A Comparison of Two Equivalent Real Formulations for Complex-Valued Linear Systems Part 1: Introduction and Method

Abnita Munankarmy and Michael A. Heroux
Department of Computer Science
College of Saint Benedict
37 South College Avenue
St. Joseph, Minnesota 56374 USA

Received May 14, 2002 *Accepted October 22, 2002*

ABSTRACT

Many iterative linear solver packages focus on real-valued systems and do not deal well with complex-valued systems, even though preconditioned iterative methods typically apply to both real and complex-valued linear systems. Instead, commonly available packages such as *PETSc* [1] and *Aztec* [2] tend to focus on the real-valued systems, while complex-valued systems are seen as a late addition. At the same time, by changing the complex problem into an equivalent real formulation (ERF), a real valued solver can be used. In this paper we consider two ERF's that can be used to solve complex-valued linear systems. We investigate the spectral properties of each and show how each can be preconditioned to move eigenvalues in a cloud around the point (1,0) in the complex plane. Finally, we consider an interleaved formulation, combining each of the previously mentioned approaches, and show that the interleaved form achieves a better outcome than either separate ERF.

I. INTRODUCTION

This paper describes an effective extension of a real-valued preconditioned iterative solver package for complex-valued linear systems

$$C w = d, \qquad (1)$$

where C is an m-by-n complex matrix, d is a known n-by-1vector and w is an unknown n-by-1 vector. There are few complex-valued solver packages available to solve this system. Here we solve this system using an equivalent real formulation (ERF). This paper explores several topics left for further research in the paper by Day and Heroux [3], who first introduced the topic.

II. PROPECTIVE ERFs

In order to derive ERFs for equation (1), we begin by writing (1) in terms of its real and imaginary parts. Specifically,

$$(A + iB)(x + iy) = b + ic$$

where $C = A + iB$, $w = x + iy$ and $d = b + ic$. Doing this, we can consider four possible 2-by-2 block formulations as described by Day and Heroux [3]. These are shown below in equations (2a) - (2d), and their solutions are equivalent to those for equation (1). We shall denote these formulations by K1, K2, K3, and K4.

K1 formulation,

$$\begin{pmatrix} A & -B \\ B & A \end{pmatrix}\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} b \\ c \end{pmatrix} \qquad (2a)$$

K2 formulation,

$$\begin{pmatrix} A & B \\ B & -A \end{pmatrix}\begin{pmatrix} x \\ -y \end{pmatrix} = \begin{pmatrix} b \\ c \end{pmatrix} \qquad (2b)$$

K3 formulation,

$$\begin{pmatrix} B & A \\ A & -B \end{pmatrix}\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} c \\ b \end{pmatrix} \qquad (2c)$$

and the K4 formulation,

$$\begin{pmatrix} B & -A \\ A & B \end{pmatrix}\begin{pmatrix} x \\ -y \end{pmatrix} = \begin{pmatrix} c \\ b \end{pmatrix} \qquad (2d)$$

For future reference, we shall denote the matrix associated with the K1 to K4 formulations by K_1 to K_4, respectively. Along with these four real formulations, there is also a combined formulation K14, which is expressed as an interleaved ERF. In the future, the matrix associated with K14 will be denoted by K_{14}. We will discuss K_{14} in Section IV.

The eigenvalues of the matrices in the K2 and K3 formulations are in problematic configuration because any hull containing the spectra necessarily contains the origin, which degrades the convergence rate of an iterative method such as GMRES [5]. The other reasons for this are beyond the scope of this paper. Matrices in the K1 and K4 formulations are more promising because they have more favorable spectral properties, and can be preconditioned to move eigenvalues in a cloud around the point (1,0) in the complex plane. This makes the convergence rate of an iterative method, such as GMRES [5], more robust. It also impacts on the reflection of the eigenvalues of the complex matrix and the ERFs K1, K4 and K14 respectively. Table 1 summarizes the spectral properties of K_1 to K_{14}. Since K1 and K4 formulations are more promising than K2 and K3, we want to determine if either K1 or K4 would be the best one to find a solution. Success with the K1 and K4 formulations depends on the quality of the preconditioner but the slower rate is due to the spectral properties of the ERF.

For the K1 formulation, the spectrum of K_1 should not present a major dilemma to an iterative method such as GMRES, given all the eigenvalues of C are on one side of the imaginary axis. However, K_1 will have twice as many problematic eigenvalues if C has eigenvalues on both sides of the imaginary axis and this is a property that degrades the GMRES convergence rate. Most importantly, the convex hull containing the eigenvalues of K_1 will also contain the origin. For the K4 formulation, the eigenvalues of K_4 will be in the right half plane as long as all the eigenvalues of C are in the upper half plane. The K4 formulation is actually a constant variation of the K1 formulation.

It has been shown [3] that if C is Hermitian then K_1 is symmetric and the convergence rate of an ERF is identical to the convergence rate with the original complex formulation. However, if this is not the case there are different results based on the eigenvalues of the complex matrix. This will be demonstrated below.

III. BASIC PROPERTIES OF ERFs

In this section we present a simple example to illustrate the important properties of each ERF. This example illustrates the distribution of the eigenvalues of a complex matrix C that has eigenvalues on one side of the imaginary axis, and then shows how that impacts the eigenvalues distribution of K_1, K_4, and K_{14}, respectively.

The original matrix C is a tridiagonal matrix of the form

$$C = \begin{bmatrix}
2+i & 2+i & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1+\frac{1}{2}i & 4+2i & 3+\frac{3}{2}i & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 2+i & 6+3i & 4+2i & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 3+\frac{3}{2}i & 8+4i & 5+\frac{5}{2}i & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 4+2i & 10+5i & 6+3i & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 5+\frac{5}{2}i & 12+6i & 7+\frac{7}{2}i & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 6+3i & 14+7i & 8+4i & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 7+\frac{7}{2}i & 16+8i & 9+\frac{9}{2}i & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 8+4i & 18+9i & 10+5i \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 9+\frac{9}{2}i & 20+10i
\end{bmatrix} \qquad (3)$$

Figures 1, 2 and 3 show scatter plots of the eigenvalues for the matrix C and its K1 and K4 formulations. Denoting the spectrum of a matrix A by $\sigma(A)$, we see from these figures that $\sigma(K_1)$ contains the $\sigma(C)$ and the reflection of $\sigma(C)$ across the real axis. $\sigma(K_4)$ is obtained from the $\sigma(C)$ by rotating $\sigma(C)$ by 90 degrees clockwise and then reflecting across the real axis. For the K14 formulation, the eigenvalues of the matrix K_{14} act the same way as those of matrix K_1.

In general, we can summarize the spectral properties for the K1 through K4 formulations using Table 1, reproduced from [3].

Matrix	**Spectral Properties**
K_1	(i) If $\lambda \in \sigma(C)$ then $\lambda, \lambda^{-1} \in \sigma(K_1)$. (ii) If C is Hermitian (positive definite) then K_1 is symmetric (positive definite).
K_2	(i) If $\lambda \in \sigma(K_2)$ then $-\lambda, \lambda^{-1}, -\lambda^{-1} \in \sigma(K_2)$. (ii) If C is symmetric then K_2 is symmetric.
K_3	(i) If $\lambda \in \sigma(K_3)$ then $-\lambda, \lambda^{-1}, -\lambda^{-1} \in \sigma(K_3)$. (ii) If C is symmetric then K_3 is symmetric. (iii) $\sigma(K_3) = \sigma(K_2)$.
K_4	(i) If $\lambda \in \sigma(C)$ then $-i\lambda, i\lambda^{-1} \in \sigma(K_4)$. (ii) If C is Hermitian (positive definite) then K_4 is skew symmetric (with eigenvalues that have positive imaginary parts).
K_{14}	(i) If $\lambda \in \sigma(C)$ then $\lambda, \lambda^{-1} \in \sigma(K_{14})$. (ii) If C is Hermitian (positive definite) and the diagonal values have larger real part, then K_{14} is symmetric (positive definite). (iii) If C is Hermitian (positive definite) and the diagonal values have larger imaginary part, then K_4 is skew symmetric (with eigenvalues that have positive imaginary parts).

Table 1. Spectral Properties of the K formulations. $\sigma(K)$ denotes the spectrum of K and $i = \sqrt{-1}$

IV. INTERLEAVED ERF

From the example in section III, we can see intuitively that the K1 and K4 formulations will be desirable ERF's in a different setting. When $\sigma(C)$ tends to have eigenvalues with large real parts, K1 would tend to cluster the eigenvalues around (1,0) in the complex plane. When these eigenvalues have large imaginary parts, K4 would be a better formulation. In this section we introduce the K14 formulation that attempts to provide a combined approach using a simple heuristic. In the K14 formulation, we interleave the individual K1 and K4 formulations for each equation of the complex matrix. Interleaving is done by permuting and scaling the real matrix. Permutation of the matrix is necessary for both theoretical and computational reasons. The process of interleaving is demonstrated below. Success with the K14 formulation depends on the real and imaginary parts of the eigenvalues. The major part of the solution depends on eigenvalues. Eigenvalues help to be familiar with the speed of the iterative method and the better the iterative method is, the better the solution. Success also depends on the quality of the preconditioner. In fact, our experience shows that for the classes of problems we are solving, in particular eigenvalues problems for computational fluid dynamics, interleaving leads to better diagonal precondtioners. The following example demonstrates an interleave process for a simple 2-by-2 complex matrix for form a K14 formulation.

Let C w = d be a 2-by-2 complex system with C explicitly defined as follows:

$$\begin{bmatrix} 1+2i & 3+i \\ 4+2i & 5+i \end{bmatrix}\begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} d_1 \\ d_2 \end{bmatrix}.$$

Then the K1 and K4 formulations of the linear system are as follows:

For K1,

$$\begin{bmatrix} 1 & -2 & 3 & -1 \\ 2 & 1 & 1 & 3 \\ 4 & -2 & 5 & -1 \\ 2 & 4 & 1 & 5 \end{bmatrix}\begin{bmatrix} x_1 \\ y_1 \\ x_2 \\ y_2 \end{bmatrix} = \begin{bmatrix} b_1 \\ c_1 \\ b_2 \\ c_2 \end{bmatrix}.$$

Writing out this matrix equation as a linear system of equations gives,

$$\begin{aligned} x_1 - 2y_1 + 3x_2 - y_2 &= b_1 \\ 2x_1 + y_1 + x_2 + 3y_2 &= c_1 \\ 4x_1 - 2y_1 + 5x_2 - y_2 &= b_2 \\ 2x_1 + 4y_1 + x_2 + 5y_2 &= c_2 \end{aligned}$$

For K4,

$$\begin{bmatrix} 2 & -1 & 1 & -3 \\ 1 & 2 & 3 & 1 \\ 2 & -4 & 1 & -5 \\ 4 & 2 & 5 & 1 \end{bmatrix}\begin{bmatrix} x_1 \\ -y_1 \\ x_2 \\ -y_2 \end{bmatrix} = \begin{bmatrix} b_1 \\ c_1 \\ b_2 \\ c_2 \end{bmatrix}.$$

The system of linear equations for the K4 matrix formulation is

$$\begin{aligned} 2x_1 + y_1 + x_2 + 3y_2 &= b_1 \\ x_1 - 2y_1 + 3x_2 - y_2 &= c_1 \\ 2x_1 + 4y_1 + x_2 + 5y_2 &= b_2 \\ 4x_1 - 2y_1 + 5x_2 - y_2 &= c_2 \end{aligned}$$

We now observe that both the formulations have the same linear solutions. It is always true that K1 and K4 formulations should have the same solutions, and it is also true that the K14 formulation should have exactly the same solutions as the K1 and the K4 formulations.

The K14 formulation interleaves the K1 and K4 formulations on an equation-by-equation basis by looking at the diagonal entries of the complex matrix C. In particular, for each diagonal entry, if the imaginary part of the entry is greater then the real part, then we use K4 for that equation. Otherwise, we use the K1 formulation. For the example above, the K14 formulation is,

$$\begin{bmatrix} 2 & -1 & 1 & 3 \\ 1 & 2 & 3 & -1 \\ 4 & 2 & 5 & -1 \\ 2 & -4 & 1 & 5 \end{bmatrix}\begin{bmatrix} x_1 \\ -y_1 \\ x_2 \\ y_2 \end{bmatrix} = \begin{bmatrix} c_1 \\ b_1 \\ b_2 \\ c_2 \end{bmatrix}.$$

Writing out the system of equations for this matrix equation gives,

$$\begin{aligned} 2x_1 + y_1 + x_2 + 3y_2 &= c_1 \\ x_1 - 2y_1 + 3x_2 - y_2 &= b_1 \\ 4x_1 - 2y_1 + 5x_2 - y_2 &= b_2 \\ 2x_1 + 4y_1 + x_2 + 5y_2 &= c_2 \end{aligned}$$

In examining the K14 formulation, we observe that the diagonal values are larger in size than the diagonal values in the individual formulations. When the diagonal values are larger, many preconditoners tend to be more effective. In fact, one could introduce a thesis for matrices with larger diagonal values and the effect of using the preconditioners such as Jacobi and Gaussian Elimination methods.

V. PRECONDITIONING ERFs

Preconditioning can be done in many different ways, and hence in trying to solve the original complex system in equation (1) via the K1 formulation in equation (2a), the most fascinating question is how to precondition K_1. Standard real-valued preconditioners such as Jacobi, Gauss-Seidel or ILU applied directly to K_1 are not robust enough for our needs. Furthermore, the ordering of the unknowns is unsuitable for sparse matrix operations related to factorization, particularly ILU preconditioning [3]. A good preconditioner makes the iterations of K formulations comparable to that of solving the original

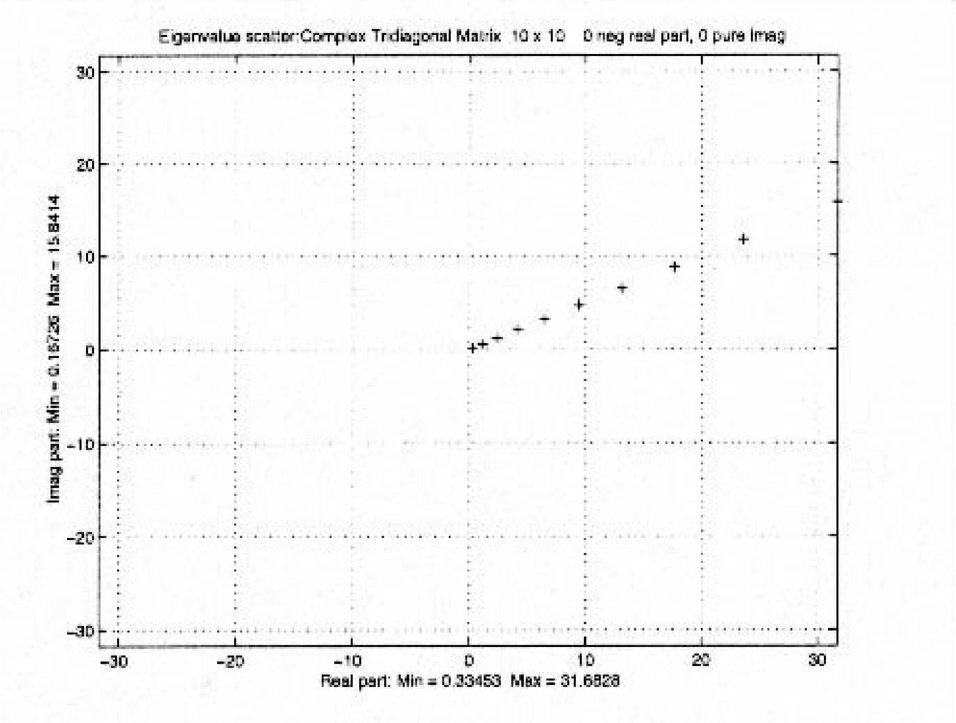

Figure 1. Eigenvalues of the complex tridiagonal matrix.

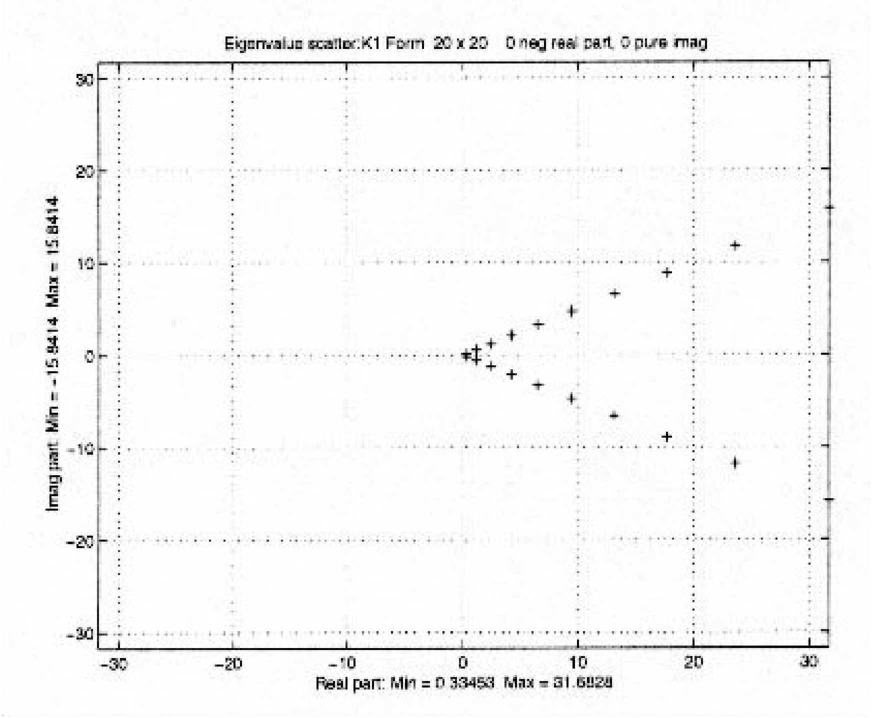

Figure 2. Eigenvalues of the K1 formulation matrix.

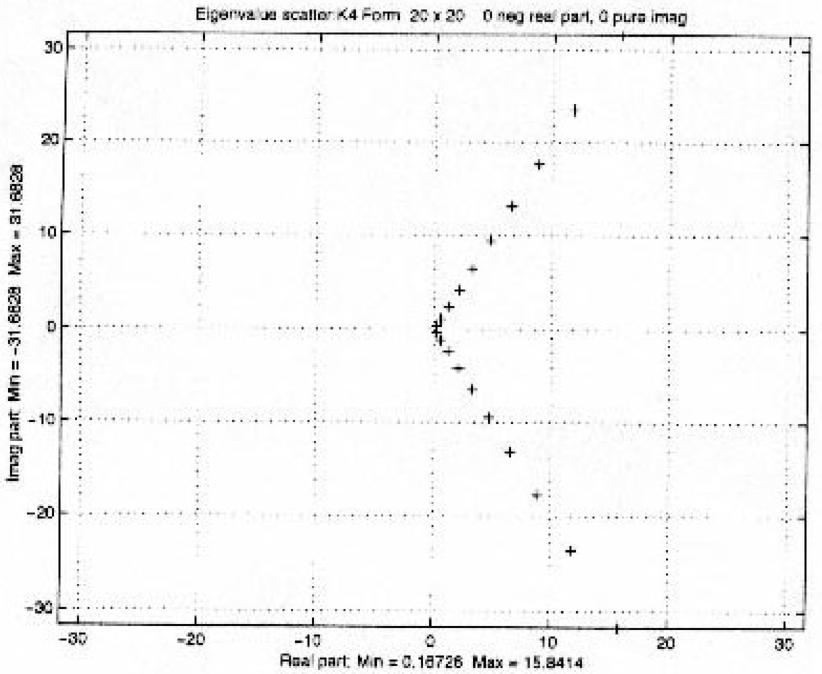

Figure 3. Eigenvalues of the K4 formulation matrix.

complex problem, with a true complex preconditioned iterative solver. Convergence often can be accelerated substantially by preconditioning, which can be thought of as implicitly multiplying matrix A by matrix M^{-1}, where matrix M is a preconditioner that we define. We start with a matrix equation

$$A\,x = b\,, \qquad (1)$$

and assume that there exist a lower triangular matrix, L, and an upper triangular matrix, U, such that A = LU. Further assuming that both L and U have inverses, equation (1) can be rewritten as follows,

$$\begin{aligned}(L\,U)\,x &= b\\ (L\,U)\,I\,x &= b\\ (L\,U)(U^{-1}\,U)\,x &= b\end{aligned}$$

where I is the identity matrix. Multiplying through by L^{-1} gives,

$$L^{-1}\,(L\,U)(U^{-1}\,U)\,x = L^{-1}\,b$$

or

$$(L^{-1}\,L\,)(U\,U^{-1}\,)U\,x = L^{-1}\,b\,.$$

Since L^{-1} L = I and U^{-1} U = I, then

$$U\,x = L^{-1}\,b\,,$$

or, multiplying through with U^{-1},

$$x = U^{-1}\,L^{-1}\,b\,. \qquad (2)$$

The goal of preconditioning is to find L and U such that A ≈ L U. The better this approximation, the better the solution, equation (2), to equation (1).

For this research, we use the Incomplete LU (ILU) factorization preconditioner discussed by Saad [4]. This produces a lower triangular matrix, an upper triangular matrix, and a permutation matrix.

VI. FORMULATION OF K_1, K_4, AND K_{14}

The ILU factorization is often used for sparse matrices, and the approach that gives a better result is one that preserves the sparsity pattern of the complete matrix C. The entries of a block entry sparse matrix are all (small) dense (sub-) matrices. The K formulation preserves the non-zero pattern of the block entries with the effect of doubling the size of each dense sub-matrix.

In the K formulation, $c_{pq} = a_{pq} + i\, b_{pq}$ corresponds, via the scalar K1 formulation, to the 2-by-2 block entry of the 2m-by-2n real matrix K given by

$$\begin{pmatrix} a_{pq} & -b_{pq} \\ b_{pq} & a_{pq} \end{pmatrix}$$

For example, if

$$C = \begin{pmatrix} c_{11} & 0 & c_{13} & 0 & c_{15} \\ 0 & c_{22} & c_{23} & 0 & 0 \\ c_{31} & 0 & c_{33} & c_{34} & 0 \\ 0 & 0 & c_{43} & c_{44} & 0 \\ c_{51} & 0 & 0 & 0 & c_{55} \end{pmatrix}$$

then,

$$K = \begin{pmatrix} a_{11} & -b_{11} & 0 & 0 & a_{13} & -b_{13} & 0 & 0 & a_{15} & -b_{15} \\ b_{11} & a_{11} & 0 & 0 & b_{13} & a_{13} & 0 & 0 & b_{15} & a_{15} \\ 0 & 0 & a_{22} & -b_{22} & a_{23} & -b_{23} & 0 & 0 & 0 & 0 \\ 0 & 0 & b_{22} & a_{22} & b_{23} & a_{23} & 0 & 0 & 0 & 0 \\ a_{31} & -b_{31} & 0 & 0 & a_{33} & -b_{33} & a_{34} & -b_{44} & 0 & 0 \\ b_{31} & a_{31} & 0 & 0 & b_{33} & a_{33} & a_{34} & -b_{34} & 0 & 0 \\ 0 & 0 & 0 & 0 & a_{43} & -b_{43} & a_{44} & -b_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & b_{43} & a_{43} & a_{44} & -b_{44} & 0 & 0 \\ a_{51} & -b_{51} & 0 & 0 & 0 & 0 & 0 & 0 & a_{55} & -b_{55} \\ b_{51} & a_{51} & 0 & 0 & 0 & 0 & 0 & 0 & b_{55} & a_{55} \end{pmatrix}$$

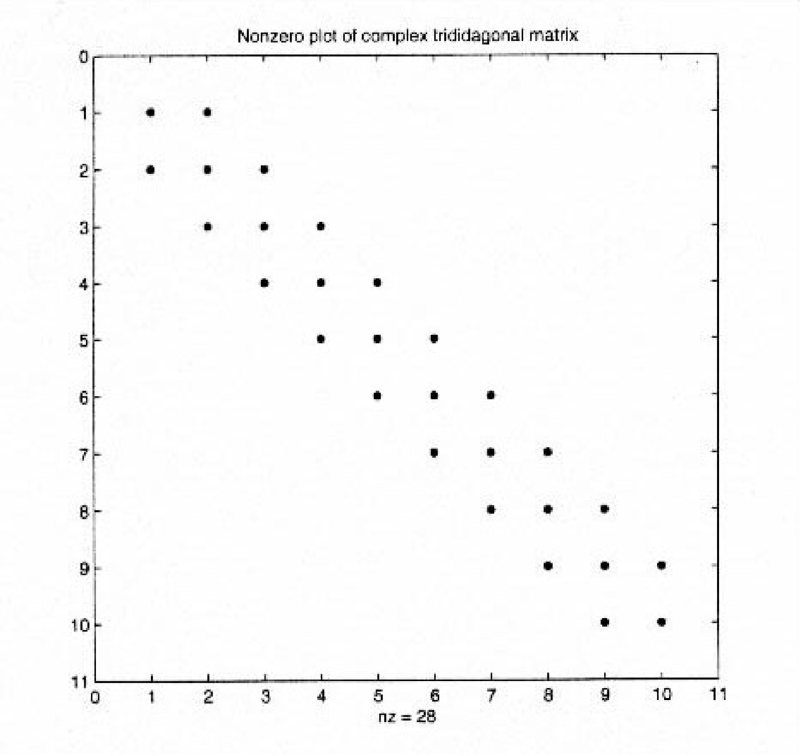

Figure 4. Non-zero pattern of a complex tridiagonal matrix.

In the related real formulations, this is a way to preserve the non-zero structure in the block entries. We can see the tridiagonal pattern in Figure 4. We can generate the matrices K_4 and K_{14} in a similar manner via the scalar K4 and K14 formulations, respectively. Figure 4 shows the non-zero pattern associated with the matrix C in Equation (3). Figure 5 shows the non-zero pattern that would be generated by any of the K formulations. Note how it preserves the sparsity pattern, but in block form.

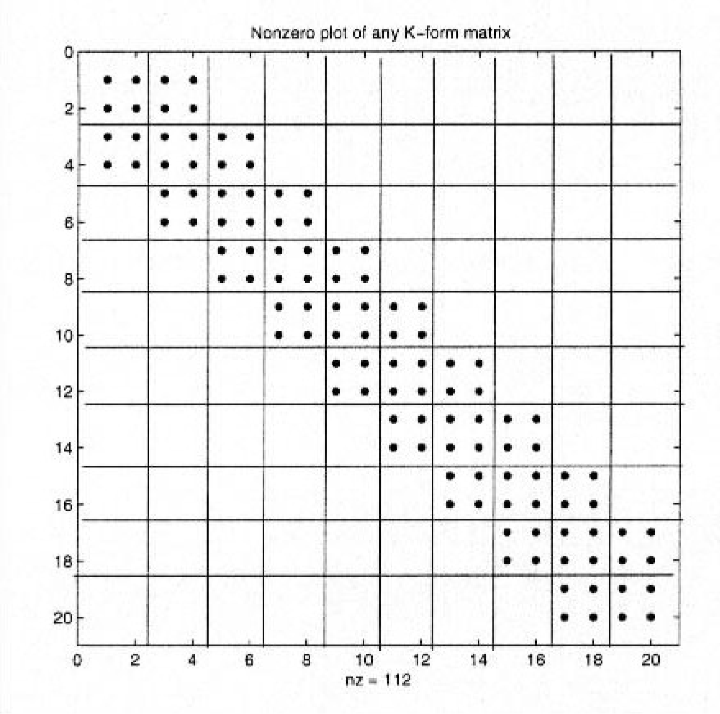

Figure 5. Non-zero pattern of any of the K formulations of the complex tridiagonal matrix.

VII. PRECONDITIONING IMPLEMENTATION

The properties of the K formulation defined in section VI enable the preconditioned iterative solvers for complex linear systems to provide a better result. We can efficiently compute and apply the exact equivalent of a complex-valued pre-conditioner. K1 and K4 formulations have nice spectral properties; hence these formulations lead to convergence that is competitive with the true complex solver. Similarly, the properties of the interleaved formulation defined in section IV enable us to re-boost the solution. The K14 formulation can be used to obtain better ILU factors.

The command luinc, from the software MATLAB, was used to produce a unit lower triangular matrix, an upper triangular matrix, and a permutation matrix. We implemented the precondition in two ways. One way was by producing L and U from the original complex matrix, and using the same preconditioning to find the solutions for the K1 and K4 formulations. For example, let *C* be the complex matrix. We get L and U from [L, U] = luinc(*C*,'0'). Transforming this into the K1 transformation of L we get L1 = K1(L). What happens here is that the new L1 is no longer a triangular matrix. We can see this in the following example:

Let L be a 2-by-2 lower triangular matrix,

$$L = \begin{pmatrix} l_{11} & 0 \\ l_{21} & l_{22} \end{pmatrix}.$$

Converting this in the K1 formulation, we have the matrix

$$K_1 = \begin{pmatrix} l_{11}^{r} & -l_{11}^{i} & 0 & 0 \\ l_{11}^{i} & l_{11}^{r} & 0 & 0 \\ l_{21}^{r} & -l_{21}^{i} & l_{22}^{r} & -l_{22}^{i} \\ l_{21}^{i} & l_{21}^{r} & l_{22}^{i} & l_{22}^{r} \end{pmatrix}$$

The matrix K_1 clearly shows that the transformation is no longer a lower triangular matrix. This is the reason why we use the same L and U. Similar reactions occur for the K4 formulation.

The second way to implement the preconditioner was to simply generate an individual preconditioner for each of the formulations K1, K4, and K14. We also tried diagonal preconditioning because some of the problems we were using for computation were not giving right answers with an ILU preconditioner. For problems like M3D2, M4D2 and vm214img45 (see Table 2), the ILU preconditioner was pivoting and permuting large matrices. Hence for these matrices we used a diagonal preconditioner, a weak preconditioner.

VIII. PREVIEW OF COMPUTATIONAL RESULTS

We have used the K1, K4, and K14 formulations to solve several complex liner systems, starting with simple 10-by-10 matrices and on up to larger dimension matrices (Table 2). Each large system comes from a real application, and most of these problems are ill conditioned. Hence, preconditioners like Jacobi or block Jacobi are not sufficient to move eigenvalues to a cloud around point (1, 0) in the complex plane in order to get a better solution. We found that ILU preconditioning was very effective in obtaining accurate results. Applying this preconditioner made the problems solvable using fewer iterations. However, in this research we came across one situation where a preconditioner could not be applied: if the diagonal values of a unit lower or upper triangular matrix has any zeros then the non-structure array prevented preconditioning.

Details of these results will be reported in a future communication [*Part 2: Results*, will be in the March 2003 issue—*Editor*].

REFERENCES

1. S. Balay, W. Gropp. L. McInnes, and B. Smith, *PETSc 2.0 Users Manual, Technical Report ANL-95/11 – Revision 2.0.22* (Argonne national Laboratory, Argonne, IL, 1998).
2. Ray S. Tuminaro, Michael A. Heroux, Scott A. Hutchinson, and J.N. Shadid, *Official Aztec User's Guide, Version 2.1* (Sandia national laboratories, Albuguerque, NM, 1999).
3. David Day and Michael A. Heroux, “solving Complex-Valued Linear systems via Equivalent Real Formulations” *SIAM J. Sci. Comput.*, Vol. 23, No. 2 (2001), pp. 480-498.
4. Y. Saad, *Iterative Methods for Sparse Linear Systems* (PWS Publishing Company, Boston, MA, USA, 1996).
5. Youcef Saad and Martin H. Schultz, ‘GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems,” *SIAM J. Sci. Statist. Comput.*, Vol. 7, No. 3 (july 1986), pp. 856-869.

Why Sigma Xi?

THE HONOR SOCIETY FOR SCIENTISTS AND ENGINEERS

Recent chapter events have included:

- Student research symposia
- Outreach and mentoring
- Workshop series
- Social events
- Resource activities for K-12 students and teachers
- Awards for outstanding teaching and research
- Lectureship series

Gain professional credentials

Being active in Sigma Xi demonstrates to employers and graduate schools that you have achieved a high standard of excellence. As a member of this prestigious society, you join the company of over 180 Nobel Prize winners and other distinguished scientists and engineers.

Receive research funding

Twice annually, the Sigma Xi Grants-in-Aid of Research program awards research grants to nearly 1,000 graduate and undergraduate students. As a Sigma Xi member you will be eligible to apply for grants from all available funds.

Build your professional network

Gain valuable leadership experience and make career contacts. Through local chapters, Sigma Xi provides a unique platform for interacting within a multi-disciplinary community of researchers, educators, administrators and others, both within and outside your institution.

Become a lifelong learner

Stay current in all fields with *American Scientist*, Sigma Xi's award winning magazine. As a member benefit, your *American Scientist* subscription is a valuable resource that offers articles in a variety of disciplines, resources and a comprehensive book review section.

Contribute to the scientific community

Sigma Xi members form a powerful international network at over 500 institutions. Through individual members and chapters, Sigma Xi addresses ethics, education, advocacy for research funding, public understanding of science, diversity and international networking.

Have fun

Build community and a sense of belonging with other scientists at your institution. Get involved in a special event with the help of others at your local chapter.

www.sigmaxi.org

Sigma Xi, The Scientific Research Society • 99 Alexander Drive • P.O. Box 13975 • Research Triangle Park, NC 27709 • 919-549-4691 • 800-243-6534