Numerical Linear Algebra with Applications, Vol. 4(5), 351–368 (1997)

Sine Transform Based Preconditioners for Elliptic Problems

Raymond H. Chan[1] and C. K. Wong[2*]

[1]*Department of Mathematics, Chinese University of Hong Kong, Shatin, Hong Kong*

[2]*Department of Mathematics, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, CA 90095-1555, USA*

We consider applying the preconditioned conjugate gradient (PCG) method to solving linear systems $Ax = b$ where the matrix A comes from the discretization of second-order elliptic operators with Dirichlet boundary conditions. Let $(L+\Sigma)\Sigma^{-1}(L^t+\Sigma)$ denote the block Cholesky factorization of A with lower block triangular matrix L and diagonal block matrix Σ. We propose a preconditioner $M = (\hat{L}+\hat{\Sigma})\hat{\Sigma}^{-1}(\hat{L}^t+\hat{\Sigma})$ with block diagonal matrix $\hat{\Sigma}$ and lower block triangular matrix $\hat{L}$. The diagonal blocks of $\hat{\Sigma}$ and the subdiagonal blocks of $\hat{L}$ are respectively the optimal sine transform approximations to the diagonal blocks of Σ and the subdiagonal blocks of L. We show that for two-dimensional domains, the construction cost of M and the cost for each iteration of the PCG algorithm are of order $O(n^2 \log n)$. Furthermore, for rectangular regions, we show that the condition number of the preconditioned system $M^{-1}A$ is of order $O(1)$. In contrast, the system preconditioned by the MILU and MINV methods are of order $O(n)$. We will also show that M can be obtained from A by taking the optimal sine transform approximations of each sub-block of A. Thus, the construction of M is similar to that of Level-1 circulant preconditioners. Our numerical results on two-dimensional square and L-shaped domains show that our method converges faster than the MILU and MINV methods. Extension to higher-dimensional domains will also be discussed.

Numer. Linear Algebra Appl., Vol. 4, 351–368 (1997)
(No. of Figures: 1 No. of Tables: 3 No. of Refs: 14)

KEY WORDS sine transform; preconditioned conjugate gradient method; elliptic partial differential equation

* Correspondence to C. K. Wong, Department of Mathematics, Chinese University of Hong Kong, Shatin, Hong Kong

CCC 1070–5325/97/050351–18 $17.50

Received 14 June 1995
Revised 11 November 1996

1. Introduction

In this paper, we study the numerical solutions of second order elliptic equations with Dirichlet boundary conditions by iterative methods. After discretization by using standard finite-difference method, such problems reduce to the solution of linear systems of the form $Ax = b$ where A is usually symmetric and positive definite. One of the most popular iterative methods for solving such systems is the conjugate gradient (CG) method, see Axelsson and Barker [1, p.18]. In general, the convergence rate of the CG method depends on the condition number $\kappa(A)$ of A. The smaller $\kappa(A)$ is, the faster the convergence of the method will be. However, if $\kappa(A)$ is large, then we can apply the preconditioned conjugate gradient (PCG) method, i.e. we apply the CG method to the preconditioned system $M^{-1}Ax = M^{-1}b$. The matrix M, called a preconditioner to the matrix A, is chosen with two criteria in mind: $Mr = d$ is easy to solve for any vector d; the spectrum of $M^{-1}A$ is clustered and/or $M^{-1}A$ is well conditioned compared with A.

One of the successful classes of preconditioners for elliptic problems is the class of modified incomplete LU (ILU, MILU) factorizations, see for instance, Axelsson and Barker [1, p.337] and Dupont, Kendall and Rachford [11]. The ILU method computes an approximate LU factorization M of A based on the Gaussian elimination in which fill-ins at the (i, j)th element are dropped if the (i, j)th entry of A is zero. In the MILU method, the dropped fill-ins are added back to the diagonal plus an additional term $1/n^2$, where $1/n$ is the mesh-size. For matrices A arising from the discretization of second-order elliptic problems, usually $\kappa(A) = O(n^2)$. However, it has been proved in Dupont, Kendall and Rachford [11] that the condition numbers $\kappa(M^{-1}A)$ of the preconditioned systems for the ILU and the MILU methods are bounded by $O(n^2)$ and $O(n)$ respectively.

Besides the ILU-type preconditioners, incomplete block Cholesky factorizations (INV, MINV) are another popular class of block preconditioners for solving two-dimensional elliptic problems. The motivation behind these preconditioners comes from the complete block Cholesky decomposition of A. For any tridiagonal matrix D_1 and non-singular tridiagonal matrix D_2, the INV preconditioner approximates the Schur complement $D_1 - D_2^{-1}$ by the band matrix $D_1 - \mathcal{T}_3(D_2^{-1})$. Here $\mathcal{T}_3(D_2^{-1})$ denotes the tridiagonal matrix with diagonals identical to the three main diagonals of D_2^{-1}. In the MINV method, the dropped bands are added back to the main diagonal. Numerical experiments in [10] indicate that the condition numbers $\kappa(M^{-1}A)$ of the INV and the MINV methods are bounded by $O(n^2)$ and $O(n)$ respectively.

In [4], Chan and Chan propose another class of preconditioners which is based on averaging the coefficients of A to form a circulant approximation. Part of the motivation is to exploit the fast inversion of circulant systems via the fast Fourier transform (FFT). They proved that circulant preconditioners can be chosen so that $\kappa(M^{-1}A) = O(n)$, just like that for the MILU and MINV type preconditioners.

The fact that the condition numbers of the Dirichlet problems are not improved by circulant preconditioners can be explained partly by a result in Manteuffel and Parter [13, Theorem 3.1]. The result states that in order to improve the condition number, the Dirichlet boundary condition of a given problem should be retained by the preconditioner. For the model problem, i.e. the Laplacian operator with Dirichlet boundary condition, its circulant preconditioner is the same Laplacian operator but with periodic boundary condition. Thus the boundary condition is changed. We note that this circulant preconditioner, being a circulant matrix, can be diagonalized by the Fourier transform matrix. For general Dirichlet

problems with discrete matrix A, the circulant preconditioner proposed in Chan and Chan [4] is defined to be the best circulant approximation to A in Frobenius norm, i.e., it is the best approximation to A in Frobenius norm amongst the class of matrices that can be diagonalized by the Fourier matrix.

However, we note that for the model problem, i.e., the Laplacian operator with Dirichlet boundary condition, its discretization matrix can be diagonalized not by the Fourier matrix, but by the sine-transform matrix. Thus, for general Dirichlet problems, a possible approach to finding a good preconditioner is to look for the best approximation in Frobenius norm amongst all matrices that can be diagonalized by the sine-transform matrix. This gives an exact approximation for the model problem. We remark that Lirkov, Margenov and Vassilevski in [12] construct a circulant preconditioner that matches boundary condition by embedding the original Dirichlet boundary value problem to become a y-periodic boundary value problem. It can be proved that the condition number of the preconditioned systems is of order $O(1)$. However, the cost per PCG iteration, which is dominated by taking FFT on the embedded problem, will be twice as much as that of the original problem, see [12].

In this paper, we propose a class of block preconditioners which is based on the idea of constructing the INV preconditioner. However, we will use matrices that can be diagonalized by the sine transform matrix to approximate the Schur complement $D_1 - D_2^{-1}$ instead of using band matrices as in the INV method or the circulant matrices. For a given matrix K, the optimal sine transform approximation to K is the minimizer of $\|B - K\|_{\mathrm{F}}$ over the set of matrices B that can be diagonalized by the discrete sine transform matrix. Here $\|\cdot\|_{\mathrm{F}}$ denotes the Frobenius norm. The minimizer, denoted by $s(K)$, will be used in constructing our preconditioners. The motivation behind our choosing the sine transform approximation rather than circulant ones is that the optimal sine transform approximation gives exact approximation to the discrete Laplacian with Dirichlet boundary conditions. Therefore we expect our preconditioners still to be good approximations of elliptic operators that are small perturbations of the Laplacian. Our theoretical and numerical results in Sections 3 and 5 verify this claim.

The construction of our block preconditioners M is similar to that of the INV, namely, both use easily invertible matrices to approximate Schur's complements. However, for elliptic problems on two-dimensional rectangular domains, we will show that our M can also be constructed from the matrix A by taking the optimal sine transform approximations of each n-by-n block of A. Thus, the construction of M is also similar to the so-called Level-1 circulant approximation of A as defined in Chan and Olkin [9]. We will see that the construction cost of M and the matrix–vector multiplication $M^{-1}v$ for any vector v can be done in $O(n^2 \log n)$ operations. Furthermore, we will show that the condition number of the preconditioned system $M^{-1}A$ is of order $O(1)$. Thus M is an efficient preconditioner. We note that both the construction (based on averaging of the coefficients of the elliptic operator) and the inversion (using fast sine transforms) of our preconditioner are highly parallelizable. Moreover, we remark that the construction approach we use can easily be extended to two-dimensional irregular domains or higher dimensional regular domains. Our numerical results on two-dimensional rectangular and L-shaped domains show that our preconditioners work better than the MILU, MINV and circulant preconditioners.

The outline of the paper is as follows. In the next section, we will describe the optimal sine transform approximation for general matrices. In section 3, we use the optimal sine transform approximation to construct a block preconditioner on a two-dimensional rectangular domain. We will show that the INV approach is the same as the Level-1 approach and we will also analyze the spectral condition number of the preconditioned system. In section

4, we extend the definition of our preconditioner to two-dimensional irregular domains and also higher dimensional regular domains. Finally, numerical results and concluding remarks are given in section 5 and section 6 respectively.

2. Sine transform approximations for general matrices

In this section we recall some of the results in approximating a given matrix by matrices that can be diagonalized by the discrete sine transform. Let S_n be the n-by-n discrete sine transform matrix. Its (i, j)th entry is given by

$$\sqrt{\frac{2}{n+1}}\sin(\frac{\pi ij}{n+1}), \quad 1 \le i, j \le n$$

We note that S_n is symmetric, orthogonal and the matrix–vector multiplication $S_n v$ can be computed in $O(n \log n)$ operations for any n-vector v, see Yip and Rao [14].

Let $\mathcal{B}_{n\times n}$ be the vector space containing all the n-by-n matrices that can be diagonalized by S_n. Given $B \in \mathcal{B}_{n\times n}$, we now show that the product $B^{-1}v$ for any vector v can be done in $O(n \log n)$ operations. We first emphasize the relationship between the first column of B and its eigenvalues. Since $B = S_n \Lambda S_n$ for some diagonal matrix Λ, if $e_n = (1, 0, \ldots, 0)^t$ and $1_n = (1, 1, \ldots, 1)^t$, then we have

$$D^{-1}S_n B e_n = \Lambda 1_n \tag{2.1}$$

where D is the diagonal matrix whose diagonal is equal to $S_n e_n$. Thus, by exploiting the fast sine transform, the matrix Λ and hence the matrix–vector multiplication $B^{-1}v = S_n \Lambda^{-1} S_n v$ can be computed in $O(n \log n)$ operations.

Given an n-by-n matrix A, we are interested in finding a matrix $B \in \mathcal{B}_{n\times n}$ which minimizes $\|A_n - B\|_F$ in the Frobenius norm $\|\cdot\|_F$. We will denote the minimizer by $s(A_n)$ and called it the optimal sine transform approximation to A_n. The following lemma gives some basic properties of $s(A_n)$.

Lemma 2.1. *Let A_n be an n-by-n symmetric matrix and $s(A_n)$ be the minimizer of $\|B_n - A_n\|_F$ over all $B_n \in \mathcal{B}_{n\times n}$. Then $s(A_n)$ is uniquely determined by A_n and is given by*

$$s(A_n) = S_n \delta(S_n A_n S_n) S_n \tag{2.2}$$

where $\delta(S_n A_n S_n)$ denotes the diagonal matrix whose diagonal is equal to the diagonal of the matrix $S_n A_n S_n$. Furthermore,

$$\lambda_{\min}(A_n) \le \lambda_{\min}(\delta(A_n)) \le \lambda_{\max}(\delta(A_n)) \le \lambda_{\max}(A_n)$$

In particular, if A_n is positive definite, then $s(A_n)$ is also positive definite.

Proof
Follows directly from Chan and Jin [7, Lemma 1]. ∎

We note that forming $s(A_n)$ by computing all the diagonal entries of $S_n A_n S_n$ as in (2.2) requires $O(n^2 \log n)$ operations. Chan, Ng and Wong [8] give another approach of constructing $s(A_n)$ which reduces the cost to $O(n^2)$ operations. Before we describe how the matrix $s(A_n)$ is formed, we need the following definitions.

Definition 2.1. *Let Q_i, $i = 1, \ldots, n$, be n-by-n matrices with the (h, k)th entry given by*

$$Q_i(h,k) = \begin{cases} 1 & \text{if } |h-k| = i-1 \\ -1 & \text{if } h+k = i-1 \\ -1 & \text{if } h+k = 2n-i+3 \\ 0 & \text{otherwise} \end{cases}$$

Boman and Koltracht in [3] showed that $\{Q_i\}_{i=1}^n$ is a basis for $\mathcal{B}_{n\times n}$. We note that each Q_i is a sparse matrix with at most $2n$ non-zero entries. Also, we let

$$r_n = (1_n^t(Q_1 \circ A_n)1_n, 1_n^t(Q_2 \circ A_n)1_n, \ldots, 1_n^t(Q_n \circ A_n)1_n)^t \tag{2.3}$$

where $\circ$ is the Hadamard product.

With the help of the above definitions, we can give explicit formula for the entries of the minimizer $s(A_n)$.

Lemma 2.2. (Chan, Ng and Wong [8]) *Let $A_n = [a_{jk}]$ be an n-by-n symmetric matrix and $s(A_n)$ be the minimizer of $\|B_n - A_n\|_F$ over all $B_n \in \mathcal{B}_{n\times n}$. Denote $\mathbf{z}$ to be the first column of $s(A_n)$. If s_o and s_e are defined respectively to be the sum of the odd and even index entries of $\mathbf{r}_n$, then we have*

$$\begin{aligned} [\mathbf{z}]_1 &= \frac{1}{2(n+1)}(2[\mathbf{r}_n]_1 - [\mathbf{r}_n]_3) \\ [\mathbf{z}]_i &= \frac{1}{2(n+1)}([\mathbf{r}_n]_i - [\mathbf{r}_n]_{i+2}) \qquad i = 2, \ldots, n-2 \end{aligned}$$

with

$$\begin{aligned} [\mathbf{z}]_{n-1} &= \frac{1}{2(n+1)}(s_o + [\mathbf{r}_n]_{n-1}) \\ [\mathbf{z}]_n &= \frac{1}{2(n+1)}(2s_e + [\mathbf{r}_n]_n) \end{aligned}$$

if n is even; and

$$\begin{aligned} [\mathbf{z}]_{n-1} &= \frac{1}{2(n+1)}(s_e + [\mathbf{r}_n]_{n-1}) \\ [\mathbf{z}]_n &= \frac{1}{2(n+1)}(2s_o + [\mathbf{r}_n]_n) \end{aligned}$$

if n is odd.

We remark that if A_n has no special structure, then r_n can be computed in $O(n^2)$ operations because Q_i are sparse with only $O(n)$ non-zero entries each. Therefore, $s(A_n)$ can be computed in $O(n^2)$ operations. However, we show below that if A_n is a band matrix, then the cost can be reduced.

Corollary 2.1. *The construction cost of the optimal sine transformation approximation $s(A_n)$ is of order $O(\ell n)$ if A_n is an n-by-n band matrix with bandwidth ℓ.*

Proof
By the definition of Q_i in Definition 1, $Q_i \circ A_n$ has at most 2ℓ non-zero entries for $i = 3, \ldots, n$. When $i = 1, 2$, the number of non-zero entries of $Q_i \circ A_n$ is less than $2n$. Therefore, the cost of forming r_n in (2.3) is bounded by $O(\ell n)$. Hence by Lemma 2, $s(A_n)$ can be calculated in $O(\ell n)$ operations. ■

It is well known that the discrete sine transform matrix S_n diagonalizes the set of symmetric tridiagonal Toeplitz matrices. By the definition of $s(A_n)$ or (2.2), the optimal sine transform preconditioner gives exact approximation to all matrices in the set, in particular to the one-dimensional discrete Laplacian: tridiag$[-1, 2, -1]$. Therefore, the system tridiag$[-1, 2, -1]x = b$ can be solved in exactly one iteration by the PCG method with $s(A_n)$ as preconditioner and the condition number of the preconditioned system is $O(1)$. We therefore expect our preconditioner still to be a good approximation for elliptic operators that are small perturbations of the Laplacian. In contrast, we remark that the condition number of the system preconditioned by the optimal circulant preconditioner is of $O(n^{3/2})$, see Chan and Chan [4]. In the following sections, we will use the optimal sine transform approximation to construct block preconditioners for elliptic problems in higher dimensional domains.

3. Two-dimensional rectangular domains

In this section, we apply the optimal sine transform approximation to construct a block preconditioner on the two-dimensional rectangular domain. The construction is given in section 3.1 and the analysis of the spectral condition number of the preconditioned system is given in section 3.2.

3.1. Construction of preconditioners

Consider the two-dimensional elliptic problems

$$-(a(x, y)u_x)_x - (b(x, y)u_y)_y = f(x, y) \tag{3.1}$$

on the unit square $[0, 1] \times [0, 1]$ with Dirichlet boundary condition. Assume that the coefficient functions $a(x, y)$, $b(x, y)$ satisfy

$$0 < c_{\min} \leq a(x, y), b(x, y) \leq c_{\max} \tag{3.2}$$

for some constants $c_{\min}$ and $c_{\max}$. Let the domain be discretized by using a uniform grid with n internal grid points in each co-ordinate direction. With the usual five-point centered differencing, the resulting discretization matrix A will be an n^2-by-n^2 symmetric positive definite matrix of the form

$$A = \begin{pmatrix} D_1 & A_2 & & & \\ A_2 & D_2 & A_3 & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & A_n \\ & & & A_n & D_n \end{pmatrix} \tag{3.3}$$

Here D_i are symmetric tridiagonal matrices for $1 \le i \le n$ and A_i are diagonal matrices for $2 \le i \le n$.

Let

$$L = \begin{pmatrix} 0 & & & & \\ A_2 & 0 & & & \\ & A_3 & \ddots & & \\ & & \ddots & \ddots & \\ & & & A_n & 0 \end{pmatrix} \tag{3.4}$$

be the block lower triangular matrix. Then the block Cholesky factorization of A can be written as

$$A = (\Sigma + L)\Sigma^{-1}(\Sigma + L^t)$$

where Σ is a symmetric block diagonal matrix with diagonal blocks Σ_i satisfying

$$\begin{aligned} \Sigma_1 &= D_1 \\ \Sigma_i &= D_i - A_i \Sigma_{i-1}^{-1} A_i, \quad 2 \le i \le n \end{aligned} \tag{3.5}$$

Because of the work and storage required in large problems for computing the Schur complements Σ_i, carrying out the complete block Cholesky factorization is not an efficient way for solving the systems $Ax = b$.

Concus, Golub and Meurant in [10] focus on sparse approximations on the matrices Σ_i. Their idea is to approximate Σ_i by band matrices consisting of the three main diagonals of Σ_i. More precisely, their preconditioner M is defined as follows:

$$M = (\Delta + L)\Delta^{-1}(\Delta + L^t)$$

where Δ is a symmetric block diagonal matrix with diagonal blocks Δ_i satisfying

$$\begin{aligned} \Delta_1 &= D_1 \\ \Delta_i &= D_i - \mathcal{T}_3(A_i \Delta_{i-1}^{-1} A_i), \quad 2 \le i \le n \end{aligned} \tag{3.6}$$

Here $\mathcal{T}_3(A_i \Delta_{i-1}^{-1} A_i)$ is the tridiagonal matrix consisting of the three main diagonals of $A_i \Delta_{i-1}^{-1} A_i$.

The preconditioner M is called the INV preconditioner in [10] and each Δ_i is proved to be positive definite by showing that the minimum row sum of matrices $\Delta_i + A_{i+1}$ are greater than zero. Hence, the factorization process (3.6) can be carried out for $2 \le i \le n$. It was also proved that the inverse of an n-by-n symmetric tridiagonal matrix is determined by two n-vectors which can be computed in $O(n)$ operations and hence forming $\mathcal{T}_3(A_i \Delta_{i-1}^{-1} A_i)$ only needs $O(n)$ operations. As a result, forming the INV preconditioner M and computing $M^{-1}v$ can be done in $O(n^2)$ operations.

In this paper, we use an optimal sine transform preconditioner to approximate each Σ_i. We follow the approach in [10] and propose, for any matrix A that satisfies (3.3), a preconditioner $\hat{s}(A)$ of the form

$$\hat{s}(A) = (\hat{\Sigma} + \hat{L})\hat{\Sigma}^{-1}(\hat{\Sigma} + \hat{L}^t) \tag{3.7}$$

Here

$$\hat{L} \equiv \begin{pmatrix} 0 & & & & \\ s(A_2) & 0 & & & \\ & s(A_3) & \cdot & & \\ & & \ddots & \ddots & \\ & & & s(A_n) & 0 \end{pmatrix} \tag{3.8}$$

is a block lower triangular matrix which approximates L and

$$\hat{\Sigma} \equiv \begin{pmatrix} \hat{\Sigma}_1 & & & 0 \\ & \hat{\Sigma}_2 & & \\ & & \ddots & \\ 0 & & & \hat{\Sigma}_n \end{pmatrix}$$

is a diagonal block matrix with diagonal blocks $\hat{\Sigma}_i$ satisfying

$$\begin{aligned} \hat{\Sigma}_1 &= s(D_1) \\ \hat{\Sigma}_i &= s(D_i) - s(A_i)\hat{\Sigma}_{i-1}^{-1} s(A_i), \quad 2 \le i \le n \end{aligned} \tag{3.9}$$

In order to ensure the above factorization can be carried out, we show in the next theorem that each $\hat{\Sigma}_i$ is positive definite. In addition, the theorem also proves that the preconditioner $\hat{s}(A)$ actually comes from the matrix A by taking the optimal sine transform approximations to each n-by-n block of A. Thus, the construction of our preconditioner is similar to that of the Level-1 circulant preconditioners proposed by Chan and Olkin [9].

Theorem 3.1. *Let A be an n^2-by-n^2 symmetric positive definite matrix of the form given in (3.3). Define*

$$G = \begin{pmatrix} s(D_1) & s(A_2) & & & \\ s(A_2) & s(D_2) & s(A_3) & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & s(A_n) \\ & & & s(A_n) & s(D_n) \end{pmatrix} \tag{3.10}$$

where $s(\cdot)$ is the optimal sine approximation. Then G is positive definite and the block Cholesky factorization of G is given by

$$(\hat{\Sigma} + \hat{L})\hat{\Sigma}^{-1}(\hat{\Sigma} + \hat{L}^t)$$

where $\hat{L}$ and $\hat{\Sigma}$ are given in (3.8) and (3.9) respectively. In particular, $G = \hat{s}(A)$ and hence $\hat{s}(A)$ and $\hat{\Sigma}_i$ are positive definite.

Proof

The positive definiteness of G follows from Chan and Jin [7, Theorem 1]. Similar to (3.5), the block Cholesky factorization of G is given by

$$G = (\Phi + \hat{L})\Phi^{-1}(\Phi + \hat{L}^t)$$

where Φ is a symmetric block diagonal matrix with diagonal blocks Φ_i satisfying

$$\begin{aligned}\Phi_1 &= s(D_1)\\ \Phi_i &= s(D_i) - s(A_i)\Phi_{i-1}^{-1}s(A_i), \quad 2 \le i \le n\end{aligned}$$

Comparing this with the definition of $\hat{\Sigma}$ in (3.9), we see that $\Phi = \hat{\Sigma}$ and hence the block preconditioner $\hat{s}(A)$ is identical to G. ∎

In the following, we apply the recursion formula (3.9) to show that both the preconditioner $\hat{s}(A)$ and the solution of the linear system $\hat{s}(A)x = b$ can be obtained in $O(n^2 \log n)$ operations. Since matrices $\hat{\Sigma}_i$, $s(D_i)$ and $s(A_i)$ all belong to $\mathcal{B}_{n\times n}$, we let Λ_i, Λ_i^d and Λ_i^a be their corresponding eigenvalue matrices. Specifically, we have

$$\hat{\Sigma}_i = S_n \Lambda_i S_n, \qquad s(D_i) = S_n \Lambda_i^d S_n \quad \text{and} \quad s(A_i) = S_n \Lambda_i^a S_n \tag{3.11}$$

As D_i and A_i are band matrices, by Corollary 2.1, forming $s(D_i)$ and $s(A_i)$ costs $O(n)$ operations. By (2.1), Λ_i^d and Λ_i^a can be computed in $O(n \log n)$ operations. Using (3.9), we have the following equality which relates the eigenvalues of matrices $\hat{\Sigma}_i$ and $\hat{\Sigma}_{i-1}$

$$\Lambda_i = \Lambda_i^d - \Lambda_i^a \Lambda_{i-1}^{-1} \Lambda_i^a$$

Therefore, Λ_i can be obtained from Λ_{i-1} in $O(n \log n)$ steps. Hence, forming $\hat{\Sigma}$ and the preconditioner $\hat{s}(A)$ require only $O(n^2 \log n)$ operations. Finally in solving the system $\hat{s}(A)x = b$, we are required to multiply each $\hat{\Sigma}_i$ by some vector and to solve systems with coefficient matrices $\hat{\Sigma}_i$. By noting the equalities in (3.11), it is easy to see that the system $\hat{s}(A)x = b$ can also be solved in $O(n^2 \log n)$ operations.

We note that, by Theorem 1, the preconditioner $\hat{s}(A)$ can also be constructed by using the approach used in constructing Level-1 circulant preconditioners in Chan and Olkin [9]. The cost will be cheaper but still requires $O(n \log n)$ operations. As we will see in section 4, the INV approach of constructing $\hat{s}(A)$ in (3.7) can be easily extended to irregular domains but the Level-1 approach in (3.10) cannot.

3.2. Convergence analysis of the preconditioners

In this subsection, we are going to show that the condition numbers of the preconditioned systems $\kappa(\hat{s}(A)^{-1}A)$ are bounded by a constant which is independent of the size of the matrix A. Hence, the convergence rate of the conjugate gradient method when applied to the preconditioned systems $\hat{s}(A)^{-1}A$ is linear, see Axelsson and Barker [1, p.26].

Before we present our proof, let us introduce the following notations. Let A_{nn} be any n^2-by-n^2 matrix partitioned as

$$A_{nn} = \begin{pmatrix} A_{1,1} & A_{1,2} & \dots & A_{1,n} \\ A_{2,1} & A_{2,2} & \dots & A_{2,n} \\ \vdots & \ddots & \ddots & \vdots \\ A_{n,1} & A_{n,2} & \dots & A_{n,n} \end{pmatrix}$$

Here $A_{i,j}$ are square matrices of order n. Then we define $\hat{\delta}(A_{nn})$ to be the matrix

$$\hat{\delta}(A_{nn}) \equiv \begin{pmatrix} \delta(A_{1,1}) & \delta(A_{1,2}) & \dots & \delta(A_{1,n}) \\ \delta(A_{2,1}) & \delta(A_{2,2}) & \dots & \delta(A_{2,n}) \\ \vdots & \ddots & \ddots & \vdots \\ \delta(A_{n,1}) & \delta(A_{n,2}) & \dots & \delta(A_{n,n}) \end{pmatrix}$$

where each block $\delta(A_{i,j})$ is the diagonal matrix of order n whose diagonal is equal to the diagonal of the matrix $A_{i,j}$. The following lemma relates eigenvalues of matrices A_{nn} and $\hat{\delta}(A_{nn})$ and is useful in our analysis of the convergence rate.

Lemma 3.1. (Chan and Jin [7]) *Given any n^2-by-n^2 symmetric matrix A_{nn}, we have*

$$\lambda_{\min}(A_{nn}) \le \lambda_{\min}(\hat{\delta}(A_{nn})) \le \lambda_{\max}(\hat{\delta}(A_{nn})) \le \lambda_{\max}(A_{nn})$$

In particular, if A_{nn} is positive definite, then $\hat{\delta}(A_{nn})$ is also positive definite.

Using the $\delta(\cdot)$ notation, we can give another formula for $\hat{s}(A)$.

Lemma 3.2. *Let A be an n^2-by-n^2 symmetric positive definite matrix of the form given in (3.3). Then*

$$\hat{s}(A) = (I \otimes S_n)\hat{\delta}((I \otimes S_n)A(I \otimes S_n))(I \otimes S_n) \tag{3.12}$$

Proof
We first observe that $\hat{\delta}((I \otimes S_n)A(I \otimes S_n))$ is equal to

$$\begin{pmatrix} \delta(S_nD_1S_n) & \delta(S_nA_2S_n) & & & & \\ \delta(S_nA_2S_n) & \delta(S_nD_2S_n) & \delta(S_nA_3S_n) & & & \\ & \ddots & \ddots & \ddots & & \\ & & \ddots & \ddots & \delta(S_nA_nS_n) & \\ & & & \delta(S_nA_nS_n) & \delta(S_nD_nS_n) & \end{pmatrix}$$

Then the lemma follows by using (2.2) and (3.10). ∎

With the help of these two lemmas, we prove the main theorem in this section.

Theorem 3.2. *Let A be the five-point discretization matrix of (3.1) on the unit square satisfying conditions (3.2). If $\hat{s}(A)$ is the preconditioner defined in (3.7), we have*

$$\kappa(\hat{s}(A)^{-1}A) \le (\frac{c_{\max}}{c_{\min}})^2 \tag{3.13}$$

Proof
Let A_L be the five-point discretization matrix of the Laplace operator on the unit square, i.e.,

$$A_L = \text{tridiag}[-1, 2, -1] \otimes I_n + I_n \otimes \text{tridiag}[-1, 2, -1]$$

Then

$$c_{\min}A_L \le A \le c_{\max}A_L \tag{3.14}$$

see [4, (4.2)]. Multiplying $(I \otimes S_n)$ on the left and on the right of the matrices in the above inequality and applying Lemma 3.1, we have

$$\begin{aligned} c_{\min}(I \otimes S_n)A_L(I \otimes S_n) &= c_{\min}\hat{\delta}((I \otimes S_n)A_L(I \otimes S_n)) \\ &\leq \hat{\delta}((I \otimes S_n)A(I \otimes S_n)) \\ &\leq c_{\max}((I \otimes S_n)A_L(I \otimes S_n)) \\ &= c_{\max}(I \otimes S_n)A_L(I \otimes S_n) \end{aligned}$$

Multiplying again the left and the right of the above matrices by $(I \otimes S_n)$ and by noting (3.12), we then have

$$c_{\min}A_L \leq \hat{s}(A) = (I \otimes S_n)\hat{\delta}((I \otimes S_n)A(I \otimes S_n))(I \otimes S_n) = \hat{s}(A) \leq c_{\max}A_L \quad (3.15)$$

This shows that $\hat{s}(A) - c_{\min}A_L$ and $c_{\max}A_L - \hat{s}(A)$ are positive semidefinite matrices. Combining (3.14) and (3.15), we have

$$0 < \frac{c_{\min}}{c_{\max}}\frac{x^tA_Lx}{x^tA_Lx} \leq \frac{x^tAx}{x^t\hat{s}(A)x} \leq \frac{c_{\max}}{c_{\min}}\frac{x^tA_Lx}{x^tA_Lx}$$

Hence, the theorem is proved. ■

4. Extension to other domains

In this section, we apply the optimal sine transform approximation to constructing preconditioners for two-dimensional irregular domains and higher dimensional regular domains. The two cases will be discussed respectively in sections 4.1 and 4.2.

4.1. Two-dimensional irregular domains

For ease of presentation, we consider irregular domains that are a union of rectangular domains. In this case, the matrix A still has the form given in (3.3) but the diagonal submatrices D_i of A are of different sizes and the submatrices A_i may not be square matrices. We note that the number of submatrices of A_i that are not square is proportional to the number of rectangular regions used in forming the given domain. For an L-shaped domain, there is only one A_i that is not square; and for a T-shaped domain, the number is two.

When the A_i are square, we can carry out the construction of the preconditioner just as we did in section 3. Therefore let us concentrate on the sub-block of A where the A_i are not square. In particular, let us consider A to be of the form

$$A = \begin{pmatrix} D_1 & A_2^t \\ A_2 & D_2 \end{pmatrix}$$

where D_1 and D_2 are respectively n_1-by-n_1 and n_2-by-n_2 symmetric tridiagonal matrices and A_2 is an n_2-by-n_1 rectangular matrix with $(A_2)_{ij} = 0$ for $i \neq j$. Without loss of generality we assume that $n_1 > n_2$.

The block preconditioner $\hat{s}(A)$ in this case will still be of the form given by (3.7) except that the $\hat{\Sigma}_i$, which are approximations to Σ_i, are defined as follows:

$$\begin{aligned}\hat{\Sigma}_1 &= s(D_1)\\ \hat{\Sigma}_2 &= s(D_2) - s(A_2 E_{n_2,n_1}^t) s(E_{n_2,n_1} \hat{\Sigma}_1^{-1} E_{n_2,n_1}^t) s(E_{n_2,n_1} A_2^t)\end{aligned}$$

where E_{n_2,n_1} is an n_2-by-n_1 matrix such that $(E_{n_2,n_1})_{i,j} = \delta_{ij}$, the Kronecker delta. We note that the matrix $E_{n_2,n_1} \hat{\Sigma}_1^{-1} E_{n_2,n_1}^t$ is the n_2-by-n_2 principal submatrix of $\hat{\Sigma}_1^{-1}$. Hence it is a dense matrix without any special algebraic structure. Constructing the optimal sine transform approximation to the matrix requires $O(n_2^2)$ operations according to (2.3).

Thus, we conclude that for a given irregular domain which is the union of m rectangular regions, the construction cost of the preconditioner $\hat{s}(A)$ is bounded by $O(mn^2) + O(n^2 \log n)$ where n is the size of the largest diagonal block of A. For an L-shaped or T-shaped domain, the cost will still be bounded by $O(n^2 \log n)$. Once $\hat{s}(A)$ is formed, the cost of solving the system $\hat{s}(A)y = v$ is the same as in the rectangular case, i.e., it is bounded by $O(n^2 \log n)$.

4.2. Higher dimensional rectangular domains

In this section, we extend the construction of our preconditioner to the three-dimensional cubic domain $[0, 1]^3$. The approach can easily be generalized to higher dimensional regular domains.

By applying the usual seven-point centered differencing with n internal grid nodes in each co-ordinate direction, the resulting discretization matrix A will be an n^3-by-n^3 symmetric positive definite matrix of the form given in (3.3) with A_i and D_i being n^2-by-n^2 diagonal matrices and n^2-by-n^2 tridiagonal block matrices respectively. If we let L be the matrix defined in (3.4), then the block Cholesky factorization of A is given by

$$A = (\Sigma + L)\Sigma^{-1}(\Sigma + L^t)$$

where Σ is a symmetric block diagonal matrix with each n^2-by-n^2 diagonal blocks Σ_i satisfying (3.5).

To emulate the approach of constructing the INV preconditioner, namely using an easily invertible matrix to approximate the Schur complement Σ_i, we introduce the so-called Level-2 circulant approximation to Σ_i, see Chan and Olkin [9]. We need the following notation first. For any n^2-by-n^2 block matrix A_{nn}, we denote $(A_{nn})_{i,j;k,l}$ to be the (i, j)th entry of the (k, l)th block of A_{nn}. Let P be the permutation matrix that satisfies

$$(P^t A_{nn} P)_{i,j;k,l} = (A_{nn})_{k,l;i,j}, \quad 1 \le i, j \le n, 1 \le k, l \le n$$

Then we define an approximation $\check{s}(A_{nn})$ to A_{nn} by

$$\check{s}(A_{nn}) = P\hat{s}(P^t \hat{s}(A_{nn}) P) P^t \tag{4.1}$$

By Theorem 3 in [7], the approximation $\check{s}(A_{nn})$ can be diagonalized by $S_n \otimes S_n$. Specifically we have,

$$\check{s}(A_{nn}) = (S_n \otimes S_n)\delta((S_n \otimes S_n) A_{nn} (S_n \otimes S_n))(S_n \otimes S_n)$$

Using this equality, we can relate the eigenvalues of $\check{s}(A_{nn})$ with its first column as in (2.1). Hence, the inverse of $\check{s}(A_{nn})$ can easily be computed.

Now we apply the approximation $\check{s}(\cdot)$ to each Σ_i and define the block preconditioner $\tilde{s}(A)$ for the matrix A as

$$\tilde{s}(A) = (\tilde{\Sigma} + \tilde{L})\tilde{\Sigma}^{-1}(\tilde{\Sigma} + \tilde{L}^t)$$

where

$$\tilde{L} = \begin{pmatrix} 0 & & & & \\ \check{s}(A_2) & 0 & & & \\ & \check{s}(A_3) & \cdot & & \\ & & \ddots & \ddots & \\ & & & \check{s}(A_n) & 0 \end{pmatrix}$$

is a block lower triangular matrix which approximates L and

$$\tilde{\Sigma} \equiv \begin{pmatrix} \tilde{\Sigma}_1 & & & 0 \\ & \tilde{\Sigma}_2 & & \\ & & \ddots & \\ 0 & & & \tilde{\Sigma}_n \end{pmatrix}$$

is a diagonal block matrix with diagonal blocks $\tilde{\Sigma}_i$ of order n^2 satisfying

$$\begin{aligned} \tilde{\Sigma}_1 &= \check{s}(D_1) \\ \tilde{\Sigma}_i &= \check{s}(D_i) - \check{s}(A_i)\tilde{\Sigma}_{i-1}^{-1}\check{s}(A_i), \quad 2 \le i \le n \end{aligned} \tag{4.2}$$

Similar to Theorem 3.1, we can prove that the preconditioner $\tilde{s}(A)$ can be obtained from A by taking the $\check{s}(\cdot)$ approximation of each n^2-by-n^2 block of A. Hence by Theorem 1 in [7], $\tilde{s}(A)$ is positive definite. Also, since matrices A_i are diagonal and D_i have the same graph structure as two-dimensional discretization matrices, by (4.1), we see that obtaining each $\check{s}(A_i)$ and $\check{s}(D_i)$ requires $O(n^2 \log n)$ operations. Hence by (4.2) the construction cost of the preconditioner $\tilde{s}(A)$ is of order $O(n^3 \log n)$.

5. Numerical experiments

In this section, we compare the performance of our method with the MILU, MINV and the circulant type preconditioning method proposed in [12]. The equation we used is

$$\frac{\partial}{\partial x}\Big[\big(1 + \epsilon e^{x+y}\big)\frac{\partial u}{\partial x}\Big] + \frac{\partial}{\partial y}\Big[\big(1 + \frac{\epsilon}{2}\sin(2\pi(x+y))\big)\frac{\partial u}{\partial y}\Big] = f(x, y) \tag{5.1}$$

with $u = 0$ on the boundary. The ϵ here is a parameter controlling the variation of the coefficient functions. We discretize the equation using the standard five-point scheme. The initial guess and the right hand side are chosen to be random vectors and are the same for all methods. All computations are done by Matlab on a SUN sparc workstation. The iterations are stopped when the residual vector r_k at the kth iteration satisfies $||r_k||_2/||r_0||_2 < 10^{-6}$.

In Tables 1(a)–2(b), we show the numbers of iterations required for convergence for (5.1) with different choices of ϵ and preconditioners. The notation in the second row of the tables indicates the type of preconditioner we used with I and C meaning no preconditioner or the circulant type preconditioner as proposed in [12]. The parameter n is equal to $1/h$

Table 1(a). Numbers of iterations for the unit square

ϵ			0.0					0.01		
n	I	$\hat{s}(A)$	MINV	MILU	C	I	$\hat{s}(A)$	MINV	MILU	C
8	22	1	5	9	1	25	3	5	9	3
16	43	1	7	13	1	47	3	7	13	3
32	82	1	11	19	1	91	3	11	19	3
64	154	1	16	27	1	159	3	16	27	3
128	306	1	23	39	1	339	3	23	39	3

Table 1(b). Numbers of iterations for the unit square

ϵ			0.1					1.0		
n	I	$\hat{s}(A)$	MINV	MILU	C	I	$\hat{s}(A)$	MINV	MILU	C
8	25	5	5	9	5	30	9	4	9	9
16	47	5	7	13	5	59	10	6	13	9
32	96	5	11	19	5	121	10	9	18	9
64	185	6	15	27	5	247	10	13	26	9
128	388	6	23	39	5	515	11	20	37	9

Table 2(a). Numbers of iterations for the L-shaped domain

ϵ			0.0				0.01		
n	I	$\hat{s}(A)$	MINV	MILU	I	$\hat{s}(A)$	MINV	MILU
8	21	3	4	9	21	3	4	9
16	39	4	6	12	41	4	6	12
32	74	4	10	18	77	4	10	18
64	144	4	14	25	153	4	14	25
128	286	4	22	37	297	4	22	37

Table 2(b). Numbers of iterations for the L-shaped domain

ϵ			0.1				1.0		
n	I	$\hat{s}(A)$	MINV	MILU	I	$\hat{s}(A)$	MINV	MILU
8	22	5	4	9	25	8	4	9
16	41	5	7	12	47	10	6	12
32	81	6	10	17	93	11	9	17
64	161	6	14	25	195	13	14	25
128	323	7	22	36	405	17	20	36

Table 3. Number of million flops per PCG iteration

	$\hat{s}(A)$	MINV	MILU	C
m	$n = 2^m - 1$	$n = 2^m - 1$	$n = 2^m - 1$	$n = 2^m$
5	0.13	0.04	0.04	0.22
6	0.56	0.18	0.17	0.94
7	2.39	0.71	0.69	4.00
8	10.23	2.86	2.79	16.95

where h is the mesh size. Tables 1(a), (b) and 2(a), (b) are the results on the unit square and on the L-shaped domain $[0, \frac{1}{2}] \times [0, 1] \cup [\frac{1}{2}, 1] \times [0, \frac{1}{2}]$ respectively. Note that the circulant preconditioner proposed in [12] is not defined in irregular domains. We see from the tables that the numbers of iterations for I, $\hat{s}(A)$, MILU, MINV and C grow as $O(n)$, $O(1)$, $O(\sqrt{n})$, $O(\sqrt{n})$ and $O(1)$ respectively. We emphasize that although the numbers of iterations for C are roughly the same as those for $\hat{s}(A)$, the cost per PCG iteration for C is approximately twice that for $\hat{s}(A)$. It is because C is the circulant preconditioner constructed for the problem that it extends the original Dirichlet boundary value problem periodically in the y-direction [12]. Hence, C is a $2n^2$-by-$2n^2$ circulant block matrix having block size $2n$-by-$2n$. We will present results to show that $\hat{s}(A)$ is better than C in terms of the overall computational complexity and the cost per PCG iteration. We also remark that the MILU and MINV preconditioning methods are less sensitive to the changes in ϵ.

We note that in Table 1(b), the numbers of iterations for the preconditioner $\hat{s}(A)$ seem to grow slightly whereas by Theorem 2 they should be bounded independent of n for large n. Owing to computational time limitation, let us illustrate the results for larger n by performing the experiment for the case $\epsilon = 1$ with the tolerance set to 10^{-4}. In this case, the numbers of iterations corresponding to $n = 32, 64, 128, 256, 512$ become 7, 7, 7, 7, 7 respectively. These numerical results agree with the theoretical results in Theorem 2.

In Table 3, we compare the number of floating point operations (flops) per PCG iteration for different types of preconditioners in Matlab. In order to exploit the fast sine transform (FST) and the fast Fourier transform (FFT) algorithms, we choose $n = 2^m - 1$ for the sine transform preconditioners and $n = 2^m$ for the circulant preconditioners. We note that C is a circulant block matrix of size $2n^2$-by-$2n^2$ with $2n$-by-$2n$ circulant block, see [12], and therefore the cost per PCG iteration for C is twice that for $\hat{s}(A)$. We observe from Table 3 that the cost per PCG step for the sine transform preconditioner is relatively more expensive than those for the MILU and MINV preconditioners. However, we find that the Matlab implementations of the FST and FFT algorithms are not optimal. For $5 \leq m \leq 8$, the Matlab implementations of the n-dimensional FST and FFT are about $9n \log n$ flops and $3.5n \log n$ flops respectively. On the other hand, the optimal implementations of both the FST and the FFT algorithms should be $2.5n \log n$ flops, see [2,14]. Thus, for our method, where the dominant cost is in taking the FST, there should be about $9/2.5 \approx 3.6$ times speed up if optimal implementation is used. Similarly, there is a factor of $3.5/2.5 \approx 1.4$ times speed up for the circulant preconditioning method. The cost per iteration of our preconditioner will therefore be comparable with those of the MINV and MILU preconditioners. In addition, we note that the FST is easier to parallelize than tridiagonal solvers. We also remark that after taking into account the optimal implementation of FST and FFT, we find that the

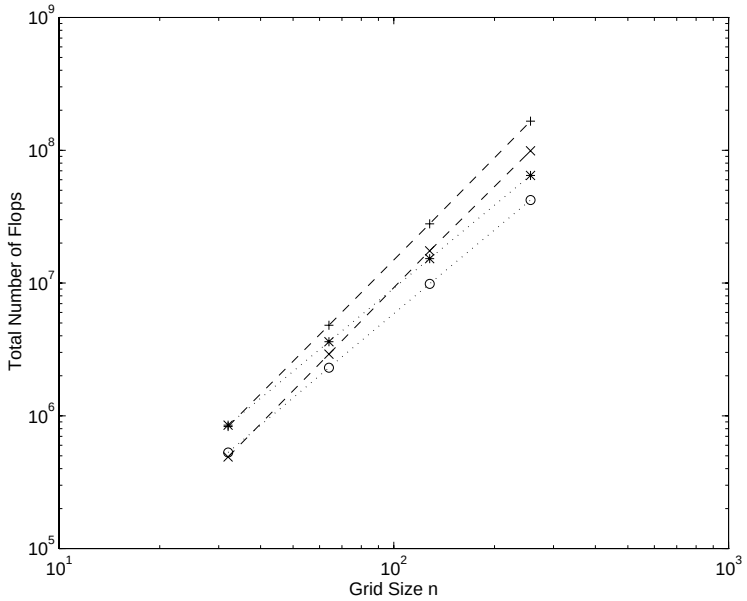

Figure 1(a). $\epsilon = 0.01$, $\circ$: $\hat{s}(\cdot)$, $\times$: MINV, $+$: MILU, $*$: C

circulant preconditioner C still costs twice that of the preconditioner $\hat{s}(A)$. This is due to the fact that computing $C^{-1}v$ by FFT involves solving a block tridiagonal system with a complex right-hand side vector while $\hat{s}(A)^{-1}v$ by FST only involves a real right-hand side vector. In terms of Matlab flops count, the circulant preconditioning method will require, in addition, roughly twice more computational work than that of our preconditioner.

In Figure 1(a), (b), we plot the total number of flops that are required for solving (5.1) by the PCG method against the grid size n. The counts are obtained by using Matlab's implementation of FST and FFT. We observe that the total number of flops for convergence, with the preconditioner $\hat{s}(\cdot)$, is always less than that of the circulant preconditioners C and grows with a rate slower than those of the MINV, MILU preconditioners. It is expected that the overall computational cost for solving (5.1) with the sine transform preconditioner will be cheaper than those of the other three preconditioners especially when n is large. As mentioned above, the cost can be further reduced by $2/3$ if optimal implementation of FST is used.

6. Concluding remarks

In this paper, we have developed preconditioners for Dirichlet problems based on a sine transform matrix. We find our preconditioner by looking for the best approximation in Frobenius norm amongst all matrices that can be diagonalized by the sine-transform matrix. This gives exact approximation to the Laplacian operator with Dirichlet boundary condition. We have since applied our idea to problems in queueing networks [5] and image processing [6] where the boundary conditions are Neumann in nature. Since the Lapla-

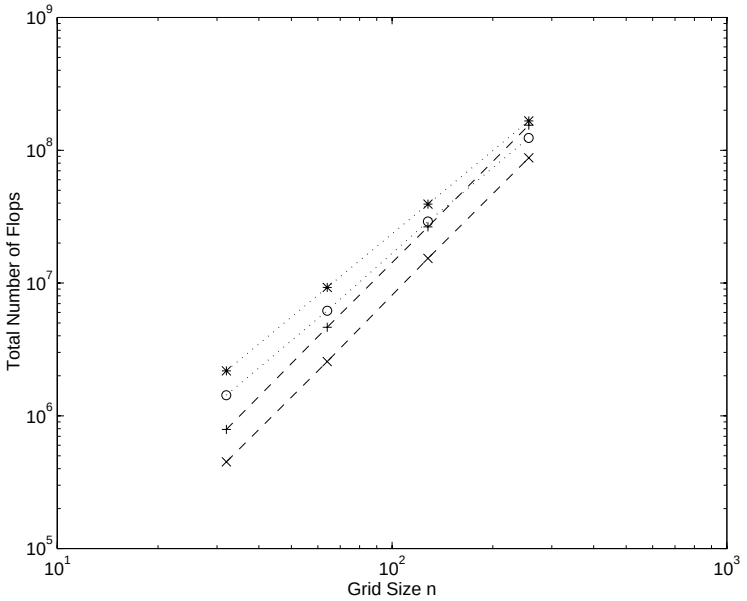

Figure 1(b). $\epsilon = 1$, $\circ$: $\hat{s}(\cdot)$, $\times$: MIN V, $+$: MILU, $*$: C

cian operator with Neumann boundary condition can be diagonalized by a cosine transform matrix, we construct our preconditioners in these cases by looking for the best approximation in Frobenius norm amongst all matrices that can be diagonalized by the cosine transform matrix. The numerical results there show great improvement over methods previously used.

Acknowledgements

The authors wish to thank Professors Tony Chan and Vassilevski for their helpful discussion in this work and also the referees for their valuable comments.

REFERENCES

1. O. Axelsson and V. Barker. *Finite Element Solution of Boundary Value Problems: Theory and Computation*. Academic Press, Orlando, FL, 1983.
2. G. Bergland. A fast Fourier transform algorithm for real-valued series. *Comm. ACM*, 11, 703–710, 1968.
3. E. Boman and I. Koltracht. Fast transform based preconditioners for Toeplitz equations. *SIAM J. Matrix Anal. Appl.*, 16, 628–645, 1995.
4. R. Chan and T. Chan. Circulant Preconditioners for elliptic problems. *J. Numer. Linear Algebra Appls.*, 1, 77–101, 1992.
5. R. Chan, W. Ching and C. Wong. Optimal trigonometric preconditioners for elliptic and queueing problems. *SEA Bull. Math.*, 1996, to appear.

6. R. Chan, T. Chan and C. Wong. Cosine transform based preconditioners for total variation minimization problems in image processing, in *Iterative Methods in Linear Algebra, II*, S. Margenov and P. Vassilevski, editors, pp. 311–329. IMACS Series in Computational and Applied Mathematics, V3, IMACS, NJ, 1996.
7. R. Chan and X. Jin. A family of block preconditioners for block systems. *SIAM J. Sci. Statist. Comput.*, 13, 1218–1235, 1992.
8. R. Chan, K. Ng and C. Wong. Sine transform based preconditioners for symmetric Toeplitz systems. *Linear Algebra Appls.*, 232, 237–260, 1996.
9. T. Chan and J. Olkin. Circulant preconditioners for Toeplitz-block matrices. *Numer. Algorithms*, 6, 89–101, 1994.
10. P. Concus, G. Golub and G. Meurant. Block preconditioning for the conjugate gradient method. *SIAM J. Sci. Statist. Comput.*, 6, 220–252, 1985.
11. T. Dupont, R. Kendall and H. Rachford. An approximate factorization procedure for solving self-adjoint elliptic difference equations. *SIAM J. Numer. Anal.*, 5, 559–573, 1968.
12. I. Lirkov, S. Margenov and P. Vassilevski. Circulant block-factorization preconditioners for elliptic problems. *Computing*, 53, 59–74, 1994.
13. T. Manteuffel and S. Parter. Preconditioning and boundary conditions. *SIAM J. Numer. Anal.*, 27, 656–694, 1990.
14. P. Yip and K. Rao. Fast decimation-in-time algorithms for a family of discrete sine and cosine transforms. *Circuits Systems Signal Process.*, 3, 387–408 1984.

