Spectral Graph Theory Lecture 19

Fast Laplacian Solvers by Sparsification

Daniel A. Spielman November 9, 2015

Disclaimer

These notes are not necessarily an accurate representation of what happened in class. The notes written before class say what I think I should say. I sometimes edit the notes after class to make them way what I wish I had said.

There may be small mistakes, so I recommend that you check any mathematically precise statement before using it in your own work.

These notes were last revised on November 9, 2015.

19.1 Overview

We will see how sparsification allows us to solve systems of linear equations in Laplacian matrices and their sub-matrices in nearly linear time. By "nearly-linear", I mean time $O(m \log^c (n\kappa^{-1}) \log \epsilon^{-1})$ for systems with m nonzero entries, n dimensions, condition number κ. and accuracy ϵ.

This algorithm comes from [PS14].

19.2 Today's notion of approximation

In today's lecture, I will find it convenient to define matrix approximations slightly differently from previous lectures. Today, I define $\boldsymbol{A} \approx_\epsilon \boldsymbol{B}$ to mean

$$e^{-\epsilon} \boldsymbol{A} \preccurlyeq \boldsymbol{B} \preccurlyeq e^{\epsilon} \boldsymbol{A}.$$

Note that this relation is symmetric in $\boldsymbol{A}$ and $\boldsymbol{B}$, and that for ϵ small $e^\epsilon \approx 1 + \epsilon$.

The advantage of this definition is that

$$\boldsymbol{A} \approx_\alpha \boldsymbol{B} \quad \text{and} \quad \boldsymbol{B} \approx_\beta \boldsymbol{C} \quad \text{implies} \quad \boldsymbol{A} \approx_{\alpha+\beta} \boldsymbol{C}.$$

19.3 The Idea

I begin by describing the idea behind the algorithm. This idea won't quite work. But, we will see how to turn it into one that does.

We will work with matrices that look like $\boldsymbol{M} = \boldsymbol{L} + \boldsymbol{X}$ where $\boldsymbol{L}$ is a Laplacian and $\boldsymbol{X}$ is a non-zero, non-negative diagonal matrix. Such matrices are called M-matrices. A *symmetric M-matrix* is a matrix $\boldsymbol{M}$ with nonpositive off-diagonal entries such that $\boldsymbol{M}\boldsymbol{1}$ is nonnegative and nonzero. We have encountered M-matrices before without naming them. If $G = (V, E)$ is a graph, $S \subset V$, and $G(S)$ is connected, then the submatrix of $\boldsymbol{L}_G$ indexed by rows and columns in S is an M-matrix. Algorithmically, the problems of solving systems of equations in Laplacians and symmetric M-matrices are equivalent.

The sparsification results that we learned for Laplacians translate over to M-matrices. Every M-matrix $\boldsymbol{M}$ can be written in the form $\boldsymbol{X} + \boldsymbol{L}$ where $\boldsymbol{L}$ is a Laplacian and $\boldsymbol{X}$ is a nonnegative diagonal matrix. If $\widehat{\boldsymbol{L}} \approx_\epsilon \boldsymbol{L}$, then it is easy to show (too easy for homework) that

$$\boldsymbol{X} + \widehat{\boldsymbol{L}} \approx_\epsilon \boldsymbol{X} + \boldsymbol{L}.$$

In Lecture 7, Lemma 7.3.1, we proved that if $\boldsymbol{X}$ has at least one nonzero entry and if $\boldsymbol{L}$ is connected, then $\boldsymbol{X} + \boldsymbol{L}$ is nonsingular. We write such a matrix in the form $\boldsymbol{M} = \boldsymbol{D} - \boldsymbol{A}$ where $\boldsymbol{D}$ is positive diagonal and $\boldsymbol{A}$ is nonnegative, and note that its being nonsingular and positive semidefinite implies

$$\boldsymbol{D} - \boldsymbol{A} \succ 0 \quad \iff \quad \boldsymbol{D} \succ \boldsymbol{A}. \tag{19.1}$$

Using the Perron-Frobenius theorem, one can also show that

$$\boldsymbol{D} \succ -\boldsymbol{A}. \tag{19.2}$$

Multiplying $\boldsymbol{M}$ by $\boldsymbol{D}^{-1/2}$ on either side, we obtain

$$\boldsymbol{I} - \boldsymbol{D}^{-1/2}\boldsymbol{A}\boldsymbol{D}^{-1/2}.$$

Define

$$\boldsymbol{B} = \boldsymbol{D}^{-1/2}\boldsymbol{A}\boldsymbol{D}^{-1/2},$$

and note that inequalities (19.1) and (19.2) imply that all eigenvalues of $\boldsymbol{B}$ have absolute value strictly less than 1.

It suffices to figure out how to solve systems of equations in $\boldsymbol{I} - \boldsymbol{B}$. One way to do this is to exploit the power series expansion:

$$(\boldsymbol{I} - \boldsymbol{B})^{-1} = \boldsymbol{I} + \boldsymbol{B} + \boldsymbol{B}^2 + \boldsymbol{B}^3 + \cdots$$

However, this series might need many terms to converge. We can figure out how many. If the largest eigenvalue of $\boldsymbol{B}$ is $(1 - \kappa) < 1$, then we need at least $1/\kappa$ terms.

We can write a series with fewer terms if we express it as a product instead of as a sum:

$$\sum_{i \geq 0} \boldsymbol{B}^i = \prod_{j \geq 1} (\boldsymbol{I} + \boldsymbol{B}^{2^j}).$$

To see why this works, look at the first few terms

$$(\boldsymbol{I}+\boldsymbol{B})(\boldsymbol{I}+\boldsymbol{B}^2)(\boldsymbol{I}+\boldsymbol{B}^4) = (\boldsymbol{I}+\boldsymbol{B}+\boldsymbol{B}^2+\boldsymbol{B}^3)(\boldsymbol{I}+\boldsymbol{B}^4) = (\boldsymbol{I}+\boldsymbol{B}+\boldsymbol{B}^2+\boldsymbol{B}^3)+\boldsymbol{B}^4(\boldsymbol{I}+\boldsymbol{B}+\boldsymbol{B}^2+\boldsymbol{B}^3).$$

We only need $O(\log \kappa^{-1})$ terms of this product to obtain a good approximation of $(\boldsymbol{I} - \boldsymbol{B})^{-1}$. The obstacle to quickly applying a series like this is that the matrices $\boldsymbol{I} + \boldsymbol{B}^{2^j}$ are probably dense. We know how to solve this problem: we can sparsify them! I'm not saying that flippantly. We actually do know how to sparify matrices of this form.

But, simply sparsifying the matrices $\boldsymbol{I} + \boldsymbol{B}^{2^j}$ does not solve our problem because approximation is not preserved by products. That is, even if $\boldsymbol{A} \approx_\epsilon \widehat{\boldsymbol{A}}$ and $\boldsymbol{B} \approx_\epsilon \widehat{\boldsymbol{B}}$, $\widehat{\boldsymbol{A}}\widehat{\boldsymbol{B}}$ could be a very poor approximation of $\boldsymbol{AB}$. In fact, since the product $\widehat{\boldsymbol{A}}\widehat{\boldsymbol{B}}$ is not necessarily symmetric, we haven't even defined what it would mean for it to approximate $\boldsymbol{AB}$.

19.4 A symmetric expansion

We will now derive a way of expanding $(\boldsymbol{I} - \boldsymbol{B})^{-1}$ that is amenable to approximation. We begin with an alternate derivation of the series we saw before. Note that

$$(\boldsymbol{I} - \boldsymbol{B})(\boldsymbol{I} + \boldsymbol{B}) = (\boldsymbol{I} - \boldsymbol{B}^2),$$

and so

$$(\boldsymbol{I} - \boldsymbol{B}) = (\boldsymbol{I} - \boldsymbol{B}^2)(\boldsymbol{I} + \boldsymbol{B})^{-1}.$$

Taking the inverse of both sides gives

$$(\boldsymbol{I} - \boldsymbol{B})^{-1} = (\boldsymbol{I} + \boldsymbol{B})(\boldsymbol{I} - \boldsymbol{B}^2)^{-1}.$$

We can then apply the same expansion to $(\boldsymbol{I} - \boldsymbol{B}^2)^{-1}$ to obtain

$$(\boldsymbol{I} - \boldsymbol{B})^{-1} = (\boldsymbol{I} + \boldsymbol{B})(\boldsymbol{I} + \boldsymbol{B}^2)(\boldsymbol{I} - \boldsymbol{B}^4)^{-1}.$$

What we need is a symmetric expansion. We use

$$(\boldsymbol{I} - \boldsymbol{B})^{-1} = \frac{1}{2}\boldsymbol{I} + \frac{1}{2}(\boldsymbol{I} + \boldsymbol{B})(\boldsymbol{I} - \boldsymbol{B}^2)^{-1}(\boldsymbol{I} + \boldsymbol{B}). \tag{19.3}$$

We will verify this by multiplying the right hand side by $(\boldsymbol{I} - \boldsymbol{B})$:

$$(\boldsymbol{I} + \boldsymbol{B})(\boldsymbol{I} - \boldsymbol{B}^2)^{-1}(\boldsymbol{I} + \boldsymbol{B})(\boldsymbol{I} - \boldsymbol{B}) = (\boldsymbol{I} + \boldsymbol{B})(\boldsymbol{I} - \boldsymbol{B}^2)^{-1}(\boldsymbol{I} - \boldsymbol{B}^2) = \boldsymbol{I} + \boldsymbol{B};$$

so

$$\frac{1}{2}\left[\boldsymbol{I} + (\boldsymbol{I} + \boldsymbol{B})(\boldsymbol{I} - \boldsymbol{B}^2)^{-1}(\boldsymbol{I} + \boldsymbol{B})\right](\boldsymbol{I} - \boldsymbol{B}) = \frac{1}{2}\left[(\boldsymbol{I} - \boldsymbol{B}) + (\boldsymbol{I} + \boldsymbol{B})\right] = \boldsymbol{I}.$$

This expression for $(\boldsymbol{I} - \boldsymbol{B})^{-1}$ plays nicely with matrix approximations. If

$$\boldsymbol{M}_1 \approx_\epsilon (\boldsymbol{I} - \boldsymbol{B}^2),$$

then you can show

$$(\boldsymbol{I} - \boldsymbol{B})^{-1} \approx_\epsilon \frac{1}{2}\left[\boldsymbol{I} + (\boldsymbol{I} + \boldsymbol{B})\boldsymbol{M}_1^{-1}(\boldsymbol{I} + \boldsymbol{B})\right].$$

If we can apply $\boldsymbol{M}_1^{-1}$ quickly and if $\boldsymbol{B}$ is sparse, then we can quickly approximate $(\boldsymbol{I} - \boldsymbol{B})^{-1}$. You may now be wondering how we will construct such an $\boldsymbol{M}_1$. The answer, in short, is "recursively".

19.5 $\boldsymbol{D}$ and $\boldsymbol{A}$

Unfortunately, we are going to need to stop writting matrices in terms of $\boldsymbol{I}$ and $\boldsymbol{B}$, and return to writing them in terms of $\boldsymbol{D}$ and $\boldsymbol{A}$. The reason this is unfortunate is that it makes for longer expressions.

The analog of (19.3) is

$$(\boldsymbol{D}-\boldsymbol{A})^{-1} = \frac{1}{2}\left[\boldsymbol{D}^{-1} + (\boldsymbol{I}+\boldsymbol{D}^{-1}\boldsymbol{A})(\boldsymbol{D}-\boldsymbol{A}\boldsymbol{D}^{-1}\boldsymbol{A})^{-1}(\boldsymbol{I}+\boldsymbol{A}\boldsymbol{D}^{-1})\right]. \tag{19.4}$$

In order to be able to work with this expression inductively, we need to check that the middle matrix is an M-matrix.

Lemma 19.5.1. *If $\boldsymbol{D}$ is a diagonal matrix and $\boldsymbol{A}$ is a nonnegative matrix so that $\boldsymbol{M} = \boldsymbol{D}-\boldsymbol{A}$ is an M-matrix, then*

$$\boldsymbol{M}_1 = \boldsymbol{D} - \boldsymbol{A}\boldsymbol{D}^{-1}\boldsymbol{A}$$

is also an M-matrix.

Proof. As the off-diagonal entries of this matrix are symmetric and nonpositive, it suffices to prove that $\boldsymbol{M}\boldsymbol{1} \geq 0$ and $\boldsymbol{M}\boldsymbol{1} \neq \boldsymbol{0}$. To compute the row sums set

$$\boldsymbol{d} = \boldsymbol{D}\boldsymbol{1} \quad \text{and} \quad \boldsymbol{a} = \boldsymbol{A}\boldsymbol{1},$$

and note that $\boldsymbol{d}-\boldsymbol{a} \geq 0$ and $\boldsymbol{d}-\boldsymbol{a} \neq \boldsymbol{0}$. For $\boldsymbol{M}_1$, we have

$$(\boldsymbol{D}-\boldsymbol{A}\boldsymbol{D}^{-1}\boldsymbol{A})\boldsymbol{1} = \boldsymbol{d} - \boldsymbol{A}\boldsymbol{D}^{-1}\boldsymbol{a} \geq \boldsymbol{d} - \boldsymbol{A}\boldsymbol{1} = \boldsymbol{d}-\boldsymbol{a},$$

which is nonnegative and not exactly zero. □

We will apply transformation like this many times during our algorithm. To keep track of progress, I say that $(\boldsymbol{D},\boldsymbol{A})$ is an (α,β)-pair if

a. $\boldsymbol{D}$ is positive diagonal,

b. $\boldsymbol{A}$ is nonnegative (and can have diagonal entries), and

c. $\alpha\boldsymbol{D} \succcurlyeq \boldsymbol{A}$ and $\beta\boldsymbol{D} \succcurlyeq -\boldsymbol{A}$.

For our initial matrix $\boldsymbol{M} = \boldsymbol{D}-\boldsymbol{A}$, we know that there is some number $\kappa > 0$ for which $(\boldsymbol{D},\boldsymbol{A})$ is a $(1-\kappa, 1-\kappa)$-pair.

At the end of our recursion we will seek a $(1/4, 1/4)$-pair. When we have such a pair, we can just approximate $\boldsymbol{D}-\boldsymbol{A}$ by $\boldsymbol{D}$.

Lemma 19.5.2. *If $\boldsymbol{M} = \boldsymbol{D}-\boldsymbol{A}$ and $(\boldsymbol{D},\boldsymbol{A})$ is a $(1/4,1/4)$-pair, then*

$$\boldsymbol{M} \approx_{1/3} \boldsymbol{D}.$$

Proof. We have
$$\boldsymbol{M} = \boldsymbol{D} - \boldsymbol{A} \preccurlyeq (1 + 1/4)\boldsymbol{D} \leq e^{1/4}\boldsymbol{D},$$
and
$$\boldsymbol{M} = \boldsymbol{D} - \boldsymbol{A} \succcurlyeq \boldsymbol{D} - (1/4)\boldsymbol{D} = (3/4)\boldsymbol{D} \succcurlyeq e^{-1/3}\boldsymbol{D}.$$
□

Lemma 19.5.3. *If $(\boldsymbol{D}, \boldsymbol{A})$ is an (α, α)-pair, then $(\boldsymbol{D}, \boldsymbol{A}\boldsymbol{D}^{-1}\boldsymbol{A})$ is an $(\alpha^2, 0)$-pair.*

Proof. From Lecture 14, Lemma 3.1, we know that the condition of the lemma is equivalent to the assertion that all eigenvalues of $\boldsymbol{D}^{-1}\boldsymbol{A}$ have absolute value at most α, and that the conclusion is equivalent to the assertion that all eigenvalues of $\boldsymbol{D}^{-1}\boldsymbol{A}\boldsymbol{D}^{-1}\boldsymbol{A}$ lie between 0 and α^2, which is immediate as they are the squares of the eigenvalues of $\boldsymbol{D}^{-1}\boldsymbol{A}$. □

So, if we start with matrices $\boldsymbol{D}$ and $\boldsymbol{A}$ that are a $(1 - \kappa, 1 - \kappa)$-pair, then after applying this transformation approximately $\log \kappa^{-1} + 2$ times we obtain a $(1/4, 0)$-pair. But, the matrices in this pair could be dense. To keep them sparse, we need to figure out how approximating $\boldsymbol{D} - \boldsymbol{A}$ degrades its quality.

Lemma 19.5.4. *If $\epsilon \leq 1/3$,*

a. $(\boldsymbol{D}, \boldsymbol{A})$ *is a* $(1 - \kappa, 0)$ *pair,*

b. $\boldsymbol{D} - \boldsymbol{A} \approx_\epsilon \widehat{\boldsymbol{D}} - \widehat{\boldsymbol{A}}$*, and*

c. $\boldsymbol{D} \approx_\epsilon \widehat{\boldsymbol{D}}$*,*

then $\widehat{\boldsymbol{D}} - \widehat{\boldsymbol{A}}$ is an $(1 - \kappa e^{-2\epsilon}, 3\epsilon)$-pair.

Proof. First observe that
$$(1 - \kappa)\boldsymbol{D} \succcurlyeq \boldsymbol{A} \iff \boldsymbol{D} - \boldsymbol{A} \succcurlyeq \kappa\boldsymbol{D}.$$
Then, compute
$$\widehat{\boldsymbol{D}} - \widehat{\boldsymbol{A}} \succcurlyeq e^{-\epsilon}(\boldsymbol{D} - \boldsymbol{A}) \succcurlyeq e^{-\epsilon}\kappa\boldsymbol{D} \succcurlyeq e^{-2\epsilon}\kappa\widehat{\boldsymbol{D}}.$$
For the other side, compute
$$e^{2\epsilon}\widehat{\boldsymbol{D}} \succcurlyeq e^{\epsilon}\boldsymbol{D} \succcurlyeq e^{\epsilon}(\boldsymbol{D} - \boldsymbol{A}) \succcurlyeq (\widehat{\boldsymbol{D}} - \widehat{\boldsymbol{A}}).$$
For $\epsilon \leq 1/3$, $3\epsilon \geq e^{2\epsilon} - 1$, so
$$3\epsilon\widehat{\boldsymbol{D}} \succcurlyeq (e^{2\epsilon} - 1)\widehat{\boldsymbol{D}} \succcurlyeq -\widehat{\boldsymbol{A}}.$$
□

It remains to confirm that sparsification satisfies the requirements of this lemma. The reason this might not be obvious is that we allow $\boldsymbol{A}$ to have nonnegative diagonal elements. While this does not interfere with condition b, you might be concerned that it would interfere with condition c. It need not.

Let $\boldsymbol{C}$ be the diagonal of $\boldsymbol{A}$, and let $\boldsymbol{L}$ be the Laplacian of the graph with adjacency matrix $\boldsymbol{A}-\boldsymbol{C}$, and set $\boldsymbol{X}$ so that $\boldsymbol{X}+\boldsymbol{L}=\boldsymbol{D}-\boldsymbol{A}$. Let $\widetilde{\boldsymbol{L}}$ be a sparse ϵ-approximation of $\boldsymbol{L}$. By computing the quadratic form in elementary unit vectors, you can check that the diagonals of $\boldsymbol{L}$ and $\widetilde{\boldsymbol{L}}$ approximate each other. If we now write $\widetilde{\boldsymbol{L}}=\widetilde{\boldsymbol{D}}-\widetilde{\boldsymbol{A}}$, where $\widetilde{\boldsymbol{A}}$ has zero diagonal, and set

$$\widehat{\boldsymbol{D}}=\widetilde{\boldsymbol{D}}+\boldsymbol{C} \quad \text{and} \quad \widehat{\boldsymbol{A}}=\widetilde{\boldsymbol{A}}+\boldsymbol{C}$$

You can now check that $\widehat{\boldsymbol{D}}$ and $\widehat{\boldsymbol{A}}$ satisfy the requirements of Lemma 19.5.4.

You might wonder why we bother to keep diagonal elements in a matrix like $\boldsymbol{A}$. It seems simpler to get rid of them. However, we want $(\boldsymbol{D},\boldsymbol{A})$ to be an (α,β) pair, and removing subtracting $\boldsymbol{C}$ from both of them would make β worse. This might not matter too much as we have good control over β. But, I don't yet see a nice way to carry out a proof that exploits this.

19.6 Sketch of the construction

We begin with an M-matrix $\boldsymbol{M}_0=\boldsymbol{D}_0-\boldsymbol{A}_0$. Since this matrix is nonsingular, there is a $\kappa_0>0$ so that $(\boldsymbol{D}_0,\boldsymbol{A}_0)$ is a $(1-\kappa_0,1-\kappa_0)$ pair.

We now know that the matrix

$$\boldsymbol{D}_0-\boldsymbol{A}_0\boldsymbol{D}_0^{-1}\boldsymbol{A}_0$$

is an M-matrix and that $(\boldsymbol{D}_0,\boldsymbol{A}_0\boldsymbol{D}_0^{-1}\boldsymbol{A}_0)$ is a $((1-\kappa_0)^2,0)$-pair. Define κ_1 so that $1-\kappa_1=(1-\kappa)_0^2$, and note that κ_1 is approximately $2\kappa_0$. Lemma 19.5.4 and the discussion following it tells us that there is a $(1-\kappa_1 e^{-2\epsilon},3\epsilon)$-pair $(\boldsymbol{D}_1,\boldsymbol{A}_1)$ so that

$$\boldsymbol{D}_1-\boldsymbol{A}_1\approx_\epsilon \boldsymbol{D}_0-\boldsymbol{A}_0\boldsymbol{D}_0^{-1}\boldsymbol{A}_0$$

and so that $\boldsymbol{A}_1$ has $O(n/\epsilon^2)$ nonzero entries.

Continuing inductively for some number k steps, we find $(1-\kappa_i,3\epsilon)$ pairs $(\boldsymbol{D}_i,\boldsymbol{A}_i)$ so that

$$\boldsymbol{M}_i=\boldsymbol{D}_i-\boldsymbol{A}_i$$

has $O(n/\epsilon^2)$ nonzero entries, and

$$\boldsymbol{M}_i\approx_\epsilon \boldsymbol{D}_i-\boldsymbol{A}_{i-1}\boldsymbol{D}_{i-1}^{-1}\boldsymbol{A}_{i-1}.$$

For the i such that κ_i is small, κ_{i+1} is approximately twice κ_i. So, for $k=2+\log_2 1/\kappa$ and ϵ close to zero, we can guarantee that $(\boldsymbol{D}_k,\boldsymbol{A}_k)$ is a $(1/4,1/4)$ pair.

We now see how this construction allows us to approximately solve systems of equations in $\boldsymbol{D}_0-\boldsymbol{A}_0$, and how we must set ϵ for it to work. For every $0\le i<k$, we have

$$(\boldsymbol{D}_i-\boldsymbol{A}_i)^{-1}\frac{1}{2}\boldsymbol{D}_i^{-1}+\frac{1}{2}(\boldsymbol{I}+\boldsymbol{D}_i^{-1}\boldsymbol{A}_i)(\boldsymbol{D}_i-\boldsymbol{A}_i\boldsymbol{D}_i^{-1}\boldsymbol{A}_i)^{-1}(\boldsymbol{I}+\boldsymbol{A}_i\boldsymbol{D}_i^{-1})\approx_\epsilon \frac{1}{2}\boldsymbol{D}_i^{-1}+\frac{1}{2}(\boldsymbol{I}+\boldsymbol{D}_i^{-1}\boldsymbol{A}_i)(\boldsymbol{D}_{i+1}-\boldsymbol{A}_{i+1})^{-1}(\boldsymbol{I}-$$

and

$$(\boldsymbol{D}_k - \boldsymbol{A}_k)^{-1} \approx_{1/3} \boldsymbol{D}_k^{-1}.$$

By substituting through each of these approximations, we obtain solutions to systems of equations in $\boldsymbol{D}_0 - \boldsymbol{A}_0$ with accuracy $1/3 + k\epsilon$. So, we should set $k\epsilon = 1/3$, and thus

$$\epsilon = 1/(2 + \log_2 \kappa^{-1}).$$

The dominant cost of the resulting algorithm will be the multiplication of vectors by $2k$ matrices of $O(n/\epsilon^2)$ entries, with a total cost of

$$O(n(\log_2(1/\kappa))^3).$$

19.7 Making the construction efficient

In the above construction, I just assumed that appropriate sparsifiers exist, rather than constructing them efficiently. To construct them efficiently, we need two ideas. The first is that we need to be able to quickly approximate effective resistances so that we can use the sampling algorithm from Lecture 17.

The second is to observe that we do not actually want to form the matrix $\boldsymbol{A}\boldsymbol{D}^{-1}\boldsymbol{A}$ before sparsifying it, as that could take too long. Instead, we express it as a product of cliques that have succinct descriptions, and we form the sum of approximations of each of those.

19.8 Improvements

The fastest known algorithms for solving systems of equations run in time $O(m\sqrt{\log n}\log \epsilon^{-1})$ [CKM+14]. The algorithm I have presented here can be substantially improved by combining it with Cholesky factorization. This both gives an efficient parallel algorithm, and proves the existence of an approximate inverse for every M-matrix that has a linear number of nonzeros [LPS15].

References

[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 mlog1/2n time. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pages 343–352, New York, NY, USA, 2014. ACM.

[LPS15] Yin Tat Lee, Richard Peng, and Daniel A. Spielman. Sparsified cholesky solvers for SDD linear systems. *CoRR*, abs/1506.08204, 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.