UCLA
COMPUTATIONAL AND APPLIED MATHEMATICS

Preconditioned Toeplitz Least Squares Iterations

Raymond H. Chan
James G. Nagy
Robert J. Plemmons

March 1992

CAM Report 92-15

Department of Mathematics
University of California, Los Angeles
Los Angeles, CA. 90024-1555

Preconditioned Toeplitz Least Squares Iterations*

Raymond H. Chan† James G. Nagy‡ and Robert J. Plemmons§

Abstract

We consider the solution of least squares problems $\min \|b - Tx\|_2$ by the preconditioned conjugate gradient method, for m-by-n complex Toeplitz matrices T of rank n. We use a circulant preconditioner C, derived using the T. Chan optimal preconditioner on n-by-n row blocks of T, such that C^*C approximates T^*T. For Toeplitz T that are generated by 2π-periodic continuous complex-valued functions without any zeros, we prove that the singular values of the preconditioned matrix TC^{-1} are clustered around 1. We show that if the condition number of T is of $O(n^\alpha)$, $\alpha > 0$, then the least squares conjugate gradient method converges in at most $O(\alpha \log n + 1)$ steps. Since each iteration requires only $O(m \log n)$ operations using the FFT, it follows that the total complexity of the algorithm is then only $O(\alpha m \log^2 n + m \log n)$. Conditions for *superlinear convergence* are given and numerical examples are provided illustrating the effectiveness of our methods.

1 Introduction

The conjugate gradient (CG) method is an iterative method for solving Hermitian positive definite systems $Ax = b$, see for instance Golub and van Loan [13]. When A is a rectangular m-by-n matrix of rank n, one can still use the CG algorithm to find the solution to the least squares problem

$$\min \|b - Ax\|_2. \tag{1}$$

This can be done by applying the algorithm to the normal equations in factored form,

$$A^*(b - Ax) = 0, \tag{2}$$

which can be solved by conjugate gradients without explicitly forming the matrix A^*A, see Bjorck [2].

The convergence of the conjugate gradient algorithm and its variations depends on the the singular values of the data matrix A, see Axelsson [1]. If the singular values cluster around a fixed point, convergence will be rapid. Thus, to make the algorithm a useful iterative method, one usually *preconditions* the system. The preconditioned conjugate gradient (PCG) algorithm then solves (1) by transforming the problem with a preconditioner M, applying the conjugate gradient method to the transformed problem, and then transforming back. More precisely, one can use the conjugate gradient method to solve $\min \|b - AM^{-1}x\|_2$, and then solve $Mx = y$.

In this paper we consider the least squares problem (1), with the data matrix $A = T$, where T is a rectangular m-by-n Toeplitz matrix of rank n. The matrix $T = (t_{jk})$ is said to be *Toeplitz* if $t_{jk} = t_{j-k}$, *i.e.*, T is constant along its diagonals. An n-by-n matrix C is said to be *circulant* if it is Toeplitz and its diagonals c_j satisfy $c_{n-j} = c_{-j}$ for $0 < j \le n-1$.

*To appear in the *Proceedings of the Copper Mountain Conference on Iterative Methods, April, 1992*

†Department of Mathematics, University of Hong Kong, Hong Kong. Research supported in part by ONR contract no. N00014-90-J-1695 and DOE grant no. DE-FG03-87ER25037.

‡Institute for Mathematics and its Applications, University of Minnesota, Minneapolis, MN 55455.

§Department of Mathematics and Computer Science, Wake Forest University, P.O. Box 7388, Winston-Salem, NC 27109. Research supported by the US Air Force under grant no. AFOSR-91-0163.

matrix T and padding zeros to the bottom rows, 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.

Recall that the solution to the least squares problem

$$\min \|b - Tx\|_2 \tag{3}$$

can be found by the preconditioned conjugate gradient method by applying the method to the normal equations (2) in factored form, that is, using T and T^* without forming T^*T. The preconditioner M considered in this paper is given by an n-by-n circulant matrix $M = C$, where C^*C is then a circulant matrix that approximates T^*T.

The version of the PCG algorithm we use is given in [2, p.102] and will be called Algorithm PCG for Least Squares.

The idea of using the preconditioned conjugate gradient method with circulant preconditioners for solving square positive definite Toeplitz systems was first proposed by Strang [19], although the application of circulant approximations to Toeplitz matrices has been used for some time in image processing. The convergence rate of the method was analyzed in R. Chan and Strang [3] for Toeplitz matrices that are generated by positive Wiener class functions. Since then, considerable research have been done in finding other good circulant preconditioners or extending the class of generating functions for which the method is effective, see T. Chan [10], R. Chan [4], Tyrtyshnikov [21], Tismenetsky [20], Huckle [14], Ku and Kuo [15], R. Chan and Yeung [7], T. Chan and Olkin [11], R. Chan, Jin and Yeung [6] and R. Chan and Yeung [8].

Recently, the idea of using circulant preconditioners has been extended to non-Hermitian square Toeplitz systems by R. Chan and Yeung [8] and to Toeplitz least squares problems by Nagy [16] and Nagy and Plemmons [17]. The main aim of this paper is to formalize and establish convergence results, and to provide applications, in the case where T is a rectangular (block) Toeplitz matrix. More precisely, we consider in this paper, kn-by-n matrices T of the form

$$T^T = [T_1^T, T_2^T, \cdots T_k^T], \tag{4}$$

where each square block T_j is a Toeplitz matrix. Notice that if T itself is a rectangular Toeplitz matrix, then each block T_j is necessarily Toeplitz.

Following [16, 17], for each block T_j, we construct a circulant approximation C_j. Then our preconditioner is defined as a square circulant matrix C, such that

$$C^*C = \sum_{j=1}^{k} C_j^* C_j.$$

Notice that each C_j is an n-by-n circulant matrix. Hence they can all be diagonalized by the Fourier matrix F, i.e. $C_j = F\Lambda_j F^*$ where Λ_j is diagonal, see Davis [12]. Therefore the spectrum of C_j, $j = 1, \cdots k$, can be computed in $O(n \log n)$ operations by using the Fast Fourier Transform (FFT). Since

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

C^*C is also circulant and its spectrum can be computed in $O(kn \log n)$ operations. Here we choose, as in [16, 17],

$$C = F(\sum_{j=1}^{k} \Lambda_j^* \Lambda_j)^{\frac{1}{2}} F^*. \tag{5}$$

The number of operations per iteration in Algorithm PCG for Least Squares depends mainly on the work of computing the matrix-vector multiplications. In our case, this amounts to computing products: $Ty, T^*z, C^{-1}y, C^{-*}y$ for some n-vectors y and m-vectors z. Since C is circulant the products $C^{-1}y$ and $C^{-*}y$ can be found efficiently by using the FFT in $O(n \log n)$ operations. For the products Ty and T^*z, with T in block form with k n-by-n blocks T_j, we have to compute n products of the form $T_j w$ where T_j is an n-by-n Toeplitz matrix and w is an n-vector. However the product $T_j w$ can be computed using the FFT by first embedding T_j into a $2n$-by-$2n$ circulant matrix. The multiplication thus requires $O(2n\log(2n))$ operations. It follows that the operations for computing Ty and T^*z are of the order $O(m \log n)$, where $m = nk$. Thus we conclude that the cost per iteration in the preconditioned conjugate gradient method is of the order $O(m \log n)$.

As already mentioned in the beginning, the convergence rate of the method depends on the distribution of the singular values of the matrix TC^{-1} which are the same as the square roots of the eigenvalues of the matrix $(C^*C)^{-1}(T^*T)$. We will show, then, that if the generating functions of the blocks T_j are 2π-periodic continuous functions and if one of these functions has no zeros, then the spectrum of $(C^*C)^{-1}(T^*T)$ will be clustered around 1. We remark that the class of 2π-periodic continuous functions contains the Wiener class of functions which in turn contains the class of rational functions considered in Ku and Kuo [15].

By using a standard error analysis of the conjugate gradient method, we then show that if the condition number $\kappa(T)$ of T is of $O(n^\alpha)$, then the number of iterations required for convergence is at most $O(\alpha \log n + 1)$ where $\alpha > 0$. 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 $O(\alpha m \log^2 n + m \log n)$. In the case when $\alpha = 0$, i.e. T is well-conditioned, the method converges in $O(1)$ steps. Hence the complexity is reduced to just $O(m \log n)$ operations.

2 Properties of the Circulant Preconditioner

In this section, we consider circulant preconditioners for least square problems and study their spectral properties. We begin by recalling some results for square Toeplitz systems.

For simplicity, we denote by $\mathcal{C}_{2\pi}$ the Banach space of all 2π-periodic continuous complex-valued functions equipped with the supremum norm $||\cdot||_\infty$. As already mentioned in §1, this class of functions contains the Wiener class of functions. For all $f \in \mathcal{C}_{2\pi}$, let

$$a_k = \frac{1}{2\pi}\int_{-\pi}^{\pi} f(\theta)e^{-ik\theta}d\theta, \; k = 0, \pm1, \pm2, \cdots,$$

be the Fourier coefficients of f. Let A be the n-by-n complex Toeplitz matrix with the (j,k)th entry given by a_{j-k}. The function f is called the generating function of the matrix A.

For a given n-by-n Toeplitz matrix A, we let C be the n-by-n circulant preconditioner of A as defined in T. Chan [10], i.e. C is the minimizer of $F(X) = ||A - X||_F$ over all circulant matrices X. We note that the (j,ℓ)th entry of C is given by the diagonal $c_{j-\ell}$ where

$$c_k = \begin{cases} \dfrac{(n-k)a_k + k a_{k-n}}{n} & 0 \le k < n, \\ c_{n+k} & 0 < -k < n. \end{cases} \tag{6}$$

The following three Lemmas are proved in R. Chan and Yeung [8]. The first two give the bounds of $||A||_2$ and $||C||_2$ and the last one shows that $A - C$ has clustered spectrum.

Lemma 1 *Let $f \in \mathcal{C}_{2\pi}$. Then we have*

$$||A||_2 \le 2||f||_\infty < \infty, \quad n = 1, 2, \cdots. \tag{7}$$

If moreover f has no zeros then there exists a constant $c > 0$ such that for all n sufficiently large, we have

$$||A||_2 > c. \tag{8}$$

Lemma 2 *Let $f \in \mathcal{C}_{2\pi}$. Then we have*

$$||C||_2 \leq 2||f||_\infty < \infty, \quad n = 1, 2, \cdots. \tag{9}$$

If moreover f has no zeros, then for all sufficiently large n, we also have

$$||C^{-1}||_2 \leq 2||\frac{1}{f}||_\infty < \infty. \tag{10}$$

Lemma 3 *Let $f \in \mathcal{C}_{2\pi}$. Then for all $\epsilon > 0$, there exist N and $M > 0$, such that for all $n > N$, $A - C = U + V$ where rank $U \leq M$ and $||V||_2 \leq \epsilon$.*

Now let us consider the general least squares problem (3) where T is an m-by-n matrix with $m \geq n$. Here we assume that $m = kn$, without loss of generality, since otherwise the final block T_k can be extended to an $n \times n$ Toeplitz matrix by extending the diagonals and padding the lower left part with zeros. (This modification is only for constructing the preconditioner. The original least squares problem (3) is not changed.) Thus we can partition T as (4), without loss of generality. We note that the solution to the least square problem (3) can be obtained by solving the normal equations $T^*Tx = T^*b$, in factored form, where

$$T^*T = \sum_{j=1}^{k} T_j^* T_j.$$

Of course one can avoid actually forming T^*T for implementing the conjugate gradient method for the normal equations [2].

We will assume in the following that k is a constant independent of n and that each square block T_j, $j = 1, \cdots, k$ is generated by a generating function f_j in $\mathcal{C}_{2\pi}$. Following Nagy [16], and Nagy and Plemmons [17], we define a preconditioner for T based upon preconditioners for the blocks T_j.

For each block T_j, let C_j be the corresponding T. Chan's circulant preconditioner as defined in (6). Then it is natural to consider the square circulant matrix

$$C^*C = \sum_{j=1}^{k} C_j^* C_j \tag{11}$$

as a circulant approximation to T^*T [17]. Note, however, that C is computed (or applied) using the equation (5). Clearly C is invertible if one of the C_j is. In fact, using Lemma 2, we have

Lemma 4 *Let $f_j \in \mathcal{C}_{2\pi}$ for $j = 1, 2, \cdots, k$. Then we have*

$$||C||_2^2 \leq 4\sum_{j=1}^{k} ||f_j||_\infty^2 < \infty, \quad n = 1, 2, \cdots. \tag{12}$$

If moreover one of the f_j, say f_ℓ, has no zeros, then for all sufficiently large n, we also have

$$||(C^*C)^{-1}||_2 \leq 4||\frac{1}{f_\ell}||_\infty^2 < \infty. \tag{13}$$

Proof: Equation (12) clearly follows from (11) and (9). To prove (13), we just note that $C_j^*C_j$ are positive semidefinite matrices for all $j = 1, \cdots, k$, hence $\lambda_{\min}(C^*C) \geq \lambda_{\min}(C_\ell^* C_\ell)$, where $\lambda_{\min}(\cdot)$ denotes the smallest eigenvalue. Thus by (10), we then have

$$||(C^*C)^{-1}||_2 \leq ||(C_\ell^* C_\ell)^{-1}||_2 = ||C_\ell^{-1}||_2^2 \leq 4||\frac{1}{f_\ell}||_\infty^2. \quad \square$$

3 Spectrum of TC^{-1}

In this section, we show that the spectrum of the matrix $(C^*C)^{-1}(T^*T)$ is clustered around 1. It will follow then, that the singular values of TC^{-1} are also clustered around 1, since $(C^*C)^{-1}(T^*T)$ is similar to $(TC^{-1})^*(TC^{-1})$. We begin by analyzing the spectrum of each block.

Lemma 5 *For $1 \le j \le k$, if $f_j \in \mathcal{C}_{2\pi}$, then for all $\epsilon > 0$, there exist N_j and $M_j > 0$, such that for all $n > N_j$, $T_j^*T_j - C_j^*C_j = U_j + V_j$ where U_j and V_j are Hermitian matrices with rank $U_j \le M_j$ and $||V_j||_2 \le \epsilon$.*

Proof: We first note that by Lemma 3, we have for all $\epsilon > 0$, there exist positive integers N_j and M_j such that for all $n > N_j$, $T_j - C_j = \tilde{U}_j + \tilde{V}_j$ where rank $\tilde{U}_j \le M_j$ and $||\tilde{V}_j||_2 \le \epsilon$. Therefore,

$$\begin{aligned} T_j^*T_j - C_j^*C_j &= T_j^*(T_j - C_j) + (T_j - C_j)^*C_j = T_j^*(T_j - C_j) - (T_j - C_j)^*(T_j - C_j) + (T_j - C_j)^*T_j \\ &= T_j^*(\tilde{U}_j + \tilde{V}_j) - (\tilde{U}_j + \tilde{V}_j)^*(\tilde{U}_j + \tilde{V}_j) + (\tilde{U}_j + \tilde{V}_j)^*T_j \equiv U_j + V_j. \end{aligned}$$

Here

$$U_j = T_j^*\tilde{U}_j + \tilde{U}_j^*T_j - \tilde{U}_j^*\tilde{U}_j - \tilde{U}_j^*\tilde{V}_j - \tilde{V}_j^*\tilde{U}_j = \tilde{U}_j^*(T_j - \tilde{U}_j - \tilde{V}_j) + (T_j - \tilde{V}_j)^*\tilde{U}_j$$

and $V_j = \tilde{V}_j^*T_j + T_j^*\tilde{V}_j - \tilde{V}_j^*\tilde{V}_j$. It is clear that both U_j and V_j are Hermitian matrices. Moreover we have rank $U_j \le 2M_j$ and $||V_j||_2 \le 2\epsilon||T_j||_2 + \epsilon^2$. By (7), we then have $||V_j||_2 \le 4\epsilon||f_j||_\infty + 2\epsilon^2$. $\square$

Using the facts that

$$T^*T - C^*C = \sum_{j=1}^{k}(T_j^*T_j - C_j^*C_j)$$

and that k is independent of n, we immediately have

Lemma 6 *Let $f_j \in \mathcal{C}_{2\pi}$ for $j = 1, \cdots k$. Then for all $\epsilon > 0$, there exist N and $M > 0$, such that for all $n > N$, $T^*T - C^*C = \tilde{U} + \tilde{V}$ where $\tilde{U}$ and $\tilde{V}$ are Hermitian matrices with*

$$\text{rank } \tilde{U} \le M \tag{14}$$

and

$$||\tilde{V}||_2 \le \epsilon. \tag{15}$$

We now show that the spectrum of the preconditioned matrix $(C^*C)^{-1}(T^*T)$ is clustered around 1. We note that this is equivalent to showing that the spectrum of $(C^*C)^{-1}(T^*T) - I$, where I is the n-by-n identity matrix, is clustered around zero.

Theorem 1 *Let $f_j \in \mathcal{C}_{2\pi}$ for all $j = 1, \cdots, k$. If one of the f_j, say f_ℓ, has no zeros, then for all $\epsilon > 0$, there exist N and $M > 0$, such that for all $n > N$, at most M eigenvalues of the matrix $(C^*C)^{-1}(T^*T) - I$ have absolute values larger than ϵ.*

Proof: By Lemma 6, we have $(C^*C)^{-1}(T^*T) - I = (C^*C)^{-1}(T^*T - C^*C) = (C^*C)^{-1}(\tilde{U} + \tilde{V})$. Therefore the spectra of the matrices $(C^*C)^{-1}(T^*T) - I$ *and* $(C^*C)^{-1/2}(\tilde{U} + \tilde{V})(C^*C)^{-1/2}$ are the same. However, by (14), we have *rank* $\left\{(C^*C)^{-1/2}\tilde{U}(C^*C)^{-1/2}\right\} \le M$ and by (15) and (13), we have

$$||(C^*C)^{-1/2}\tilde{V}(C^*C)^{-1/2}||_2 \le ||\tilde{V}||_2||(C^*C)^{-1}||_2 \le 4\epsilon||\frac{1}{f_\ell}||_\infty^2.$$

Thus by applying Cauchy's interlace theorem (see Wilkinson [22]) to the Hermitian matrix

$$(C^*C)^{-1/2}\tilde{U}(C^*C)^{-1/2} + (C^*C)^{-1/2}\tilde{V}(C^*C)^{-1/2},$$

we see that its spectrum is clustered around zero. Hence the spectrum of the matrix $(C^*C)^{-1}(T^*T)$ is clustered around 1. $\square$

From Theorem 1, we have the desired clustering result; namely, if $f_j \in \mathcal{C}_{2\pi}$ for all $j = 1, \cdots, k$ and if one of the f_j has no zeroes, then the *singular values of the preconditioned matrix* TC^{-1} *are clustered around* 1.

4 Convergence Rate of the Method

In this section, we analyze the convergence rate of Algorithm PCG for Least Squares, for our circulant preconditioned Toeplitz matrix TC^{-1}. We show first that the method converges in at most $O(\alpha \log n + 1)$ steps where $O(n^\alpha)$ is the condition number of T . We begin by noting the following error estimate of the conjugate gradient method, see [8].

Lemma 7 *Let G be a positive definite matrix and x be the solution to $Gx = b$. Let x_j be the jth iterant of the ordinary conjugate gradient method applied to the equation $Gx = b$. If the eigenvalues $\{\delta_k\}$ of G are such that*

$$0 < \delta_1 \le \ldots \le \delta_p \le b_1 \le \delta_{p+1} \le \ldots \le \delta_{n-q} \le b_2 \le \delta_{n-q+1} \le \ldots \le \delta_n,$$

then

$$\frac{||x - x_j||_G}{||x - x_0||_G} \le 2\left(\frac{b-1}{b+1}\right)^{j-p-q} \cdot \max_{\delta \in [b_1, b_2]} \left\{ \prod_{k=1}^{p} \left(\frac{\delta - \delta_k}{\delta_k} \right) \right\}. \tag{16}$$

Here $b \equiv (b_2/b_1)^{\frac{1}{2}} \ge 1$ and $||v||_G \equiv v^* G v$.

For the system

$$(C^*C)^{-1}(T^*T)x = (C^*C)^{-1}T^*b, \tag{17}$$

the iteration matrix G is given by $G = (C^*C)^{-1/2}(T^*T)(C^*C)^{-1/2}$. By Theorem 1, we can choose $b_1 = 1 - \epsilon$ and $b_2 = 1 + \epsilon$. Then p and q are constants that depend only on ϵ but not on n. By choosing $\epsilon < 1$, we have

$$\frac{b-1}{b+1} = \frac{1 - \sqrt{1 - \epsilon^2}}{\epsilon} < \epsilon.$$

In order to use (16), we need a lower bound for δ_k, $1 \le k \le p$. We first note that

$$||G^{-1}||_2 = ||(T^*T)^{-1}(C^*C)||_2 \le \frac{||C||_2^2}{||T||_2^2} \kappa(T^*T).$$

If one of the f_ℓ has no zeros, then by (8), we have for n sufficiently large $||T||_2^2 \ge ||T_\ell||_2^2 \ge c$ for some $c > 0$ independent of n. Combining this with (12), we then see that for all n sufficiently large,

$$||G^{-1}||_2 \le \tilde{c} \cdot \kappa(T^*T) \le \tilde{c} n^\alpha,$$

for some constant $\tilde{c}$ that does not depend on n. Therefore,

$$\delta_k \ge \min_\ell \delta_\ell = \frac{1}{||G^{-1}||_2} \ge c n^{-\alpha}, \quad 1 \le k \le n.$$

Thus for $1 \le k \le p$ and $\delta \in [1 - \epsilon, 1 + \epsilon]$, we have,

$$0 \le \frac{\delta - \delta_k}{\delta_k} \le c n^\alpha.$$

Hence (16) becomes

$$\frac{||x - x_j||_G}{||x - x_0||_G} < c^p n^{p\alpha} \epsilon^{j-p-q}.$$

Given arbitrary tolerance $\tau > 0$, an upper bound for the number of iterations required to make

$$\frac{||x - x_j||_G}{||x - x_0||_G} < \tau$$

is therefore given by

$$j_0 \equiv p + q - \frac{p \log c + \alpha p \log n - \log \tau}{\log \epsilon} = O(\alpha \log n + 1).$$

Since by using FFTs, the matrix-vector products in Algorithm PCG for Least Squares can be done in $O(m \log n)$ operations for any n-vector v, the cost per iteration of the conjugate gradient method is of $O(m \log n)$. Thus we conclude that the work of solving (17) to a given accuracy τ is $O(\alpha m \log^2 n + m \log n)$ when $\alpha > 0$.

The convergence analysis given above can be further strengthened. For T an m-by-n matrix of the form (4) with $m = kn$, let $\lambda_{\min}(T_j^* T_j) = O(n^{-\alpha_j})$ for $j = 1, \cdots k$. By Lemma 1, we already know that

$$\lambda_{\min}(T_j^* T_j) \leq \lambda_{\max}(T_j^* T_j) \leq 2||f||_\infty^2,$$

therefore $\alpha_j \geq 0$. By the Cauchy interlace theorem, we see that

$$\lambda_{\min}(T^* T) \geq \sum_{j=1}^{k} \lambda_{\min}(T_j^* T_j) \geq O(n^{-\alpha}),$$

where $\alpha = \min_j \alpha_j \geq 0$. Therefore

$$\kappa(T^* T) \leq \frac{\lambda_{\max}(T^* T)}{\lambda_{\min}(T^* T)} \leq O(n^\alpha).$$

In the case when one of the $\alpha_j = 0$, i.e. the block T_j is well-conditioned independent of n, we see that the least squares problem is also well-conditioned, so that $\kappa(T) = O(1)$.

When at least one $\alpha_j = 0$, *i.e.*, $\kappa(T) = O(1)$, the number of iterations required for convergence is of $O(1)$. Hence the complexity of the algorithm reduces to $O(m \log n)$. We remark that in this case, one can show further that the *method converges superlinearly* for the preconditioned least squares problem due to the clustering of the singular values for sufficiently large n (See R. Chan and Strang [3] or R. Chan [5] for details). In contrast, the method converges just linearly for the non-preconditioned case. This contrast is illustrated very well in the section on numerical tests.

5 Numerical Tests

In this section we report on some numerical experiments which use the preconditioner C given by equation (5) in §1 for the conjugate gradient algorithm PCG for solving Toeplitz and block Toeplitz least squares problems. Here the preconditioner C is based on the T. Chan optimal preconditioner C_i, for each block T_i of T, as in §2. The experiments are designed to illustrate the performance of the preconditioner on a variety of problems, including some in which one or more Toeplitz blocks are very ill-conditioned.

For all numerical tests given in this section we use the stopping criteria $||s^{(j)}||_2/||s^{(0)}||_2 < 10^{-7}$, where $s^{(j)}$ is the (normal equations) residual after j iterations, and the zero vector is our initial guess. (Observe that the value $||s^{(j)}||_2$ is computed as part of the conjugate gradient algorithm.) All experiments were performed using the Pro-Matlab software on our workstations. The machine epsilon for Pro-Matlab on this system is approximately 2.2×10^{-16}.

To describe most of the Toeplitz matrices used in the examples below, we use the following notation. Let the m-vector c be the first column of T, and the n-vector r^T be the first row of T. Then $T = \text{Toep}(c, r)$. The right hand side vector b is generally chosen to be the vector of all ones.

Example 1: In this example we use the following three generating functions in the Wiener class to construct a $3n \times n$ block Toeplitz matrix.

(*i*) $c_1(j) = r_1(j) = (|j-1|+1)^{-1.1} + \sqrt{-1}(|j-1|+1)^{-1.1}, \quad j = 1, 2, \ldots, n.$

(*ii*) $c_2(i) = (|i-1|+1)^{-1.1}, \quad i = 1, 2, \ldots, n,$
$r_2(j) = \sqrt{-1}(|j-1|+1)^{-1.1}, \quad j = 1, 2, \ldots, n.$

(*iii*) $c_3(1) = r_3(1) = \frac{1}{5}\pi^4$
$c_3(j) = r_3(j) = 4(-1)^{(j-1)}(\frac{\pi^2}{(j-1)^2} - \frac{6}{(j-1)^4}), \quad j = 2, 3, \ldots, n.$

The matrix T is defined as $T^T = [T_1^T,\ T_2^T,\ T_3^T]$, where T_1 =Toep(c_1, r_1), T_2 =Toep(c_2, r_2) and T_3 =Toep(c_3, r_3). For $n \times n$ systems R. Chan and Yeung [8] show that $\kappa_2(T_3) = O(n^4)$, while T_1 and T_2 are well-conditioned. They also show that T. Chan's preconditioner works well for T_1 and T_2, but not well for T_3.

In Table 1 we show the convergence results for this example, using no preconditioner and C as a preconditioner, for several values of m and n. Figure 1 shows the singular values of T and TC^{-1} for $m = 210$ and $n = 70$. These results illustrate the good convergence properties using the preconditioner C for this example containing an ill-conditioned block. Moreover, our computations verify the fact that $\kappa_2(T)$ remains almost constant as n increases from 40 to 80.

Example 2: In this example we form a $2n \times n$ block Toeplitz matrix using generating functions from R. Chan and Yeung [8] which construct ill-conditioned $n \times n$ Toeplitz matrices. Here $T_1 = T_2$ and thus both blocks of T are ill-conditioned. The generating function, which is in the Wiener class, is:
$c_1(1) = r_1(1) = 0$
$c_1(j) = r_1(j) = (|j-1|+1)^{-1.1} + \sqrt{-1}(|j-1|+1)^{-1.1}, \quad j = 2, \ldots, n.$
Using the above generating functions, we let $T^T = [T_1,\ T_2]^T$, where $T_1 = T_2$ =Toep(c_1, r_1).

In Table 1 we show the convergence results for this example, using no preconditioner and C as a preconditioner, for several values of m and n. Figure 2 shows the singular values of T and TC^{-1} for $m = 140$ and $n = 70$. These results illustrate the good convergence properties of C for this example even though it contains all ill-conditioned blocks.

In summary, we have shown how to construct circulant preconditioners for the efficient solution of a wide class of Toeplitz least squares problems. The numerical experiments given collaborate our convergence analysis. Examples 1 and 2 both illustrate superlinear convergence for the PCG algorithm preconditioned by C, even when in Example 1 the matrix T contains an ill-conditioned block. In addition, even though the matrix T in Example 2 contains *all* ill-conditioned blocks, the scheme works well for the computations we performed.

2–dimensional signal or image restoration computations often lead to very large least squares problems where the coefficient matrix is block Toeplitz with Toeplitz blocks. Block circulant preconditioners for this case are considered elsewhere [9].

In this paper we have used the T. Chan [10] preconditioner for the Toeplitz blocks. Other circulant preconditioners such as ones studied by R. Chan [5], Huckle [14], Ku and Kuo [15], Strang [19], Tismenetsky [20], or Tyrtyshnikov [21], can be used, but the class of generating functions may need to be restricted for the convergence analysis to hold.

References

[1] O. Axelsson and G. Lindskog, *On the rate of convergence of the preconditioned conjugate gradient algorithm*, Numer. Math. V48 (1986), pp. 499–523.

[2] A. Bjorck, *Least Squares Methods*, in Handbook of Numerical Methods, ed. P. Ciarlet and J. Lions, Elsevier/North Holland Vol. 1, 1989.

[3] R. Chan and G. Strang, *Toeplitz Equations by Conjugate Gradients with Circulant Preconditioner*, SIAM J. Sci. Statist. Comput., V10 (1989), pp. 104–119.

[4] R. Chan, *The Spectrum of a Family of Circulant Preconditioned Toe-plitz Systems*, SIAM J. Numer. Anal., V26 (1989), pp. 503-506.

[5] R. Chan, *Circulant Preconditioners for Hermitian Toeplitz Systems*, SIAM J. Matrix Anal. Appl., V10 (1989), pp. 542–550.

[6] R. Chan, X. Jin and M. Yeung, *The Spectra of Super-optimal Circulant Preconditioned Systems*, SIAM J. Numer. Anal., V28 (1991), pp. 871–879.

[7] R. Chan and M. Yeung, *Circulant Preconditioners for Toeplitz Matrices with Positive Continuous Generating Functions*, Math. Comp., V58 (1992), pp. 233-240.

[8] R. Chan and M. Yeung, *Circulant Preconditioners for Complex Toeplitz Matrices*, Research Report 91-6, Dept. of Math., Univ. of Hong Kong.

[9] R. Chan, J. Nagy and R. Plemmons, *Block Circulant Preconditioners for 2–dimensional Deconvolution Problems*, in preparation.

[10] T. Chan, *An Optimal Circulant Preconditioner for Toeplitz Systems*, SIAM J. Sci. Statist. Comput., V9 (1988), pp. 766–771.

[11] T. Chan and J. Olkin, *Preconditioners for Toeplitz–block matrices*, preprint, (1991).

[12] P. Davis, *Circulant Matrices*, John Wiley & Sons, Inc., New York, 1979.

[13] G. Golub and C. van Loan, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, MD, 2nd Edition, 1989.

[14] T. Huckle, *Circulant and Skew-circulant Matrices for Solving Toeplitz Matrix Problems*, in Proc. Copper Mountain Conference on Iterative Methods, Copper Mountain, Colorado, 1990.

[15] T. Ku and C. Kuo, *Design and Analysis of Toeplitz Preconditioners*, IEEE Trans. Acoust. Speech Signal Process., to appear.

[16] J. Nagy, *Toeplitz Least Squares Computations*, Ph.D. thesis, North Carolina State University, directed by R. J. Plemmons.

[17] J. Nagy and R. Plemmons, *Some Fast Toeplitz Least Squares Algorithms*, Proc. SPIE Conference on Advanced Signal Processing Algorithms, Architectures, and Implementations II, V1566, San Diego, CA, July 1991.

[18] A. Oppenheim and R. Schafer, *Discrete-time Signal Processing*, Prentice-Hall, Englewood Cliffs, NJ 1989.

[19] G. Strang, *A Proposal for Toeplitz Matrix Calculations*, Stud. Appl. Math., V74 (1986), pp. 171–176.

[20] M. Tismenetsky, *A Decomposition of Toeplitz Matrices and Optimal Circulant Preconditioning*, Lin. Alg. and Applic. Special Issue on Iterative Methods, V154–156 (1991), pp. 105–121.

[21] E. Tyrtyshnikov, *Optimal and Super-optimal Circulant Preconditioners*, SIAM J. Matrix Anal. Appl., to appear.

[22] J. Wilkinson, *The Algebraic Eigenvalue Problem*, Clarendon Press, Oxford, 1965.

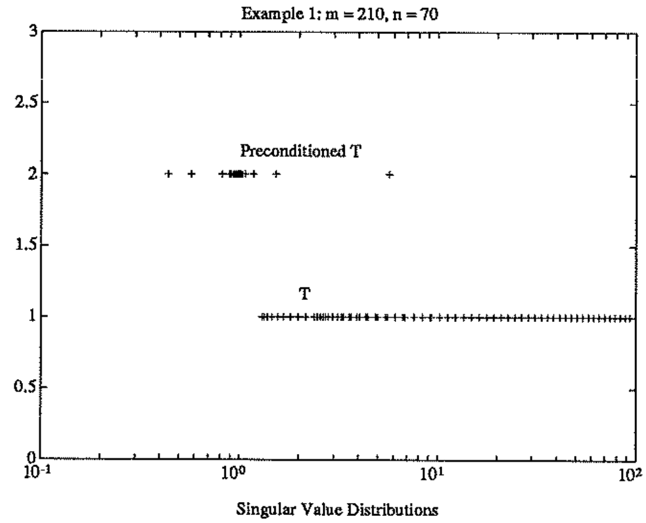

Figure 1. Singular values for T and TC^{-1} in Example 1.

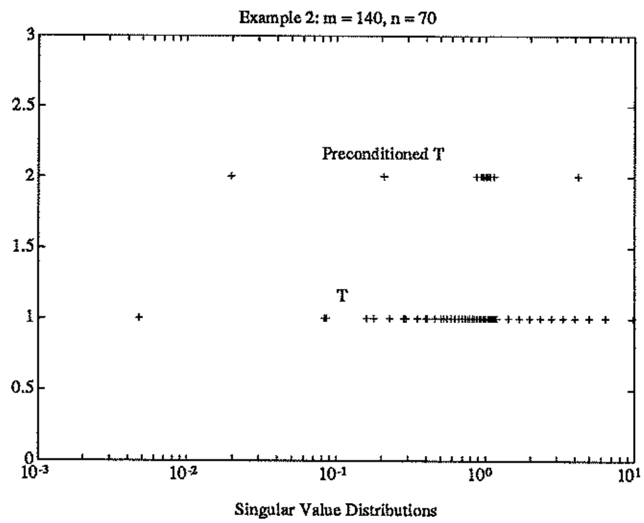

Figure 2. Singular values for T and TC^{-1} in Example 2.

	Example 1 ($m = 3n$)		Example 2 ($m = 2n$)	
n	no prec.	with prec.	no prec.	with prec.
40	96	14	29	11
50	126	14	33	15
60	155	13	44	13
70	167	13	52	12
80	186	13	65	14

Table 1. Numbers of iterations for convergence in Examples 1 - 2.