


# HP MLIB SCILIB User's Guide

Second Edition



**Hewlett-Packard Company**  
Convex Division  
3000 Waterview Parkway  
P.O. Box 833851  
Richardson, TX 75083-3851  
United States of America



---

# HP MLIB SCILIB User's Guide

---

B5649-90002

Second Edition

January 1997

Hewlett-Packard Company  
Convex Division  
Richardson, Texas .  
United States of America

---

## HP MLIB SCILIB User's Guide

B5649-90002

© Copyright Hewlett-Packard Company 1997. All Rights Reserved. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

### Notice

The information contained in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance or use of this material.




This entire book is recyclable.

Printed in the United States of America

# Revision Information for HP MLIB SCILIB User's Guide

---

Edition	Document No.	Description
Second	B5649-90002	Released January 1997. Includes general updates.
First	720-005630-004	Initial release October 1994.



# Contents

<b>Preface</b> .....	<b>xv</b>
Purpose and Audience .....	xv
Organization .....	xv
Notational Conventions .....	xvi
Associated Documents .....	xvii
Ordering Documentation .....	xix
Technical Assistance .....	xix
<b>1 Introduction to SCILIB</b> .....	<b>1</b>
Overview .....	1
Chapter Objectives .....	2
What You Need to Know to Use SCILIB .....	2
Standardization .....	2
Accessing SCILIB .....	2
Interactions Between VECLIB, SCILIB, and LAPACK .....	4
Performance Value .....	5
Optimization .....	5
Parallel Processing .....	5
Profiling SCILIB Applications .....	7
Floating-Point Formats .....	8
Roundoff Effects .....	8
Required Data Item Byte Lengths and How to Get Them .....	8
Error Handling .....	9
HP MLIB Man Pages .....	9
Support Services .....	10
Supplemental Reading .....	11
<b>2 Basic Vector Operations</b> .....	<b>13</b>
Overview .....	13
Chapter Objectives .....	13

## Table of contents

What You Need to Know to Use These Subprograms .....	13
BLAS Storage Conventions .....	14
Supplemental Reading .....	17
CLUSEQ/.../CLUSILT – Find Clusters of Selected Vector Elements .....	18
GATHER – Gather Sparse Vector .....	21
IILZ – Count Initial Zero Elements .....	23
ILLZ – Count Initial Positive Elements .....	25
ILSUM – Count TRUE Vector Elements .....	27
INFLMAX – Index of Maximum Element of Vector .....	29
INFLMIN – Index of Minimum Element of Vector .....	31
ISAMAX/ICAMAX – Index of Maximum of Magnitudes .....	33
ISAMIN – Index of Minimum of Magnitudes .....	36
ISMAX/INTMAX – Index of Maximum Element of Vector .....	38
ISMIN/INTMIN – Index of Minimum Element of Vector .....	40
ISRCHEQ/ISRCHNE/.../ISRCHILT – Search Vector for Element .....	42
ISRCHMEQ/ISRCHMGE/.../ISRCHMNE – Search Vector for Element .....	45
OSRCHF/OSRCHI – Search Ordered Vector for Element .....	47
OSRCHM – Search Ordered Vector for Element .....	50
SASUM/SCASUM – Sum of Magnitudes .....	53
SAXPY/CAXPY – Elementary Vector Operation .....	55
SCATTER – Scatter Sparse Vector .....	58
SCOPY/CCOPY – Copy Vector .....	60
SDOT/CDOTC/CDOTU – Dot Product .....	63
SNRM2/SCNRM2 – Euclidean Norm .....	66
SPAXPY – Sparse Elementary Vector Operation .....	68
SPDOT – Sparse Dot Product .....	70
SROT/CROT – Apply Givens Rotation .....	72
SROTG/CROTG – Construct Givens Rotation .....	75
SROTM – Apply Modified Givens Rotation .....	77
SROTMG – Construct Modified Givens Rotation .....	80
SSCAL/CSCAL/CSSCAL – Scale Vector .....	83
SSUM/CSUM – Vector Sum .....	85
SSWAP/CSWAP – Swap Two Vectors .....	87
WHENEQ/WHENNE/.../WHENILT – Find Selected Vector Elements .....	90
WHENMEQ/WHENMGE/.../WHENMNE – Find Selected Vector Elements .....	94

<b>3 Basic Matrix Operations.....</b>	<b>97</b>
Overview .....	97
Chapter Objectives .....	97
What You Need to Know to Use These Subprograms .....	98
Subroutine Naming Convention .....	98
Supplemental Reading .....	99
MXM – Specialized Matrix-Matrix Multiply .....	101
MXMA – Generalized Matrix-Matrix Multiply .....	103
MXV – Specialized Matrix-Vector Multiply .....	108
MXVA – Generalized Matrix-Vector Multiply .....	110
SGBMV/CGBMV – Matrix-Vector Multiply .....	114
SGEMM/CGEMM – Matrix-Matrix Multiply .....	119
SGEMMS/CGEMMS – Strassen Matrix-Matrix Multiply .....	123
SGEMV/CGEMV – Matrix-Vector Multiply .....	128
SGER/CGERC/CGERU – Rank-1 Update .....	132
SMXPY – Matrix-Vector Multiply and Add .....	135
SSBMV/CHBMV – Matrix-Vector Multiply .....	137
SSPMV/CHPMV – Matrix-Vector Multiply .....	142
SSPR/CHPR – Rank-1 Update .....	146
SSPR2/CHPR2 – Rank-2 Update .....	150
SSYMM/CHEMM/CSYMM – Matrix-Matrix Multiply .....	154
SSYMV/CHEMV – Matrix-Vector Multiply .....	158
SSYR/CHER – Rank-1 Update .....	161
SSYR2/CHER2 – Rank-2 Update .....	164
SSYR2K/CHER2K/CSYR2K – Rank-2k Update .....	168
SSYRK/CHERK – Rank-k Update .....	172
STBMV/CTBMV – Matrix-Vector Multiply .....	176
STBSV/CTBSV – Solve Triangular Band System .....	181
STPMV/CTPMV – Matrix-Vector Multiply .....	187
STPSV/CTPSV – Solve Triangular System .....	191
STRMM/CTRMM – Triangular Matrix-Matrix Multiply .....	195
STRMV/CTRMV – Matrix-Vector Multiply .....	199
STRSM/CTRSM – Solve Triangular Systems .....	202
STRSV/CTRSV – Solve Triangular System .....	206
SXMPY – Vector-Matrix Multiply and Add .....	210

## Table of contents

XERBLA – Error Handler .....	212
<b>4 Linear Equations .....</b>	<b>215</b>
Overview .....	215
Chapter Objectives .....	216
What You Need to Know to Use These Subprograms .....	216
Subroutine Naming Convention .....	216
Condition Number .....	218
Determinant and Inverse .....	219
Supplemental Reading .....	219
MINV – Invert Matrix or Solve Linear Equations .....	220
OPFILT – Solve Symmetric Toeplitz Linear Equations .....	223
SGBCO/CGBCO – Estimate Condition .....	225
SGBDI/CGBDI – Determinant .....	229
SGBFA/CGBFA – Band LU Factorization .....	232
SGBSL/CGBSL – Solve Band Linear Equations .....	236
SGECO/CGECO – Estimate Condition .....	239
SGEDI/CGEDI – Determinant and Inverse .....	242
SGEFA/CGEFA – LU Factorization .....	246
SGESL/CGESL – Solve Linear Equations .....	249
SGTSL/CGTSL – Solve Tridiagonal Linear Equations .....	252
SPBCO/CPBCO – Estimate Condition .....	254
SPBDI/CPBDI – Determinant .....	258
SPBFA/CPBFA – Cholesky Factorization .....	261
SPBSL/CPBSL – Solve Linear Equations .....	265
SPOCO/CPOCO – Estimate Condition .....	267
SPODI/CPODI – Determinant and Inverse .....	270
SPOFA/CPOFA – Cholesky Factorization .....	274
SPOSL/CPOSL – Solve Linear Equations .....	277
SPTSL/CPTSL – Solve Positive Definite Tridiagonal Linear Equations .....	279
LINPACK Subprograms not in this Guide – .....	281
<b>5 Eigenvalues and Eigenvectors .....</b>	<b>283</b>
Overview .....	283
Chapter Objectives .....	283

What You Need to Know to Use These Subprograms .....	284
Supplemental Reading .....	284
RS – Eigenvalues and Eigenvectors of a Real Symmetric Matrix .....	285
TQL2 – Eigenvalues and Eigenvectors of a Real Symmetric Matrix .....	288
TQLRAT – Eigenvalues of a Real Symmetric Matrix .....	291
TRED1 – Reduce Real Symmetric Matrix to Tridiagonal Form .....	294
TRED2 – Reduce Real Symmetric Matrix to Tridiagonal Form .....	296
EISPACK Subprograms not in this Guide – .....	298
<b>6 Fast Fourier Transforms.....</b>	<b>301</b>
Overview .....	301
Chapter Objectives .....	301
What You Need to Know to Use These Subprograms .....	301
Supplemental Reading .....	302
CFFT2 – One-Dimensional Complex-to-Complex FFT .....	303
CFTFAX/CFFTMLT – Simultaneous One-Dimensional FFT .....	306
CRFFT2 – Complex-to-Real One-Dimensional FFT .....	311
RCFFT2 – Real-to-Complex One-Dimensional FFT .....	314
FFTFAX/RFFTMLT – Simultaneous One-Dimensional FFT .....	317
<b>7 Correlation and Convolution Subprograms.....</b>	<b>323</b>
Overview .....	323
Chapter Objectives .....	323
What You Need to Know to Use These Subprograms .....	323
Supplemental Reading .....	323
FILTERG – Discrete Correlation .....	324
FILTERS – Discrete Correlation .....	327
<b>8 Linear Recurrences .....</b>	<b>331</b>
Overview .....	331
Chapter Objectives .....	331
What You Need to Know to Use These Subprograms .....	331
FOLR/FOLRP – First Order Linear Recurrence .....	332

## Table of contents

FOLR2/FOLR2P – First Order Linear Recurrence .....	335
FOLRC – First Order Linear Recurrence .....	339
FOLRN/FOLRNP – Last Term of First Order Linear Recurrence .....	342
RECPP – Compute Vector of Partial Products .....	346
RECPS – Compute Vector of Partial Sums .....	349
SOLR – Second Order Linear Recurrence .....	352
SOLR3 – Second Order Linear Recurrence .....	355
SOLRN – Last Term of Second Order Linear Recurrence .....	358
<b>9 Miscellaneous Routines. ....</b>	<b>363</b>
Overview .....	363
Chapter Objectives .....	363
What You Need to Know to Use These Subprograms .....	363
XERSCI – SCILIB Error Handler .....	364

# Tables

Data Item Byte Length vs. Declaration and Compiler Option .....	9
Extended BLAS Naming Convention — Data Type .....	98
Extended BLAS Naming Convention — Matrix Form .....	98
Extended BLAS Naming Convention — Computation .....	98
Extended BLAS Naming Convention — Subprogram Names .....	99
LINPACK Naming Convention — Data Type .....	216
LINPACK Naming Convention — Form or Decomposition .....	217
LINPACK Naming Convention — Computation .....	217
LINPACK Naming Convention — Subprogram Names .....	218



## Preface

---

## Purpose and Audience

This guide describes the SCILIB software library and shows how to use it. SCILIB, a component of MLIB, is a collection of Fortran-callable subprograms identical in name and operation to those found in the Cray Research Incorporated's UNICOS Math and Scientific Library, V5.0. SCILIB subprograms have been optimized for use on the Hewlett-Packard Exemplar family of supercomputers.

The *HP MLIB SCILIB User's Guide* addresses experienced Fortran programmers who:

- convert programs written in Cray Fortran to Fortran.
- optimize existing software to improve performance and increase productivity on Exemplar supercomputers.
- use Fortran to develop new programs that rely heavily on matrix operations.

---

## Organization

To learn fundamental information necessary for using the SCILIB library, read Chapter 1 and the introductory sections of the other chapters. These sections of background information will help you efficiently use the SCILIB library subprograms.

To learn more about the subject of any given chapter, refer to the literature cited in the "Supplemental Reading" section of each chapter.

This guide is organized into the following chapters:

- Chapter 1 introduces general concepts about SCILIB.
- Chapter 2 describes basic vector operations included in SCILIB.

- Chapter 3 explains basic matrix operations.
- Chapter 4 describes linear equation subprograms in SCILIB.
- Chapter 5 explains the eigenanalysis capabilities available to SCILIB users.
- Chapter 6 describes the discrete Fourier transforms in SCILIB.
- Chapter 7 describes SCILIB subprograms that compute convolutions and correlations of data sets.
- Chapter 8 describes SCILIB subprograms that deal with linear recurrences.
- Chapter 9 describes miscellaneous subprograms to sort elements of a vector and to report errors detected in the usage of SCILIB routines.
- An index is included at the back of the manual.

---

## Notational Conventions

The following conventions are used in this manual:

- *Italics* within text indicate mathematical entities used or manipulated by the program: for example, solve the  $n$ -by- $n$  system of linear equations  $Ax = b$ .  
*Italics* within command lines indicate generic commands, file names, or subprogram names. Substitute actual commands, file names, or subprograms for the *italicized* words. For example, the command line  
`fc prog_name.o`  
instructs you to type the command *fc*, followed by the name of a program or subprogram object file.
- **UPPERCASE BOLDFACE** within text and in prototype Fortran statements indicates Fortran keywords and subprogram names that must be typed just as they appear: for example, **CALL SGESL**.
- Type in **lowercase boldface** indicates Fortran generic variable or array names. You should substitute actual variable or array names. The *italicized* mathematical entities and the **lowercase boldface** variable and array names usually correspond. For example,  $A$  will be a matrix and `a` will be the Fortran array containing the matrix:

**CALL SGESL (a, lda, n, ipvt, b, job)**

- UPPERCASE CONSTANT WIDTH represents Fortran programs.
- Brackets ( [ ] ) enclose optional entries.
- Many SCILIB subprogram names are prefixed to indicate the type of data they operate on.

For example, the subprograms to copy a vector are SCOPY and CCOPY, for REAL and COMPLEX vector types, respectively.

Note that in Cray Fortran, the single precision REAL type is 64 bits (one Cray word) in length. This is essentially equivalent to the type DOUBLE PRECISION in Fortran. Similarly, the Cray Fortran type COMPLEX is 128 bits long; this is equivalent to the Fortran type DOUBLE COMPLEX. The Cray double precision versions of REAL and COMPLEX types (128 and 256 bit, respectively) are not supported in SCILIB because of processing time considerations.

---

## Associated Documents

Using this guide successfully may require information not specific to the tasks described herein or not within the scope of this guide. The following documents are provided to help you:

- *HP MLIB VECLIB User's Guide* (B5649-90003). This guide provides definitions for many additional subprograms available to SCILIB users through inclusion of VECLIB, but not documented in the *HP MLIB SCILIB User's Guide*.
- *HP MLIB LAPACK User's Guide* (B5649-90001). This guide provides information on the subprograms provided with the LAPACK library.
- *Exemplar Programming Guide: S-Class and X-Class Servers* (B5600-90001). This manual describes efficient programming techniques for the Exemplar family of computers.
- *Exemplar C and Fortran 77 Programmer's Guide* (B5600-90002), *Release Notice, Fortran 77 Compiler V1.0*, and *Release Notice, C Compiler V1.0*. These documents describe the Exemplar Fortran 77 and C compilers.
- *HP C/HP-UX Reference Manual* (92453-90024). This manual presents reference information on the C programming language, as implemented by Hewlett-Packard.

Associated Documents

- *HP C/HP-UX Programmer's Guide* (92434-90002). This guide contains detailed discussions of selected C topics.
- *Fortran/9000 Programmer's Reference* (B3906-90002). This book is a language reference for Hewlett-Packard Fortran 77.
- *Fortran/9000 Programmer's Guide* (B3906-90001). This manual is a task reference. It describes features and requirements in terms of the tasks a programmer might perform. These tasks include how to compile, link, run, debug, and optimize programs.
- *Programming on HP-UX* (B2355-90652). This book describes how to develop software on HP-UX using the HP compilers, assemblers, linker, libraries, and object files.
- *Exemplar SPP1000 Series Architecture* (DHW-014). This book describes the SPP1200 and SPP1600 architectures.
- *Exemplar Architecture: S-Class and X-Class Servers* (A4716-90001). This book describes the architectures of the S2000 and X2000 servers.
- *CXpa Reference: Exemplar S-Class and X-Class Servers* (B5639-90002). This guide explains the operation of the CXpa Performance Analyzer and the steps needed to create and interpret a CXpa profile.
- *CXdb Quick Reference: Exemplar S-Class and X-Class Servers* (B5639-90001). This book describes prominent features of the CXdb visual debugger.
- *CXdb Online Help*. This document describes the CXdb visual debugger. Also, see the xcdb(1) man page.
- *HP MPI User's Guide: S-, X-, D-, and K-Class Servers* (B6011-90001). This book discusses message-passing programming using the Message-Passing Interface library.
- *HP PVM User's Guide: S-Class and X-Class Servers* (B5885-90001). This book discusses message-passing programming using the Parallel Virtual Machine library.
- *HP Fortran 90 Programmer's Reference* (B5876-90001). This book is a complete Fortran 90 language reference. It also covers compiler options, compiler directives, and library information.
- *HP Fortran 90 Programmer's Notes* (B5876-90002). This book provides extensive usage information, including how to compile and link, suggestions and tools for migrating to HP Fortran 90, and how to call C and HP-UX routines from HP Fortran 90.
- *Exemplar C++ Programming Guide: S-Class and X-Class Servers* (B5630-90001). This book describes the Exemplar C++ compiler.

- *HP-UX Assembly Language Reference Manual* (92432-90001). This manual describes the HP-UX Assembler for the PA-RISC processor.
- *HP PA-RISC 2.0 Architecture Reference* (B5655-90009). This manual describes the architecture of the Hewlett-Packard PA-RISC 2.0 processor.
- *HP PA-RISC 1.1 Architecture and Instruction Set Reference Manual* (DHP-181). This manual describes the architecture and the instruction set of the Hewlett-Packard PA-RISC 1.1 processor.
- *PA-RISC Procedure Calling Conventions Reference* (09740-90015). This manual describes the conventions for creating PA-RISC assembly language procedure calls.

---

## Ordering Documentation

To order additional copies of this document or other documents listed in 'Associated Documents,' send requests to

Hewlett-Packard Company  
Convex Division  
Customer Service  
P.O. Box 833851  
Richardson, TX 75083-3851 USA

Please include the order number (xxxxx-9xxxx number) or the exact title of the document.

---

## Technical Assistance

If you have questions that are not answered in this book, contact the Hewlett-Packard Convex Technical Assistance Center (TAC) at the following locations:

Preface

**Technical Assistance**

- Within the continental U.S., call 1 (800) 952-0379.
- From Canada, call 1 (800) 345-2384.
- All other locations, contact your local Hewlett-Packard office.

You can also use the contact utility, if you would like to report any problems you may have with HP MLIB LAPACK or its associated documentation.

# 1 Introduction to SCILIB

---

## Overview

SCILIB is a collection of Fortran-callable mathematical subprograms which provides a look-alike implementation of the Scientific Library portion of Cray Research Incorporated's UNICOS Math and Scientific Library, V5.0.

Many SCILIB subroutines are optimized for use on the Hewlett-Packard computer systems. This general library addresses a variety of linear algebra operations on multiple data types. It contains subprograms for:

- dense vector operations
- sparse vector operations
- matrix operations
- linear equation solution
- discrete Fourier transforms
- convolution and correlation
- linear recurrences
- error reporting

Although SCILIB was designed for use with Fortran programs, C programs can call SCILIB subprograms, as described in appendix A of the *HP MLIB VECLIB User's Guide*, which is included in this documentation set.

This chapter provides information necessary for efficient use of SCILIB, including discussions of SCILIB software standardization, how to access SCILIB subprograms, optimizations, including parallel processing and interactions with other HP analysis and optimization products, supported floating-point formats, roundoff effects, SCILIB subprogram capabilities, how to use the library and various compiler options, error handling, online documentation, and HP support services.

## Chapter Objectives

After reading this chapter you will:

- know how to access SCILIB
- understand how SCILIB works in a parallel computing environment
- know how SCILIB interacts with the CXpa Performance Analyzer and other profilers, and architecture-specific features.
- understand roundoff effects
- be able to use Fortran type declarations and compiler options
- understand how SCILIB handles errors
- know how to access the online *HP MLib Man Pages*
- know what to do if you are having trouble using SCILIB subprograms

---

## What You Need to Know to Use SCILIB

You should be familiar with the following sections to make efficient use of SCILIB.

### Standardization

SCILIB is designed to provide HP users with a look-alike implementation of UNICOS Math and Scientific Library subroutines. This allows programs written for Cray machines to be easily ported to HP, and it provides a common programmer interface between the two machines, which can ease a Cray programmer's transition to writing code for HP computer systems, especially when used in conjunction with the Fortran 90 compiler's `+autodbl` or `+autodbl4` command line flag.

### Accessing SCILIB

The SCILIB library consists of compiled subprograms ready for you to incorporate into your programs with the linker. Simply include the appropriate declarations and CALL statements in your Fortran source program and specify that SCILIB be used as an object library at link time.

MLIB libraries are installed in the `/opt/mlib/` directory. The entire path depends on your system type, as follows:

System type	CPU Type	Installation Directory
C-, D-, or K-Class	PA-8000	<code>/opt/mlib/lib/pa2.0</code>
S- or X-Class	PA-8000	<code>/opt/mlib/lib/pa2.0parallel</code>
C-, D-, J-, or K-Class	PA-7200	<code>/opt/mlib/lib/pa1.1</code>
SPP-1200 or SPP-1600	PA-7200	<code>/opt/mlib/lib/pa1.1parallel</code>

The file names of the SCILIB library is *libscilib.a*.

There are several ways to specify the linking of your program with SCILIB:

- specify the entire path of the library file on the `f90` command line that links your program. For example, on a C-Class PA-8000 system, use:

```
f90 [options] file /opt/mlib/lib/pa2.0/libscilib.a
```

- use the `-l` option on the `f90` command line that links your program, preceded by the option `-Wl,-L<path>`, where `<path>` is the appropriate one of the above installation directories. For example, the above `f90` command line could be written as:

```
f90 [options] file -Wl, -L/opt/mlib/lib/pa2.0 -lscilib
```

- set the `LDOPTS` environment variable to include `-L<path>`, where `<path>` is the appropriate one of the above installation directories, and use the `-l` option on the `f90` command line that links your program. For example, use:

```
f90 [options] file -lscilib
```

---

**NOTE**

If you are running MLIB on an S- or X-Class machine, the SCILIB, VECLIB, and LAPACK libraries contain parallelized subprograms. See “Parallel Processing” for additional information about linking your program.

VECLIB is documented in the *HP MLIB VECLIB User's Guide*. If your program uses subprograms from both SCILIB and VECLIB, specify both `-lscilib` and `-lveclib8`. For example, using method 2 above to specify the location of the libraries, link with

```
f90 [options] file -lscilib -lveclib8
```

See “Interactions Between VECLIB, SCILIB, and LAPACK” for details about how to order the two `-l` options. Do not try to use subprograms from both `-lscilib` and `-lveclib` in the same program.

LAPACK is documented in the *HP MLIB LAPACK User's Guide*. If you use subprograms from both SCILIB and LAPACK, specify both `-lscilib` and `-llapack8` when you link your program. For example, using method 2 above on an S-Class machine, link with

```
f90 [options] file -Wl, -L/opt/mlib/lib/pa2.0parallel -lscilib -llapack8
```

Add the linker option `-lveclib8` if VECLIB subprograms are also used. See "Interactions Between VECLIB, SCILIB, and LAPACK" for details about the order of the `-l` options. Do not try to use subprograms from both `-lscilib` and `-llapack` in the same program.

## Interactions Between VECLIB, SCILIB, and LAPACK

Each of the five library files in VECLIB, SCILIB, and LAPACK is complete in itself, meaning that you will not need to load one library merely because you have used subprograms from another. This is accomplished by including various subprograms in more than one library. For example, subroutine SGEMV is in all of these products, but with identical functionality. Thus, in general, you have to load only the libraries you need, and you may list them in any order on your load command line, as described in the previous section. However, there are a few differences between the libraries that may force you to put the libraries into a specific order to obtain the results you expect.

## Differences between VECLIB8 and SCILIB

Five subprograms common to VECLIB8 and SCILIB differ slightly in functionality.

Subprograms ICAMAX, ISAMAX, ISAMIN, ISMAX, and ISMIN in VECLIB8 handle a negative `incx` argument by taking its absolute value and searching the `x` vector in forward order, while in SCILIB, a negative `incx` argument results in searching the array `x` in backward order. No VECLIB8 subprograms call any of these subprograms with a negative `incx` argument, so you may safely load SCILIB before VECLIB8 if you need the SCILIB functionality.

Two other subprograms in both VECLIB8 and SCILIB have the same functionality but different numbers of arguments.

Subroutines SGEMMS and CGEMMS from the two libraries implement Strassen's method for matrix multiplication, but the SCILIB versions have an extra argument, for working storage, that is not needed in the VECLIB8 versions. Be certain that your calls to these subprograms have 14 arguments if you load SCILIB before VECLIB8.

## Differences between LAPACK8 and SCILIB

Two subprograms common to LAPACK8 and SCILIB differ slightly in functionality.

Subprograms ICAMAX and ISAMAX in LAPACK8 handle a negative **incx** argument by taking its absolute value and searching the **x** vector in forward order, while in SCILIB, a negative **incx** argument results in searching the array **x** in backward order. No LAPACK8 subprograms call either of these subprograms with a negative **incx** argument, so you may safely load SCILIB before LAPACK8 if you need the SCILIB functionality.

## Performance Value

As computer architectures have become more complicated, it has become more important to know the architecture of the target computer to maximize program performance. When a program is moved from one computer to another, architectural considerations on which the program was based may no longer be valid. If, however, the computationally intensive part of the program is based on highly tuned subprograms from a vendor-supplied library, the vendor's knowledge of the architecture is transferred to the program. SCILIB provides this feature, enabling you to achieve good performance at low cost.

## Optimization

Keep in mind that while SCILIB subroutines are identical in name and purpose to subroutines found in the UNICOS Math and Scientific Library, many have been optimized for use on Hewlett-Packard machines, and this required somewhat different implementations. Since the two machines use different architectures and different methods of carrying out various mathematical operations in hardware and software, SCILIB subroutines cannot be expected to give identical answers to their Cray counterparts in all cases. However, SCILIB subroutines have been tested to insure that they give essentially equivalent answers in all circumstances.

## Parallel Processing

This section applies only to Hewlett-Packard S- and X-Class computer systems.

Parallel processing is available on Hewlett-Packard S- and X-Class computer systems. These systems can divide a single computational process into small streams of execution, called *threads*. The result is that you can have more than one processor executing on behalf of the same process.

You can permit or disable parallel processing at link time or at run time. To permit parallel processing at link time, your link step must produce a multithreaded executable. The S-Class f77, C and C++ compilers always pass **+tm S2000** to the linker, so you always get a multithreaded executable from the Fortran 77 and C compilers. Similarly, the X-Class f77, C, and C++ compilers always pass **+tm X2000**. Fortran 90, on the other hand, does not

pass the **+tm** string to the linker so you must include the **-Wl,+tm,S2000** or **-Wl,+tm,X2000** option on the f90 command line that links your program.

A multithreaded executable will go parallel when the **+parallel** flag is passed to the linker. The executable will go parallel without further ado if compiled with **+O3 +Oparallel**, or simply linked with **+Oparallel**:

**f77** [options including **+O3 +Oparallel** ] *file* **-lscilib**

To disable MLIB's automatic parallelism at link time, you omit the **+Oparallel** option and include the following required runtime libraries:

**f77** [options] *file* **-lscilib -lpthread -lcps -lpthread -lail**

The f90 command lines corresponding to these f77 command lines are:

**f90** [options] including **-Wl,+tm,S2000,+parallel**] *file* **-lscilib -lpthread -lcps -lpthread -lail**

and

**f90** [options including **-Wl,+tm,S2000**] *file* **-lscilib -lpthread -lcps -lpthread -lail**

At run time, you can use the *mpa*(1) utility or the **MP\_NUMBER\_OF\_THREADS** environment variable to control parallelism. Refer to the *mpa*(1) or *f77*(1) man page, respectively, for details.

If parallelism is permitted, parallelized LAPACK subprograms determine at run time whether multiple processors are available. If so, it detects whether the program is already using multiple threads. It uses this information to automatically choose between a single- or parallel-processor algorithm.

If you are using an S- or X-Class system, you can realize the performance benefits of parallel processing in three ways. First, you can call any parallelized LAPACK subprogram and let it use parallelism internally if it determines that it is appropriate to do so, based on such factors as problem size, system configuration, and user environment. Alternatively, you can call LAPACK subprograms in a parallelized loop or region. To use this mechanism, you must be familiar with the techniques of parallel processing; refer to the *Exemplar Programming Guide*. The third mechanism is to use the MPI explicit parallel model. See the *MPI*(1) man page for details.

LAPACK subprograms are reentrant. This means that they may be called several times in parallel to do independent computations without one call interfering with another. You can use this feature to call LAPACK subprograms in a parallelized loop or region. The compiler does not automatically parallelize loops containing a function reference or subroutine call. You can force it to parallelize such a loop by inserting compiler directives before the loop.

For example, the following Fortran code makes parallel calls to subprogram SGETRS:

```
C$DIR LOOP_PRIVATE (J)
C$DIR LOOP_PARALLEL
  DO 10 J=1, N
    CALL SGETRS ('N', N, 1, A, LDA, IPIV, B(1,J), LDB, INFO(J))
  10 CONTINUE
```

While optimizing a parallel program, you might want to make parallel calls to a SCILIB subprogram to execute independent operations where the call statements are not in a loop. The Fortran compiler does not automatically parallelize code outside a loop, but you can use the `BEGIN_TASKS`, `NEXT_TASK`, and `END_TASKS` compiler directives to tell the compiler to parallelize such code.

If a parallelized SCILIB subprogram is called from a parallelized loop or region, the internal parallelism is disabled.

## Profiling SCILIB Applications

The CXpa Performance Analyzer is an interactive tool for HP computer systems that gathers and analyzes program execution timing (profiling) data. CXpa provides the programmer with the means to study the timing behavior of a program for the purposes of optimizing, benchmarking, and debugging. To use the performance analyzer, you must first compile your Fortran program with the `+pa` compiler option. These options instrument the compiled program so that its performance can be measured at the subprogram level, the loop level, the block level, or the region level.

SCILIB has been instrumented at the subprogram level so that the performance of SCILIB subprograms can be included in the analysis. This instrumentation is nonintrusive, so it is not necessary to use a different version of SCILIB when you desire to profile your program. Also, the CXpa instrumentation does not interfere with the *prof* or *gprof* instrumentation in your program. However, you may not profile your program with both CXpa and *prof* or *gprof* at the same time.

Subprogram-level profiling produces summary information about the subprograms that are called during profiled execution of the program. This information includes:

- the number of times each subprogram is called
- the CPU time in each subprogram and the percentage of the program total, either including or excluding the cumulative time in called subprograms
- a dynamic call graph, listing the subprogram calls that take place within a computer program

Introduction to SCILIB  
What You Need to Know to Use SCILIB

CXpa is an optional product. For more information about CXpa, refer to the *CXpa Reference: Exemplar S-Class and X-Class Servers* (or contact your Hewlett-Packard sales representative).

## Floating-Point Formats

Hewlett-Packard Exemplar systems and other computers with PA-RISC-based architectures operate only on floating-point data in the IEEE format. For further information on Hewlett-Packard floating-point formats, refer to the *Fortran/9000 Programmer's Guide*.

## Roundoff Effects

SCILIB subprograms may use a different arithmetic order of evaluation than that employed by the UNICOS Math and Scientific Library or other mathematical software. Different roundoff characteristics may result. Accuracy of results is usually about the same, so using SCILIB should not materially affect the accumulation of roundoff errors in a complete application program. If it does, you should examine the mathematical analysis of the problem, which will likely show that the problem is ill-conditioned. Ill-conditioned means that the small roundoff errors that are inadvertently introduced into any computation are magnified out of proportion to the desired result. Similarly, if results with and without SCILIB differ materially, both sets of answers are probably inaccurate and you should investigate further. If the program correctly applies stable computational algorithms, the problem itself is probably ill-posed.

## Required Data Item Byte Lengths and How to Get Them

In SCILIB subprograms all INTEGER, REAL, and LOGICAL arguments must be 64-bit quantities, and all COMPLEX arguments must occupy 128 bits. Table 1-1 shows the correspondence between the lengths of data items declared in various ways.

Note that if the Fortran data types are not given length specifiers (for example, REAL is used instead of REAL\*8) then either of the compiler options **+autodbl** or **+autodbl4** are compatible with the data types required by SCILIB. If any DOUBLE PRECISION declarations or double precision constants occur in the program, the **+autodbl4** compiler option causes them to be treated as 64-bit quantities. This leads us to recommend that you use either **+autodbl4**.

SCILIB subroutines will not accept INTEGERS or REALS of length 4 bytes, which is the default for these types in Fortran. If you call SCILIB subroutines and *do not* compile with **+autodbl** or **+autodbl4**, you must be very careful that all variables and constants passed to SCILIB subroutines are of the proper length.

For more information on Fortran data types and Fortran compiler options refer to a Fortran language reference.

**Table 1-1 Data Item Byte Length vs. Declaration and Compiler Option**

Fortran Declaration	Fortran Compiler Option		
	none	+autodbl	+autodbl4
<b>INTEGER</b>	4	8	8
<b>INTEGER*4</b>	4	4	4
<b>INTEGER*8</b>	8	8	8
<b>Integer by default</b>	4	8	8
<b>REAL</b>	4	8	8
<b>REAL*4</b>	4	4	4
<b>REAL*8</b>	8	8	8
<b>Real by default</b>	4	8	8
<b>DOUBLE PRECISION</b>	8	16	8
<b>Double Precision constant</b>	8	16	8
<b>COMPLEX</b>	8	16	16
<b>COMPLEX*8</b>	8	8	8
<b>COMPLEX*16</b>	16	16	16
<b>Complex constant</b>	8	16	16
<b>LOGICAL</b>	4	8	8
<b>LOGICAL*4</b>	4	4	4
<b>LOGICAL*8</b>	8	8	8
<b>Logical constant</b>	4	8	8

## Error Handling

Some SCILIB subprograms do not have a success/error code in their argument lists, but instead call another SCILIB subprogram to process the error condition. Two error handlers are provided: XERBLA and XERSCI; these are documented in Chapter 3 and Chapter 9 of this guide, respectively. The documentation for each SCILIB subprogram indicates if either of these error handlers is used. The standard versions of XERBLA and XERSCI write an error message onto the standard error file. If the main program is in Fortran, a call traceback is also written onto the standard error file. Execution is then terminated with a nonzero exit status. You may supply a version of XERBLA or XERSCI that alters this action; see the documentation for these subprograms for more information.

## HP MLIB Man Pages

The *HP MLIB Man Pages* contain online documentation that includes information from the user's guides. This reference contains an introduction to

Introduction to SCILIB  
What You Need to Know to Use SCILIB

SCILIB and to each set of subprograms in SCILIB, and reference entries for each subprogram. Subprogram entries include descriptions and examples of usage.

This reference is provided for users to easily and efficiently obtain online information on SCILIB. Because of the limited number of fonts supported and the difficulty of presenting mathematical equations in the *man*(1) system, the *HP MLIB Man Pages* is not a substitute for the user's guides; the most detailed information on SCILIB will be in the user's guide.

The *HP MLIB Man Pages* are installed in the directory */opt/mlib/share/man*. You must have this directory in your MANPATH environment variable to access man pages for VECLIB, SCILIB, or LAPACK. For example, you will add */opt/mlib/share/man* to your MANPATH environment variable. Refer to the *mllib*(3m) master man page for details about the MANPATH variable and the other HP MLIB man pages.

To access the *HP MLIB Man Pages*, use the shell command

```
man mlib
```

For further explanation and a table of contents of reference entries, refer to the *scilib*(3m) entry by typing

```
man scilib.
```

## Support Services

Hewlett-Packard maintains a staff to provide technical help if you have difficulty. Located in the Hewlett-Packard Technical Assistance Center (TAC), these people are the primary link between you and the company, and they stand ready to assist you with any difficulties. Note, though, that SCILIB has been tested extensively and is very reliable. Therefore, before contacting the TAC about a SCILIB problem, follow this procedure to isolate the cause of the trouble and to simplify the job of resolving it:

- Check any error response provided by the subprogram in question. The subprogram descriptions in this manual describe how to check an error response. If the answer is wrong because an error has been detected, correct the cause of the error and run the job again.
- Verify that the subprogram usage in the program matches the subprogram specifications in this manual. Pay special attention to the number of arguments in the **CALL** statement and to the declarations of arrays and integer constants or variables that describe them. If everything is in order, write out all the arguments immediately before and after the **CALL** statement.
- Make sure there really is a problem. For example, if an apparently incorrect answer is being computed, check to see if the answer does satisfy

the problem as defined in the program. Also, for problems with more than one answer, SCILIB may produce a different answer or give the answers in a different order than expected. If the problem is ill-conditioned, SCILIB may not be able to compute a reliable answer at all. Again, error messages often suggest the cause of the problem.

- Isolate the problem. If possible, write a small test program that encounters the same difficulty. Perhaps data causing the problem may be written out from the original program and read into the small one. Try to remove the problem area from a large program and concentrate it in a small program. In this way, you eliminate extraneous code from suspicion. If the problem area is large, try to pare it to a manageable size. For example, if a 50-by-50 linear system fails, try to produce a 2-by-2 system that fails in the same way. Clearly, this is not always possible, but the process often leads to insight.

You will frequently discover a usage error and resolve the problem by following the steps above. If the trouble persists, contact the TAC for help. Providing a small test program and expected answers will help the TAC further analyze the problem. To report a software or documentation problem to the TAC, use the *contact* utility. The *contact* utility allows you to submit a problem or suggest an enhancement directly to the TAC from your own system.

For information about *contact*, use the shell command

```
man contact
```

---

## Supplemental Reading

The SCILIB documentation set includes the *HP MLIB VECLIB User's Guide*, the *HP MLIB SCILIB User's Guide*, and the *HP MLIB LAPACK User's Guide*. The following additional documents provide supplemental help.

Dongarra, J.J., J.R. Bunch, C.B. Moler, and G.W. Stewart. *LINPACK Users' Guide*. Philadelphia, PA: SIAM Publications. 1979.

Garbow, B.S., *et al* "Matrix Eigensystem Routines—EISPACK Guide Extension." *Lecture Notes in Computer Science*, Vol. 51. New York: Springer-Verlag. 1977.

Smith, B.T., *et al* "Matrix Eigensystem Routines—EISPACK Guide." *Lecture Notes in Computer Science*, Vol. 6, 2nd edition. New York: Springer-Verlag. 1976.

Introduction to SCILIB  
**Supplemental Reading**

## 2 Basic Vector Operations

---

### Overview

This chapter explains how to use the SCILIB vector subprograms that serve as building blocks for many user programs. It describes subprograms for performing dense and sparse vector operations, and it includes the Fortran equivalent of each subprogram. This set of SCILIB subprograms includes:

- the Basic Linear Algebra Subprograms (BLAS)
- Cray extensions to the BLAS

The term BLAS, as used in this section, refers to the standard BLAS operations and the Cray extensions to the BLAS.

---

### Chapter Objectives

After reading this chapter you will:

- understand BLAS storage conventions
  - know how to specify array sections
  - know how to handle backward storage
  - know how to use increment (also called stride) arguments
- 

### What You Need to Know to Use These Subprograms

This section discusses commonly used or computationally expensive operations of linear algebra. Even though you can code most of these operations in fewer than 10 lines of Fortran, using SCILIB subprograms can

improve program performance, as well as program modularity and readability. Note, however, that in some situations you can achieve better computational performance by entering Fortran code than by calling one of these subprograms.

## BLAS Storage Conventions

The Basic Linear Algebra Subprograms (BLAS) were developed to enhance the portability of published linear algebra codes. In particular, LINPACK, the high-level public-domain linear equation package, uses the BLAS. Thus, if you use LINPACK from SCILIB you will normally be replacing the standard Fortran BLAS with the SCILIB BLAS and increasing the efficiency of LINPACK.

You need not limit your use of the SCILIB BLAS to LINPACK. Because these subprograms are portable, modular, self-documenting, and efficient, you can incorporate them into your programs. To realize the full power of the BLAS, you must understand the following three subjects:

- Fortran storage of arrays
- Fortran array argument association
- BLAS indexing conventions

## Fortran Storage of Arrays

Two-dimensional arrays in Fortran are stored by columns. Consider the following specifications:

```
DIMENSION A(N1, N2), B(N3)
EQUIVALENCE (A, B)
```

where  $N3 = N1 \times N2$ . Then  $A(I, J)$  is associated with the same memory location as  $B(K)$  where

$$K = I + (J-1) \times N1$$

Successive elements of a column of  $A$  are adjacent in memory, while successive elements of a row of  $A$  are stored with a difference of  $N1$  storage units between them. Remember that the size of a storage unit depends on the data type.

## Fortran Array Argument Association

When a Fortran subprogram is called with an array element as an argument, the value is not passed. Instead, the subprogram receives the address in memory of the element. Consider the following code segment:

```

REAL A(10,10)
J = 3
L = 10
CALL SUBR (A(1,J),L)
.
.
.
SUBROUTINE SUBR (X,N)
REAL X(N)
.
.
.

```

SUBR is given the address of the first element of the third column of A. Since it treats that argument as a one-dimensional array, successive elements  $X(1)$ ,  $X(2)$ , . . . ., occupy the same memory locations as the successive elements of the third column of A, that is,  $A(1,3)$ ,  $A(2,3)$ , . . . . Hence, the entire third column of A is available to the subprogram.

## BLAS Indexing Conventions

The rest of this section describes dealing with stride arguments and handling forward and backward storage.

A vector in the BLAS is defined by three quantities:

1. The vector length.
2. The array or starting element within an array.
3. The increment, sometimes called the *stride*, which defines the number of storage units between successive vector elements.

*Forward Storage.* Suppose that  $x$  is a real array. Let  $N$  be the vector length and let  $INCX$  be the increment. Suppose that a vector  $x$  with components  $x_i$ ,  $i = 1, 2, \dots, N$  is stored in  $x$ . If  $INCX \geq 0$ , then  $x_i$  is stored in  $x(1 + (i-1) \times INCX)$ . This is forward storage starting from  $x(1)$  with stride equal to  $INCX$ , ending with  $x(1 + (N-1) \times INCX)$ . Thus, if  $N = 4$  and  $INCX = 2$ , the vector components  $x_1, x_2, x_3$ , and  $x_4$  are stored in the array elements  $x(1)$ ,  $x(3)$ ,  $x(5)$ , and  $x(7)$ , respectively.

*Backward Storage.* Some BLAS routines permit the backward storage of vectors, which is specified by using a negative  $INCX$ . If  $INCX < 0$ , then  $x_i$  is stored in  $x(1 + (N-i) \times |INCX|)$  or equivalently in  $x(1 - (i-1) \times INCX)$ . This is backward storage starting from  $x(1 - (N-1) \times INCX)$  with stride equal to  $INCX$ , ending with  $x(1)$ . Thus, if  $N = 4$  and  $INCX = -2$ , the vector components  $x_1, x_2, x_3$ , and  $x_4$  are stored in the array elements  $x(7)$ ,  $x(5)$ ,  $x(3)$ , and  $x(1)$ , respectively.

## Basic Vector Operations

### What You Need to Know to Use These Subprograms

$INCX = 0$  is permitted by some BLAS routines and is not permitted by others. When it is allowed, it means that  $x$  is a vector of length  $N$ , whose components all equal the value of  $X(1)$ .

The notation  $(N, X, INCX)$  describes a BLAS vector. For example, if  $X$  is an array of dimension  $N$ , then  $(N, X, 1)$  represents forward storage and  $(N, X, -1)$  represents backward storage. If  $A$  is an  $M$ -by- $N$  array, then  $(M, A(1, J), 1)$  represents column  $J$  and  $(N, A(I, 1), M)$  represents row  $I$ . Finally, if an  $M$ -by- $N$  matrix is embedded in the upper left-hand corner of an array  $B$  of size  $LDB$  by  $NMAX$ , then column  $J$  is  $(M, B(1, J), 1)$  and row  $I$  is  $(N, B(I, 1), LDB)$ .

## Examples

The following examples illustrate how to use increment arguments to perform different operations with the same subprogram. These examples use the function `SDOT` with the following usage:

```
REAL*8 SDOT,S,X(1+(N-1)*|INCX|),Y(1+(N-1)*|INCX|)
S = SDOT (N, X,INCX, Y,INCX)
```

This sets  $S$  to the dot product of the vectors  $(N, X, INCX)$  and  $(N, Y, INCY)$ .

### Example 1

Compute the dot product  $T = X(1)*Y(1) + X(2)*Y(2) + X(3)*Y(3) + X(4)*Y(4)$ :

```
REAL*8 SDOT,T,X(4),Y(4)
T = SDOT (4, X,1, Y,1)
```

### Example 2

Compute the convolution  $T = X(1)*Y(4) + X(2)*Y(3) + X(3)*Y(2) + X(4)*Y(1)$ :

```
REAL*8 SDOT,T,X(4),Y(4)
T = SDOT (4, X,1, Y,-1)
```

### Example 3

Compute the dot product  $Y(2) = A(2,1)*X(1) + A(2,2)*X(2) + A(2,3)*X(3)$ , which is the dot product of the second row of an  $M$  by 3 matrix  $A$ , stored in a 10-by-3 array, with a 3-vector  $X$ :

```
PARAMETER (LDA = 10)
REAL*8 SDOT,A(LDA,3),X(3),Y(LDA)
N = 3
Y(2) = SDOT (N, A(2,1),LDA, X,1)
```

## Supplemental Reading

Lawson, C., R. Hanson, D. Kincaid, and F. Krogh. "Basic Linear Algebra Subprograms for Fortran Usage." *ACM Transactions on Mathematical Software*. September, 1979. Vol. 5, No. 3.

**NAME** CLUSEQ/.../CLUSILT – Find Clusters of Selected Vector Elements

### Purpose

Given a real or integer vector  $x$  of length  $n$ , these subprograms search sequentially through the vector and fill an array with the beginning and ending indices  $i$  of the maximal contiguous groups (clusters) of elements which satisfy a specified relationship with a given scalar  $a$ .

A cluster is a set of one or more elements  $\{x_i, x_{i+1}, \dots, x_j\}$ , with the following properties:

- 1) for every  $k$  with  $i \leq k \leq j$ ,  $x_k$  satisfies the specified relationship with  $a$ ,
- 2) either  $i = 1$  or  $x_{i-1}$  does not satisfy the relationship, and
- 3) either  $j = n$  or  $x_{j+1}$  does not satisfy the relationship.

At most, there are  $\lceil n/2 \rceil$  clusters, where  $\lceil x \rceil$  represents the smallest integer greater than or equal to  $x$ .

The last two characters of the subprogram name specify the relationship of interest between the elements of the vector and the scalar. These characters and the corresponding cluster relationship may be

xx	Cluster relationship
EQ	$\{i : x_i = a\}$
GE	$\{i : x_i \geq a\}$
GT	$\{i : x_i > a\}$
LE	$\{i : x_i \leq a\}$
LT	$\{i : x_i < a\}$
NE	$\{i : x_i \neq a\}$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

### Usage

SCILIB:

```

INTEGER*8      n, incx, indx(2, (n+1)/2), nindx
REAL*8         x(lenx), a
CALL CLUSEQ(n, x, incx, a, indx, nindx)
INTEGER*8      n, x(lenx), incx, a, indx(2, (n+1)/2), nindx

```

## CLUSEQ/.../CLUSILT – Find Clusters of Selected Vector Elements

```

CALL CLUSEQ(n, x, incx, a, indx, nindx)
INTEGER*8      n, incx, indx(2, (n+1)/2), nindx
REAL*8        x(lenx), a
CALL CLUSNE(n, x, incx, a, indx, nindx)
INTEGER*8      n, x(lenx), incx, a, indx(2, (n+1)/2), nindx
CALL CLUSNE(n, x, incx, a, indx, nindx)
INTEGER*8      n, incx, indx(2, (n+1)/2), nindx
REAL*8        x(lenx), a
CALL CLUSFxx(n, x, incx, a, indx, nindx)
INTEGER*8      n, x(lenx), incx, a, indx(2, (n+1)/2), nindx
CALL CLUSLxx(n, x, incx, a, indx, nindx)

```

## Input

- n** Number of elements of vector  $x$  to be compared to  $a$ . If  $n \leq 0$ , the subprograms do not reference  $x$  or  $\mathbf{indx}$ .
- x** Array of length  $\mathbf{lenx} = (n-1) \times |\mathbf{incx}| + 1$  containing the  $n$ -vector  $x$ .
- incx** Increment for the array  $x$ :
- incx**  $\geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \mathbf{incx} + 1)$ .
- incx**  $< 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \mathbf{incx} + 1)$ .
- Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.
- a** The scalar  $a$ .

## Output

- indx**  $\mathbf{indx}(1,k)$  and  $\mathbf{indx}(2,k)$  contain the indices of the beginning and ending elements of the  $k$ th cluster of  $x$  respectively. Only the first **nindx** elements of  $\mathbf{indx}$  are changed.
- nindx** If  $n \leq 0$ , then **nindx** = 0. Otherwise, **nindx** is the number of clusters of elements of  $x$  that satisfy the relationship with  $a$  specified by the subprogram name.

## Notes

These subprograms are sometimes useful for optimizing a loop containing an IF statement.

## Basic Vector Operations

### CLUSEQ/.../CLUSILT – Find Clusters of Selected Vector Elements

#### Fortran Equivalent

```
SUBROUTINE CLUSEQ (N,X,INCX,A,INDX,NINDX)
INTEGER*8 N,X(*),INCX,A,INDX(2,*),NINDX
LOGICAL*8 INCLUS
IX = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
NINDX = 0
INCLUS = .FALSE.
DO 10 I = 1, N
  IF ( .NOT. INCLUS ) THEN
    IF ( X(IX) .EQ. A ) THEN
      NINDX = NINDX + 1
      INDX(1,NINDX) = I
      INCLUS = .TRUE.
    END IF
  ELSE
    IF ( X(IX) .NE. A ) THEN
      INDX(2,NINDX) = I-1
      INCLUS = .FALSE.
    END IF
  END IF
  IX = IX + INCX
10 CONTINUE
IF ( INCLUS ) INDX(2,NINDX) = N
RETURN
END
```

#### Example

Find the clusters of positive elements of a REAL\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```
INTEGER*8 N,INCX,INDX(2,11),NINDX
REAL*8    A,X(20)
N = 10
INCX = 1
A = 0.0
CALL CLUSGT (N,X,INCX,A,INDX,NINDX)
```

**NAME**            GATHER – Gather Sparse Vector

**Purpose**

Given a dense real vector  $y$  stored in full storage form, and a set of indices of *interesting* elements of  $y$ , this subprogram gathers those elements into a sparse vector  $x$  stored in compact form via the set of indices.

More precisely, let  $\{k_1, k_2, \dots, k_m\}$  be the indices of the interesting elements. If  $x$  is represented by arrays  $\mathbf{x}$  and  $\mathbf{indx}$  such that  $\mathbf{indx}(i) = k_i$  and  $\mathbf{x}(i) = x_{k_i}$ , then

$$\mathbf{x}_i = y_{k_i}, \quad i = 1, 2, \dots, m.$$

**Usage**

SCILIB:

```

      INTEGER*8      m, indx(m)
      REAL*8        y(n), x(m)
      CALL GATHER(m, x, y, indx)
  
```

**Input**

- m**                    Number of interesting elements,  $\mathbf{m} \leq \mathbf{n}$ , where  $\mathbf{n}$  is the length of  $\mathbf{y}$ . If  $\mathbf{m} \leq 0$ , the subprogram does not reference  $\mathbf{x}$ ,  $\mathbf{indx}$ , or  $\mathbf{y}$ .
- y**                    Array containing the elements of  $\mathbf{y}$ ,  $\mathbf{y}(i) = y_i$ . Only the elements of  $\mathbf{y}$  whose indices are included in  $\mathbf{indx}$  are accessed.
- indx**                Array containing the indices  $\{k_i\}$  of the interesting elements of  $\mathbf{y}$ . The indices must satisfy
 
$$1 \leq \mathbf{indx}(i) \leq \mathbf{n}, \quad i = 1, 2, \dots, \mathbf{m},$$
 where  $\mathbf{n}$  is the length of  $\mathbf{y}$ .

**Output**

- x**                    If  $\mathbf{m} \leq 0$ , then  $\mathbf{x}$  is unchanged. Otherwise, the  $\mathbf{m}$  interesting elements of  $\mathbf{y}$ :  $\mathbf{x}(j) = y_i$  if  $\mathbf{indx}(j) = i$ .

**Notes**

Cray Research, Inc. has declared this subprogram obsolete in release 6.0 of the UNICOS Math and Scientific Library.

Basic Vector Operations  
**GATHER – Gather Sparse Vector**

The result is unspecified if any element of **indx** is out of range or if **x**, **indx**, and **y** overlap such that any element of **y** or any index shares a memory location with any element of **x**.

**Fortran Equivalent**

```
SUBROUTINE GATHER (M, X, Y, INDX)
  INTEGER*8 M, INDX(*)
  REAL*8 X(*), Y(*)
  DO 10 I = 1, M
    X(I) = Y(INDX(I))
10 CONTINUE
  RETURN
END
```

**Example**

Gather **y** into **x**, where **y** is a vector with interesting elements  $y_1, y_4, y_5,$  and  $y_9$  stored in one-dimensional array **Y** of dimension 20, and **x** is a vector stored in compact form in a one-dimensional array **X**.

```
INTEGER*8 M, INDX(4)
REAL*8    Y(20), X(4)
DATA      INDX / 1, 4, 5, 9 /
M = 4
CALL GATHER (M, X, Y, INDX)
```

**NAME** IILZ – Count Initial Zero Elements

### Purpose

Given an integer vector  $x$  of length  $n$ , this subprogram counts the number of zero elements of  $x$  before the first nonzero element. Given a logical vector  $x$ , IILZ counts the number of .FALSE. elements of  $x$  before the first .TRUE. element.

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

### Usage

SCILIB:

```

      INTEGER*8      i, IILZ, n, x(lenx), incx
      i = IILZ(n, x, incx)
      INTEGER*8      i, IILZ, n, incx
      LOGICAL*8      x(lenx)
      i = IILZ(n, x, incx)
  
```

### Input

**n** Number of elements of vector  $x$ . If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx** Increment for the array  $x$ :

**incx**  $\geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx**  $< 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

### Output

**i** If  $n \leq 0$ , then  $i = 0$ . If  $n > 0$  and all elements of  $x$  are zero or .TRUE., then  $i = n$ . Otherwise,  $i$  is the number of the zero or .FALSE. elements  $x_i$  of  $x$  before the first nonzero or .TRUE. element.

Basic Vector Operations  
IILZ – Count Initial Zero Elements

**Fortran Equivalent**

```
INTEGER*8 FUNCTION IILZ (N,X,INCX)
INTEGER*8 N,X(*),INCX
IILZ = 0
IX = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
DO 10 I = 1, N
    IF ( X(IX) .NE. 0 ) RETURN
    IX = IX + INCX
    IILZ = I
10 CONTINUE
RETURN
END
```

**Example**

Determine the number of initial zero elements of an INTEGER\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array X of dimension 20.

```
INTEGER*8 I,IILZ,N,X(20),INCX
N = 10
INCX = 1
I = IILZ (N,X,INCX)
```

**NAME** ILLZ – Count Initial Positive Elements

**Purpose**

Given a vector  $x$  of length  $n$ , this subprogram counts the number of positive elements of  $x$  before the first negative element. In this context, positive means that the leftmost or high-order bit is zero, and negative means that the leftmost bit is one.

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```
INTEGER*8      i, ILLZ, n, x(lenx), incx
i = ILLZ(n, x, incx)
```

**Input**

**n** Number of elements of vector  $x$ . If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx** Increment for the array  $x$ :

**incx**  $\geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx**  $< 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

**Output**

**i** If  $n \leq 0$ , then  $i = 0$ . If  $n > 0$  and all elements of  $x$  are zero, then  $i = n$ . Otherwise,  $i$  is the number of the positive elements  $x_i$  of  $x$  before the first negative element.

Basic Vector Operations  
**ILLZ – Count Initial Positive Elements**

**Fortran Equivalent**

```
INTEGER*8 FUNCTION ILLZ (N,X, INCX)
INTEGER*8 N,X(*), INCX
ILLZ = 0
IX = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
DO 10 I = 1, N
    IF ( X(IX) .LT. 0 ) RETURN
    IX = IX + INCX
    ILLZ = I
10 CONTINUE
RETURN
END
```

**Example**

Count the number of initial positive elements of an INTEGER\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array X of dimension 20.

```
INTEGER*8 I, ILLZ, N, X(20), INCX
N = 10
INCX = 1
I = ILLZ (N,X, INCX)
```

**NAME** ILSUM – Count TRUE Vector Elements

**Purpose**

Given a logical vector  $x$  of length  $n$ , this subprogram counts the number of elements of the vector that have the logical value `.TRUE`.

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```
INTEGER*8      i, ILSUM, n, incx
LOGICAL*8      x(lenx)
i = ILSUM(n, x, incx)
```

**Input**

**n** Number of elements of vector  $x$ . If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx** Increment for the array  $x$ :

**incx**  $\geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx**  $< 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

**Output**

**i** If  $n \leq 0$ , then  $i = 0$ . Otherwise,  $i$  is the number of elements of  $x$  that have the logical value `.TRUE`.

Basic Vector Operations  
**ILSUM – Count TRUE Vector Elements**

**Fortran Equivalent**

```
INTEGER*8 FUNCTION ILSUM (N,X,INCX)
INTEGER*8 N,INCX
LOGICAL*8 X(*)
ILSUM = 0
IX = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
DO 10 I = 1, N
    IF ( X(IX) ) ILSUM = ILSUM + 1
    IX = IX + INCX
10 CONTINUE
RETURN
END
```

**Example**

Count the number of **.TRUE.** elements of a **LOGICAL\*8** vector *x*, where *x* is a vector 10 elements long stored in a one-dimensional array *X* of dimension 20.

```
INTEGER*8    I, ILSUM, N, INCX
LOGICAL*8    X(20)
N = 10
INCX = 1
I = ILSUM (N,X, INCX)
```

## INFLMAX – Index of Maximum Element of Vector

**NAME** INFLMAX – Index of Maximum Element of Vector

**Purpose**

Given a vector  $x$  of length  $n$ , these subprograms determine the index of the first element  $x_i$  in which a specified group of bits attains its maximum value in the vector. Specifically, the subprograms determine the smallest index  $i$  such that

$$\text{AND}(\text{SHIFTR}(x_i, \text{rshift}), \text{mask}) = \max(\text{AND}(\text{SHIFTR}(x_j, \text{rshift}), \text{mask}) : j = 1, 2, \dots, n).$$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

INTEGER\*8            **i, INFLMAX, n, x(lenx), incx, mask, rshift**  
**i = INFLMAX(n, x, incx, mask, rshift)**

**Input**

**n**            Number of elements of vector  $x$  to be used. If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x**            Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx**        Increment for the array  $x$ :

**incx**  $\geq 0$       $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx**  $< 0$       $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

          Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

**mask**        Mask of 1-bits to extract desired group of bits from the shifted elements of  $x$  with a bitwise logical product operation. Refer to "Purpose."

**rshift**      Number of bits by which to right shift each element of  $x$  so as to align the specified group of bits with **a**,  $0 \leq \text{rshift} \leq 63$ . Refer to "Purpose."

## Basic Vector Operations

### INFLMAX – Index of Maximum Element of Vector

#### Output

**i**                      If  $n \leq 0$ , then  $i = 0$ . Otherwise,  $i$  is the index of the maximum element of  $x$ .

#### Fortran Equivalent

```
INTEGER*8 FUNCTION INFLMAX (N,X,INCX,MASK,RSHIFT)
INTEGER*8 N,X(*),INCX,MASK,RSHIFT,TEMP,XMAX
INFLMAX = 1
IF ( N .GT. 1 ) THEN
  IX = 1 + INCX
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  XMAX = AND(SHIFTR(X(IX-INCX),RSHIFT),MASK)
  DO 10 I = 2, N
    TEMP = AND(SHIFTR(X(I),RSHIFT),MASK)
    IF ( TEMP .GT. XMAX ) THEN
      INFLMAX = I
      XMAX = TEMP
    END IF
    IX = IX + INCX
10  CONTINUE
ELSE IF ( N .LT. 1 ) THEN
  INFLMAX = 0
END IF
RETURN
END
```

#### Example

Locate the element of an INTEGER\*8 vector  $x$  in which the field consisting of the rightmost 8 bits achieves its maximum value, where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```
INTEGER*8 I,INFLMAX,N,X(20),INCX,MASK,RSHIFT
N = 10
INCX = 1
MASK = 'FF'X
RSHIFT = 0
I = INFLMAX (N,X,INCX,MASK,RSHIFT)
```

**NAME** INFLMIN – Index of Minimum Element of Vector

### Purpose

Given a vector  $x$  of length  $n$ , these subprograms determine the index of the first element  $x_i$  in which a specified group of bits attains its minimum value in the vector. Specifically, the subprograms determine the smallest index  $i$  such that

$$\text{AND}(\text{SHIFTR}(x_i, \text{rshift}), \text{mask}) = \min(\text{AND}(\text{SHIFTR}(x_j, \text{rshift}), \text{mask}) : j = 1, 2, \dots, n).$$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

### Usage

SCILIB:

```
INTEGER*8          i, INFLMIN, n, x(lenx), incx, mask, rshift
i = INFLMIN(n, x, incx, mask, rshift)
```

### Input

**n** Number of elements of vector  $x$  to be used. If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx** Increment for the array  $x$ :  
**incx**  $\geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .  
**incx**  $< 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

**mask** Mask of 1-bits to extract desired group of bits from the shifted elements of  $x$  with a bitwise logical product operation. Refer to “Purpose.”

**rshift** Number of bits by which to right shift each element of  $x$  so as to align the specified group of bits with **a**,  $0 \leq \text{rshift} \leq 63$ . Refer to “Purpose.”

## Basic Vector Operations

### INFLMIN – Index of Minimum Element of Vector

#### Output

**i**                      If  $n \leq 0$ , then  $i = 0$ . Otherwise,  $i$  is the index of the minimum element of  $x$ .

#### Fortran Equivalent

```
INTEGER*8 FUNCTION INFLMIN (N,X,INCX,MASK,RSHIFT)
INTEGER*8 N,X(*),INCX,MASK,RSHIFT,TEMP,XMIN
INFLMIN = 1
IF ( N .GT. 1 ) THEN
  IX = 1 + INCX
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  XMIN = AND(SHIFTR(X(IX-INCX),RSHIFT),MASK)
  DO 10 I = 2, N
    TEMP = AND(SHIFTR(X(IX),RSHIFT),MASK)
    IF ( TEMP .LT. XMIN ) THEN
      INFLMIN = I
      XMIN = TEMP
    END IF
    IX = IX + INCX
10  CONTINUE
ELSE IF ( N .LT. 1 ) THEN
  INFLMIN = 0
END IF
RETURN
END
```

#### Example

Locate the element of an INTEGER\*8 vector  $x$  in which the field consisting of the rightmost 8 bits achieves its minimum value, where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```
INTEGER*8 I,INFLMIN,N,X(20),INCX,MASK,RSHIFT
N = 10
INCX = 1
MASK = 'FF'X
RSHIFT = 0
I = INFLMIN (N,X,INCX,MASK,RSHIFT)
```

## ISAMAX/ICAMAX – Index of Maximum of Magnitudes

**NAME** ISAMAX/ICAMAX – Index of Maximum of Magnitudes

**Purpose**

Given a real or integer vector  $x$  of length  $n$ , ISAMAX determines the index of the element of the vector of maximum magnitude. Specifically, the subprograms determine the smallest index  $i$  such that

$$|x_i| = \max(|x_j| : j = 1, 2, \dots, n).$$

Given a complex vector  $x$  of length  $n$ , ICAMAX determines the smallest index  $i$  such that

$$|Re(x_i)| + |Im(x_i)| = \max(|Re(x_j)| + |Im(x_j)| : j = 1, 2, \dots, n)$$

where  $Re(x_i)$  and  $Im(x_i)$  are the real and imaginary parts of  $x_i$ , respectively. The usual definition of complex magnitude is

$$\left\{ Re(x_i)^2 + Im(x_i)^2 \right\}^{1/2}.$$

This definition is not used because of computational speed. If the index  $i$  is used for pivot selection in matrix factorization, no significant difference in numerical stability should result.

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```

INTEGER*8      i, ISAMAX, n, incx
REAL*8         x(lenx)
i = ISAMAX(n, x, incx)

INTEGER*8      i, ICAMAX, n, incx
COMPLEX*16     x(lenx)
i = ICAMAX(n, x, incx)

```

**Input**

**n** Number of elements of vector  $x$  to be used. If  $n \leq 0$ , the subprograms do not reference  $x$ .

## Basic Vector Operations

### ISAMAX/ICAMAX – Index of Maximum of Magnitudes

- x**                    Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .
- incx**                  Increment for the array  $x$ :
- $\text{incx} \geq 0$          $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .
- $\text{incx} < 0$          $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .
- Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

## Output

- i**                      If  $n \leq 0$ , then  $i = 0$ . Otherwise,  $i$  is the index of the element of  $x$  of maximum magnitude.

## Notes

The handling of  $\text{incx} < 0$  differs between these subprograms and ISAMAX/ICAMAX in VECLIB.

## Fortran Equivalent

```
INTEGER*8 FUNCTION ISAMAX (N,X,INCX)
INTEGER*8 N,INCX
REAL*8 X(*),TEMP,XMAX
ISAMAX = 1
IF ( N .GT. 1 ) THEN
  IX = 1 + INCX
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  XMAX = ABS ( X(IX - INCX) )
  DO 10 I = 2, N
    TEMP = ABS ( X(IX) )
    IF ( TEMP .GT. XMAX ) THEN
      ISAMAX = I
      XMAX = TEMP
    END IF
    IX = IX + INCX
10  CONTINUE
  ELSE IF ( N .LT. 1 ) THEN
    ISAMAX = 0
  END IF
  RETURN
END
```

## Example

Locate the largest element of a REAL\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

Basic Vector Operations  
**ISAMAX/ICAMAX – Index of Maximum of Magnitudes**

```
INTEGER*8 I, ISAMAX, N, INCX  
REAL*8    X(20)  
N = 10  
INCX = 1  
I = ISAMAX (N, X, INCX)
```

**NAME** ISAMIN – Index of Minimum of Magnitudes

### Purpose

Given a real or integer vector  $x$  of length  $n$ , ISAMIN determines the index of element of the vector of minimum magnitude. Specifically, the subprogram determines the smallest index  $i$  such that

$$|x_i| = \min(|x_j| : j = 1, 2, \dots, n).$$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

### Usage

SCILIB:

INTEGER\*8            **i, ISAMIN, n, incx**

REAL\*8                **x(lenx)**

**i = ISAMIN(n, x, incx)**

### Input

- n**                    Number of elements of vector  $x$  to be used. If  $n \leq 0$ , the subprogram does not reference  $x$ .
- x**                    Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .
- incx**                Increment for the array  $x$ :
- incx**  $\geq 0$          $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .
- incx**  $< 0$          $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .
- Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

### Output

- i**                    If  $n \leq 0$ , then **i** = 0. Otherwise, **i** is the index of the element of  $x$  of minimum magnitude.

### Notes

The handling of **incx**  $< 0$  differs between ISAMIN in SCILIB and VECLIB.

## Fortran Equivalent

```
INTEGER*8 FUNCTION ISAMIN (N,X,INCX)
INTEGER*8 N,INCX
REAL*8 X(*),TEMP,XMIN
ISAMIN = 1
IF ( N .GT. 1 ) THEN
  IX = 1 + INCX
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  XMIN = ABS ( X(IX-INCX) )
  DO 10 I = 2, N
    TEMP = ABS ( X(IX) )
    IF ( TEMP .LT. XMIN ) THEN
      ISAMIN = I
      XMIN = TEMP
    END IF
    IX = IX + INCX
10  CONTINUE
  ELSE IF ( N .LT. 1 ) THEN
    ISAMIN = 0
  END IF
RETURN
END
```

## Example

Locate the smallest element of a REAL\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```
INTEGER*8 I,ISAMIN,N,INCX
REAL*8 X(20)
N = 10
INCX = 1
I = ISAMIN (N,X,INCX)
```

Basic Vector Operations  
ISMAX/INTMAX – Index of Maximum Element of Vector

**NAME** ISMAX/INTMAX – Index of Maximum Element of Vector

**Purpose**

Given a real or integer vector  $x$  of length  $n$ , these subprograms determine the index of maximum element of the vector. Specifically, the subprograms determine the smallest index  $i$  such that

$$x_i = \max(x_j : j = 1, 2, \dots, n).$$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```
INTEGER*8      i, ISMAX, n, incx
REAL*8         x(lenx)
i = ISMAX(n, x, incx)
INTEGER*8      i, INTMAX, n, x(lenx), incx
i = INTMAX(n, x, incx)
```

**Input**

**n** Number of elements of vector  $x$  to be used. If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx** Increment for the array  $x$ :

**incx**  $\geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx**  $< 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

**Output**

**i** If  $n \leq 0$ , then **i** = 0. Otherwise, **i** is the index of the maximum element of  $x$ .

## Notes

The handling of `incx < 0` differs between ISMAX in SCILIB and VECLIB.

## Fortran Equivalent

```
INTEGER*8 FUNCTION ISMAX (N,X, INCX)
INTEGER*8 N, INCX
REAL*8 X(*), XMAX
ISMAX = 1
IF ( N .GT. 1 ) THEN
  IX = 1 + INCX
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  XMAX = X(IX - INCX)
  DO 10 I = 2, N
    IF ( X(IX) .GT. XMAX ) THEN
      ISMAX = I
      XMAX = X(IX)
    END IF
    IX = IX + INCX
10  CONTINUE
ELSE IF ( N .LT. 1 ) THEN
  ISMAX = 0
END IF
RETURN
END
```

## Example

Locate the largest element of a REAL\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```
INTEGER*8 I, ISMAX, N, INCX
REAL*8 X(20)
N = 10
INCX = 1
I = ISMAX (N,X, INCX)
```

Basic Vector Operations  
ISMIN/INTMIN – Index of Minimum Element of Vector

**NAME** ISMIN/INTMIN – Index of Minimum Element of Vector

**Purpose**

Given a real or integer vector  $x$  of length  $n$ , these subprograms determine the index of minimum element of the vector. Specifically, the subprograms determine the smallest index  $i$  such that

$$x_i = \min(x_j : j = 1, 2, \dots, n).$$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```
INTEGER*8      i, ISMIN, n, incx
REAL*8        x(lenx)
i = ISMIN(n, x, incx)
INTEGER*8      i, INTMIN, n, x(lenx), incx
i = INTMIN(n, x, incx)
```

**Input**

**n** Number of elements of vector  $x$  to be used. If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx** Increment for the array  $x$ :

**incx**  $\geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx**  $< 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

**Output**

**i** If  $n \leq 0$ , then  $i = 0$ . Otherwise,  $i$  is the index of the minimum element of  $x$ .

## Notes

The handling of  $incx < 0$  differs between ISMIN in SCILIB and VECLIB.

## Fortran Equivalent

```
INTEGER*8 FUNCTION ISMIN (N,X,INCX)
INTEGER*8 N,INCX
REAL*8 X(*),XMIN
ISMIN = 1
IF ( N .GT. 1 ) THEN
  IX = 1 + INCX
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  XMIN = X(IX - INCX)
  DO 10 I = 2, N
    IF ( X(IX) .LT. XMIN ) THEN
      ISMIN = I
      XMIN = X(IX)
    END IF
    IX = IX + INCX
10  CONTINUE
ELSE IF ( N .LT. 1 ) THEN
  ISMIN = 0
END IF
RETURN
END
```

## Example

Locate the smallest element of a REAL\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```
INTEGER*8 I, ISMIN, N, INCX
REAL*8 X(20)
N = 10
INCX = 1
I = ISMIN (N,X,INCX)
```

**NAME** ISRCHEQ/ISRCHNE/.../ISRCHILT – Search Vector for Element

### Purpose

Given a real or integer vector  $x$  of length  $n$ , these subprograms search sequentially through the vector for the first element  $x_i$  that satisfies a specified relationship with a given scalar  $a$  and return the index  $i$  of that element.

The last two characters of the subprogram name specify the relationship of interest between the element of the vector and the scalar. These characters and the corresponding function values may be

xx	Function value
EQ	$\min\{i : x_i = a\}$
GE	$\min\{i : x_i \geq a\}$
GT	$\min\{i : x_i > a\}$
LE	$\min\{i : x_i \leq a\}$
LT	$\min\{i : x_i < a\}$
NE	$\min\{i : x_i \neq a\}$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

### Usage

SCILIB:

```

INTEGER*8      i, ISRCHEQ, n, incx
REAL*8        x(lenx), a
i = ISRCHEQ(n, x, incx, a)

INTEGER*8      i, ISRCHEQ, n, x(lenx), incx, a
i = ISRCHEQ(n, x, incx, a)

INTEGER*8      i, ISRCHNE, n, incx
REAL*8        x(lenx), a
i = ISRCHNE(n, x, incx, a)

INTEGER*8      i, ISRCHNE, n, x(lenx), incx, a
i = ISRCHNE(n, x, incx, a)

INTEGER*8      i, ISEARCH, n, incx
REAL*8        x(lenx), a
i = ISEARCH(n, x, incx, a)

INTEGER*8      i, ISEARCH, n, x(lenx), incx, a
i = ISEARCH(n, x, incx, a)

INTEGER*8      i, ISRCHFxx, n, incx

```

## ISRCHQ/ISRCHNE/.../ISRCHILT – Search Vector for Element

```

REAL*8          x(lenx), a
i = ISRCHFxx(n, x, incx, a)
INTEGER*8       i, ISRCHLxx, n, x(lenx), incx, a
i = ISRCHLxx(n, x, incx, a)

```

**Input**

**n**                    Number of elements of vector  $x$  to be compared to  $a$ . If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x**                    Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx**                Increment for the array  $x$ :

$\text{incx} \geq 0$       $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

$\text{incx} < 0$       $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

                      Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

**a**                    The scalar  $a$ .

**Output**

**i**                    If  $n \leq 0$ , then  $i = 0$ . If  $n > 0$  and no element of  $x$  satisfies the relationship with  $a$  specified by the subprogram name, then  $i = n + 1$ . Otherwise,  $i$  is the index  $i$  of the first element  $x_i$  of  $x$  that satisfies the relationship with  $a$  specified by the subprogram name.

## Basic Vector Operations

### ISRCHQ/ISRCHNE/.../SRCHILT – Search Vector for Element

#### Fortran Equivalent

```
INTEGER*8 FUNCTION ISRCHQ (N,X, INCX,A)
INTEGER*8 N,X(*), INCX,A
ISRCHQ = 0
IF ( N .LE. 0 ) RETURN
IX = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
DO 10 I = 1, N
  IF ( X(IX) .EQ. A ) THEN
    ISRCHQ = I
    RETURN
  END IF
  IX = IX + INCX
10 CONTINUE
ISRCHQ = N+1
RETURN
END
```

#### Example

Search for the first positive element of a REAL\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```
INTEGER*8 I,ISRCHFGT,N, INCX
REAL*8 X(20),A
N = 10
INCX = 1
A = 0.0
I = ISRCHFGT (N,X, INCX,A)
```

## ISRCHMEQ/ISRCHMGE/.../ISRCHMNE – Search Vector for Element

**NAME** ISRCHMEQ/ISRCHMGE/.../ISRCHMNE – Search Vector for Element

**Purpose**

Given a vector  $x$  of length  $n$ , these subprograms search sequentially through the vector for the first element  $x_i$  which contains a specified group of bits that satisfies a specified relationship with a given scalar  $a$ , and return the index  $i$  of that element.

The last two characters of the subprogram name specify the relationship of interest between the element of the vector and the scalar. These characters and the corresponding function values, may be

<b>xx</b>	Function value
<b>EQ</b>	$\min\{i : \text{AND}(\text{SHIFTR}(x_i, rshift), mask) = a\}$
<b>GE</b>	$\min\{i : \text{AND}(\text{SHIFTR}(x_i, rshift), mask) \geq a\}$
<b>GT</b>	$\min\{i : \text{AND}(\text{SHIFTR}(x_i, rshift), mask) > a\}$
<b>LE</b>	$\min\{i : \text{AND}(\text{SHIFTR}(x_i, rshift), mask) \leq a\}$
<b>LT</b>	$\min\{i : \text{AND}(\text{SHIFTR}(x_i, rshift), mask) < a\}$
<b>NE</b>	$\min\{i : \text{AND}(\text{SHIFTR}(x_i, rshift), mask) \neq a\}$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

INTEGER\*8            **i, ISRCHMxx, n, x(lenx), incx, a, mask, rshift**  
**i = ISRCHMxx(n, x, incx, a, mask, rshift)**

**Input**

**n**                    Number of elements of vector  $x$  to be compared to  $a$ . If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x**                    Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx**                Increment for the array  $x$ :

**incx**  $\geq 0$          $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx**  $< 0$          $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

## Basic Vector Operations

### ISRCHMEQ/ISRCHMGE/.../ISRCHMNE – Search Vector for Element

	Use $\text{incx} = 1$ if the vector $x$ is stored contiguously in array $x$ , i.e., if $x_i$ is stored in $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.
<b>a</b>	The scalar $a$ .
<b>mask</b>	Mask of 1-bits to extract desired group of bits from the shifted elements of $x$ with a bitwise logical product operation. Refer to “Purpose.”
<b>rshift</b>	Number of bits by which to right shift each element of $x$ so as to align the specified group of bits with $a$ , $0 \leq \text{rshift} \leq 63$ . Refer to “Purpose.”

## Output

<b>i</b>	If $n \leq 0$ , then $i = 0$ . If $n > 0$ and no element of $x$ satisfies the relationship with $a$ specified by the subprogram name, then $i = n + 1$ . Otherwise, $i$ is the index $i$ of the first element $x_i$ of $x$ that satisfies the relationship with $a$ specified by the subprogram name.
----------	---

## Fortran Equivalent

```
INTEGER*8 FUNCTION ISRCHMEQ (N,X, INCX,A,MASK,RSHIFT)
INTEGER*8 N,X(*), INCX,A,MASK,RSHIFT
ISRCHMEQ = 0
IF ( N .LE. 0 ) RETURN
IX = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
DO 10 I = 1, N
  IF ( AND(SHIFTR(X(IX),RSHIFT),MASK) .EQ. A ) THEN
    ISRCHMEQ = I
    RETURN
  END IF
  IX = IX + INCX
10 CONTINUE
ISRCHMEQ = N+1
RETURN
END
```

## Example

Search for the first odd element of an INTEGER\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```
INTEGER*8 I, ISRCHMGT,N,X(20), INCX,A,MASK,RSHIFT
N = 10
INCX = 1
A = 1
MASK = 1
RSHIFT = 0
I = ISRCHMEQ (N,X, INCX,A,MASK,RSHIFT)
```

**NAME** OSRCHF/OSRCHI – Search Ordered Vector for Element

**Purpose**

Given an ordered real or integer vector  $x$  of length  $n$ , these subprograms search sequentially through the vector for the first element  $x_i$  that equals a given scalar  $a$  and return the index  $i$  of that element. They also return the number of elements of the vector that are equal to the scalar, and the index of the location within the vector where the scalar should fit in the array, whether they find it or not.

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```

INTEGER*8      n, incx, indx, ifound, iwould, icount
REAL*8        x(lenx), a
CALL OSRCHF(n, x, incx, a, ifound, iwould, icount)
INTEGER*8      n, x(lenx), incx, a, ifound, iwould, icount
CALL OSRCHI(n, x, incx, a, ifound, iwould, icount)
    
```

**Input**

**n** Number of elements of vector  $x$  to be compared to  $a$ . If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ . The elements of  $x$  are assumed to be in ascending order.

**incx** Increment for the array  $x$ :  
 $\text{incx} \geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .  
 $\text{incx} < 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

**a** The scalar  $a$ .

**icount** Flag indicating whether to count the number of occurrences of  $a$  in  $x$ :  
 $\text{icount} \neq 0$  do count the number of occurrences of  $a$  in  $x$ .

**icount** = 0    do not count the number of occurrences of **a** in  $x$ .

## Output

<b>ifound</b>	If $n \leq 0$ , then <b>ifound</b> = 0. If $n > 0$ and no element of $x$ equals <b>a</b> , then <b>ifound</b> = $nf + 1$ . Otherwise, <b>ifound</b> is the index $i$ of the first element $x_i$ of $x$ that equals <b>a</b> .
<b>iwould</b>	If $n \leq 0$ , then <b>iwould</b> = 0. If $n > 0$ , then <b>iwould</b> is the index $i$ of the first element $x_i$ of $x$ such that $x_i \leq a \leq x_{i-1}$ . If <b>a</b> is found in $x$ , then <b>iwould</b> = <b>ifound</b> .
<b>icount</b>	If <b>icount</b> $\neq 0$ on entry, the number of occurrences of <b>a</b> in $x$ . Not used as output if <b>icount</b> = 0.

## Notes

No check is made to ensure that the elements of  $x$  are in ascending order. The output is undefined if the elements are unordered.

**Fortran Equivalent**

```

SUBROUTINE ORSCHF (N,X, INCX,A, IFOUND, IWOULD, ICOUNT)
INTEGER*8 N, INCX, IFOUND, IWOULD, ICOUNT
REAL*8 X(*),A
IF ( N .LE. 0 ) RETURN
    IFOUND = 0
    IWOULD = 0
    IF ( ICOUNT .NE. 0 ) ICOUNT = 0
    RETURN
END IF
IX = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
DO 20 I = 1, N
    IF ( X(IX) .GE. A ) THEN
        IF ( X(IX) .EQ. A ) THEN
            IFOUND = I
            IWOULD = I
            IF ( ICOUNT .NE. 0 ) THEN
                ICOUNT = 1
                IX = IX + INCX
                DO 10 J = I+1, N
                    IF ( X(IX) .EQ. A ) THEN
                        ICOUNT = ICOUNT + 1
                        IX = IX + INCX
                    ELSE
                        RETURN
                    END IF
                CONTINUE
            END IF
        ELSE
            IFOUND = N+1
            IWOULD = I
            IF ( ICOUNT .NE. 0 ) ICOUNT = 0
        END IF
        RETURN
    END IF
    IX = IX + INCX
20 CONTINUE
    IFOUND = N+1
    IWOULD = N+1
    IF ( ICOUNT .NE. 0 ) ICOUNT = 0
    RETURN
END

```

**Example**

Search for the first element of a REAL\*8 vector  $x$  equal to 10, where  $x$  is a ordered vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```

INTEGER*8 N, INCX, IFOUND, IWOULD, ICOUNT
REAL*8    X(20),A
N = 10
INCX = 1
A = 10.0
CALL ORSCHF (N,X, INCX,A, IFOUND, IWOULD, ICOUNT)

```

**NAME** OSRCHM – Search Ordered Vector for Element

**Purpose**

Given an ordered integer vector  $x$  of length  $n$ , this subprogram searches sequentially through the vector for the first element  $x_i$  which contains a specified group of bits that equals a given scalar  $a$  and returns the index  $i$  of that element. The index is such that

$$\text{AND}(\text{SHIFTR}(x_i, \text{rshift}), \text{mask}) = a$$

It also returns the number of elements of the vector that are equal to the scalar, and the index of the location within the vector where the scalar should fit in the array, whether it finds it or not.

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```

      INTEGER*8      n, x(lenx), incx, a, mask, rshift, ifound, iwould,
                    icount
      CALL OSRCHM(n, x, incx, a, mask, rshift, ifound, iwould, icount)
  
```

**Input**

**n** Number of elements of vector  $x$  to be compared to  $a$ . If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ . The elements of  $x$  are assumed to be in ascending order.

**incx** Increment for the array  $x$ :

**incx**  $\geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx**  $< 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

**a** The scalar  $a$ .

<b>mask</b>	Mask of 1-bits to extract desired group of bits from the shifted elements of $x$ with a bitwise logical product operation. Refer to “Purpose.”
<b>rshift</b>	Number of bits by which to right shift each element of $x$ so as to align the specified group of bits with $a$ , $0 \leq \text{rshift} \leq 63$ . Refer to “Purpose.”
<b>icount</b>	Flag indicating whether to count the number of occurrences of $a$ in $x$ : <b>icount</b> $\neq 0$ do count the number of occurrences of $a$ in $x$ . <b>icount</b> = 0    do not count the number of occurrences of $a$ in $x$ .

## Output

<b>ifound</b>	If $n \leq 0$ , then <b>ifound</b> = 0. If $n > 0$ and no element of $x$ equals $a$ , then <b>ifound</b> = $n + 1$ . Otherwise, <b>ifound</b> is the index $i$ of the first element $x_i$ of $x$ that equals $a$ .
<b>iwould</b>	If $n \leq 0$ , then <b>iwould</b> = 0. If $n > 0$ , then <b>iwould</b> is the index $i$ of the first element $x_i$ of $x$ such that $x_i \leq a \leq x_{i+1}$ . If $a$ is found in $x$ , then <b>iwould</b> = <b>ifound</b> .
<b>icount</b>	If <b>icount</b> $\neq 0$ on entry, the number of occurrences of $a$ in $x$ . Not used as output if <b>icount</b> = 0.

## Notes

No check is made to ensure that the specified group of bits of the elements of  $x$  are in ascending order. The output is undefined if the elements are unordered.

Basic Vector Operations  
**OSRCHM – Search Ordered Vector for Element**

**Fortran Equivalent**

```

SUBROUTINE ORSCHM (N,X, INCX,A,MASK,RSHIFT, IFOUND, IWOULD, ICOUNT)
INTEGER*8 N,X(*), INCX,A, IFOUND, IWOULD, ICOUNT
IF ( N .LE. 0 ) THEN
    IFOUND = 0
    IWOULD = 0
    IF ( ICOUNT .NE. 0 ) ICOUNT = 0
    RETURN
END IF
IX = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
DO 20 I = 1, N
    IF ( AND(SHIFTR(X(IX),RSHIFT),MASK) .GE. A ) THEN
        IF ( AND(SHIFTR(X(IX),RSHIFT),MASK) .EQ. A ) THEN
            IFOUND = I
            IWOULD = I
            IF ( ICOUNT .NE. 0 ) THEN
                ICOUNT = 1
                IX = IX + INCX
                DO 10 J = I+1, N
                    IF ( AND(SHIFTR(X(IX),RSHIFT),MASK) .EQ. A ) THEN
                        ICOUNT = ICOUNT + 1
                        IX = IX + INCX
                    ELSE
                        RETURN
                    END IF
                CONTINUE
            END IF
        ELSE
            IFOUND = N+1
            IWOULD = I
            IF ( ICOUNT .NE. 0 ) ICOUNT = 0
        END IF
        RETURN
    END IF
    IX = IX + INCX
20 CONTINUE
IFOUND = N+1
IWOULD = N+1
IF ( ICOUNT .NE. 0 ) ICOUNT = 0
RETURN
END

```

**Example**

Search for the first element of an INTEGER\*8 vector  $x$  such that the second group of eight bits from the right contain the value 13, where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```

INTEGER*8 N,X(20), INCX,A,MASK,RSHIFT, IFOUND, IWOULD, ICOUNT
N = 10
INCX = 1
A = 13
MASK = 'FF'X
RSHIFT = 8
CALL OSRCHM (N,X, INCX,A,MASK,RSHIFT, IFOUND, IWOULD, ICOUNT)

```

**NAME** SASUM/SCASUM – Sum of Magnitudes

**Purpose**

Given a real or integer vector  $x$  of length  $n$ , SASUM computes the  $l_1$  norm of  $x$ , i.e., the sum of magnitudes of the elements of the vector

$$s = \|x\|_1 = \sum_{i=1}^n |x_i|.$$

Given a complex vector  $x$  of length  $n$ , SCASUM computes

$$s = \sum_{i=1}^n |Re(x_i)| + |Im(x_i)|$$

where  $Re(x_i)$  and  $Im(x_i)$  are the real and imaginary parts of  $x_i$ , respectively. The usual definition of sum of magnitudes of a complex vector is

$$t = \|x\|_1 = \sum_{i=1}^n \left\{ Re(x_i)^2 + Im(x_i)^2 \right\}^{1/2}.$$

$s$  is computed instead of  $t$  since it is faster because it does not require the  $n$  square roots. Since  $t \leq s \leq \sqrt{2}t$ ,  $s$  will often be an acceptable substitute for  $t$ .

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```

INTEGER*8      n, incx
REAL*8        s, SASUM, x(lenx)
s = SASUM(n, x, incx)

INTEGER*8      n, incx
REAL*8        s, SCASUM
COMPLEX*16    x(lenx)
s = SCASUM(n, x, incx)

```

**Input**

**n** Number of elements of vector  $x$  to be used in the sum of magnitudes. If  $n \leq 0$ , the subprograms do not reference  $x$ .

## Basic Vector Operations

### SASUM/SCASUM – Sum of Magnitudes

- x**                    Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .
- incx**                Increment for the array  $x$ .  $x$  is stored forward in array  $x$  with increment  $|\text{incx}|$ , i.e.,  $x_i$  is stored in  $x((i-1) \times |\text{incx}| + 1)$ .
- Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

### Output

- s**                    If  $n \leq 0$ , then  $s = 0$ . Otherwise,  $s$  is the sum of magnitudes of the elements of  $x$ .

### Fortran Equivalent

```
REAL*8 FUNCTION SASUM (N, X, INCX)
INTEGER*8 N, INCX
REAL*8 X(*)
SASUM = 0.0
IF ( N .LE. 0 ) RETURN
IX = 1
INCXA = ABS ( INCX )
DO 10 I = 1, N
    SASUM = SASUM + ABS ( X(IX) )
    IX = IX + INCXA
10 CONTINUE
RETURN
END
```

### Example

Compute the sum of magnitudes of the elements of a REAL\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```
INTEGER*8 N, INCX
REAL*8    S, SASUM, X(20)
N = 10
INCX = 1
S = SASUM (N, X, INCX)
```

**NAME** SAXPY/CAXPY – Elementary Vector Operation

**Purpose**

Given a real or complex scalar  $a$  and real or complex vectors  $x$  and  $y$  of length  $n$ , these subprograms perform the elementary vector operations

$$y \leftarrow ax+y$$

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

INTEGER*8      n, incx, incy
REAL*8         a, x(lenx), y(leny)
CALL SAXPY(n, a, x, incx, y, incy)

INTEGER*8      n, incx, incy
COMPLEX*16     a, x(lenx), y(leny)
CALL CAXPY(n, a, x, incx, y, incy)
    
```

**Input**

- n**                    Number of elements of vectors  $x$  and  $y$  to be used in the elementary vector operation. If  $n \leq 0$ , the subprograms do not reference  $x$  or  $y$ .
- a**                    The scalar  $a$ .
- x**                    Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ . Refer to "Purpose."
- incx**                Increment for the array  $x$ :
  - $\text{incx} \geq 0$          $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .
  - $\text{incx} < 0$          $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.
- y**                    Array of length  $\text{leny} = (n-1) \times |\text{incy}| + 1$  containing the  $n$ -vector  $y$ .
- incy**                Increment for the array  $y$ ,  $\text{incy} \neq 0$ :

## Basic Vector Operations

### SAXPY/CAXPY – Elementary Vector Operation

$\text{incy} > 0$   $y$  is stored forward in array  $y$ , i.e.,  $y_i$  is stored in  $y((i-1)\times\text{incy}+1)$ .

$\text{incy} < 0$   $y$  is stored backward in array  $y$ , i.e.,  $y_i$  is stored in  $y((i-n)\times\text{incy}+1)$ .

Use  $\text{incy} = 1$  if the vector  $y$  is stored contiguously in array  $y$ , i.e., if  $y_i$  is stored in  $y(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

### Output

$y$  If  $n \leq 0$  or  $a = 0$ , then  $y$  is unchanged. Otherwise,  $ax+y$  overwrites the input.

### Notes

If  $\text{incx} = 0$ , then  $x_i = x(1)$  for all  $i$ .

The result is unspecified if  $\text{incy} = 0$  or if  $x$  and  $y$  overlap such that any element of  $x$  shares a memory location with any element of  $y$ .

### Fortran Equivalent

```
SUBROUTINE SAXPY (N, A, X, INCX, Y, INCY)
  INTEGER*8 N, INCX, INCY
  REAL*8 X(*), Y(*), A
  IF ( N .LE. 0 ) RETURN
  IF ( A .EQ. 0.0 ) RETURN
  IX = 1
  IY = 1
  IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
  IF ( INCY .LT. 0 ) IY = 1 - (N-1) * INCY
  DO 10 I = 1, N
    Y(IY) = A * X(IX) + Y(IY)
    IX = IX + INCX
    IY = IY + INCY
  10 CONTINUE
  RETURN
END
```

### Example 1

Compute the REAL\*8 elementary vector operation

$$y \leftarrow 2x+y,$$

where  $x$  and  $y$  are vectors 10 elements long stored in one-dimensional arrays  $X$  and  $Y$  of dimension 20.

```
INTEGER*8 N, INCX, INCY
REAL*8 A, X(20), Y(20)
      N = 10
```

```
A = 2.0  
INCX = 1  
INCY = 1  
CALL SAXPY (N,A,X,INCX,Y,INCY)
```

## Example 2

Subtract 3 times the 4th row of a 10-by-10 matrix from the 5th row. The matrix is stored in a two-dimensional array B of dimension 20 by 21.

```
INTEGER*8 N, INCX, INCY  
REAL*8    A,B(20,21)  
N = 10  
A = -3.0  
INCX = 20  
INCY = 20  
CALL SAXPY (N,A,B(4,1),INCX,B(5,1),INCY)
```

Basic Vector Operations  
**SCATTER – Scatter Sparse Vector**

**NAME** SCATTER – Scatter Sparse Vector

**Purpose**

Given a sparse vector  $x$  stored in compact form via a set of indices, this subprogram scatters those elements into the corresponding elements of a dense vector  $y$  stored in full storage form.

More precisely, let  $x$  be a sparse  $n$ -vector with  $m \leq n$  interesting (usually nonzero) elements, and let  $\{k_1, k_2, \dots, k_m\}$  be the indices of these elements. If  $x$  is represented by arrays  $\mathbf{x}$  and  $\mathbf{indx}$  such that  $\mathbf{indx}(i) = k_i$  and  $\mathbf{x}(i) = x_{k_i}$ , then

$$y_{k_i} = \mathbf{x}_i, \quad i = 1, 2, \dots, m.$$

**Usage**

SCILIB:

```
INTEGER*8      m, indx(m)
REAL*8         y(n), x(m)
CALL SCATTER(m, y, indx, x)
```

**Input**

**m** Number of interesting elements,  $\mathbf{m} \leq \mathbf{n}$ , where  $\mathbf{n}$  is the length of  $\mathbf{y}$ . If  $\mathbf{m} \leq 0$ , the subprogram does not reference  $\mathbf{x}$ ,  $\mathbf{indx}$ , or  $\mathbf{y}$ .

**indx** Array containing the indices  $\{k_i\}$  of the interesting elements of  $x$ . The indices must satisfy

$$1 \leq \mathbf{indx}(i) \leq \mathbf{n}, \quad i = 1, 2, \dots, \mathbf{m}$$

and

$$\mathbf{indx}(i) \neq \mathbf{indx}(j), \quad 1 \leq i \neq j \leq \mathbf{m},$$

where  $\mathbf{n}$  is the length of  $\mathbf{y}$ .

**x** Array of length  $\mathbf{m}$  containing the interesting elements of  $x$ .  $\mathbf{x}(j) = x_i$  if  $\mathbf{indx}(j) = i$ .

**Output**

**y** Array containing the elements of  $\mathbf{y}$ ,  $\mathbf{y}(i) = y_i$ . If  $\mathbf{m} \leq 0$ , then  $\mathbf{y}$  is unchanged. Otherwise, only the elements of  $\mathbf{y}$  whose indices are included in  $\mathbf{indx}$  are changed.

## Notes

Cray Research, Inc. has declared this subprogram obsolete in release 6.0 of the UNICOS Math and Scientific Library.

The result is unspecified if any element of **indx** is out of range, if any two elements of **indx** have the same value, or if **x**, **indx**, and **y** overlap such that any element of **x** or any index shares a memory location with any element of **y**.

## Fortran Equivalent

```
      SUBROUTINE SCATTER (M, Y, INDX, X)
      INTEGER*8 M,INDX(*)
      REAL*8 Y(*),X(*)
      DO 10 I = 1, M
         Y(INDX(I)) = X(I)
10 CONTINUE
      RETURN
      END
```

## Example

Scatter **x** into **y**, where **x** is a sparse vector with interesting elements  $x_1$ ,  $x_4$ ,  $x_5$ , and  $x_9$  stored in one-dimensional array **X**, and **y** is stored in a one-dimensional array **Y** of dimension 20.

```
      INTEGER*8 M,INDX(4)
      REAL*8     Y(20),X(4)
      DATA      INDX / 1, 4, 5, 9 /
      M = 4
      CALL SCATTER (M,Y,INDX,X)
```

Basic Vector Operations  
SCOPY/CCOPY – Copy Vector

**NAME** SCOPY/CCOPY – Copy Vector

**Purpose**

Given real, integer, or complex vectors  $x$  and  $y$  of length  $n$ , these subprograms perform the vector copy operations

$$y \leftarrow x$$

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays. Indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```
INTEGER*8      n, incx, incy
REAL*8         x(lenx), y(leny)
CALL SCOPY(n, x, incx, y, incy)

INTEGER*8      n, incx, incy
COMPLEX*16     x(lenx), y(leny)
CALL CCOPY(n, x, incx, y, incy)
```

**Input**

**n** Number of elements of vectors  $x$  and  $y$  to be used in the copy operation. If  $n \leq 0$ , the subprograms do not reference  $x$  or  $y$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ . Refer to "Purpose."

**incx** Increment for the array  $x$ :

$\text{incx} \geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

$\text{incx} < 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "Notes" for use of  $\text{incx} = 0$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

**incy** Increment for the array  $y$ ,  $\text{incy} \neq 0$ :

$\text{incy} > 0$   $y$  is stored forward in array  $y$ , i.e.,  $y_i$  is stored in  $y((i-1) \times \text{incy} + 1)$ .

**incy** < 0      *y* is stored backward in array *y*, i.e.,  $y_i$  is stored in  $y((i-n) \times \text{incy} + 1)$ .

Use **incy** = 1 if the vector *y* is stored contiguously in array *y*, i.e., if  $y_i$  is stored in  $y(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

## Output

**y**                      Array of length **leny** =  $(n-1) \times |\text{incy}| + 1$  containing the *n*-vector *y*. If **n** ≤ 0, then *y* is unchanged. Otherwise,  $y \leftarrow x$ .

## Notes

If **incx** = 0, then  $x_i = x(1)$  for all *i*. This can be used to initialize all elements of *y* to a constant. Refer to “Example 2.”

The result is unspecified if *x* and *y* overlap such that any element of *x* shares a memory location with any element of *y*.

## Fortran Equivalent

```

SUBROUTINE SCOPY (N, X, INCX, Y, INCY)
  INTEGER*8 N, INCX, INCY
  REAL*8 X(*), Y(*)
  IF ( N .LE. 0 ) RETURN
  IX = 1
  IY = 1
  IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
  IF ( INCY .LT. 0 ) IY = 1 - (N-1) * INCY
  DO 10 I = 1, N
    Y(IY) = X(IX)
    IX = IX + INCX
    IY = IY + INCY
  10 CONTINUE
  RETURN
END

```

## Example 1

Copy the REAL\*8 vector *x* into *y*, where *x* and *y* are vectors 10 elements long, stored in one-dimensional arrays X and Y of dimension 20.

```

INTEGER*8 N, INCX, INCY
REAL*8 X(20), Y(20)
N = 10
INCX = 1
INCY = 1
CALL SCOPY (N, X, INCX, Y, INCY)

```

Basic Vector Operations  
**SCOPY/CCOPY – Copy Vector**

**Example 2**

Initialize a one-dimensional array to zero.

```
INTEGER*8 N, INCX, INCY
REAL*8    Y(20)
N = 10
INCX = 0
INCY = 1
CALL SCOPY (N, 0.0, INCX, Y, INCY)
```

**NAME** SDOT/CDOTC/CDOTU – Dot Product

**Purpose**

Given real or complex data vectors  $x$  and  $y$  of length  $n$ , these subprograms compute the dot products

$$s = \sum_{i=1}^n x_i y_i \quad \text{and} \quad s = \sum_{i=1}^n \bar{x}_i y_i$$

where  $\bar{x}$  is the complex conjugate of  $x$ . The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays. Indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

INTEGER*8      n, incx, incy
REAL*8        s, SDOT, x(lenx), y(leny)
s = SDOT(n, x, incx, y, incy)

INTEGER*8      n, incx, incy
COMPLEX*16    s, CDOTC, x(lenx), y(leny)
s = CDOTC(n, x, incx, y, incy)

INTEGER*8      n, incx, incy
COMPLEX*16    s, CDOTU, x(lenx), y(leny)
s = CDOTU(n, x, incx, y, incy)

```

**Input**

- n**                    Number of elements of vectors  $x$  and  $y$  to be used in the dot product. If  $n \leq 0$ , the subprograms do not reference  $x$  or  $y$ .
- x**                    Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .  $x$  is used in unconjugated form by the subprograms. Refer to "Purpose."
- incx**                Increment for the array  $x$ :
  - incx**  $\geq 0$          $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .
  - incx**  $< 0$          $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

Basic Vector Operations  
SDOT/CDOTC/CDOTU – Dot Product

**y**                    Array of length  $\text{leny} = (n-1) \times |\text{incy}| + 1$  containing the  $n$ -vector  $y$ .

**incy**                Increment for the array  $y$ :

**incy**  $\geq 0$          $y$  is stored forward in array  $y$ , i.e.,  $y_i$  is stored in  $y((i-1) \times \text{incy} + 1)$ .

**incy**  $< 0$          $y$  is stored backward in array  $y$ , i.e.,  $y_i$  is stored in  $y((i-n) \times \text{incy} + 1)$ .

Use **incy** = 1 if the vector  $y$  is stored contiguously in array  $y$ , i.e., if  $y_i$  is stored in  $y(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

**Output**

**s**                    The resulting value of the dot product. If  $n \leq 0$ , then  $s = 0$ . Otherwise,

$$s = \sum_{i=1}^n x_i y_i$$

unless the subprogram name is CDOTC, in which case

$$s = \sum_{i=1}^n \bar{x}_i y_i.$$

**Notes**

If **incx** = 0, then  $x_i = \mathbf{x}(1)$  for all  $i$ . If **incy** = 0, then  $y_i = \mathbf{y}(1)$  for all  $i$ . In either of these cases, another SCILIB subprogram would be more efficient.

## Fortran Equivalent

```

REAL*8 FUNCTION SDOT (N, X, INCX, Y, INCY)
INTEGER*8 N, INCX, INCY
REAL*8 X(*), Y(*)
SDOT = 0.0
IF ( N .LE. 0 ) RETURN
IX = 1
IY = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
IF ( INCY .LT. 0 ) IY = 1 - (N-1) * INCY
DO 10 I = 1, N
    SDOT = SDOT + X(IX) * Y(IY)
    IX = IX + INCX
    IY = IY + INCY
10 CONTINUE
RETURN
END

```

## Example 1

Compute the REAL\*8 dot product

$$s = \sum_{i=1}^{10} x_i y_i,$$

where  $x$  and  $y$  are vectors 10 elements long stored in one-dimensional arrays  $X$  and  $Y$  of dimension 20.

```

INTEGER*8 N, INCX, INCY
REAL*8 S, SDOT, X(20), Y(20)
N = 10
INCX = 1
INCY = 1
S = SDOT (N, X, INCX, Y, INCY)

```

## Example 2

Compute the REAL\*8 dot product

$$s = \sum_{i=1}^{10} x_i y_i,$$

where  $x$  is the 4th row of a 10-by-10 matrix stored in a two-dimensional array  $X$  of dimension 20 by 21, and  $y$  is a vector 10 elements long stored in one-dimensional array  $Y$  of dimension 20.

```

INTEGER*8 N, INCX, INCY
REAL*8 S, SDOT, X(20,21), Y(20)
N = 10
S = SDOT (N, X(4,1), 20, Y, 1)

```

Basic Vector Operations  
SNRM2/SCNRM2 – Euclidean Norm

**NAME** SNRM2/SCNRM2 – Euclidean Norm

**Purpose**

Given a real, integer, or complex vector  $x$  of length  $n$ , these subprograms compute the Euclidean (i.e.,  $l_2$ ) norm of the vector

$$s = \|x\|_2 = \left\{ \sum_{i=1}^n |x_i|^2 \right\}^{1/2}.$$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```
INTEGER*8      n, incx
REAL*8         s, SNRM2, x(lenx)
s = SNRM2(n, x, incx)

INTEGER*8      n, incx
REAL*8         s, SCNRM2
COMPLEX*16     x(lenx)
s = SCNRM2(n, x, incx)
```

**Input**

**n** Number of elements of vector  $x$  to be used in the Euclidean norm. If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx** Increment for the array  $x$ .  $x$  is stored forward in array  $x$  with increment  $|\text{incx}|$ , i.e.,  $x_i$  is stored in  $x((i-1) \times |\text{incx}| + 1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

**Output**

**s** If  $n \leq 0$ , then  $s = 0$ . Otherwise,  $s$  is the Euclidean norm of  $x$ .

## Fortran Equivalent

```

REAL*8 FUNCTION SNRM2 ( N, X, INCX )
INTEGER*8 INCX, INCXA, IX, N
REAL*8 ABSXI, SCALE, SSQ, X(*)
IF ( N .GT. 1 ) THEN
    INCXA = ABS ( INCX )
    SCALE = 0.0
    SSQ = 1.0
    DO 10 IX = 1, 1 + (N-1)*INCA, INCA
        IF ( X(IX) .NE.0.0 ) THEN
            ABSXI = ABS ( X(IX) )
            IF ( SCALE .LT. ABSXI ) THEN
                SSQ = 1.0 + SSQ * (SCALE/ABSXI) ** 2
                SCALE = ABSXI
            ELSE
                SSQ = SSQ + (ABSXI/SCALE) ** 2
            END IF
        END IF
    END DO
10 CONTINUE
    SNRM2 = SCALE * SQRT ( SSQ )
ELSE IF ( N .EQ. 1 ) THEN
    SNRM2 = ABS ( X(1) )
ELSE
    SNRM2 = 0.0
END IF
RETURN
END

```

## Example

Compute the Euclidean norm of the REAL\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```

INTEGER*8 N, INCX
REAL*8 S, SNRM2, X(20)
N = 10
INCX = 1
S = SNRM2 (N, X, INCX)

```

**SPAXPY – Sparse Elementary Vector Operation****NAME** SPAXPY – Sparse Elementary Vector Operation**Purpose**

Given a real scalar  $a$ , a sparse vector  $x$  stored in compact form via a set of indices, and a dense vector  $y$  stored in full storage form, this subprogram performs the elementary vector operation

$$y \leftarrow ax+y.$$

More precisely, let  $x$  be a sparse  $n$ -vector with  $m \leq n$  interesting (usually nonzero) elements, and let  $\{k_1, k_2, \dots, k_m\}$  be the indices of these elements. All *uninteresting* elements of  $x$  are assumed to be zero. Let  $y$  be an ordinary  $n$ -vector. If  $x$  is represented by arrays  $\mathbf{x}$  and  $\mathbf{indx}$  such that  $\mathbf{indx}(i) = k_i$  and  $\mathbf{x}(i) = x_{k_i}$  then these subprograms compute

$$y_{k_i} \leftarrow a\mathbf{x}_i + y_{k_i}, \quad i = 1, 2, \dots, m.$$

**Usage**

SCILIB:

```

INTEGER*8      m, indx(m)
REAL*8        a, x(m), y(n)
CALL SPAXPY(m, a, x, y, indx)

```

**Input**

- m** Number of interesting elements of  $x$ ,  $\mathbf{m} \leq \mathbf{n}$ , where  $\mathbf{n}$  is the length of  $y$ . If  $\mathbf{m} \leq 0$ , the subprogram does not reference  $\mathbf{x}$ ,  $\mathbf{indx}$ , or  $y$ .
- a** The scalar  $a$ .
- x** Array of length  $\mathbf{m}$  containing the interesting elements of  $x$ .  $\mathbf{x}(j) = x_i$  if  $\mathbf{indx}(j) = i$ .
- y** Array containing the elements of  $y$ ,  $\mathbf{y}(i) = y_i$ .
- indx** Array containing the indices  $\{k_i\}$  of the interesting elements of  $x$ . The indices must satisfy
- $$1 \leq \mathbf{indx}(i) \leq \mathbf{n}, \quad i = 1, 2, \dots, \mathbf{m}$$
- and
- $$\mathbf{indx}(i) \neq \mathbf{indx}(j), \quad 1 \leq i \neq j \leq \mathbf{m},$$
- where  $\mathbf{n}$  is the length of  $y$ .

## Output

$y$  If  $m \leq 0$  or  $a = 0$ , then  $y$  is unchanged. Otherwise,  $ax+y$  overwrites the input. Only the elements of  $y$  whose indices are included in  $indx$  are changed.

## Notes

The result is unspecified if any element of  $indx$  is out of range, if any two elements of  $indx$  have the same value, or if  $x$ ,  $indx$ , and  $y$  overlap such that any element of  $x$  or any index shares a memory location with any element of  $y$ .

## Fortran Equivalent

```
SUBROUTINE SPAXPY (M, A, X, Y, INDX)
  INTEGER*8 M, INDX(*)
  REAL*8 A, X(*), Y(*)
  IF ( A .EQ. 0.0 ) RETURN
  DO 10 I = 1, M
    Y(INDX(I)) = A * X(I) + Y(INDX(I))
  10 CONTINUE
  RETURN
END
```

## Example

Compute the REAL\*8 elementary vector operation

$$y \leftarrow 2x+y,$$

where  $x$  is a sparse vector with interesting elements  $x_1, x_4, x_5$ , and  $x_9$  stored in one-dimensional array  $X$ , and  $y$  is stored in a one-dimensional array  $Y$  of dimension 20.

```
INTEGER*8 M, INDX(4)
REAL*8    A, X(4), Y(20)
DATA      INDX / 1, 4, 5, 9 /
M = 4
A = 2.0
CALL SPAXPY (M, A, X, INDX, Y)
```

Basic Vector Operations  
SPDOT – Sparse Dot Product

**NAME** SPDOT – Sparse Dot Product

**Purpose**

Given a real sparse vector  $x$  stored in compact form via an index vector, and a dense vector  $y$  stored in full storage form, this subprogram computes the sparse dot product

$$s = \sum_{i=1}^n x_i y_i$$

More precisely, let  $x$  be a sparse  $n$ -vector with  $m \leq n$  interesting (usually nonzero) elements, let  $\{k_1, k_2, \dots, k_m\}$  be the indices of these elements. (While some interesting elements of  $x$  may be zero, all *uninteresting* elements are assumed to be zero.) Let  $y$  be an ordinary  $n$ -vector. If  $x$  is represented by arrays  $\mathbf{x}$  and  $\mathbf{indx}$  such that  $\mathbf{indx}(i) = k_i$  and  $\mathbf{x}(i) = x_{k_i}$ , then these subprograms compute

$$s = \sum_{i=1}^m \mathbf{x}_i y_{k_i}$$

**Usage**

SCILIB:

```
INTEGER*8      m, indx(m)
REAL*8        s, SPDOT, y(n), x(m)
s = SPDOT(m, y, indx, x)
```

**Input**

**m** Number of interesting elements of  $x$ ,  $m \leq n$ . If  $m \leq 0$ , the subprogram does not reference  $\mathbf{x}$ ,  $\mathbf{indx}$ , or  $\mathbf{y}$ .

**y** Array containing the elements of  $y$ ,  $\mathbf{y}(i) = y_i$ .

**indx** Array containing the indices  $\{k_i\}$  of the interesting elements of  $x$ . The indices must satisfy  
 $1 \leq \mathbf{indx}(i) \leq n$ ,  $i = 1, 2, \dots, m$ ,  
where  $n$  is the length of  $\mathbf{y}$ .

**x** Array of length  $m$  containing the interesting elements of  $x$ .

**Output**

**s** The resulting value of the dot product. If  $m \leq 0$ , then  $s = 0$ . Otherwise,

$$s = \sum_{i=1}^m x(i) \times y(\text{indx}(i)).$$

### Fortran Equivalent

```
REAL*8 FUNCTION SPDOT (M, Y, INDX, X)
INTEGER*8 M, INDX(*)
REAL*8 Y(*), X(*)
SPDOT = 0.0
DO 10 I = 1, M
    SPDOT = SPDOT + X(I) * Y(INDX(I))
10 CONTINUE
RETURN
END
```

### Example

Compute the REAL\*8 sparse dot product

$$s = \sum_{i=1}^{10} x_i y_i,$$

where  $x$  is a sparse vector with interesting elements  $x_1, x_4, x_5,$  and  $x_9$  stored in one-dimensional array  $X$ , and  $y$  is a vector 10 elements long stored in a one-dimensional array  $Y$  of dimension 20.

```
INTEGER*8 M, INDX(4)
REAL*8 S, SPDOT, Y(20), X(4)
DATA INDX / 1, 4, 5, 9 /
M = 4
S = SPDOT (M, Y, INDX, X)
```

Basic Vector Operations  
**SROT/CROT – Apply Givens Rotation**

**NAME**           SROT/CROT – Apply Givens Rotation

**Purpose**

Given a real scalar  $c$ , a real or complex scalar,  $s$  and real or complex vectors  $x$  and  $y$  of length  $n$ , these subprograms apply the Givens rotation

$$\begin{bmatrix} x_i \\ y_i \end{bmatrix} \leftarrow \begin{bmatrix} c & s \\ -\bar{s} & c \end{bmatrix} \cdot \begin{bmatrix} x_i \\ y_i \end{bmatrix} \quad \text{for } i = 1, \dots, n$$

where  $\bar{s}$  is the complex conjugate of  $s$ ;  $\bar{s} = s$  if  $s$  is real. The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

Usually,  $c$  and  $s$  have been determined by the companion subprogram SROTG, or CROTG.

**Usage**

SCILIB:

```

INTEGER*8            n, incx, incy
REAL*8               x(lenx), y(leny), c, s
CALL SROT(n, x, incx, y, incy, c, s)
INTEGER*8            n, incx, incy
REAL*8               c
COMPLEX*16          x(lenx), y(leny), s
CALL CROT(n, x, incx, y, incy, c, s)

```

**Input**

**n**                    Number of elements of vectors  $x$  and  $y$  to be used in the Givens rotation. If  $n \leq 0$ , the subprograms do not reference  $x$  or  $y$ .

**x**                    Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx**                Increment for the array  $x$ ,  $\text{incx} \neq 0$ :

**incx > 0**         $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx < 0**         $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $\mathbf{x}$ , i.e., if  $x_i$  is stored in  $\mathbf{x}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

**y**                    Array of length  $\text{leny} = (n-1) \times |\text{incy}| + 1$  containing the  $n$ -vector  $y$ .

**incy**                Increment for the array  $\mathbf{y}$ ,  $\text{incy} \neq 0$ :

$\text{incy} > 0$        $y$  is stored forward in array  $\mathbf{y}$ , i.e.,  $y_i$  is stored in  $\mathbf{y}((i-1) \times \text{incy} + 1)$ .

$\text{incy} < 0$        $y$  is stored backward in array  $\mathbf{y}$ , i.e.,  $y_i$  is stored in  $\mathbf{y}((i-n) \times \text{incy} + 1)$ .

Use  $\text{incy} = 1$  if the vector  $y$  is stored contiguously in array  $\mathbf{y}$ , i.e., if  $y_i$  is stored in  $\mathbf{y}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

**c**                    The scalar  $c$ .

**s**                    The scalar  $s$ .

## Output

**x and y**            If  $n \leq 0$  or if  $c = 1$  and  $s = 0$ , then  $\mathbf{x}$  and  $\mathbf{y}$  are unchanged. Otherwise, the resulting vectors overwrite the input.

## Notes

The result is unspecified if  $\text{incx} = 0$  or  $\text{incy} = 0$  or if  $\mathbf{x}$  and  $\mathbf{y}$  overlap such that any element of  $x$  shares a memory location with any element of  $y$ .

SCILIB also contains subprograms that construct and apply modified Givens rotations. They are documented elsewhere in this chapter. The modified Givens subprograms are a little more difficult to use, but are more efficient.

Basic Vector Operations  
**SROT/CROT – Apply Givens Rotation**

**Fortran Equivalent**

```
SUBROUTINE SROT (N, X, INCX, Y, INCY, C, S)
REAL*8 C, S, TEMP, X(*), Y(*)
INTEGER*8 N, INCX, INCY
IF ( N .LE. 0 ) RETURN
IF ( C .EQ. 1.0 .AND. S .EQ. 0.0 ) RETURN
IX = 1
IY = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
IF ( INCY .LT. 0 ) IY = 1 - (N-1) * INCY
DO 10 I = 1, N
    TEMP = C * X(IX) + S * Y(IY)
    Y(IY) = C * Y(IY) - S * X(IX)
    X(IX) = TEMP
    IX = IX + INCX
    IY = IY + INCY
10 CONTINUE
RETURN
END
```

**Example 1**

Apply a Givens rotation to  $x$  and  $y$ , vectors 10 elements long stored in one-dimensional arrays  $X$  and  $Y$  of dimension 20.

```
INTEGER*8 N, INCX, INCY
REAL*8    X(20), Y(20), C, S
N = 10
INCX = 1
INCY = 1
CALL SROT (N, X, INCX, Y, INCY, C, S)
```

**Example 2**

Reduce 10-by-10 matrix  $a$  stored in two-dimensional array  $A$  of dimension 20 by 21 to upper-triangular form via Givens rotations (compare with “Example 2” in the description of SROTM).

```
INTEGER*8 INCA, I, J, N
REAL*8    A(20,21), C, S
INCA = 20
DO 20 I = 1, 9
    N = 10 - I
    DO 10 J = I+1, 10
        CALL SROTG (A(I, I), A(J, I), C, S)
        CALL SROT (N, A(I, I+1), INCA, A(J, I+1), INCA, C, S)
10 CONTINUE
20 CONTINUE
```

**NAME** SROTG/CROTG – Construct Givens Rotation

**Purpose**

Given real or complex scalars  $a$  and  $b$ , these subprograms construct a Givens plane rotation matrix that annihilates  $b$ . Specifically, they determine scalars  $c$  and  $s$  such that

$$\begin{bmatrix} c & s \\ -\bar{s} & c \end{bmatrix} \cdot \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} r \\ 0 \end{bmatrix}$$

where  $c$  is real,  $r$  and  $s$  are of the same type as  $a$  and  $b$ , and  $\bar{s}$  is the complex conjugate of  $s$ .

Usually,  $c$  and  $s$  are passed to companion subprogram SROT, or CROT to apply the Givens rotation to a pair of vectors.

SROTG also determines a quantity  $z$  that permits the later stable reconstruction of  $c$  and  $s$  from a single quantity.

**Usage**

SCILIB:

```
REAL*8          a, b, c, s
CALL SROTG(a, b, c, s)
REAL*8          c
COMPLEX*16      a, b, s
CALL CROTG(a, b, c, s)
```

**Input**

**a** The scalar  $a$ .  
**b** The scalar  $b$ .

**Output**

**a** The rotated result  $r$  overwrites  $a$ .  
**b** Not used as output by CROTG. In SROTG, the reconstruction quantity  $z$  overwrites  $b$ . The reconstruction quantity  $z$  is useful if a matrix is being transformed by a sequence of Givens rotations that must be saved to be applied again. Since each  $z$  overwrites an element that has been reduced to zero, the transformations can be saved without using any additional storage.

**SROTG/CROTG – Construct Givens Rotation**

The quantities  $c$  and  $s$  may be reconstructed from  $z$  as follows:

if  $|z| = 0$ ,                    set  $c = 0$  and  $s = 1$ .

if  $|z| < 0$ ,                  set  $c = \sqrt{(1-z^2)}$  and  $s = z$ .

if  $|z| > 0$ ,                  set  $c = 1/z$  and  $s = \sqrt{(1-c^2)}$ .

**c**                    The rotation scalar  $c$ .  
**s**                    The rotation scalar  $s$ .

**Notes**

SCILIB also contains subprograms that construct and apply modified Givens rotations. They are documented elsewhere in this chapter. The modified Givens subprograms are a little more difficult to use, but are more efficient.

**Example**

Construct a Givens plane rotation that will rotate vectors  $x$  and  $y$  in such a way as to annihilate  $y_1$ .  $x$  and  $y$  are vectors 10 elements long stored in one-dimensional arrays  $X$  and  $Y$  of dimension 20.

```
REAL*8 X(20),Y(20),C,S
CALL SROTG (X(1),Y(1),C,S)
```

$X(1)$  is the rotated result and  $Y(1)$  is the reconstruction quantity, so these elements should not be rotated by a subsequent call to SROT.

## SROTM – Apply Modified Givens Rotation

**NAME** SROTM – Apply Modified Givens Rotation

**Purpose**

Given a modified Givens rotation matrix  $H = \{h_{ij}\}$  as constructed by SROTM, and real vectors  $x$  and  $y$  of length  $n$ , these subprograms apply the modified rotation

$$\begin{bmatrix} x_i \\ y_i \end{bmatrix} \leftarrow \begin{bmatrix} h_{11} & h_{12} \\ h_{21} & h_{22} \end{bmatrix} \cdot \begin{bmatrix} x_i \\ y_i \end{bmatrix} \quad \text{for } i = 1, \dots, n.$$

Refer to the description of the companion subprogram SROTMG for more details about the modified Givens rotation.

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

INTEGER*8      n, incx, incy
REAL*8        x(lenx), y(leny), param(5)
CALL SROTM(n, x, incx, y, incy, param)

```

**Input**

- n** Number of elements of vectors  $x$  and  $y$  to be used. If  $n \leq 0$ , the subprograms do not reference  $x$  or  $y$ .
- x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .
- incx** Increment for the array  $x$ ,  $\text{incx} \neq 0$ :
- incx** > 0  $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .
- incx** < 0  $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .
- Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.
- y** Array of length  $\text{leny} = (n-1) \times |\text{incy}| + 1$  containing the  $n$ -vector  $y$ .

**SROTM – Apply Modified Givens Rotation**

<b>incy</b>	<p>Increment for the array <b>y</b>, <b>incy</b> <math>\neq</math> 0:</p> <p><b>incy</b> &gt; 0      <b>y</b> is stored forward in array <b>y</b>, i.e., <math>y_i</math> is stored in <math>y((i-1) \times \text{incy} + 1)</math>.</p> <p><b>incy</b> &lt; 0      <b>y</b> is stored backward in array <b>y</b>, i.e., <math>y_i</math> is stored in <math>y((i-n) \times \text{incy} + 1)</math>.</p> <p>Use <b>incy</b> = 1 if the vector <b>y</b> is stored contiguously in array <b>y</b>, i.e., if <math>y_i</math> is stored in <math>y(i)</math>. Refer to “BLAS Indexing Conventions” in the introduction to this chapter.</p>
<b>param</b>	<p>Array containing the matrix elements of the modified Givens rotation matrix <b>H</b> and a flag indicating which form the rotation matrix takes, and therefore which of the elements of <b>param</b> are significant. <b>param</b> will usually have been set by the companion subprogram SROTMG; refer to the description of this companion subprogram for the specific contents of <b>param</b>.</p>

**Output**

<b>x and y</b>	<p>If <math>n \leq 0</math> or if <b>param</b>(1) = -2, <b>x</b> and <b>y</b> are unchanged. Otherwise, the resulting vectors overwrite the input.</p>
----------------	--

**Notes**

The result is unspecified if **incx** = 0 or **incy** = 0 or if **x** and **y** overlap such that any element of **x** shares a memory location with any element of **y**.

SCILIB also contains subprograms that construct and apply regular Givens rotations. They are documented elsewhere in this chapter. The modified Givens subprograms are a little more difficult to use, but are more efficient.

**Example 1**

Apply a modified Givens rotation to **x** and **y**, vectors 10 elements long stored in one-dimensional arrays **X** and **Y** of dimension 20.

```

INTEGER*8 N, INCX, INCY
REAL*8    X(20), Y(20), PARAM(5)
N = 10
INCX = 1
INCY = 1
CALL DROTM (N, X, INCX, Y, INCY, PARAM)

```

## Example 2

Reduce 10-by-10 matrix **a** stored in two-dimensional array **A** of dimension 20 by 21 to upper-triangular form via modified Givens rotations (compare with “Example 2” in the description of SROT.)

```
      INTEGER*8 INCA, I, J, N
      REAL*8    A(20,21), D(20), PARAM(5)
      INCA = 20
      DO 10 I = 1, 10
         D(I) = 1.0
10     CONTINUE
      DO 30 I = 1, 9
         N = 10 - I
         DO 20 J = I+1, 10
            CALL SROTMG (D(I), D(J), A(I, I), A(J, I), PARAM)
            CALL SROTM  (N, A(I, I+1), INCA, A(J, I+1), INCA, PARAM)
20     CONTINUE
30     CONTINUE
      DO 40 I = 1, 10
         N = 11 - I
         CALL SSCAL  (N, SQRT(D(I)), A(I, I), INCA)
40     CONTINUE
```

**NAME** SROTMG – Construct Modified Givens Rotation

**Purpose**

The Givens rotation,  $G$ , that annihilates  $z_1$ , if  $z_1 \neq 0$ , is

$$GW = \begin{bmatrix} c & s \\ -s & c \end{bmatrix} \cdot \begin{bmatrix} w_1 & \dots & w_n \\ z_1 & \dots & z_n \end{bmatrix},$$

where  $c = w_1/r$ ,  $s = z_1/r$ , and  $r = \pm(w_1^2 + z_1^2)^{1/2}$ . Computing  $G$  and applying it to a pair of  $n$  vectors requires  $\sim 4n$  floating-point multiplications,  $\sim 2n$  floating-point additions, and one square root.

The modified Givens rotation is a device for reducing this operation count. Suppose that  $W$  above is available in factored form

$$W = D^{1/2}X \equiv \begin{bmatrix} d_1^{1/2} & 0 \\ 0 & d_2^{1/2} \end{bmatrix} \cdot \begin{bmatrix} x_1 & \dots & x_n \\ y_1 & \dots & y_n \end{bmatrix}.$$

These subprograms construct  $\bar{d}_1$ ,  $\bar{d}_2$ , and  $H$  such that  $GW$  is obtained in the same factored form in which  $W$  was given

$$GW = \begin{bmatrix} \bar{d}_1^{1/2} & 0 \\ 0 & \bar{d}_2^{1/2} \end{bmatrix} \cdot \begin{bmatrix} h_{11} & h_{12} \\ h_{21} & h_{22} \end{bmatrix} \cdot \begin{bmatrix} x_1 & \dots & x_n \\ y_1 & \dots & y_n \end{bmatrix}.$$

$H$  is chosen to have the same numerical stability as the standard Givens rotation but better computational efficiency. Thus,  $H$  will usually have two elements equal to  $\pm 1$ . When this is true, computing  $H$  and applying it to a pair of  $n$ -vectors requires  $\sim 2n$  floating-point multiplications,  $\sim 2n$  floating-point additions, and no square roots. Companion SCILIB subprograms SROTMG are provided to apply the modified Givens rotation to a pair of vectors.

In most applications,  $d_1$  and  $d_2$  are initially set to 1, manipulated by SROTMG as the modified Givens rotations are constructed, and then applied to the vectors as the final step in the computation. For example, the reduction of an  $n$ -by- $n$  matrix to upper-triangular form via modified Givens rotations requires  $O(n)$  square roots compared to the  $O(n^2)$  required by ordinary Givens rotations. Refer to "Example 2" in the description of SROTMG.

**Usage**

SCILIB:

```

REAL*8          d1, d2, x1, y1, param(5)
CALL SROTMG(d1, d2, x1, y1, param)

```

**Input**

**d1**            The scale factor for the “x” row.  
**d2**            The scale factor for the “y” row.  
**x1**            The first element of the “x” row.  
**y1**            The first element of the “y” row. This is the element that will be annihilated by the rotation.

**Output**

**d1**            The updated scale factor for the “x” row.  
**d2**            The updated scale factor for the “y” row.  
**x1**            The rotated first element of the “x” row.  
**param**        Array containing the matrix elements of the modified Givens rotation matrix  $H$  and a flag indicating which form the rotation matrix  $H$  takes and, therefore, which elements of **param** are significant. **param** will usually be an argument to the companion subprogram SROTM.

**param**(1) specifies the form of the rotation matrix  $H$ , as follows:

**param**(1) = -2

$$H = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

**param**(1) = -1

$$H = \begin{bmatrix} \text{param}(2) & \text{param}(4) \\ \text{param}(3) & \text{param}(5) \end{bmatrix}$$

**param**(1) = 0

$$H = \begin{bmatrix} 1 & \text{param}(4) \\ \text{param}(3) & 1 \end{bmatrix}$$

**param**(1) = 1

$$H = \begin{bmatrix} \text{param}(2) & 1 \\ -1 & \text{param}(5) \end{bmatrix}$$

## Basic Vector Operations

### SROTMG – Construct Modified Givens Rotation

For each of the four values of **param**(1), only the indicated values of **param**(2) through **param**(5) are defined. The 0, 1, and -1 elements are not stored in **param**.

## Notes

SCILIB also contains subprograms that construct and apply ordinary Givens rotations. They are documented elsewhere in this chapter. The modified Givens subprograms are a little more difficult to use, but are more efficient.

## Example

Construct a modified Givens plane rotation that will rotate vectors  $d_1x$  and  $d_2y$  in such a way as to annihilate  $d_2y_1$ .  $x$  and  $y$  are vectors 10 elements long stored in one-dimensional arrays X and Y of dimension 20.

```
REAL*8 D1,D2,X(20),Y(20),PARAM(5)
CALL SROTMG (D1,D2,X(1),Y(1),PARAM)
```

X(1) is the rotated result, so it should not be rotated by a subsequent call to SROTM.

**NAME**           SSCAL/CSCAL/CSSCAL – Scale Vector

**Purpose**

Given a real or complex scalar  $a$  and a real or complex vector  $x$  of length  $n$ , these subprograms perform the vector scaling operations

$$x \leftarrow ax \text{ and } x \leftarrow a\bar{x}$$

where  $\bar{x}$  is the complex conjugate of  $x$ . The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```

INTEGER*8      n, incx
REAL*8         a, x(lenx)
CALL SSCAL(n, a, x, incx)

INTEGER*8      n, incx
COMPLEX*16     a, x(lenx)
CALL CSCAL(n, a, x, incx)

INTEGER*8      n, incx
REAL*8         a
COMPLEX*16     x(lenx)
CALL CSSCAL(n, a, x, incx)

```

**Input**

- n**           Number of elements of vector  $x$  to be used in the scaling operation. If  $n \leq 0$ , the subprograms do not reference  $x$ .
- a**           The scalar  $a$ .
- x**           Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .  $x$  is used in unconjugated form by the subprograms. Refer to "Purpose."
- incx**       Increment for the array  $x$ ,  $\text{incx} \neq 0$ .  $x$  is stored forward in array  $x$  with increment  $|\text{incx}|$ , i.e.,  $x_i$  is stored in  $x((i-1) \times |\text{incx}| + 1)$ .  
  
Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

Basic Vector Operations  
**SSCAL/CSCAL/CSSCAL – Scale Vector**

## Output

**x**                      If  $n \leq 0$ , then **x** is unchanged. Otherwise,  $ax$  replaces the input.

## Notes

The result is unspecified if **incx** = 0.

## Fortran Equivalent

```
SUBROUTINE SSCAL (N,A, X, INCX)
REAL*8 A,X(*)
INTEGER*8 N, INCX
IF ( N .LE. 0 ) RETURN
IX = 1
INCXA = ABS ( INCX )
DO 10 I = 1, N
    X(IX) = A * X(IX)
    IX = IX + INCXA
10 CONTINUE
RETURN
END
```

## Example

Scale the REAL\*8 vector *x* by 2, where *x* is a vector 10 elements long stored in a one-dimensional array X of dimension 20.

```
INTEGER*8 N, INCX
REAL*8    A,X(20)
N = 10
INCX = 1
A = 2.0
CALL SSCAL (N,A,X, INCX)
```

**NAME** SSUM/CSUM – Vector Sum

**Purpose**

Given a real, integer, or complex vector  $x$  of length  $n$ , these subprograms compute the sum of the elements of the vector

$$s = \sum_{i=1}^n x_i.$$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

**Usage**

SCILIB:

```

INTEGER*8      n, incx
REAL*8        s, SSUM, x(lenx)
s = SSUM(n, x, incx)

INTEGER*8      n, incx
COMPLEX*16    s, CSUM, x(lenx)
s = CSUM(n, x, incx)

```

**Input**

**n** Number of elements of vector  $x$  to be used in the sum. If  $n \leq 0$ , the subprograms do not reference  $x$ .

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx** Increment for the array  $x$ .  $x$  is stored forward in array  $x$  with increment  $|\text{incx}|$ , i.e.,  $x_i$  is stored in  $x((i-1) \times |\text{incx}| + 1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.

**Output**

**s** If  $n \leq 0$ , then  $s = 0$ . Otherwise,  $s$  is the sum of the elements of  $x$ .

Basic Vector Operations  
**SSUM/CSUM – Vector Sum**

**Fortran Equivalent**

```
REAL*8 FUNCTION SSUM (N, X, INCX)
INTEGER*8 N, INCX
REAL*8 X(*)
SSUM = 0.0
IF ( N .LE. 0 ) RETURN
IX = 1
INCXA = ABS ( INCX )
DO 10 I = 1, N
    SSUM = SSUM + X(IX)
    IX = IX + INCXA
10 CONTINUE
RETURN
END
```

**Example**

Compute the sum of the elements of a REAL\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array X of dimension 20.

```
INTEGER*8 N, INCX
REAL*8 S, SSUM, X(20)
N = 10
INCX = 1
S = SSUM (N, X, INCX)
```

**NAME** SSWAP/CSWAP – Swap Two Vectors

**Purpose**

Given real, integer, or complex vectors  $x$  and  $y$  of length  $n$ , these subprograms perform the vector interchange operation

$$x \longleftrightarrow y.$$

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

INTEGER*8      n, incx, incy
REAL*8        x(lenx), y(leny)
CALL SSWAP(n, x, incx, y, incy)

INTEGER*8      n, incx, incy
COMPLEX*16    x(lenx), y(leny)
CALL CSWAP(n, x, incx, y, incy)

```

**Input**

- n**            Number of elements of vectors  $x$  and  $y$  to be used in the swap operation. If  $n \leq 0$ , the subprograms do not reference  $x$  or  $y$ .
- x**            Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .
- incx**        Increment for the array  $x$ ,  $\text{incx} \neq 0$ :
  - $\text{incx} > 0$      $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .
  - $\text{incx} < 0$      $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to this chapter.
- y**            Array of length  $\text{leny} = (n-1) \times |\text{incy}| + 1$  containing the  $n$ -vector  $y$ .
- incy**        Increment for the array  $y$ ,  $\text{incy} \neq 0$ :
  - $\text{incy} > 0$      $y$  is stored forward in array  $y$ , i.e.,  $y_i$  is stored in  $y((i-1) \times \text{incy} + 1)$ .

## Basic Vector Operations

### SSWAP/CSWAP – Swap Two Vectors

$\text{incy} < 0$   $y$  is stored backward in array  $y$ , i.e.,  $y_i$  is stored in  $y((i-n)\times\text{incy}+1)$ .

Use  $\text{incy} = 1$  if the vector  $y$  is stored contiguously in array  $y$ , i.e., if  $y_i$  is stored in  $y(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

## Output

$x$  and  $y$  If  $n \leq 0$ , then  $x$  and  $y$  are unchanged. Otherwise,  $x$  and  $y$  are interchanged in  $x$  and  $y$ .

## Notes

The result is unspecified if  $\text{incx} = 0$  or  $\text{incy} = 0$  or if  $x$  and  $y$  overlap such that any element of  $x$  shares a memory location with any element of  $y$ .

## Fortran Equivalent

```
SUBROUTINE SSWAP (N, X, INCX, Y, INCY)
REAL*8 TEMP, X(*), Y(*)
INTEGER*8 N, INCX, INCY
IF ( N .LE. 0 ) RETURN
IX = 1
IY = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
IF ( INCY .LT. 0 ) IY = 1 - (N-1) * INCY
DO 10 I = 1, N
    TEMP = X(IX)
    X(IX) = Y(IY)
    Y(IY) = TEMP
    IX = IX + INCX
    IY = IY + INCY
10 CONTINUE
RETURN
END
```

## Example 1

Interchange REAL\*8 vectors  $x$  and  $y$ , where  $x$  and  $y$  are vectors 10 elements long stored in one-dimensional arrays  $X$  and  $Y$  of dimension 20.

```
INTEGER*8 N, INCX, INCY
REAL*8    X(20), Y(20)
N = 10
INCX = 1
INCY = 1
CALL SSWAP (N, X, INCX, Y, INCY)
```

## Example 2

Interchange rows 3 and 6 of a 10-by-10 matrix a stored in two-dimensional array A of dimension 20 by 21.

```
INTEGER*8 N, INCA  
REAL*8    A(20,21)  
N = 10  
INCA = 20  
CALL SSWAP (N,A(3,1),INCA,A(6,1),INCA)
```

**NAME** WHENEQ/WHENNE/.../WHENILT – Find Selected Vector Elements

### Purpose

Given a real or integer vector  $x$  of length  $n$ , these subprograms search sequentially through the vector and fill an array with a list of the indices  $i$  for which the elements  $x_i$  satisfy a specified relationship with a given scalar  $a$ .

The last two characters of the subprogram name specify the relationship of interest between the elements of the vector and the scalar. These characters and the corresponding list contents may be

xx	List contents
EQ	$\{i : x_i = a\}$
GE	$\{i : x_i \geq a\}$
GT	$\{i : x_i > a\}$
LE	$\{i : x_i \leq a\}$
LT	$\{i : x_i < a\}$
NE	$\{i : x_i \neq a\}$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

### Usage

#### SCILIB:

```

INTEGER*8      n, incx, indx(n), nindx
REAL*8        x(lenx), a
CALL WHENEQ(n, x, incx, a, indx, nindx)

INTEGER*8      n, x(lenx), incx, a, indx(n), nindx
CALL WHENEQ(n, x, incx, a, indx, nindx)

INTEGER*8      n, incx, indx(n), nindx
REAL*8        x(lenx), a
CALL WHENNE(n, x, incx, a, indx, nindx)

INTEGER*8      n, x(lenx), incx, a, indx(n), nindx
CALL WHENNE(n, x, incx, a, indx, nindx)

INTEGER*8      n, incx, indx(n), nindx
REAL*8        x(lenx), a
CALL WHENFxx(n, x, incx, a, indx, nindx)

INTEGER*8      n, x(lenx), incx, a, indx(n), nindx
CALL WHENIxx(n, x, incx, a, indx, nindx)

```

## WHENEQ/WHENNE/.../WHENILT – Find Selected Vector Elements

## Input

<b>n</b>	Number of elements of vector $x$ to be compared to $a$ . If $n \leq 0$ , the subprograms do not reference $x$ or $\mathbf{indx}$ .
<b>x</b>	Array of length $\mathbf{lenx} = (n-1) \times  \mathbf{incx}  + 1$ containing the $n$ -vector $x$ .
<b>incx</b>	Increment for the array $x$ : $\mathbf{incx} \geq 0$ $x$ is stored forward in array $x$ , i.e., $x_i$ is stored in $x((i-1) \times \mathbf{incx} + 1)$ . $\mathbf{incx} < 0$ $x$ is stored backward in array $x$ , i.e., $x_i$ is stored in $x((i-n) \times \mathbf{incx} + 1)$ .  Use $\mathbf{incx} = 1$ if the vector $x$ is stored contiguously in array $x$ , i.e., if $x_i$ is stored in $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.
<b>a</b>	The scalar $a$ .

## Output

<b>indx</b>	Array filled with the list of indices $i$ of the elements $x_i$ of $x$ that satisfy the relationship with $a$ specified by the subprogram name. Only the first $\mathbf{nindx}$ elements of $\mathbf{indx}$ are changed.
<b>nindx</b>	If $n \leq 0$ , then $\mathbf{nindx} = 0$ . Otherwise, $\mathbf{nindx}$ is the number of elements of $x$ that satisfy the relationship with $a$ specified by the subprogram name.

## Notes

These subprograms are sometimes useful for optimizing a loop containing an **IF** statement. Refer to “Example 2.”

## Basic Vector Operations

### WHENEQ/WHENNE/.../WHENILT – Find Selected Vector Elements

#### Fortran Equivalent

```

SUBROUTINE WHENEQ (N,X, INCX,A, INDX,NINDX)
INTEGER*8 N,X(*), INCX,A, INDX(*),NINDX
IX = 1
IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
NINDX = 0
DO 10 I = 1, N
  IF ( X(IX) .EQ. A ) THEN
    NINDX = NINDX + 1
    INDX(NINDX) = I
  END IF
  IX = IX + INCX
10 CONTINUE
RETURN
END
```

#### Example 1

Find the zero elements of a REAL\*8 vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

```

INTEGER*8 N, INCX, INDX(20), NINDX
REAL*8    A, X(20)
N = 10
INCX = 1
A = 0.0
CALL WHENEQ (N,X, INCX,A, INDX,NINDX)
```

#### Example 2

Optimize the following program segment, where the THEN clause of the IF statement is much more likely than the ELSE clause.

```

INTEGER*8 I,N
REAL*8    A,B,D,DLIM,R
REAL*8    F(20000),X(20000),Y(20000),Z(20000)
N = 20000
DO 10 I = 1, N
  D = SQRT( X(I)**2 + Y(I)**2 + Z(I)**2 ) - R
  IF ( D .GT. DLIM ) THEN
    F(I) = A * EXP( B * D )
  ELSE
    CALL FORCE (D,F(I))
  END IF
10 CONTINUE
```

Change  $D$  to an array and introduce array  $INDX$  to hold the indices corresponding to the ELSE clause. Split the body of the DO loop into two parts. The first part corresponds to the body of the loop before the IF statement and the THEN clause. It fully optimizes, so even though it computes a few more exponentials than the original code, it is still considerably faster. WHENFLE is then called to determine the indices for which the ELSE clause must be executed, and the second DO loop executes the ELSE clause for those indices. The resulting program segment is

**WHENEQ/WHENNE/.../WHENILT – Find Selected Vector Elements**

```
INTEGER*8 I,J,N,INDX(20000),NINDX
REAL*8    A,B,DLIM,R
REAL*8    D(20000),F(20000),X(20000),Y(20000),Z(20000)
N = 20000
DO 10 I = 1, N
    D(I) = SQRT( X(I)**2 + Y(I)**2 + Z(I)**2 ) - R
    F(I) = A * EXP( B * D(I) )
10 CONTINUE
CALL WHENFLE (N,D,1,DLIM,INDX,NINDX)
DO 20 J = 1, NINDX
    I = INDX(J)
    CALL FORCE (D(I),F(I))
20 CONTINUE
```

**NAME** WHENMEQ/WHENMGE/.../WHENMNE – Find Selected Vector Elements

### Purpose

Given a vector  $x$  of length  $n$ , these subprograms search sequentially through the vector and fill an array with a list of the indices of the elements  $x_i$  which contain a specified group of bits that satisfy a specified relationship with a given scalar  $a$ .

The last two characters of the subprogram name specify the relationship of interest between the elements of the vector and the scalar. These characters and the corresponding list contents may be

xx	List contents
EQ	$\{i : \text{AND}(\text{SHIFTR}(x_i, \text{rshift}), \text{mask}) = a\}$
GE	$\{i : \text{AND}(\text{SHIFTR}(x_i, \text{rshift}), \text{mask}) \geq a\}$
GT	$\{i : \text{AND}(\text{SHIFTR}(x_i, \text{rshift}), \text{mask}) > a\}$
LE	$\{i : \text{AND}(\text{SHIFTR}(x_i, \text{rshift}), \text{mask}) \leq a\}$
LT	$\{i : \text{AND}(\text{SHIFTR}(x_i, \text{rshift}), \text{mask}) < a\}$
NE	$\{i : \text{AND}(\text{SHIFTR}(x_i, \text{rshift}), \text{mask}) \neq a\}$

The vector may be stored in a one-dimensional array or in either a row or a column of a two-dimensional array.

### Usage

SCILIB:

```
INTEGER*8      n, x(lenx), incx, a, indx(n), nindx, mask, rshift
CALL WHENMxx(n, x, incx, a, indx, nindx, mask, rshift)
```

### Input

<b>n</b>	Number of elements of vector $x$ to be compared to $a$ . If $n \leq 0$ , the subprograms do not reference $x$ or $\text{indx}$ .
<b>x</b>	Array of length $\text{lenx} = (n-1) \times  \text{incx}  + 1$ containing the $n$ -vector $x$ .
<b>incx</b>	Increment for the array $x$ : $\text{incx} \geq 0$ $x$ is stored forward in array $x$ , i.e., $x_i$ is stored in $x((i-1) \times \text{incx} + 1)$ . $\text{incx} < 0$ $x$ is stored backward in array $x$ , i.e., $x_i$ is stored in $x((i-n) \times \text{incx} + 1)$ .

**WHENMEQ/WHENMGE/.../WHENMNE – Find Selected Vector Elements**

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $\mathbf{x}$ , i.e., if  $x_i$  is stored in  $\mathbf{x}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to this chapter.

<b>a</b>	The scalar $a$ .
<b>mask</b>	Mask of 1-bits to extract desired group of bits from the shifted elements of $x$ with a bitwise logical product operation. Refer to “Purpose.”
<b>rshift</b>	Number of bits by which to right shift each element of $x$ so as to align the specified group of bits with <b>a</b> , $0 \leq \text{rshift} \leq 63$ . Refer to “Purpose.”

**Output**

<b>indx</b>	Array filled with the list of indices $i$ of the elements $x_i$ of $x$ that satisfy the relationship with <b>a</b> specified by the subprogram name. Only the first <b>nindx</b> elements of <b>indx</b> are changed.
<b>nindx</b>	If $n \leq 0$ , then <b>nindx</b> = 0. Otherwise, <b>nindx</b> is the number of elements of $x$ that satisfy the relationship with <b>a</b> specified by the subprogram name.

**Fortran Equivalent**

```

SUBROUTINE WHENMEQ (N,X,INCX,A,INDX,NINDX,MASK,RSHIFT)
  INTEGER*8 N,X(*),INCX,A,INDX(*),NINDX,MASK,RSHIFT
  IX = 1
  IF ( INCX .LT. 0 ) IX = 1 - (N-1) * INCX
  NINDX = 0
  DO 10 I = 1, N
    IF ( AND(SHIFTR(X(IX),RSHIFT),MASK) .EQ. A ) THEN
      NINDX = NINDX + 1
      INDX(NINDX) = I
    END IF
    IX = IX + INCX
  10 CONTINUE
  RETURN
  END

```

**Example**

Find the odd elements of an **INTEGER\*8** vector  $x$ , where  $x$  is a vector 10 elements long stored in a one-dimensional array  $X$  of dimension 20.

## Basic Vector Operations

### WHENMEQ/WHENMGE/.../WHENMNE – Find Selected Vector Elements

```
INTEGER*8 N, X(20), INCX, A, INDX(20), NINDX, MASK, RSHIFT
N = 10
INCX = 1
A = 1
MASK = 1
RSHIFT = 0
CALL WHENMEQ (N, X, INCX, A, INDX, NINDX, MASK, RSHIFT)
```

## 3 Basic Matrix Operations

---

### Overview

This chapter describes the subprograms in the Level 2 (two-loop) BLAS and the Level 3 (three-loop) BLAS. Collectively, these two sets of subprograms are called the Extended BLAS.

This chapter explains how to use the SCILIB matrix subprograms, which perform common computationally-intensive linear algebra operations. The operations covered are:

- basic matrix/vector operations
- basic matrix/matrix operations

Chapter 4 discusses matrix inverse operations.

---

### Chapter Objectives

After reading this chapter you will:

- be familiar with the Extended BLAS subroutine naming convention
- know what operations the Extended BLAS performs
- know how to use the described subprograms

## What You Need to Know to Use These Subprograms

### Subroutine Naming Convention

The Extended BLAS uses a subroutine naming convention that encodes the function of each subroutine into its name. Extended BLAS subprogram names consist of four, five, or six characters in the form TXXY, TXXYY, or TXXYYY. The first letter in the naming convention indicates one of the four Fortran data types, as shown in Table 3-1:

**Table 3-1 Extended BLAS Naming Convention — Data Type**

T	Data Type
S	Single Precision REAL
C	Single Precision COMPLEX

The next two letters in the naming convention indicate the form of the matrix, as presented in Table 3-2:

**Table 3-2 Extended BLAS Naming Convention — Matrix Form**

<b>XX</b>	Form of Matrix
GE	General
GB	General band
HE	Hermitian
HB	Hermitian band
HP	Hermitian packed
SY	Symmetric
SB	Symmetric band
SP	Symmetric packed
TR	Triangular
TB	Triangular band
TP	Triangular packed

Table 3-3 lists the final one, two, or three characters in the naming convention, indicating the computation of a particular subroutine:

**Table 3-3 Extended BLAS Naming Convention — Computation**

<b>YY</b>	Subroutine Computation
MM	Matrix-Matrix multiply
MV	Matrix-Vector multiply

R	Rank-1 update
R2	Rank-2 update
RK	Rank-k update
R2K	Rank-2k update
SM	Solve multiple systems of linear equations
SV	Solve a system of linear equations

For example, SGBMV multiplies a vector (MV) by a general band matrix (GB) using the single precision REAL data type (S). CTRSM solves a system of linear equations with one triangular coefficient matrix and a matrix of right-hand sides, using the single precision COMPLEX data type. Table 3-4 shows the valid combinations of T, XX, and Y, YY, or YYY. Each line indicates the allowable T prefixes and Y, YY, or YYY suffixes for a particular root name XX.

**Table 3-4 Extended BLAS Naming Convention — Subprogram Names**

Valid T	XX	Valid Y, YY, or YYY						
S	GE	MM	MV	R				
C	GE	MM	MV			RC	RU	
S	C	GB	MV					
C	HE	MM	MV	R	R2	RK	R2K	
C	HB		MV					
C	HP		MV	R	R2			
S	C	SY	MM	MV	R	R2	RK	R2K
C	SY	MM				RK	R2K	
S	SB		MV					
S	SP		MV	R	R2			
S	C	TR	MM	MV			SM	SV
S	C	TB		MV				SV
S	C	TP		MV				SV

## Supplemental Reading

Dongarra, J.J., J. DuCroz, S. Hammarling, and R. Hanson. "An Extended Set of Fortran Basic Linear Algebra Subprograms." *ACM Transactions on Mathematical Software*. March, 1988. Vol. 14, No. 1.

Basic Matrix Operations  
Supplemental Reading

Dongarra, J.J., J. DuCroz, S. Hammarling, and I. Duff. "A Set of Level 3 Basic Linear Algebra Subprograms." *ACM Transactions on Mathematical Software*. March, 1990. Vol. 16, No. 1.

Higham, Nicholas J. "Is Fast Matrix Multiplication of Practical Use?" *SIAM News*. November, 1990. Vol. 23, No. 6.

**NAME** MXM – Specialized Matrix-Matrix Multiply

**Purpose**

This subprogram computes the matrix-matrix product  $C = AB$ , where  $A$  is an  $m$ -by- $k$  matrix,  $B$  is a  $k$ -by- $n$  matrix, and  $C$  is an  $m$ -by- $n$  matrix. The elements of the matrices must be stored in consecutive memory locations in two-dimensional arrays of size  $m$  by  $k$ ,  $k$  by  $n$ , and  $m$  by  $n$ , respectively. SCILIB subprograms SGEMM, SGEMMS, and MXMA allow more general matrix storage and also admit the transposes of  $A$ ,  $B$ , and, in the case of MXMA,  $C$ .

**Usage**

SCILIB:

```
INTEGER*8      m, k, n
REAL*8         a(m, k), b(k, n), c(m, n)
CALL MXM(a, m, b, k, c, n)
```

**Input**

- a** Array containing the  $m$ -by- $k$  matrix  $A$ .
- m** Number of rows in matrices  $A$  and  $C$ ,  $m \geq 0$ . If  $m = 0$ , the subprogram does not reference **a**, **b**, or **c**.
- b** Array containing the  $k$ -by- $n$  matrix  $B$ .
- k** Number of columns in matrix  $A$  and number of rows in matrix  $B$ ,  $k \geq 0$ . If  $k = 0$ , the subprogram computes  $C \leftarrow 0$  without referencing **a** or **b**.
- n** Number of columns in matrices  $B$  and  $C$ ,  $n \geq 0$ . If  $n = 0$ , the subprogram does not reference **a**, **b**, or **c**.

**Output**

- c** The result  $C$  matrix.

**Notes**

Cray Research, Inc. has declared this subprogram obsolete in release 6.0 of the UNICOS Math and Scientific Library.

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are

$m < 0$ ,

Basic Matrix Operations  
**MXM – Specialized Matrix-Matrix Multiply**

**n** < 0, and  
**k** < 0.

### Fortran Equivalent

Except for the argument error checking, the following Fortran subroutine is equivalent to **MXMA**, and illustrates the meanings of the six increment arguments.

```
SUBROUTINE MXM (A,M,B,K,C,N)
  INTEGER*8 M,K,N
  REAL*8 A(M,K),B(K,N),C(M,N)
  DO 110 J = 1, N
    DO 120 I = 1, M
      C(I,J) = 0.0
110  CONTINUE
120  CONTINUE
    DO 150 J = 1, N
      DO 140 L = 1, K
        DO 130 I = 1, M
          C(I,J) = C(I,J) + A(I,L) * B(L,J)
130  CONTINUE
140  CONTINUE
150  CONTINUE
      RETURN
    END
```

### Example

Form the **REAL\*8** matrix product  $C = AB$ , where  $A$  is a 9-by-6 real matrix stored in an array  $A$  whose dimensions are 9 by 6,  $B$  is a 6 by 8 real matrix stored in an array  $B$  of dimension 6 by 8, and  $C$  is a 9 by 8 real matrix stored in an array  $C$ , also of dimension 9 by 8.

```
INTEGER*8 M,K,N
REAL*8 A(9,6),B(6,8),C(9,8)
M = 9
N = 8
K = 6
CALL MXM (A,M,B,K,C,N)
```

**NAME** MXMA – Generalized Matrix-Matrix Multiply

### Purpose

This subprogram computes the matrix-matrix product  $C = AB$ , where  $A$  is an  $m$ -by- $k$  matrix,  $B$  is a  $k$ -by- $n$  matrix, and  $C$  is an  $m$ -by- $n$  matrix. The rows and columns of the matrices may be stored with unit or non-unit strides, effectively allowing  $A$ ,  $B$ , and  $C$ , or their transposes, to be stored in two-dimensional arrays via the storage-association rules of Fortran.

SCILIB subprograms SGEMM and SGEMMS allow matrix and transposed matrix storage without resorting to storage association, also admitting the ability to add or subtract the product matrix from the original contents of the result matrix.

### Usage

SCILIB:

```

INTEGER*8      ia, ja, ib, jb, ic, jc, m, k, n
REAL*8         a(lena), b(lenb), c(lenc)
CALL MXMA(a, ia, ja, b, ib, jb, c, ic, jc, m, k, n)

```

### Input

**a** Array containing the  $m$ -by- $k$  matrix  $A$ . Typically, **a** will be a two-dimensional array with the rows and columns of  $A$  comprising one-dimensional array sections of **a**. Refer to “Notes” for suggested usages. Treating **a** as a one-dimensional array results in

$$\text{lena} = (\mathbf{m}-1) \times |\mathbf{ia}| + (\mathbf{k}-1) \times |\mathbf{ja}| + 1.$$

$A_{ij}$ ,  $1 \leq i \leq \mathbf{m}$ ,  $1 \leq j \leq \mathbf{k}$ , is stored in

$$\mathbf{a}((i-1) \times \mathbf{ia} + (j-1) \times \mathbf{ja} + 1).$$

Note that negative **ia** or **ja** will result in subscript values that lie outside the **a** array as declared above. This need not be an error; see “Example 3” for details.

**ia** Storage increment between successive elements in the same column of matrix  $A$  in array **a**. Refer to “Notes” for suggested values.

**ja** Storage increment between successive elements in the same row of matrix  $A$  in array **a**. Refer to “Notes” for suggested values.

**b** Array containing the  $k$ -by- $n$  matrix  $B$ . Typically, **b** will be a two-dimensional array with the rows and columns of  $B$

## Basic Matrix Operations

### MXMA – Generalized Matrix-Matrix Multiply

comprising one-dimensional array sections of **b**. Refer to “Notes” for suggested usages. Treating **b** as a one-dimensional array results in

$$\mathbf{lenb} = (\mathbf{k}-1) \times |\mathbf{ib}| + (\mathbf{n}-1) \times |\mathbf{jb}| + 1.$$

$B_{ij}$ ,  $1 \leq i \leq \mathbf{k}$ ,  $1 \leq j \leq \mathbf{n}$ , is stored in

$$\mathbf{b}((i-1) \times \mathbf{ib} + (j-1) \times \mathbf{jb} + 1).$$

Note that negative **ib** or **jb** will result in subscript values that lie outside the **b** array as declared above. This need not be an error; see “Example 3” for details.

- ib** Storage increment between successive elements in the same column of matrix *B* in array **b**. Refer to “Notes” for suggested values.
- jb** Storage increment between successive elements in the same row of matrix *B* in array **b**. Refer to “Notes” for suggested values.
- ic** Storage increment between successive elements in the same column of matrix *C* in array **c**. Refer to “Notes” for suggested values.
- jc** Storage increment between successive elements in the same row of matrix *C* in array **c**. Refer to “Notes” for suggested values.
- m** Number of rows in matrices *A* and *C*,  $\mathbf{m} \geq 0$ . If  $\mathbf{m} = 0$ , the subprogram does not reference **a**, **b**, or **c**.
- k** Number of columns in matrix *A* and number of rows in matrix *B*,  $\mathbf{k} \geq 0$ . If  $\mathbf{k} = 0$ , the subprogram computes  $C \leftarrow 0$  without referencing **a** or **b**.
- n** Number of columns in matrices *B* and *C*,  $\mathbf{n} \geq 0$ . If  $\mathbf{n} = 0$ , the subprogram does not reference **a**, **b**, or **c**.

## Output

- c** The result *C* matrix. Typically, **c** will be a two-dimensional array with the rows and columns of *C* comprising one-dimensional array sections of **c**. Refer to “Notes” for suggested usages. Treating **c** as a one-dimensional array results in

$$\mathbf{lenc} = (\mathbf{m}-1) \times |\mathbf{ic}| + (\mathbf{n}-1) \times |\mathbf{jc}| + 1.$$

$C_{ij}$ ,  $1 \leq i \leq \mathbf{m}$ ,  $1 \leq j \leq \mathbf{n}$ , is stored in

$$c((i-1) \times ic + (j-1) \times jc + 1).$$

Note that negative  $ic$  or  $jc$  will result in subscript values that lie outside the  $c$  array as declared above. This need not be an error; see “Example 3” for details.

## Notes

Cray Research, Inc. has declared this subprogram obsolete in release 6.0 of the UNICOS Math and Scientific Library.

Typically,  $a$ ,  $b$ , and  $c$  will be two-dimensional arrays with the rows and columns comprising one-dimensional sections of the arrays, i.e., one subscript will vary within a row or column of the matrix, and the other will be constant.

If  $a$ , for example, is a two-dimensional array of dimension  $lda$  by  $mda$ , and  $A$  is stored in untransposed form in  $a$ , then  $ia = 1$  and  $ja = lda$ , while if  $A$  is stored in transposed form ( $A^T$  is stored), then  $ia = lda$  and  $ja = 1$ .

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are

$m < 0$ ,  
 $n < 0$ ,  
 $k < 0$ ,  
 $ia = 0$ ,  
 $ja = 0$ ,  
 $ib = 0$ ,  
 $jb = 0$ ,  
 $ic = 0$ , and  
 $jc = 0$ .

## Fortran Equivalent

Except for the argument error checking, the following Fortran subroutine is equivalent to MXMA, and illustrates the meanings of the six increment arguments.

## Basic Matrix Operations

### MXMA – Generalized Matrix-Matrix Multiply

```

SUBROUTINE MXMA (A, IA, JA, B, IB, JB, C, IC, JC, M, K, N)
INTEGER*8 IA, JA, IB, JB, IC, JC, M, K, N
REAL*8 A(*), B(*), C(*)
DO 120 J = 1, N
  DO 110 I = 1, M
    C((I-1)*IC+(J-1)*JC+1) = 0.0      ! C(I,J) = 0.0
110  CONTINUE
120  CONTINUE
  DO 150 J = 1, N
    DO 140 L = 1, K
      DO 130 I = 1, M
        C((I-1)*IC+(J-1)*JC+1) =      ! C(I,J) =
1      C((I-1)*IC+(J-1)*JC+1) +      ! C(I,J) +
2      A((I-1)*IA+(L-1)*JA+1) *      ! A(I,L) *
3      B((L-1)*IB+(J-1)*JB+1)        ! B(L,J)
130  CONTINUE
140  CONTINUE
150  CONTINUE
      RETURN
END
```

#### Example 1

Form the REAL\*8 matrix product  $C = AB$ , where  $A$  is a 9-by-6 real matrix stored in an array  $A$  of dimension 10 by 11,  $B$  is a 6-by-8 real matrix stored in an array  $B$  of dimension 12 by 13, and  $C$  is a 9-by-8 real matrix stored in an array  $C$ , of dimension 14 by 15.

```

INTEGER*8 IA, JA, IB, JB, IC, JC, M, K, N
REAL*8 A(10,11), B(12,13), C(14,15)
IA = 1
JA = 10
IB = 1
JB = 12
IC = 1
JC = 14
M = 9
N = 8
K = 6
CALL MXMA (A, IA, JA, B, IB, JB, C, IC, JC, M, K, N)
```

#### Example 2

Form the REAL\*8 matrix product  $C = A^T B$ , where  $A$  is a 6 by 9 real matrix stored in an array  $A$  of dimension 10 by 11,  $B$  is a 6 by 8 real matrix stored in an array  $B$  of dimension 12 by 13, and  $C$  is a 9-by-8 real matrix stored in an array  $C$ , of dimension 14 by 15.

```
INTEGER*8 IA,JA,IB,JB,IC,JC,M,K,N
REAL*8    A(10,11),B(12,13),C(14,15)
IA = 10
JA = 1
IB = 1
JB = 12
IC = 1
JC = 14
M = 9
N = 8
K = 6
CALL MXMA (A,IA,JA, B,IB,JB, C,IC,JC, M,K,N)
```

### Example 3

Form the REAL\*8 matrix product  $C = AB$ , where  $A$  is a 9-by-6 real matrix stored “upside-down and backwards”, i.e., with the row and column subscripts decreasing to the right and bottom, in an array  $A$  of dimension 10 by 11,  $B$  is a 6-by-8 real matrix stored in an array  $B$  of dimension 12 by 13, and  $C$  is a 9-by-8 real matrix stored in an array  $C$ , of dimension 14 by 15.

```
INTEGER*8 IA,JA,IB,JB,IC,JC,M,K,N
REAL*8    A(10,11),B(12,13),C(14,15)
IA = -1
JA = -10
IB = 1
JB = 12
IC = 1
JC = 14
M = 9
N = 8
K = 6
CALL MXMA (A(M,K),IA,JA, B,IB,JB, C,IC,JC, M,K,N)
```

**MXV – Specialized Matrix-Vector Multiply****NAME** MXV – Specialized Matrix-Vector Multiply**Purpose**

This subprogram computes the matrix-vector product  $y = Ax$ , where  $A$  is an  $m$ -by- $n$  matrix,  $x$  is an  $n$ -vector, and  $y$  is an  $m$ -vector. The elements of  $A$  must be stored in consecutive memory locations in a two-dimensional array of size  $m$  by  $n$ , and the elements of the  $x$  and  $y$  must be stored in consecutive memory locations in one-dimensional arrays of size  $n$  and  $m$ , respectively. SCILIB subprograms SGEMV and MXVA allow more general storage and also admit the transpose of  $A$ . SGEMV also admits the ability to add or subtract the product from the original contents of the result vector.

**Usage**

SCILIB:

```

INTEGER*8      m, n
REAL*8        a(m, n), x(n), y(m)
CALL MXV(a, m, x, n, y)

```

**Input**

**a**            Array containing the  $m$ -by- $n$  matrix  $A$ .

**m**            Number of rows in matrix  $A$  and length of vector  $y$ ,  $m > 0$ . If  $m = 0$ , the subprogram does not reference  $a$ ,  $x$ , or  $y$ .

**x**            Array containing the  $n$ -vector  $x$ .

**n**            Number of columns in matrix  $A$  and length of vector  $x$ ,  $n > 0$ . If  $n = 0$ , the subprogram does not reference  $a$ ,  $x$ , or  $y$ .

**Output**

**y**            The result  $y$  vector.

**Notes**

Cray Research, Inc. has declared this subprogram obsolete in release 6.0 of the UNICOS Math and Scientific Library.

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are

```

m ≤ 0, and
n ≤ 0.

```

## Fortran Equivalent

Except for the argument error checking, the following Fortran subroutine is equivalent to MXVA, and illustrates the meanings of the four increment arguments.

```
      SUBROUTINE MXV (A,M,X,N,Y)
      INTEGER*8 M,N
      REAL*8 A(M,N),X(N),Y(M)
      DO 110 I = 1, M
          Y(I) = 0.0
110  CONTINUE
      DO 130 J = 1, N
          DO 120 I = 1, M
              Y(I) = Y(I) + A(I,J) * X(J)
120  CONTINUE
130  CONTINUE
      RETURN
      END
```

## Example

Form the REAL\*8 matrix-vector product  $y = Ax$ , where  $A$  is a 9 by 6 real matrix stored in an array  $A$  whose dimensions are 9 by 6,  $x$  is a real vector 6 elements long stored in an array  $X$  of dimension 6, and  $y$  is a real vector 9 elements long stored in an array  $Y$  of dimension 9.

```
      INTEGER*8 M,K,N
      REAL*8    A(9,6),X(6),Y(9)
      M = 9
      N = 6
      CALL MXV (A,M,X,N,Y)
```

Basic Matrix Operations  
MXVA – Generalized Matrix-Vector Multiply

**NAME** MXVA – Generalized Matrix-Vector Multiply

**Purpose**

This subprogram computes the matrix-vector product  $y = Ax$ , where  $A$  is an  $m$ -by- $n$  matrix,  $x$  is a  $n$ -vector, and  $y$  is an  $m$ -vector. The rows and columns of  $A$  may be stored with unit or non-unit stride, effectively allowing  $A$  or its transpose to be stored in a two-dimensional array via the storage-association rules of Fortran.

SCILIB subprogram SGEMV allows matrix and transposed matrix storage without resorting to storage association, also admitting the ability to add or subtract the product from the original contents of the result vector.

**Usage**

SCILIB:

```
INTEGER*8      ia, ja, incx, incy, m, n
REAL*8        a(lena), x(lenx), y(leny)
CALL MXVA(a, ia, ja, x, incx, y, incy, m, n)
```

**Input**

- a** Array containing the  $m$ -by- $n$  matrix  $A$ . Typically, **a** will be a two-dimensional array with the rows and columns of  $A$  comprising one-dimensional array sections of **a**. Refer to “Notes” for suggested usages. Treating **a** as a one-dimensional array results in
- $$\text{lena} = (m-1) \times |\text{ia}| + (n-1) \times |\text{ja}| + 1.$$
- $A_{ij}$ ,  $1 \leq i \leq m$ ,  $1 \leq j \leq n$ , is stored in
- $$\text{a}((i-1) \times \text{ia} + (j-1) \times \text{ja} + 1).$$
- Note that negative **ia** or **ja** will result in subscript values that lie outside the **a** array as declared above. This need not be an error; refer to “Example 3” for details.
- ia** Storage increment between successive elements in the same column of matrix  $A$  in array **a**. Refer to “Notes” for suggested values.
- ja** Storage increment between successive elements in the same row of matrix  $A$  in array **a**. Refer to “Notes” for suggested values.
- x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

<b>incx</b>	Storage increment between successive elements of vector $x$ in array $\mathbf{x}$ . $x_i$ is stored in $\mathbf{x}((i-1) \times \mathbf{incx} + 1)$ . Use $\mathbf{incx} = 1$ if the vector $x$ is stored contiguously in array $\mathbf{x}$ , i.e., if $x_i$ is stored in $\mathbf{x}(i)$ .  Note that negative $\mathbf{incx}$ will result in subscript values that lie outside the $\mathbf{x}$ array as declared above. This need not be an error; refer to "Example 3" for details.
<b>incy</b>	Storage increment between successive elements of vector $y$ in array $\mathbf{y}$ . $y_i$ is stored in $\mathbf{y}((i-1) \times \mathbf{incy} + 1)$ . Use $\mathbf{incy} = 1$ if the vector $y$ is stored contiguously in array $\mathbf{y}$ , i.e., if $y_i$ is stored in $\mathbf{y}(i)$ .  Note that a negative $\mathbf{incy}$ will result in subscript values that lie outside the $\mathbf{y}$ array as declared above. This need not be an error; refer to "Example 3" for details.
<b>m</b>	Number of rows in matrix $A$ and vector $y$ , $\mathbf{m} > 0$ .
<b>n</b>	Number of columns in matrix $A$ and length of vector $x$ , $\mathbf{n} > 0$ . $\mathbf{a}$ or $\mathbf{x}$ .

## Output

<b>y</b>	Array of length $\mathbf{leny} = (\mathbf{m}-1) \times  \mathbf{incy}  + 1$ containing the resulting $y$ vector.
----------	--

## Notes

Cray Research, Inc. has declared this subprogram obsolete in release 6.0 of the UNICOS Math and Scientific Library.

Typically,  $\mathbf{a}$  will be a two-dimensional array with rows and columns comprising one-dimensional sections of the array, i.e., one subscript will vary within a row or column of the matrix, and the other will be constant.

If  $\mathbf{a}$  is a two-dimensional array of dimension  $\mathbf{l da}$  by  $\mathbf{m da}$ , and  $A$  is stored in untransposed form in  $\mathbf{a}$ , then  $\mathbf{ia} = 1$  and  $\mathbf{ja} = \mathbf{l da}$ , while if  $A$  is stored in transposed form ( $A^T$  is stored), then  $\mathbf{ia} = \mathbf{l da}$  and  $\mathbf{ja} = 1$ .

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are

$$\begin{aligned} \mathbf{m} &\leq 0, \\ \mathbf{n} &\leq 0, \end{aligned}$$

Basic Matrix Operations  
MXVA – Generalized Matrix-Vector Multiply

**ia = 0,**  
**ja = 0,**  
**incx = 0, and**  
**incy = 0.**

### Fortran Equivalent

Except for the argument error checking, the following Fortran subroutine is equivalent to MXVA, and illustrates the meanings of the four increment arguments.

```
      SUBROUTINE MXVA (A, IA, JA, X, INCX, Y, INCY, M, N)
      INTEGER*8 IA, JA, INCX, INCY, M, N
      REAL*8 A(*), X(*), Y(*)
      DO 110 I = 1, M
          Y((I-1)*INCY+1) = 0.0           ! Y(I) = 0.0
110  CONTINUE
      DO 130 J = 1, N
          DO 120 I = 1, M
              Y((I-1)*INCY+1) =
1          Y((I-1)*INCY+1) +           ! Y(I) =
2          A((I-1)*IA+(J-1)*JA+1) *   ! Y(I) +
3          X((J-1)*INCX+1)           ! A(I, J) *
120  CONTINUE
130  CONTINUE
      RETURN
      END
```

### Example 1

Form the REAL\*8 matrix-vector product  $y = Ax$ , where  $A$  is a 9 by 6 real matrix stored in an array  $A$  of dimension 10 by 11,  $x$  is a real vector 6 elements long stored in an array  $X$  of dimension 12, and  $y$  is a real vector 9 elements long stored in an array  $Y$ , of dimension 13.

```
      INTEGER*8 IA, JA, INCX, INCY, M, N
      REAL*8     A(10, 11), B(12), Y(13)
      IA = 1
      JA = 10
      INCX = 1
      INCY = 1
      M = 9
      N = 6
      CALL MXVA (A, IA, JA, X, INCX, Y, INCY, M, N)
```

### Example 2

Form the REAL\*8 matrix-vector product  $y = A^T x$ , where  $A$  is a 6-by-9 real matrix stored in an array  $A$  of dimension 10 by 11,  $x$  is a real vector 6 elements long stored in an array  $X$  of dimension 12, and  $y$  is a real vector 9 elements long stored in an array  $Y$ , of dimension 13.

```
INTEGER*8 IA,JA, INCX, INCY, M, N
REAL*8    A(10,11), X(12), Y(13)
IA = 10
JA = 1
INCX = 1
INCY = 1
M = 9
N = 6
CALL MXVA (A, IA, JA, X, INCX, Y, INCY, M, N)
```

### Example 3

Form the REAL\*8 matrix-vector product  $y = Ax$ , where  $A$  is a 9 by 6 real matrix stored “upside-down and backwards”, i.e., with the row and column subscripts decreasing to the right and bottom, in an array  $A$  of dimension 10 by 11,  $x$  is a real vector 6 elements long stored in an array  $X$  of dimension 12, and  $y$  is a real vector 9 elements long stored in an array  $Y$ , of dimension 13.

```
INTEGER*8 IA,JA, INCX, INCY, M, N
REAL*8    A(10,11), X(12), Y(13)
IA = -1
JA = -10
INCX = 1
INCY = 1
M = 9
N = 6
CALL MXVA (A(M,K), IA, JA, X, INCX, Y, INCY, M, N)
```

Basic Matrix Operations  
**SGBMV/CGBMV – Matrix-Vector Multiply**

**NAME** SGBMV/CGBMV – Matrix-Vector Multiply

**Purpose**

These subprograms compute the matrix-vector products  $Ax$ ,  $A^T x$ , and  $A^* x$ , where  $A$  is an  $m$ -by- $n$  band matrix stored in a two-dimensional array,  $A^T$  is the transpose of  $A$ , and  $A^*$  is the conjugate transpose of  $A$ .

A band matrix is a matrix whose nonzero elements all are near the principal diagonal. Specifically,  $a_{ij} = 0$  if  $i-j > kl$  or  $j-i > ku$  for some integers  $kl$  and  $ku$ . The smallest such  $kl$  and  $ku$  for a given matrix are called the lower and upper bandwidths, respectively, and  $k = kl+ku+1$  is the total bandwidth.

The product may be stored in the result array, or optionally added to or subtracted from it. This is handled in a convenient, but general, way by two scalar arguments,  $\alpha$  and  $\beta$ , which are used as multipliers of the matrix-vector product and the result vector. Specifically, these subprograms compute matrix-vector products of the forms

$$y \leftarrow \alpha Ax + \beta y, \quad y \leftarrow \alpha A^T x + \beta y, \quad \text{and} \quad y \leftarrow \alpha A^* x + \beta y.$$

**Matrix Storage**

Because it is not necessary to store or operate on the zeros outside the band of  $A$ , you need only provide the elements within the band of  $A$ . The subprograms for general band matrices use less storage than the subprograms for general full matrices if  $kl+ku < n$ .

The following example illustrates the storage of general band matrices. Consider the following matrix  $A$  of size  $m = 9$  by  $n = 8$ , with lower and upper bandwidths  $kl = 2$  and  $ku = 3$ , respectively:

11	12	13	14	0	0	0	0
21	22	23	24	25	0	0	0
31	32	33	34	35	36	0	0
0	42	43	44	45	46	47	0
0	0	53	54	55	56	57	58
0	0	0	64	65	66	67	68
0	0	0	0	75	76	77	78
0	0	0	0	0	86	87	88
0	0	0	0	0	0	97	98

$A$  is given in an array  $ab$  with at least  $kl+ku+1 = 6$  rows and  $n = 8$  columns as follows:

*	*	*	14	25	36	47	58
*	*	13	24	35	46	57	68
*	12	23	34	45	56	67	78
11	22	33	44	55	66	77	88
21	32	43	54	65	76	87	98
31	42	53	64	75	86	97	*

The asterisks in the  $ku$ -by- $ku$  triangle at the upper left corner and in the  $(kl+n-m)$ -by- $(kl+n-m)$  triangle at the lower right corner represent elements of **ab** that are not referenced. Thus, if  $a_{ij}$  is an element within the band of  $A$ , then it is stored in  $ab(ku+1+i-j,j)$ . Therefore, the columns of  $A$  are stored in the columns of **ab**, and the diagonals of  $A$  are stored in the rows of **ab**, such that the principal diagonal is stored in row  $ku+1$  of **ab**.

## Usage

### SCILIB:

```

CHARACTER*1      trans
INTEGER*8        m, n, kl, ku, ldab, incx, incy
REAL*8           alpha, beta, ab(ldab, n), x(lenx), y(leny)
CALL SGBMV(trans, m, n, kl, ku, alpha, ab, ldab, x, incx, beta, y, incy)

CHARACTER*1      trans
INTEGER*8        m, n, kl, ku, ldab, incx, incy
COMPLEX*16       alpha, beta, ab(ldab, n), x(lenx), y(leny)
CALL CGBMV(trans, m, n, kl, ku, alpha, ab, ldab, x, incx, beta, y, incy)

```

## Input

**trans**            Transposition option for  $A$ :

    'N' or 'n'      Compute  $y \leftarrow \alpha Ax + \beta y$

    'T' or 't'      Compute  $y \leftarrow \alpha A^T x + \beta y$

    'C' or 'c'      Compute  $y \leftarrow \alpha A^* x + \beta y$

where  $A^T$  is the transpose of  $A$  and  $A^*$  is the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.

**m**                Number of rows in matrix  $A$ ,  $m \geq 0$ . If  $m = 0$ , the subprograms do not reference **ab**, **x**, or **y**.

**n**                Number of columns in matrix  $A$ ,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference **ab**, **x**, or **y**.

<b>kl</b>	The lower bandwidth of $A$ , i.e., the number of nonzero diagonals below the principal diagonal in the band, $0 \leq \mathbf{kl} < \mathbf{n}$ .						
<b>ku</b>	The upper bandwidth of $A$ , i.e., the number of nonzero diagonals above the principal diagonal in the band, $0 \leq \mathbf{ku} < \mathbf{n}$ .						
<b>alpha</b>	The scalar $\alpha$ . If <b>alpha</b> = 0, the subprograms compute $y \leftarrow \beta y$ without referencing <b>ab</b> or <b>x</b> .						
<b>ab</b>	Array containing the $m$ -by- $n$ band matrix $A$ in the compressed form described above. If $a_{ij}$ is in the band, it is stored in <b>ab</b> ( $\mathbf{ku}+1+i-j, j$ ). The columns of $A$ are stored in the columns of <b>ab</b> , and the diagonals of $A$ are stored in rows 1 through $\mathbf{kl}+\mathbf{ku}+1$ .						
<b>ldab</b>	The leading dimension of array <b>ab</b> as declared in the calling program unit, with $\mathbf{ldab} \geq \mathbf{kl}+\mathbf{ku}+1$ .						
<b>x</b>	Array containing the vector $x$ . The number of elements of $x$ and the value of <b>lenx</b> , the dimension of the array <b>x</b> , depend on <b>trans</b> : <table style="margin-left: 40px;"> <tbody> <tr> <td style="padding-right: 20px;"><math>\mathbf{N}</math> or <math>\mathbf{n}</math></td> <td><math>x</math> has <math>n</math> elements</td> <td><b>lenx</b> = <math>(\mathbf{n}-1) \times  \mathbf{incx}  + 1</math></td> </tr> <tr> <td>otherwise</td> <td><math>x</math> has <math>m</math> elements</td> <td><b>lenx</b> = <math>(\mathbf{m}-1) \times  \mathbf{incx}  + 1</math></td> </tr> </tbody> </table>	$\mathbf{N}$ or $\mathbf{n}$	$x$ has $n$ elements	<b>lenx</b> = $(\mathbf{n}-1) \times  \mathbf{incx}  + 1$	otherwise	$x$ has $m$ elements	<b>lenx</b> = $(\mathbf{m}-1) \times  \mathbf{incx}  + 1$
$\mathbf{N}$ or $\mathbf{n}$	$x$ has $n$ elements	<b>lenx</b> = $(\mathbf{n}-1) \times  \mathbf{incx}  + 1$					
otherwise	$x$ has $m$ elements	<b>lenx</b> = $(\mathbf{m}-1) \times  \mathbf{incx}  + 1$					
<b>incx</b>	Increment for the array <b>x</b> , <b>incx</b> $\neq$ 0: <table style="margin-left: 40px;"> <tbody> <tr> <td><b>incx</b> &gt; 0</td> <td><math>x</math> is stored forward in array <b>x</b>, i.e., <math>x_i</math> is stored in <math>\mathbf{x}((i-1) \times \mathbf{incx} + 1)</math>.</td> </tr> <tr> <td><b>incx</b> &lt; 0</td> <td><math>x</math> is stored backward in array <b>x</b>, i.e., if <b>trans</b> = <math>\mathbf{N}</math> or <math>\mathbf{n}</math>, then <math>x_i</math> is stored in <math>\mathbf{x}((i-\mathbf{n}) \times \mathbf{incx} + 1)</math>; otherwise, <math>x_i</math> is stored in <math>\mathbf{x}((i-\mathbf{m}) \times \mathbf{incx} + 1)</math>.</td> </tr> </tbody> </table> <p>Use <b>incx</b> = 1 if the vector <math>x</math> is stored contiguously in array <b>x</b>, i.e., if <math>x_i</math> is stored in <math>\mathbf{x}(i)</math>. Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.</p>	<b>incx</b> > 0	$x$ is stored forward in array <b>x</b> , i.e., $x_i$ is stored in $\mathbf{x}((i-1) \times \mathbf{incx} + 1)$ .	<b>incx</b> < 0	$x$ is stored backward in array <b>x</b> , i.e., if <b>trans</b> = $\mathbf{N}$ or $\mathbf{n}$ , then $x_i$ is stored in $\mathbf{x}((i-\mathbf{n}) \times \mathbf{incx} + 1)$ ; otherwise, $x_i$ is stored in $\mathbf{x}((i-\mathbf{m}) \times \mathbf{incx} + 1)$ .		
<b>incx</b> > 0	$x$ is stored forward in array <b>x</b> , i.e., $x_i$ is stored in $\mathbf{x}((i-1) \times \mathbf{incx} + 1)$ .						
<b>incx</b> < 0	$x$ is stored backward in array <b>x</b> , i.e., if <b>trans</b> = $\mathbf{N}$ or $\mathbf{n}$ , then $x_i$ is stored in $\mathbf{x}((i-\mathbf{n}) \times \mathbf{incx} + 1)$ ; otherwise, $x_i$ is stored in $\mathbf{x}((i-\mathbf{m}) \times \mathbf{incx} + 1)$ .						
<b>beta</b>	The scalar $\beta$ .						
<b>y</b>	Array containing the vector $y$ . The number of elements of $y$ and the value of <b>leny</b> , the dimension of the array <b>y</b> , depend on <b>trans</b> : <table style="margin-left: 40px;"> <tbody> <tr> <td style="padding-right: 20px;"><math>\mathbf{N}</math> or <math>\mathbf{n}</math></td> <td><math>y</math> has <math>m</math> elements</td> <td><b>leny</b> = <math>(\mathbf{m}-1) \times  \mathbf{incy}  + 1</math></td> </tr> <tr> <td>otherwise</td> <td><math>y</math> has <math>n</math> elements</td> <td><b>leny</b> = <math>(\mathbf{n}-1) \times  \mathbf{incy}  + 1</math></td> </tr> </tbody> </table> <p>Not used as input if <b>beta</b> = 0.</p>	$\mathbf{N}$ or $\mathbf{n}$	$y$ has $m$ elements	<b>leny</b> = $(\mathbf{m}-1) \times  \mathbf{incy}  + 1$	otherwise	$y$ has $n$ elements	<b>leny</b> = $(\mathbf{n}-1) \times  \mathbf{incy}  + 1$
$\mathbf{N}$ or $\mathbf{n}$	$y$ has $m$ elements	<b>leny</b> = $(\mathbf{m}-1) \times  \mathbf{incy}  + 1$					
otherwise	$y$ has $n$ elements	<b>leny</b> = $(\mathbf{n}-1) \times  \mathbf{incy}  + 1$					

**incy**            Increment for the array **y**, **incy**  $\neq$  0:

**incy** > 0        **y** is stored forward in array **y**, i.e.,  $y_i$  is stored in  $y((i-1) \times \text{incy} + 1)$ .

**incy** < 0        **y** is stored backward in array **y**, i.e., if **trans** = 'N' or 'n', then  $y_i$  is stored in  $y((i-m) \times \text{incy} + 1)$ ; otherwise,  $y_i$  is stored in  $y((i-n) \times \text{incy} + 1)$ .

Use **incy** = 1 if the vector **y** is stored contiguously in array **y**, i.e., if  $y_i$  is stored in  $y(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.

## Output

**y**                    The updated **y** vector replaces the input.

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**trans**  $\neq$  'N' or 'n' or 'T' or 't' or 'C' or 'c',  
**m** < 0,  
**n** < 0,  
**kl** < 0,  
**ku** < 0,  
**ldab** < **kl**+**ku**+1,  
**incx** = 0, and  
**incy** = 0.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **trans** argument as 'NORMAL' or 'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

## Example 1

Form the REAL\*8 matrix-vector product  $y = Ax$ , where **A** is a 9 by 6 real band matrix whose lower bandwidth is 2 and whose upper bandwidth is 3. **A** is stored in an array **AB** whose dimensions are 10 by 10, **x** is a real vector 6

## Basic Matrix Operations

### SGBMV/CGBMV – Matrix-Vector Multiply

elements long stored in an array X of dimension 10, and y is a real vector 9 elements long stored in an array Y, also of dimension 10.

```
CHARACTER*1 TRANS
INTEGER*8 M,N,KL,KU,LDAB, INCX, INCY
REAL*8 ALPHA,BETA, AB(10,10), X(10), Y(10)
TRANS = 'N'
M = 9
N = 6
KL = 2
KU = 3
ALPHA = 1.0
BETA = 0.0
LDAB = 10
INCX = 1
INCY = 1
CALL SGBMV (TRANS, M, N, KL, KU, ALPHA, AB, LDAB, X, INCX, BETA, Y,
            INCY)
```

### Example 2

Form the REAL\*8 matrix-vector product  $y = \frac{1}{2}y - \rho A^T x$ , where  $\rho$  is a real scalar, A is a 6-by-9 real band matrix whose lower bandwidth is 1 and whose upper bandwidth is 2. A is stored in an array AB whose dimensions are 10 by 10, x is a real vector 6 elements long stored in an array X of dimension 10, and y is a real vector 9 elements long stored in an array Y, also of dimension 10.

```
INTEGER*8 M,N,KL,KU,LDAB
REAL*8 RHO, AB(10,10), X(10), Y(10)
M = 9
N = 6
KL = 1
KU = 2
LDAB = 10
CALL SGBMV ('TRANSPOSE', M, N, KL, KU, -RHO, AB, LDAB, X, 1, 0.5,
            Y, 1)
```

**NAME**           SGEMM/CGEMM – Matrix-Matrix Multiply

**Purpose**

These subprograms compute the matrix-matrix product  $AB$ , where  $A$  is an  $m$ -by- $k$  matrix, and  $B$  is a  $k$ -by- $n$  matrix. Optionally,  $A$  may be replaced by  $A^T$  or  $A^*$ , where  $A$  is a  $k$ -by- $m$  matrix, and  $B$  may be replaced by  $B^T$  or  $B^*$ , where  $B$  is an  $n$ -by- $k$  matrix. Here,  $A^T$  and  $B^T$  are the transposes and  $A^*$  and  $B^*$  are the conjugate-transposes of  $A$  and  $B$ , respectively. The product may be stored in the result matrix (which is always of size  $m$  by  $n$ ) or optionally may be added to or subtracted from it. This is handled in a convenient, but general, way by two scalar arguments,  $\alpha$  and  $\beta$ , which are used as multipliers of the matrix product and the result matrix. Specifically, these subprograms compute matrix products of the forms

$$\begin{array}{lll}
 C \leftarrow \alpha AB + \beta C, & C \leftarrow \alpha A^T B + \beta C, & C \leftarrow \alpha A^* B + \beta C, \\
 C \leftarrow \alpha AB^T + \beta C, & C \leftarrow \alpha A^T B^T + \beta C, & C \leftarrow \alpha A^* B^T + \beta C, \\
 C \leftarrow \alpha AB^* + \beta C, & C \leftarrow \alpha A^T B^* + \beta C, & C \leftarrow \alpha A^* B^* + \beta C.
 \end{array}$$

**Usage**

SCILIB:

```

CHARACTER*1      transa, transb
INTEGER*8        m, n, k, lda, ldb, ldc
REAL*8          alpha, beta, a(lda, *), b(ldb, *), c(ldc, n)
CALL SGEMM(transa, transb, m, n, k, alpha, a, lda, b, ldb, beta, c, ldc)

CHARACTER*1      transa, transb
INTEGER*8        m, n, k, lda, ldb, ldc
COMPLEX*16      alpha, beta, a(lda, *), b(ldb, *), c(ldc, n)
CALL CGEMM(transa, transb, m, n, k, alpha, a, lda, b, ldb, beta, c, ldc)

```

**Input**

**transa**           Transposition option for  $A$ :  
                  'N' or 'n'     Use  $m$ -by- $k$  matrix  $A$   
                  'T' or 't'     Use  $A^T$  where  $A$  is a  $k$ -by- $m$  matrix  
                  'C' or 'c'     Use  $A^*$  where  $A$  is a  $k$ -by- $m$  matrix

where  $A^T$  is the transpose of  $A$  and  $A^*$  is the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.

**transb**           Transposition option for  $B$ :  
                  'N' or 'n'     Use  $k$ -by- $n$  matrix  $B$

## Basic Matrix Operations

### SGEMM/CGEMM – Matrix-Matrix Multiply

'T' or 't'      Use  $B^T$  where  $B$  is an  $n$ -by- $k$  matrix

'C' or 'c'      Use  $B^*$  where  $B$  is an  $n$ -by- $k$  matrix

where  $B^T$  is the transpose of  $B$  and  $B^*$  is the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.

- m**      Number of rows in matrix  $C$ ,  $m \geq 0$ . If  $m = 0$ , the subprograms do not reference  $a$ ,  $b$ , or  $c$ .
- n**      Number of columns in matrix  $C$ ,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference  $a$ ,  $b$ , or  $c$ .
- k**      The *middle* dimension of the matrix multiply,  $k \geq 0$ . If  $k = 0$ , the subprograms compute  $C \leftarrow \beta C$  without referencing  $a$  or  $b$ .
- alpha**      The scalar  $\alpha$ . If **alpha** = 0, the subprograms compute  $C \leftarrow \beta C$  without referencing  $a$  or  $b$ .
- a**      Array containing the matrix  $A$ , whose size is indicated by **transa**:  
'N' or 'n'       $A$  is an  $m$ -by- $k$  matrix  
otherwise       $A$  is a  $k$ -by- $m$  matrix
- lda**      The leading dimension of array  $a$  as declared in the calling program unit, with  $lda \geq \max(\text{the number of rows of } A, 1)$ .
- b**      Array containing the matrix  $B$ , whose size is indicated by **transb**:  
'N' or 'n'       $B$  is a  $k$ -by- $n$  matrix  
otherwise       $B$  is an  $n$ -by- $k$  matrix
- ldb**      The leading dimension of array  $b$  as declared in the calling program unit, with  $ldb \geq \max(\text{the number of rows of } B, 1)$ .
- beta**      The scalar  $\beta$ .
- c**      Array containing the  $m$ -by- $n$  matrix  $C$ . Not used as input if **beta** = 0.
- ldc**      The leading dimension of array  $c$  as declared in the calling program unit, with  $ldc \geq \max(m, 1)$ .

## Output

- c**      The updated  $C$  matrix replaces the input.

## Notes

These subprograms conform to specifications of the Level 3 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

```

transa ≠ 'N' or 'n' or 'T' or 't' or 'C' or 'c',
transb ≠ 'N' or 'n' or 'T' or 't' or 'C' or 'c',
m < 0,
n < 0,
k < 0,
lda too small,
ldb too small, and
ldc < max(m,1).

```

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved, for example, by coding the **transa** and **transb** arguments as 'NORMAL' or 'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

### Example 1

Form the REAL\*8 matrix product  $C = AB$ , where  $A$  is a 9-by-6 real matrix stored in an array  $A$  whose dimensions are 10 by 10,  $B$  is a 6-by-8 real matrix stored in an array  $B$  of dimension 10 by 10, and  $C$  is a 9-by-8 real matrix stored in an array  $C$ , also of dimension 10 by 10.

```

CHARACTER*1  TRANSA, TRANSB
INTEGER*8    M, N, K, LDA, LDB, LDC
REAL*8      ALPHA, BETA, A(10,10), B(10,10), C(10,10)
TRANSA = 'N'
TRANSB = 'N'
M = 9
N = 8
K = 6
ALPHA = 1.0
BETA = 0.0
LDA = 10
LDB = 10
LDC = 10
CALL SGEMM (TRANSA, TRANSB, M, N, K, ALPHA, A, LDA, B, LDB, BETA, C,
            LDC)

```

### Example 2

Form the REAL\*8 matrix product  $C = \frac{1}{2}C + \rho A^T B$ , where  $\rho$  is a real scalar,  $A$  is a 6-by-9 real matrix stored in an array  $A$  whose dimensions are 10 by 10,  $B$  is

## Basic Matrix Operations

### SGEMM/CGEMM – Matrix-Matrix Multiply

a 6-by-8 real matrix stored in an array B of dimension 10 by 10, and C is a 9-by-8 real matrix stored in an array C, also of dimension 10 by 10.

```
INTEGER*8 M,N,K,LDA,LDB,LDC
REAL*8    RHO,A(10,10),B(10,10),C(10,10)
M = 9
N = 8
K = 6
LDA = 10
LDB = 10
LDC = 10
CALL SGEMM ('TRAN', 'NONTRAN', M,N,K,RHO,A,LDA,B,LDB,0.5,C,
           LDC)
```

## SGEMMS/CGEMMS – Strassen Matrix-Matrix Multiply

**NAME** SGEMMS/CGEMMS – Strassen Matrix-Matrix Multiply

**Purpose**

These subprograms use Strassen's method to compute the matrix-matrix product  $AB$ , where  $A$  is an  $m$ -by- $k$  matrix, and  $B$  is a  $k$ -by- $n$  matrix. Strassen's method is an algorithm for matrix multiplication which, under certain circumstances, uses fewer than  $mnk$  multiplications and additions. These subprograms are functionally equivalent to the SCILIB Level 3 BLAS subprograms SGEMM and CGEMM, and differ in usage only by the extra character in the subprogram name and the additional argument, **work**. By using Strassen's method, these subprograms may be considerably faster than their SCILIB counterparts. Refer to "Notes" for details.

In addition to computing the matrix-matrix product  $AB$ ,  $A$  may be replaced by  $A^T$  or  $A^*$ , where  $A$  is a  $k$ -by- $m$  matrix, and  $B$  may be replaced by  $B^T$  or  $B^*$ , where  $B$  is an  $n$ -by- $k$  matrix. Here,  $A^T$  and  $B^T$  are the transposes and  $A^*$  and  $B^*$  are the conjugate-transposes of  $A$  and  $B$ , respectively. The product may be stored in the result matrix (which is always of size  $m$  by  $n$ ) or optionally may be added to or subtracted from it. This is handled in a convenient, but general, way by two scalar arguments,  $\alpha$  and  $\beta$ , which are used as multipliers of the matrix product and the result matrix. Specifically, these subprograms compute matrix products of the forms

$$\begin{array}{lll} C \leftarrow \alpha AB + \beta C, & C \leftarrow \alpha A^T B + \beta C, & C \leftarrow \alpha A^* B + \beta C, \\ C \leftarrow \alpha AB^T + \beta C, & C \leftarrow \alpha A^T B^T + \beta C, & C \leftarrow \alpha A^* B^T + \beta C, \\ C \leftarrow \alpha AB^* + \beta C, & C \leftarrow \alpha A^T B^* + \beta C, & C \leftarrow \alpha A^* B^* + \beta C. \end{array}$$

**Usage**

SCILIB:

CHARACTER\*1      transa, transb  
 INTEGER\*8        m, n, k, lda, ldb, ldc  
 REAL\*8           alpha, beta, a(lda, \*), b(ldb, \*), c(ldc, n),  
                   work(lwork)  
 CALL SGEMMS(transa, transb, m, n, k, alpha, a, lda, b, ldb, beta, c, ldc,  
               work)

CHARACTER\*1      transa, transb  
 INTEGER\*8        m, n, k, lda, ldb, ldc  
 COMPLEX\*16       alpha, beta, a(lda, \*), b(ldb, \*), c(ldc, n),  
                   work(lwork)  
 CALL CGEMMS(transa, transb, m, n, k, alpha, a, lda, b, ldb, beta, c, ldc,  
               work)

**Input**

<b>transa</b>	<p>Transposition option for <math>A</math>:</p> <p>'N' or 'n'    Use <math>m</math>-by-<math>k</math> matrix <math>A</math></p> <p>'T' or 't'    Use <math>A^T</math> where <math>A</math> is a <math>k</math>-by-<math>m</math> matrix</p> <p>'C' or 'c'    Use <math>A^*</math> where <math>A</math> is a <math>k</math>-by-<math>m</math> matrix</p> <p>where <math>A^T</math> is the transpose of <math>A</math> and <math>A^*</math> is the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.</p>
<b>transb</b>	<p>Transposition option for <math>B</math>:</p> <p>'N' or 'n'    Use <math>k</math>-by-<math>n</math> matrix <math>B</math></p> <p>'T' or 't'    Use <math>B^T</math> where <math>B</math> is an <math>n</math>-by-<math>k</math> matrix</p> <p>'C' or 'c'    Use <math>B^*</math> where <math>B</math> is an <math>n</math>-by-<math>k</math> matrix</p> <p>where <math>B^T</math> is the transpose of <math>B</math> and <math>B^*</math> is the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.</p>
<b>m</b>	Number of rows in matrix $C$ , $m \geq 0$ . If $m = 0$ , the subprograms do not reference $a$ , $b$ , or $c$ .
<b>n</b>	Number of columns in matrix $C$ , $n \geq 0$ . If $n = 0$ , the subprograms do not reference $a$ , $b$ , or $c$ .
<b>k</b>	The <i>middle</i> dimension of the matrix multiply, $k \geq 0$ . If $k = 0$ , the subprograms compute $C \leftarrow \beta C$ without referencing $a$ or $b$ .
<b>alpha</b>	The scalar $\alpha$ . If <b>alpha</b> = 0, the subprograms compute $C \leftarrow \beta C$ without referencing $a$ or $b$ .
<b>a</b>	<p>Array containing the matrix <math>A</math>, whose size is indicated by <b>transa</b>:</p> <p>'N' or 'n'    <math>A</math> is an <math>m</math>-by-<math>k</math> matrix</p> <p>otherwise    <math>A</math> is a <math>k</math>-by-<math>m</math> matrix</p>
<b>lda</b>	The leading dimension of array <b>a</b> as declared in the calling program unit, with <b>lda</b> $\geq$ max (the number of rows of $A$ , 1).
<b>b</b>	<p>Array containing the matrix <math>B</math>, whose size is indicated by <b>transb</b>:</p> <p>'N' or 'n'    <math>B</math> is a <math>k</math>-by-<math>n</math> matrix</p> <p>otherwise    <math>B</math> is an <math>n</math>-by-<math>k</math> matrix</p>
<b>ldb</b>	The leading dimension of array <b>b</b> as declared in the calling program unit, with <b>ldb</b> $\geq$ max (the number of rows of $B$ , 1).

## SGEMMS/CGEMMS – Strassen Matrix-Matrix Multiply

<b>beta</b>	The scalar $\beta$ .
<b>c</b>	Array containing the $m$ -by- $n$ matrix $C$ . Not used as input if <b>beta</b> = 0.
<b>ldc</b>	The leading dimension of array <b>c</b> as declared in the calling program unit, with <b>ldc</b> $\geq$ $\max(m, 1)$ .

**Working Storage**

<b>work</b>	An array of size <b>lwork</b> = $2.34 \times \max(m, k) \times \max(n, k)$ , used for work space.
-------------	---

**Output**

<b>c</b>	The updated $C$ matrix replaces the input.
----------	--

**Notes**

Except for the extra character in the subprogram name and the additional working storage argument, these subprograms conform to specifications of the Level 3 BLAS subprograms SGEMM and CGEMM.

Because of their use of Strassen's method, CGEMMS and SGEMMS are asymptotically faster than standard matrix multiply methods such as those employed in the SCILIB routines CGEMM and SGEMM. In practice these particular implementations are faster than their standard counterparts. If  $\min(m, n, k) > 200$  for CGEMMS and  $\min(m, n, k) > 512$  for SGEMMS. The speedup in the complex case is much more pronounced. That is due in large part to the complex bilinear reduction technique (implemented underneath Strassen's method) that allows two complex matrices to be multiplied using only 3/4 of the multiplications required by the traditional method. Also, the relative cost of data motion is lower in the complex case. In this first release, the gains in the real case are marginal.

- If  $\min(m, n, k) > 200$  for CGEMMS and  $\min(m, n, k) > 512$  for SGEMMS. The speedup in the complex case is much more pronounced. That is due in large part to the complex bilinear reduction technique (implemented underneath Strassen's method) that allows two complex matrices to be multiplied using only 3/4 of the multiplications required by the traditional method. Also, the relative cost of data motion is lower in the complex case. In this first release, the gains in the real case are marginal.

In the operator norm, Strassen's method is slightly less stable than traditional matrix multiplication, and the computation of individual elements is unstable. The emerging consensus seems to be that Strassen's method is sufficiently stable for most applications.

For a good overview and bibliography of this subject, see (Higham).

## Basic Matrix Operations

### SGEMMS/CGEMMS – Strassen Matrix-Matrix Multiply

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

```
transa ≠ 'N' or 'n' or 'T' or 't' or 'C' or 'c',
transb ≠ 'N' or 'n' or 'T' or 't' or 'C' or 'c',
m < 0,
n < 0,
k < 0,
lda too small,
ldb too small, and
ldc < max(m,1).
```

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the CALL statement may be improved, for example, by coding the **transa** and **transb** arguments as 'NORMAL' or 'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

#### Example 1

Form the REAL\*8 matrix product  $C = AB$ , where  $A$  is a 900-by-600 real matrix stored in an array  $A$  whose dimensions are 1000 by 1000,  $B$  is a 600-by-800 real matrix stored in an array  $B$  of dimension 1000 by 1000, and  $C$  is a 900-by-800 real matrix stored in an array  $C$ , also of dimension 1000 by 1000. WORK is declared large enough to handle all matrices that will fit in the arrays.

```
CHARACTER*1  TRANSA, TRANSB
INTEGER*8    M, N, K, LDA, LDB, LDC
REAL*8      ALPHA, BETA, A(1000,1000), B(1000,1000),
1           C(1000,1000), WORK(2340000)
TRANSA = 'N'
TRANSB = 'N'
M = 900
N = 800
K = 600
ALPHA = 1.0
BETA = 0.0
LDA = 1000
LDB = 1000
LDC = 1000
CALL SGEMMS (TRANSA, TRANSB, M, N, K, ALPHA, A, LDA, B, LDB, BETA,
1           C, LDC, WORK)
```

**Example 2**

Form the COMPLEX\*16 matrix product  $C = \frac{1}{2}C + \rho A*B$ , where  $\rho$  is a complex scalar,  $A$  is a 600-by-900 complex matrix stored in an array  $A$  whose dimensions are 1000 by 1000,  $B$  is a 600-by-800 complex matrix stored in an array  $B$  of dimension 1000 by 1000, and  $C$  is a 900-by-800 complex matrix stored in an array  $C$ , also of dimension 1000 by 1000.  $WORK$  is declared large enough to handle all matrices that will fit in the arrays.

```
      INTEGER*8  M,N,K, LDA, LDB, LDC
      COMPLEX*16 RHO, A(1000,1000), B(1000,1000), C(1000,1000),
1      WORK(2340000)
      M = 900
      N = 800
      K = 600
      LDA = 1000
      LDB = 1000
      LDC = 1000
      CALL CGEMMS ('CONJ', 'NORMAL', M, N, K, RHO, A, LDA, B, LDB,
1      (0.5, 0.0), C, LDC, WORK)
```

Basic Matrix Operations  
SGEMV/CGEMV – Matrix-Vector Multiply

**NAME** SGEMV/CGEMV – Matrix-Vector Multiply

**Purpose**

These subprograms compute the matrix-vector products  $Ax$ ,  $A^T x$ , and  $A^* x$ , where  $A$  is an  $m$ -by- $n$  matrix,  $A^T$  is the transpose of  $A$ , and  $A^*$  is the conjugate transpose of  $A$ . The product may be stored in the result array, or optionally added to or subtracted from it. This is handled in a convenient, but general, way by two scalar arguments,  $\alpha$  and  $\beta$ , which are used as multipliers of the matrix-vector product and the result vector. Specifically, these subprograms compute matrix-vector products of the forms

$$y \leftarrow \alpha Ax + \beta y, \quad y \leftarrow \alpha A^T x + \beta y, \quad \text{and} \quad y \leftarrow \alpha A^* x + \beta y.$$

**Usage**

SCILIB:

```
CHARACTER*1      trans
INTEGER*8        m, n, lda, incx, incy
REAL*8           alpha, beta, a(lda, n), x(lenx), y(leny)
CALL SGEMV(trans, m, n, alpha, a, lda, x, incx, beta, y, incy)

CHARACTER*1      trans
INTEGER*8        m, n, lda, incx, incy
COMPLEX*16       alpha, beta, a(lda, n), x(lenx), y(leny)
CALL CGEMV(trans, m, n, alpha, a, lda, x, incx, beta, y, incy)
```

**Input**

**trans** Transposition option for  $A$ :  
‘N’ or ‘n’ Compute  $y \leftarrow \alpha Ax + \beta y$   
‘T’ or ‘t’ Compute  $y \leftarrow \alpha A^T x + \beta y$   
‘C’ or ‘c’ Compute  $y \leftarrow \alpha A^* x + \beta y$

where  $A^T$  is the transpose of  $A$  and  $A^*$  is the conjugate transpose. In the real subprograms, ‘C’ and ‘c’ have the same meaning as ‘T’ and ‘t’.

**m** Number of rows in matrix  $A$ ,  $m \geq 0$ . If  $m = 0$ , the subprograms do not reference  $a$ ,  $x$ , or  $y$ .

**n** Number of columns in matrix  $A$ ,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference  $a$ ,  $x$ , or  $y$ .

**alpha** The scalar  $\alpha$ . If **alpha** = 0, the subprograms compute  $y \leftarrow \beta y$  without referencing  $A$  or  $x$ .

**a** Array containing the  $m$ -by- $n$  matrix  $A$ .

**lda** The leading dimension of array **a** as declared in the calling program unit, with  $lda \geq \max(m, 1)$ .

**x** Array containing the vector  $x$ . The number of elements of  $x$  and the value of **lenx**, the dimension of the array **x**, depend on **trans**:

$\text{'N' or 'n'}$	$x$ has $n$ elements	$\text{lenx} = (n-1) \times  \text{incx}  + 1$
otherwise	$x$ has $m$ elements	$\text{lenx} = (m-1) \times  \text{incx}  + 1$

**incx** Increment for the array **x**,  $\text{incx} \neq 0$ :

$\text{incx} > 0$   $x$  is stored forward in array **x**, i.e.,  $x_i$  is stored in  $\mathbf{x}((i-1) \times \text{incx} + 1)$ .

$\text{incx} < 0$   $x$  is stored backward in array **x**, i.e., if **trans** =  $\text{'N' or 'n'}$ , then  $x_i$  is stored in  $\mathbf{x}((i-n) \times \text{incx} + 1)$ ; otherwise,  $x_i$  is stored in  $\mathbf{x}((i-m) \times \text{incx} + 1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array **x**, i.e., if  $x_i$  is stored in  $\mathbf{x}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

**beta** The scalar  $\beta$ .

**y** Array containing the vector  $y$ . The number of elements of  $y$  and the value of **leny**, the dimension of the array **y**, depend on **trans**:

$\text{'N' or 'n'}$	$y$ has $m$ elements	$\text{leny} = (m-1) \times  \text{incy}  + 1$
otherwise	$y$ has $n$ elements	$\text{leny} = (n-1) \times  \text{incy}  + 1$

Not used as input if  $\text{beta} = 0$ .

**incy** Increment for the array **y**,  $\text{incy} \neq 0$ :

$\text{incy} > 0$   $y$  is stored forward in array **y**, i.e.,  $y_i$  is stored in  $\mathbf{y}((i-1) \times \text{incy} + 1)$ .

$\text{incy} < 0$   $y$  is stored backward in array **y**, i.e., if **trans** =  $\text{'N' or 'n'}$ , then  $y_i$  is stored in  $\mathbf{y}((i-m) \times \text{incy} + 1)$ ; otherwise,  $y_i$  is stored in  $\mathbf{y}((i-n) \times \text{incy} + 1)$ .

Use  $\text{incy} = 1$  if the vector  $y$  is stored contiguously in array **y**, i.e., if  $y_i$  is stored in  $\mathbf{y}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

## Output

**y** The updated  $y$  vector replaces the input.

Basic Matrix Operations  
SGEMV/CGEMV – Matrix-Vector Multiply

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**trans** ≠ 'N' or 'n' or 'T' or 't' or 'C' or 'c',  
**m** < 0,  
**n** < 0,  
**lda** < max(**m**,1),  
**incx** = 0, and  
**incy** = 0.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **trans** argument as 'NORMAL' or 'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

## Example 1

Form the REAL\*8 matrix-vector product  $y = Ax$ , where  $A$  is a 9-by-6 real matrix stored in an array  $A$  whose dimensions are 10 by 10,  $x$  is a real vector 6 elements long stored in an array  $X$  of dimension 10, and  $y$  is a real vector 9 elements long stored in an array  $Y$ , also of dimension 10.

```
CHARACTER*1 TRANS
INTEGER*8 M,N,LDA, INCX, INCY
REAL*8 ALPHA,BETA,A(10,10),X(10),Y(10)
TRANS = 'N'
M = 9
N = 6
ALPHA = 1.0
BETA = 0.0
LDA = 10
INCX = 1
INCY = 1
CALL SGEMV (TRANS,M,N,ALPHA,A,LDA,X,INCX,BETA,Y,INCY)
```

## Example 2

Form the REAL\*8 matrix-vector product  $y = \frac{1}{2}y - \rho A^T x$ , where  $\rho$  is a real scalar,  $A$  is a 6-by-9 real matrix stored in an array  $A$  whose dimensions are 10 by 10,  $x$  is a real vector 6 elements long stored in an array  $X$  of dimension 10, and  $y$  is a real vector 9 elements long stored in an array  $Y$ , also of dimension 10.

Basic Matrix Operations  
**SGEMV/CGEMV – Matrix-Vector Multiply**

```
INTEGER*8 M,N,LDA  
REAL*8    RHO,A(10,10),X(10),Y(10)  
M = 9  
N = 6  
LDA = 10  
CALL SGEMV ('TRANSPOSE',M,N,-RHO,A,LDA,X,1,0.5,Y,1)
```

Basic Matrix Operations  
SGER/CGERC/CGERU – Rank-1 Update

**NAME** SGER/CGERC/CGERU – Rank-1 Update

**Purpose**

These subprograms compute the rank-1 updates

$$A \leftarrow \alpha xy^T + A \quad \text{and} \quad A \leftarrow \alpha y^* + A,$$

where  $A$  is an  $m$ -by- $n$  matrix,  $\alpha$  is a scalar,  $x$  is an  $m$ -vector,  $y$  is an  $n$ -vector, and  $y^T$  and  $y^*$  are the transpose and conjugate transpose of  $y$ , respectively.

**Usage**

SCILIB:

```
INTEGER*8          m, n, lda, incx, incy
REAL*8             alpha, a(lda, n), x(lenx), y(leny)
CALL SGER(m, n, alpha, x, incx, y, incy, a, lda)

INTEGER*8          m, n, lda, incx, incy
COMPLEX*16         alpha, a(lda, n), x(lenx), y(leny)
CALL CGERC(m, n, alpha, x, incx, y, incy, a, lda)

INTEGER*8          m, n, lda, incx, incy
COMPLEX*16         alpha, a(lda, n), x(lenx), y(leny)
CALL CGERU(m, n, alpha, x, incx, y, incy, a, lda)
```

**Input**

**m** Number of rows in matrix  $A$  and elements of vector  $x$ ,  $m \geq 0$ .  
If  $m = 0$ , the subprograms do not reference  $a$ ,  $x$ , or  $y$ .

**n** Number of columns in matrix  $A$  and elements of vector  $y$ ,  $n \geq 0$ .  
If  $n = 0$ , the subprograms do not reference  $a$ ,  $x$ , or  $y$ .

**alpha** The scalar  $\alpha$ . If  $\alpha = 0$ , the subprograms do not reference  $A$ ,  $x$ , or  $y$ .

**x** Array of length  $\text{lenx} = (m-1) \times |\text{incx}| + 1$  containing the  $m$ -vector  $x$ .

**incx** Increment for the array  $x$ ,  $\text{incx} \neq 0$ :

**incx** > 0  $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx** < 0  $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-m) \times \text{incx} + 1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.

- y** Array of length  $\text{leny} = (n-1) \times |\text{incy}| + 1$  containing the  $n$ -vector  $y$ .  $y$  is used in conjugated form by CGERC, and in unconjugated form by the other subprograms. Refer to “Purpose.”
- incy** Increment for the array  $y$ ,  $\text{incy} \neq 0$ :
- incy** > 0  $y$  is stored forward in array  $y$ , i.e.,  $y_i$  is stored in  $y((i-1) \times \text{incy} + 1)$ .
- incy** < 0  $y$  is stored backward in array  $y$ , i.e.,  $y_i$  is stored in  $y((i-n) \times \text{incy} + 1)$ .
- Use  $\text{incy} = 1$  if the vector  $y$  is stored contiguously in array  $y$ , i.e., if  $y_i$  is stored in  $y(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.
- a** Array containing the  $m$ -by- $n$  matrix  $A$ .
- lda** The leading dimension of array  $a$  as declared in the calling program unit, with  $\text{lda} \geq \max(m, 1)$ .

## Output

- a** The updated  $A$  matrix replaces the input.

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

- $m < 0$ ,
- $n < 0$ ,
- $\text{lda} < \max(m, 1)$ ,
- $\text{incx} = 0$ , and
- $\text{incy} = 0$ .

## Example 1

Apply a REAL\*8 rank-1 update  $xy^T$  to  $A$ , where  $A$  is a 6-by-9 real matrix stored in an array  $A$  whose dimensions are 10 by 10,  $x$  is a real vector 6 elements long stored in an array  $X$  of dimension 10, and  $y$  is a real vector 9 elements long stored in an array  $Y$ , also of dimension 10.

Basic Matrix Operations  
SGER/CGERC/CGERU – Rank-1 Update

```
INTEGER*8 M,N,LDA, INCX, INCY
REAL*8    ALPHA,A(10,10),X(10),Y(10)
M = 6
N = 9
ALPHA = 1.0
LDA = 10
INCX = 1
INCY = 1
CALL SGER (M,N,ALPHA,X, INCX,Y, INCY,A,LDA)
```

**Example 2**

Apply a COMPLEX\*16 conjugated rank-1 update  $-2xy^*$  to  $A$ , where  $A$  is a 6-by-9 complex matrix stored in an array  $A$  whose dimensions are 10 by 10,  $x$  is a complex vector 6 elements long stored in an array  $X$  of dimension 10, and  $y$  is a complex vector 9 elements long stored in an array  $Y$ , also of dimension 10.

```
INTEGER*8 M,N,LDA
COMPLEX*16 A(10,10),X(10),Y(10)
M = 6
N = 9
LDA = 10
CALL CGERC (M,N,(-2.0E0,0.0E0),X,1,Y,1,A,LDA)
```

**NAME** SMXPY – Matrix-Vector Multiply and Add

**Purpose**

This subprogram computes the matrix-vector product  $Ax$ , and adds the result to another vector  $y$ , where  $A$  is an  $m$ -by- $n$  matrix,  $x$  is an  $n$ -vector, and  $y$  is an  $m$ -vector. SCILIB subprogram SGEMV allows more general storage of  $x$  and  $y$  and also admits scaling, subtraction, and transposing  $A$ .

**Usage**

SCILIB:

```
INTEGER*8      m, n, lda
REAL*8        a(lda, n), x(n), y(m)
CALL SMXPY(m, y, n, lda, x, a)
```

**Input**

<b>m</b>	Number of rows in matrix $A$ and length of vector $y$ , $m \geq 0$ . If $m = 0$ , the subprogram does not reference $a$ , $x$ , or $y$ .
<b>y</b>	Array containing the vector $y$ .
<b>n</b>	Number of columns in matrix $A$ and length of vector $x$ , $n \geq 0$ . If $n = 0$ , the subprogram does not reference $a$ , $x$ , or $y$ .
<b>lda</b>	The leading dimension of array $a$ as declared in the calling program unit.
<b>x</b>	Array containing the $n$ -vector $x$ .
<b>a</b>	Array containing the $m$ -by- $n$ matrix $A$ .

**Output**

<b>y</b>	The updated $y$ vector replaces the input.
----------	--

**Notes**

Cray Research, Inc. has declared this subprogram obsolete in release 6.0 of the UNICOS Math and Scientific Library.

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are

```
m < 0,
n < 0, and
lda < m.
```

Basic Matrix Operations  
SMXPY – Matrix-Vector Multiply and Add

### Fortran Equivalent

Except for the argument error checking, the following Fortran subroutine is equivalent to SMXPY.

```
SUBROUTINE SMXPY (M,Y,N,LDA,X,A)
INTEGER*8 M,N,LDA
REAL*8 A(LDA,N),X(N),Y(M)
DO 120 J = 1, N
    DO 110 I = 1, M
        Y(I) = A(I,J) * X(J) + Y(I)
110     CONTINUE
120 CONTINUE
RETURN
END
```

### Example

Form the REAL\*8 matrix-vector product  $y = Ax$ , where  $A$  is a 9 by 6 real matrix stored in an array  $A$  whose dimensions are 10 by 10,  $x$  is a real vector 6 elements long stored in an array  $X$  of dimension 10, and  $y$  is a real vector 9 elements long stored in an array  $Y$  of dimension 10.

```
INTEGER*8 M,N,LDA
REAL*8 A(10,10),X(10),Y(10)      M = 9
N = 6
LDA = 10
CALL SMXPY (M,Y,N,LDA,X,A)
```

**NAME**           SSBMV/CHBMV – Matrix-Vector Multiply

**Purpose**

These subprograms compute the matrix-vector product  $Ax$  where  $A$  is an  $n$  by  $n$  real symmetric or complex Hermitian band matrix and  $x$  is a real or complex  $n$ -vector. The product may be stored in the result array, or, optionally, be added to or subtracted from it. This is handled in a convenient, but general, way by two scalar arguments,  $\alpha$  and  $\beta$ , which are used as multipliers of the matrix-vector product and the result vector. Specifically, these subprograms compute the matrix-vector product of the form

$$y \leftarrow \alpha Ax + \beta y.$$

The structure of  $A$  is indicated by the name of the subprogram used:

SSBMV             $A$  is a real symmetric band matrix  
CHBMV             $A$  is a complex Hermitian band matrix

A symmetric or Hermitian band matrix is a symmetric or Hermitian matrix whose nonzero elements all are on or fairly near the principal diagonal. Specifically,  $a_{ij} \neq 0$  only if  $|i-j| \leq kd$  for some integer  $kd$ , called the half bandwidth.

**Matrix Storage**

Because it is not necessary to store or operate on the zeros outside the band of  $A$ , and since either triangle of  $A$  may be obtained from the other, you only need to provide the band within one triangle of  $A$ . Compared to storing the entire matrix, this can save memory in two ways: only the elements within the band are stored, and of them, only the upper or the lower triangle.

The following examples illustrate the storage of symmetric band matrices. Consider the following matrix  $A$  of order  $n = 7$  and half bandwidth  $kd = 2$ :

11	12	13	0	0	0	0
12	22	23	24	0	0	0
13	23	33	34	35	0	0
0	24	34	44	45	46	0
0	0	35	45	55	56	57
0	0	0	46	56	66	67
0	0	0	0	57	67	77

**Upper triangular storage**

The upper triangle of  $A$  is stored in an array  $ab$  with at least  $kd+1 = 3$  rows and 7 columns as follows:

Basic Matrix Operations  
**SSBMV/CHBMV – Matrix-Vector Multiply**

```

      *   *   13  24  35  46  57
      *  12  23  34  45  56  67
     11  22  33  44  55  66  77

```

The asterisks represent elements in the  $kd$ -by- $kd$  triangle at the upper left corner of **ab** that are not referenced. Thus, if  $a_{ij}$  is an element within the band of the upper triangle of  $A$ , it is stored in **ab**( $kd+1+i-j,j$ ). Therefore, the columns of the upper triangle of  $A$  are stored in the columns of **ab**, and the diagonals of the upper triangle of  $A$  are stored in the rows of **ab**, with the principal diagonal in row  $kd+1$ , the first superdiagonal starting in the second position in row  $kd$ , and so on.

### Lower triangular storage

The lower triangle of  $A$  is stored in the array **ab** as follows:

```

     11  22  33  44  55  66  77
     12  23  34  45  56  67  *
     13  24  35  46  57  *  *

```

The asterisks represent elements in the  $kd$ -by- $kd$  triangle at the lower right corner of **ab** that are not referenced. Thus, if  $a_{ij}$  is an element within the band of the lower triangle of  $A$ , it is stored in **ab**( $1+i-j,j$ ). Therefore, the columns of the lower triangle of  $A$  are stored in the columns of **ab**, and the diagonals of the lower triangle of  $A$  are stored in the rows of **ab**, with the principal diagonal in the first row, the first subdiagonal in the second row, and so on.

### Usage

#### SCILIB:

```

CHARACTER*1      uplo
INTEGER*8        n, kd, ldab, incx, incy
REAL*8          alpha, beta, ab(ldab, n), x(lenx), y(leny)
CALL SSBMV(uplo, n, kd, alpha, ab, ldab, x, incx, beta, y, incy)

CHARACTER*1      uplo
INTEGER*8        n, kd, ldab, incx, incy
COMPLEX*16      alpha, beta, ab(ldab, n), x(lenx), y(leny)
CALL CHBMV(uplo, n, kd, alpha, ab, ldab, x, incx, beta, y, incy)

```

### Input

**uplo**            Upper/lower triangular option for  $A$ :  
           'L' or 'l'    The lower triangle of  $A$  is stored.  
           'U' or 'u'    The upper triangle of  $A$  is stored.

<b>n</b>	Number of rows and columns in matrix $A$ , $n \geq 0$ . If $n = 0$ , the subprograms do not reference $\mathbf{ab}$ or $\mathbf{x}$ .
<b>kd</b>	The number of nonzero diagonals above or below the principal diagonal.
<b>alpha</b>	The scalar $\alpha$ . If <b>alpha</b> = 0, the subprograms compute $y \leftarrow \beta y$ without referencing $\mathbf{ab}$ or $\mathbf{x}$ .
<b>ab</b>	Array containing the $n$ -by- $n$ symmetric band matrix $A$ in the compressed form described above. The columns of the band of $A$ are stored in the columns of $\mathbf{ab}$ , and the diagonals of the band of $A$ are stored in the rows of $\mathbf{ab}$ .
<b>ldab</b>	The leading dimension of array $\mathbf{ab}$ as declared in the calling program unit, with $\mathbf{ldab} \geq \mathbf{kd}+1$ .
<b>x</b>	Array of length $\mathbf{lenx} = (n-1) \times  \mathbf{incx}  + 1$ containing the input vector $x$ .
<b>incx</b>	Increment for the array $\mathbf{x}$ , $\mathbf{incx} \neq 0$ : $\mathbf{incx} > 0$ $x$ is stored forward in array $\mathbf{x}$ , i.e., $x_i$ is stored in $\mathbf{x}((i-1) \times \mathbf{incx} + 1)$ . $\mathbf{incx} < 0$ $x$ is stored backward in array $\mathbf{x}$ , i.e., $x_i$ is stored in $\mathbf{x}((i-n) \times \mathbf{incx} + 1)$ .  Use $\mathbf{incx} = 1$ if the vector $x$ is stored contiguously in array $\mathbf{x}$ , i.e., if $x_i$ is stored in $\mathbf{x}(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.
<b>beta</b>	The scalar $\beta$ .
<b>y</b>	Array of length $\mathbf{leny} = (n-1) \times  \mathbf{incy}  + 1$ containing the $n$ -vector $y$ . Not used as input if <b>beta</b> = 0.
<b>incy</b>	Increment for the array $\mathbf{y}$ , $\mathbf{incy} \neq 0$ : $\mathbf{incy} > 0$ $y$ is stored forward in array $\mathbf{y}$ , i.e., $y_i$ is stored in $\mathbf{y}((i-1) \times \mathbf{incy} + 1)$ . $\mathbf{incy} < 0$ $y$ is stored backward in array $\mathbf{y}$ , i.e., $y_i$ is stored in $\mathbf{y}((i-n) \times \mathbf{incy} + 1)$ .  Use $\mathbf{incy} = 1$ if the vector $y$ is stored contiguously in array $\mathbf{y}$ , i.e., if $y_i$ is stored in $\mathbf{y}(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.

## Output

<b>y</b>	The updated $y$ vector replaces the input.
----------	--

Basic Matrix Operations  
SSBMV/CHBMV – Matrix-Vector Multiply

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo** ≠ 'L' or 'l' or 'U' or 'u',  
**n** < 0,  
**kd** < 0,  
**ldab** < **kd**+1,  
**incx** = 0, and  
**incy** = 0.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **uplo** argument as 'LOWER' for 'L' or 'UPPER' for 'U'. Refer to "Example 2."

## Example 1

Form the REAL\*8 matrix-vector product  $y = Ax$ , where  $A$  is a 75-by-75 real symmetric band matrix with half bandwidth 15 whose lower triangular part is stored in an array AB whose dimensions are 25 by 100, and  $x$  and  $y$  are real vectors 75 elements long stored in arrays X and Y of dimension 100, respectively.

```
CHARACTER*1 UPLO
INTEGER*8   N,KD,LDAB,INCX,INCY
REAL*8     AB(25,100),X(100),Y(100)
UPLO = 'L'
N = 75
KD = 15
LDAB = 25
INCX = 1
INCY = 1
CALL SSBMV (UPLO, N, KD, 1.0, AB, LDAB, X, INCX, 0.0, Y, INCY)
```

## Example 2

Form the REAL\*8 matrix-vector product  $y = Ax$ , where  $A$  is a 75-by-75 real symmetric band matrix with half bandwidth 15 whose upper triangle is stored in an array AB whose dimensions are 25 by 100, and  $x$  and  $y$  are real vectors 75 elements long stored in arrays X and Y of dimension 100, respectively.

Basic Matrix Operations  
**SSBMV/CHBMV – Matrix-Vector Multiply**

```
INTEGER*8 N,KD,LDAB  
REAL*8    AB(25,100),X(100),Y(100)  
N = 75  
KD = 15  
LDAB = 25  
CALL SSBMV ('UPPER', N, KD, 1.0, AB, LDAB, X, 1, 1.0, Y, 1)
```

Basic Matrix Operations  
**SSPMV/CHPMV – Matrix-Vector Multiply**

**NAME** SSPMV/CHPMV – Matrix-Vector Multiply

**Purpose**

These subprograms compute the matrix-vector product  $Ax$  where  $A$  is an  $n$  by  $n$  real symmetric or complex Hermitian matrix stored in packed form as described in “Matrix Storage,” and  $x$  is a real or complex  $n$ -vector. The product may be stored in the result array, or, optionally, be added to or subtracted from it. This is handled in a convenient, but general, way by two scalar arguments,  $\alpha$  and  $\beta$ , which are used as multipliers of the matrix-vector product and the result vector. Specifically, these subprograms compute the matrix-vector product of the form

$$y \leftarrow \alpha Ax + \beta y.$$

The structure of  $A$  is indicated by the name of the subprogram used:

SSPMV	$A$ is a real symmetric matrix
CHPMV	$A$ is a complex Hermitian matrix

**Matrix Storage**

Because either triangle of  $A$  may be obtained from the other, you only need to provide one triangle of  $A$ , either the upper or the lower triangle. Compared to storing the entire matrix, you save memory by supplying that triangle stored column-by-column in packed form in a 1-dimensional array.

The following examples illustrate the packed storage of symmetric or Hermitian matrices.

**Upper triangular storage**

If the upper triangle of  $A$  is

11	12	13	14
	22	23	24
		33	34
			44

then  $A$  is packed column-by-column into an array  $ap$  as follows:

$k$	1	2	3	4	5	6	7	8	9	10
$ap(k)$	11	12	22	13	23	33	14	24	34	44

Upper triangular matrix element  $a_{ij}$  is stored in array element  $ap(i+j \times (j-1)/2)$ .

### Lower triangular storage

If the lower triangle of  $A$  is

```

      11
      21  22
      31  32  33
      41  42  43  44
  
```

then  $A$  is packed column-by-column into an array  $\mathbf{ap}$  as follows:

$k$	1	2	3	4	5	6	7	8	9	10
$\mathbf{ap}(k)$	11	21	31	41	22	32	42	33	43	44

Lower triangular matrix element  $a_{ij}$  is stored in array element  $\mathbf{ap}(i+(j-1)\times(2n-j)/2)$ .

### Usage

SCILIB:

```

CHARACTER*1      uplo
INTEGER*8        n, incx, incy
REAL*8          alpha, beta, ap(lenap), x(lenx), y(leny)
CALL SSPMV(uplo, n, alpha, ap, x, incx, beta, y, incy)

CHARACTER*1      uplo
INTEGER*8        n, incx, incy
COMPLEX*16      alpha, beta, ap(lenap), x(lenx), y(leny)
CALL CHPMV(uplo, n, alpha, ap, x, incx, beta, y, incy)
  
```

### Input

**uplo**            Upper/lower triangular option for  $A$ :  
           'L' or 'l'     The lower triangle of  $A$  is stored in the packed array.  
           'U' or 'u'     The upper triangle of  $A$  is stored in the packed array.

**n**                Number of rows and columns in matrix  $A$ ,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference  $\mathbf{ap}$ ,  $\mathbf{x}$ , or  $\mathbf{y}$ .

**alpha**            The scalar  $\alpha$ . If  $\mathbf{alpha} = 0$ , the subprograms compute  $y \leftarrow \beta y$  without referencing  $\mathbf{ap}$  or  $\mathbf{x}$ .

**ap**                Array of length  $\mathbf{lenap} = n \times (n+1)/2$  containing the upper or lower triangle, as specified by **uplo**, of an  $n$ -by- $n$  real symmetric or complex Hermitian matrix  $A$ , stored by columns in the packed form described above.

Basic Matrix Operations  
SSPMV/CHPMV – Matrix-Vector Multiply

**x** Array of length  $\text{lenx} = (\mathbf{n}-1) \times |\mathbf{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx** Increment for the array  $\mathbf{x}$ ,  $\mathbf{incx} \neq 0$ :

**incx** > 0  $x$  is stored forward in array  $\mathbf{x}$ , i.e.,  $x_i$  is stored in  $\mathbf{x}((i-1) \times \mathbf{incx} + 1)$ .

**incx** < 0  $x$  is stored backward in array  $\mathbf{x}$ , i.e.,  $x_i$  is stored in  $\mathbf{x}((i-\mathbf{n}) \times \mathbf{incx} + 1)$ .

Use  $\mathbf{incx} = 1$  if the vector  $x$  is stored contiguously in array  $\mathbf{x}$ , i.e., if  $x_i$  is stored in  $\mathbf{x}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

**beta** The scalar  $\beta$ .

**y** Array of length  $\text{leny} = (\mathbf{n}-1) \times |\mathbf{incy}| + 1$  containing the  $n$ -vector  $y$ . Not used as input if  $\mathbf{beta} = 0$ .

**incy** Increment for the array  $\mathbf{y}$ ,  $\mathbf{incy} \neq 0$ :

**incy** > 0  $y$  is stored forward in array  $\mathbf{y}$ , i.e.,  $y_i$  is stored in  $\mathbf{y}((i-1) \times \mathbf{incy} + 1)$ .

**incy** < 0  $y$  is stored backward in array  $\mathbf{y}$ , i.e.,  $y_i$  is stored in  $\mathbf{y}((i-\mathbf{n}) \times \mathbf{incy} + 1)$ .

Use  $\mathbf{incy} = 1$  if the vector  $y$  is stored contiguously in array  $\mathbf{y}$ , i.e., if  $y_i$  is stored in  $\mathbf{y}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

## Output

**y** The updated  $y$  vector replaces the input.

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo**  $\neq$  'L' or 'l' or 'U' or 'u',  
**n** < 0,  
**incx** = 0, and  
**incy** = 0.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **uplo** argument as 'LOWER' for 'L' or 'UPPER' for 'U'. Refer to "Example 2."

### Example 1

Form the REAL\*8 matrix-vector product  $y = Ax$ , where  $A$  is a 9-by-9 real symmetric matrix whose upper triangle is stored in packed form in an array AP of dimension 55,  $x$  is a real vector 9 elements long stored in an array X of dimension 10, and  $y$  is a real vector 9 elements long stored in an array Y, also of dimension 10.

```
CHARACTER*1 UPLO
INTEGER*8   N, INCX, INCY
REAL*8     ALPHA, BETA, AP(55), X(10), Y(10)
UPLO = 'U'
N = 9
ALPHA = 1.0
BETA = 0.0
INCX = 1
INCY = 1
CALL SSPMV (UPLO, N, ALPHA, AP, X, INCX, BETA, Y, INCY)
```

### Example 2

Form the COMPLEX\*16 matrix-vector product  $y = \frac{1}{2}y - \rho Ax$ , where  $\rho$  is a complex scalar,  $A$  is a 9-by-9 complex Hermitian matrix whose lower triangle is stored in packed form in an array AP of dimension 55,  $x$  is a complex vector 9 elements long stored in an array X of dimension 10, and  $y$  is a complex vector 9 elements long stored in an array Y, also of dimension 10.

```
INTEGER*8   N
COMPLEX*16 RHO, AP(55), X(10), Y(10)
N = 9
CALL CHPMV ('LOWER', N, -RHO, AP, X, 1, (0.5, 0.0), Y, 1)
```

Basic Matrix Operations  
**SSPR/CHPR – Rank-1 Update**

**NAME** SSPR/CHPR – Rank-1 Update

**Purpose**

These subprograms compute the real symmetric or complex Hermitian rank-1 update

$$A \leftarrow \alpha x x^* + A,$$

where  $A$  is an  $n$ -by- $n$  real symmetric or complex Hermitian matrix stored in packed form as described in “Matrix Storage,”  $\alpha$  is a real scalar,  $x$  is a real or complex  $n$ -vector, and  $x^*$  is the conjugate transpose of  $x$ . (The conjugate transpose of a real vector is simply the transpose.)

The structure of  $A$  is indicated by the name of the subprogram used:

SSPR         $A$  is a real symmetric matrix  
 CHPR         $A$  is a complex Hermitian matrix

**Matrix Storage**

Because either triangle of  $A$  may be obtained from the other, you only need to provide one triangle of  $A$ , either the upper or the lower triangle. Compared to storing the entire matrix, you save memory by supplying that triangle stored column-by-column in packed form in a 1-dimensional array.

The following examples illustrate the packed storage of symmetric or Hermitian matrices.

**Upper triangular storage**

If the upper triangle of  $A$  is

11	12	13	14
	22	23	24
		33	34
			44

then  $A$  is packed column-by-column into an array **ap** as follows:

$k$	1	2	3	4	5	6	7	8	9	10
<b>ap</b> ( $k$ )	11	12	22	13	23	33	14	24	34	44

Upper triangular matrix element  $a_{ij}$  is stored in array element **ap**( $i+j \times (j-1)/2$ ).

**Lower triangular storage**

If the lower triangle of  $A$  is

```

11
21 22
31 32 33
41 42 43 44

```

then  $A$  is packed column-by-column into an array  $\mathbf{ap}$  as follows:

$k$	1	2	3	4	5	6	7	8	9	10
$\mathbf{ap}(k)$	11	21	31	41	22	32	42	33	43	44

Lower triangular matrix element  $a_{ij}$  is stored in array element  $\mathbf{ap}(i+(j-1)\times(2n-j)/2)$ .

## Usage

SCILIB:

```

CHARACTER*1      uplo
INTEGER*8        n, incx
REAL*8           alpha, ap(lenap), x(lenx)
CALL SSPR(uplo, n, alpha, x, incx, ap)

CHARACTER*1      uplo
INTEGER*8        n, incx
REAL*8           alpha
COMPLEX*16       ap(lenap), x(lenx)
CALL CHPR(uplo, n, alpha, x, incx, ap)

```

## Input

**uplo**            Upper/lower triangular option for  $A$ :  
    'L' or 'l'        The lower triangle of  $A$  is stored in the packed array.  
    'U' or 'u'        The upper triangle of  $A$  is stored in the packed array.

**n**                Number of rows and columns in matrix  $A$  and elements of vector  $x$ ,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference  $\mathbf{ap}$  or  $\mathbf{x}$ .

**alpha**            The scalar  $\alpha$ . If  $\mathbf{alpha} = 0$ , the subprograms do not reference  $\mathbf{ap}$  or  $\mathbf{x}$ .

**x**                Array of length  $\mathbf{lenx} = (n-1) \times |\mathbf{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx**            Increment for the array  $\mathbf{x}$ ,  $\mathbf{incx} \neq 0$ :

Basic Matrix Operations  
SSPR/CHPR – Rank-1 Update

**incx** > 0      $x$  is stored forward in array  $\mathbf{x}$ , i.e.,  $x_i$  is stored in  $\mathbf{x}((i-1)\times\mathbf{incx}+1)$ .

**incx** < 0      $x$  is stored backward in array  $\mathbf{x}$ , i.e.,  $x_i$  is stored in  $\mathbf{x}((i-\mathbf{n})\times\mathbf{incx}+1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $\mathbf{x}$ , i.e., if  $x_i$  is stored in  $\mathbf{x}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

**ap**     Array of length **lenap** =  $\mathbf{n}\times(\mathbf{n}+1)/2$  containing the upper or lower triangle, as specified by **uplo**, of an  $n$ -by- $n$  real symmetric or complex Hermitian matrix  $A$ , stored by columns in the packed form described above.

## Output

**ap**     The upper or lower triangle of the updated  $A$  matrix, as specified by **uplo**, replaces the input.

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo** ≠ ‘L’ or ‘l’ or ‘U’ or ‘u’,  
**n** < 0, and  
**incx** = 0.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the CALL statement may be improved by coding the **uplo** argument as ‘LOWER’ for ‘L’ or ‘UPPER’ for ‘U’. Refer to “Example 2.”

## Example 1

Apply a REAL\*8 symmetric rank-1 update  $xx^T$  to  $A$ , where  $A$  is a 9-by-9 real symmetric matrix whose upper triangle is stored in packed form in an array AP of dimension 55, and  $x$  is a real vector 9 elements long stored in an array X of dimension 10.

```
CHARACTER*1 UPLO
INTEGER*8   N, INCX
REAL*8     ALPHA, AP(55), X(10)
UPLO = 'U'
N = 9
ALPHA = 1.0
INCX = 1
CALL SSPR (UPLO, N, ALPHA, X, INCX, AP)
```

## Example 2

Apply a COMPLEX\*16 Hermitian rank-1 update  $-2xx^*$  to  $A$ , where  $A$  is a 9-by-9 complex Hermitian matrix whose lower triangle is stored in packed form in an array AP of dimension 55, and  $x$  is a complex vector 9 elements long stored in an array X of dimension 10.

```
INTEGER*8   N
COMPLEX*16 AP(55), X(10)
N = 9
CALL CHPR ('LOWER', N, -2.0, X, 1, AP)
```

Basic Matrix Operations  
SSPR2/CHPR2 – Rank-2 Update

**NAME**           SSPR2/CHPR2 – Rank-2 Update

**Purpose**

These subprograms compute the real symmetric or complex Hermitian rank-2 update

$$A \leftarrow \alpha xy^* + \bar{\alpha}yx^* + A,$$

where  $A$  is an  $n$ -by- $n$  real symmetric or complex Hermitian matrix stored in packed form as described in “Matrix Storage,”  $\alpha$  is a complex scalar,  $\bar{\alpha}$  is the complex conjugate of  $\alpha$ ,  $x$  and  $y$  are real or complex  $n$ -vectors, and  $x^*$  and  $y^*$  are the conjugate transposes of  $x$  and  $y$ , respectively. (The conjugate of a real scalar is just the scalar, and the conjugate transpose of a real vector is simply the transpose.)

The structure of  $A$  is indicated by the name of the subprogram used:

SSPR2	$A$ is a real symmetric matrix
CHPR2	$A$ is a complex Hermitian matrix

**Matrix Storage**

Because either triangle of  $A$  may be obtained from the other, you only need to provide one triangle of  $A$ , either the upper or the lower triangle. Compared to storing the entire matrix, you save memory by supplying that triangle stored column-by-column in packed form in a 1-dimensional array.

The following examples illustrate the packed storage of symmetric or Hermitian matrices.

**Upper triangular storage**

If the upper triangle of  $A$  is

11	12	13	14
	22	23	24
		33	34
			44

then  $A$  is packed column-by-column into an array  $\mathbf{ap}$  as follows:

$k$	1	2	3	4	5	6	7	8	9	10
$\mathbf{ap}(k)$	11	12	22	13	23	33	14	24	34	44

Upper triangular matrix element  $a_{ij}$  is stored in array element  $\mathbf{ap}(i+j \times (j-1)/2)$ .

## Lower triangular storage

If the lower triangle of  $A$  is

11										
21	22									
31	32	33								
41	42	43	44							

then  $A$  is packed column-by-column into an array  $\mathbf{ap}$  as follows:

$k$	1	2	3	4	5	6	7	8	9	10
$\mathbf{ap}(k)$	11	21	31	41	22	32	42	33	43	44

Lower triangular matrix element  $a_{ij}$  is stored in array element  $\mathbf{ap}(i+(j-1)\times(2n-j)/2)$ .

## Usage

SCILIB:

```
CHARACTER*1      uplo
INTEGER*8        n, incx, incy
REAL*8           alpha, ap(lenap), x(lenx), y(leny)
CALL SSPR2(uplo, n, alpha, x, incx, y, incy, ap)
```

```
CHARACTER*1      uplo
INTEGER*8        n, incx, incy
COMPLEX*16       alpha, ap(lenap), x(lenx), y(leny)
CALL CHPR2(uplo, n, alpha, x, incx, y, incy, ap)
```

## Input

<b>uplo</b>	Upper/lower triangular option for $A$ : 'L' or 'l'   The lower triangle of $A$ is stored in the packed array. 'U' or 'u'   The upper triangle of $A$ is stored in the packed array.
<b>n</b>	Number of rows and columns in matrix $A$ and elements of vectors $x$ and $y$ , $n \geq 0$ . If $n = 0$ , the subprograms do not reference $\mathbf{ap}$ , $\mathbf{x}$ , or $\mathbf{y}$ .
<b>alpha</b>	The scalar $\alpha$ . If $\mathbf{alpha} = 0$ , the subprograms do not reference $\mathbf{ap}$ , $\mathbf{x}$ , or $\mathbf{y}$ .
<b>x</b>	Array of length $\mathbf{lenx} = (n-1)\times \mathbf{incx} +1$ containing the $n$ -vector $x$ .

Basic Matrix Operations  
SSPR2/CHPR2 – Rank-2 Update

<b>incx</b>	Increment for the array <b>x</b> , <b>incx</b> $\neq$ 0: <b>incx</b> > 0 $x$ is stored forward in array <b>x</b> , i.e., $x_i$ is stored in $\mathbf{x}((i-1)\times\mathbf{incx}+1)$ . <b>incx</b> < 0 $x$ is stored backward in array <b>x</b> , i.e., $x_i$ is stored in $\mathbf{x}((i-\mathbf{n})\times\mathbf{incx}+1)$ .  Use <b>incx</b> = 1 if the vector $x$ is stored contiguously in array <b>x</b> , i.e., if $x_i$ is stored in $\mathbf{x}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.
<b>y</b>	Array of length <b>leny</b> = $(\mathbf{n}-1)\times \mathbf{incy} +1$ containing the $n$ -vector $y$ .
<b>incy</b>	Increment for the array <b>y</b> , <b>incy</b> $\neq$ 0: <b>incy</b> > 0 $y$ is stored forward in array <b>y</b> , i.e., $y_i$ is stored in $\mathbf{y}((i-1)\times\mathbf{incy}+1)$ . <b>incy</b> < 0 $y$ is stored backward in array <b>y</b> , i.e., $y_i$ is stored in $\mathbf{y}((i-\mathbf{n})\times\mathbf{incy}+1)$ .  Use <b>incy</b> = 1 if the vector $y$ is stored contiguously in array <b>y</b> , i.e., if $y_i$ is stored in $\mathbf{y}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.
<b>ap</b>	Array of length <b>lenap</b> = $\mathbf{n}\times(\mathbf{n}+1)/2$ containing the upper or lower triangle, as specified by <b>uplo</b> , of an $n$ -by- $n$ real symmetric or complex Hermitian matrix $A$ , stored by columns in the packed form described above.

## Output

<b>ap</b>	The upper or lower triangle of the updated $A$ matrix, as specified by <b>uplo</b> , replaces the input.
-----------	--

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo**  $\neq$  'L' or 'l' or 'U' or 'u',  
**n** < 0,  
**incx** = 0, and

**incy = 0.**

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **uplo** argument as 'LOWER' for 'L' or 'UPPER' for 'U'. Refer to "Example 2."

### Example 1

Apply a REAL\*8 symmetric rank-2 update  $xy^T + x^T y$  to  $A$ , where  $A$  is a 9-by-9 real symmetric matrix whose upper triangle is stored in packed form in an array  $AP$  of dimension 55,  $x$  is a real vector 9 elements long stored in an array  $X$  of dimension 10, and  $y$  is a real vector 9 elements long stored in an array  $Y$  also of dimension 10.

```

CHARACTER*1 UPLO
INTEGER*8   N, INCX, INCY
REAL*8     ALPHA, AP(55), X(10), Y(10)
UPLO = 'U'
N = 9
ALPHA = 1.0
INCX = 1
INCY = 1
CALL SSPR2 (UPLO, N, ALPHA, X, INCX, Y, INCY, AP)

```

### Example 2

Apply a COMPLEX\*16 Hermitian rank-2 update  $\alpha xy^* + \bar{\alpha} yx^*$  to  $A$ , where  $A$  is a 9-by-9 complex Hermitian matrix whose lower triangle is stored in packed form in an array  $AP$  of dimension 55,  $\alpha$  is a complex scalar,  $x$  is a complex vector 9 elements long stored in an array  $X$  of dimension 10, and  $y$  is a complex vector 9 elements long stored in an array  $Y$  of dimension 10.

```

INTEGER*8   N
COMPLEX*16 ALPHA, AP(55), X(10), Y(10)
N = 9
CALL CHPR2 ('LOWER', N, ALPHA, X, 1, Y, 1, AP)

```

Basic Matrix Operations  
SSYMM/CHEMM/CSYMM – Matrix-Matrix Multiply

**NAME** SSYMM/CHEMM/CSYMM – Matrix-Matrix Multiply

**Purpose**

These subprograms compute the matrix-matrix products  $AB$  and  $BA$ , where  $A$  is a real symmetric, complex symmetric, or complex Hermitian matrix and  $B$  is an  $m$ -by- $n$  matrix. The size of  $A$ , either  $m$  by  $m$  or  $n$  by  $n$ , depends on which matrix product is requested. The product may be stored in the result matrix (which is always of size  $m$  by  $n$ ) or, optionally, may be added to or subtracted from it. This is handled in a convenient, but general, way by two scalar arguments,  $\alpha$  and  $\beta$ , which are used as multipliers of the matrix product and the result matrix. Specifically, these subprograms compute matrix products of the forms

$$C \leftarrow \alpha AB + \beta C \quad \text{and} \quad C \leftarrow \alpha BA + \beta C.$$

The structure of  $A$  is indicated by the name of the subprogram used:

SSYMM	$A$ is a real symmetric matrix
CHEMM	$A$ is a complex Hermitian matrix
CSYMM	$A$ is a complex symmetric matrix

**Matrix Storage**

Because either triangle of  $A$  may be obtained from the other, you only need to provide one triangle of  $A$ . You may supply either the upper or the lower triangle of  $A$ , in a two-dimensional array large enough to hold the entire matrix. The other triangle of the array is not referenced.

**Usage**

SCILIB:

CHARACTER*1	side, uplo
INTEGER*8	m, n, lda, ldb, ldc
REAL*8	alpha, beta, a(lda, *), b(ldb, *), c(ldc, n)
CALL SSYMM(side, uplo, m, n, alpha, a, lda, b, ldb, beta, c, ldc)	
CHARACTER*1	side, uplo
INTEGER*8	m, n, lda, ldb, ldc
COMPLEX*16	alpha, beta, a(lda, *), b(ldb, *), c(ldc, n)
CALL CHEMM(side, uplo, m, n, alpha, a, lda, b, ldb, beta, c, ldc)	
CHARACTER*1	side, uplo
INTEGER*8	m, n, lda, ldb, ldc
COMPLEX*16	alpha, beta, a(lda, *), b(ldb, *), c(ldc, n)
CALL CSYMM(side, uplo, m, n, alpha, a, lda, b, ldb, beta, c, ldc)	

## Input

<b>side</b>	Specifies whether symmetric or Hermitian matrix $A$ is the left or right matrix operand:  'L' or 'l' $A$ is the left matrix operand, i.e. compute $C \leftarrow \alpha AB + \beta C$ 'R' or 'r' $A$ is the right matrix operand, i.e. compute $C \leftarrow \alpha BA + \beta C$
<b>uplo</b>	Upper/lower triangular storage option for $A$ : 'L' or 'l'      Reference only the lower triangle of $A$ 'U' or 'u'      Reference only the upper triangle of $A$
<b>m</b>	Number of rows in matrix $C$ , $m \geq 0$ . If $m = 0$ , the subprograms do not reference $a$ , $b$ , or $c$ .
<b>n</b>	Number of columns in matrix $B$ , $n \geq 0$ . If $n = 0$ , the subprograms do not reference $a$ , $b$ , or $c$ .
<b>alpha</b>	The scalar $\alpha$ . If <b>alpha</b> = 0, the subprograms compute $C \leftarrow \beta C$ without referencing $a$ or $b$ .
<b>a</b>	Array whose upper or lower triangle, as specified by <b>uplo</b> , contains the upper or lower triangle of the matrix $A$ . The other triangle of $a$ is not referenced. The size of $A$ is indicated by <b>side</b> :  'L' or 'l' $A$ is $m$ by $m$ 'R' or 'r' $A$ is $n$ by $n$
<b>lda</b>	The leading dimension of array $a$ as declared in the calling program unit, with $lda \geq \max$ (the number of rows of $A$ , 1).
<b>b</b>	Array containing the $m$ -by- $n$ matrix $B$ .
<b>ldb</b>	The leading dimension of array $b$ as declared in the calling program unit, with $ldb \geq \max(m, 1)$ .
<b>beta</b>	The scalar $\beta$ .
<b>c</b>	Array containing the $m$ -by- $n$ matrix $C$ . Not used as input if <b>beta</b> = 0.
<b>ldc</b>	The leading dimension of array $c$ as declared in the calling program unit, with $ldc \geq \max(m, 1)$ .

## Output

<b>c</b>	The updated $C$ matrix replaces the input.
----------	--

## Notes

These subprograms conform to specifications of the Level 3 BLAS.

**SSYMM/CHEMM/CSYMM – Matrix-Matrix Multiply**

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**side** ≠ 'L' or 'l' or 'R' or 'r',  
**uplo** ≠ 'L' or 'l' or 'U' or 'u',  
**m** < 0,  
**n** < 0,  
**lda** too small,  
**ldb** < max(**m**,1), and  
**ldc** < max(**m**,1).

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved, for example, by coding the **side** argument as 'LEFT' for 'L' or 'RIGHT' for 'R'. Refer to "Example 2."

**Example 1**

Form the **REAL\*8** matrix product  $C = AB$ , where  $A$  is a 6-by-6 real symmetric real matrix whose upper triangle is stored in the upper triangle of an array  $A$  of dimension 10 by 10,  $B$  is a 6-by-8 real matrix stored in an array  $B$  of dimension 10 by 10, and  $C$  is a 6-by-8 real matrix stored in an array  $C$ , also of dimension 10 by 10.

```
CHARACTER*1 SIDE,UPLO
INTEGER*8   M,N,LDA,LDB,LDC
REAL*8     ALPHA,BETA,A(10,10),B(10,10),C(10,10)
SIDE = 'L'
UPLO = 'U'
M = 6
N = 8
ALPHA = 1.0
BETA = 0.0
LDA = 10
LDB = 10
LDC = 10
CALL SSYMM (SIDE,UPLO,M,N,ALPHA,A,LDA,B,LDB,BETA,C,LDC)
```

**Example 2**

Form the **COMPLEX\*16** matrix-matrix product  $C = \frac{1}{2}BA - \rho C$ , where  $\rho$  is a scalar,  $A$  is an 8-by-8 complex Hermitian matrix whose lower triangle is stored in the lower triangle of an array  $A$  of dimension 10 by 10,  $B$  is a 6-by-8 complex matrix stored in an array whose dimensions are 10 by 10, and  $C$  is a 6-by-8 complex matrix stored in an array  $C$ , also of dimension 10 by 10.

Basic Matrix Operations  
**SSYMM/CHEMM/CSYMM – Matrix-Matrix Multiply**

```
INTEGER*8  M,N,LDA,LDB,LDC
COMPLEX*16 HALF,RHO,A(10,10),B(10,10),C(10,10)
M = 6
N = 8
LDA = 10
LDB = 10
LDC = 10
HALF = (0.5,0.0)
CALL CHEMM ('RIGHT', 'LOWER', M, N, -RHO, A, LDA, B, LDB, HALF, C,
           LDC)
```

Basic Matrix Operations  
SSYMV/CHEMV – Matrix-Vector Multiply

**NAME** SSYMV/CHEMV – Matrix-Vector Multiply

**Purpose**

These subprograms compute the matrix-vector product  $Ax$  where  $A$  is an  $n$ -by- $n$  real symmetric or complex Hermitian matrix and  $x$  is a real or complex  $n$ -vector. The product may be stored in the result array, or, optionally, be added to or subtracted from it. This is handled in a convenient, but general, way by two scalar arguments,  $\alpha$  and  $\beta$ , which are used as multipliers of the matrix-vector product and the result vector. Specifically, these subprograms compute the matrix-vector product of the form

$$y \leftarrow \alpha Ax + \beta y.$$

The structure of  $A$  is indicated by the name of the subprogram used:

SSYMV       $A$  is a real symmetric matrix  
CHEMV       $A$  is a complex Hermitian matrix

**Matrix Storage**

Because either triangle of  $A$  may be obtained from the other, you only need to provide one triangle of  $A$ . You may supply either the upper or the lower triangle of  $A$ , in a two-dimensional array large enough to hold the entire matrix. The other triangle of the array is not referenced.

**Usage**

SCILIB:

```
CHARACTER*1      uplo
INTEGER*8        n, lda, incx, incy
REAL*8          alpha, beta, a(lda, n), x(lenx), y(leny)
CALL SSYMV(uplo, n, alpha, a, lda, x, incx, beta, y, incy)

CHARACTER*1      uplo
INTEGER*8        n, lda, incx, incy
COMPLEX*16      alpha, beta, a(lda, n), x(lenx), y(leny)
CALL CHEMV(uplo, n, alpha, a, lda, x, incx, beta, y, incy)
```

**Input**

**uplo**      Upper/lower triangular option for  $A$ :  
            'L' or 'l'      Reference only the lower triangle of  $A$ .  
            'U' or 'u'      Reference only the upper triangle of  $A$ .

**n**          Number of rows and columns in matrix  $A$ ,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference  $a$ ,  $x$ , or  $y$ .

<b>alpha</b>	The scalar $\alpha$ . If <b>alpha</b> = 0, the subprograms compute $y \leftarrow \beta y$ without referencing <b>a</b> or <b>x</b> .
<b>a</b>	Array whose upper or lower triangle, as specified by <b>uplo</b> , contains the upper or lower triangle of an $n$ -by- $n$ real symmetric or complex Hermitian matrix $A$ . The other triangle of <b>a</b> is not referenced.
<b>lda</b>	The leading dimension of array <b>a</b> as declared in the calling program unit, with $lda \geq \max(n,1)$ .
<b>x</b>	Array of length $lenx = (n-1) \times  incx  + 1$ containing the $n$ -vector $x$ .
<b>incx</b>	Increment for the array <b>x</b> , <b>incx</b> $\neq$ 0: <div style="margin-left: 20px;"> <b>incx</b> &gt; 0     <math>x</math> is stored forward in array <b>x</b>, i.e., <math>x_i</math> is stored in <math>x((i-1) \times incx + 1)</math>.  <b>incx</b> &lt; 0     <math>x</math> is stored backward in array <b>x</b>, i.e., <math>x_i</math> is stored in <math>x((i-n) \times incx + 1)</math>. </div> <p>Use <b>incx</b> = 1 if the vector <math>x</math> is stored contiguously in array <b>x</b>, i.e., if <math>x_i</math> is stored in <math>x(i)</math>. Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.</p>
<b>beta</b>	The scalar $\beta$ .
<b>y</b>	Array of length $leny = (n-1) \times  incy  + 1$ containing the $n$ -vector $y$ . Not used as input if <b>beta</b> = 0.
<b>incy</b>	Increment for the array <b>y</b> , <b>incy</b> $\neq$ 0: <div style="margin-left: 20px;"> <b>incy</b> &gt; 0     <math>y</math> is stored forward in array <b>y</b>, i.e., <math>y_i</math> is stored in <math>y((i-1) \times incy + 1)</math>.  <b>incy</b> &lt; 0     <math>y</math> is stored backward in array <b>y</b>, i.e., <math>y_i</math> is stored in <math>y((i-n) \times incy + 1)</math>. </div> <p>Use <b>incy</b> = 1 if the vector <math>y</math> is stored contiguously in array <b>y</b>, i.e., if <math>y_i</math> is stored in <math>y(i)</math>. Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.</p>

## Output

**y**            The updated  $y$  vector replaces the input.

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

Basic Matrix Operations  
SSYMV/CHEMV – Matrix-Vector Multiply

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo** ≠ 'L' or 'l' or 'U' or 'u',  
**n** < 0,  
**lda** < max(**n**,1),  
**incx** = 0, and  
**incy** = 0.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **uplo** argument as 'LOWER' for 'L' or 'UPPER' for 'U'. Refer to "Example 2."

### Example 1

Form the REAL\*8 matrix-vector product  $y = Ax$ , where A is a 9-by-9 real symmetric matrix whose upper triangle is stored in the upper triangle of an array A whose dimensions are 10 by 10, x is a real vector 9 elements long stored in an array X of dimension 10, and y is a real vector 9 elements long stored in an array Y, also of dimension 10.

```
CHARACTER*1 UPLO
INTEGER*8   N,LDA, INCX, INCY
REAL*8     ALPHA, BETA, A(10,10), X(10), Y(10)
UPLO = 'U'
N = 9
ALPHA = 1.0
BETA = 0.0
LDA = 10
INCX = 1
INCY = 1
CALL SSYMV (UPLO, N, ALPHA, A, LDA, X, INCX, BETA, Y, INCY)
```

### Example 2

Form the COMPLEX\*16 matrix-vector product  $y = \frac{1}{2}y - \rho Ax$ , where  $\rho$  is a complex scalar, A is a 9-by-9 complex Hermitian matrix whose lower triangle is stored in the lower triangle of an array A whose dimensions are 10 by 10, x is a complex vector 9 elements long stored in an array X of dimension 10, and y is a complex vector 9 elements long stored in an array Y, also of dimension 10.

```
INTEGER*8   N,LDA
COMPLEX*16 RHO, A(10,10), X(10), Y(10)
N = 9
LDA = 10
CALL CHEMV ('LOWER', N, -RHO, A, LDA, X, 1, (0.5, 0.0), Y, 1)
```

**NAME** SSYR/CHER – Rank-1 Update

**Purpose**

These subprograms compute the real symmetric or complex Hermitian rank-1 update

$$A \leftarrow \alpha x x^* + A,$$

where  $A$  is an  $n$ -by- $n$  real symmetric or complex Hermitian matrix,  $\alpha$  is a real scalar,  $x$  is a real or complex  $n$ -vector, and  $x^*$  is the conjugate transpose of  $x$ . (The conjugate transpose of a real vector is simply the transpose.)

The structure of  $A$  is indicated by the name of the subprogram used:

SSYR       $A$  is a real symmetric matrix  
CHER       $A$  is a complex Hermitian matrix

**Matrix Storage**

Because either triangle of  $A$  may be obtained from the other, these subprograms reference and apply the update to only one triangle of  $A$ . You may supply either the upper or the lower triangle of  $A$ , in a two-dimensional array large enough to hold the entire matrix, and the same triangle of the updated matrix is returned in the array. The other triangle of the array is not referenced.

**Usage**

SCILIB:

```

CHARACTER*1      uplo
INTEGER*8        n, lda, incx
REAL*8           alpha, a(lda, n), x(lenx)
CALL SSYR(uplo, n, alpha, x, incx, a, lda)

CHARACTER*1      uplo
INTEGER*8        n, lda, incx
REAL*8           alpha
COMPLEX*16       a(lda, n), x(lenx)
CALL CHER(uplo, n, alpha, x, incx, a, lda)

```

**Input**

**uplo**            Upper/lower triangular option for  $A$ :  
                  'L' or 'l'      Reference and update only the lower triangle of  $A$ .  
                  'U' or 'u'      Reference and update only the upper triangle of  $A$ .

Basic Matrix Operations  
SSYR/CHER – Rank-1 Update

<b>n</b>	Number of rows and columns in matrix $A$ and elements of vector $x$ , $n \geq 0$ . If $n = 0$ , the subprograms do not reference $a$ or $x$ .
<b>alpha</b>	The scalar $\alpha$ . If <b>alpha</b> = 0, the subprograms do not reference $a$ or $x$ .
<b>x</b>	Array of length $\text{lenx} = (n-1) \times  \text{incx}  + 1$ containing the $n$ -vector $x$ .
<b>incx</b>	Increment for the array $x$ , <b>incx</b> $\neq 0$ : <b>incx</b> > 0 $x$ is stored forward in array $x$ , i.e., $x_i$ is stored in $x((i-1) \times \text{incx} + 1)$ . <b>incx</b> < 0 $x$ is stored backward in array $x$ , i.e., $x_i$ is stored in $x((i-n) \times \text{incx} + 1)$ .  Use <b>incx</b> = 1 if the vector $x$ is stored contiguously in array $x$ , i.e., if $x_i$ is stored in $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.
<b>a</b>	Array whose upper or lower triangle, as specified by <b>uplo</b> , contains the upper or lower triangle of an $n$ -by- $n$ real symmetric or complex Hermitian matrix $A$ . The other triangle of $a$ is not referenced.
<b>lda</b>	The leading dimension of array $a$ as declared in the calling program unit, with $\text{lda} \geq \max(n, 1)$ .

## Output

<b>a</b>	The upper or lower triangle of the updated $A$ matrix, as specified by <b>uplo</b> , replaces the upper or lower triangle of the input, respectively. The other triangle of $a$ is unchanged.
----------	---

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo**  $\neq$  'L' or 'l' or 'U' or 'u',  
**n** < 0,  
**lda** <  $\max(n, 1)$ , and  
**incx** = 0.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **uplo** argument as **'LOWER'** for **'L'** or **'UPPER'** for **'U'**. Refer to “Example 2.”

### Example 1

Apply a **REAL\*8** symmetric rank-1 update  $xx^T$  to **A**, where **A** is a 9-by-9 real symmetric matrix whose upper triangle is stored in the upper triangle of an array **A** whose dimensions are 10 by 10, and **x** is a real vector 9 elements long stored in an array **X** of dimension 10.

```

CHARACTER*1 UPLO
INTEGER*8   N,LDA, INCX
REAL*8     ALPHA,A(10,10),X(10)
UPLO = 'U'
N = 9
ALPHA = 1.0
LDA = 10
INCX = 1
CALL SSYR (UPLO,N,ALPHA,X, INCX,A,LDA)

```

### Example 2

Apply a **COMPLEX\*16** Hermitian rank-1 update  $-2xx^*$  to **A**, where **A** is a 9-by-9 complex Hermitian matrix whose lower triangle is stored in the lower triangle of an array **A** whose dimensions are 10 by 10, and **x** is a complex vector 9 elements long stored in an array **X** of dimension 10.

```

INTEGER*8   N,LDA
COMPLEX*16 A(10,10),X(10)
N = 9
LDA = 10
CALL CHER ('LOWER',N,-2.0,X,1,A,LDA)

```

Basic Matrix Operations  
SSYR2/CHER2 – Rank-2 Update

**NAME** SSYR2/CHER2 – Rank-2 Update

**Purpose**

These subprograms compute the real symmetric or complex Hermitian rank-2 update

$$A \leftarrow \alpha xy^* + \bar{\alpha}yx^* + A,$$

where  $A$  is an  $n$ -by- $n$  real symmetric or complex Hermitian matrix,  $\alpha$  is a complex scalar,  $\bar{\alpha}$  is the complex conjugate of  $\alpha$ ,  $x$  and  $y$  are real or complex  $n$ -vectors, and  $x^*$  and  $y^*$  are the conjugate transposes of  $x$  and  $y$ , respectively. (The conjugate of a real scalar is just the scalar, and the conjugate transpose of a real vector is simply the transpose.)

The structure of  $A$  is indicated by the name of the subprogram used:

SSYR2       $A$  is a real symmetric matrix  
CHER2       $A$  is a complex Hermitian matrix

**Matrix Storage**

Because either triangle of  $A$  may be obtained from the other, these subprograms reference and apply the update to only one triangle of  $A$ . You may supply either the upper or the lower triangle of  $A$ , in a two-dimensional array large enough to hold the entire matrix, and the same triangle of the updated matrix is returned in the array. The other triangle of the array is not referenced.

**Usage**

SCILIB:

```
CHARACTER*1      uplo
INTEGER*8        n, lda, incx, incy
REAL*8           alpha, a(lda, n), x(lenx), y(leny)
CALL SSYR2(uplo, n, alpha, x, incx, y, incy, a, lda)

CHARACTER*1      uplo
INTEGER*8        n, lda, incx, incy
COMPLEX*16       alpha, a(lda, n), x(lenx), y(leny)
CALL CHER2(uplo, n, alpha, x, incx, y, incy, a, lda)
```

**Input**

uplo            Upper/lower triangular option for  $A$ :  
                 'L' or 'l'      Reference and update only the lower triangle  
                                    of  $A$ .

	<b>‘U’ or ‘u’</b>	Reference and update only the upper triangle of $A$ .
<b>n</b>		Number of rows and columns in matrix $A$ and elements of vectors $x$ and $y$ , $n \geq 0$ . If $n = 0$ , the subprograms do not reference $a$ , $x$ , or $y$ .
<b>alpha</b>		The scalar $\alpha$ . If <b>alpha</b> = 0, the subprograms do not reference $a$ , $x$ , or $y$ .
<b>x</b>		Array of length <b>lenx</b> = $(n-1) \times  \mathbf{incx}  + 1$ containing the $n$ -vector $x$ .
<b>incx</b>		Increment for the array $x$ , <b>incx</b> $\neq 0$ : <b>incx</b> > 0 $x$ is stored forward in array $x$ , i.e., $x_i$ is stored in $x((i-1) \times \mathbf{incx} + 1)$ . <b>incx</b> < 0 $x$ is stored backward in array $x$ , i.e., $x_i$ is stored in $x((i-n) \times \mathbf{incx} + 1)$ .  Use <b>incx</b> = 1 if the vector $x$ is stored contiguously in array $x$ , i.e., if $x_i$ is stored in $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.
<b>y</b>		Array of length <b>leny</b> = $(n-1) \times  \mathbf{incy}  + 1$ containing the $n$ -vector $y$ .
<b>incy</b>		Increment for the array $y$ , <b>incy</b> $\neq 0$ : <b>incy</b> > 0 $y$ is stored forward in array $y$ , i.e., $y_i$ is stored in $y((i-1) \times \mathbf{incy} + 1)$ . <b>incy</b> < 0 $y$ is stored backward in array $y$ , i.e., $y_i$ is stored in $y((i-n) \times \mathbf{incy} + 1)$ .  Use <b>incy</b> = 1 if the vector $y$ is stored contiguously in array $y$ , i.e., if $y_i$ is stored in $y(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.
<b>a</b>		Array whose upper or lower triangle, as specified by <b>uplo</b> , contains the upper or lower triangle of an $n$ -by- $n$ real symmetric or complex Hermitian matrix $A$ . The other triangle of $a$ is not referenced.
<b>lda</b>		The leading dimension of array $a$ as declared in the calling program unit, with <b>lda</b> $\geq \max(n, 1)$ .

Basic Matrix Operations  
SSYR2/CHER2 – Rank-2 Update

## Output

- a** The upper or lower triangle of the updated  $A$  matrix, as specified by **uplo**, replaces the upper or lower triangle of the input, respectively. The other triangle of **a** is unchanged.

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo**  $\neq$  'L' or 'l' or 'U' or 'u',  
**n** < 0,  
**lda** < max(**n**,1),  
**incx** = 0, and  
**incy** = 0.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **uplo** argument as 'LOWER' for 'L' or 'UPPER' for 'U'.

## Example 1

Apply a REAL\*8 symmetric rank-2 update  $xy^T + x^T y$  to  $A$ , where  $A$  is a 9-by-9 real symmetric matrix whose upper triangle is stored in the upper triangle of an array  $A$  whose dimensions are 10 by 10,  $x$  is a real vector 9 elements long stored in an array  $X$  of dimension 10, and  $y$  is a real vector 9 elements long stored in an array  $Y$  also of dimension 10.

```
CHARACTER*1 UPLO
INTEGER*8   N, LDA, INCX, INCY
REAL*8     ALPHA, A(10,10), X(10), Y(10)
UPLO = 'U'
N = 9
ALPHA = 1.0
LDA = 10
INCX = 1
INCY = 1
CALL SSYR2 (UPLO, N, ALPHA, X, INCX, Y, INCY, A, LDA)
```

## Example 2

Apply a COMPLEX\*16 Hermitian rank-2 update  $\alpha xy^* + \bar{\alpha} yx^*$  to  $A$ , where  $A$  is a 9-by-9 complex Hermitian matrix whose lower triangle is stored in the lower

triangle of an array A whose dimensions are 10 by 10,  $\alpha$  is a complex scalar,  $x$  is a complex vector 9 elements long stored in an array X of dimension 10, and  $y$  is a complex vector 9 elements long stored in an array Y of dimension 10.

```
INTEGER*8  N,LDA
COMPLEX*16 ALPHA,A(10,10),X(10),Y(10)
N = 9
LDA = 10
CALL CHER2 ('LOWER',N,ALPHA,X,1,Y,1,A,LDA)
```

Basic Matrix Operations  
**SSYR2K/CHER2K/CSYR2K – Rank-2k Update**

**NAME** SSYR2K/CHER2K/CSYR2K – Rank-2k Update

**Purpose**

These subprograms apply a symmetric or Hermitian rank-2k update to a real symmetric, complex symmetric, or complex Hermitian matrix; specifically they compute the following operations:

for symmetric  $C$ :  $C \leftarrow \alpha AB^T + \bar{\alpha} BA^T + \beta C$  and  $C \leftarrow \alpha A^T B + \bar{\alpha} B^T A + \beta C$   
 for Hermitian  $C$ :  $C \leftarrow \alpha AB^* + \bar{\alpha} BA^* + \beta C$  and  $C \leftarrow \alpha A^* B + \bar{\alpha} B^* A + \beta C$

where  $\alpha$  and  $\beta$  are scalars,  $\bar{\alpha}$  is the complex conjugate of  $\alpha$ ,  $C$  is an  $n$ -by- $n$  real symmetric, complex symmetric, or complex Hermitian matrix, and  $A$  and  $B$  are matrices whose size, either  $n$  by  $k$  or  $k$  by  $n$ , depends on which form of the update is requested. Here,  $A^T$  and  $B^T$  are the transposes and  $A^*$  and  $B^*$  are the conjugate transposes of  $A$  and  $B$ , respectively. (The conjugate of a real scalar is just the scalar, and the conjugate transpose of a real matrix is simply the transpose.)

The structure of  $C$  is indicated by the name of the subprogram used:

SSYR2K	$C$ is a real symmetric matrix
CHER2K	$C$ is a complex Hermitian matrix
CSYR2K	$C$ is a complex symmetric matrix

**Matrix Storage**

Because either triangle of  $C$  may be obtained from the other, these subprograms reference and apply the update to only one triangle of  $C$ . You may supply either the upper or the lower triangle of  $C$ , in a two-dimensional array large enough to hold the entire matrix, and the same triangle of the updated matrix is returned in the array. The other triangle of the array is not referenced.

**Usage**

SCILIB:

CHARACTER*1	uplo, trans
INTEGER*8	n, k, lda, ldb, ldc
REAL*8	alpha, beta, a(lda, *), b(ldb, *), c(ldc, n)
CALL SSYR2K(uplo, trans, n, k, alpha, a, lda, b, ldb, beta, c, ldc)	
CHARACTER*1	uplo, trans
INTEGER*8	n, k, lda, ldb, ldc
REAL*8	alpha, beta
COMPLEX*16	a(lda, *), b(ldb, *), c(ldc, n)

```
CALL CHER2K(uplo, trans, n, k, alpha, a, lda, b, ldb, beta, c, ldc)
CHARACTER*1      uplo, trans
INTEGER*8        n, k, lda, ldb, ldc
COMPLEX*16       alpha, beta, a(lda, *), b(ldb, *), c(ldc, n)
CALL CSYR2K(uplo, trans, n, k, alpha, a, lda, b, ldb, beta, c, ldc)
```

## Input

<b>uplo</b>	Upper/lower triangular storage option for $C$ : 'L' or 'l'     Reference and update only the lower triangle of $C$ 'U' or 'u'     Reference and update only the upper triangle of $C$
<b>trans</b>	Specifies the operation to be performed: 'N' or 'n'     Compute $C \leftarrow \alpha AB^T + \alpha BA^T + \beta C$ 'T' or 't'     Compute $C \leftarrow \alpha A^T B + \alpha B^T A + \beta C$ 'C' or 'c'     Compute $C \leftarrow \alpha A * B + \alpha B * A + \beta C$ 'T' and 't' are invalid in subprogram CHER2K, and 'C' and 'c' are invalid in subprogram CSYR2K. In subprogram SSYR2K, 'C' and 'c' have the same meaning as 'T' and 't'.
<b>n</b>	Number of rows and columns in matrix $C$ , $n \geq 0$ . If $n = 0$ , the subprograms do not reference $a$ , $b$ , or $c$ .
<b>k</b>	Number of rows or columns in matrices $A$ and $B$ , depending on <b>trans</b> ; refer to the description of $a$ for details. $k \geq 0$ ; if $k = 0$ , the subprograms do not reference $a$ or $b$ .
<b>alpha</b>	The scalar $\alpha$ . If <b>alpha</b> = 0, the subprograms compute $C \leftarrow \beta C$ without referencing $a$ or $b$ .
<b>a</b>	Array containing the matrix $A$ , whose size is indicated by <b>trans</b> : 'N' or 'n' $A$ is $n$ by $k$ otherwise $A$ is $k$ by $n$
<b>lda</b>	The leading dimension of array $a$ as declared in the calling program unit, with $lda \geq \max$ (the number of rows of $A, 1$ ).
<b>b</b>	Array containing matrix $B$ , which is the same size as matrix $A$ . Refer to the description of $a$ above for details.
<b>ldb</b>	The leading dimension of array $b$ as declared in the calling program unit, with $ldb \geq \max$ (the number of rows of $B, 1$ ).
<b>beta</b>	The scalar $\beta$ .

Basic Matrix Operations  
SSYR2K/CHER2K/CSYR2K – Rank-2k Update

- c** Array whose upper or lower triangle, as specified by **uplo**, contains the upper or lower triangle of the  $n$ -by- $n$  symmetric or Hermitian matrix  $C$ . Not used as input if **beta** = 0.
- ldc** The leading dimension of array **c** as declared in the calling program unit, with **ldc**  $\geq$  **max(n,1)**.

## Output

- c** The upper or lower triangle of the updated matrix  $C$ , as specified by **uplo**, replaces the upper or lower triangle of the input, respectively. The other triangle of **c** is unchanged.

## Notes

These subprograms conform to specifications of the Level 3 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo**  $\neq$  'L' or 'l' or 'U' or 'u',  
**trans**  $\neq$  'N' or 'n' or 'T' or 't' or 'C' or 'c',  
**n** < 0,  
**k** < 0,  
**lda** too small,  
**ldb** too small, and  
**ldc** < **max(m,1)**.

Also, note that some of the values of **trans** listed above are invalid in subprograms CHER2K and CSYR2K. Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the CALL statement may be improved, for example, by coding the **uplo** argument as 'LOWER' for 'L' or 'UPPER' for 'U'. Refer to "Example 2."

## Example 1

Apply a REAL\*8 rank-6 update  $AB^T + BA^T$  to an 8-by-8 real symmetric matrix  $C$  whose upper triangle is stored in the upper triangle of an array **C** of dimension 10 by 10, where  $A$  is an 8-by-3 real matrix stored in an array **A**, also of dimension 10 by 10.

```
CHARACTER*1 UPLO, TRANS
INTEGER*8   N, K, LDA, LDB, LDC
REAL*8     ALPHA, BETA, A(10,10), B(10,10), C(10,10)
UPLO = 'U'
TRANS = 'N'
N = 8
K = 3
ALPHA = 1.0
BETA = 1.0
LDA = 10
LDB = 10
LDC = 10
CALL SSYR2K (UPLO, TRANS, N, K, ALPHA, A, LDA, B, LDB, BETA, C, LDC)
```

## Example 2

Apply a COMPLEX\*16 Hermitian rank-4 update  $-2AB^*-2BA^*$  to a 9-by-9 complex Hermitian matrix  $C$  whose lower triangle is stored in the lower triangle of an array  $C$  of dimension 10 by 10, where  $A$  is a 9-by-2 complex matrix stored in an array  $A$  of dimension 10 by 10.

```
INTEGER*8   N, K, LDA, LDB, LDC
COMPLEX*16 A(10,10), B(10,10), C(10,10)
N = 9
K = 2
LDA = 10
LDB = 10
LDC = 10
CALL CHER2K ('LOWER', 'NONTRANS', N, K, -2.0, A, LDA, B, LDB,
&          1.0, C, LDC)
```

Basic Matrix Operations  
SSYRK/CHERK – Rank-k Update

**NAME** SSYRK/CHERK – Rank-k Update

**Purpose**

These subprograms apply a rank- $k$  update to a real symmetric, complex symmetric, or complex Hermitian matrix; specifically they compute:

$$\begin{aligned} C &\leftarrow \alpha AA^T + \beta C, & C &\leftarrow \alpha A^T A + \beta C, \\ C &\leftarrow \alpha AA^* + \beta C, & C &\leftarrow \alpha A^* A + \beta C, \end{aligned}$$

where  $\alpha$  and  $\beta$  are scalars,  $C$  is an  $n$ -by- $n$  real symmetric, complex symmetric, or complex Hermitian matrix, and  $A$  is a matrix whose size, either  $n$  by  $k$  or  $k$  by  $n$ , depends on which form of the update is requested. Here,  $A^T$  and  $A^*$  are the transpose and conjugate transpose of  $A$ , respectively.

The structure of  $C$  is indicated by the name of the subprogram used:

SSYRK	$C$ is a real symmetric matrix
CHERK	$C$ is a complex Hermitian matrix
CSYRK	$C$ is a complex symmetric matrix

**Matrix Storage**

Because either triangle of  $C$  may be obtained from the other, these subprograms reference and apply the update to only one triangle of  $C$ . You may supply either the upper or the lower triangle of  $C$ , in a two-dimensional array large enough to hold the entire matrix, and the same triangle of the updated matrix is returned in the array. The other triangle of the array is not referenced.

**Usage**

SCILIB:

CHARACTER*1	uplo, trans
INTEGER*8	n, k, lda, ldc
REAL*8	alpha, beta, a(lda, *), c(ldc, n)
CALL SSYRK(uplo, trans, n, k, alpha, a, lda, beta, c, ldc)	
CHARACTER*1	uplo, trans
INTEGER*8	n, k, lda, ldc
REAL*8	alpha, beta
COMPLEX*16	a(lda, *), c(ldc, n)
CALL CHERK(uplo, trans, n, k, alpha, a, lda, beta, c, ldc)	
CHARACTER*1	uplo, trans
INTEGER*8	n, k, lda, ldc
COMPLEX*16	alpha, beta, a(lda, *), c(ldc, n)

CALL CSYRK(**uplo**, **trans**, **n**, **k**, **alpha**, **a**, **lda**, **beta**, **c**, **ldc**)

## Input

<b>uplo</b>	Upper/lower triangular storage option for <i>C</i> : 'L' or 'l'     Reference and update only the lower triangle of <i>C</i> 'U' or 'u'     Reference and update only the upper triangle of <i>C</i>
<b>trans</b>	Specifies the operation to be performed: 'N' or 'n'     Compute $C \leftarrow \alpha AA^T + \beta C$ 'T' or 't'     Compute $C \leftarrow \alpha A^T A + \beta C$ 'C' or 'c'     Compute $C \leftarrow \alpha A^* A + \beta C$ 'T' and 't' are invalid in subprogram CHER2K, and 'C' and 'c' are invalid in subprogram CSYR2K. In subprogram SSYRK, 'C' and 'c' have the same meaning as 'T' and 't'.
<b>n</b>	Number of rows and columns in matrix <i>C</i> , $n \geq 0$ . If $n = 0$ , the subprograms do not reference <b>a</b> or <b>c</b> .
<b>k</b>	Number of rows or columns in matrix <i>A</i> , $k \geq 0$ , depending on <b>trans</b> ; refer to description of <b>A</b> for details. If $k = 0$ , the subprograms do not reference <b>a</b> .
<b>alpha</b>	The scalar $\alpha$ . If <b>alpha</b> = 0, the subprograms compute $C \leftarrow \beta C$ without referencing <b>a</b> .
<b>a</b>	Array containing the matrix <i>A</i> , whose size is indicated by <b>trans</b> : 'N' or 'n' <i>A</i> is $n$ by $k$ otherwise <i>A</i> is $k$ by $n$
<b>lda</b>	The leading dimension of array <b>a</b> as declared in the calling program unit, with $lda \geq \max(n, k)$ (the number of rows of <i>A</i> , 1).
<b>beta</b>	The scalar $\beta$ .
<b>c</b>	Array whose upper or lower triangle, as specified by <b>uplo</b> , contains the upper or lower triangle of the $n$ -by- $n$ symmetric or Hermitian matrix <i>C</i> . Not used as input if <b>beta</b> = 0.
<b>ldc</b>	The leading dimension of array <b>c</b> as declared in the calling program unit, with $ldc \geq \max(n, k)$ .

Basic Matrix Operations  
SSYRK/CHERK – Rank-k Update

## Output

- c** The upper or lower triangle of the updated  $C$  matrix, as specified by **uplo**, replaces the upper or lower triangle of the input, respectively. The other triangle of  $c$  is unchanged.

## Notes

These subprograms conform to specifications of the Level 3 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo**  $\neq$  'L' or 'l' or 'U' or 'u',  
**trans**  $\neq$  'N' or 'n' or 'T' or 't' or 'C' or 'c',  
**n** < 0,  
**k** < 0,  
**lda** too small, and  
**ldc** < max(**m**,1).

Also, some values of **trans** listed above are invalid in subprograms CHERK and CSYRK.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the CALL statement may be improved, for example, by coding the **uplo** argument as 'LOWER' for 'L' or 'UPPER' for 'U'. Refer to "Example 2."

## Example 1

Apply a REAL\*8 rank-6 update  $AA^T$  to an 8-by-8 real symmetric matrix  $C$  whose upper triangle is stored in the upper triangle of an array  $C$  of dimension 10 by 10, where  $A$  is an 8-by-6 real matrix stored in an array  $A$ , also of dimension 10 by 10.

```
CHARACTER*1 UPLO, TRANS
INTEGER*8   N, K, LDA, LDC
REAL*8     ALPHA, BETA, A(10,10), C(10,10)
UPLO = 'U'
TRANS = 'N'
N = 8
K = 6
ALPHA = 1.0
BETA = 1.0
LDA = 10
LDC = 10
CALL SSYRK (UPLO, TRANS, N, K, ALPHA, A, LDA, BETA, C, LDC)
```

## Example 2

Apply a COMPLEX\*16 Hermitian rank-2 update  $-2AA^*$  to a 9-by-9 complex Hermitian matrix  $C$  whose lower triangle is stored in the lower triangle of an array  $C$  of dimension 10 by 10, where  $A$  is a 9-by-2 complex matrix stored in an array  $A$  of dimension 10 by 10.

```
INTEGER*8  N, K, LDA, LDC
COMPLEX*16 A(10,10), C(10,10)
N = 9
K = 2
LDA = 10
LDC = 10
CALL CHERK ( 'LOWER', 'NONTRANS', N, K, -2.0, A, LDA, 1.0, C, LDC)
```

Basic Matrix Operations  
STBMV/CTBMV – Matrix-Vector Multiply

**NAME** STBMV/CTBMV – Matrix-Vector Multiply

**Purpose**

Given an  $n$ -by- $n$  upper- or lower-triangular band matrix  $A$ , and an  $n$ -vector  $x$ , these subprograms compute the matrix-vector products  $Ax$ ,  $A^T x$ , and  $A^* x$ , where  $A^T$  is the transpose of  $A$ , and  $A^*$  is the conjugate transpose. Specifically, these subprograms compute matrix-vector products of the forms

$$x \leftarrow Ax, \quad x \leftarrow A^T x, \quad \text{and} \quad x \leftarrow A^* x.$$

A lower-triangular band matrix is a matrix whose strict upper triangle is zero, and whose nonzero lower-triangular elements all are on or fairly near the principal diagonal. Specifically,  $a_{ij} \neq 0$  only if  $0 \leq i-j \leq kd$  for some integer  $kd$ . In contrast, an upper-triangular band matrix is a matrix whose strict lower triangle is zero, and whose nonzero upper-triangular elements all are on or fairly near the principal diagonal, i.e., with  $a_{ij} \neq 0$  only if  $0 \leq j-i \leq kd$ .

**Matrix Storage**

Triangular band matrices are stored in a compressed form that takes advantage of knowing the positions of the only elements that may be nonzero. The following examples illustrate the storage of triangular band matrices.

**Lower triangular storage**

If  $A$  is a 9-by-9 lower-triangular band matrix with bandwidth  $kd = 3$ , for example,

11	0	0	0	0	0	0	0	0
21	22	0	0	0	0	0	0	0
31	32	33	0	0	0	0	0	0
41	42	43	44	0	0	0	0	0
0	52	53	54	55	0	0	0	0
0	0	63	64	65	66	0	0	0
0	0	0	74	75	76	77	0	0
0	0	0	0	85	86	87	88	0
0	0	0	0	0	96	97	98	99

the lower triangular band part of  $A$  is stored in an array **ab** with at least  $kd+1 = 4$  rows and 9 columns:

```

11  22  33  44  55  66  77  88  99
21  32  43  54  65  76  87  98  *
31  42  53  64  75  86  97  *  *
41  52  63  74  85  96  *  *  *

```

where asterisks represent elements in the  $kd$ -by- $kd$  triangle at the lower-right corner of  $\mathbf{ab}$  that are not referenced. Thus, if  $a_{ij}$  is an element within the band of  $A$ , it is stored in  $\mathbf{ab}(1+i-j,j)$ . Therefore, the columns of  $A$  are stored in the columns of  $\mathbf{ab}$ , and the diagonals of  $A$  are stored in the rows of  $\mathbf{ab}$ , with the principal diagonal in the first row, the first subdiagonal in the second row, and so on.

### Upper triangular storage

If  $A$  is a 9-by-9 upper-triangular band matrix with bandwidth  $kd = 3$ , for example,

```

11  12  13  14  0  0  0  0  0
0  22  23  24  25  0  0  0  0
0  0  33  34  35  36  0  0  0
0  0  0  44  45  46  47  0  0
0  0  0  0  55  56  57  58  0
0  0  0  0  0  66  67  68  69
0  0  0  0  0  0  77  78  79
0  0  0  0  0  0  0  88  89
0  0  0  0  0  0  0  0  99

```

the upper triangular band part of  $A$  is stored in an array  $\mathbf{ab}$  with at least  $kd+1 = 4$  rows and 9 columns:

```

*  *  *  14  25  36  47  58  69
*  *  13  24  35  46  57  68  79
*  12  23  34  45  56  67  78  89
11  22  33  44  55  66  77  88  99

```

where asterisks represent elements in the  $kd$ -by- $kd$  triangle at the upper-left corner of  $\mathbf{ab}$  that are not referenced. Thus, if  $a_{ij}$  is an element within the band of  $A$ , it is stored in  $\mathbf{ab}(kd+1+i-j,j)$ . Therefore, the columns of  $A$  are stored in the columns of  $\mathbf{ab}$ , and the diagonals of  $A$  are stored in the rows of  $\mathbf{ab}$ , with the principal diagonal in row  $kd+1$ , the first superdiagonal starting in the second position in row  $kd$ , and so on.

Basic Matrix Operations  
**STBMV/CTBMV – Matrix-Vector Multiply**

**Usage**

SCILIB:

```

CHARACTER*1      uplo, trans, diag
INTEGER*8       n, kd, ldab, incx
REAL*8          ab(ldab, n), x(lenx)
CALL STBMV(uplo, trans, diag, n, kd, ab, ldab, x, incx)

CHARACTER*1      uplo, trans, diag
INTEGER*8       n, kd, ldab, incx
COMPLEX*16     ab(ldab, n), x(lenx)
CALL CTBMV(uplo, trans, diag, n, kd, ab, ldab, x, incx)

```

**Input**

**uplo**           Upper/lower triangular option for A:  
           'L' or 'l'    A is lower triangular  
           'U' or 'u'    A is upper triangular

**trans**          Transposition option for A:  
           'N' or 'n'    Compute  $x \leftarrow Ax$   
           'T' or 't'    Compute  $x \leftarrow A^T x$   
           'C' or 'c'    Compute  $x \leftarrow A^* x$

where  $A^T$  is the transpose of A, and  $A^*$  is the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.

**diag**           Specifies whether the matrix is unit triangular, i.e.,  $a_{ii} = 1$ , or not:  
           'N' or 'n'    The diagonal of A is stored in the array  
           'U' or 'u'    The diagonal of A consists of unstored ones

When **diag** is supplied as 'U' or 'u', diagonal elements of A are not referenced, but space must be reserved for them.

**n**               Number of rows and columns in matrix A,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference **ab** or **x**.

**kd**             The number of nonzero diagonals above or below the principal diagonal. If **uplo** is supplied as 'U' or 'u', **kd** specifies the number of nonzero diagonals above the principal diagonal. If **uplo** is supplied as 'L' or 'l', **kd** specifies the number of nonzero diagonals below the principal diagonal.

<b>ab</b>	Array containing the $n$ -by- $n$ triangular band matrix $A$ in the compressed form described above. The columns of the band of $A$ are stored in the columns of <b>ab</b> , and the diagonals of the band of $A$ are stored in the rows of <b>ab</b> .
<b>ldab</b>	The leading dimension of array <b>ab</b> as declared in the calling program unit, with $\text{ldab} \geq \text{kd}+1$ .
<b>x</b>	Array of length $\text{lenx} = (n-1) \times  \text{incx}  + 1$ containing the input vector $x$ .
<b>incx</b>	Increment for the array <b>x</b> , $\text{incx} \neq 0$ : $\text{incx} > 0$ $x$ is stored forward in array <b>x</b> , i.e., $x_i$ is stored in $\mathbf{x}((i-1) \times \text{incx} + 1)$ . $\text{incx} < 0$ $x$ is stored backward in array <b>x</b> , i.e., $x_i$ is stored in $\mathbf{x}((i-n) \times \text{incx} + 1)$ .  Use $\text{incx} = 1$ if the vector $x$ is stored contiguously in array <b>x</b> , i.e., if $x_i$ is stored in $\mathbf{x}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

## Output

**x**     The updated  $x$  vector replaces the input.

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo**  $\neq$  'L' or 'l' or 'U' or 'u',  
**trans**  $\neq$  'N' or 'n' or 'T' or 't' or 'C' or 'c',  
**diag**  $\neq$  'N' or 'n' or 'U' or 'u',  
**n**  $< 0$ ,  
**kd**  $< 0$ ,  
**ldab**  $< \text{kd}+1$ , and  
**incx**  $= 0$ .

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **trans** argument as 'NORMAL' or

## Basic Matrix Operations

### STBMV/CTBMV – Matrix-Vector Multiply

'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

#### Example 1

Form the REAL\*8 matrix-vector product  $Ax$ , where  $A$  is a 75-by-75 unit-diagonal, lower-triangular real band matrix with bandwidth 15 that is stored in an array AB whose dimensions are 25 by 100, and  $x$  is a real vector 75 elements long stored in an array X of dimension 100.

```
CHARACTER*1  UPLO, TRANS, DIAG
INTEGER*8    N, KD, LDAB, INCX
REAL*8       AB(25,100), X(100)
UPLO = 'L'
TRANS = 'N'
DIAG = 'U'
N = 75
KD = 15
LDAB = 25
INCX = 1
CALL STBMV (UPLO, TRANS, DIAG, N, KD, AB, LDAB, X, INCX)
```

#### Example 2

Form the REAL\*8 matrix-vector product  $A^T x$ , where  $A$  is a 75-by-75 nonunit-diagonal, upper-triangular real band matrix with bandwidth 15 that is stored in an array AB whose dimensions are 25 by 100, and  $x$  is a real vector 75 elements long stored in an array X of dimension 100.

```
INTEGER*8    N, KD, LDAB
REAL*8       AB(25,100), X(100)
N = 75
KD = 15
LDAB = 25
CALL STBMV ('UPPER', 'TRANSPOSE', 'NONUNIT', N, KD, AB, LDAB,
           X, 1)
```

**NAME** STBSV/CTBSV – Solve Triangular Band System

**Purpose**

Given an  $n$ -by- $n$  upper- or lower-triangular band matrix  $A$ , and an  $n$ -vector  $x$ , these subprograms overwrite  $x$  with the solution  $y$  to the system of linear equations  $Ay = x$ . This is the forward elimination or back substitution step of Gaussian elimination for band matrices. Optionally,  $A$  may be replaced by  $A^T$ , the transpose of  $A$ , or by  $A^*$ , the conjugate transpose of  $A$ .

A lower-triangular band matrix is a matrix whose strict upper triangle is zero, and whose nonzero lower-triangular elements all are on or fairly near the principal diagonal. Specifically,  $a_{ij} \neq 0$  only if  $0 \leq i-j \leq kd$  for some integer  $kd$ . In contrast, an upper-triangular band matrix is a matrix whose strict lower triangle is zero, and whose nonzero upper-triangular elements all are on or fairly near the principal diagonal, but with  $a_{ij} \neq 0$  only if  $0 \leq j-i \leq kd$ .

Specifically, these subprograms compute

$$x \leftarrow A^{-1}x, \quad x \leftarrow A^{-T}x, \quad \text{and} \quad x \leftarrow A^{-*}x$$

where  $A^{-T}$  is the inverse of the transpose of  $A$ , and  $A^{-*}$  is the inverse of the conjugate transpose of  $A$ .

These subprograms are more primitive than the LINPACK band equation solvers. As such, they are intended to supplement but not replace the equation solvers, serving instead as building blocks in constructing optimized linear algebra software. In fact, many of the LINPACK subprograms have been recoded to call these routines.

**Matrix Storage**

Triangular band matrices are stored in a compressed form that takes advantage of knowing the positions of the only elements that may be nonzero. The following examples illustrate the storage of triangular band matrices.

**Lower triangular storage**

If  $A$  is a 9-by-9 lower-triangular band matrix with bandwidth  $kd = 3$ , for example,

Basic Matrix Operations

**STBSV/CTBSV – Solve Triangular Band System**

11	0	0	0	0	0	0	0	0
21	22	0	0	0	0	0	0	0
31	32	33	0	0	0	0	0	0
41	42	43	44	0	0	0	0	0
0	52	53	54	55	0	0	0	0
0	0	63	64	65	66	0	0	0
0	0	0	74	75	76	77	0	0
0	0	0	0	85	86	87	88	0
0	0	0	0	0	96	97	98	99

the lower triangular band part of  $A$  is stored in an array  $\mathbf{ab}$  with at least  $kd+1 = 4$  rows and 9 columns:

11	22	33	44	55	66	77	88	99
21	32	43	54	65	76	87	98	*
31	42	53	64	75	86	97	*	*
41	52	63	74	85	96	*	*	*

where asterisks represent elements in the  $kd$ -by- $kd$  triangle at the lower-right corner of  $\mathbf{ab}$  that are not referenced. Thus, if  $a_{ij}$  is an element within the band of  $A$ , it is stored in  $\mathbf{ab}(1+i-j,j)$ . Therefore, the columns of  $A$  are stored in the columns of  $\mathbf{ab}$ , and the diagonals of  $A$  are stored in the rows of  $\mathbf{ab}$ , with the principal diagonal in the first row, the first subdiagonal in the second row, and so on.

**Upper triangular storage**

If  $A$  is a 9-by-9 upper-triangular band matrix with bandwidth  $kd = 3$ , for example,

11	12	13	14	0	0	0	0	0
0	22	23	24	25	0	0	0	0
0	0	33	34	35	36	0	0	0
0	0	0	44	45	46	47	0	0
0	0	0	0	55	56	57	58	0
0	0	0	0	0	66	67	68	69
0	0	0	0	0	0	77	78	79
0	0	0	0	0	0	0	88	89
0	0	0	0	0	0	0	0	99

the upper triangular band part of  $A$  is stored in an array  $\mathbf{ab}$  with at least  $kd+1 = 4$  rows and 9 columns:

## STBSV/CTBSV – Solve Triangular Band System

```

*      *      *      14      25      36      47      58      69
*      *      13      24      35      46      57      68      79
*      12      23      34      45      56      67      78      89
11     22      33      44      55      66      77      88      99

```

where asterisks represent elements in the  $kd$ -by- $kd$  triangle at the upper-left corner of  $\mathbf{ab}$  that are not referenced. Thus, if  $a_{ij}$  is an element within the band of  $A$ , it is stored in  $\mathbf{ab}(kd+1+i-j,j)$ . Therefore, the columns of  $A$  are stored in the columns of  $\mathbf{ab}$ , and the diagonals of  $A$  are stored in the rows of  $\mathbf{ab}$ , with the principal diagonal in row  $kd+1$ , the first superdiagonal starting in the second position in row  $kd$ , and so on.

## Usage

SCILIB:

```

CHARACTER*1      uplo, trans, diag
INTEGER*8        n, kd, ldab, incx
REAL*8           ab(ldab, n), x(lenx)
CALL STBSV(uplo, trans, diag, n, kd, ab, ldab, x, incx)

CHARACTER*1      uplo, trans, diag
INTEGER*8        n, kd, ldab, incx
COMPLEX*16       ab(ldab, n), x(lenx)
CALL CTBSV(uplo, trans, diag, n, kd, ab, ldab, x, incx)

```

## Input

<b>uplo</b>	Upper/lower triangular option for A: 'L' or 'l'      Solve lower-triangular band system (forward elimination) 'U' or 'u'      Solve upper-triangular band system (back substitution)
<b>trans</b>	Transposition option for A: 'N' or 'n'      Compute $x \leftarrow A^{-1}x$ 'T' or 't'      Compute $x \leftarrow A^{-T}x$ 'C' or 'c'      Compute $x \leftarrow A^{-*}x$

where  $A^{-T}$  is the inverse of the transpose of  $A$ , and  $A^{-*}$  is the inverse of the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.

**diag**      Specifies whether the matrix is unit triangular, i.e.,  $a_{ii} = 1$ , or not:

## Basic Matrix Operations

### STBSV/CTBSV – Solve Triangular Band System

**'N' or 'n'** The diagonal of  $A$  is stored in the array

**'U' or 'u'** The diagonal of  $A$  consists of unstored ones

When **diag** is supplied as **'U'** or **'u'**, diagonal elements of  $A$  are not referenced, but space must be reserved for them.

**n** Number of rows and columns in matrix  $A$ ,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference **ab** or **x**.

**kd** The number of nonzero diagonals above or below the principal diagonal. If **uplo** is supplied as **'U'** or **'u'**, **kd** specifies the number of nonzero diagonals above the principal diagonal. If **uplo** is supplied as **'L'** or **'l'**, **kd** specifies the number of nonzero diagonals below the principal diagonal.

**ab** Array containing the  $n$ -by- $n$  triangular band matrix  $A$  in the compressed form described above. The columns of the band of  $A$  are stored in the columns of **ab**, and the diagonals of the band of  $A$  are stored in the rows of **ab**.

**ldab** The leading dimension of array **ab** as declared in the calling program unit, with **ldab**  $\geq$  **kd**+1.

**x** Array of length **lenx** =  $(n-1) \times |\text{incx}| + 1$  containing the right-hand side  $n$ -vector  $x$ .

**incx** Increment for the array **x**, **incx**  $\neq$  0:

**incx** > 0  $x$  is stored forward in array **x**, i.e.,  $x_i$  is stored in  $\mathbf{x}((i-1) \times \text{incx} + 1)$ .

**incx** < 0  $x$  is stored backward in array **x**, i.e.,  $x_i$  is stored in  $\mathbf{x}((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array **x**, i.e., if  $x_i$  is stored in  $\mathbf{x}(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.

## Output

**x** The solution vector of the triangular band system replaces the input.

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

The subprograms do not check for singularity of matrix  $A$ .  $A$  is singular if **diag** = **'N'** or **'n'** and some  $a_{ii} = 0$ . This condition causes a division by zero to occur.

Therefore, the program must detect singularity and take appropriate action to avoid a problem before calling any of these subprograms.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

```

uplo ≠ 'L' or 'l' or 'U' or 'u',
trans ≠ 'N' or 'n' or 'T' or 't' or 'C' or 'c',
diag ≠ 'N' or 'n' or 'U' or 'u',
n < 0,
kd < 0,
ldab < kd+1, and
incx = 0.

```

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the CALL statement may be improved by coding the **trans** argument as 'NORMAL' or 'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

### Example 1

Perform REAL\*8 forward elimination using the 75-by-75 unit-diagonal lower-triangular real band matrix with bandwidth 15 that is stored in an array AB whose dimensions are 25 by 100, and *x* is a real vector 75 elements long stored in an array X of dimension 100.

```

CHARACTER*1  UPLO, TRANS, DIAG
INTEGER*8    N, KD, LDAB, INCX
REAL*8       AB(25,100), X(100)
UPLO = 'L'
TRANS = 'N'
DIAG = 'U'
N = 75
KD = 15
LDAB = 25
INCX = 1
CALL STBSV (UPLO, TRANS, DIAG, N, KD, AB, LDAB, X, INCX)

```

### Example 2

Perform REAL\*8 back substitution using the 75-by-75 nonunit-diagonal, upper-triangular real band matrix with bandwidth 15 that is stored in an array AB whose dimensions are 25 by 100, and *x* is a real vector 75 elements long stored in an array X of dimension 100.

## Basic Matrix Operations

### STBSV/CTBSV – Solve Triangular Band System

```
INTEGER*8 N, KD, LDAB
REAL*8    AB(25,100), X(100)
N = 75
KD = 15
LDAB = 25
CALL STBSV ('UPPER', 'NONTRANS', 'NONUNIT', N, KD, AB, LDAB, X, 1)
```

**NAME** STPMV/CTPMV – Matrix-Vector Multiply

**Purpose**

Given an  $n$ -by- $n$  upper- or lower-triangular matrix  $A$  stored in packed form as described in “Matrix Storage,” and an  $n$ -vector  $x$ , these subprograms compute the matrix-vector products  $Ax$ ,  $A^T x$ , and  $A^* x$ , where  $A^T$  is the transpose of  $A$ , and  $A^*$  is the conjugate transpose of  $A$ . Specifically, these subprograms compute matrix-vector products of the forms

$$x \leftarrow Ax, \quad x \leftarrow A^T x, \quad \text{and} \quad x \leftarrow A^* x.$$

**Matrix Storage**

You supply the upper or lower triangle of  $A$ , stored column-by-column in packed form in a 1-dimensional array. This saves memory compared to storing the entire matrix.

The following examples illustrate the packed storage of a triangular matrix.

**Upper triangular matrix**

If  $A$  is

$$\begin{matrix} 11 & 12 & 13 & 14 \\ 0 & 22 & 23 & 24 \\ 0 & 0 & 33 & 34 \\ 0 & 0 & 0 & 44 \end{matrix}$$

then  $A$  is packed column by column into an array  $ap$  as follows:

$k$	1	2	3	4	5	6	7	8	9	10
$ap(k)$	11	12	22	13	23	33	14	24	34	44

Upper-triangular matrix element  $a_{ij}$  is stored in array element  $ap(i+j \times (j-1)/2)$ .

**Lower triangular matrix**

If  $A$  is

$$\begin{matrix} 11 & 0 & 0 & 0 \\ 21 & 22 & 0 & 0 \\ 31 & 32 & 33 & 0 \\ 41 & 42 & 43 & 44 \end{matrix}$$

then  $A$  is packed column by column into an array  $ap$  as follows:

Basic Matrix Operations  
**STPMV/CTPMV – Matrix-Vector Multiply**

$k$	1	2	3	4	5	6	7	8	9	10
$ap(k)$	11	21	31	41	22	32	42	33	43	44

Lower-triangular matrix element  $a_{ij}$  is stored in array element  $ap(i+(j-1)\times(2n-j)/2)$ .

## Usage

SCILIB:

```

CHARACTER*1      uplo, trans, diag
INTEGER*8       n, incx
REAL*8         ap(lenap), x(lenx)
CALL STPMV(uplo, trans, diag, n, ap, x, incx)

CHARACTER*1      uplo, trans, diag
INTEGER*8       n, incx
COMPLEX*16     ap(lenap), x(lenx)
CALL CTPMV(uplo, trans, diag, n, ap, x, incx)

```

## Input

**uplo**            Upper/lower triangular option for A:  
           'L' or 'l'    A is lower triangular  
           'U' or 'u'    A is upper triangular

**trans**            Transposition option for A:  
           'N' or 'n'    Compute  $x \leftarrow Ax$   
           'T' or 't'    Compute  $x \leftarrow A^T x$   
           'C' or 'c'    Compute  $x \leftarrow A^* x$

where  $A^T$  is the transpose of A and  $A^*$  is the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.

**diag**            Specifies whether the matrix is unit triangular, i.e.,  $a_{ii} = 1$ , or not:  
           'N' or 'n'    The diagonal of A is stored in the array  
           'U' or 'u'    The diagonal of A consists of unstored ones

When **diag** is supplied as 'U' or 'u', the diagonal elements are not referenced.

**n**                Number of rows and columns in matrix A,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference **ap** or **x**.

**ap**                    Array of length  $\text{lenap} = n \times (n+1)/2$  containing the  $n$ -by- $n$  triangular matrix  $A$ , stored by columns in the packed form described above. Space must be left for the diagonal elements of  $A$  even when **diag** is supplied as 'U' or 'u'.

**x**                        Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the input vector  $x$ .

**incx**                    Increment for the array  $x$ ,  $\text{incx} \neq 0$ :

**incx** > 0             $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx** < 0             $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.

## Output

**x**                        The updated  $x$  vector replaces the input.

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo**  $\neq$  'L' or 'l' or 'U' or 'u',  
**trans**  $\neq$  'N' or 'n' or 'T' or 't' or 'C' or 'c',  
**diag**  $\neq$  'N' or 'n' or 'U' or 'u',  
**n** < 0, and  
**incx** = 0.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **trans** argument as 'NORMAL' or 'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

**Example 1**

Form the REAL\*8 matrix-vector product  $Ax$ , where  $A$  is a 9-by-9 unit-diagonal, lower-triangular real matrix stored in packed form in an array AP of dimension 55 and  $x$  is a real vector 9 elements long stored in an array X of dimension 10.

```

CHARACTER*1 UPLO,TRANS,DIAG
INTEGER*8   N, INCX
REAL*8      AP(55),X(10)
UPLO = 'L'
TRANS = 'N'
DIAG = 'U'
N = 9
INCX = 1
CALL STPMV (UPLO,TRANS,DIAG,N,AP,X,INCX)

```

**Example 2**

Form the REAL\*8 matrix-vector product  $A^T x$ , where  $A$  is a 9-by-9 nonunit-diagonal, upper-triangular real matrix stored in packed form in an array AP of dimension 55 and  $x$  is a real vector 9 elements long stored in an array X of dimension 10.

```

INTEGER*8 N
REAL*8    AP(55),X(10)
N = 6
CALL STPMV ('UPPER', 'TRANSPOSE', 'NONUNIT', N, AP, X, 1)

```

**NAME** STPSV/CTPSV – Solve Triangular System

**Purpose**

Given an  $n$ -by- $n$  upper- or lower-triangular matrix  $A$  stored in packed form as described in “Matrix Storage,” and an  $n$ -vector  $x$ , these subprograms overwrite  $x$  with the solution  $y$  to the system of linear equations  $Ay = x$ . This is the forward elimination or back substitution step of Gaussian elimination.

Optionally,  $A$  may be replaced by  $A^T$ , the transpose of  $A$ , or by  $A^*$ , the conjugate transpose of  $A$ . Specifically, these subprograms compute

$$x \leftarrow A^{-1}x, \quad x \leftarrow A^{-T}x, \quad \text{and} \quad x \leftarrow A^*x$$

where  $A^{-T}$  is the inverse of the transpose of  $A$ , and  $A^*$  is the inverse of the conjugate transpose of  $A$ .

These subprograms are more primitive than the LINPACK linear equation solvers. As such, they are intended to supplement but not replace the equation solvers, serving instead as building blocks in constructing optimized linear algebra software. In fact, many of the LINPACK subprograms have been recoded to call these subprograms.

**Matrix Storage**

You supply the upper or lower triangle of  $A$ , stored column-by-column in packed form in a 1-dimensional array. This saves memory compared to storing the entire matrix.

The following examples illustrate the packed storage of a triangular matrix.

**Upper triangular matrix**

If  $A$  is

$$\begin{array}{cccc} 11 & 12 & 13 & 14 \\ 0 & 22 & 23 & 24 \\ 0 & 0 & 33 & 34 \\ 0 & 0 & 0 & 44 \end{array}$$

then  $A$  is packed column by column into an array  $\mathbf{ap}$  as follows:

$k$	1	2	3	4	5	6	7	8	9	10
$\mathbf{ap}(k)$	11	12	22	13	23	33	14	24	34	44

Upper-triangular matrix element  $a_{ij}$  is stored in array element  $\mathbf{ap}(i+j \times (j-1)/2)$ .

Basic Matrix Operations  
 STPSV/CTPSV – Solve Triangular System

**Lower triangular matrix**

If  $A$  is

11	0	0	0
21	22	0	0
31	32	33	0
41	42	43	44

then  $A$  is packed column by column into an array  $ap$  as follows:

$k$	1	2	3	4	5	6	7	8	9	10
$ap(k)$	11	21	31	41	22	32	42	33	43	44

Lower-triangular matrix element  $a_{ij}$  is stored in array element  $ap(i+(j-1)\times(2n-j)/2)$ .

**Usage**

SCILIB:

```

CHARACTER*1      uplo, trans, diag
INTEGER*8        n, incx
REAL*8           ap(lenap), x(lenx)
CALL STPSV(uplo, trans, diag, n, ap, x, incx)

CHARACTER*1      uplo, trans, diag
INTEGER*8        n, incx
COMPLEX*16       ap(lenap), x(lenx)
CALL CTPSV(uplo, trans, diag, n, ap, x, incx)
  
```

**Input**

<b>uplo</b>	Upper/lower triangular option for $A$ : 'L' or 'l'    Solve lower-triangular system (forward elimination) 'U' or 'u'    Solve upper-triangular system (back substitution)
<b>trans</b>	Transposition option for $A$ : 'N' or 'n'    Compute $x \leftarrow A^{-1}x$ 'T' or 't'    Compute $x \leftarrow A^{-T}x$ 'C' or 'c'    Compute $x \leftarrow A^{-*}x$

where  $A^{-T}$  is the inverse of the transpose of  $A$ , and  $A^{-*}$  is the inverse of the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.

**diag** Specifies whether the matrix is unit triangular, i.e.,  $a_{ii} = 1$ , or not:  
     'N' or 'n'     The diagonal of  $A$  is stored in the array  
     'U' or 'u'     The diagonal of  $A$  consists of unstored ones

When **diag** is supplied as 'U' or 'u', the diagonal elements are not referenced.

**n**       Number of rows and columns in matrix  $A$ ,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference **ap** or **x**.

**ap**       Array of length  $\text{lenap} = n \times (n+1)/2$  containing the  $n$ -by- $n$  triangular matrix  $A$ , stored by columns in the packed form described above. Space must be left for the diagonal elements of  $A$  even when **diag** is supplied as 'U' or 'u'.

**x**        Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the right-hand side  $n$ -vector  $x$ .

**incx**     Increment for the array **x**,  $\text{incx} \neq 0$ :  
      $\text{incx} > 0$       $x$  is stored forward in array **x**, i.e.,  $x_i$  is stored in  $\mathbf{x}((i-1) \times \text{incx} + 1)$ .  
      $\text{incx} < 0$       $x$  is stored backward in array **x**, i.e.,  $x_i$  is stored in  $\mathbf{x}((i-n) \times \text{incx} + 1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array **x**, i.e., if  $x_i$  is stored in  $\mathbf{x}(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.

## Output

**x**        The solution vector of the triangular system replaces the input.

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

The subprograms do not check for singularity of matrix  $A$ .  $A$  is singular if **diag** = 'N' or 'n' and some  $a_{ii} = 0$ . This condition will cause a division by zero to occur. Therefore, the program must detect singularity and take appropriate action to avoid a problem before calling any of these subprograms.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of

Basic Matrix Operations  
STPSV/CTPSV – Solve Triangular System

this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo** ≠ 'L' or 'l' or 'U' or 'u',  
**trans** ≠ 'N' or 'n' or 'T' or 't' or 'C' or 'c',  
**diag** ≠ 'N' or 'n' or 'U' or 'u',  
**n** < 0, and  
**incx** = 0.

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **trans** argument as 'NORMAL' or 'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

### Example 1

Perform REAL\*8 forward elimination using a 75-by-75 unit-diagonal, lower-triangular real matrix stored in packed form in an array AP of dimension 5500, and x is a real vector 75 elements long stored in an array X of dimension 100.

```
CHARACTER*1  UPLO, TRANS, DIAG
INTEGER*8    N, INCX
REAL*8       AP(5500), X(100)
UPLO = 'L'
TRANS = 'N'
DIAG = 'U'
N = 75
INCX = 1
CALL STPSV (UPLO, TRANS, DIAG, N, AP, X, INCX)
```

### Example 2

Perform REAL\*8 back substitution using a 75-by-75 nonunit-diagonal, upper-triangular real matrix stored in packed form in an array AP of dimension 5500, and x is a real vector 75 elements long stored in an array X of dimension 100.

```
INTEGER*8 N
REAL*8     AP(5500), X(100)
N = 75
CALL STPSV ('UPPER', 'NONTRANS', 'NONUNIT', N, AP, X, 1)
```

**NAME** STRMM/CTRMM – Triangular Matrix-Matrix Multiply

### Purpose

Given a scalar  $\alpha$ , an  $m$ -by- $n$  matrix  $B$ , and an upper- or lower-triangular matrix  $A$ , these subprograms compute either of the matrix-matrix products  $\alpha AB$  or  $\alpha BA$ . The size of  $A$ , either  $m$  by  $m$  or  $n$  by  $n$ , depends on which matrix product is requested. Optionally,  $A$  may be replaced by  $A^T$ , the transpose of  $A$ , or by  $A^*$ , the conjugate transpose of  $A$ . The resulting matrix product overwrites the input  $B$  matrix. Specifically, these subprograms compute matrix products of the forms

$$\begin{aligned} B &\leftarrow \alpha AB, & B &\leftarrow \alpha A^T B, & B &\leftarrow \alpha A^* B, \\ B &\leftarrow \alpha BA, & B &\leftarrow \alpha B A^T, & B &\leftarrow \alpha B A^*. \end{aligned}$$

### Matrix Storage

For these subprograms, you supply  $A$  in a two-dimensional array large enough to hold a square matrix. The other triangle of the array is not referenced. If  $A$  has an unstored unit diagonal (see input argument **diag**), then the diagonal elements of the array also will not be referenced.

### Usage

SCILIB:

```
CHARACTER*1      side, uplo, transa, diag
INTEGER*8        m, n, lda, ldb
REAL*8           alpha, a(lda, *), b(ldb, *)
CALL STRMM(side, uplo, transa, diag, m, n, alpha, a, lda, b, ldb)

CHARACTER*1      side, uplo, transa, diag
INTEGER*8        m, n, lda, ldb
COMPLEX*16       alpha, a(lda, *), b(ldb, *)
CALL CTRMM(side, uplo, transa, diag, m, n, alpha, a, lda, b, ldb)
```

### Input

**side** Specifies whether triangular matrix  $A$  is the left or right matrix operand:

- 'L' or 'l'  $A$  is the left matrix operand: for example,  $B \leftarrow \alpha AB$
- 'R' or 'r'  $A$  is the right matrix operand: for example,  $B \leftarrow \alpha BA$

**uplo** Upper/lower triangular option for  $A$ :

- 'L' or 'l'  $A$  is a lower-triangular matrix
- 'U' or 'u'  $A$  is an upper-triangular matrix

## STRMM/CTRMM – Triangular Matrix-Matrix Multiply

<b>transa</b>	<p>Transposition option for <i>A</i>:</p> <p>'N' or 'n'      Use matrix <i>A</i> directly</p> <p>'T' or 't'      Use <math>A^T</math>, the transpose of <i>A</i></p> <p>'C' or 'c'      Use <math>A^*</math>, the conjugate transpose of <i>A</i></p> <p>In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.</p>				
<b>diag</b>	<p>Specifies whether the <i>A</i> matrix is unit triangular, i.e., <math>a_{ii} = 1</math>, or not:</p> <p>'N' or 'n'      The diagonal of <i>A</i> is stored in the array</p> <p>'U' or 'u'      The diagonal of <i>A</i> consists of unstored ones</p> <p>When <b>diag</b> is supplied as 'U' or 'u', the diagonal elements of <i>A</i> are not referenced.</p>				
<b>m</b>	Number of rows in matrix <i>B</i> , $m \geq 0$ . If $m = 0$ , the subprograms do not reference <b>a</b> or <b>b</b> .				
<b>n</b>	Number of columns in matrix <i>B</i> , $n \geq 0$ . If $n = 0$ , the subprograms do not reference <b>a</b> or <b>b</b> .				
<b>alpha</b>	The scalar $\alpha$ . If <b>alpha</b> = 0, the subprograms compute $B \leftarrow 0$ without referencing <b>a</b> .				
<b>a</b>	<p>Array whose upper or lower triangle, as specified by <b>uplo</b>, contains the upper- or lower-triangular matrix <i>A</i>, whose size is indicated by <b>side</b>:</p> <table> <tr> <td>'L' or 'l'</td> <td><i>A</i> is <math>m</math> by <math>m</math></td> </tr> <tr> <td>'R' or 'r'</td> <td><i>A</i> is <math>n</math> by <math>n</math></td> </tr> </table> <p>The other triangle of <b>a</b> is not referenced. Not used as input if <b>alpha</b> = 0.</p>	'L' or 'l'	<i>A</i> is $m$ by $m$	'R' or 'r'	<i>A</i> is $n$ by $n$
'L' or 'l'	<i>A</i> is $m$ by $m$				
'R' or 'r'	<i>A</i> is $n$ by $n$				
<b>lda</b>	The leading dimension of array <b>a</b> as declared in the calling program unit, with $lda \geq \max(\text{the number of rows of } A, 1)$ .				
<b>b</b>	Array containing the $m$ -by- $n$ matrix <i>B</i> . Not used as input if <b>alpha</b> = 0.				
<b>ldb</b>	The leading dimension of array <b>b</b> as declared in the calling program unit, with $ldb \geq \max(m, 1)$ .				
<b>Output</b>					
<b>b</b>	The indicated matrix product replaces the input.				

## Notes

These subprograms conform to specifications of the Level 3 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**side** ≠ 'L' or 'l' or 'R' or 'r',  
**uplo** ≠ 'U' or 'u' or 'L' or 'l',  
**transa** ≠ 'N' or 'n' or 'T' or 't' or 'C' or 'c',  
**diag** ≠ 'N' or 'n' or 'U' or 'u',  
**m** < 0,  
**n** < 0,  
**lda** too small, and  
**ldb** < max(**m**,1).

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved, for example, by coding the **transa** argument as 'NORMAL' or 'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

### Example 1

Form the REAL\*8 matrix product  $AB$ , where  $A$  is a 6-by-6 nonunit-diagonal, upper-triangular real matrix stored in an array  $A$  whose dimensions are 10 by 10, and  $B$  is a 6-by-8 real matrix stored in an array  $B$  of dimension 10 by 10. The matrix product will overwrite the input  $B$  matrix.

```

CHARACTER*1 SIDE,UPLO,TRANSA,DIAG
INTEGER*8   M,N,LDA,LDB
REAL*8     ALPHA,A(10,10),B(10,10)
SIDE = 'L'
UPLO = 'U'
TRANSA = 'N'
DIAG = 'N'
M = 6
N = 8
ALPHA = 1.0
LDA = 10
LDB = 10
CALL STRMM (SIDE,UPLO,TRANSA,DIAG,M,N,ALPHA,A,LDA,B,LDB)
  
```

### Example 2

Form the REAL\*8 matrix product  $qBA^T$ , where  $q$  is a real scalar,  $B$  is a 6-by-8 real matrix stored in an array  $B$  of dimension 10 by 10, and  $A$  is a 8-by-8

## Basic Matrix Operations

### STRMM/CTRMM – Triangular Matrix-Matrix Multiply

unit-diagonal lower-triangular real matrix stored in an array *A* whose dimensions are 10 by 10. The matrix product will overwrite the input *B* matrix.

```
INTEGER*8 M,N,LDA,LDB
REAL*8    Q,A(10,10),B(10,10)
M = 6
N = 8
LDA = 10
LDB = 10
CALL STRMM ( 'RIGHT', 'LOWER', 'TRANS', 'UNIT', M,N,Q,A,LDA,B,
             LDB)
```

**NAME** STRMV/CTRMV – Matrix-Vector Multiply

**Purpose**

Given an  $n$ -by- $n$  upper- or lower-triangular matrix  $A$ , and an  $n$ -vector  $x$ , these subprograms compute the matrix-vector products  $Ax$ ,  $A^T x$ , and  $A^* x$ , where  $A^T$  is the transpose of  $A$ , and  $A^*$  is the conjugate transpose of  $A$ . Specifically, these subprograms compute matrix-vector products of the forms

$$x \leftarrow Ax, \quad x \leftarrow A^T x, \quad \text{and} \quad x \leftarrow A^* x.$$

**Matrix Storage**

For these subprograms, you supply  $A$  in a two-dimensional array large enough to hold a square matrix. The other triangle of the array is not referenced. If  $A$  has an unstored unit diagonal (see input argument **diag**), then the diagonal elements of the array also will not be referenced.

**Usage**

SCILIB:

```

CHARACTER*1      uplo, trans, diag
INTEGER*8        n, lda, incx
REAL*8           a(lda, n), x(lenx)
CALL STRMV(uplo, trans, diag, n, a, lda, x, incx)

CHARACTER*1      uplo, trans, diag
INTEGER*8        n, lda, incx
COMPLEX*16       a(lda, n), x(lenx)
CALL CTRMV(uplo, trans, diag, n, a, lda, x, incx)

```

**Input**

**uplo** Upper/lower triangular option for  $A$ :  
'L' or 'l'      $A$  is lower triangular  
'U' or 'u'      $A$  is upper triangular

The other triangle of the array  $a$  is not referenced.

**trans** Transposition option for  $A$ :  
'N' or 'n'     Compute  $x \leftarrow Ax$   
'T' or 't'     Compute  $x \leftarrow A^T x$   
'C' or 'c'     Compute  $x \leftarrow A^* x$

where  $A^T$  is the transpose of  $A$  and  $A^*$  is the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.

## Basic Matrix Operations

### STRMV/CTRMV – Matrix-Vector Multiply

<b>diag</b>	Specifies whether the matrix is unit triangular, i.e., $a_{ii} = 1$ , or not: <b>‘N’ or ‘n’</b> The diagonal of $A$ is stored in the array <b>‘U’ or ‘u’</b> The diagonal of $A$ consists of unstored ones When <b>diag</b> is supplied as <b>‘U’ or ‘u’</b> , the diagonal elements are not referenced.
<b>n</b>	Number of rows and columns in matrix $A$ , $n \geq 0$ . If $n = 0$ , the subprograms do not reference <b>a</b> or <b>x</b> .
<b>a</b>	Array containing the $n$ -by- $n$ triangular matrix $A$ .
<b>lda</b>	The leading dimension of array <b>a</b> as declared in the calling program unit, with $lda \geq \max(n,1)$ .
<b>x</b>	Array of length $lenx = (n-1) \times  incx  + 1$ containing the input vector $x$ .
<b>incx</b>	Increment for the array <b>x</b> , $incx \neq 0$ : <b>incx &gt; 0</b> $x$ is stored forward in array <b>x</b> , i.e., $x_i$ is stored in $x((i-1) \times incx + 1)$ . <b>incx &lt; 0</b> $x$ is stored backward in array <b>x</b> , i.e., $x_i$ is stored in $x((i-n) \times incx + 1)$ . Use <b>incx = 1</b> if the vector $x$ is stored contiguously in array <b>x</b> , i.e., if $x_i$ is stored in $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

### Output

<b>x</b>	The updated $x$ vector replaces the input.
----------	--

### Notes

These subprograms conform to specifications of the Level 2 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo**  $\neq$  **‘L’ or ‘l’ or ‘U’ or ‘u’**,  
**trans**  $\neq$  **‘N’ or ‘n’ or ‘T’ or ‘t’ or ‘C’ or ‘c’**,  
**diag**  $\neq$  **‘N’ or ‘n’ or ‘U’ or ‘u’**,  
**n**  $< 0$ ,  
**lda**  $< \max(n,1)$ , and

**incx = 0.**

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **trans** argument as 'NORMAL' or 'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

### Example 1

Form the **REAL\*8** matrix-vector product  $Ax$ , where  $A$  is a 9-by-9 unit-diagonal lower-triangular real matrix stored in an array  $A$  whose dimensions are 10 by 10, and  $x$  is a real vector 9 elements long stored in an array  $X$  of dimension 10.

```

CHARACTER*1 UPLO,TRANS,DIAG
INTEGER*8   N,LDA,INCX
REAL*8     A(10,10),X(10)
UPLO = 'L'
TRANS = 'N'
DIAG = 'U'
N = 9
LDA = 10
INCX = 1
CALL STRMV (UPLO,TRANS,DIAG,N,A,LDA,X,INCX)

```

### Example 2

Form the **REAL\*8** matrix-vector product  $A^T x$ , where  $A$  is a 9-by-9 nonunit-diagonal, upper-triangular real matrix stored in an array  $A$  whose dimensions are 10 by 10, and  $x$  is a real vector 9 elements long stored in an array  $X$  of dimension 10.

```

INTEGER*8 N,LDA
REAL*8   A(10,10),X(10)
N = 6
LDA = 10
CALL STRMV ('UPPER','TRANSPOSE','NONUNIT',N,A,LDA,X,1)

```

Basic Matrix Operations  
**STRSM/CTRSM – Solve Triangular Systems**

**NAME** STRSM/CTRSM – Solve Triangular Systems

**Purpose**

Given a scalar  $\alpha$ , an upper- or lower-triangular matrix  $A$ , and an  $m$ -by- $n$  matrix  $B$ , these subprograms compute either of the matrix solutions  $\alpha A^{-1}B$  or  $\alpha BA^{-1}$ . The size of  $A$ , either  $m$  by  $m$  or  $n$  by  $n$ , depends on which matrix solution is requested. Optionally,  $A^{-1}$  may be replaced by  $A^{-T}$ , the inverse of the transpose of  $A$ , or by  $A^*$ , the inverse of the conjugate transpose of  $A$ . The resulting matrix solution overwrites the input  $B$  matrix. Specifically, these subprograms compute matrix solutions of the forms

$$\begin{aligned}
 B &\leftarrow \alpha A^{-1}B, & B &\leftarrow \alpha A^{-T}B, & B &\leftarrow \alpha A^*B, \\
 B &\leftarrow \alpha BA^{-1}, & B &\leftarrow \alpha BA^{-T}, & B &\leftarrow \alpha BA^*.
 \end{aligned}$$

**Matrix Storage**

For these subprograms, you supply  $A$  in a two-dimensional array large enough to hold a square matrix. The other triangle of the array is not referenced. If  $A$  has an unstored unit diagonal (see input argument **diag**), then the diagonal elements of the array also will not be referenced.

**Usage**

SCILIB:

```

CHARACTER*1      side, uplo, transa, diag
INTEGER*8        m, n, lda, ldb
REAL*8          alpha, a(lda, *), b(ldb, *)
CALL STRSM(side, uplo, transa, diag, m, n, alpha, a, lda, b, ldb)

CHARACTER*1      side, uplo, transa, diag
INTEGER*8        m, n, lda, ldb
COMPLEX*16      alpha, a(lda, *), b(ldb, *)
CALL CTRSM(side, uplo, transa, diag, m, n, alpha, a, lda, b, ldb)

```

**Input**

**side** Specifies whether triangular matrix  $A$  is the left or right matrix operand:  
 $L'$  or  $l'$   $A$  is the left matrix operand: for example,  $B \leftarrow \alpha A^{-1}B$   
 $R'$  or  $r'$   $A$  is the right matrix operand: for example,  $B \leftarrow \alpha BA^{-1}$

**uplo** Upper/lower triangular option for  $A$ :

	<p>‘L’ or ‘l’      <math>A</math> is a lower-triangular matrix</p> <p>‘U’ or ‘u’      <math>A</math> is an upper-triangular matrix</p>				
<b>transa</b>	<p>Transposition option for <math>A</math>:</p> <p>‘N’ or ‘n’      Use matrix <math>A^{-1}</math></p> <p>‘T’ or ‘t’      Use <math>A^{-T}</math>, the inverse of the transpose of <math>A</math></p> <p>‘C’ or ‘c’      Use <math>A^*</math>, the inverse of the conjugate transpose of <math>A</math></p>				
	<p>In the real subprograms, ‘C’ and ‘c’ have the same meaning as ‘T’ and ‘t’.</p>				
<b>diag</b>	<p>Specifies whether the <math>A</math> matrix is unit triangular, i.e., <math>a_{ii} = 1</math>, or not:</p> <p>‘N’ or ‘n’      The diagonal of <math>A</math> is stored in the array</p> <p>‘U’ or ‘u’      The diagonal of <math>A</math> consists of unstored ones</p> <p>When <b>diag</b> is supplied as ‘U’ or ‘u’, the diagonal elements of <math>A</math> are not referenced.</p>				
<b>m</b>	<p>Number of rows in matrix <math>B</math>, <math>m \geq 0</math>. If <math>m = 0</math>, the subprograms do not reference <b>a</b> or <b>b</b>.</p>				
<b>n</b>	<p>Number of columns in matrix <math>B</math>, <math>n \geq 0</math>. If <math>n = 0</math>, the subprograms do not reference <b>a</b> or <b>b</b>.</p>				
<b>alpha</b>	<p>The scalar <math>\alpha</math>. If <b>alpha</b> = 0, the subprograms compute <math>B \leftarrow 0</math> without referencing <b>a</b>.</p>				
<b>a</b>	<p>Array whose upper or lower triangle, as specified by <b>uplo</b>, contains the upper- or lower-triangular matrix <math>A</math>, whose size is indicated by <b>side</b>:</p> <table style="margin-left: 40px; border: none;"> <tr> <td style="padding-right: 20px;">‘L’ or ‘l’</td> <td><math>A</math> is <math>m</math> by <math>m</math></td> </tr> <tr> <td>‘R’ or ‘r’</td> <td><math>A</math> is <math>n</math> by <math>n</math></td> </tr> </table> <p>The other triangle of <b>a</b> is not referenced. Not used as input if <b>alpha</b> = 0.</p>	‘L’ or ‘l’	$A$ is $m$ by $m$	‘R’ or ‘r’	$A$ is $n$ by $n$
‘L’ or ‘l’	$A$ is $m$ by $m$				
‘R’ or ‘r’	$A$ is $n$ by $n$				
<b>lda</b>	<p>The leading dimension of array <b>a</b> as declared in the calling program unit, with <b>lda</b> <math>\geq</math> max (the number of rows of <math>A</math>, 1).</p>				
<b>b</b>	<p>Array containing the <math>m</math>-by-<math>n</math> matrix <math>B</math>. Not used as input if <b>alpha</b> = 0.</p>				
<b>ldb</b>	<p>The leading dimension of array <b>b</b> as declared in the calling program unit, with <b>ldb</b> <math>\geq</math> max(<b>m</b>, 1).</p>				

Basic Matrix Operations  
STRSM/CTRSM – Solve Triangular Systems

## Output

- b** The indicated matrix solution replaces the input.

## Notes

These subprograms conform to specifications of the Level 3 BLAS.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**side** ≠ 'L' or 'l' or 'R' or 'r',  
**uplo** ≠ 'L' or 'l' or 'U' or 'u',  
**transa** ≠ 'N' or 'n' or 'T' or 't' or 'C' or 'c',  
**diag** ≠ 'N' or 'n' or 'U' or 'u',  
**m** < 0,  
**n** < 0,  
**lda** too small, and  
**ldb** < max(m,1).

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the CALL statement may be improved, for example, by coding the **transa** argument as 'NORMAL' or 'NONTRANS' for 'N', 'TRANSPOSE' for 'T', or 'CTRANS' for 'C'. Refer to "Example 2."

## Example 1

Form the REAL\*8 matrix solution  $A^{-1}B$ , where  $A$  is a 6-by-6 nonunit-diagonal, upper-triangular real matrix stored in an array  $A$  whose dimensions are 10 by 10 and  $B$  is a 6-by-8 real matrix stored in an array  $B$  of dimension 10 by 10. The matrix solution will overwrite the input  $B$  matrix.

```
CHARACTER*1 SIDE,UPLO,TRANSA,DIAG
INTEGER*8   M,N,LDA,LDB
REAL*8     ALPHA,A(10,10),B(10,10)
SIDE = 'L'
UPLO = 'U'
TRANSA = 'N'
DIAG = 'N'
M = 6
N = 8
ALPHA = 1.0
LDA = 10
LDB = 10
CALL STRSM (SIDE,UPLO,TRANSA,DIAG,M,N,ALPHA,A,LDA,B,LDB)
```

## Example 2

Form the REAL\*8 matrix solution  $qBA^{-T}$ , where  $q$  is a real scalar,  $B$  is a 6-by-8 real matrix stored in an array B of dimension 10 by 10, and  $A$  is a 8-by-8 unit-diagonal lower-triangular real matrix stored in an array A whose dimensions are 10 by 10. The matrix solution will overwrite the input  $B$  matrix.

```
INTEGER*8 M,N,LDA,LDB
REAL*8    Q,A(10,10),B(10,10)
M = 6
N = 8
LDA = 10
LDB = 10
CALL STRSM ('RIGHT', 'LOWER', 'TRANS', 'UNIT', M,N,Q,A,LDA,B,LDB)
```

**STRSV/CTRSV – Solve Triangular System****NAME** STRSV/CTRSV – Solve Triangular System**Purpose**

Given an  $n$ -by- $n$  upper- or lower-triangular matrix  $A$ , and an  $n$ -vector  $x$ , these subprograms overwrite  $x$  with the solution  $y$  to the system of linear equations  $Ay = x$ . This is the forward elimination or back substitution step of Gaussian elimination. Optionally,  $A$  may be replaced by  $A^T$ , the transpose of  $A$ , or by  $A^*$ , the conjugate transpose of  $A$ . Specifically, these subprograms compute

$$x \leftarrow A^{-1}x, \quad x \leftarrow A^{-T}x, \quad \text{and} \quad x \leftarrow A^{-*}x$$

where  $A^{-T}$  is the inverse of the transpose of  $A$ , and  $A^{-*}$  is the inverse of the conjugate transpose of  $A$ .

These subprograms are more primitive than the LINPACK linear equation solvers. As such, they are intended to supplement but not replace the equation solvers, serving instead as building blocks in constructing optimized linear algebra software. In fact, many of the LINPACK subprograms have been recoded to call these subprograms.

**Matrix Storage**

For these subprograms, you supply  $A$  in a two-dimensional array large enough to hold a square matrix. The other triangle of the array is not referenced. If  $A$  has an unstored unit diagonal (see input argument **diag**), then the diagonal elements of the array also will not be referenced.

**Usage**

SCILIB:

```

CHARACTER*1      uplo, trans, diag
INTEGER*8        n, lda, incx
REAL*8           a(lda, n), x(lenx)
CALL STRSV(uplo, trans, diag, n, a, lda, x, incx)

CHARACTER*1      uplo, trans, diag
INTEGER*8        n, lda, incx
COMPLEX*16       a(lda, n), x(lenx)
CALL CTRSV(uplo, trans, diag, n, a, lda, x, incx)

```

**Input**

<b>uplo</b>	Upper/lower triangular option for $A$ :
	‘L’ or ‘l’      Solve lower-triangular system (forward elimination)

'U' or 'u'      Solve upper-triangular system (back substitution)

The other triangle of the array **a** is not referenced.

**trans**

Transposition option for **A**:

'N' or 'n'      Compute  $x \leftarrow A^{-1}x$

'T' or 't'      Compute  $x \leftarrow A^{-T}x$

'C' or 'c'      Compute  $x \leftarrow A^*x$

where  $A^{-T}$  is the inverse of the transpose of **A**, and  $A^*x$  is the inverse of the conjugate transpose. In the real subprograms, 'C' and 'c' have the same meaning as 'T' and 't'.

**diag**

Specifies whether the matrix is unit triangular, i.e.,  $a_{ii} = 1$ , or not:

'N' or 'n'      The diagonal of **A** is stored in the array

'U' or 'u'      The diagonal of **A** consists of unstored ones

When **diag** is supplied as 'U' or 'u', the diagonal elements are not referenced.

**n**

Number of rows and columns in matrix **A**,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference **a** or **x**.

**a**

Array containing the  $n$ -by- $n$  triangular matrix **A**.

**lda**

The leading dimension of array **a** as declared in the calling program unit, with  $lda \geq \max(n, 1)$ .

**x**

Array of length  $lenx = (n-1) \times |incx| + 1$  containing the right-hand side  $n$ -vector  $x$ .

**incx**

Increment for the array **x**,  $incx \neq 0$ :

$incx > 0$        $x$  is stored forward in array **x**, i.e.,  $x_i$  is stored in  $x((i-1) \times incx + 1)$ .

$incx < 0$        $x$  is stored backward in array **x**, i.e.,  $x_i$  is stored in  $x((i-n) \times incx + 1)$ .

Use  $incx = 1$  if the vector  $x$  is stored contiguously in array **x**, i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.

## Output

**x**

The solution vector of the triangular system replaces the input.

Basic Matrix Operations  
STRSV/CTRSV – Solve Triangular System

## Notes

These subprograms conform to specifications of the Level 2 BLAS.

The subprograms do not check for singularity of matrix  $A$ .  $A$  is singular if  $\text{diag} = \text{'N'}$  or  $\text{'n'}$  and some  $a_{ii} = 0$ . This condition will cause a division by zero to occur. Therefore, the program must detect singularity and take appropriate action to avoid a problem before calling any of these subprograms.

If an error in the arguments is detected, the subprograms call error handler XERBLA, which writes an error message onto the standard error file and terminates execution. The standard version of XERBLA (refer to the end of this chapter) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**uplo**  $\neq$   $\text{'L'}$  or  $\text{'l'}$  or  $\text{'U'}$  or  $\text{'u'}$ ,  
**trans**  $\neq$   $\text{'N'}$  or  $\text{'n'}$  or  $\text{'T'}$  or  $\text{'t'}$  or  $\text{'C'}$  or  $\text{'c'}$ ,  
**diag**  $\neq$   $\text{'N'}$  or  $\text{'n'}$  or  $\text{'U'}$  or  $\text{'u'}$ ,  
**n**  $< 0$ ,  
**lda**  $< \max(\text{n}, 1)$ , and  
**incx**  $= 0$ .

Actual character arguments in a subroutine call may be longer than the corresponding dummy arguments. Therefore, readability of the **CALL** statement may be improved by coding the **trans** argument as  $\text{'NORMAL'}$  or  $\text{'NONTRANS'}$  for  $\text{'N'}$ ,  $\text{'TRANSPOSE'}$  for  $\text{'T'}$ , or  $\text{'CTRANS'}$  for  $\text{'C'}$ . Refer to "Example 2."

## Example 1

Perform **REAL\*8** forward elimination using the 75-by-75 unit-diagonal lower-triangular real matrix stored in an array **A** whose dimensions are 100 by 100, and **x** is a real vector 75 elements long stored in an array **X** of dimension 100.

```
CHARACTER*1  UPLO, TRANS, DIAG
INTEGER*8    N, LDA, INCX
REAL*8       A(100,100), X(100)
UPLO = 'L'
TRANS = 'N'
DIAG = 'U'
N = 75
LDA = 100
INCX = 1
CALL STRSV (UPLO, TRANS, DIAG, N, A, LDA, X, INCX)
```

## Example 2

Perform **REAL\*8** back substitution using the 75-by-75 nonunit-diagonal, upper-triangular real matrix stored in an array **A** whose dimensions are 100 by

Basic Matrix Operations  
**STRSV/CTRSV – Solve Triangular System**

100, and  $x$  is a real vector 75 elements long stored in an array  $x$  of dimension 100.

```
INTEGER*8 N,LDA
REAL*8    A(100,100),X(100)
N = 75
LDA = 100
CALL STRSV ('UPPER','NONTRANS','NONUNIT',N,A,LDA,X,1)
```

Basic Matrix Operations  
**SXMPY – Vector-Matrix Multiply and Add**

**NAME** SXMPY – Vector-Matrix Multiply and Add

**Purpose**

This subprogram computes the matrix-vector product  $xA$ , and adds the result to another vector  $y$ , where  $A$  is an  $m$ -by- $n$  matrix,  $x$  is an  $m$ -dimensional row vector, and  $y$  is an  $n$ -dimensional row vector. SCILIB subprogram SGEMV allows more general storage of  $x$  and  $y$  and also admits scaling, subtraction, and transposing  $A$ .

**Usage**

SCILIB:

```
INTEGER*8      n, m, lda, incx, incy
REAL*8         a(lda, n), x(lenx), y(leny)
CALL SXMPY(n, incy, y, m, incx, x, lda, a)
```

**Input**

<b>n</b>	Number of columns in matrix $A$ and length of row vector $y$ , $n \geq 0$ . If $n = 0$ , the subprogram does not reference $a$ , $x$ , or $y$ .
<b>incy</b>	Storage increment between successive elements of vector $y$ in array $y$ . $y_i$ is stored in $y((i-1) \times \text{incy} + 1)$ . Use <b>incy</b> = 1 if the vector $y$ is stored contiguously in array $y$ , i.e., if $y_i$ is stored in $y(i)$ .
<b>y</b>	Array of length <b>leny</b> = $(n-1) \times \text{incy} + 1$ containing the row vector $y$ .
<b>m</b>	Number of rows in matrix $A$ and length of row vector $x$ , $m \geq 0$ . If $m = 0$ , the subprogram does not reference $a$ , $x$ , or $y$ .
<b>incx</b>	Storage increment between successive elements of vector $x$ in array $x$ . $x_i$ is stored in $x((i-1) \times \text{incx} + 1)$ . Use <b>incx</b> = 1 if the vector $x$ is stored contiguously in array $x$ , i.e., if $x_i$ is stored in $x(i)$ .
<b>x</b>	Array of length <b>lenx</b> = $(m-1) \times \text{incx} + 1$ containing the $n$ -vector $x$ .
<b>lda</b>	The leading dimension of array $a$ as declared in the calling program unit.
<b>a</b>	Array containing the $m$ -by- $n$ matrix $A$ .

**Output**

<b>y</b>	The updated $y$ vector replaces the input.
----------	--

## Notes

Cray research, Inc. has declared this subprogram obsolete in release 6.0 of the UNICOS Math and Scientific Library.

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**m** < 0,  
**n** < 0,  
**lda** < **m**,  
**incx** = 0, and  
**incy** = 0.

## Fortran Equivalent

Except for the argument error checking, the following Fortran subroutine is equivalent to SXMPY.

```

SUBROUTINE SXMPY (N, INCY, Y, M, INCX, X, LDA, A)
  INTEGER*8 M, N, LDA, INCX, INCY
  REAL*8 A(LDA, N), X(*), Y(*)
  DO 120 J = 1, M
    DO 110 I = 1, N
      Y((I-1)*INCY+1) =           ! Y(I) =
1      Y((I-1)*INCY+1) +           ! Y(I) +
2      X((J-1)*INCX+1) * A(J, I) ! X(J) * A(J, I)
110   CONTINUE
120  CONTINUE
      RETURN
      END

```

## Example

Form the REAL\*8 vector-matrix product  $y = y + xA$ , where  $A$  is a 9-by-6 real matrix stored in an array  $A$  whose dimensions are 10 by 10,  $x$  is a real vector 9 elements long stored in row 3 of an array  $X$  of dimension 10 by 10, and  $y$  is a real vector 6 elements long stored in row 7 of an array  $Y$  of dimension 10 by 10.

```

INTEGER*8 M, N, LDA, INCX, INCY
REAL*8 A(10, 10), X(10, 10), Y(10, 10)
M = 9
N = 6
LDA = 10
INCX = 10
INCY = 10
CALL SXMPY (N, INCY, Y(7, 1), M, INCX, X(3, 1), LDA, A)

```

Basic Matrix Operations  
XERBLA – Error Handler

**NAME** XERBLA – Error Handler

**Purpose**

This subprogram is the error handler for many of the subprograms in this chapter, as indicated in the “Notes” section in the applicable subprogram descriptions. As supplied in SCILIB, XERBLA writes the following error message onto the standard error file:

```
*****  
* XERBLA: subprogram name called with invalid value of argument number iarg  
*  
*****
```

where *name* is the name of the subprogram in which the error was detected, and *iarg* is the argument number of the offending argument. For example, in SGEMV, *trans* is argument number 1 and *m* is argument number 2. If the main program is in Fortran and is executed under SPP-UX, a call traceback is also written onto the standard error file (XERBLA does not write a call traceback when used on HP-UX systems). XERBLA then terminates execution with a nonzero exit status.

You may supply a version of XERBLA that alters this action. Be aware that other subprograms, including many in LAPACK, also call XERBLA. All BLAS, VECLIB, SCILIB, and LAPACK subprograms that call XERBLA follow the **CALL XERBLA** statement with a **RETURN** statement, so your version of XERBLA can exit with a **RETURN** statement. However, many of those subprograms do not have a status response variable in their argument list that could be used to alert the caller. If you write a XERBLA that does not end with a **STOP** statement, you need some other mechanism to detect errors occurring in those subprograms. One such mechanism is a flag in a common block that is set by your XERBLA and tested by the calling program after calls where errors could be detected.

**Usage**

SCILIB:

```
CHARACTER*6    name  
INTEGER*8     iarg  
CALL XERBLA(name, iarg)
```

**Input**

<b>name</b>	The name of the subprogram in which the error was detected.
<b>iarg</b>	The number of the argument that was found to be in error.

## Notes

This subprogram conforms to specifications of the Level 2 and 3 BLAS and LAPACK.

Basic Matrix Operations  
XERBLA – Error Handler

## 4 Linear Equations

---

### Overview

This chapter describes the LINPACK software library included with SCILIB.

The most important subprograms in this library have been upgraded by incorporating the Level 2 and Level 3 BLAS and other algorithmic changes.

Although SCILIB includes all LINPACK subprograms, only those subprograms optimized for use on Hewlett-Packard supercomputers are described in this chapter. Table 4-5 at the end of this chapter lists the LINPACK subprograms that are included in SCILIB but are not documented in the *HP MLIB SCILIB User's Guide*. You may find information for these subprograms in the *LINPACK Users' Guide* included in the SCILIB documentation set.

The LAPACK software library included with SCILIB is a comprehensive collection of linear equation solvers and subprograms for other linear algebra computations. This software is documented in the *HP MLIB LAPACK User's Guide*. We recommend that you use LAPACK subprograms rather than LINPACK subprograms in new programs. Future optimization efforts will be directed to LAPACK rather than LINPACK.

This chapter explains how to use SCILIB subprograms to solve systems of linear equations. The operations covered are:

- solution of a system of linear equations
- calculation of the inverse of a matrix
- evaluation of the determinant of a matrix

These operations are performed for a variety of types of matrices, including:

- real and complex general dense matrices
- real and complex general band matrices
- real and complex positive definite dense matrices
- real and complex positive definite band matrices
- real and complex general tridiagonal matrices
- real and complex positive definite tridiagonal matrices

Refer to Chapter 6 for software to solve sparse symmetric linear equations.

---

## Chapter Objectives

After you read this chapter you will:

- be familiar with the LINPACK subroutine naming convention
- understand the role of the condition number in solving linear equations
- know how to compute the determinant or inverse of a matrix
- know when not to compute the determinant or inverse of a matrix
- be able to locate documentation for LINPACK subroutines not documented here
- know how to use the described subprograms

---

## What You Need to Know to Use These Subprograms

### Subroutine Naming Convention

LINPACK uses a subroutine naming convention that encodes the function of each subroutine into its name. LINPACK subprogram names consist of five letters in the form TXXYY.

The first letter in the naming convention indicates one of the four Fortran data types, as shown in Table 4-1.

Table 4-1 LINPACK Naming Convention — Data Type

T	Data Type
S	Single Precision REAL
C	Single Precision COMPLEX

Table 4-2 shows the next two letters in the naming convention which indicate the form of the matrix or its decomposition.

**Table 4-2 LINPACK Naming Convention — Form or Decomposition**

XX	Form or Decomposition
GE	General
GB	General band
PO	Positive definite
PB	Positive definite band
PP	Positive definite packed
SI	Symmetric indefinite
SP	Symmetric indefinite packed
HI	Hermitian indefinite
HP	Hermitian indefinite packed
TR	Triangular
GT	General tridiagonal
PT	Positive definite tridiagonal
CH	Cholesky decomposition
QR	Orthogonal-triangular decomposition
SV	Singular value decomposition

Table 4-3 lists the final two letters in the naming convention which indicate the computation of a particular subroutine.

**Table 4-3 LINPACK Naming Convention — Computation**

YY	Subroutine Computation
FA	Factor
CO	Factor and estimate condition
SL	Solve
DI	Determinant and/or inverse and/or inertia
DC	Decompose
UD	Update
DD	Downdate
EX	Exchange

For example, SGBCO factors a general band (GB) matrix and estimates its condition number (CO) using the single precision REAL data type (S). CGEFA calculates the factorization (FA) of a general dense matrix (GE) using the single precision COMPLEX data type (C).

Table 4-4 shows the valid combinations of T, XX, and YY. Each line indicates the allowable T prefixes and YY suffixes for a particular root name XX.

Table 4-4 LINPACK Naming Convention — Subprogram Names

Valid T		XX	Valid YY			
S	C	GE	CO	FA	SL	DI
S	C	GB	CO	FA	SL	DI
S	C	PO	CO	FA	SL	DI
S	C	PB	CO	FA	SL	DI
S	C	PP	CO	FA	SL	DI
S	C	SI	CO	FA	SL	DI
S	C	SP	CO	FA	SL	DI
	C	HI	CO	FA	SL	DI
	C	HP	CO	FA	SL	DI
S	C	TR	CO		SL	DI
S	C	GT			SL	
S	C	PT			SL	
S	C	CH	DC		UD	DD EX
S	C	QR	DC	SL		
S	C	SV	DC			

LINPACK is organized so that it is usually necessary to call two subprograms to perform the above operations. One subprogram is called to process the matrix and another is called to process a particular right-hand side. This division of labor significantly reduces computer time when there is a sequence of problems involving the same matrix but different right-hand sides. It also allows you the flexibility to choose between subprograms that are fast but use a less reliable, elementary test for singularity and subprograms that are slightly slower but use a significantly more reliable test involving an estimate of the condition number of the coefficient matrix.

### Condition Number

The condition number,  $\kappa(A)$ , of the coefficient matrix  $A$  measures the sensitivity of the solution  $x$  of the system of linear equations  $Ax = b$  to errors in the matrix  $A$  and the right-hand side  $b$ . If  $\delta A$  and  $\delta b$  represent the errors in  $A$  and  $b$ , respectively, and if  $\| \cdot \|$  represents any vector norm and its subordinate matrix norm, the error  $\delta x$  in  $x$  that results from solving  $(A+\delta A)(x+\delta x) = b+\delta b$  instead of  $Ax = b$  is bounded by

$$\frac{\|\delta x\|}{\|x\|} \leq \frac{\kappa(A)}{1 - \|A^{-1}\| \|\delta A\|} \left( \frac{\|\delta A\|}{\|A\|} + \frac{\|\delta b\|}{\|b\|} \right)$$

A standard result of numerical analysis shows that the roundoff error introduced by the solution process may be modeled by taking  $\|\delta A\|/\|A\|$  and  $\|\delta b\|/\|b\|$  to be small multiples of the computer's machine epsilon. Computational singularity of  $A$  results in  $\kappa(A) = \infty$ . A more common situation occurs when  $A$  is not numerically singular but is ill-conditioned. When a matrix is ill-conditioned,  $\kappa(A)$  is large, so small errors in the matrix and right-hand side and small roundoff errors introduced during the solution process itself are magnified greatly in the solution.

Since  $1 < \kappa(A) \leq \infty$ , it is more convenient to compute the reciprocal condition number,  $1/\kappa(A)$ , than  $\kappa(A)$  itself. The reciprocal condition number has the interpretation that if  $1/\kappa(A)$  approximately equals  $10^{-d}$ , elements of  $x$  can be expected to have  $d$  fewer significant digits of accuracy than the elements of  $A$  or  $b$ . Consequently, if errors in the coefficient matrix and right-hand side exceed  $1/\kappa(A)$ , or if  $1/\kappa(A)$  is negligible compared to 1.0, then  $x$  may have no significant digits at all.

## Determinant and Inverse

Subprograms for computing the determinant and inverse of a matrix are provided, although it is almost never necessary to compute either one.

While papers and reference books extensively use the notation " $\det(A) \neq 0$ " to mean "A is nonsingular," SCILIB includes both more efficient and more reliable subprograms for detecting singularity. Similarly, references frequently use " $A^{-1}b$ " to mean "the solution  $x$  of the system of linear equations  $Ax = b$ ." Again, it is more efficient and accurate to compute the solution directly than to invert the coefficient matrix and multiply the inverse times the right-hand side vector. This is true even if there are many systems of equations, all using the same coefficient matrix; the matrix may be factored once and the systems may be solved from the factors just as efficiently—and more accurately—than by matrix multiplication by the inverse.

---

## Supplemental Reading

Dongarra, J.J., J.R. Bunch, C.B. Moler, and G.W. Stewart. *LINPACK Users' Guide*. Philadelphia, PA: SIAM Publications. 1979.

Forsythe, G., and C.B. Moler. *Computer Solution of Linear Algebraic Systems*. Englewood Cliffs, NJ: Prentice-Hall, Inc. 1967.

**NAME** MINV – Invert Matrix or Solve Linear Equations

### Purpose

Given a general dense  $n$ -by- $n$  matrix  $A$ , this subprogram evaluates the determinant of  $A$ , optionally solves one or more systems of linear equations  $Ax_j=b_j$ , and optionally computes  $A^{-1}$ .

Computational singularity of  $A$  results in  $\det(A)=0$ . The partial product of pivot elements is computed as  $A$  is factored, and  $A$  is declared singular if the absolute value of the partial product ever fails to exceed a user-supplied tolerance. If this condition is detected during factorization, the computation is terminated and the “small” partial determinant is returned to indicate its occurrence. A more common situation, however, is that  $A$  is not numerically singular but happens to be ill-conditioned. When a matrix is ill-conditioned, small errors in the matrix and right-hand side and small roundoff errors introduced during the solution process itself are magnified greatly in the solution.

Although singular matrices are characterized by having zero determinants, a small nonzero determinant is unrelated to computational singularity or ill-conditioning. Therefore, the stopping criteria is artificial, and this subprogram may give a completely unreliable indication of the singularity of  $A$ . The SCILIB subprograms SGECO, SGEDI, and SGESL may be combined to perform the same functions as MINV, while providing a more reliable indication of singularity.

### Usage

SCILIB:

```

INTEGER*8      n, ldab, m, job
REAL*8        ab(ldab, n+m), work(2*n), det, tol
CALL MINV(ab, n, ldab, work, det, tol, m, job)

```

### Input

**ab** Array containing the  $n$ -by- $n$  matrix  $A$  in columns 1 through  $n$ , and  $m \geq 0$  right-hand side vectors  $b_j$  in columns  $n+1$  through  $n+m$ .

**n** The order of matrix  $A$ ,  $n > 0$ .

**ldab** The leading dimension of array **ab** as declared in the calling program unit, with  $ldab \geq n$ .

**tol** Lower limit for the partial product of pivot elements,  $tol \geq 0$ .  $A$  is considered singular if the magnitude of the partial product of pivot elements ever fails to exceed **tol**.

## MINV – Invert Matrix or Solve Linear Equations

<b>m</b>	Number of right-hand side vectors, $m \geq 0$ . If $m = 0$ , no right-hand sides are solved.				
<b>job</b>	Option flag: <table style="margin-left: 40px;"> <tr> <td><b>job</b> = 0</td> <td>do not compute <math>A^{-1}</math></td> </tr> <tr> <td><b>job</b> <math>\neq</math> 0</td> <td>compute <math>A^{-1}</math></td> </tr> </table>	<b>job</b> = 0	do not compute $A^{-1}$	<b>job</b> $\neq$ 0	compute $A^{-1}$
<b>job</b> = 0	do not compute $A^{-1}$				
<b>job</b> $\neq$ 0	compute $A^{-1}$				

**Working Storage**

<b>work</b>	An array of size $2n$ , used for work space.
-------------	--

**Output**

<b>ab</b>	If $ \det  > \text{tol}$ on return and $\text{job} \neq 0$ , then $A^{-1}$ overwrites $A$ in columns 1 to $n$ of <b>ab</b> . If $ \det  \leq \text{tol}$ on return or $\text{job} = 0$ , then columns 1 to $n$ of <b>ab</b> may have been destroyed.  If $ \det  > \text{tol}$ on return and $m \neq 0$ , then solution vectors $x_j$ of the systems of linear equations $Ax_j = b_j$ overwrite the right-hand side vectors in columns $n+1$ to $n+m$ of <b>ab</b> . If $ \det  \leq \text{tol}$ on return or $m = 0$ , then no solution vectors have been computed.
<b>det</b>	The determinant of $A$ if $ \det  > \text{tol}$ . Otherwise, <b>det</b> is the last partial product computed before the matrix was declared singular and the computation was terminated. Caution: an overflow may be produced in the computation of <b>det</b> .

**Notes**

Cray Research, Inc. has declared this subprogram obsolete in release 6.0 of the UNICOS Math and Scientific Library.

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

$n < 0$ ,  
 $m < 0$ ,  
 $\text{ldab} < n$ , and  
 $\text{tol} < 0$ .

**It is almost never necessary to compute either the determinant or the inverse of a matrix.** While papers and reference books extensively use the notation “ $\det(A) \neq 0$ ” to mean “ $A$  is nonsingular,” SCILIB includes both more efficient

## Linear Equations

### MINV – Invert Matrix or Solve Linear Equations

and more reliable subprograms for detecting singularity. Similarly, references frequently use “ $A^{-1}b$ ” to mean “the solution  $x$  of the system of linear equations  $Ax = b$ .” Again, it is more efficient and accurate to compute the solution directly than to invert the coefficient matrix and multiply the inverse times the right-hand side vector. This is true even if there are many systems of equations, all using the same coefficient matrix; the matrix may be factored once, and the systems may be solved from the factors as efficiently, but more accurately, as by matrix multiplication by the inverse.

### Example

Invert the 6-by-6 REAL\*8 matrix  $A$  stored in array  $AB$  whose dimensions are 10 by 50 and solve 25 systems of linear equations whose right-hand sides  $b_j$  are stored in columns 7 to 31 of  $AB$ . Consider  $A$  to be singular if  $|\det(A)| < 10^{-10}$ .

```
INTEGER*8 N, LDAB, M, JOB
REAL*8    AB(10,50), WORK(20), DET, TOL
N = 6
LDAB = 10
M = 25
TOL = 1.0E-10
JOB = 1
CALL MINV (AB,N,LDAB,WORK,DET,TOL,M,JOB)
IF ( ABS(DET) .LE. TOL ) THEN
    handle singular matrix
END IF
```

**NAME** OPFILT – Solve Symmetric Toeplitz Linear Equations

### Purpose

Given a real symmetric  $n$ -by- $n$  Toeplitz matrix  $A$  and a right-hand side vector  $b$ , this subprogram solves the system of linear equations  $Ax=b$  by means of the Weiner-Levinson algorithm. A matrix  $A$  is a symmetric Toeplitz matrix if its elements  $a_{ij}$  are given by  $a_{ij} = q_{|i-j|+1}$ . The following illustrates a 3-by-3 symmetric Toeplitz matrix.

$$\begin{array}{ccc} q_1 & q_2 & q_3 \\ q_2 & q_1 & q_2 \\ q_3 & q_2 & q_1 \end{array}$$

OPFILT is designed for use only on matrices which do not require pivoting to maintain numerical stability.

### Usage

SCILIB:

```
INTEGER*8      n
REAL*8        x(n), b(n), work(2*n), q(n)
CALL OPFILT(n, x, b, work, q)
```

### Input

**n**            The order of matrix  $A$ ,  $n \geq 0$ .  
**b**            The right-hand side vector  $b$ .  
**q**            The vector that generates the  $A$  matrix via the relationship  
 $a_{ij} = q_{|i-j|+1}$ .

### Working Storage

**work**        An array of size  $2n$ , used for work space.

### Output

**x**            The solution vector  $x$ .

### Notes

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

Linear Equations

**OPFILT – Solve Symmetric Toeplitz Linear Equations**

$n < 0$ .

An overflow or divide by zero may be produced if the matrix is not suitable for solution with the Weiner-Levinson algorithm.

### Example

Solve the 6-by-6 REAL\*8 symmetric Toeplitz matrix system  $Ax = b$  where  $A$  is represented by array  $Q$  of dimension 10, and  $x$  and  $b$  are stored in arrays  $X$  and  $B$ , also of dimension 10 by 10, respectively.

```
INTEGER*8 N
REAL*8    X(10),B(10),WORK(20),Q(10)
N = 6
CALL OPFILT (N,X,B,WORK,Q)
```

**NAME** SGBCO/CGBCO – Estimate Condition

**Purpose**

These subprograms compute the triangular factorization and estimate the condition number of a general nonsymmetric  $n$ -by- $n$  band matrix  $A$  stored in a two-dimensional array. A band matrix is a matrix whose nonzero elements all are near the principal diagonal. Specifically,  $a_{ij} = 0$  if  $i-j > kl$  or  $j-i > ku$  for some integers  $kl$  and  $ku$ . The smallest such  $kl$  and  $ku$  for a given matrix are called the lower and upper bandwidths, respectively, and  $k = kl+ku+1$  is the total bandwidth. The subprograms for band matrices use less storage than the subprograms for full matrices if  $2kl+ku < n$ .

Tridiagonal matrices are the special case  $kl = ku = 1$ . They can be handled more efficiently by the subprograms SGTSL and CGTSL. SCILIB also contains subprograms designed to handle positive definite band matrices. These subprograms are documented elsewhere in this chapter.

Specifically, given  $A$ , these subprograms determine an upper-triangular band matrix  $U$ , and a matrix  $L$  that is the product of elementary lower-triangular band matrices and permutation matrices such that

$$A = LU$$

and compute an estimate of  $\kappa(A)$ , the condition number of  $A$ . Refer to “Condition Number” in the introduction to this chapter for a discussion of  $\kappa(A)$ . When a matrix is ill-conditioned,  $\kappa(A)$  is large, so small errors in the matrix and right-hand side and small roundoff errors introduced during the solution process itself are magnified greatly in the solution.

Since  $1 < \kappa(A) \leq \infty$ , these subprograms actually compute the reciprocal condition number,  $1/\kappa(A)$ . The reciprocal condition number has the interpretation that if  $1/\kappa(A)$  approximately equals  $10^{-d}$ , elements of  $x$  can be expected to have  $d$  fewer significant digits of accuracy than the elements of  $A$  or  $b$ . Consequently, if errors in the coefficient matrix and right-hand side exceed  $1/\kappa(A)$ , or if  $1/\kappa(A)$  is negligible compared to 1.0, then  $x$  may have no significant digits at all.

A set of companion subprograms computes the triangular factorization of a general band matrix without estimating its condition number. These companion subprograms are faster but provide a less reliable indication of singularity.

The triangular factors may be used to solve a system of linear equations,  $Ax = b$ , by successively solving  $L(Ux) = b$ . The determinant of  $A$  can be computed as  $\det(A) = \det(L)\det(U)$ . These operations are performed by a set of companion SCILIB subprograms whose names depend on the data type:

Linear Equations

**SGBCO/CGBCO – Estimate Condition**

Data Type	Estimate Condition	Factor	Solve	Determinant
REAL*8	SGBCO	SGBFA	SGBSL	SGBDI
COMPLEX*16	CGBCO	CGBFA	CGBSL	CGBDI

The inverse of  $A$  will usually be a full  $n$ -by- $n$  matrix that cannot be stored in the band storage of  $A$ . Therefore, no direct provision is made for computing  $A^{-1}$ . Calculations formulated in terms of matrix inverses are invariably more efficient when expressed in terms of the solution of sets of linear equations.

**Matrix Storage**

Because it is not necessary to store or operate on the zeros outside the band of  $A$ , you need only provide the elements within the band of  $A$ . Compared to storing the entire matrix, this can save memory if  $2kl+ku+1 < n$ .

The following example illustrates the storage of general band matrices. Consider the following matrix  $A$  of order  $n = 9$  and lower and upper bandwidths  $kl = 2$  and  $ku = 3$ , respectively:

11	12	13	14	0	0	0	0	0
21	22	23	24	25	0	0	0	0
31	32	33	34	35	36	0	0	0
0	42	43	44	45	46	47	0	0
0	0	53	54	55	56	57	58	0
0	0	0	64	65	66	67	68	69
0	0	0	0	75	76	77	78	79
0	0	0	0	0	86	87	88	89
0	0	0	0	0	0	97	98	99

When Gaussian elimination is performed on a general band matrix, pivoting introduces nonzero elements outside the band.  $L$  can be stored with a lower bandwidth of  $kl$ , but  $U$  requires an upper bandwidth of  $kl+ku$ . You must, therefore, provide storage for the extra  $kl$  diagonals. This is done by presenting the original matrix to the subprogram in an array large enough to satisfy the additional storage requirements. Thus, for the above matrix,  $A$  is given in an array  $ab$  with at least  $2kl+ku+1 = 8$  rows and  $n = 9$  columns as follows:

*	*	*	*	*	+	+	+	+
*	*	*	*	+	+	+	+	+
*	*	*	14	25	36	47	58	69
*	*	13	24	35	46	57	68	79
*	12	23	34	45	56	67	78	89
11	22	33	44	55	66	77	88	99
21	32	43	54	65	76	87	98	*
31	42	53	64	75	86	97	*	*

The asterisks in the  $(kl+ku)$ -by- $(kl+ku)$  triangle at the upper left corner and in the  $ku$ -by- $ku$  triangle at the lower right corner represent elements of **ab** that are not referenced, and the plus signs in the first  $kl$  rows indicate elements that may be filled in during the factorization. Thus, if  $a_{ij}$  is an element within the band of  $A$ , then it is stored in  $\mathbf{ab}(kl+ku+1+i-j,j)$ . Therefore, the columns of  $A$  are stored in the columns of **ab**, and the diagonals of  $A$  are stored in the rows of **ab**, such that the principal diagonal is stored in row  $kl+ku+1$  of **ab**.

## Usage

### SCILIB:

```

INTEGER*8          ldab, n, kl, ku, ipvt(n)
REAL*8            ab(ldab, n), rcond, work(n)
CALL SGBCO(ab, ldab, n, kl, ku, ipvt, rcond, work)

INTEGER*8          ldab, n, kl, ku, ipvt(n)
COMPLEX*16        ab(ldab, n), work(n)
REAL*8            rcond
CALL CGBCO(ab, ldab, n, kl, ku, ipvt, rcond, work)

```

## Input

**ab**            Array containing the **n**-by-**n** band matrix  $A$  in the compressed form described above. If  $a_{ij}$  is in the band, it is stored in  $\mathbf{ab}(kl+ku+1+i-j,j)$ . The columns of  $A$  are stored in the columns of **ab**, and the diagonals of  $A$  are stored in rows  $kl+1$  through  $2kl+ku+1$ . The first  $kl$  rows are used for work space and output.

**ldab**          The leading dimension of array **ab** as declared in the calling program unit, with  $ldab \geq 2kl+ku+1$ .

**n**             The order of matrix  $A$ ,  $n > 0$ .

**kl**            The lower bandwidth of  $A$ , i.e., the number of nonzero diagonals below the principal diagonal in the band,  $0 \leq kl < n$ .

Linear Equations  
SGBCO/CGBCO – Estimate Condition

**ku** The upper bandwidth of  $A$ , i.e., the number of nonzero diagonals above the principal diagonal in the band,  $0 \leq \text{ku} < \mathbf{n}$ . These subprograms are more efficient if  $\text{kl} \leq \text{ku}$ . This usually can be arranged since the factors used by these subprograms can be used to solve either  $Ax = b$  or  $A^*x = b$ .

### Working Storage

**work** An array of size  $\mathbf{n}$ , used for work space.

### Output

**ab** The triangular factors replace the input matrix. **ab** must be preserved between the condition number estimation call and any solve or determinant call.

**ipvt** The pivot information necessary to construct the permutations in the lower-triangular factor,  $L$ . **ipvt** must be preserved between the condition number estimation call and any solve or determinant call.

**rcond** An estimate of the reciprocal condition number,  $1/\kappa(A)$ . If **rcond** is small enough so that the logical expression

$$1.0 + \text{rcond} \text{ .EQ. } 1.0$$

is true, then  $A$  can be regarded as singular to working precision.

### Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

### Example

Factor and estimate the reciprocal condition number of the 9-by-9 REAL\*8 general band matrix  $A$  whose lower bandwidth is 2 and whose upper bandwidth is 3.  $A$  is stored as illustrated above in array **AB** whose dimensions are 8 by 10.

```
INTEGER*8 LDAB,N,KL,KU,IPVT(10)
REAL*8 AB(8,10),RCOND,WORK(10)
LDAB = 8
N = 9
KL = 2
KU = 3
CALL SGBCO (AB,LDAB,N,KL,KU,IPVT,RCOND,WORK)
IF ( 1.0 + RCOND .EQ. 1.0 ) THEN
    handle singular matrix
END IF
```

**NAME** SGBDI/CGBDI – Determinant

**Purpose**

Given the triangular factorization of a general  $n$ -by- $n$  band matrix  $A$ , these subprograms evaluate the determinant of  $A$ . No provision is made to compute  $A^{-1}$  since it will usually be a full  $n$ -by- $n$  matrix that cannot be stored in the band storage of  $A$ . Moreover, it is almost never necessary to compute the inverse of a matrix. Mathematical references frequently use " $A^{-1}b$ " to mean "the solution  $x$  of the system of linear equations  $Ax = b$ ." It is more efficient and accurate to compute the solution directly than to invert the coefficient matrix and multiply the inverse times the right-hand side vector. This is true even if there are many systems of equations, all using the same coefficient matrix; the matrix may be factored once and the systems may be solved from the factors just as efficiently, and more accurately, than by matrix multiplication by the inverse.

Specifically, given an  $n$ -by- $n$  upper-triangular band matrix  $U$ , and a matrix  $L$  which is the product of elementary lower-triangular band matrices and permutation matrices such that

$$A = LU,$$

the subprograms compute

$$\det(A) = \det(L)\det(U).$$

The triangular factorization of the coefficient matrix may be computed by either of two companion subprograms. One computes only the factorization, using an elementary test for singularity of the coefficient matrix; it is slightly faster. The other not only computes the factorization, but also estimates the condition number of the matrix. This takes a little more time, but is considerably more reliable. The names of the companion subprograms depend on the data type:

Data Type	Estimate Condition	Factor	Solve
REAL*8	SGBCO	SGBFA	SGBDI
COMPLEX*16	CGBCO	CGBFA	CGBDI

The companion subprograms are documented elsewhere in this chapter.

**Usage**

SCILIB:

INTEGER\*8            **ldab, n, kl, ku, ipvt(n)**  
 REAL\*8                **ab(ldab, n), det(2)**

Linear Equations  
SGBDI/CGBDI – Determinant

```
CALL SGBDI(ab, ldab, n, kl, ku, ipvt, det)
INTEGER*8      ldab, n, kl, ku, ipvt(n)
COMPLEX*16    ab(ldab, n), det(2)
CALL CGBDI(ab, ldab, n, kl, ku, ipvt, det)
```

### Input

**ab** Array containing the triangular factors of the  $n$ -by- $n$  general band matrix  $A$  as computed by the companion factorization or condition number estimation subprogram. **ab** must have been preserved between the factorization or condition number call and the determinant call.

**ldab** The leading dimension of array **ab** as declared in the calling program unit, with  $ldab \geq 2kl + ku + 1$ .

**n** The order of matrix  $A$ ,  $n \geq 0$ .

**kl** The lower bandwidth of  $A$ , i.e., the number of nonzero diagonals below the principal diagonal in the band,  $0 \leq kl < n$ .

**ku** The upper bandwidth of  $A$ , i.e., the number of nonzero diagonals above the principal diagonal in the band,  $0 \leq ku < n$ .

**ipvt** The pivot information necessary to construct the permutations in the lower-triangular factor,  $L$ . **ipvt** must have been preserved between the factorization or condition number call and the determinant call.

### Output

**det** The determinant of  $A$ , in the form  $\det(A) = \mathbf{det}(1) \times 10^{\mathbf{det}(2)}$ . This expression may underflow or overflow if evaluated; on the Hewlett-Packard supercomputer, underflows automatically flush to zero, but overflows normally terminate execution. For REAL\*8 and COMPLEX\*16, overflow cannot occur if  $\mathbf{det}(2) \leq 306$ . If evaluation is safe, an efficient way to do it is with the statement

$$\det(A) = \mathbf{det}(1) * 10.0 ** \mathbf{INT}(\mathbf{det}(2))$$

The value stored in  $\mathbf{det}(2)$  is an integer in REAL or COMPLEX form.  $\mathbf{det}(1)$  is normalized so that either  $\mathbf{det}(1) = 0$  or  $1 \leq |Re(\mathbf{det}(1))| + |Im(\mathbf{det}(1))| < 10$ , where  $Re(z)$  and  $Im(z)$  are the real and imaginary parts of  $z$ ;  $Re(z) = z$  and  $Im(z) = 0$  if  $z$  is real.

## Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

**It is almost never necessary to compute the determinant of a matrix** While it is true that papers and reference books make extensive use of the notation “ $\det(A) \neq 0$ ” to mean “A is nonsingular,” SCILIB includes more efficient and more reliable subprograms for detecting singularity.

## Example

Compute the determinant of a 9-by-9 REAL\*8 general band matrix A whose lower bandwidth is 2 and whose upper bandwidth is 3. A is stored in array AB whose dimensions are 8 by 10. The less reliable, but slightly faster, factorization subprogram is used to factor the coefficient matrix.

```

INTEGER*8 LDAB,N,KL,KU,IPVT(10),IER
REAL*8    AB(8,10),DET(2),DETA
LDAB = 8
N = 9
KL = 2
KU = 3
CALL SGBFA (AB,LDAB,N,KL,KU,IPVT,IER)
IF ( IER .EQ. 0 ) THEN
  CALL SGBDI (AB,LDAB,N,KL,KU,IPVT,DET)
  IF ( DET(1) .EQ. 0.0 ) THEN
    DETA = 0.0
  ELSE IF ( DET(2) .LE. 306 ) THEN
    DETA = DET(1) * 10.0 ** INT(DET(2))
  ELSE
    the determinant of A is too large to evaluate
    without overflow
  END IF
ELSE
  DETA = 0.0
END IF

```

Linear Equations  
**SGBFA/CGBFA – Band LU Factorization**

**NAME** SGBFA/CGBFA – Band LU Factorization

**Purpose**

These subprograms compute the triangular factorization of a general nonsymmetric  $n$ -by- $n$  band matrix  $A$  stored in a two-dimensional array. A band matrix is a matrix whose nonzero elements all are near the principal diagonal. Specifically,  $a_{ij} = 0$  if  $i-j > kl$  or  $j-i > ku$  for some integers  $kl$  and  $ku$ . The smallest such  $kl$  and  $ku$  for a given matrix are called the lower and upper bandwidths, respectively, and  $m = kl+ku+1$  is the total bandwidth. The subprograms for band matrices use less storage than the subprograms for full matrices if  $2kl+ku < n$ .

Tridiagonal matrices are the special case  $kl = ku = 1$ . They can be handled more efficiently by the subprograms SGTSL or CGTSL. SCILIB also contains subprograms designed to handle positive definite band matrices. These subprograms are documented elsewhere in this chapter.

Specifically, given  $A$ , these subprograms determine an upper-triangular band matrix  $U$ , and a matrix  $L$  that is the product of elementary lower triangular band matrices and permutation matrices such that

$$A = LU.$$

Computational singularity of  $A$  results in one or more zero diagonal elements of  $U$ . This condition is detected during factorization, and a status response is returned to indicate its occurrence. A more common situation, however, is that  $A$  is not numerically singular but happens to be ill-conditioned. When a matrix is ill-conditioned, small errors in the matrix and right-hand side and small roundoff errors introduced during the solution process itself are magnified greatly in the solution. A set of companion subprograms computes the triangular factorization of a general band matrix and also estimates its condition number. These companion subprograms provide a more reliable indication of singularity. The small amount of additional time they require is usually worthwhile, especially when developing a program or encountering stability or convergence problems.

The triangular factors may be used to solve a system of linear equations,  $Ax = b$ , by successively solving  $L(Ux) = b$ . The determinant of  $A$  can be computed as  $\det(A) = \det(L)\det(U)$ . These operations are performed by a set of companion SCILIB subprograms whose names depend on the data type:

Data Type	Factor	Estimate Condition	Solve	Determinant
REAL*8	SGBFA	SGBCO	SGBSL	SGBDI
COMPLEX*16	CGBFA	CGBCO	CGBSL	CGBDI

The companion subprograms are documented elsewhere in this chapter.

The inverse of  $A$  will usually be a full  $n$ -by- $n$  matrix that cannot be stored in the band storage of  $A$ . Therefore, no direct provision is made for computing  $A^{-1}$ . Calculations formulated in terms of matrix inverses are invariably more efficient when expressed in terms of the solution of sets of linear equations.

### Matrix Storage

Because it is not necessary to store or operate on the zeros outside the band of  $A$ , you need only provide the elements within the band of  $A$ . Compared to storing the entire matrix, this can save memory if  $2kl+ku+1 < n$ .

The following example illustrates the storage of general band matrices. Consider the following matrix  $A$  of order  $n = 9$  and lower and upper bandwidths  $kl = 2$  and  $ku = 3$ , respectively:

11	12	13	14	0	0	0	0	0
21	22	23	24	25	0	0	0	0
31	32	33	34	35	36	0	0	0
0	42	43	44	45	46	47	0	0
0	0	53	54	55	56	57	58	0
0	0	0	64	65	66	67	68	69
0	0	0	0	75	76	77	78	79
0	0	0	0	0	86	87	88	89
0	0	0	0	0	0	97	98	99

When Gaussian elimination is performed on a general band matrix, pivoting introduces nonzero elements outside the band.  $L$  can be stored with a lower bandwidth of  $kl$ , but  $U$  requires an upper bandwidth of  $kl+ku$ . You must, therefore, provide storage for the extra  $kl$  diagonals. This is done by presenting the original matrix to the subprogram in an array large enough to satisfy the additional storage requirements. Thus, for the above matrix,  $A$  is given in an array  $ab$  with at least  $2kl+ku+1 = 8$  rows and  $n = 9$  columns as follows:

*	*	*	*	*	+	+	+	+
*	*	*	*	+	+	+	+	+
*	*	*	14	25	36	47	58	69
*	*	13	24	35	46	57	68	79
*	12	23	34	45	56	67	78	89
11	22	33	44	55	66	77	88	99
21	32	43	54	65	76	87	98	*
31	42	53	64	75	86	97	*	*

The asterisks in the  $(kl+ku)$ -by- $(kl+ku)$  triangle at the upper left corner and in the  $ku$ -by- $ku$  triangle at the lower right corner represent elements of  $ab$  that

Linear Equations  
 SGBFA/CGBFA – Band LU Factorization

are not referenced, and the plus signs in the first  $kl$  rows indicate elements that may be filled in during the factorization. Thus, if  $a_{ij}$  is an element within the band of  $A$ , then it is stored in  $ab(kl+ku+1+i-j,j)$ . Therefore, the columns of  $A$  are stored in the columns of  $ab$ , and the diagonals of  $A$  are stored in the rows of  $ab$ , such that the principal diagonal is stored in row  $kl+ku+1$  of  $ab$ .

## Usage

SCILIB:

```

INTEGER*8      ldab, n, kl, ku, ipvt(n), ier
REAL*8        ab(ldab, n)
CALL SGBFA(ab, ldab, n, kl, ku, ipvt, ier)
INTEGER*8      ldab, n, kl, ku, ipvt(n), ier
COMPLEX*16    ab(ldab, n)
CALL CGBFA(ab, ldab, n, kl, ku, ipvt, ier)
  
```

## Input

**ab**            Array containing the  $n$ -by- $n$  band matrix  $A$  in the compressed form described above. If  $a_{ij}$  is in the band, it is stored in  $ab(kl+ku+1+i-j,j)$ . Columns of  $A$  are stored in the columns of  $ab$ , and the diagonals of  $A$  are stored in rows  $kl+1$  through  $2kl+ku+1$ . The first  $kl$  rows are used for work space and output.

**ldab**          The leading dimension of array  $ab$  as declared in the calling program unit, with  $ldab \geq 2kl+ku+1$ .

**n**             The order of matrix  $A$ ,  $n > 0$ .

**kl**            The lower bandwidth of  $A$ , i.e., the number of nonzero diagonals below the principal diagonal in the band,  $0 \leq kl < n$ .

**ku**            The upper bandwidth of  $A$ , i.e., the number of nonzero diagonals above the principal diagonal in the band,  $0 \leq ku < n$ . These subprograms are more efficient if  $kl \leq ku$ . This usually can be arranged because factors used by these subprograms can be used to solve either  $Ax = b$  or  $A^*x = b$ .

## Output

**ab**            The triangular factors replace the input matrix.  $ab$  must be preserved between the factorization call and any solve or determinant call.

**ipvt**          The pivot information necessary to construct the permutations in the lower triangular factor,  $L$ .  $ipvt$  must be

preserved between the factorization call and any solve or determinant call.

**ier**

Status response:

**ier** = 0 Normal return.

**ier** =  $k \neq 0$  if  $u_{kk}=0$ . ( $u_{kk}$  is the  $k$ -th element on the diagonal of upper triangular matrix  $U$ ). Technically, this is not an error condition for these subprograms, but it does indicate that  $A$  is computationally singular and that a division by zero will occur if the factorization is used to solve a system of linear equations.

## Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

## Example

Factor the 9-by-9 REAL\*8 general band matrix  $A$  whose lower bandwidth is 2 and whose upper bandwidth is 3.  $A$  is stored, as illustrated above, in array AB whose dimensions are 8 by 10.

```
INTEGER*8 LDAB,N,KL,KU,IPVT(10),IER
REAL*8    AB(8,10)
LDAB = 8
N = 9
KL = 2
KU = 3
CALL SGBFA (AB,LDAB,N,KL,KU,IPVT,IER)
IF ( IER .NE. 0 ) THEN
    handle singular matrix
END IF
```

Linear Equations  
**SGBSL/CGBSL – Solve Band Linear Equations**

**NAME** SGBSL/CGBSL – Solve Band Linear Equations

**Purpose**

Given the triangular factorization of a general  $n$ -by- $n$  band matrix  $A$ , and a right-hand side  $n$ -vector  $b$ , these subprograms solve the system of linear equations  $Ax = b$ . Optionally, these subprograms will solve the system  $A^*x = b$ , where  $A^*$  is the conjugate transpose of  $A$  (the conjugate transpose of a real matrix is simply the transpose). Specifically, given an  $n$ -by- $n$  upper-triangular band matrix  $U$ , and a matrix  $L$  that is the product of elementary lower-triangular band matrices and permutation matrices such that

$$A = LU,$$

and an  $n$ -vector  $b$ , to find  $x$  satisfying  $Ax = b$ , the subprograms successively solve

$$Lw = b$$

and

$$Ux = w,$$

while to solve  $A^*x = b$ , the subprograms successively solve

$$U^*v = b$$

and

$$L^*x = v.$$

Triangular factors of the coefficient matrix may be computed by either of two companion subprograms. One computes only the factorization, using an elementary test for singularity of the coefficient matrix; it is slightly the faster. The other not only computes the factorization, but also estimates the condition number of the matrix. This takes a little more time, but is considerably more reliable. Names of companion subprograms depend on the data type:

Data Type	Estimate Condition	Factor	Solve
REAL*8	SGBCO	SGBFA	SGBSL
COMPLEX*16	CGBCO	CGBFA	CGBSL

The companion subprograms are documented elsewhere in this chapter.

**Usage**

SCILIB:

INTEGER\*8            **ldab, n, kl, ku, ipvt(n), job**

```

REAL*8          ab(ldab, n), b(n)
CALL SGBSL(ab, ldab, n, kl, ku, ipvt, b, job)
INTEGER*8       ldab, n, kl, ku, ipvt(n), job
COMPLEX*16      ab(ldab, n), b(n)
CALL CGBSL(ab, ldab, n, kl, ku, ipvt, b, job)

```

## Input

<b>ab</b>	Array containing the triangular factors of the $n$ -by- $n$ general band matrix $A$ as computed by the companion factorization or condition number estimation subprogram. <b>ab</b> must have been preserved between the factorization or condition number call and the solve call.				
<b>ldab</b>	The leading dimension of array <b>ab</b> as declared in the calling program unit, with $ldab \geq 2kl+ku+1$ .				
<b>n</b>	The order of matrix $A$ , $n \geq 0$ .				
<b>kl</b>	The lower bandwidth of $A$ , i.e., the number of nonzero diagonals below the principal diagonal in the band, $0 \leq kl < n$ .				
<b>ku</b>	The upper bandwidth of $A$ , i.e., the number of nonzero diagonals above the principal diagonal in the band, $0 \leq ku < n$ .				
<b>ipvt</b>	Pivot information necessary to construct permutations in the lower-triangular factor, $L$ . <b>ipvt</b> must have been preserved between the factorization or condition number estimation call and the solve call.				
<b>b</b>	The right-hand side vector $b$ .				
<b>job</b>	Option flag: <table style="margin-left: 40px;"> <tr> <td><b>job</b> = 0</td> <td>solve <math>Ax = b</math></td> </tr> <tr> <td><b>job</b> <math>\neq</math> 0</td> <td>solve <math>A^*x = b</math></td> </tr> </table>	<b>job</b> = 0	solve $Ax = b$	<b>job</b> $\neq$ 0	solve $A^*x = b$
<b>job</b> = 0	solve $Ax = b$				
<b>job</b> $\neq$ 0	solve $A^*x = b$				

## Output

<b>b</b>	The solution vector $x$ overwrites the right-hand side vector $b$ .
----------	---

## Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

**Example**

Solve a system of linear equations  $Ax = b$ , where  $A$  is a 9-by-9 REAL\*8 general band matrix whose lower bandwidth is 2 and whose upper bandwidth is 3.  $A$  is stored in array AB whose dimensions are 8 by 10.  $b$  is a vector 9 elements long stored in an array B of dimension 10. The more robust, but slightly slower, condition number estimation subprogram is used to factor the coefficient matrix.

```
INTEGER*8 LDAB,N,KL,KU,IPVT(10),JOB
REAL*8 AB(8,10),B(10),RCOND,WORK(10)
LDAB = 8
N = 9
KL = 2
KU = 3
JOB = 0
CALL SGBCO (AB,LDAB,N,KL,KU,IPVT,RCOND,WORK)
IF ( 1.0 + RCOND .NE. 1.0 ) THEN
  CALL SGBSL (AB,LDAB,N,KL,KU,IPVT,B,JOB)
ELSE
  handle singular matrix
END IF
```

**NAME** SGEKO/CGECO – Estimate Condition

### Purpose

These subprograms compute the triangular factorization and estimate the condition number of a general dense  $n$ -by- $n$  matrix  $A$ . Specifically, given  $A$ , these subprograms determine an  $n$ -by- $n$  permutation matrix  $P$ , an  $n$ -by- $n$  unit lower-triangular matrix  $L$ , and an  $n$ -by- $n$  upper-triangular matrix  $U$ , such that

$$PA = LU$$

and compute an estimate of  $\kappa(A)$ , the condition number of  $A$ . Refer to “Condition Number” in the introduction to this chapter for a discussion of  $\kappa(A)$ . When a matrix is ill-conditioned,  $\kappa(A)$  is large, so small errors in the matrix and right-hand side and small roundoff errors introduced during the solution process itself are magnified greatly in the solution.

Since  $1 < \kappa(A) \leq \infty$ , these subprograms actually compute the reciprocal condition number,  $1/\kappa(A)$ . The reciprocal condition number has the interpretation that if  $1/\kappa(A)$  approximately equals  $10^{-d}$ , elements of  $x$  can be expected to have  $d$  fewer significant digits of accuracy than the elements of  $A$  or  $b$ . Consequently, if errors in the coefficient matrix and right-hand side exceed  $1/\kappa(A)$ , or if  $1/\kappa(A)$  is negligible compared to 1.0, then  $x$  may have no significant digits at all.

A set of companion subprograms computes the triangular factorization of a matrix without estimating its condition number. These companion subprograms are faster but provide a less reliable indication of singularity.

The triangular factors may be used to solve a system of linear equations,  $Ax = b$ , by successively solving  $L(Ux) = Pb$ . The determinant of  $A$  can be computed as  $\det(A) = \det(P) \times \det(L) \times \det(U)$ . The inverse of  $A$  may be formed as  $A^{-1} = U^{-1}L^{-1}P$ . These operations are performed by a set of companion SCILIB subprograms whose names depend on the data type:

Data Type	Estimate Condition	Factor	Solve	Determinant or inverse
REAL*8	SGECO	SGEFA	SGESL	SGEDI
COMPLEX*16	CGECO	CGEFA	CGESL	CGEDI

The companion subprograms are documented elsewhere in this chapter.

### Usage

SCILIB:

INTEGER\*8             $lda, n, ipvt(n)$

Linear Equations  
SGECO/CGECO – Estimate Condition

```
REAL*8          a(lda, n), rcond, work(n)
CALL SGECO(a, lda, n, ipvt, rcond, work)
INTEGER*8       lda, n, ipvt(n)
COMPLEX*16      a(lda, n), work(n)
REAL*8          rcond
CALL CGECO(a, lda, n, ipvt, rcond, work)
```

### Input

**a**                    Array containing the  $n$ -by- $n$  matrix  $A$ .

**lda**                  The leading dimension of array  $a$  as declared in the calling program unit, with  $lda \geq \max(n,1)$ .

**n**                     The order of matrix  $A$ ,  $n \geq 0$ .

### Working storage

**work**                 An array of size  $n$ , used for work space.

### Output

**a**                     The triangular factors replace the input matrix: the strict lower triangle of  $a$  contains the strict lower triangle of  $L$  and the upper triangle of  $a$  contains  $U$ .  $a$  must be preserved between the condition number estimation call and any solve, determinant, or inverse call.

**ipvt**                 The pivot information necessary to construct the permutation matrix  $P$ . **ipvt** must be preserved between the condition number estimation call and any solve, determinant, or inverse call.

**rcond**                An estimate of the reciprocal condition,  $1/\kappa(A)$ . If **rcond** is small enough so that the logical expression

$$1.0 + 1\text{rcond} \text{ .EQ. } 1.0$$

is true, then  $A$  can be regarded as singular to working precision. If **rcond** is zero, then the companion subprograms for solving and computing the inverse may divide by zero.

### Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

The triangular factors are stored in a different format from the format used by the standard LINPACK subprograms, but are compatible with the SCILIB factorization, solve, and determinant and inverse subprograms.

### Example

Factor the 6-by-6 REAL\*8 matrix *A* stored in array *A* whose dimensions are 10 by 10 and estimate its reciprocal condition number.

```
INTEGER*8 LDA,N,IPVT(10)
REAL*8    A(10,10),RCOND,WORK(10)
LDA = 10
N = 6
CALL SGECO (A,LDA,N,IPVT,RCOND,WORK)
IF ( 1.0 + RCOND .EQ. 1.0 ) THEN
    handle singular matrix
END IF
```

**NAME** SGEDI/CGEDI – Determinant and Inverse

**Purpose**

Given the triangular factorization of a general dense  $n$ -by- $n$  coefficient matrix  $A$ , these subprograms evaluate the determinant of  $A$  and/or compute  $A^{-1}$ . Specifically, given an  $n$ -by- $n$  permutation matrix  $P$ , an  $n$ -by- $n$  unit lower-triangular matrix  $L$ , and a nonsingular  $n$ -by- $n$  upper-triangular matrix  $U$ , such that

$$PA = LU,$$

the subprograms compute

$$\det(A) = \det(P) \times \det(L) \times \det(U)$$

and/or

$$A^{-1} = U^{-1}L^{-1}P.$$

The triangular factors of the coefficient matrix may be computed by either of two companion subprograms. One computes only the factorization, using an elementary test for singularity of the coefficient matrix; it is slightly faster. The other not only computes the factorization, but also estimates the condition number of the matrix. This process takes a little more time, but is considerably more reliable, especially when  $A^{-1}$  is desired. The names of the companion subprograms depend on the data type:

Data Type	EstimateCondit ion	Factor	Determinant or inverse
REAL*8	SGECO	SGEFA	SGEDI
COMPLEX*16	CGECO	CGEFA	CGEDI

The companion subprograms are documented elsewhere in this chapter.

**Usage**

SCILIB:

```

INTEGER*8      lda, n, ipvt(n), job
REAL*8         a(lda, n), det(2), work(n)
CALL SGEDI(a, lda, n, ipvt, det, work, job)

INTEGER*8      lda, n, ipvt(n), job
COMPLEX*16     a(lda, n), det(2), work(n)
CALL CGEDI(a, lda, n, ipvt, det, work, job)
    
```

## Input

- a** Array containing the triangular factors  $L$  and  $U$  of the  $n$ -by- $n$  coefficient matrix  $A$  as computed by the companion factorization or condition number estimation subprogram. **a** must have been preserved between the factorization or condition number call and the determinant or inverse call.
- lda** The leading dimension of array **a** as declared in the calling program unit, with  $lda \geq \max(n,1)$ .
- n** The order of matrix  $A$ ,  $n \geq 0$ .
- ipvt** The pivot information necessary to construct the permutation matrix  $P$  as computed by the companion factorization or condition number estimation subprogram. **ipvt** must have been preserved between the factorization or condition number call and the determinant or inverse call.
- job** Option flag:
- |              |    |                                     |
|--------------|----|-------------------------------------|
| <b>job</b> = | 1  | compute only $A^{-1}$               |
| <b>job</b> = | 10 | compute only $\det(A)$              |
| <b>job</b> = | 11 | compute both $A^{-1}$ and $\det(A)$ |

## Working storage

- work** An array of size  $n$ , used for work space if  $A^{-1}$  is requested.

## Output

- a** Unchanged if  $A^{-1}$  is not requested. Otherwise,  $A^{-1}$  overwrites the triangular factors of the coefficient matrix.
- det** Not referenced if the determinant is not requested. Otherwise, the determinant of  $A$ , in the form  $\det(A) = \mathbf{det}(1) \times 10^{\mathbf{det}(2)}$ . This expression may underflow or overflow if evaluated; on the Hewlett-Packard supercomputer, underflows automatically flush to zero, but overflows normally terminate execution. For REAL\*8 and COMPLEX\*16, overflow cannot occur if  $\mathbf{det}(2) \leq 306$ . If evaluation is safe, an efficient way to do it is with the statement

$$\mathit{det}(A) = \mathbf{det}(1) * 10.0 ** \mathbf{INT}(\mathbf{det}(2))$$

Refer to "Example 2."

The value stored in **det(2)** is an integer in REAL or COMPLEX form. **det(1)** is normalized so that either

$\det(1) = 0$  or  $1 \leq |Re(\det(1))| + |Im(\det(1))| < 10$ ,  
 where  $Re(z)$  and  $Im(z)$  are the real and imaginary parts of  $z$ ;  
 $Re(z) = z$  and  $Im(z) = 0$  if  $z$  is real.

## Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

**It is almost never necessary to compute either the determinant or the inverse of a matrix** While papers and reference books extensively use the notation “ $\det(A) \neq 0$ ” to mean “ $A$  is nonsingular,” SCILIB includes both more efficient and more reliable subprograms for detecting singularity. Similarly, references frequently use “ $A^{-1}b$ ” to mean “the solution  $x$  of the system of linear equations  $Ax = b$ .” Again, it is more efficient and accurate to compute the solution directly than to invert the coefficient matrix and multiply the inverse times the right-hand side vector. This is true even if there are many systems of equations, all using the same coefficient matrix; the matrix may be factored once and the systems may be solved from the factors just as efficiently, and more accurately, than by matrix multiplication by the inverse.

## Example 1

Compute only the inverse of a 6-by-6 REAL\*8 matrix  $A$  stored in array  $A$  whose dimensions are 10 by 10. The more robust, but slightly slower condition number estimation subprogram is used to factor the coefficient matrix.

```

INTEGER*8 LDA,N,IPVT(10),JOB
REAL*8 A(10,10),DET(2),RCOND,WORK(10)
LDA = 10
N = 6
JOB = 1
CALL SGECO (A,LDA,N,IPVT,RCOND,WORK)
IF ( 1.0 + RCOND .NE. 1.0 ) THEN
    CALL SGEDI (A,LDA,N,IPVT,DET,WORK,JOB)
ELSE
    handle singular matrix
END IF

```

If the coefficient matrix  $A$  is determined to be nonsingular,  $A^{-1}$  overwrites the coefficient matrix  $A$  in array  $A$ .

## Example 2

Compute only the determinant of a 6-by-6 REAL\*8 matrix  $A$  stored in array  $A$  whose dimensions are 10 by 10. The less reliable, but slightly faster factorization subprogram is used to factor the coefficient matrix.

Linear Equations  
**SGED/CGEDI – Determinant and Inverse**

```
INTEGER*8 LDA,N,IPVT(10),IER,JOB
REAL*8    A(10,10),DET(2),DETA,WORK(10)
LDA = 10
N = 6
JOB = 10
CALL SGEFA (A,LDA,N,IPVT,IER)
IF ( IER .EQ. 0 ) THEN
  CALL SGEDI (A,LDA,N,IPVT,DET,WORK,JOB)
  IF ( DET(1) .EQ. 0.0 ) THEN
    DETA = 0.0
  ELSE IF ( DET(2) .LE. 306 ) THEN
    DETA = DET(1) * 10.0 ** INT(DET(2))
  ELSE
    the determinant of A is too large to evaluate
    without overflow
  END IF
ELSE
  DETA = 0.0
END IF
```

Linear Equations  
**SGEFA/CGEFA – LU Factorization**

**NAME** SGEFA/CGEFA – LU Factorization

**Purpose**

These subprograms compute the triangular factorization of a general dense  $n$  by  $n$  matrix  $A$ . Specifically, given  $A$ , these subprograms determine an  $n$ -by- $n$  permutation matrix  $P$ , an  $n$ -by- $n$  unit lower-triangular matrix  $L$ , and an  $n$ -by- $n$  upper-triangular matrix  $U$ , such that

$$PA = LU.$$

Computational singularity of  $A$  results in one or more zero diagonal elements of  $U$ . This condition is detected during factorization, and a status response is returned to indicate its occurrence. A more common situation, however, is that  $A$  is not numerically singular but happens to be ill-conditioned. When a matrix is ill-conditioned, small errors in the matrix and right-hand side and small roundoff errors introduced during the solution process itself are magnified greatly in the solution. A set of companion subprograms computes the triangular factorization of a matrix and also estimates its condition number. These companion subprograms provide a more reliable indication of singularity. The small amount of additional time they require is usually worthwhile, especially when developing a program or encountering stability or convergence problems.

The triangular factors may be used to solve a system of linear equations,  $Ax = b$ , by successively solving  $L(Ux) = Pb$ . The determinant of  $A$  can be computed as  $\det(A) = \det(P) \times \det(L) \times \det(U)$ . The inverse of  $A$  may be formed as  $A^{-1} = U^{-1}L^{-1}P$ . These operations are performed by a set of companion SCILIB subprograms whose names depend on the data type:

Data Type	Factor	EstimateCondition	Solve	Determinant or inverse
REAL*8	SGEFA	SGECO	SGESL	SGEDI
COMPLEX*16	CGEFA	CGECO	CGESL	CGEDI

The companion subprograms are documented elsewhere in this chapter.

**Usage**

**SCILIB:**

```

INTEGER*8      lda, n, ipvt(n), ier
REAL*8        a(lda, n)
CALL SGEFA(a, lda, n, ipvt, ier)

INTEGER*8      lda, n, ipvt(n), ier
COMPLEX*16    a(lda, n)

```



## Linear Equations

### SGEFA/CGEFA – LU Factorization

```
INTEGER*8 LDA,N,IPVT(10),IER
REAL*8    A(10,10)
LDA = 10
N = 6
CALL SGEFA (A,LDA,N,IPVT,IER)
IF ( IER .NE. 0 ) THEN
    handle singular matrix
END IF
```

**NAME** SGESL/CGESL – Solve Linear Equations

**Purpose**

Given the triangular factorization of a general dense  $n$ -by- $n$  coefficient matrix  $A$ , and a right-hand side  $n$ -vector  $b$ , these subprograms solve the system of linear equations  $Ax = b$ . Optionally, these subprograms will solve the system  $A^*x = b$ , where  $A^*$  is the conjugate transpose of  $A$  (the conjugate transpose of a real matrix is simply the transpose). Specifically, given an  $n$ -by- $n$  permutation matrix  $P$ , an  $n$ -by- $n$  unit lower-triangular matrix  $L$ , and a nonsingular  $n$ -by- $n$  upper-triangular matrix  $U$ , such that

$$PA = LU,$$

and an  $n$ -vector  $b$ , to find  $x$  satisfying  $Ax = b$ , the subprograms compute

$$v = Pb,$$

then successively solve

$$Lw = v$$

and

$$Ux = w,$$

To solve  $A^*x = b$ , the subprograms successively solve

$$U^*v = b$$

and

$$L^*w = v,$$

and then compute

$$x = P^*w.$$

The triangular factors of the coefficient matrix may be computed by either of two companion subprograms. One computes only the factorization, using an elementary test for singularity of the coefficient matrix; it is slightly faster. The other not only computes the factorization, but also estimates the condition number of the matrix. This process takes a little more time, but is considerably more reliable. The names of the companion subprograms depend on the data type:

Data Type	EstimateCon dition	Factor	Solve
REAL*8	SGECO	SGEFA	SGESL
COMPLEX*16	CGECO	CGEFA	CGESL

The companion subprograms are documented elsewhere in this chapter.

Linear Equations  
SGESL/CGESL – Solve Linear Equations

## Usage

SCILIB:

```
INTEGER*8      lda, n, ipvt(n), job
REAL*8         a(lda, n), b(n)
CALL SGESL(a, lda, n, ipvt, b, job)

INTEGER*8      lda, n, ipvt(n), job
COMPLEX*16     a(lda, n), b(n)
CALL CGESL(a, lda, n, ipvt, b, job)
```

## Input

**a**            Array containing the triangular factors  $L$  and  $U$  of the  $n$ -by- $n$  coefficient matrix  $A$  as computed by the companion factorization or condition number estimation subprogram. **a** must have been preserved between the factorization or condition number call and the solve call.

**lda**           The leading dimension of array **a** as declared in the calling program unit, with  $lda \geq \max(n, 1)$ .

**n**            The order of matrix  $A$ ,  $n \geq 0$ .

**ipvt**        The pivot information necessary to construct the permutation matrix  $P$  as computed by the companion factorization or condition number estimation subprogram. **ipvt** must have been preserved between the factorization or condition number call and the solve call.

**b**            The right-hand side vector  $b$ .

**job**         Option flag:

**job** = 0   solve  $Ax = b$

**job**  $\neq$  0   solve  $A^*x = b$

## Output

**b**            The solution vector  $x$  overwrites the right-hand side vector  $b$ .

## Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

## Example 1

Solve a system of linear equations  $Ax = b$ , where  $A$  is a 6-by-6 REAL\*8 matrix stored in array **A** whose dimensions are 10 by 10, and  $b$  is a vector 6 elements long stored in an array **B** of dimension 10. The more robust, but slightly slower

condition number estimation subprogram is used to factor the coefficient matrix.

```

INTEGER*8 LDA,N,IPVT(10),JOB
REAL*8    A(10,10),B(10),RCOND,WORK(10)
LDA = 10
N = 6
JOB = 0
CALL SGECCO (A,LDA,N,IPVT,RCOND,WORK)
IF ( 1.0 + RCOND .NE. 1.0 ) THEN
    CALL SGESL (A,LDA,N,IPVT,B,JOB)
ELSE
    handle singular matrix
END IF

```

If the coefficient matrix  $A$  is determined to be nonsingular, the solution vector  $x$  overwrites the right-hand side  $b$  in array  $b$ .

## Example 2

Solve a system of linear equations  $A^T x = b$ , where  $A$  is a 6-by-6 REAL\*8 matrix stored in array  $A$  whose dimensions are 10 by 10, and  $b$  is a vector 6 elements long stored in an array  $B$  of dimension 10. The less reliable, but slightly faster, factorization subprogram is used to factor the coefficient matrix.

```

INTEGER*8 LDA,N,IPVT(10),IER,JOB
REAL*8    A(10,10),B(10)
LDA = 10
N = 6
JOB = 1
CALL SGEFA (A,LDA,N,IPVT,IER)
IF ( IER .EQ. 0 ) THEN
    CALL SGESL (A,LDA,N,IPVT,B,JOB)
ELSE
    handle singular matrix
END IF

```

If the coefficient matrix  $A$  is determined to be nonsingular, the solution vector  $x$  overwrites the right-hand side  $b$  in array  $b$ .

Linear Equations  
**SGTSL/CGTSL – Solve Tridiagonal Linear Equations**

**NAME** SGTSL/CGTSL – Solve Tridiagonal Linear Equations

**Purpose**

Given an  $n$ -by- $n$  tridiagonal matrix  $A$ , and a right-hand side  $n$ -vector  $b$ , these subprograms solve the system of linear equations  $Ax = b$ . A tridiagonal matrix  $A = \{a_{ij}\}$  is a matrix whose nonzero elements lie only on the principal diagonal ( $i = j$ ), the subdiagonal ( $i = j+1$ ), and the superdiagonal ( $i = j-1$ ) of the matrix.

**Matrix Storage**

The following example illustrates the storage of general tridiagonal matrices. Consider the following tridiagonal matrix of order  $n = 7$ :

11	12	0	0	0	0	0
21	22	23	0	0	0	0
0	32	33	34	0	0	0
0	0	43	44	45	0	0
0	0	0	54	55	56	0
0	0	0	0	65	66	67
0	0	0	0	0	76	77

The subdiagonal is stored in array **dl**, the principal diagonal is stored in array **d**, and the superdiagonal is stored in array **du**, as follows:

$i$	<b>dl</b> ( $i$ )	<b>d</b> ( $i$ )	<b>du</b> ( $i$ )
1	*	11	12
2	21	22	23
3	32	33	34
4	43	44	45
5	54	55	56
6	65	66	67
7	76	77	*

The asterisks represent elements whose initial contents are not used.

**Usage**

SCILIB:

```

INTEGER*8          n, ier
REAL*8            dl(n), d(n), du(n), b(n)
CALL SGTSL(n, dl, d, du, b, ier)
INTEGER*8          n, ier

```

## SGTSL/CGTSL – Solve Tridiagonal Linear Equations

```

COMPLEX*16      dl(n), d(n), du(n), b(n)
CALL CGTSL(n, dl, d, du, b, ier)

```

**Input**

**n** The order of matrix  $A$ ,  $n > 0$ .

**dl** Array containing the subdiagonal of the tridiagonal matrix,  $dl(i) = a_{i,i-1}$ ,  $i = 2, 3, \dots, n$ . On return, **dl** is destroyed, including **dl(1)**.

**d** Array containing the principal diagonal of the tridiagonal matrix,  $d(i) = a_{ii}$ ,  $i = 1, 2, \dots, n$ . On return, **d** is destroyed.

**du** Array containing the superdiagonal of the tridiagonal matrix,  $du(i) = a_{i,i+1}$ ,  $i = 1, 2, \dots, n-1$ . On return, **du** is destroyed, including **du(n)**.

**b** The right-hand side vector  $b$ .

**Output**

**b** The solution vector  $x$  overwrites the right-hand side vector  $b$  if **ier** = 0 is returned.

**ier** Status response:

**ier** = 0 Normal return.

**ier** =  $k \neq 0$  if the  $k$ -th element of the diagonal becomes zero.

**Notes**

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

**Example**

Solve a system of linear equations  $Ax = b$ , where  $A$  is a 7-by-7 REAL\*8 tridiagonal matrix. The subdiagonal of  $A$  is stored in array DL, the principal diagonal is stored in array D, and the superdiagonal is stored in array DU.  $b$  is a vector 7 elements long stored in an array B.

```

INTEGER*8 N, IER
REAL*8    DL(10), D(10), DU(10), B(10)
N = 7
CALL SGTSL (N, DL, D, DU, B, IER)
IF ( IER .NE. 0 ) THEN
    handle error condition
END IF

```

**NAME** SPBCO/CPBCO – Estimate Condition

**Purpose**

These subprograms compute the Cholesky factorization and estimate the condition number of an  $n$ -by- $n$  positive definite band matrix  $A$  stored in a two-dimensional array. A matrix  $A$  is positive definite if and only if it is Hermitian; that is,  $A$  is equal to  $A^*$ , its conjugate transpose, and the quadratic form  $x^*Ax$  is positive for all nonzero vectors  $x$ . (The conjugate transpose of a real matrix or vector is simply the transpose.)

A positive definite band matrix is a positive definite matrix whose nonzero elements all are fairly near the principal diagonal. Specifically,  $a_{ij} = 0$  if  $|i-j| > kd$  for some integer  $kd$ . The smallest such  $kd$  for a given matrix is called the half bandwidth, and  $2kd+1$  is called the total bandwidth.

Tridiagonal matrices are the special case  $kd = 1$ . They can be handled more efficiently by the subprograms SPTSL and CPTSL.

Specifically, given  $A$ , these subprograms determine an  $n$ -by- $n$  upper-triangular band matrix  $R$ , such that

$$A = R^*R$$

and compute an estimate of  $\kappa(A)$ , the condition number of  $A$ . Refer to "Condition Number" in the introduction to this chapter for a discussion of  $\kappa(A)$ . When a matrix is ill-conditioned,  $\kappa(A)$  is large, so small errors in the matrix and right-hand side and small roundoff errors introduced during the solution process itself are magnified greatly in the solution.

Since  $1 < \kappa(A) \leq \infty$ , these subprograms actually compute the reciprocal condition number,  $1/\kappa(A)$ . The reciprocal condition number has the interpretation that if  $1/\kappa(A)$  approximately equals  $10^{-d}$ , elements of  $x$  can be expected to have  $d$  fewer significant digits of accuracy than the elements of  $A$  or  $b$ . Consequently, if errors in the coefficient matrix and right-hand side exceed  $1/\kappa(A)$ , or if  $1/\kappa(A)$  is negligible compared to 1.0, then  $x$  may have no significant digits at all.

A set of companion subprograms computes the Cholesky factorization of a matrix without estimating its condition number. These companion subprograms are faster but provide a less reliable indication of singularity.

The triangular factors may be used to solve a system of linear equations,  $Ax = b$ , by successively solving  $R^*(Rx) = b$ . The determinant of  $A$  can be computed as  $\det(A) = \det(R)^2$ . These operations are performed by a set of companion SCILIB subprograms whose names depend on the data type:

Data Type	EstimateCondition	Factor	Solve	Determinant
REAL*8	SPBCO	SPBFA	SPBSL	SPBDI
COMPLEX*16	CPBCO	CPBFA	CPBSL	CPBDI

The inverse of  $A$  will usually be a full  $n$ -by- $n$  matrix, which cannot be stored in the band storage of  $A$ . Therefore, no direct provision is made for computing  $A^{-1}$ . Calculations formulated in terms of matrix inverses are invariably more efficient when expressed in terms of the solution of sets of linear equations.

### Matrix Storage

Because it is not necessary to store or operate on the zeros outside the band of  $A$ , and since the Cholesky factorization of  $A$  may be computed from either triangle of  $A$ , you need only provide the band within the upper triangle. Compared to storing the entire matrix, this can save memory in two ways: only the elements within the band are stored, and of them, only the upper triangle.

The following examples illustrate the storage of positive definite band matrices. Consider the following matrix  $A$  of order  $n = 7$  and half bandwidth  $kd = 2$ :

11	12	13	0	0	0	0
12	22	23	24	0	0	0
13	23	33	34	35	0	0
0	24	34	44	45	46	0
0	0	35	45	55	56	57
0	0	0	46	56	66	67
0	0	0	0	57	67	77

The upper triangle of  $A$  is stored in an array **ab** with at least  $kd+1 = 3$  rows and 7 columns as follows:

*	*	13	24	35	46	57
*	12	23	34	45	56	67
11	22	33	44	55	66	77

The asterisks represent elements in the  $kd$ -by- $kd$  triangle at the upper-left corner of **ab** that are not referenced. Thus, if  $a_{ij}$  is an element within the band of the upper triangle of  $A$ , it is stored in  $ab(kd+1+i-j, j)$ . Therefore, the columns of the upper triangle of  $A$  are stored in the columns of **ab**, and the diagonals of the upper triangle of  $A$  are stored in the rows of **ab**.

Linear Equations  
SPBCO/CPBCO – Estimate Condition

## Usage

SCILIB:

```
INTEGER*8      ldab, n, kd, ier
REAL*8         ab(ldab, n), rcond, work(n)
CALL SPBCO(ab, ldab, n, kd, rcond, work, ier)

INTEGER*8      ldab, n, kd, ier
COMPLEX*16     ab(ldab, n), work(n)
REAL*8         rcond
CALL CPBCO(ab, ldab, n, kd, rcond, work, ier)
```

## Input

**ab** Array containing the upper triangle of the  $n$ -by- $n$  positive definite band matrix  $A$  in the compressed form described above. If  $0 \leq j-i \leq kd$ , then  $a_{ij}$  is stored in **ab**( $kd+1+i-j, j$ ). Columns of the upper triangle of  $A$  are stored in the columns of **ab**, and diagonals of the upper triangle of  $A$  are stored in the rows of **ab**.

**ldab** The leading dimension of array **ab** as declared in the calling program unit, with  $ldab \geq kd+1$ .

**n** The order of matrix  $A$ ,  $n > 0$ .

**kd** The half bandwidth of  $A$ , i.e., the number of diagonals above the principal diagonal in the band,  $0 \leq kd < n$ .

## Working storage

**work** An array of size  $n$ , used for work space.

## Output

**ab** The Cholesky factor  $R$  replaces the input matrix. The factorization is not complete if **ier** is nonzero. **ab** must be preserved between the condition number estimation call and any solve or determinant call.

**rcond** An estimate of the reciprocal condition number,  $1/\kappa(A)$ , if **ier** is zero; unchanged from its input value if **ier** is nonzero. If **ier** is zero and **rcond** is so small that the logical expression

$$1.0 + \mathbf{rcond} \text{ .EQ. } 1.0$$

is true,  $A$  can be regarded as singular to working precision.

**ier** Status response:

<code>ier = 0</code>	Normal return—factorization complete.
<code>ier = k ≠ 0</code>	The leading submatrix of order $k$ is not computationally positive definite, possibly because of roundoff error.

## Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

## Example

Factor the 7-by-7 REAL\*8 positive definite band matrix  $A$  whose half bandwidth is 2 and whose upper triangle is stored in the upper triangle of array  $AB$  whose dimensions are 5 by 10, and estimate its reciprocal condition number.

```
INTEGER*8 LDAB,N,KD,IER
REAL*8 AB(5,10),RCOND,WORK(10)
LDAB = 5
N = 7
M = 2
CALL SPBCO (AB,LDAB,N,KD,RCOND,WORK,IER)
IF ( IER .NE. 0 ) THEN
    handle indefinite matrix
ELSE IF ( 1.0 + RCOND .EQ. 1.0 ) THEN
    handle singular matrix
END IF
```

Linear Equations  
**SPBDI/CPBDI – Determinant**

**NAME** SPBDI/CPBDI – Determinant

**Purpose**

Given the Cholesky factorization of an  $n$ -by- $n$  positive definite band matrix  $A$ , these subprograms evaluate the determinant of  $A$ . No provision is made to compute  $A^{-1}$  because it will usually be a full  $n$ -by- $n$  matrix, which cannot be stored in the band storage of  $A$ . Moreover, it is almost never necessary to compute the inverse of a matrix. Mathematical references frequently use “ $A^{-1}b$ ” to mean “the solution  $x$  of the system of linear equations  $Ax = b$ .” It is more efficient and accurate to compute the solution directly than to invert the coefficient matrix and multiply the inverse times the right-hand side vector. This is true even if there are many systems of equations, all using the same coefficient matrix; the matrix may be factored once and the systems may be solved from the factors just as efficiently, and more accurately, than by matrix multiplication by the inverse.

Specifically, given an  $n$ -by- $n$  upper-triangular band matrix  $R$ , such that

$$A = R^*R,$$

where  $R^*$  is the conjugate transpose of  $R$ , the subprograms compute

$$\det(A) = \det(R)^2.$$

The Cholesky factorization of the coefficient matrix may be computed by either of two companion subprograms. One computes only the factorization, using an elementary test for singularity of the coefficient matrix; it is slightly faster. The other not only computes the factorization, but also estimates the condition number of the matrix. This process takes a little more time, but is considerably more reliable. The names of the companion subprograms depend on the data type:

Data Type	EstimateCondition	Factor	Determinant
REAL*8	SPBCO	SPBFA	SPBDI
COMPLEX*16	CPBCO	CPBFA	CPBDI

The companion subprograms are documented elsewhere in this chapter.

**Usage**

SCILIB:

```

INTEGER*8          ldab, n, kd
REAL*8             ab(ldab, n), det(2)
CALL SPBDI(ab, ldab, n, kd, det)

```

```

INTEGER*8      ldab, n, kd
COMPLEX*16     ab(ldab, n)
REAL*8         det(2)
CALL CPBDI(ab, ldab, n, kd, det)

```

## Input

**ab**            Array containing the Cholesky factor  $R$  of the  $n$ -by- $n$  positive definite band matrix  $A$  as computed by the companion factorization or condition number estimation subprogram. **ab** must have been preserved between the factorization or condition number call and the determinant call.

**ldab**          The leading dimension of array **ab** as declared in the calling program unit, with  $ldab \geq n+1$ .

**n**             The order of matrix  $A$ ,  $n \geq 0$ .

**kd**            The half bandwidth of  $A$ , i.e., the number of diagonals above the principal diagonal in the band,  $0 \leq kd < n$ .

## Output

**det**            The determinant of  $A$ , in the form  $det(A) = \mathbf{det}(1) \times 10^{\mathbf{det}(2)}$ . This expression may underflow or overflow if evaluated; on the Hewlett-Packard supercomputer, underflows automatically flush to zero, but overflows normally terminate execution. For REAL\*8 and COMPLEX\*16, overflow cannot occur if  $\mathbf{det}(2) \leq 306$ . If evaluation is safe, an efficient way to do it is with the statement

$$det(A) = \mathbf{det}(1) * 10.0 ** INT(\mathbf{det}(2))$$

The value stored in **det(2)** is an integer in REAL form. **det(1)** is normalized so that  $\mathbf{det}(1) = 0$  or  $1 \leq \mathbf{det}(1) < 10$ .

## Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

**It is almost never necessary to compute the determinant of a matrix** While it is true that papers and reference books make extensive use of the notation “ $det(A) \neq 0$ ” to mean “ $A$  is nonsingular,” SCILIB includes both more efficient and more reliable subprograms for detecting singularity.

Linear Equations  
SPBDI/CPBDI – Determinant

**Example**

Compute the determinant of a 7-by-7 REAL\*8 matrix *A* whose half bandwidth is 2 and whose upper triangle is stored in the upper triangle of array *AB* whose dimensions are 5 by 10. The less reliable, but slightly faster factorization subprogram is used to factor the coefficient matrix.

```
INTEGER*8 LDAB,N,KD,IER
REAL*8 AB(5,10),DET(2),DETA
LDAB = 5
N = 7
KD = 2
CALL SPBFA (AB,LDAB,N,KD,IER)
IF ( IER .EQ. 0 ) THEN
  CALL SPBDI (AB,LDAB,N,KD,DET)
  IF ( DET(1) .EQ. 0.0 ) THEN
    DETA = 0.0
  ELSE IF ( DET(2) .LE. 306 ) THEN
    DETA = DET(1) * 10.0 ** INT(DET(2))
  ELSE
    the determinant of A is too large to evaluate
    without overflow
  END IF
ELSE
  DETA = 0.0
END IF
```

**NAME** SPBFA/CPBFA – Cholesky Factorization

**Purpose**

These subprograms compute Cholesky factorization of an  $n$ -by- $n$  positive definite band matrix  $A$  stored in a two-dimensional array. A matrix  $A$  is positive definite if and only if it is Hermitian; that is,  $A$  is equal to  $A^*$ , its conjugate transpose, and the quadratic form  $x^*Ax$  is positive for all nonzero vectors  $x$ . (The conjugate transpose of a real matrix or vector is simply the transpose.)

A positive definite band matrix is a positive definite matrix whose nonzero elements all are fairly near the principal diagonal. Specifically,  $a_{ij} = 0$  if  $|i-j| > kd$  for some integer  $kd$ . The smallest such  $kd$  for a given matrix is called the half bandwidth, and  $2m+1$  is called the total bandwidth.

Tridiagonal matrices are the special case  $kd = 1$ . They can be handled more efficiently by the subprograms SPTSL and CPTSL.

Specifically, given  $A$ , these subprograms determine an  $n$ -by- $n$  upper-triangular band matrix  $R$ , such that

$$A = R^*R.$$

Computational singularity of  $A$  results in one or more zero diagonal elements of  $R$ , or, more frequently, in the loss of positive definiteness as evidenced by a negative diagonal element. This condition is detected during factorization, and a status response is returned to indicate its occurrence. A more common situation, however, is that  $A$  is not numerically singular but is ill-conditioned. When a matrix is ill-conditioned, small errors in the matrix and right-hand side and small roundoff errors introduced during the solution process itself are magnified greatly in the solution. A set of companion subprograms computes Cholesky factorization of a matrix and estimates its condition number. These companion subprograms provide a more reliable indication of singularity. The small amount of additional time they require is usually worthwhile, especially when developing a program or encountering stability or convergence problems.

The triangular factors may be used to solve a system of linear equations,  $Ax = b$ , by successively solving  $R^*(Rx) = b$ . The determinant of  $A$  can be computed as  $\det(A) = \det(R)^2$ . These operations are performed by a set of companion SCILIB subprograms whose names depend on the data type:

Data Type	Factor	EstimateCondition	Solve	Determinant
REAL*8	SPBFA	SPBCO	SPBSL	SPBDI
COMPLEX*16	CPBFA	CPBCO	CPBSL	CPBDI

Linear Equations  
SPBFA/CPBFA – Cholesky Factorization

The companion subprograms are documented elsewhere in this chapter.

The inverse of  $A$  will usually be a full  $n$ -by- $n$  matrix, which cannot be stored in the band storage of  $A$ . Therefore, no direct provision is made for computing  $A^{-1}$ . Calculations formulated in terms of matrix inverses are invariably more efficient when expressed in terms of the solution of sets of linear equations.

### Matrix Storage

Because it is not necessary to store or operate on the zeros outside the band of  $A$ , and since the Cholesky factorization of  $A$  may be computed from either triangle of  $A$ , you need only provide the band within the upper triangle. Compared to storing the entire matrix, this can save memory in two ways: only the elements within the band are stored, and of them, only the upper triangle.

The following examples illustrate the storage of positive definite band matrices. Consider the following matrix  $A$  of order  $n = 7$  and half bandwidth  $kd = 2$ :

11	12	13	0	0	0	0
12	22	23	24	0	0	0
13	23	33	34	35	0	0
0	24	34	44	45	46	0
0	0	35	45	55	56	57
0	0	0	46	56	66	67
0	0	0	0	57	67	77

The upper triangle of  $A$  is stored in an array **ab** with at least  $kd+1 = 3$  rows and 7 columns:

*	*	13	24	35	46	57
*	12	23	34	45	56	67
11	22	33	44	55	66	77

The asterisks represent elements in the  $kd$ -by- $kd$  triangle at the upper-left corner of **ab** that are not referenced. Thus, if  $a_{ij}$  is an element within the band of the upper triangle of  $A$ , it is stored in **ab**( $kd+1+i-j,j$ ). Therefore, the columns of the upper triangle of  $A$  are stored in the columns of **ab**, and the diagonals of the upper triangle of  $A$  are stored in the rows of **ab**.

### Usage

SCILIB:

```
INTEGER*8      ldab, n, kd, ier
REAL*8         ab(ldab, n)
CALL SPBFA(ab, ldab, n, kd, ier)
```

```

INTEGER*8      ldab, n, kd, ier
COMPLEX*16    ab(ldab, n)
CALL CPBFA(ab, ldab, n, kd, ier)

```

### Input

**ab**            Array containing the upper triangle of the **n**-by-**n** positive definite band matrix *A* in the compressed form described above. If  $0 \leq j-i \leq kd$ , then  $a_{ij}$  is stored in **ab(kd+1+i-j,j)**. The columns of the upper triangle of *A* are stored in the columns of **ab** and the diagonals of the upper triangle of *A* are stored in the rows of **ab**.

**ldab**          The leading dimension of array **ab** as declared in the calling program unit, with **ldab**  $\geq$  **kd+1**.

**n**             The order of matrix *A*, **n**  $>$  0.

**kd**            The half bandwidth of *A*, i.e., the number of diagonals above the principal diagonal in the band,  $0 \leq kd < n$ .

### Output

**ab**            The Cholesky factor *R* replaces the input matrix. The factorization is not complete if **ier** is nonzero. **ab** must be preserved between the condition number estimation call and any solve or determinant call.

**ier**           Status response:

<b>ier</b> =	0	Normal return—factorization complete.
<b>ier</b> =	<b>k</b> $\neq$ 0	The leading submatrix of order <i>k</i> is not computationally positive definite, possibly because of roundoff error.

### Notes

These subprograms are usage-compatible with the standard LINPACK subprograms with the same names.

### Example

Factor the 7-by-7 REAL\*8 positive definite band matrix *A* whose half bandwidth is 2 and whose upper triangle is stored in the upper triangle of array **AB** whose dimensions are 5 by 10, and estimate its reciprocal condition number.

## Linear Equations

### SPBFA/CPBFA – Cholesky Factorization

```
INTEGER*8 LDAB,N,KD,IER
REAL*8    AB(5,10)
LDAB = 5
N = 7
KD = 2
CALL SPBFA (AB,LDAB,N,KD,IER)
IF ( IER .NE. 0 ) THEN
    handle singular or indefinite matrix
END IF
```

**NAME** SPBSL/CPBSL – Solve Linear Equations

**Purpose**

Given the Cholesky factorization of an  $n$ -by- $n$  positive definite band matrix  $A$ , and a right-hand side  $n$ -vector  $b$ , these subprograms solve the system of linear equations  $Ax = b$ . Specifically, given an  $n$ -by- $n$  upper-triangular band matrix  $R$ , such that

$$A = R^*R,$$

where  $R^*$  is the conjugate transpose of  $R$ , and an  $n$ -vector  $b$ , to find  $x$  satisfying  $Ax = b$ , the subprograms successively solve

$$R^*w = b$$

and

$$Rx = w.$$

Cholesky factorization of the coefficient matrix may be computed by either of two companion subprograms. One computes only the factorization, using an elementary test for singularity of the coefficient matrix; it is slightly faster. The other not only computes factorization, but also estimates the condition number of the matrix. This process takes a little more time, but is considerably more reliable. The names of the companion subprograms depend on the data type:

Data Type	EstimateCon dition	Factor	Solve
REAL*8	SPBCO	SPBFA	SPBSL
COMPLEX*16	CPBCO	CPBFA	CPBSL

The companion subprograms are documented elsewhere in this chapter.

**Usage**

SCILIB:

```

INTEGER*8          ldab, n, kd
REAL*8            ab(ldab, n), b(n)
CALL SPBSL(ab, ldab, n, kd, b)

INTEGER*8          ldab, n, kd
COMPLEX*16        ab(ldab, n), b(n)
CALL CPBSL(ab, ldab, n, kd, b)

```

**Input**

**ab**                    Array containing the Cholesky factor  $R$  of the  $n$ -by- $n$  positive definite band matrix  $A$  as computed by the companion

	factorization or condition number estimation subprogram. <b>ab</b> must have been preserved between the factorization or condition number call and the solve call.
<b>ldab</b>	The leading dimension of array <b>ab</b> as declared in the calling program unit, with $ldab \geq kd+1$ .
<b>n</b>	The order of matrix $A$ , $n \geq 0$ .
<b>kd</b>	The half bandwidth of $A$ , i.e., the number of diagonals above the principal diagonal in the band, $0 \leq kd < n$ .
<b>b</b>	The right-hand side vector $b$ .

**Output**

<b>b</b>	The solution vector $x$ overwrites the right-hand side vector $b$ .
----------	---

**Notes**

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

**Example**

Solve a system of linear equations  $Ax = b$ , where  $A$  is a 7-by-7 REAL\*8 positive definite band matrix with half bandwidth 2. The upper triangle of  $A$  is stored in array AB whose dimensions are 5 by 10.  $b$  is a vector 7 elements long stored in an array B of dimension 10. The more robust, but slightly slower condition number estimation subprogram is used to factor the coefficient matrix.

```

INTEGER*8 LDAB,N,KD,IER
REAL*8 AB(5,10),B(10),RCOND,WORK(10)
LDAB = 5
N = 7
KD = 2
CALL SPBCO (AB,LDAB,N,KD,RCOND,WORK,IER)
IF ( IER .NE. 0 ) THEN
    handle indefinite matrix
ELSE IF ( 1.0 + RCOND .NE. 1.0 ) THEN
    CALL SPBSL (AB,LDAB,N,KD,B)
ELSE
    handle singular matrix
END IF

```

If the coefficient matrix  $A$  is determined to be positive definite and nonsingular, the solution vector  $x$  overwrites the right-hand side  $b$  in array **b**.

**NAME** SPOCO/CPOCO – Estimate Condition

**Purpose**

These subprograms compute Cholesky factorization and estimate the condition number of an  $n$ -by- $n$  positive definite matrix  $A$  stored in a two-dimensional array and estimate its condition number. A matrix  $A$  is positive definite if and only if it is Hermitian; that is,  $A$  is equal to  $A^*$ , its conjugate transpose, and the quadratic form  $x^*Ax$  is positive for all nonzero vectors  $x$ . (The conjugate transpose of a real matrix or vector is simply the transpose.)

Specifically, given  $A$ , these subprograms determine an  $n$ -by- $n$  upper-triangular matrix  $R$ , such that

$$A = R^*R$$

and compute an estimate of  $\kappa(A)$ , the condition number of  $A$ . Refer to “Condition Number” in the introduction to this chapter for a discussion of  $\kappa(A)$ . When a matrix is ill-conditioned,  $\kappa(A)$  is large, so small errors in the matrix and right-hand side and small roundoff errors introduced during the solution process itself are magnified greatly in the solution.

Since  $1 < \kappa(A) \leq \infty$ , these subprograms actually compute the reciprocal condition number,  $1/\kappa(A)$ . The reciprocal condition number has the interpretation that if  $1/\kappa(A)$  approximately equals  $10^{-d}$ , elements of  $x$  can be expected to have  $d$  fewer significant digits of accuracy than the elements of  $A$  or  $b$ . Consequently, if errors in the coefficient matrix and right-hand side exceed  $1/\kappa(A)$ , or if  $1/\kappa(A)$  is negligible compared to 1.0, then  $x$  may have no significant digits at all.

A set of companion subprograms computes Cholesky factorization of a matrix without estimating its condition number. These companion subprograms are faster but provide a less reliable indication of singularity.

The triangular factors may be used to solve a system of linear equations,  $Ax = b$ , by successively solving  $R^*(Rx) = b$ . The determinant of  $A$  can be computed as  $\det(A) = \det(R)^2$ . The inverse of  $A$  may be formed as

$A^{-1} = R^{-1}R^*$ , where  $R^*$  is the conjugate transpose of the inverse of  $R$ . These operations are performed by a set of companion SCILIB subprograms whose names depend on the data type:

Data Type	EstimateCon dition	Factor	Solve	Determinantor inverse
REAL*8	SPOCO	SPOFA	SPOSL	SPODI
COMPLEX*16	CPOCO	CPOFA	CPOSL	CPODI

The companion subprograms are documented elsewhere in this chapter.

## Matrix Storage

Because the Cholesky factorization of  $A$  may be computed from either triangle of  $A$ , you need only provide the upper triangle. Provide it in a two-dimensional array large enough to hold the entire array. The lower triangle of the array is not referenced.

## Usage

SCILIB:

```
INTEGER*8      lda, n, ier
REAL*8         a(lda, n), rcond, work(n)
CALL SPOCO(a, lda, n, rcond, work, ier)
INTEGER*8      lda, n, ier
COMPLEX*16     a(lda, n), work(n)
REAL*8         rcond
CALL CPOCO(a, lda, n, rcond, work, ier)
```

## Input

**a** Array containing the diagonal and upper triangle of the  $n$ -by- $n$  positive definite matrix  $A$ . The elements in the strict lower triangle of **a** are not referenced.

**lda** The leading dimension of array **a** as declared in the calling program unit, with  $lda \geq \max(n, 1)$ .

**n** The order of matrix  $A$ ,  $n \geq 0$ .

## Working storage

**work** An array of size **n**, used for work space.

## Output

**a** The Cholesky factor  $R$  replaces the input matrix  $A$  in the upper triangle of **a**. The strict lower triangle of **a** is unchanged. The factorization is not complete if **ier** is nonzero. **a** must be preserved between the condition number estimation call and any solve, determinant, or inverse call.

**rcond** An estimate of the reciprocal condition number,  $1/\kappa(A)$ , if **ier** is zero; unchanged from its input value if **ier** is nonzero. If **ier** is zero and **rcond** is so small that the logical expression

$$1.0 + \text{rcond} \text{ .EQ. } 1.0$$

is true, then  $A$  can be regarded as singular to working precision.

**ier**

Status response:

<b>ier</b> =	0	Normal return—factorization complete.
<b>ier</b> =	$k \neq 0$	The leading submatrix of order $k$ is not computationally positive definite, possibly because of roundoff error.

## Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

## Example

Factor the 6-by-6 REAL\*8 positive definite matrix  $A$  whose upper triangle is stored in the upper triangle of array  $A$  whose dimensions are 10 by 10, and estimate its reciprocal condition number.

```
INTEGER*8 LDA,N,IER
REAL*8    A(10,10),RCOND,WORK(10)
LDA = 10
N = 6
CALL SPOCO (A,LDA,N,RCOND,WORK,IER)
IF ( IER .NE. 0 ) THEN
    handle indefinite matrix
ELSE IF ( 1.0 + RCOND .EQ. 1.0 ) THEN
    handle singular matrix
END IF
```

Linear Equations  
**SPODI/CPODI – Determinant and Inverse**

**NAME** SPODI/CPODI – Determinant and Inverse

**Purpose**

Given the Cholesky factorization of an  $n$ -by- $n$  positive definite coefficient matrix  $A$ , these subprograms evaluate the determinant of  $A$  and/or compute  $A^{-1}$ . Specifically, given an  $n$ -by- $n$  upper-triangular matrix  $R$ , such that

$$A = R^*R,$$

where  $R^*$  is the conjugate transpose of  $R$ , the subprograms compute

$$\det(A) = \det(R)^2$$

and/or

$$A^{-1} = R^{-1}R^*$$

where  $R^*$  is the conjugate transpose of the inverse of  $R$ .

The Cholesky factorization of the coefficient matrix may be computed by either of two companion subprograms. One computes only the factorization, using an elementary test for singularity of the coefficient matrix; it is slightly faster. The other not only computes factorization, but also estimates the condition number of the matrix. This process takes a little more time, but is considerably more reliable, especially when  $A^{-1}$  is desired. The names of the companion subprograms depend on the data type:

Data Type	EstimateCondit ion	Factor	Determinantor inverse
REAL*8	SPOCO	SPOFA	SPODI
COMPLEX*16	CPOCO	CPOFA	CPODI

The companion subprograms are documented elsewhere in this chapter.

**Usage**

SCILIB:

```

INTEGER*8          lda, n, job
REAL*8            a(lda, n), det(2)
CALL SPODI(a, lda, n, det, job)

INTEGER*8          lda, n, job
COMPLEX*16        a(lda, n)
REAL*8            det(2)
CALL CPODI(a, lda, n, det, job)

```

## Input

- a** Array containing the Cholesky factor  $R$  of the  $n$ -by- $n$  positive definite coefficient matrix  $A$  in its upper triangle, as computed by the companion factorization or condition number estimation subprogram. The upper triangle of **a** must have been preserved between the factorization or condition number call and the determinant or inverse call.
- lda** The leading dimension of array **a** as declared in the calling program unit, with  $lda \geq \max(n,1)$ .
- n** The order of matrix  $A$ ,  $n \geq 0$ .
- job** Option flag:
- |              |    |                                     |
|--------------|----|-------------------------------------|
| <b>job</b> = | 1  | compute only $A^{-1}$               |
| <b>job</b> = | 10 | compute only $\det(A)$              |
| <b>job</b> = | 11 | compute both $A^{-1}$ and $\det(A)$ |

## Output

- a** Unchanged if  $A^{-1}$  is not requested. Otherwise, the upper triangle of  $A^{-1}$  overwrites the Cholesky factor of the coefficient matrix. The strict lower triangle of **a** is never changed.
- det** Not referenced if the determinant is not requested. Otherwise, the determinant of  $A$ , in the form  $\det(A) = \mathbf{det}(1) \times 10^{\mathbf{det}(2)}$ . This expression may underflow or overflow if evaluated; on the Hewlett-Packard supercomputer, underflows automatically flush to zero, but overflows normally terminate execution. For REAL\*8 and COMPLEX\*16, overflow cannot occur if  $\mathbf{det}(2) \leq 306$ . If evaluation is safe, an efficient way to do it is with the statement
- $$\mathbf{det}(A) = \mathbf{det}(1) * 10.0 ** \mathbf{INT}(\mathbf{det}(2))$$
- Refer to "Example 2."
- The value stored in **det(2)** is an integer in REAL form. **det(1)** is normalized so that  $\mathbf{det}(1) = 0$  or  $1 \leq \mathbf{det}(1) < 10$ .

## Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

It is almost never necessary to compute either the determinant or the inverse of a matrix. While papers and reference books extensively use the notation “ $\det(A) \neq 0$ ” to mean “ $A$  is nonsingular,” SCILIB includes more efficient and more reliable subprograms for detecting singularity. Similarly, references frequently use “ $A^{-1}b$ ” to mean “the solution  $x$  of the system of linear equations  $Ax = b$ .” Again, it is more efficient and accurate to compute the solution directly than to invert the coefficient matrix and multiply the inverse times the right-hand side vector. This is true even if there are many systems of equations, all using the same coefficient matrix; the matrix may be factored once and the systems may be solved from the factors just as efficiently, and more accurately, than by matrix multiplication by the inverse.

### Example 1

Compute only the inverse of a 6-by-6 REAL\*8 positive definite matrix  $A$  whose upper triangle is stored in the upper triangle of array  $A$  whose dimensions are 10 by 10. The more robust, but slightly slower condition number estimation subprogram is used to factor the coefficient matrix.

```

INTEGER*8 LDA,N,IER,JOB
REAL*8    A(10,10),DET(2),RCOND,WORK(10)
LDA = 10
N = 6
JOB = 1
CALL SPOCO (A,LDA,N,RCOND,WORK,IER)
IF ( IER .NE. 0 ) THEN
    handle indefinite matrix
ELSE IF ( 1.0 + RCOND .NE. 1.0 ) THEN
    CALL SPODI (A,LDA,N,DET,JOB)
ELSE
    handle singular matrix
END IF

```

If the coefficient matrix  $A$  is determined to be nonsingular,  $A^{-1}$  overwrites the coefficient matrix  $A$  in array  $a$ .

### Example 2

Compute only the determinant of a 6-by-6 REAL\*8 matrix  $A$  stored in array  $A$  whose dimensions are 10 by 10. The less reliable, but slightly faster factorization subprogram is used to factor the coefficient matrix.

Linear Equations  
SPODI/CPODI – Determinant and Inverse

```
INTEGER*8 LDA,N,IER,JOB
REAL*8    A(10,10),DET(2),DETA
LDA = 10
N = 6
JOB = 10
CALL SPOFA (A,LDA,N,IER)
IF ( IER .EQ. 0 ) THEN
  CALL SPODI (A,LDA,N,DET,JOB)
  IF ( DET(1) .EQ. 0.0 ) THEN
    DETA = 0.0
  ELSE IF ( DET(2) .LE. 306 ) THEN
    DETA = DET(1) * 10.0 ** INT(DET(2))
  ELSE
    the determinant of A is too large to evaluate
    without overflow
  END IF
ELSE
  DETA = 0.0
END IF
```

Linear Equations  
SPOFA/CPOFA – Cholesky Factorization

**NAME** SPOFA/CPOFA – Cholesky Factorization

**Purpose**

These subprograms compute the Cholesky factorization of an  $n$ -by- $n$  positive definite matrix  $A$  stored in a two-dimensional array. A matrix  $A$  is positive definite if and only if it is Hermitian; that is,  $A$  is equal to  $A^*$ , its conjugate transpose, and the quadratic form  $x^*Ax$  is positive for all nonzero vectors  $x$ . (The conjugate transpose of a real matrix or vector is simply the transpose.) Specifically, given  $A$ , these subprograms determine an  $n$ -by- $n$  upper-triangular matrix  $R$ , such that

$$A = R^*R.$$

Computational singularity of  $A$  results in one or more zero diagonal elements of  $R$ , or, more frequently, in the loss of positive definiteness as evidenced by a negative diagonal element. This condition is detected during the factorization, and a status response is returned to indicate its occurrence. A more common situation, however, is that  $A$  is not numerically singular but happens to be ill-conditioned. When a matrix is ill-conditioned, small errors in the matrix and right-hand side and small roundoff errors introduced during the solution process itself are magnified greatly in the solution. A set of companion subprograms computes the Cholesky factorization of a matrix and also estimates its condition number. These companion subprograms provide a more reliable indication of singularity. The small amount of additional time they require is usually worthwhile, especially when developing a program or encountering stability or convergence problems.

The triangular factors may be used to solve a system of linear equations,  $Ax = b$ , by successively solving  $R^*(Rx) = b$ . The determinant of  $A$  can be computed as  $\det(A) = \det(R)^2$ . The inverse of  $A$  may be formed as  $A^{-1} = R^{-1}R^*$ , where  $R^*$  is the conjugate transpose of the inverse of  $R$ . These operations are performed by a set of companion SCILIB subprograms whose names depend on the data type:

Data Type	Factor	EstimateCon dition	Solve	Determinantor inverse
REAL*8	SPOFA	SPOCO	SPOSL	SPODI
COMPLEX*16	CPOFA	CPOCO	CPOSL	CPODI

The companion subprograms are documented elsewhere in this chapter.

**Matrix Storage**

Because the Cholesky factorization of  $A$  may be computed from either triangle of  $A$ , you need only provide the upper triangle. Provide it in a two-dimensional

array large enough to hold the entire array. The lower triangle of the array is not referenced.

## Usage

SCILIB:

```

      INTEGER*8      lda, n, ier
      REAL*8        a(lda, n)
      CALL SPOFA(a, lda, n, ier)

      INTEGER*8      lda, n, ier
      COMPLEX*16    a(lda, n)
      CALL CPOFA(a, lda, n, ier)
  
```

## Input

**a**            Array containing the diagonal and upper triangle of the  $n$ -by- $n$  positive definite matrix  $A$ . The elements of the strict lower triangle are not referenced.

**lda**          The leading dimension of array  $a$  as declared in the calling program unit, with  $lda \geq \max(n, 1)$ .

**n**            The order of matrix  $A$ ,  $n \geq 0$ .

## Output

**a**            The Cholesky factor  $R$  replaces the input matrix  $A$  in the upper triangle of  $a$ . The strict lower triangle of  $a$  is unchanged. The factorization is not complete if  $ier$  is nonzero.  $a$  must be preserved between the factorization call and any solve, determinant, or inverse call.

**ier**          Status response:

<b>ier</b> = 0	Normal return—factorization complete.
<b>ier</b> = $k \neq 0$	The leading submatrix of order $k$ is not computationally positive definite, possibly because of roundoff error.

## Notes

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

## Example

Factor the 6-by-6 REAL\*8 positive definite matrix  $A$  whose upper triangle is stored in the upper triangle of array  $A$  whose dimensions are 10 by 10.

## Linear Equations

### SPOFA/CPOFA – Cholesky Factorization

```
INTEGER*8 LDA,N,IER
REAL*8    A(10,10)
LDA = 10
N = 6
CALL SPOFA (A,LDA,N,IER)
IF ( IER .NE. 0 ) THEN
    handle indefinite matrix
END IF
```

**NAME** SPOSL/CPOSL – Solve Linear Equations

**Purpose**

Given the Cholesky factorization of an  $n$ -by- $n$  positive definite coefficient matrix  $A$ , and a right-hand side  $n$ -vector  $b$ , these subprograms solve the system of linear equations  $Ax = b$ . Specifically, given an  $n$ -by- $n$  upper-triangular matrix  $R$ , such that

$$A = R^*R,$$

where  $R^*$  is the conjugate transpose of  $R$ , and an  $n$ -vector  $b$ , to find  $x$  satisfying  $Ax = b$ , the subprograms successively solve

$$R^*w = b$$

and

$$Rx = w.$$

The Cholesky factorization of the coefficient matrix may be computed by either of two companion subprograms. One computes only the factorization, using an elementary test for singularity of the coefficient matrix; it is slightly faster. The other not only computes the factorization, but also estimates the condition number of the matrix. This process takes a little more time, but is considerably more reliable. The names of the companion subprograms depend on the data type:

Data Type	Estimatecondition	Factor	Solve
REAL*8	SPOCO	SPOFA	SPOSL
COMPLEX*16	CPOCO	CPOFA	CPOSL

The companion subprograms are documented elsewhere in this chapter.

**Usage**

SCILIB:

```

INTEGER*8      lda, n
REAL*8        a(lda, n), b(n)
CALL SPOSL(a, lda, n, b)

INTEGER*8      lda, n
COMPLEX*16    a(lda, n), b(n)
CALL CPOSL(a, lda, n, b)

```

**Input**

- a** Array containing the Cholesky factor  $R$  of the  $n$ -by- $n$  positive definite coefficient matrix  $A$  in its upper triangle, as computed by the companion factorization or condition number estimation subprogram. The upper triangle of  $a$  must have been preserved between the factorization or condition number call and the solve call.
- lda** The leading dimension of array  $a$  as declared in the calling program unit, with  $lda \geq \max(n,1)$ .
- n** The order of matrix  $A$ ,  $n \geq 0$ .
- b** The right-hand side vector  $b$ .

**Output**

- b** The solution vector  $x$  overwrites the right-hand side vector  $b$ .

**Notes**

These subprograms are usage compatible with the standard LINPACK subprograms with the same names.

**Example**

Solve a system of linear equations  $Ax = b$ , where  $A$  is a 6-by-6 REAL\*8 positive definite matrix whose upper triangle is stored in array A whose dimensions are 10 by 10, and where  $b$  is a vector 6 elements long stored in an array B of dimension 10. The more robust, but slightly slower condition number estimation subprogram is used to factor the coefficient matrix.

```

INTEGER*8 LDA,N,IER
REAL*8    A(10,10),B(10),RCOND,WORK(10)
LDA = 10
N = 6
CALL SPOCO (A,LDA,N,RCOND,WORK,IER)
IF ( IER .NE. 0 ) THEN
    handle indefinite matrix
ELSE IF ( 1.0 + RCOND .NE. 1.0 ) THEN
    CALL SPOSL (A,LDA,N,B)
ELSE
    handle singular matrix
END IF

```

If the coefficient matrix  $A$  is determined to be positive definite and nonsingular, the solution vector  $x$  overwrites the right-hand side  $b$  in array  $b$ .

## SPTSL/CPTSL – Solve Positive Definite Tridiagonal Linear Equations

**NAME** SPTSL/CPTSL – Solve Positive Definite Tridiagonal Linear Equations

**Purpose**

Given an  $n$ -by- $n$  positive definite tridiagonal matrix  $A$ , and a right-hand side  $n$ -vector  $b$ , these subprograms solve the system of linear equations  $Ax = b$ . A matrix  $A$  is positive definite if and only if it is Hermitian; that is,  $A$  is equal to  $A^*$ , its conjugate transpose, and the quadratic form  $x^*Ax$  is positive for all nonzero vectors  $x$ . (The conjugate transpose of a real matrix or vector is simply the transpose.)

A positive definite tridiagonal matrix is a positive definite matrix  $A = \{a_{ij}\}$  whose nonzero elements lie only on the principal diagonal ( $i = j$ ), the subdiagonal ( $i = j+1$ ), and the superdiagonal ( $i = j-1$ ) of the matrix. Because of conjugate symmetry, the principal diagonal is always real, and the subdiagonal and superdiagonal are complex conjugates of each other. Thus, it is not necessary to store both the subdiagonal and the superdiagonal.

**Matrix Storage**

The following example illustrates the storage of a real symmetric or complex Hermitian tridiagonal matrix. Consider the following symmetric tridiagonal matrix of order  $n = 7$

11	12	0	0	0	0	0
12	22	23	0	0	0	0
0	23	33	34	0	0	0
0	0	34	44	45	0	0
0	0	0	45	55	56	0
0	0	0	0	56	66	67
0	0	0	0	0	67	77

then the principal diagonal is stored in array  $d$  and the superdiagonal is stored in array  $e$  as follows:

$i$	$d(i)$	$e(i)$
1	11	12
2	22	23
3	33	34
4	44	45
5	55	56
6	66	67
7	77	*

The asterisk represents an element that is not referenced.

## Usage

SCILIB:

```

INTEGER*8      n
REAL*8         d(n), e(n-1), b(n)
CALL SPTSL(n, d, e, b)

INTEGER*8      n
COMPLEX*16     d(n), e(n-1), b(n)
CALL CPTSL(n, d, e, b)

```

## Input

**n** The order of matrix  $A$ ,  $n > 0$ .

**d** Array containing the principal diagonal of the tridiagonal matrix,  $d(i) = a_{ii}$ ,  $i = 1, 2, \dots, n$ . For CPTSL and ZPTSL, only the real parts of **d** are used. On return, **d** is destroyed.

**e** Array containing the superdiagonal of the tridiagonal matrix,  $e(i) = a_{i,i+1}$ ,  $i = 1, 2, \dots, n-1$ .

**b** The right-hand side vector  $b$ .

## Output

**b** The solution vector  $x$  overwrites the right-hand side vector  $b$ .

## Notes

These subprograms are usage-compatible with the standard LINPACK subprograms with the same names.

Caution is necessary since these subprograms do not detect error conditions. An inaccurate solution may be computed or a division by zero may occur if the matrix is indefinite or singular.

## Example

Solve a system of linear equations  $Ax = b$ , where  $A$  is a 7-by-7 REAL\*8 positive definite tridiagonal matrix. The principal diagonal of  $A$  is stored in array D, and the superdiagonal is stored in array E.  $b$  is a vector 7 elements long stored in an array B.

```

INTEGER*8 N
REAL*8    D(10),E(10),B(10)
N = 7
CALL SPTSL (N,D,E,B)

```

## NAME LINPACK Subprograms not in this Guide –

Although SCILIB includes all LINPACK subprograms, the following nonoptimized subprograms are not documented in the *HP MLIB SCILIB User's Guide*. The *LINPACK Users' Guide*, included in the SCILIB documentation set, documents these subprograms.

## LINPACK Subprograms not in this Guide

Name	Function
SCHDC	Cholesky Decomposition of a Symmetric Matrix
CCHDC	Cholesky Decomposition of a Hermitian Matrix
SCHDD	Recompute the Cholesky Decomposition of a DOWNDATED Symmetric Matrix
CCHDD	Recompute the Cholesky Decomposition of a DOWNDATED Hermitian Matrix
SCHEX	Recompute the Cholesky Decomposition of a Permuted Symmetric Matrix
CCHEX	Recompute the Cholesky Decomposition of a Permuted Hermitian Matrix
SCHUD	Recompute the Cholesky Decomposition of a Updated Symmetric Matrix
CCHUD	Recompute the Cholesky Decomposition of a Updated Hermitian Matrix
CHICO	Factor a Hermitian Indefinite Matrix and Estimate its Condition Number
CHIDI	Determinant, Inverse, and Inertia of a Hermitian Indefinite Matrix
CHIFA	Factor a Hermitian Indefinite Matrix
CHISL	Solve Linear Equations with a Hermitian Indefinite Matrix
CHPCO	Factor a Hermitian Indefinite Packed Matrix and Estimate its Condition Number
CHPDI	Determinant, Inverse, and Inertia of a Hermitian Indefinite Packed Matrix
CHPFA	Factor a Hermitian Indefinite Packed Matrix
CHPSL	Solve Linear Equations with a Hermitian Indefinite Packed Matrix
SPPCO	Factor a Positive Definite Packed Matrix and Estimate its Condition Number
CPPCO	Factor a Positive Definite Packed Matrix and Estimate its Condition Number
SPPDI	Determinant and Inverse of a Positive Definite Packed Matrix
CPPDI	Determinant and Inverse of a Positive Definite Packed Matrix
SPPFA	Factor a Positive Definite Packed Matrix
CPPFA	Factor a Positive Definite Packed Matrix
SPPSL	Solve Linear Equations with a Positive Definite Packed Matrix
CPPSL	Solve Linear Equations with a Positive Definite Packed Matrix

Linear Equations

LINPACK Subprograms not in this Guide –

Name	Function
SQRDC	QR Decomposition of a General Rectangular Matrix
CQRDC	QR Decomposition of a General Rectangular Matrix
SQRSL	Solve Linear Equations using the QR Decomposition
CQRSL	Solve Linear Equations using the QR Decomposition
SSICO	Factor a Symmetric Indefinite Matrix and Estimate its Condition Number
CSICO	Factor a Symmetric Indefinite Matrix and Estimate its Condition Number
SSIDI	Determinant, Inverse, and Inertia of a Symmetric Indefinite Matrix
CSIDI	Determinant, Inverse, and Inertia of a Symmetric Indefinite Matrix
SSIFA	Factor a Symmetric Indefinite Matrix
CSIFA	Factor a Symmetric Indefinite Matrix
SSISL	Solve Linear Equations with a Symmetric Indefinite Matrix
CSISL	Solve Linear Equations with a Symmetric Indefinite Matrix
SSPCO	Factor a Symmetric Indefinite Packed Matrix and Estimate its Condition Number
CSPCO	Factor a Symmetric Indefinite Packed Matrix and Estimate its Condition Number
SSPDI	Determinant, Inverse, and Inertia of a Symmetric Indefinite Packed Matrix
CSPDI	Determinant, Inverse, and Inertia of a Symmetric Indefinite Packed Matrix
SSPFA	Factor a Symmetric Indefinite Packed Matrix
CSPFA	Factor a Symmetric Indefinite Packed Matrix
SSPSL	Solve Linear Equations with a Symmetric Indefinite Packed Matrix
CSPSL	Solve Linear Equations with a Symmetric Indefinite Packed Matrix
SSVDC	Singular Value Decomposition of a General Rectangular Matrix
CSVDC	Singular Value Decomposition of a General Rectangular Matrix
STRCO	Estimate the Condition Number of a Triangular Matrix
CTRCO	Estimate the Condition Number of a Triangular Matrix
STRDI	Determinant and Inverse of a Triangular Matrix
CTRDI	Determinant and Inverse of a Triangular Matrix
STRSL	Solve Linear Equations with a Triangular Matrix
CTRSL	Solve Linear Equations with a Triangular Matrix

## 5 Eigenvalues and Eigenvectors

---

### Overview

This chapter describes the EISPACK library included with SCILIB.

Some subprograms in this library have been upgraded by incorporating Level 2 and Level 3 BLAS and other algorithmic improvements.

Although all EISPACK subprograms are included in SCILIB, only upgraded ones are described in this chapter. Table 5-1 at the end of this chapter lists the subprograms that are included in SCILIB but not documented in the *HP MLIB SCILIB User's Guide*. You may find information for these subprograms in the *EISPACK Guide* and the *EISPACK Guide Extension*.

The LAPACK software library included with SCILIB is a comprehensive collection of eigenvalue and eigenvector solvers and subprograms for other linear algebra computations. This software is documented in the *HP MLIB LAPACK User's Guide*. We recommend that you use LAPACK subprograms rather than EISPACK subprograms in new programs. Future optimization efforts will be directed to LAPACK rather than EISPACK.

This chapter explains how to use SCILIB subprograms to compute eigenvalues or eigenvalues and eigenvectors of matrices. The operations covered are:

- dense Hermitian eigenproblems,  $Ax = \lambda x$ , with  $A = A^*$
- dense general eigenproblems,  $Ax = \lambda x$ , for arbitrary  $A$
- dense generalized eigenproblems,  $Ax = \lambda Bx$
- banded eigenproblems,  $Ax = \lambda x$

Refer to Chapter 7 for software to compute the eigenvalues or eigenvectors of a real, symmetric, sparse, ordinary or generalized eigenproblem.

---

### Chapter Objectives

After reading this chapter you will:

- know which version of EISPACK is included in the SCILIB library
- know how to use the described subprograms

---

## What You Need to Know to Use These Subprograms

EISPACK exists in single- and double-precision versions. Only the single-precision (64-bit) version is included in SCILIB.

---

## Supplemental Reading

Garbow, B.S., *et al* "Matrix Eigensystem Routines—EISPACK Guide Extension." *Lecture Notes in Computer Science*, Vol. 51. New York: Springer-Verlag. 1977.

Parlett, B.N. *The Symmetric Eigenproblem*. Englewood Cliffs, NJ: Prentice-Hall, Inc. 1980.

Smith, B.T., *et al* "Matrix Eigensystem Routines—EISPACK Guide." *Lecture Notes in Computer Science*, Vol. 6, 2nd edition. New York: Springer-Verlag. 1976.

Wilkinson, J.H. *The Algebraic Eigenproblem*. New York: Oxford University Press. 1965.

## RS – Eigenvalues and Eigenvectors of a Real Symmetric Matrix

**NAME** RS – Eigenvalues and Eigenvectors of a Real Symmetric Matrix

**Purpose**

This subprogram computes eigenvalues or eigenvalues and eigenvectors of a full real symmetric  $n$ -by- $n$  matrix  $A$ . Specifically, given  $A$ , this subprogram determines  $n$  scalars,  $\lambda_i, i = 1, 2, \dots, n$ , for which there exist corresponding nonzero vectors,  $x_i$ , such that

$$Ax_i = \lambda_i x_i.$$

Optionally, the  $x_i$  also may be computed.

**Matrix Storage**

Because the upper triangle of  $A$  may be obtained from the lower triangle, you need only provide the lower triangle of  $A$ , in a two-dimensional array large enough to hold the entire matrix. The upper triangle of the array is not referenced.

**Usage**

SCILIB:

```

INTEGER*8      ldax, n, job, ier
REAL*8         a(ldax, n), w(n), x(ldax, n), work1(n), work2(n)
CALL RS(ldax, n, a, w, job, x, work1, work2, ier)

```

**Input**

<b>ldax</b>	The leading dimension of arrays <b>a</b> and <b>x</b> as declared in the calling program unit, with $ldax \geq \max(n, 1)$ .				
<b>n</b>	The order of matrix $A$ , $n \geq 0$ .				
<b>a</b>	Array containing the diagonal and lower triangle of the $n$ -by- $n$ matrix $A$ . Elements in the strict upper triangle are not referenced.				
<b>job</b>	Option flag: <table> <tr> <td><b>job</b> = 0</td> <td>compute eigenvalues only</td> </tr> <tr> <td><b>job</b> <math>\neq</math> 0</td> <td>compute eigenvalues and eigenvectors</td> </tr> </table>	<b>job</b> = 0	compute eigenvalues only	<b>job</b> $\neq$ 0	compute eigenvalues and eigenvectors
<b>job</b> = 0	compute eigenvalues only				
<b>job</b> $\neq$ 0	compute eigenvalues and eigenvectors				

**Working storage**

<b>work1</b>	Array of size $n$ , used for work space.
<b>work2</b>	Array of size $n$ , used for work space.

**Output**

<b>a</b>	The lower triangle is destroyed if <b>job</b> = 0. Not modified if <b>job</b> ≠ 0.									
<b>w</b>	The eigenvalues $\lambda_i$ of <i>A</i> in ascending order if <b>ier</b> = 0 is returned.									
<b>x</b>	Not referenced if eigenvectors are not requested. In this case, <b>x</b> can be a dummy variable. Otherwise, eigenvectors of <i>A</i> if <b>ier</b> = 0 is returned. The <i>j</i> -th column of <b>x</b> contains the eigenvector $x_j$ of <i>A</i> corresponding to the eigenvalue in <b>w</b> ( <i>j</i> ), <i>j</i> = 1, 2, ..., <b>n</b> . Eigenvectors are normalized to have Euclidean length = 1.									
<b>ier</b>	Status response: <table style="margin-left: 40px;"> <tr> <td><b>ier</b> =</td> <td>0</td> <td>Normal return.</td> </tr> <tr> <td><b>ier</b> =</td> <td><math>k, 1 \leq k \leq n</math></td> <td>if calculation of the <i>k</i>-th eigenvalue failed to converge. <b>w</b>(1), <b>w</b>(2), ..., <b>w</b>(<i>k</i>-1) are eigenvalues, but are not necessarily the smallest and are not necessarily sorted. If eigenvectors are requested, the first <i>k</i>-1 columns of <b>x</b> are eigenvectors corresponding to the first <i>k</i>-1 elements of <b>w</b>.</td> </tr> <tr> <td><b>ier</b> =</td> <td>10n</td> <td>if <b>n</b> &gt; <b>ldax</b>. No eigenvalues or eigenvectors are returned.</td> </tr> </table>	<b>ier</b> =	0	Normal return.	<b>ier</b> =	$k, 1 \leq k \leq n$	if calculation of the <i>k</i> -th eigenvalue failed to converge. <b>w</b> (1), <b>w</b> (2), ..., <b>w</b> ( <i>k</i> -1) are eigenvalues, but are not necessarily the smallest and are not necessarily sorted. If eigenvectors are requested, the first <i>k</i> -1 columns of <b>x</b> are eigenvectors corresponding to the first <i>k</i> -1 elements of <b>w</b> .	<b>ier</b> =	10n	if <b>n</b> > <b>ldax</b> . No eigenvalues or eigenvectors are returned.
<b>ier</b> =	0	Normal return.								
<b>ier</b> =	$k, 1 \leq k \leq n$	if calculation of the <i>k</i> -th eigenvalue failed to converge. <b>w</b> (1), <b>w</b> (2), ..., <b>w</b> ( <i>k</i> -1) are eigenvalues, but are not necessarily the smallest and are not necessarily sorted. If eigenvectors are requested, the first <i>k</i> -1 columns of <b>x</b> are eigenvectors corresponding to the first <i>k</i> -1 elements of <b>w</b> .								
<b>ier</b> =	10n	if <b>n</b> > <b>ldax</b> . No eigenvalues or eigenvectors are returned.								

**Notes**

This subprogram is usage-compatible with the standard single-precision EISPACK subprogram with the same name. It calls EISPACK subprograms TRED1 and TQLRAT or TRED2 and TQL2, which are documented elsewhere in this chapter.

**Example 1**

Compute eigenvalues of a 6-by-6 REAL\*8 symmetric matrix *A* whose diagonal and lower triangle are stored in array *A* whose dimensions are 10 by 10. Eigenvalues are stored in array *W* of dimension 10.

```

INTEGER*8 LDA,N,JOB,IER
REAL*8    A(10,10),W(10),X,WORK1(10),WORK2(10)
LDA = 10
N = 6
JOB = 0
CALL RS (LDA,N,A,W,JOB,X,WORK1,WORK2,IER)
IF ( IER .NE. 0 ) THEN
    handle convergence failure
END IF

```

**Example 2**

Compute eigenvalues and eigenvectors of a 6-by-6 REAL\*8 symmetric matrix *A* whose diagonal and lower triangle are stored in array *A* whose dimensions are 10 by 10. Eigenvalues are stored in array *W* of dimension 10; eigenvectors are stored in the first six columns of array *X* of dimension 10 by 10.

```
INTEGER*8 LDAX,N,JOB,IER
REAL*8    A(10,10),W(10),X(10,10),WORK1(10),WORK2(10)
LDAX = 10
N = 6
JOB = 1
CALL RS (LDAX,N,A,W,JOB,X,WORK1,WORK2,IER)
IF ( IER .NE. 0 ) THEN
    handle convergence failure
END IF
```

**NAME** TQL2 – Eigenvalues and Eigenvectors of a Real Symmetric Matrix

**Purpose**

This subprogram computes eigenvalues and eigenvectors of a tridiagonal real symmetric  $n$ -by- $n$  matrix. Eigenvalues and eigenvectors of a full real symmetric matrix can also be computed by this subprogram if TRED2 has been used to reduce the full matrix to tridiagonal form.

Specifically, given a tridiagonal real symmetric matrix  $A$  or output of TRED2 applied to a full real symmetric matrix  $A$ , this subprogram determines scalars,  $\lambda_i, i = 1, 2, \dots, n$ , and nonzero vectors,  $x_i, i = 1, 2, \dots, n$ , such that

$$Ax_i = \lambda_i x_i.$$

**Matrix Storage**

The following example illustrates the storage of symmetric tridiagonal matrices. Consider the following symmetric tridiagonal matrix of order  $n = 7$ :

11	21	0	0	0	0	0
21	22	32	0	0	0	0
0	32	33	43	0	0	0
0	0	43	44	54	0	0
0	0	0	54	55	65	0
0	0	0	0	65	66	76
0	0	0	0	0	76	77

The subdiagonal is stored in array  $e$ , and the principal diagonal is stored in array  $d$ , as follows:

$i$	$e(i)$	$d(i)$
1	*	11
2	21	22
3	32	33
4	43	44
5	54	55
6	65	66
7	76	77

The asterisk represents an element whose initial contents are not used.

**Usage**

SCILIB:



## Eigenvalues and Eigenvectors

### TQL2 – Eigenvalues and Eigenvectors of a Real Symmetric Matrix

#### Example 1

Compute eigenvalues and eigenvectors of a 6-by-6 tridiagonal REAL\*8 symmetric matrix  $A$  whose diagonal and lower subdiagonal are stored in arrays  $D$  and  $E$  of dimension 10. Eigenvalues are returned in array  $D$ ; eigenvectors are placed in the first six columns of array  $X$  of dimension 10 by 10.

```
INTEGER*8 LDX,N,IER
REAL*8    D(10),E(10),X(10,10)
LDX = 10
N = 6
DO J = 1, N
  DO I = 1, N
    X(I,J) = 0.0
  END DO
  X(J,J) = 1.0
END DO
CALL TQL2 (LDX,N,D,E,X,IER)
IF ( IER .NE. 0 ) THEN
  handle convergence failure
END IF
```

#### Example 2

Compute eigenvalues and eigenvectors of a 6-by-6 REAL\*8 symmetric matrix  $A$  whose diagonal and lower triangle are stored in array  $A$  whose dimensions are 10 by 10. Eigenvalues are stored in array  $W$  of dimension 10; eigenvectors are stored in the first six columns of array  $X$  of dimension 10 by 10. (Compare with “Example 2” in the description of RS.)

```
INTEGER*8 LDAX,N,IER
REAL*8    A(10,10),W(10),X(10,10),WORK(10)
LDAX = 10
N = 6
CALL TRED2 (LDAX,N,A,W,WORK,X)
CALL TQL2 (LDAX,N,W,WORK,X,IER)
IF ( IER .NE. 0 ) THEN
  handle convergence failure
END IF
```

## TQLRAT – Eigenvalues of a Real Symmetric Matrix

**NAME** TQLRAT – Eigenvalues of a Real Symmetric Matrix

**Purpose**

This subprogram computes the eigenvalues of a tridiagonal real symmetric  $n$ -by- $n$  matrix. Eigenvalues of a full real symmetric matrix can also be computed by this subprogram if TRED1 has been used to reduce the full matrix to tridiagonal form.

Specifically, given a tridiagonal real symmetric matrix  $A$  or output of TRED1 applied to a full real symmetric matrix  $A$ , this subprogram determines scalars,  $\lambda_i, i = 1, 2, \dots, n$ , for which there exist corresponding nonzero vectors,  $x_i, i = 1, 2, \dots, n$ , such that

$$Ax_i = \lambda_i x_i.$$

**Matrix Storage**

The following example illustrates the storage of symmetric tridiagonal matrices. Consider the following symmetric tridiagonal matrix of order  $n = 7$ :

11	21	0	0	0	0	0
21	22	32	0	0	0	0
0	32	33	43	0	0	0
0	0	43	44	54	0	0
0	0	0	54	55	65	0
0	0	0	0	65	66	76
0	0	0	0	0	76	77

The squares of the subdiagonal elements are stored in array **e2**, and the principal diagonal is stored in array **d**, as follows:

$i$	<b>e2</b> ( $i$ )	<b>d</b> ( $i$ )
1	*	11
2	$21^2$	22
3	$32^2$	33
4	$43^2$	44
5	$54^2$	55
6	$65^2$	66
7	$76^2$	77

The asterisk represents an element whose initial contents are not used.



## Example 2

Compute eigenvalues of a 6-by-6 REAL\*8 symmetric matrix *A* whose diagonal and lower triangle are stored in array *A* whose dimensions are 10 by 10. Eigenvalues are stored in array *W* of dimension 10. (Compare with “Example 1” in the description of RS.)

```
INTEGER*8 LDA,N,IER
REAL*8    A(10,10),W(10),WORK1(10),WORK2(10)
LDA = 10
N = 6
CALL TRED1 (LDA,N,A,W,WORK1,WORK2)
CALL TQLRAT (N,W,WORK2,IER)
IF ( IER .NE. 0 ) THEN
    handle convergence failure
END IF
```

Eigenvalues and Eigenvectors  
TRED1 – Reduce Real Symmetric Matrix to Tridiagonal Form

**NAME** TRED1 – Reduce Real Symmetric Matrix to Tridiagonal Form

**Purpose**

This subprogram uses orthogonal-similarity transformations to reduce a full real symmetric  $n$ -by- $n$  matrix  $A$  to symmetric tridiagonal form without accumulating reduction transformations. The reduced form may be passed to subprogram TQLRAT, documented elsewhere in this chapter, to find the eigenvalues of  $A$ .

Specifically, given  $A$ , this subprogram determines an  $n$ -by- $n$  tridiagonal matrix  $T$  that is orthogonally similar to  $A$ , i.e., such that there exists an  $n$ -by- $n$  orthogonal matrix  $Q$  for which

$$Q^T A Q = T.$$

**Matrix Storage**

Because the upper triangle of  $A$  may be obtained from the lower triangle, you need only provide the lower triangle of  $A$ , in a two-dimensional array large enough to hold the entire matrix. The upper triangle of the array is not referenced.

**Usage**

SCILIB:

```
INTEGER*8      lda, n
REAL*8         a(lda, n), d(n), e(n), e2(n)
CALL TRED1(lda, n, a, d, e, e2)
```

**Input**

**lda** The leading dimension of array **a** as declared in the calling program unit, with  $lda \geq \max(n, 1)$ .

**n** The order of matrix  $A$ ,  $n \geq 0$ .

**a** Array containing the diagonal and lower triangle of the  $n$ -by- $n$  matrix  $A$ . Elements in the strict upper triangle are not referenced.

**Output**

**a** The diagonal and lower triangle are destroyed.

**d** Array containing diagonal elements of the tridiagonal matrix  $T$ .

**e** Array containing the subdiagonal elements of  $T$  in elements  $e(2)$  through  $e(n)$ .  $e(1) = 0$ .

## TRED1 – Reduce Real Symmetric Matrix to Tridiagonal Form

**e2**                    Array containing squares of subdiagonal elements of  $T$  in elements  $e(2)$  through  $e(n)$ .  $e2(1) = 0$ .

**Notes**

This subprogram is usage-compatible with the standard single-precision EISPACK subprogram with the same name.

Output arrays **e** and **e2** are redundant. Some EISPACK subprograms that can be used following TRED1 require **e** as input and some require **e2**.

**Example 1**

Reduce the 6-by-6 REAL\*8 symmetric matrix  $A$  whose diagonal and lower triangle are stored in array  $A$  whose dimensions are 10 by 10 to tridiagonal form.

```
INTEGER*8 LDA,N
REAL*8    A(10,10),D(10),E(10),E2(10)
LDA = 10
N = 6
CALL TRED1 (LDA,N,A,D,E,E2)
```

**Example 2**

Compute eigenvalues of a 6-by-6 REAL\*8 symmetric matrix  $A$  whose diagonal and lower triangle are stored in array  $A$  whose dimensions are 10 by 10. Eigenvalues will be stored in array  $W$  of dimension 10. (Compare with “Example 1” in the description of RS.)

```
INTEGER*8 LDA,N,IER
REAL*8    A(10,10),W(10),WORK1(10),WORK2(10)
LDA = 10
N = 6
CALL TRED1 (LDA,N,A,W,WORK1,WORK2)
CALL TQLRAT (N,W,WORK2,IER)
IF ( IER .NE. 0 ) THEN
    handle convergence failure
END IF
```

**NAME** TRED2 – Reduce Real Symmetric Matrix to Tridiagonal Form

### Purpose

This subprogram uses orthogonal similarity transformations to reduce a full real symmetric  $n$ -by- $n$  matrix  $A$  to symmetric tridiagonal form and accumulates reduction transformations. This reduced form and the transformation matrix may be passed to subprogram TQL2, documented elsewhere in this chapter, to find eigenvalues and eigenvectors of  $A$ .

Specifically, given  $A$ , this subprogram determines an  $n$ -by- $n$  orthogonal matrix  $X$  and an  $n$ -by- $n$  symmetric tridiagonal matrix  $T$  such that

$$X^T A X = T.$$

### Matrix Storage

Because the upper triangle of  $A$  may be obtained from the lower triangle, you need only provide the lower triangle of  $A$ , in a two-dimensional array large enough to hold the entire matrix. The upper triangle of the array is not referenced.

### Usage

SCILIB:

```
INTEGER*8      ldax, n
REAL*8         a(ldax, n), d(n), e(n), x(ldax, n)
CALL TRED2(ldax, n, a, d, e, x)
```

### Input

**ldax** The leading dimension of arrays **a** and **x** as declared in the calling program unit, with  $ldax \geq \max(n, 1)$ .

**n** The order of matrix  $A$ ,  $n \geq 0$ .

**a** Array containing the diagonal and lower triangle of the  $n$ -by- $n$  matrix  $A$ . Elements in the strict upper triangle are not referenced.

### Output

**d** Array containing diagonal elements of the tridiagonal matrix  $T$ .

**e** Array containing subdiagonal elements of  $T$  in elements  $e(2)$  through  $e(n)$ .  $e(1) = 0$ .

**x** The transformation matrix  $X$  that reduces  $A$  to tridiagonal form.

**TRED2 – Reduce Real Symmetric Matrix to Tridiagonal Form****Notes**

This subprogram is usage-compatible with the standard single-precision EISPACK subprogram with the same name.

**Example 1**

Reduce the 6-by-6 REAL\*8 symmetric matrix *A* whose diagonal and lower triangle are stored in array *A* whose dimensions are 10 by 10 to tridiagonal form and accumulate the transformation matrix.

```

INTEGER*8 LDAX,N
REAL*8    A(10,10),D(10),E(10),X(10,10)
LDAX = 10
N = 6
CALL TRED2 (LDAX,N,A,D,E,X)

```

**Example 2**

Compute eigenvalues and eigenvectors of a 6-by-6 REAL\*8 symmetric matrix *A* whose diagonal and lower triangle are stored in array *A* whose dimensions are 10 by 10. Eigenvalues will be stored in array *W* of dimension 10; eigenvectors will be stored in the first six columns of array *X* of dimension 10 by 10. (Compare with “Example 2” in the description of RS.)

```

INTEGER*8 LDAX,N,IER
REAL*8    A(10,10),W(10),X(10,10),WORK(10)
LDAX = 10
N = 6
CALL TRED2 (LDAX,N,A,W,WORK,X)
CALL TQL2 (LDAX,N,W,WORK,X,IER)
IF ( IER .NE. 0 ) THEN
    handle convergence failure
END IF

```

Eigenvalues and Eigenvectors  
**EISPACK Subprograms not in this Guide -**

**NAME** EISPACK Subprograms not in this Guide -

Although SCILIB includes all EISPACK subprograms, the following nonoptimized routines are not documented in the *HP MLIB SCILIB User's Guide*. The *EISPACK Guide* and the *EISPACK Guide Extension* document these subprograms.

EISPACK Subprograms not in this Guide

Name	Function
BAKVEC	Back Transform Eigenvectors following FIGI
BALANC	Balance a Real General Matrix
BALBAK	Back Transform Eigenvectors following BALANC
BANDR	Reduce a Real Symmetric Band Matrix to Real Symmetric Tridiagonal Form
BANDV	Determine Some Eigenvectors of a Real Symmetric Band Matrix
BISECT	Determine Some Eigenvectors of a Real Symmetric Tridiagonal Matrix
BQR	Determine Some Eigenvalues of a Real Symmetric Band Matrix
CBABK2	Back Transform Eigenvectors following CBAL
CBAL	Balance a Complex General Matrix
CG	Determine Eigenvalues/vectors of a Complex General Matrix
CH	Determine Eigenvalues/vectors of a Complex Hermitian Matrix
CINVIT	Determine Some Eigenvectors of a Complex Upper Hessenberg Matrix
COMBAK	Back Transform Eigenvectors following COMHES
COMHES	Reduce a Complex General Matrix to Complex Upper Hessenberg Form
COMLR	Determine the Eigenvalues of a Complex Upper Hessenberg Matrix
COMLR2	Determine the Eigenvalues/vectors of a Complex Hessenberg Matrix
COMQR	Determine the Eigenvalues of a Complex Upper Hessenberg Matrix
COMQR2	Determine the Eigenvalues/vectors of a Complex Upper Hessenberg Matrix
CORTB	Back Transform Eigenvectors following CORTH
CORTH	Reduce a Complex General Matrix to Complex Upper Hessenberg Form
ELMBAK	Back Transform Eigenvectors following ELMHES
ELMHES	Reduce a Real General Matrix to Real Upper Hessenberg Form
ELTRAN	Accumulate the Transformations in the Reduction by ELMHES
FIGI	Transform a Real Non-symmetric Tridiagonal Matrix to Real Symmetric Form
FIGI2	Transform a Real Non-symmetric Tridiagonal Matrix to Real Symmetric Form
HQR	Determine the Eigenvalues of a Real Upper Hessenberg Matrix

Eigenvalues and Eigenvectors  
EISPACK Subprograms not in this Guide –

Name	Function
HQR2	Determine the Eigenvalues/vectors of a Real Upper Hessenberg Matrix
HTRIB3	Back Transform Eigenvectors following HTRID3
HTRIBK	Back Transform Eigenvectors following HTRIDI
HTRID3	Reduce a Complex Hermitian Matrix to Real Symmetric Tridiagonal Form
HTRIDI	Reduce a Complex Hermitian Matrix to Real Symmetric Tridiagonal Form
IMTQL1	Determine the Eigenvalues of a Real Symmetric Tridiagonal Matrix
IMTQL2	Determine the Eigenvalues/vectors of a Real Symmetric Tridiagonal Matrix
IMTQLV	Determine the Eigenvalues of a Real Symmetric Tridiagonal Matrix
INVIT	Determine Some Eigenvectors of a Real Upper Hessenberg Matrix
MINFIT	Solve a Least Squares Problem with a Real Rectangular Coefficient Matrix
ORTBAK	Back Transform Eigenvectors following ORTHES
ORTHES	Reduce a Real General Matrix to Real Upper Hessenberg Form
ORTRAN	Accumulate the Transformations in the Reduction by ORTHES
QZHES	Partially Reduce a Real General Generalized Eigenproblem
QZIT	Complete the Reduction of a Real General Generalized Eigenproblem
QZVAL	Determine the Eigenvalues of a Reduced Real General Generalized Eigenproblem
QZVEC	Determine the Eigenvectors of a Reduced Real General Generalized Eigenproblem
RATQR	Determine Some Extreme Eigenvalues of a Real Symmetric Tridiagonal Matrix
REBAK	Back Transform Eigenvectors following REDUC or REDUC2
REBAKB	Back Transform Eigenvectors following REDUC2
REDUC	Reduce a Real Symmetric Generalized Eigenproblem to Standard Form
REDUC2	Reduce a Real Symmetric Generalized Eigenproblem to Standard Form
RG	Determine the Eigenvalues/vectors of a Real General Matrix
RGG	Determine the Eigenvalues/vectors of a Real General Generalized Eigenproblem
RSB	Determine the Eigenvalues/vectors of a Real Symmetric Band Matrix
RSG	Determine the Eigenvalues/vectors of a Real Symmetric Generalized Eigenproblem
RSGAB	Determine the Eigenvalues/vectors of a Real Symmetric Generalized Eigenproblem
RSGBA	Determine the Eigenvalues/vectors of a Real Symmetric Generalized Eigenproblem

Eigenvalues and Eigenvectors  
 EISPACK Subprograms not in this Guide –

Name	Function
RSM	Determine All Eigenvalues and Some Eigenvectors of a Real Symmetric Matrix
RSP	Determine the Eigenvalues/vectors of a Real Symmetric Packed Matrix
RST	Determine the Eigenvalues/vectors of a Real Symmetric Tridiagonal Matrix
RT	Determine the Eigenvalues/vectors of a Real Tridiagonal Matrix
SVD	Compute the Singular Value Decomposition of a Real Rectangular Matrix
TINVIT	Determine Some Eigenvectors of a Real Symmetric Tridiagonal Matrix
TQL1	Determine the Eigenvalues of a Real Symmetric Tridiagonal Matrix
TRBAK1	Back Transform Eigenvectors following TRED1
TRBAK3	Back Transform Eigenvectors following TRED3
TRED3	Reduce a Real Symmetric Matrix to Real Symmetric Tridiagonal Form
TRIDIB	Determine Some Eigenvalues of a Real Symmetric Tridiagonal Matrix
TSTURM	Determine Some Eigenvalues/vectors of a Real Symmetric Tridiagonal Matrix

## 6 Fast Fourier Transforms

---

### Overview

This chapter explains how to use the SCILIB Fast Fourier Transform (FFT) subprograms. The operations covered are

- one-dimensional complex-to-complex FFT subprograms
- one-dimensional real-to-complex and complex-to-real FFT subprograms
- one-dimensional simultaneous complex-to-complex FFT subprograms
- one-dimensional simultaneous real-to-complex and complex-to-real FFT subprograms

---

### Chapter Objectives

After reading this chapter you will

- understand the SCILIB FFT subprogram restrictions
- know how to augment subprograms with zero-value data points
- know how to use the described subprograms

---

### What You Need to Know to Use These Subprograms

Strictly speaking, an FFT is not a type of transform but a class of algorithms for efficiently computing the discrete Fourier transform (DFT).

Although the DFT is defined for any number of data points, the SCILIB FFT subprograms restrict the number of points to certain forms. For single one-dimensional transforms, the number of points must be a power of two:

Fast Fourier Transforms  
Supplemental Reading

$$l = 2^p, \quad p \geq 0.$$

For simultaneous transforms, the number of points in each direction must be a product of powers of two, three, and five:

$$l = 2^p 3^q 5^r, \quad p, q, r \geq 0.$$

While these restrictions limit the utility of the subprograms, the gain in speed is enormous. You can frequently adapt your data set to the SCILIB FFT subprograms by augmenting it with enough zero-value data points to reach the next acceptable number of points. Doing so slightly changes the problem, which may or may not be important, depending on the problem. For example, adding zero-value points to a time series changes the implied sampling frequency, but adding zero-value points to data sets before using FFT subprograms to compute convolutions does not change the result.

---

## Supplemental Reading

Brigham, E.O. *The Fast Fourier Transform*. Englewood Cliffs, NJ: Prentice-Hall, Inc. 1974.

Rabiner, L.R., and B. Gold. *Theory and Application of Digital Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, Inc. 1975.

**NAME** CFFT2 – One-Dimensional Complex-to-Complex FFT

### Purpose

Given an array of complex data, this subprogram computes the one-dimensional unscaled forward or inverse discrete Fourier transform using a radix 2 fast Fourier transform (FFT) algorithm.

The one-dimensional unscaled forward discrete Fourier transform of  $z(n)$ , for  $n = 1, 2, \dots, l$ , is defined by

$$Z(m) = \sum_{n=1}^l z(n) e^{-2\pi i(m-1)(n-1)/l}$$

for  $m = 1, 2, \dots, l$  and  $i = \sqrt{-1}$ .

Alternatively, the one-dimensional unscaled inverse discrete Fourier transform of  $Z(m)$ , for  $m = 1, 2, \dots, l$ , is defined by

$$z(n) = \sum_{m=1}^l Z(m) e^{+2\pi i(m-1)(n-1)/l}$$

for  $n = 1, 2, \dots, l$ .

CFFT2 requires that  $l$  be a power of 2, i.e., of the form  $l = 2^p$ , where  $p \geq 0$ .

### Usage

Because it is common to use one data set length repetitively, this subprogram has a separate initialization call so that the setup can be performed only once for each different transform size. You will, therefore, always have at least two **CALL** statements to the FFT subprogram, using the same working storage array. Refer to "Example."

SCILIB:

```

      INTEGER*8      init, isgn, l
      COMPLEX*16    zin(l), zout(l)
      REAL*8        work(5*l)
      CALL CFFT2(init, isgn, l, zin, work, zout)
  
```

### Input

**init** Initialization flag:  
       **init**  $\neq$  0 initialize **work** for subsequent transforms of length **l**.  
       **init** = 0 compute transform.

**isgn** Operation flag: if **init** = 0,

**isgn** < 0                    compute unscaled forward transform.  
**isgn** > 0                    compute unscaled inverse transform.

The sign of the exponential is the same as the sign of **isgn**.

**l**                                Number of data points, of the form  $l = 2^p$ , with  $p \geq 0$ .  
**zin**                             Array of data to be transformed. Not used if **init**  $\neq$  0. **zin** may be equivalenced to **work**, in which case the input values may be overwritten.

### Working Storage

**work**                            If **init**  $\neq$  0, **work** is initialized for computing transforms of length **l**.  
                                       If **init** = 0, **work** must have been initialized by a previous call with this value of **l** in which **init**  $\neq$  0.

### Output

**zout**                            Array of transformed data if **init** = 0. Not used if **init**  $\neq$  0.

### Notes

It is usual to scale the inverse transform by multiplying the summation by  $1/l$  so that the inverse transform of the forward transform of a data set returns the original data. This subprogram omits this scaling, meaning that the inverse transform of the forward transform of a data set is the original data multiplied by **l**.

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

**init** = 0 and **isgn** = 0;  
**l** not a power of 2.

### Example

Compute the forward discrete Fourier transform of two COMPLEX\*16 data sets of length 1024. The length of working storage is  $5 \times 1024 = 5120$ .

Fast Fourier Transforms  
**CFFT2 – One-Dimensional Complex-to-Complex FFT**

```
INTEGER*8  INIT, ISGN, L
COMPLEX*16 ZIN1(1024), ZOUT1(1024), ZIN2(1024), ZOUT2(1024)
REAL*8     WORK(5120)
L = 1024
INIT = 1
CALL CFFT2 (INIT, ISGN, L, ZIN1, WORK, ZOUT1) ! INITIALIZE
INIT = 0
ISGN = -1
CALL CFFT2 (INIT, ISGN, L, ZIN1, WORK, ZOUT1) ! FIRST TRANSFORM
CALL CFFT2 (INIT, ISGN, L, ZIN2, WORK, ZOUT2) ! SECOND TRANSFORM
```

Fast Fourier Transforms  
CFTFAX/CFFTMLT – Simultaneous One-Dimensional FFT

**NAME** CFTFAX/CFFTMLT – Simultaneous One-Dimensional FFT

**Purpose**

Given a number of sets of one-dimensional complex data with real and imaginary parts in separate real arrays, subroutine CFFTMLT computes all of their one-dimensional forward or inverse discrete Fourier transforms using a radix 2-3-5 fast Fourier transform (FFT) algorithm.

The one-dimensional forward discrete Fourier transform of a complex set of data  $z(n)$ , for  $n = 1, 2, \dots, l$ , is defined by

$$Z(m) = \sum_{n=1}^l z(n) e^{-2\pi i(m-1)(n-1)/l}$$

for  $m = 1, 2, \dots, l$  and  $i = \sqrt{-1}$ .

Alternatively, the one-dimensional unscaled inverse discrete Fourier transform of  $Z(m)$ , for  $m = 1, 2, \dots, l$ , is defined by

$$z(n) = \sum_{m=1}^l Z(m) e^{+2\pi i(m-1)(n-1)/l}$$

for  $n = 1, 2, \dots, l$ .

These subprograms perform forward or unscaled inverse transform operations simultaneously on a number of data sets. They require that the length  $l$  of the data sets be a product of powers of 2, 3, and 5, i.e., of the form

$$l = 2^p 3^q 5^r,$$

where  $p, q, r \geq 0$ . Refer to "Notes" for a partial list of permissible values of  $l$ .

The complex data,  $z$  or  $Z$ , are stored with real and imaginary parts in separate real arrays,  $x$  and  $y$ , respectively.

**Usage**

Because it is common to use one data set length repetitively, subroutine CFFTMLT has a separate initialization subprogram, CFTFAX, so that the setup can be performed only once for each different transform size. You will, therefore, always have at least one **CALL CFTFAX** statement and at least one **CALL CFFTMLT** statement, using the same working storage arrays. Refer to "Example."

**SCILIB:**

INTEGER\*8                   work3(19), incl, incn, l, n, isgn  
REAL\*8                      x(lenxy), y(lenxy), work1(4\*1\*n), work2(2\*1)

Initialization call: CALL CFTFAX(l, work3, work2)

Transform call: CALL CFFTMLT(x, y, work1, work2, work3, incl, incn, l, n, isgn)

## Input

**x and y**                      Arrays containing **n** data sets, each consisting of **l** data points, to be transformed. Typically, **x** and **y** will be two- or three-dimensional arrays with each data set being a one-dimensional array section. Refer to “Notes” for suggested usages.

Treating **x** and **y** as one-dimensional arrays results in

$$\text{lenxy} = (l-1) \times \text{incl} + (n-1) \times \text{incn} + 1.$$

The real and imaginary parts of the *i*-th data point of the *j*-th data set,  $1 \leq i \leq l$ ,  $1 \leq j \leq n$ , are stored in

$$\mathbf{x}((i-1) \times \text{incl} + (j-1) \times \text{incn} + 1)$$

and

$$\mathbf{y}((i-1) \times \text{incl} + (j-1) \times \text{incn} + 1),$$

respectively.

**incl**                              Storage increment between successive elements of the same data set, **incl** > 0. Use **incl** = 1 if each data set is stored contiguously in **x** and **y**.

**incn**                              Storage increment between corresponding data points of successive data sets, **incn** > 0.

**l**                                      Number of data points in each data set, of the form  $l = 2^p 3^q 5^r$ , with  $p, q, r \geq 0$ .

**n**                                      The number of data sets, **n** > 0.

**isgn**                                Operation flag:

<b>isgn</b> = -1	compute forward transform.
<b>isgn</b> = +1	compute inverse transform.

The sign of the exponential is the same as the sign of **isgn**.

## Working Storage

**work1**                              Array used for work space.

**work2**                              Array, initialized by CFTFAX for use as work space in CFFTMLT.

**work3**            Array, initialized by CFTFAX for use as work space in CFFTMLT. CFTFAX returns **work3(1) = -99** if **l** is not factorable as specified above.

## Output

**x and y**            The transformed data replaces the input.

## Notes

Typically, **x** and **y** will be two- or three-dimensional arrays with each data set being a one-dimensional section of the arrays, i.e., all but one subscript will be constant within a data set.

If **x** and **y** are two-dimensional arrays of dimension **ldxy** by **mdxy**, and if the data sets are stored in the columns of **x** and **y**, then  $1 \leq l \leq ldxy$ ,  $n \leq mdxy$ , **incl = 1**, and **incn = ldxy**. For example:

**CALL CFFTMLT (x, y, work1, work2, work3, 1, ldxy, l, n, isgn)**

If **x** and **y** are two-dimensional arrays as above and data sets are stored in rows of **x** and **y**,  $1 \leq mdxy$ ,  $n \leq ldxy$ , **incl = ldxy**, and **incn = 1**. For example:

**CALL CFFTMLT (x, y, work1, work2, work3, ldxy, 1, l, n, isgn)**

The subprograms generally are faster if the data sets are the rows of the arrays, so that **incn = 1**, rather than the columns. Otherwise, it is generally better to have an odd value of **incn**.

If **x** and **y** are three-dimensional arrays of dimension **ldxy** by **mdxy**-by-**ndxy**, then **incl** and **incn** will usually be 1, **ldxy**, or **ldxy**×**mdxy**, depending on which of the subscripts of the three-dimensional array varies within a data set, which subscript varies between data sets, and which remains constant. Specifically, if the subscript that varies within a data set is the

1st	subscript, use	<b>incl = 1.</b>
2nd	subscript, use	<b>incl = ldxy.</b>
3rd	subscript, use	<b>incl = ldxy×mdxy.</b>

Similarly, if the subscript that varies between data sets is the

1st	subscript, use	<b>incn = 1.</b>
2nd	subscript, use	<b>incn = ldxy.</b>
3rd	subscript, use	<b>incn = ldxy×mdxy.</b>

**incl**, **incn**, **l**, and **n** must be such that no two points of any data sets occupy the same element of **x** and **y**. CFFTMLT detects this situation if

**incl < n × gcd(incl,incn)**

and

$\text{incn} < l \times \text{gcd}(\text{incl}, \text{incn})$

where  $\text{gcd}(\cdot, \cdot)$  is the greatest common divisor.

If an error in the arguments is detected, CFFTMLT calls error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are

```

work3(1) = -99,
incl ≤ 0,
incn ≤ 0,
l not of the form 2p 3q 5r for p,q,r ≥ 0,
n ≤ 0, and
incl, incn, l, and n are incompatible.

```

The following list indicates some of the permissible values of  $l$ . Although  $l$  can be any value of the form  $2^p 3^q 5^r$  where  $p, q, r \geq 0$ , this list only shows values not exceeding 1000:

1, 2, 3, 4, 5, 6, 8, 9, 10, 12, 15, 16, 18, 20, 24, 25, 27, 30, 32, 36, 40, 45, 48, 50, 54, 60, 64, 72, 75, 80, 81, 90, 96, 100, 108, 120, 125, 128, 135, 144, 150, 160, 162, 180, 192, 200, 216, 225, 240, 243, 250, 256, 270, 288, 300, 320, 324, 360, 375, 384, 400, 405, 432, 450, 480, 486, 500, 512, 540, 576, 600, 625, 640, 648, 675, 720, 729, 750, 768, 800, 810, 864, 900, 960, 972, and 1000.

### Example 1

Compute the forward discrete Fourier transform of 256 complex data sets of length 1024. Real and imaginary parts of data sets are stored as columns of arrays  $X$  and  $Y$  whose dimensions are 1025 by 256.

```

INTEGER*8 WORK3(19), INCL, INCN, L, N, ISGN
REAL*8     X(1025,256), Y(1025,256), WORK1(1048576),
           WORK2(512)

L = 1024
INCL = 1
N = 256
INCN = 1025
ISGN = -1
CALL CFTFAX (L, WORK3, WORK2)
IF ( WORK3(1) .EQ. -99 ) THEN
    handle error condition
END IF
CALL CFFTMLT (X, Y, WORK1, WORK2, WORK3, INCL, INCN, L, N, ISGN)

```

Fast Fourier Transforms  
CFTFAX/CFFTMLT – Simultaneous One-Dimensional FFT

**Example 2**

Compute the inverse discrete Fourier transform of 1024 complex data sets of length 256. Real and imaginary parts of data sets are stored as rows of arrays X and Y whose dimensions are 1025 by 256.

```
INTEGER*8 WORK3(19),INCL,INCN,L,N,ISGN
REAL*8    X(1025,256),Y(1025,256),WORK1(1048576),
          WORK2(512)

L = 256
INCL = 1025
N = 1024
INCN = 1
ISGN = 1
CALL CFTFAX (L,WORK3,WORK2)
IF ( WORK3(1) .EQ. -99 ) THEN
    handle error condition
END IF
CALL CFFTMLT (X,Y,WORK1,WORK2,WORK3,INCL,INCN,L,N,ISGN)
```

**NAME** CRFFT2 – Complex-to-Real One-Dimensional FFT

**Purpose**

A complex sequence  $Z(m)$ , for  $m = 1, 2, \dots, l$ , is conjugate-symmetric about  $Z(l/2+1)$  if

$$\text{Im}(Z(1)) = \text{Im}(Z(l/2 + 1)) = 0$$

and

$$Z(l/2+1 + m) = \bar{Z}(l/2+1 - m), \quad m = 1, 2, \dots, l/2-1,$$

where  $\bar{Z}$  is the complex conjugate of  $Z$ .

Given an array of conjugate-symmetric complex data, this subprogram computes the one-dimensional, complex-to-real, unscaled forward or unscaled inverse, discrete Fourier transform using a radix 2 fast Fourier transform (FFT) algorithm optimized for real output.

The one-dimensional unscaled forward discrete Fourier transform of a data set,  $z(n)$ , for  $n = 1, 2, \dots, l$ , is defined by

$$Z(m) = \sum_{n=1}^l z(n) e^{-2\pi i(m-1)(n-1)/l}$$

for  $m = 1, 2, \dots, l$  and  $i = \sqrt{-1}$ .

Alternatively, the one-dimensional unscaled inverse discrete Fourier transform of the data set  $Z(m)$ , for  $m = 1, 2, \dots, l$ , is defined by

$$z(n) = \sum_{m=1}^l Z(m) e^{+2\pi i(m-1)(n-1)/l}$$

for  $n = 1, 2, \dots, l$ .

This subprogram requires that  $l$  be a power of 2, i.e., of the form  $l = 2^p$ , where  $p \geq 0$ .

**Usage**

Because it is common to use one data set length repetitively, this subprogram has a separate initialization call so that the setup can be performed only once for each different transform size. Therefore, you will always have at least two **CALL** statements to the FFT subprogram, using the same working storage array.

SCILIB:

INTEGER\*8            **init, isgn, l**



## Example

Compute the forward discrete Fourier transform of two conjugate-symmetric COMPLEX\*16 data sets of length 1024. Only the first 513 elements of the input data sets are stored. The length of working storage is  $3*1024+4 = 3076$ .

```
INTEGER*8  INIT, ISGN, L
COMPLEX*16 ZIN1(513), ZIN2(513)
REAL*8     WORK(3076), XOUT1(1024), XOUT2(1024)
L = 1024
INIT = 1
CALL CRFFT2 (INIT, ISGN, L, ZIN1, WORK, XOUT1)  !  INITIALIZE
INIT = 0
ISGN = -1
CALL CRFFT2 (INIT, ISGN, L, ZIN1, WORK, XOUT1)  !
          FIRST TRANSFORM
CALL CRFFT2 (INIT, ISGN, L, ZIN2, WORK, XOUT2)  !
          SECOND TRANSFORM
```

Fast Fourier Transforms  
**RCFFT2 – Real-to-Complex One-Dimensional FFT**

**NAME** RCFFT2 – Real-to-Complex One-Dimensional FFT

**Purpose**

Given an array of real data, this subprogram computes the one-dimensional, real-to-complex, unscaled forward or inverse, discrete Fourier transform using a radix 2 fast Fourier transform (FFT) algorithm optimized for real input.

The one-dimensional unscaled forward discrete Fourier transform of a data set,  $z(n)$ , for  $n = 1, 2, \dots, l$ , is defined by

$$Z(m) = \sum_{n=1}^l z(n) e^{-2\pi i(m-1)(n-1)/l}$$

for  $m = 1, 2, \dots, l$  and  $i = \sqrt{-1}$ .

When the sequence  $z(n)$  is real, the sequence  $Z(m)$  is conjugate-symmetric about  $Z(l/2+1)$ , i.e.,

$$\text{Im}(Z(1)) = \text{Im}(Z(l/2 + 1)) = 0$$

and

$$Z(l/2+1 + m) = \bar{Z}(l/2+1 - m), \quad m = 1, 2, \dots, l/2-1,$$

where  $\bar{Z}$  is the complex conjugate of  $Z$ .

This subprogram actually computes twice the above quantity:

$$Z(m) = 2 \sum_{n=1}^l z(n) e^{-2\pi i(m-1)(n-1)/l}$$

Alternatively, the one-dimensional unscaled inverse discrete Fourier transform of the data set  $Z(m)$ , for  $m = 1, 2, \dots, l$ , is defined by

$$z(n) = \sum_{m=1}^l Z(m) e^{+2\pi i(m-1)(n-1)/l}$$

for  $n = 1, 2, \dots, l$ .

Again, twice the above quantity is what is actually computed:

$$z(n) = 2 \sum_{m=1}^l Z(m) e^{+2\pi i(m-1)(n-1)/l}$$

This subprogram requires that  $l$  be a power of 2, i.e., of the form  $l = 2^p$ , where  $p \geq 0$ .



**RCFFT2 – Real-to-Complex One-Dimensional FFT**

the original data. This subprogram replaces this scaling with scaling both forward and inverse transforms by a factor of 2, meaning that the inverse transform of the forward transform of a data set does not return the original data.

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are

**isgn** = 0, and  
**I** not a power of 2.

**Example**

Compute the forward discrete Fourier transform of two REAL\*8 data sets of length 1024. Only the first 513 elements of the transformed data are computed. The length of working storage is  $3 \times 1024 + 4 = 3076$ .

```

INTEGER*8  INIT, ISGN, L
REAL*8     XIN1(1024), XIN2(1024), WORK(3076)
COMPLEX*16 ZOUT1(513), ZOUT2(513)
L = 1024
INIT = 1
CALL RCFFT2 (INIT, ISGN, L, XIN1, WORK, ZOUT1) ! INITIALIZE
INIT = 0
ISGN = -1
CALL RCFFT2 (INIT, ISGN, L, XIN1, WORK, ZOUT1) !
      FIRST TRANSFORM
CALL RCFFT2 (INIT, ISGN, L, XIN2, WORK, ZOUT2) !
      SECOND TRANSFORM

```

**NAME** FFTFAX/RFFTMLT – Simultaneous One-Dimensional FFT

**Purpose**

Given a number of one-dimensional real data sets, subroutine RFFTMLT computes the nonredundant portion of all of their one-dimensional forward real-to-complex discrete Fourier transforms using a radix 2-3-5 fast Fourier transform (FFT) algorithm optimized for real input. Alternatively, given the nonredundant parts of a number of conjugate-symmetric one-dimensional complex data sets, subroutine RFFTMLT computes the inverse complex-to-real discrete Fourier transform using a radix 2-3-5 FFT algorithm optimized for real output.

The one-dimensional forward, scaled, real-to-complex Fourier transform of a real set of data  $z(n)$ , for  $n = 1, 2, \dots, l$ , is defined by:

$$Z(m) = \frac{1}{l} \sum_{n=1}^l z(n) e^{-2\pi i(m-1)(n-1)/l}$$

for  $m = 1, 2, \dots, l$ , where  $i = \sqrt{-1}$ .

The sequence  $Z(m)$  is conjugate-symmetric about  $Z(l/2+1)$ , i.e.,

$$\text{Im}(Z(1)) = \text{Im}(Z(l/2 + 1)) = 0$$

and

$$Z(l/2+1 + m) = \bar{Z}(l/2+1 - m), \quad m = 1, 2, \dots, l/2-1,$$

where  $\bar{Z}$  is the complex conjugate of  $Z$ . Therefore, the nonredundant part consists of the first  $l/2+1$  elements of  $Z$ , which is all of  $Z$  that is computed or stored.

Alternatively, if  $Z(m)$ , for  $m = 1, 2, \dots, l$ , is a conjugate-symmetric complex data set, the one-dimensional, inverse, unscaled complex-to-real Fourier transform of  $Z(m)$  is defined by:

$$z(n) = \sum_{m=1}^l Z(m) e^{+2\pi i(m-1)(n-1)/l}$$

for  $n = 1, 2, \dots, l$ . Only the nonredundant part of  $Z$  is used.

These subprograms perform forward or inverse transform operations simultaneously on a number of data sets. They require that the length  $l$  of the data sets be a product of powers of 2, 3, and 5, i.e., of the form:

$$l = 2^p 3^q 5^r,$$

where  $p, q, r \geq 0$ , and where either  $l = 1$  or  $l$  is even. Refer to “Notes” for a partial list of permissible values of  $l$ .

## Usage

Because it is common to use one data set length repetitively, subroutine RFFTMLT has a separate initialization subprogram, FFTFAX, so that the setup can be performed only once for each different transform size. You will, therefore, always have at least one CALL FFTFAX statement and at least one CALL RFFTMLT statement, using the same working storage arrays. Refer to “Example 1.”

### SCILIB:

```

INTEGER*8      work3(19), incl, incn, l, n, isgn
REAL*8        x(lenx), work1(2*1*n), work2(2*1)
Initialization call: CALL FFTFAX(l, work3, work2)
Transform call: CALL RFFTMLT(x, work1, work2, work3, incl, incn, l, n,
                             isgn)

```

## Input

**x**            Array containing  $n$  one-dimensional data sets, each consisting of  $l$  real data points or the first  $l/2+1$  complex data points of a conjugate-symmetric complex data set of length  $l$ , to be transformed. Typically,  $x$  is a two- or three-dimensional array with each set of data being a one-dimensional array section. Refer to “Notes” for suggested usages.

Treating  $x$  as a one-dimensional array results in:

$$\text{lenx} = (l+1) \times \text{incl} + (n-1) \times \text{incn} + 1.$$

For a forward real-to-complex transform, the  $i$ -th real data point of the  $j$ -th data set,  $1 \leq i \leq l$ ,  $1 \leq j \leq n$ , is stored in:

$$x((i-1) \times \text{incl} + (j-1) \times \text{incn} + 1).$$

For an inverse complex-to-real transform, the real part of the  $i$ -th data point of the  $j$ -th data set,  $1 \leq i \leq l/2+1$ ,  $1 \leq j \leq n$ , is stored in:

$$x((2 \times i - 2) \times \text{incl} + (j-1) \times \text{incn} + 1)$$

and the imaginary part is stored in:

$$x((2 \times i - 1) \times \text{incl} + (j-1) \times \text{incn} + 1),$$

respectively.

<b>incl</b>	Storage increment between successive elements of the same data set, <b>incl</b> > 0. Use <b>incl</b> = 1 if each data set is stored contiguously in <b>x</b> .				
<b>incn</b>	Storage increment between corresponding data points of successive data sets, <b>incn</b> > 0.				
<b>l</b>	Number of data points in each complete data set, of the form $l = 2^p 3^q 5^r$ , with $q, r \geq 0$ and either $l = 1$ or $p \geq 1$ .				
<b>n</b>	The number of data sets, <b>n</b> > 0.				
<b>isgn</b>	Option flag: <table border="0" style="margin-left: 2em;"> <tr> <td><b>isgn</b> = -1</td> <td>compute real-to-complex forward transform.</td> </tr> <tr> <td><b>isgn</b> = +1</td> <td>compute complex-to-real inverse transform.</td> </tr> </table>	<b>isgn</b> = -1	compute real-to-complex forward transform.	<b>isgn</b> = +1	compute complex-to-real inverse transform.
<b>isgn</b> = -1	compute real-to-complex forward transform.				
<b>isgn</b> = +1	compute complex-to-real inverse transform.				

### Working Storage

<b>work1</b>	Array used for work space.
<b>work2</b>	Array, initialized by FFTFAX for use as work space in RFFTMLT.
<b>work3</b>	Array, initialized by FFTFAX for use as work space in RFFTMLT. FFTFAX returns <b>work3</b> (1) = -99 if <b>l</b> is not factorable as specified above.

### Output

<b>x</b>	<p>The transformed data replaces the input.</p> <p>For a forward real-to-complex transform, the real part of the <math>i</math>-th output point of the <math>j</math>-th data set, <math>1 \leq i \leq l/2+1</math>, <math>1 \leq j \leq n</math>, is stored in:</p> $\mathbf{x}((2 \times i - 2) \times \mathbf{incl} + (j - 1) \times \mathbf{incn} + 1)$ <p>and the imaginary part is stored in</p> $\mathbf{x}((2 \times i - 1) \times \mathbf{incl} + (j - 1) \times \mathbf{incn} + 1),$ <p>respectively. If needed, the remaining <math>(l/2 - 1) \times n</math> complex output values may be formed by using the conjugate-symmetry condition.</p> <p>For an inverse complex-to-real transform, the <math>i</math>-th real output point of the <math>j</math>-th data set, <math>1 \leq i \leq l</math>, <math>1 \leq j \leq n</math>, is stored in:</p> $\mathbf{x}((i - 1) \times \mathbf{incl} + (j - 1) \times \mathbf{incn} + 1).$
----------	--

Fast Fourier Transforms  
FFTFAX/RFFTMLT – Simultaneous One-Dimensional FFT

## Notes

Typically,  $\mathbf{x}$  will be a two- or three-dimensional array with each set of data being a one-dimensional section of the array, i.e., all but one subscript will be constant within a data set.

If  $\mathbf{x}$  is a two-dimensional array of dimension  $l\mathbf{d}\mathbf{x}$  by  $m\mathbf{d}\mathbf{x}$ , and if the data sets are stored in the columns of  $\mathbf{x}$ , then  $l+2 \leq l\mathbf{d}\mathbf{x}$ ,  $n \leq m\mathbf{d}\mathbf{x}$ ,  $incl = 1$ , and  $incn = l\mathbf{d}\mathbf{x}$ . For example:

**CALL RFFTMLT ( $\mathbf{x}$ , work1, work2, work3, 1,  $l\mathbf{d}\mathbf{x}$ , 1,  $n$ , isgn)**

If  $\mathbf{x}$  is a two-dimensional array as above and the data sets are stored in the rows of  $\mathbf{x}$ , then  $l+2 \leq m\mathbf{d}\mathbf{x}$ ,  $n \leq l\mathbf{d}\mathbf{x}$ ,  $incl = l\mathbf{d}\mathbf{x}$ , and  $incn = 1$ . For example:

**CALL RFFTMLT ( $\mathbf{x}$ , work1, work2, work3,  $l\mathbf{d}\mathbf{x}$ , 1, 1,  $n$ , isgn)**

The subprograms generally are faster if the data sets are the rows of the arrays, so that  $incn = 1$ , rather than the columns. Otherwise, it is generally better to have an odd value of  $incn$ .

If  $\mathbf{x}$  is a three-dimensional array of dimension  $l\mathbf{d}\mathbf{x}$  by  $m\mathbf{d}\mathbf{x}$  by  $n\mathbf{d}\mathbf{x}$ , then  $incl$  and  $incn$  will usually be 1,  $l\mathbf{d}\mathbf{x}$ , or  $l\mathbf{d}\mathbf{x} \times m\mathbf{d}\mathbf{x}$ , depending on which of the subscripts of the three-dimensional array varies within a data set, which subscript varies between data sets, and which remains constant. Specifically, if the subscript that varies within a data set is the

1st	subscript, use	$incl = 1$ .
2nd	subscript, use	$incl = l\mathbf{d}\mathbf{x}$ .
3rd	subscript, use	$incl = l\mathbf{d}\mathbf{x} \times m\mathbf{d}\mathbf{x}$ .

Similarly, if the subscript that varies between data sets is the

1st	subscript, use	$incn = 1$ .
2nd	subscript, use	$incn = l\mathbf{d}\mathbf{x}$ .
3rd	subscript, use	$incn = l\mathbf{d}\mathbf{x} \times m\mathbf{d}\mathbf{x}$ .

$incl$ ,  $incn$ ,  $l$ , and  $n$  must be such that no two points of any data sets occupy the same element of  $\mathbf{x}$ . RFFTMLT detects this situation if

$incl < n \times \text{gcd}(incl, incn)$

and

$incn < l \times \text{gcd}(incl, incn)$

where  $\text{gcd}(\cdot, \cdot)$  is the greatest common divisor.

If an error in the arguments is detected, RFFTMLT calls error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9,

Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

```

work3(1) = -99,
incl ≤ 0,
incn ≤ 0,
l not of the required form  $2^p 3^q 5^r$ , with  $l = 1$  or  $l$  even,
n ≤ 0, and
incl, incn, l, and n are incompatible.

```

The following list indicates some of the permissible values of  $l$ . Although  $l$  can be any even value of the form  $2^p 3^q 5^r$  where  $q, r \geq 0$  and either  $l = 1$  or  $p \geq 1$ , this list only shows the values not exceeding 1000:

1, 2, 4, 6, 8, 10, 12, 16, 18, 20, 24, 30, 32, 36, 40, 48, 50, 54, 60, 64, 72, 80, 90, 96, 100, 108, 120, 128, 144, 150, 160, 162, 180, 192, 200, 216, 240, 250, 256, 270, 288, 300, 320, 324, 360, 384, 400, 432, 450, 480, 486, 500, 512, 540, 576, 600, 640, 648, 720, 750, 768, 800, 810, 864, 900, 960, 972, and 1000.

### Example 1

Compute the forward real-to-complex discrete Fourier transform of 256 real data sets of length 1024. The real input data sets are stored as columns of REAL\*8 array X whose dimensions are 1027 by 256. The complex output data sets are stored as columns of array X, with the real parts in rows 1, 3, 5, ..., 1025, and the imaginary parts in rows 2, 4, 6, ..., 1026.

```

INTEGER*8 WORK3(19), INCL, INCN, L, N, ISGN
REAL*8 X(1027,256), WORK1(1048576), WORK2(512)
L = 1024
INCL = 1
N = 256
INCN = 1027
ISGN = -1
CALL FFTFAX (L, WORK3, WORK2)
IF ( WORK3(1) .EQ. -99 ) THEN
    handle error condition
END IF
CALL RFFTMLT (X, WORK1, WORK2, WORK3, INCL, INCN, L, N, ISGN)

```

### Example 2

Compute the inverse complex-to-real discrete Fourier transform of 1024 sets of conjugate-symmetric complex data of length 256. The real and imaginary parts of the first 129 complex data points of the input data sets are stored as the rows of array X whose dimensions are 1025 by 258, with the real parts in columns 1, 3, 5, ..., 257, and the imaginary parts in columns 2, 4, 6, ..., 258. The real output data sets will be stored by row in the first 256 columns of X.

## Fast Fourier Transforms

### FFTFAX/RFFTMLT - Simultaneous One-Dimensional FFT

```
INTEGER*8 WORK3(19),INCL,INCN,L,N,ISGN
REAL*8    X(1025,258),WORK1(1048576),WORK2(512)
L = 256
INCL = 1025
N = 1024
INCN = 1
ISGN = 1
CALL FFTFAX (L,WORK3,WORK2)
IF ( WORK3(1) .EQ. -99 ) THEN
    handle error condition
END IF
CALL RFFTMLT (X,WORK1,WORK2,WORK3,INCL,INCN,L,N,ISGN)
```

# 7 Correlation and Convolution Subprograms

---

## Overview

This chapter explains how to use the SCILIB subprograms available for correlations, convolutions, and related operations such as filtering by means of convolutions.

---

## Chapter Objectives

After reading this chapter, you will know how to use the described subprograms to compute correlation and convolution

---

## What You Need to Know to Use These Subprograms

The subprograms presented here can be used to compute both discrete correlations and discrete convolutions. See the specific subprogram descriptions for details.

---

## Supplemental Reading

Rabiner, L.R., and B. Gold. *Theory and Application of Digital Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, Inc. 1975.

---

Correlation and Convolution Subprograms  
**FILTERG – Discrete Correlation**

**NAME** FILTERG – Discrete Correlation

**Purpose**

This subprogram computes the fully engaged portion of the discrete correlation of a data vector and a filter vector. It can be used to compute the complete discrete correlation (the fully engaged portion plus the *tails*) by appending zeros to the ends of the data vector. Refer to “Example 2.”

If  $f_j, j = 1, 2, \dots, m$ , and  $x_i, i = 1, 2, \dots, n$ , are a filter vector and a data vector, respectively, their discrete correlation  $\tilde{y}_i$  is defined by

$$\tilde{y}_i = \sum_j f_j x_{i+j-1},$$

for  $i = -m+2, -m+3, \dots, n$ , where the sum is taken over all indices  $j$  for which both  $f_j$  and  $x_{i+j-1}$  are defined.

This subprogram computes only the fully engaged portion of the correlation, i.e., the part where the sums have exactly  $\min(m,n)$  terms. Hence, if  $m \leq n$ , it computes

$$y_i = \sum_{j=1}^m f_j x_{i+j-1},$$

for  $i = 1, 2, \dots, n-m+1$ .

The discrete convolution  $\tilde{z}_i$  of the vectors  $f$  and  $x$  is defined by

$$\tilde{z}_i = \sum_j f_{i-j+1} x_j,$$

for  $i = 1, 2, \dots, m+n-1$ , where the sum is taken over all indices  $j$  for which both  $f_{i-j+1}$  and  $x_j$  are defined.

This subprogram computes only the fully engaged portion of the convolution, i.e., the part where the sums have exactly  $\min(m,n)$  terms. Hence, if  $m \leq n$ , it computes

$$z_i = \sum_{j=1}^m f_{m+1-j} x_{i+j-1}$$

for  $i = 1, 2, \dots, n-m+1$ . A comparison of the definitions of  $y_i$  and  $z_i$  shows that the convolution may be computed by storing the  $f$  vectors in reverse order before calling FILTERG.

## Usage

SCILIB:

```
INTEGER*8      m, n
REAL*8         f(m), x(n), y(n-m+1)
CALL FILTERG(f, m, x, n, y)
```

## Input

**f**            Array containing the filter vector  $f$  of length  $m$ .  
**m**            The length of the  $f$  vector,  $m > 0$ .  
**x**            Array containing the data vector  $x$  of length  $n$ .  
**n**            The length of the  $x$  vector,  $n \geq m$ .

## Output

**y**            The fully engaged correlation vector  $y$  of length  $n-m+1$ .

## Notes

To compute the complete correlation vector, including both tails as well as the fully engaged portion, append  $m-1$  zeros to each end of the  $x$  vector. The fully engaged portion of the correlation of the resulting vector is the complete correlation corresponding to the original  $x$  vector. Refer to "Example 2."

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure.

$m \leq 0$ , and  
 $n \leq m$ .

## Example 1

Compute the fully engaged portion of the discrete correlation of the REAL\*8 vectors  $f = (2, 1)$  and  $x = (4, 1, 3, 5, 2)$  stored in arrays F and X, respectively. In this instance,  $m = 2$  and  $n = 5$ .

```
INTEGER*8 M,N
REAL*8    F(2),X(5),Y(4)
DATA      F / 2.0 , 1.0 /
DATA      X / 4.0 , 1.0 , 3.0 , 5.0, 2.0 /
M = 2
N = 4
CALL FILTERG (F,M,X,N,Y)
```

The result is  $y = (9, 5, 11, 12)$ .

Correlation and Convolution Subprograms  
**FILTERG – Discrete Correlation**

**Example 2**

Compute the complete discrete correlation of the REAL\*8 vectors  $f = (1, 2)$  and  $x = (1, 3, 5)$  stored in arrays F and X, respectively. Thus  $m = 2$ , and to get the complete correlation, we append  $m-1 = 1$  zero to each end of  $x$ , getting  $\bar{x} = (0, 1, 3, 5, 0)$  and  $n = 5$ .

```
INTEGER*8 M,N
REAL*8    F(2),X(5),Y(4)
DATA      F / 1.0 , 2.0 /
DATA      X / 0.0 , 1.0 , 3.0 , 5.0, 0.0 /
M = 2
N = 4
CALL FILTERG (F,M,X,N,Y)
```

The result is  $y = (2, 7, 13, 5)$ .

**NAME** FILTERS – Discrete Correlation

**Purpose**

This subprogram computes the fully engaged portion of the discrete correlation of a data vector and a symmetric filter vector. It can be used to compute the complete discrete correlation (the fully engaged portion plus the *tails*) by appending zeros to the ends of the data vector. Refer to "Example 2."

If  $f_i, i = 1, 2, \dots, m$ , and  $x_i, i = 1, 2, \dots, n$ , are a filter vector and a data vector, respectively, their discrete correlation  $\tilde{y}_i$  is defined by

$$\tilde{y}_i = \sum_j f_j x_{i+j-1},$$

for  $i = -m+2, -m+3, \dots, n$ , where the sum is taken over all indices  $j$  for which both  $f_j$  and  $x_{i+j-1}$  are defined.

The filter vector is assumed to be symmetric, i.e.,  $f_i = f_{m+1-i}$ ,  $i = 1, 2, \dots, [m/2]$ , so only the first  $[m/2]$  elements of  $f$  must be stored.

This subprogram computes only the fully engaged portion of the correlation, i.e., the part where the sums have exactly  $\min(m,n)$  terms. Hence, if  $m \leq n$ , it computes

$$y_i = \begin{cases} f_{(m+1)/2} x_{i+(m+1)/2} + \sum_{j=1}^{(m-1)/2} f_j (x_{i+j-1} + x_{i+m-j}), & \text{if } m \text{ odd} \\ \sum_{j=1}^{m/2} f_j (x_{i+j-1} + x_{i+m-j}), & \text{if } m \text{ even} \end{cases}$$

for  $i = 1, 2, \dots, n-m+1$ .

The fully engaged portion of the discrete convolution  $z_i$  of the vectors  $f$  and  $x$  is defined by

$$z_i = \sum_{j=1}^m f_{m+1-j} x_{i+j-1}$$

for  $i = 1, 2, \dots, n-m+1$ . Because  $f$  is a symmetric vector,  $y_i = z_i$  for  $i = 1, 2, \dots, n-m+1$ ; i.e., the fully engaged discrete correlation is identical to the fully engaged discrete convolution.

Correlation and Convolution Subprograms  
FILTERS – Discrete Correlation

## Usage

SCILIB:

```
INTEGER*8      m, n
REAL*8         f((m+1)/2), x(n), y(n-m+1)
CALL FILTERS(f, m, x, n, y)
```

## Input

**f**                    Array containing the first  $(m+1)/2$  elements of the filter vector  $f$  of length  $m$ .

**m**                    The length of the  $f$  vector,  $m > 0$ .

**x**                    Array containing the data vector  $x$  of length  $n$ .

**n**                    The length of the  $x$  vector,  $n \geq m$ .

## Output

**y**                    The fully engaged correlation vector  $y$  of length  $n-m+1$ .

## Notes

To compute the complete correlation vector, including both tails as well as the fully engaged portion, append  $m-1$  zeros to each end of the  $x$  vector. The fully engaged portion of the correlation of the resulting vector is the complete correlation corresponding to the original  $x$  vector.

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

$m \leq 0$ , and  
 $n \leq m$ .

Refer to "Example 2."

## Example 1

Compute the fully engaged portion of the discrete correlation of the REAL\*8 vectors  $f = (-1, 2, -1)$  and  $x = (4, 1, 3, 5, 2)$  stored in arrays F and X, respectively. In this instance,  $m = 3$  and  $n = 5$ .

```
INTEGER*8 M,N
REAL*8    F(2),X(5),Y(3)
DATA      F / -1.0 , 4.0 /
DATA      X / 4.0 , 1.0 , 3.0 , 5.0, 2.0 /
M = 3
N = 4
CALL FILTERS (F,M,X,N,Y)
```

The result is  $y = (-3, 6, 15)$ .

### Example 2

Compute the complete discrete correlation of the REAL\*8 vectors  $f = (1, 2, 1)$  and  $x = (1, 3, 5)$  stored in arrays F and X, respectively. Thus  $m = 3$ , and to get the complete correlation, we append  $m-1 = 2$  zeros to each end of  $x$ , getting  $\bar{x} = (0, 0, 1, 3, 5, 0, 0)$  and  $n = 7$ .

```
INTEGER*8 M,N
REAL*8    F(2),X(7),Y(5)
DATA      F / 1.0 , 2.0 /
DATA      X / 0.0, 0.0 , 1.0 , 3.0 , 5.0, 0.0, 0.0 /
M = 2
N = 4
CALL FILTERS (F,M,X,N,Y)
```

The result is  $y = (1, 5, 12, 13, 5)$ .

Correlation and Convolution Subprograms  
**FILTERS – Discrete Correlation**

## 8 Linear Recurrences

---

### Overview

This chapter explains how to use SCILIB subprograms for a variety of linear recurrence operations.

---

### Chapter Objectives

After reading this chapter you will

- be able to recognize a recurrence
  - know how to use the described subprograms
- 

### What You Need to Know to Use These Subprograms

Recurrence subprograms were written for optimizing tridiagonal and pentadiagonal linear equation solvers. Consider using the LINPACK tridiagonal or band solvers described in Chapter 4 or the LAPACK solvers described in the HP MLIB LAPACK User's Guide.

Linear Recurrences  
FOLR/FOLRP – First Order Linear Recurrence

**NAME** FOLR/FOLRP – First Order Linear Recurrence

**Purpose**

Given real vectors  $a$  and  $x$  of length  $n$ , these subprograms solve the first order linear recurrence

$$y_1 = x_1$$
$$y_i = x_i \pm a_i y_{i-1}, \quad i = 2, 3, \dots, n,$$

overwriting the input  $x$  vector with the resulting  $y$  vector.

The operation indicated by  $\pm$  above is specified by the subprogram name used:

FOLR      $\pm$  represents  $-$ :  $y_i = x_i - a_i y_{i-1}$

FOLRP     $\pm$  represents  $+$ :  $y_i = x_i + a_i y_{i-1}$

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```
INTEGER*8      n, inca, incx
REAL*8         a(lena), x(lenx)
CALL FOLR(n, a, inca, x, incx)
INTEGER*8      n, inca, incx
REAL*8         a(lena), x(lenx)
CALL FOLRP(n, a, inca, x, incx)
```

**Input**

- n**            Number of elements of vectors  $a$  and  $x$  to be used in the recurrence,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference  $a$  or  $x$ .
- a**            Array of length  $\text{lena} = (n-1) \times |\text{inca}| + 1$  containing the  $n$ -vector  $a$ .
- inca**        Increment for the array  $a$ :
- inca** > 0      $a$  is stored forward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-1) \times \text{inca} + 1)$ .
- inca** < 0      $a$  is stored backward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-n) \times \text{inca} + 1)$ .

Use **inca** = 1 if the vector  $a$  is stored contiguously in array  $a$ , i.e., if  $a_i$  is stored in  $a(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

**x**                    Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx**                Increment for the array  $x$ , **incx**  $\neq 0$ :

**incx** > 0       $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx** < 0       $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

## Output

**x**                    If **n** = 0, then  $x$  is unchanged. Otherwise, the recurrence’s solution vector overwrites the input.

## Notes

The result is unspecified if  $a$  and  $x$  overlap such that any element of  $a$  shares a memory location with any element of  $x$ .

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

**n** < 0,  
**inca** = 0, and  
**incx** = 0.

Linear Recurrences  
FOLR/FOLRP – First Order Linear Recurrence

**Fortran Equivalent**

```
SUBROUTINE FOLR (N, A, INCA, X, INCX)
  INTEGER*8 N, INCA, INCX
  REAL*8 A(*), X(*)
  IF ( N .LT. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCA .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCX .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IA = 1 + INCA
  IX = 1 + INCX
  IF ( INCA .LT. 0 ) IA = 1 - (N-2) * INCA
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  DO 10 I = 2, N
    X(IX) = X(IX) - A(IA) * X(IX-INCX)
    IA = IA + INCA
    IX = IX + INCX
10 CONTINUE
  RETURN
  END
```

**Example**

Solve the REAL\*8 first order linear recurrence

$$y_1 = x_1$$
$$y_i = x_i - a_i y_{i-1}, \quad i = 2, 3, \dots, n,$$

where  $a$  and  $x$  are vectors 10 elements long stored in one-dimensional arrays  $A$  and  $X$  of dimension 20, and  $y$  overwrites  $x$ .

```
INTEGER*8 N, INCA, INCX
REAL*8 A(20), X(20)
N = 10
INCA = 1
INCX = 1
CALL FOLR (N, A, INCA, X, INCX)
```

**NAME** FOLR2/FOLR2P – First Order Linear Recurrence

**Purpose**

Given real vectors  $a$  and  $x$  of length  $n$ , these subprograms solve the first order linear recurrence

$$y_1 = x_1$$

$$y_i = x_i \pm a_i y_{i-1}, \quad i = 2, 3, \dots, n,$$

for the  $y$  vector.

The operation indicated by  $\pm$  above is specified by the subprogram name used:

FOLR2      $\pm$  represents  $-$ :  $y_i = x_i - a_i y_{i-1}$   
 FOLR2P    $\pm$  represents  $+$ :  $y_i = x_i + a_i y_{i-1}$

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

INTEGER*8      n, inca, incx, incy
REAL*8        a(lena), x(lenx), y(leny)
CALL FOLR2(n, a, inca, x, incx, y, incy)

INTEGER*8      n, inca, incx, incy
REAL*8        a(lena), x(lenx), y(leny)
CALL FOLR2P(n, a, inca, x, incx, y, incy)
  
```

**Input**

- n**            Number of elements of vectors  $a$ ,  $x$ , and  $y$  to be used in the recurrence,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference  $a$ ,  $x$ , or  $y$ .
- a**            Array of length  $lena = (n-1) \times |inca| + 1$  containing the  $n$ -vector  $a$ .
- inca**        Increment for the array  $a$ :
  - inca** > 0     $a$  is stored forward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-1) \times inca + 1)$ .
  - inca** < 0     $a$  is stored backward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-n) \times inca + 1)$ .

- Use  $\mathbf{inca} = 1$  if the vector  $a$  is stored contiguously in array  $\mathbf{a}$ , i.e., if  $a_i$  is stored in  $\mathbf{a}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.
- x**      Array of length  $\mathbf{lenx} = (\mathbf{n}-1) \times |\mathbf{incx}| + 1$  containing the  $n$ -vector  $x$ .
- incx**    Increment for the array  $\mathbf{x}$ ,  $\mathbf{incx} \neq 0$ :
- $\mathbf{incx} > 0$        $x$  is stored forward in array  $\mathbf{x}$ , i.e.,  $x_i$  is stored in  $\mathbf{x}((i-1) \times \mathbf{incx} + 1)$ .
- $\mathbf{incx} < 0$        $x$  is stored backward in array  $\mathbf{x}$ , i.e.,  $x_i$  is stored in  $\mathbf{x}((i-\mathbf{n}) \times \mathbf{incx} + 1)$ .
- Use  $\mathbf{incx} = 1$  if the vector  $x$  is stored contiguously in array  $\mathbf{x}$ , i.e., if  $x_i$  is stored in  $\mathbf{x}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.
- incy**    Increment for the array  $\mathbf{y}$ ,  $\mathbf{incy} \neq 0$ :
- $\mathbf{incy} > 0$        $y$  is stored forward in array  $\mathbf{y}$ , i.e.,  $y_i$  is stored in  $\mathbf{y}((i-1) \times \mathbf{incy} + 1)$ .
- $\mathbf{incy} < 0$        $y$  is stored backward in array  $\mathbf{y}$ , i.e.,  $y_i$  is stored in  $\mathbf{y}((i-\mathbf{n}) \times \mathbf{incy} + 1)$ .
- Use  $\mathbf{incy} = 1$  if the vector  $y$  is stored contiguously in array  $\mathbf{y}$ , i.e., if  $y_i$  is stored in  $\mathbf{y}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

## Output

- y**      Array of length  $\mathbf{leny} = (\mathbf{n}-1) \times |\mathbf{incy}| + 1$  containing the  $n$ -vector  $y$ . If  $\mathbf{n} = 0$ , then  $\mathbf{y}$  is unchanged. Otherwise,  $\mathbf{y}$  is the recurrence’s solution vector.

## Notes

The result is unspecified if  $\mathbf{a}$ ,  $\mathbf{x}$ , or  $\mathbf{y}$  overlap such that any element of  $a$ ,  $x$ , or  $y$  shares a memory location with any other element of  $a$ ,  $x$ , or  $y$ .

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

$$\begin{aligned} \mathbf{n} &< 0, \\ \mathbf{inca} &= 0, \end{aligned}$$

**incx = 0, and  
 incy = 0.**

### Fortran Equivalent

```

SUBROUTINE FOLR2 (N, A, INCA, X, INCX, Y, INCY)
  INTEGER*8 N, INCA, INCX, INCY
  REAL*8 A(*), X(*)
  IF ( N .LT. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCA .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCX .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCY .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IA = 1 + INCA
  IX = 1 + INCX
  IY = 1 + INCY
  IF ( INCA .LT. 0 ) IA = 1 - (N-2) * INCA
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  IF ( INCY .LT. 0 ) IY = 1 - (N-2) * INCY
  Y(IY-INCY) = X(IX-INCX)
  DO 10 I = 2, N
    Y(IY) = X(IX) - A(IA) * Y(IY-INCY)
    IA = IA + INCA
    IX = IX + INCX
    IY = IY + INCY
  10 CONTINUE
  RETURN
  END
  
```

### Example

Solve the REAL\*8 first order linear recurrence

$$y_1 = x_1$$

$$y_i = x_i - a_i y_{i-1}, \quad i = 2, 3, \dots, n,$$

where  $a$ ,  $x$  and  $y$  are vectors 10 elements long stored in one-dimensional arrays  $A$ ,  $X$ , and  $Y$  of dimension 20.

## Linear Recurrences

### FOLR2/FOLR2P – First Order Linear Recurrence

```
INTEGER*8 N, INCA, INCX, INCY
REAL*8    A(20), X(20), Y(20)
N = 10
INCA = 1
INCX = 1
INCY = 1
CALL FOLR2 (N, A, INCA, X, INCX, Y, INCY)
```

**NAME** FOLRC – First Order Linear Recurrence

**Purpose**

Given a real coefficient  $\alpha$  and a real vector  $a$  of length  $n$ , this subprogram solves the first order linear recurrence

$$x_1 = a_1$$

$$x_i = a_i + \alpha x_{i-1}, \quad i = 2, 3, \dots, n.$$

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

INTEGER*8      n, incx, inca
REAL*8        x(lenx), a(lena), alpha
CALL FOLRC(n, x, incx, a, inca, alpha)

```

**Input**

- n** Number of elements of vectors  $a$  and  $x$  to be used in the recurrence,  $n \geq 0$ . If  $n = 0$ , the subprogram does not reference  $a$  or  $x$ .
- incx** Increment for the array  $x$ ,  $incx \neq 0$ :
- incx > 0**  $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times incx + 1)$ .
- incx < 0**  $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times incx + 1)$ .
- Use **incx = 1** if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.
- a** Array of length  $lena = (n-1) \times |inca| + 1$  containing the  $n$ -vector  $a$ .
- inca** Increment for the array  $a$ :
- inca > 0**  $a$  is stored forward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-1) \times inca + 1)$ .
- inca < 0**  $a$  is stored backward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-n) \times inca + 1)$ .

Use **inca** = 1 if the vector  $a$  is stored contiguously in array **a**, i.e., if  $a_i$  is stored in **a**( $i$ ). Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

**alpha**                      The scalar  $\alpha$ .

## Output

**x**                              Array of length **lenx** =  $(n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ . If **n** = 0, then **x** is unchanged. Otherwise, the recurrence’s solution vector is returned.

## Notes

The result is unspecified if **a** and **x** overlap such that any element of  $a$  shares a memory location with any element of  $x$ .

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

**n** < 0,  
**incx** = 0, and  
**inca** = 0.

### Fortran Equivalent

```

SUBROUTINE FOLRC (N, X, INCX, A, INCA, ALPHA)
  INTEGER*8 N, INCX, INCA
  REAL*8 X(*), A(*), ALPHA
  IF ( N .LT. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCX .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCA .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IA = 1 + INCA
  IX = 1 + INCX
  IF ( INCA .LT. 0 ) IA = 1 - (N-2) * INCA
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  X(IX-INCX) = A(IA-INCA)
  DO 10 I = 2, N
    X(IX) = A(IA) + ALPHA * X(IX-INCX)
    IA = IA + INCA
    IX = IX + INCX
  10 CONTINUE
  RETURN
END

```

### Example

Solve the REAL\*8 first order linear recurrence

$$x_1 = a_1$$

$$x_i = a_i + 4x_{i-1}, \quad i = 2, 3, \dots, n,$$

where  $a$  and  $x$  are vectors 10 elements long stored in one-dimensional arrays A and X of dimension 20.

```

INTEGER*8 N, INCX, INCA
REAL*8 X(20), A(20), ALPHA
N = 10
INCX = 1
INCA = 1
ALPHA = 4.0
CALL FOLRC (N, X, INCX, A, INCA, ALPHA)

```

**NAME** FOLRN/FOLRNP – Last Term of First Order Linear Recurrence

**Purpose**

Given real vectors  $a$  and  $x$  of length  $n$ , these subprograms solve for the last term of the first order linear recurrence

$$y_1 = x_1$$

$$y_i = x_i \pm a_i y_{i-1}, \quad i = 2, 3, \dots, n,$$

i.e., returning  $y_n$ .

The operation indicated by  $\pm$  above is specified by the subprogram name used:

FOLRN       $\pm$  represents  $-$ :  $y_i = x_i - a_i y_{i-1}$   
 FOLRNP     $\pm$  represents  $+$ :  $y_i = x_i + a_i y_{i-1}$

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

INTEGER*8      n, inca, incx
REAL*8         yn, FOLRN, a(lena), x(lenx)
yn = FOLRN(n, a, inca, x, incx)
INTEGER*8      n, inca, incx
REAL*8         yn, FOLRNP, a(lena), x(lenx)
yn = FOLRNP(n, a, inca, x, incx)
  
```

**Input**

**n**            Number of elements of vectors  $a$  and  $x$  to be used in the recurrence,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference  $a$  or  $x$ .

**a**            Array of length  $\text{lena} = (n-1) \times |\text{inca}| + 1$  containing the  $n$ -vector  $a$ .

**inca**        Increment for the array  $a$ :

**inca** > 0       $a$  is stored forward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-1) \times \text{inca} + 1)$ .

**inca** < 0       $a$  is stored backward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-n) \times \text{inca} + 1)$ .

## FOLRN/FOLRNP – Last Term of First Order Linear Recurrence

Use  $\text{inca} = 1$  if the vector  $a$  is stored contiguously in array  $a$ , i.e., if  $a_i$  is stored in  $a(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

**x** Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx** Increment for the array  $x$ :

$\text{incx} \geq 0$   $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

$\text{incx} < 0$   $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

If  $\text{incx} = 0$ , then  $x_i = x(1)$  for all  $i$ . Refer to “Notes” to see how to use FOLRNP with  $\text{incx} = 0$  to evaluate a polynomial. Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

## Output

**yn** If  $n = 0$ , then  $\text{yn} = 0$ . Otherwise, the last term of the recurrence’s solution is returned.

## Notes

The result is unspecified if  $a$  and  $x$  overlap such that any element of  $a$  shares a memory location with any element of  $x$ .

FOLRNP may be used to evaluate a polynomial  $p(x) = \sum_{i=0}^n a_i x^i$  by recognizing that  $p(x)$  is the last term of the recurrence

$$y_0 = a_n$$

$$y_i = a_{n-i} + y_{i-1}x, \quad i = 1, 2, \dots, n.$$

Refer to “Example 2.”

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

$$\mathbf{n} < 0,$$

$$\mathbf{inca} = 0, \text{ and}$$

Linear Recurrences  
**FOLRN/FOLRNP – Last Term of First Order Linear Recurrence**

**incx = 0.**

**Fortran Equivalent**

```

REAL*8 FUNCTION FOLRN (N, A, INCA, X, INCX)
INTEGER*8 N, INCA, INCX
REAL*8 A(*), X(*)
FOLRN = 0.0
IF ( N .LT. 0 ) THEN
    CALL XERSCI (...)
    RETURN
END IF
IF ( INCA .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
IF ( INCX .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
END IF
IA = 1 + INCA
IX = 1 + INCX
IF ( INCA .LT. 0 ) IA = 1 - (N-2) * INCA
IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
FOLRN = X(IX-INCX)
DO 10 I = 2, N
    FOLRN = X(IX) - A(IA) * FOLRN
    IA = IA + INCA
    IX = IX + INCX
10 CONTINUE
RETURN
END

```

**Example 1**

Solve for the last term of the REAL\*8 first order linear recurrence

$$y_1 = x_1$$

$$y_i = x_i - a y_{i-1}, \quad i = 2, 3, \dots, n,$$

where  $a$  and  $x$  are vectors 10 elements long stored in one-dimensional arrays A and X of dimension 20.

```

INTEGER*8 N, INCA, INCX
REAL*8    YN, FOLRN, A(20), X(20)
N = 10
INCA = 1
INCX = 1
YN = FOLRN(N, A, INCA, X, INCX)

```

### Example 2

Evaluate the REAL\*8 polynomial  $p(x) = \sum_{i=0}^{10} a_i x^i$ , where  $a$  is a vector 11 elements long stored in one-dimensional array A of dimension 0:20.

```
INTEGER*8 N, INCA, INCX
REAL*8    P, FOLRNP, A(0:20), X
N = 11
INCA = -1
INCX = 0
P = FOLRNP(N, A, INCA, X, INCX)
```

**NAME** RECPP – Compute Vector of Partial Products

**Purpose**

Given real vector  $a$  of length  $n$ , this subprogram computes the  $n$ -vector  $x$  of partial products

$$x_i = \prod_{j=1}^i a_j, \quad i = 1, 2, \dots, n.$$

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

      INTEGER*8      n, incx, inca
      REAL*8        x(lenx), a(lena)
      CALL RECPP(n, x, incx, a, inca)
  
```

**Input**

- n** Number of elements of vectors  $a$  and  $x$ ,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference  $a$  or  $x$ .
- incx** Increment for the array  $x$ ,  $\text{incx} \neq 0$ :
- incx** > 0  $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .
- incx** < 0  $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .
- Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.
- a** Array of length  $\text{lena} = (n-1) \times |\text{inca}| + 1$  containing the  $n$ -vector  $a$ .
- inca** Increment for the array  $a$ :
- inca** > 0  $a$  is stored forward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-1) \times \text{inca} + 1)$ .
- inca** < 0  $a$  is stored backward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-n) \times \text{inca} + 1)$ .

Use **inca** = 1 if the vector  $a$  is stored contiguously in array **a**, i.e., if  $a_i$  is stored in **a**( $i$ ). Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

## Output

**x**            Array of length  $\text{lenx} = (\mathbf{n}-1) \times |\mathbf{incx}| + 1$  containing the  $n$ -vector  $x$ . If  $\mathbf{n} = 0$ , then **x** is unchanged. Otherwise, the vector of partial products replaces the input.

## Notes

The result is unspecified if **a** and **x** overlap such that any element of  $a$  shares a memory location with any element of  $x$ .

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

**n** < 0,  
**incx** = 0, and  
**inca** = 0.

## Linear Recurrences

### RECPP – Compute Vector of Partial Products

#### Fortran Equivalent

```
SUBROUTINE RECPP (N, X, INCX, A, INCA)
  INTEGER*8 N, INCX, INCA
  REAL*8 X(*), A(*)
  IF ( N .LT. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCX .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCA .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IA = 1 + INCA
  IX = 1 + INCX
  IF ( INCA .LT. 0 ) IA = 1 - (N-2) * INCA
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  X(IX-INCX) = A(IA-INCA)
  DO 10 I = 2, N
    X(IX) = X(IX-INCX) * A(IA)
    IA = IA + INCA
    IX = IX + INCX
  10 CONTINUE
  RETURN
END
```

#### Example

Compute the REAL\*8 vector of partial products of the vector  $a$ , where  $a$  is a vector 10 elements long stored in a one-dimensional array A of dimension 20. The result is stored in array X, also of dimension 20.

```
INTEGER*8 N, INCX, INCA
REAL*8 X(20), A(20)
N = 10
INCA = 1
INCX = 1
CALL RECPP (N, X, INCX, A, INCA)
```

**NAME** RECPS – Compute Vector of Partial Sums

**Purpose**

Given real vector  $a$  of length  $n$ , this subprogram computes the  $n$ -vector  $x$  of partial sums

$$x_i = \sum_{j=1}^i a_j, \quad i = 1, 2, \dots, n.$$

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

      INTEGER*8      n, incx, inca
      REAL*8        x(lenx), a(lena)
      CALL RECPS(n, x, incx, a, inca)
  
```

**Input**

- n**            Number of elements of vectors  $a$  and  $x$ ,  $n \geq 0$ . If  $n = 0$ , the subprograms do not reference  $a$  or  $x$ .
- incx**        Increment for the array  $x$ ,  $incx \neq 0$ :
- incx** > 0      $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times incx + 1)$ .
- incx** < 0      $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times incx + 1)$ .
- Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.
- a**            Array of length  $lena = (n-1) \times |inca| + 1$  containing the  $n$ -vector  $a$ .
- inca**        Increment for the array  $a$ :
- inca** > 0      $a$  is stored forward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-1) \times inca + 1)$ .
- inca** < 0      $a$  is stored backward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-n) \times inca + 1)$ .

**RECPS – Compute Vector of Partial Sums**

Use  $\mathbf{inca} = 1$  if the vector  $a$  is stored contiguously in array  $\mathbf{a}$ , i.e., if  $a_i$  is stored in  $\mathbf{a}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

**Output**

$\mathbf{x}$                       Array of length  $\mathbf{lenx} = (\mathbf{n}-1) \times |\mathbf{incx}| + 1$  containing the  $n$ -vector  $x$ . If  $\mathbf{n} = 0$ , then  $\mathbf{x}$  is unchanged. Otherwise, the vector of partial sums replaces the input.

**Notes**

The result is unspecified if  $\mathbf{a}$  and  $\mathbf{x}$  overlap such that any element of  $a$  shares a memory location with any element of  $x$ . If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

$\mathbf{n} < 0$ ,  
 $\mathbf{incx} = 0$ , and  
 $\mathbf{inca} = 0$ .

### Fortran Equivalent

```

SUBROUTINE RECPS (N, X, INCX, A, INCA)
INTEGER*8 N, INCX, INCA
REAL*8 X(*), A(*)
IF ( N .LT. 0 ) THEN
    CALL XERSCI (...)
    RETURN
END IF
IF ( INCX .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
END IF
IF ( INCA .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
END IF
IA = 1 + INCA
IX = 1 + INCX
IF ( INCA .LT. 0 ) IA = 1 - (N-2) * INCA
IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
X(IX-INCX) = A(IA-INCA)
DO 10 I = 2, N
    X(IX) = X(IX-INCX) + A(IA)
    IA = IA + INCA
    IX = IX + INCX
10 CONTINUE
RETURN
END

```

### Example

Compute the **REAL\*8** vector of partial sums of the vector  $a$ , where  $a$  is a vector 10 elements long stored in a one-dimensional array  $A$  of dimension 20. The result is stored in array  $X$ , also of dimension 20.

```

INTEGER*8 N, INCX, INCA
REAL*8    X(20), A(20)
N = 10
INCA = 1
INCX = 1
CALL RECPS (N, X, INCX, A, INCA)

```

**NAME** SOLR – Second Order Linear Recurrence

### Purpose

Given real vectors  $a$  and  $b$  of length  $n$  and the first two elements of  $n$ -vector  $x$ , this subprogram solves the second order linear recurrence

$$x_i = a_{i-2}x_{i-1} + b_{i-2}x_{i-2}, \quad i = 3, 4, \dots, n.$$

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

### Usage

SCILIB:

```

INTEGER*8      n, inca, incb, incx
REAL*8        a(lena), b(lenb), x(lenx)
CALL SOLR(n, a, inca, b, incb, x, incx)

```

### Input

- n** Number of elements of vectors  $a$ ,  $b$ , and  $x$  to be used in the recurrence,  $n \geq 0$ . If  $n = 0$ , the subprogram does not reference  $a$ ,  $b$ , or  $x$ .
- a** Array of length  $\text{lena} = (n-1) \times |\text{inca}| + 1$  containing the  $n$ -vector  $a$ .
- inca** Increment for the array  $a$ :
- inca**  $> 0$   $a$  is stored forward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-1) \times \text{inca} + 1)$ .
- inca**  $< 0$   $a$  is stored backward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-n) \times \text{inca} + 1)$ .
- Use **inca** = 1 if the vector  $a$  is stored contiguously in array  $a$ , i.e., if  $a_i$  is stored in  $a(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.
- b** Array of length  $\text{lenb} = (n-1) \times |\text{incb}| + 1$  containing the  $n$ -vector  $b$ .
- incb** Increment for the array  $b$ :
- incb**  $> 0$   $b$  is stored forward in array  $b$ , i.e.,  $b_i$  is stored in  $b((i-1) \times \text{incb} + 1)$ .
- incb**  $< 0$   $b$  is stored backward in array  $b$ , i.e.,  $b_i$  is stored in  $b((i-n) \times \text{incb} + 1)$ .

- Use  $\text{incb} = 1$  if the vector  $b$  is stored contiguously in array  $\mathbf{b}$ , i.e., if  $b_i$  is stored in  $\mathbf{b}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.
- x**            Array of length  $\text{lenx} = (\mathbf{n}-1) \times |\text{incx}| + 1$  containing the first two elements,  $x_1$  and  $x_2$  of the  $n$ -vector  $x$ .
- incx**        Increment for the array  $\mathbf{x}$ ,  $\text{incx} \neq 0$ :
- $\text{incx} > 0$       $x$  is stored forward in array  $\mathbf{x}$ , i.e.,  $x_i$  is stored in  $\mathbf{x}((i-1) \times \text{incx} + 1)$ .
- $\text{incx} < 0$       $x$  is stored backward in array  $\mathbf{x}$ , i.e.,  $x_i$  is stored in  $\mathbf{x}((i-\mathbf{n}) \times \text{incx} + 1)$ .
- Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $\mathbf{x}$ , i.e., if  $x_i$  is stored in  $\mathbf{x}(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

## Output

- x**            If  $\mathbf{n} = 0$ , then  $\mathbf{x}$  is unchanged. Otherwise, the recurrence’s solution vector replaces the input.

## Notes

The result is unspecified if  $\mathbf{a}$ ,  $\mathbf{b}$ , or  $\mathbf{x}$  overlap such that any element of  $a$ ,  $b$ , or  $x$  shares a memory location with any other element of  $a$ ,  $b$ , or  $y$ .

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

- $\mathbf{n} < 0$ ,
- $\text{inca} = 0$ ,
- $\text{incb} = 0$ , and
- $\text{incx} = 0$ .

## Linear Recurrences

### SOLR – Second Order Linear Recurrence

#### Fortran Equivalent

```
SUBROUTINE SOLR (N, A, INCA, B, INCB, X, INCX)
  INTEGER*8 N, INCA, INCB, INCX
  REAL*8 A(*), B(*), X(*)
  IF ( N .LT. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCA .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCB .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCX .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IA = 1
  IB = 1
  IX = 1 + INCX
  IF ( INCA .LT. 0 ) IA = 1 - (N-3) * INCA
  IF ( INCB .LT. 0 ) IB = 1 - (N-3) * INCB
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  DO 10 I = 3, N
    X(IX+INCX) = A(IA) * X(IX) + B(IB) * X(IX-INCX)
    IA = IA + INCA
    IB = IB + INCB
    IX = IX + INCX
  10 CONTINUE
  RETURN
END
```

#### Example

Solve the REAL\*8 second order linear recurrence

$$x_i = a_{i-2}x_{i-1} + b_{i-2}x_{i-2}, \quad i = 3, 4, \dots, n,$$

where  $a$ ,  $b$ , and  $x$  are vectors 10 elements long stored in one-dimensional arrays  $A$ ,  $B$ , and  $X$  of dimension 20.

```
INTEGER*8 N, INCA, INCB, INCX
REAL*8 A(20), B(20), X(20)
N = 10
INCA = 1
INCB = 1
INCX = 1
X(1) = ...
X(2) = ...
CALL SOLR (N, A, INCA, B, INCB, X, INCX)
```

**NAME** SOLR3 – Second Order Linear Recurrence

**Purpose**

Given real vectors  $a$ ,  $b$ , and  $x$  of length  $n$ , this subprogram solves the second order linear recurrence

$$y_1 = x_1$$

$$y_2 = x_2$$

$$y_i = x_i + a_{i-2}y_{i-1} + b_{i-2}y_{i-2}, \quad i = 3, 4, \dots, n.$$

overwriting the input  $x$  vector with the resulting  $y$  vector.

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

INTEGER*8      n, inca, incb, incx
REAL*8        a(lena), b(lenb), x(lenx)
CALL SOLR3(n, a, inca, b, incb, x, incx)

```

**Input**

- n**                    Number of elements of vectors  $a$ ,  $b$ , and  $x$  to be used in the recurrence,  $n \geq 0$ . If  $n = 0$ , the subprogram does not reference  $a$ ,  $b$ , or  $x$ .
- a**                    Array of length  $lena = (n-1) \times |inca| + 1$  containing the  $n$ -vector  $a$ .
- inca**                Increment for the array  $a$ :  
                        $inca > 0$       $a$  is stored forward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-1) \times inca + 1)$ .  
                        $inca < 0$       $a$  is stored backward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-n) \times inca + 1)$ .  
                       Use  $inca = 1$  if the vector  $a$  is stored contiguously in array  $a$ , i.e., if  $a_i$  is stored in  $a(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.
- b**                    Array of length  $lenb = (n-1) \times |incb| + 1$  containing the  $n$ -vector  $b$ .
- incb**                Increment for the array  $b$ :

**incb** > 0       $b$  is stored forward in array  $b$ , i.e.,  $b_i$  is stored in  $b((i-1) \times \text{incb} + 1)$ .

**incb** < 0       $b$  is stored backward in array  $b$ , i.e.,  $b_i$  is stored in  $b((i-n) \times \text{incb} + 1)$ .

Use **incb** = 1 if the vector  $b$  is stored contiguously in array  $b$ , i.e., if  $b_i$  is stored in  $b(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

**x**      Array of length  $\text{lenx} = (n-1) \times |\text{incx}| + 1$  containing the  $n$ -vector  $x$ .

**incx**      Increment for the array  $x$ , **incx**  $\neq$  0:

**incx** > 0       $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1) \times \text{incx} + 1)$ .

**incx** < 0       $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n) \times \text{incx} + 1)$ .

Use **incx** = 1 if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

## Output

**x**      If  $n = 0$ , then  $x$  is unchanged. Otherwise, the recurrence’s solution vector replaces the input.

## Notes

The result is unspecified if  $a$ ,  $b$ , or  $x$  overlap such that any element of  $a$ ,  $b$ , or  $x$  shares a memory location with any other element of  $a$ ,  $b$ , or  $x$ .

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

**n** < 0,  
**inca** = 0,  
**incb** = 0, and  
**incx** = 0.

## Fortran Equivalent

```

SUBROUTINE SOLR3 (N, A, INCA, B, INCB, X, INCX)
  INTEGER*8 N, INCA, INCB, INCX
  REAL*8 A(*), B(*), X(*)
  IF ( N .LT. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCA .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCB .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IF ( INCX .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
  END IF
  IA = 1
  IB = 1
  IX = 1 + INCX
  IF ( INCA .LT. 0 ) IA = 1 - (N-3) * INCA
  IF ( INCB .LT. 0 ) IB = 1 - (N-3) * INCB
  IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
  DO 10 I = 3, N
    X(IX+INCX) = X(IX+INCX) + A(IA) * X(IX) + B(IB) * X(IX-INCX)
    IA = IA + INCA
    IB = IB + INCB
    IX = IX + INCX
  10 CONTINUE
  RETURN
END

```

## Example

Solve the REAL\*8 second order linear recurrence

$$y_1 = x_1$$

$$y_2 = x_2$$

$$y_i = x_i + a_{i-2}y_{i-1} + b_{i-2}y_{i-2}, \quad i = 3, 4, \dots, n.$$

where  $a$ ,  $b$ , and  $x$  are vectors 10 elements long stored in one-dimensional arrays A, B, and X of dimension 20, and  $y$  overwrites  $x$ .

```

INTEGER*8 N, INCA, INCB, INCX
REAL*8 A(20), B(20), X(20)
N = 10
INCA = 1
INCB = 1
INCX = 1
CALL SOLR3 (N, A, INCA, B, INCB, X, INCX)

```

**NAME** SOLRN – Last Term of Second Order Linear Recurrence

**Purpose**

Given real vectors  $a$  and  $b$  of length  $n$  and the first two elements of  $n$ -vector  $x$ , this subprogram solves for the last term of the second order linear recurrence

$$x_i = a_{i-2}x_{i-1} + b_{i-2}x_{i-2}, \quad i = 3, 4, \dots, n,$$

i.e., returning  $x_n$ .

The vectors may be stored in one-dimensional arrays or in either rows or columns of two-dimensional arrays, and the indexing through the arrays may be either forward or backward.

**Usage**

SCILIB:

```

INTEGER*8      n, inca, incb, incx
REAL*8        xn, SOLRN, a(lena), b(lenb), x(lenx)
xn = SOLRN(n, a, inca, b, incb, x, incx)

```

**Input**

- n** Number of elements of vectors  $a$ ,  $b$ , and  $x$  to be used in the recurrence,  $n \geq 0$ . If  $n = 0$ , the subprogram does not reference  $a$ ,  $b$ , or  $x$ .
- a** Array of length  $\text{lena} = (n-1) \times |\text{inca}| + 1$  containing the  $n$ -vector  $a$ .
- inca** Increment for the array  $a$ :  
**inca** > 0  $a$  is stored forward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-1) \times \text{inca} + 1)$ .  
**inca** < 0  $a$  is stored backward in array  $a$ , i.e.,  $a_i$  is stored in  $a((i-n) \times \text{inca} + 1)$ .
- Use **inca** = 1 if the vector  $a$  is stored contiguously in array  $a$ , i.e., if  $a_i$  is stored in  $a(i)$ . Refer to "BLAS Indexing Conventions" in the introduction to Chapter 2.
- b** Array of length  $\text{lenb} = (n-1) \times |\text{incb}| + 1$  containing the  $n$ -vector  $b$ .
- incb** Increment for the array  $b$ :  
**incb** > 0  $b$  is stored forward in array  $b$ , i.e.,  $b_i$  is stored in  $b((i-1) \times \text{incb} + 1)$ .

## SOLRN – Last Term of Second Order Linear Recurrence

$\text{incb} < 0$       $b$  is stored backward in array  $b$ , i.e.,  $b_i$  is stored in  $b((i-n)\times\text{incb}+1)$ .

Use  $\text{incb} = 1$  if the vector  $b$  is stored contiguously in array  $b$ , i.e., if  $b_i$  is stored in  $b(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

$x$      Array of length  $\text{lenx} = (n-1)\times|\text{incx}|+1$  containing the first two elements,  $x_1$  and  $x_2$  of the  $n$ -vector  $x$ .

$\text{incx}$      Increment for the array  $x$ ,  $\text{incx} \neq 0$ :

$\text{incx} > 0$       $x$  is stored forward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-1)\times\text{incx}+1)$ .

$\text{incx} < 0$       $x$  is stored backward in array  $x$ , i.e.,  $x_i$  is stored in  $x((i-n)\times\text{incx}+1)$ .

Use  $\text{incx} = 1$  if the vector  $x$  is stored contiguously in array  $x$ , i.e., if  $x_i$  is stored in  $x(i)$ . Refer to “BLAS Indexing Conventions” in the introduction to Chapter 2.

## Output

$xn$      If  $n = 0$ , then  $xn = 0$ . Otherwise, the last term of the recurrence’s solution is returned.

$x$      If  $n = 0$ , then  $x$  is unchanged. Otherwise,  $x$  is overwritten with intermediate results. These intermediate results are not necessarily what is shown in “Fortran Equivalent”.

## Notes

The result is unspecified if  $a$ ,  $b$ , or  $x$  overlap such that any element of  $a$ ,  $b$ , or  $x$  shares a memory location with any other element of  $a$ ,  $b$ , or  $y$ .

If an error in the arguments is detected, the subprograms call error handler XERSCI, which writes an error message onto the standard error file and terminates execution. The standard version of XERSCI (refer to Chapter 9, Miscellaneous Routines) may be replaced with a user-supplied version to change the error procedure. Error conditions are:

$n < 0$ ,  
 $\text{inca} = 0$ ,  
 $\text{incb} = 0$ , and  
 $\text{incx} = 0$ .

## Linear Recurrences

### SOLRN – Last Term of Second Order Linear Recurrence

#### Fortran Equivalent

```
REAL*8 FUNCTION SOLRN (N, A, INCA, B, INCB, X, INCX)
INTEGER*8 N, INCA, INCB, INCX
REAL*8 A(*), B(*), X(*)
SOLRN = 0.0
IF ( N .LT. 0 ) THEN
    CALL XERSCI (...)
    RETURN
END IF
IF ( INCA .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
END IF
IF ( INCB .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
END IF
IF ( INCX .EQ. 0 ) THEN
    CALL XERSCI (...)
    RETURN
END IF
IA = 1
IB = 1
IX = 1 + INCX
IF ( INCA .LT. 0 ) IA = 1 - (N-3) * INCA
IF ( INCB .LT. 0 ) IB = 1 - (N-3) * INCB
IF ( INCX .LT. 0 ) IX = 1 - (N-2) * INCX
DO 10 I = 3, N      ! CAUTION: X NOT NECESSARILY RETURNED AS SHOWN
    X(IX+INCX) = A(IA) * X(IX) + B(IB) * X(IX-INCX)
    IA = IA + INCA
    IB = IB + INCB
    IX = IX + INCX
10 CONTINUE
IF ( N .EQ. 1 ) THEN
    SOLRN = X(IX-INCX)
ELSE
    SOLRN = X(IX)
END IF
RETURN
END
```

#### Example

Solve for the last term of the REAL\*8 second order linear recurrence

$$x_i = a_{i-2}x_{i-1} + b_{i-2}x_{i-2}, \quad i = 3, 4, \dots, n,$$

where  $a$ ,  $b$ , and  $x$  are vectors 10 elements long stored in one-dimensional arrays A, B, and X of dimension 20.

**SOLRN – Last Term of Second Order Linear Recurrence**

```
INTEGER*8 N, INCA, INCB, INCX
REAL*8    XN, SOLRN, A(20), B(20), X(20)
N = 10
INCA = 1
INCB = 1
INCX = 1
X(1) = ...
X(2) = ...
XN = SOLRN (N, A, INCA, B, INCB, X, INCX)
```

Linear Recurrences

**SOLRN – Last Term of Second Order Linear Recurrence**

## 9 Miscellaneous Routines

---

### Overview

This chapter explains how to use SCILIB subprograms for operations that do not fit in the categories covered by other chapters. It includes a description of subprograms that report errors detected in the usage of SCILIB subprograms.

---

### Chapter Objectives

After reading this chapter you will

- know how to use the described subprogram
  - know how to change the method of error reporting in SCILIB subprograms
- 

### What You Need to Know to Use These Subprograms

Subprograms described in this chapter do not normally perform vector operations.

Miscellaneous Routines  
**XERSCI – SCILIB Error Handler**

**NAME** XERSCI – SCILIB Error Handler

**Purpose**

XERSCI is the error handler for many of the subprograms in the SCILIB library, as indicated in the “Notes” section in the subprogram descriptions. As supplied in SCILIB, XERSCI writes one of the following error messages onto the standard error file:

```
*****  
* XERSCI: subprogram name called with invalid value of argument number iarg  
*  
*****
```

or

```
*****  
* XERSCI: error detected by subprogram name: text of error message *  
*****
```

or

```
*****  
* XERSCI: error iarg detected by subprogram name: text of error message *  
*****
```

where *name* is the name of the subprogram in which the error was detected, *iarg* is the argument number of the offending argument, and *text of error message* is a character string. If the main program is in Fortran, a call traceback is also written onto the standard error file. XERSCI then terminates execution with a nonzero exit status.

You may supply a version of XERSCI that alters this action. All SCILIB subprograms that call XERSCI have a RETURN statement after the CALL statement, so your version could perhaps set a flag in a common block and RETURN. The flag could be tested in the program unit that calls the SCILIB subprogram.

**Usage**

SCILIB:

```
CHARACTER*(*)    name, messag  
INTEGER*8        iarg  
CALL XERSCI(name, iarg, messag)
```

**Input**

<b>name</b>	The name of the subprogram in which the error was detected.
<b>iarg</b>	If <i>iarg</i> > 0, the error message is printed in the first form given above and <i>iarg</i> is the number of the argument that was

found to be in error. If *iarg* = 0, the error message is printed in the second form given above and *iarg* is not part of the error message. If *iarg* < 0, the error message is printed in the third form given above and *iarg* is the error number.

**messag**

The text of the error message to be printed if *iarg* ≤ 0. Not used as input if *iarg* > 0.

Miscellaneous Routines  
**XERSCI – SCILIB Error Handler**

---

# Index

---

## A

accessing SCILIB 2  
arithmetic format 8  
ASAP, automatic self allocating processors 5  
automatic self allocating processors (ASAP) 5

---

## B

backward storage 15  
BAKVEC 298  
BALANC 298  
BALBAK 298  
band matrix 114, 137, 176, 181, 225, 229, 232, 236,  
252, 254, 258, 261, 265, 279  
BANDR 298  
BANDV 298  
Basic Linear Algebra Subprograms 13  
bibliography xvii  
BISECT 298  
BLAS 13, 212  
BLAS indexing conventions 15  
BLAS, Extended 97, 215, 283  
BLAS, Level 2 97, 215, 283  
BLAS, Level 3 97, 215, 283  
BQR 298

---

## C

C, calling SCILIB from 1  
calling SCILIB from C 1  
CAXPY 55  
CBABK2 298  
CBAL 298  
CCHDC 281  
CCHDD 281  
CCHEX 281  
CCHUD 281  
CCOPY 60  
CDOTC 63  
CDOTU 63  
-cfc compiler option 9  
CFFT2 303  
CFFTMLT 306  
CFTFAX 306  
CG 298  
CGBCO 225  
CGBDI 229

CGBFA 232  
CGBMV 114  
CGBSL 236  
CGECO 239  
CGEDI 242  
CGEFA 246  
CGEMM 119  
CGEMMS 4, 123  
CGEMV 128  
CGERC 132  
CGERU 132  
CGESL 249  
CGTSL 252  
CH 298  
CHBMV 137  
CHEMM 154  
CHEMV 158  
CHER 161  
CHER2 164  
CHER2K 168  
CHERK 172  
CHICO 281  
CHIDI 281  
CHIFA 281  
CHISL 281  
Cholesky factorization 254, 261, 267, 274  
CHPCO 281  
CHPDI 281  
CHPFA 281  
CHPMV 142  
CHPR 146  
CHPR2 150  
CHPSL 281  
CINVT 298  
CLUSEQ 18  
CLUSFGE 18  
CLUSFGT 18  
CLUSFLE 18  
CLUSFLT 18  
CLUSIGE 18  
CLUSIGT 18  
CLUSILE 18  
CLUSILT 18  
CLUSNE 18  
cluster 18  
COMBAK 298  
COMHES 298  
COMLR 298  
COMLR2 298  
COMQR 298  
COMQR2 298  
condition number 218, 225, 239, 254, 267

ConvexMLIB Man Pages 9  
convolution 324, 327  
correlation 324, 327  
correlation and convolution subprograms 323  
CORTB 298  
CORTH 298  
count vector elements 23, 25, 27  
CPBCO 254  
CPBDI 258  
CPBFA 261  
CPBSL 265  
CPOCO 267  
CPODI 270  
CPOFA 274  
CPOSL 277  
CPPCO 281  
CPPDI 281  
CPPFA 281  
CPPSL 281  
CPTSL 279  
CQRDC 281  
CQRSL 281  
Cray SCILIB 1  
CRFFT2 311  
CROT 72  
CROTG 75  
CSCAL 83  
CSICO 281  
CSIDI 281  
CSIFA 281  
CSISL 281  
CSPCO 281  
CSPDI 281  
CSPFA 281  
CSPSL 281  
CSSCAL 83  
CSUM 85  
CSVDC 281  
CSWAP 87  
CSYMM 154  
CSYR2K 168  
CSYRK 172  
CTBMV 176  
CTBSV 181  
CTPMV 187  
CTPSV 191  
CTRCO 281  
CTRDI 281  
CTRMM 195  
CTRMV 199  
CTRSL 281  
CTRSM 202  
CTRSV 206  
CXpa 7

---

## D

determinant 219, 229, 242, 258, 270  
DFT 301  
discrete Fourier transform 301  
documentation, online 9  
documentation, ordering xix  
dot product 63, 70

---

## E

eigenvalues 283, 285, 288, 291, 294, 296  
eigenvectors 283, 285, 288, 296  
EISPACK 283  
EISPACK Guide 11  
EISPACK Guide Extension 11  
ELMBAK 298  
ELMHES 298  
ELTRAN 298  
error handler 212  
error handling, SCILIB 9  
Euclidean norm 66  
Extended BLAS 97, 215, 283

---

## F

fast Fourier transforms 301  
FFT, complex-to-real 311  
FFT, one-dimensional 303, 311, 314  
FFT, one-dimensional, simultaneous 306, 317  
FFT, real-to-complex 314, 317  
FFTFAX 317  
FIGI 298  
FIGI2 298  
FILTERG 324  
filtering 324, 327  
FILTERS 327  
find 18, 42, 45, 47, 50, 90, 94  
first order recurrence 332, 335, 339, 342  
floating-point format 8  
FOLR 332  
FOLR2 335  
FOLR2P 335  
FOLRC 339  
FOLRN 342  
FOLRNP 342  
FOLRP 332  
Fortran array argument association 14  
Fortran storage of arrays 14  
forward storage 15  
further reference xvii

---

## G

GATHER 21  
gather 21  
Givens rotation 72, 75  
Givens rotation, modified 77, 80

---

## H

HQR 298  
HQR2 298  
HTRIB3 298  
HTRIBK 298  
HTRID3 298  
HTRIDI 298

---

## I

ICAMAX 4, 33  
IEEE arithmetic format 8  
IILZ 23  
ILLZ 25  
ILSUM 27  
IMTQL1 298  
IMTQL2 298  
IMTQLV 298  
increment 15  
increment arguments 16  
INCX 15  
index 42, 45, 47, 50, 90, 94  
index of maximum 29, 38  
index of maximum absolute value 33  
index of minimum 31, 40  
index of minimum absolute value 36  
INFLMAX 29  
INFLMIN 31  
inner product 63, 70  
INTMAX 38  
INTMIN 40  
inverse 219, 220, 242, 270  
INVIT 298  
ISAMAX 4, 33  
ISAMIN 4, 36  
ISEARCH 42  
ISMAX 4, 38  
ISMIN 4, 40  
ISRCHEQ 42  
ISRCHFGE 42  
ISRCHFGT 42  
ISRCHFLE 42  
ISRCHFLT 42  
ISRCHIGE 42  
ISRCHIGT 42

ISRCHILE 42  
ISRCHILT 42  
ISRCHMEQ 45  
ISRCHMGE 45  
ISRCHMGT 45  
ISRCHMLE 45  
ISRCHMLT 45  
ISRCHMNE 45  
ISRCHNE 42

---

## L

LAPACK 3, 4, 212  
Level 2 BLAS 97, 212, 215, 283  
Level 3 BLAS 97, 212, 215, 283  
linear equations 215, 220, 223  
LINPACK 215  
LINPACK Users' Guide 11  
-llapack8 compiler option 4  
locate 18, 23, 25, 42, 45, 47, 50, 90, 94  
-lscilib compiler option 3  
LU factorization 225, 232, 239, 246  
-lveclib8 compiler option 3

---

## M

man pages 9  
matrix inverse 219, 220, 242, 270  
matrix-matrix multiply 101, 103, 119, 123, 154  
matrix-matrix multiply, triangular 195  
matrix-vector multiply 108, 110, 114, 128, 137, 142,  
158, 176, 187, 199  
matrix-vector multiply and add 135  
maximum 29, 33, 38  
MINFIT 298  
minimum 31, 36, 40  
MINV 220  
modified Givens rotation 77, 80  
MXM 101  
MXMA 103  
MXV 108  
MXVA 110

---

## N

native arithmetic format 8  
negative 25  
nonzeros 23  
norm 53, 66

---

---

## O

online documentation 9  
OPFILT 223  
ordered vector 47, 50  
ordering documentation xix  
ORTBAK 298  
ORTHES 298  
ORTRAN 298  
OSRCHF 47  
OSRCHI 47  
OSRCHM 50

---

## P

-p8 compiler option 9  
packed matrix 142, 146, 150, 187, 191  
parallel processing 5  
partial products 346  
partial sums 349  
-pd8 compiler option 9  
performance analysis 7  
polynomial, evaluate 343  
positive 25  
positive definite matrix 254, 258, 261, 265, 267, 270,  
274, 277, 279  
products, partial 346  
profiling 7  
programmer's reference 9

---

## Q

QZHES 298  
QZIT 298  
QZVAL 298  
QZVEC 298

---

## R

rank-2k update 168  
rank-k update 172  
rank-one update 132, 146, 161  
rank-two update 150, 164  
RATQR 298  
RCFFT2 314  
REBAK 298  
REBAKB 298  
RECPP 346  
RECPS 349  
recurrence, first order 332, 335, 339, 342  
recurrence, second order 352, 355, 358  
REDUC 298

REDUC2 298  
RFFTMLT 317  
RG 298  
RGG 298  
RS 285  
RSB 298  
RSG 298  
RSGAB 298  
RSGBA 298  
RSM 298  
RSP 298  
RST 298  
RT 298

---

## S

SASUM 53  
SAXPY 55  
SCASUM 53  
SCATTER 58  
scatter 58  
SCHDC 281  
SCHDD 281  
SCHEX 281  
SCHUD 281  
SCILIB 1, 4  
SCILIB error handling 9  
SCILIB man pages 9  
SCNRM2 66  
SCOPY 60  
SDOT 63  
search vector 18, 42, 45, 47, 50, 90, 94  
second order recurrence 352, 355, 358  
SGBCO 225  
SGBDI 229  
SGBFA 232  
SGBMV 114  
SGBSL 236  
SGECO 239  
SGEDI 242  
SGEFA 246  
SGEMM 119  
SGEMMS 4, 123  
SGEMV 128  
SGER 132  
SGESL 249  
SGTSL 252  
SMXPY 135  
SNRM2 66  
SOLR 352  
SOLR3 355  
SOLRN 358  
sparse 18, 21, 58, 68, 70, 90, 94  
SPAXPY 68  
SPBCO 254  
SPBDI 258

SPBFA 261  
SPBSL 265  
SPDOT 70  
SPOCO 267  
SPODI 270  
SPOFA 274  
SPOSL 277  
SPPCO 281  
SPPDI 281  
SPPFA 281  
SPPSL 281  
SPTSLS 279  
SQRDC 281  
SQRSL 281  
SROT 72  
SROTG 75  
SROTM 77  
SROTMG 80  
SSBMV 137  
SSCAL 83  
SSICO 281  
SSIDI 281  
SSIFA 281  
SSISL 281  
SSPCO 281  
SSPDI 281  
SSPFA 281  
SSPMV 142  
SSPR 146  
SSPR2 150  
SSPSL 281  
SSUM 85  
SSVDC 281  
SSWAP 87  
SSYMM 154  
SSYMV 158  
SSYR 161  
SSYR2 164  
SSYR2K 168  
SSYRK 172  
standardization 2  
STBMV 176  
STBSV 181  
storage, backward 15  
storage, forward 15  
STPMV 187  
STPSV 191  
STRCO 281  
STRDI 281  
stride 15  
stride arguments 16  
STRMM 195  
STRMV 199  
STRSL 281  
STRSM 202  
STRSV 206  
sums, partial 349

supplemental reading xvii, 11, 17, 99, 219, 284, 302,  
323  
SVD 298  
SXMPY 210

---

## T

TAC, technical assistance center xix  
Technical Assistance Center 10  
technical assistance center, TAC xix  
thread, definition 5  
TINVIT 298  
Toeplitz matrix 223  
TQL1 298  
TQL2 288  
TQLRAT 291  
TRBAK1 298  
TRBAK3 298  
TRED1 294  
TRED2 296  
TRED3 298  
triangular factorization 225, 232, 239, 246  
triangular matrix-matrix multiply 195  
triangular solve 181, 191, 202, 206  
tridiagonal linear equations 252, 279  
TRIDIB 298  
TSTURM 298

---

## U

UNICOS Math and Scientific Library 1

---

## V

VECLIB 3, 4, 212  
vector-matrix multiply and add 210

---

## W

Weiner-Levinson algorithm 223  
WHENEQ 90  
WHENFGE 90  
WHENFGT 90  
WHENFLE 90  
WHENFLT 90  
WHENIGE 90  
WHENIGT 90  
WHENILE 90  
WHENILT 90  
WHENMEQ 94  
WHENMGE 94

WHENMGT 94  
WHENMLE 94  
WHENMLT 94  
WHENMNE 94  
WHENNE 90

---

## X

XERBLA 9, 212  
XERSCI 9, 364

---

## Z

zeros 23









CONVEX  
PRESS

B5649-90002

