LAPACK Working Note 9 A Test Matrix Generation Suite

J. Demmel
A. McKenney

Courant Institute
251 Mercer Str.
New York, NY 10012

February 28, 1989

AUG 13 1992

Abstract

We discuss the design and implementation of a suite of test matrix generators for testing linear algebra software. These routines generate random matrices with certain properties which are useful for testing linear equation solving, least squares, and eigen-decomposition software. These properties include the spectrum, symmetry, bandwidth, norm, sparsity, conditioning (with respect to inversion or for the eigenproblem), type (real or complex), and storage scheme (dense, packed or banded).

1 Introduction

Testing a large numerical linear algebra library such as LAPACK [3] requires generating many different kinds of test matrices. For example, LAPACK contains linear equation solving routines for real and complex matrices which may be symmetric, Hermitian or non-symmetric, banded or dense, packed or unpacked, positive definite or indefinite, and so on. Thorough testing requires test matrices with a range of condition numbers and scalings (i.e., with norms near the overflow and underflow thresholds). There are similar requirements for testing the eigendecomposition (and singular value decomposition) routines.

To meet this need we have developed a suite of test matrix generation software written in standard FORTRAN 77 which generates random matrices with various controlled properties. For example, one option permits the generation of random real rectangular matrices with singular values forming a geometric sequence between 1 and a user specified condition number, and with user specified upper and lower bandwidths. Of course, not all options can be specified independently; for example, there is no known way to generate nontrivial random nonsymmetric matrices with a given spectrum and arbitrary given upper and lower bandwidths. The software we describe checks for consistency among the specified matrix properties. The software can also generate consistent random examples in the sense that one can generate random matrices differing only in their storage scheme (dense vs. packed vs. banded), or in the order of their rows and columns.

The three main routines we have developed are called xLATMR, xLATMS and xLATME. The leading 'x' may be 'S' (for single precision real (REAL)), 'D' (for double precision real (DOUBLE PRECISION)), 'C' (for single precision complex (COMPLEX)) and 'Z' (for double precision complex (DOUBLE COMPLEX)). xLATMR generates matrices with random offdiagonal entries, xLATMS generates both random symmetric matrices with given eigenvalues and bandwidth, and random nonsymmetric matrices with given singular values and bandwidth. xLATME generates random nonsymmetric matrices with given eigenvalues and either a given Jordan form (with certain restrictions) or given condition number for its eigenproblem. The routines can generate output in any legal LAPACK storage scheme (dense, packed, or banded).

The rest of this paper is organized as follows: Sections 2, 3 and 4 present high level descriptions of xLATMR, xLATMS and xLATME, respectively. Section 5 discusses naming and semantic conventions for common arguments of the three routines. Sections 6, 7 and 8 present the detailed calling sequence of the three routines. Finally, section 9 has implementation notes.

All variable names and FORTRAN fragments will appear in typewriter font.

2 xLATMR — High Level Description

xLATMR can generate matrices with random offdiagonal entries and given diagonal. It is the simplest (and fastest) of the three routines in the suite, and permits no direct control over the eigenvalues or singular values of the generated matrix. It is also fast and space efficient because the both the time and space required to generate a band matrix are proportional to the number of entries inside the desired bandwidth rather than the square of the matrix dimension; this efficiency is not generally possible if we wish to specify the spectrum of the resulting matrix. A high level description of its operation is as follows:

- (1) Generate a matrix A with random entries. The probability distribution used for real random number generation may be chosen from uniform on $(0, 1)$, uniform on $(-1, 1)$, and normal with zero mean and unit variance ($\text{normal}(0, 1)$). For generating complex random numbers we may chose to have real and imaginary parts independently chosen from uniform on $(0, 1)$, from uniform on $(-1, 1)$, or from normal $(0, 1)$, as well as to have the complex number chosen uniformly from the unit disk in the complex plane. A may be nonsymmetric, symmetric (real or complex), or Hermitian (complex only).
- (2) Set the diagonal of A to D (an array), where the entries of D may be computed as follows (N is the matrix dimension):
 1. Input by the user.
 2. $D(1) = 1$ and the other $D(i) = 1/\text{COND}$, where $\text{COND} \geq 1$ is a user input.
 3. $D(N) = 1$ and the other $D(i) = 1/\text{COND}$.
 4. The $D(i)$ form a geometric sequence from 1 to $1/\text{COND}$.
 5. The $D(i)$ form an arithmetic sequence from 1 to $1/\text{COND}$.
 6. The $D(i)$ are random in the range $[1/\text{COND}, 1]$ with uniformly distributed logarithms.

7. The $D(i)$ are random with the same distribution as the other matrix entries.

In addition, each $D(i)$ may optionally be multiplied by a random number with absolute value 1.

- (3) Grade A , if desired, by pre- and postmultiplying it by diagonal matrices DL and DR , respectively. The entries of DL and DR are chosen just as the entries of D above.
- (4) Permute the rows and columns of A , if desired.
- (5) Set random entries of A to zero, if desired, to get a matrix with a given fraction of zero entries.
- (6) Make A a band matrix, if desired, by zeroing out its entries outside given upper and lower bandwidths.
- (7) Scale A , if desired, to have a given maximum absolute entry.
- (8) Pack A , if desired. The allowable options are no packing, zeroing out the upper or lower half (if symmetric or Hermitian), storing the upper or lower half columnwise (if symmetric, Hermitian or triangular), using triangular band storage (upper half or lower half, and only if the matrix is symmetric, Hermitian or triangular), and full band storage.

If two calls to xLATMR differ only in the desired packing, they will generate mathematically equivalent matrices; this is convenient for testing routines which accept matrices in different storage schemes. If two calls to xLATMR both specify matrices with full bandwidth, and differ only in the order in which they permute the rows and columns (and possibly in the packing), then the matrices generated will differ only in the order of the rows and columns, and otherwise contain the same data. This consistency (which clearly cannot be attained for banded matrices, since permutations generally destroy any band structure) is useful for testing linear system solvers which perform pivoting.

3 xLATMS — High Level Description

xLATMS can generate both random symmetric (or Hermitian) matrices with given eigenvalues and bandwidth, and random nonsymmetric matrices with given singular values and upper and lower bandwidths. The space required to generate an n by n matrix is $n^2 + O(n)$ even if the matrix has a small bandwidth, and generally takes time $O(n^3)$. The reason for these time and space requirements is that a dense matrix must be reduced to banded form in a process analogous to reducing a symmetric matrix to tridiagonal form by Householder transformations [6]. The same packing options as for xLATMR are also available.

A high level description of its operation is as follows:

- (1) Set the diagonal of the matrix A to D , where the entries of D can be chosen using the same options as for xLATMR. The entries of D will be the eigenvalues (and/or singular values) of the final matrix.

- (2) Pre- and postmultiply A by random orthogonal matrices (or unitary matrices, if complex). This may be done so that the resulting A is nonsymmetric, symmetric (real or complex), or Hermitian (complex only). The resulting matrix is generically dense.
- (3) Reduce A to have the desired bandwidth using Householder transformations.
- (4) Pack A , if desired. The same options as for xLATMR are available.

4 xLATME — High Level Description

xLATME can generate random nonsymmetric matrices with given eigenvalues, a given condition number for the eigenvalues and/or Jordan form (with certain restrictions), and a given one-sided bandwidth. Thus, for example, we can generate random Hessenberg matrices with given eigenvalues and sensitivities; this is useful for testing QR iteration algorithms for the nonsymmetric eigenproblem. Only dense storage is provided, since the nonsymmetric eigenroutines only work on matrices in dense storage format. xLATME uses $n^2 + O(n)$ space and $O(n^3)$ time for the same reasons as xLATMS.

A high level description of its operation is as follows:

- (1) Set the diagonal of the matrix A to D , where D is specified as for xLATMR and xLATMS.
- (2) If A is real and complex conjugate pairs of eigenvalues are desired, certain pairs of adjacent elements of D are interpreted as the real and imaginary parts of a complex conjugate pair of eigenvalues, and A is made block diagonal with 2 by 2 blocks in the corresponding locations. If D is input by the user, the user also specifies which entries are to be interpreted as real and imaginary parts of complex conjugate eigenpairs. If D is not input by the user, each pair of adjacent entries may be randomly designated either as a pair of real eigenvalues or as real and imaginary parts of an eigenpair.
- (3) The upper triangle of A may be filled with random numbers if desired, which may be chosen with the same range of options as for xLATMR. This option may be used to partially control the Jordan form of A as follows: If A has any multiple eigenvalues (as determined by the last two steps), and the upper triangle is filled in, then there will be exactly one Jordan block per distinct eigenvalue; such a matrix is called defective. If the upper triangle is not filled in, there will only be 1 by 1 blocks in the Jordan form, even if there are multiple eigenvalues; such a matrix is called derogatory.
- (4) A may be premultiplied by a random matrix S and postmultiplied by S^{-1} , if desired. Here, S is a random dense nonsymmetric matrix whose singular values DS may be chosen with the same options as D in xLATMR. This option may be used to control the condition of A 's eigenproblem as follows: We assume the upper triangle has *not* been filled in in the last step. Then the most sensitive eigenvalue of A will have sensitivity approximately equal to $\kappa \equiv \max_i |DS(i)| / \min_i |DS(i)|$, where sensitivity means that a perturbation of norm ϵ in A will cause a change an eigenvalue by at most approximately $\kappa \cdot \epsilon$ [6,2]. The approximation arises because the true sensitivity need only be within a factor N (the dimension of A) of the condition number of SX , where

X is a diagonal matrix chosen so that the columns of SX have equal norm. In general, the columns of S will not differ too much in norm, so that the condition number of the eigenproblem may indeed be approximately controlled with this approach.

- (5) Either the upper or lower bandwidth (but not both) may be reduced to any positive value desired.
- (6) Scale A , if desired, to have a given maximum absolute entry.

5 Common Argument Conventions

As can be seen from the last three sections, the three test matrix generators share various arguments and argument conventions. To simplify the later description of the detailed calling sequences, we collect those common conventions in this section. They are

1. Output matrix — M, N, A, LDA
2. Probability Distribution — $DIST$
3. Random Number Generator Seed — $ISEED(4)$
4. Symmetry — SYM
5. Diagonal Matrix Specification — $D, MODE, COND$ (and sometimes $DMAX$ and $RSIGN$)
6. Bandwidths — KL, KU
7. Norm — $ANORM$
8. Packing Option — $PACK$
9. Error Flag — $INFO$

In addition, we follow the convention that character arguments are one character ($CHARACTER*1$) and case independent. In this document we will always use upper case, although the software recognizes lower case as well.

5.1 Output Matrix — M, N, A, LDA

The output matrix is denoted A , and has M rows and N columns. If A must be square, only N is given. LDA is the leading dimension of A in the calling routine; therefore within the test matrix generator A is dimensioned as $A(LDA, *)$. A may be $REAL$, $COMPLEX$, $DOUBLE\ PRECISION$ or $DOUBLE\ COMPLEX$. M, N and LDA are integers. All arguments are read-only.

5.2 Probability Distribution — $DIST$

$DIST$ is a read-only $CHARACTER*1$ variable. If A is real ($REAL$ or $DOUBLE\ PRECISION$), the following options are available for $DIST$:

$DIST = 'U'$ — Uniform distribution on $(0, 1)$.

DIST = 'S' — Uniform distribution on $(-1, 1)$.

DIST = 'N' — Normal $(0, 1)$.

If A is complex (COMPLEX or DOUBLE COMPLEX), the following options are available for DIST:

DIST = 'U' — Both real and imaginary parts are independent and uniformly distributed on $(0, 1)$.

DIST = 'S' — Both real and imaginary parts are independent and uniformly distributed on $(-1, 1)$.

DIST = 'N' — Both real and imaginary parts are independent and normally distributed with zero mean and unit variance.

DIST = 'D' — The complex number is uniformly distributed inside the unit disk in the complex plane.

5.3 Random Number Generator Seed — ISEED

ISEED is an array of 4 integers, which are the seeds of the random number generator. The random number generator will operate identically on any machine with at least 24 bit integers, and generate a random number sequence with period 2^{48} . By specifying ISEED on entry, a specific random sequence can be generated. ISEED is modified by the routines.

5.4 Symmetry — SYM

SYM is a read-only CHARACTER*1 variable. It has the following interpretations (in all cases A may be real or complex):

SYM = 'N' — Nonsymmetric.

SYM = 'S' — Symmetric.

SYM = 'H' — Hermitian.

SYM = 'P' — Positive semidefinite.

5.5 Diagonal Matrix Specification — D, MODE, COND, (DMAX, RSIGN)

There are various array arguments like D which are meant to be the diagonal entries of matrices. All of them may be computed using the options MODE and COND; only some use DMAX and RSIGN as well. In all cases, D is an array of n REAL, DOUBLE PRECISION, COMPLEX or DOUBLE COMPLEX numbers which may be modified by the routine (n may equal either N or M, depending on the situation). MODE is a read-only integer. COND is a read-only REAL or DOUBLE PRECISION variable, which must be at least 1 if it is referenced. DMAX is a read-only variable of the same type as D. RSIGN is a read-only CHARACTER*1 variable. MODE and COND are interpreted as follows:

MODE = 0 — D is supplied on entry by the user, in which case it is not modified by the program.

- MODE = 1 — $D(1)$ is set to 1 and the other $D(i)$ are set to $1/\text{COND}$.
- MODE = 2 — $D(n)$ is set to $1/\text{COND}$ and the other $D(i)$ are set to 1.
- MODE = 3 — $D(i)$ is set to $\text{COND}^{-(i-1)/(n-1)}$, i.e., a geometric sequence ranging from 1 to $1/\text{COND}$.
- MODE = 4 — $D(i)$ is set to $1 - (i-1)/(n-1) \cdot (1 - 1/\text{COND})$, i.e., an arithmetic sequence ranging from 1 to $1/\text{COND}$.
- MODE = 5 — Each $D(i)$ is an independent random number in the range from $1/\text{COND}$ to 1 with a uniformly distributed logarithm. This guarantees that the ratio $\max_i D(i) / \min_i D(i)$ is close to COND .
- MODE = 6 — Each $D(i)$ is an independent random number with distribution DIST .

If $\text{MODE} < 0$, the meaning is the same as for $|\text{MODE}|$, except that the order of the entries of $D(i)$ is reversed. Thus, if $1 \leq \text{MODE} \leq 4$, the $D(i)$ range from 1 down to $1/\text{COND}$, and if $-4 \leq \text{MODE} \leq -1$, the $D(i)$ range from $1/\text{COND}$ up to 1.

If DMAX is specified as well, each D is scaled by $\text{DMAX} / \max_i |D(i)|$, so that the maximum absolute entry is $|\text{DMAX}|$. Note that DMAX may be negative (or complex).

If RSIGN is specified, it has the following meaning:

$\text{RSIGN} = \text{'F'}$ — D is unchanged.

$\text{RSIGN} = \text{'T'}$ — If D is real, each $D(i)$ is multiplied by $+1$ or -1 with a .5 probability. If D is complex, each $D(i)$ is multiplied by an independent random number r uniformly distributed on the unit circle (thus r has unit absolute value).

5.6 Bandwidths — KL, KU

KL and KU are read-only integers, specifying the lower and upper bandwidths, respectively. Thus, if $\text{KL} = 0$ (or $\text{KU} = 0$), the matrix is upper (or lower) triangular. If $\text{KL} = 1$ (or $\text{KU} = 1$), the matrix is upper (or lower) Hessenberg. If $\text{KL} \geq M - 1$ (or $\text{KU} \geq N - 1$), the matrix is full below (or above) the diagonal.

5.7 Norm — ANORM

This read-only REAL or DOUBLE PRECISION variable specifies the maximum absolute entry of the output A matrix.

5.8 Packing Option — PACK

This read-only CHARACTER*1 variable specifies the storage scheme of the output A matrix. All storage schemes in LAPACK (which includes all storage schemes in LINPACK [1] and the BLAS [4,5]) are accounted for. In this section we list the options and describe these storage schemes. PACK is interpreted as follows:

PACK = 'N' — No packing (i.e., dense storage of all entries).

PACK = 'U' — Zero out all the subdiagonal entries, but store A in dense format. This is allowed only if A is symmetric or Hermitian.

PACK = 'L' — Zero out all the superdiagonal entries, but store A in dense format. This is allowed only if A is symmetric or Hermitian.

PACK = 'C' — Store the upper triangle columnwise. This is allowed only if A is symmetric, Hermitian or upper triangular. This is a packed format, requiring about half the space of dense storage.

PACK = 'R' — Store the lower triangle columnwise. This is allowed only if A is symmetric, Hermitian or lower triangular. This is a packed format, requiring about half the space of dense storage.

PACK = 'B' — Store the lower triangle in band storage format. This is allowed only if A is symmetric, Hermitian or lower triangular. To illustrate this storage scheme, suppose the original matrix is

$$\begin{bmatrix} 11 & 12 & 13 & 0 & 0 & 0 \\ 21 & 22 & 23 & 24 & 0 & 0 \\ 31 & 32 & 33 & 34 & 35 & 0 \\ 0 & 42 & 43 & 44 & 45 & 46 \\ 0 & 0 & 53 & 54 & 55 & 56 \\ 0 & 0 & 0 & 64 & 65 & 66 \end{bmatrix} \quad (1)$$

where we have labeled each nonzero entry by its concatenated indices. In lower band format this matrix would be stored as

$$\begin{bmatrix} 11 & 22 & 33 & 44 & 55 & 66 \\ 21 & 32 & 43 & 54 & 65 & 0 \\ 31 & 42 & 53 & 64 & 0 & 0 \end{bmatrix} \quad (2)$$

Thus, the columns of the band storage format correspond to columns of the original matrix, and rows of the band storage format correspond to diagonals of the original matrix. This is true for the PACK = 'Q' and PACK = 'Z' options below as well.

PACK = 'Q' — Store the upper triangle in band storage format. This is allowed only if A is symmetric, Hermitian or upper triangular. To illustrate, the matrix in (1) would be stored as

$$\begin{bmatrix} 0 & 0 & 13 & 24 & 35 & 46 \\ 0 & 12 & 23 & 34 & 45 & 56 \\ 11 & 22 & 33 & 44 & 55 & 66 \end{bmatrix} \quad (3)$$

PACK = 'Z' — Store the matrix in band storage format. To illustrate, the matrix in (1) would be stored as

$$\begin{bmatrix} 0 & 0 & 13 & 24 & 35 & 46 \\ 0 & 12 & 23 & 34 & 45 & 56 \\ 11 & 22 & 33 & 44 & 55 & 66 \\ 21 & 32 & 43 & 54 & 65 & 0 \\ 31 & 42 & 53 & 64 & 0 & 0 \end{bmatrix} \quad (4)$$

5.9 Error Flag — INFO

Following the convention used in LAPACK, the variable INFO will be set to zero to indicate successful completion. A negative value of INFO on return means that argument number $-$ INFO in the calling sequence is incorrect; in this case the comments in the front of the program explain the error in more detail. If INFO is positive on return, then a numerical or other error was encountered during execution; again the comments at the front of the routines describe the situation in more detail.

6 xLATMR — Detailed Calling Sequence

Here is the specification of the calling sequence of xLATMR:

```
SUBROUTINE xLATMR( M, N, DIST, ISEED, SYM, D, MODE, COND, DMAX,  
$                RSIGN, GRADE, DL, MODEL, CONDL, DR, MODER,  
$                CONDR, PIVTNG, IPIVOT, KL, KU, SPARSE, ANORM,  
$                PACK, A, LDA, IWORK, INFO )  
*  
*    .. Scalar arguments ..  
*  
    CHARACTER*1    DIST, SYM, RSIGN, GRADE, PIVTNG, PACK  
    INTEGER        MODE, MODEL, MODER, KL, KU, M, N, LDA, INFO  
    type1          COND, CONDL, CONDR, SPARSE, ANORM, DMAX  
    type2          DMAX  
*  
*    .. Array arguments ..  
*  
    INTEGER        ISEED( 4 ), IPIVOT( * ), IWORK( * )  
    type2          D( * ), DL( * ), DR( * ), A( LDA, * )
```

Here, type1 and type2 are defined as follows:

Type of output matrix A	x in xLATMR	type1	type2
REAL	S	REAL	REAL
DOUBLE PRECISION	D	DOUBLE PRECISION	DOUBLE PRECISION
COMPLEX	C	REAL	COMPLEX
DOUBLE COMPLEX	Z	DOUBLE PRECISION	DOUBLE COMPLEX

Many of the arguments were described in the last section, so we only describe the new ones and variations here.

6.1 Symmetry — SYM

SYM was described in the last section, but we describe it here again because one of its possible values does not apply to xLATMR: 'P' (for positive semidefinite). The only allowable options are 'N' for nonsymmetric, 'S' for symmetric (real or complex), and 'H' for Hermitian (same as 'S' for real matrices).

6.2 Grading — GRADE, DL, MODEL, CONDL, DR, MODER, CONDR

GRADE is a read-only CHARACTER*1 variable which specifies if the matrix A is to be pre- or postmultiplied by diagonal matrices with diagonal entries DL and DR, respectively.

GRADE = 'N' — No grading.

GRADE = 'L' — A is premultiplied by $\text{diag}(\text{DL})$. This is allowed only if A is nonsymmetric.

GRADE = 'R' — A is postmultiplied by $\text{diag}(\text{DR})$. This is allowed only if A is nonsymmetric.

GRADE = 'B' — A is premultiplied by $\text{diag}(\text{DL})$ and postmultiplied by $\text{diag}(\text{DR})$. This is allowed only if A is nonsymmetric.

GRADE = 'S' — A is pre- and postmultiplied by $\text{diag}(\text{DL})$. This is allowed only if A is symmetric or nonsymmetric, but not complex Hermitian.

GRADE = 'H' — A is premultiplied by $\text{diag}(\text{DL})$ and postmultiplied by the complex conjugate of $\text{diag}(\text{DL})$. This is allowed only if A is nonsymmetric or Hermitian, but not complex symmetric.

GRADE = 'E' — A is premultiplied by $\text{diag}(\text{DL})$ and postmultiplied by the inverse of $\text{diag}(\text{DL})$. This is allowed only if A is nonsymmetric. This preserves the original eigenvalues.

DL is specified by MODEL and CONDL just as D is specified by MODE and COND. DR is specified by MODER and CONDR the same way.

6.3 Pivoting — PIVTNG and IPIVOT

PIVTNG is a read-only CHARACTER*1 variable which specifies how the rows and columns of A are to be permuted. IPIVOT is a read-only integer array which specified the order itself. PIVTNG is interpreted as follows:

PIVTNG = 'N' — No pivoting.

PIVTNG = 'L' — Left or row pivoting. A must be nonsymmetric.

PIVTNG = 'R' — Right or column pivoting. A must be nonsymmetric.

PIVTNG = 'B' or 'F' — Both or Full pivoting, i.e., on both sides. A must be square.

The IPIVOT array specifies the permutation used. After the basic matrix is generated, the rows, columns, or both are permuted. If, say, row pivoting is selected, xLATMR starts with the *last* row and interchanges the M-th and IPIVOT(M)-th rows, then moves to the next to last row, interchanging the (M-1)-st and the IPIVOT(M-1)-st rows, and so on. In terms of '2-cycles', the permutation is (1 IPIVOT(1)) (2 IPIVOT(2)) ... (M IPIVOT(M)) where the rightmost cycle is applied first. This is the *inverse* of the effect of pivoting in LINPACK or LAPACK. The idea is that factoring (with pivoting) an identity matrix which has been inverse-pivoted in this way should result in a pivot vector output from LAPACK or LINPACK identical to IPIVOT.

6.4 Sparsifying — SPARSE

This read-only variable must be between 0 and 1, inclusive. For each matrix entry an independent random variable r uniformly distributed on $(0,1)$ is generated and compared to SPARSE. If $r > \text{SPARSE}$, the matrix entry is unchanged, but if $r \leq \text{SPARSE}$, the matrix entry is set to zero (preserving symmetry if the matrix is symmetric or Hermitian). Thus, on average a fraction SPARSE of the entries will be set to zero. If SPARSE = 0, no entries will be set to zero.

6.5 Workspace — IWORK

This integer workspace array must have dimension as large as IPIVOT, if IPIVOT is referenced (PIVOTING $\neq$ 'N').

7 xLATMS — Detailed Calling Sequence

Here is the specification of the calling sequence for xLATMS:

```
SUBROUTINE xLATMS( M, N, DIST, ISEED, SYM, D, MODE, COND,
$                DMAX, KL, KU, PACK, A, LDA, WORK, INFO )
*
*   .. Scalar arguments ..
*
  CHARACTER*1      DIST, SYM, PACK
  INTEGER          MODE, KL, KU, M, N, LDA, INFO
  type1            COND, DMAX
*
*   .. Array arguments ..
*
  INTEGER          ISEED( 4 )
  type1            D( * ),
  type2            A( LDA, * ), WORK( * )
```

Here, type1 and type2 are defined just as for xLATMR. As before, we only describe the arguments which were not completely described in section 5.

7.1 Symmetry — SYM

The use of SYM depends on whether A is real or complex. If A is real then

SYM = 'N' means that A is nonsymmetric and the vector D, as determined by MODE, COND and DMAX, determines its singular values.

SYM = 'P' means that A is positive semidefinite and the vector D, as determined by MODE, COND and DMAX, determines its eigenvalues; in this case the eigenvalues will be non-negative (unless the user inputs negative eigenvalues in D with MODE = 0).

SYM = 'S' or 'H' means that A is symmetric and the vector D, as determined by **MODE**, **COND**, and **DMAX**, determines the eigenvalues. In addition, when **MODE** $\neq$ 0, each entry of D is multiplied at random by either +1 or -1. Thus A will have both positive and negative eigenvalues (unless the user inputs eigenvalues of all one sign in D with **MODE** = 0).

If A is complex then

SYM = 'N' means that A is nonsymmetric and the vector D, as determined by **MODE**, **COND** and **DMAX**, determines the singular values.

SYM = 'P' means that A is positive semidefinite and the vector D, as determined by **MODE**, **COND** and **DMAX**, determines its eigenvalues; in this case the eigenvalues will be non-negative (unless the user inputs negative eigenvalues in D with **MODE** = 0).

SYM = 'S' means that A is complex symmetric and the vector D, as determined by **MODE**, **COND** and **DMAX**, determines the singular values.

SYM = 'H' means that A is Hermitian the vector D, as determined by **MODE**, **COND**, and **DMAX** determines the eigenvalues. In addition, when **MODE** $\neq$ 0, each entry of D is multiplied at random by either +1 or -1. Thus A will have both positive and negative eigenvalues (unless the user inputs eigenvalues of all one sign in D with **MODE** = 0).

7.2 Workspace — WORK

The array **WORK**, which has the same type as A, must be dimensioned at least $3 \cdot \max(M, N)$. It is modified by the program.

8 xLATME — Detailed Calling Sequence

Here is the specification of the calling sequence for xLATME:

```

SUBROUTINE xLATME( N, DIST, ISEED, D, MODE, COND, DMAX, EI,
$                RSIGN, UPPER, SIM, DS, MODES, CONDS,
$                KL, KU, ANORM, A, LDA, WORK, INFO )
*
*   .. Scalar arguments ..
*
CHARACTER*1      DIST, RSIGN, UPPER, SIM, EI
INTEGER          MODE, MODES, KL, KU, N, LDA, INFO
type1            COND, CONDS, ANORM
type2            DMAX
*
*   .. Array arguments ..
*
INTEGER          ISEED( 4 )
type1            DS( * )
type2            D( * ), A( LDA, * ), WORK( * )

```

Here, type1 and type2 are determined just as for xLATMR and xLATMS. As before, we only discuss the arguments not completely described in section 5.

8.1 Eigenvalues — D, MODE, COND, DMAX, RSIGN, EI

These variables determine the eigenvalues of A. Their interpretations depend on whether A is real or complex.

If A is real, the eigenvalues are determined as follows. Recall that a real matrix either has real eigenvalues or complex eigenvalues appearing in complex conjugate pairs. D is first computed from MODE, COND, DMAX and RSIGN as described in section 5. In order to get complex conjugate eigenvalue, the following additional computations are done when MODE = 0 or MODE = 5.

When MODE = 0, both D and the array EI of CHARACTER*1 variables must be supplied by the user; EI specifies which entries of D are real eigenvalues, and which are the real and imaginary parts of a pair of complex conjugate eigenvalue. If EI(1) = ' ' (a blank), then all D(i) are taken to be real eigenvalues. Otherwise, all entries EI(i) must either be 'R' (for real part) or 'I' (for imaginary part). Each 'I' must follow an 'R'; thus if EI(i) = 'R' and EI(i + 1) = 'I', then D(i) is the real part and D(i + 1) is the imaginary part of a complex conjugate pair of eigenvalue. The assembly of A is begun by putting 1 by 1 blocks D(i) corresponding to real eigenvalues and 2 by 2 blocks

$$\begin{bmatrix} D(i) & D(i+1) \\ -D(i+1) & D(i) \end{bmatrix}, \quad (5)$$

which have complex conjugate eigenvalues $D(i) \pm i \cdot D(i + 1)$, on the diagonal of A. When MODE $\neq$ 0, EI is ignored.

When MODE = 5, complex conjugate pairs of eigenvalues are obtained as follows. For each adjacent pair (D(2i - 1), D(2i)) of entries of D, an independent random number r is chosen which has takes the values ± 1 with probability .5. If r = +1, the pair is treated as two real eigenvalues, and if r = -1, the pair is treated as the real and imaginary parts of a complex conjugate pair of eigenvalues, and incorporated into A as in (5).

When A is complex, the eigenvalues are determined as follows. D is first computed from MODE, COND, DMAX as described in section 5. Then, if RSIGN = 'T', and MODE is neither 0 (input by the user) nor ± 6 (random with distribution DIST), each D(i) is multiplied by an independent random complex number r uniformly distributed on the unit circle. These D(i) are the eigenvalues; EI is ignored.

8.2 Jordan Form — UPPER

If UPPER = 'T', the upper triangle of A (above the eigenvalues in the 1 by 1 and 2 by 2 diagonal blocks) is filled in with random numbers from distribution DIST. This means there will be exactly one Jordan block per distinct eigenvalue. If UPPER = 'F', the upper triangular is left zero. Thus A's Jordan form will only have 1 by 1 blocks, even if the eigenvalues are multiple.

8.3 Similarity Transform — SIM, DS, MODES, CONDS

If SIM = 'T', then A is premultiplied by S and postmultiplied by S^{-1} , where $S = U \cdot \text{diag}(\text{DS}) \cdot V$ is a random matrix. Here, U and V are random orthogonal (or unitary) matrices, and DS is constructed from MODES and CONDS just as D is constructed from MODE and COND (see section 3). If SIM = 'F', this pre- and postmultiplication is not performed. The condition number $\kappa \equiv \max_i |\text{DS}(i)| / \min_i |\text{DS}(i)|$ of S approximately determines the sensitivity of A's eigenproblem as described in section 4.

8.4 Bandwidth — KL, KU

As before, KL is the lower bandwidth and KU is the upper bandwidth. However, at most one of them may be less than $N - 1$, i.e., less than full bandwidth, and both must be positive. This is because no way is known to generate nontrivial random nonsymmetric matrices with fixed spectrum and arbitrary bandwidth. If a triangular matrix ($\text{KL} = 0$ or $\text{KU} = 0$) is desired, set SIM = 'F' so that no multiplication by S and S^{-1} is performed.

8.5 Workspace — WORK

The array WORK, which has the same type as A, must be dimensioned at least $3 \cdot N$. It is modified by the program.

9 Implementation Notes

9.1 Generating Random Orthogonal and Unitary Matrices

We use a method developed by Stewart [7], which we just outline here. The following pseudocode shows how to postmultiply a given matrix A by a random orthogonal (unitary) matrix U:

```
for i = 2 to n
  generate a vector v with i entries each of which is an independent normal (0, 1)
    random real (complex) number
  let x be an n-vector whose first n - i entries are zero and the remainder equal v
  Let H be an n by n Householder transformation which reduces the vector x to
    one whose only nonzero entry is at location n - i + 1
  apply the Householder transformation to A:  $A \leftarrow AH$ 
endfor
for i = 1 to n
  multiply the i-th column of A by an independent real (complex) random
    number which equals  $\pm 1$  with probability .5 (is uniformly distributed on the
    unit circle)
endfor
```

Thus, U is represented as a product of random Householder transformations and a

random diagonal orthogonal (unitary) matrix. The random orthogonal (unitary) matrix generated is distributed according to Haar measure [7], which is analogous to the uniform distribution. In other words, if U is a set of orthogonal (unitary) matrices with probability μ , then the sets $P \cdot U$ and $U \cdot P$, where P is any fixed orthogonal (unitary) matrix, also have probability μ .

The subroutine, xLAROR, also permits premultiplication $U \cdot A$, as well as $U \cdot A \cdot U^T$ and $U \cdot A \cdot U^*$ (conjugate transpose).

9.2 Accuracy Limitations

Any option which requires multiplication by orthogonal, unitary or other dense matrices within the test suite will introduce rounding errors which may obscure tiny eigenvalues or singular values. In particular, if ϵ is the machine precision, then the true eigenvalues or singular values of the A computed by xLATMS will differ from their prescribed values by as much as $O(\epsilon \|A\|)$ ($\|A\|$ is the largest absolute eigenvalue or singular value). In particular, tiny eigenvalues or singular values may have large absolute errors. This effect is even more pronounced in the case of xLATME, since eigenvalues of a nonsymmetric matrix may be extremely sensitive to perturbations, as discussed in section 4. In fact, if the matrix is chosen to have all zero eigenvalues, UPPER = 'T' so there is just one Jordan block, and SIM = 'T' so that there is pre- and postmultiplication by dense matrices, the true eigenvalues may be as large as $O(\epsilon^{1/n})$, where n is the dimension. This sensitivity is inherent in the problem [6].

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

- [1] James Bunch, Jack Dongarra, Cleve Moler, and G. W. Stewart. *LINPACK User's Guide*. SIAM, Philadelphia, PA, 1979.
- [2] James Demmel. On condition numbers and the distance to the nearest ill-posed problem. *Numerische Mathematik*, 51(3):251–289, July 1987.
- [3] James Demmel, Jack Dongarra, Jeremy Du Croz, Anne Greenbaum, Sven Hammarling, and Danny Sorensen. *Prospectus for the Development of a Linear Algebra Library for High-Performance Computers*. Mathematics and Computer Science Division Report ANL/MCS-TM-97, Argonne National Laboratory, Argonne, IL, September 1987.
- [4] Jack Dongarra, Jeremy Du Croz, Sven Hammarling, and Richard J. Hanson. An extended set of fortran basic linear algebra subroutines. *ACM Transactions on Mathematical Software*, 14(1):1–17, March 1988.
- [5] Jack Dongarra, Jeremy Du Croz, Iain Duff, and Sven Hammarling. A proposal for a set of level 3 basic linear algebra subprograms. *SIGNUM Newsletter*, 22(3):2–14, February 1987.
- [6] Gene Golub and Charles Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, 1983.
- [7] G. W. Stewart. On efficient generation of random orthogonal matrices with an application to condition estimation. *SIAM Journal of Numerical Analysis*, 17(3):403–409, 1980.

END

**DATE
FILMED**

10 / 1 / 92