

LOCALIZED SPECTRUM SLICING

LIN LIN

ABSTRACT. Given a sparse Hermitian matrix A and a real number μ , we construct a set of sparse vectors, each approximately spanned only by eigenvectors of A corresponding to eigenvalues near μ . This set of vectors spans the column space of a localized spectrum slicing (LSS) operator, and is called an LSS basis set. The sparsity of the LSS basis set is related to the decay properties of matrix Gaussian functions. We present a divide-and-conquer strategy with controllable error to construct the LSS basis set. This is a purely algebraic process using only submatrices of A , and can therefore be applied to general sparse Hermitian matrices. The LSS basis set leads to sparse projected matrices with reduced sizes, which allows the projected problems to be solved efficiently with techniques using sparse linear algebra. As an example, we demonstrate that the LSS basis set can be used to solve interior eigenvalue problems for a discretized second order partial differential operator in one-dimensional and two-dimensional domains, as well as for a matrix of general sparsity pattern.

1. INTRODUCTION

Let A be an $n \times n$ large, sparse, Hermitian matrix. In many applications in science and engineering, one would like to find eigenvalues and eigenfunctions of A near a given real number μ . As a motivating problem, we consider A to be obtained from a certain discretization (e.g., finite difference or finite element discretization) of a second order partial differential operator of the form $-\Delta + V(x)$, where Δ is the Laplacian operator, and $V(x)$ is a potential function. Depending on the context and the choice of V , this type of problem can arise from quantum mechanics, wave propagation, electromagnetism, etc.

When μ locates inside the spectrum of A , the eigenvalues to be computed are called interior eigenvalues. These interior eigenvalues and corresponding eigenfunctions are in general difficult to compute. Since n is large and A is sparse, iterative methods such as inverse power method [13], preconditioned conjugate gradient type of methods [5, 6, 17], and shift-inverse Lanczos type of methods [19, 29] are desirable. The effectiveness of such methods often depends on the availability of a good preconditioner that can approximately apply $(A - \mu I)^{-1}$ to vectors, and such a preconditioner can be difficult to construct.

Another type of method that has recently received an increasing amount of attention is based on the construction of a matrix function $f_\mu(A)$, where the corresponding scalar function $f_\mu(z)$ only takes significant values on a small interval near

Received by the editor November 20, 2014 and, in revised form, October 10, 2015 and March 16, 2016.

2010 *Mathematics Subject Classification*. Primary 65F60, 65F50, 65F15, 65N22.

Key words and phrases. Spectrum slicing, localization, decay properties, basis set, interior eigenvalue problem.

μ on the real line. Such a matrix function $f_\mu(A)$ can be called a *spectrum slicing* operator, since for any vector $v \in \mathbb{C}^n$, $f_\mu(A)v$ is approximately only spanned by eigenvectors of A corresponding to eigenvalues near μ , and the vector $f_\mu(A)v$ is said to be *spectrally localized*. The spectrum slicing operator can be simultaneously applied to a set of random vectors $V = [v_1, \dots, v_p]$. When p is large enough but is still small compared to n , the subspace spanned by

$$W = f_\mu(A)V$$

will approximately contain the subspace of all eigenvectors corresponding to eigenvalues near μ . Let

$$A_W = W^*AW, \quad B_W = W^*W,$$

then the desired eigenvalues and eigenvectors can be computed via the solution of a generalized eigenvalue problem,

$$(1.1) \quad A_W C = B_W C \Theta.$$

In practice $f_\mu(A)$ can be constructed through relatively high order Chebyshev polynomials [28], or contour integral based methods [25, 27]. It should be noted that contour integral based methods still require solving equations of the form $(A - zI)^{-1}v$ where z is close to μ in the complex plane, either through direct methods or iterative methods.

In general the spectrum slicing operator $f_\mu(A)$ is a dense matrix. Therefore the matrix $W = f_\mu(A)V$ is in general a dense matrix, regardless of how the initial matrix $V \in \mathbb{C}^{n \times p}$ is chosen. Furthermore, the matrices A_W, B_W are in general dense matrices that do not reveal much structure to be further exploited, and the solution of the projected problem (1.1) may still be expensive when p is large.

1.1. Contribution. In this paper, we consider the use of a simple choice of Gaussian function with a positive number σ

$$(1.2) \quad f_{\sigma, \mu}(z) = e^{-\frac{(z-\mu)^2}{\sigma^2}},$$

and the corresponding matrix Gaussian function $f_{\sigma, \mu}(A)$ is spectrally localized near μ with width proportional to σ . We demonstrate that under a proper choice of σ , $f_{\sigma, \mu}(A)$ can have many entries that are small in magnitude, so that after truncating these small entries the resulting matrix is close to being a spectrum slicing operator but is also sparse. In this sense, $f_{\sigma, \mu}(A)$ is called a *localized spectrum slicing (LSS) operator*. The parameter σ plays a crucial role in terms of the computational efficiency, since it balances the spatial locality (measured by the sparsity of the truncated $f_{\sigma, \mu}(A)$), and spectral locality (measured by the approximate rank of $f_{\sigma, \mu}(A)$). We also remark that the use of the Gaussian function in (1.2) is not essential, and other smooth approximation of the Dirac- δ function can be used as well.

We demonstrate that the LSS operator $f_{\sigma, \mu}(A)$ can be constructed in a divide-and-conquer method with controllable error using only a sequence of submatrices of A with $\mathcal{O}(n)$ cost, under certain assumptions of the behavior of the sparsity, spectral radius, and sizes of submatrices of A as n increases. The column space of the LSS operator is spanned by a sparse matrix $U \in \mathbb{C}^{n \times p}$, and the subspace spanned by U will approximately contain the subspace of eigenvectors to be computed. As a result, the projected matrices

$$(1.3) \quad A_U = U^*AU, \quad B_U = U^*U$$

are sparse matrices. In this aspect, the matrix U can be regarded as a specially tailored basis set for representing the subspace approximately spanned by eigenvectors of A near μ , and each column of U is localized both spectrally and spatially. In the following text U is called a *localized spectrum slicing (LSS) basis set*. The LSS basis set can be constructed without explicitly constructing the LSS operator. The generalized eigenvalue problem for the sparse projected matrices A_U, B_U may be solved both by direct methods and by methods using sparse linear algebra techniques. During the construction of the LSS operator and/or the LSS basis set, a good global preconditioner for $(A - \mu I)^{-1}$ is not needed. We demonstrate the construction of the LSS basis set and its use for solving interior eigenvalues problems for matrices obtained from discretizing second order partial differential operators, and find that the use of the LSS basis set can be more efficient than solving the global problem directly for matrices of large sizes. We also apply the LSS method to a general matrix from the University of Florida matrix collection [8].

1.2. Related work. The spectral locality of the LSS operator is valid by construction. Comparatively the spatial locality of the LSS operator is less obvious, and is given more precisely by the *decay properties* of matrix functions that are analytic in a certain region in the complex plane (see e.g. [1–3]). The decay properties of matrix functions were first realized for matrix inverse A^{-1} (i.e., $f(z) = z^{-1}$), where A is a banded, positive definite matrix [9,10]. The method for showing decay properties relies on whether $f(z)$ can be well approximated by a low order Chebyshev polynomial evaluated at the eigenvalues of A , and this method is therefore generalizable to any analytic function $f(z)$ for banded matrices A . In order to generalize from banded matrices to general sparse matrices, decay properties should be defined using geodesic distances of the graph induced by A . These techniques have been shown in [1,3] and the references therein, for demonstrating the decay properties of, e.g., Fermi-Dirac operators in electronic structure theory. These techniques are directly used for showing the decay properties of the LSS operator in this work, which then allows the construction of the divide-and-conquer method. In physics literature, such decay property is dubbed the “near-sightedness property” and is vastly studied using various models (see e.g. [18,24,26]). The decay property is also used for constructing linear scaling algorithms [4,12] for density functional theory calculations.

1.3. Contents. The rest of this paper is organized as follows. We introduce the decay properties of matrix functions and in particular the localized spectrum slicing operator in section 2. Based on the decay properties, section 3 describes a divide-and-conquer algorithm for constructing the LSS operator and the LSS basis set, and provides the error bound and computational complexity. The computation of interior eigenvalues and a domain partitioning strategy for general sparse matrices are also discussed. We demonstrate numerical results using the LSS basis set for solving interior eigenvalue problems in section 4, and discuss the conclusion and future work in section 5.

2. PRELIMINARIES

2.1. Notation. The (i, j) -th element of a matrix $A \in \mathbb{C}^{n \times n}$ is denoted by A_{ij} . The submatrix of A corresponding to a set of row indices \mathcal{I} and a set of column

indices \mathcal{I} is denoted by $A_{\mathcal{I},\mathcal{J}}$. Using MATLAB notation, all elements in the i -th row of A are denoted by $A_{i,:}$, and all elements in a set of rows \mathcal{I} are denoted by $A_{\mathcal{I},:}$. Similarly, all elements in the j -th column of A are denoted by $A_{:,j}$, and all elements for a set of columns \mathcal{J} are denoted by $A_{:,\mathcal{J}}$. The k -th power of A is denoted by A^k . The matrix p -norm of A is denoted by $\|A\|_p$, and the vector p -norm of a vector u is denoted by $\|u\|_p$ ($p \geq 1$). The max norm of a matrix is denoted by $\|A\|_{\max} \equiv \max_{i,j}\{|A_{ij}|\}$, which is the same as the ∞ -norm of a vector of length n^2 , formed by all the elements of A . The Hermitian conjugate of A is denoted by A^* . Depending on the context, we may also refer to a matrix as an *operator*.

A Hermitian matrix A induces an undirected graph $G = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \{i | i = 1, \dots, n\}$, and $\mathcal{E} = \{(i, j) | A_{ij} \neq 0, 1 \leq i, j \leq n\}$. Each element in \mathcal{V} is called a vertex, and each element in \mathcal{E} is called an edge. The cardinality of a set of indices \mathcal{I} is denoted by $|\mathcal{I}|$.

A Hermitian matrix $A \in \mathbb{C}^{n \times n}$ has the eigendecomposition

$$(2.1) \quad AX = X\Lambda.$$

Here $\Lambda = \text{diag}[\lambda_1, \dots, \lambda_n]$ is a diagonal matrix containing the (real) eigenvalues of A and we assume $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ are ordered nondecreasingly. $X = [x_1, \dots, x_n]$ and x_i is the eigenvector corresponding to the eigenvalue λ_i . If all eigenvalues (and corresponding eigenvectors) to be computed are within a small interval $(\mu - c, \mu + c)$ on the real line with $\lambda_1 < \mu - c < \mu + c < \lambda_n$, then this problem is called an interior eigenvalue problem.

2.2. Decay property of matrix functions. In this section, we provide a short but self-contained description of the decay properties of $f_{\sigma,\mu}(A)$. More details on the description of the decay properties of general matrix functions can be found in [1] and the references therein.

Let k be a nonnegative integer, and let \mathbb{P}_k be the set of all polynomials of degrees less than or equal to k with real coefficients. Without loss of generality we assume the eigenvalues of A are within the interval $(-1, 1)$. For a real continuous function f on $[-1, 1]$, the best approximation error is defined as

$$(2.2) \quad E_k(f) = \min_{p \in \mathbb{P}_k} \left\{ \|f - p\|_{\infty} \equiv \max_{-1 \leq x \leq 1} |f(x) - p(x)| \right\}.$$

Consider an ellipse in the complex plane \mathbb{C} with foci in -1 and 1 , and let $a > 1, b > 0$ be the half axes so that the vertices of the ellipse are $a, -a, ib, -ib$, respectively. Let the sum of the half axes be $\chi = a + b$, then using the identity $a^2 - b^2 = 1$ we have

$$a = \frac{\chi^2 + 1}{2\chi}, \quad b = \frac{\chi^2 - 1}{2\chi}.$$

Thus the ellipse is determined only by χ , and such ellipse is denoted by \mathcal{E}_{χ} . Then Bernstein's theorem [23] is stated as follows.

Theorem 2.1 (Bernstein). *Let $f(z)$ be analytic in \mathcal{E}_{χ} with $\chi > 1$, and $f(z)$ is a real-valued function for real z . Then*

$$(2.3) \quad E_k(f) \leq \frac{2M(\chi)}{\chi^k(\chi - 1)},$$

where

$$(2.4) \quad M(\chi) = \sup_{z \in \mathcal{E}_{\chi}} |f(z)|.$$

Using Theorem 2.1, a more quantitative description of the approximation properties for $f_{\sigma,\mu}(z)$ in (1.2) is given in Theorem 2.2.

Theorem 2.2. *Let $f_{\sigma,\mu}(z)$ be a Gaussian function defined in (1.2), then for any $\alpha > 0$,*

$$(2.5) \quad E_k(f_{\sigma,\mu}) \leq \frac{2}{\alpha\sigma} e^{\alpha^2} (1 + \alpha\sigma)^{-k}.$$

Proof. For any $\mu \in (-1, 1), \sigma > 0$, the Gaussian function $f_{\sigma,\mu}$ is analytic in any ellipse \mathcal{E}_χ with $\chi > 1$, then

$$M(\chi) = \sup_{z \equiv x+iy \in \mathcal{E}_\chi} |f_{\sigma,\mu}(x+iy)| \leq \sup_{z \equiv x+iy \in \mathcal{E}_\chi} e^{\frac{y^2}{\sigma^2}} \leq e^{\frac{(x-\frac{1}{\chi})^2}{4\sigma^2}}.$$

For any $\alpha > 0$, let

$$(2.6) \quad \chi = 1 + \alpha\sigma,$$

then $\chi - \frac{1}{\chi} \leq 2\alpha\sigma$, and

$$(2.7) \quad M(1 + \alpha\sigma) \leq e^{\alpha^2}.$$

Using Theorem 2.1, (2.5) is the direct consequence of (2.7) and the choice of χ in (2.6). \square

For the graph $G = (\mathcal{V}, \mathcal{E})$ associated with the matrix A and vertices $i, j \in \mathcal{V}$, a path linking i, j is given by a sequence of edges $p = \{(i_0 \equiv i, i_1), (i_1, i_2), \dots, (i_l, i_{l+1} \equiv j)\}$ where $i_1, \dots, i_l \in \mathcal{V}$, and each element in p is an edge in \mathcal{E} . The length of the path p is defined to be $l + 1$. If $p = \{(i, j)\}$, then the length of p is 1. The *geodesic distance* $d(i, j)$ between vertices i and j is defined as the length of the shortest path between i and j . It should be noted that for structurally symmetric matrices, i.e., $A_{ij} \neq 0$ implies $A_{ji} \neq 0$ for all indices i, j , the geodesic distance is symmetric, i.e., $d(i, j) = d(j, i)$. In particular, Hermitian matrices are structurally symmetric. If $d(i, j) > 1$, then $A_{ij} = 0$. If $d(i, j) = \infty$, then there is no path connecting i and j . More generally, for any positive integer k , if $d(i, j) > k$, then $(A^k)_{ij} = 0$, where A^k is the k -th power of the matrix A .

The precise statement of the spatial locality of the matrix function $f_{\sigma,\mu}(A)$ is given by the decay properties along the off-diagonal direction in Theorem 2.3. For a given column j , the magnitude of each element $f_{\sigma,\mu}(A)_{i,j}$ decays exponentially with respect to the geodesic distance $d(i, j)$.

Theorem 2.3. *Let A be a sparse and Hermitian matrix with all eigenvalues contained in the interval $(-1, 1)$. For any $\alpha > 0, \sigma > 0$, let*

$$(2.8) \quad \rho = (1 + \alpha\sigma)^{-1}, \quad K = \frac{2}{\rho\alpha\sigma} e^{\alpha^2},$$

then for all $d(i, j) \geq 1, i, j = 1, \dots, n$,

$$(2.9) \quad |f_{\sigma,\mu}(A)_{ij}| \leq K\rho^{d(i,j)},$$

where $d(i, j)$ is the geodesic distance between vertices i and j .

Proof. For any integer $k \geq 0$, there exists a polynomial $p_k \in \mathbb{P}_k$ such that

$$\|f_{\sigma,\mu}(A) - p_k(A)\|_2 = \|f_{\sigma,\mu} - p_k\|_\infty = E_k(f_{\sigma,\mu}) \leq K\rho^{k+1}.$$

The last inequality follows from Theorem 2.2. Now consider all edges (i, j) such that the geodesic distance $d(i, j) = k + 1$, and then $p_k(A)_{ij} = 0$. Therefore,

$$|f_{\sigma, \mu}(A)_{ij}| = |f_{\sigma, \mu}(A)_{ij} - p_k(A)_{ij}| \leq \|f_{\sigma, \mu}(A) - p_k(A)\|_2 \leq K\rho^{k+1} = K\rho^{d(i, j)}.$$

□

Remark 2.4. As suggested in (2.8), ρ, K only depend on σ but not on μ . Therefore the decay properties of the matrix function $f_{\sigma, \mu}(A)$ seem to be independent of the shift μ . This is because an upper bound for $M(\chi)$ is given in Theorem 2.2 that is valid for all μ . Numerical results in section 4 indicate that the preconstant of the exponential decay may have a strong dependency on μ , and such dependency can be specific to the matrix under study.

Remark 2.5. In Theorem 2.3 there is an arbitrary positive constant α . For any given $\alpha > 0$, the off-diagonal entries of $|f_{\sigma, \mu}(A)_{ij}|$ should decay exponentially with respect to the geodesic distance. By optimizing α together with the degree of the Chebyshev polynomial k , the actual decay rate can be slightly faster than exponential. Figure 2.1 gives an example of the magnitude of the first column $|f_{\sigma, \mu}(A)_{:,1}|$ where A is a discretized Laplacian operator in 1D with periodic boundary conditions, with $\sigma = 1.0, \mu = 2.0$ and $\sigma = 1.0, \mu = 10.0$, respectively. Although the discretized 1D Laplacian matrix is a banded matrix, all its eigenfunctions are plane waves which are fully delocalized in the global domain. Nonetheless the upper bound of the decay rate of the LSS operator is clearly exponential as shown in Figure 2.1.




Figure 2.1. Log-scale plot of the magnitude of the first column $|f_{\sigma, \mu}(A)_{:,1}|$. A is a discretized Laplacian operator in 1D with periodic boundary conditions, with $\sigma = 1.0$ and $\mu = 2.0, 10.0$, respectively.

Remark 2.6. In order to limit the numerical rank of $f_{\sigma, \mu}(A)$ in practice, it is desirable to use a small σ . With fixed α and assume $\alpha\sigma < 1$, we have

$$\rho^{d(i, j)} = (1 + \alpha\sigma)^{-d(i, j)} \leq e^{-\frac{1}{2}\alpha\sigma d(i, j)}.$$

Here σ reflects the spectral locality, and $d(i, j)$ reflects the spatial locality. This reveals the balance between the spectral and spatial locality, tuned by one parameter σ .

3. LOCALIZED SPECTRUM SLICING

3.1. Algorithm. Using the decay properties of the LSS operator $f_{\sigma,\mu}(A)$ in Theorem 2.3, a set of basis functions called the LSS basis set can be constructed in a divide-and-conquer fashion. Below we demonstrate that if the smearing parameter σ is large enough, then the localized spectrum slicing operator $f_{\sigma,\mu(A)}$ can be approximately computed using submatrices of A . The size of each submatrix is independent of the size of A . This is important for reducing the computational complexity and for parallel computation.

Proposition 3.1. *Let A, B be $n \times n$ Hermitian matrices. The graph induced by B is a spanning subgraph of the graph G induced by A , and the geodesic distance $d(i, j)$ is defined using the graph G . We assume for a given integer j, m ($1 \leq j, m \leq n$),*

$$A_{il} = B_{il}, \quad \forall i, l \quad \text{s.t.} \quad d(j, i) \leq m, \quad d(j, l) \leq m.$$

Then for any integer k ($1 \leq k \leq m$),

$$(A^k)_{il} = (B^k)_{il}, \quad \forall i, l \quad \text{s.t.} \quad d(j, i) \leq m - k + 1, \quad d(j, l) \leq m - k + 1.$$

Proof. The statement is apparently correct for $k = 1$. Assume the statement for $k - 1$ is proved, and we prove the statement is true for k . First,

$$(A^k)_{il} = \sum_p A_{ip}(A^{k-1})_{pl}, \quad (B^k)_{il} = \sum_p B_{ip}(B^{k-1})_{pl}.$$

In the summation above, $A_{ip}(A^{k-1})_{pl}$ is nonzero only if $A_{ip} \neq 0$. Similarly, $B_{ip}(B^{k-1})_{pl}$ is nonzero only if $B_{ip} \neq 0$. Since the graph induced by B is a subgraph of the graph induced by A , $B_{ip} \neq 0$ implies $A_{ip} \neq 0$, and therefore we only need to consider p such that $A_{ip} \neq 0$, i.e., $d(i, p) = 1$. Consider i such that $d(j, i) \leq m - k + 1$, then p satisfies

$$d(j, p) \leq d(j, i) + d(i, p) = m - k + 2 = m - (k - 1) + 1.$$

Also, for any l such that $d(j, l) \leq m - k + 1 < m - (k - 1) + 1$, by the assumption that the statement for $k - 1$ is proved, we have $(A^{k-1})_{pl} = (B^{k-1})_{pl}$. Together with $d(j, i) \leq m, d(j, p) \leq m$, we have $A_{ip} = B_{ip}$. Therefore,

$$(A^k)_{il} = \sum_p A_{ip}(A^{k-1})_{pl} = \sum_p B_{ip}(B^{k-1})_{pl} = (B^k)_{il}$$

is valid for all i, l such that $d(j, i) \leq m - k + 1, d(j, l) \leq m - k + 1$. □

Using Proposition 3.1, Theorem 3.2 shows that the j -th column of $f_{\sigma,\mu}(A)$ can be accurately computed from $f_{\sigma,\mu}(B)$, as long as the columns of A and B are sufficiently close in the vicinity of j in the sense of the geodesic distance.

Theorem 3.2. *Let A, B be $n \times n$ Hermitian matrices with eigenvalues in $(-1, 1)$. For a given j and an even integer m ($1 \leq j, m \leq n$),*

$$A_{il} = B_{il}, \quad \forall i, l \quad \text{s.t.} \quad d(j, i) \leq m, \quad d(j, l) \leq m.$$

Then

$$|f_{\sigma,\mu}(A)_{ij} - f_{\sigma,\mu}(B)_{ij}| \leq 2K\rho^{\frac{m}{2}+1},$$

for all i such that $d(j, i) \leq m/2 + 1$, where the constants K, ρ are given in (2.8).

Proof. For any i, j and $k \geq 0$ we have

$$\begin{aligned} |f_{\sigma,\mu}(A)_{ij} - f_{\sigma,\mu}(B)_{ij}| &\leq |f_{\sigma,\mu}(A)_{ij} - p_k(A)_{ij}| + |p_k(A)_{ij} - p_k(B)_{ij}| \\ &\quad + |f_{\sigma,\mu}(B)_{ij} - p_k(B)_{ij}|. \end{aligned}$$

Take $k = \frac{m}{2}$. For any i such $d(j, i) \leq m - k + 1 = \frac{m}{2} + 1$, by Proposition 3.1, $p_k(A)_{ij} = p_k(B)_{ij}$. Also from Theorem 2.2, we have

$$\begin{aligned} |f_{\sigma,\mu}(A)_{ij} - p_k(A)_{ij}| &\leq \|f_{\sigma,\mu}(A) - p_k(A)\|_2 \leq K\rho^{\frac{m}{2}+1}, \\ |f_{\sigma,\mu}(B)_{ij} - p_k(B)_{ij}| &\leq \|f_{\sigma,\mu}(B) - p_k(B)\|_2 \leq K\rho^{\frac{m}{2}+1}, \end{aligned}$$

and hence the result. \square

Theorem 3.2 shows that in order to compute any column j of the matrix $f_{\sigma,\mu}(A)$ up to certain accuracy, it is only necessary to have a matrix that shares the same columns as A up to a certain distance away from j . Together with the decay property of each column of $f_{\sigma,\mu}(A)$, this allows the j -th column of $f_{\sigma,\mu}(A)$ to be constructed in a *divide-and-conquer* manner. For instance, for a given integer m we can define

$$(3.1) \quad B_{il} = \begin{cases} A_{il}, & \forall i, l \text{ s.t. } d(j, i) \leq m, d(j, l) \leq m, \\ 0, & \text{otherwise,} \end{cases}$$

which is simply a submatrix of A . As a submatrix, $\|B\|_2 \leq \|A\|_2$ and the assumption of the spectral radius in Theorem 3.2 is satisfied.

In practice it would be very time consuming to construct an approximate matrix for each column of j , since the rank of the LSS operator $f_{\sigma,\mu}(A)$ is often much smaller compared to n . For structured matrices such as matrices obtained from finite difference or finite element discretization of PDE operators, it is often possible to partition the domain into well structured disjoint columns sets, and apply the truncated LSS operator to each column set. The cost for generating such a partition can be very small if the structure of the matrix is known *a priori*. For the discussion below, we assume that the partition $\mathcal{V} = \{1, \dots, n\}$ into M simply connected disjoint sets $\{E_\kappa\}_{\kappa=1}^M$ is given, i.e.,

$$\mathcal{V} = \bigcup_{\kappa=1}^M E_\kappa \quad \text{and} \quad E_\kappa \cap E_{\kappa'} = \emptyset, \quad \kappa \neq \kappa'.$$

For general sparse matrices, such partitions may not be readily available. We discuss the choice of domain partitioning strategy in section 3.4.

For each E_κ and an integer m , we define an associated set

$$(3.2) \quad Q_\kappa = \{i | d(i, j) \leq m, \forall j \in E_\kappa\}.$$

Theorem 3.2 implies that the submatrix $(f_{\sigma,\mu}(A))_{:,E_\kappa}$ can be constructed by a submatrix of A defined as

$$(3.3) \quad (A_\kappa)_{ij} = \begin{cases} A_{ij}, & i, j \in Q_\kappa, \\ 0, & \text{otherwise.} \end{cases}$$

In the following discussion, we refer to E_κ as an *element*, and to Q_κ as an *extended element* associated with E_κ . It should be noted that the zero entries of A_κ outside the index set Q_κ do not need to be explicitly stored.

Remark 3.3. The choice in (3.1) takes a submatrix of A to compute the localized spectrum slicing operator. From the point of view of partial differential operators, this is similar to imposing the zero Dirichlet boundary condition on some local domains. However, the zero Dirichlet boundary condition may not be the optimal choice. Generalization of other types of boundary conditions can be considered as well to construct the local matrices A_κ .

Since A is Hermitian and sparse, and so is A_κ . The latter has the eigendecomposition

$$(3.4) \quad A_\kappa X_\kappa = X_\kappa D_\kappa.$$

Here D_κ is a diagonal matrix. Note that A_κ only takes nonzero values on the extended element Q_κ . The entries of each column of X_κ outside the index set Q_κ can be set to zero, and such zero entries do not need to be explicitly stored. This is equivalent to solving an eigenvalue problem of size $|Q_\kappa| \times |Q_\kappa|$. Define

$$(3.5) \quad f_{\sigma,\mu}(A_\kappa) \equiv X_\kappa f_{\sigma,\mu}(D_\kappa) X_\kappa^*.$$

Using Theorem 3.2, $f_{\sigma,\mu}(A)_{Q_\kappa, E_\kappa}$ can be approximated by $f_{\sigma,\mu}(A_\kappa)$, in the sense that

$$|f_{\sigma,\mu}(A)_{ij} - f_{\sigma,\mu}(A_\kappa)_{ij}| \leq 2K\rho^{\frac{m}{2}+1}, \quad \forall i \in Q_\kappa, \quad j \in E_\kappa.$$

Since $f_{\sigma,\mu}$ is spectrally localized, in practice not all eigenvalues and eigenvectors of A_κ as in (3.4) are needed. Instead only a *partial eigendecomposition* is needed to compute all eigenvalues of A_κ in the interval $(\mu - c\sigma, \mu + c\sigma)$. Due to the fast decay properties of Gaussian functions, in practice c can be chosen to be $2 \sim 4$ to be sufficiently accurate. We denote by s_κ the column dimension of X_κ in the partial eigendecomposition of A_κ .

The factorized representation in (3.5) also allows the computation of a set of vectors approximately spanning the column space of $f_{\sigma,\mu}(A_\kappa)$, through a local singular value decomposition (SVD) procedure, i.e.,

$$(3.6) \quad \|f_{\sigma,\mu}(D_\kappa)((X_\kappa)_{Q_\kappa, :})^* - \tilde{U}_\kappa \tilde{S}_\kappa \tilde{V}_\kappa^*\|_2 \leq \tilde{\tau}.$$

Here $\tilde{\tau}$ is SVD truncation criterion. The size of the matrix for the SVD decomposition is $s_\kappa \times |Q_\kappa|$. In practice $\tilde{\tau}$ may also be chosen using a relative criterion as $\tilde{\tau} = \tau(\tilde{S}_\kappa)_{1,1}$ is used in our numerical experiment, where we assume $(\tilde{S}_\kappa)_{1,1}$ is the largest singular value in (3.6). In practice this can be performed by only keeping the singular values in the diagonal matrix \tilde{S}_κ that are larger than $\tilde{\tau}$. Then we can define

$$(3.7) \quad U_\kappa = X_\kappa \tilde{U}_\kappa, \quad V_\kappa = \tilde{S}_\kappa \tilde{V}_\kappa^*.$$

We combine all U_κ together

$$(3.8) \quad U \equiv [U_1, \dots, U_M],$$

and U is the LSS basis set that is both spectrally localized and spatially localized. We denote by n_b the total number of columns of U , which is also referred to as the size of the LSS basis set. Using the LSS basis set, an approximation to the LSS operator is defined as

$$(3.9) \quad \tilde{f}_{ij} = \begin{cases} (U_\kappa)_{i,:} (V_\kappa)_{:,j}, & i \in Q_\kappa, j \in E_\kappa, \quad \text{for some } \kappa, \\ 0, & \text{otherwise.} \end{cases}$$

Then \tilde{f} is an $n \times n$ sparse matrix, and the error in the max norm for approximating the LSS operator $f_{\sigma,\mu}(A)$ is given in Theorem 3.4.

Theorem 3.4. *Let A be an $n \times n$ Hermitian matrix with eigenvalues in $(-1, 1)$, and the induced graph is partitioned into M elements $\{E_\kappa\}$. For each element E_κ , there is an extended element Q_κ given in (3.2), a submatrix A_κ given in (3.3), and matrices U_κ, V_κ satisfying (3.6) and (3.7). Let \tilde{f} be an $n \times n$ matrix defined in (3.9), then*

$$(3.10) \quad \|f_{\sigma,\mu}(A) - \tilde{f}\|_{\max} \leq 2K\rho^{\frac{m}{2}+1} + \tilde{\tau}.$$

Proof. For each element κ , from (3.6) we have

$$(3.11) \quad \begin{aligned} \max_{i \in Q_\kappa, j \in E_\kappa} |f_{\sigma,\mu}(A_\kappa)_{ij} - \tilde{f}_{ij}| &= \max_{i \in Q_\kappa, j \in E_\kappa} |f_{\sigma,\mu}(A_\kappa)_{ij} - (U_\kappa)_{i,:} (V_\kappa)_{:,j}| \\ &\leq \|f_{\sigma,\mu}(A_\kappa) - U_\kappa V_\kappa\|_2 \leq \|X_\kappa\|_2 \tilde{\tau} = \tilde{\tau}. \end{aligned}$$

Using Theorem 3.2 and the definition of the extended element (3.2),

$$(3.12) \quad \max_{i \in Q_\kappa, j \in E_\kappa} |f_{\sigma,\mu}(A)_{ij} - f_{\sigma,\mu}(A_\kappa)_{ij}| \leq 2K\rho^{\frac{m}{2}+1}.$$

For vertices $i \notin Q_\kappa, j \in E_\kappa$, $\tilde{f}_{ij} = 0$. Then from Theorem 2.3 and using $\rho < 1$, we get

$$(3.13) \quad \max_{i \notin Q_\kappa, j \in E_\kappa} |f_{\sigma,\mu}(A)_{ij} - \tilde{f}_{ij}| = |f_{\sigma,\mu}(A)_{ij}| \leq K\rho^{m+1} \leq K\rho^{m/2+1}.$$

Combining equations (3.11), (3.12), (3.13), we have

$$\begin{aligned} \|f_{\sigma,\mu}(A) - \tilde{f}\|_{\max} &= \max_{1 \leq i, j \leq n} |f_{\sigma,\mu}(A)_{ij} - \tilde{f}_{ij}| \\ &= \max_{\kappa} \left\{ \max \left\{ \max_{i \in Q_\kappa, j \in E_\kappa} |f_{\sigma,\mu}(A)_{ij} - \tilde{f}_{ij}|, \max_{i \notin Q_\kappa, j \in E_\kappa} |f_{\sigma,\mu}(A)_{ij} - \tilde{f}_{ij}| \right\} \right\} \\ &= \max_{\kappa} \left\{ \max \{ 2K\rho^{\frac{m}{2}+1} + \tilde{\tau}, K\rho^{m/2+1} \} \right\} = 2K\rho^{\frac{m}{2}+1} + \tilde{\tau}. \quad \square \end{aligned}$$

Remark 3.5. Theorem 3.4 indicates that in order to accurately approximate the LSS operator, the SVD truncation criterion $\tilde{\tau}$ must be small enough. However, the result in Theorem 3.4 is only an upper bound of the error. Numerical results in section 4 indicate that accurate eigenvalues may be obtained even when $\tilde{\tau}$ is relatively large.

Finally, we summarize the algorithm for finding the divide-and-conquer method for constructing the LSS basis set in Algorithm 1.

3.2. Complexity. In order to simplify the analysis of the complexity of Algorithm 1 for finding the LSS basis set, we make the assumption that the set of n vertices is equally divided into M elements, so that $|E_\kappa| = \frac{n}{M} \equiv |E|$. As n increases we assume $|E|$ can be kept as a constant, i.e., the number of elements M increases proportionally with respect to n . $|Q_\kappa| = \frac{c_Q n}{M} \equiv c_Q |E|$, where c_Q is a small number denoting the ratio between the size of the extended element and the size of the element. For instance, for the discretized 1D and 2D Laplacian operators in the numerical examples, c_Q is set to be 3 and 9, respectively.

Denote by s_κ the column dimension of X_κ in the partial eigendecomposition of A_κ , and by t_κ the column dimension of U_κ with $t_\kappa \leq s_\kappa$. For simplicity we assume $\{s_\kappa\}, \{t_\kappa\}$ are uniform, i.e., $s_\kappa = s, t_\kappa = t, \kappa = 1, \dots, M$. If A_κ is treated as a dense

Algorithm 1: Localized spectrum slicing basis set.

(1) Sparse Hermitian matrix A , center μ , width σ , SVD truncation tolerance $\tilde{\tau}$.

Input: (2) Number of elements M , partition of elements $\{E_\kappa\}_{\kappa=1}^M$ and extended elements $\{Q_\kappa\}_{\kappa=1}^M$.

Output: LSS basis set $\{U_\kappa\}_{\kappa=1}^M$.

for $\kappa = 1, \dots, M$ **do**

- Compute the (partial) eigendecomposition according to (3.4);
- Compute the local SVD decomposition according to (3.6) and only keep singular vectors with singular values larger than $\tilde{\tau}$;
- Compute U_κ with matrix multiplication according to (3.7);

end

matrix for the computation of the local eigendecomposition of A_κ , then the cost is $c_{\text{Eig,d}}|Q_\kappa|^3$. The cost of the SVD decomposition is $c_{\text{SVD}}|E_\kappa|s_\kappa^2$. The cost of matrix multiplication to obtain U_κ is $c_{\text{MM}}|Q_\kappa|s_\kappa t_\kappa$. So the total cost for finding the LSS basis set is proportional to

$$(3.14) \quad \sum_{\kappa=1}^M c_{\text{Eig,d}}|Q_\kappa|^3 + c_{\text{SVD}}|E_\kappa|s_\kappa^2 + c_{\text{MM}}|Q_\kappa|s_\kappa t_\kappa = n (c_{\text{Eig,d}}c_Q^3|E|^2 + c_{\text{SVD}}s^2 + c_{\text{MM}}c_Qst).$$

If we assume that as n increases, the spectral radius of A does not increase, then all constants in the parenthesis in the right-hand side of (3.14) are independent of n , and the overall computational complexity for finding the LSS basis set is $\mathcal{O}(n)$.

In practice the constant for the finding the local eigendecomposition can be large due to the term $|E|^2$ in (3.14). Since A_κ is still a sparse matrix on Q_κ , iterative methods can be used to reduce the computational cost to $c_{\text{Eig,i}}|Q_\kappa|s_\kappa^2$. This modifies the overall complexity to be

$$n (c_{\text{Eig,i}}c_Qs^2 + c_{\text{SVD}}s^2 + c_{\text{MM}}c_Qst).$$

However, it should be noted that the preconstant $c_{\text{Eig,i}}$ might be larger than $c_{\text{Eig,d}}$. Whether a direct or an iterative method should be used to solve the local eigenvalue problem may depend on a number of practical factors such as the size of the local problem, and the availability of efficient preconditioner on the local domain, etc.

3.3. Compute interior eigenvalues. Using the LSS basis set in (3.8), one may compute the interior eigenvalues near μ together with its associated eigenvectors. This can be done by using the projected matrices A_U, B_U according to (1.3). Due to the spatial sparsity of U , A_U, B_U are also sparse matrices, and can be assembled efficiently with local computation. First, the matrix multiplication $Z = AU$ can be performed locally. This is because each column of U_κ is localized in Q_κ , then

$$(3.15) \quad Z_\kappa = AU_\kappa \approx A_\kappa U_\kappa.$$

Second, denote by

$$(A_U)_{\kappa',\kappa} = U_{\kappa'}^* Z_\kappa, \quad (B_U)_{\kappa',\kappa} = U_{\kappa'}^* U_\kappa,$$

then for each κ it is sufficient to loop over elements $E_{\kappa'}$ so that $Q_{\kappa'} \cap Q_\kappa$ is nonempty. The details for constructing the projected matrices are given in Algorithm 2.

Algorithm 2: Assembly of the projected matrices.

(1) Sparse Hermitian matrix A .

Input: (2) Number of elements M , partition of elements $\{E_\kappa\}_{\kappa=1}^M$, extended elements $\{Q_\kappa\}_{\kappa=1}^M$, submatrices $\{A_\kappa\}_{\kappa=1}^M$, LSS basis set $\{U_\kappa\}_{\kappa=1}^M$ with total number of basis functions n_b .

Output: Projected matrices A_U, B_U .

Let A_U, B_U be zero matrices of size $n_b \times n_b$.

for $\kappa = 1, \dots, M$ **do**

Compute $Z_\kappa \leftarrow A_\kappa U_\kappa$;

for κ' so that $Q_{\kappa'} \cap Q_\kappa \neq \emptyset$ **do**

Compute $(A_U)_{\kappa', \kappa} \leftarrow U_{\kappa'}^* Z_\kappa$;

Compute $(B_U)_{\kappa', \kappa} \leftarrow U_{\kappa'}^* U_\kappa$;

end

end

Symmetrize $A_U \leftarrow \frac{1}{2}(A_U + A_U^*)$, $B_U \leftarrow \frac{1}{2}(B_U + B_U^*)$.

After A_U, B_U are assembled, the eigenvalues and corresponding eigenvectors near μ can be solved in various ways. When the size of the LSS basis set n_b is small, one can treat A_U, B_U as dense matrices and solve the generalized eigenvalue problem

$$(3.16) \quad A_U C = B_U C \Theta,$$

and only keep the Ritz values $\Theta = \text{diag}[\theta_1, \dots, \theta_{n_b}]$ and corresponding Ritz vectors C near μ . Each column of the Ritz vector C_j can be partitioned according to the element partition $\{E_\kappa\}$ as

$$C_j = [C_{1,j}, \dots, C_{M,j}]^T.$$

Then an approximate eigenvector for A can be computed as

$$(3.17) \quad \tilde{X}_j = U C_j = \sum_{\kappa} U_\kappa C_{\kappa,j}.$$

We remark that in the computation of interior eigenvalues, spurious eigenvalues may appear. A spurious eigenvalue is a Ritz value θ_j near the vicinity of μ as obtained from (3.16), but the corresponding vector \tilde{X}_j as given in (3.17) is not an approximate eigenvector. The appearance of a spurious eigenvalue is also referred to as spectral pollution [14, 16], and could be identified by computing the residual

$$(3.18) \quad R_j = A \tilde{X}_j - \tilde{X}_j \theta_j.$$

A Ritz value θ_j corresponding to large residual norm $\|R_j\|_2$ should be removed. Note that the residual can also be computed with local computation

$$(3.19) \quad R_j = \sum_{\kappa} (Z_\kappa C_{\kappa,j} - U_\kappa C_{\kappa,j} \theta_j),$$

where Z_κ is given in (3.15). Our numerical experience indicates that the use of the residual is an effective way for identifying spurious eigenvalues when the LSS basis set is accurate enough for approximating the subspace spanned by the eigenvectors to be computed. In such a case the norm of the residual for most Ritz values is small and the norm of the residual for the spurious eigenvalue stands out. When the basis set cannot accurately capture all the eigenvalues in the prescribed interval

especially for those clustered near the boundary of the interval, it becomes more difficult to identify all the spurious eigenvalues.

3.4. Domain partitioning for general sparse matrices. For a general sparse matrix A , we discuss here the strategy to partition the associated undirected graph $G = (\mathcal{V}, \mathcal{E})$ into M elements $\{E_\kappa\}_{\kappa=1}^M$. Intuitively we would like to choose a partition that keeps all E_κ to have similar sizes, while minimizing the number of edges that is being cut by the partition, i.e., $\sum_{\kappa, \kappa'=1}^M \sum_{i \in E_\kappa, j \in E_{\kappa'}} w_{ij}$. Here $w_{ij} = 1$ if $A_{ij} \neq 0$ and 0 otherwise. This is called a minimal M -cut problem. It is known that the minimal M -cut problem is NP-hard. Various heuristic methods have been developed. Here we use the nested dissection approach [11] as implemented in the METIS [15] package. The nested dissection approach can find an approximate minimal 2-cut of the graph, and then recursively partitions each part of the graph, with iterative adjustment of the size of E_κ . For each κ we define a neighbor list N_κ , which consists of κ itself, as well as other element indices κ' such that there exists at least one pair of indices $i \in E_\kappa, j \in E_{\kappa'}$ and $A_{ij} \neq 0$. Then the extended element Q_κ is defined as the collection of all indices in $E_{\kappa'}$ such that $\kappa \in N_{\kappa'}$. Algorithm 3 gives a pseudo-code for generating the elements $\{E_\kappa\}$, the neighbor lists $\{N_\kappa\}$, and the extended elements $\{Q_\kappa\}$. In terms of implementation, the partition of the graph is given by a *graph partition map* ξ such that $E_\kappa = \{i \in \mathcal{V} | \xi(i) = \kappa\}$, and ξ can be directly returned from a graph partitioning package such as METIS.

Algorithm 3: Generating the set of elements $\{E_\kappa\}$ and neighboring elements for a general sparse matrix.

Input: Sparse Hermitian matrix A of size $n \times n$. Number of elements M .

Output: $E_\kappa, N_\kappa, Q_\kappa, \kappa = 1, \dots, M$.

$\xi = \text{GraphPartition}(A)$.

$E_\kappa = \{i \in \mathcal{V} | \xi(i) = \kappa\}, \kappa = 1, \dots, M$.

$N_\kappa = \{\kappa\} \cup \{\kappa' | \exists i \in E_{\kappa'}, j \in E_\kappa, A_{ij} \neq 0\}, \kappa = 1, \dots, M$.

$Q_\kappa = \{i \in \mathcal{V} | i \in E_{\kappa'}, \kappa' \in N_\kappa\}, \kappa = 1, \dots, M$.

4. NUMERICAL RESULTS

In this section we demonstrate the accuracy and efficiency of the divide-and-conquer procedure for computing the LSS operator and the LSS basis set, and for computing interior eigenvalues. All the computation is performed on a single computational thread of an Intel i7 CPU processor with 64 gigabytes (GB) of memory using MATLAB. The matrix A is obtained from a discretized second order partial differential operator $-\Delta + V$ in one dimension (1D) and in two dimension (2D) with periodic boundary conditions, and a general matrix from the University of Florida matrix collection.

4.1. One-dimensional case. In the 1D case, the global domain is $\Omega = [0, L]$. The Laplacian operator is discretized using a 3-point finite difference stencil. The domain is uniformly discretized into $n = c_n M$ grid points so that $x_i = (i - 1)h$, with the grid spacing $h \equiv L/n = 0.1$. All the n grid points (vertices) are uniformly and contiguously partitioned into M elements $\{E_\kappa\}_{\kappa=1}^M$. For simplicity let Q_κ be

the union of E_κ and its two neighbors taking into account the periodic boundary condition, i.e.,

$$Q_\kappa = \begin{cases} E_M \cup E_1 \cup E_2, & \kappa = 1, \\ \bigcup_{\kappa'=\kappa-1}^{\kappa+1} E_{\kappa'}, & \kappa = 2, \dots, M-1, \\ E_{M-1} \cup E_M \cup E_1, & \kappa = M. \end{cases}$$

The potential $V(x)$ is given by the sum of n_w exponential functions as

$$(4.1) \quad V(x) = - \sum_{i=1}^{n_w} a_i e^{-\frac{\text{dist}(x, R_i)}{\delta_i}}.$$

Here $\{R_i\}$ are a set of equally spaced points. The distance between two points x and x' is defined to be the minimal distance between x and all the periodic images of x' , i.e.,

$$\text{dist}(x, x') = \min_{\tilde{x}'=x'+kL, k \in \mathbb{Z}} |x - \tilde{x}'|.$$

In order to study the performance of the algorithm for systems of increasing sizes, we set $L = 20n_w$ so that the length of the computational domain is proportional to the number of potential wells n_w . To show that we do not take advantage of the periodicity of the potential, we introduce some randomness in each exponential function. We choose $a_i \sim \mathcal{N}(5.0, 1.0)$, which is a Gaussian random variable with a mean value 5.0 and a standard deviation 1.0. Similarly, the width of the exponential function $\delta_i \sim \mathcal{N}(2.0, 0.2)$. One realization of the potential with $n_w = 8$ is given in Figure 4.1(a), with the partition of elements indicated by black dashed lines. For the choice of parameter $\mu = 2.0$ and $\sigma = 1.0$, Figure 4.1(b) shows the function $f_{\sigma, \mu}(\lambda)$ evaluated on the eigenvalues of A plotted in log-scale in the interval $(-5, 10)$, and the LSS operator $f_{\sigma, \mu}(A)$ is spectrally localized. Figure 4.1(c) demonstrates the histogram of the eigenvalues (unnormalized spectral density) for all eigenvalues of A .




Figure 4.1. (a) One realization of the 1D potential with $n_w = 8$. The domain is partitioned into 8 equally sized elements separated by black dashed lines. (b) The function $f_{\sigma, \mu}(A)$ with $\sigma = 1.0$ and $\mu = 2.0$ viewed spectrally in the interval $(-5, 15)$ plotted in the log scale. The spectral radius of A is 199.89. (c) The histogram of the eigenvalues of A .

Figure 4.2(a)-(c) demonstrates the behavior of the exact LSS operator $f_{\sigma,\mu}(A)$ with $\sigma = 1.0$ and increasing value of μ . In Figure 4.2, $[f_{\sigma,\mu}(A)](x, y)$ should be interpreted using its discretized matrix element $[f_{\sigma,\mu}(A)]_{ij}$ for $x = (i - 1)h, y = (j - 1)h$. We find that as μ increases, the off-diagonal elements of f decay rapidly and remain to be well approximated by a banded (and therefore sparse) matrix with increasing bandwidth. Figure 4.2(d)-(f) demonstrates the quality of the divide-and-conquer approximation \tilde{f} to the LSS operator. Here we first demonstrate the accuracy of \tilde{f} without the truncation using SVD decomposition (i.e., the SVD truncation criterion $\tilde{\tau} = 0$ as in (3.6)). When $\mu = -2.0$, the approximation is nearly exact. When μ increases to 20.0 the relative error is around 10% since the support size of each column of f already extends beyond each extended element Q_{κ} .

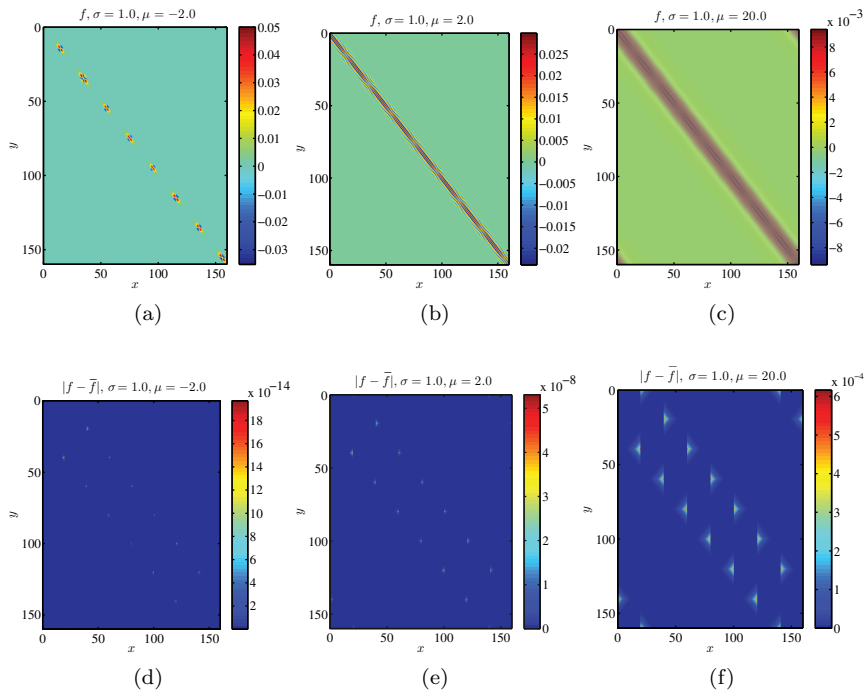


Figure 4.2. The LSS operator $f_{\sigma,\mu}(A)$ with $\sigma = 1.0$ for (a) $\mu = -2.0$; (b) $\mu = 2.0$; (c) $\mu = 20.0$. The max error between the LSS operator and its divide-and-conquer approximation $\tilde{f}_{\sigma,\mu}(A)$ for (d) $\mu = -2.0$; (e) $\mu = 2.0$; (f) $\mu = 20.0$.

A more complete picture of the μ -dependence for approximating the LSS operator is given in Figure 4.3. Figure 4.3(a) shows the max norm error of the divide-and-conquer approximation to the LSS operator for μ traversing the entire spectrum of A from -3.0 to 200.0 . The error increases rapidly as μ initially increases, achieves its maximum at $\mu = 100$ and then starts to decrease. Figure 4.3(b) shows the same picture but zooms into the interval near $\mu = 0$. As μ increases above 10.0, the vectors spanning columns of $f_{\sigma,\mu}(A)$ are approximately linear combination of

high frequency Fourier modes, and Figure 4.3(a) shows that the Fourier modes are increasingly more difficult to localize as the frequency increases. Figure 4.3(c)-(d) shows similar behavior for $\sigma = 2.0$. The profile of the error with respect to μ closely resembles a Gaussian function. Compared to the case with $\sigma = 1.0$ the error significantly reduces for all μ , indicating the balance between spatial locality and spectral locality with varying σ .

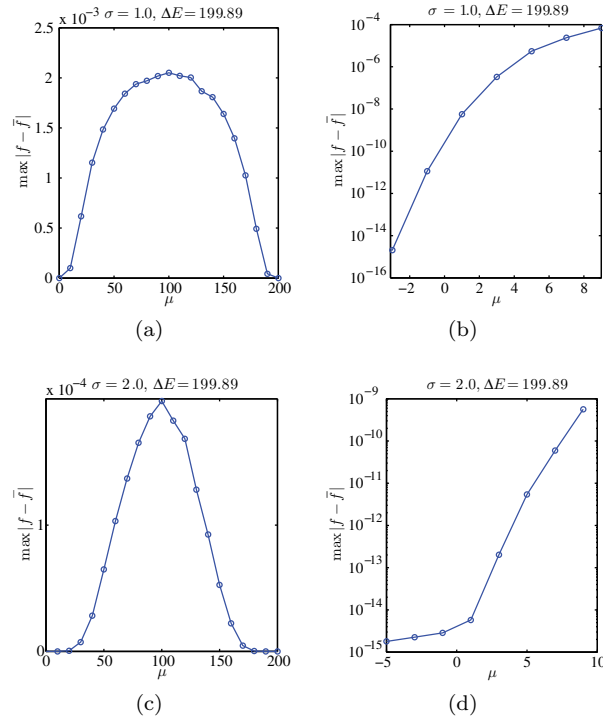


Figure 4.3. Max norm error of the LSS operator traversing the entire spectrum of A for (a) $\sigma = 1.0$; (b) $\sigma = 1.0$, zoom-in; and (c) $\sigma = 2.0$; (d) $\sigma = 2.0$, zoom-in.

Figure 4.4(a) demonstrates the max norm error of the LSS operator for $\mu = 2.0$ with increasing value of σ . When σ is less than 0.25 the LSS operator is very localized spectrally, but the matrix is almost dense. Therefore the divide-and-conquer approximation leads to large error. As σ increases above 0.25, the max norm error decreases exponentially with the increase of σ . We observe that the choice of σ is crucial: by varying σ from 0.5 to 1.5, the error is reduced by over 6 orders of magnitude from 10^{-4} to below 10^{-10} .

Next we study the effect of grid refinement by varying the grid size from $h = 0.20$ to $h = 0.033$. For 3-point finite difference stencil the spectral radius of A , denoted by ΔE is proportional to $1/h^2$, and in practice ΔE increases from 50 to 1800. We note that Theorem 3.2 indicates that the error should be determined by the ratio $\sigma/\Delta E$, and therefore the size of the extended element as characterized by the geodesic distance m should increase proportionally to ΔE to preserve accuracy. Here instead we fix the number of elements to be 8 as the grid refines. Therefore

$m \sim 1/h \sim \sqrt{\Delta E}$, and we should expect that the error increases as the grid refines. Figure 4.4(b) shows that max norm error of the LSS operator for $\mu = 2.0, \sigma = 1.0$, with increasing ΔE . As the ratio $\sigma/\Delta E$ decreases over one order of magnitude, the max norm error does not increase, but rather decreases by more than a factor of 2. We note that this numerical result does not contradict the theoretical prediction, since Theorem 2.2 only provides an *upper bound* of the decay rate, and the *actual* decay rate can be faster. This behavior has been recently verified analytically for discretized Schrödinger type operators [20]. Note that as the grid refines, the change towards the high end of the spectrum is often larger than the change at the low end of the spectrum. Figure 4.4 indicates that the accuracy of the LSS operator is relatively insensitive to the change in the high end of the spectrum, and it may be possible to construct the LSS operator with improved discretization scheme, without sacrificing too much in terms of the spatial locality.




Figure 4.4. Max norm error of the LSS operator for (a) $\mu = 2.0, \Delta E = 199.89$ and increasing value of σ ; (b) $\mu = 2.0, \sigma = 1.0$ and increasing value of ΔE .

So far the numerical results are obtained for the divide-and-conquer approximation to the LSS operator with $\tilde{\tau} = 0$. Next we apply the SVD truncation to obtain the LSS basis set $\{U_\kappa\}_{\kappa=1}^M$ for varying SVD relative truncation criterion. In our numerical experiments, we use τ as the *relative* SVD truncation criterion with respect to the largest singular value of \tilde{S}_κ . Figure 4.5 shows the error of the approximation to the LSS operator with τ being 0.001, 0.01, 0.1, respectively. As indicated in (3.6), the max norm error of the approximation of the LSS operator is approximately proportional to τ , as τ becomes dominant in (3.10).

The LSS basis set comes from the SVD decomposition of \tilde{f} on each element. Figure 4.6(a) shows the 1-st LSS basis function on two elements $\kappa = 2$ and $\kappa = 6$, respectively, and Figure 4.6(b) shows the 5-th LSS basis function on the same two elements for $\mu = 2.0, \sigma = 1.0$. It is clear that each LSS basis function is well localized in each extended element Q_κ and its center is in E_κ .

Figure 4.5 seems to suggest that in order to accurately compute the interior eigenvalues, a very tight SVD criterion τ is needed. However, we note that many of the LSS basis functions associated with the small singular values actually corresponds to the tail of the Gaussian function in (1.2) which are away from μ . Therefore, in order to compute the interior eigenvalues near μ accurately, it is possible to use a




Figure 4.5. Error of the divide-and-conquer approximation to the LSS operator with $\sigma = 1.0, \mu = 2.0$ and different SVD relative truncation criterion (a) $\tau = 10^{-3}$; (b) $\tau = 10^{-2}$; (c) $\tau = 10^{-1}$.

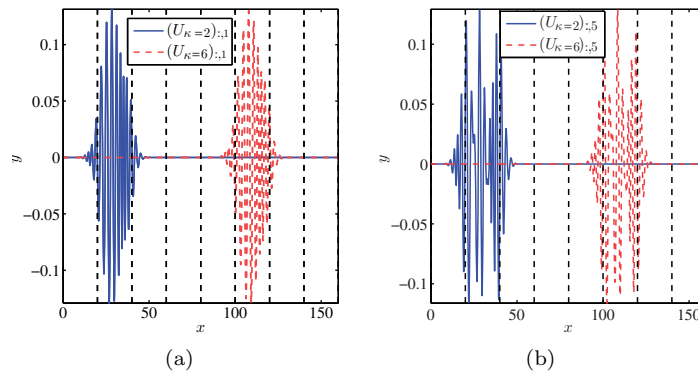


Figure 4.6. Example of the LSS basis function on two elements $\kappa = 2$ and $\kappa = 6$ for (a) the 1-st LSS basis function and (b) the 5-th LSS basis function.

much larger value of τ . Figure 4.7(a) shows the difference between the 24 eigenvalues of A within the interval $(\mu - 0.5\sigma, \mu + 0.5\sigma)$ and the corresponding Ritz values of A with $\tau = 0.1$. The computed Ritz values are highly accurate and the maximum error is under 5×10^{-6} even though a large SVD truncation criterion τ is used. Section 3.3 discusses the identification of spurious eigenvalues using the residual for each computed Ritz value. Indeed, within the interval $(\mu - 0.5\sigma, \mu + 0.5\sigma)$ we find 25 Ritz values, and the 1 additional Ritz value should be a spurious eigenvalue. Figure 4.7(b) shows $\|R_j\|_2$ for each Ritz value, and we identify that the 11-th Ritz value has a much larger residual than the rest and should be removed. After removing this spurious eigenvalue, the remaining Ritz values become an accurate approximation to the eigenvalues as indicated in Figure 4.7(a).

While the accuracy of the divide-and-conquer approximation to the LSS operator improves as the SVD truncation criterion τ decreases, using a very small value of τ may result in ill-conditioned projection matrices A_U and B_U , i.e., some of the LSS basis functions can be approximately represented as the linear combination




Figure 4.7. (a) Difference between the 24 eigenvalues and corresponding Ritz values within the interval $(\mu - 0.5\sigma, \mu + 0.5\sigma)$ with $\sigma = 1.0, \mu = 2.0$. (b) The 2-norm of the residual for each of the 25 Ritz values. The 11-th Ritz value has a large residual norm and is a spurious eigenvalue.




Figure 4.8. (a) The condition number of A_U and B_U ; and (b) the number of basis functions, with varying SVD relative truncation criterion τ . Here $\mu = 2.0, \sigma = 1.0$.

of other LSS basis functions. Figure 4.8(a) shows the condition number of A_U, B_U with respect to τ . The condition numbers are below 10^4 when $\tau \geq 0.1$, and increase very rapidly to 10^{13} for $\tau = 10^{-3}$. In the latter case, numerical results obtained from the generalized eigenvalue solver cannot be trusted. Decreasing τ also leads to increasing the size of the LSS basis set. As τ decreases from 10^{-1} to 10^{-3} , the number of LSS basis functions increases from 87 to 173. The accuracy of the LSS basis set for different values of τ is given in Table 4.1. When τ is too small, the number of computed Ritz values is less than 24 due to the very large condition number of the generalized eigenvalue problem, and the difference between the eigenvalues and the Ritz values is not a meaningful quantity to report and is reported as N/A. The error of the Ritz values reaches its minimum near $\tau = 0.032$ at only 7.59×10^{-8} , and then starts to increase as τ increases. We observe that even if $\tau = 0.316$, the absolute (and relative) error of the Ritz values is still within

0.2%. For this case the dimension of the projected generalized eigenvalue problem is 62, which is much smaller compared to the dimension of A which is 1600.

Table 4.1. The number of computed Ritz values in the interval $(\mu - 0.5\sigma, \mu + 0.5\sigma)$ with $\sigma = 1.0, \mu = 2.0$ (spurious eigenvalues removed). If the number of Ritz values match the number of eigenvalues in the interval (24), then the third column gives the maximum difference between the eigenvalues and the Ritz values. Otherwise the third column gives N/A.

τ	# Ritz values	$\max_j \lambda_j - \theta_j $
0.001	1	N/A
0.003	19	N/A
0.010	24	2.49×10^{-6}
0.032	24	7.59×10^{-8}
0.100	24	4.40×10^{-6}
0.316	24	1.50×10^{-3}

Even for the 1D simple example, the LSS basis set can be an efficient way to compute interior eigenvalue problems compared to the solution of the eigenvalue problem directly. For comparison of efficiency and accuracy, MATLAB's sparse eigenvalue solver `eigs` is used for the matrix A . We acknowledge that `eigs` may not be the best eigensolver to use for large interior eigenvalue problems, and other choices such as preconditioned conjugate gradient type of solvers, or Jacobi-Davidson type of solvers may give better results. We also remark that the current implementation of the LSS solver is only for proof of principle, and many of its components can be further optimized before a more thorough performance study is to be performed. Here we consider systems of increasing size by changing n_w in the potential function in (4.1) from 8 to 256. Correspondingly the number of grid points n increases from 1600 to 51200, and the number of elements increases proportionally from 8 to 256. $\mu = 2.0, \sigma = 1.0, \tau = 3 \times 10^{-2}$ is used for all systems to compute the eigenvalues within the interval $(\mu - 0.5\sigma, \mu + 0.5\sigma)$. Figure 4.9 shows the time for computing the interior eigenvalues near μ using MATLAB's sparse eigenvalue solver `eigs` ("Global total"), and the time using the LSS basis set ("LSS total"). The tolerance for `eigs` is set to 10^{-5} . The breakdown of the time cost for the LSS solver includes the time for constructing the LSS basis set ("LSS basis"), the time for assembling the projected matrix ("Assembly"), and the time for solving the projected eigenvalue problem ("LSS solve"). Figure 4.10 shows the sparsity pattern of A_U for $n = 6400$, and the sparsity pattern of B_U is by definition the same. The number of nonzero elements is 15.6% of the total number of elements in A_U . The sparsity of the projected matrices is not used in our example here, but can be exploited using alternative methods.

Since the size of the local problem is small, the local eigenvalue problem on each Q_κ is performed using MATLAB's dense eigenvalue solver `eig`, and so is the solution of the generalized eigenvalue problem for the projected matrix. The time for the global solver scales cubically with respect to n , and the constructing the LSS basis and the assembly of the projected matrix increases linearly with respect to n . The solution of the generalized eigenvalue problem also scales cubically with respect to

n , and therefore does not dominate in the LSS solver until $n = 51200$. The cross-over time between the LSS solver and the global solver is around $n = 10000$. For $n = 51200$, the LSS solver costs 46.6 sec, which is 11.2 times faster than the global solver which costs 520.8 sec.

Figure 4.9(b) shows the accuracy of the LSS solver. The Ritz values remain as accurate approximation to the eigenvalues as the number of eigenvalues in the interval increases from 24 to 706.




Figure 4.9. (a) Comparison of time cost between the global solver and the LSS solver for 1D interior eigenvalue problem with increasing system size. See text for details of the comparison. (b) Maximum error of the Ritz values.




Figure 4.10. Sparsity pattern for A_U for $n = 6400$.

4.2. Two-dimensional case. The setup of the 2D example is similar to that in 1D. The global domain is $\Omega = [0, L] \times [0, L]$, and the Laplacian operator is discretized using a 5-point finite difference stencil. The grid spacing is chosen to be $h = 1.0$. The potential function $V(x, y)$ is given by the sum of periodized exponential functions with random perturbation in terms of heights, widths and positions of the exponential functions. This can be viewed as a model potential for a crystal under thermal noise. One realization of this potential is given in Figure 4.11. Let the number of elements M is a square number and the number of grid points n

is divisible by M . Then all n grid points (vertices) are uniformly partitioned into $\sqrt{M} \times \sqrt{M}$ elements. We also assume each extended element Q_κ contains E_κ and its 8 nearest neighbor elements. Figure 4.11 shows the partition of the 2D domain into $8 \times 8 = 64$ elements separated by black dashed lines.




Figure 4.11. One realization of the 2D potential. The domain is partitioned into $8 \times 8 = 64$ elements separated by black dashed lines.

We compare the accuracy of the LSS basis set by comparing the eigenvalues within the interval $(\mu - \sigma, \mu + \sigma)$ with $\mu = -1.0, \sigma = 1.0$. The SVD relative truncation criterion τ is set to be 10^{-1} . Figure 4.12(a) shows the error of Ritz values compared to all the 828 eigenvalues within the interval, and the error of all Ritz values is very small, within 7×10^{-5} . Figure 4.12(b) shows the residual of the Ritz values. For all the Ritz values the residual are below 7×10^{-3} and no spurious eigenvalue is identified for this case.




Figure 4.12. (a) Error of the Ritz values. (b) The 2-norm of residual corresponding to Ritz values for the 2D problem with $\mu = -1.0, \sigma = 1.0, \tau = 10^{-1}$.

Finally we demonstrate the performance of the LSS solver for a 2D problem with increasing size. The number of grid points n increases from 1600 to 25600, and the number of elements increases proportionally from 16 to 256. Figure 4.13 shows the time for computing the interior eigenvalues near μ using MATLAB's sparse eigenvalue solver `eigs` ("Global total"), and the time using the LSS basis set




Figure 4.13. (a) Comparison of time cost between the global solver and the LSS solver for 2D interior eigenvalue problem with increasing system size. See text for details of the comparison. (b) Maximum error of the Ritz values.

(“LSS total”). The tolerance for `eigs` is set to 10^{-5} . The breakdown of the LSS solver includes the time for constructing the LSS basis set (“LSS basis”), the time for assembling the projected matrix (“Assembly”), and the time for solving the projected eigenvalue problem (“LSS solve”). Again the local eigenvalue problem on each Q_κ is performed using MATLAB’s dense eigenvalue solver `eig`, and so is the solution of the generalized eigenvalue problem for the projected matrix. The crossover point between the global solver and the LSS solver is around $n = 3000$. For $n = 25600$, the LSS solver costs 143 sec, which is 8.3 times faster than the global solver which costs 1183 sec.

Figure 4.13(b) shows the accuracy of the LSS solver. The Ritz values remain as an accurate approximation to the eigenvalues as the number of eigenvalues in the interval increases with respect to the system size and no spurious eigenvalue is observed for all cases.

4.3. Sparse matrix with general sparsity pattern. For a general sparse matrix, we take the `turon-m` matrix from the University of Florida matrix collection [8]. The dimension of the matrix is 189924, with 1690876 number of nonzeros. The LU factorization procedure for this matrix is relatively expensive. Using the approximate minimum degree (AMD) ordering strategy provided through the `symamd` command in MATLAB [7]. The number of nonzeros in L and U are 364176421 with a fill-in ratio (i.e., the ratio between the number of nonzeros in L, U and the number of nonzeros in A) is 215. The LU factorization takes 952 sec, and each triangular solve $U^{-1}(L^{-1}b)$ for a random right-hand side vector b takes 0.52 sec, compared to each matrix vector multiplication Ax which takes 0.006 sec. The spectral radius of this matrix is 86. The sparsity pattern of this matrix, together with the histogram of the eigenvalues (unnormalized spectral density) in the interval $(1, 7)$ is given in Figure 4.14(a), (b), respectively.

In order to apply the LSS method to this unstructured matrix, we use the strategy in section 3.4 and use the METIS [15] package interfaced by the `metis mex` program¹ with MATLAB for generating the graph partitioning map ξ .

¹<https://github.com/dgleich/metismex>




Figure 4.14. (a) Sparsity pattern and (b) histogram of the eigenvalues in the interval $(1, 7)$. (c) Accuracy of the Ritz values for the interior eigenvalues in the interval $(3.5, 4.5)$ of the `turon-m` matrix.

We set $\mu = 4.0, \sigma = 0.5$. As in Figure 4.14(b), $\mu = 4.0$ indeed corresponds to interior eigenvalues. We select this region mainly because the spectral density is relatively low so that the computation can be treated on a single computational core. The matrix is partitioned into 16 elements using METIS. The matrix partition routine is efficient and only takes 0.55 sec. Due to the large size of the submatrix on a single extended element, we use `eigs` to solve 500 eigenvalues on each extended element with tolerance set to 10^{-5} , and set the SVD relative truncation criterion τ to be 0.05. The size of the projected matrix is 8000, which is much reduced compared to the dimension of A . The projected generalized eigenvalue problem is solved with the dense eigenvalue solver `eig`.

We compare the accuracy of the LSS basis set by comparing the eigenvalues within the interval $(\mu - \sigma, \mu + \sigma) = (3.5, 4.5)$. There are 914 eigenvalues in this interval, and `eigs` takes 1886 sec to converge to tolerance with 10^{-5} . For LSS, the time for computing the basis functions for all 16 elements is 3989 sec. The time for constructing the projected matrix is 12 sec, and the time for solving the projected matrix is 93 sec. For the projected matrix, we find 919 eigenvalues in total, and identified 5 spurious eigenvalues. After removing the spurious eigenvalues with the largest residual, the accuracy of the Ritz values compared to the true eigenvalues are given in Figure 4.14(c). In this case, the LSS method is more expensive. This is mainly due to the cost for constructing the LSS basis functions. However, this part can be potentially performed independently for each element and without inter-element communication on parallel computers.

5. CONCLUSION

In this paper, we present a method for constructing a novel basis set called the localized spectrum slicing (LSS) basis set. Each function in the LSS basis set is localized both spectrally and spatially, and therefore can be used as an efficient way for representing eigenvectors of a general sparse Hermitian matrix corresponding to a relatively narrow range of eigenvalues. The LSS basis set uses the decay

properties of analytic matrix functions, and can be constructed in a divide-and-conquer method. We show that by carefully tuning one parameter σ , spatial locality and spectral locality of the basis functions can be balanced. The projected matrices can be sparse matrices with reduced sizes.

In terms of the future work, the Gaussian function used in the LSS operator is a smooth approximation to the Dirac- δ function. The same concept of locality can be used to approximate other matrix functions, such as matrix sign functions. This aspect is, e.g., closely related to the recently developed adaptive local basis functions [21] and element orbitals [22] for constructing efficient basis functions for solving the Kohn-Sham density functional theory. The LSS basis set can also be used to efficiently characterize the eigenvectors close to the null space of A , which could potentially be used to construct preconditioners to accelerate linear solves for indefinite problems.

From an efficiency point of view, in the current implementation, the local eigenvalue problem is solved mostly using a dense eigenvalue solver. This is still feasible for the 1D and 2D model problems presented in the numerical section in this paper, but for 3D problems this is going to be too expensive. Efficient iterative solvers, or local Chebyshev expansion based schemes should be used instead. Another practical issue is to control the condition number of the LSS basis set when the SVD truncation criterion is small. An efficient way to identify a subset of well conditioned LSS basis functions will be useful to improve the robustness of the algorithm.

The balance between spatial and spectral locality is an important topic in Fourier analysis and multi-resolution analysis. Because the construction of the LSS basis set is completely algebraic and can be applied to any sparse Hermitian matrix, it is possible to extend the current work to construct multi-resolution basis functions tailored for given matrices, or multi-resolution basis functions for operators on graphs. On the other hand, since the LSS basis set is generated from numerical computation rather than from closed form representation, additional difficulty may arise due to the presence of numerical noise. These topics will be studied in the future.

ACKNOWLEDGMENTS

This work was supported by Laboratory Directed Research and Development (LDRD) funding from Berkeley Lab, provided by the Director, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, the DOE Scientific Discovery through the Advanced Computing (SciDAC) program and the DOE Center for Applied Mathematics for Energy Research Applications (CAMERA) program, and the Alfred. P. Sloan fellowship.

REFERENCES

- [1] M. Benzi, P. Boito, and N. Razouk, *Decay properties of spectral projectors with applications to electronic structure*, SIAM Rev. **55** (2013), no. 1, 3–64, DOI 10.1137/100814019. MR3032838
- [2] M. Benzi and G. H. Golub, *Bounds for the entries of matrix functions with applications to preconditioning*, BIT **39** (1999), no. 3, 417–438, DOI 10.1023/A:1022362401426. MR1708693
- [3] M. Benzi and N. Razouk, *Decay bounds and $O(n)$ algorithms for approximating functions of sparse matrices*, Electron. Trans. Numer. Anal. **28** (2007/08), 16–39. MR2455657
- [4] D. R. Bowler and T. Miyazaki, *$O(N)$ methods in electronic structure calculations*, Rep. Prog. Phys. **75** (2012), 036503.

- [5] W. W. Bradbury and R. Fletcher, *New iterative methods for solution of the eigenproblem*, Numer. Math. **9** (1966), 259–267. MR0219230
- [6] E. R. Davidson, *The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices*, J. Computational Phys. **17** (1975), 87–94. MR0381271
- [7] T. A. Davis, J. R. Gilbert, S. I. Larimore, and E. G. Ng, *A column approximate minimum degree ordering algorithm*, ACM Trans. Math. Software **30** (2004), no. 3, 353–376, DOI 10.1145/1024074.1024079. MR2124396
- [8] T. A. Davis and Y. Hu, *The University of Florida sparse matrix collection*, ACM Trans. Math. Software **38** (2011), no. 1, Art. 1, 25, DOI 10.1145/2049662.2049663. MR2865011
- [9] S. Demko, *Inverses of band matrices and local convergence of spline projections*, SIAM J. Numer. Anal. **14** (1977), no. 4, 616–619. MR0455281
- [10] S. Demko, W. F. Moss, and P. W. Smith, *Decay rates for inverses of band matrices*, Math. Comp. **43** (1984), no. 168, 491–499, DOI 10.2307/2008290. MR758197
- [11] A. George, *Nested dissection of a regular finite element mesh*, SIAM J. Numer. Anal. **10** (1973), 345–363. MR0388756
- [12] S. Goedecker, *Linear scaling electronic structure methods*, Rev. Mod. Phys. **71** (1999), 1085–1123.
- [13] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed., Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, 1996. MR1417720
- [14] R. Gruber and J. Rappaz, *Finite Element Methods in Linear Ideal Magnetohydrodynamics*, Springer Series in Computational Physics, Springer-Verlag, Berlin, 1985. MR800851
- [15] G. Karypis and V. Kumar, *A fast and high quality multilevel scheme for partitioning irregular graphs*, SIAM J. Sci. Comput. **20** (1998), no. 1, 359–392 (electronic), DOI 10.1137/S1064827595287997. MR1639073
- [16] A. V. Knyazev, *New estimates for Ritz vectors*, Math. Comp. **66** (1997), no. 219, 985–995, DOI 10.1090/S0025-5718-97-00855-7. MR1415802
- [17] A. V. Knyazev, *Toward the optimal preconditioned eigensolver: locally optimal block preconditioned conjugate gradient method*, SIAM J. Sci. Comput. **23** (2001), no. 2, 517–541 (electronic), DOI 10.1137/S1064827500366124. MR1861263
- [18] W. Kohn, *Density functional and density matrix method scaling linearly with the number of atoms*, Phys. Rev. Lett. **76** (1996), 3168–3171.
- [19] R. Lehoucq and D. Sorensen, *Implicitly restarted Lanczos method (section 4.5)*, Templates for the Solution of Algebraic Eigenvalue Problems: a Practical Guide (Philadelphia) (Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, eds.), SIAM, 2000, pp. 67–81.
- [20] L. Lin and J. Lu, *Sharp decay estimates of discretized Green's functions for Schrödinger type operators*, arXiv:1511.07957 (2015).
- [21] L. Lin, J. Lu, L. Ying, and W. E, *Adaptive local basis set for Kohn-Sham density functional theory in a discontinuous Galerkin framework I: Total energy calculation*, J. Comput. Phys. **231** (2012), 2140–2154.
- [22] L. Lin and L. Ying, *Element orbitals for Kohn-Sham density functional theory*, Phys. Rev. B **85** (2012), 235144–235153.
- [23] G. Meinardus, *Approximation of Functions: Theory and Numerical Methods*, Expanded translation of the German edition. Translated by Larry L. Schumaker. Springer Tracts in Natural Philosophy, Vol. 13, Springer-Verlag New York, Inc., New York, 1967. MR0217482
- [24] G. Nenciu, *Existence of the exponentially localised Wannier functions*, Comm. Math. Phys. **91** (1983), no. 1, 81–85. MR719812
- [25] E. Polizzi, *Density-matrix-based algorithm for solving eigenvalue problems*, Phys. Rev. B **79** (2009), 115112–115117.
- [26] E. Prodan and W. Kohn, *Nearsightedness of electronic matter*, Proc. Natl. Acad. Sci. **102** (2005), 11635–11638.
- [27] T. Sakurai and H. Sugiura, *A projection method for generalized eigenvalue problems using numerical integration*, Proceedings of the 6th Japan-China Joint Seminar on Numerical Mathematics (Tsukuba, 2002), J. Comput. Appl. Math. **159** (2003), no. 1, 119–128, DOI 10.1016/S0377-0427(03)00565-X. MR2022322
- [28] G. Schofield, J. R. Chelikowsky, and Y. Saad, *A spectrum slicing method for the Kohn-Sham problem*, Comput. Phys. Commun. **183** (2012), no. 3, 497–505, DOI 10.1016/j.cpc.2011.11.005. MR2875139

- [29] H. Zhang, B. Smith, M. Sternberg, and P. Zapol, *SIPs: shift-and-invert parallel spectral transformations*, ACM Trans. Math. Software **33** (2007), no. 2, Art. 9, 19, DOI 10.1145/1236463.1236464. MR2326953

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA BERKELEY AND COMPUTATIONAL RESEARCH DIVISION, LAWRENCE BERKELEY NATIONAL LABORATORY, BERKELEY, CALIFORNIA 94720

E-mail address: `linlin@math.berkeley.edu`