i-1)}, C_i\right).$$

The algorithm APPLYCHAIN, described in Section C, applies an SCC to solve equations in $\mathbf{M}^{(0)}$ in the natural way and satisfies the following guarantee.

Lemma 3.3. Consider an ϵ -SCC for $\mathbf{M}^{(0)}$ where $\mathbf{M}^{(i)}$ and $\mathbf{Z}^{(i)}$ can be applied to a vector in work $W_{\mathbf{M}^{(i)}}$, $W_{\mathbf{D}^{(i)}}$ and depth $D_{\mathbf{M}^{(i)}}$, $D_{\mathbf{Z}^{(i)}}$ respectively.

The algorithm APPLYCHAIN $\left((\mathbf{M}^{(1)}, \mathbf{Z}^{(1)}), \dots, (\mathbf{M}^{(d)}, \mathbf{Z}^{(d)}); F_1, \dots, F_d; \mathbf{b} \right)$ corresponds to a linear operator \mathbf{W} acting on \mathbf{b} such that

1. $\mathbf{W}^{-1} \approx_{\sum_{i=1}^d 2\epsilon_i} \mathbf{M}^{(0)}$, and
2. for any vector \mathbf{b} , APPLYCHAIN $\left((\mathbf{M}^{(1)}, \mathbf{Z}^{(1)}), \dots, (\mathbf{M}^{(d)}, \mathbf{Z}^{(d)}); F_1, \dots, F_d; \mathbf{b} \right)$ runs in $O\left(\sum_{i=1}^d (D_{\mathbf{M}^{(i)}} + D_{\mathbf{Z}^{(i)}})\right)$ depth and $O\left(\sum_{i=1}^d (W_{\mathbf{M}^{(i)}} + W_{\mathbf{Z}^{(i)}})\right)$ work.

For an ϵ -SCC chain for $\mathbf{M}^{(0)}$, we define the depth and work of the chain to be the work and depth required by APPLYCHAIN to apply the SCC.

3.2 Choosing F_i : α -bDD matrices

We must choose the set of vertices F_i so that we can approximate the inverse of $\mathbf{M}_{[F_i, F_i]}^{(i)}$ by an operator $\mathbf{Z}^{(i)}$ that is efficiently computable. We do this by requiring that the matrix $\mathbf{M}_{[F_i, F_i]}^{(i)}$ be α -block diagonally dominant (α -bDD), a term that we now define.

Definition 3.4. A Hermitian block-matrix \mathbf{M} is α -bDD if

$$\forall i, \quad \mathbf{M}_{[i, i]} \succcurlyeq (1 + \alpha) \mathbb{I}_r \cdot \sum_{j: j \neq i} \|\mathbf{M}_{[i, j]}\|. \quad (2)$$

We remark that a 0-bDD matrix is simply a bDD matrix. In particular, for $r = 1$, Laplacian matrices are 0-bDD.

By picking a subset of rows at random and discarding those that violate condition (2), the algorithm `BDDSUBSET` (described in Section D) finds a linear sized subset F of the block-rows of a bDD matrix \mathbf{M} so that $\mathbf{M}_{[F,F]}$ is α -bDD.

Lemma 3.5. *Given a bDD matrix \mathbf{M} with n block-rows, and an $\alpha \geq 0$, `BDDSUBSET` computes a subset F of size least $n/(8(1 + \alpha))$ such that $\mathbf{M}_{[F,F]}$ is α -bDD. It runs in $O(m)$ expected work and $O(\log n)$ expected depth, where m is the number of nonzero blocks in \mathbf{M} .*

We can express an α -bDD matrix as a sum of a block diagonal matrix and a bDD matrix so that it is well-approximated by the diagonal.

Lemma 3.6. *Every α -bDD matrix \mathbf{M} can be written in the form $\mathbf{X} + \mathbf{L}$ where \mathbf{X} is block-diagonal, \mathbf{L} is bDD, and $\mathbf{X} \succcurlyeq \frac{\alpha}{2}\mathbf{L}$.*

As \mathbf{L} is positive semidefinite, $(1 + 2/\alpha)\mathbf{X} \succcurlyeq \mathbf{M} \succcurlyeq \mathbf{X}$, which means that \mathbf{X} is a good approximation of \mathbf{M} when α is reasonably big. As block-diagonal matrices like \mathbf{X} are easy to invert, systems in these well-conditioned matrices can be solved rapidly using preconditioned iterative methods. In Section E, we show that a variant of Jacobi iteration provides an operator that satisfies the requirements of Definition 3.2.

Theorem 3.7. *Let \mathbf{M} be a bDD matrix with index set V , and let $F \subseteq V$ such that $\mathbf{M}_{[F,F]}$ is α -bDD for some $\alpha \geq 4$, and has m_{FF} nonzero blocks. The algorithm `JACOBI`($\epsilon, \mathbf{M}_{[F,F]}, \mathbf{b}$) acts as a linear operator \mathbf{Z} on \mathbf{b} that satisfies*

$$0 \preceq (\mathbf{Z})^{-1} - \mathbf{M}_{[F,F]} \preceq \epsilon \cdot \text{Sc}(\mathbf{M}, V \setminus F).$$

The algorithm takes $O(m_{FF} \log(\frac{1}{\epsilon}))$ work and $O(\log n \log(\frac{1}{\epsilon}))$ depth.

We now explain how Theorems 3.7 and 3.1 allow us to construct Schur complement chains that can be applied in nearly linear time. We optimize the construction in the next section.

Theorem 3.1 tells us that there is a bDD matrix $\mathbf{M}^{(0)}$ with $O(n/\epsilon^2)$ nonzero blocks that ϵ -approximates \mathbf{M} , and that for every i there is a bDD matrix $\mathbf{M}^{(i+1)}$ with $O(|C_i|/\epsilon^2)$ nonzero blocks that is an ϵ -approximation of $\text{Sc}(\mathbf{M}^{(i)}, F_i)$. We will pick ϵ later. Lemma 3.5 provides a set F_i containing a constant fraction of the block-rows of $\mathbf{M}^{(i)}$ so that $\mathbf{M}_{[F_i,F_i]}^{(i)}$ is 4-bDD. Theorem 3.7 then provides an operator that solves systems in $\mathbf{M}_{[F_i,F_i]}^{(i)}$ to ϵ accuracy in time $O(\log \epsilon^{-1})$ times the number of nonzero entries in $\mathbf{M}_{[F_i,F_i]}^{(i)}$. This is at most the number of nonzero entries in $\mathbf{M}^{(i)}$, and thus at most $O(|C_i|/\epsilon^2)$. As each F_i contains at least a constant fraction of the rows of $\mathbf{M}^{(i)}$, the depth of the recursion, d , is $O(\log n)$. Thus, we can obtain constant accuracy by setting $\epsilon = \Theta(1/\log n)$. The time required to apply the resulting SCC would thus be $O(n \log^2 n \log \log n)$.

We can reduce the running time by setting ϵ_1 to a constant and allowing it to shrink as i increases. For example, setting $\epsilon_i = 1/2(i + 1)^2$ results in a linear time algorithm that produces a constant-factor approximation of the inverse of \mathbf{M} . We refine this idea in the next section.

3.3 Linear Sized Approximate Inverses

In this section we sketch the proof of Theorem 1.3, which tells us that every bDD matrix has a linear-sized approximate inverse. The rest of the details appear in Section F.

The linear-sized approximate inverse of a bDD matrix \mathbf{M} with n block rows and columns is provided by a 3/4-approximate of the form $\mathbf{U}^T \mathbf{D} \mathbf{U}$ where \mathbf{D} is block diagonal and \mathbf{U} is block

upper-triangular and has $O(n)$ nonzero blocks. As systems of linear equations in block-triangular matrices like \mathbf{U} and \mathbf{U}^T can be solved in time proportional to their number of nonzero blocks, this provides a linear time algorithm for computing a 3/4 approximation of the inverse of \mathbf{M} . By iteratively refining the solutions provided by the approximate inverse, this allows one to find ϵ -accurate solutions to systems of equations in \mathbf{M} in time $O(m \log \epsilon^{-1})$.

The matrix \mathbf{U} that we construct has meta-block structure that allows solves in \mathbf{U} and \mathbf{U}^T to be performed with linear work and depth $O(\log^2 n)$. This results in a parallel algorithm for solving equations in \mathbf{M} in work $O(m \log \epsilon^{-1})$ and depth $O(\log^2 n \log \epsilon^{-1})$. We remark that with some additional work (analogous to Section 7 of [LPS15]), one could reduce this depth to $O(\log n \log \log n \log \epsilon^{-1})$.

The key to constructing \mathbf{U} is realizing that the algorithm JACOBI corresponds to multiplying by the matrix $\mathbf{Z}^{(k)}$ defined in equation (7). Moreover, this matrix is a polynomial in \mathbf{X} and \mathbf{L} of degree $\log(3/\epsilon)$, where $\mathbf{M}_{[F_i, F_i]}^{(i)} = \mathbf{X} + \mathbf{L}$ where \mathbf{L} is bDD, and \mathbf{X} block diagonal such that $\mathbf{X} \succcurlyeq 2\mathbf{L}$. To force \mathbf{Z}^k to be a sparse matrix, we require that \mathbf{L} be sparse.

If we use the algorithm BSDDSUBSET to choose F_i , then \mathbf{L} need not be sparse. However, this problem is easily remedied by forbidding algorithm BSDDSUBSET from choosing any vertex of more than twice average degree. Thus, we can ensure that all $\mathbf{Z}^{(i)}$ are sparse.

Lemma 3.8. *For every bDD matrix \mathbf{M} and every $\alpha \geq 0$, there is a subset F of size at least $\frac{n}{16(1+\alpha)}$ such that $\mathbf{M}_{[F, F]}$ is α -bDD and the number of nonzeros blocks in each block-row of F at most twice the average number of nonzero blocks in a block-row of \mathbf{M} .*

Proof. Discard every block-row of \mathbf{M} that has more than twice the average number of nonzeros blocks per row-block. Then remove the corresponding row blocks. The remaining matrix has dimension at least $n/2$. We can now use Lemma 3.5 to find an α -bDD submatrix. \square

We obtain the $\mathbf{U}^T \mathbf{D} \mathbf{U}$ factorization by applying the inverse of the factorization (1):

$$\mathbf{M} = \begin{bmatrix} \mathbf{I} & 0 \\ \mathbf{M}_{[F, F]}^{-1} \mathbf{M}_{[F, C]} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{M}_{[F, F]} & 0 \\ 0 & \text{Sc}(\mathbf{M}, F) \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{M}_{[C, F]} \mathbf{M}_{[F, F]}^{-1} \\ 0 & \mathbf{I} \end{bmatrix}.$$

In the left and right triangular matrices we replace $\mathbf{M}_{[F, F]}^{-1}$ with the polynomial we obtain from JACOBI. In the middle matrix, it suffices to approximate $\mathbf{M}_{[F, F]}$ by a block diagonal matrix, and $\text{Sc}(\mathbf{M}, F)$ by a factorization of its sparse approximation given by Theorem 3.1. The details, along with a careful setting of ϵ_i , are carried out in Section F.

3.4 Spectral Vertex Sparsification Algorithm

In this section, we outline a procedure APPROXSCHUR that efficiently approximates the Schur complement of a bDD matrix \mathbf{M} w.r.t. a set of indices F s.t. $\mathbf{M}_{[F, F]}$ is α -bDD. The following lemma summarizes the guarantees of APPROXSCHUR.

Lemma 3.9. *Let \mathbf{M} be a bDD matrix with index set V , and m nonzero blocks. Let $F \subseteq V$ be such that $\mathbf{M}_{[F, F]}$ is α -bDD for some $\alpha \geq 4$. The algorithm APPROXSCHUR(\mathbf{M}, F, ϵ), returns a matrix $\widetilde{\mathbf{M}}_{SC}$ s.t.*

1. $\widetilde{\mathbf{M}}_{SC}$ has $O(m(\epsilon^{-1} \log \log \epsilon^{-1})^{O(\log \log \epsilon^{-1})})$ nonzero blocks, and
2. $\widetilde{\mathbf{M}}_{SC} \approx_{\epsilon} \text{Sc}(\mathbf{M}, F)$,

in $O(m(\epsilon^{-1} \log \log \epsilon^{-1})^{O(\log \log \epsilon^{-1})})$ work and $O(\log n(\log \log \epsilon^{-1}))$ depth.

We sketch a proof of the above lemma in this section. A complete proof and pseudocode for APPROXSCHUR are given in Section G.

First consider a very simple special case: where F is a singleton, $F = \{i\}$. Let $C = V \setminus F$ be the remaining indices.

$$Sc(\mathbf{M}, i) = \mathbf{M}_{[C,C]} - \mathbf{M}_{[C,i]} \mathbf{M}_{[i,i]}^{-1} \mathbf{M}_{[i,C]}$$

Thus, if $\mathbf{M}_{[C,i]}$ has k nonzero blocks, $Sc(\mathbf{M}, i)$ could have k^2 additional nonzero blocks compared to $\mathbf{M}_{[C,C]}$, potentially making it dense. If \mathbf{M} were a graph Laplacian, then $\mathbf{M}_{[C,i]} \mathbf{M}_{[i,i]}^{-1} \mathbf{M}_{[i,C]}$ would represent the adjacency matrix of a weighted clique. In Section J, we construct weighted expanders that allow us to ϵ -approximate $Sc(\mathbf{M}, i)$ in this case using $m + O(k\epsilon^{-4})$ edges. In Section G.4, we show how to use such weighted expanders to sparsify $Sc(\mathbf{M}, i)$ when \mathbf{M} is a bDD matrix.

This reduction can also be performed in parallel. If F is such that $\mathbf{M}_{[F,F]}$ is block diagonal, we can approximate $Sc(\mathbf{M}, F)$ by expressing $\mathbf{M}_{[C,i]} \mathbf{M}_{[i,i]}^{-1} \mathbf{M}_{[i,C]}$ as $\sum_{i \in F} \mathbf{M}_{[C,i]} \mathbf{M}_{[i,i]}^{-1} \mathbf{M}_{[i,C]}$, and using $|F|$ weighted expanders. However, $\mathbf{M}_{[F,F]}$ may not be diagonal. Instead, we give a procedure SCHURSQUARE that generates \mathbf{M}' that is better approximated by its diagonal.

Invoking SCHURSQUARE a few times leads a sequence of matrices $\mathbf{M}_{[F,F]}^{(0)}, \mathbf{M}_{[F,F]}^{(1)}, \dots, \mathbf{M}_{[F,F]}^{(i)}$. We will show that $\mathbf{M}_{[F,F]}^{(i)}$ is ϵ -approximated by its diagonal and we call the procedure LASTSTEP to approximate $Sc(\mathbf{M}^{(i)}, F)$. An additional caveat is that replacing $\mathbf{M}_{[F,F]}^{(i)}$ by its diagonal at this step gives errors that are difficult to bound. We discuss the correct approximation below

SCHURSQUARE is based on a *squaring* identity for matrix inverse developed in [PS14]. Given a splitting of $\mathbf{M}_{[F,F]}$ into $\mathbf{D} - \mathbf{A}$, where \mathbf{D} is block-diagonal, and \mathbf{A} has its diagonal blocks as zero, it relies on the fact that the matrix

$$\mathbf{M}_2 = \frac{1}{2} \begin{bmatrix} \mathbf{D} - \mathbf{A}\mathbf{D}^{-1}\mathbf{A} & \mathbf{M}_{[F,C]} + \mathbf{A}\mathbf{D}^{-1}\mathbf{M}_{[F,C]} \\ \mathbf{M}_{[C,F]} + \mathbf{M}_{[C,F]}\mathbf{D}^{-1}\mathbf{A} & 2\mathbf{M}_{[C,C]} - \mathbf{M}_{[C,F]}\mathbf{D}^{-1}\mathbf{M}_{[F,C]} \end{bmatrix} \quad (3)$$

satisfies $Sc(\mathbf{M}, F) = Sc(\mathbf{M}_2, F)$. Furthermore, we can show that if \mathbf{M} is α -bDD, $\mathbf{D} - \mathbf{A}\mathbf{D}^{-1}\mathbf{A}$ is α^2 -bDD, which indicates that the block on $[F, F]$ rapidly approaches being diagonal.

As \mathbf{M}_2 may be dense, we construct a sparse approximation to it. Since \mathbf{D} is diagonal, we can construct sparse approximations to the blocks $(\mathbf{M}_2)_{[F,F]}$ and $(\mathbf{M}_2)_{[C,C]}$ in a manner analogous to the case of diagonal $\mathbf{M}_{[F,F]}$. Similarly, we use bipartite expanders to construct sparse approximations to $(\mathbf{M}_2)_{[C,F]}$ and $(\mathbf{M}_2)_{[F,C]}$.

Our sequence of calls to SCHURSQUARE terminates with $\mathbf{M}_{[F,F]}^{(i)}$ being roughly ϵ^{-1} -bDD. We then return LASTSTEP($\mathbf{M}^{(i)}, F, \epsilon^{-1}, \epsilon$). As mentioned above, we cannot just replace $\mathbf{M}_{[F,F]}^{(i)}$ by its diagonal. Instead, LASTSTEP performs one step of *squaring* similar to SCHURSQUARE with a key difference: Rather than expressing $\mathbf{M}_{[F,F]}^{(i)}$ as $\mathbf{D} - \mathbf{A}$, it expresses it as $\mathbf{X} + \mathbf{L}$, where \mathbf{X} is block-diagonal, and \mathbf{L} is just barely bDD. With this splitting, it constructs $\mathbf{M}^{(last)}$ after performing one iteration similar to Eq. (3). After this step, it replaces the $\mathbf{M}_{[F,F]}^{(last)}$ block with the block-diagonal matrix \mathbf{X} . Again, we directly produce sparse approximations to $\mathbf{M}^{(last)}$ and its Schur complement via weighted (bipartite) expanders. A precise description and proofs are given in Section G.2.

3.5 Sparsifying bDD matrices

The main technical hurdle left to address is how we sparsify bDD matrices. We to do this both to approximate \mathbf{M} by $\mathbf{M}^{(0)}$, if \mathbf{M} is not already sparse, and to ensure that all the matrices $\mathbf{M}^{(i)}$ remain sparse. While the spectral vertex sparsification algorithm described in the previous section allows us to compute an approximation to a Schur complement $Sc(\mathbf{M}^{(i)}, F_i)$, it is sparse only

when $\mathbf{M}^{(i)}$ is already sparse. As we iteratively apply this procedure, the density of the matrices produced will grow unacceptably. We overcome this problem by occasionally applying another sparsification routine that substantially decreases the number of nonzero blocks. The cost of this sparsification routine is that it requires solving systems of equations in sparse *bDD* matrices. We, of course, do this recursively.

Our sparsification procedure begins by generalizing the observation that graph Laplacians can be sparsified by sampling edges with probabilities determined by their effective resistances [SS11b]. There is a block analog of leverage scores (37) that provides probabilities of sampling blocks so that the resulting sampled matrix approximates the original and has $O(n \log n)$ nonzero blocks with high probability. To compute this block analog of leverage scores we employ a recently introduced procedure of Cohen *et. al.* [CLM⁺14].

Once we generalize their results to block matrices, they show that we can obtain sufficiently good estimates of the block leverage scores by computing leverage scores in a *bDD* matrix obtained by randomly subsampling blocks of the original. The block leverage scores in this matrix are obtained by solving a logarithmic number of linear equations in this subsampled matrix. Thus, our sparsification procedure requires constructing a solver for a subsampled matrix and then applying that solver a logarithmic number of times. We compute this solver recursively.

There is a tradeoff between the number of nonzero blocks in the subsampled system and in the resulting approximation of the original matrix. If the original matrix has m nonzero blocks and we subsample to a system of m/K nonzero blocks, then we obtain an ϵ -approximation of the original matrix with $O(K\epsilon^{-2}n \log n)$ nonzero blocks.

The details of the analysis of the undersampling procedure appear in Section H.

3.6 The main algorithm

We now explain how we prove Theorem 1.2. The details supporting this exposition appear in Section I. Our main goal is to control the density of the Schur complement chain as we repeatedly invoke Lemma 3.9.

Starting from some $\mathbf{M}^{(0)}$, we compute sets F_i (via calls to `BDDSUBSET`), approximate solvers $\mathbf{Z}^{(i)}$ (via `JACOBI`), and approximations of Schur complements $\mathbf{M}^{(i)}$ (via `APPROXSCHUR`), until we obtain a matrix $\mathbf{M}^{(i)}$ such that its dimension is a smaller than that of $\mathbf{M}^{(0)}$ by a large constant factor (like 4). While the dimension of $\mathbf{M}^{(i)}$ is much smaller, its number of nonzero blocks is potentially larger by an even larger factor. We use the procedure described in the previous section to sparsify it. This sparsification procedure produces a sparse approximation of the matrix at the cost of solving systems of equations in a subsampled version of that matrix. Some care is required to balance the cost of the resulting recursion.

We now sketch an analysis of a nearly linear time algorithm that results from a simple choice of parameters. We optimize the parameter choice and analysis in Section I. Let n be the dimension of $\mathbf{M}^{(0)}$ and let m be its number of nonzero blocks. To begin, assume that $m \leq n\Delta \log^3 n$, for a Δ to be specified later. We call this the *sparse* case, and address the case of dense \mathbf{M} later.

We consider fixing $\epsilon_i = c/\log n$ for all i , for some constant c . As the depth of the Schur complement chain is $O(\log n)$, this results in a solver with constant accuracy. A constant number of iterations of the procedure described above are required to produce an $\mathbf{M}^{(i)}$ whose dimension is a factor of 4 smaller than $\mathbf{M}^{(0)}$. Lemma 3.9 tells us that the edge density of this $\mathbf{M}^{(i)}$ is potentially higher than that of $\mathbf{M}^{(0)}$ by a factor of

$$O\left(\left(\epsilon^{-1} \log \log \epsilon^{-1}\right)^{O(\log \log \epsilon^{-1})}\right) = \exp(O(\log \log^2 n)) = n^{o(1)}.$$

Set Δ to be this factor. We use the sparsification procedure from the previous section to guarantee

that no matrix in the chain has density higher than that of this matrix, which is upper bounded by $(m/n)\Delta = \Delta^2 \log^3 n$.

Setting $K = 2\Delta$, the subsampling produces a matrix of density half that of $\mathbf{M}^{(0)}$, and it produces a sparse approximation of $\mathbf{M}^{(i)}$ of density $O(K\epsilon_i^{-2} \log n) \leq O(\Delta \log^3 n)$, which, by setting constants appropriately, we can also force to be half that of $\mathbf{M}^{(0)}$. In order to perform the sparsification procedure, we need to construct a Schur complement chain for the subsampled matrix, and then use it to solve $O(\log n)$ systems of linear equations. The cost of using this chain to solve equations in the subsampled system is at most $O(n\Delta^2 \log^4 n)$, and the cost of using the solutions to these equations to sparsify $\mathbf{M}^{(i)}$ is $O(m \log n)$. The cost of the calls to `BDDSUBSET` and `APPROXSCHUR` are proportional to the number of edges in the matrices, which is $O(n\Delta \log^4 n)$.

We repeat this procedure all the way down the chain, only using sparsification when the dimension of $\mathbf{M}^{(i)}$ shrinks by a factor of 4. Since none of the matrices that we generate have density higher than $\Delta \log^3 n$, we remain in the sparse case. Let $T_{sparse}(n)$ be the time required to construct a solver chain on systems of size n with $m \leq n\Delta \log^3 n$. We obtain the following recurrence

$$T_{sparse}(n) \leq 2T_{sparse}(n/4) + n\Delta^2 \log^4 n + m \log n + m\Delta \leq 2T_{sparse}(n/4) + n\Delta^2 \log^4 n + n\Delta \log n + n\Delta^2,$$

which gives

$$T(n)_{sparse} \leq O(n\Delta^2 + n\Delta^2 \log^4 n) \leq n^{1+o(1)}.$$

To handle the case of dense \mathbf{M} , we repeatedly sparsify while keeping n fixed until we obtain a matrix with fewer than $n\Delta \log^3 n$ edges, at which point we switch to the algorithm described above. The running time of this algorithm on a graph with m edges, $T_{dense}(m)$, satisfies the recurrence

$$T_{dense}(m) \leq \begin{cases} T_{sparse}(n) & \text{if } m \leq n\Delta \log^3 n, \text{ and} \\ 2T_{dense}(m/2) + n\Delta^2 \log^4 n + m \log n + m\Delta & \text{otherwise.} \end{cases}$$

Thus $T_{dense}(m)$ is upper bounded bounded by $O(mn^{o(1)} + n^{1+o(1)})$.

We tighten this bound in Section I to prove Theorem 1.2 by carefully choosing the parameters to accompany a sequence ϵ that starts constant and decreases slowly.

4 Summary

We introduce a new approach to solving systems of linear equations that gives the first nearly linear time algorithms for solving systems in connection Laplacians and the first proof that connection Laplacians have linear-sized approximate inverses. This was unknown even for graph Laplacians.

Our algorithms build on ideas introduced in [PS14] and are a break from those used in the previous work on solving systems of equations in graph Laplacians [Vai90, ST14, KMP10, KMP11, KOSZ13, CKM⁺14]. Those algorithms all rest on *support theory* [BGH⁺06], originally introduced by Vaidya [Vai90], and rely on the fact that the Laplacian of one edge is approximated by the Laplacian of a path between its endpoints. No analogous fact is true for connection Laplacians, even those with complex entries for $r = 1$.

Instead, our algorithms rely on many new ideas, the first being that of sparsifying the matrices that appear during elimination. Other critical ideas are finding α -bDD subsets of vertices to eliminate in bulk, approximating Schur complements without computing them explicitly, and the use of sub-sampling to sparsify in a recursive fashion. To efficiently compute approximations of the Schur complements, we introduce a new operation that transforms a matrix into one with the same Schur complement but a much better conditioned upper block (3). To obtain the sharp bounds in Theorem 1.2, we exploit a new linear-time algorithm for constructing linear-sized sparse approximations to implicitly represented weighted cliques whose edge weights are products of weights at vertices (Section J), and extend this to the analog for bDD matrices (Section G.4).

References

- [ABFM14] Boris Alexeev, Afonso S Bandeira, Matthew Fickus, and Dustin G Mixon. Phase retrieval with polarization. *SIAM Journal on Imaging Sciences*, 7(1):35–66, 2014.
- [ANKKS⁺12] Mica Arie-Nachimson, Shahar Z Kovalsky, Ira Kemelmacher-Shlizerman, Amit Singer, and Ronen Basri. Global motion estimation from point matches. In *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2012 Second International Conference on*, pages 81–88. IEEE, 2012.
- [AT11] Haim Avron and Sivan Toledo. Effective stiffness: Generalizing effective resistance sampling to finite element matrices. *arXiv preprint arXiv:1110.4437*, 2011.
- [AW02] Rudolf Ahlswede and Andreas Winter. Strong converse for identification via quantum channels. *Information Theory, IEEE Transactions on*, 48(3):569–579, 2002.
- [BGH⁺06] M. Bern, J. Gilbert, B. Hendrickson, N. Nguyen, and S. Toledo. Support-graph preconditioners. *SIAM J. Matrix Anal. & Appl.*, 27(4):930–951, 2006.
- [BSS13] Afonso S Bandeira, Amit Singer, and Daniel A Spielman. A cheeger inequality for the graph connection laplacian. *SIAM Journal on Matrix Analysis and Applications*, 34(4):1611–1630, 2013.
- [CKM⁺14] Michael B. Cohen, Rasmus Kyng, Gary L. Miller, Jakub W. Pachocki, Richard Peng, Anup B. Rao, and Shen Chen Xu. Solving sdd linear systems in nearly $m \log 1/2n$ time. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC '14*, pages 343–352, New York, NY, USA, 2014. ACM.
- [CLM⁺14] Michael B Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, and Aaron Sidford. Uniform sampling for matrix approximation. *arXiv preprint arXiv:1408.5099*, 2014.
- [dCSHS11] Marcel K. de Carli Silva, Nicholas J. A. Harvey, and Cristiane M. Sato. Sparse sums of positive semidefinite matrices. *CoRR*, abs/1107.0088, 2011.
- [KFS13] Dilip Krishnan, Raanan Fattal, and Richard Szeliski. Efficient preconditioning of laplacian matrices for computer graphics. *ACM Transactions on Graphics (TOG)*, 32(4):142, 2013.
- [KMP10] I. Koutis, G.L. Miller, and R. Peng. Approaching optimality for solving SDD linear systems. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 235–244, 2010.
- [KMP11] I. Koutis, G.L. Miller, and R. Peng. A nearly- $m \log n$ time solver for SDD linear systems. In *Foundations of Computer Science (FOCS), 2011 52nd Annual IEEE Symposium on*, pages 590–598, 2011.
- [KOSZ13] Jonathan A Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. A simple, combinatorial algorithm for solving sdd systems in nearly-linear time. In *Proceedings of the 45th annual ACM symposium on Symposium on theory of computing*, pages 911–920. ACM, 2013.

- [LMP13] Mu Li, Gary L Miller, and Rongkun Peng. Iterative row sampling. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 127–136. IEEE, 2013.
- [LPS88] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.
- [LPS15] Yin Tat Lee, Richard Peng, and Daniel A. Spielman. Sparsified cholesky solvers for SDD linear systems. *CoRR*, abs/1506.08204, 2015.
- [Mar88] G. A. Margulis. Explicit group theoretical constructions of combinatorial schemes and their application to the design of expanders and concentrators. *Problems of Information Transmission*, 24(1):39–46, July 1988.
- [MP13] Gary L. Miller and Richard Peng. Approximate maximum flow on separable undirected graphs. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1151–1170. SIAM, 2013.
- [MSS15] Adam W Marcus, Nikhil Srivastava, and Daniel A Spielman. Interlacing families IV: Bipartite Ramanujan graphs of all sizes. *arXiv preprint arXiv:1505.08010*, 2015. to appear in FOCS 2015.
- [MTW14] Stefano Marchesini, Yu-Chao Tu, and Hau-tieng Wu. Alternating projection, ptychographic imaging and phase synchronization. *arXiv preprint arXiv:1402.0550*, 2014.
- [MV77] J. A. Meijerink and H. A. van der Vorst. An iterative solution method for linear systems of which the coefficient matrix is a symmetric m -matrix. *Mathematics of Computation*, 31(137):148–162, 1977.
- [OSB15] Onur Ozyesil, Amit Singer, and Ronen Basri. Stable camera motion estimation using convex programming. *SIAM Journal on Imaging Sciences*, 8(2):1220–1262, 2015.
- [PS14] Richard Peng and Daniel A. Spielman. An efficient parallel solver for SDD linear systems. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 333–342, 2014.
- [SS11a] Amit Singer and Yoel Shkolnisky. Three-dimensional structure determination from common lines in cryo-em by eigenvectors and semidefinite programming. *SIAM journal on imaging sciences*, 4(2):543–572, 2011.
- [SS11b] D. Spielman and N. Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.
- [SS11c] Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.
- [SS12] Yoel Shkolnisky and Amit Singer. Viewing direction estimation in cryo-em using synchronization. *SIAM journal on imaging sciences*, 5(3):1088–1110, 2012.
- [ST14] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM. J. Matrix Anal. & Appl.*, 35:835885, 2014.

- [SW12] Amit Singer and H-T Wu. Vector diffusion maps and the connection laplacian. *Communications on pure and applied mathematics*, 65(8), 2012.
- [Tch36] Nikolai Tchudakoff. On the difference between two neighbouring prime numbers. *Rec. Math. [Mat. Sbornik] N.S.*, 1(6):799–814, 1936.
- [TOS00] Ulrich Trottenberg, Cornelius W Oosterlee, and Anton Schuller. *Multigrid*. Academic press, 2000.
- [Tro12] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.
- [Vai90] Pravin M. Vaidya. Solving linear equations with symmetric diagonally dominant matrices by constructing good preconditioners. Unpublished manuscript UIUC 1990. A talk based on the manuscript was presented at the IMA Workshop on Graph Theory and Sparse Matrix Computation, October 1991, Minneapolis., 1990.
- [ZS14] Zhizhen Zhao and Amit Singer. Rotationally invariant image representation for viewing direction classification in cryo-em. *Journal of structural biology*, 186(1):153–166, 2014.

Contents

1	Introduction	1
1.1	Connection Laplacians and Block DD Matrices	2
2	Background	3
3	Overview of the algorithms	4
3.1	Schur Complement Chains	5
3.2	Choosing F_i : α -bDD matrices	5
3.3	Linear Sized Approximate Inverses	6
3.4	Spectral Vertex Sparsification Algorithm	7
3.5	Sparsifying bDD matrices	8
3.6	The main algorithm	9
4	Summary	10
A	Background	15
B	Block Diagonally Dominant Matrices	16
C	Schur Complement Chains	18
D	Finding α-bDD Subsets	21
E	Jacobi Iteration on α-bDD Matrices	23
F	Existence of Linear-Sized Approximate Inverses	25
G	Spectral Vertex Sparsification Algorithm	27
G.1	Iterative Squaring and Sparsification	29
G.2	Schur Complement w.r.t. Highly α -bDD Submatrices	32
G.3	Deferred Proofs from Section G.2	36
G.4	Sparsifying Product Demand Block-Laplacians	39
G.5	Constructing sparsifiers for lifts of graphs	42
H	Estimating Leverage Scores by Undersampling	42
I	Recursive Construction of Schur Complement Chains	45
J	Weighted Expander Constructions	49
J.1	Sparsifying Complete Product Demand Graphs	53
J.2	Product demand graphs with most weights maximal	56
J.3	Weighted Bipartite Expanders	58

A Background

Proof. (of Claim 2.4) Note that

$$\|\mathbf{x} - \mathbf{Z}\mathbf{b}\|_M^2 = \mathbf{x}^* \mathbf{M} \mathbf{x} - 2\mathbf{x}^* \mathbf{M} \mathbf{Z} \mathbf{M} \mathbf{x} + \mathbf{x}^* \mathbf{M} \mathbf{Z} \mathbf{M} \mathbf{Z} \mathbf{M} \mathbf{x}$$

Since all nonzero eigenvalues of \mathbf{M} are at least μ , the eigenvalues of $\mathbf{M}^{1/2}(\mathbf{M} + \epsilon\mu\mathbf{I})^{-1}\mathbf{M}^{1/2}$ lie between $1/(1 + \epsilon)$ and 1. Using $\mathbf{Z} \approx_\epsilon (\mathbf{M} + \epsilon\mu\mathbf{I})^{-1}$, we see that the eigenvalues of $\mathbf{M}^{1/2}\mathbf{Z}\mathbf{M}^{1/2}$ lie between $e^{-2\epsilon}$ and e^ϵ . Using $0 < \epsilon < 1/2$, we have

$$\|\mathbf{x} - \mathbf{Z}\mathbf{b}\|_M^2 \leq (1 - 2e^{-2\epsilon} + e^{2\epsilon})\mathbf{x}^* \mathbf{M} \mathbf{x} \leq 6\epsilon \|\mathbf{x}\|_M^2.$$

□

Claim A.1. Let \mathbf{A} be a matrix of condition number κ and let $\mathbf{A} \approx_\epsilon \mathbf{B}$ for $\epsilon \leq (56\kappa^3)^{-1}$. Then, $\mathbf{A}^3 \approx_{28\kappa^3\epsilon} \mathbf{B}^3$.

Proof. First, observe that $\mathbf{A} \approx_\epsilon \mathbf{B}$ implies that $\|\mathbf{B}\| \leq (1 + \epsilon^\epsilon)\|\mathbf{A}\| \leq 2\|\mathbf{A}\|$. It also implies that $\|\mathbf{A} - \mathbf{B}\| \leq 2\epsilon\|\mathbf{A}\|$. As

$$\mathbf{A}^2 - \mathbf{B}^2 = \frac{1}{2}(\mathbf{A} - \mathbf{B})(\mathbf{A} + \mathbf{B}) + \frac{1}{2}(\mathbf{A} + \mathbf{B})(\mathbf{A} - \mathbf{B}),$$

$\|\mathbf{A}^2 - \mathbf{B}^2\| \leq 6\epsilon\|\mathbf{A}\|^2$. Similarly, as

$$\mathbf{A}^3 - \mathbf{B}^3 = \frac{1}{2}(\mathbf{A} - \mathbf{B})(\mathbf{A}^2 + \mathbf{B}^2) + \frac{1}{2}(\mathbf{A} + \mathbf{B})(\mathbf{A}^2 - \mathbf{B}^2),$$

$\|\mathbf{A}^3 - \mathbf{B}^3\| \leq 28\epsilon\|\mathbf{A}\|^3$.

Let $\kappa = \|\mathbf{A}\|/\lambda_{\min}(\mathbf{A})$. The above relation implies that

$$\|\mathbf{A}^3 - \mathbf{B}^3\| \leq 28\epsilon\kappa^3\lambda_{\min}(\mathbf{A})^3 = 28\epsilon\kappa^3\lambda_{\min}(\mathbf{A}^3).$$

Thus, as $28\epsilon\kappa^3 \leq 1/2$,

$$\mathbf{A}^3 \approx_{56\epsilon\kappa^3} \mathbf{B}^3.$$

□

Proof. (of Claim 2.5)

By Claim A.1, using $\delta = \frac{\epsilon}{56\kappa^3}$,

$$\mathbf{M} \mathbf{Z}^3 \mathbf{M} \approx_\epsilon \mathbf{M}(\mathbf{M} + \epsilon\mu\mathbf{I})^{-3} \mathbf{M}.$$

\mathbf{M} has an eigendecomposition in the same basis as $(\mathbf{M} + \epsilon\mu\mathbf{I})^{-3}$, and so it follows that $\mathbf{M}(\mathbf{M} + \epsilon\mu\mathbf{I})^{-3}\mathbf{M}$ has the same eigenbasis, and the same null space as \mathbf{M} .

When λ^{-1} is the eigenvalue of \mathbf{M}^+ of an eigenvector \mathbf{v} , the corresponding eigenvalue of $\mathbf{M}(\mathbf{M} + \epsilon\mu\mathbf{I})^{-3}\mathbf{M}$ is $\beta \stackrel{\text{def}}{=} \frac{\lambda^2}{(\lambda + \epsilon\mu)^3}$ and

$$\lambda^{-1} > \beta \geq \frac{\lambda^2}{(1 + \epsilon)^3 \lambda^3} = e^{-3\epsilon} \lambda^{-1}.$$

So $\mathbf{M}(\mathbf{M} + \epsilon\mu\mathbf{I})^{-3}\mathbf{M} \approx_{3\epsilon} \mathbf{M}^+$, and $\mathbf{M} \mathbf{Z}^3 \mathbf{M} \approx_{4\epsilon} \mathbf{M}^+$. □

Fact A.2. For every $\mathbf{d} \in \mathbb{C}^{r \times r}$, we can find $\mathbf{Q}^{(1)}, \mathbf{Q}^{(2)} \in \mathbb{C}^{r \times r}$ where

$$(\mathbf{Q}^{(1)})(\mathbf{Q}^{(1)})^* = (\mathbf{Q}^{(1)})^*(\mathbf{Q}^{(1)}) = (\mathbf{Q}^{(2)})(\mathbf{Q}^{(2)})^* = (\mathbf{Q}^{(2)})^*(\mathbf{Q}^{(2)}) = \mathbb{I}_r,$$

such that

$$\mathbf{d} = \frac{1}{2} \|\mathbf{d}\| (\mathbf{Q}^{(1)} + \mathbf{Q}^{(2)}).$$

Proof. Let $\hat{\mathbf{d}} = \frac{1}{\|\mathbf{d}\|}\mathbf{d}$. Thus, $\|\hat{\mathbf{d}}\| = 1$. Write $\hat{\mathbf{d}}$ using its singular value decomposition as $\mathbf{U}\mathbf{D}\mathbf{V}^*$, where $\mathbf{D}, \mathbf{U}, \mathbf{V} \in \mathbb{C}^{r \times r}$, \mathbf{D} is a real diagonal matrix with the singular values of \mathbf{d} on the diagonal, and $\mathbf{U}\mathbf{U}^* = \mathbf{U}^*\mathbf{U} = \mathbf{V}\mathbf{V}^* = \mathbf{V}^*\mathbf{V} = \mathbb{I}_r$. Since $\|\hat{\mathbf{d}}\| = 1$, we have $D_{j,j} \in [0, 1]$ for all $j \in [r]$. Thus, there exists a real θ_j such that $\cos \theta_j = D_{j,j}$. If we let $\mathbf{D}^{(1)}, \mathbf{D}^{(2)}$ be diagonal matrices such that $D_{j,j}^{(1)} = \exp(i\theta_j), D_{j,j}^{(2)} = \exp(-i\theta_j)$, we have $\mathbf{D} = \frac{1}{2}(\mathbf{D}^{(1)} + \mathbf{D}^{(2)})$. Moreover,

$$(\mathbf{D}^{(1)})(\mathbf{D}^{(1)})^* = (\mathbf{D}^{(1)})^*(\mathbf{D}^{(1)}) = (\mathbf{D}^{(2)})(\mathbf{D}^{(2)})^* = (\mathbf{D}^{(2)})^*(\mathbf{D}^{(2)}) = \mathbb{I}_r.$$

Letting $\mathbf{Q}^{(k)} = \mathbf{U}\mathbf{D}^{(k)}\mathbf{V}^*$ for $k = 1, 2$, we get $\hat{\mathbf{d}} = \frac{1}{2}(\mathbf{Q}^{(1)} + \mathbf{Q}^{(2)})$ and hence $\mathbf{d} = \frac{1}{2}\|\mathbf{d}\|(\mathbf{Q}^{(1)} + \mathbf{Q}^{(2)})$. Moreover,

$$(\mathbf{Q}^{(1)})(\mathbf{Q}^{(1)})^* = (\mathbf{Q}^{(1)})^*(\mathbf{Q}^{(1)}) = (\mathbf{Q}^{(2)})(\mathbf{Q}^{(2)})^* = (\mathbf{Q}^{(2)})^*(\mathbf{Q}^{(2)}) = \mathbb{I}_r.$$

□

Fact A.3 (Lemma B.1. from [MP13]). *If \mathbf{M} and $\widetilde{\mathbf{M}}$ are positive semidefinite matrices satisfying $\mathbf{M} \preceq \widetilde{\mathbf{M}}$, then*

$$Sc(\mathbf{M}, F) \preceq Sc(\widetilde{\mathbf{M}}, F).$$

This fact can be proven via an energy minimization definition of Schur complement. More details on this formulation can be found in [MP13].

B Block Diagonally Dominant Matrices

In this section, we prove a few basic facts about bDD matrices. The following lemma gives an equivalent definition of bDD matrices.

Lemma B.1. *A Hermitian block-matrix $\mathbf{M} \in (\mathbb{C}^{r \times r})^{n \times n}$ is bDD iff it can be written as $\mathbf{D} - \mathbf{A}$ where \mathbf{D} is block-diagonal, \mathbf{A} is Hermitian, and $D_{[i,i]} \succeq \mathbb{I}_r \sum_j \|\mathbf{A}_{[i,j]}\|$, for all $i \in V$.*

Proof. The only if direction is easy. For bDD matrix \mathbf{M} , if we let \mathbf{D} be the block diagonal matrix such that $D_{[i,i]} = \mathbf{M}_{[i,i]}$, and \mathbf{A} be the matrix such that

$$\mathbf{A}_{[i,j]} = \begin{cases} 0 & i = j \\ \mathbf{M}_{[i,j]} & i \neq j \end{cases}$$

It is immediate that \mathbf{A} is Hermitian. Moreover, \mathbf{M} is bDD implies that $D_{[i,i]} \succeq \mathbb{I}_r \sum_j \|\mathbf{A}_{[i,j]}\|$.

For the if direction. Suppose we have \mathbf{D}, \mathbf{A} such that \mathbf{D} is block-diagonal, and for all i $D_{[i,i]} \succeq \mathbb{I}_r \sum_j \|\mathbf{A}_{[i,j]}\|$. Thus, letting $\mathbf{M} = \mathbf{D} - \mathbf{A}$, we have, for all $i \in V$,

$$\mathbb{I}_r \cdot \sum_{j \neq i} \|\mathbf{M}_{[i,j]}\| = \mathbb{I}_r \cdot \sum_{j \neq i} \|\mathbf{A}_{[i,j]}\| \preceq D_{[i,i]} - \mathbb{I}_r \cdot \|\mathbf{A}_{[i,i]}\| \preceq D_{[i,i]} - \mathbf{A}_{[i,i]} = \mathbf{M}_{[i,i]},$$

where we used that since \mathbf{A}, \mathbf{M} are Hermitian, \mathbf{D} is also Hermitian, and thus $D_{[i,i]}, \mathbf{A}_{[i,i]}$ are Hermitian. □

This immediately implies the corollary that flipping the sign of off-diagonal blocks preserves bDD-ness.

Corollary B.2. *Given a bDD matrix \mathbf{M} , write it as $\mathbf{D} - \mathbf{A}$, where \mathbf{D} is a block-diagonal, and \mathbf{A} has its diagonal blocks as zero. Then, $\mathbf{D} + \mathbf{A}$ is also PSD.*

Proof. First observe that for all i , $(\mathbf{D} + \mathbf{A})_{[i,i]} = (\mathbf{D} - \mathbf{A})_{[i,i]} = \mathbf{M}_{[i,i]}$, i.e., their diagonal blocks are identical. Moreover, for all $i \neq j$, we have $\|(\mathbf{D} + \mathbf{A})_{[i,j]}\| = \|(\mathbf{D} - \mathbf{A})_{[i,j]}\| = \|\mathbf{M}_{[i,j]}\|$. Thus $\mathbf{D} + \mathbf{A}$ is also bDD, and hence PSD. \square

Next, we show that the class of bDD matrices is closed under Schur complement.

Lemma B.3. *The class of bDD matrices is closed under Schur complement.*

Proof. Since Schur complementation does not depend on the order of indices eliminated, it suffices to prove that for any bDD matrix $\mathbf{M} \in (\mathbb{C}^{r \times r})^{n \times n}$, $Sc(\mathbf{M}, 1)$ is a bDD matrix. Let $C = V \setminus \{1\}$.

We have $Sc(\mathbf{M}, 1) = \mathbf{M}_{[C,C]} - \mathbf{M}_{[C,1]}\mathbf{M}_{[1,1]}^{-1}\mathbf{M}_{[1,C]}$. Let $\mathbf{D} \in (\mathbb{C}^{r \times r})^{C \times C}$ be the block diagonal matrix such that $\mathbf{D}_{[i,i]} = \mathbf{M}_{[i,i]}$ for $i \in C$. Expressing $Sc(\mathbf{M}, 1)$ as $\mathbf{D} + (\mathbf{M}_{[C,C]} - \mathbf{D} + \mathbf{M}_{[C,1]}\mathbf{M}_{[1,1]}^{-1}\mathbf{M}_{[1,C]})$, we have for any $i \in C$,

$$\begin{aligned} \sum_{j \in C} \left\| (\mathbf{M}_{[C,C]} - \mathbf{D} + \mathbf{M}_{[C,1]}\mathbf{M}_{[1,1]}^{-1}\mathbf{M}_{[1,C]})_{[i,j]} \right\| &\leq \left(\sum_{j \in C: j \neq i} \|\mathbf{M}_{[i,j]}\| \right) + \sum_{j \in C} \left\| \mathbf{M}_{[i,1]}\mathbf{M}_{[1,1]}^{-1}\mathbf{M}_{[1,j]} \right\| \\ &\leq \left(\sum_{j \in C: j \neq i} \|\mathbf{M}_{[i,j]}\| \right) + \|\mathbf{M}_{[i,1]}\| \left\| \mathbf{M}_{[1,1]}^{-1} \right\| \sum_{j \in C} \|\mathbf{M}_{[1,j]}\| \\ &\leq \left(\sum_{j \in C: j \neq i} \|\mathbf{M}_{[i,j]}\| \right) + \|\mathbf{M}_{[i,1]}\| = \sum_{j \in V: j \neq i} \|\mathbf{M}_{[i,j]}\|, \end{aligned}$$

where the last inequality uses $\left\| \mathbf{M}_{[1,1]}^{-1} \right\| \left(\sum_{j \neq 1} \|\mathbf{M}_{[1,j]}\| \right) \leq 1$, since $\mathbf{M}_{[1,1]} \succcurlyeq \mathbb{I}_r \cdot \sum_{j \neq 1} \|\mathbf{M}_{[1,j]}\|$. Thus using Lemma B.1, we have $Sc(\mathbf{M}, 1) = \mathbf{D} - (-\mathbf{M}_{[C,C]} - \mathbf{D} + \mathbf{M}_{[C,1]}\mathbf{M}_{[1,1]}^{-1}\mathbf{M}_{[1,C]})$ is bDD. \square

The next definition describes a special form that we can express any bDD matrix in, which will occasionally be useful.

Definition B.4. A matrix $\mathbf{B} \in (\mathbb{C}^{r \times r})^{n \times m}$ is called a *unitary edge-vertex transfer matrix*, when each block column of $\mathbf{B}_{[e]}$ has exactly two nonzero blocks $\mathbf{U}_e, \mathbf{V}_e \in \mathbb{C}^{r \times r}$ s.t. $\mathbf{U}_e \mathbf{U}_e^* = \mathbf{V}_e \mathbf{V}_e^* = w_e \mathbb{I}_r$ where $w_e \geq 0$.

Lemma B.5. *Every bDD matrix $\mathbf{M} \in (\mathbb{C}^{r \times r})^{n \times n}$ with m nonzero off-diagonal blocks can be written as $\mathbf{X} + \mathbf{B}\mathbf{B}^*$ where $\mathbf{B} \in (\mathbb{C}^{r \times r})^{n \times 2m}$ is a unitary edge-vertex transfer matrix and \mathbf{X} is a block diagonal PSD matrix. This implies that every bDD matrix is PSD. Furthermore, for every block diagonal matrix \mathbf{Y} s.t. $\mathbf{M} - \mathbf{Y}$ is bDD, we have $\mathbf{X} \succcurlyeq \mathbf{Y}$. This decomposition can be found in $O(m)$ time and $O(\log n)$ depth.*

Proof. Consider a pair $\{i, j\} \in V \times V$ such that $i \neq j$, and $\mathbf{M}_{[i,j]} \neq 0$. Using Fact A.2, we can write such a $\mathbf{M}_{[i,j]}$ as $\frac{1}{2} \|\mathbf{M}_{[i,j]}\| (\mathbf{Q}_{\{i,j\}}^{(1)} + \mathbf{Q}_{\{i,j\}}^{(2)})$, where $\mathbf{Q}_{\{i,j\}}^{(1)} (\mathbf{Q}_{\{i,j\}}^{(1)})^* = \mathbf{Q}_{\{i,j\}}^{(2)} (\mathbf{Q}_{\{i,j\}}^{(2)})^* = \mathbb{I}_r$. we construct two vectors vector $\mathbf{B}_{\{i,j\}}^{(1)}, \mathbf{B}_{\{i,j\}}^{(2)} \in (\mathbb{C}^{r \times r})^n$ such that for $k = 1, 2$, $\left(\mathbf{B}_{\{i,j\}}^{(k)} \right)_{[i]} = \frac{1}{\sqrt{2}} \|\mathbf{M}_{[i,j]}\|^{1/2} \mathbb{I}_r$, $\left(\mathbf{B}_{\{i,j\}}^{(k)} \right)_{[j]} = \frac{1}{\sqrt{2}} \|\mathbf{M}_{[i,j]}\|^{1/2} \left(\mathbf{Q}_{\{i,j\}}^{(k)} \right)^*$, and all other blocks are zero. We can

verify that for all $k, \ell \in V$,

$$\left(\mathbf{B}_{\{i,j\}}^{(1)} (\mathbf{B}_{\{i,j\}}^{(1)})^* + \mathbf{B}_{\{i,j\}}^{(2)} (\mathbf{B}_{\{i,j\}}^{(2)})^* \right)_{[k,\ell]} = \begin{cases} \|\mathbf{M}_{[i,j]}\| \mathbb{I}_r & k = \ell = i, \\ \|\mathbf{M}_{[i,j]}\| \mathbb{I}_r & k = \ell = j, \\ \mathbf{M}_{[i,j]} & k = i, \ell = j, \\ \mathbf{M}_{[i,j]}^* = \mathbf{M}_{[j,i]} & k = j, \ell = i, \\ 0 & \text{otherwise.} \end{cases}$$

We let $\mathbf{X} = \mathbf{M} - \sum_{\{i,j\}:i \neq j} \left(\mathbf{B}_{\{i,j\}}^{(1)} (\mathbf{B}_{\{i,j\}}^{(1)})^* + \mathbf{B}_{\{i,j\}}^{(2)} (\mathbf{B}_{\{i,j\}}^{(2)})^* \right)$, which must be block-diagonal. We now show that for all $i \in V$, the block $\mathbf{X}_{[i,i]}$ is PSD. We have for all $i \in V$,

$$\mathbf{X}_{[i,i]} = \mathbf{M}_{[i,i]} - \mathbb{I}_r \cdot \sum_{j:j \neq i} \|\mathbf{M}_{[i,j]}\| \succcurlyeq 0,$$

where the last inequality holds since \mathbf{M} is bDD.

Thus, if we define $\mathbf{B} \in (\mathbb{C}^{r \times r})^{n \times 2m}$ such that its columns are all the vectors $\mathbf{B}_{\{i,j\}}^{(1)}, \mathbf{B}_{\{i,j\}}^{(2)}$ defined above, we have $\mathbf{M} = \mathbf{X} + \mathbf{B}\mathbf{B}^*$, and every column of \mathbf{B} has exactly 2 nonzero blocks.

To show that for every block diagonal \mathbf{Y} s.t. $\mathbf{M} - \mathbf{Y}$ is bDD, $\mathbf{X} \succcurlyeq \mathbf{Y}$, first consider applying the decomposition described above to $\mathbf{M} - \mathbf{Y}$ instead of \mathbf{M} . Since the construction of \mathbf{B} only depends on the off-diagonal blocks, we get $\mathbf{M} - \mathbf{Y} = \mathbf{\Lambda} + \mathbf{B}\mathbf{B}^*$, where $\mathbf{\Lambda}$ is block diagonal and PSD. So, $\mathbf{X} - \mathbf{Y} = \mathbf{M} - \mathbf{B}\mathbf{B}^* - \mathbf{Y} = \mathbf{\Lambda} \succcurlyeq 0$.

It is immediate that the decomposition can be found in $O(m)$ time and $O(\log n)$ depth. \square

C Schur Complement Chains

In this section, we give a proof of Lemma 3.3. We restate the lemma here for convenience.

Lemma 3.3. *Consider an ϵ -SCC for $\mathbf{M}^{(0)}$ where $\mathbf{M}^{(i)}$ and $\mathbf{Z}^{(i)}$ can be applied to a vector in work $W_{\mathbf{M}^{(i)}}, W_{\mathbf{D}^{(i)}}$ and depth $D_{\mathbf{M}^{(i)}}, D_{\mathbf{Z}^{(i)}}$ respectively.*

The algorithm $\text{APPLYCHAIN} \left((\mathbf{M}^{(1)}, \mathbf{Z}^{(1)}), \dots, (\mathbf{M}^{(d)}, \mathbf{Z}^{(d)}); F_1, \dots, F_d; \mathbf{b} \right)$ corresponds to a linear operator \mathbf{W} acting on \mathbf{b} such that

1. $\mathbf{W}^{-1} \approx_{\sum_{i=1}^d 2\epsilon_i} \mathbf{M}^{(0)}$, and
2. for any vector \mathbf{b} , $\text{APPLYCHAIN} \left((\mathbf{M}^{(1)}, \mathbf{Z}^{(1)}), \dots, (\mathbf{M}^{(d)}, \mathbf{Z}^{(d)}); F_1, \dots, F_d; \mathbf{b} \right)$ runs in $O \left(\sum_{i=1}^d (D_{\mathbf{M}^{(i)}} + D_{\mathbf{Z}^{(i)}}) \right)$ depth and $O \left(\sum_{i=1}^d (W_{\mathbf{M}^{(i)}} + W_{\mathbf{Z}^{(i)}}) \right)$ work.

The pseudocode for procedure APPLYCHAIN that uses an ϵ -vertex sparsifier chain to approximately solve a system of equations in $\mathbf{M}^{(0)}$ is given in Figure 1.

Proof. We begin by observing that the output vector $\mathbf{x}^{(0)}$ is a linear transformation of the input vector $\mathbf{b}^{(0)}$. Let $\mathbf{W}^{(0)}$ be the matrix that realizes this transformation. Similarly, for $1 \leq i \leq d$, define $\mathbf{W}^{(i)}$ to be the matrix so that

$$\mathbf{x}^{(i)} = \mathbf{W}^{(i)} \mathbf{b}^{(i)}.$$

An examination of the algorithm reveals that

$$\mathbf{W}^{(d)} = \left(\mathbf{M}^{(d)} \right)^{-1}, \tag{4}$$

$$\mathbf{x}^{(0)} = \text{APPLYCHAIN} \left((\mathbf{M}^{(1)}, \mathbf{Z}^{(1)}), \dots, (\mathbf{M}^{(d)}, \mathbf{Z}^{(d)}); F_1, \dots, F_d \right)$$

1. Initialize $\mathbf{b}^{(0)} \leftarrow \mathbf{b}$.
2. For $i = 1, \dots, d$
 - (a) $\mathbf{x}_{[F_i]}^{(i-1)} \leftarrow \mathbf{Z}^{(i)} \mathbf{b}_{[F_i]}^{(i-1)}$,
 - (b) $\mathbf{b}^{(i)} \leftarrow \mathbf{b}_{[C_i]}^{(i-1)} - \mathbf{M}_{[C_i, F_i]}^{(i-1)} \mathbf{x}_{[F_i]}^{(i-1)}$.
3. $\mathbf{x}^{(d)} \leftarrow (\mathbf{M}^{(d)})^{-1} \mathbf{b}^{(d)}$.
4. For $i = d, \dots, 1$
 - (a) $\mathbf{x}_{[C_i]}^{(i-1)} \leftarrow \mathbf{x}^{(i)}$.
 - (b) $\mathbf{x}_{[F_i]}^{(i-1)} \leftarrow \mathbf{x}_{[F_i]}^{(i-1)} - \mathbf{Z}^{(i)} \mathbf{M}_{[F_i, C_i]}^{(i-1)} \mathbf{x}^{(i)}$.

Figure 1: Solver Algorithm using Vertex Sparsifier Chain

and for $1 \leq i \leq d$,

$$\mathbf{W}^{(i-1)} = \begin{bmatrix} \mathbf{I} & -\mathbf{Z}^{(i)} \mathbf{M}_{[F_i, C_i]}^{(i-1)} \\ 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{Z}^{(i)} & 0 \\ 0 & \mathbf{W}^{(i)} \end{bmatrix} \begin{bmatrix} \mathbf{I} & 0 \\ -\mathbf{M}_{[C_i, F_i]}^{(i-1)} \mathbf{Z}^{(i)} & \mathbf{I} \end{bmatrix}. \quad (5)$$

We will now prove by backwards induction on i that

$$\left(\mathbf{W}^{(i)} \right)^{-1} \approx_{\sum_{j=i+1}^d 2\epsilon_j} \mathbf{M}^{(i)}.$$

The base case of $i = d$ follows from (4). Using the definition of an ϵ -SCC, we know that $0 \preceq (\mathbf{Z}^{(d)})^{-1} - \mathbf{M}_{[F_d, F_d]}^{(d-1)} \preceq \epsilon_d \cdot \text{Sc} \left(\mathbf{M}^{(d-1)}, C_d \right)$. We show in Lemma C.1 that this implies

$$\begin{bmatrix} \mathbf{I} & -\mathbf{Z}^{(d)} \mathbf{M}_{[F_d, C_d]}^{(d-1)} \\ 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{Z}^{(d)} & 0 \\ 0 & \text{Sc} \left(\mathbf{M}^{(d-1)}, F_d \right)^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & 0 \\ -\mathbf{M}_{[C_d, F_d]}^{(d-1)} \mathbf{Z}^{(d)} & \mathbf{I} \end{bmatrix} \approx_{\epsilon_d} \left(\mathbf{M}^{(d-1)} \right)^{-1}.$$

As $\mathbf{M}^{(i)} \approx_{\epsilon_i} \text{Sc} \left(\mathbf{M}^{(i-1)}, F_i \right)$,

$$\begin{bmatrix} \mathbf{I} & -\mathbf{Z}^{(i)} \mathbf{M}_{[F_i, C_i]}^{(i-1)} \\ 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{Z}^{(i)} & 0 \\ 0 & \left(\mathbf{M}^{(i-1)} \right)^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & 0 \\ -\mathbf{M}_{[C_i, F_i]}^{(i-1)} \mathbf{Z}^{(i)} & \mathbf{I} \end{bmatrix} \approx_{2\epsilon_i} \left(\mathbf{M}^{(i-1)} \right)^{-1}.$$

By combining this identity with (5) and our inductive hypothesis, we obtain

$$\mathbf{W}^{(i-1)} \approx_{\sum_{j=i}^d 2\epsilon_j} \left(\mathbf{M}^{(i-1)} \right)^{-1}.$$

Thus, by induction, we obtain

$$\mathbf{W}^{(0)} \approx_{\sum_{j=1}^d 2\epsilon_j} \left(\mathbf{M}^{(0)} \right)^{-1}.$$

The whole algorithm involves a constant number of applications of $\mathbf{Z}^{(i)}$, $\mathbf{M}_{[F_i, C_i]}^{(i-1)}$, and $\mathbf{M}_{[C_i, F_i]}^{(i-1)}$. We observe that in order to compute $\mathbf{M}_{[F_i, C_i]}^{(i-1)} x_{C_i}$, using a multiplication procedure for $\mathbf{M}^{(i-1)}$, we can pad x_{C_i} with zeros, multiply by $\mathbf{M}^{(i-1)}$, and read off the answer on the indices in C_i . Similarly, we can multiply vectors with $\mathbf{M}_{[C_i, F_i]}^{(i-1)}$. This immediately gives the claimed bounds on work and depth. \square

We now prove the deferred claims from the above proof.

Lemma C.1. *Let $\mathbf{M} \in (\mathbb{C}^{r \times r})^{n \times n}$ be a bDD matrix, $F \subseteq V$ be a subset of the indices, and $\mathbf{Z} \in (\mathbb{C}^{r \times r})^{|F| \times |F|}$ be an hermitian operator satisfying $0 \preceq \mathbf{Z}^{-1} - \mathbf{M}_{[F, F]} \preceq \epsilon \cdot \text{Sc}(\mathbf{M}, C)$. Then,*

$$\begin{bmatrix} \mathbf{I} & -\mathbf{Z}\mathbf{M}_{[F, C]} \\ 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{Z} & 0 \\ 0 & \text{Sc}(\mathbf{M}, F)^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & 0 \\ -\mathbf{M}_{[C, F]}\mathbf{Z} & \mathbf{I} \end{bmatrix} \approx_{\epsilon} \mathbf{M}^{-1}.$$

Proof. Define

$$\widehat{\mathbf{M}} = \begin{bmatrix} (\mathbf{Z})^{-1} & \mathbf{M}_{[F, C]} \\ \mathbf{M}_{[C, F]} & \mathbf{M}_{[C, C]} \end{bmatrix}.$$

Using Lemma C.2, we know that the assumption on \mathbf{Z} is equivalent to

$$\mathbf{M} \preceq \widehat{\mathbf{M}} \preceq (1 + \epsilon) \mathbf{M}.$$

By Eq. (1) and Fact 2.1, this implies

$$\mathbf{M}^{-1} \succcurlyeq \begin{bmatrix} \mathbf{I} & -\mathbf{Z}\mathbf{M}_{[F, C]} \\ 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{Z} & 0 \\ 0 & \text{Sc}(\widehat{\mathbf{M}}, F)^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & 0 \\ -\mathbf{M}_{[C, F]}\mathbf{Z} & \mathbf{I} \end{bmatrix} \succcurlyeq (1 + \epsilon)^{-1} \mathbf{M}^{-1}.$$

From Facts A.3 and 2.1, we know that

$$\text{Sc}(\mathbf{M}, F)^{-1} \succcurlyeq \text{Sc}(\widehat{\mathbf{M}}, F)^{-1} \succcurlyeq (1 + \epsilon)^{-1} \text{Sc}(\mathbf{M}, F)^{-1}.$$

When we use Fact 2.2 to substitute this inequality into the one above, we obtain

$$(1 + \epsilon) \mathbf{M}^{-1} \succcurlyeq \begin{bmatrix} \mathbf{I} & -\mathbf{Z}\mathbf{M}_{[F, C]} \\ 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{Z} & 0 \\ 0 & \text{Sc}(\mathbf{M}, F)^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & 0 \\ -\mathbf{M}_{[C, F]}\mathbf{Z} & \mathbf{I} \end{bmatrix} \succcurlyeq (1 + \epsilon)^{-1} \mathbf{M}^{-1},$$

which implies the lemma. \square

Lemma C.2. *Given a bDD matrix \mathbf{M} , a partition of its indices (F, C) such that $\mathbf{M}_{[F, F]}$, $\mathbf{M}_{[C, C]}$ are invertible, and an invertible hermitian operator $\mathbf{Z} \in (\mathbb{C}^{r \times r})^{|F| \times |F|}$, the following two conditions are equivalent:*

1. $0 \preceq (\mathbf{Z})^{-1} - \mathbf{M}_{[F, F]} \preceq \epsilon \cdot \text{Sc}(\mathbf{M}, C)$.

- 2.

$$\mathbf{M} \preceq \begin{bmatrix} (\mathbf{Z})^{-1} & \mathbf{M}_{[F, C]} \\ \mathbf{M}_{[C, F]} & \mathbf{M}_{[C, C]} \end{bmatrix} \preceq (1 + \epsilon) \mathbf{M}.$$

Proof. Writing

$$\mathbf{M} = \begin{bmatrix} \mathbf{M}_{[F, F]} & \mathbf{M}_{[F, C]} \\ \mathbf{M}_{[C, F]} & \mathbf{M}_{[C, C]} \end{bmatrix},$$

condition 2 is equivalent to

$$0 \preceq \begin{bmatrix} (\mathbf{Z})^{-1} - \mathbf{M}_{[F, F]} & 0 \\ 0 & 0 \end{bmatrix} \preceq \epsilon \mathbf{M}.$$

The left inequality in this statement is equivalent to the left inequality in condition 1. Thus, it suffices to prove the right sides are equivalent.

To this end, using $\mathbf{M}_{[F,C]} = \mathbf{M}_{[C,F]}^*$, it suffices to prove that $\forall x \in (\mathbb{C}^r)^{|F|}, y \in (\mathbb{C}^r)^{|C|}$,

$$x^*((\mathbf{Z})^{-1} - \mathbf{M}_{[F,F]})x \leq \epsilon(x^* \mathbf{M}_{[F,F]}x + 2x^* \mathbf{M}_{[F,C]}y + y^* \mathbf{M}_{[C,C]}y).$$

This is equivalent to proving $\forall x \in (\mathbb{C}^r)^{|F|}$,

$$x^*((\mathbf{Z})^{-1} - \mathbf{M}_{[F,F]})x \leq \epsilon \inf_{y \in (\mathbb{C}^r)^{|C|}} (x^* \mathbf{M}_{[F,F]}x + 2x^* \mathbf{M}_{[F,C]}y + y^* \mathbf{M}_{[C,C]}y).$$

Since $\mathbf{M}_{[C,C]} \succcurlyeq 0$, the rhs is a convex function of y . The minimum is achieved at $y = -\mathbf{M}_{[C,C]}^{-1} \mathbf{M}_{[F,C]}^* x$, and we obtain the equivalent condition

$$x^*((\mathbf{Z})^{-1} - \mathbf{M}_{[F,F]})x \leq \epsilon x^*(\mathbf{M}_{[F,F]} - \mathbf{M}_{[F,C]} \mathbf{M}_{[C,C]}^{-1} \mathbf{M}_{[F,C]}^*)x.$$

Since $Sc(\mathbf{M}, C) = \mathbf{M}_{[F,F]} - \mathbf{M}_{[F,C]} \mathbf{M}_{[C,C]}^{-1} \mathbf{M}_{[C,F]} = \mathbf{M}_{[F,F]} - \mathbf{M}_{[F,C]} \mathbf{M}_{[C,C]}^{-1} \mathbf{M}_{[F,C]}^*$, we obtain our claim. \square

D Finding α -bDD Subsets

In this section we check that a simple randomized sampling procedure leads to α -bDD subsets. Specifically we will prove Lemma 3.5:

Lemma 3.5. *Given a bDD matrix \mathbf{M} with n block-rows, and an $\alpha \geq 0$, BDDSUBSET computes a subset F of size least $n/(8(1+\alpha))$ such that $\mathbf{M}_{[F,F]}$ is α -bDD. It runs in runs in $O(m)$ expected work and $O(\log n)$ expected depth, where m is the number of nonzero blocks in \mathbf{M} .*

Pseudocode for this routine is given in Figure 2. We first show that the set returned is guaranteed

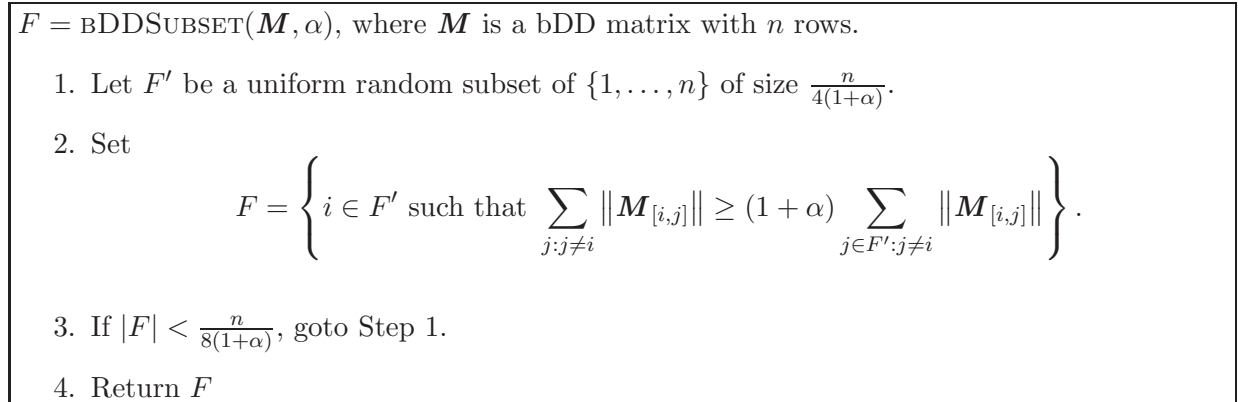


Figure 2: Routine for finding an α -strongly block diagonally dominant submatrix to be α -strongly block diagonally dominant.

Lemma D.1. *If BDDSUBSET terminates, it returns F such that $\mathbf{M}_{[F,F]}$ is α -bDD.*

Proof. Consider some $i \in F$, the criteria for including i in F in Step 2 gives:

$$\sum_{j:j \neq i} \|\mathbf{M}_{[i,j]}\| \geq (1+\alpha) \sum_{j \in F', j \neq i} \|\mathbf{M}_{[i,j]}\| \geq (1+\alpha) \sum_{j \in F, j \neq i} \|\mathbf{M}_{[i,j]}\|,$$

where the last inequality follows since F is a subset of F' .

Incorporating this into the definition of \mathbf{M} being bDD gives

$$\mathbf{M}_{[i,i]} \succcurlyeq \mathbb{I}_r \cdot \sum_{j:j \neq i} \|\mathbf{M}_{[i,j]}\| \geq (1 + \alpha)\mathbb{I}_r \cdot \sum_{j \in F, j \neq i} \|\mathbf{M}_{[i,j]}\|.$$

which means $\mathbf{M}_{[F,F]}$ is α -bDD. \square

It remains to show that the algorithm finds a big F quickly. This can be done by upper bounding the expected size of F , or the probability of a single index being in F .

Lemma D.2. *For any i , we have*

$$\Pr [i \in F' \text{ and } i \notin F] \leq \frac{1}{16(1 + \alpha)}.$$

Proof. This event only happens if $i \in F'$ and

$$\sum_{j \in F', j \neq i} \|\mathbf{M}_{[i,j]}\| > \frac{1}{1 + \alpha} \sum_{j:j \neq i} \|\mathbf{M}_{[i,j]}\|. \quad (6)$$

Conditioning on i being selected initially, or $i \in F'$, the probability that each other $j \neq i$ is in F' is

$$\frac{1}{n-1} \left(\frac{n}{4(1 + \alpha)} - 1 \right) < \frac{1}{4(\alpha + 1)},$$

which gives:

$$\mathbb{E} \left[\sum_{j \in F', j \neq i} \|\mathbf{M}_{[i,j]}\| \mid i \in F' \right] < \frac{1}{4(1 + \alpha)} \sum_{j:j \neq i} \|\mathbf{M}_{[i,j]}\|.$$

Markov's inequality then gives:

$$\Pr \left[\sum_{j \in F', j \neq i} \|\mathbf{M}_{[i,j]}\| > \frac{1}{1 + \alpha} \sum_{j:j \neq i} \|\mathbf{M}_{[i,j]}\| \mid i \in F' \right] < 1/4,$$

and thus

$$\Pr [i \in F' \text{ and } i \notin F] = \Pr [i \notin F \mid i \in F'] \Pr [i \in F'] < \frac{1}{4} \frac{1}{4(1 + \alpha)} = \frac{1}{16(1 + \alpha)}.$$

\square

Combining these two bounds gives Lemma 3.5.

Proof. (of Lemma 3.5) Applying Linearity of Expectation to Lemma D.2 gives

$$\mathbb{E} [|F' \setminus F|] \leq \frac{n}{16(1 + \alpha)}.$$

Markov's inequality then gives

$$\Pr \left[|F' \setminus F| \geq \frac{n}{8(1 + \alpha)} \right] < 1/2.$$

So, with probability at least $1/2$, $|F| \geq n/(8(1 + \alpha))$, and the algorithm will pass the test in line 3. Thus, the expected number of iterations made by the algorithm is at most 2. The claimed bounds on the expected work and depth of the algorithm follow. \square

E Jacobi Iteration on α -bDD Matrices

From an α -bDD set F , we will construct an operator \mathbf{Z} that approximates $\mathbf{M}_{[F,F]}^{-1}$ and that can be applied quickly. Specifically, we will show:

Theorem 3.7. *Let \mathbf{M} be a bDD matrix with index set V , and let $F \subseteq V$ such that $\mathbf{M}_{[F,F]}$ is α -bDD for some $\alpha \geq 4$, and has m_{FF} nonzero blocks. The algorithm $\text{JACOBI}(\epsilon, \mathbf{M}_{[F,F]}, \mathbf{b})$ acts as a linear operator \mathbf{Z} on \mathbf{b} that satisfies*

$$0 \preceq (\mathbf{Z})^{-1} - \mathbf{M}_{[F,F]} \preceq \epsilon \cdot \text{Sc}(\mathbf{M}, V \setminus F).$$

The algorithm takes $O(m_{FF} \log(\frac{1}{\epsilon}))$ work and $O(\log n \log(\frac{1}{\epsilon}))$ depth.

Pseudocode of this routine is given in Figure 3.

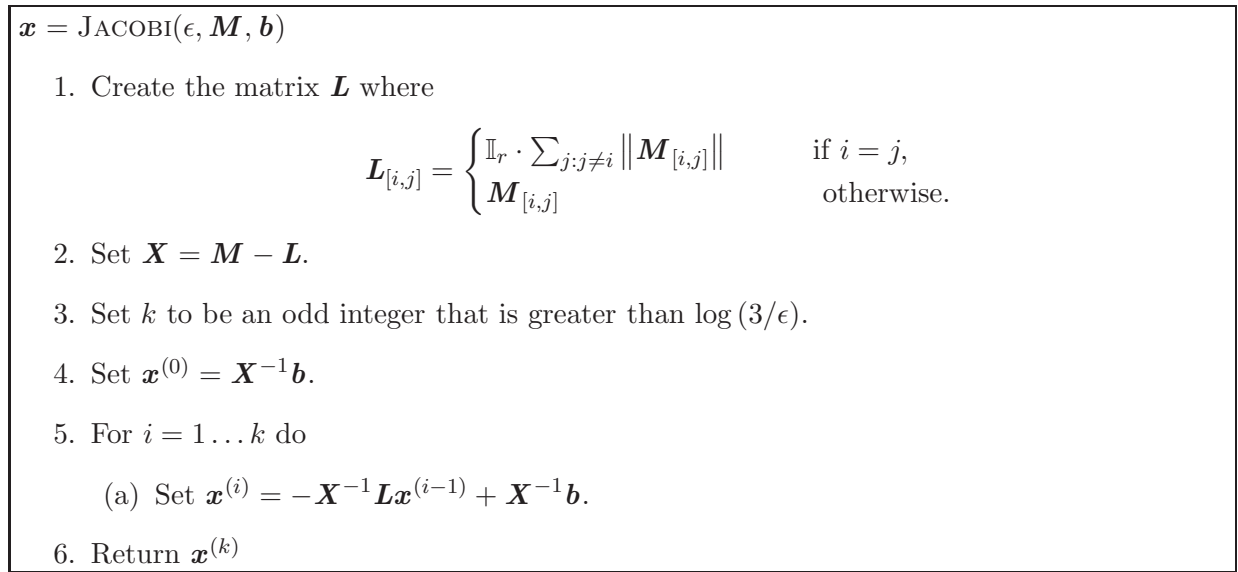


Figure 3: Jacobi Iteration for Solving Linear Systems in an α -bDD Matrix

We first verify that any α -bDD matrix has a good block-diagonal preconditioner.

Lemma 3.6. *Every α -bDD matrix \mathbf{M} can be written in the form $\mathbf{X} + \mathbf{L}$ where \mathbf{X} is block-diagonal, \mathbf{L} is bDD, and $\mathbf{X} \succcurlyeq \frac{\alpha}{2}\mathbf{L}$.*

Proof. Write $\mathbf{L} = \mathbf{Y} - \mathbf{A}$ where \mathbf{Y} is block-diagonal and \mathbf{A} has its diagonal blocks as zeros. Note that \mathbf{L} is bDD by definition. Thus, by Corollary B.5, $\mathbf{L} = \mathbf{Y} - \mathbf{A} \succcurlyeq 0$.

Using Lemma B.2, we know that $\mathbf{Y} + \mathbf{A} \succcurlyeq 0$, or $\mathbf{Y} \succcurlyeq -\mathbf{A}$. This implies $2\mathbf{Y} \succcurlyeq \mathbf{L}$.

As \mathbf{M} is α -strongly diagonally dominant and $\mathbf{M}_{[i,i]} = \mathbf{X}_{[i,i]} + \mathbf{Y}_{[i,i]}$, we have

$$(\mathbf{X} + \mathbf{Y})_{[i,i]} \succcurlyeq (1 + \alpha) \mathbb{I}_r \cdot \sum_{j \neq i} \|\mathbf{M}_{[i,j]}\| = (1 + \alpha) \mathbf{Y}_{[i,i]}.$$

Manipulating this then gives:

$$\mathbf{X} \succcurlyeq \alpha \mathbf{Y} \succcurlyeq \frac{\alpha}{2} \mathbf{L}.$$

□

We now move on to measuring the quality of the operator generated by JACOBI. It can be checked that running it k steps gives the operator

$$\mathbf{Z}^{(k)} \stackrel{\text{def}}{=} \sum_{i=0}^k \mathbf{X}_{[F,F]}^{-1} \left(-\mathbf{L}_{[F,F]} \mathbf{X}_{[F,F]}^{-1} \right)^i, \quad (7)$$

which is equivalent to evaluating a truncation of the Neumann series for \mathbf{M}^{-1}

Lemma E.1. *Let \mathbf{M} be a matrix with splitting $\mathbf{M} = \mathbf{X} + \mathbf{L}$ where $0 \preceq \mathbf{L} \preceq \beta \mathbf{X}$ for some parameter $1 > \beta > 0$. Then, for odd k and for $\mathbf{Z}^{(k)}$ as defined in (7) we have:*

$$\mathbf{X} + \mathbf{L} \preceq (\mathbf{Z}^{(k)})^{-1} \preceq \mathbf{X} + (1 + \delta) \mathbf{L} \quad (8)$$

where

$$\delta = \beta^k \frac{1 + \beta}{1 - \beta^{k+1}}.$$

Proof. The left-hand inequality is equivalent to the statement that all the eigenvalues of $\mathbf{Z}^{(k)}(\mathbf{X} + \mathbf{L})$ are at most 1 (see [BGH⁺06, Lemma 2.2] or [ST14, Proposition 3.3]). To see that this is the case, expand

$$\begin{aligned} \mathbf{Z}^{(k)}(\mathbf{X} + \mathbf{L}) &= \left(\sum_{i=0}^k \mathbf{X}^{-1} (-\mathbf{L} \mathbf{X}^{-1})^i \right) (\mathbf{X} + \mathbf{L}) \\ &= \sum_{i=0}^k (-\mathbf{X}^{-1} \mathbf{L})^i - \sum_{i=1}^{k+1} (-\mathbf{X}^{-1} \mathbf{L})^i \\ &= \mathbf{I} - (\mathbf{X}^{-1} \mathbf{L})^{k+1}. \end{aligned}$$

As all the eigenvalues of an even power of a matrix are nonnegative, all of the eigenvalues of this last matrix are at most 1.

Similarly, the other inequality is equivalent to the assertion that all of the eigenvalues of $\mathbf{Z}^{(k)}(\mathbf{X} + (1 + \delta)\mathbf{L})$ are at least one. Expanding this product yields

$$\left(\sum_{i=0}^k \mathbf{X}^{-1} (-\mathbf{L} \mathbf{X}^{-1})^i \right) (\mathbf{X} + (1 + \delta)\mathbf{L}) = \mathbf{I} - (\mathbf{X}^{-1} \mathbf{L})^{k+1} + \delta \sum_{i=0}^k (-1)^i (\mathbf{X}^{-1} \mathbf{L})^{i+1}$$

The eigenvalues of this matrix are precisely the numbers

$$1 - \lambda^{k+1} + \delta \sum_{i=0}^k (-1)^i \lambda^{i+1}, \quad (9)$$

where λ ranges over the eigenvalues of $\mathbf{X}^{-1} \mathbf{L}$. The assumption $\mathbf{L} \preceq \beta \mathbf{X}$ implies that the eigenvalues of $\mathbf{X}^{-1} \mathbf{L}$ are at most β , so $0 \leq \lambda \leq \beta$. We have chosen the value of δ precisely to guarantee that, under this condition on λ , the value of (9) is at least 1. \square

This error crucially depends only on \mathbf{L} , which for any choice of F can be upper bounded by \mathbf{M} . Propagating the error this way allows us to prove the guarantees for JACOBI

Proof. (of Theorem 3.7) Consider the matrix $\mathbf{L}_{[F,F]}$ generated when calling JACOBI with $\mathbf{M}_{[F,F]}$. \mathbf{M} being bDD means for each i we have

$$\mathbf{M}_{[i,i]} \succcurlyeq \mathbb{I}_r \cdot \sum_{j \neq i} \|\mathbf{M}_{[i,j]}\|,$$

which means for all $i \in F$

$$\mathbf{M}_{[i,i]} - \mathbb{I}_r \cdot \sum_{j \neq i, j \in F} \|\mathbf{M}_{[i,j]}\| \succcurlyeq \mathbb{I}_r \cdot \sum_{j \notin F} \|\mathbf{M}_{[i,j]}\|.$$

Therefore if we extend $\mathbf{L}_{[F,F]}$ onto the full matrix by putting zeros everywhere else, we have $\mathbf{L} \preccurlyeq \mathbf{M}$. Fact A.3 then gives $\mathbf{L}_{[F,F]} \preccurlyeq \text{Sc}(\mathbf{M}, V \setminus F)$.

Lemma 3.6 gives that $\mathbf{X}_{[F,F]} \succcurlyeq \frac{\alpha}{2} \mathbf{L}_{[F,F]}$. As $\alpha \geq 4$, we can invoke Lemma E.1 with $\beta = \frac{1}{2}$. Since $\frac{1+\beta}{1-\beta^{k+1}} \leq (\frac{3}{2})/(\frac{1}{2}) \leq 3$, our choice of $k = \log(3/\epsilon)$ gives the desired error. Each of these steps involve a matrix vector multiplication in $\mathbf{L}_{[F,F]}$ and two linear system solves in $\mathbf{X}_{[F,F]}$. The former takes $O(m)$ work and $O(\log n)$ depth since the blocks of $\mathbf{L}_{[F,F]}$ are a subset of the blocks of \mathbf{M} , while the latter takes $O(n)$ work and $O(\log n)$ depth due to $\mathbf{X}_{[F,F]}$ being block-diagonal. \square

F Existence of Linear-Sized Approximate Inverses

In this section we prove Theorem 1.3, which tells us that every *bDD* matrix has a linear-sized approximate inverse. In particular, this implies that for every *bDD* matrix there is a linear-time algorithm that approximately solves systems of equations in that matrix. To save space, we will not dwell on this algorithm, but rather will develop the linear-sized approximate inverses directly. There is some cost in doing so: there are very large constants in the linear-sized approximate inverses that are not present in a simpler linear-time solver.

To obtain a $\mathbf{U}^T \mathbf{D} \mathbf{U}$ factorization from an ϵ -SCC in which each $\mathbf{M}^{(i)}_{[F_i F_i]}$ is 4-bDD, we employ the procedure in Figure 4.




Figure 4: Converting a vertex sparsifier chain into \mathbf{U} and \mathbf{D} .

Lemma F.1. *On input an ϵ -SCC of $\mathbf{M}^{(0)}$ in which each $\mathbf{M}^{(i)}_{[F_i F_i]}$ is 4-bDD, the algorithm DECOMPOSE produces matrices \mathbf{D} and \mathbf{U} such that*

$$\mathbf{U}^T \mathbf{D} \mathbf{U} \approx_\gamma \mathbf{M},$$

where

$$\gamma \leq 2 \sum_{i=0}^{d-1} \epsilon_i + \max_i \epsilon_i + 1/2.$$

Proof. Consider the inverse of the operator $\mathbf{W} = \mathbf{W}^{(1)}$ realized by the algorithm APPLYCHAIN, and the operators $\mathbf{W}^{(i)}$ that appear in the proof of Lemma 3.3.

We have

$$\left(\mathbf{W}^{(i)}\right)^{-1} = \begin{bmatrix} \mathbf{I} & 0 \\ \mathbf{M}_{[C_i, F_i]} \mathbf{Z}^{(i)} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \left(\mathbf{Z}^{(i)}\right)^{-1} & 0 \\ 0 & \left(\mathbf{W}^{(i+1)}\right)^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{Z}^{(i)} \mathbf{M}_{[F_i, C_i]} \\ 0 & \mathbf{I} \end{bmatrix},$$

and

$$\left(\mathbf{W}^{(d)}\right)^{-1} = \mathbf{M}^{(d)} = \widehat{\mathbf{U}}^T \widehat{\mathbf{U}}.$$

After expanding and multiplying the matrices in this recursive factorization, we obtain

$$\left(\mathbf{W}^{(1)}\right)^{-1} = \mathbf{U}^T \begin{bmatrix} \left(\mathbf{Z}^{(1)}\right)^{-1} & \dots & 0 & 0 \\ 0 & \ddots & 0 & 0 \\ 0 & \dots & \left(\mathbf{Z}^{(d-1)}\right)^{-1} & 0 \\ 0 & \dots & 0 & \mathbb{I}_{|C_{d-1}|} \end{bmatrix} \mathbf{U}.$$

Moreover, we know that this latter matrix is a $2 \sum_{i=0}^{d-1} \epsilon_i$ approximation of \mathbf{M} . It remains to determine the impact of replacing the matrix in the middle of this expression with \mathbf{D} .

Lemma E.1 implies that each $\mathbf{M}_{[F_i, F_i]} \approx_{\epsilon_i} \left(\mathbf{Z}^{(i)}\right)^{-1}$ and Lemma 3.6 implies that $\mathbf{X}^{(i)} \approx_{1/2} \mathbf{M}_{[F_i, F_i]}$. So, the loss in approximation quality when we substitute the diagonal matrices is $\max_i \epsilon_i + 1/2$. \square

Invoking this decomposition procedure in conjunction with the the near-optimal sparsification routine from Theorem 3.1 gives a nearly-linear work routine. Repeatedly picking subsets using Lemma 3.8 gives then gives the linear sized decomposition.

Theorem 1.3 (Sparsified Cholesky). *For every bDD matrix \mathbf{M} with n block-rows there exists a diagonal matrix \mathbf{D} and an upper triangular matrix \mathbf{U} with $O(n)$ nonzero blocks so that*

$$\mathbf{U}^T \mathbf{D} \mathbf{U} \approx_{3/4} \mathbf{M}.$$

Moreover, linear equations in \mathbf{U} , \mathbf{U}^T , and \mathbf{D} can be solved with linear work in depth $O(\log^2 n)$, and these matrices can be computed in polynomial time.

Proof. We set $\alpha = 4$ throughout and $\epsilon_i = 1/8(i+2)^2$. Theorem 3.1 then guarantees that the average number of nonzero blocks in each column of $\mathbf{M}^{(i)}$ is at most $10r/\epsilon_i^2 = 640r(i+2)^4$. If we now apply Lemma 3.8 to find 4-diagonally dominant subsets F_i of each $\mathbf{M}^{(i)}$, we find that each such subset contains at least a 1/80 fraction of the block columns of its matrix and that each column and row of $\mathbf{M}^{(i)}$ indexed by F has at most $1280r(i+2)^4$ nonzero entries. This implies that each row of $\mathbf{Z}^{(i)} \mathbf{M}_{[F_i, C_i]}$ has at most $(1280r(i+2)^4)^{k_i+1}$ nonzero entries.

Let n_i denote the number of block columns of $\mathbf{M}^{(i)}$. By induction, we know that

$$n_i \leq n \left(1 - \frac{1}{80}\right)^{i-1}.$$

So, the total number of nonzero blocks in \mathbf{U} is at most

$$\sum_{i=1}^d n_i (1280r(i+2)^4)^{k_i+1} \leq n \sum_{i=1}^d \left(1 - \frac{1}{80}\right)^{i-1} (1280r(i+2)^4)^{k_i+1}.$$

We will show that the term multiplying n in this later expression is upper bounded by a constant. To see this, note that $k_i \leq 1 + \log(2\epsilon_i^{-1}) \leq \nu \log(i+1)$ for some constant ν . So, there is some other constant μ for which

$$(1280r(i+2)^4)^{k_i+1} \leq \exp(\mu \log^2(i+1)).$$

This implies that the sum is at most

$$\sum_{i \geq 1} \exp(\mu \log^2(i+1) - i/80),$$

which is bounded by a constant.

To bound the quality of the approximation, we compute

$$2 \sum_i \epsilon_i + \max_i \epsilon_i + 1/2 = 2 \sum_i 1/8(i+2)^2 + 1/72 + 1/2 < 3/4.$$

The claimed bound on the work to perform backwards and forwards substitution with \mathbf{U} is standard: these operations require work linear in the number of nonzero entries of \mathbf{U} . The bound on the depth follows from the fact that the substitutions can be performed level-by-level, take depth $O(\log n)$ for each level, and the number of levels, d , is logarithmic in n . \square

G Spectral Vertex Sparsification Algorithm

In this section, we give a proof of the following lemma that immediately implies Lemma 3.9.

Lemma G.1. *Let \mathbf{M} be a bDD matrix with index set V , and m nonzero blocks. Let $F \subseteq V$ be such that $\mathbf{M}_{[F,F]}$ is α -bDD for some $\alpha \geq 4$. The algorithm $\text{APPROXSCHUR}(\mathbf{M}, F, \epsilon)$, returns a matrix $\widetilde{\mathbf{M}}_{SC}$ s.t.*

1. $\widetilde{\mathbf{M}}_{SC}$ has $O(m(\epsilon^{-1} \log \log \epsilon^{-1})^{O(\log \log \epsilon^{-1})})$ nonzero blocks, and
2. $\widetilde{\mathbf{M}}_{SC} \approx_{\epsilon} Sc(\mathbf{M}, F)$,

in $O(m(\epsilon^{-1} \log \log \epsilon^{-1})^{O(\log \log \epsilon^{-1})})$ work and $O(\log n(\log \log \epsilon^{-1}))$ depth.

Moreover, if \mathbf{T} is a matrix such that only the submatrix $\mathbf{T}_{[C,C]}$ is nonzero, and $\mathbf{M} + \mathbf{T}$ is bDD, then $\text{APPROXSCHUR}(\mathbf{M} + \mathbf{T}, F, \epsilon) = \text{APPROXSCHUR}(\mathbf{M}, F, \epsilon) + \mathbf{T}_{[C,C]}$.

We show how to sparsify the Schur complement of \mathbf{M} after eliminating a set of indices F such that $\mathbf{M}_{[F,F]}$ is an α -bDD matrix. The procedure APPROXSCHUR is described in Figure 5 and uses two key subroutines SCHUR SQUARE and LASTSTEP . SCHUR SQUARE allows us to approximate $Sc(\mathbf{M}, F)$ as the Schur complement of another matrix \mathbf{M}_1 such that $\mathbf{M}_1_{[F,F]}$ is roughly α^2 -bDD. LASTSTEP allows us to approximate $Sc(\mathbf{M}, F)$ with $O(\epsilon)$ error, when $\mathbf{M}_{[F,F]}$ is roughly $1/\epsilon$ -bDD.

Lemma G.2. *Let \mathbf{M} be an bDD matrix with index set V , and m nonzero blocks and let $F \subseteq V$ be such that $\mathbf{M}_{[F,F]}$ is α -bDD for some $\alpha \geq 4$. Given $\epsilon < 1/2$, the algorithm $\text{SCHUR SQUARE}(\mathbf{M}, F, \epsilon)$, returns a bDD matrix \mathbf{M}_1 in $O(m\epsilon^{-4})$ work and $O(\log n)$ depth, such that*

1. $Sc(\mathbf{M}, F) \approx_{\epsilon} Sc(\mathbf{M}_1, F)$, and




Figure 5: Pseudocode for Computing Spectral Vertex Sparsifiers

2. $(\mathbf{M}_1)_{[F,F]}$ is $\alpha^2/2$ -bDD.

3. \mathbf{M}_1 has $O(m\epsilon^{-4})$ nonzero blocks,

If \mathbf{T} is a matrix such that only the submatrix $\mathbf{T}_{[C,C]}$ is nonzero, and $\mathbf{M} + \mathbf{T}$ is bDD, then $\text{SCHURSQUARE}(\mathbf{M} + \mathbf{T}, F, \epsilon) = \text{SCHURSQUARE}(\mathbf{M}, F, \epsilon) + \mathbf{T}_{[C,C]}$.

We can repeatedly applying the above lemma, to approximate $Sc(\mathbf{M}, F)$ as $Sc(\mathbf{M}_1, F)$, where \mathbf{M}_1 is $O(\epsilon^{-1})$ -bDD. LASTSTEP allows us to approximate the Schur complement for such a strongly block diagonally dominant matrix \mathbf{M}_1 . The guarantees of LASTSTEP are given by the following lemma.

Lemma G.3. *Let \mathbf{M} be an bDD matrix with index set V , and m nonzero blocks and let $F \subseteq V$ be such that $\mathbf{M}_{[F,F]}$ is α -bDD for some $\alpha \geq 4$. There exist a procedure LASTSTEP such that $\text{LASTSTEP}(\mathbf{M}, F, \alpha, \epsilon)$ returns in $O(m\epsilon^{-8})$ work and $O(\log n)$ depth a matrix $\widetilde{\mathbf{M}}_{SC}$ s.t. $\widetilde{\mathbf{M}}_{SC}$ has $O(m\epsilon^{-8})$ nonzero blocks and $\widetilde{\mathbf{M}}_{SC} \approx_{\epsilon+2/\alpha} Sc(\mathbf{M}, F)$. If \mathbf{T} is a matrix such that only the submatrix $\mathbf{T}_{[C,C]}$ is nonzero, and $\mathbf{M} + \mathbf{T}$ is bDD, then $\text{LASTSTEP}(\mathbf{M} + \mathbf{T}, F, \alpha, \epsilon) = \text{LASTSTEP}(\mathbf{M}, F, \alpha, \epsilon) + \mathbf{T}_{[C,C]}$.*

Combining the above two lemmas, we obtain a proof of Lemma 3.9.

Proof. (of Lemma G.1). By induction, after i steps of the main loop in APPROXSCHUR,

$$Sc(\mathbf{M}, F) \approx_{\frac{\epsilon}{2^i}} Sc(\mathbf{M}^{(i)}, F).$$

Lemma G.2 also implies that $\mathbf{M}^{(i)}$ is $2(\frac{\alpha}{2})^{2^i}$ -bDD. Thus, we have that $\mathbf{M}_{FF}^{(d)}$ is $8\epsilon^{-1}$ -strongly diagonally dominant at the last step. Hence, Lemma G.3 then gives

$$Sc(\mathbf{M}^{(d)}, F) \approx_{\frac{1}{2}\epsilon} \text{LASTSTEP}\left(\mathbf{M}^{(d)}, F, 4\epsilon^{-1}, \frac{\epsilon}{4}\right).$$

Composing this bound with the guarantees of the iterations then gives the bound on overall error.

The property that if \mathbf{T} is a matrix such that only the submatrix $\mathbf{T}_{[C,C]}$ is nonzero, and $\mathbf{M} + \mathbf{T}$ is bDD, then $\text{APPROXSCHUR}(\mathbf{M} + \mathbf{T}, F, \epsilon) = \text{APPROXSCHUR}(\mathbf{M}, F, \epsilon) + \mathbf{T}_{[C,C]}$ follows from Lemma G.2 and G.3, which ensure that this property holds for all our calls to SCHURSQUARE and LASTSTEP.

The work of these steps, and the size of the output graph follow from Lemma G.2 and G.3. \square

G.1 Iterative Squaring and Sparsification

In this section, we give a proof of Lemma G.2. At the core of procedure is a *squaring* identity that preserves Schur complements, and efficient sparsification of special classes of bDD matrices that we call product demand block-Laplacians.

Definition G.4. The product demand block-Laplacian of a vector $\mathbf{d} \in (\mathbb{C}^{r \times r})^n$, is a bDD matrix $\mathbf{L}_{G(\mathbf{d})} \in (\mathbb{C}^{r \times r})^{n \times n}$, defined as

$$(\mathbf{L}_{G(\mathbf{d})})_{[i,j]} = \begin{cases} -\mathbf{d}_{[i]} \mathbf{d}_{[j]}^* & i \neq j, \\ \mathbb{I}_r \cdot \|\mathbf{d}_{[i]}\| \sum_{k:k \neq i} \|\mathbf{d}_{[k]}\| & \text{otherwise.} \end{cases}$$

Definition G.5. Given a vector $\mathbf{d} \in (\mathbb{C}^{r \times r})^n$ and an index set $F \subseteq V$, let $C = V \setminus F$. The bipartite product demand block-Laplacian of (\mathbf{d}, F) is a bDD matrix $\mathbf{L}_{G(\mathbf{d}, F)} \in (\mathbb{C}^{r \times r})^{n \times n}$, defined as

$$(\mathbf{L}_{G(\mathbf{d}, F)})_{[i,j]} = \begin{cases} \mathbb{I}_r \cdot \|\mathbf{d}_{[i]}\| \sum_{k \in C} \|\mathbf{d}_{[k]}\| & i = j, \text{ and } i \in F \\ \mathbb{I}_r \cdot \|\mathbf{d}_{[i]}\| \sum_{k \in F} \|\mathbf{d}_{[k]}\| & i = j, \text{ and } i \in C \\ 0 & i \neq j, \text{ and } (i, j \in F \text{ or } i, j \in C) \\ -\mathbf{d}_{[i]} \mathbf{d}_{[j]}^* & i \neq j, \text{ otherwise.} \end{cases}$$

In Section G.4, we prove the following lemmas that allows us to efficiently construct sparse approximations to these matrices.

Lemma G.6. *There is a routine CLIQUESPARSIFICATION(\mathbf{d}, ϵ) such that for any demand vector $\mathbf{d} \in (\mathbb{C}^{r \times r})^n$, and $\epsilon > 0$, CLIQUESPARSIFICATION(\mathbf{d}, ϵ) returns in $O(n\epsilon^{-4})$ work and $O(\log n)$ depth a bDD matrix \mathbf{L}_H with $O(n\epsilon^{-4})$ nonzero blocks such that*

$$\mathbf{L}_H \approx_\epsilon \mathbf{L}_{G(\mathbf{d})}.$$

Lemma G.7. *There is a routine BIPARTITECLIQUESPARSIFICATION(\mathbf{d}, F, ϵ) such that for any demand vector $\mathbf{d} \in (\mathbb{C}^{r \times r})^n$, $\epsilon > 0$, and $F \subseteq V$, BIPARTITECLIQUESPARSIFICATION(\mathbf{d}, ϵ) returns in $O(n\epsilon^{-4})$ work and $O(\log n)$ depth a bDD matrix \mathbf{L}_H with $O(n\epsilon^{-4})$ nonzero blocks such that*

$$\mathbf{L}_H \approx_\epsilon \mathbf{L}_{G(\mathbf{d}, F)}.$$

Moreover, $(\mathbf{L}_H)_{[F,F]}$, and $(\mathbf{L}_H)_{[C,C]}$ are block-diagonal, where $C = V \setminus F$.

We now use these efficient sparse approximations to give a proof of the guarantees of the procedure SCHUR SQUARE.

Proof. (of Lemma G.2) Let C denote $V \setminus F$. Write $\mathbf{M}_{[F,F]}$ as $\mathbf{D} - \mathbf{A}$, where \mathbf{D} is a block-diagonal, and \mathbf{A} has its diagonal blocks as zero. The proof is based on the identity $Sc(\mathbf{M}, F) = Sc(\mathbf{M}_2, F)$, where \mathbf{M}_2 is the following matrix.

$$\mathbf{M}_2 = \frac{1}{2} \begin{bmatrix} \mathbf{D} - \mathbf{A} \mathbf{D}^{-1} \mathbf{A} & \mathbf{M}_{[F,C]} + \mathbf{A} \mathbf{D}^{-1} \mathbf{M}_{[F,C]} \\ \mathbf{M}_{[C,F]} + \mathbf{M}_{[C,F]} \mathbf{D}^{-1} \mathbf{A} & 2\mathbf{M}_{[C,C]} - \mathbf{M}_{[C,F]} \mathbf{D}^{-1} \mathbf{M}_{[F,C]} \end{bmatrix}. \quad (10)$$

It is straightforward to prove that \mathbf{M}_2 satisfies $Sc(\mathbf{M}, F) = Sc(\mathbf{M}_2, F)$ (see Lemma G.8). It is also straightforward to show that $\mathbf{D} - \mathbf{A} \mathbf{D}^{-1} \mathbf{A}$ is $((\alpha + 1)^2 - 1)$ -bDD. Thus, \mathbf{M}_2 satisfies the first two requirements.

However, \mathbf{M}_2 is likely to be a dense matrix, and we will not construct it in full. The key observation is that \mathbf{M}_2 can be written as a sum of an explicit sparse bDD matrix, and several

$\mathbf{M}_1 = \text{SCHUR SQUARE}(\mathbf{M}, F, \epsilon)$

1. Construct sparse representations of $\mathbf{d}_F^{(j)}$, $\mathbf{d}_C^{(j)}$, $\mathbf{f}^{(j)}$ defined by Eqs. (11), (12), and (13).
2. Construct \mathbf{M}_4 as given by Eq. (14).
3. Compute \mathbf{M}_3 as given by Eq. (15).
4. For all $j \in F$, let $\tilde{\mathbf{L}}_{G(\mathbf{d}_F^{(j)})} \leftarrow \text{CLIQUE SPARSIFICATION}(\mathbf{d}_F^{(j)}, \epsilon)$.
5. For all $j \in F$, let $\tilde{\mathbf{L}}_{G(\mathbf{d}_C^{(j)})} \leftarrow \text{CLIQUE SPARSIFICATION}(\mathbf{d}_C^{(j)}, \epsilon)$.
6. For all $j \in F$, let $\tilde{\mathbf{L}}_{G(\mathbf{f}^{(j)}, F)} \leftarrow \text{BIPARTITE CLIQUE SPARSIFICATION}(\mathbf{f}^{(j)}, \epsilon)$.
7. Construct and return \mathbf{M}_1 defined by Eq. (16).

Figure 6: Pseudocode for procedure SCHUR SQUARE Iterative Squaring and Sparsification

product demand block-Laplacians. Formally, for every $j \in F$, we define $\mathbf{d}_F^{(j)}, \mathbf{d}_C^{(j)} \in (\mathbb{C}^{r \times r})^n$ as follows

$$(\mathbf{d}_F^{(j)})_{[i]} = \begin{cases} \mathbf{A}_{[i,j]} \mathbf{D}_{[j,j]}^{-1/2} & i \in F, \\ 0 & i \in C. \end{cases} \quad (11)$$

$$(\mathbf{d}_C^{(j)})_{[i]} = \begin{cases} 0 & i \in F, \\ \mathbf{M}_{[i,j]} \mathbf{D}_{[j,j]}^{-1/2} & i \in C. \end{cases} \quad (12)$$

For every $j \in F$, we need to define a bipartite product demand block-Laplacian given by $\mathbf{f}^{(j)} \in (\mathbb{C}^{r \times r})^n$ defined as

$$\mathbf{f}_{[i]}^{(j)} = \begin{cases} \mathbf{A}_{[i,j]} \mathbf{D}_{[j,j]}^{-1/2} & i \in F, \\ -\mathbf{M}_{[i,j]} \mathbf{D}_{[j,j]}^{-1/2} & i \in C. \end{cases} \quad (13)$$

Letting deg_j denote the number of nonzero blocks in $\mathbf{M}_{[j]}$, the number of nonzero blocks in each of $\mathbf{d}_F^{(j)}, \mathbf{d}_C^{(j)}, \mathbf{f}^{(j)}$ is at most deg_j . We construct a sparse representation of each of them using $O(\text{deg}_j)$ work. Thus, the total number of nonzero blocks in all these block-vectors is at most $O(m)$ and we can explicitly construct sparse representations for them using $O(m)$ work and $O(\log n)$ depth.

We can now express \mathbf{M}_2 as

$$\mathbf{M}_2 = \mathbf{M}_3 + \frac{1}{2} \sum_{j \in F} \mathbf{L}_{G(\mathbf{d}_F^{(j)})} + \frac{1}{2} \sum_{j \in F} \mathbf{L}_{G(\mathbf{d}_C^{(j)})} + \frac{1}{2} \sum_{j \in F} \mathbf{L}_{G(\mathbf{f}^{(j)}, F)}.$$

For all $i, j \in V$, we compute $\beta_{i,j} = \|\mathbf{f}_i^{(j)}\|$. Define \mathbf{M}_4 to be the block diagonal matrix such that

$$\mathbf{M}_{4[i,i]} = \frac{1}{2} \mathbb{I}_r \cdot \sum_{j \in F} \beta_{i,j} \left(\left(\sum_{k \in V} \beta_{k,j} \right) - \beta_{i,j} \right). \quad (14)$$

Since at most m of $\beta_{i,j}$ are nonzero, and we can compute $\beta_{i,j}$ and construct \mathbf{M}_4 using using $O(m)$ work and $O(\log n)$ depth. It is easy to verify that we can express \mathbf{M}_3 explicitly as

$$\mathbf{M}_3 = \frac{1}{2} \begin{bmatrix} \mathbf{D} & \mathbf{M}_{[F,C]} \\ \mathbf{M}_{[C,F]} & 2\mathbf{M}_{[C,C]} \end{bmatrix} - \mathbf{M}_4. \quad (15)$$

For each j , we use Lemmas G.6 and G.7 to construct, in $O(\deg_j \epsilon^{-4})$ work and $O(\log n)$ depth, bDD matrices $\tilde{\mathbf{L}}_{G(\mathbf{d}_F^{(j)})}, \tilde{\mathbf{L}}_{G(\mathbf{d}_C^{(j)}), \tilde{\mathbf{L}}_{G(\mathbf{f}^{(j)}, F)}$, each with $O(\deg_j \epsilon^{-4})$ nonzero blocks such that

$$\tilde{\mathbf{L}}_{G(\mathbf{d}_F^{(j)})} \approx_\epsilon \mathbf{L}_{G(\mathbf{d}_F^{(j)})}, \quad \tilde{\mathbf{L}}_{G(\mathbf{d}_C^{(j)})} \approx_\epsilon \mathbf{L}_{G(\mathbf{d}_C^{(j)})}, \quad \text{and} \quad \tilde{\mathbf{L}}_{G(\mathbf{f}^{(j)}, F)} \approx_\epsilon \mathbf{L}_{G(\mathbf{f}^{(j)}, F)}.$$

We can now construct the matrix

$$\mathbf{M}_1 = \mathbf{M}_3 + \frac{1}{2} \sum_{j \in F} \tilde{\mathbf{L}}_{G(\mathbf{d}_F^{(j)})} + \frac{1}{2} \sum_{j \in F} \tilde{\mathbf{L}}_{G(\mathbf{d}_C^{(j)})} + \frac{1}{2} \sum_{j \in F} \tilde{\mathbf{L}}_{G(\mathbf{f}^{(j)}, F)}, \quad (16)$$

in $O(m\epsilon^{-4})$ time and $O(\log n)$ depth. \mathbf{M}_1 has $O(m\epsilon^{-4})$ nonzero blocks and is our required matrix.

We first show that \mathbf{M}_3 is bDD. Using $\beta_{k,j} \leq \left\| \mathbf{D}_{[j,j]}^{-1/2} \mathbf{M}_{[i,k]} \right\|$ for all $j \in F, k \in V$, we have for all $i \in F$,

$$2 \left\| \mathbf{M}_{4[i,i]} \right\| \leq \sum_{j \in F} \beta_{i,j} \sum_{k \in V} \beta_{k,j} \leq \sum_{j \in F} \left\| \mathbf{A}_{[i,j]} \right\| \left\| \mathbf{D}_{[j,j]}^{-1/2} \right\|^2 \sum_{k \in V} \left\| \mathbf{M}_{[j,k]} \right\| \leq \sum_{j \in F} \left\| \mathbf{A}_{[i,j]} \right\|, \quad (17)$$

where the last inequality uses that \mathbf{M} is bDD. Again, using \mathbf{M} is bDD, for all $i \in F$, we have

$$\mathbf{D}_{[i,i]} \succcurlyeq \mathbb{I}_r \sum_{j \in F} \left\| \mathbf{A}_{[i,j]} \right\| + \mathbb{I}_r \sum_{k \in C} \left\| \mathbf{M}_{[i,k]} \right\|.$$

Now combining with Eqs. (15) and (17), we obtain $\mathbf{M}_{3[i,i]} \succcurlyeq \mathbb{I}_r \sum_{k \in C} \left\| \mathbf{M}_{[i,k]} \right\|$. Thus, \mathbf{M}_3 is bDD, and hence \mathbf{M}_1 is bDD.

We also have $\mathbf{M}_1 \approx_\epsilon \mathbf{M}_2$. Thus, using Fact A.3, we obtain $Sc(\mathbf{M}_1, F) \approx_\epsilon Sc(\mathbf{M}_2, F)$. It remains to show that $\mathbf{M}_{1[F,F]}$ is $\alpha^2/2$ -bDD.

The key observation is that only the matrices $\tilde{\mathbf{L}}_{G(\mathbf{d}_C^{(j)})}$ contribute to off-diagonal blocks in $\mathbf{M}_{1[F,F]}$. By construction,

$$\forall i, j \in F, \quad \mathbf{L}_{G(\mathbf{d}_F^{(j)})_{[i,i]}} = \mathbb{I}_r \sum_{k \in F: k \neq i} \beta_{i,j} \beta_{k,j}.$$

Thus, for all $i \in F$

$$\begin{aligned} \sum_{j \in F} \left\| \mathbf{L}_{G(\mathbf{d}_F^{(j)})_{[i,i]}} \right\| &\leq \sum_{j \in F} \sum_{k \in F} \beta_{i,j} \beta_{k,j} \leq \sum_{j \in F} \left\| \mathbf{A}_{[i,j]} \right\| \left\| \mathbf{D}_{[j,j]}^{-1/2} \right\|^2 \sum_{k \in F} \left\| \mathbf{A}_{[j,k]} \right\| \\ &\leq (1 + \alpha)^{-1} \mathbb{I}_r \cdot \sum_{j \in F} \left\| \mathbf{A}_{[i,j]} \right\|, \end{aligned} \quad (18)$$

where the last inequality follows since $\mathbf{D} - \mathbf{A} = \mathbf{M}_{[F,F]}$ is α -bDD. Moreover $\mathbf{D} - \mathbf{A}$ being α -bDD implies that for each $i \in F$, $\mathbf{D}_{[i,i]} \succcurlyeq (\alpha + 1) \mathbb{I}_r \sum_{j \in F} \left\| \mathbf{A}_{[i,j]} \right\|$. Thus, using Eqs. (15) and (17), we obtain, $2\mathbf{M}_{3[i,i]} \succcurlyeq \alpha \mathbb{I}_r \sum_{j \in F} \left\| \mathbf{A}_{[i,j]} \right\|$.

Since for each j , $\tilde{\mathbf{L}}_{G(\mathbf{d}^{(j)})} \approx_\epsilon \tilde{\mathbf{L}}_{G(\mathbf{d}^{(j)})}$, we have $(\tilde{\mathbf{L}}_{G(\mathbf{d}^{(j)})})_{[i,i]} \preccurlyeq e^\epsilon (\mathbf{L}_{G(\mathbf{d}^{(j)})})_{[i,i]}$. Thus, $\sum_j \tilde{\mathbf{L}}_{G(\mathbf{d}_F^{(j)})}$ is bDD with sums of norms of off-diagonal blocks at most $e^\epsilon (1 + \alpha)^{-1} \sum_{j \in F} \left\| \mathbf{A}_{[i,j]} \right\|$. Since these are the only matrices that contribute to off-diagonal blocks in $\mathbf{M}_{1[F,F]}$, we get that $\mathbf{M}_{1[F,F]}$ is at least $\frac{\alpha}{e^\epsilon(1+\alpha)^{-1}}$ -bDD. Using $\epsilon < 1/2$ gives us our claim.

It is easy to verify that our transformations maintain that if \mathbf{T} is a matrix that is only nonzero inside the submatrix $\mathbf{T}_{[C,C]}$, and $\mathbf{M} + \mathbf{T}$ is bDD, then $SCHURSQUARE(\mathbf{M} + \mathbf{T}, F, \epsilon) = SCHURSQUARE(\mathbf{M}, F, \epsilon) + \mathbf{T}_{[C,C]}$. \square

We now prove the claims deferred from the above proof.

Lemma G.8. *The matrix \mathbf{M}_2 defined by Eq. (10) satisfies $Sc(\mathbf{M}_2, F) = Sc(\mathbf{M}, F)$.*

Proof. We need the following identity from [PS14]:

$$(\mathbf{D} - \mathbf{A})^{-1} = 1/2 \cdot (\mathbf{D}^{-1} + (\mathbf{I} + \mathbf{D}^{-1}\mathbf{A})(\mathbf{D} - \mathbf{A}\mathbf{D}^{-1}\mathbf{A})^{-1}(\mathbf{I} + \mathbf{A}\mathbf{D}^{-1})). \quad (19)$$

We have,

$$\begin{aligned} Sc(\mathbf{M}_2, F) &= \mathbf{M}_{[C,C]} - 1/2 \cdot \mathbf{M}_{[C,F]} \mathbf{D}^{-1} \mathbf{M}_{[F,C]} \\ &\quad - 1/2 \cdot \mathbf{M}_{[C,F]} (\mathbf{I} + \mathbf{D}^{-1}\mathbf{A})(\mathbf{D} - \mathbf{A}\mathbf{D}^{-1}\mathbf{A})^{-1} (\mathbf{I} + \mathbf{A}\mathbf{D}^{-1}) \mathbf{M}_{[F,C]} \\ &= \mathbf{M}_{[C,C]} \\ &\quad - 1/2 \cdot \mathbf{M}_{[C,F]} (\mathbf{D}^{-1} + (\mathbf{I} + \mathbf{D}^{-1}\mathbf{A})(\mathbf{D} - \mathbf{A}\mathbf{D}^{-1}\mathbf{A})^{-1} (\mathbf{I} + \mathbf{A}\mathbf{D}^{-1})) \mathbf{M}_{[F,C]} \\ &= \mathbf{M}_{[C,C]} - \mathbf{M}_{[C,F]} (\mathbf{D} - \mathbf{A})^{-1} \mathbf{M}_{[F,C]} \\ &= Sc(\mathbf{M}, F). \end{aligned}$$

□

G.2 Schur Complement w.r.t. Highly α -bDD Submatrices

$\widetilde{\mathbf{M}}_{SC} = \text{LASTSTEP}(\mathbf{M}, F, \alpha, \epsilon)$

1. Compute \mathbf{X} , \mathbf{D} , and \mathbf{A} as given by equations (20), (21), and (22).
2. Construct \mathbf{Y} as defined by equations (34), (35), and (36).
3. Construct sparse vectors $\mathbf{d}^{(j)}$, $\mathbf{g}^{(j)}$ for all $j \in F$ as defined by equations (26) and (27).
4. For all $j \in F$, let $\widetilde{\mathbf{L}}_{G(\mathbf{d}^{(j)}, F)} \leftarrow \text{BIPARTITECLIQUESPARSIFICATION}(\mathbf{d}^{(j)}, F, \epsilon/2)$
5. For all $j \in F$, let $\widetilde{\mathbf{L}}_{G(\mathbf{g}^{(j)})} \leftarrow \text{CLIQUESPARSIFICATION}(\mathbf{g}^{(j)}, \epsilon/2)$.
6. Compute \mathbf{R} as given by equation (29).
7. Construct sparse vectors $\mathbf{r}^{(j)}$ for all $j \in F$ as defined by equations (30).
8. For all $j \in F$, let $\widetilde{\mathbf{L}}_{G(\mathbf{r}^{(j)})} \leftarrow \text{CLIQUESPARSIFICATION}(\mathbf{r}^{(j)}, \epsilon/2)$.
9. Compute \mathbf{S} as given by equation (31).
10. Return

$$\widetilde{\mathbf{M}}_{SC} = \mathbf{S} + \sum_i \widetilde{\mathbf{L}}_{G(\mathbf{r}^{(i)})}.$$

Figure 7: Pseudocode for procedure LASTSTEP: Computing an approximate Schur complement w.r.t. a highly α -bDD submatrices.

In this section, we describe the LASTSTEP procedure for computing an approximate Schur complement of a bDD matrix \mathbf{M} with a highly α -bDD submatrix $\mathbf{M}_{[F,F]}$. LASTSTEP is the final step of the APPROXSCHUR algorithm. The key element of the procedure is a formula for

approximating the inverse of $\mathbf{M}_{[F,F]}$ that is leveraged to approximate the Schur complement of \mathbf{M} as the Schur complement of matrix with the FF submatrix being block diagonal.

We prove guarantees for the LASTSTEP algorithm as stated in Lemma G.3.

One could attempt to deal with the highly α -bDD matrix at the last step by directly replacing it with its diagonal, but this is problematic. Consider the case where F contains u and v with a weight ϵ edge between them, and u and v are connected to u' and v' in C by weight 1 edges respectively. Keeping only the diagonal results in a Schur complement that disconnects u' and v' . This however can be fixed by taking a step of random walk within F .

Given a bDD matrix \mathbf{M} , s.t. $\mathbf{M}_{[F,F]}$ is α -bDD we define a block diagonal matrix $\mathbf{X} \in (\mathbb{C}^{r \times r})^{|F| \times |F|}$ s.t. for each $i \in F$

$$\mathbf{X}_{[i,i]} = \frac{\alpha}{\alpha + 1} \mathbf{M}_{[i,i]} \quad (20)$$

and another block diagonal matrix $\mathbf{D} \in (\mathbb{C}^{r \times r})^{|F| \times |F|}$ s.t. for each $i \in F$

$$\mathbf{D}_{[i,i]} = \frac{1}{\alpha + 1} \mathbf{M}_{[i,i]}, \quad (21)$$

and we define a matrix $\mathbf{A} \in (\mathbb{C}^{r \times r})^{|F| \times |F|}$

$$\mathbf{A}_{[i,j]} = \begin{cases} 0 & i = j \\ -\mathbf{M}_{[i,j]} & \text{otherwise.} \end{cases} \quad (22)$$

Thus $\mathbf{M}_{[F,F]} = \mathbf{X} + \mathbf{D} - \mathbf{A}$. One can check that because \mathbf{M} is bDD and $\mathbf{M}_{[F,F]}$ is α -bDD, it follows that $\mathbf{D} - \mathbf{A}$ is bDD and the matrix

$$\begin{bmatrix} \mathbf{X} & \mathbf{M}_{[F,C]} \\ \mathbf{M}_{[C,F]} & \mathbf{M}_{[C,C]} \end{bmatrix}$$

is also bDD.

We will consider the linear operator

$$\mathbf{Z}^{(last)} \stackrel{\text{def}}{=} \frac{1}{2} \mathbf{X}^{-1} + \frac{1}{2} \mathbf{X}^{-1} (\mathbf{X} - \mathbf{D} + \mathbf{A}) \mathbf{X}^{-1} (\mathbf{X} - \mathbf{D} + \mathbf{A}) \mathbf{X}^{-1}. \quad (23)$$

We define

$$\mathbf{M}^{(last)} = \begin{pmatrix} \left(\mathbf{Z}^{(last)} \right)^{-1} & \mathbf{M}_{[F,C]} \\ \mathbf{M}_{[F,C]} & \mathbf{M}_{[C,C]} \end{pmatrix}$$

Lemma G.9.

$$\mathbf{M} \preceq \mathbf{M}^{(last)} \preceq \left(1 + \frac{2}{\alpha} \right) \mathbf{M}.$$

We defer the proof of Lemma G.9 to Section G.3.

To utilize $\mathbf{Z}^{(last)}$, define

$$\mathbf{M}_2^{(last)} \stackrel{\text{def}}{=} \begin{bmatrix} \frac{1}{2} \mathbf{X} & \frac{1}{2} (\mathbf{X} - \mathbf{D} + \mathbf{A}) \mathbf{X}^{-1} \mathbf{M}_{[F,C]} \\ \frac{1}{2} \mathbf{M}_{[C,F]} \mathbf{X}^{-1} (\mathbf{X} - \mathbf{D} + \mathbf{A}) & \mathbf{M}_{[C,C]} - \frac{1}{2} \mathbf{M}_{[C,F]} \mathbf{X}^{-1} \mathbf{M}_{[F,C]} \end{bmatrix} \quad (24)$$

and note

$$S_C \left(\mathbf{M}^{(last)}, F \right) = S_C \left(\mathbf{M}_2^{(last)}, F \right) \quad (25)$$

Lemma G.9 tells us that for large enough α , we can approximate the Schur complement of \mathbf{M} by approximating the the Schur complement of $\mathbf{M}_2^{(last)}$.

The next lemma tells us that M_2 is bDD and that we can write the matrix as a sum of an explicit bDD matrix and sparse implicitly represented product demand block-Laplacians and bipartite product demand block-Laplacians.

Lemma G.10. *Consider a bDD matrix M , where $M_{[F,F]}$ is α -bDD for some $\alpha \geq 4$, and let $M_2^{(last)}$ be the associated matrix defined by equation (24). Let m be the number of nonzero blocks of M .*

For $j \in F$, we define $\mathbf{d}^{(j)} \in (\mathbb{C}^{r \times r})^n$

$$\mathbf{d}_{[i]}^{(j)} = \begin{cases} \mathbf{A}_{[i,j]} \mathbf{X}_{[j,j]}^{-1/2} & \text{for } i \in F \\ \mathbf{M}_{[i,j]} \mathbf{X}_{[j,j]}^{-1/2} & \text{for } i \in C \end{cases} \quad (26)$$

For $j \in F$, we define $\mathbf{g}^{(j)} \in (\mathbb{C}^{r \times r})^n$

$$\mathbf{g}_{[i]}^{(j)} = \begin{cases} 0 & \text{for } i \in F \\ \mathbf{M}_{[i,j]} \mathbf{X}_{[j,j]}^{-1/2} & \text{for } i \in C \end{cases} \quad (27)$$

Then

$$\mathbf{M}_2^{(last)} = \mathbf{Y} + \frac{1}{2} \sum_{j \in F} \mathbf{L}_{G(\mathbf{g}^{(j)})} + \frac{1}{2} \sum_{j \in F} \mathbf{L}_{G(\mathbf{d}^{(j)}, F)}$$

where \mathbf{Y} is bDD and has $O(m)$ nonzero blocks, and the total number of nonzero blocks in $\mathbf{d}^{(i)}$ and $\mathbf{g}^{(i)}$ for all i combined is also $O(m)$.

\mathbf{Y} as well as $\mathbf{d}^{(i)}$ and $\mathbf{g}^{(i)}$ for all i can be computed in $O(m)$ time and $O(\log n)$ depth.

If \mathbf{T} is a matrix that is only nonzero inside the submatrix $\mathbf{T}_{[C,C]}$, then if we apply the transformation of Eq. 24, to $\mathbf{M} + \mathbf{T}$ instead of \mathbf{M} , we find $(\mathbf{M} + \mathbf{T})_2^{(last)} = \mathbf{M}_2^{(last)} + \mathbf{T}_{[C,C]}$, and in particular $\mathbf{Y}(\mathbf{M} + \mathbf{T}) = \mathbf{Y}(\mathbf{M}) + \mathbf{T}_{[C,C]}$.

We defer the proof of Lemma G.10 to Section G.3.

Proof. (of Lemma G.3) The procedure LASTSTEP(M, F, α, ϵ) first computes \mathbf{Y} and $\mathbf{d}^{(i)}$ and $\mathbf{g}^{(i)}$ for all i s.t.

$$\mathbf{M}_2^{(last)} = \mathbf{Y} + \frac{1}{2} \sum_{j \in F} \mathbf{L}_{G(\mathbf{g}^{(j)})} + \frac{1}{2} \sum_{j \in F} \mathbf{L}_{G(\mathbf{d}^{(j)}, F)}, \quad (28)$$

where $\mathbf{M}_2^{(last)} \approx_{2/\alpha} M$. Let $n_{\mathbf{d}^{(i)}}$ and $n_{\mathbf{g}^{(i)}}$ denote the number of nonzero blocks in each $\mathbf{d}^{(i)}$ and $\mathbf{g}^{(i)}$. By Lemma G.6, we may call CLIQUESPARSIFICATION($\mathbf{g}^{(i)}, \epsilon/2$) in $O(n_{\mathbf{g}^{(i)}} \epsilon^{-4})$ time to sparsify $\mathbf{L}_{G(\mathbf{g}^{(i)})}$, producing a bDD matrix $\tilde{\mathbf{L}}_{G(\mathbf{g}^{(i)})} \approx_{\epsilon/2} \mathbf{L}_{G(\mathbf{g}^{(i)})}$ with $O(n_{\mathbf{g}^{(i)}} \epsilon^{-4})$ nonzero blocks. By Lemma G.7, we may call BIPARTITECLIQUESPARSIFICATION($\mathbf{d}^{(i)}, F, \epsilon/2$) in $O(n_{\mathbf{d}^{(i)}} \epsilon^{-4})$ time to sparsify $\mathbf{L}_{G(\mathbf{d}^{(i)}, F)}$, producing a bDD matrix $\tilde{\mathbf{L}}_{G(\mathbf{d}^{(i)}, F)} \approx_{\epsilon/2} \mathbf{L}_{G(\mathbf{d}^{(i)}, F)}$ with $O(n_{\mathbf{d}^{(i)}} \epsilon^{-4})$ nonzero blocks. The total running time of this is $O(\epsilon^{-4} \sum_i (n_{\mathbf{d}^{(i)}} + n_{\mathbf{g}^{(i)}})) = O(\epsilon^{-4} m)$, and the total number of nonzero blocks in

$$\frac{1}{2} \sum_{j \in F} \tilde{\mathbf{L}}_{G(\mathbf{g}^{(j)})} + \frac{1}{2} \sum_{j \in F} \tilde{\mathbf{L}}_{G(\mathbf{d}^{(j)}, F)},$$

is $O(\epsilon^{-4} (\sum_i n_{\mathbf{d}^{(i)}} + n_{\mathbf{g}^{(i)}})) = O(\epsilon^{-4} m)$. We define

$$\mathbf{R} = \mathbf{Y} + \frac{1}{2} \sum_{j \in F} \tilde{\mathbf{L}}_{G(\mathbf{g}^{(j)})} + \frac{1}{2} \sum_{j \in F} \tilde{\mathbf{L}}_{G(\mathbf{d}^{(j)}, F)}. \quad (29)$$

which we can compute in $O(\epsilon^{-4}m)$ time and $O(\log n)$ depth. We have $\mathbf{M} \approx_{2/\alpha} \mathbf{M}_2^{(last)}$ and $\mathbf{M}_2^{(last)} \approx_{\epsilon/2} \mathbf{R}$, so that $\mathbf{M} \approx_{2/\alpha+\epsilon/2} \mathbf{R}$. It follows from Fact A.3 that $Sc(\mathbf{M}, F) \approx_{2/\alpha+\epsilon/2} Sc(\mathbf{R}, F)$.

Because the sparsifiers computed by BIPARTITECLIQUESPARSIFICATION preserve the graph bipartition, $\mathbf{R}_{[F,F]}$ is block diagonal.

We can use the block diagonal structure of $\mathbf{R}_{[F,F]}$ quickly compute a sparse approximation to $Sc(\mathbf{R}, F)$.

For $j \in F$, we define $\mathbf{g}^{(j)} \in (\mathbb{C}^{r \times r})^{|C|}$

$$\mathbf{r}_i^{(j)} = \mathbf{R}_{[i,j]} \mathbf{R}_{[j,j]}^{-1/2} \text{ for } i \in C \quad (30)$$

Then

$$Sc(\mathbf{R}, F) = \mathbf{R}_{[C,C]} - \mathbf{R}_{[C,F]} \mathbf{R}_{[F,F]}^{-1} \mathbf{R}_{[F,C]} = \mathbf{S} + \sum_{j \in F} \mathbf{L}_{G(\mathbf{r}^{(j)})}.$$

where

$$\mathbf{S} = \mathbf{R}_{[C,C]} - \text{Diag}_{i \in C} \left(\mathbb{I}_r \sum_{j \in F} \sum_{k \in C \setminus \{i\}} \left\| \mathbf{R}_{[i,j]} \mathbf{R}_{[j,j]}^{-1/2} \right\| \left\| \mathbf{R}_{[k,j]} \mathbf{R}_{[j,j]}^{-1/2} \right\| \right) \quad (31)$$

$$- \text{Diag}_{i \in C} \sum_{j \in F} \left(\mathbf{R}_{[i,j]} \mathbf{R}_{[j,j]}^{-1} \mathbf{R}_{[j,i]} \right). \quad (32)$$

Let us define, for each block row $i \in F$, $\rho_i = \sum_{j \in C} \|\mathbf{R}_{[i,j]}\|$, and for each block row $i \in C$, $\rho_i = \sum_{j \in F} \|\mathbf{R}_{[i,j]}\|$. From \mathbf{R} being bDD, we then conclude for each $i \in F$

$$\mathbf{R}_{[i,i]} \succcurlyeq \mathbb{I}_r \rho_i$$

and for each $i \in C$

$$\mathbf{R}_{[i,i]} \succcurlyeq \mathbb{I}_r \left(\rho_i + \sum_{j \in C} \|\mathbf{R}_{[i,j]}\| \right).$$

With this in mind, we check that each for $i \in C$ of \mathbf{S} is bDD.

The diagonal block satisfies

$$\begin{aligned} \mathbf{S}_{[i,i]} &\succcurlyeq \mathbb{I}_r \left(\sum_{j \in C} \|\mathbf{R}_{[i,j]}\| + \rho_i - \sum_{j \in F} \sum_{k \in C} \left\| \mathbf{R}_{[i,j]} \mathbf{R}_{[j,j]}^{-1/2} \right\| \left\| \mathbf{R}_{[k,j]} \mathbf{R}_{[j,j]}^{-1/2} \right\| \right) \\ &\succcurlyeq \mathbb{I}_r \left(\sum_{j \in C} \|\mathbf{R}_{[i,j]}\| + \rho_i - \sum_{j \in F} \sum_{k \in C} \frac{1}{\rho_j} \|\mathbf{R}_{[i,j]}\| \|\mathbf{R}_{[k,j]}\| \right) \\ &\succcurlyeq \mathbb{I}_r \left(\sum_{j \in C} \|\mathbf{R}_{[i,j]}\| + \rho_i - \sum_{j \in F} \frac{1}{\rho_j} \|\mathbf{R}_{[i,j]}\| \right) \\ &\succcurlyeq \mathbb{I}_r \sum_{j \in C} \|\mathbf{R}_{[i,j]}\| \end{aligned}$$

We can compute \mathbf{S} and all $\mathbf{r}_i^{(j)}$ in $O(m\epsilon^{-4})$ time and $O(\log n)$ depth, since this is an upper bound to the number of nonzero blocks in \mathbf{R} .

Let $n_{\mathbf{r}^{(i)}}$ denote the number of nonzero blocks in $\mathbf{r}^{(i)}$. By Lemma G.6, we may call CLIQUESPARSIFICATION($\mathbf{r}^{(i)}$), in $O(n_{\mathbf{r}^{(i)}}\epsilon^{-4})$ time to sparsify $\mathbf{L}_{G(\mathbf{r}^{(i)})}$, producing a bDD matrix $\tilde{\mathbf{L}}_{G(\mathbf{r}^{(i)})} \approx_{\epsilon/2} \mathbf{L}_{G(\mathbf{r}^{(i)})}$ with $O(n_{\mathbf{r}^{(i)}}\epsilon^{-4})$ nonzero blocks.

The total number of nonzero blocks in

$$\mathbf{S} + \sum_i \tilde{\mathbf{L}}_{G(r^{(i)})}$$

is also $O(\epsilon^{-8}m)$, and $\mathbf{S} + \sum_i \tilde{\mathbf{L}}_{G(r^{(i)})} \approx_{\epsilon/2} \text{Sc}(\mathbf{R}, F)$. So

$$\widetilde{\mathbf{M}}_{SC} = \mathbf{S} + \sum_i \tilde{\mathbf{L}}_{G(r^{(i)})} \approx_{\epsilon+2/\alpha} \text{Sc}(\mathbf{M}, F)$$

The total amount of work to compute these $\tilde{\mathbf{L}}_{G(r^{(i)})}$ is $O(\epsilon^{-4}(\sum_i n_{r^{(i)}})) = O(\epsilon^{-8}m)$. The depth for this computation is $O(\log n)$.

Suppose \mathbf{T} is a matrix that is only nonzero inside the submatrix $\mathbf{T}_{[C,C]}$, and $\mathbf{M} + \mathbf{T}$ is bDD. We can show that $\text{LASTSTEP}(\mathbf{M} + \mathbf{T}, F, \alpha, \epsilon) = \text{LASTSTEP}(\mathbf{M}, F, \alpha, \epsilon) + \mathbf{T}_{[C,C]}$, by first noting that this type of property holds for \mathbf{Y} by Lemma G.10, and from this concluding that similarly if we consider \mathbf{R} as a function of \mathbf{M} then $\mathbf{R}_{[C,C]}(\mathbf{M} + \mathbf{T}) = \mathbf{R}_{[C,C]}(\mathbf{M}) + \mathbf{T}_{[C,C]}$, and finally considering $\widetilde{\mathbf{M}}_{SC}$ as a function of \mathbf{M} , we can then easily show that $(\mathbf{M} + \mathbf{T})_{SC} = \widetilde{\mathbf{M}}_{SC} + \mathbf{T}_{[C,C]}$. \square

G.3 Deferred Proofs from Section G.2

To help us prove Lemma G.9, we first prove the next lemma.

Lemma G.11. *If $\mathbf{M}_{[F,F]} = \mathbf{X} + \mathbf{D} - \mathbf{A}$ be a α -bDD matrix for some $\alpha \geq 4$, then the operator $\mathbf{Z}^{(last)}$ as defined in Equation 23 satisfies:*

$$\mathbf{M}_{[F,F]} \preceq \left(\mathbf{Z}^{(last)} \right)^{-1} \preceq \mathbf{M}_{[F,F]} + \frac{2}{\alpha} (\mathbf{D} - \mathbf{A}).$$

Proof. Composing both sides by $\mathbf{X}^{-1/2}$ and substituting in $\mathcal{L} = \mathbf{X}^{-1/2} (\mathbf{D} - \mathbf{A}) \mathbf{X}^{-1/2}$ means it suffices to show

$$\mathbf{I} + \mathcal{L} \preceq \left(\frac{1}{2} \mathbf{I} + \frac{1}{2} (\mathbf{I} - \mathcal{L})^2 \right)^{-1} \preceq \mathbf{I} + \mathcal{L} + \frac{2}{\alpha} \mathcal{L}.$$

We can use the fact that $\mathbf{M}_{[F,F]}$ is α -strongly diagonally dominant to show $0 \preceq \mathbf{L} \preceq \frac{2}{\alpha} \mathbf{X}$, and equivalently $0 \preceq \mathcal{L} \preceq \frac{2}{\alpha} \mathbf{I}$, as follows:

Firstly, $\mathbf{D} - \mathbf{A} \succcurlyeq 0$ as $\mathbf{D} - \mathbf{A}$ is bDD, similarly, $\mathbf{D} + \mathbf{A} \succcurlyeq 0$ as $\mathbf{D} + \mathbf{A}$ is also bDD. From the latter $\mathbf{D} \succcurlyeq -\mathbf{A}$, so $2\mathbf{D} \succcurlyeq \mathbf{D} - \mathbf{A}$. Finally, $\mathbf{X} \succcurlyeq \alpha \mathbf{D} \succcurlyeq \frac{\alpha}{2} (\mathbf{D} - \mathbf{A})$.

As \mathcal{L} and \mathbf{I} commute, the spectral theorem means it suffices to show this for any scalar $0 \leq t \leq \frac{2}{\alpha}$. Note that

$$\frac{1}{2} + \frac{1}{2} (1 - t)^2 = 1 - t + \frac{1}{2} t^2$$

Taking the difference between the inverse of this and the ‘true’ value of $1 + t$ gives:

$$\left(1 - t + \frac{1}{2} t^2 \right)^{-1} - (1 + t) = \frac{1 - (1 + t) (1 - t + \frac{1}{2} t^2)}{1 - t + \frac{1}{2} t^2} = \frac{\frac{1}{2} t^2 (1 - t)}{1 - t + \frac{1}{2} t^2}$$

Incorporating the assumption that $0 \leq t \leq \frac{2}{\alpha}$ and $\alpha \geq 4$ gives that the denominator is at least

$$1 - \frac{2}{\alpha} \geq \frac{1}{2},$$

and the numerator term can be bounded by

$$0 \leq \frac{t^2}{2} (1 - t) \leq \frac{t}{\alpha}.$$

Combining these two bounds then gives the result. \square

Lemma G.9 allows us to extend the approximation of $\mathbf{M}_{[F,F]}$ by the inverse of $\mathbf{Z}^{(last)}$ to the entire matrix \mathbf{M} .

Proof. (of Lemma G.9) Recall that when a matrix \mathbf{T} is PSD,

$$\begin{pmatrix} \mathbf{T} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \succcurlyeq \mathbf{0}. \quad (33)$$

The left-hand inequality of our lemma follows immediately from Eq. 33 and the left-hand side of the guarantee of Lemma G.11. To prove the right-hand inequality we apply Eq. 33 and the right-hand side of the guarantee of Lemma G.11. to conclude

$$\begin{pmatrix} \left(\mathbf{Z}^{(last)}\right)^{-1} & \mathbf{M}_{[F,C]} \\ \mathbf{M}_{[F,C]} & \mathbf{M}_{[C,C]} \end{pmatrix} \preccurlyeq \begin{pmatrix} \mathbf{M}_{[F,F]} + \frac{2}{\alpha}\mathbf{L} & \mathbf{M}_{[F,C]} \\ \mathbf{M}_{[F,C]} & \mathbf{M}_{[C,C]} \end{pmatrix} = \mathbf{M} + \frac{2}{\alpha} \begin{pmatrix} \mathbf{L} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}.$$

The matrix

$$\begin{bmatrix} \mathbf{X} & \mathbf{M}_{[F,C]} \\ \mathbf{M}_{[C,F]} & \mathbf{M}_{[C,C]} \end{bmatrix}$$

is bDD and hence PSD. It follows that

$$\begin{pmatrix} \mathbf{D} - \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \preccurlyeq \mathbf{M},$$

by which we may conclude that

$$\mathbf{M} + \frac{2}{\alpha} \begin{pmatrix} \mathbf{D} - \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \preccurlyeq \mathbf{M} + \frac{2}{\alpha}\mathbf{M}.$$

□

Proof. (of Lemma G.10)

Each product demand clique $\mathbf{L}_{G(g^{(i)})}$ and bipartite product demand clique $\mathbf{L}_{G(d^{(i)},F)}$ is bDD.

We now have to find an expression for \mathbf{Y} and show that \mathbf{Y} is bDD. Let us write \mathbf{Y} in terms of its blocks From

$$\mathbf{Y} = \mathbf{M} - \left(\frac{1}{2} \sum_{j \in F} \mathbf{L}_{G(g^{(j)})} + \frac{1}{2} \sum_{j \in F} \mathbf{L}_{G(d^{(j)},F)} \right)$$

it then follows by a simple check that

$$\mathbf{Y}_{[C,F]}^* = \mathbf{Y}_{[F,C]} = \frac{1}{2}(\mathbf{X} - \mathbf{D})\mathbf{X}^{-1}\mathbf{M}_{[F,C]} = \frac{1}{2} \left(1 - \frac{1}{\alpha} \right) \mathbf{M}_{[F,C]} \quad (34)$$

and that

$$\mathbf{Y}_{[F,F]} = \frac{1}{2}\mathbf{X} - \frac{1}{2} \text{Diag} \left(\mathbb{I}_r \sum_{j \in F} \sum_{k \in C} \left\| \mathbf{A}_{[i,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \left\| \mathbf{M}_{[k,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \right) \quad (35)$$

and

$$\begin{aligned} \mathbf{Y}_{[C,C]} = & \mathbf{M}_{[C,C]} - \frac{1}{2} \text{Diag} \left(\mathbf{M}_{[i,F]} \mathbf{X}^{-1} \mathbf{M}_{[F,i]} \right) \\ & - \frac{1}{2} \text{Diag} \left(\mathbb{I}_r \sum_{j \in F} \sum_{k \in C \setminus \{i\}} \left\| \mathbf{M}_{[i,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \left\| \mathbf{M}_{[k,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \right) \end{aligned}$$

$$- \frac{1}{2} \text{Diag}_{i \in C} \left(\mathbb{I}_r \sum_{j \in F} \sum_{k \in F} \left\| \mathbf{M}_{[i,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \left\| \mathbf{A}_{[k,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \right). \quad (36)$$

Next, we check that \mathbf{Y} is bDD. First, let us define, for each block row $i \in F$, $\rho_i = \sum_{j \in F} \|\mathbf{A}_{[i,j]}\|$. From $\mathbf{M}_{[F,F]}$ being α -bDD, we then conclude $\mathbf{M}_{[i,i]} \succcurlyeq \mathbb{I}_r(1 + \alpha)\rho_i$. And from \mathbf{M} being bDD, we find that $\sum_{j \in C} \|\mathbf{M}_{[i,j]}\| \leq \alpha\rho_i$. We now check that each block row $i \in F$ of \mathbf{Y} is bDD. The off-diagonal block norm sum is at most $\frac{1}{2}(1 - \frac{1}{\alpha})\alpha\rho_i$.

The diagonal satisfies

$$\begin{aligned} \mathbf{Y}_{[i,i]} &\succcurlyeq \mathbb{I}_r \left(\frac{1}{2}\alpha\rho_i - \frac{1}{2} \left(\sum_{j \in F} \sum_{k \in C} \left\| \mathbf{A}_{[i,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \left\| \mathbf{M}_{[k,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \right) \right) \\ &\succcurlyeq \mathbb{I}_r \left(\frac{1}{2}\alpha\rho_i - \frac{1}{2} \left(\sum_{j \in F} \sum_{k \in C} \frac{1}{\alpha\rho_j} \left\| \mathbf{A}_{[i,j]} \right\| \left\| \mathbf{M}_{[k,j]} \right\| \right) \right) \\ &\succcurlyeq \mathbb{I}_r \left(\frac{1}{2}\alpha\rho_i - \frac{1}{2} \left(\sum_{j \in F} \left\| \mathbf{A}_{[i,j]} \right\| \right) \right) = \mathbb{I}_r \left(\frac{1}{2} \left(1 - \frac{1}{\alpha} \right) \alpha\rho_i \right) \end{aligned}$$

So the block rows with $i \in F$ are bDD.

Next we check the block rows for $i \in C$. Let us define for each $i \in C$, $\rho_i = \sum_{j \in F} \|\mathbf{M}_{[i,j]}\|$. Thus for \mathbf{Y} , the sum of block norms of row i over columns $j \in F$

$$\sum_{j \in F} \|\mathbf{Y}_{[i,j]}\| = \frac{1}{2} \left(1 - \frac{1}{\alpha} \right) \sum_{j \in F} \|\mathbf{M}_{[i,j]}\| = \frac{1}{2} \left(1 - \frac{1}{\alpha} \right) \rho_i.$$

For \mathbf{Y} , the sum of block norms of row i over columns $j \in C$ is $\sum_{j \in C \setminus \{i\}} \|\mathbf{M}_{[i,j]}\|$. So the total sum of the blocks norms of off-diagonals is

$$\frac{1}{2}\rho_i \left(1 - \frac{1}{\alpha} \right) + \sum_{j \in C \setminus \{i\}} \|\mathbf{M}_{[i,j]}\|$$

The diagonal block satisfies

$$\begin{aligned} \mathbf{Y}_{[i,i]} &\succcurlyeq \mathbb{I}_r \left(\rho_i - \frac{1}{2} \left\| \mathbf{M}_{[i,F]} \mathbf{X}^{-1} \mathbf{M}_{[F,i]} \right\| \right. \\ &\quad \left. - \frac{1}{2} \left(\sum_{j \in F} \sum_{k \in C \setminus \{i\}} \left\| \mathbf{M}_{[i,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \left\| \mathbf{M}_{[k,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \right) \right. \\ &\quad \left. - \frac{1}{2} \left(\sum_{j \in F} \sum_{k \in F} \left\| \mathbf{M}_{[i,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \left\| \mathbf{A}_{[k,j]} \mathbf{X}_{[j,j]}^{-1/2} \right\| \right) \right) \\ &\succcurlyeq \mathbb{I}_r \left(\rho_i + \sum_{j \in C \setminus \{i\}} \|\mathbf{M}_{[i,j]}\| \right. \\ &\quad \left. - \frac{1}{2} \left(\sum_{j \in F} \sum_{k \in C} \frac{1}{\alpha\rho_j} \|\mathbf{M}_{[i,j]}\| \|\mathbf{M}_{[k,j]}\| \right) \right) \end{aligned}$$

$$\begin{aligned}
& -\frac{1}{2} \left(\sum_{j \in F} \sum_{k \in F} \frac{1}{\alpha \rho_j} \|\mathbf{M}_{[i,j]}\| \|\mathbf{A}_{[k,j]}\| \right) \\
& \asymp \mathbb{I}_r \left(\rho_i + \sum_{j \in C \setminus \{i\}} \|\mathbf{M}_{[i,j]}\| - \frac{1}{2} \rho_i - \frac{1}{2} \frac{1}{\alpha} \rho_i \right) \\
& = \mathbb{I}_r \left(\frac{1}{2} \rho_i \left(1 - \frac{1}{\alpha} \right) + \sum_{j \in C \setminus \{i\}} \|\mathbf{M}_{[i,j]}\| \right).
\end{aligned}$$

For these block rows are also bDD, and hence \mathbf{Y} is bDD. It is clear from the definitions that \mathbf{Y} as well as $\mathbf{d}^{(i)}$ and $\mathbf{g}^{(i)}$ for all i can be computed in $O(m)$ time and $O(\log n)$ depth.

It is easy to verify that if \mathbf{T} is a matrix that is only nonzero inside the submatrix $\mathbf{T}_{[C,C]}$, then if we apply the transformation of Eq. 24, to $\mathbf{M} + \mathbf{T}$ instead of \mathbf{M} , we find $(\mathbf{M} + \mathbf{T})_2^{(last)} = \mathbf{M}_2^{(last)} + \mathbf{T}_{[C,C]}$, and in particular $\mathbf{Y}(\mathbf{M} + \mathbf{T}) = \mathbf{Y}(\mathbf{M}) + \mathbf{T}_{[C,C]}$. \square

G.4 Sparsifying Product Demand Block-Laplacians

In this section, we show how to efficiently sparsify product demand block-Laplacians and their bipartite analogs. We prove the following key lemma later in this section that allows us to transfer results on graph sparsification to sparsifying these product block-Laplacians.

Lemma G.12. *Suppose we have two graphs on n vertices, $H^{(1)}$ and $H^{(2)}$ such that $H^{(1)} \approx_\epsilon H^{(2)}$. Given $\mathbf{d} \in (\mathbb{C}^{r \times r})^n$, define the bDD matrix $\mathbf{L}^{(\ell)} \in (\mathbb{C}^{r \times r})^{n \times n}$ for each $\ell = 1, 2$, as*

$$\mathbf{L}^{(\ell)}_{[i,j]} = \begin{cases} -\frac{h_{i,j}^{(\ell)}}{\|\mathbf{d}_{[i]}\| \|\mathbf{d}_{[j]}\|} \mathbf{d}_{[i]} \mathbf{d}_{[j]}^* & i \neq j, \\ \mathbb{I}_r \cdot \sum_{k:k \neq i} h_{i,k}^{(\ell)} & \text{otherwise,} \end{cases}$$

where $h_{i,j}^{(\ell)}$ denotes the weight of the edge i, j in $H^{(\ell)}$. Then, $\mathbf{L}^{(1)} \approx_\epsilon \mathbf{L}^{(2)}$.

We now introduce scalar versions of product block-Laplacian matrices that will be useful.

Definition G.13. The product demand graph of a vector $\mathbf{d} \in (\mathfrak{R}_{>0})^n$, $G(\mathbf{d})$, is a complete weighted graph on n vertices whose weight between vertices i and j is given by $w_{ij} = \mathbf{d}_i \mathbf{d}_j$.

The Laplacian of $G(\mathbf{d})$, denoted $\mathbf{L}_{G(\mathbf{d})}$ is called the product demand Laplacian of \mathbf{d} .

Definition G.14. The bipartite product demand graph of two vectors $\mathbf{d}^A \in (\mathfrak{R}_{>0})^{|A|}$, $\mathbf{d}^B \in (\mathfrak{R}_{>0})^{|B|}$, $G(\mathbf{d}^A, \mathbf{d}^B)$, is a weighted bipartite graph on vertices $A \cup B$, whose weight between vertices $i \in A$ and $j \in B$ is given by $w_{ij} = \mathbf{d}_i^A \mathbf{d}_j^B$.

The Laplacian of $G(\mathbf{d}^A, \mathbf{d}^B)$, denoted $\mathbf{L}_{G(\mathbf{d}^A, \mathbf{d}^B)}$ is called the bipartite product demand Laplacian of $(\mathbf{d}^A, \mathbf{d}^B)$.

In Section J, we give results on efficiently constructing approximations to product demand Laplacians that can be summarized as follows:

Lemma G.15. *There is a routine WEIGHTEDEXPANDER(\mathbf{d}, ϵ) such that for any $\epsilon > 0$, and a demand vector $\mathbf{d} \in (\mathfrak{R}_{>0})^n$, WEIGHTEDEXPANDER(\mathbf{d}, ϵ) returns in $O(n\epsilon^{-4})$ work and $O(\log n)$ depth a graph H with $O(n\epsilon^{-4})$ edges such that*

$$\mathbf{L}_H \approx_\epsilon \mathbf{L}_{G(\mathbf{d})}.$$

Lemma G.16. *There is a routine $\text{WEIGHTEDBIPARTITEEXPANDER}(\mathbf{d}^A, \mathbf{d}^B, \epsilon)$ such that for any demand vectors \mathbf{d}^A and \mathbf{d}^B of total length n , and a parameter ϵ , it returns in $O(n\epsilon^{-4})$ work and $O(\log n)$ depth a bipartite graph H between A and B with $O(n\epsilon^{-4})$ edges such that*

$$\mathbf{L}_H \approx_\epsilon \mathbf{L}_{G(\mathbf{d}^A, \mathbf{d}^B)}.$$

Finally, we need to define an operation on graphs: Given a graph G , define $G^{(K^2)}$ to be the graph obtained by duplicating each vertex in G , and for each edge (i, j) in G , add a 2×2 bipartite clique between the two copies of i and j .

$\mathbf{L} = \text{CLIQUESPARSIFICATION}(\mathbf{d}, \epsilon)$

1. Initialize $\mathbf{L} \in (\mathbb{C}^{r \times r})^{n \times n}$ to 0.
2. Construct \mathbf{w} where $\mathbf{w}_i = \|\mathbf{d}_{[i]}\|$.
3. $H \leftarrow \text{WEIGHTEDEXPANDER}(\mathbf{w}, \epsilon)$.
4. For each edge $(i, j) \in H$, with weight $h_{i,j}$,
 - (a) $\mathbf{L}_{[i,i]} \leftarrow \mathbf{L}_{[i,i]} + h_{i,j} \mathbb{I}_r$.
 - (b) $\mathbf{L}_{[j,j]} \leftarrow \mathbf{L}_{[j,j]} + h_{i,j} \mathbb{I}_r$.
 - (c) $\mathbf{L}_{[i,j]} \leftarrow \mathbf{L}_{[i,j]} - \frac{h_{i,j}}{\mathbf{w}_i \mathbf{w}_j} \mathbf{d}_{[i]} \mathbf{d}_{[j]}^*$.
 - (d) $\mathbf{L}_{[j,i]} \leftarrow \mathbf{L}_{[j,i]} - \frac{h_{i,j}}{\mathbf{w}_i \mathbf{w}_j} \mathbf{d}_{[j]} \mathbf{d}_{[i]}^*$.
5. Return \mathbf{L} .

Figure 8: Pseudocode for Sparsifying Product Demand Block-Laplacians

We now combine the above construction of sparsifiers for product demand graphs with Lemma G.12 to efficiently construct sparse approximations to product demand block-Laplacians.

Proof. (of Lemma G.6) The procedure $\text{CLIQUESPARSIFICATION}(\mathbf{d}, \epsilon)$ returns the matrix \mathbf{L} where

$$\mathbf{L}_{[i,j]} = \begin{cases} -\frac{h_{i,j}}{\|\mathbf{d}_{[i]}\| \|\mathbf{d}_{[j]}\|} \mathbf{d}_{[i]} \mathbf{d}_{[j]}^* & i \neq j, \\ \mathbb{I}_r \cdot \sum_{k:k \neq i} h_{i,k} & \text{otherwise.} \end{cases}$$

By construction $\mathbf{L}_H \approx_\epsilon \mathbf{L}_{G(\mathbf{w})}$. Thus, applying Lemma G.12, with $H^{(1)} = H$ and $H^{(2)} = G(\mathbf{w})$, we know that $\mathbf{L} \approx_\epsilon \mathbf{L}^{(2)}$, where $\mathbf{L}^{(2)}$ is given by

$$\mathbf{L}^{(2)}_{[i,j]} = \begin{cases} -\frac{\mathbf{w}_i \mathbf{w}_j}{\|\mathbf{d}_{[i]}\| \|\mathbf{d}_{[j]}\|} \mathbf{d}_{[i]} \mathbf{d}_{[j]}^* & i \neq j, \\ \mathbb{I}_r \cdot \mathbf{w}_i \sum_{k:k \neq i} \mathbf{w}_k & \text{otherwise.} \end{cases}$$

Since $\mathbf{w}_i = \|\mathbf{d}_{[i]}\|$, we have $\mathbf{L}^{(2)} = \mathbf{L}_{G(\mathbf{d})}$, proving our claim. \square

Proof. (of Lemma G.7) The procedure $\text{BIPARTITECLIQUESPARSIFICATION}(\mathbf{d}, F, \epsilon)$ returns the matrix \mathbf{L} where

$$\mathbf{L}_{[i,j]} = \begin{cases} -\frac{h_{i,j}}{\|\mathbf{d}_{[i]}\| \|\mathbf{d}_{[j]}\|} \mathbf{d}_{[i]} \mathbf{d}_{[j]}^* & i \neq j, \\ \mathbb{I}_r \cdot \sum_{k:k \neq i} h_{i,k} & \text{otherwise.} \end{cases}$$

$\mathbf{L} = \text{BIPARTITECLIQUESPARSIFICATION}(\mathbf{b}, F, \epsilon)$

1. Initialize $\mathbf{L} \in (\mathbb{C}^{r \times r})^{n \times n}$ to 0.
2. Construct \mathbf{w} where $w_i = \|\mathbf{d}_{[i]}\|$. Let $C = V \setminus F$.
3. $H \leftarrow \text{WEIGHTEDBIPARTITEEXPANDER}(\mathbf{w}|_F, \mathbf{w}|_C, \epsilon)$.
4. For each edge $(i, j) \in H$, with weight $h_{i,j}$,
 - (a) $\mathbf{L}_{[i,i]} \leftarrow \mathbf{L}_{[i,i]} + h_{i,j} \mathbb{I}_r$.
 - (b) $\mathbf{L}_{[j,j]} \leftarrow \mathbf{L}_{[j,j]} + h_{i,j} \mathbb{I}_r$.
 - (c) $\mathbf{L}_{[i,j]} \leftarrow \mathbf{L}_{[i,j]} - \frac{h_{i,j}}{w_i w_j} \mathbf{d}_{[i]} \mathbf{d}_{[j]}^*$.
 - (d) $\mathbf{L}_{[j,i]} \leftarrow \mathbf{L}_{[j,i]} - \frac{h_{i,j}}{w_i w_j} \mathbf{d}_{[j]} \mathbf{d}_{[i]}^*$.
5. Return \mathbf{L} .

Figure 9: Pseudocode for Sparsifying Bipartite Product Demand Block-Laplacians

Since H returned by `WEIGHTEDBIPARTITEEXPANDER` is guaranteed to be bipartite using Lemma G.16, we obtain that $(\mathbf{L}_H)_{[F,F]}$, and $(\mathbf{L}_H)_{[C,C]}$ are block-diagonal.

By construction $\mathbf{L}_H \approx_\epsilon \mathbf{L}_{G(\mathbf{w}|_F, \mathbf{w}|_C)}$. Thus, applying Lemma G.12, with $H^{(1)} = H$ and $H^{(2)} = G(\mathbf{w}|_F, \mathbf{w}|_C)$, we know that $\mathbf{L} \approx_\epsilon \mathbf{L}^{(2)}$, where $\mathbf{L}^{(2)}$ is given by

$$\mathbf{L}^{(2)}_{[i,j]} = \begin{cases} \mathbb{I}_r \cdot \mathbf{w}_i \sum_{k \in C} \mathbf{w}_k & i = j, \text{ and } i \in F, \\ \mathbb{I}_r \cdot \mathbf{w}_i \sum_{k \in F} \mathbf{w}_k & i = j, \text{ and } i \in C, \\ 0 & i \neq j, \text{ and } (i, j \in F \text{ or } i, j \in C), \\ -\frac{w_i w_j}{\|\mathbf{d}_{[i]}\| \|\mathbf{d}_{[j]}\|} \mathbf{d}_{[i]} \mathbf{d}_{[j]}^* & i \neq j, \text{ otherwise.} \end{cases}$$

Since $w_i = \|\mathbf{d}_{[i]}\|$, we have $\mathbf{L}^{(2)} = \mathbf{L}_{G(\mathbf{d}|_F, \mathbf{d}|_C)}$, proving our claim. \square

Finally, we give a proof of Lemma G.12.

Proof. (of Lemma G.12) Using Fact A.2, we can write each $\mathbf{d}_{[i]}$ as $\frac{1}{2} \|\mathbf{d}_{[i]}\| (\mathbf{Q}_i^{(1)} + \mathbf{Q}_i^{(2)})$, where $\mathbf{Q}_i^{(1)} (\mathbf{Q}_i^{(1)})^* = \mathbf{Q}_i^{(2)} (\mathbf{Q}_i^{(2)})^* = \mathbb{I}_r$. Construct the matrix $\mathbf{F} \in (\mathbb{C}^{r \times r})^{2n \times 2n}$ as follows:

$$\mathbf{F} = \begin{bmatrix} \mathbf{Q}_1^{(1)} & \mathbf{Q}_1^{(2)} & 0 & 0 & \dots \\ 0 & 0 & \mathbf{Q}_2^{(1)} & \mathbf{Q}_2^{(2)} & \dots \\ \vdots & \vdots & & & \ddots \end{bmatrix}.$$

Now, observe that for each $\ell = 1, 2$,

$$\mathbf{L}^{(\ell)}_{[i,i]} = \frac{1}{2} \left(\sum_{k: k \neq i} h_{i,k}^{(\ell)} \right) \cdot (\mathbf{Q}_i^{(1)} (\mathbf{Q}_i^{(1)})^* + \mathbf{Q}_i^{(2)} (\mathbf{Q}_i^{(2)})^*) = \frac{1}{4} \left(\mathbf{F} (\mathbf{L}_{(H^{(\ell)})^{(K2)}} \otimes \mathbb{I}_r) \mathbf{F}^* \right)_{[i,i]},$$

where $\mathbf{L}_{(H^{(\ell)})^{(K2)}}$ is the Laplacian of the graph $(H^{(\ell)})^{(K2)}$ defined above, and \otimes tensor product. Also, for $i \neq j$,

$$\mathbf{L}^{(\ell)}_{[i,j]} = -\frac{h_{i,j}^{(\ell)}}{\mathbf{w}_i \mathbf{w}_j} \mathbf{d}_{[i]} \mathbf{d}_{[j]}^* = -\frac{h_{i,j}^{(\ell)}}{4} \cdot \left(\mathbf{Q}_i^{(1)} + \mathbf{Q}_i^{(2)} \right) \left(\mathbf{Q}_j^{(1)} + \mathbf{Q}_j^{(2)} \right)^* = \frac{1}{4} \left(\mathbf{F} \left(\mathbf{L}_{(H^{(\ell)})^{(K2)}} \otimes \mathbb{I}_r \right) \mathbf{F}^* \right)_{[i,j]},$$

Thus,

$$\mathbf{L}^{(\ell)} = \frac{1}{4} \mathbf{F} \left(\mathbf{L}_{(H^{(\ell)})^{(K2)}} \otimes \mathbb{I}_r \right) \mathbf{F}^*.$$

By assumption, we have $\mathbf{L}_{H^{(1)}} \approx_\epsilon \mathbf{L}_{H^{(2)}}$. Using Lemma G.17, we get that $\mathbf{L}_{H^{(1)}^{(K2)}} \approx_\epsilon \mathbf{L}_{H^{(2)}^{(K2)}}$. This implies

$$\mathbf{L}_{(H^{(1)})^{(K2)}} \otimes \mathbb{I}_r \approx_\epsilon \mathbf{L}_{(H^{(2)})^{(K2)}} \otimes \mathbb{I}_r,$$

and thus, $\mathbf{L}^{(1)} \approx_\epsilon \mathbf{L}^{(2)}$. □

G.5 Constructing sparsifiers for lifts of graphs

Given a graph G , define $G^{(K2)}$ to be the graph obtained by duplicating each vertex in G , and for each edge (i, j) in G , add a 2×2 bipartite clique between the two copies of i and j .

Lemma G.17. *If H is a sparsifier for G , i.e., $L_G \approx_\epsilon L_H$, then $H^{(K2)}$ is a sparsifier for $G^{(K2)}$, i.e. $L_{G^{(K2)}} \approx_\epsilon L_{H^{(K2)}}$.*

Proof. Since $L_G \approx_\epsilon L_H$, we have $e_i^\top L_G e_i \approx_\epsilon e_i^\top L_H e_i$. Thus, if D_G denotes the diagonal matrix of degrees of G , we have $D_G \approx_\epsilon D_H$. The Laplacian for $G^{(K2)}$ is

$$L_{G^{(K2)}} = L_G \otimes \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} + D_G \otimes \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}.$$

Similarly,

$$L_{H^{(K2)}} = L_H \otimes \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} + D_H \otimes \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}.$$

Observe that we can write

$$L_G \otimes \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 & 0 & \dots \\ 0 & 0 & 1 & 1 & \dots \\ \vdots & \vdots & & \ddots & \end{bmatrix} L_G \begin{bmatrix} 1 & 1 & 0 & 0 & \dots \\ 0 & 0 & 1 & 1 & \dots \\ \vdots & \vdots & & \ddots & \end{bmatrix}^\top.$$

Since $L_G \approx_\epsilon L_H$, this implies

$$L_G \otimes \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \approx_\epsilon L_H \otimes \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}.$$

Similarly, we get,

$$D_G \otimes \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \approx_\epsilon D_H \otimes \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}.$$

Adding the above two, we get $L_{G^{(K2)}} \approx_\epsilon L_{H^{(K2)}}$. □

H Estimating Leverage Scores by Undersampling

We will control the densities of all intermediate bDD matrices using the uniform sampling technique introduced by [CLM⁺14]. It relies on sampling columns³ of matrices by upper bounds of their true

³The randomized numerical linear algebra literature, e.g. [CLM⁺14], typically samples rows instead of columns of matrices. We sample columns instead in order to use a more natural set of notations.

leverage scores,

$$\tau_i(\mathbf{A}) = \mathbf{a}_i^T (\mathbf{A}\mathbf{A}^*)^{-1} \mathbf{a}_i.$$

These upper bounds are measured w.r.t. a different matrix, giving generalized leverage scores of the form:

$$\tau_i^{\mathbf{B}}(\mathbf{A}) = \begin{cases} \mathbf{a}_i^* (\mathbf{B}\mathbf{B}^*)^{-1} \mathbf{a}_i & \text{if } \mathbf{a}_i \perp \ker(\mathbf{B}), \\ 1 & \text{otherwise.} \end{cases}$$

We introduced *unitary edge-vertex transfer matrices* in Definition B.4. Sparsifying bDD matrices can be transformed into the more general setting described above via a unitary edge-vertex transfer matrix, which is analogous to the edge-vertex incidence matrix.

Lemma B.5 proves that every bDD matrix $\mathbf{M} \in (\mathbb{C}^{r \times r})^{n \times n}$ with m nonzero off-diagonal blocks can be written as $\mathbf{X} + \mathbf{B}\mathbf{B}^*$ where $\mathbf{B} \in (\mathbb{C}^{r \times r})^{n \times 2m}$ is a unitary edge-vertex transfer matrix and \mathbf{X} is a block diagonal PSD matrix. Additionally, for every block diagonal matrix \mathbf{Y} s.t. $\mathbf{M} - \mathbf{Y}$ is bDD, we have $\mathbf{X} \succcurlyeq \mathbf{Y}$. This decomposition can be found in $O(m)$ time and $O(\log n)$ depth. We rely on this \mathbf{X} to detect some cases of high leverage scores as samples of \mathbf{B} may have lower rank.

We will reduce the number of nonzero blocks in \mathbf{M} by sampling columns blocks from this matrix. This is more restrictive than sampling individual columns. Nonetheless, it can be checked via matrix concentration bounds [AW02, Tro12] that it suffices to sample the block by analogs of leverage scores:

$$\tau_{[i]} = \text{tr} \left(\mathbf{B}_{[i]}^* (\mathbf{X} + \mathbf{B}\mathbf{B}^*)^{-1} \mathbf{B}_{[i]} \right) \quad (37)$$

As in [CLM⁺14], we recursively estimate upper bounds for these scores, leading to the pseudocode given in Figure 10.

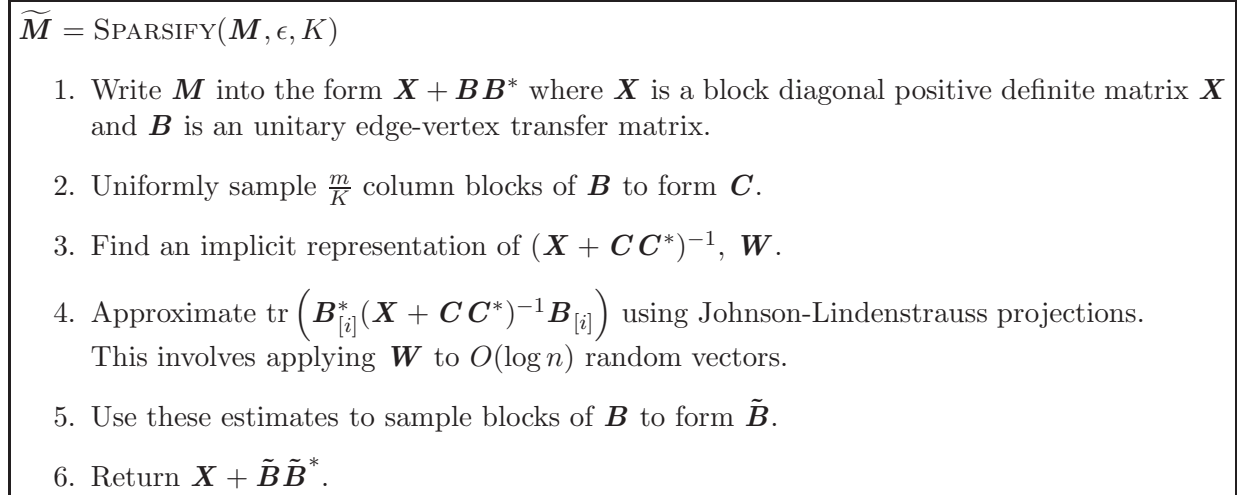


Figure 10: Sparsification via. Uniform Sampling

Lemma H.1. *Given a positive definite bDD matrix $\mathbf{M} \in (\mathbb{C}^{r \times r})^{n \times n}$ with m nonzero blocks.*

Assume that for any positive definite bDD matrix $\mathbf{M}' \in (\mathbb{C}^{r \times r})^{n \times n}$ with \widehat{m} nonzero blocks, we can find an implicit representation of a matrix \mathbf{W} such that $\mathbf{W} \approx_1 (\mathbf{M}')^{-1}$ in depth $d_{\text{con}}(\widehat{m}, n)$ and work $w_{\text{con}}(\widehat{m}, n)$ and for any vector \mathbf{b} , we can evaluate $\mathbf{W}\mathbf{b}$ in depth $d_{\text{eval}}(\widehat{m}, n)$ and work $w_{\text{eval}}(\widehat{m}, n)$.

For any $K \geq 1$, $1 \geq \epsilon > 0$, the algorithm $\text{SPARSIFY}(\mathbf{M}, \epsilon, K)$ outputs an explicit positive definite bDD matrix $\widetilde{\mathbf{M}}$ with $O(Kn \log n / \epsilon^2)$ nonzero blocks and $\widetilde{\mathbf{M}} \approx_\epsilon \mathbf{M}$.

Also, this algorithm runs in

$$d_{\text{con}}\left(\frac{m}{K}, n\right) + O\left(d_{\text{eval}}\left(\frac{m}{K}, n\right) + \log n\right)$$

depth and

$$w_{\text{con}}\left(\frac{m}{K}, n\right) + O\left(w_{\text{eval}}\left(\frac{m}{K}, n\right) \log n + m \log n\right)$$

work.

The guarantees of this process can be obtained from the main result from [CLM⁺14]:

Theorem H.2 (Theorem 3 from [CLM⁺14]). *Let \mathbf{A} be a n by m matrix, and $\alpha \in (0, 1]$ a density parameter. Consider the matrix \mathbf{A}' consisting of a random α fraction of the columns of \mathbf{A} . Then, with high probability we have $\mathbf{u}_i = \tau_i^{\mathbf{A}'}(\mathbf{A})$ is a vector of leverage score overestimates, i.e. $\tau_i(\mathbf{A}) \leq \mathbf{u}_i$, and*

$$\sum_{i=1}^m \mathbf{u}_i \leq \frac{3n}{\alpha}.$$

Proof. (Sketch of Lemma H.1) Instead of sampling m/K column blocks, consider sampling m/K individual columns to form $\widehat{\mathbf{C}}$. Theorem H.2 gives that the total leverage scores of all the columns of \mathbf{B} w.r.t. $\widehat{\mathbf{C}}$ is at most $O(nKr^2)$.

As \mathbf{C} is formed by taking blocks instead of columns, we have

$$\mathbf{X} + \mathbf{C}\mathbf{C}^* \succcurlyeq \mathbf{X} + \widehat{\mathbf{C}}\widehat{\mathbf{C}}^*,$$

so the individual leverage scores computed w.r.t. \mathbf{C} (and in turn \mathbf{W}) sums up to less than the ones computed w.r.t. $\widehat{\mathbf{C}}$.

Let us use $\mathbf{b}_{[i],j} \in \mathbb{C}^n$ to denote the j^{th} column of the $\mathbf{B}_{[i]}$ block column. Note that

$$\mathbf{B}_{[i]}\mathbf{B}_{[i]}^* = \sum_j \mathbf{b}_{[i],j}\mathbf{b}_{[i],j}^*.$$

Define $\tau_{[i],j}^{\mathbf{X}+\mathbf{C}\mathbf{C}^*} = \text{tr}\left(\mathbf{b}_{[i],j}^*(\mathbf{X} + \mathbf{C}\mathbf{C}^*)^{-1}\mathbf{b}_{[i],j}\right)$.

We then have

$$\text{tr}\left(\mathbf{B}_{[i]}^*(\mathbf{X} + \mathbf{C}\mathbf{C}^*)^{-1}\mathbf{B}_{[i]}\right) = \sum_j \tau_{[i],j}^{\mathbf{X}+\mathbf{C}\mathbf{C}^*},$$

so the total sum of the estimates we obtain is at most $O(nr^2K)$, which gives $\widetilde{\mathbf{B}}$ with $O(Kn \log n \epsilon^{-2})$ blocks.

To approximate $\sum_j \tau_{[i],j}^{\mathbf{X}+\mathbf{C}\mathbf{C}^*}$, we apply a standard technique given by [SS11c, AT11, LMP13]. The rough idea is to write

$$\tau_{[i],j}^{\mathbf{X}+\mathbf{C}\mathbf{C}^*} = \left\| \mathbf{C}^*(\mathbf{X} + \mathbf{C}\mathbf{C}^*)^{-1}\mathbf{b}_{[i],j} \right\|^2 + \left\| \sqrt{\mathbf{X}}(\mathbf{X} + \mathbf{C}\mathbf{C}^*)^{-1}\mathbf{b}_{[i],j} \right\|^2$$

By Johnson-Lindenstrauss Lemma, for a random Gaussian matrix \mathbf{G} with $\Theta(\log(n))$ rows, we know that

$$\tau_{[i],j}^{\mathbf{X}+\mathbf{C}\mathbf{C}^*} \approx_2 \left\| \mathbf{G}\mathbf{C}^*(\mathbf{X} + \mathbf{C}\mathbf{C}^*)^{-1}\mathbf{b}_{[i],j} \right\|^2 + \left\| \mathbf{G}\sqrt{\mathbf{X}}(\mathbf{X} + \mathbf{C}\mathbf{C}^*)^{-1}\mathbf{b}_{[i],j} \right\|^2$$

for high probability. Since \mathbf{G} has $O(\log(n))$ rows, this can be approximated by applying the approximate inverse \mathbf{W} to $O(\log(n))$ vectors. Hence, step 3 runs in $O(d_{\text{eval}}(\frac{m}{K}, n) + \log n)$ depth and $O(w_{\text{eval}}(\frac{m}{K}, n) \log n + m \log n)$ work. \square

We remark that the extra factor of r^2 can likely be improved by modifying the proof of Theorem H.2 to work with blocks. We omit this improvement for simplicity, especially since we're already treating r as a constant.

I Recursive Construction of Schur Complement Chains

We now give the full details for the recursive algorithm that proves the running times as stated in Theorem 1.2:

Theorem 1.2 (Sparsified Multigrid). *There is an algorithm that, when given a bDD matrix \mathbf{M} with n block rows and m nonzero blocks, produces a solver for \mathbf{M} in $O(m \log n + n \log^{2+o(1)} n)$ work and $O(n^{o(1)})$ depth so that the solver finds ϵ -approximate solutions to systems of equations in \mathbf{M} in $O((m + n \log^{1+o(1)} n) \log(1/\epsilon))$ work and $O(\log^2 n \log \log n \log(1/\epsilon))$ depth.*

This argument can be viewed as a more sophisticated version of the one in Section 3.6. The main idea is to invoke routines for reducing edges and vertices in a slightly unbalanced recursion, where several steps of vertex reductions take place before a single edge reduction. The routines that we will call are:

1. BDDSUBSET given by Lemma 3.5 proven in Section D for finding a large set of α -bDD subset.
2. APPROXSCHUR given by Lemma G.1 proven in Section G that gives sparse approximations to Schur complements.
3. SPARSIFY given by Lemma H.1 proven in Section H that allows us to sparsify a bDD matrix by recursing on a uniform subsample of its non-zero blocks.

An additional level of complication comes from the fact that the approximation guarantees of our constructions rely on gradually decreasing errors down the Schur complement chain. This means that the density increases faster and larger reduction factors K are required as the iteration goes on. Pseudocode of our algorithm is given in Figure 11.

Note that since $\mathbf{M}^{(0)}$ may be initially dense, we first make a recursive call to SPARSIFY before computing the approximate Schur complements.

In our analysis, we will use $n^{(i)}$ to denote the number of non-zero column/row blocks in $\mathbf{M}^{(i)}$, and $m^{(i)}$ to denote the number of non-zero blocks in $\mathbf{M}^{(i)}$. These are analogous to dimension and number of non-zeros in the matrix. It is also useful to refer to the steps between calls to SPARSIFY as phases. Specifically, phase j consists of iterations $(j-1)k$ to $jk-1$. We will also use K_j to denote the reduction factor used to perform the sparsification at the start of phase j , aka.

$$K_j = 2^{2ck \log^2((j-1)k+1)}.$$

A further technicality is that Lemma H.1 require a strictly positive definite block-diagonal part to facilitate the detection of vectors in the null space of the sample. We do so by padding \mathbf{M} with a small copy of the identity, and will check below that this copy stays throughout the course of this algorithm.

Lemma I.1. *For any bDD matrix $\mathbf{M}^{(0)}$ expressible as $\xi \mathbb{I} + \mathbf{B}^{(0)}(\mathbf{B}^{(0)})^*$ for some $\xi > 0$ and unitary edge-vertex transfer matrix \mathbf{B} , all intermediate matrices \mathbf{M}' generated during the algorithm RECURSIVECONSTRUCT($\mathbf{M}^{(0)}$) can be expressed as $\xi \mathbb{I} + \mathbf{B}'(\mathbf{B}')^*$.*

Proof. There are two mechanisms by which this recursive algorithm generates new matrices: through uniform sampling within SPARSIFY and via APPROXSCHUR. We will show inductively down the algorithmic calls that all matrices satisfy this property.

$(\mathbf{M}^{(1)}, \mathbf{M}^{(2)}, \dots; F_1, F_2, \dots) = \text{RECURSIVECONSTRUCT}(\mathbf{M}^{(0)})$

1. Initialize:
 - (a) $i \leftarrow 0$, and
 - (b) k and c are parameters to be decided.
2. While $\mathbf{M}^{(i)}$ has more than $\Theta(1)$ blocks,
 - (a) If $i \bmod k = 0$, Then
 - i. $\mathbf{M}^{(i)} \leftarrow \text{SPARSIFY}(\mathbf{M}^{(i)}, (i+8)^{-2}, 2^{2ck \log^2(i+1)})$.
 - (b) $F_{i+1} \leftarrow \text{BDDSUBSET}(\mathbf{M}^{(i)}, 4)$.
 - (c) $\mathbf{M}^{(i+1)} \leftarrow \text{APPROXSCHUR}(\mathbf{M}^{(i)}, F_{i+1}, 4, (i+8)^{-2})$.
 - (d) $i \leftarrow i + 1$.

Figure 11: Pseudocode for Recursively Constructing Schur Complement Chains, the recursive calls happen through SPARSIFY, which constructs its own Schur Complement Chains to perform the sparsification

Lemma H.1 gives that this \mathbf{X} is preserved in the sample.

For calls to $\text{APPROXSCHUR}(\mathbf{M}, C)$, the inductive hypothesis gives that $\mathbf{M}_{[C,C]}$ is expressible as $\widehat{\mathbf{M}} + \xi \mathbb{I}_C$. The last condition in Lemma 3.9 then gives that the result also has a $\xi \mathbb{I}_C$ part, which gives the inductive hypothesis for it as well. \square

As the interaction between this padding and our recursion is minimal, we will simply state the input conditions as $\mathbf{M} = \xi \mathbb{I} + \mathbf{B}\mathbf{B}^*$ for some $\xi > 0$. We will first bound the sizes of the matrices in the Schur complement chain produced by RECURSIVECONSTRUCT.

Lemma I.2. *For any bDD matrix $\mathbf{M}^{(0)} = \xi \mathbb{I} + \mathbf{B}^{(0)}(\mathbf{B}^{(0)})^*$ for some $\xi > 0$, the algorithm $\text{RECURSIVECONSTRUCT}(\mathbf{M}^{(0)})$ returns a Schur complement chain $(\mathbf{M}^{(1)}, \mathbf{M}^{(2)}, \dots; F_1, F_2, \dots)$ such that:*

1. $n^{(i)} \leq \beta^i n^{(0)}$ for some absolute constant $\beta < 1$.
2. The number of non-zero blocks in any iteration i of phase j is at most $2^{3ck \log^2(jk)} n^{((j-1)k)} \log n^{((j-1)k)}$.

Proof. Lemma 3.5 and the choice of $\alpha = 4$ ensures $|F_k| = \Omega(n^{(i)})$, which means there is constant $\beta < 1$ such that $n^{(i)} \leq \beta^k n$. Furthermore, we do not increase vertex count at any point in this recursion, and all recursive calls are made to smaller graphs. Therefore, the recursion terminates.

Lemma G.1 shows that after computing each approximate Schur complement,

$$\begin{aligned} m^{(i+1)} &= O(m^{(i)}(i^2 \log(i+8))^{O(\log(i+8))}) \\ &\leq 2^{O(\log^2(i+1))} m^{(i)}. \end{aligned}$$

Hence, by picking c appropriately we can guarantee that the density increases by a factor of most $2^{c \log^2(i+1)}$ during each iteration. This size increase is controlled by calls to SPARSIFY. Specifically,

Lemma H.1 gives that at the start of phase j we have:

$$\frac{m^{((j-1)k)}}{n^{((j-1)k)} \log n^{((j-1)k)}} \leq 2^{2ck \log^2((j-1)k)}.$$

Then, since we go at most k steps without calling SPARSIFY, this increase in density can be bounded by:

$$2^{2ck \log^2((j-1)k)} 2^{c \log^2(jk-1)} \dots 2^{c \log^2((j-1)k+1)} \leq 2^{3ck \log^2(jk)}.$$

□

Lemma I.3. *For any bDD matrix $\mathbf{M}^{(0)} = \xi \mathbb{I} + \mathbf{B}^{(0)} (\mathbf{B}^{(0)})^*$ for some $\xi > 0$ that has n blocks, the algorithm RECURSIVECONSTRUCT($\mathbf{M}^{(0)}$) returns a Schur Complement Chain $(\mathbf{M}^{(1)}, \mathbf{M}^{(2)}, \dots; F_1, F_2, \dots)$ whose corresponding the linear operator \mathbf{W} satisfies*

$$\mathbf{W} \approx_{O(1)} \left(\mathbf{M}^{(0)} \right)^{-1}.$$

Also, we can evaluate $\mathbf{W}\mathbf{b}$ in $O(\log^2 n \log \log n)$ depth and $2^{O(k \log^2 k)} n \log n$ work for any vector \mathbf{b} .

Proof. We first bound the quality of approximation between \mathbf{W} and $\left(\mathbf{M}^{(0)} \right)^{-1}$. The approximate Schur complement $\mathbf{M}^{(i+1)}$ was constructed so that $\mathbf{M}^{(i+1)} \approx_{(i+8)^{-2}} \text{Sc} \left(\mathbf{M}^{(i)}, F_i \right)$. The other source of error, SPARSIFY, is called only for some i . In those iterations, Lemma H.1 guarantee that $\mathbf{M}^{(i)}$ changes only by $(i+8)^{-2}$ factor. This means overall we have $\mathbf{M}^{(i+1)} \approx_{2(i+8)^{-2}} \text{Sc} \left(\mathbf{M}^{(i)}, F_i \right)$. By Lemma 3.3, we have:

$$\mathbf{W} \approx_{1/2+4 \sum_i (i+8)^{-2}} \left(\mathbf{M}^{(0)} \right)^{-1},$$

and it can be checked that $\sum_i (i+8)^{-2}$ is a constant.

The cost of APPLYCHAIN is dominated by the sequence of calls to JACOBI. As each F_i is chosen to be 4-bDD, the number of iterations required is $O(\log(\epsilon_i)) = O(\log i)$. As matrix-vector multiplications take $O(\log n)$ depth, the total depth can be bounded by.

$$O \left(\sum_i \log i \log n^{(i)} \right) = O(\log^2 n \log \log n)$$

The total work of these steps depend on the number of non-zero blocks, $m^{(i)}$. Substituting the bounds from Lemma I.2 into JACOBI gives a total of:

$$O \left(\sum_i 2^{3ck \log^2(i+k)} n^{(i)} \log n^{(i)} \log i \right) \leq 2^{O(k \log^2 k)} n \log n$$

where the inequality follows from the fact that the $n^{(i)}$ s are geometrically decreasing. □

This allows us to view the additional overhead of SPARSIFY as a black box, and analyze the total cost incurred by the non-recursive parts of RECURSIVECONSTRUCT during each phase.

Lemma I.4. *The total cost of RECURSIVECONSTRUCT during phase j , including the linear system solves made by SPARSIFY at iteration $(j-1)k$ (but not its recursive invocation to RECURSIVECONSTRUCT) is*

$$2^{O(k \log^2(jk))} n^{((j-1)k)} \log^2 n^{((j-1)k)}$$

in work and

$$O\left(k \log(jk) \log^2 n^{((j-1)k)}\right) + O\left(\log^2 n^{((j-1)k)} \log \log n^{((j-1)k)}\right)$$

in depth.

Proof. Let $m^{(i)}$ and $n^{(i)}$ be the number of non zeros and variables in $\mathbf{M}^{(i)}$ before the SPARSIFY call if there is. Lemmas 3.5 and G.1 show that the i^{th} iteration takes $O(m^{(i)} + m^{(i+1)})$ work and $O(\log i \log n^{(i)})$ depth. By Lemma I.2 the cost during these iterations excluding the call to SPARSIFY is:

$$\begin{aligned} \sum_{i=(j-1)k}^{jk-1} O\left(m^{(i)} + m^{(i+1)}\right) &\leq \sum_{i=(j-1)k}^{jk-1} 2^{O(k \log^2 i)} n^{(i)} \log n^{(i)} \\ &\leq 2^{O(k \log^2(jk))} n^{((j-1)k)} \log n^{((j-1)k)} \end{aligned}$$

work and

$$\sum_{i=(j-1)k}^{jk-1} O\left(\log i \log n^{(i)}\right) \leq O(k \log((j-1)k) \log n^{((j-1)k)})$$

depth.

We now consider the call to SPARSIFY made during iteration $(j-1)k$. Access to a fast solver for the sampled bDD matrix is obtained via recursive calls to RECURSIVECONSTRUCT. The guarantees of the chain given by Lemma I.3 above means each solve takes depth

$$d_{eval} = \log^2 n^{((j-1)k)} \log \log n^{((j-1)k)},$$

and work

$$w_{eval} = n^{((j-1)k)} \log n^{((j-1)k)}$$

Incorporating these parameters into Lemma H.1 allows us to bound the overhead from these solves by

$$2^{O(k \log^2(jk))} n^{((j-1)k)} \log^2 n^{((j-1)k)}$$

work and

$$O\left(k \log(jk) \log n^{(jk)}\right) + O\left(\log^2 n^{((j-1)k)} \log \log \left(n^{((j-1)k)}\right)\right)$$

depth. □

Note that at the end of the j^{th} phase, the time required to construct an extra Schur complement chain for the SPARSIFY call is less than the remaining cost after the j^{th} phase. This is the reason why we use $2^{2ck \log^2(i+1)}$ as the reduction factor for the SPARSIFY call. The following theorem takes account for the recursive call and show the total running time for the algorithm.

Lemma I.5. *With high probability, RECURSIVECONSTRUCT $\left(\mathbf{M}^{(0)}\right)$ takes $2^{O(\log n/k)}$ depth and $m \log n + 2^{O(k \log^2 k)} n \log^2 n$ work.*

Proof. Lemma I.2 shows that the call to SPARSIFY at the start of each phase j requires the construction of an extra Schur complement chain on a matrix with $n^{((j-1)k)}$ row/column blocks and at most $2^{ck \log^2((j-1)k)} n^{((j-2)k)} \log n^{((j-2)k)}$ non-zeros blocks. The guarantees of Lemma H.1 gives that the number of non-zero block in the sparsified matrix is at most

$$C 2^{2ck \log^2((j-1)k)} n^{((j-1)k)} \log n^{((j-1)k)}$$

for some absolute constant C . Therefore the cost of this additional call is less than the cost of constructing the rest of the phases during the construction process. Therefore, every recursive call except the first one can be viewed as an extra branching factor of 2 at each subsequent phase.

Depth can be bounded by the total number of recursive invocations to `RECURSIVECONSTRUCT`. The fact that $n^{(i)}$ is geometrically decreasing means there are $O(\log n/k)$ phases. Choosing k so that $k \log^2 k = o(\log n)$ gives a depth of:

$$\begin{aligned}
& \sum_{j=1}^{O(\log n/k)} 2^j \left(k \log(jk) \log n^{(jk)} + \log^2 n^{(jk)} \log \log n^{(jk)} \right) \\
&= 2^{O(\log n/k)} O \left(k \log \log n \log n^{(last)} + \log^2 n^{(last)} \log \log n^{(last)} \right) \\
&= 2^{O(\log n/k)} O \left(k^2 \log \log n + k^2 \log k \right) \\
&= 2^{O(\log n/k)} k^2 \log \log(n) \\
&= 2^{O(\log n/k)}.
\end{aligned}$$

For bounding work, we start with the sparse case since all intermediate matrices generated during the construction process have density at most $2^{O(k \log^2 k)}$. In this case, the extra branching factor of 2 at each phase can be accounted for by weighting the cost of iteration j by 2^j , giving:

$$\begin{aligned}
& \sum_{j=1}^{O(\log n/k)} 2^j \left(2^{O(k \log^2(jk))} n^{(jk)} \log^2 n^{(jk)} \right) \\
&= 2^{O(k \log^2 k)} n \log^2 n.
\end{aligned}$$

For the dense case, the first recursive call to `SPARSIFY`($\mathbf{M}^{(0)}, O(1), 2^{2ck}$) is made to a graph whose edge count is much less. This leads to a geometric series, and an overhead of $O(m \log n)$ work at each step. As this can happen at most $O(\log n)$ times, it gives an additional factor of $O(\log n)$ in depth, which is still $2^{O(\log n/k)}$. The work obeys the recurrence:

$$W(m)_{dense} \leq \begin{cases} 2^{O(k \log^2 k)} n \log^2 n & \text{if } m \leq 2^{O(k \log^2 k)} n \log n, \text{ and} \\ W_{dense}(m/2) + m \log n + 2^{O(k \log^2 k)} n \log^2 n & \text{otherwise.} \end{cases}$$

which solves to:

$$W_{dense}(m) \leq O(m \log n) + 2^{O(k \log^2 k)} n \log^2 n.$$

□

To obtain Theorem 1.2, we simply choose an appropriate initial padding and set the parameter k .

Proof. (of Theorem 1.2) Lemma 2.4 allows us to solve the linear system $\mathbf{M} + \epsilon\mu\mathbb{I}$ instead. Invoking Lemma I.5 with $k = \log \log \log n$ gives a depth of $2^{O(\log n / \log \log \log n)}$, and work of $m \log n + 2^{O(\log \log \log n \log \log \log^4 n)} n \log^2 n$. The depth term can be simplified to $n^{o(1)}$ while the work term can be simplified to $2^{O(\log \log \log^2 n)} = \log^{o(1)} n$. □

J Weighted Expander Constructions

In this section, we give a linear time algorithm for computing linear sized spectral sparsifiers of complete and bipartite product demand graphs. These routines give Lemmas G.16 and G.15, which are crucial for controlling the densities of intermediate matrices in the spectral vertex sparsification

algorithm from Section G. Recall that the *product demand graph* with vertex set V and demands $\mathbf{d} : V \rightarrow \mathbb{R}_{>0}$ is the complete graph in which the weight of edge (u, v) is the product $d_u d_v$. Similarly, the *bipartite demand graph* with vertex set $U \cup V$ and demands $\mathbf{d} : U \cup V \rightarrow \mathbb{R}_{>0}$ is the complete bipartite graph on which the weight of the edge (u, v) is the product $d_u d_v$. Our routines are based on reductions to the unweighted, uniform case. In particular, we

1. Split all of the high demand vertices into many vertices that all have the same demand. This demand will still be the highest.
2. Given a graph in which almost all of the vertices have the same highest demand, we
 - a. drop all of the edges between vertices of lower demand,
 - b. replace the complete graph between the vertices of highest demand with an expander, and
 - c. replace the bipartite graph between the high and low demand vertices with a union of stars.
3. To finish, we merge back together the vertices that split off from each original vertex.

We start by showing how to construct the expanders that we need for step (2b). We state formally and analyze the rest of the algorithm for the complete case in the following two sections. We explain how to handle the bipartite case in Section J.3.

Expanders give good approximations to unweighted complete graphs, and our constructions will use the spectrally best expanders—Ramanujan graphs. These are defined in terms of the eigenvalues of their adjacency matrices. We recall that the adjacency matrix of every d -regular graph has eigenvalue d with multiplicity 1 corresponding to the constant eigenvector. If the graph is bipartite, then it also has an eigenvalue of $-d$ corresponding to an eigenvector that takes value 1 on one side of the bipartition and -1 on the other side. These are called the *trivial* eigenvalues. A d -regular graph is called a Ramanujan graph if all of its non-trivial eigenvalues have absolute value at most $2\sqrt{d-1}$. Ramanujan graphs were constructed independently by Margulis [Mar88] and Lubotzky, Phillips, and Sarnak [LPS88]. The following theorem and proposition summarizes part of their results.

Theorem J.1. *Let p and q be unequal primes congruent to 1 modulo 4. If p is a quadratic residue modulo q , then there is a non-bipartite Ramanujan graph of degree $p+1$ with $q^2(q-1)/2$ vertices. If p is not a quadratic residue modulo q , then there is a bipartite Ramanujan graph of degree $p+1$ with $q^2(q-1)$ vertices.*

The construction is explicit.

Proposition J.2. *If $p < q$, then the graph guaranteed to exist by Theorem J.1 can be constructed in parallel depth $O(\log n)$ and work $O(n)$, where n is its number of vertices.*

Sketch of proof. When p is a quadratic residue modulo q , the graph is a Cayley graph of $PSL(2, Z/qZ)$. In the other case, it is a Cayley graph of $PGL(2, Z/qZ)$. In both cases, the generators are determined by the $p+1$ solutions to the equation $p = a_0^2 + a_1^2 + a_2^2 + a_3^2$ where $a_0 > 0$ is odd and a_1, a_2 , and a_3 are even. Clearly, all of the numbers a_0, a_1, a_2 and a_3 must be at most \sqrt{p} . So, we can compute a list of all sums $a_0^2 + a_1^2$ and all of the sums $a_2^2 + a_3^2$ with work $O(p)$, and thus a list of all $p+1$ solutions with work $O(p^2) < O(n)$.

As the construction requires arithmetic modulo q , it is convenient to compute the entire multiplication table modulo q . This takes time $O(q^2) < O(n)$. The construction also requires the

computation of a square root of -1 modulo q , which may be computed from the multiplication table. Given this data, the list of edges attached to each vertex of the graph may be produced using linear work and logarithmic depth. \square

For our purposes, there are three obstacles to using these graphs:

1. They do not come in every degree.
2. They do not come in every number of vertices.
3. Some are bipartite and some are not.

We handle the first two issues by observing that the primes congruent to 1 modulo 4 are sufficiently dense. To address the third issue, we give a procedure to convert a non-bipartite expander into a bipartite expander, and *vice versa*.

An upper bound on the gaps between consecutive primes congruent to 1 modulo 4 can be obtained from the following theorem of Tchudakoff.

Theorem J.3 ([Tch36]). *For two integers a and b , let p_i be the i th prime congruent to a modulo b . For every $\epsilon > 0$,*

$$p_{i+1} - p_i \leq O(p_i^{3/4+\epsilon}).$$

Corollary J.4. *There exists an n_0 so that for all $n \geq n_0$ there is a prime congruent to 1 modulo 4 between n and $2n$.*

We now explain how we convert between bipartite and non-bipartite expander graphs. To convert a non-bipartite expander into a bipartite expander, we take its double-cover. We recall that if $G = (V, E)$ is a graph with adjacency matrix \mathbf{A} , then its double-cover is the graph with adjacency matrix

$$\begin{pmatrix} 0 & \mathbf{A} \\ \mathbf{A}^T & 0 \end{pmatrix}.$$

It is immediate from this construction that the eigenvalues of the adjacency matrix of the double-cover are the union of the eigenvalues of \mathbf{A} with the eigenvalues of $-\mathbf{A}$.

Proposition J.5. *Let G be a connected, d -regular graph in which all matrix eigenvalues other than d are bounded in absolute value by λ . Then, all non-trivial adjacency matrix eigenvalues of the double-cover of G are also bounded in absolute value by λ .*

To convert a bipartite expander into a non-bipartite expander, we will simply collapse the two vertex sets onto one another. If $G = (U \cup V, E)$ is a bipartite graph, we specify how the vertices of V are mapped onto U by a permutation $\pi : V \rightarrow U$. We then define the *collapse* of G induced by π to be the graph with vertex set U and edge set

$$\{(u, \pi(v)) : (u, v) \in E\}.$$

Note that the collapse will have self-loops at vertices u for which $(u, v) \in E$ and $u = \pi(v)$. We assign a weight of 2 to every self loop. When a double-edge would be created, that is when $(\pi(v), \pi^{-1}(u))$ is also an edge in the graph, we give the edge a weight of 2. Thus, the collapse can be a weighted graph.

Proposition J.6. *Let G be a d -regular bipartite graph with all non-trivial adjacency matrix eigenvalues bounded by λ , and let H be a collapse of G . Then, every vertex in H has weighted degree $2d$ and all adjacency matrix eigenvalues of H other than d are bounded in absolute value by 2λ .*

Proof. To prove the bound on the eigenvalues, let G have adjacency matrix

$$\begin{pmatrix} 0 & \mathbf{A} \\ \mathbf{A}^T & 0 \end{pmatrix}.$$

After possibly rearranging rows and columns, we may assume that the adjacency matrix of the collapse is given by

$$\mathbf{A} + \mathbf{A}^T.$$

Note that the self-loops, if they exist, correspond to diagonal entries of value 2. Now, let \mathbf{x} be a unit vector orthogonal to the all-1s vector. We have

$$\mathbf{x}^T(\mathbf{A} + \mathbf{A}^T)\mathbf{x} = \begin{pmatrix} \mathbf{x} \\ \mathbf{x} \end{pmatrix}^T \begin{pmatrix} 0 & \mathbf{A} \\ \mathbf{A}^T & 0 \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{x} \end{pmatrix} \leq \lambda \left\| \begin{pmatrix} \mathbf{x} \\ \mathbf{x} \end{pmatrix} \right\|^2 \leq 2\lambda,$$

as the vector $[\mathbf{x}; \mathbf{x}]$ is orthogonal to the eigenvectors of the trivial eigenvalues of the adjacency matrix of G . \square

We now state how bounds on the eigenvalues of the adjacency matrices of graphs lead to approximations of complete graphs and complete bipartite graphs.

Proposition J.7. *Let G be a graph with n vertices, possibly with self-loops and weighted edges, such that every vertex of G has weighted degree d and such that all non-trivial eigenvalues of the adjacency matrix of G have absolute value at most $\lambda \leq d/2$. If G is not bipartite, then $(n/d)\mathbf{L}_G$ is an ϵ -approximation of K_n for $\epsilon = (2 \ln 2)(\lambda)/d$. If G is bipartite, then $(n/d)\mathbf{L}_G$ is an ϵ -approximation of $K_{n,n}$ for $\epsilon = (2 \ln 2)(\lambda)/d$.*

Proof. Let \mathbf{A} be the adjacency matrix of G . Then,

$$\mathbf{L}_G = d\mathbf{I} - \mathbf{A}.$$

In the non-bipartite case, we observe that all of the non-zero eigenvalues of \mathbf{L}_{K_n} are n , so for all vectors x orthogonal to the constant vector,

$$x^T \mathbf{L}_{K_n} x = nx^T x.$$

As all of the non-zero eigenvalues of \mathbf{L}_G are between $d - \lambda$ and $d + \lambda$, for all vectors x orthogonal to the constant vector

$$n \left(1 - \frac{\lambda}{d}\right) x^T x \leq x^T (n/d)\mathbf{L}_G x \leq n \left(1 + \frac{\lambda}{d}\right) x^T x.$$

Thus,

$$\left(1 - \frac{\lambda}{d}\right) \mathbf{L}_{K_n} \preceq \mathbf{L}_G \preceq \left(1 + \frac{\lambda}{d}\right) \mathbf{L}_{K_n}.$$

In the bipartite case, we naturally assume that the bipartition is the same in both G and $K_{n,n}$. Now, let \mathbf{x} be any vector on the vertex set of G . Both the graphs $K_{n,n}$ and $(n/d)G$ have Laplacian matrix eigenvalue 0 with the constant eigenvector, and eigenvalue $2n$ with eigenvector $[1; -1]$. The other eigenvalues of the Laplacian of $K_{n,n}$ are n , while the other eigenvalues of the Laplacian of $(n/d)G$ are between

$$n \left(1 - \frac{\lambda}{d}\right) \quad \text{and} \quad n \left(1 + \frac{\lambda}{d}\right).$$

Thus,

$$\left(1 - \frac{\lambda}{d}\right) \mathbf{L}_{K_{n,n}} \preceq \mathbf{L}_G \preceq \left(1 + \frac{\lambda}{d}\right) \mathbf{L}_{K_{n,n}}.$$

The proposition now follows from our choice of ϵ , which guarantees that

$$e^{-\epsilon} \leq 1 - \lambda/d \quad \text{and} \quad 1 + \lambda/d \leq e^\epsilon,$$

provided that $\lambda/d \leq 1/2$. □

Lemma J.8. *There are algorithms that on input n and $\epsilon > n^{-1/6}$ produce a graph having $O(n/\epsilon^2)$ edges that is an $O(\epsilon)$ approximation of $K_{n'}$ or $K_{n',n'}$ for some $n \leq n' \leq 8n$. These algorithms run in $O(\log n)$ depth and $O(n/\epsilon^2)$ work.*

Proof. We first consider the problem of constructing an approximation of $K_{n',n'}$. By Corollary J.4 there is a constant n_0 so that if $n > n_0$, then there is a prime q that is equivalent to 1 modulo 4 so that $q^2(q-1)$ is between n and $8n$. Let q be such a prime and let $n' = q^2(q-1)$. Similarly, for ϵ sufficiently small, there is a prime p equivalent to 1 modulo 4 that is between $\epsilon^{-2}/2$ and ϵ^{-2} . Our algorithm should construct the corresponding Ramanujan graph, as described in Theorem J.1 and Proposition J.2. If the graph is bipartite, then Proposition J.7 tells us that it provides the desired approximation of $K_{n',n'}$. If the graph is not bipartite, then we form its double cover to obtain a bipartite graph and use Proposition J.5 and Proposition J.7 to see that it provides the desired approximation of $K_{n',n'}$.

The non-bipartite case is similar, except that we require a prime q so that $q^2(q-1)/2$ is between n and $8n$, and we use a collapse to convert a bipartite expander to a non-bipartite one, as analyzed in Proposition J.6. □

For the existence results in Section F, we just need to know that there exist graphs of low degree that are good approximations of complete graphs. We may obtain them from the recent theorem of Marcus, Spielman and Srivastava that there exist bipartite Ramanujan graphs of every degree and number of vertices [MSS15].

Lemma J.9. *For every integer n and even integer d , there is a weighted graph on n vertices of degree at most d that is a $4/\sqrt{d}$ approximation of K_n .*

Proof. The main theorem of [MSS15] tells us that there is a bipartite Ramanujan graph on $2n$ vertices of degree k for every $k \leq n$. By Propositions J.6 and J.7, a collapse of this graph is a weighted graph of degree at most $2k$ that is a $(4 \ln 2)/\sqrt{k}$ approximation of $K_{n,n}$. The result now follows by setting $d = 2k$. □

J.1 Sparsifying Complete Product Demand Graphs

In the rest of the section, we will adapt these expander constructions to weighted settings. We start with product demand graphs.

Lemma G.15. *There is a routine $\text{WEIGHTEDEXPANDER}(\mathbf{d}, \epsilon)$ such that for any $\epsilon > 0$, and a demand vector $\mathbf{d} \in (\mathbb{R}_{>0})^n$, $\text{WEIGHTEDEXPANDER}(\mathbf{d}, \epsilon)$ returns in $O(n\epsilon^{-4})$ work and $O(\log n)$ depth a graph H with $O(n\epsilon^{-4})$ edges such that*

$$\mathbf{L}_H \approx_\epsilon \mathbf{L}_{G(\mathbf{d})}.$$

Our algorithm for sparsifying complete product demand graphs begins by splitting the vertices of highest demands into many vertices. By *splitting* a vertex, we mean replacing it by many vertices whose demands sum to its original demand. In this way, we obtain a larger product demand graph. We observe that we can obtain a sparsifier of the original graph by sparsifying the larger graph, and then collapsing back together the vertices that were split.

Proposition J.10. *Let G be a product demand graph with vertex set $\{1, \dots, n\}$ and demands \mathbf{d} , and let $\widehat{G} = (\widehat{V}, \widehat{E})$ be a product demand graph with demands $\widehat{\mathbf{d}}$. If there is a partition of \widehat{V} into sets S_1, \dots, S_n so that for all $i \in V$, $\sum_{j \in S_i} \widehat{d}_j = d_i$, then \widehat{G} is a splitting of G and there is a matrix \mathbf{M} so that*

$$\mathbf{L}_G = \mathbf{M} \mathbf{L}_{\widehat{G}} \mathbf{M}^T.$$

Proof. The (i, j) entry of matrix \mathbf{M} is 1 if and only if $j \in S_i$. Otherwise, it is zero. \square

We now show that we can sparsify G by sparsifying \widehat{G} .

Proposition J.11. *Let \widehat{G}_1 and \widehat{G}_2 be graphs on the same vertex set \widehat{V} such that $\widehat{G}_1 \approx_\epsilon \widehat{G}_2$ for some ϵ . Let S_1, \dots, S_n be a partition of \widehat{V} , and let G_1 and G_2 be the graphs obtained by collapsing together all the vertices in each set S_i and eliminating any self loops that are created. Then*

$$G_1 \approx_\epsilon G_2.$$

Proof. Let \mathbf{M} be the matrix introduced in Proposition J.10. Then,

$$\mathbf{L}_{G_1} = \mathbf{M} \mathbf{L}_{\widehat{G}_1} \mathbf{M}^T \quad \text{and} \quad \mathbf{L}_{G_2} = \mathbf{M} \mathbf{L}_{\widehat{G}_2} \mathbf{M}^T.$$

The proof now follows from Fact 2.2. \square

For distinct vertices i and j , we let (i, j) denote the graph with an edge of weight 1 between vertex i and vertex j . If $i = j$, we let (i, j) be the empty graph. With this notation, we can express the product demand graph as

$$\sum_{i < j} d_i d_j (i, j) = \frac{1}{2} \sum_{i, j \in V} d_i d_j (i, j).$$

This notation also allows us to precisely express our algorithm for sparsifying product demand graphs.

$G' = \text{WEIGHTEDEXPANDER}(\mathbf{d}, \epsilon)$

1. Let \hat{n} be the least integer greater than $2n/\epsilon^2$ such that the algorithm described in Lemma J.8 produces an ϵ -approximation of $K_{\hat{n}}$.
2. Let $t = \frac{\sum_k d_k}{\hat{n}}$.
3. Create a new product demand graph \widehat{G} with demand vector $\widehat{\mathbf{d}}$ by splitting each vertex i into a set of $\lceil d_i/t \rceil$ vertices, S_i :
 - (a) $\lfloor d_i/t \rfloor$ vertices with demand t .
 - (b) one vertex with demand $d_i - t \lfloor d_i/t \rfloor$.
4. Let H be a set of \hat{n} vertices in \widehat{G} with demand t , and let L contain the other vertices. Set $k = |L|$.
5. Partition H arbitrarily into sets V_1, \dots, V_k , so that $|V_i| \geq \lfloor \hat{n}/k \rfloor$ for all $1 \leq i \leq k$.
6. Use the algorithm described in Lemma J.8 to produce \widetilde{K}_{HH} , an ϵ -approximation of the complete graph on H . Set

$$\widetilde{G} = t^2 \widetilde{K}_{HH} + \sum_{l \in L} \frac{|H|}{|V_l|} \sum_{h \in V_l} \widehat{d}_l \widehat{d}_h(l, h).$$

7. Let G' be the graph obtained by collapsing together all vertices in each set S_i .

This section and the next are devoted to the analysis of this algorithm. Given Proposition J.11, we just need to show that \widetilde{G} is a good approximation to \widehat{G} .

Proposition J.12. *The number of vertices in \widehat{G} is at most $n + \hat{n}$.*

Proof. The number of vertices in \widehat{G} is

$$\sum_{i \in V} \lceil d_i/t \rceil \leq n + \sum_{i \in V} d_i/t = n + \hat{n}.$$

□

So, $k \leq n$ and $\hat{n} \geq 2k/\epsilon^2$. That is, $|H| \geq 2|L|/\epsilon^2$. In the next section, we prove the lemmas that show that for these special product demand graphs \widehat{G} in which almost all weights are the maximum, our algorithm produces a graph \widetilde{G} that is a good approximation of \widehat{G} .

Proof. (of Lemma G.15) The number of vertices in the graph \widehat{G} will be between $n + 2n/\epsilon^2$ and $n + 16n/\epsilon^2$. So, the algorithm described in Lemma J.8 will take $O(\log n)$ depth and $O(n/\epsilon^4)$ work to produce an ϵ approximation of the complete graph on \hat{n} vertices. This dominates the computational cost of the algorithm.

Proposition J.11 tells us that G' approximates G at least as well as \widetilde{G} approximates \widehat{G} . To bound how well \widetilde{G} approximates \widehat{G} , we use two lemmas that are stated in the next section. Lemma J.14 shows that

$$\widehat{G}_{HH} + \widehat{G}_{LH} \approx_{O(\epsilon^2)} \widehat{G}.$$

Lemma J.16 shows that

$$\widehat{G}_{HH} + \widehat{G}_{LH} \approx_{4\epsilon} \widehat{G}_{HH} + \sum_{l \in L} \frac{|H|}{|V_l|} \sum_{h \in V_l} \widehat{d}_l \widehat{d}_h(l, h).$$

And, we already know that $t^2 \widetilde{K}$ is an ϵ -approximation of \widehat{G}_{HH} . Fact 2.3 says that we can combine these three approximations to conclude that \widetilde{G} is an $O(\epsilon)$ -approximation of \widehat{G} . \square

J.2 Product demand graphs with most weights maximal

In this section, we consider product demand graphs in which almost all weights are the maximum. For simplicity, we make a slight change of notation from the previous section. We drop the hats, we let n be the number of vertices in the product demand graph, and we order the demands so that

$$d_1 \leq d_2 \leq \dots \leq d_k \leq d_{k+1} = \dots = d_n = 1.$$

We let $L = \{1, \dots, k\}$ and $H = \{k+1, \dots, n\}$ be the set of low and high demand vertices, respectively. Let G be the product demand graph corresponding to \mathbf{d} , and let G_{LL} , G_{HH} and G_{LH} be the subgraphs containing the low-low, high-high and low-high edges respectively. We now show that little is lost by dropping the edges in G_{LL} when k is small.

Our analysis will make frequent use of the following Poincare inequality:

Lemma J.13. *Let $c(u, v)$ be an edge of weight c and let P be a path from u to v consisting of edges of weights c_1, c_2, \dots, c_k . Then*

$$c(u, v) \preceq c \left(\sum c_i^{-1} \right) P.$$

As the weights of the edges we consider in this section are determined by the demands of their vertices, we introduce the notation

$$(i, j)_2 = d_i d_j (i, j).$$

With this notation, we can express the product demand graph as

$$\sum_{i < j} (i, j)_2 = \frac{1}{2} \sum_{i, j \in V} (i, j)_2.$$

Lemma J.14. *If $|L| \leq |H|$, then*

$$G_{HH} + G_{LH} \approx_{3 \frac{|L|}{|H|}} G.$$

Proof. The lower bound $G_{HH} + G_{LH} \preceq G_{HH} + G_{LH} + G_{LL}$ follows from $G_{LL} \succeq 0$.

Using lemma J.13 and the assumptions $d_l \leq 1$ for $l \in L$ and $d_h = 1$ for $h \in H$, we derive for every $l_1, l_2 \in L$,

$$(l_1, l_2)_2 = \frac{1}{|H|^2} \sum_{h_1, h_2 \in H} (l_1, l_2)_2$$

(by Lemma J.13)

$$\preceq \frac{1}{|H|^2} \sum_{h_1, h_2 \in H} d_{l_1} d_{l_2} \left(\frac{1}{d_{l_1} d_{h_1}} + \frac{1}{d_{h_1} d_{h_2}} + \frac{1}{d_{h_2} d_{l_2}} \right) ((l_1, h_1)_2 + (h_1, h_2)_2 + (h_2, l_2)_2)$$

$$\begin{aligned}
&\preceq \frac{3}{|H|^2} \sum_{h_1, h_2 \in H} ((l_1, h_1)_2 + (h_1, h_2)_2 + (h_2, l_2)_2) \\
&= \frac{3}{|H|} \sum_{h \in H} ((l_1, h)_2 + (l_2, h)_2) + \frac{6}{|H|^2} G_{HH}.
\end{aligned}$$

So,

$$\begin{aligned}
G_{LL} &= \frac{1}{2} \sum_{l_1, l_2 \in L} (l_1, l_2)_2 \\
&\preceq \frac{1}{2} \sum_{l_1, l_2} \left(\frac{3}{|H|} \sum_{h \in H} ((l_1, h)_2 + (l_2, h)_2) + \frac{6}{|H|^2} G_{HH} \right) \\
&= \frac{3|L|}{|H|} G_{LH} + \frac{3|L|^2}{|H|^2} G_{HH}.
\end{aligned}$$

The assumption $|L| \leq |H|$ then allows us to conclude

$$G_{HH} + G_{LH} + G_{LL} \preceq \left(1 + 3\frac{|L|}{|H|}\right) (G_{HH} + G_{LH}).$$

□

Using a similar technique, we will show that the edges between L and H can be replaced by the union of a small number of stars. In particular, we will partition the vertices of H into k sets, and for each of these sets we will create one star connecting the vertices in that set to a corresponding vertex in L .

We employ the following consequence of the Poincare inequality in Lemma J.13.

Lemma J.15. *For any $\epsilon \leq 1$, $l \in L$ and $h_1, h_2 \in H$,*

$$\epsilon(h_1, l)_2 + (1/2)(h_1, h_2)_2 \approx_{4\sqrt{\epsilon}} \epsilon(h_2, l)_2 + (1/2)(h_1, h_2)_2.$$

Proof. By applying Lemma J.13 and recalling that $d_{h_1} = d_{h_2} = 1$ and $d_l \leq 1$, we compute

$$\begin{aligned}
(h_1, l)_2 &\preceq d_{h_1} d_l \left(\frac{\sqrt{\epsilon}}{d_{h_1} d_{h_2}} + \frac{1}{d_{h_2} d_l} \right) \left(\frac{1}{\sqrt{\epsilon}} (h_1, h_2)_2 + (h_2, l)_2 \right) \\
&\preceq \frac{1 + \sqrt{\epsilon}}{\sqrt{\epsilon}} (h_1, h_2)_2 + (1 + \sqrt{\epsilon})(h_2, l)_2 \\
&\preceq (1 + \sqrt{\epsilon})(h_2, l)_2 + \frac{2}{\sqrt{\epsilon}} (h_1, h_2)_2.
\end{aligned}$$

Multiplying both sides by ϵ and adding $(1/2)(h_1, h_2)_2$ then gives

$$\begin{aligned}
\epsilon(h_1, l)_2 + (1/2)(h_1, h_2)_2 &\preceq (1 + \sqrt{\epsilon})\epsilon(h_2, l)_2 + (2\sqrt{\epsilon} + 1/2)(h_1, h_2)_2 \\
&\preceq (1 + 4\sqrt{\epsilon}) (\epsilon(h_2, l)_2 + (1/2)(h_1, h_2)_2) \\
&\preceq e^{4\sqrt{\epsilon}} (\epsilon(h_2, l)_2 + (1/2)(h_1, h_2)_2).
\end{aligned}$$

By symmetry, we also have

$$\epsilon(h_2, l)_2 + (1/2)(h_1, h_2)_2 \preceq e^{4\sqrt{\epsilon}} (\epsilon(h_1, l)_2 + (1/2)(h_1, h_2)_2).$$

□

Lemma J.16. *Recall that $L = \{1, \dots, k\}$ and let V_1, \dots, V_k be a partition of $H = \{k+1, \dots, n\}$ so that $|V_l| \geq s$ for all l . Then,*

$$G_{HH} + G_{LH} \approx_{4/\sqrt{s}} G_{HH} + \sum_{l \in L} \frac{|H|}{|V_l|} \sum_{h \in V_l} (l, h)_2.$$

Proof. Observe that

$$G_{LH} = \sum_{l \in L} \sum_{h \in H} (l, h)_2.$$

For each $l \in L$, $h_1 \in H$ and $h_2 \in V_l$ we apply Lemma J.15 to show that

$$\frac{1}{|V_l|} (l, h_1)_2 + \frac{1}{2} (h_1, h_2)_2 \approx_{4/\sqrt{s}} \frac{1}{|V_l|} (l, h_2)_2 + \frac{1}{2} (h_1, h_2)_2.$$

Summing this approximation over all $h_2 \in V_l$ gives

$$(l, h_1)_2 + \sum_{h_2 \in V_l} \frac{1}{2} (h_1, h_2)_2 \approx_{4/\sqrt{s}} \sum_{h_2 \in V_l} \left(\frac{1}{|V_l|} (l, h_2)_2 + \frac{1}{2} (h_1, h_2)_2 \right).$$

Summing the left-hand side of this approximation over all $l \in L$ and $h_1 \in H$ gives

$$\sum_{l \in L, h_1 \in H} (l, h_1)_2 + \sum_{h_2 \in V_l} \frac{1}{2} (h_1, h_2)_2 = \sum_{l \in L, h_1 \in H} (l, h_1)_2 + \frac{1}{2} \sum_{h_1 \in H, l \in L} \sum_{h_2 \in V_l} (h_1, h_2)_2 = G_{LH} + G_{HH}.$$

On the other hand, the sum of the right-hand terms gives

$$G_{HH} + \sum_{l \in L, h_1 \in H} \sum_{h_2 \in V_l} \frac{1}{|V_l|} (l, h_2)_2 = G_{HH} + \sum_{l \in L} \sum_{h_2 \in V_l} \frac{|H|}{|V_l|} (l, h_2)_2.$$

□

J.3 Weighted Bipartite Expanders

This construction extends analogously to bipartite product graphs. The bipartite product demand graph of vectors $(\mathbf{d}^A, \mathbf{d}^B)$ is a complete bipartite graph whose weight between vertices $i \in A$ and $j \in B$ is given by $w_{ij} = d_i^A d_j^B$. The main result that we will prove is:

Lemma G.16. *There is a routine `WEIGHTEDBIPARTITEEXPANDER` $(\mathbf{d}^A, \mathbf{d}^B, \epsilon)$ such that for any demand vectors \mathbf{d}^A and \mathbf{d}^B of total length n , and a parameter ϵ , it returns in $O(n\epsilon^{-4})$ work and $O(\log n)$ depth a bipartite graph H between A and B with $O(n\epsilon^{-4})$ edges such that*

$$\mathbf{L}_H \approx_\epsilon \mathbf{L}_{G(\mathbf{d}^A, \mathbf{d}^B)}.$$

Without loss of generality, we will assume $d_1^A \geq d_2^A \geq \dots \geq d_{n^A}^A$ and $d_1^B \geq d_2^B \geq \dots \geq d_{n^B}^B$. As the weights of the edges we consider in this section are determined by the demands of their vertices, we introduce the notation

$$(i, j)_2 = d_i^A d_j^B (i, j).$$

Our construction is based on a similar observation that if most vertices on A side have d_i^A equaling to d_1^A and most vertices on B side have d_i^B equaling to d_1^B , then the uniform demand graph on these vertices dominates the graph.

$G' = \text{WEIGHTEDBIPARTITEEXPANDER}(\mathbf{d}^A, \mathbf{d}^B, \epsilon)$

1. Let $n' = \max(n^A, n^B)$ and \hat{n} be the least integer greater than $2n'/\epsilon^2$ such that the algorithm described in Lemma J.8 produces an ϵ -approximation of $K_{\hat{n}, \hat{n}}$.
2. Let $t^A = \frac{\sum_k d_k^A}{\hat{n}}$ and $t^B = \frac{\sum_k d_k^B}{\hat{n}}$.
3. Create a new bipartite demand graph \hat{G} with demands $\hat{\mathbf{d}}^A$ and $\hat{\mathbf{d}}^B$ follows:
 - (a) On the side A of the graph, for each vertex i , create a subset S_i consisting of $\lceil d_i^A/t^A \rceil$ vertices:
 - i. $\lfloor d_i^A/t^A \rfloor$ with demand t^A .
 - ii. one vertex with demand $d_i^A - t^A \lfloor d_i^A/t^A \rfloor$.
 - (b) Let H^A contain \hat{n} vertices of A of with demand t^A , and let L^A contain the rest. Set $k^A = |L^A|$.
 - (c) Create the side B of the graph with partition H^B, L^B and demand vector $\hat{\mathbf{d}}^B$ similarly.
4. Partition H^A into sets of size $|V_i^A| \geq \lceil \hat{n}/k^A \rceil$, one corresponding to each vertex $l \in L^A$. Partition V_B similarly.
5. Let $\tilde{K}_{H^A H^B}$ be a bipartite expander produced by Lemma J.8 that ϵ -approximates $K_{\hat{n}, \hat{n}}$, identified with the vertices H^A and H^B .

Set

$$\tilde{G} = t^A t^B \tilde{K} + \sum_{l \in L^A} \frac{|H^B|}{|V_l^B|} \sum_{h \in V_l^B} \hat{d}_l^A \hat{d}_h^B(l, h) + \sum_{l \in L^B} \frac{|H^A|}{|V_l^A|} \sum_{h \in V_l^A} \hat{d}_l^B \hat{d}_h^A(l, h).$$

6. Let G' be the graph obtained by collapsing together all vertices in each set S_i^A and S_i^B .

Similarly to the non-bipartite case, the Poincare inequality show that the edges between low demand vertices can be completely omitted if there are many high demand vertices which allows the demand routes through high demand vertices.

Lemma J.17. *Let G be the bipartite product demand graph of the demand $(\mathbf{d}_i^A, \mathbf{d}_j^B)$. Let H^A a subset of vertices on A side with demand higher than the set of remaining vertices L^A on A side. Define H^B, L^B similarly. Assume that $|L^A| \leq |H^A|$ and $|L^B| \leq |H^B|$, then*

$$G_{H^A H^B} + G_{H^A L^B} + G_{L^A H^B} \approx_{3 \max\left(\frac{|L^A|}{|H^A|}, \frac{|L^B|}{|H^B|}\right)} G.$$

Proof. The proof is analogous to Lemma J.14, but with the upper bound modified for bipartite graphs.

For every edge l_A, l_B , we embed it evenly into paths of the form l_A, h_B, h_A, l_B over all choices of h_A and h_B . The support of this embedding can be calculated using Lemma J.13, and the overall accounting follows in the same manner as Lemma J.14. □

It remains to show that the edges between low demand and high demand vertices can be compressed into a few edges. The proof here is also analogous to Lemma J.15: we use the Poincare inequality to show that all demands can routes through high demand vertices. The structure of

the bipartite graph makes it helpful to further abstract these inequalities via the following Lemma for four edges.

Lemma J.18. *Let G be the bipartite product demand graph of the demand $(\mathbf{d}_i^A, \mathbf{d}_j^B)$. Given $h_A, l_A \in A$ and $h_{B,1}, h_{B,2} \in B$. Assume that $d_{h_A}^A = d_{h_{B,1}}^B = d_{h_{B,2}}^B \geq d_{l_A}^A$. For any $\epsilon < 1$, we have*

$$\epsilon(l_A, h_{B,1})_2 + (h_A, h_{B,2})_2 + (h_A, h_{B,1})_2 \approx_{3\sqrt{\epsilon}} \epsilon(l_A, h_{B,2})_2 + (h_A, h_{B,2})_2 + (h_A, h_{B,1})_2.$$

Proof. Using Lemma J.13 and $d_{h_A}^A = d_{h_{B,1}}^B = d_{h_{B,2}}^B \geq d_{l_A}^A$, we have

$$\begin{aligned} & (l_A, h_{B,1})_2 \\ & \preceq d_{l_A}^A d_{h_{B,1}}^B \left(\frac{1}{d_{l_A}^A d_{h_{B,2}}^B} + \frac{\sqrt{\epsilon}}{d_{h_A}^A d_{h_{B,2}}^B} + \frac{\sqrt{\epsilon}}{d_{h_A}^A d_{h_{B,1}}^B} \right) \left((l_A, h_{B,2})_2 + \frac{1}{\sqrt{\epsilon}}(h_A, h_{B,2})_2 + \frac{1}{\sqrt{\epsilon}}(h_A, h_{B,1})_2 \right) \\ & \preceq (1 + 2\sqrt{\epsilon})(l_A, h_{B,2})_2 + \frac{1 + 2\sqrt{\epsilon}}{\sqrt{\epsilon}}(h_A, h_{B,2})_2 + \frac{1 + 2\sqrt{\epsilon}}{\sqrt{\epsilon}}(h_A, h_{B,1})_2. \end{aligned}$$

Therefore,

$$\epsilon(l_A, h_{B,1})_2 + (h_A, h_{B,2})_2 + (h_A, h_{B,1})_2 \preceq (1 + 3\sqrt{\epsilon}) (\epsilon(l_A, h_{B,2})_2 + (h_A, h_{B,2})_2 + (h_A, h_{B,1})_2).$$

The other side is similar due to the symmetry. \square

Proof. (of Lemma G.16) The proof is analogous to Lemma G.15. After the splitting, the demands in H^A are higher than the demands in L^A and so is H^B to L^B . Therefore, Lemma J.17 shows that that

$$\widehat{G}_{H^A H^B} + \widehat{G}_{H^A L^B} + \widehat{G}_{L^A H^B} \approx_{3\epsilon^{2/2}} \widehat{G}.$$

By a proof analogous to Lemma J.16, one can use Lemma J.18 to show that

$$\widehat{G}_{H^A H^B} + \widehat{G}_{H^A L^B} + \widehat{G}_{L^A H^B} \approx_{O(\epsilon)} \widehat{G}_{H^A H^B} + \frac{|H^B|}{|V_l^B|} \sum_{h \in V_l^B} \hat{d}_l^A \hat{d}_h^B(l, h) + \sum_{l \in L^B} \frac{|H^A|}{|V_l^A|} \sum_{h \in V_l^A} \hat{d}_l^B \hat{d}_h^A(l, h).$$

And, we already know that $t^A t^B \tilde{K}$ is an ϵ -approximation of $\widehat{G}_{H^A H^B}$. Fact 2.3 says that we can combine these three approximations to conclude that \tilde{G} is an $O(\epsilon)$ -approximation of \widehat{G} . \square