



The Australian National University

TR-CS-95-07

**Conjugate Gradient Methods for
Toeplitz Systems**

Raymond H. Chan and Michael K. Ng

September 1995

Joint Computer Science Technical Report Series

Department of Computer Science
Faculty of Engineering and Information Technology

Computer Sciences Laboratory
Research School of Information Sciences and Engineering

This technical report series is published jointly by the Department of Computer Science, Faculty of Engineering and Information Technology, and the Computer Sciences Laboratory, Research School of Information Sciences and Engineering, The Australian National University.

Please direct correspondence regarding this series to:

Technical Reports
Department of Computer Science
Faculty of Engineering and Information Technology
The Australian National University
Canberra ACT 0200
Australia

or send email to:

`Technical.Reports@cs.anu.edu.au`

A list of technical reports, including some abstracts and copies of some full reports may be found at:

<http://cs.anu.edu.au/techreports/>

Recent reports in this series:

- TR-CS-95-06 Oscar Bosman and Heinz W. Schmidt. *Object test coverage using finite state machines*. September 1995.
- TR-CS-95-05 Jeffrey X. Yu, Kian-Lee Tan, and Xun Qu. *On balancing workload in a highly mobile environment*. August 1995.
- TR-CS-95-04 Department of Computer Science. *Annual report 1994*. August 1995.
- TR-CS-95-03 Douglas R. Sweet and Richard P. Brent. *Error analysis of a partial pivoting method for structured matrices*. June 1995.
- TR-CS-95-02 Jeffrey X. Yu and Kian-Lee Tan. *Scheduling issues in partitioned temporal join*. May 1995.
- TR-CS-95-01 Craig Eldershaw and Richard P. Brent. *Factorization of large integers on some vector and parallel computers*. January 1995.

Conjugate Gradient Methods for Toeplitz Systems *

Raymond H. Chan[†] Michael K. Ng[‡]

25 September 1995

Abstract

In this expository paper, we survey some of the latest developments on using preconditioned conjugate gradient methods for solving Toeplitz systems. One of the main results is that the complexity of solving a large class of n -by- n Toeplitz systems is reduced to $O(n \log n)$ operations as compared to $O(n \log^2 n)$ operations required by fast direct Toeplitz solvers. Different preconditioners proposed for Toeplitz systems are reviewed. Applications to Toeplitz-related systems arising from partial differential equations, queueing networks, signal and image processing, integral equations, and time series analysis are given.

Key words. Toeplitz matrices, preconditioners, preconditioned conjugate gradient methods, differential equations, signal and image processing, time series, queueing problems, integral equations

AMS Subject classifications. 45E10, 62M10, 65F10, 65N22, 65P05, 68U10, 93E11

*To appear in SIAM Review.

[†]Department of Mathematics, The Chinese University of Hong Kong, Shatin, Hong Kong. Research supported in part by HKRGC grants no. CUHK 178/93E and CUHK 316/94E.

[‡]Computer Sciences Laboratory, The Australian National University, Canberra ACT 0200, Australia. Research supported by the Cooperative Research Centre for Advanced Computational Systems.

1 Introduction

1.1 Background

An n -by- n matrix A_n is said to be *Toeplitz* if

$$A_n = \begin{bmatrix} a_0 & a_{-1} & \cdots & a_{2-n} & a_{1-n} \\ a_1 & a_0 & a_{-1} & & a_{2-n} \\ \vdots & a_1 & a_0 & \ddots & \vdots \\ a_{n-2} & & \ddots & \ddots & a_{-1} \\ a_{n-1} & a_{n-2} & \cdots & a_1 & a_0 \end{bmatrix}, \quad (1.1)$$

i.e., A_n is constant along its diagonals. The name Toeplitz originates from the work of Otto Toeplitz [180] in the early 1900's on bilinear forms related to Laurent series, see Grenander and Szegő [91] for details. We are interested in solving the Toeplitz system $A_n \mathbf{x} = \mathbf{b}$.

Toeplitz systems arise in a variety of applications in mathematics and engineering. In signal processing, solutions of Toeplitz systems are required in order to obtain the filter coefficients in the design of recursive digital filters, see Chui and A. Chan [61] and Haykin [99, pp. 167–168]. Time series analysis involves solutions of Toeplitz systems for the unknown parameters of stationary auto-regressive models, see King et al. [125, pp. 368–379]. By using discrete time and spatial sampling of the domain and the sinc function $\sin(\pi x)/(\pi x)$ as basis function for approximating the initial data, the numerical solutions of inverse heat problems can be obtained by solving a Toeplitz system, see Gilliam et al. [83]. Other applications involve solutions of partial differential equations, solutions of convolution-type integral equations, Padé approximations, and minimum realization problems in control theory, see Bunch [23] and the references therein. These applications have motivated both mathematicians and engineers to develop specific algorithms catering to solving Toeplitz systems. We will call these algorithms *Toeplitz solvers*.

Most of the early works on Toeplitz solvers were focused on direct methods. A straightforward application of the Gaussian elimination method will result in an algorithm of $O(n^3)$ complexity. However, since n -by- n Toeplitz matrices are determined by only $(2n - 1)$ entries rather than n^2 entries, it is expected that the solution of Toeplitz systems can be obtained in less than $O(n^3)$ operations. There are a number of Toeplitz solvers that decrease the complexity to $O(n^2)$ operations, see, for instance, Levinson (1946) [134], Baxter (1961) [8], Trench (1964) [182], and Zohar (1974) [199]. These algorithms require the invertibility of the $(n - 1)$ -by- $(n - 1)$ principal submatrix of A_n . Around 1980, fast direct Toeplitz solvers of complexity $O(n \log^2 n)$ were developed, see, for instance, Brent, Gustavson, and Yun (1980) [20], Bitmead and Anderson (1980) [16], Morf (1980) [143], de Hoog (1987) [108], and Ammar and Gragg (1988) [4]. These algorithms require the invertibility of the $\lfloor n/2 \rfloor$ -by- $\lfloor n/2 \rfloor$ principal submatrix of A_n .

The stability properties of these direct methods for symmetric positive definite Toeplitz systems are discussed in Bunch [23]. It is noted that if A_n has a singular or ill-conditioned principal submatrix, then a *breakdown* or *near-breakdown* can occur in these algorithms. Such

breakdowns will cause numerical instabilities in subsequent steps of the algorithms and result in inaccurately computed solutions. The question of how to avoid breakdowns or near-breakdowns by skipping over singular submatrices or ill-conditioned submatrices has been studied extensively, and various such algorithms have been proposed, see [58, 70, 89, 93, 102, 157, 175, 198]. In particular, T. Chan and Hansen (1992) [58] were the first to derive a *look-ahead* variant of the Levinson algorithm. The basic idea is to relax the inverse triangular decomposition slightly and to compute an inverse block factorization of the Toeplitz matrices with a block diagonal matrix instead of a scalar diagonal matrix. Other look-ahead extensions of fast Toeplitz solvers can be found in [77, 80, 95].

Recent research on using the preconditioned conjugate gradient method as an iterative method for solving Toeplitz systems has brought much attention. One of the main important results of this methodology is that the complexity of solving a large class of Toeplitz systems can be reduced to $O(n \log n)$ operations as compared to the $O(n \log^2 n)$ operations required by fast direct Toeplitz solvers, provided that a suitable preconditioner is chosen under certain conditions on the Toeplitz operator. Besides the reduction of the arithmetic complexity, there are large classes of important Toeplitz matrices where the fast direct Toeplitz solvers are notoriously unstable, e.g., indefinite and certain non-Hermitian Toeplitz matrices. Therefore, iterative methods provide alternatives to solving these Toeplitz systems. In this paper, we will survey results for these *iterative Toeplitz solvers* and give some insight in how to construct effective preconditioners for them. Applications of these Toeplitz solvers to some practical problems will also be given.

1.2 Toeplitz Matrices and Circulant Matrices

Let us begin by introducing the notation that will be used throughout the paper. Let $\mathbf{C}_{2\pi}$ be the set of all 2π -periodic continuous real-valued functions defined on $[-\pi, \pi]$. For all $f \in \mathbf{C}_{2\pi}$, let

$$a_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(\theta) e^{-ik\theta} d\theta, \quad k = 0, \pm 1, \pm 2, \dots$$

be the Fourier coefficients of f . For all $n \geq 1$, let A_n be the n -by- n Toeplitz matrix with entries $a_{j,k} = a_{j-k}$, $0 \leq j, k < n$. The function f is called the *generating function* of the sequence of Toeplitz matrices A_n , see Grenander and Szegö [91]. Since f is a real-valued function, we have

$$a_{-k} = \bar{a}_k, \quad k = 0, \pm 1, \pm 2, \dots$$

It follows that A_n are Hermitian matrices. Note that when f is an even function, the matrices A_n are real symmetric. We emphasize that in practical applications, the functions f are readily available. Typical examples of generating functions are the kernels of Wiener-Hopf equations, see Gohberg and Fel'dman [85, p.82], the functions which give the amplitude characteristics of recursive digital filters, see Chui and A. Chan [61], the spectral density functions in stationary stochastic processes, see Grenander and Szegö [91, p.171], and the point-spread functions in image deblurring, see Jain [116, p.269].

We will solve the systems $A_n \mathbf{x} = \mathbf{b}$ by conjugate gradient methods. The convergence rate of the methods depends partly on how clustered the spectra of the sequence of matrices A_n are, see Axelsson and Barker [7, p.24]. The clustering of the spectra of a sequence of matrices is defined as follows:

Definition 1 A sequence of matrices $\{A_n\}_{n=1}^\infty$ is said to have clustered spectra around 1 if for any given $\epsilon > 0$, there exist positive integers n_1 and n_2 such that for all $n > n_1$, at most n_2 eigenvalues of the matrix $A_n - I_n$ have absolute value larger than ϵ .

For Toeplitz matrices, we note that there is a close relationship between the spectrum of A_n and its generating function f .

Theorem 1 (Grenander and Szegö [91, pp.63–65]) Let $f \in \mathbf{C}_{2\pi}$. Then the spectrum $\lambda(A_n)$ of A_n satisfies

$$\lambda(A_n) \subseteq [f_{\min}, f_{\max}], \quad \forall n \geq 1, \quad (1.2)$$

where f_{\min} and f_{\max} are the minimum and maximum values of f , respectively. Moreover, the eigenvalues $\lambda_j(A_n)$, $j = 0, 1, \dots, n-1$, are equally distributed as $f(2\pi j/n)$, i.e.,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=0}^{n-1} \left[g(\lambda_j(A_n)) - g\left(f\left(\frac{2\pi j}{n}\right)\right) \right] = 0 \quad (1.3)$$

for any continuous function g defined on $[-\pi, \pi]$.

The equal distribution of eigenvalues of Toeplitz matrices indicates that the eigenvalues will not be clustered in general. To illustrate this, consider the 1-dimensional discrete Laplacian matrix

$$A_n = \text{tridiag}[-1, 2, -1].$$

Its eigenvalues are given by

$$\lambda_j(A_n) = 4 \sin^2\left(\frac{\pi j}{n+1}\right), \quad 1 \leq j \leq n.$$

For $n = 32$, the eigenvalues of A_n are depicted in Figure 1.




Figure 1: Spectrum of 1-D discrete Laplacian: $A_{32} = \text{tridiag}[-1, 2, -1]$.

An n -by- n matrix C_n is said to be *circulant* if

$$C_n = \begin{bmatrix} c_0 & c_{-1} & \cdots & c_{2-n} & c_{1-n} \\ c_1 & c_0 & c_{-1} & & c_{2-n} \\ \vdots & c_1 & c_0 & \ddots & \vdots \\ c_{n-2} & & \ddots & \ddots & c_{-1} \\ c_{n-1} & c_{n-2} & \cdots & c_1 & c_0 \end{bmatrix},$$

where $c_{-k} = c_{n-k}$ for $1 \leq k \leq n-1$. Circulant matrices are diagonalized by the Fourier matrix F_n , i.e.,

$$C_n = F_n^* \Lambda_n F_n \quad (1.4)$$

where the entries of F_n are given by

$$[F_n]_{j,k} = \frac{1}{\sqrt{n}} e^{2\pi i j k / n}, \quad 0 \leq j, k \leq n-1,$$

and Λ_n is a diagonal matrix holding the eigenvalues of C_n , see for instance Davis [69, p.73]. We note that Λ_n can be obtained in $O(n \log n)$ operations by taking the fast Fourier transform (FFT) of the first column of C_n . For the fast Fourier transform algorithm, we refer to Cooley and Tukey [66]. In fact, the diagonal entries λ_k of Λ_n are given by

$$\lambda_k = \sum_{j=0}^{n-1} c_j e^{2\pi i j k / n}, \quad k = 0, \dots, n-1. \quad (1.5)$$

Once Λ_n is obtained, the products $C_n \mathbf{y}$ and $C_n^{-1} \mathbf{y}$ for any vector \mathbf{y} can be computed by FFTs in $O(n \log n)$ operations using (1.4).

1.3 The Conjugate Gradient Method for Toeplitz Matrices

The conjugate gradient method is an iterative method for solving Hermitian positive definite matrix systems. The algorithm of the method can be found in Golub and Van Loan [88, pp. 516–527]. In each iteration, it requires two inner products of n -vectors and one multiplication of the coefficient matrix with an n -vector. Storage for four temporary n -vectors is needed but there is no explicit storage required for the coefficient matrix.

Let $f \in \mathbf{C}_{2\pi}$. For simplicity, let us assume for the moment that f is positive, i.e., $f_{\min} > 0$. Then by (1.2), A_n are positive definite for all n . Consider applying the conjugate gradient method to solve these symmetric positive definite Toeplitz systems $A_n \mathbf{x} = \mathbf{b}$. In each iteration, besides the two inner products required, one matrix-vector multiplication $A_n \mathbf{y}$ is also needed. That can be computed by FFTs by first embedding A_n into a $2n$ -by- $2n$ circulant matrix, i.e.,

$$\begin{bmatrix} A_n & \times \\ \times & A_n \end{bmatrix} \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} A_n \mathbf{y} \\ \dagger \end{bmatrix}, \quad (1.6)$$

see Strang [172], and then carrying out the multiplication by using the decomposition as in (1.4). The matrix-vector multiplication thus requires $O(2n \log(2n))$ operations. It follows that the total number of operations per iteration is of $O(n \log n)$ operations. As for the storage required, besides the four temporary n -vectors, we need an extra $2n$ -vector for storing the eigenvalues of the embedded circulant matrix given in (1.6).

The convergence rate of the conjugate gradient method is well studied, see Axelsson and Barker [7, p.24]. It depends on the condition number of the matrix A_n and how clustered the spectrum of A_n is. If the spectrum is not clustered, as is usually the case for Toeplitz matrices (cf. Theorem 1), a good estimate of the convergence rate is given in terms of the largest and smallest eigenvalues of A_n . Using (1.2), this estimate can be expressed as

$$\frac{\|\mathbf{e}_q\|_{A_n}}{\|\mathbf{e}_0\|_{A_n}} < 2 \left(\frac{\sqrt{f_{\max}} - \sqrt{f_{\min}}}{\sqrt{f_{\max}} + \sqrt{f_{\min}}} \right)^q,$$

where \mathbf{e}_q is the error vector at the q th iteration and $\|\mathbf{x}\|_{A_n}^2 \equiv \mathbf{x}^* A_n \mathbf{x}$. This indicates that the rate of convergence is linear. Thus, the method will converge in a constant number of iterations, and hence the complexity of solving the Toeplitz system is $O(n \log n)$. However, we remark that if f_{\max}/f_{\min} is large, the constant in the operation count will be large and hence the convergence will be very slow.

One way to speed up the convergence rate of the method is to precondition the Toeplitz system. Thus, instead of solving $A_n \mathbf{x} = \mathbf{b}$, we solve the preconditioned system

$$P_n^{-1} A_n \mathbf{x} = P_n^{-1} \mathbf{b}. \quad (1.7)$$

The matrix P_n , called the preconditioner, should be chosen according to the following criteria:

- P_n should be constructed within $O(n \log n)$ operations.
- $P_n \mathbf{v} = \mathbf{y}$ should be solved in $O(n \log n)$ operations.
- The spectrum of $P_n^{-1} A_n$ should be clustered.

The first two criteria are to keep the operation count per iteration within $O(n \log n)$ as that is the count for the non-preconditioned system. The third criterion comes from the fact that the more clustered the eigenvalues are, the faster the convergence of the method will be, see for instance [139, pp. 249-251] and [7, pp. 27-28]. If $P_n^{-1} A_n$ has a clustered spectrum as defined in Definition 1, then the conjugate gradient method, when applied to solving the preconditioned system (1.7), converges superlinearly for large n , see [49]. More precisely, for any given $\epsilon > 0$, there exists a constant $c(\epsilon) > 0$ such that the error vector \mathbf{e}_q of the preconditioned conjugate gradient method at the q th iteration satisfies

$$\frac{\|\|\mathbf{e}_q\|\|}{\|\|\mathbf{e}_0\|\|} \leq c(\epsilon) \epsilon^q \quad (1.8)$$

where

$$\|\|\mathbf{v}\|\|^2 \equiv \mathbf{v}^* P_n^{-1/2} A_n P_n^{-1/2} \mathbf{v}.$$

The main aim of this paper is to review different preconditioners developed for Toeplitz systems that satisfy the three criteria mentioned earlier. We will also study applications to Toeplitz-related systems arising from partial differential equations, queueing networks, signal and image processing, integral equations, and time series analysis. For simplicity, we will drop the subscripts on matrices when their dimensions are apparent. The outline of the paper is as follows. In §2, we survey the use of circulant matrices as preconditioners for Toeplitz matrices. In §3, other useful and successful non-circulant preconditioners for Toeplitz matrices are also considered. Applications of preconditioned conjugate gradient methods for Toeplitz-related systems are discussed in §4. Finally, concluding remarks are given in §5.

2 Circulant Preconditioners for Toeplitz Systems

2.1 Circulant Preconditioners

In 1986, Strang [172] and Olkin [155] independently proposed the use of circulant matrices to precondition Toeplitz matrices in conjugate gradient iterations. Part of their motivation is to exploit the fast inversion of circulant matrices. Numerical results in [173, 155] suggest that the method converges very fast for a wide range of Toeplitz matrices. This has later been proved theoretically in [49] and in other papers for other circulant preconditioners. In this subsection, we will give a brief account of these developments.

With circulant matrices as preconditioners, in each iteration, we have to solve a circulant system. From (1.4), we see that circulant matrices can be diagonalized by discrete Fourier matrices, and hence the inversion of n -by- n circulant systems can be done in $O(n \log n)$ operations by using FFTs of size n . In contrast, by (1.6), we see that the cost of computing $A\mathbf{y}$, which is also required in each iteration whether the system is preconditioned or not, is done by using FFTs of size $2n$. Notice that if FFT is used to compute the discrete Fourier transform of a $2n$ -vector for which the even discrete Fourier transform components are already known, then the cost is the same as carrying out a length n FFT, see Linzer [136] for details. Thus, the cost per iteration of the circulant preconditioned conjugate gradient method is roughly 1.25 times of that required by the method without using preconditioners. We remark that Huckle [113] has also discussed different ways to reduce the number of FFTs in the iterative scheme, even when n is not a power of 2. In particular, he has proposed a way to compute $A\mathbf{y}$ such that the computational cost per iteration of the preconditioned system is nearly the same as that required by the non-preconditioned system.

We emphasize that the use of circulant matrices as preconditioners for Toeplitz systems allows the use of FFT throughout the computations; and FFT is highly parallelizable and has been implemented on multiprocessors efficiently [2, p.238] and [174]. Since conjugate gradient methods are easily parallelizable too [12, p.165], the circulant preconditioned conjugate gradient method is well-adapted for parallel computing.

We remark that circulant approximations to Toeplitz matrices have been considered and used for some time in signal processing (e.g. [158] and [171, pp.75–86]), time series analysis (e.g. [21, p.133] and [169]), and image processing (e.g. [13], [123, p.147], and [5, p.136]). However,

in these applications, the circulant approximations thus obtained were used to replace the given Toeplitz matrices in subsequent computations. In contrast, circulant approximations are used here only as preconditioners for Toeplitz systems and the solutions to the Toeplitz systems are unchanged. In the following, we review some successful circulant preconditioners proposed for Toeplitz matrices.

2.1.1 Strang's Preconditioner

The first circulant preconditioner is proposed by Strang [172] in 1986 and is defined to be the matrix that copies the central diagonals of A and reflects them around to complete the circulant requirement. For A given by (1.1), the diagonals s_j of the Strang preconditioner $S = [s_{k-\ell}]_{0 \leq k, \ell < n}$ are given by

$$s_j = \begin{cases} a_j, & 0 < j \leq \lfloor n/2 \rfloor, \\ a_{j-n}, & \lfloor n/2 \rfloor < j < n, \\ s_{n+j}, & 0 < -j < n. \end{cases} \quad (2.1)$$

One of the interesting properties of S is that S minimizes

$$\|C - A\|_1 \quad \text{and} \quad \|C - A\|_\infty$$

over all Hermitian circulant matrices C , see [27]. The spectra of these circulant preconditioned matrices have been analyzed by R. Chan and Strang [49].

Theorem 2 (R. Chan and Strang (1989) [49]) *Let f be an even positive function in the Wiener class, i.e., its Fourier coefficients are absolutely summable,*

$$\sum_{k=0}^{\infty} |a_k| < \infty.$$

Let A be generated by f . Then the spectra of $S^{-1}A$ are clustered around 1 for large n .

The main idea of their proof is to use an orthogonal transformation to transform $S - A$ into a Hankel matrix. Then Nehari's Theorem [62, p.120] is used to show that the limiting Hankel operator is compact. Using the theory of collectively compact sets of operators [6, pp.65–70], the spectra of the finite Hankel matrices are then shown to be clustered. However, this proof cannot be readily generalized to real-valued f , i.e., to Hermitian Toeplitz systems. Thus R. Chan in [27] developed a purely linear algebra technique to extend the results in Theorem 2. The approach is to decompose, for a given $\epsilon > 0$, the matrix $S - A$ into a sum of two matrices L and V , where $\text{rank } L \leq c(\epsilon)$ and $\|V\|_2 \leq \epsilon$ and then apply Cauchy's interlace theorem [88].

Theorem 3 (R. Chan (1989) [27]) *Let f be a positive function in the Wiener class. Then the spectra of $S^{-1}A$ are clustered around 1 for large n .*

Using standard error analysis of the conjugate gradient method, we can then show that the convergence rate of the method is superlinear, see (1.8). If extra smoothness conditions are imposed on the generating function f , we can get more precise estimates on how $\|\mathbf{e}_q\|$ in (1.8) goes to zero.

Theorem 4 (Trefethen (1990) [181], Ku and Kuo (1993) [130, 131]) *Suppose f is a rational function of the form $f = p/q$ where p and q are polynomials of degrees μ and ν , respectively. Then the number of outlying eigenvalues of $S^{-1}A$ is exactly equal to $2 \max\{\mu, \nu\}$. Hence, the method converges in at most $1 + 2 \max\{\mu, \nu\}$ steps for large n . If however $f(z) = \sum_{j=0}^{\infty} a_j z^j$ is only analytic in a neighborhood of $|z| = 1$, then there exist constants c and $0 \leq r < 1$ such that*

$$\frac{\|\mathbf{e}_{q+1}\|}{\|\mathbf{e}_0\|} \leq c^q r^{q^2/4+q/2}. \quad (2.2)$$

The idea of Trefethen's proof is to use rational approximation to bound the singular values of the Hankel matrix considered in the proof of Theorem 2. It follows from (2.2) that

$$\frac{\|\mathbf{e}_{q+1}\|}{\|\mathbf{e}_q\|} \approx cr^q \rightarrow 0.$$

For generating functions f with Fourier coefficients a_j decaying at a slower rate, we have the following two theorems.

Theorem 5 (R. Chan (1989) [27]) *Let f be a ν -times differentiable function with $f^{(\nu)} \in L^1[-\pi, \pi]$ where $\nu > 1$ (i.e. $|a_j| \leq \tilde{c}/j^{\nu+1}$ for some constant \tilde{c} .) Then there exists a constant c which depends only on f and ν , such that for large n ,*

$$\frac{\|\mathbf{e}_{2q}\|}{\|\mathbf{e}_0\|} \leq \frac{c^q}{((q-1)!)^{2\nu-2}}.$$

Theorem 5 was proved by using Cauchy's interlace theorem [88]. R. Chan and Yeung later used Jackson's theorem [60] in approximation theory to prove a stronger result than that in Theorem 5.

Theorem 6 (R. Chan and Yeung (1992) [55]) *Let f be a Lipschitz function of order ν , $0 < \nu \leq 1$, or f has a continuous ν th derivative, $\nu \geq 1$. Then there exists a constant c which depends only on f and ν , such that for large n ,*

$$\frac{\|\mathbf{e}_{2q}\|}{\|\mathbf{e}_0\|} \leq \prod_{k=2}^q \frac{c \log^2 k}{k^{2\nu}}.$$

Theorems 2–6 give the rate at which the error goes to zero in terms of the rate of decay of $|a_j|$. To see if solving Toeplitz systems by preconditioned conjugate gradient methods is more efficient than by fast direct Toeplitz solvers, Linzer [137] has performed tests for Toeplitz

matrices with different condition numbers and coefficients with different decaying rate. His results show that the iterative methods have the edge if $|a_j|$ decays like $O(j^{-0.5})$ or faster for matrices with condition number about 10, and the rate of decay required increases to $O(j^{-2})$ when the condition number is about 10^5 .

We also remark that Huckle [114] has recently compared the number of floating point operations for iterative methods with that of direct Toeplitz solvers and superfast Toeplitz solvers. He derived an upper bound of the number of PCG iterations such that the iterative method will be better than direct and superfast Toeplitz solvers. His finding shows that for positive definite Toeplitz systems, PCG methods are competitive for large matrices with small number of PCG iterations. In the indefinite or near-singular case, iterative methods may give a higher accuracy. In the unsymmetric case, only classical $O(n^2)$ direct Toeplitz solvers are available and therefore the iterative methods will have the edge if a good preconditioner can be found, see §2.3.

2.1.2 T. Chan's Preconditioner

For an n -by- n Toeplitz matrix A , T. Chan's circulant preconditioner $c(A)$ is defined to be the minimizer of

$$\|C - A\|_F \tag{2.3}$$

over all n -by- n circulant matrices C , see T. Chan (1988) [57]. Here $\|\cdot\|_F$ denotes the Frobenius norm. In [57], the matrix $c(A)$ is called an *optimal* circulant preconditioner because it minimizes (2.3). The j th diagonals of $c(A)$ are shown to be equal to

$$c_j = \begin{cases} \frac{(n-j)a_j + ja_{j-n}}{n}, & 0 \leq j < n, \\ c_{n+j}, & 0 < -j < n, \end{cases} \tag{2.4}$$

which are just the average of the diagonals of A , with the diagonals being extended to length n by a wrap-around. By using (1.5) and (2.4), we see that the eigenvalues $\lambda_k(c(A))$ of $c(A)$ are given by

$$\lambda_k(c(A)) = \sum_{j=-n+1}^{n-1} a_j \left(1 - \frac{|j|}{n}\right) e^{2\pi ijk/n}, \quad k = 0, \dots, n-1. \tag{2.5}$$

As for the performance of $c(A)$ as a preconditioner for A , R. Chan [28] proved that under the Wiener class assumptions in Theorem 3 (i.e., f is a positive function with absolutely summable Fourier coefficients), the spectra of $c(A)^{-1}A$ and $S^{-1}A$ are asymptotically the same as n tends to infinity, i.e., $\lim_{n \rightarrow \infty} \|c(A)^{-1}A - S^{-1}A\|_2 = 0$. Hence, $c(A)$ works as well for Wiener class functions as S does. Using Weierstrass' theorem to approximate 2π -periodic continuous generating functions by Wiener class functions, we have the following theorem.

Theorem 7 (R. Chan and Yeung (1992) [53]) *Let f be a positive function in $\mathbf{C}_{2\pi}$. Then the spectra of $c(A)^{-1}A$ are clustered around 1 for large n .*

However, the Weierstrass approach used in proving this theorem does not work for Strang's preconditioner. From Theorem 6, we see that the class of generating functions where Strang's preconditioner works is the class of 2π -periodic Lipschitz continuous functions. We will discuss this discrepancy in §2.2.

When A is not a Toeplitz matrix, the circulant minimizer $c(A)$ of (2.3) can still be obtained easily by taking the arithmetic average of the entries of A , i.e., its diagonals are given by

$$c_\ell = \frac{1}{n} \sum_{j-k=\ell(\bmod n)} a_{j,k}, \quad \ell = 0, \dots, n-1, \quad (2.6)$$

see [183]. Therefore, T. Chan's preconditioner is particularly useful in solving *non-Toeplitz* systems arising from the numerical solutions of elliptic partial differential equations [31] and Toeplitz least squares problems arising from signal and image processing [42, 43, 44, 59, 97, 152, 153, 159]. Convergence results for T. Chan's preconditioner have been established for these problems, see §4.

Another interesting spectral property of $c(A)$ is that if A is positive definite, then $c(A)$ is also positive definite. In fact, for an arbitrary Hermitian matrix A , we have

$$\lambda_{\min}(A) \leq \lambda_{\min}(c(A)) \leq \lambda_{\max}(c(A)) \leq \lambda_{\max}(A), \quad (2.7)$$

see Tyrtyshnikov [183] and R. Chan, Jin, and Yeung [38]. We remark that Strang's preconditioner does not satisfy (2.7) even for Toeplitz matrices, see R. Chan and Yeung [54]. In addition, R. Chan and Wong [51] recently proved that for some Toeplitz matrices A , T. Chan's preconditioner $c(A)$ minimizes $\kappa(C^{-1}A)$ over all n -by- n non-singular circulant matrices C .

Huckle in 1992 [111] proposed a preconditioner that is an extension of the T. Chan's preconditioner. Let $1 \leq p \leq n$. Huckle's circulant preconditioner H is defined to be the circulant matrix with eigenvalues

$$\lambda_k(H) = \sum_{j=-p+1}^{p-1} a_j \left(1 - \frac{|j|}{p}\right) e^{2\pi i j k / n}, \quad k = 0, \dots, n-1, \quad (2.8)$$

cf. (2.5). Thus, when $p = n$, it is the T. Chan preconditioner. Besides Wiener class functions, Huckle [110] has shown that H also works for generating functions with Fourier coefficients a_k that satisfy

$$\sum_{k=-\infty}^{\infty} |k| |a_k|^2 < \infty.$$

2.1.3 Preconditioners by Embedding

Let the Toeplitz matrix A be embedded into a $2n$ -by- $2n$ circulant matrix

$$\begin{bmatrix} A & B^* \\ B & A \end{bmatrix} \quad (2.9)$$

R. Chan's circulant preconditioner R is defined as $R = A + B$, see R. Chan (1989) [27]. We remark that T. Chan's and Huckle's preconditioners to A are just equal to R. Chan's preconditioner obtained from the Toeplitz matrix with diagonals $(1 - |j|/n)a_j$ and $\max\{0, 1 - |j|/p\}a_j$ respectively; i.e., the diagonals a_j of A are damped by the factor $(1 - |j|/n)$ and $\max\{0, 1 - |j|/p\}$ respectively, see (2.5) and (2.8).

Using the embedding (2.9), Ku and Kuo (1991) [128] constructed four different preconditioners K_i , $1 \leq i \leq 4$, based on different combinations of the matrices A and B . They are

$$K_1 = A + B; \quad K_2 = A - B; \quad K_3 = A + JB; \quad K_4 = A - JB$$

where J is the n -by- n anti-identity (reversal) matrix. We note that K_2 , K_3 , and K_4 are not circulant matrices.

2.1.4 Preconditioners by Minimization of Norms

Besides using the minimizer of $\|C - A\|_F$ as preconditioners for Toeplitz systems, minimizers of other approximations have also been proposed and used. For instance, Tyrtshnikov's circulant preconditioner T (1992) [183] is defined to be the minimizer of

$$\|I - C^{-1}A\|_F \tag{2.10}$$

over all non-singular circulant matrices C . In [183], T is called the *superoptimal* circulant preconditioner because it minimizes (2.10) instead of (2.3) and is shown to be equal to

$$T = c(AA^*)c(A)^{-1},$$

see also [38]. Tyrtshnikov [183] showed that for a general positive definite Toeplitz matrix A , T is also positive definite, cf. (2.7). Tismenetsky [179] and Sayed and Kailath [165] independently proposed the same preconditioner.

Huckle [112] has considered the minimizer M by minimization of

$$\|I - C^{-1/2}AC^{-1/2}\|_F,$$

over all non-singular circulant matrices C . The constructions of T and M require $O(n \log n)$ operations, see R. Chan, Jin, and Yeung [38], Tismenetsky [179], and Huckle [112], respectively.

Finally, we compare the performance of these preconditioners (R , $\{K_i\}_{i=1}^4$, T , and M) with Strang's and T. Chan's preconditioners. It has been proved in R. Chan [27], R. Chan, Jin, and Yeung [39], Ku and Kuo [128], and Huckle [112] that under the same Wiener class assumptions, these circulant preconditioned systems have spectra that are asymptotically the same as Strang's and T. Chan's circulant preconditioned systems. In particular, all these preconditioned systems converge at the same rate for large n .

In the following, we illustrate the effectiveness of circulant preconditioners for Toeplitz systems by a numerical example. We use the continuous generating function

$$f(\theta) = \theta^4 + 1, \quad -\pi \leq \theta \leq \pi,$$

in the test. Table 1 shows the numbers of iterations required to solve non-preconditioned systems $A\mathbf{x} = \mathbf{b}$ and circulant preconditioned systems $C^{-1}A\mathbf{x} = C^{-1}\mathbf{b}$ for different preconditioners. The right hand side vector \mathbf{b} is the vector of all ones. The zero vector is the initial guess. The stopping criterion is when the residual vector \mathbf{r}_q at the q th iteration satisfies $\|\mathbf{r}_q\|_2/\|\mathbf{r}_0\|_2 < 10^{-7}$. All computations are done by Matlab.

n	I	S	$c(A)$	R	K_2	H	T
16	8	8	8	6	6	8	8
32	20	8	7	5	5	10	16
64	37	6	7	5	5	7	18
128	56	5	6	5	5	7	13
256	67	5	6	5	5	6	10
512	70	5	6	5	5	6	8

Table 1: Number of iterations for Different Preconditioned Systems.

In the table, I denotes that no preconditioner is used and H is Huckle's preconditioner with $p = n/2$ (see (2.8)). We see from Table 1 that the number of iterations required for convergence for non-preconditioned systems is much greater than those for circulant preconditioned systems. Figure 2 depicts the spectra of the non-preconditioned matrix and the circulant preconditioned matrices for different circulant preconditioners. We note that the spectra of the circulant preconditioned matrices are indeed clustered around 1.

2.2 Circulant Preconditioners from Kernel Functions

A unifying approach of constructing circulant preconditioners is given in R. Chan and Yeung [54], where it is shown that most of the above mentioned circulant preconditioners can be derived by using the convolution products of some well-known kernels with the generating function f . For example, Strang's and T. Chan's circulant preconditioners are constructed by using the Dirichlet and Fejér kernels, respectively. To see this, we start by noting that the eigenvalues of Strang's preconditioner S are given by

$$\lambda_j(S) = (\hat{D}_{\lfloor \frac{n}{2} \rfloor} * f)\left(\frac{2\pi j}{n}\right), \quad 0 \leq j < n,$$

where the convolution of the Dirichlet kernel with f is given by

$$(\hat{D}_{\lfloor \frac{n}{2} \rfloor} * f)(\theta) \equiv \frac{1}{2\pi} \int_{-\pi}^{\pi} \hat{D}_{\lfloor \frac{n}{2} \rfloor}(\theta - \phi) f(\phi) d\phi$$

and

$$\hat{D}_k(\theta) = \frac{\sin(k + \frac{1}{2})\theta}{\sin \frac{1}{2}\theta}, \quad k = 1, 2, \dots$$

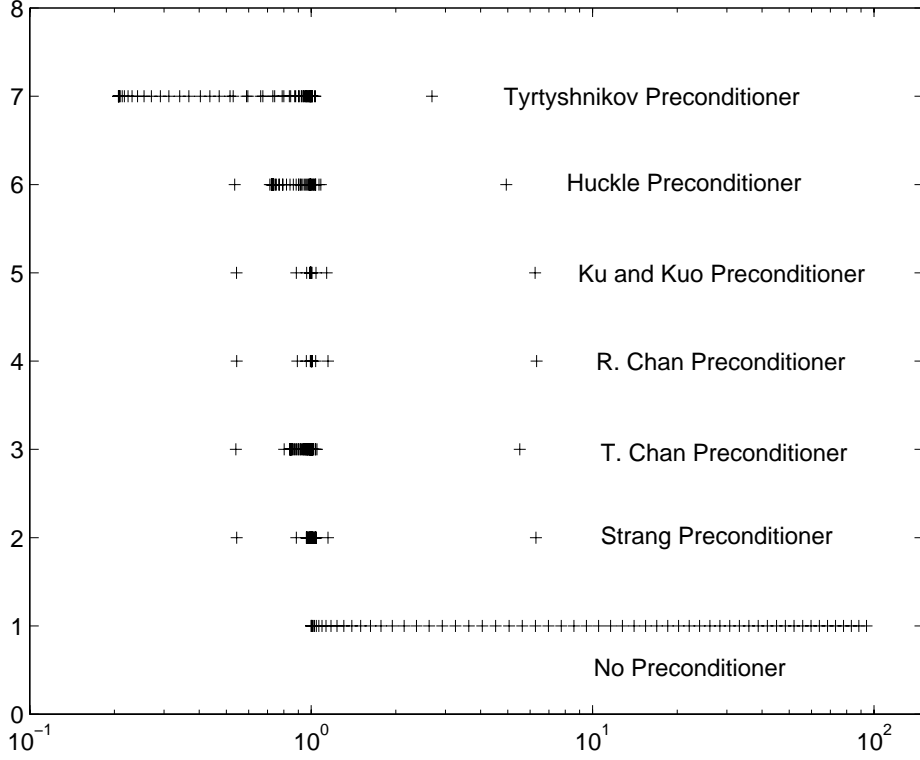


Figure 2: Spectra of Preconditioned Matrices for $n = 64$.

The eigenvalues of T. Chan's preconditioner $c(A)$ are given by

$$\lambda_j(c(A)) = (\hat{F}_n * f)\left(\frac{2\pi j}{n}\right), \quad 0 \leq j < n,$$

where the Fejér kernels are given by

$$\hat{F}_k(\theta) = \frac{1}{k} \left(\frac{\sin \frac{k}{2}\theta}{\sin \frac{1}{2}\theta} \right)^2, \quad k = 1, 2, \dots$$

Similarly, the eigenvalues of R. Chan's preconditioner R and Huckle's preconditioner H are given by

$$\lambda_j(R) = (\hat{D}_{n-1} * f)\left(\frac{2\pi j}{n}\right), \quad 0 \leq j < n$$

and

$$\lambda_j(H) = (\hat{F}_p * f)\left(\frac{2\pi j}{n}\right), \quad 0 \leq j < n,$$

respectively. The idea can be applied to design other circulant preconditioners C from kernels \hat{C}_n such as the von Hann kernel, Hamming kernel, and Bernstein kernel that are commonly used

in function theory [190], signal processing [96], and time series analysis [160]. In [54], several circulant preconditioners were constructed using this approach.

We remark that the convolution products of these kernels with f are just smooth approximations of the generating function f itself. R. Chan and Yeung proved that if the convolution product is a good approximation of f , then the correspondingly constructed circulant matrix will be a good preconditioner.

Theorem 8 (R. Chan and Yeung (1992) [54]) *Let $f \in \mathbf{C}_{2\pi}$ be positive. Let \hat{C}_n be a kernel such that $\hat{C}_n * f$ tends to f uniformly on $[-\pi, \pi]$. If C is the circulant matrix with eigenvalues given by*

$$\lambda_j(C) = (\hat{C}_n * f)\left(\frac{2\pi j}{n}\right), \quad 0 \leq j < n,$$

then the spectra of $C^{-1}A$ are clustered around 1 for large n .

We have mentioned in §2.1 that Strang's and T. Chan's preconditioners are fundamentally different, where Strang's preconditioners work for Lipschitz continuous functions (Theorem 6) and T. Chan's preconditioners work for 2π -periodic continuous functions (Theorem 7). That can be explained by the associations of Strang's preconditioner with the Dirichlet kernel and T. Chan's preconditioner with the Fejér kernel. It is well-known in Fourier analysis that if f is 2π -periodic continuous (or respectively Lipschitz continuous), then the convolution product of the Fejér kernel (or respectively the Dirichlet kernel) with f will converge to f uniformly on $[-\pi, \pi]$, see Walker [190, p.79, p.52].

In addition, it is interesting to note that for a piecewise continuous function f , the convolution product with the Fejér kernel will no longer converge to f uniformly on $[-\pi, \pi]$. Therefore, for generating functions that are only piecewise continuous, we don't expect the spectra of $c(A)^{-1}A$ to be clustered around 1.

Theorem 9 (Yeung and R. Chan (1993) [195]) *Let f be a non-negative piecewise continuous function on $[-\pi, \pi]$. Then for any given $\epsilon > 0$, the number of eigenvalues of $c(A)^{-1}A$ that lie outside the interval $(1 - \epsilon, 1 + \epsilon)$ is at least of $O(\log n)$ for n sufficiently large. If moreover f is strictly positive, then the number of outlying eigenvalues is exactly of $O(\log n)$.*

The theorem is established by noting that $A - c(A)$ is orthogonally similar to a Hankel matrix. Then Widom's theorem [192], which gives an estimate of the eigenvalues of the Hilbert matrices, is used to estimate the number of outlying eigenvalues of $A - c(A)$. Numerical examples are given in [195] to verify that the convergence rate of the method will no longer be superlinear in general. In fact, the numbers of iterations required for convergence do increase like $O(\log n)$. These results have recently been extended by Tyrtshnikov [184, 186]. In [184], he has established a generalized Szegő theorem that if f is in \mathbf{L}_2 , then the singular values of A_n is equally distributed (in a generalized sense) as $|f(x)|$ (cf. (1.3)). He then used the result to prove that if f is in \mathbf{L}_2 , then the number of outlying eigenvalues of the preconditioned system grows no more than $o(n)$. In [186], he further extended the results of circulant preconditioners for products of Toeplitz matrices.

2.3 Non-Hermitian Type Toeplitz Systems

In this subsection, we study Toeplitz matrices A generated by complex-valued functions. We note that such A are complex non-Hermitian matrices. In general the fast direct Toeplitz solvers are not applicable, and neither is the conjugate gradient method when applied to the system $A\mathbf{x} = \mathbf{b}$. For such matrices A , one can apply the conjugate gradient method to the normal equations $A^*A\mathbf{x} = A^*\mathbf{b}$. Another way of solving non-Hermitian Toeplitz systems is to employ some CG-like method [78] such as restarted GMRES [164] or TFQMR [76]. To speed up the CG-like methods, we can choose a matrix C such that the singular values of the preconditioned matrices $C^{-1}A$ are clustered.

Let us begin with skew-Hermitian type Toeplitz matrices, i.e., Toeplitz matrices of the form

$$A = \begin{bmatrix} a_0 & -\overline{a_1} & \cdots & -\overline{a_{n-2}} & -\overline{a_{n-1}} \\ a_1 & a_0 & -\overline{a_1} & & -\overline{a_{n-2}} \\ \vdots & a_1 & a_0 & \ddots & \vdots \\ a_{n-2} & & \ddots & \ddots & -\overline{a_1} \\ a_{n-1} & a_{n-2} & \cdots & a_1 & a_0 \end{bmatrix},$$

where a_0 is a real number. Obviously, $A = a_0I + A_S$, where I is the identity matrix and A_S is a skew-Hermitian Toeplitz matrix. These systems or low rank perturbations of such systems often appear in solving discretized hyperbolic differential equations, see Buckley [22] and Holmgren and Otto [104, 105].

In [35], R. Chan and Jin used R. Chan's circulant preconditioners R defined in §2.1.3 to precondition these skew-Hermitian Toeplitz matrices. Under the Wiener class assumptions on the entries of the first column of A , they proved that the singular values of $R^{-1}A$ are clustered around 1.

For Toeplitz matrices A generated by a complex-valued function, R. Chan and Yeung [56] have proved that if the generating function is 2π -periodic continuous with no zeros on $[-\pi, \pi]$, then the spectra of the iteration matrices $(c(A)^{-1}A)^*(c(A)^{-1}A)$ are clustered around 1. From that they showed that if the condition number $\kappa(A)$ of A is of $O(n^\alpha)$, $\alpha > 0$, then the number of iterations required for convergence is at most $O(\alpha \log n)$. Hence the total complexity for solving non-Hermitian type Toeplitz systems is of $O(n \log^2 n)$. When $\alpha = 0$, i.e., A is well-conditioned, the method converges in $O(1)$ steps and the complexity is reduced to $O(n \log n)$.

Numerical results in [56] shows that the requirements on f , namely that f has no zeros and $\kappa(A) = O(n^\alpha)$ are indispensable in order to get the said convergence rate. Moreover, these two conditions are mutually exclusive. For instance, if $f(\theta) = e^{i\theta}$, then f has no zeros on $[-\pi, \pi]$ but A is singular for all n . On the other hand, if $f(\theta) = 4 \sin^2 \theta$, then A is just the 1-dimensional discrete Laplacian with $\kappa(A) = O(n^2)$.

We remark that Ku and Kuo [130, 131, 132] have also considered solving non-symmetric Toeplitz matrix systems by the preconditioned conjugate gradient method. In their papers, A is assumed to be generated by complex-valued rational functions in the Wiener class, which happens to be a subclass of the class of 2π -periodic continuous functions considered in [56].

2.4 $\{\omega\}$ -Circulant Preconditioners

Circulant matrices belong to the class of $\{\omega\}$ -circulant matrices which are defined as follows:

Definition 2 Let $\omega = e^{i\theta_0}$ with $\theta_0 \in [-\pi, \pi]$. An n -by- n matrix W is said to be an $\{\omega\}$ -circulant matrix if it has the spectral decomposition

$$W = \Omega^* F^* \Lambda F \Omega.$$

Here $\Omega = \text{diag} [1, \omega^{-1/n}, \dots, \omega^{-(n-1)/n}]$ and Λ is a diagonal matrix containing the eigenvalues of W .

Notice that $\{\omega\}$ -circulant matrices are Toeplitz matrices with the first entry of each row obtained by multiplying the last entry of the preceding row by ω . In particular, $\{1\}$ -circulant matrices are circulant matrices, while $\{-1\}$ -circulant matrices are skew-circulant matrices. Huckle [111] and R. Chan and Jin [35] have used skew-circulant matrices as preconditioners for Toeplitz matrices and proved that under the Wiener class assumptions, the spectra of these preconditioned matrices are clustered around 1. Performances of general $\{\omega\}$ -circulant matrices as preconditioners for Toeplitz matrices are discussed in R. Chan and Ng [46] and Huckle [113].

2.5 General Remarks on Circulant Preconditioners

In this section, we have discussed many different kinds of circulant preconditioners. We note from Theorem 8 that most of them can be derived from the convolution approach. Moreover, the theorem changes the problem of finding a preconditioner to a problem in approximation theory. In particular, using results in approximation theory, the theorem can give us a guideline as to which preconditioner is better for a given generating function. For example, for 2π -periodic continuous functions that are not in the Wiener class, T. Chan's preconditioner is better than the Strang preconditioner. Also if we use a positive kernel to construct the preconditioner, then the preconditioner retains positive definiteness of the given Toeplitz matrix.

We emphasize that the assumptions on the generating functions f in the theorems in this section are to simplify the arguments. The main thing required in the proof is not an explicit form of f but a bound on the rate of decay of the diagonals $\{a_j\}_{j=0}^{\infty}$, see the definition of Wiener class functions and also the statement of Theorem 5. Thus, for the circulant preconditioning methods to work, there is no need to know the exact form of f , but just an estimate of the decay rate of a_j . Along this line, we remark further that Zygmund [200, p.183] has shown that if the diagonals a_j are convex, i.e., the second-order differences $a_{j+1} - 2a_j + a_{j-1} \geq 0$ for all j , then f is non-negative. Moreover, if one of the second-order differences is positive, then f is positive.

3 Non-Circulant Preconditioners for Toeplitz Systems

3.1 Dense Toeplitz Preconditioners

As alternatives to circulant preconditioners, Toeplitz matrices have also been proposed and analyzed as preconditioners for Toeplitz systems. We recall that Toeplitz matrix-vector products

can be computed in $O(n \log n)$ operations by using FFTs. Thus, Toeplitz matrices themselves will be good candidates for preconditioners.

In [46], R. Chan and Ng used the Toeplitz matrix \tilde{A} generated by $1/f$ to approximate the inverse of the Toeplitz matrix A generated by f , i.e., the preconditioned matrix is $\tilde{A}A$. We remark that the inverse of a Toeplitz matrix is non-Toeplitz in general, but is closely related to Toeplitz matrices, see Friedlander et al. [81]. It has been proved in [46] that the spectrum of the preconditioned matrix $\tilde{A}A$ is clustered around 1. However, in general it may be difficult to compute the Fourier coefficients of $1/f$ explicitly, and hence \tilde{A} cannot be formed efficiently. R. Chan and Ng [46] thus have derived families of Toeplitz preconditioners $P^{(s)}$ by using different kernel functions mentioned in §2.2 and different levels of approximation for the Fourier coefficients of $1/f$. For the first level of approximation, $s = 1$, $P^{(1)}$ is the circulant preconditioner mentioned in §2.2, depending on which kernel function is used. For integers $s > 1$, the preconditioner $P^{(s)}$ thus constructed can be written as a sum of $\{\omega\}$ -circulant matrices. More precisely,

$$P^{(s)} = \sum_{t=0}^{s-1} W^{(t)},$$

where $W^{(t)}$ are $\{e^{-2\pi it/s}\}$ -circulant matrices.

Using this fact, we showed that given any Toeplitz matrix A and integer $s > 1$, we can decompose A as

$$A = \sum_{t=0}^{s-1} U^{(t)},$$

where $U^{(t)}$ are $\{\omega\}$ -circulant matrices. For $s = 2$, the formula just states that any Toeplitz matrices can be written as the sum of a circulant matrix and a skew-circulant matrix, a fact first discovered by Pustyl'nikov [161]. Using the decomposition, we showed further that if all $U^{(t)}$ are invertible, then

$$P^{(s)} = \sum_{t=0}^{s-1} (U^{(t)})^{-1}.$$

We recall that in the additive Schwarz method for elliptic problems, a matrix A is first decomposed into the sum of individual projection matrices,

$$A = \sum_{t=0}^{s-1} A^{(t)},$$

and then the generalized inverses of these matrices are added back together to form a preconditioner P for the original matrix A , i.e.,

$$P = \sum_{t=0}^{s-1} A^{(t)+},$$

see Dryja and Widlund [72]. In this respect, our construction of the Toeplitz preconditioner $P^{(s)}$ is very similar to the approach used in the additive Schwarz-type preconditioners.

As for the convergence rate, R. Chan and Ng [46] proved that these preconditioned matrices $P^{(s)}A$ have clustered spectra around 1, and numerical results show that under the Wiener class assumption on the Toeplitz matrices, this method converges faster than those preconditioned by circulant preconditioners.

Recall that in each iteration of the preconditioned conjugate gradient method, we require the matrix-vector product $P^{(s)}\mathbf{y}$. Since $P^{(s)}$ is a Toeplitz matrix, the product can be computed by using FFTs of size $2n$ in $O(2n \log(2n))$ operations. As circulant matrix-vector multiplication can be done by using FFTs of size n , the cost per iteration of this method is roughly 4/3 times of that required by circulant preconditioned systems.

Another way of looking at these Toeplitz preconditioners is by embedding. For a given n -by- n Toeplitz matrix A , we embed A into a sn -by- sn circulant matrix C_{sn} with its first column given by

$$[C_{sn}]_{k,0} = \begin{cases} a_k & 0 < k < n, \\ 0 & n \leq k \leq sn - n, \\ a_{k-sn} & sn - n < k < sn. \end{cases}$$

The Toeplitz preconditioner $P^{(s)}$ is just equal to the leading n -by- n principal submatrix of C_{sn}^{-1} .

Recently, Hanke and Nagy [98] independently considered this approach of constructing Toeplitz preconditioners for band-Toeplitz matrices with bandwidth $2\beta + 1$. The given band-Toeplitz matrix is first embedded into an $(n + \beta)$ -by- $(n + \beta)$ circulant matrix $C_{n+\beta}$ by dragging down the diagonals to fill the northeast and southwest corners of the expanded matrix. Then the inverse of $C_{n+\beta}$ is partitioned as

$$C_{n+\beta}^{-1} = \begin{bmatrix} B^{(1)} & B^{(2)} \\ B^{(3)} & B^{(4)} \end{bmatrix}$$

where $B^{(1)}$ is a square matrix of size n . We note that since $C_{n+\beta}^{-1}$ is circulant, $B^{(1)}$ is a Toeplitz matrix. Hanke and Nagy's preconditioned system is defined to be $B^{(1)}A$.

In another development, Linzer [137] proposed using the Schur complement

$$B = B^{(1)} - B^{(2)}(B^{(4)})^{-1}B^{(3)}$$

rather than $B^{(1)}$ in the preconditioning of band-Toeplitz matrices. We note that B is not a Toeplitz matrix in general. However, since a circulant matrix can be diagonalized by the discrete Fourier transform matrix, it follows that $B\mathbf{y}$ can be computed by using FFTs of size $(n + \beta)$ and by solving the β -by- β Toeplitz system $B^{(4)}\mathbf{z} = \mathbf{v}$. Linzer [137] showed that B performs better than $B^{(1)}$ theoretically and numerically. We remark that Schur's complement formula was used by Jain [117], Morf [143], and Bitmead and Anderson [16] to develop direct methods for solving Toeplitz systems in signal processing.

3.2 Optimal Transform Based Preconditioners

From (1.4), we see that circulant matrices are precisely those matrices that can be diagonalized by the discrete Fourier transform, a transform which has a fast algorithm for its computations.

However, there are other transforms (for instance, the Hartley transform and the sine and cosine transforms) with fast algorithms too, see [188]. It is therefore natural to consider using these fast transforms to construct new classes of preconditioners for solving Toeplitz systems.

Recall that the optimal circulant preconditioners $c(A)$, which are the minimizers of $\|C - A\|_F$ over all circulant matrices C , are good preconditioners for Toeplitz systems, see §2.1.2. One may therefore consider preconditioners that are minimizers of $\|Q - A\|_F$ over a set of matrices Q that can be diagonalized by a fast discrete transform matrix Ψ . According to the terminology used in T. Chan [57] for optimal circulant preconditioners (where the preconditioner is based on the fast Fourier transform), we call these minimizers the *optimal transform based* preconditioners.

As the Frobenius norm is a unitary-invariant norm, the minimum of $\|Q - A\|_F$ over all Q of the form $Q = \Psi\Lambda\Psi^*$, Λ a diagonal matrix, is attained at $\Psi\Delta\Psi^*$. Here Δ is a diagonal matrix with diagonal entries

$$\Delta_{j,j} = [\Psi^* A \Psi]_{j,j}, \quad j = 1, \dots, n. \quad (3.1)$$

For Toeplitz matrices, (3.1) can be computed directly for optimal transform based preconditioners. However, computing Δ using (3.1) is costly even when the matrix-vector product $\Psi\mathbf{y}$ can be done efficiently. We emphasize that to construct T. Chan's minimizers economically, we exploit the fact that the class of circulant matrices has a very nice basis, namely the shift operator (i.e. the circulant matrix with $[0, \dots, 0, 1]$ as its first row) and its powers. Given another transform matrix Ψ , in order to construct its minimizer efficiently, we need to find matrices having special algebraic structures to characterize all matrices that can be diagonalized by Ψ . This is the crucial step in finding a fast algorithm for obtaining the minimizers.

In the following, optimal sine transform based and optimal Hartley transform based preconditioners for symmetric Toeplitz matrices A are considered. The constructions of such preconditioners require $O(n)$ operations for Toeplitz matrices, the same count as that for the optimal circulant preconditioner $c(A)$. Similar to T. Chan's circulant preconditioner, these optimal transform based preconditioners are also defined for arbitrary matrices. In general, the construction of such optimal approximations for any given n -by- n matrix is of $O(n^2)$ operations, see R. Chan, Ng, and Wong [48].

3.2.1 Optimal Sine Transform Based Preconditioner

The (j, k) entry of the n -by- n discrete sine transform matrix $\Psi^{(s)}$ is given by

$$\sqrt{\frac{2}{n+1}} \sin\left(\frac{\pi jk}{n+1}\right), \quad 1 \leq j, k \leq n. \quad (3.2)$$

For any n -vector \mathbf{v} , the matrix-vector product $\Psi^{(s)}\mathbf{v}$ can be done in $O(n \log n)$ operations by the fast sine transforms, see for instance Yip and Rao [196]. We define the optimal sine transform based preconditioner $s(A)$ to be the minimizer of $\|Q - A\|_F$ over the set of matrices Q that can be diagonalized by $\Psi^{(s)}$.

Let \mathbf{S} be the vector space over \mathbb{R} containing all n -by- n matrices that can be diagonalized by the discrete sine transform matrix $\Psi^{(s)}$, i.e.

$$\mathbf{S} = \left\{ \Psi^{(s)} \Lambda \Psi^{(s)} \mid \Lambda \text{ is a diagonal matrix} \right\}.$$

Boman and Koltracht [19], Bini and Di Benedetto [15], and Huckle [114] independently proved that a matrix belongs to \mathbf{S} if and only if the matrix can be expressed as a special sum of a Toeplitz matrix and a Hankel matrix. The idea of their proof is to exhibit a basis for \mathbf{S} with each element in the basis being a sparse matrix. The following theorem gives the basis that Boman and Koltracht have derived.

Theorem 10 (Boman and Koltracht (1995) [19]) *Let Z_i , $i = 1, \dots, n$, be n -by- n matrices with the (h, k) entry given by*

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

Then $\{Z_i\}_{i=1}^n$ is a basis for \mathbf{S} .

Thus, the vector space \mathbf{S} can be identified as follows.

Theorem 11 (Boman and Koltracht (1995) [19], Bini and Di Benedetto (1990) [15], and Huckle (1994) [114]) *Any matrix Q in \mathbf{S} can be written as $Q = X - Y$, where X is a symmetric Toeplitz matrix with first column $\mathbf{x} = [x_1, x_2, \dots, x_n]^T$, and Y is a Hankel matrix with first column $[0, 0, x_n, \dots, x_3]^T$ and last column $[x_3, \dots, x_n, 0, 0]^T$.*

Similar to the case of circulant matrices (cf. (1.5)), there is a relationship between the first column of matrices Q in \mathbf{S} and their eigenvalues. For any Q in \mathbf{S} , its eigenvalues can be obtained by multiplying the first column of Q by $\Psi^{(s)}$, which can be done in $O(n \log n)$ operations by fast sine transforms. Hence, any matrix in \mathbf{S} is determined by its first column. The following theorem gives the explicit formula for the entries of the first column of $s(A)$.

Theorem 12 (R. Chan, Ng, and Wong [48]) *Let A be an n -by- n symmetric Toeplitz matrix. Let $[s_1, s_2, \dots, s_n]^T$ be the first column of $s(A)$. Then*

$$s_k = \begin{cases} a_0 - \left(\frac{n-2}{n+1}\right) a_2, & k = 1, \\ a_1 - \left(\frac{n-3}{n+1}\right) a_3, & k = 2, \\ \left(\frac{n-k+3}{n+1}\right) a_{k-1} - \left(\frac{n-k-1}{n+1}\right) a_{k+1}, & k = 3, 4, \dots, n-2, \\ \left(\frac{4}{n+1}\right) a_{n-2} & k = n-1, \\ \left(\frac{3}{n+1}\right) a_{n-1} & k = n. \end{cases}$$

We see that the cost of constructing the minimizer $s(A)$ is $O(n)$ operations.

3.2.2 Optimal Hartley Transform Based Preconditioner

The (j, k) entry of the n -by- n discrete Hartley transform matrix $\Psi^{(h)}$ is given by

$$\frac{1}{\sqrt{n}} \cos\left(\frac{2\pi jk}{n}\right) + \frac{1}{\sqrt{n}} \sin\left(\frac{2\pi jk}{n}\right), \quad 0 \leq j, k \leq n-1.$$

The discrete Hartley transform of any n -vector can be computed in $O(n \log n)$ real operations. Bini and Favati [14] first characterized the class of matrices that can be diagonalized by the discrete Hartley transform, and showed that the class is the set of matrices that can be expressed as a special sum of a circulant matrix and a Hankel matrix.

Theorem 13 (Bini and Favati (1993) [14]) *Any matrix Q that can be diagonalized by $\Psi^{(h)}$ can be expressed as*

$$Q = W + XY, \tag{3.3}$$

where W is a circulant matrix, Y is a skew-circulant matrix with the first entry of its first column being zero and

$$X = \begin{bmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & & \ddots & \ddots & 1 \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \ddots & \ddots & \ddots & \vdots \\ 0 & 1 & 0 & \cdots & 0 \end{bmatrix}.$$

Using Theorem 13, Bini and Favati [14] then determined the minimizer $h(A)$ of $\|Q - A\|_F$ over all matrices that can be diagonalized by $\Psi^{(h)}$. As before the cost of constructing $h(A)$ is $O(n)$ operations and its eigenvalues can be computed in $O(n \log n)$ operations by using fast Hartley transforms.

Theorem 14 (Bini and Favati (1993) [14]) *Let A be a symmetric Toeplitz matrix. Let \mathbf{w} and \mathbf{y} be the first columns of W and Y of $h(A)$ as defined in (3.3). Then we have*

$$w_k = \frac{(n-k)a_k + ka_{n-k}}{n}, \quad k = 0, 1, \dots, n-1$$

and

$$y_k = \begin{cases} 0, & k = 0, \\ \frac{a_k - a_{n-k}}{n}, & k = 1, 2, \dots, n-1. \end{cases}$$

We note that the optimal circulant preconditioner $c(A)$ is just equal to the circulant part of $h(A)$. As for how good the optimal sine transform based and Hartley transform based preconditioners are as preconditioners for Toeplitz systems, R. Chan, Ng, and Wong [48], Bini and Favati [14], Bini and Di Benedetto [15], and Jin [119] proved that they both have the same convergence properties as the optimal circulant preconditioners. More precisely, if a given Toeplitz matrix A is generated by a 2π -periodic positive continuous function, then both the spectra of $s(A)^{-1}A$ and $h(A)^{-1}A$ are clustered around 1. However, numerical results given in R. Chan, Ng, and Wong [48], Bini and Favati [14], and Jin [119] show that the convergence performance of these new transform based preconditioners are better in terms of the number of iterations than that of the optimal circulant preconditioner.

In addition, Boman and Koltracht [19] have constructed bases for the vector spaces of matrices that can be diagonalized by some common fast transform matrices in engineering. These transform matrices include different cosine transform matrices and another sine transform matrix defined in [188]. Thus, one can construct optimal transform based preconditioners corresponding to these transform matrices. Boman and Koltracht [19] and R. Chan, Ching, and Wong [33] have studied optimal cosine transform based preconditioners. The cost of construction is the same as those of $s(A)$ and $h(A)$.

3.3 Band-Toeplitz Preconditioners

In this subsection, we consider Toeplitz matrices A generated by non-negative 2π -periodic real-valued functions. We first recall that a function f is said to have a ν th order zero at θ_0 if $f(\theta_0) = 0$ and ν is the smallest positive integer such that $f^{(\nu)}(\theta_0) \neq 0$ and $f^{(\nu+1)}(\theta)$ is continuous in a neighborhood of θ_0 . With the knowledge of the order of f at its minimum, we can give a better estimate of the spectrum of A than that in (1.2).

Theorem 15 (R. Chan (1991) [29]) *Suppose that $f(\theta) - f_{\min}$ has a unique zero of order 2ν at $\theta = \theta_0$. Then for all $n > 0$, we have*

$$\lambda_{\min}(A) \leq d_1 f_{\min} + d_2 n^{-2\nu},$$

and

$$\kappa(A) \geq \frac{d_3 n^{2\nu}}{d_4 + f_{\min} n^{2\nu}},$$

where $\{d_i\}_{i=1}^4$ are some constants independent of n .

Thus when $f_{\min} = 0$, the condition number of A is not uniformly bounded and the Toeplitz matrix A is ill-conditioned. Tyrtyshnikov has proved theoretically [185] that Strang's and T. Chan's preconditioners will fail in this case. In fact, he showed that the numbers of outlying eigenvalues of $S^{-1}A$ and $c(A)^{-1}A$ are of $O(n^{\nu/(\nu+\mu)})$ and $O(n^{\nu/(\nu+1)})$, respectively. Here, μ is the degree of smoothness of the function f , and ν is the order of f at the zeros. These results were numerically verified in Tyrtyshnikov and Strela [187].

Instead of finding other possible circulant preconditioners, R. Chan [29] resorted to using band-Toeplitz matrices as preconditioners. The motivation behind using band-Toeplitz matrices is to approximate the generating function f by trigonometric polynomials of fixed degree rather than by convolution products of f with some kernels. The advantage here is that trigonometric polynomials can be chosen to match the zeros of f , so that the preconditioned method still works when f has zeros.

Theorem 16 (R. Chan (1991) [29] and R. Chan and Ng (1993) [45]) *Let f be a non-negative piecewise continuous real-valued function defined on $[-\pi, \pi]$. Suppose that $f(\theta) - f_{\min}$ has a unique zero of order 2ν at $\theta = \theta_0$. Let B be the Toeplitz matrix generated by the function*

$$b_\nu(\theta) = [2 - 2 \cos(\theta - \theta_0)]^\nu + f_{\min}. \quad (3.4)$$

Then $\kappa(B^{-1}A)$ is uniformly bounded for all $n > 0$.

We note that B is a band matrix with bandwidth $2\nu + 1$ and its diagonals can be obtained by using Pascal's triangle. The band system $B\mathbf{y} = \mathbf{z}$ can be solved by using any band matrix solver, see Golub and Van Loan [88], or Wright [193] for a parallel one. The cost of factorizing B is about $\frac{1}{2}\nu^2 n$ operations, and then each subsequent solve requires an extra $(2\nu + 1)n$ operations. Hence, the total number of operations per iteration is of $O(n \log n)$ as ν is independent of n .

When $f_{\min} = 0$, the band preconditioner has improved the condition number from $\kappa(A) = O(n^{2\nu})$ to $\kappa(B^{-1}A) = O(1)$. Since the number of iterations required to attain a given tolerance ϵ is bounded by

$$\frac{1}{2} \sqrt{\kappa(B^{-1}A)} \log\left(\frac{2}{\epsilon}\right) + 1,$$

see for instance [7, p.26], the overall work required to attain the given tolerance is reduced from $O(n^{\nu+1} \log n)$ to $O(n \log n)$ operations. As for the storage, we just need an n -by- $(2\nu + 1)$ matrix to hold the factors of the preconditioner B . Thus, the overall storage requirement in the conjugate gradient method is about $(8 + \nu)n$. Finally, we remark that similar results hold when there are multiple points on $[-\pi, \pi]$ where f takes on its minimum value, see R. Chan [29].

The main drawback of using these band-Toeplitz matrices as preconditioners is that when f is positive, these preconditioned systems converge much slower than those preconditioned

by circulant preconditioners. Therefore, R. Chan and Tang [50] designed other kinds of band-Toeplitz preconditioners such that their preconditioned systems converge at the same rate as the circulant preconditioned systems even when f is positive. Their idea is to increase the band-width of the band-Toeplitz preconditioner to get extra degrees of freedom, which enable them not only to match the zeros in f , but also to minimize the relative error $\|(f - g)/f\|_\infty$ in approximating f by trigonometric polynomials g . The minimizer, which is a trigonometric polynomial, is found by a version of the Remez algorithm proposed by Tang [177].

Theorem 17 (R. Chan and Tang (1994) [50]) *Let f be the generating function of A and g_ℓ be the minimizer of $\|(f - g)/f\|_\infty$ over all trigonometric polynomials of degree ℓ . If*

$$\left\| \frac{f - g_\ell}{f} \right\|_\infty = \alpha < 1,$$

then the Toeplitz matrix B_ℓ generated by g_ℓ is positive definite and

$$\kappa(B_\ell^{-1}A) \leq \frac{1 + \alpha}{1 - \alpha}, \quad n = 1, 2, 3, \dots$$

The parameter α is given explicitly in the Remez algorithm. It gives an a priori bound on the number of iterations required for convergence.

The main idea behind Theorems 16 and 17 is to approximate the given non-negative generating function f by trigonometric polynomials that match the zeros of f . Clearly, any function g that matches the zeros of f and gives rise to Toeplitz matrices that are easily invertible can be considered too. This idea is exploited in Di Benedetto [9], Di Benedetto, Fiorentino, and Serra [11], and Serra [168]. In [11], f is first approximated by b_ν as in (3.4), then the quotient f/b_ν is further approximated by a trigonometric polynomial or rational function to enhance the convergence rate. In [168], $|f|$ is used to generate a Toeplitz preconditioner for the Toeplitz matrices generated by non-definite functions f .

From the above discussion, we see that if $f_{\min} = 0$, then the circulant preconditioners fail because they cannot match the zeros of the given f , while the band-Toeplitz preconditioners give only linear convergence because the preconditioned matrices do not have clustered spectra. In R. Chan and Ching [32], they considered using products of circulant matrices and band-Toeplitz matrices as preconditioners for Toeplitz systems generated by non-negative functions. The band-Toeplitz part of these *circulant-Toeplitz preconditioners* is to match the zeros of the given function, and the circulant part is to speed up the convergence rate of the algorithm. Instead of using powers of $2 - 2 \cos \theta$ as in (3.4) to generate the band-Toeplitz part of the preconditioner, they considered using powers of $1 - e^{i\theta}$ instead. This results in preconditioners that can handle complex-valued generating functions with zeros of arbitrary orders. We remark that Freund and Huckle [79] also considered using the band Toeplitz and circulant preconditioning techniques via displacement-based formula to solve these Toeplitz systems without the knowledge of the underlying generating function.

Another approach to handle ill-conditioned Toeplitz matrices A is developed in Concus and Saylor [67]. The Trench algorithm [182] is performed on A . When breakdowns or near-breakdowns occur, the matrix is perturbed so that the algorithm can proceed to produce an

approximate inverse B . The resulting preconditioned system is BA , and they showed that it has a clustered spectrum.

One application of band-Toeplitz preconditioners is in the solving of Toeplitz-plus-band systems $(A + D)\mathbf{x} = \mathbf{b}$. Here A is an n -by- n Hermitian Toeplitz matrix and D is an n -by- n Hermitian band matrix with band-width independent of n . These systems appear in solving Fredholm integro-differential equations of the form

$$\mathcal{D}\{x(\theta)\} + \int_{\alpha}^{\beta} a(\theta - \phi)x(\phi)d\phi = b(\theta),$$

where $x(\theta)$ is the unknown function to be found, $a(\theta)$ is a convolution kernel and \mathcal{D} is a differential operator. After discretization, the integral will lead to a Toeplitz matrix and \mathcal{D} a band matrix, see Delves and Mohamed [71, p.343]. Toeplitz-plus-band matrices also appear in signal processing literature and have been referred to as peripheral innovation matrices, see Carayannis et al. [24].

Unlike Toeplitz systems, there exist no fast direct solvers for solving Toeplitz-plus-band systems. It is mainly because the displacement rank of the matrix $A + D$ can take any value between 0 and n . Hence, fast Toeplitz solvers that are based on small displacement rank of the matrices cannot be applied. Conjugate gradient methods with circulant preconditioners do not work for Toeplitz-plus-band systems either. In fact, Strang's circulant preconditioner is not even defined for non-Toeplitz matrices. T. Chan's circulant preconditioner, while defined for $A + D$, does not work well when the eigenvalues of D are not clustered, see [45]. Also, the matrix $c(A) + D$ cannot be used as a preconditioner for it cannot be inverted easily. In [45], R. Chan and Ng proposed using the matrix $B + D$ to precondition $A + D$ where B is the band-Toeplitz preconditioner given in Theorem 16. Both theoretical and numerical results show that the convergence performance of the preconditioner is better in the number of iterations than those of non-preconditioned systems and circulant preconditioned systems. Clearly, instead of B , the band-Toeplitz preconditioners in Theorem 17 could be used.

4 Applications to Toeplitz-Related Systems

In this section, we discuss applications of optimal transform based preconditioners to Toeplitz-related systems arising from partial differential equations, queueing problems, signal and image processing, integral equations, and time series analysis. Part of the motivation of using optimal transform based preconditioners is to exploit their fast inversion via their transform matrices. In most of the applications, we will simply use the optimal circulant-type preconditioners. We therefore start by extending the results for point optimal circulant preconditioners in §2.1.2 to block circulant preconditioners.

4.1 Block Circulant Preconditioners

Let us consider a general system $A\mathbf{x} = \mathbf{b}$, where A is an mn -by- mn matrix partitioned as

$$A = \begin{bmatrix} A_{1,1} & A_{1,2} & \cdots & A_{1,m} \\ A_{2,1} & A_{2,2} & \cdots & A_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ A_{m,1} & A_{m,2} & \cdots & A_{m,m} \end{bmatrix}. \quad (4.1)$$

Here the blocks $A_{i,j}$ are square matrices of order n . Given such a matrix A , one obvious choice is to use the mn -by- mn point-circulant matrix $c(A)$, defined by (2.6), as a circulant approximation to A . However, this in general will spoil the block structure of A .

T. Chan and Olkin [59] and Holmgren and Otto [105], in solving noise reduction problems and hyperbolic differential equations, independently proposed using circulant-block (CB) matrices (cf. Davis [69, p.181]) to approximate A . Since $c(\cdot)$ is well-defined for any square matrix, it is natural to define the circulant-block approximation to A as:

$$c_1(A) = \begin{bmatrix} c(A_{1,1}) & c(A_{1,2}) & \cdots & c(A_{1,m}) \\ c(A_{2,1}) & c(A_{2,2}) & \cdots & c(A_{2,m}) \\ \vdots & \vdots & \ddots & \vdots \\ c(A_{m,1}) & c(A_{m,2}) & \cdots & c(A_{m,m}) \end{bmatrix}. \quad (4.2)$$

Some of the spectral properties of $c(A)$ can be extended to $c_1(A)$ (cf. (2.7)).

Theorem 18 (R. Chan and Jin (1992) [36]) *Given any mn -by- mn Hermitian matrix A partitioned as in (4.1), we have*

$$\lambda_{\min}(A) \leq \lambda_{\min}(c_1(A)) \leq \lambda_{\max}(c_1(A)) \leq \lambda_{\max}(A).$$

In particular, if A is positive definite, then $c_1(A)$ is also positive definite. Moreover, the operator $c_1(\cdot)$ is a linear projection operator with operator norm $\|c_1\|_2 = \|c_1\|_F = 1$.

It is interesting to note that the matrix $c_1(A)$ is just the minimizer of $\|A - C\|_F$ over all matrices C that are m -by- m block matrices with n -by- n circulant blocks. It can also be viewed as the approximation of A along one specific direction. It is natural to consider the preconditioner that results from approximation along the other direction. In this case, the preconditioners are block-circulant (BC) matrices (cf. Davis [69, pp. 176–177]). Both circulant-block and block-circulant preconditioners are called *level-1* preconditioners in T. Chan and Olkin [59].

Clearly one can do approximations in both directions and get a preconditioner $c_2(A)$, which is based on circulant approximations within each block and also on each block level. The resulting preconditioners are block-circulant-circulant-block (BCCB) matrices (cf. Davis [69, p.184]) and are called *level-2* preconditioners in [59]. BCCB preconditioners for block-Toeplitz-Toeplitz-block (BTTB) matrices (cf. Chan and Olkin [59]) and low rank perturbations thereof have been

investigated by Holmgren and Otto [104], Ku and Kuo [129], Tyrtysnikov [184] and R. Chan and Jin [36].

BTTB matrices occur in many applications. Let us consider the cost of constructing circulant-block and BCCB preconditioners for them. Let

$$A = \begin{bmatrix} A_0 & A_1 & \cdots & A_{m-1} \\ A_1 & A_0 & & A_{m-2} \\ \vdots & & \ddots & \vdots \\ A_{m-1} & A_{m-2} & \cdots & A_0 \end{bmatrix}, \quad (4.3)$$

where the blocks A_k are themselves symmetric Toeplitz matrices of order n . By (4.2), the blocks of $c_1(A)$ are just $c(A_k)$. From (1.4) and (3.1), we see that $c(A_k) = F^* \delta(F A_k F^*) F$, where $\delta(F A_k F^*)$ is the diagonal matrix whose diagonal equals that of $F A_k F^*$. Since A_k are Toeplitz matrices, each $\delta(F A_k F^*)$ can be computed in $O(n \log n)$ operations. Therefore, we need $O(mn \log n)$ operations to form

$$\Delta = (I \otimes F) c_1(A) (I \otimes F^*) = \begin{bmatrix} \delta(F A_0 F^*) & \delta(F A_1 F^*) & \cdots & \delta(F A_{m-1} F^*) \\ \delta(F A_1 F^*) & \delta(F A_0 F^*) & & \delta(F A_{m-2} F^*) \\ \vdots & & \ddots & \vdots \\ \delta(F A_{m-1} F^*) & \delta(F A_{m-2} F^*) & \cdots & \delta(F A_0 F^*) \end{bmatrix}. \quad (4.4)$$

Here \otimes is the Kronecker tensor product and I is the m -by- m identity matrix.

To solve $c_1(A)\mathbf{y} = \mathbf{d}$, we permute the diagonals in (4.4) to form a block diagonal matrix

$$\tilde{A} = P^* \Delta P = \begin{bmatrix} \tilde{A}_{1,1} & 0 & \cdots & 0 \\ 0 & \tilde{A}_{2,2} & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & \tilde{A}_{n,n} \end{bmatrix} \quad (4.5)$$

where

$$[\tilde{A}_{k,k}]_{ij} = [\delta(F A_{i,j} F^*)]_{kk} = [\delta(F A_{|i-j|} F^*)]_{kk}, \quad 1 \leq i, j \leq m, \quad 1 \leq k \leq n.$$

Notice that the diagonal blocks $\tilde{A}_{k,k}$ are still symmetric Toeplitz matrices of order m . We note that the linear equations with $\tilde{A}_{k,k}$ have to be solved in every iteration step. Therefore, it is convenient to use $O(m \log^2 m)$ operations algorithm (cf. [4]) to compute a Gohberg-Semencul formula [87]. Thus, the matrix-vector product $\tilde{A}_{k,k}^{-1} \mathbf{v}$ for any vector \mathbf{v} , can be computed in $O(m \log m)$ operations. Hence the system $c_1(A)\mathbf{y} = \mathbf{d}$ can be solved in $O(mn(\log m + \log n))$ operations.

For the BCCB preconditioner $c_2(A)$, we have to take the level-1 approximation of \tilde{A} in (4.5). That will add another $O(nm \log m)$ operations to the construction cost. However, the cost of solving $c_2(A)\mathbf{y} = \mathbf{d}$ is reduced from $O(nm \log^2 m + mn(\log m + \log n))$ to $O(mn \log(mn))$ as the preconditioner $c_2(A)$ can be completely diagonalized by FFTs.

Next we consider the cost of the matrix-vector multiplication $A\mathbf{v}$. We recall that the matrix-vector multiplication $A_k\mathbf{w}$ for any n -vector \mathbf{w} can be computed by FFTs of size $2n$ by first embedding A_k into a $2n$ -by- $2n$ circulant matrix, see (1.6). For the matrix-vector product $A\mathbf{v}$, we use the same trick. We first embed A into a (block-wise) $2m$ -by- $2m$ block circulant matrix, where each block itself is a $2n$ -by- $2n$ circulant matrix. Then we extend \mathbf{v} to a $4mn$ -vector by putting zeros in the appropriate places. We note that $A\mathbf{v}$ can be obtained in $O(mn \log(mn))$ operations.

Thus, we conclude that for level-1 circulant preconditioners $c_1(A)$, the initialization cost is $O(mn \log n)$ and the cost per iteration is $O(nm \log^2 m + mn \log n)$. For level-2 preconditioners $c_2(A)$, both costs are $O(mn \log(mn))$. Hence, it is cheaper to use level-2 preconditioners. As for the convergence rate, we have the following theorem for $c_2(A)$ when A is generated by a positive Wiener class function. (See also Ku and Kuo [129] for matrices generated by rational functions instead.)

Theorem 19 (Tyrtyshnikov [184]) *Let A be given by (4.3) with the entries of the block A_j denoted by $a_{pq}^{(j)} = a_{p-q}^{(j)}$, for $0 \leq p, q < n, 0 \leq j < m$. If $a_k^{(j)}$ are Fourier coefficients of a positive function in the Wiener class, i.e.,*

$$\sum_{j=0}^{\infty} \sum_{k=0}^{\infty} |a_k^{(j)}| < \infty$$

and

$$\sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} a_k^{(j)} e^{-ijx} e^{-iky} > 0, \quad \forall x, y \in [-\pi, \pi],$$

then the spectrum of $c_2(A)^{-1}A$ is clustered around 1 for large n and m .

As a consequence, when the conjugate gradient method is applied to solving the system $c_2(A)^{-1}A\mathbf{x} = c_2(A)^{-1}\mathbf{b}$, we expect superlinear convergence as in the point case. We recall that the algorithm requires $O(mn \log(mn))$ operations both in the initialization step and also in each iteration. Thus, the total complexity of the algorithm is bounded above by $O(mn \log(mn))$.

Besides extending results in the point-circulant case to the block-circulant case, one can also extend the results for other types of preconditioners as discussed in §3 to the block case. Works in this direction can be found in Di Benedetto [10] where sine transform type preconditioners are extended, in Serra [167] where Theorem 16 is extended, in Jin [120] where Theorem 17 is extended, in Hemmingsson [100] where Toeplitz-block (TB) and BTTB matrices as preconditioners are considered, and finally in Tyrtyshnikov [183, 184, 185, 186] where results on superoptimal circulant preconditioners and product preconditioners are generalized to the block case.

4.2 Applications to Partial Differential Equations

In this subsection, we consider using preconditioned conjugate gradient methods with optimal transform based preconditioners to solve different types of partial differential equations.

4.2.1 Elliptic Problems

Consider the elliptic problem

$$-(a(x, y)u_x)_x - (b(x, y)u_y)_y = g(x, y) \quad (4.6)$$

on the unit square $[0, 1] \times [0, 1]$ with Dirichlet boundary conditions. After discretization with a mesh-size of $1/(n+1)$, such a problem reduces to the solution of an n^2 -by- n^2 linear system of the form $A\mathbf{x} = \mathbf{b}$. With a standard 5-point stencil and lexicographical ordering, the discretization matrix A is a block tridiagonal matrix, where the diagonal blocks are tridiagonal matrices and the off-diagonal blocks are diagonal matrices. The linear system is often solved by iterative methods such as the preconditioned conjugate gradient method.

R. Chan and T. Chan [31] proposed two choices of circulant preconditioners for these discretization matrices A . The first one is $C_P = c(A) + \rho n^{-2}I$, where $c(\cdot)$ is defined by (2.6). The diagonals of this point-circulant preconditioner is thus obtained as the simple averages of the coefficients $a(x, y)$ and $b(x, y)$ over the whole grid. The second choice is a BCCB preconditioner which preserves the block structure of A and is defined as $C_B = c_2(A) + \rho n^{-2}I$, where $c_2(A)$ is the level-2 preconditioner of A as defined in §4.1. Thus, the diagonals are obtained as the simple averages of the coefficients along the lines of the grid. The constant ρn^{-2} added to the main diagonal is to minimize the condition number of the resulting preconditioned systems, an approach similar in idea to that used in modified incomplete LU factorizations, see [73].

We note that the product $C_B^{-1}\mathbf{y}$ can be computed by using $2n$ FFTs of size n and the solution of n intermediate circulant tridiagonal systems, each requiring $O(n)$ complexity, whereas $C_P^{-1}\mathbf{y}$ requires two FFTs of size n^2 . Similar circulant preconditioners can be defined for more general elliptic operators with more complicated difference stencils and also in higher dimensions.

Theorem 20 (R. Chan and T. Chan (1992) [31]) *Assume that in (4.6),*

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

for some constants c_{\min} and c_{\max} . Then we have

$$\kappa(C_B^{-1}A) \leq O(n) \quad \text{and} \quad \kappa(C_P^{-1}A) \leq O(n \log n).$$

We remark that for second-order elliptic problems, the condition number $\kappa(A)$ of A is $O(n^2)$. It follows from the above theorem that the condition number of the system is reduced from $O(n^2)$ to $O(n)$ by circulant preconditioning. Works in this direction can be found in Huckle [115] where skew circulant preconditioner are used for these discretization matrices. Finally, we note that the application of the circulant preconditioners requires $O(n^2 \log n)$ flops per iteration, which is slightly more expensive than the $O(n^2)$ flops for the incomplete LU-type preconditioners studied in [7, 73, 94]. However, FFTs can be computed in $O(\log n)$ parallel steps with $O(n^2)$ processors, whereas the incomplete LU preconditioners require at least $O(n)$ steps regardless of how many processors are available. The computation (based on averaging of the coefficients of the elliptic

operator) of these circulant preconditioners is also highly parallelizable across a wide variety of architectures.

One can view the circulant preconditioners as approximations to the original given equation, but with the given boundary conditions being replaced by periodic ones. It is thus natural to consider using other fast transform based preconditioners to precondition elliptic problems. The discrete sine transform matrix $\Psi^{(s)}$ with its entries defined by (3.2) diagonalizes all symmetric tridiagonal Toeplitz matrices, in particular the 1-dimensional discrete Laplacian with Dirichlet boundary conditions: $\text{tridiag}[-1, 2, -1]$. Therefore, one expects that the optimal sine transform based preconditioners discussed in §3.2.1 will give better approximations to elliptic problems with Dirichlet boundary conditions.

In [52], R. Chan and Wong proposed using these optimal sine transform based preconditioners for matrices A that come from the discretization of second-order elliptic operators. For simplicity, we let $(\hat{\Sigma} + \Sigma)\Sigma^{-1}(\hat{\Sigma} + \Sigma)$ be the block Cholesky factorization of A with lower block triangular matrix $\hat{\Sigma}$, and diagonal block matrix Σ . For such factorizations, the preconditioner is defined to be the matrix

$$P = (\hat{\Sigma} + \Phi)\Phi^{-1}(\hat{\Sigma} + \Phi)$$

with block diagonal matrix Φ and lower block triangular matrix $\hat{\Sigma}$. Here the diagonal blocks of Φ and the subdiagonal blocks of $\hat{\Sigma}$ are respectively the optimal sine transform approximations to the diagonal blocks of Σ and the subdiagonal blocks of $\hat{\Sigma}$. R. Chan and Wong showed that for 2-dimensional domains, the construction cost of P and the cost for each iteration of the preconditioned conjugate gradient algorithm are of $O(n^2 \log n)$. They also showed that for rectangular domains, P can be obtained from A by taking the optimal sine transform approximations of each sub-block of A . Thus, the construction of P is similar to the level-1 circulant preconditioners for rectangular domains, except that the construction process can now be extended to irregular regions as well.

For rectangular regions, the condition number of the preconditioned system $P^{-1}A$ is proved to be of $O(1)$. In contrast, the system preconditioned by the MILU, MINV, and optimal circulant preconditioners are of $O(n)$. We remark that a similar construction of optimal circulant approximations on L-shaped domains has recently been considered by Lirkov and Margenov [138].

In the following, we compare the performance of preconditioned conjugate gradient methods with the optimal circulant preconditioners, optimal sine transform based preconditioners, MILU, and MINV types preconditioners. The equation we used is

$$\frac{\partial}{\partial x}[(1 + \varepsilon e^{x+y})\frac{\partial u}{\partial x}] + \frac{\partial}{\partial y}[(1 + \frac{\varepsilon}{2} \sin(2\pi(x+y)))\frac{\partial u}{\partial y}] = g(x, y). \quad (4.7)$$

The parameter ε controls the variation of the coefficient functions. We remark that for this equation, there are other alternative competitive methods, such as multigrid methods.

The initial guess and the right hand side are chosen to be random vectors and are the same for all methods. The iteration stops when the residual vector \mathbf{r}_q at the q th iteration satisfies $\|\mathbf{r}_q\|_2 / \|\mathbf{r}_0\|_2 < 10^{-6}$. All computations are done with double precision on a VAX 6420.

We remark again that the costs per iteration of sine transform based or circulant preconditioners are $O(\log n)$ times more than those required by MINV or MILU type preconditioners. However, most of the computations involving fast transform based preconditioners can be done in parallel. Thus, it is difficult to just compare either the time or flop counts. For simplicity, we here just compare the iteration numbers. Tables 2–3 show the numbers of iterations required for convergence for the equation (4.7) on the unit square and the L-shaped domain $[0, 1]^2 \setminus [1/2, 1] \times [0, 1/2]$, respectively, see [31, 52]. In the tables, the notation I means that no preconditioner is used and the parameter h is the mesh-size. These results show that the optimal sine transform based preconditioner requires less number of iterations than that of the optimal circulant, MILU, and MINV preconditioners for small ε , i.e., when the variation of the coefficients is small.

$1/h$	$\varepsilon = 0.01$					$\varepsilon = 1.0$				
	I	C_B	P	MILU	MINV	I	C_B	P	MILU	MINV
4	12	9	3	7	4	15	10	5	6	3
8	25	12	3	9	5	29	13	7	9	4
16	47	15	3	13	7	54	18	9	14	6
32	90	20	3	20	11	107	25	11	20	10
64	186	25	3	28	16	209	35	12	28	15
128	363	33	3	41	24	419	50	13	41	22

Table 2: Number of Iterations for the Unit Square.

$1/h$	$\varepsilon = 0.01$				$\varepsilon = 1.0$			
	I	P	MILU	MINV	I	P	MILU	MINV
8	22	3	9	4	24	7	9	4
16	40	3	12	6	45	9	13	6
32	80	4	17	9	86	10	18	8
64	155	4	25	14	169	12	26	12
128	311	4	36	21	338	14	37	19

Table 3: Number of Iterations for the L-Shaped Domain.

4.2.2 Domain Decomposition

Let $\Omega \subset \mathbb{R}^2$ be a bounded domain with a piecewise smooth boundary $\partial\Omega$. Given $f \in L^2(\Omega)$, we are interested in finding the weak solution $u \in H^1(\Omega)$ of the following Neumann problem:

$$a(u, v) = \int_{\Omega} f v dx, \quad \forall v \in H^1(\Omega),$$

where

$$a(u, v) = \int_{\Omega} \left(\sum_{i,j=1}^2 a_{ij}(x) \frac{\partial u}{\partial x_i} \frac{\partial v}{\partial x_j} + c(x) uv \right) dx,$$

and $a_{ij}(x), c(x) \in L^\infty(\Omega)$. We assume that the bilinear form $a(u, v)$ is symmetric and there exist positive constants α_1, α_2 such that

$$\alpha_1 \|v\|_{H^1(\Omega)}^2 \leq a(u, v) \leq \alpha_2 \|v\|_{H^1(\Omega)}^2.$$

Suppose that Ω is a disjoint union of two subdomains Ω_1 and Ω_2 , and that a fast solver is available on each subdomain. Domain decomposition methods for elliptic problems defined on union of subdomains have been studied in many papers, see for instance [72, 191]. The idea of substructuring is to reduce the problem in Ω to one on the interface $\partial\Omega_1 \cap \partial\Omega_2$. Let $H^{1/2}(\partial\Omega_k)$ be the Sobolev space of order one half on the boundary $\partial\Omega_k$, $k = 1, 2$, with weighted norm:

$$\|v\|_{H^{1/2}(\partial\Omega_k)} \equiv \int_{\partial\Omega_k} \int_{\partial\Omega_k} \frac{|v(x(s), y(s)) - v(x(r), y(r))|^2}{\| (x(s), y(s)) - (x(r), y(r)) \|^2} ds dr + \frac{1}{d_k} \int_{\partial\Omega_k} |v(x(s), y(s))|^2 ds \quad (4.8)$$

where $\|\cdot\|$ is the Euclidean distance in \mathbb{R}^2 . In [126, 127], Kiss and Molnárka proposed using circulant matrices as preconditioners for these elliptic problems. Their idea is to approximate the Euclidean norm $\|\cdot\|$ in (4.8) by

$$|s - r|_k \equiv \min\{|s - r|, \text{length}(\partial\Omega_k) - |s - r|\}.$$

It turns out that the matrix representation C of the resulting approximated bilinear form is circulant and for the Neumann problem it is a good approximation.

Theorem 21 (Kiss and Molnárka (1991,1992) [127, 126]) *For the Neumann problem, the original bilinear form and the approximated one are spectrally equivalent, and hence $\kappa(C^{-1}A) = O(1)$.*

We remark that Dirichlet and mixed boundary problems have also been considered in [126, 127].

4.2.3 Hyperbolic and Parabolic Problems

The idea of circulant preconditioners has also been applied to systems arising from implicit time-marching methods for first-order hyperbolic equations of the form

$$\frac{\partial u}{\partial t} + a(x, y) \frac{\partial u}{\partial x} + b(x, y) \frac{\partial u}{\partial y} = g(x, y),$$

see Holmgren and Otto [103, 104, 105, 106, 107] and Otto [156]. For such problems, the discretization matrix A is non-symmetric and often highly non-diagonally dominant, and hence many classical preconditioning techniques are not effective (and sometimes not well-defined). For these problems, the circulant preconditioners are often the only ones that work.

In [104], comparisons of circulant and BCCB preconditioners with incomplete LU (ILU) and block ILU are done. Some of the circulant-type preconditioners are obtained by changing the boundary conditions from Dirichlet type to periodic ones. It is found that if κ , the ratio between the time-step and the spatial mesh size, is large, then ILU-type preconditioners fail to converge, while the rate of convergence for circulant-type preconditioners remains constant independent of the mesh size. In [105], a framework of CB and BCCB preconditioners is developed, and an analysis of the convergence rate for a BCCB preconditioner is performed for problems that have both periodic and Dirichlet boundary conditions.

Analysis of a circulant-block preconditioner for problems with only Dirichlet boundary conditions is done in [156], which also includes a Fourier analysis of ILU, MILU, and block-MILU preconditioned systems. The analysis indicates that the condition number of the MILU preconditioned system remains $O(\kappa)$, whereas the convergence rate of the circulant-block preconditioned systems remains independent of κ and mesh size h . The convergence analysis was later extended to include a time-independent equation with a weak artificial viscosity, see [106]. Applications of circulant-type preconditioners to the computation of flow in a driven cavity governed by the Navier-Stokes equations can be found in [107]. Implementations of the method on shared memory vector computers and distributed memory computers are studied in [103].

Table 4 below, taken from [104, Table 8], shows the effectiveness of using BCCB matrices as preconditioners for first-order hyperbolic systems when κ is large.

$1/h$	None	MILU	Block MILU	BCCB
8	48	18	198	17
16	223	32	>500	20
32	334	51	>500	19
64	423	69	>500	20
128	473	89	>500	20
256	>500	135	>500	21
512	>500	>500	>500	19

Table 4: Number of CGS Iterations When $\kappa = 100$.

In [100], Hemmingsson has considered using Toeplitz-block and BTTB preconditioners for the same first-order hyperbolic PDE. The Toeplitz blocks are obtained by minimization as is done in (2.3), except that the minimization is now taken over a restricted set of Toeplitz matrices. The convergence rate of the resulting preconditioned methods is also favorable and thoroughly analyzed in [101].

Circulant preconditioners for second-order hyperbolic equations have been considered by Jin

and R. Chan [121]. In [121], the equation

$$u_{tt} = (a(x, y)u_x)_x + (b(x, y)u_y)_y + g(x, y)$$

is considered. The coefficient matrix, obtained from an implicit time-marching scheme, has condition number of $O(\kappa^2) + O(h^{-2})$. With circulant-type preconditioners, Jin and R. Chan [121] proved that the condition number is reduced to $O(\kappa) + O(h^{-1})$. The same idea of proof can also be applied to parabolic equations of the form

$$u_t = (a(x, y)u_x)_x + (b(x, y)u_y)_y + g(x, y),$$

see [118]. In [107], Holmgren and Otto investigated circulant-type solvers for mixed hyperbolic-parabolic equations. The second-order terms were considered small or used as artificial viscosity, such as for the discretized Euler equations in computational fluid dynamics.

4.3 Applications to Queueing Problems

4.3.1 Overflow Queueing Networks

Consider a 2-queue Markovian network with overflow permitted only from queue 1 to queue 2 when queue 1 is full, see [124]. We are interested in finding the steady-state probability distribution vector of the network. Let λ_i , μ_i , n_i , and s_i be the input rate, output rate of a single server, buffer size and number of servers for queue i , respectively. If the traffic density, defined as $\lambda_i/(s_i\mu_i)$ is close to 1, i.e.,

$$\frac{\lambda_i}{s_i\mu_i} = 1 + O(n_i^{-\alpha}), \tag{4.9}$$

for some $\alpha > 0$, then the queueing problem resembles a second-order elliptic equation on a rectangle with an oblique boundary condition on one side (the side with overflow) and Neumann boundary conditions on the others, see [26].

The SOR method is one of the standard methods for solving this problem, see [124]. However, in [26], the preconditioned conjugate gradient method has also been considered, with the preconditioner being constructed by changing the oblique boundary condition to Neumann boundary condition. This preconditioner will be referred to as ‘‘Neumann’’ preconditioner below. The convergence rate of the preconditioned conjugate gradient method with Neumann preconditioner is much better than the SOR method, ranging from 10 times faster for small n_i to about 100 times faster for $n_i = 128$. However, inversion of Neumann preconditioners is expensive when the number of servers $s_i > 1$, see [26].

Since the 1-dimensional discrete Laplacian with Neumann boundary conditions can be diagonalized by the cosine transform matrix, one is naturally lead to consider optimal cosine transform based preconditioners mentioned in §3.2 for such queueing systems. By using FFTs, each inversion of the preconditioner requires $O(n_1n_2 \log(n_1n_2))$ operations independent of s_i . Thus, the cost per iteration is comparable to the $O(n_1n_2)$ operations required by the SOR method.

Numerical results in R. Chan, Ching, and Wong [33] show that the optimal cosine transform based preconditioner performs even better than the Neumann preconditioner. Table 5 gives the number of iterations required for convergence when the tolerance is 10^{-6} . In the test, the parameters λ_i , $i = 1, 2$, are set to 1 and μ_i are computed according to (4.9) with $\alpha = 1$. For the SOR method, the optimal relaxation factor, obtained numerically to 4 significant digits, is used. The results show that the new method is about 10 to 200 times faster than the SOR method.

$n_1 = n_2$	$s_1 = s_2 = 1$			$s_1 = s_2 = 5$		
	Cosine	Neumann	SOR	Cosine	Neumann	SOR
8	6	7	210	7	7	70
16	8	9	512	10	9	196
32	9	11	> 1000	10	12	533
64	10	14	> 1000	12	14	> 1000
128	10	14	> 1000	13	16	> 1000

Table 5: Number of Iterations for Different Preconditioners.

4.3.2 Queueing Networks with Batch Arrivals

Queueing systems with batch arrivals occur in many applications, such as the telecommunication networks [154] and the loading dock models [166]. Again we are interested in finding the stationary distribution vector of the network. It will be the normalized null-vector of the generator matrix of the problem.

For this problem, the generator matrix can be written as a sum of a Toeplitz matrix and a rank s matrix, where s is the number of servers. Since the generator matrix is singular, the Toeplitz matrix will have a generating function f that has zeros. In R. Chan and Ching [32], they have considered using circulant-Toeplitz preconditioners discussed in §3.3 to precondition such systems. The singularity of the generator matrix is canceled by the band-Toeplitz part of the preconditioner, and the circulant part of the preconditioner is used to speed up the convergence of the algorithm. They proved superlinear convergence of the method when s is independent of n , the size of the queue. Numerical results in [32] verified the fast convergence for small s , and show moreover that the convergence is still linear when $s = n$.

4.4 Applications to Signal and Image Restoration

Image restoration refers to the removal or reduction of degradations (or blur) in an image using a priori knowledge about the degradation phenomena. Applications of image restoration can be found in remote sensing, where details about the photographed terrain are needed to be resolved; in medical imaging where the diagnosis is based on the clarity of the x-ray radiographs taken; and in space exploration where images transmitted back to earth by spacecrafts are analyzed.

In ground-based imaging, astronomers seek to remove the degradation of astronomical images caused by the atmospheric turbulence, which are in part due to the mixing of the warm and cold air layers. When the quality of the images is degraded by blurring and noise, important information remains hidden and cannot be directly interpreted without numerical processing. Our presentation of the image restoration problem here will be brief, and we refer the interested readers to Andrews and Hunt [5] and Jain [116] for more detail discussions.

We begin with a mathematical model of the image restoration problem. The image of an object can be modeled as

$$g(\xi, \delta) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} a(\xi, \delta; \alpha, \beta) h(\alpha, \beta) d\alpha d\beta + \eta(\xi, \delta), \quad (4.10)$$

where $g(\xi, \delta)$ is the recorded (or degraded) image, $h(\alpha, \beta)$ is the original image, the vector $\eta(\xi, \delta)$ represents an additive noise. The function $a(\xi, \delta; \alpha, \beta)$ is called the *point spread function* (PSF) and represents the degradation of the image. In digital implementation, (4.10) is discretized to obtain the discrete scalar model

$$g(i, j) = \sum_{k=1}^n \sum_{\ell=1}^n a(i, j; k, \ell) h(k, \ell) + \eta(i, j).$$

In matrix-vector notation, we obtain the linear algebraic form of the image restoration problem,

$$\mathbf{g} = A\mathbf{h} + \mathbf{n}, \quad (4.11)$$

where \mathbf{g} , \mathbf{h} , and \mathbf{n} are n^2 -vectors and A is an n^2 -by- n^2 matrix. This is the square image formulation. Often the discretization is chosen so that \mathbf{g} is a longer vector than \mathbf{h} . In this case, A is a rectangular m^2 -by- n^2 matrix with $m > n$. The image restoration problem can be stated as follows. Given the observed image \mathbf{g} , the matrix A which represents the degradation, and possibly, the statistics of the noise vector \mathbf{n} , compute an approximation to the original signal \mathbf{h} .

Writing the PSF as $a(\xi, \delta; \alpha, \beta)$ provides the most general description of the imaging system. This representation allows the PSF to vary with position in both the image and object planes. In this case the PSF is said to be *spatially variant*, and the matrix A in (4.11) will have no special structure. Thus, computing a solution to (4.11) can be very expensive. In many practical applications, though, the PSF is *spatially invariant*, i.e., it acts uniformly across the image and object planes. In particular, it can be written as

$$a(\xi, \delta; \alpha, \beta) = a(\xi - \alpha, \delta - \beta),$$

and the matrix A it generates is a block-Toeplitz-Toeplitz-block (BTTB) matrix.

Because of the ill-conditioning of A , naively solving $A\mathbf{h} = \mathbf{g}$ will lead to extreme instability with respect to perturbations in \mathbf{g} , see [5]. Thus, one cannot consider the noise vector \mathbf{n} insignificant in solving the problem. The method of *regularization* can be used to achieve stability for these problems [1, 17, 162]. In the classical *Tikhonov regularization* [92], stability is attained by introducing a stabilizing operator D (called a regularization operator), which restricts the

set of admissible solutions. Since this causes the regularized solution to be biased, a scalar μ , called a regularization parameter, is introduced to control the degree of bias. More specifically, the regularized solution is computed as

$$\min \left\| \begin{bmatrix} \mathbf{g} \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} A \\ \mu D \end{bmatrix} \mathbf{h}(\mu) \right\|_2. \quad (4.12)$$

The term $\|D\mathbf{h}\|_2^2$ is added in order to smooth the solution \mathbf{h} . Choosing D as a k th order difference operator matrix forces the solution to have a small k th order derivative. Notice that if D is a Toeplitz matrix, then (4.12) reduces to a block-Toeplitz least squares problem.

Thus in general, let us consider the least squares problem

$$\min_{\mathbf{x}} \|\mathbf{b} - A\mathbf{x}\|_2, \quad (4.13)$$

where A is a rectangular Toeplitz or block-Toeplitz matrix. Besides image restoration problems, Toeplitz least squares problems also arise in other important areas, such as the denoising problems, see T. Chan and Olkin [59]. Considerable effort has been devoted to developing fast algorithms for them and most works have been focused on direct methods, such as the fast QR factorization algorithms of Bojanczyk, Brent, and de Hoog [18], Chun and Kailath [63], Cybenko [68], and Sweet [176]. Here we will consider using the preconditioned conjugate gradient method.

Although the classical conjugate gradient algorithm applies only to square Hermitian positive definite systems, one can still use it to find the solution to (4.13) by applying it to the normal equations in factored form,

$$A^*(\mathbf{b} - A\mathbf{x}) = 0.$$

The method can be applied without explicitly forming the normal equations matrix A^*A , see Björck [17]. As in the square case, we can precondition the equation to speed up the convergence. Given a non-singular matrix C , one can use the conjugate gradient method to solve

$$\min \|\mathbf{b} - AC^{-1}\mathbf{y}\|_2,$$

and then set $\mathbf{x} = C^{-1}\mathbf{y}$. The cost per iteration of the preconditioned conjugate gradient method is dominated by matrix vector multiplies with A and A^* , and by linear system solves with C as coefficient matrix. If A is an m -by- n Toeplitz matrix, then matrix vector multiplies with A and A^* can be accomplished in $O(m \log n)$ operations using FFTs. As discussed in §1.3, the preconditioner matrix C should be chosen such that the singular values of AC^{-1} are clustered around 1, and the linear system with coefficient matrix C can be easily solved. In the following subsections, we survey some of the possible preconditioners.

4.4.1 Block-Based Preconditioners

We start with the 1-dimensional Toeplitz least squares problems. For the purpose of constructing the preconditioner, we extend the Toeplitz structure of the matrix A in (4.13) by padding zeros to the bottom left-hand side. In doing so, we may assume without loss of generality that $m = kn$ for

some positive integer k . This padding is only for convenience in constructing the preconditioner and does not alter the original least squares problem. In the material to follow, we consider the case where k is a constant independent of n . More precisely, we consider kn -by- n matrices A of the form

$$A = \begin{bmatrix} A_1 \\ A_2 \\ \vdots \\ A_k \end{bmatrix},$$

where each square block A_j is a Toeplitz matrix. Notice that if A itself is a rectangular Toeplitz matrix, then each block A_j is necessarily Toeplitz. Also the matrix in (4.12) is of this form if the regularization operator D is a Toeplitz matrix.

Following [144, 42], for each block A_j , the optimal circulant approximation $c(A_j)$ is constructed. Then the block-based preconditioner is defined to be the square circulant matrix C such that

$$C^*C = \sum_{j=1}^k c(A_j)^* c(A_j).$$

Notice that each $c(A_j)$ is an n -by- n circulant matrix. Hence, they can all be diagonalized by the n -by- n discrete Fourier matrix F , i.e., $c(A_j) = F^* \Lambda_j F$ where Λ_j is diagonal. Thus, the spectrum of $c(A_j)$, $j = 1, \dots, k$, can be computed in $O(n \log n)$ operations by using FFT and we have

$$C^*C = F^* \sum_{j=1}^k (\Lambda_j^* \Lambda_j) F.$$

Clearly C^*C is a circulant matrix and its spectrum can be computed in $O(kn \log n)$ operations. The preconditioner is then given by

$$C = F^* \left(\sum_{j=1}^k \Lambda_j^* \Lambda_j \right)^{\frac{1}{2}} F.$$

Recall that the matrix vector multiplies with A and A^* can be done in $O(m \log n)$ operations by FFT, the cost per iteration in the preconditioned conjugate gradient method is therefore of $O(m \log n)$. The convergence rate of the method depends on the distribution of the singular values of the matrix AC^{-1} , which are the same as the square roots of the eigenvalues of the matrix $(C^*C)^{-1}(A^*A)$.

Theorem 22 (R. Chan, Nagy, and Plemmons (1994) [42]) *Suppose the generating functions of the blocks A_j are 2π -periodic continuous functions and if one of these functions has no zeros, then the spectrum of $(C^*C)^{-1}(A^*A)$ is clustered around 1, for sufficiently large n .*

Thus, if the condition number $\kappa(A)$ of A is of $O(n^\alpha)$, then for sufficiently large n , the number of iterations required for convergence is at most $O(\alpha \log n)$ when $\alpha > 0$. Since the

number of operations per iteration in the conjugate gradient method is of $O(m \log n)$, the total complexity of the algorithm is therefore of $O(\alpha m \log^2 n)$. In the case when $\alpha = 0$, i.e., A is well-conditioned, the method converges in $O(1)$ steps. Hence, the complexity is reduced to just $O(m \log n)$ operations.

Next we consider the 2-dimensional Toeplitz least squares problems. Let A be a kmn -by- mn matrix of the form

$$A = \begin{bmatrix} A^{(1)} \\ A^{(2)} \\ \vdots \\ A^{(k)} \end{bmatrix}, \quad (4.14)$$

where each block $A^{(i)}$, $i = 1, \dots, k$, is a Toeplitz-block matrix. More precisely, $A^{(i)}$ can be partitioned as

$$A^{(i)} = \begin{bmatrix} A_{1,1}^{(i)} & A_{1,2}^{(i)} & \cdots & A_{1,m}^{(i)} \\ A_{2,1}^{(i)} & A_{2,2}^{(i)} & \cdots & A_{2,m}^{(i)} \\ \vdots & \vdots & \ddots & \vdots \\ A_{m,1}^{(i)} & A_{m,2}^{(i)} & \cdots & A_{m,m}^{(i)} \end{bmatrix}, \quad i = 1, \dots, k,$$

where each $A_{\alpha,\beta}^{(i)}$, $1 \leq \alpha, \beta \leq m$, is an n -by- n Toeplitz matrix.

In [43], R. Chan, Nagy, and Plemmons considered preconditioners based on the level-1 and level-2 circulant approximations proposed by T. Chan and Olkin [59]. For $j = 1, 2$, they construct the level- j approximation $c_j(A^{(i)})$ to each $A^{(i)}$, see §4.1. Then the preconditioner P_j is defined by

$$P_j^* P_j = \sum_{i=1}^k c_j(A^{(i)})^* c_j(A^{(i)}).$$

It can be shown that $P_1 = (I \otimes F_n)^* \tilde{P} (I \otimes F_n)$, where \tilde{P} is a block upper triangular matrix with diagonal blocks. Thus, P_1 is block upper triangular with circulant blocks, and hence computing $P_1^{-1} \mathbf{y}$ involves application of FFTs of size n together with backward solves involving \tilde{P} . The computation of $P_1^{-1} \mathbf{y}$ is parallelizable and involves $O(nm^2 + mn \log n)$ operations. For the level-2 preconditioner, we have $P_2 = (F_m \otimes F_n)^* \Phi (F_m \otimes F_n)$ where Φ is a diagonal matrix. Hence, $P_2^{-1} \mathbf{y}$ can be computed in $O(mn \log(mn))$ operations by using FFTs in both directions.

Convergence analysis has been given when all the blocks $A^{(i)}$ in the matrix A are BTTB, i.e.,

$$A^{(i)} = \begin{bmatrix} A_0^{(i)} & A_{-1}^{(i)} & \cdots & A_{2-m}^{(i)} & A_{1-m}^{(i)} \\ A_1^{(i)} & A_0^{(i)} & \ddots & & A_{2-m}^{(i)} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ A_{m-2}^{(i)} & & \ddots & \ddots & A_{-1}^{(i)} \\ A_{m-1}^{(i)} & A_{m-2}^{(i)} & \cdots & A_1^{(i)} & A_0^{(i)} \end{bmatrix}, \quad i = 1, \dots, k,$$

where each $A_\mu^{(i)}$, $|\mu| < m$, is a Toeplitz matrix with entries

$$[A_\mu^{(i)}]_{\gamma,\delta} = a_{\mu,\gamma-\delta}^{(i)}, \quad 0 \leq \gamma, \delta < n.$$

Notice that the generating function of $A^{(i)}$ is given by

$$f^{(i)}(x, y) = \sum_{\mu=-\infty}^{\infty} \sum_{\nu=-\infty}^{\infty} a_{\mu,\nu}^{(i)} e^{-i\mu x} e^{-i\nu y}, \quad \forall x, y \in [-\pi, \pi].$$

Theorem 23 (R. Chan, Nagy, and Plemmons (1993) [43]) *Let A be given as in (4.14) with the entries of each $A^{(i)}$ satisfying*

$$\sum_{\mu=-\infty}^{\infty} \sum_{\nu=-\infty}^{\infty} |a_{\mu,\nu}^{(i)}| \leq K_i < \infty, \quad i = 1, \dots, k.$$

If one of the $f^{(i)}(x, y)$ is a positive function, then for any given $\epsilon > 0$, there exist positive integers n_1 and n_2 such that for all $m > n_1$ and $n > n_2$, at most $O(m) + O(n)$ eigenvalues of the matrix

$$(P_j^* P_j)^{-1} (A^* A) - I, \quad j = 1, 2,$$

have absolute values larger than ϵ .

Numerical experiments in [43] show the effectiveness of the preconditioners proposed for block Toeplitz least squares problems and actual image restoration problems. We remark that by using the new clustering result in Theorem 19 for BCCB preconditioners, we can conclude further that the preconditioned matrix $(P_2^* P_2)^{-1} (A^* A)$ will have a clustered spectrum around 1.

4.4.2 Displacement-Based Preconditioners

We first briefly review relevant definitions and results on displacement structure representation of Toeplitz matrices. We introduce the n -by- n lower shift matrix Z , whose entries are zero everywhere except for 1's on the first subdiagonal. The displacement operator ∇ is defined by

$$\nabla A = A - ZAZ^*,$$

where ∇A is called the displacement of A , see Chun and Kailath [63]. Let $L(\mathbf{w})$ denote the n -by- n lower triangular Toeplitz matrix with the vector \mathbf{w} as its first column. Using these definitions, we have the following lemma.

Theorem 24 (Chun, Kailath, and Lev-Ari (1987) [64]) *An arbitrary n -by- n matrix A can be written in the form*

$$A = \sum_{i=1}^{\rho} L(\mathbf{u}_i) L(\mathbf{v}_i)^*,$$

where $\rho = \text{rank}(\nabla A)$ and \mathbf{u}_i and \mathbf{v}_i are n -vectors.

The sum given in the theorem above is called the *displacement representation* of the given matrix A , and the scalar ρ is called the *displacement rank* of A . Square Toeplitz matrices and Toeplitz-related matrices have small displacement rank [64]. For example, if A is an n -by- n Hermitian Toeplitz matrix, then

$$A = L(\mathbf{x}_+)L(\mathbf{x}_+)^* - L(\mathbf{x}_-)L(\mathbf{x}_-)^*,$$

where

$$\mathbf{x}_\pm = \left[\frac{1}{2}(a_0 \pm 1), a_1, \dots, a_{n-1}\right]^T.$$

If A is an m -by- n ($m \geq n$) Toeplitz matrix, then A^*A is in general not a Toeplitz matrix. However, A^*A does have a small displacement rank $\rho \leq 4$, and a displacement representation in the form

$$A^*A = \tilde{A} + L(\mathbf{y}_1)L(\mathbf{y}_1)^* - L(\mathbf{y}_2)L(\mathbf{y}_2)^*, \quad (4.15)$$

where \tilde{A} is a Hermitian Toeplitz matrix and

$$\mathbf{y}_1 = [0, a_{-1}, \dots, a_{1-n}]^T \quad \text{and} \quad \mathbf{y}_2 = [0, a_{m-1}, \dots, a_{m-n+1}]^T.$$

When the matrix A is an n -by- n Hermitian Toeplitz matrix, R. Chan, Nagy, and Plemmons [44] and Freund and Huckle [79] defined the *displacement preconditioner* to be

$$C = c(L(\mathbf{x}_+))c(L(\mathbf{x}_+))^* - c(L(\mathbf{x}_-))c(L(\mathbf{x}_-))^*. \quad (4.16)$$

Clearly, C is a Hermitian circulant matrix. R. Chan, Nagy, and Plemmons [44] proved that C so defined in (4.16) is equal to the optimal circulant approximation of A , i.e., $C = c(A)$. In Freund and Huckle [79], displacement based preconditioners for general matrices of low displacement ranks are also considered.

For general rectangular Toeplitz matrices, one can define, according to (4.15), the preconditioner

$$c(\tilde{A}) + c(L(\mathbf{y}_1))c(L(\mathbf{y}_1))^* - c(L(\mathbf{y}_2))c(L(\mathbf{y}_2))^*.$$

However, R. Chan, Nagy, and Plemmons [44] proved that the last term $L(\mathbf{y}_2)L(\mathbf{y}_2)^*$ in (4.15) is not significant as far as the conjugate gradient method is concerned. Thus, they define the *displacement preconditioner* C as

$$C = c(\tilde{A}) + c(L(\mathbf{y}_1))c(L(\mathbf{y}_1))^*.$$

Under the Wiener class assumptions on the generating function of A , clustering results have been established for C , see [44].

4.4.3 Strang’s Generalized Preconditioner

Recall that T. Chan’s circulant preconditioner is defined for general square matrices, not necessarily of Toeplitz form, see (2.6). Most circulant preconditioners including Strang’s do not enjoy this property. As a result, T. Chan’s preconditioners have been used extensively in solving non-Toeplitz systems in least squares problems.

Recently, R. Chan, Ng, and Plemmons [47] proposed a method to generalize the construction of Strang’s circulant preconditioner to arbitrary n -by- n matrices A , and to employ this new circulant approximation in deconvolution applications in signal and image processing. For a general n -by- n non-Toeplitz matrix $A = [a_{j,k}]$, we define its generalized Strang circulant preconditioner S to be

$$[S]_{j, \lfloor \frac{n}{2} \rfloor} = a_{j, \lfloor \frac{n}{2} \rfloor}, \quad 0 \leq j \leq n - 1,$$

i.e., the $\lfloor \frac{n}{2} \rfloor$ th column of S is given by the $\lfloor \frac{n}{2} \rfloor$ th column of A . Note that if A is Toeplitz, then this definition coincides with that in (2.1).

This idea of constructing circulant preconditioners is similar to the forward-backward projection method used in [65] for image reconstruction. In [65], they tried to estimate the PSF involving a given, but not necessarily Toeplitz, matrix A by forward projecting and backprojecting a one-pixel point source located at the center of the field-of-view. In matrix terms, the approximate PSF is just equal to the $\lfloor \frac{n}{2} \rfloor$ th column of A . The circulant matrix thus obtained was used in [65] as a preconditioner in the steepest descent method to speed up the convergence rate. The convergence rate of the preconditioned conjugate gradient method with generalized Strang preconditioner has been analyzed in [47], where again, under the Wiener class assumptions, the preconditioned matrix will have a clustered spectrum.

4.4.4 Inverse Filter Preconditioners

Let us consider the convolution of a 1-dimensional discrete signal \mathbf{x} of length n with a convolution vector \mathbf{a} of the form

$$\mathbf{a} = [a_{-m+1}, a_{-m+2}, \dots, a_0, \dots, a_{m-2}, a_{m-1}]^T.$$

The resulting vector \mathbf{b} is of length $2m + n - 2$, and the convolution operation can be expressed in matrix notation as $\mathbf{b} = A\mathbf{x}$, where A is a column circulant matrix of the form

$$A = \begin{bmatrix} a_{-m+1} & & & 0 \\ a_{-m+2} & a_{-m+1} & & \vdots \\ \vdots & \vdots & \ddots & 0 \\ a_0 & \vdots & & a_{-m+1} \\ \vdots & a_0 & & \vdots \\ a_{m-2} & \vdots & \ddots & \vdots \\ a_{m-1} & \vdots & & a_0 \\ 0 & a_{m-1} & & \vdots \\ \vdots & & \ddots & \vdots \\ 0 & & & a_{m-1} \end{bmatrix}. \quad (4.17)$$

In deconvolution problems, the aim is to compute \mathbf{x} given the matrix A and the vector \mathbf{b} . The first column of A is usually the discrete point spread function of the blurring operator and is usually obtained empirically by blurring a point source image. In many applications of interest $m \ll n$, see [116].

In signal and image processing, one common technique used to compute an approximate solution to the deconvolution problem is the inverse filter method. There the rectangular Toeplitz matrix A is embedded into an n -by- n circulant matrix C , and the solution to the circulant system is used as an approximation to the solution of the Toeplitz system. This approach is attractive since n -by- n circulant systems can be solved in $O(n \log n)$ operations by FFTs. However, the computed solution will not be very accurate.

Nagy, Plemmons, and Torgersen [145, 146] employed the inverse filter technique to construct a new preconditioner for Toeplitz least squares deconvolution problems. Their preconditioner M is based on the partition of the inverse of the circulant extension C . Thus writing

$$C^{-1} = F^* \Lambda F = \begin{bmatrix} M^* \\ \tilde{M}^* \end{bmatrix},$$

their preconditioner is defined to be M , and they showed that M^*M is a good approximate inverse of A^*A .

For 2-dimensional deconvolution problems, one is still concerned with solving a least squares problem as in (4.13). But the matrix A will be a block column circulant matrix with column

circulant blocks, i.e.,

$$A = \begin{pmatrix} A^{(-\ell+1)} & & & 0 \\ A^{(-\ell+2)} & A^{(-\ell+1)} & & \vdots \\ \vdots & \vdots & \ddots & 0 \\ A^{(0)} & \vdots & & A^{(-\ell+1)} \\ \vdots & A^{(0)} & & \vdots \\ A^{(\ell-2)} & \vdots & \ddots & \vdots \\ A^{(\ell-1)} & \vdots & & A^{(0)} \\ 0 & A^{(\ell-1)} & & \vdots \\ \vdots & & \ddots & \vdots \\ 0 & & & A^{(\ell-1)} \end{pmatrix} \quad (4.18)$$

with each subblock $A^{(j)}$ being a $(2m + n - 2)$ -by- n matrix of the form given by (4.17). We note that A^*A will be an n -block-by- n -block block Toeplitz matrix with n -by- n Toeplitz blocks. The construction of inverse filter preconditioners M for the 2-dimensional case proceeds as in the 1-dimensional case. The matrix A is extended to a block circulant matrix with circulant blocks. In [145, 146], Nagy, Plemmons, and Torgersen proved that the preconditioned matrices can also be written as a sum of the identity matrix and a matrix with rank dependent on m and ℓ .

4.4.5 Numerical Example

Restoration of real images by using the preconditioned conjugate gradient algorithm with the preconditioners mentioned above has been carried out in [43, 44, 47, 145, 146]. Here we report the results found in [47]. We use the preconditioned conjugate gradient algorithm with the generalized Strang circulant preconditioner to remove the blurring in a 2-dimensional image arising from atmospheric turbulence. The problem consists of a 256-by-256 image of an ocean reconnaissance satellite observed by a simulated ground-based imaging system together with a 256-by-256 image of a guide star observed under similar circumstances (see Figure 3.) The data are provided by the Phillips Air Force Laboratory at Kirkland AFB, NM [25]. The imaging system estimates the atmospheric distortions using the natural guide star image. A wavefront sensor measures the optical distortions, which can then be digitized into a blurred image of the guide star pixel. To form the discrete point spread function a , the rows of the blurred pixel image are stacked into a column vector. Then the point spread function matrix A is given in block form as in (4.18) with the stacked vector as its first column. We note that A satisfies the Wiener class assumptions, since the guide star for the atmospheric imaging problem yields a Gaussian point spread function [146].

In Figures 4 and 5, we present restorations without and with using the generalized Strang preconditioner described in §4.4.3. The regularization parameter μ in both cases is set to .018 as suggested in [145]. From the figures, we observe that when no preconditioner is used, an




Figure 3: Observed Image (left) and Guide Star Image (right).




Figure 4: Restored Images without Preconditioning: 10 (left) and 50 Iterations (right).




Figure 5: Restored Image with Preconditioning: 2 (left) and 10 Iterations (right).

acceptable restoration is achieved after 50 iterations. Essentially, the same restoration is achieved in 10 iterations when preconditioning is used. Finally, we remark that roughly 0.61×10^8 floating point operations (done by Matlab) per iteration are used for the circulant based deconvolution, while the count for no preconditioning is 0.50×10^8 .

4.5 Applications to Integral Equations

4.5.1 Inverse Heat Problem in \mathbb{R}^m

The inverse heat problem in \mathbb{R}^m is the problem of recovering the initial data $\Upsilon(\mathbf{y})$ for all $\mathbf{y} \in \mathbb{R}^m$ in

$$u(\mathbf{x}, t) = \frac{1}{\sqrt{4\pi t}} \int_{\mathbb{R}^m} \exp\left(-\frac{(\mathbf{x} - \mathbf{y})^2}{4t}\right) \Upsilon(\mathbf{y}) d\mathbf{y},$$

when for some $t > 0$, $u(\mathbf{x}, t)$ is given for all $\mathbf{x} \in \mathbb{R}^m$. Using discrete time and spatial sampling of the domain and sinc expansion for approximating the initial data, the problem is reduced to solving a linear system with block Toeplitz coefficient matrices:

$$(A \otimes A \otimes \cdots \otimes A)\Upsilon = \mathbf{u}, \quad (4.19)$$

where A is a Toeplitz matrix, see Gilliam, Martin, and Lund [83]. The generating function of the Toeplitz matrix A is

$$f(\theta) = \exp\left(-\frac{\theta^2}{4\pi^2}\right), \quad \theta \in [-\pi, \pi],$$

which is positive and in the Wiener class. In [30], the system (4.19) is preconditioned by $C \otimes C \otimes \cdots \otimes C$ with different circulant preconditioners C discussed in §2.1. Numerical results in [30] show that for solving block Toeplitz systems of order 65536, the time required by using the circulant preconditioned conjugate gradient method is half the time required by the direct Toeplitz solver available in the IMSL package.

4.5.2 Wiener-Hopf Equations

Half-line Wiener-Hopf integral equations

$$y(t) + \int_0^\infty a(t-s)y(s)ds = g(t), \quad 0 \leq t < \infty, \quad (4.20)$$

in which $a(t) \in L_1(\mathbb{R})$ and $g(t) \in L_2[0, \infty)$ are given functions, arise in a variety of practical applications in mathematics and engineering, such as linear prediction problems for stationary stochastic processes [90, pp.145–146], diffusion problems and scattering problems [90, pp.186–189].

One way of solving (4.20) is by the projection method [85], where the solution $y(t)$ of (4.20) is approximated by the solution $y_\tau(t)$ of the finite-section equation

$$y_\tau(t) + \int_0^\tau a(t-s)y_\tau(s)ds = g(t), \quad 0 \leq t \leq \tau. \quad (4.21)$$

It is shown in [85, Theorem 3.1] that

$$\lim_{\tau \rightarrow \infty} \|y_\tau - y\|_{L^p[0, \tau]} = 0, \quad 1 \leq p < \infty.$$

The finite-section equation (4.21) can be solved numerically by either direct or iterative methods. For a fixed τ , the finite-section operator \mathcal{A}_τ defined by

$$(\mathcal{A}_\tau x)(t) = \begin{cases} \int_0^\tau a(t-s)x(s)ds, & 0 \leq t \leq \tau, \\ 0, & t > \tau. \end{cases} \quad (4.22)$$

is a compact operator. Therefore, the spectrum of the Wiener-Hopf operator $\mathcal{I} + \mathcal{A}_\tau$ (where \mathcal{I} is the identity operator) is clustered around 1, and hence solving (4.21) by iterative methods such as the conjugate gradient method is less expensive than direct methods.

However, as τ tends to ∞ , the spectrum of this finite-section operator \mathcal{A}_τ becomes dense in the spectrum of the half-line operator defined by (4.20), and hence the convergence rate of the conjugate gradient method deteriorates, see the numerical results in [86] for instance. A standard way of speeding up the convergence rate of the conjugate gradient method is to apply a preconditioner. Thus, instead of solving (4.21), one solves a preconditioned operator equation. We remark that there is a close relation between Wiener-Hopf integral equations and semi-infinite Toeplitz operators, see Gohberg and Fel'dman [85, p.5]. The methodology of preconditioning “discrete” Toeplitz systems can be modified and applied to the “continuous” Wiener-Hopf equations.

In [86], Gohberg, Hanke, and Koltracht proposed using *circulant integral operators* to precondition (4.21). Circulant integral operators are operators of the form

$$(\mathcal{C}_\tau x)(t) = \int_0^\tau c_\tau(t-s)x(s)ds, \quad 0 \leq t \leq \tau, \quad (4.23)$$

where c_τ is a τ -periodic conjugate symmetric function in $L_1[-\tau, \tau]$, i.e.,

$$c_\tau(t + \tau) = c_\tau(t) \quad \text{and} \quad c_\tau(-t) = \overline{c_\tau(t)}, \quad \forall t \in [-\tau, \tau].$$

We remark that \mathcal{C}_τ is a compact, self-adjoint operator on $L_2[-\tau, \tau]$. The preconditioned equation is given by

$$(\mathcal{I} + \mathcal{C}_\tau)^{-1}(\mathcal{I} + \mathcal{A}_\tau)y_\tau(t) = (\mathcal{I} + \mathcal{C}_\tau)^{-1}g(t), \quad 0 \leq t \leq \tau. \quad (4.24)$$

It has been shown in [86] that for large τ , the spectra of the circulant preconditioned operators $(\mathcal{I} + \mathcal{C}_\tau)^{-1}(\mathcal{I} + \mathcal{A}_\tau)$ are clustered around 1.

Theorem 25 (Gohberg, Hanke, and Koltracht (1994) [86]) *If*

$$\lim_{\tau \rightarrow \infty} \|a - c_\tau\|_{L_1[-\tau/2, \tau/2]} = 0, \quad (4.25)$$

then for any given $\epsilon > 0$, there exist a positive integer ρ and a number $\tau^ > 0$ such that for all $\tau > \tau^*$, $\mathcal{I} + \mathcal{A}_\tau$ are positive definite, and the spectra of $(\mathcal{I} + \mathcal{C}_\tau)^{-1/2}(\mathcal{I} + \mathcal{A}_\tau)(\mathcal{I} + \mathcal{C}_\tau)^{-1/2}$ have at most ρ eigenvalues outside the interval $(1 - \epsilon, 1 + \epsilon)$.*

Thus under the assumption (4.25), the preconditioned conjugate gradient method converges superlinearly for large τ . In addition, Gohberg, Hanke, and Koltracht [86] proposed “wrap-round” and “optimal” circulant integral operators for Wiener-Hopf equations.

1) “Wrap-round” Circulant Integral Operator \mathcal{S}_τ

Given the operator \mathcal{A}_τ as in (4.22), it is defined as

$$(\mathcal{S}_\tau x)(t) = \begin{cases} \int_0^\tau s_\tau(t-s)x(s)ds, & 0 \leq t \leq \tau, \\ 0, & t > \tau. \end{cases} \quad (4.26)$$

Here the function s_τ is a τ -periodic function defined by

$$s_\tau(t) = a(t), \quad -\tau/2 \leq t \leq \tau/2.$$

2) “Optimal” Circulant Integral Operator \mathcal{F}_τ

Given the operator \mathcal{A}_τ , it is defined as

$$(\mathcal{F}_\tau x)(t) = \begin{cases} \int_0^\tau f_\tau(t-s)x(s)ds, & 0 \leq t \leq \tau, \\ 0, & t > \tau. \end{cases}$$

Here $f_\tau(t)$ is a τ -periodic function defined by

$$f_\tau(t) = \left(\frac{\tau-t}{\tau}\right)a(t) + \left(\frac{t}{\tau}\right)a(t-\tau), \quad 0 \leq t \leq \tau.$$

Gohberg et al. [86] showed that \mathcal{F}_τ minimizes the Hilbert-Schmidt norm

$$\|\mathcal{A}_\tau - \mathcal{C}_\tau\|^2 \equiv \int_0^\tau \int_0^\tau |a(t-s) - c_\tau(t-s)|^2 ds dt,$$

over all circulant integral operators \mathcal{C}_τ .

We note that the above choices of \mathcal{S}_τ and \mathcal{F}_τ are continuous analogs of Strang’s and T. Chan’s circulant preconditioners, respectively. In §2.2, a unified approach of constructing circulant matrix preconditioners for finite Toeplitz matrices is derived from the viewpoint of convolution products. Using the same viewpoint, R. Chan, Jin, and Ng [37] have derived an easy and general scheme for constructing circulant integral preconditioners for Wiener-Hopf equations. For ease of presentation, let us denote by \hat{q} the Fourier transform of any function q . The first step is to relate $c_\tau(t)$ in (4.23) to a sequence of conjugate symmetric functions $\{C_\tau(t)\}_\tau$, i.e.,

$$c_\tau(t) = \begin{cases} C_\tau(t)a(t) + C_\tau(t-\tau)a(t-\tau), & 0 \leq t \leq \tau, \\ C_\tau(t+\tau)a(t+\tau) + C_\tau(t)a(t), & -\tau \leq t \leq 0, \end{cases} \quad (4.27)$$

and translate the convergence requirement (4.25) on $c_\tau(t)$ to conditions on $\{C_\tau(t)\}_\tau$. Basically, $C_\tau(t)$ should be such that the convolution product $\hat{C}_\tau * \hat{a}$ converge uniformly to \hat{a} on \mathbb{R} .

As examples, the “wrap-round” circulant integral operator \mathcal{S}_τ can be constructed by setting C_τ in (4.27) to be the function

$$S_\tau(t) = \begin{cases} 1, & |t| \leq \tau/2, \\ 0, & |t| > \tau/2. \end{cases}$$

The operator \mathcal{F}_τ can be also be derived by setting C_τ in (4.27) to be the function

$$F_\tau(t) = \begin{cases} \frac{\tau - |t|}{\tau}, & |t| \leq \tau, \\ 0, & |t| > \tau. \end{cases}$$

Moreover, R. Chan, Jin, and Ng proved that for sufficiently large τ , if $\{C_\tau\}_\tau$ is uniformly bounded on the real line and the convolution product of \hat{C}_τ with $\hat{a}(\omega)$ converges to $\hat{a}(\omega)$ uniformly on \mathbb{R} , then the spectra of the circulant preconditioned operators are clustered around 1. They also showed that $\{C_\tau(t)\}_\tau$ can be derived easily from the Dirac delta function or from *approximate convolution identities* commonly used in Fourier analysis [190].

As in the Toeplitz matrix case, there are other ways of constructing operators as preconditioners for (4.21), see Ng and Lin [150], Ng, Lin, and R. Chan [151] and R. Chan and Lin [41]. In R. Chan and Lin [40], optimal and super-optimal circulant integral preconditioners are constructed for general integral equations of the second kind,

$$y(t) + \int_0^\infty a(t, s)y(s)ds = g(t), \quad 0 \leq t < \infty.$$

Here $a(t, s)$ is not necessarily a convolution kernel.

When (4.21) is discretized with the rectangular quadrature rule, we get an n -by- n matrix system $\mathbf{A}\mathbf{y} = \mathbf{g}$, where A is a Toeplitz matrix. Here n is the number of quadrature points used in the discretization. We note that if the rectangular quadrature rule is used to discretize (4.24), then we get a matrix system:

$$(I + C)^{-1}(I + A)\mathbf{y} = (I + C)^{-1}\mathbf{g}, \tag{4.28}$$

where the matrices $(I + C)$ and $(I + A)$ are n -by- n circulant and Toeplitz matrices, respectively. We see that (4.28) is basically a circulant-preconditioned Toeplitz system, which requires only $O(n \log n)$ operations in each iteration.

One main drawback of using the rectangular rule is that the order of accuracy of the discretized solution \mathbf{y} depends only linearly on the number of quadrature points. Thus, in order to obtain a reasonably accurate solution for (4.21), a small step-size has to be used and hence the dimension of the resulting matrix system will be large. In order to obtain high order of accuracy, one can use higher order quadrature rules such as the trapezoidal rule or Simpson’s rule, which have second and fourth order of accuracy, respectively. In these cases, the discretization

matrices of the Wiener-Hopf integral operators are *non-Toeplitz* matrices. Moreover, the corresponding discretization matrices of the circulant integral operators are in general *not* circulant, and therefore their inversion cannot be computed by using FFT. Hence, the cost per iteration of the preconditioned conjugate gradient method exceeds $O(n \log n)$ operations.

Instead of constructing matrix preconditioners for these discretization matrices, we can consider preconditioners from the operator point of view. In [135], Lin, Ng, and R. Chan constructed preconditioners \mathcal{B}_τ for (4.21) such that the preconditioned operators $(\mathcal{I} - \mathcal{B}_\tau)(\mathcal{I} + \mathcal{A}_\tau)$ have clustered spectra and that only $O(n \log n)$ operations are required in each iteration of the preconditioned conjugate gradient method, even when higher order quadrature rules are employed. Their idea is to use convolution operators rather than circulant operators to precondition $\mathcal{I} + \mathcal{A}_\tau$. More precisely, given \mathcal{A}_τ , they constructed the convolution operator \mathcal{B}_τ , whose kernel function $b(t)$ satisfies

$$\hat{b}(\omega) = \frac{\hat{a}(\omega)}{1 + \hat{a}(\omega)}. \quad (4.29)$$

Then they precondition (4.21) as

$$(\mathcal{I} - \mathcal{B}_\tau)(\mathcal{I} + \mathcal{A}_\tau)y_\tau(t) = (\mathcal{I} - \mathcal{B}_\tau)g(t), \quad 0 \leq t \leq \tau. \quad (4.30)$$

Theorem 26 (Lin, Ng, and R. Chan [135]) *Let $a(t) \in L_1(\mathbb{R})$ be conjugate symmetric and $\hat{a}(\omega) \geq 0$. Then for any given $\epsilon > 0$, there exist a positive integer ρ and a real number $\tau^* > 0$ such that for all $\tau > \tau^*$, the spectrum of $(\mathcal{I} - \mathcal{B}_\tau)(\mathcal{I} + \mathcal{A}_\tau)$ has at most ρ eigenvalues outside the interval $(1 - \epsilon, 1 + \epsilon)$.*

According to this theorem, if we apply the conjugate gradient method to solve (4.30), the convergence rate will be superlinear for large τ .

In the following, we test the effectiveness of different integral operator preconditioners by using the kernel function $a(t) = (1 + t^2)^{-1}/\mu$. In practical applications, the parameter μ is the regularization parameter and is usually a small positive number. In the test, we arbitrarily set $\mu = 0.01$. To discretize (4.29) we partition the interval $[-\tau, \tau]$ into $2n$ equal subintervals of step-size h , and then compute $\hat{a}(\omega)$ and $\hat{b}(\omega)$ by using the formula

$$\hat{a}(\omega) \approx h \sum_{k=-n}^n a(kh)e^{-i\omega kh},$$

which can be computed easily by FFTs. The right hand side function $g(t)$ is chosen such that the corresponding solution for the Wiener-Hopf equation (4.20) is

$$y(t) = \begin{cases} (16 - t)^2, & 0 \leq t \leq 16, \\ 0, & t > 16. \end{cases} \quad (4.31)$$

We use the same random vector as initial guess for all preconditioners, and the tolerance is set to 10^{-6} . All computations are done by Matlab. We define the error of the numerical solution:

$$e \equiv \left\{ h \sum_{j=0}^n |\tilde{y}(jh) - y(jh)|^2 \right\}^{1/2} \approx \left\{ \int_0^\tau |\tilde{y}(t) - y(t)|^2 dt \right\}^{1/2},$$

where $\{\tilde{y}(jh)\}_{j=0}^n$ is the computed solution and $y(t)$ is the true solution given by (4.31).

Table 6 gives the numbers of mega-flops used to achieve a given accuracy ϵ (i.e., the error of the numerical solution $e \leq \epsilon$) for different quadrature rules and preconditioners. The symbol B denotes that we are solving (4.30) with the quadrature rule listed. The symbol S denotes that we are solving (4.24) with the “wrap-round” circulant integral operator \mathcal{S}_τ defined in (4.26), and that \mathcal{S}_τ is discretized by the rectangular rule (so as to make it a circulant matrix) but the operator \mathcal{A}_τ is still discretized according to the rule listed in the table. The discretization matrix for \mathcal{A}_τ will be a product of a Toeplitz matrix and a diagonal matrix, where the diagonal matrix depends on the discretization rule used. The symbol I represents that (4.21) is solved without using any preconditioner and discretized according to the rule listed. In the table, ** denotes that the corresponding number exceeds 1000.

Accuracy	Rectangular			Trapezoidal			Simpson's		
	B	S	I	B	S	I	B	S	I
10^0	65.1	61.2	492.5	1.49	1.69	5.50	0.50	5.15	5.76
10^{-1}	**	**	**	5.62	6.09	51.28	0.84	9.27	9.89
10^{-2}	**	**	**	16.94	18.35	157.51	1.18	13.44	14.48
10^{-3}	**	**	**	89.48	98.62	930.19	2.68	27.04	29.21
10^{-4}	**	**	**	318.02	332.33	**	5.61	57.49	61.88
10^{-5}	**	**	**	**	**	**	11.64	109.98	116.66
10^{-6}	**	**	**	**	**	**	28.18	270.56	285.09

Table 6: Number of Mega-flops for Different Quadrature Rules and Preconditioners

From the table, we see that without using any preconditioner, the number of flops for achieving the required accuracy is very large. We note that for the trapezoidal rule, the Toeplitz structure is disturbed only in the two rows corresponding to the boundary. Therefore, the circulant preconditioner works well for the trapezoidal case. However, it does not work well if \mathcal{A}_τ is discretized using Simpson's rule. The performance of our proposed preconditioner for Simpson's rule is the best one in terms of the accuracy and the computational work. The accuracy of the computed solution depends only on the quadrature rule used in discretizing \mathcal{A}_τ . However, the convergence rate of the preconditioned systems and the cost per iteration of the PCG method depend on how we discretize the preconditioning operators. Therefore, it is advantageous to use a higher order quadrature rule to discretize the operator equation because of the accuracy concern. But to speed up the convergence rate of the method and to minimize the cost per iteration, one needs to use our proposed preconditioner rather than circulant ones.

4.6 Applications to Time Series Analysis

4.6.1 Finite Impulse Response Filters

Least squares estimations have been used extensively in a wide variety of scientific applications, such as equalization [84, p.139], system identification [140, 141, 142], adaptive signal processing [3, 170, 99], active noise control [59], and speech processing [84, p.343]. In these applications, we usually need to estimate the transmitted signal from a sequence of received signal samples or to model an unknown system by using a linear system model.

To present the problem properly, let us introduce some terminology used in signal processing. Let $x(i)$ be a discrete-time stationary zero-mean complex-valued process (see Fuller [82, pp. 10-11]). A *finite impulse response* (FIR) linear filter of order n is of the form

$$\tilde{d}(i) = \sum_{k=1}^n b_k x(i - k + 1),$$

where $\tilde{d}(i)$ is the filter output based on the data $\{x(k)\}_{k=i}^{i-n+1}$ and $\{b_k\}_{k=1}^n$ are the impulse responses of the filter. The difference between the desired response $d(i)$ of the process and the filter output $\tilde{d}(i)$ is called the estimation error of order n . Since we are interested in estimating the desired response based on the input measurements, the impulse responses $\{b_k\}_{k=1}^n$ should be chosen to make the estimation error as small as possible.

For the case of known statistics, i.e., the autocovariances of the stationary process are known, the optimal least mean squares predictor coefficients $\{b_k\}_{k=1}^n$ are given by the solution of the linear system of equations

$$A\mathbf{b} = \mathbf{d}, \tag{4.32}$$

see Giordano and Hsu [84, pp.41-43]. Here A is an n -by- n Hermitian Toeplitz matrix and its entries $\{a_j\}_{j=0}^{n-1}$ are the autocovariances of a discrete-time stationary process and are given by

$$a_k = \mathcal{E}[x(j)\overline{x(j-k)}],$$

where \mathcal{E} is the expectation operator.

The matrix A is called the covariance matrix of the stationary process, and the Toeplitz system (4.32) is commonly called the Yule-Walker equations, see Yule [197].

For a discrete-time stationary process, if the autocovariances of the process are absolutely summable, i.e., $\sum_{k=-\infty}^{\infty} |a_k| < \infty$, then the function $f(\theta)$ with a_k as Fourier coefficients is called the spectral density function of the stationary process, see [21, p.118]. The covariance matrix A is then a Toeplitz matrix generated by $f(\theta)$. As examples, we consider the following stationary processes:

1. *First order auto-regressive process AR(1)* [160, p.238]: The process is given by

$$x(t) = \rho x(t-1) + v(t),$$

where $\{v(t)\}$ is a white noise process with variance η^2 . The autocovariances of the process are given by

$$a_k = \frac{\eta^2 \rho^{|k|}}{1 - \rho^2}, \quad k = 0, \pm 1, \pm 2, \dots,$$

where $|\rho| < 1$. The corresponding spectral density function $f(\theta)$ is given by

$$f(\theta) = \frac{\eta^2}{2\pi(1 - 2\rho \cos \theta + \rho^2)}. \quad (4.33)$$

The covariance matrix is a scalar multiple of the Kac-Murdock Szegő matrix, see Kac et al. [122].

2. *First order moving-average process MA(1)* [21, p.123]: The process is given by

$$x(t) = v(t) + \chi v(t - 1),$$

where $|\chi| < 1$ and $\{v(t)\}$ is a white noise process with variance η^2 . The autocovariances of the process are given by

$$a_k = \begin{cases} \eta^2(1 + \chi^2), & k = 0, \\ \eta^2\chi, & k = 1, \\ 0, & \text{otherwise.} \end{cases}$$

We see that the covariance matrix is a tridiagonal Toeplitz matrix with

$$f(\theta) = \frac{\eta^2}{2\pi}(1 + 2\chi \cos \theta + \chi^2).$$

If one assumes that the spectral density function of the stationary process exists and is positive (that can be guaranteed by the causality of the process [21, p.85]), then the Yule-Walker equations (4.32) can be solved in $O(n \log n)$ operations by using the preconditioned conjugate gradient method with circulant preconditioners discussed in §2.1, see [149]. We remark that all the above results are derived deterministically. In the least squares estimation algorithms, we always deal with data samples from random processes and the convergence rate should be considered in a probabilistic way.

4.6.2 Least Squares Filters

We note that in practical cases, no prior knowledge is usually available on the autocovariances of the process. If m data samples have been taken, then all the information we have is contained in the finite number of data points $\{x(t)\}_{t=1}^m$. In this case, one can still formulate a well-defined least squares prediction problem by estimating the autocovariances from the data samples $\{x(t)\}_{t=1}^m$ with various types of windowing methods; such as the *correlation*, *covariance*, *pre-windowed*, and

post-windowed methods, see for instance Giordano and Hsu [84, pp.65-66]. The least squares estimators can then be found by solving for the n -vector \mathbf{b} in

$$\min \|\mathbf{d} - A\mathbf{b}\|_2. \quad (4.34)$$

Here A is an m -by- n Toeplitz matrix with full column rank n , obtained by applying various types of windowing methods on the data samples $\{x(t)\}_{t=1}^m$. The solution \mathbf{b} of (4.34) can be obtained by solving the normal equation $(A^*A)\mathbf{b} = A^*\mathbf{d}$. We note that if the correlation method is employed, the normal matrix A^*A is Toeplitz. The other three windowing methods lead to non-Toeplitz normal matrices A^*A of the form

$$A^*A = \tilde{A} - L^*L - U^*U, \quad (4.35)$$

where \tilde{A} is an n -by- n Toeplitz matrix, L and U are lower triangular and upper triangular Toeplitz matrices, respectively (cf. (4.15)).

To prove convergence, the following practical signal processing assumptions of the random process were made in [149, 148].

- The process is stationary with constant mean μ .
- The spectral density function of the process is positive and in the Wiener class.
- There exist positive constants β_1 and β_2 such that

$$\text{Var} \left(\frac{1}{m} \sum_{j=1}^{m-k} x(j) \right) \leq \frac{\beta_1}{m}, \quad k = 0, 1, 2, \dots, m-1$$

and

$$\text{Var} \left(\frac{1}{m} \sum_{j=1}^{m-k} [x(j) - \mu][x(j+k) - \mu] \right) \leq \frac{\beta_2}{m}, \quad k = 0, 1, 2, \dots, m-1.$$

We note that the positiveness of the spectral density function can be guaranteed by the causality of the process [21, p.85], whereas the absolute summability of the autocovariances can be assured by the invertibility of the process [21, p.86]. With these assumptions, we proved that the spectra of the preconditioned matrices $c(\tilde{A})^{-1}(A^*A)$ are clustered around 1 with probability 1, provided that a sufficiently large number of data samples are taken.

Theorem 27 (Ng and R. Chan (1994) [149]) *Let the discrete-time process satisfy the above assumptions. Then for any given $\epsilon > 0$ and $0 < \eta < 1$, there exist positive integers ρ_1 and ρ_2 such that for $n > \rho_1$, the probability that at most ρ_2 eigenvalues of the matrix $I - c(\tilde{A})^{-1}(A^*A)$ have absolute value greater than ϵ is greater than $1 - \eta$, provided that $m = O(n^{3+\nu})$ with $\nu > 0$.*

Hence, when we apply the conjugate gradient method to the preconditioned system, the method converges superlinearly with probability 1. Since the data matrix A is an m -by- n rectangular Toeplitz matrix, the normal equations and the circulant preconditioner can be formed in $O(m \log n)$ operations. Once they are formed, the cost per iteration of the preconditioned conjugate gradient method is $O(n \log n)$ operations. Therefore, the total work of obtaining the predictor coefficients to a given accuracy is of $O((m+n) \log n)$.

4.6.3 Linear-Phase Least Squares Filters

Besides FIR linear filters, FIR linear-phase filters are also commonly used in signal processing. Such filters are especially important for applications where frequency dispersion due to non-linear phase is harmful, such as in speech processing. In this case, the impulse responses can be found by solving the Toeplitz-plus-Hankel least squares problem

$$\min \|\mathbf{d} - (A + H)\mathbf{b}\|_2,$$

where $A + H$ is a rectangular Toeplitz-plus-Hankel matrix, see [140, 141, 142, 194, 109, 147]. By exploiting the structure of the normal equations matrices, it can be written as

$$(A + H)^*(A + H) = \tilde{A} + \tilde{H} - V_1 - V_2 - V_3 - V_4,$$

where \tilde{A} is a Toeplitz matrix, \tilde{H} is a Hankel matrix, and $\{V_i\}_{i=1}^4$ are non-Toeplitz and non-Hankel matrices. In [147], $c(\tilde{A})$ is used as a preconditioner for the problem. Under the same assumptions as in Theorem 27, it has been shown that the spectra of the matrices \tilde{H} and $\{V_i\}_{i=1}^4$ are all clustered around 0. Hence, the spectrum of the preconditioned matrix $c(\tilde{A})^{-1}(A + H)^*(A + H)$ will also be clustered around 1 with probability 1.

Ku and Kuo [133] have also proposed a preconditioner for Toeplitz-plus-Hankel systems. Their approach is basically to take circulant approximations of the Toeplitz matrix and the Hankel matrix separately and then combine them together to form a preconditioner. The motivation behind Ng's preconditioner is that the Toeplitz matrix \tilde{A} is the sample autocorrelation matrix, which intuitively should be a good estimation to the autocorrelation matrix of the discrete-time stationary process, provided that a sufficiently large number of data samples are taken. Moreover, under practical signal processing assumptions, the spectrum of the Hankel matrix \tilde{H} is clustered around zero. Hence, it suffices to approximate the Toeplitz part by circulant preconditioners.

4.6.4 Recursive Least Squares Filters

So far we have discussed only block-processing type least squares estimations, i.e., data samples are collected over a finite time interval and the resulting linear system is solved. Recently, Plemmons [159] and Ng and Plemmons [152, 153] proposed to use circulant preconditioners for the recursive least squares (RLS) estimations, which are used extensively in many signal processing and control applications. They considered the RLS computations where the data matrices are assumed to have a Toeplitz (displacement) structure. Their new algorithm computes least squares estimators recursively by using sliding data windows involving multiple (rank k) updating and downdating computations for superior tracking capabilities. When $A(t)$ is an m -by- n rectangular data matrix (m is the length of the sliding window) with full column rank, then the least squares estimator $\mathbf{b}(t)$ at step t can be obtained by solving the normal equations

$$A(t)^* A(t)\mathbf{b}(t) = A(t)^* \mathbf{d}(t).$$

We note that although $A(t)^*A(t)$ is generally not Toeplitz, it can still be written in the form

$$A(t)^*A(t) = \tilde{A}(t) - L(t)^*L(t) - U(t)^*U(t),$$

where $\tilde{A}(t)$ is Toeplitz, and $L(t)$ and $U(t)$ are lower triangular and upper triangular Toeplitz matrices, (cf. (4.15) and (4.35)). In [152, 153], Ng and Plemmons employed the preconditioned conjugate gradient method with circulant preconditioners to solve such systems at each step. In the case of point-processing ($k = 1$), the method requires $O(n \log n)$ operations per adaptive filter input, where n is the number of least squares estimators. In the case of block-processing ($k \geq n$), the method requires only $O(\log n)$ operations per adaptive filter input. These FFT-based iterative RLS algorithms, with reasonable complexity for computing least squares estimators recursively, may also avoid some of the instability problems associated with direct fast RLS methods.

In the following, we test the convergence performance of the FFT-based sliding window RLS algorithm and compare it with fast transversal filter algorithms [99, pp.586–600] and standard recursive least squares [99, p.485]. We remark that they are $O(n)$ and $O(n^2)$ algorithms, respectively. An exponential weighting factor γ is generally used in these two algorithms. The inverse of $1 - \gamma$ is approximately a “measure” of the memory of the algorithm. Therefore, the length of the sliding window m used in the FFT-based sliding window RLS algorithm is related to γ by $m \approx 1/(1 - \gamma)$.

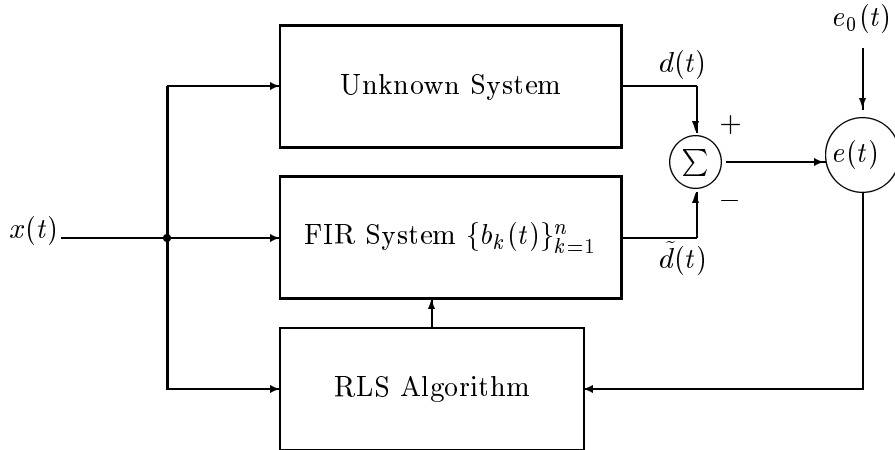


Figure 6: Adaptive FIR System Identification Model

Figure 6 is a block diagram for the implementation of the FFT-based RLS algorithm in an adaptive FIR system identification model. The input signal $x(t)$ drives the unknown system to produce the output sequence $d(t)$. We model the unknown system as an FIR filter. As input stochastic process, we used the first order autoregressive AR(1) process given by

$$x(t) + \rho x(t - 1) = v(t),$$

where $\rho = -0.9999$ and $\{v(t)\}$ is a white noise process with variance $\eta^2 = 1$. We note that the spectral density function of the process is positive and in the Wiener class, see (4.33). The reference (unknown) system is an n -th order linear phase FIR filter with uncorrelated Gaussian white noise added. The FIR's impulse responses $\{\hat{b}_k\}_{k=1}^n$ of the unknown system used are

$$\hat{b}_k = 1.1 - \frac{|2k - n - 1|}{n - 1}, \quad k = 1, 2, \dots, n.$$

In the numerical tests, the stopping criterion for the preconditioned conjugate gradient method is when the ℓ_2 norm of the residual vector is less than 10^{-7} . All the computations are done by Matlab.

Figures 7–9 show the prior average least squares error and the ℓ_2 norm of the weight-error vector for different algorithms when different levels of variances of Gaussian white noise are added to the reference system. The size n of the filters used is 32. We see from the figures that the fast transversal filter algorithm does not converge. However, both the standard recursive least squares and the FFT-based sliding window algorithms converge very fast.




Figure 7: Fast Transversal Filter Algorithm with $\gamma = 0.9922$.

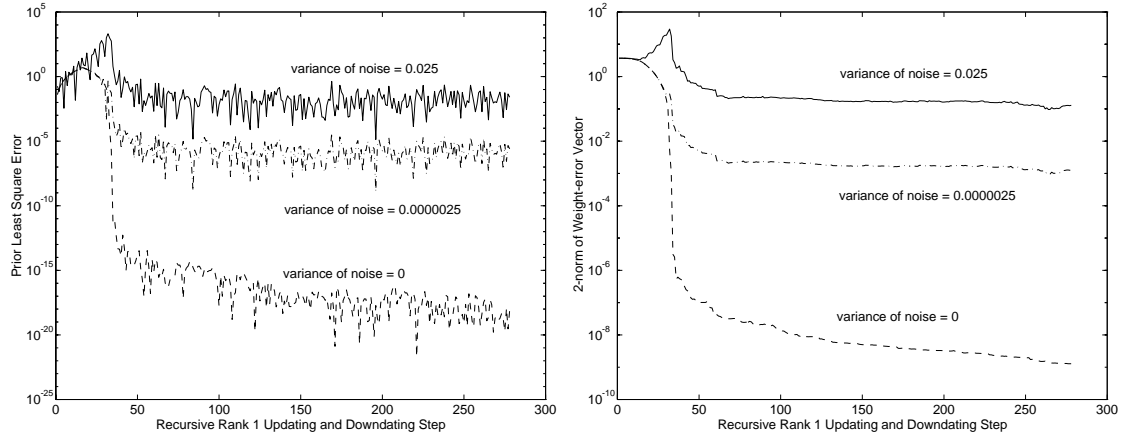


Figure 8: Standard Recursive Least Squares Algorithm with $\gamma = 0.9922$.




Figure 9: FFT-based Sliding Window RLS Algorithm with $m = 128$.

5 Concluding Remarks

In 1986, Strang addressed the question of whether iterative methods can compete with direct methods for solving symmetric positive definite Toeplitz systems. The answer has turned out to be an unqualified yes. The conjugate gradient method coupled with a suitable preconditioner can solve a large class of n -by- n Toeplitz systems in $O(n \log n)$ operations, as compared to the $O(n \log^2 n)$ operations required by fast direct Toeplitz solvers. This paper summarizes some of the developments of this iterative method in the past few years. Applications of the method to partial differential equations, queueing networks, integral equations, image restoration, and time series analysis are also given. The results show that the method in some instances works better than traditional methods used specifically for these problems.

Possible directions of future works are to apply the method to other areas where solutions of Toeplitz or Toeplitz-related systems are sought, such as in control theory and image processing. For instance, the restoration of images in nonlinear space-invariant systems involves the solutions of Toeplitz-like systems, see [178]; and the total variation method for restoring noisy, blurred images involves operators which are the sum of second-order elliptic operators and Toeplitz operators [163, 189]. Also in the numerical solution of the biharmonic equation by conformal mapping, we have to solve Toeplitz-plus-Hankel systems [34].

Other directions consist of a thorough comparison of different preconditioners proposed and more importantly with fast direct Toeplitz solvers (cf. [112, 136]); and the use of other iterative methods, such as multigrid methods, for Toeplitz systems, see [74, 75].

Acknowledgments: We are deeply grateful to Professors Tony Chan and Robert Plemmons for their help in the preparation of this paper. Also we would like to thank the referees for their valuable comments, which made the paper more updated and presentable.

References

- [1] J. Abbiss and P. Earwicker, *Compact Operator Equations, Regularization and Super-resolution*, in Mathematics in Signal Processing, T. Durrani et al., eds., Clarendon Press, Oxford, 1987.
- [2] S. Akl, *The Design and Analysis of Parallel Algorithms*, Prentice-Hall, Englewood Cliffs, NJ, 1989.
- [3] T. Alexander, *Adaptive Signal Processing, Theory and Applications*, Springer-Verlag, New York, 1986.
- [4] G. Ammar and W. Gragg, *Superfast Solution of Real Positive Definite Toeplitz Systems*, SIAM J. Matrix Anal. Appl., 9 (1988), pp. 61–76.
- [5] H. Andrews and B. Hunt, *Digital Image Restoration*, Prentice-Hall, Englewood Cliffs, NJ, 1977.
- [6] P. Anselone, *Collectively Compact Operator Approximation Theory and Applications to Integral Equations*, Prentice-Hall, Englewood Cliffs, NJ, 1971.
- [7] O. Axelsson and V. Barker, *Finite Element Solution of Boundary Value Problems, Theory and Computation*, Academic Press, Orlando, FL, 1984.
- [8] G. Baxter, *Polynomials Defined By a Difference System*, J. Math. Anal. Appl., 2 (1961), pp. 223–263.
- [9] F. Di Benedetto, *Analysis of Preconditioning Techniques for Ill-Conditioned Toeplitz Matrices*, SIAM J. Sci. Comput., 16 (1995), pp. 682–697.
- [10] F. Di Benedetto, *Preconditioning of Block Toeplitz Matrices by Sine Transforms*, SIAM J. Sci. Comput., to appear.
- [11] F. Di Benedetto, G. Fiorentino, and S. Serra, *C.G. Preconditioning for Toeplitz Matrices*, Computers Math. Applic., 25 (1993), No. 6, pp. 35–45.
- [12] D. Bertsekas and J. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*, Prentice-Hall, Englewood Cliffs, NJ, 1989.
- [13] J. Biemond, R. Lagendijk, and R. Mersereau, *Iterative Methods for Image Deblurring*, in Proc. IEEE, 78 (1990), pp. 856–883.
- [14] D. Bini and P. Favati, *On a Matrix Algebra Related to the Discrete Hartley Transform*, SIAM J. Matrix Anal. Appl., 14 (1993), pp. 500–507.
- [15] D. Bini and F. Di Benedetto, *A New Preconditioner for the Parallel Solution of Positive Definite Toeplitz Systems*, in Proc. Second ACM Symp. on Parallel Algorithms and Architectures, Crete, Greece, 1990, pp. 220–223.

- [16] R. Bitmead and B. Anderson, *Asymptotically Fast Solution of Toeplitz and Related Systems of Linear Equations*, Linear Algebra Appl., 34 (1980), pp. 103–116.
- [17] Å. Björck, *Least Squares Methods*, in Handbook of Numerical Analysis, Vol. 1, P. Ciarlet and J. Lions, eds., Elsevier, Amsterdam, 1989.
- [18] A. Bojanczyk, R. Brent, and F. de Hoog, *QR Factorization of Toeplitz Matrices*, Numer. Math., 49 (1986), pp. 81–94.
- [19] E. Boman and I. Koltracht, *Fast Transform Based Preconditioners for Toeplitz Equations*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 628–645.
- [20] R. Brent, F. Gustavson, and D. Yun, *Fast Solution of Toeplitz Systems of Equations and Computation of Padé Approximants*, J. Algo., 1 (1980), pp. 259–295.
- [21] P. Brockwell and R. Davis, *Time Series: Theory and Methods*, 2nd ed. Springer-Verlag, New York, 1991.
- [22] A. Buckley, *On the Solution of Certain Skew Symmetric Linear Systems*, SIAM J. Numer. Anal., 14 (1977), pp. 566–570.
- [23] J. Bunch, *Stability of Methods for Solving Toeplitz Systems of Equations*, SIAM J. Sci. Stat. Comput., 6 (1985), pp. 349–364.
- [24] G. Carayannis, N. Kalouptsidis, and D. Manolakis, *Fast Recursive Algorithms for a Class of Linear Equations*, IEEE Trans. Acoust., Speech, Signal Process., 30 (1982), pp. 227–239.
- [25] R. Carreras, *Personal Correspondence*, 1993.
- [26] R. Chan, *Iterative Methods for Overflow Queueing Models I*, Numer. Math., 51 (1987), pp. 143–180.
- [27] R. Chan, *Circulant Preconditioners for Hermitian Toeplitz Systems*, SIAM J. Matrix Anal. Appl., 10 (1989), pp. 542–550.
- [28] R. Chan, *The Spectrum of a Family of Circulant Preconditioned Toeplitz Systems*, SIAM J. Numer. Anal., 26 (1989), pp. 503–506.
- [29] R. Chan, *Toeplitz Preconditioners for Toeplitz Systems with Nonnegative Generating Functions*, IMA J. Numer. Anal., 11 (1991), pp. 333–345.
- [30] R. Chan, *Numerical Solutions for the Inverse Heat Problems in \mathbb{R}^N* , SEA Bull. Math., 16 (1992), pp. 97–105.
- [31] R. Chan and T. Chan, *Circulant Preconditioners for Elliptic Problems*, Numer. Linear Algebra Appl., 1 (1992), pp. 77–101.

- [32] R. Chan and W. Ching, *Toeplitz-Circulant Preconditioners for Toeplitz Systems and Their Applications to Queueing Networks with Batch Arrivals*, SIAM J. Sci. Comput., to appear.
- [33] R. Chan, W. Ching, and C. Wong, *Optimal Trigonometric Preconditioners for Elliptic Problems and Queueing Problems*, SEA Bull. Math., 20 (1996), to appear.
- [34] R. Chan, T. DeLillo, and M. Horn, *The Numerical Solution of the Biharmonic Equation by Conformal Mapping*, preprint, 1995.
- [35] R. Chan and X. Jin, *Circulant and Skew-circulant Preconditioners for Skew-Hermitian Type Toeplitz Systems*, BIT, 31 (1991), pp. 632–646.
- [36] R. Chan and X. Jin, *A Family of Block Preconditioners for Block Systems*, SIAM J. Sci. Stat. Comput., 13 (1992), pp. 1218–1235.
- [37] R. Chan, X. Jin, and M. Ng, *Circulant Integral Operators as Preconditioners for Wiener-Hopf Equations*, Integr. Equat. Oper. Theory, 21 (1995), pp. 12–23.
- [38] R. Chan, X. Jin, and M. Yeung, *The Circulant Operator in the Banach Algebra of Matrices*, Linear Algebra Appl., 149 (1991), pp. 41–53.
- [39] R. Chan, X. Jin, and M. Yeung, *The Spectra of Super-Optimal Circulant Preconditioned Toeplitz Systems*, SIAM J. Numer. Anal., 28 (1991), pp. 871–879.
- [40] R. Chan and F. Lin, *Preconditioned Conjugate Gradient Methods for Integral Equations of the Second Kind Defined on the Half-line*, J. Comput. Math., to appear.
- [41] R. Chan and F. Lin, *Circulant Integral Operators and Circulant Preconditioners*, in Proc. Cornelius Lanczos Intl. Centenary Conf., Raleigh, NC, 1993, pp. 407–409.
- [42] R. Chan, J. Nagy, and R. Plemmons, *Circulant Preconditioned Toeplitz Least Squares Iterations*, SIAM J. Matrix Anal. Appl., 15 (1994), pp. 80–97.
- [43] R. Chan, J. Nagy, and R. Plemmons, *FFT-Based Preconditioners for Toeplitz-Block Least Squares Problems*, SIAM J. Numer. Anal., 30 (1993), pp. 1740–1768.
- [44] R. Chan, J. Nagy, and R. Plemmons, *Displacement Preconditioner for Toeplitz Least Squares Iterations*, Elec. Trans. Numer. Anal., 2 (1994), pp. 44–56.
- [45] R. Chan and M. Ng, *Fast Iterative Solvers for Toeplitz-Plus-Band Systems*, SIAM J. Sci. Comput., 14 (1993), pp. 1013–1019.
- [46] R. Chan and M. Ng, *Toeplitz Preconditioners for Hermitian Toeplitz Systems*, Linear Algebra Appl., 190 (1993), pp. 181–208.
- [47] R. Chan, M. Ng, and R. Plemmons, *Generalization of Strang’s Preconditioner with Applications to Toeplitz Least Squares Problems*, Numer. Linear Algebra Appl., to appear.

- [48] R. Chan, M. Ng, and C. Wong, *Sine Transform Based Preconditioners for Symmetric Toeplitz Systems*, Linear Algebra Appl., to appear.
- [49] R. Chan and G. Strang, *Toeplitz Equations by Conjugate Gradients with Circulant Preconditioner*, SIAM J. Sci. Stat. Comput., 10 (1989), pp. 104–119.
- [50] R. Chan and P. Tang, *Fast Band-Toeplitz Preconditioners for Hermitian Toeplitz Systems*, SIAM J. Sci. Comput., 15 (1994), pp. 164–171.
- [51] R. Chan and C. Wong, *Best-Conditioned Circulant Preconditioners*, Linear Algebra Appl., 218 (1995), pp. 205–212.
- [52] R. Chan and C. Wong, *Sine Transform Based Preconditioners for Elliptic Problems*, submitted.
- [53] R. Chan and M. Yeung, *Circulant Preconditioners for Toeplitz Matrices with Positive Continuous Generating Functions*, Math. Comp., 58 (1992), pp. 233–240.
- [54] R. Chan and M. Yeung, *Circulant Preconditioners Constructed from Kernels*, SIAM J. Numer. Anal., 29 (1992), pp. 1093–1103.
- [55] R. Chan and M. Yeung, *Jackson's Theorem and Circulant Preconditioned Toeplitz Systems*, J. Approx. Theory, 70 (1992), pp. 191–205.
- [56] R. Chan and M. Yeung, *Circulant Preconditioners for Complex Toeplitz Matrices*, SIAM J. Numer. Anal., 30 (1993), pp. 1193–1207.
- [57] T. Chan, *An Optimal Circulant Preconditioner for Toeplitz Systems*, SIAM J. Sci. Stat. Comput., 9 (1988), pp. 766–771.
- [58] T. Chan and P. Hansen, *A Look-Ahead Levinson Algorithm for General Toeplitz Systems*, IEEE Trans. Signal Process., 40 (1992), pp. 1079–1090.
- [59] T. Chan and J. Olkin, *Circulant Preconditioners for Toeplitz-block Matrices*, Numer. Algo., 6 (1994), pp. 89–101.
- [60] E. Cheney, *Introduction to Approximation Theory*, McGraw-Hill, New York, 1966.
- [61] C. Chui and A. Chan, *Application of Approximation Theory Methods to Recursive Digital Filter Design*, IEEE Trans. Acoust., Speech, Signal Process., 30 (1982), pp. 18–24.
- [62] C. Chui and G. Chen, *Signal Processing and Systems Theory*, Selected Topics, Springer-Verlag, Berlin, 1992.
- [63] J. Chun and T. Kailath, *Divide-and-Conquer Solutions of Least-Squares Problems for Matrices with Displacement Structure*, SIAM J. Matrix Anal. Appl., 12 (1991), pp. 128–145.

- [64] J. Chun, T. Kailath, and H. Lev-Ari, *Fast Parallel Algorithms for QR and Triangular Factorization*, SIAM J. Sci. Stat. Comput., 8 (1987), pp. 899–913.
- [65] N. Clinthorne, T. Pan, P. Chiao, L. Rogers, and J. Stamos, *Preconditioning Methods for Improved Convergence Rates in Iterative Reconstructions*, IEEE Trans. Med. Imag., 12 (1993), pp. 78–83.
- [66] J. Cooley and J. Tukey, *An Algorithm for the Machine Calculation of Complex Fourier Series*, Math. Comp., 19 (1965), pp. 297–301.
- [67] P. Concus and P. Saylor, *A Modified Direct Preconditioner for Indefinite Symmetric Toeplitz Systems*, Numer. Linear Algebra Appl., to appear.
- [68] G. Cybenko, *Fast Toeplitz Orthogonalization Using Inner Products*, SIAM J. Sci. Stat. Comput., 8 (1987), pp. 734–740.
- [69] P. Davis, *Circulant Matrices*, John Wiley & Sons, New York, 1979.
- [70] P. Delsarte, Y. Genin, and Y. Kamp, *A Generalization of the Levinson Algorithm for Hermitian Toeplitz Matrices with Any Rank Profile*, IEEE Trans. Acoust., Speech, Signal Process., 33 (1985), pp. 964–971.
- [71] L. Delves and J. Mohamed, *Computational Methods for Integral Equations*, Cambridge University Press, Cambridge, 1985.
- [72] M. Dryja and O. Widlund, *Some Domain Decomposition Algorithms for Elliptic Problems*, in Proc. Conf. on Iterative Methods for Large Linear Systems, Austin, TX, 1989.
- [73] T. Dupont, R. Kendall, and H. Rachford, *An Approximate Factorization Procedure for Solving Self-adjoint Elliptic Difference Equations*, SIAM J. Numer. Anal., 5 (1968), pp. 559–573.
- [74] G. Fiorentino and S. Serra, *Multigrid Methods for Toeplitz Matrices*, Calcolo, 28 (1991), pp. 283–305.
- [75] G. Fiorentino and S. Serra, *Multigrid Methods for Symmetric Positive Definite Block Toeplitz Matrices with Nonnegative Generating Functions*, SIAM J. Sci. Comput., to appear.
- [76] R. Freund, *A Transpose-free Quasi-minimal Residual Algorithm for Non-Hermitian Linear Systems*, SIAM J. Sci. Comput., 14 (1993), pp. 470–482.
- [77] R. Freund, *A Look-ahead Bareiss Algorithm for General Toeplitz Matrices*, Numer. Math., 68 (1994), pp. 35–69.
- [78] R. Freund, G. Golub, and N. Nachtigal, *Iterative Solution of Linear Systems*, Acta Numerica, (1992), pp. 57–100.

- [79] R. Freund and T. Huckle, *Iterative Solution of Linear Systems with Low Displacement Rank by Conjugate Gradient-type Algorithms*, XII Householder Symp., Lake Arrowhead, CA, 1993, in prepration.
- [80] R. Freund and H. Zha, *Formally Biorthogonal Polynomials and a Look-ahead Levinson Algorithm for General Toeplitz Systems*, Linear Algebra Appl., 188,189 (1993), pp. 255–303.
- [81] B. Friedlander, M. Morf, T. Kailath, and L. Ljung, *New Inversion Formulas for Matrices Classified in Terms of Their Distance from Toeplitz Matrices*, Linear Algebra Appl., 27 (1979), pp. 31–60.
- [82] W. Fuller, *Introduction to Statistical Time Series*, John Wiley & Sons, New York, 1976.
- [83] D. Gilliam, C. Martin, and J. Lund, *Analytic and Numerical Aspects of the Observation of the Heat Equation*, in Proc. 26th IEEE Conf. on Decision and Control, 2 (1987), pp. 975–976.
- [84] A. Giordano and F. Hsu, *Least Square Estimation with Applications to Digital Signal Processing*, Wiley, New York, 1985.
- [85] I. Gohberg and I. Fel'dman, *Convolution Equations and Projection Methods for Their Solution*, Transl. Math. Monographs, Vol. 41, Amer. Math. Soc., Providence, RI, 1974.
- [86] I. Gohberg, M. Hanke, and I. Koltracht, *Fast Preconditioned Conjugate Gradient Algorithms for Wiener-Hopf Integral Equations*, SIAM J. Numer. Anal., 31 (1994), pp. 429–443.
- [87] I. Gohberg and A. Semencul, *On the Inversion of Finite Toeplitz Matrices and Their Continuous Analogs*, Mat. Issled., 2 (1972), pp. 201–233.
- [88] G. Golub and C. Van Loan, *Matrix Computations*, 2nd ed., The Johns Hopkins University Press, Baltimore, MD, 1989.
- [89] M. Gover and S. Barnett, *Inversion of Toeplitz Matrices which are not Strongly Non-singular*, IMA J. Numer. Anal., 5 (1985), pp. 101–110.
- [90] C. Green, *Integral Equation Methods*, Barnes & Noble, New York, 1969.
- [91] U. Grenander and G. Szegö, *Toeplitz Forms and Their Applications*, 2nd ed., Chelsea Publishing, New York, 1984.
- [92] C. Groetsch, *The Theory of Tikhonov Regularization for Fredholm Equations of the First Kind*, Pitman Publishing, London, 1984.
- [93] C. Gueguen, *Linear Prediction in the Singular Case and the Stability of Eigen Models*, in Proc. IEEE Intl. Conf. on Acoust., Speech, and Signal Process., 2 (1981), pp. 881–885.

- [94] I. Gustafsson, *A Class of First Order Factorization Methods*, BIT, 18 (1978), pp. 142–156.
- [95] M. Gutknecht, *Stable Row Recurrences for the Padé Table and Generically Superfast Look-ahead Solvers for Non-Hermitian Toeplitz Systems*, Linear Algebra Appl., 188,189 (1993), pp. 351–421.
- [96] R. Hamming, *Digital Filters*, 3rd ed., Prentice Hall, Englewood Cliffs, NJ, 1989.
- [97] M. Hanke, J. Nagy, and R. Plemmons, *Preconditioned Iterative Regularization for Ill-posed Problems*, in Numerical Linear Algebra, L. Reichel, A. Ruttan, R. Varga, eds., de Gruyter, Berlin, 1993.
- [98] M. Hanke and J. Nagy, *Toeplitz Approximate Inverse Preconditioner for Banded Toeplitz Matrices*, Numer. Algo., 7 (1994), pp. 183–199.
- [99] S. Haykin, *Adaptive Filter Theory*, 2nd ed., Prentice-Hall, Englewood Cliffs, NJ, 1991.
- [100] L. Hemmingsson, *Toeplitz Preconditioners with Block Structure for First-order PDEs*, Numer. Linear Algebra Appl., to appear.
- [101] L. Hemmingsson and K. Otto, *Analysis of Semi-Toeplitz Preconditioners for First-order PDE*, SIAM J. Sci. Comput., to appear.
- [102] W. Henkel, *An Extended Berlekamp-Massey Algorithm for the Inversion of Toeplitz Matrices*, IEEE Trans. Comm., 40 (1992), pp. 1557–1561.
- [103] S. Holmgren, *CG-like Iterative Methods and Semicirculant Preconditioners on Vector and Parallel Computers*, Report No. 148, Dept. of Scientific Computing, Uppsala Univ., Uppsala, Sweden, 1992.
- [104] S. Holmgren and K. Otto, *A Comparison of Preconditioned Iterative Methods for Nonsymmetric Block-tridiagonal Systems of Equations*, Report No. 123 (revised), Dept. of Scientific Computing, Uppsala Univ., Uppsala, Sweden, 1990.
- [105] S. Holmgren and K. Otto, *Iterative Solution Methods and Preconditioners for Block-tridiagonal Systems of Equations*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 863–886.
- [106] S. Holmgren and K. Otto, *Semicirculant Preconditioners for First-order Partial Differential Equations*, SIAM J. Sci. Comput., 15 (1994), pp. 385–407.
- [107] S. Holmgren and K. Otto, *Semicirculant Solvers and Boundary Corrections for First-order Partial Differential Equations*, SIAM J. Sci. Comput., to appear.
- [108] F. de Hoog, *A New Algorithm for Solving Toeplitz Systems of Equations*, Linear Algebra Appl., 88/89 (1987), pp. 123–138.
- [109] J. Hsue and A. Yagle, *Fast Algorithms for Close-to-Toeplitz-Plus-Hankel Systems and Two-Sided Linear Prediction*, IEEE Trans. Signal Process., 41 (1993), pp. 2349–2361.

- [110] T. Huckle, *Circulant/Skew Circulant Matrices as Preconditioners for Hermitian Toeplitz Systems*, in Proc. IMACS Conf. on Iterative Methods in Linear Algebra, Brussels, 1991.
- [111] T. Huckle, *Circulant and Skew Circulant Matrices for Solving Toeplitz Matrix Problems*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 767–777.
- [112] T. Huckle, *Some Aspects of Circulant Preconditioners*, SIAM J. Sci. Comput., 14 (1993), pp. 531–541.
- [113] T. Huckle, *Iterative Methods for Toeplitz-like Matrices*, Report SCCM-94-05, Computer Science Dept., Stanford Univ., Stanford, CA, 1994.
- [114] T. Huckle, *Fast Transforms for Tridiagonal Linear Equations*, BIT, 34 (1994), pp. 99–112.
- [115] T. Huckle, *A Note on Skew Circulant Preconditioners for Elliptic Problems*, Numer. Algo., 2 (1992), pp. 279–286.
- [116] A. Jain, *Fundamentals of Digital Image Processing*, Prentice-Hall, Englewood Cliffs, NJ, 1989.
- [117] A. Jain, *Fast Inversion of Banded Toeplitz Matrices by Circular Decompositions*, IEEE Trans. Acoust., Speech, Signal Process., 26 (1978), pp. 121–126.
- [118] X. Jin, *A Note on Circulant Preconditioners for Hyperbolic and Parabolic Equations*, Chinese J. Math., 21 (1993), pp. 129–142.
- [119] X. Jin, *Hartley Preconditioners for Toeplitz Systems Generated by Positive Continuous Functions*, BIT, 34 (1994), pp. 367–371.
- [120] X. Jin, *Band-Toeplitz Preconditioners for Block Toeplitz Systems*, submitted.
- [121] X. Jin and R. Chan, *Circulant Preconditioners for Second Order Hyperbolic Equations*, BIT, 32 (1992), pp. 650–664.
- [122] M. Kac, W. Murdock, and G. Szegő, *On the Eigenvalues of Certain Hermitian Forms*, J. Ration. Mech. Anal., 2 (1953), pp. 767–800.
- [123] A. Katsaggelos, *Digital Image Restoration*, Springer-Verlag, Berlin, 1991.
- [124] L. Kaufman, *Matrix Methods for Queueing Problems*, SIAM J. Sci. Stat. Comput., 4 (1983), pp. 525–552.
- [125] R. King et al., *Digital Filtering in One and Two Dimensions: Design and Applications*, Plenum Press, New York, 1989.
- [126] B. Kiss and G. Molnárka, *A Preconditioned Domain Decomposition Algorithm for the Solution of the Elliptic Neumann Problem*, Period. Math. Hungar., 24 (1992), pp. 151–165.

- [127] B. Kiss and G. Molnárka, *A Preconditioned Domain Decomposition Algorithm for the Solution of the Elliptic Neumann Problem*, in *Parallel Algorithms for Partial Differential Equations*, W. Hackbusch, ed., Vieweg, Braunschweig, Germany, 1991.
- [128] T. Ku and C. Kuo, *Design and Analysis of Toeplitz Preconditioners*, *IEEE Trans. Signal Process.*, 40 (1992), pp. 129–141.
- [129] T. Ku and C. Kuo, *On the Spectrum of a Family Preconditioned Block Toeplitz Matrices*, *SIAM J. Sci. Stat. Comput.*, 13 (1992), pp. 948–966.
- [130] T. Ku and C. Kuo, *Spectral Properties of Preconditioned Rational Toeplitz Matrices*, *SIAM J. Matrix Anal. Appl.*, 14 (1993), pp. 146–165.
- [131] T. Ku and C. Kuo, *Spectral Properties of Preconditioned Rational Toeplitz Matrices: the Nonsymmetric Case*, *SIAM J. Matrix Anal. Appl.*, 14 (1993), pp. 521–544.
- [132] T. Ku and C. Kuo, *A Minimum-phase LU Factorization Preconditioner for Toeplitz Matrices*, *SIAM J. Sci. Stat. Comput.*, 13 (1992), pp. 1470–1487.
- [133] T. Ku and C. Kuo, *Preconditioned Iterative Methods for Solving Toeplitz-plus-Hankel Systems*, *SIAM J. Numer. Anal.*, 30 (1993), pp. 824–845.
- [134] N. Levinson. *The Wiener RMS (Root Mean Square) Error Criterion in Filter Design and Prediction*, *J. Math. and Phys.*, 25 (1946), pp. 261–278.
- [135] F. Lin, M. Ng, and R. Chan, *Preconditioners for Wiener-Hopf Equations with Higher Order Quadrature Rules*, submitted.
- [136] E. Linzer, *Arithmetic Complexity of Iterative Toeplitz Solvers*, submitted.
- [137] E. Linzer, *Extended Circulant Conditioning of Toeplitz Systems*, Tech. Report R.C, IBM Research, Yorktown Heights, NY, October 1992.
- [138] I. Lirkov and S. Margenov, *On Circulant Preconditioning of Elliptic Problems in L-shaped Domain*, in *Advances in Numerical Methods and Applications*, I. Dimov, Bl. Sendov, P. Vassilevski, eds., World Scientific, Singapore, 1994.
- [139] D. Luenberger, *Linear and Nonlinear Programming*, 2nd ed., Addison-Wesley, Reading, MA, 1984.
- [140] L. Marple, *A New Autoregressive Spectrum Analysis Algorithm*, *IEEE Trans. Acoust., Speech, Signal Process.*, 28 (1980), pp. 441–454.
- [141] L. Marple, *Efficient Least Squares FIR System Identification*, *IEEE Trans. Acoust., Speech, Signal Process.*, 29 (1981), pp. 62–73.
- [142] L. Marple, *Fast Algorithms for Linear Prediction and System Identification Filters with Linear Phase*, *IEEE Trans. Acoust., Speech, Signal Process.*, 30 (1982), pp. 942–953.

- [143] M. Morf, *Doubling Algorithms for Toeplitz and Related Equations*, in Proc. IEEE Intl. Conf. on Acoust., Speech, and Signal Process., 3 (1980), pp. 954–959.
- [144] J. Nagy and R. Plemmons, *Some Fast Toeplitz Least Squares Algorithms*, in Proc. SPIE Conf. on Advanced Signal Processing Algorithms, Architectures, and Implementations II, Vol. 1566, (1991), pp. 35–46.
- [145] J. Nagy, R. Plemmons, and T. Torgersen, *Preconditioned Iterative Deconvolution*, submitted.
- [146] J. Nagy, R. Plemmons, and T. Torgersen, *Fast Restoration of Atmospherically Blurred Images*, in Proc. SPIE Symp. on Advanced Signal Processing: Algorithms, Architectures, and Implementations, Vol. 2296, (1994), to appear.
- [147] M. Ng, *Fast Iterative Methods for Solving Toeplitz-plus-Hankel Least Squares Problems*, Elec. Trans. Numer. Anal., 2 (1994), pp. 154–170.
- [148] M. Ng, *Fast FFT-based Conjugate Gradient Algorithm for Applications in Signal Processing*, in Proc. Cornelius Lanczos Intl. Centenary Conf., Raleigh, NC, 1993, pp. 404–406.
- [149] M. Ng and R. Chan, *Fast Iterative Methods for Least Squares Estimations*, Numer. Algo., 6 (1994), pp. 353–378.
- [150] M. Ng and F. Lin, *Decomposition and Construction of Preconditioners for Wiener-Hopf Equations*, SEA Bull. Math., 20 (1996), to appear.
- [151] M. Ng, F. Lin, and R. Chan, *Construction of Preconditioners for Wiener-Hopf Equations by Operator Splitting*, Appl. Math. Comp., to appear.
- [152] M. Ng and R. Plemmons, *Fast Recursive Least Squares Using the FFT*, in Mathematics in Signal Processing III, J. McWhirter, ed., Clarendon Press, Oxford, 1994.
- [153] M. Ng and R. Plemmons, *Fast RLS Adaptive Filtering by FFT-based Conjugate Gradient Iterations*, SIAM J. Sci. Comput., (1996), to appear.
- [154] T. Oda, *Moment Analysis for Traffic Associated with Markovian Queueing Systems*, IEEE Trans. Comm., 39 (1991), pp. 737–745.
- [155] J. Olkin, *Linear and Nonlinear Deconvolution Problems*, Ph.D. thesis, Rice Univ., Houston, TX, 1986.
- [156] K. Otto, *Analysis of Preconditioners for Hyperbolic Partial Differential Equations*, SIAM J. Numer. Anal., to appear.
- [157] D. Pal, *Fast Algorithms for Structured Matrices with Arbitrary Rank Profile*, Ph.D. thesis, Stanford Univ., Stanford, CA, 1990.

- [158] J. Pearl, *On Coding and Filtering Stationary Signals by Discrete Fourier Transforms*, IEEE Trans. Inform. Theory, 19 (1973), pp. 229–232.
- [159] R. Plemmons, *FFT-based RLS in Signal Processing*, in Proc. IEEE Intl. Conf. on Acoust., Speech, and Signal Process., 3 (1993), pp. 571–574.
- [160] M. Priestley, *Spectral Analysis and Time Series*, Academic Press, London, 1981.
- [161] L. Pustyl'nikov, *On the Algebraic Structure of the Spaces of Toeplitz and Hankel Matrices*, Soviet Math. Dokl., 21 (1980), pp. 141–144.
- [162] S. Reeves and R. Mersereau, *Optimal Regularization Parameter Estimation for Image Restoration*, in Proc. SPIE Conf. on Image Processing Algorithms and Techniques II, Vol. 1452 (1991), pp. 127–138.
- [163] L. Rudin, S. Osher, and C. Fu, *Total Variation Based Restoration of Noisy, Blurred Images*, preprint, 1993.
- [164] Y. Saad and M. Schultz, *GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.
- [165] A. Sayed and T. Kailath, *Optimal Circulant Preconditioning for Arbitrary Matrices*, submitted.
- [166] A. Seila, *Multivariate Estimation of Conditional Performance Measure in Regenerative Simulation*, Amer. J. Math. Manag. Sci., 10 (1990), pp. 17–50.
- [167] S. Serra, *Preconditioning Strategies for Asymptotically Ill-Conditioned Block Toeplitz Systems*, BIT, 34 (1994), pp. 579–594.
- [168] S. Serra, *Preconditioning Strategies for Hermitian Toeplitz Systems with Non-Definite Generating Functions*, submitted.
- [169] P. Sherman, *Circulant Approximations of the Inverses of Toeplitz Matrices and Related Quantities with Applications to Stationary Random Processes*, IEEE Trans. Acoust., Speech, Signal Process., 33 (1985), pp. 1630–1632.
- [170] L. Sibul, *Adaptive Signal Processing*, IEEE Press, New York, 1987.
- [171] U. Steimel, *Fast Computation of Toeplitz Forms under Narrow Band Conditions with Applications to Spectral Estimation*, in Digital Signal Processing, V. Cappellini and A. G. Constantinides, eds., Academic Press, London, 1980.
- [172] G. Strang, *A Proposal for Toeplitz Matrix Calculations*, Stud. Appl. Math., 74 (1986), pp. 171–176.

- [173] G. Strang and A. Edelman, *The Toeplitz-Circulant Eigenvalue Problem $Ax = \lambda Cx$* , in Oakland Conf. on PDE and Appl. Math., L. Bragg and J. Dettman, eds., Longman Sci. Tech., New York, 1987.
- [174] P. Swarztrauber, *Multiprocessor FFTs*, Parallel Comput., 5 (1987), pp. 197–210.
- [175] D. Sweet, *The Use of Pivoting to Improve the Numerical Performance of Algorithms for Toeplitz Matrices*, SIAM J. Matrix Anal. Appl., 14 (1993), pp. 468–493.
- [176] D. Sweet, *Fast Toeplitz Orthogonalization*, Numer. Math., 43 (1984), pp. 1–21.
- [177] P. Tang, *A Fast Algorithm for Linear Complex Chebyshev Approximations*, Math. Comp., 51 (1988), pp. 721–739.
- [178] A. Tekalp and G. Pavlović, *Restoration of Scanned Photographic Images*, in Digital Image Restoration, A. Katsaggelos, ed., Springer-Verlag, Berlin, 1991.
- [179] M. Tismenetsky, *A Decomposition of Toeplitz Matrices and Optimal Circulant Preconditioning*, Linear Algebra Appl., 154–156 (1991), pp. 105–121.
- [180] O. Toeplitz, *Zur Theorie der quadratischen und bilinearen Formen von unendlichvielen Veränderlichen. I. Teil: Theorie der L-Formen.*, Math. Annal., 70 (1911), pp. 351–376.
- [181] N. Trefethen, *Approximation Theory and Numerical Linear Algebra*, in Algorithms for Approximation II, J. Mason and M. Cox, eds., Chapman and Hall, London, 1990.
- [182] W. Trench, *An Algorithm for the Inversion of Finite Toeplitz Matrices*, SIAM J. Appl. Math., 12 (1964), pp. 515–522.
- [183] E. Tyrtyshnikov, *Optimal and Superoptimal Circulant Preconditioners*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 459–473.
- [184] E. Tyrtyshnikov, *A Unifying Approach to Some Old and New Theorems on Distribution and Clustering*, Linear Algebra. Appl., to appear.
- [185] E. Tyrtyshnikov, *Circulant Preconditioners with Unbounded Inverses*, Linear Algebra Appl., 216 (1995), pp. 1–24.
- [186] E. Tyrtyshnikov, *Influence of Matrix Operations on the Distribution of Eigenvalues and Singular Values of Toeplitz Matrices*, Linear Algebra Appl., 207 (1994), pp. 225–249.
- [187] E. Tyrtyshnikov and V. Strela, *Which Circulant Preconditioners are Better?*, Math. Comp., to appear.
- [188] C. Van Loan, *Computational Frameworks for the Fast Fourier Transform*, SIAM, Philadelphia, PA, 1992.
- [189] C. Vogel, *Total Variation Regularization of Ill-posed Problems*, preprint, 1993.

- [190] J. Walker, *Fourier Analysis*, Oxford University Press, New York, 1988.
- [191] O. Widlund, *Iterative Substructuring Methods: Algorithms and Theory for Elliptic Problems in the Plane*, in First Intl. Symp. on Domain Decomposition Methods for Partial Differential Equations, R. Glowinski et al., eds., SIAM, Philadelphia, PA, 1988.
- [192] H. Widom, *Hankel Matrices*, Trans. Amer. Math. Soc., 121 (1966), pp. 1–35.
- [193] S. Wright, *Parallel Algorithms for Banded Linear Systems*, SIAM J. Sci. Stat. Comput., 12 (1991), pp. 824–842.
- [194] A. Yagle, *New Analogs of Split Algorithms for Arbitrary Toeplitz-plus-Hankel Matrices*, IEEE Trans. Signal Process., 39 (1991), pp. 2457–2463.
- [195] M. Yeung and R. Chan, *Circulant Preconditioners for Toeplitz Matrices with Piecewise Continuous Generating Functions*, Math. Comp., 61 (1993), pp. 701–718.
- [196] P. Yip and K. Rao, *Fast Decimation-in-time Algorithms for a Family of Discrete Sine and Cosine Transforms*, Circuits, Systems, Signal Process., 3 (1984), pp. 387–408.
- [197] U. Yule, *On a Method of Investigating Periodicities in Disturbed Series, with Special Reference to Wolfer's Sunspot Numbers*, Phil. Trans. Roy. Soc. London, A 226 (1927), pp. 267–298.
- [198] C. Zarowski, *Schur Algorithms for Hermitian Toeplitz, and Hankel Matrices with Singular Leading Principal Submatrices*, IEEE Trans. Signal Process., 39 (1991), pp. 2464–2480.
- [199] S. Zohar, *The Solution of a Toeplitz Set of Linear Equations*, J. Assoc. Comput. Mach., 21 (1974), pp. 272–276.
- [200] A. Zygmund, *Trigonometric Series*, Cambridge University Press, London, 1968.