Numerical Methods for Gremban's Expansion of Signed Graphs

Alyson L. Fox *

January 23, 2016

Abstract

Signed graphs contain both positive and negative relationships between vertices. Data scientists are interested in solving linear systems associated with signed graphs. Fairly robust solvers for unsigned, undirected graph Laplacians have been developed but these solvers are not directly applicable to general signed graphs. Gremban's expansion [2] is used to transform the signed, undirected graph Laplacian into an unsigned, undirected graph Laplacian. The solution to the linear system of the expanded matrix yields the solution of the original linear system. Thus, using Gremban's expansion we can extend the current Laplacian solvers' robustness to signed graph Laplacians. This paper delves into the numerical stability and applicability of Gremban's expansion and proves that the error of the solution of the original linear system can be tightly bounded by the error of the expanded system. Gremban's expansion was originally only for symmetric matrices. However, this paper demonstrates that the expansion is applicable to the nonsymmetric matrices associated with signed, directed graphs. Both manufactured and real-world signed, undirected graph Laplacians are tested with various solvers to show that the expansion is numerically stable.

1 Introduction

For graph-ranking applications, data scientists are interested in solving $L\mathbf{x} = \mathbf{b}$, where L is a graph-associated matrix (e.g. a graph Laplacian). Four types of graphs that are of interest are: unsigned-undirected (UU), signed-undirected (SU), unsigned-directed (UD), and signed-directed (SD). Every edge in an undirected graph is bidirectional, where as some edges in a directed graph are directed from a source vertex to a target vertex. The graph Laplacian of an undirected graph is symmetric due to the symmetry of the edges, which is not the case for directed graphs. A signed graph, however, has both positive edges and negative edges. A positive edge of a signed graph can be seen as a similarity or closeness, where a negative edge can show dissimilarity or distance. A social network, for example, may be represented as a signed graph where positive edges represent "friends" and negative edges represent "foes." When a graph is used to cluster or group a social network, a negative edge would represent two people whose grouping together should be avoided as it may cause a disruption. Data scientists employ various models that may involve any of these four graph types and they require efficient solvers for the associated linear systems.

There exists a multitude of fast solvers for UU graphs, one such solver is Lean Algebraic Multigrid (LAMG) [5]. UU graph systems are known to be the simplest of the four problems since they result in a symmetric, diagonally dominant Laplacian with a known null-space. Directed graph Laplacians are nonsymmetric and there is no efficient, general-purpose solver for nonsymmetric M-matrices. Therefore, there is a lack of robust UD graph solvers. This paper explores Gremban's expansion for signed graphs, which decomposes any symmetric diagonally dominant matrix into a diagonally dominant Z-matrix and a nonpositive matrix, $S = M + P$. The user can solve $S\mathbf{x} = \mathbf{b}$, by using an expanded UU graph Laplacian system,

$$G\mathbf{w} = \begin{bmatrix} M & -P \\ -P & M \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ -\mathbf{x} \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ -\mathbf{b} \end{bmatrix} = \mathbf{z},$$

which is amenable to fast solvers like LAMG. Under suitable hypotheses on S, the system G is either nonsingular or has a trivial null space. In either case, the solution to the original system can be found by solving the expanded system and then extracting $\mathbf{x}$. (See Section 3 for details). This paper shows that, for SU graphs, Gremban's expansion is stable such that the error of the original linear system is less than the error of the linear system for the expansion. The expanded system and a robust UU solver (e.g. LAMG) yield a robust SU solver, as demonstrated in the results section. This paper shows how Gremban's expansion can be generalized to any (possibly nonsymmetric) diagonally dominant matrix. Gremban's expansion could then expand SD graphs to UD graphs. However, since robust, optimal solvers for directed graphs are not readily available, the main focus of the paper will be SU graphs.

*Department of Applied Mathematics, University of Colorado at Boulder, Boulder, Colorado

The paper is organized as follows. Section 2 describes the notation, important definitions, and extends useful concepts of unsigned graphs to signed graphs. Section 3 introduces Gremban's expansion and demonstrates its numerical stability. It is also shown that the expansion can be also used for linear systems involving SD graphs. Section 4 gives a brief description of LAMG. Section 5 presents numerical results demonstrating the use of the expansion with LAMG.

2 Preliminaries

A graph, $\mathcal{G}$, may belong to many different classes. This work extends the linear solver technology that works well on UU graphs to SU graphs. A UU graph, $\mathcal{G}(\mathcal{V}, \mathcal{E}, w)$, relates a set of n vertices, $\mathcal{V}$, by m connections or edges in the set $\mathcal{E}$ with weights, w. An edge $(i,j) \in \mathcal{E}$ between two vertices i and j is an undirected or symmetric relationship, meaning (j,i) is also in $\mathcal{E}$ and $w_{ij} = w_{ji}$. Edge (i,j) is also a positive relationship, $w_{ij} > 0$, as its size is proportional to the strength of affinity between i and j. Let us assume $\mathcal{G}$ is connected with no self loops, that is $(i,i) \notin \mathcal{E}$. Typically, a UU $\mathcal{G}$ is represented by its symmetric adjacency matrix,

$$A_{ij} = \begin{cases} w_{ij} & (i,j) \in \mathcal{E} \\ 0 & otherwise \end{cases}. \tag{1}$$

This and other matrix representations are commonly analyzed. Define the diagonal degree matrix,

$$D_{ij} = \begin{cases} d_i & i = j \\ 0 & otherwise \end{cases}, \tag{2}$$

where $d_i = \sum_i w_{ij}$ is the total weight of edges incident to vertex i. The combinatorial graph Laplacian of $\mathcal{G}$ is $L = D - A$.

Data analysts often desire the solution to $L\mathbf{x} = \mathbf{b}$. Matrix L is singular with known kernel, $L\mathbf{1} = \mathbf{0}$. If L is symmetric, $\mathbf{b}$ is in the range of L only if $\mathbf{1}^t\mathbf{b} = 0$. If $\mathbf{b}$ is not in the range, we are interested in $L^\dagger\mathbf{b}$ which can be calculated by using projections and a SPD solve,

$$\mathbf{x} = \left(L + \frac{\alpha}{n}\mathbf{1}\mathbf{1}^t\right)^{-1}\left(I - \frac{1}{n}\mathbf{1}\mathbf{1}^t\right)\mathbf{b}, \tag{3}$$

for any $\alpha > 0$. This is an important discussion for the signed case, as the matrices of interest are not always singular.

For UU graphs, matrix L can be symmetrically factored using its incidence matrix, $E \in \mathbb{R}^{n\times m}$, which maps the edges to the vertices. Let the e-th edge be $(i,j) \in \mathcal{E}$ and *orient* e as follows. If $i < j$, then let $E_{i,e} = +1$, $E_{j,e} = -1$, $E_{k,e} = 0$ for $k \neq i, j$. Let W be a diagonal matrix such that $W_{e,e} = w_{ij}$. This yields the factorization, $L = EWE^t$. This factorization and the quadratic form, $\mathbf{x}^t L\mathbf{x} = \sum_{(i,j)\in\mathcal{E}} w_{ij}(x_i - x_j)^2$, are powerful tools used for deriving properties of solutions to $L\mathbf{x} = \mathbf{b}$, designing numerical solvers, and performing numerical analysis of such solvers. This can be generalized to the signed case.

For SU graphs, which we denote by $\mathcal{G}^{\pm}(\mathcal{V}, \mathcal{E}^+ \cup \mathcal{E}^-, w)$, there are some positively weighted edges, $w_{ij} > 0$ for $(i,j) \in \mathcal{E}^+$, and some negatively weighted edges, $w_{ij} < 0$ for $(i,j) \in \mathcal{E}^-$. For simplicity, we assume there are no contradictory edges, that is, $(i,j) \in \mathcal{E}^+ \Rightarrow (i,j) \notin \mathcal{E}^-$, and, conversely, $(p,q) \in \mathcal{E}^- \Rightarrow (p,q) \notin \mathcal{E}^+$. Thus, $\mathcal{E}^+ \cap \mathcal{E}^- = \emptyset$ and $\mathcal{E}^+ \cup \mathcal{E}^- = \mathcal{E}$. The matrix forms associated with $\mathcal{G}^{\pm}$ are easily generalized with the same formulae we provide above. The signed adjacency matrix, A, will have some negative entries, D is the diagonal matrix of the *total degree* ($d_i = \sum_j |a_{i,j}|$), and the signed combinatorial Laplacian L has some positive off-diagonal entries. We extend the incident factorization to signed graphs.The incidence matrix, E, for a signed graph must have the product of the entries in each column be equal to the negative of the sign of each edge weight. Let the e-th edge be $(i,j) \in \mathcal{E}$ and *orient* e as follows. If $i < j$, then let $E_{i,e} = +1$, $E_{j,e} = \pm 1$ such that $E_{ie}E_{je} = -sign(w_{ij})$, $E_{k,e} = 0$ for $k \neq i, j$. Let $W_{e,e} = |w_{ij}|$, then the Laplacian is factor as before, $L = EWE^t$. Note that that L of a SU graph is a symmetric diagonally dominant matrix.

The goal is to apply tools that work well for combinatorial Laplacians of UU graphs to solve $L\mathbf{x} = \mathbf{b}$ for SU graphs. For $\mathcal{G}^{\pm}$, we will show that, in special situations, L is singular. In real-world datasets, L is usually nonsingular. We address the singular case first. When L is singular, the eigendecomposition is

directly related to that of the combinatorial Laplacian generated by reversing the sign of all the negative edge weights, $\tilde{L}$. To formalize this we require the concept of a *balanced* signed graph. A bipartition $\mathbf{y}$ of the set of vertices $\mathcal{V}$, is defined as, $y_i = 1$, $i \in \mathcal{U}$, and $y_j = -1$, $j \in \mathcal{W}$, where $\mathcal{U} \cup \mathcal{W} = \mathcal{V}$ and $\mathcal{U} \cap \mathcal{W} = \emptyset$.

Definition 2.1. *A connected signed graph, $\mathcal{G}^{\pm}(\mathcal{V}, \mathcal{E}^+ \cup \mathcal{E}^-, w)$, with nonzero edge weights is balanced if $\mathcal{V}$ can be partitioned into two groups $\mathcal{U}$ and $\mathcal{W}$, such that $(i,j) \in \mathcal{E}^+$ implies either both vertices i, j are in $\mathcal{U}$ or both are in $\mathcal{W}$ and $(i,j) \in \mathcal{E}^-$ implies one of the two vertices is in $\mathcal{U}$ and the other is in $\mathcal{W}$.*

Equivalently, unbalanced graphs can be defined as graphs containing a cycle with an odd number of negative edges. If a graph is balanced then the spectrum is easily relatable to the corresponding unsigned Laplacian.

Theorem 2.1. *Let L be a signed Laplacian matrix of the balanced connected graph, $\mathcal{G}$, with bipartition $\mathbf{y}$ and eigenvalue decomposition $L = U\Lambda U^t$. Let $\tilde{L} = diag(\mathbf{y})Ldiag(\mathbf{y})$, then $\tilde{L}$ is the corresponding Laplacian matrix of the unsigned graph Laplacian of $\mathcal{G}$. The eigenvalue decomposition of the Laplacian matrix, $\tilde{L}$, is similar to L, by $\tilde{L} = \tilde{U}\Lambda\tilde{U}^t$ where $\tilde{U} = diag(\mathbf{y})U$. Moreover, $\mathbf{y}$ is the kernel of L.*

Proof. Order the vertices so that $\mathcal{U}$ is first and $\mathcal{W}$ is second and let block L_{12} represent the edges between $i \in \mathcal{U}$ and $j \in \mathcal{W}$. A balanced graph implies that all positive off-diagonal entries in L occur in L_{12} and $L_{12} \geq O$. The similarity transform $\tilde{L} = \mathrm{d}iag(\mathbf{y})Ldiag(\mathbf{y})$ makes these entries negative (with their symmetric counterparts) and leaves all entries in the other blocks of L alone. The rest of the results are derived easily via the similarity transform and the fact that $\mathbf{1}$ is the kernel of $\tilde{L}$. □

The partition of a balanced graph can be found by a depth-first traversal, assigning each vertex to a partition such that the balance property is fulfilled [3]. Additionally, the same process can be used to determine if the graph is unbalanced. If $\mathcal{G}$ is balanced, then one would solve the spectrally equivalent Laplacian system $\tilde{L}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}$, where $\tilde{\mathbf{b}} = \mathrm{diag}(\mathbf{y})\mathbf{b}$ and the solution to the original linear system is then found by $\mathbf{x} = \mathrm{diag}(\mathbf{y})^{-1}\tilde{\mathbf{x}}$. Therefore, in what follows, we assume $\mathcal{G}$ is unbalanced. It is unlikely that a real-world graph is balanced unless it is a byproduct of the underlying dataset.

The incidence factorization for $\mathcal{G}^{\pm}$, shown above, is used to show that L is always at least positive semi-definite. From the factorization, we can derived the quadratic form,

$$\mathbf{x}^t L\mathbf{x} = \sum_{(i,j)\in\mathcal{E}^+} |w_{ij}|(x_i - x_j)^2 + \sum_{(p,q)\in\mathcal{E}^-} |w_{pq}|(x_p + x_q)^2, \tag{4}$$

which is greater or equal to 0 for any $\mathbf{x}$. For an unbalanced $\mathcal{G}^{\pm}$, we show this quadratic form is strictly positive, implying matrix L is strictly positive definite. For the first term to be 0, we require $\mathbf{x}$ is constant on each connected component of $\mathcal{G}^+(\mathcal{V}, \mathcal{E}^+, w)$. For the second term to be 0, any cycle in $\mathcal{G}^{\pm}$ must have an even number of edges in $\mathcal{E}^-$. Thus, the terms $(x_p + x_q)^2$ of the connecting components cancel out. This contradicts the definition of an unbalanced graph [3]. Thus, the signed Laplacian of an unbalanced $\mathcal{G}^{\pm}$ is nonsingular and $L\mathbf{x} = \mathbf{b}$ can be solved without projection for any $\mathbf{b}$.

Remark 2.1. (SD Graphs) *As mentioned in the introduction, data scientists are also interested in SD graphs. In a SD graph, each vertex of the edge is given a direction and a sign. The following section will show that Gremban's expansion can be generalized to nonsymmetric diagonally dominant matrices. This would allow any SD graph Laplacian to be expanded into an UD graph Laplacian, and the solution of the linear system of the expanded matrix will return a solution to the original system.*

The following section contains a detailed description of Gremban's expansion and how it can be used with a robust solver for UU graph Laplacian.

3 Numerical Stability of Gremban's Expansion

In [2], the author shows that any symmetric diagonally dominant (SDD) matrix can be expanded into a undirected graph Laplacian, and solving a linear system involving the expanded matrix yields the solution to a linear system involving the original matrix. An important distinction needs to be made between

signed and unsigned graph Laplacians. Signed Laplacians have some positive off-diagonal elements while a unsigned Laplacian has strictly nonpositive off-diagonal elements. The following defines a *diagonally dominant Z-matrix*, which is a generalization of an unsigned Laplacian:

Definition 3.1. *A matrix M as a* Diagonally Dominant Z-matrix *if it is diagonally dominant, $m_{ii} \geq \sum_{j=1}^{n} |m_{ij}|$, with positive diagonal, $m_{ii} > 0$ for every i, and has nonpositive off-diagonal elements, $m_{ij} \leq 0$ for every $i \neq j$.*

With this definition, M is possibly strongly diagonally dominant, $m_{ii} > \sum_{j=1}^{n} |m_{ij}|$, for some i. Also, M is not required to be symmetric. Any diagonally dominant matrix can be decomposed into a diagonally dominant Z-matrix and a nonpositive matrix. The *Gremban expansion* is then defined as:

Definition 3.2. *Let S be diagonally dominant with positive diagonal. We decompose, $S = M + P$, such that M is a diagonally dominant Z-matrix, and P has purely nonnegative entries and zero diagonal. Define the* Gremban Expansion *matrix, G, as*

$$G = \begin{bmatrix} M & -P \\ -P & M \end{bmatrix}.$$

Note that G is a diagonally dominant Z-matrix. If S is a signed Laplacian, then G would be an unsigned Laplacian. All unsigned Laplacians are diagonally dominant Z-matrics by definition. The following demonstrates the relationships between the spectra of the expansion and the original matrix. We will also show that the expansion is numerically stable for SU graphs, meaning that a small residual for the expanded system implies a small residual for the original system. Although not the primary focus of this paper, we show that the expansion and some of the theory is relevant for the nonsymmetric matrices associated with SD graphs. The following theorem relates the spectrum of the expansion, G, with the matrices that define the diagonally dominant matrix, S:

Theorem 3.1. *Let S be diagonally dominant with positive diagonal, and let G be a Gremban expansion of S. Employ the decomposition, $S = M + P$, with M a diagonally dominant Z-matrix, and $P \geq O$ and zero diagonal. Then,*

$$\sigma(G) \;=\; \sigma(S) \cup \sigma(M - P).$$

Proof. Let $(\lambda, \mathbf{u})$ be any eigenpair of S. Then, $(\lambda, [\mathbf{u}^t, -\mathbf{u}^t]^t)$ is an eigenpair for G,

$$\begin{bmatrix} M & -P \\ -P & M \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ -\mathbf{u} \end{bmatrix} = \begin{bmatrix} (M+P)\mathbf{u} \\ -(M+P)\mathbf{u} \end{bmatrix} = \lambda \begin{bmatrix} \mathbf{u} \\ -\mathbf{u} \end{bmatrix}.$$

Similarly, let $(\theta, \mathbf{v})$ be any eigenpair of $(M - P)$. Then, $(\theta, [\mathbf{v}^t, \mathbf{v}^t]^t)$ is an eigenpair for G,

$$\begin{bmatrix} M & -P \\ -P & M \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{v} \end{bmatrix} = \begin{bmatrix} (M-P)\mathbf{v} \\ (M-P)\mathbf{v} \end{bmatrix} = \theta \begin{bmatrix} \mathbf{v} \\ \mathbf{v} \end{bmatrix}.$$

It has been shown that $\sigma(G) \supset \sigma(S) \cup \sigma(M - P)$. To show, $\sigma(G) \subset \sigma(S) \cup \sigma(M - P)$, let $\lambda \in \sigma(G)$ with the corresponding eigenvector $\mathbf{v} = [\mathbf{v}_1^t, \mathbf{v}_2^t]^t$. Then

$$\begin{bmatrix} M & -P \\ -P & M \end{bmatrix} \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix} = \begin{bmatrix} M\mathbf{v}_1 - P\mathbf{v}_2 \\ -P\mathbf{v}_1 + M\mathbf{v}_2 \end{bmatrix} = \lambda \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix} = \begin{bmatrix} \lambda\mathbf{v}_1 \\ \lambda\mathbf{v}_2 \end{bmatrix}.$$

This implies, $M\mathbf{v}_1 - P\mathbf{v}_2 = \lambda\mathbf{v}_1$ and $-P\mathbf{v}_1 + M\mathbf{v}_2 = \lambda\mathbf{v}_2$. Then, $(\lambda, \mathbf{v}_1 + \mathbf{v}_2)$ is an eigenpair of $(M - P)$,

$$(M-P)(\mathbf{v}_1 + \mathbf{v}_2) = (M\mathbf{v}_1 - P\mathbf{v}_2) + (M\mathbf{v}_2 - P\mathbf{v}_1) = \lambda\mathbf{v}_1 + \lambda\mathbf{v}_2 = \lambda(\mathbf{v}_1 + \mathbf{v}_2)$$

and $(\lambda, \mathbf{v}_1 - \mathbf{v}_2)$ as an eigenpair of S,

$$S(\mathbf{v}_1 - \mathbf{v}_2) = (M+P)(\mathbf{v}_1 - \mathbf{v}_2) = (M\mathbf{v}_1 - P\mathbf{v}_2) - (M\mathbf{v}_2 - P\mathbf{v}_1) = \lambda\mathbf{v}_1 - \lambda\mathbf{v}_2 = \lambda(\mathbf{v}_1 - \mathbf{v}_2).$$

Therefore, $\sigma(G) \subset \sigma(S) \cup \sigma(M-P)$ and thus, $\sigma(G) = \sigma(S) \cup \sigma(M-P)$. For S symmetric, let $S = U\Lambda U^t$ and $(M-P) = V\Theta V^t$ be the eigenvalue decompositions of the smaller matrices, then the eigendecomposition of G is,

$$G = \frac{1}{\sqrt{2}} \begin{bmatrix} U & V \\ -U & V \end{bmatrix} \begin{bmatrix} \Lambda & \\ & \Theta \end{bmatrix} \left(\frac{1}{\sqrt{2}} \begin{bmatrix} U & V \\ -U & V \end{bmatrix} \right)^t.$$

□

It is often the case that G is singular even though S is nonsingular. If S is a signed graph Laplacian associated with an unbalanced $\mathcal{G}^{\pm}$, then it is nonsingular, as shown in the previous section. However, G is singular with constant kernel. Note that nonsingular S does not imply an unbalanced sign structure of the off-diagonal elements. For example, consider a signed Laplacian of a balanced $\mathcal{G}^{\pm}$ shifted by a non-negative diagonal matrix. In such cases, the graph of G is not connected, but each of the two components are strictly diagonally dominant and G is nonsingular.

Lemma 3.1. *Assume S is nonsingular, strongly connected, and diagonally dominant. Let G be the Gremban expansion of S. Consider the systems $S\mathbf{x} = \mathbf{b}$ and $G\mathbf{w} = \mathbf{z}$ where $\mathbf{z} = [\mathbf{b}^t, -\mathbf{b}^t]^t$. If G is nonsingular then $\mathbf{w} = [\mathbf{x}^t, -\mathbf{x}^t]^t$. If G is singular then the null-space of G is the constant vector, $\mathbf{1}$, and the solution computed by the process described in (3) applied to $G\mathbf{w} = \mathbf{z}$ yields $\mathbf{w} = [\mathbf{x}^t, -\mathbf{x}^t]^t$.*

Proof. If S is nonsingular and strongly connected, then by Theorem 3.1, G is only singular if $(M-P)$ is singular. By the Peron-Frobenius theorem, for $(M-P)$ to be singular, S needs to be only weakly diagonally dominant, and the constant vector is the unique kernel component, $(M-P)\mathbf{1} = \mathbf{0}$. In this case, $[\mathbf{1}^t, \mathbf{1}^t]^t$ is the unique kernel component of G, and the solution can be found as described in equation (3) where $\mathbf{z} = [\mathbf{b}^t, -\mathbf{b}^t]^t$ is orthogonal to the kernel. We see that $G\mathbf{w} = \mathbf{z}$,

$$\begin{bmatrix} M & -P \\ -P & M \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ -\mathbf{x} \end{bmatrix} = \begin{bmatrix} (M+P)\mathbf{x} \\ -(M+P)\mathbf{x} \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ -\mathbf{b} \end{bmatrix}. \tag{5}$$

In the case of nonsingular G, $\mathbf{w} = [\mathbf{x}^t, -\mathbf{x}^t]^t$ is the unique solution. □

The above proofs are extremely powerful. Given a good linear solver for Laplacians of UU graphs, we can solve any symmetric diagonally dominant system, $S\mathbf{x} = \mathbf{b}$, by transforming it into the associated diagonally dominant Z-matrix system, $G\mathbf{w} = \mathbf{z}$, solving for $\mathbf{w}$, and extracting $\mathbf{x}$.

Remark 3.1. (Gremban's Expansion for SD graphs) *We note that several of the results in this section provide theoretical underpinnings for the analysis of applying Gremban's expansion to systems involving SD graphs, where S is nonsymmetric. First, there is no assumption of symmetry in Lemma 3.1, and the expansion applies to linear systems involving nonsymmetric S. Moreover, the eigenvalue containment result in Theorem 3.1 holds for the nonsymmetric G associated with nonsymmetric S. It is also possible to show that the singular values and vectors of nonsymmetric G are directly related to those of S and $(M-P)$. Both the eigenpairs and singular values of G are important for analyzing the numerical stability of the expansion and the design of LAMG-like solvers for nonsymmetric diagonally dominant Z-matrix systems.*

For the rest of this paper we will assume that S is symmetric. Note that the solvers will be using the expanded matrix G, however, the actual interest is for the solution to the smaller problem S. The following theorem shows that if given an approximate solution to the Gremban expansion system, we can tightly bound the norm of the residual of the original system by the norm of the residual of the expansion.

Theorem 3.2. *Let S be symmetric and diagonally dominant and G be the Gremban expansion of S. Let $\mathbf{y}$ be an approximate solution to $G\mathbf{w} = \mathbf{z}$. Define $P_1 = [I \quad O]$ and $P_2 = [O \quad I]$. Let $\mathbf{v} = \alpha P_1\mathbf{y} + \beta P_2\mathbf{y}$, where α, β are chosen so that $\|S\mathbf{v} - \mathbf{b}\|$ is minimized. Then,*

$$\|S\mathbf{v} - \mathbf{b}\| \leq C_S \|G\mathbf{y} - \mathbf{z}\|$$

where $C_s = \frac{1}{\sqrt{2}}$. This bound is tight.

Proof. Let $R = \frac{1}{2}[I, -I]$ be a restriction operator, thus $R\mathbf{z} = \mathbf{b}$ and $R\mathbf{w} = \mathbf{x}$. Also $R\mathbf{y} = \frac{1}{2}P_1\mathbf{y} - \frac{1}{2}P_2\mathbf{y}$.

$$\|S\mathbf{z} - \mathbf{b}\| \quad \leq \quad \|SR\mathbf{y} - \mathbf{b}\| = \frac{\|SR(\mathbf{y} - \mathbf{w})\|}{\|G(\mathbf{y} - \mathbf{w})\|}\|G\mathbf{y} - \mathbf{z}\| = \frac{\|SR\mathbf{e}\|}{\|G\mathbf{e}\|}\|G\mathbf{y} - \mathbf{z}\| = C_s\|G\mathbf{y} - \mathbf{z}\|$$

where $C_s = \max_{\mathbf{e}\neq 0, \mathbf{e}\perp 1}\sqrt{\frac{\langle R^tS^2R\mathbf{e},\mathbf{e}\rangle}{\langle G^2\mathbf{e},\mathbf{e}\rangle}}$ which is equivalent to the maximum eigenvalue of the generalized eigenvalue problem $R^tS^2R\mathbf{x} = \lambda G^2\mathbf{x}$. Let $S = U\Lambda U^t$ and $(M - P) = V\Theta V^t$ be the eigendecompositions of the smaller matrices. Then from Theorem 3.1 we have,

$$\begin{aligned} R^tS^2R\mathbf{x} &= \lambda G^2\mathbf{x} \\ \frac{1}{2}\begin{bmatrix} I \\ -I \end{bmatrix} U\Lambda^2U^t\frac{1}{2}\begin{bmatrix} I & -I \end{bmatrix}\mathbf{x} &= \lambda\frac{1}{\sqrt{2}}\begin{bmatrix} U & V \\ -U & V \end{bmatrix}\begin{bmatrix} \Lambda^2 & \\ & \Theta^2 \end{bmatrix}\left(\frac{1}{\sqrt{2}}\begin{bmatrix} U & V \\ -U & V \end{bmatrix}\right)^t\mathbf{x} \end{aligned}$$

Let $\mathbf{y} = \left(\frac{1}{\sqrt{2}}\begin{bmatrix} U & V \\ -U & V \end{bmatrix}\right)^t\mathbf{x}$, then,

$$\begin{aligned} \frac{1}{4}\begin{bmatrix} U \\ -U \end{bmatrix}\Lambda^2\begin{bmatrix} U^t & -U^t \end{bmatrix}\frac{1}{\sqrt{2}}\begin{bmatrix} U & V \\ -U & V \end{bmatrix}\left(\frac{1}{\sqrt{2}}\begin{bmatrix} U & V \\ -U & V \end{bmatrix}\right)^t\mathbf{x} &= \lambda\frac{1}{\sqrt{2}}\begin{bmatrix} U & V \\ -U & V \end{bmatrix}\begin{bmatrix} \Lambda^2 & \\ & \Theta^2 \end{bmatrix}\mathbf{y} \\ \left(\frac{1}{\sqrt{2}}\begin{bmatrix} U & V \\ -U & V \end{bmatrix}\right)^t\frac{1}{4\sqrt{2}}\begin{bmatrix} U \\ -U \end{bmatrix}\Lambda^2\begin{bmatrix} U^t & -U^t \end{bmatrix}\begin{bmatrix} U & V \\ -U & V \end{bmatrix}\mathbf{y} &= \lambda\begin{bmatrix} \Lambda^2 & \\ & \Theta^2 \end{bmatrix}\mathbf{y} \\ \frac{1}{2}\begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix}\mathbf{y} &= \lambda\begin{bmatrix} I & \\ & \Theta^2 \end{bmatrix}\mathbf{y} \end{aligned}$$

Thus, the eigenvalues of the generalized eigenvalue problem, $R^tS^2R\mathbf{x} = \lambda G^2\mathbf{x}$,are 1/2 and 0, and $C_s = \frac{1}{\sqrt{2}}$. □

The above theorem states that in any norm the residual of the original system is smaller that the residual of the expanded system. We will validate this result by empirically demonstrating numerical stability of Gremban's expansion with LAMG (a known robust multilevel solver for diagonally dominant Z-matrix systems) as well as preconditioned conjugate gradient. The next section will give a brief description of LAMG.

4 Lean Algebraic Multigrid for Undirected Graphs

Multigrid is an iterative solver that relies on two main components: fine-grid smoothing and coarse-grid correction. Fine-grid smoothing should be cheap and quickly attenuate the oscillatory components of the error. This is usually done using weighted Jacobi or Gauss-Seidel. The coarse-grid correction projects the current smooth error down to a smaller grid. The error then becomes oscillatory and the process can be repeated until the current grid is small enough to solve directly. The solution is then projected back to the fine-grid. This is a called a V-cycle iteration. The coarse-grid correction can either be found geometrically, in which the aggregates are chosen based off of the underlying geometry of the problem, or algebraically, in which the aggregates are chosen based on a strength of connection. The latter case is known as algebraic multigird (AMG) and is a powerful solver for PDE-type problems in which the geometry of the underlying mesh is unknown. AMG is a attractive option for graphs since the geometry of a graph is highly complex. Livne and Brandt modified AMG specifically for graph Laplacians to create Lean Algebraic Multigrid (LAMG). The three main contributions of LAMG are low-degree elimination, caliber-1 interpolation with an energy correction step, and the use of the adaptivity to compute the strength of connection matrix.

These key modifications have been shown to be efficient and robust for most topologies of unsigned, undirected graphs. Our own implementation of LAMG in Julia has shown to be an efficient and robust solver for unsigned, undirected and weighted graphs. In a restricted subset of graph Laplacian stemming from data analysis application, we found that V-cycles were sufficient, including scale-free graphs and bipartite graphs. In a wider set including mesh-like graphs, Livne and Brandt, claimed W-cycles were necessary. A detailed explanation can be found the technical report [6].

5 Numerical Results

A Julia [5] implementation of the LAMG algorithm was developed to serve as a platform for the investigations described in this paper. There are several measurements that are used to asses the efficiency of a multigrid scheme: convergence factor, cycle complexity and effective convergence factor. The *convergence factor*, ρ, measures the asymptotic reduction in the 2-norm of the residual for a single V-cycle. To estimate ρ, we calculate the convergence factor after each cycle, ρ_k, and then report the geometric mean for a solve containing K multigrid cycles,

$$\rho_k = \frac{\|G\mathbf{y}_{k+1} - \mathbf{z}\|_2}{\|G\mathbf{y}_k - \mathbf{z}\|_2} \quad \text{and} \quad \rho \approx \left(\prod_{k=1}^{K} \rho_i\right)^{1/K}.$$

The *cycle complexity*, γ, is the number of work units (the cost of one fine-gird matrix multiplication) required for a single multigrid cycle. The *effective convergence factor* (ECF) is the average reduction in residual per work unit cost, $\text{ECF} = \rho^{1/\gamma}$.

To get an idea of how the Gremban expansion works for signed graphs, the team gathered a set of undirected, unsigned graphs from the SNAP graph database[4]. For each graph, a randomly chosen percentage, $p = [0.1, 0.8]$, of the edges were turned into negatively signed edges. For a random vector $\mathbf{b}$ the goal is to solve $L\mathbf{x} = \mathbf{b}$. The Gremban expansion, G, of L was formed and we let $\mathbf{z} = [\mathbf{b}^t, -\mathbf{b}^t]^t$. LAMG then solved the system $G\mathbf{w} = \mathbf{z}$ with V-cycles with $(1, 2)$ pre- and post- Gauss-Seidel smoothing iterations, used a cut off of a 100 nodes for the coarsest graph and a residual tolerance of 10^{-7}. Figure 1a, 1b, and 1d show the convergence factor, cycle complexity and effective convergence factor respectively. Note that due to the Gremban expansion, the size of the problem is twice the size of the original problem. Table 1 displays the graphs that had an effective convergence factor greater than .85.

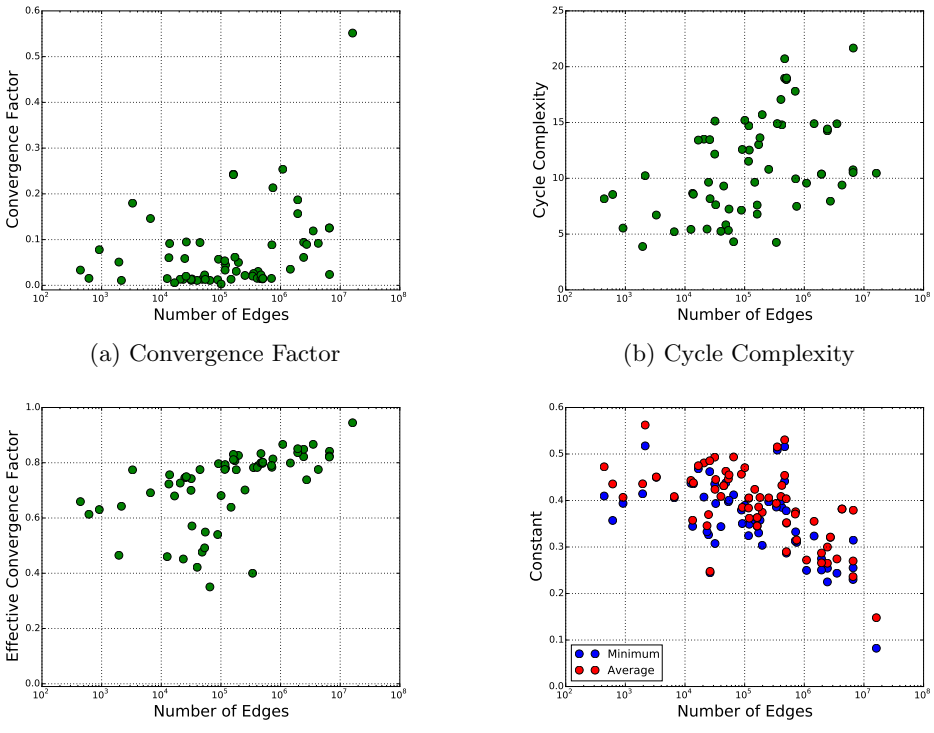

(a) Convergence Factor

(b) Cycle Complexity

(c) Effective Conference Factor

(d) Ratio of Relative Residuals

Figure 1: Performance of LAMG on Constructed Signed Graphs

After the system $G\mathbf{w} = \mathbf{z}$ was solved, an approximation to $\mathbf{x}$ was found as a weighted average, $\mathbf{x} = \alpha P_1 \mathbf{w} + \beta P_2 \mathbf{w}$, as in Theorem 3.2. The solution was found in two different ways: the true minimum of the

Graph	$2n$	$2m$	n_{level}	Solve Time(s)	Setup Time(s)	ρ	γ	ECF
eu-2005($p = .20$)	862664	16138468	6	232.1919	175.9230	0.5163	10.4607	0.9388
in-2004($p = .20$)	1353703	13126172	29	301.4328	63.1952	0.1866	13.2205	0.8808
amazon-2008($p = .30$)	735323	3523472	14	117.6931	28.8957	0.1240	14.5854	0.8666
wb-edu($p = .80$)	8863287	44185251	32	1458.5940	212.9845	0.2508	9.4042	0.8632
amazon0601($p = .60$)	403364	2443311	16	80.2968	18.8181	0.1135	14.3926	0.8597
web-Stanford($p = .60$)	255265	1941926	23	48.3831	12.3606	0.1945	10.3860	0.8541

Table 1: Constructed Signed Graphs with $ECF > .85$

residual, $\mathbf{x} = \min_{\alpha,\beta} ||L\mathbf{x} - \mathbf{b}||_2$, and a simple averaging, $\mathbf{x} = \frac{1}{2}P_1\mathbf{w} + \frac{1}{2}P_2\mathbf{w}$. Define the *ratio of relative residuals*, τ, as $\tau = ||L\mathbf{x} - \mathbf{b}||_2 \,/\, ||G\mathbf{w} - \mathbf{z}||_2$. Figure 1(d) shows the ratio of relative residuals for both the minimum constant using the optimal α and β for $\mathbf{x}$ and the averaging constant.

The average solution still produces ratio of relative residuals that are less than one. From a data-science perspective, the average solution will give a good approximation with out the extra cost of finding the minimum.

Note that PCG was not used in conjunction with LAMG. To see how PCG performed with Gremban expansion, and without the expansion on signed graphs (without LAMG), we have the following results. With the same set of graphs, we perform three different tests, PCG with Jacobi without Gremban's expansion and without LAMG, PCG with Jacobi with the expansion and without LAMG, and LAMG with the expansion and without PCG. We ran each test to a tolerance of 10^{-7} and a maximum number of iterations of $10,000$. Figure 2a, 2b, 2c and 2d compares the convergence factor, number of V-cycle verses number of iterations, the total time to setup and solve, and only solve time for LAMG and PCG.

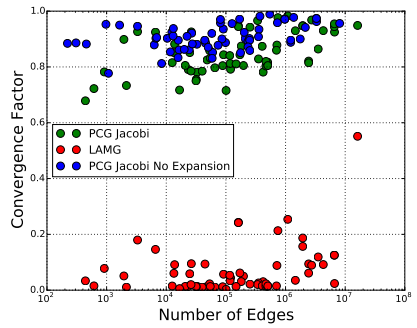

(a) Convergence Factor

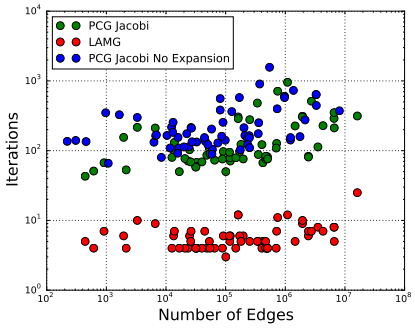

(b) Iterations

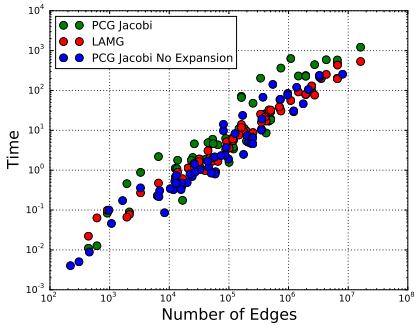

(c) Total Time

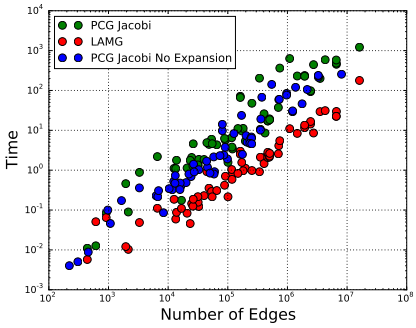

(d) Solve Time

Figure 2: LAMG vs PCG

We see that LAMG provides better convergence factors and fewer iterations for most graphs. For the total time, LAMG was 12.7% faster than PCG without expansion and 63.49% faster than PCG with the

expansion for all graphs tested. For only solve time, LAMG was faster for 92.06% of graphs than for PCG without the expansion and 98.41% of graphs than for PCG with the expansion. Thus, if solving the same problem for various right hand sides, LAMG would be the better choice since the hierarchy could be saved.

To get a better idea on how the expansion works for real-world graphs, the second test that was performed were on bipartite user-movie rating matrices, R, from the Grouplens research*. The adjacency matrix of the data is written as,

$$A = \begin{bmatrix} 0 & R \\ R^t & 0 \end{bmatrix}. \tag{6}$$

Grouplens provides four different size data sets, 100k, 1M, 10M and 20M. Each user rated the movies between 1 and 5. We scaled the data in four different ways to get a resulting SU graph.

- **User Average:** If a user rates every movie high then ratings are not proportional thus we subtract the users average rating, μ_i, from every movie that they have rated:

$$A_{ij}^{(user)} = A_{ji}^{(user)} = R_{ij} - \mu_i.$$

- **Movie Average:** Similar to above, with μ_j as a movie's average rating:

$$A_{ij}^{(movie)} = A_{ji}^{(movie)} = R_{ij} - \mu_j.$$

- **Full Average:** Scales by the user average and movie average:

$$A_{ij}^{(full)} = A_{ji}^{(full)} = R_{ij} - \frac{\mu_j + \mu_i}{2}.$$

- **Shifted:** Maps the ratings from: $[1, 2, 3, 4, 5] \to [-2, -1, 1, 2, 3]$.

Table 2 presents the convergence factor, cycle complexity, and effective convergence factor for all shifts. In the last two columns we have the ratio of relative residuals for the solution using the minimum α and β, and the average. LAMG performed well on all the matrices with effective convergence factors less than 0.70.

Graph	$2m$	Setup Time(s)	Solve Time(s)	ρ	γ	ECF	τ_{min}	τ_{max}
ml-100k_full	99997	2.179957151	0.191529036	0.0038	10.3139	0.5832	0.3656	0.4347
ml-100k_movie	1.00E+05	1.960252047	0.177932978	0.0019	10.2738	0.5434	0.3674	0.4576
ml-100k_shift	1.00E+05	3.095001936	0.165970087	0.0024	14.0861	0.6519	0.2830	0.2933
ml-100k_user	1.00E+05	1.951591969	0.18362093	0.0019	10.3179	0.5456	0.3257	0.5143
ml-1m_full	1000068	1.729801178	0.523294926	9.9600E-13	0.9997	9.8682E-13	0.5979	0.6401
ml-1m_movie	1000086	1.585048914	0.492215157	6.4000E-12	0.9997	6.3446E-12	0.4036	0.4104
ml-1m_shift	1000086	34.1253221	1.945924997	0.0022	15.5254	0.6751	0.2047	0.3667
ml-1m_user	1000086	1.656098127	0.493187904	1.8600E-12	0.9997	1.8471E-12	0.4304	0.5378
ml-10m_full	9971967	125.957454	19.60491014	0.0017	5.8705	0.3384	0.5016	0.5019
ml-10m_movie	9971981.5	124.7016339	18.84954786	0.0016	5.8630	0.3338	0.4528	0.4901
ml-10m_shift	9971981.5	131.9069471	22.86774802	0.0021	5.8612	0.3484	0.5181	0.5205
ml-10m_user	9971981.5	126.364444	18.90600181	0.0015	5.8626	0.3301	0.4811	0.4907
ml-20m_full	19947969	480.011148	33.49572492	0.0015	12.3430	0.5904	0.2928	0.4803
ml-20m_movie	19948004	573.897356	31.06468916	0.0024	14.4520	0.6590	0.1794	0.4982
ml-20m_shift	19948004	589.2875509	36.9788301	0.0025	13.9342	0.6502	0.4016	0.5092
ml-20m_user	19948004	217.217633	32.85207701	0.0020	5.7631	0.3388	0.3280	0.4796

Table 2: LAMG performance on Movielens

This indicates that the expansion, though it doubles the size of the problem, resulted in systems that were solvable. Table 3 compares the convergence factor as well as the solve time for LAMG and PCG with and without Gremban's expansion. For PCG, using Gremban's expansion resulted in a lower convergence factor.

*http://grouplens.org/datasets/movielens/

Graph	LAMG(ρ)	PCG with Expansion(ρ)	PCG w/o Expansion(ρ)	LAMG (SolveTime)	PCG w/ Expansion (SolveTime)	PCG w/o Expansion (SolveTime)
ml-100k_full	0.0038	0.7132	0.8844	0.1915	0.8981	0.8548
ml-100k_movie	0.0019	0.7180	0.9087	0.1779	0.8376	1.0951
ml-100k_shift	0.0024	0.7150	0.7740	0.1660	1.0382	0.3889
ml-100k_user	0.0019	0.7155	0.8971	0.1836	0.8027	0.8812
ml-1m_full	9.9600E-13	0.7127	0.8941	0.5233	7.5420	9.2640
ml-1m_movie	6.4000E-12	0.7193	0.9005	0.4922	6.8716	12.4039
ml-1m_shift	0.0022	0.7222	0.7880	1.9459	9.3258	4.9282
ml-1m_user	1.8600E-12	0.7191	0.9123	0.4932	6.6885	11.3779
ml-10m_full	0.0017	0.7005	0.9021	19.6049	58.3471	301.9563
ml-10m_movie	0.0016	0.6944	0.8973	18.8495	55.5631	479.5299
ml-10m_shift	0.0021	0.7017	0.7858	22.8677	66.9099	213.9651
ml-10m_user	0.0015	0.7018	0.8951	18.9060	91.6121	357.2703
ml-20m_full	0.0015	0.6988	0.8950	33.4957	192.2985	370.1953
ml-20m_movie	0.0024	0.6976	0.8963	31.0647	172.5820	186.1037
ml-20m_shift	0.0025	0.7029	0.7905	36.9788	197.5440	109.1963
ml-20m_user	0.0020	0.7036	0.9120	32.8521	187.4787	175.2370

Table 3: LAMG and PCG with Jacobi performance on Movielens

However, LAMG without PCG had convergence factors were all under .05, which is much lower than any of the PCG results. LAMG also resulted in lower solve time. Even though the problem size is doubled for Gremban's expansion, PCG with the expansion was at most 7% slower than without the expansion. For the problems tested, it can be concluded that Gremban's expansion is numerically stable. It must be noted that the condition number of the expansion cannot be bounded by the condition number for the original graph Laplacian. Thus, when using the expansion, a robust UU solver is required to be confident about efficiency across a diverse set of graphs.

6 Conclusions

Tests suggest that the Gremban expansion for signed undirected graphs performed well in conjunction with an adequate solver. Using the Gremban expansion with LAMG and V-cycles, it performed well on the constructed signed graphs as well as the movielens user-movie rating matrices. The movielens matrices give an indication of real-world signed, bipartite graphs. The expansion when used with a robust solver gave excellence convergence rates. We can conclude that the Gremban expansion is numerically stable for SU graphs when a robust solver is used. We showed a tight bound on the expansion for not only undirected, unsigned graphs but also any symmetric diagonally dominant matrix and proved that Gremban's expansion can be generalized to any diagonally dominant matrix. This gives the focus to future work, which is to create a solver for directed graph Laplacians. We can then use Gremban's generalized expansion to solve any diagonally dominant matrix.

References

[1] Jeff Bezanson, Stefan Karpinski, Viral B. Shah, and Alan Edelman. Julia: A fast dynamic language for technical computing. *CoRR*, abs/1209.5145, 2012.

[2] Keith Gremban. *Combinatorial Preconditioners for Sparse, Symmetric, Diagonally Dominant Linear Systems*. PhD thesis, Carnegie Mellon University, Pittsburgh, October 1996. CMU CS Tech Report CMU-CS-96-123.

[3] Jrme Kunegis, Stephan Schmidt, Andreas Lommatzsch, Jrgen Lerner, Ernesto W. De, and Luca Sahin Albayrak. Spectral analysis of signed graphs for clustering, prediction and visualization.

[4] Jure Leskovec and Rok Sosič. SNAP: A general purpose network analysis and graph mining library in C++. `http://snap.stanford.edu/snap`, June 2014.

[5] O.E. Livne and A. Brandt. Lean algebraic multigrid (lamg): Fast graph laplacian linear solver. *SIAM Journal of Scientific Computing*, 2011. accepted.

[6] Geoff Sanders, Tom Manteuffel, Alyson Fox, and Nate Monning. An investigation into lamg, 2015.