LAPACK++ V. 1.0

High Performance Linear Algebra
Users' Guide

April 1994

Jack Dongarra
Roldan Pozo
David Walker

Oak Ridge National Laboratory
University of Tennessee, Knoxville

Contents

1 Introduction **3**

2 Overview **5**
2.1 Contents of LAPACK++ . . . 5
2.2 A simple code example . . . 6
2.3 Basic Linear Algebra Subroutines (BLAS++) . . . 6
2.4 Performance . . . 7

3 LAPACK++ Matrix Objects **9**
3.1 Fundamental Matrix Operations . . . 9
3.2 General Rectangular Matrices . . . 10
3.2.1 Declarations . . . 10
3.3 Triangular Matrices . . . 12
3.3.1 Symmetric, Hermitian and SPD Matrices . . . 12
3.4 Banded Matrices . . . 12
3.4.1 Triangular Banded Matrices . . . 13
3.4.2 Symmetric and Hermitian Banded Matrices . . . 14
3.5 Tridiagonal Matrices . . . 14

4 Driver Routines **15**
4.1 Linear Equations . . . 15
4.2 Eigenvalue Problems . . . 18
4.3 Memory Optimizations: Factorizing in place . . . 18

5 Factorization Classes **19**
5.1 Optimizations: Factoring in Place . . . 19

A Programming Examples **21**
A.1 Polynomial data fitting . . . 21

B Release Notes for v. 1.0 **23**
B.1 How to report bugs . . . 23
B.2 Tips and Suggestions . . . 23
B.3 Reducing Compilation time . . . 23
B.3.1 Minimize the number of included header files . . . 23
B.3.2 Header file madness . . . 23
B.4 Performance Considerations . . . 23
B.4.1 Array Bounds Checking . . . 23
B.4.2 Improving $A(i, j)$ efficiency . . . 24
B.5 Exception Handling . . . 24

B.6 Frequently asked questions 24
B.6.1 What is the performance of LAPACK++? 24
B.6.2 Do I need a Fortran compiler to use LAPACK++? 24
B.6.3 Why are the LAPACK++ matrix classes not templated? 24
B.6.4 Can LAPACK++ work with built-in C/C++ arrays and other matrix classes? 25
B.6.5 Can the matrix and vector objects be used independently of BLAS++ and LAPACK++? 25
B.6.6 Is there an existing standard for C++ matrix classes? 25

C Direct Interfaces to LAPACK and BLAS Routines **27**

Chapter 1

Introduction

LAPACK++ is an object-oriented C++ extension to the LAPACK [1] library for numerical linear algebra. This package includes state-of-the-art numerical algorithms for the more common linear algebra problems encountered in scientific and engineering applications: solving linear equations, linear least squares, and eigenvalue problems for dense and banded systems.

Traditionally, such libraries have been available only in Fortran; however, with an increasing number of programmers using C and C++ for scientific software development, there is a need to have high-quality numerical libraries to support these platforms as well. LAPACK++ provides the speed and efficiency competitive with native Fortran codes (see section 2.4), while allowing programmers to capitalize on the software engineering benefits of object oriented programming.

The overall design of LAPACK++ includes support for distributed and shared memory architectures [6] . Version 1.0 includes support only for uniprocessor and shared memory platforms. Distributed memory architectures will be supported in Version 2.0.

Replacing the Fortran 77 interface of LAPACK with an object-oriented framework simplifies the coding style and allows for a more flexible and extendible software platform. The design goals of LAPACK++ include

- Maintain performance competitive with Fortran.
- Provide a simple interface that hides implementation details of various matrix storage schemes and their corresponding factorization structures.
- Provide a universal interface and open system design for integration into user-defined data structures and third-party matrix packages.
- Replace static work array limitations of Fortran with more flexible and type-safe dynamic memory allocation schemes.
- Provide an efficient indexing scheme for matrix elements that has minimal overhead and can be optimized for in most application code loops.
- Utilize function and operator overloading in C++ to simplify and reduce the number of interface entry points to LAPACK.
- Utilize exception error handling in C++ for intelligent managing of error situations without cluttering up application codes.
- Provide the capability to access submatrices by **reference**, rather than by value, and perform factorizations "in place". This is vital for implementing blocked algorithms efficiently.

- Provide more meaningful naming conventions for variables and function names. (Names no longer limited to six alphanumeric characters.)

LAPACK++ also provides an object-oriented interface to the Basic Linear Algebra Subprograms (BLAS) [4], [8] (see section 2.3), allowing programmers to utilize these optimized computational kernels in their own C++ applications.

Chapter 2

Overview

The underlying philosophy of the LAPACK++ is to provide an interface which is relatively simple, yet powerful enough to express all complex and subtle tasks within LAPACK, including those which optimize performance and/or storage. Following the framework of LAPACK, the C++ extension contains **driver routines** for solving standard types of problems, **computational routines** to perform a distinct computational task, and **auxiliary routines** to perform a certain subtask or common low-level computation. Each driver routine typically calls a sequence of computational routines. Taken as a whole, the computational routines can perform a wider range of tasks than are covered by the driver routines. Currently, dense and band matrices are supported. General sparse matrices are handled in [7].

2.1 Contents of LAPACK++

With over 1,000 subroutines in the original f77 LAPACK, not every routine is implemented in LAPACK++. Instead, source code examples in the various major areas are provided, allowing users to easily extend the package for their particular needs. LAPACK++ provides source code for

- Algorithms
  - LU Factorization
  - Cholesky (LL^T) Factorization
  - QR Factorization
  - Eigenvalue problems
- Storage Classes
  - rectangular matrices
  - symmetric and symmetric positive definite (SPD)
  - banded matrices
  - tri/bidiagonal matrices
- Element Data Types
  - int, long int, float, double, (double precision) complex,
  - arbitrary `Vector` data types via templates (section B.6.3)

2.2 A simple code example

To provide a first glimpse at how LAPACK++ simplifies the user interface, this section presents a few simple code fragments. The examples are incomplete and are meant to merely illustrate the interface style. The next few sections will further discuss the details of matrix classes and their operations.

The first example illustrates a code fragment to solve a linear system $Ax = b$ using LU factorization:

```
#include <lapack++.h>                          // 1

LaGenMatDouble A(N,N);                         // 2
LaVectorDouble x(N), b(N);                     // 3
...

LaLinSolve(A,x,b);                             // 4
```

Line (1) includes all of the LAPACK++ object and function declarations. Line (2) declares A to be a square $N \times N$ coefficient matrix, while line (3) declares the right-hand-side and solution vectors. Finally, the `LaLinSolve()` function in line (4) calls the underlying LAPACK driver routine `SGESV()` for solving linear equations.

Consider now solving a similar system with a tridiagonal coefficient matrix:

```
#include <lapack++.h>

LaSPDMatDouble A(N,N);
LaVectorDouble x(N), b(N);
...

LaLinSolve(A,x,b);
```

The only code modfication is in the declaration of A. In this case `LaLinSolve()` calls the Cholesky driver routine for solving symmetric, positive-definite linear systems. The `LaLinSolve()` function has been **overloaded** to perform different tasks depending on the type of the input matrix A. If the matrix types are known at compile time, as in this example, then there is no runtime overhead associated with this.

2.3 Basic Linear Algebra Subroutines (BLAS++)

The Basic Linear Algebra Subprograms (BLAS) [4] has been the key to obtaining good performance on a wide variety of computer architectures. The BLAS define a common interface for low-level operations often found in computational kernels. These operations, such as matrix/matrix multiply and triangular solves, typically comprise of most of the computatioal workload found in dense and banded linear algebra algorithms. The Level 3 BLAS obtains good performance on a wide variety of architectures by keeping data used most often in the closest level of memory hierarhcy (registers, cache, etc.).

The BLAS++ interface simplifies many of the calling sequences to the traditional f77 BLAS interface, by using the LAPACK++ matrix classes. These routines are called within the LAPACK++ algorithms, or can be called directly at the user-level within applications.

There are two levels of the BLAS++ interface. The first is a direct interface, as shown in Table C, are essentially inlined to call BLAS directly in eliminating any overhead. The other, more elegant interface, overloads the binary opertors * and + for simple expressions such as `C=A*B`. Having these two interfaces gives the users the choice between simplicity and performance in their application codes. See Appendix B for a list of BLAS++ interface routines.

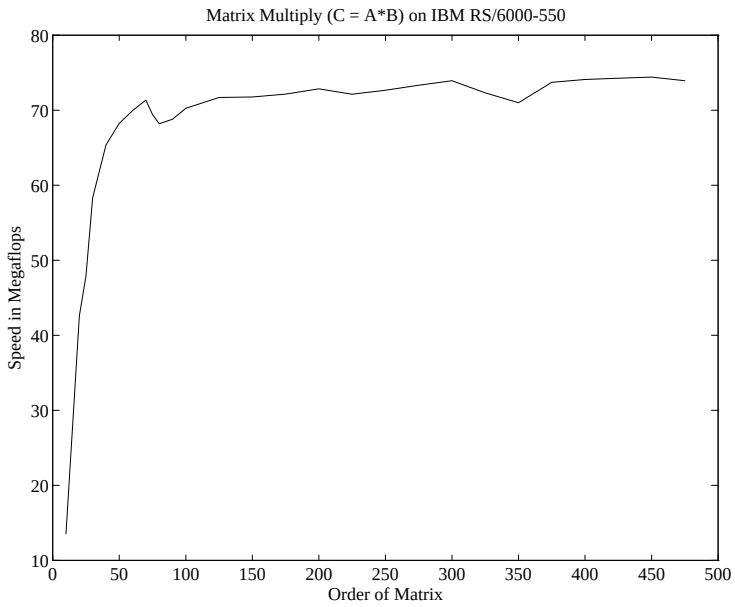

Figure 2.1: Performance of matrix multiply in LAPACK++ on the IBM RS/6000 Model 550 workstation. GNU g++ v. 2.3.1 was used together with the ESSL BLAS-3 routine dgemm.

2.4 Performance

The performance of LAPACK++ is almost indistinguishable from optimized Fortran. Figure 2.1, for example, illustrates the performance (Megaflop) rating of the simple code

```
C = A*B;
```

for square matrices of various sizes on the IBM RS/6000 Model 550 workstation. This particular implementation used GNU g++ v. 2.3.1 and utlized the BLAS-3 routines from the native ESSL library. The performance results are nearly identical with those of optimized Fortran calling the same library. This is accomplished by *inlining* the LAPACK++ BLAS kernels directly into the unerlying function call. This occurs at *compile* time, without any runtime overhead. The performance numbers are very near the machine peak and illustrate that using C++ with optimized computational kernels provides an elegant high-level interface without sacrificing performance.

The performance difference between optimized BLAS called from Fortran and the BLAS++ called from C++ is barely measureable. Figure 2.2 illustrates performance characteristics of the LU factorization of for square matrices of various sizes on the IBM RS/6000 Model 550 workstation. This particular implementation used GNU g++ v. 2.3.1 and utlized the BLAS-3 routines from the native ESSL library. The performance results are nearly identical with those of optimized Fortran calling the same library. This is accomplished by *inlining* the LAPACK++ BLAS kernels directly into the unerlying function call.

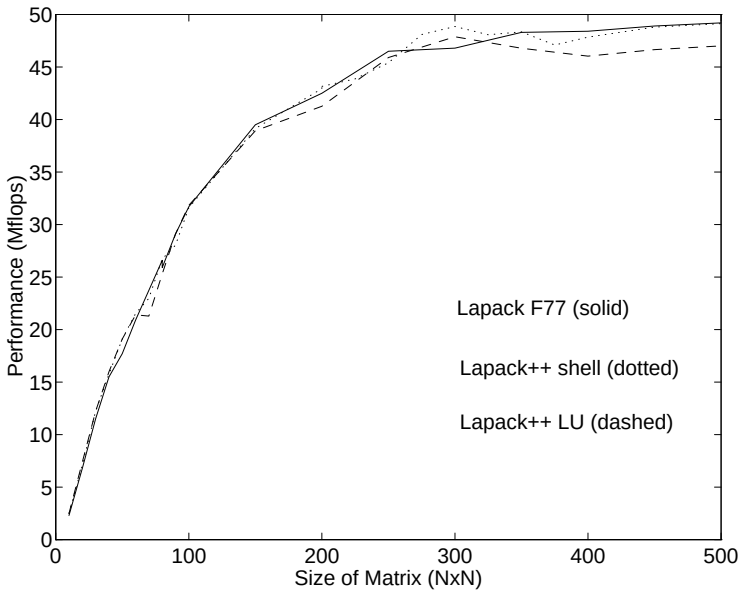

Figure 2.2: Performance of LAPACK++ LU factorization on a IBM RS/6000 Model 550 workstation, using GNU g++ v. 2.3.1 and BLAS routines from the IBM ESSL library. The results are nearly identical to the Fortran LAPACK performance. The LAPACK++ shell is call the Fortan `dgetrf()`, the LAPACK++ LU implements the right-looking LU algorithm in C++.

Chapter 3

LAPACK++ Matrix Objects

The fundamental objects in LAPACK++ are numerical vectors and matrices; however, LAPACK++ is **not** a general-purpose array package. Rather, LAPACK++ is a self-contained interface consistsing of only the minimal number of classes to support the functionality of the LAPACK algorithms and data strucutres.

LAPACK++ matrices can be referenced, assigned, and used in mathematical expressions as naturally as if they were an integral part of C++; the matrix element a_{ij}, for example, is referenced as `A(i,j)`. In keeping with the indexing convention of C++, matrix subscripts begin at zero. Thus, `A(0,0)` denotes the first element of `A`. Internally, LAPACK++ matrices are typically stored in column-order for compatibility with Fortran subroutines and libraries.

Various types of matrix structures are supported: banded, symmetric, Hermitian, packed, triangular, tridiagonal, bidiagonal, and nonsymmetric. The following sections describe these in more detail.

Matrix classes andother related data types specific to LAPACK++ begin with the prefix "`La`" to avoid naming conflicts with user-defined or third-party matrix packages. The list of valid names is a subset of the nomenclature shown in Figure 3.1.

3.1 Fundamental Matrix Operations

In its most general formulation, a matrix is a data structure representing an ordered collection of elements that can be accessed by an integer pair. In addition to their most common representation as rectangular arrays, matrices can also be represented as linked lists (unstructured sparse), in packed formats (triangular), as a series of Householder transformations (orthogonal), or stored by diagonals. Nevertheless, these various representations exhibit some commonality. The following operations are available for all matrix types:

- **Declarations.** Matrices can be constructed dynamically by specifying their size as a pair of non-negative integers, as in

```
LaGenMatDouble A(M,N)
```

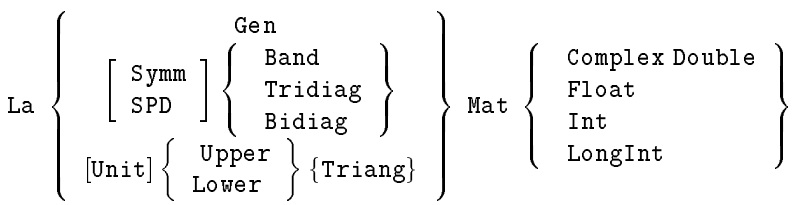

Figure 3.1: LAPACK++ matrix nomenclature. Items in square brackets are optional.

- **Indexing.** Individual elements of the matrix can be accessed by an integer pair (i, j) where $0 \leq i < M$ and $0 \leq j < N$. The indexing syntax for a matrix `A` is the natural

  ```
  A(i,j)
  ```

  notation, and can also be used as a destination for an assignment, as in `A(0,3) = 3.14`. A runtime error is generated if the index is out of the matrix bounds or outside of specified storage scheme (for example, accessing outside the non-zero portion of a banded matrix). This index range checking can be turned off by defining `LA_NO_BOUNDS_CHECK`.

- **Assignment.** The basic matrix assignment is by copying, and is denoted by the `=` operator, as in

  ```
  A = B
  ```

  The sizes of `A` and `B` must conform.

- **Referencing.** The unnecessary copying of matrix data elements has been avoided by employing the `reference()` method, as in

  ```
  A.ref(B)
  ```

  to assign the data space of B to the matrix A. The matrix elements are shared by both A and B, and changes to the data elements of one matrix will effect the other.

In addition to these basic operations, all matrix objects employ a destructor to free their memory when leaving their scope of visibility. By utilizing a reference counting scheme that keeps track of the number of aliases utilizing all or part of its data space, a matrix object can safely recycle unused data space.

3.2 General Rectangular Matrices

One of the fundamental matrix types in LAPACK++ is a general dense rectangular matrix. This corressponds to the most common notion of a matrix as a two-dimensional array. The corresponding LAPACK++ names are given as `LaGenMat`*Type* for `La`pack `Gen`eral `Mat`rix. The *Type* can be `Int`, `LongInt`, `Float`, `Double`, or `Complex`. Matrices in this category have the added property that submatrices can be efficiently accessed and referenced in matrix expressions. This is a necessity for describing block-structured algorithms.

3.2.1 Declarations

General LAPACK++ matrices may be declared (constructed) in various ways:

```
#include <lapack++.h>
float d[4] = {1.0, 2.0, 3.0, 4.0};

LaGenMatDouble    A(200,100)            // 1
LaGenMatComplex   B;                    // 2
LaGenMatDouble    D(A);                 // 3
LaGenMatFloat     E(d, 2, 2)            // 4
```

Line (1) declares `A` to be a rectangular 200x100 matrix. The elements are **uninitialized**. Line (2) declares `B` to be an empty (uninitialized) matrix of complex numbers. Until B becomes initialized, any attempt to reference its elements will result in a run time error. Line (3) illustrates the copy constructor; `D` is a copy of `A`. Finally, line (4) demonstrates how one can initialize a 2x2 matrix with the data from a standard C++ vector. The values are initalized in column-major form, so that the first column of `E` contains $\{1.0, 2.0\}$, and the second column contains $\{3.0, 4.0\}$.

Submatrices

Blocked linear algebra algorithms utilize submatrices as their basic unit of computation. It is crucial that submatrix operations be highly optimized. Because of this, LAPACK++ provides mechanisms for accessing rectangular subregions of a general matrix. These regions are accessed by *reference*, that is, without copying data, and can be used in any matrix expression.

Submatrices in LAPACK++ are denoted by specifying a subscript range through the `LaIndex()` function. For example, the 3x3 matrix in the upper left corner of `A` is denoted as

```
A( LaIndex(0,2), LaIndex(0,2) )
```

This references $A_{ij}, i = 0, 1, 2\ j = 0, 1, 2$, and is equivalent to the `A(0:2,0:2)` colon notation used in Fortran 90 and MATLAB^{TM}. Submatrix expressions may be also be used as a destination for assignment, as in

```
A( LaIndex(0,2), LaIndex(0,2) ) = 0.0;
```

which sets the 3x3 submatrix of A to zero. The index notation has an optional third argument denoting the stride value,

```
LaIndex(start, end, increment)
```

If the `increment` value is not specified it is assumed to be one. The expression `LaIndex(s, e, i)` is equivalent to the index sequence

$$s,\ \ s+i,\ \ s+2i,\ \ \ldots\ \ s+\lfloor\frac{e-s}{i}\rfloor i$$

Of course, the internal representation of an index is not expanded to a full vector, but kept in its compact (start,increment,end) format. The increment values may be negative and allow one to traverse a subscript range in the opposite direction, such as in `(10,7,-1)` which denotes the sequence $\{10, 9, 8, 7\}$. Indices can be named and used in expressions, as in the following submatrix assignments,

```
LaGenMatDouble A(10,10), B, C;                    // 1
LaIndex I(1,9,2),                                 // 2
LaIndex J(1,3,2);                                 // 3


B = A(I,I);                                       // 4
B(2,3) = 3.1;                                     // 5
C = B(LaIndex(2,2,4), J);                         // 6
```

In lines (2) and (3) we declare indices `I` $= \{1, 3, 5, 7, 9\}$, and `J` $=\{1, 3\}$. Line (4) sets `B` to the specified 5x5 submatrix of `A`. The submatrix B can used in any matrix expression, including accessing its individual elements, in line (5). Note that `B(2,3)` is the memory location as `A(5,7)`. Line (6) assigns the 2x2 submatrix of B to C. Note that C can also be thought of as `A( LaIndex(5,2,9), LaIndex(3,2,7))`.

Although LAPACK++ submatrix expressions allow one to access non-contiguous row or columns, many of the LAPACK routines only allow submatrices with unit stride in the column direction. Calling an LAPACK++ routine with a non-contiguous submatrix columns will cause data to be copied into a contiguous submatrix, if necessary. (The alternative, as done in the Fortran LAPACK, is not to allow this type of call at all.)

Storage	Triangular matrix A	Storage in array A
`Upper`	$\begin{pmatrix} a_{00} & a_{01} & a_{02} & a_{03} \\ & a_{11} & a_{12} & a_{13} \\ & & a_{22} & a_{23} \\ & & & a_{33} \end{pmatrix}$	$\begin{matrix} a_{00} & a_{01} & a_{02} & a_{03} \\ * & a_{11} & a_{12} & a_{13} \\ * & * & a_{22} & a_{23} \\ * & * & * & a_{33} \end{matrix}$
`Lower`	$\begin{pmatrix} a_{00} & & & \\ a_{10} & a_{11} & & \\ a_{20} & a_{21} & a_{22} & \\ a_{30} & a_{31} & a_{32} & a_{33} \end{pmatrix}$	$\begin{matrix} a_{00} & * & * & * \\ a_{10} & a_{11} & * & * \\ a_{20} & a_{21} & a_{22} & * \\ a_{30} & a_{31} & a_{32} & a_{33} \end{matrix}$

Figure 3.2: Internal storage pattern for an example 5x5 triangular matrix.

3.3 Triangular Matrices

Triangular matrices are denoted having only zero-valued entires below (lower triangular), or above (upper triangular) the main diagonal. (See Figure 3.2.) The triangular matrix types in LAPACK++ consist of

$$\texttt{La} \left\{ \begin{array}{l} \texttt{Upper} \\ \texttt{Lower} \\ \texttt{Unit} \end{array} \right\} \texttt{Triang} \left\{ \begin{array}{l} \texttt{Float} \\ \texttt{Double} \end{array} \right\}$$

3.3.1 Symmetric, Hermitian and SPD Matrices

Symmetric matrices are square and have $a_{ij} = a_{ji}$; Hermitian matrices have $a_{ij} = a^*_{ji}$, where * denotes complex conjugation; a symmetric positive definite matrix A has the added property that $x^T Ax > 0$ for any nonzero x.

Matrix types in each class of this group consists of two components: the mathematical characteristic (e.g. symmetric, Hermitian), and the underlying matrix element type (e.g., float, complex). The possible matrix types include

$$\texttt{La} \begin{array}{ll} \texttt{Symm} & \left\{ \begin{array}{l} \texttt{Double} \\ \texttt{Complex} \end{array} \right\} \\ \texttt{SPD} & \texttt{Double} \\ \texttt{HPD} & \texttt{Complex} \end{array}$$

Internally, symmetric and Hermitian matrices may be stored in an `Upper` or `Lower` format as shown in figure 3.3.

3.4 Banded Matrices

Square matrices whose sparsity pattern has nonzeros close to the diagonal can often be efficiently represented as banded matrices. An $N \times N$ banded matrix is specified by its size N, the number of subdiagonals kl, and the number of superdiagonals ku. In practice, this storage scheme should be used only when kl, and ku are much smaller than N. These matrices can be declared as

```
LaBandedMatDouble B(N, kl, ku);
LaBandedMatDouble A;
```

A reference to element B_{ij} is expressed as `B(i,j)`, and if the macro `LA_NO_BOUNDS_CHECK` is undefined, an error is generated if `(i,j)` lies outside the bands. Banded matrices may be copied and assigned as

Storage	Hermitian matrix A	Storage in array A
`Upper`	$\begin{pmatrix} a_{00} & a_{01} & a_{02} & a_{03} \\ \bar{a}_{01} & a_{11} & a_{12} & a_{13} \\ \bar{a}_{02} & \bar{a}_{12} & a_{22} & a_{23} \\ \bar{a}_{03} & \bar{a}_{13} & \bar{a}_{23} & a_{33} \end{pmatrix}$	$\begin{matrix} a_{00} & a_{01} & a_{02} & a_{03} \\ * & a_{11} & a_{12} & a_{13} \\ * & * & a_{22} & a_{23} \\ * & * & * & a_{33} \end{matrix}$
`Lower`	$\begin{pmatrix} a_{00} & \bar{a}_{10} & \bar{a}_{20} & \bar{a}_{30} \\ a_{10} & a_{11} & \bar{a}_{21} & \bar{a}_{31} \\ a_{20} & a_{21} & a_{22} & \bar{a}_{32} \\ a_{30} & a_{31} & a_{32} & a_{33} \end{pmatrix}$	$\begin{matrix} a_{00} & * & * & * \\ a_{10} & a_{11} & * & * \\ a_{20} & a_{21} & a_{22} & * \\ a_{30} & a_{31} & a_{32} & a_{33} \end{matrix}$

Figure 3.3: Storage pattern for Symmetric and Hermitian matrices.

Band matrix A	Band storage in array AB
$\begin{pmatrix} a_{00} & a_{01} & & & \\ a_{10} & a_{11} & a_{12} & & \\ a_{20} & a_{21} & a_{22} & a_{23} & \\ & a_{31} & a_{32} & a_{33} & a_{34} \\ & & a_{42} & a_{43} & a_{44} \end{pmatrix}$	$\begin{matrix} * & a_{01} & a_{12} & a_{23} & a_{34} \\ a_{00} & a_{11} & a_{22} & a_{33} & a_{44} \\ a_{10} & a_{21} & a_{32} & a_{43} & * \\ a_{20} & a_{31} & a_{42} & * & * \end{matrix}$

Figure 3.4: Storage pattern for banded 5x5 matrix with 2 superdiagonals and 1 subdiagonal.

```
A.ref(B);          // shallow assignment
A = B;             // copy
```

similar to the `LaGenMat` classes. The ith diagonal of `B` is accessed as `B(i)`, and `B(-i)` is the ith subdiagonal ($i > 0$). Thus `B(0)` is the main diagonal, The result is a vector (with possible non-unit stride),

```
LaVectorDouble d = B(-2);      // 2nd subdiagonal of B
```

whose size is the length of the particular diagonal, not the matrix size N. Note that this diagonal indexing expression can be used as any `LaVector`. For example, it is perfectly legal to write

```
B(-2)(LaIndex(0,3)) = 3.1;
```

to set the first four elements of the second subdiagonal of B to the value of 3.1.

Accessing out of the matrix band generates a run-time error, unless `LA_NO_BOUNDS_CHECK` is set.

3.4.1 Triangular Banded Matrices

Triangular banded matrices are declared, stored, and accessed in a similar format, except that $kl = 0$ for upper triangular, and $ku = 0$ for lower triangular matrices:

```
LaUpperTriangBandedMatDouble  U(N, kl);
LaLowerTriangBandedMatDouble  L(N, ku);
```

Diagonals and individual elements are accessed the same way as general banded matrices. Triangular banded matrices can also be aliases for the upper or lower region of general banded matrices in the following example

```
LaBandedMatDouble B(N, kl, ku) = 0.0;
LaUpperTriangBandedMatDouble U;
```

```
LaLowerTriangBandedMatDouble L;

U.ref(B);
L.ref(B);

U(0,0) = 1.0;           // B(0,0) and L(0,0) also set to 1.0
                        //    (diagonal is shared by L and U.)
```

3.4.2 Symmetric and Hermitian Banded Matrices

Symmetric and Hermitian banded matrices with kd subdiagonals and superdiagonals are specified as

```
LaSymmBandedMatDouble S(N, kd);
LaHermBandedMatComplex C(N, kd);
```

Diagonals and individual elements are accessed the same way as general baned matrices. If `LA_NO_BOUNDS_CHECK` is undefined, accessing out of the range of diagonal bands generates a run-time error.

3.5 Tridiagonal Matrices

Although bidiagonal and tridiagonal matrices are special cases of the more general banded matrix structure, their occurrence is so common in numerical codes that it is advantageous to treat them as a special case. Unlike general banded matrices, tridiagonal matrices are stored by diagonals rather than columns. A tridiagonal matrix of order N is stored in three one-dimensional arrays, one of length N containing the diagonal elements and two of length $N-1$ containing the subdiagonal and superdiagonal elements in elements 0 through $N-2$. This ensures that diagonal elements are in consecutive memory locations.

Tridiagonal $N \times N$ matrices are constructed as

```
LaTridiagMatDouble B(N,N);
LaTridiagMatDouble A;
```

reserving a space of 3N (see [1]) elements, or from user data as

```
LaTridiagMatDouble T(N, d, dl, du)
```

where `*d`, `*dl`, and `*du` are contiguous vectors that point to the diagonal, subdiagonal, and superdiagonals, respectively. Note that `*d` must point to N elements, but `*du` and `*dl` can point to only N-1.

Matrix elements can be accessed as `T(i,j)`, where $|(i-j)| \leq 2$, (i.e. elements on the diagonal, sub or superdiagonal) otherwise there is a bounds-check error generated at run time.

Since the largest dense submatrix of a tridiagonal matrix has only four elements (2×2), index ranges are of limited use and are therefore not implemented with tridiagonal matrices.

Chapter 4

Driver Routines

4.1 Linear Equations

This section provides LAPACK++ routines for solving linear systems of the form

$$Ax = b. \tag{4.1}$$

where A is the **coefficient matrix**, b is the **right hand side**, and x is the **solution**. A is assumed to be square matrix of order n, although underlying computational routines allow for A to be rectangular. For several right hand sides, we write

$$AX = B, \tag{4.2}$$

where the columns of B are individual right hand sides, and the columns of X are the corresponding solutions. The task is to find X, given A and B. The coefficient matrix A can be one of the types show in Figure 3.1. Note that for real (non-complex) matrices, *symmetric* and *Hermitian* are equivalent.

The basic syntax for a linear equation driver in LAPACK++ is given by

```
LaLinSolve(A, X, B);
```

The matrices `A` and `B` are input, and `X` is the output. `A` is an $M \times N$ matrix of one of the above types. Letting *nrhs* denote the number of right hand sides in eq. 4.2, `X` and `B` are both rectangular matrices of size $N \times nrhs$.

This version requires intermediate storage of $\approx M * (N + nrhs$ elements. Section 5 describes how to use factorization classes to reduce this storage requirement at the expense of overwriting A and B.

In cases where no additional information is supplied, the LAPACK++ routines will attempt to follow an intelligent course of action. For example, if `LaLinSolve(A,X,B)` is called with a non-square MxN matrix, the solution returned will be the linear least square that minimizes $||Ax - b||_2$ using QR factorization. Or, if A is SPD, then the Cholesky factorization will be used. Alternatively, one can directly specify the exact factorization method, such as `LU_factor(F, A)`. In this case, if A is non-square, only the factors return represent only a partial factorization of the upper square portion of A.

Error conditions in performing the `LaLinSolve()` operations can be retrieved via the `LaLinSolveInfo()` function, which returns information about the last called `LaLinSolve()`. A zero value denotes a successful completion. A negative value of $-i$ denotes that the ith argument was somehow invalid or inappropriate. A positive value of i denotes that in the LU decomposition, $U(i,i) = 0$; the factorization has been completed but the factor U is exactly singular, so the solution could not be computed. In this case, the value returned by `LaLinSolve()` is a null (0x0) matrix.

Table 4.1: LAPACK++ Drivers for Linear Equations: Rectangular Matrices.

General	`LaGenMat<t> A(M,N);`
	`LaGenMat<t> B(M,nrhs), X(M,nrhs);`
	`LaGenFact<t> F;`
	
	`LaLinSolve(A, X, B);`
	`LaLinSolveIP(A, X, B);`
	`LaLinSolve(F, X, B);`

Symmetric	`LaUpperSymmMat<t> A(N,N);`
	`LaLowerSymmMat<t> A(N,N);`
	`LaGenMat<t> B(N, nrhs), X(N, nrhs);`
	`LaSymmFact<t> F;`
	
	`LaLinSolve(A, X, B);`
	`LaLinSolveIP(A, X, B);`
	`LaLinSolve(F, X, B);`

Symmetric Positive Definite	`LaUpperSPDMat<t> A(N,N);`
	`LaLowerSPDMat<t> A(N,N);`
	`LaGenMat<t> B(N, nrhs), X(N, nrhs);`
	`LaSPDFact<t> F;`
	
	`LaLinSolve(A, X, B);`
	`LaLinSolveIP(A, X, B);`
	`LaLinSolve(F, X, B);`

Complex Symmetric	`LaUpperSymmMat<t> A(N,N);`
	`LaLowerSymmMat<t> A(N,N);`
	`LaGenMat<t> B(N, nrhs), X(N, nrhs);`
	`LaSymmFact<t> F;`
	
	`LaLinSolve(A, X, B);`
	`LaLinSolveIP(A, X, B);`
	`LaLinSolve(F, X, B);`

Table 4.2: LAPACK++ Drivers for Linear Equations: Tridiagonal Matrices.

General	`LaTriadMat<T> A(N,N);`
	`LaGenMat<T> B(N,nrhs), X(N,nrhs);`
	`LaGenFact<T> F;`
	
	`LaLinSolve(A, X, B);`
	`LaLinSolveIP(A, X, B);`
	`LaLinSolve(F, X, B);`

Symmetric Positive Definite	`LaUpperTriadSPDMat<T> A(N,N);`
	`LaLowerTriadSPDMat<T> A(N,N);`
	`LaGenMat<T> B(N, nrhs), X(N, nrhs);`
	`LaSPDFact<T> F;`
	
	`LaLinSolve(A, X, B);`
	`LaLinSolveIP(A, X, B);`
	`LaLinSolve(F, X, B);`

Table 4.3: LAPACK++ Drivers for Linear Equations: Banded Matrices.

General	`LaBandedMat<T> A(N,N);`
	`LaGenMat<T> B(N,nrhs), X(N,nrhs);`
	`LaBandedFact<T> F;`
	
	`LaLinSolve(A, X, B);`
	`LaLinSolve(F, X, B);`

Symmetric Positive Definite	`LaUpperBandedSPDMat<T> A(N,N);`
	`LaLowerBandedSPDMat<T> A(N,N);`
	`LaGenMat<T> B(N, nrhs), X(N, nrhs);`
	`LaSPDFact<T> F;`
	
	`LaLinSolve(A, X, B);`
	`LaLinSolve(F, X, B);`

4.2 Eigenvalue Problems

Example routines are provided in `LAPACK++/SRC/eigslv.cc` for the solution of symmetric eigenvalue problems. The function

```
LaEigSolve(A, v);
```

for computing the eigenvalues `v` of symmetric matrix `A`, and

```
LaEigSolve(A, v, V);
```

for computing the eigenvectors `V`. Similarly,

```
LaEigSolveIP(A, v);
```

overwrites `A` with the eigenvectors.

4.3 Memory Optimizations: Factorizing in place

When using large matrices that consume a significant portion of available memory, it may be beneficial to remove the requirement of separately storing intermediate factorizaton representations at the expense of destroying the contents of the input matrix A. For most matrix factorizations we require temporary data structures roughly equal to the size of the original input matrix. (For general banded matrices, one may need even slight more, see section 3.4.) For example, for a square $N \times N$ dense nonsymmetric factorization, the temporary memory requirement can be reduced from $N \times (N + nrhs + 1)$ elements to $N \times 1$ Such memory-efficient factorizations are accomplished with the `LaLinSolveIP()` routine:

```
LaLinSolveIP(A, X, B);
```

Here the contents of A are overwritten (with the respective factorization), and B is overwritten by the soultion. It is also explicitly returned by the function so that it matches the return type of `LaLinSolve()`. In the line above, both `X` and `B` refer to the same memory locations.

Chapter 5

Factorization Classes

Factorization classes are used to describe the various types of matrix decompositions: LU, Cholesky (LL^T), QR, and singular-value decompositions (SVD). The driver routines of LAPACK++ typically choose an appropriate factorization, but the advanced user can express specific factorization algorithms and their variants for finer control of their application or for meeting strict memory storage and performance requirements.

In an object-oriented paradigm it is natural to encapsulate the factored representation of a matrix in a single object. An LU factorization, for example, returns the upper and unit-lower triangular factors, L and U, as well the pivoting information that describes how the rows were permuted during the factorization process. The representation of the L and U factors is incomplete without this information. Rather than store and manage these components separately, a factorization class is used as follows,

```
LaGenMatDouble A, B, X;
LaGenFactDouble F;

LaLUFactor(F, A);
LaLinSolve(F, X, B);
```

More importantly, the various factorization components can be extracted from `F`, as in,

```
LaUnitLowerTriangMatDouble L;
LaUpperTriangMatDouble U;
LaGenMatDouble Y;

L = F.L();
U = F.U();

Y = LaLinSolve(L, B);
X = LaLinSolve(U, Y);
```

Here we solve $AX = B$ by first solving the lower triangular system, $LY = B$, and then the upper triangular system, $UX = Y$. (The pivot information is stored in both L and U.) The `LaGenFact` object can also be used directly to solve $(LU)X = B$ by calling

```
LaLinSolve(F, X, B);
```

5.1 Optimizations: Factoring in Place

By default, the matrix factorization does not alter the contents of the original matrix or overwrite it. Nevertheless, the ability to "factor in place" is crucial when dealing with realistic memory constraints on large matrices, and is a necessity to implement blocked linear algebra algorithms efficiently.

Here we utilize the "in place" versions of computational routines (see section 4.3), which overwrite the matrix A to conserve space:

```
LaGenMatDouble A;
LaGenFactDouble F;

LaLUFactorIP(F, A);
```

Appendix A

Programming Examples

To illustrate what programming with LAPACK++ looks like to scientific and engineering code developers, this section provides a few code examples. These examples are presented here for illustrative purposes, yet we have tried to use realistic examples that accurately display a level of sophistication encountered in realistic applications.

A.1 Polynomial data fitting

This code example solves the liner least squares problem of fitting N data points (x_i, y_i) to a dth degree polynomial equation

$$p(x) = a_0 + a_1 x + a_2 x^2 + \ldots + a_d x^d$$

using QR factorization. Given the two vectors x and y it returns the vector of coefficients $\mathtt{a} = \{a_0, a_1, a_2, \ldots a_{d-1}\}$. It is assumed that $N \gg d$. The solution arises from solving the overdetermined Vandermonde system $Xa = y$:

$$\begin{bmatrix} 1 & x_0^1 & x_0^2 & \ldots & x_0^d \\ 1 & x_1^1 & x_1^2 & \ldots & x_1^d \\ \vdots & & & & \vdots \\ 1 & x_{N-1}^1 & x_{N-1}^2 & \ldots & x_{N-1}^d \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ \vdots \\ a_d \end{bmatrix} = \begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_{N-1} \end{bmatrix}$$

in the least squares sense, i.e., minimizing $||Xa - y||_2$. The resulting code is shown in figure A.1.

```
LaVectorDouble poly_fit(LaVectorDouble x, LaVectorDouble y, int d)
{
    int N = min(x.size(), y.size());

    LaGenMatDouble P(N,d);
    LaVectorDouble a(d);

    for (i=0; i<N; i++)     // construct Vandermonde matrix
    {
        x_to_the_j = 1;
        for (j=0; j<d; j++)
        {
            P(i,j) = x_to_the_j;
            x_to_the_j *= x(i);
        }
    }
    a = LaQRLinSolveIP(P, y);   // solve Pa = y using linear least squares

    return a;
}
```

Figure A.1: Code Example: polynomial data fitting.

Appendix B

Release Notes for v. 1.0

B.1 How to report bugs

Report bugs, comments, questions and suggestions to `lapackpp@cs.utk.edu`.

B.2 Tips and Suggestions

B.3 Reducing Compilation time

B.3.1 Minimize the number of included header files

LAPACK++ header files can be quite large, particularly when one takes into the various data types (e.g. double precision, complex) and different matrix types (e.g. symmetric, upper triangular, tridiagonal).

If your compiler does not support the ability pre-compile header files, it could be spending a significant portion of its time processing LAPACK++ object and function declarations.

The main header `lapack++.h` is a "catch-all" file containing complete declarations for over fifty matrix classes and several hundred inlined matrix functions. While this file is most general, it is unlikely that one will need all of its declarations. It is much more efficient to include a specific header file for a matrix storage class or matrix data type.

B.3.2 Header file madness

The LAPACK++ header files will typical only include those function declarations for the matrix types currently defined. For example, if `LaSpdMatDouble` matrices are used, then the compiler macro `_LA_SPD_MAT_DOUBLE_H_` is defined. Thus, the LAPACK++ can select to include only those header files which are relevant.

The alternative is to include every possible matrix type and LAPACK++ function, for every program, even a tiny program which uses a small portion of the LAPACK++ / MATRIX++ classes.

Therefore, one should declare all of their matrix and vector types **before** including BLAS++ and LAPACK++ header files.

B.4 Performance Considerations

B.4.1 Array Bounds Checking

LAPACK++ allows the application developer to determine the level of run-time index range checking performed when accessing individual $A(i, j)$ elements. It is strongly recommend for developers to use the

`LA_BOUNDS_CHECK` whenever possible, thus guarding against out-of-bounds references. This is typically used in the design and debugging phases of the program development, and can be turned off in production runs.

B.4.2 Improving $A(i,j)$ efficiency

The LAPACK++ matrix classes were optimized for use with BLAS and LAPACK routines. These routines require matrices to be column-ordered and adhere to Fortran-like accessibility.

The code for `A(i,j)` inlines integer multiplications to compute the address offset; however, most C++ compilers cannot optimize this out of a loop. (See note below.) An alternate implementation uses an indirect-addressing scheme similar to $A[i][j]$ to access individual elements. Tests have shown this to be more successful on various C++ compilers: Sun, Borland.

B.5 Exception Handling

Although exception handling has been officially accepted into the language by the ANSI committee (November 1990) it is not yet supported by all C++ compilers. LAPACK++ will use C++ exception handling as soon as it becomes widely available. Currently, Version 0.9 uses a macro mechanism to simulate `throw` expressions so that the library code will work correctly when exception handling becomes supported:

```
#define throw throw_
inline void throw(const char *message)
{
   cerr << "Exception: " << message << endl;
   exit(1);
}
```

Transfer of control does not correspond to a "try" block, but rather exits the program with the proper diagnostic message. This is a similar behavior to the real exception handling if there was no explicit user-supplied try block.

B.6 Frequently asked questions

B.6.1 What is the performance of LAPACK++?

For medium-to-large large matrices (say $n > 100$) where performance is a consideration, LAPACK++ performance is identical to the native Fortran LAPACK codes. For example, a recent test on IBM RS/6000 workstation showed both packages obtaining speeds of 50 Mflops for matrices of size 300x300 and larger. This is not surprising, since the both utilize the same optimized Level 3 BLAS routines for the computation-intensive sections. See [6] for more performance details.

B.6.2 Do I need a Fortran compiler to use LAPACK++?

No, you can use the C version of LAPACK available from `netlib`.

B.6.3 Why are the LAPACK++ matrix classes not templated?

There is a templated vector example `template_v.h` in the `LAPACK++/INCLUDE` subdirectory. This should provide provide an example for those who wish to create vectors and matrices of arbitrary types.

LAPACK++ does not support generic matrix types, since the underlying BLAS code supports only support real and complex floating point numbers.

B.6.4 Can LAPACK++ work with built-in C/C++ arrays and other matrix classes?

Yes. LAPACK++ routines utilize contiguous storage for optimum data reuse (Level 3 BLAS) and C++ arrays must be explicitly transformed into this data storage before integrating with LAPACK++. This is easily accomplished with the various matrix constructors, e.g. `LaGenMatDouble::LaGenMatDouble( double **, int, int )`.

For integrating with external C++ matrix classes, all that is required of a user-specific matrix class is the ability to access individual elements, i.e. something similar to `A(i,j)`. Any C++ matrix library should support this. If your matrix class uses column-major ordering, then the conversion can be as simple as a copying a pointer to the first element. Consult the *LAPACK++ Users' Manual* for more details.

B.6.5 Can the matrix and vector objects be used independently of BLAS++ and LAPACK++?

Yes, the MATRIX++ subdirectory `LAPACK++/MATRIX++` subdirectory (see figure ??) can be used as an independent package. This may be useful if to develop matrices and vectors of user-defined types. The codes in this subdirectory define only access functions, assignments, construction, and basic non-arithmetic relationship between matrices.

B.6.6 Is there an existing standard for C++ matrix classes?

As of early 1994, no. Researchers in the LAPACK++ project have been working with various groups in C++ numerics to establish such a standard. Some efforts are being made to bring an multi-dimensional array class proposal to the ANSI C++ committee, but as far as matrix classes (e.g. banded, symmetric, sparse, orthogonal) are concerned, a formal standard has not been presented. Some of this will rely on user experiences and feeback with matrix packages.

Appendix C

Direct Interfaces to LAPACK and BLAS Routines

Filename	*Function*	*Description*
blas++.h	Vec * Vec	Vector * Vector
	Vec + Vec	Vector + Vector
	Vec - Vec	Vector - Vector
	Mat * Vec	General Matrix * Vector
	Mat * Vec	Banded Matrix * Vector
	Mat * Vec	Symmetric Matrix * Vector
	Mat * Vec	Symmetric Banded Matrix * Vector
	Mat * Vec	SPD Matrix * Vector
	Mat * Vec	Lower Triang Matrix * Vector
	Mat * Vec	Upper Triang Matrix * Vector
	Mat * Vec	Unit Lower Triang Matrix * Vector
	Mat * Vec	Unit Upper Triang Matrix * Vector
	Mat + Mat	General Matrix + General Matrix
	Mat - Mat	General Matrix - General Matrix
	Mat * Mat	General Matrix * General Matrix
	Mat * Mat	Symmetric Matrix * General Matrix
	Mat * Mat	Lower Triang Matrix * General Matrix
	Mat * Mat	Upper Triang Matrix * General Matrix
	Mat * Mat	Unit Lower Triang Matrix * General Matrix
	Mat * Mat	Unit Upper Triang Matrix * General Matrix
	Norm_Inf(Vec)	Infinity Norm of a Vector
	Norm_Inf(Mat)	Infinity Norm of a Matrix
	Mach_eps_float()	Returns Machine Epsilon - float
	Mach_eps_double()	Returns Machine Epsilon - double

Filename	*Function*	*Description*
blas1++.h	Blas_Norm1(Vec)	One Norm of a Vector
	Blas_Norm2(Vec)	$nrm2 \leftarrow \|x\|_2$
	Blas_Add_Mult(double,Vec,Vec)	$y \leftarrow \alpha x + y$
	Blas_Add_Mult_IP(double,Vec,Vec)	$y \leftarrow \alpha y += x$
	Blas_Copy(Vec,Vec)	$y \leftarrow x$
	Blas_Dot_Prod(Vec,Vec)	$dot \leftarrow \alpha + x^T y$
	Blas_Apply_Plane_Rot(Vec,Vec)	Apply Plane Rotation
	Blas_Gen_Plane_Rot(double x 4)	Generate Plane Rotation
	Blas_Scale(double,Vec)	$x \leftarrow \alpha x$
	Blas_Swap(Vec,Vec)	$x \leftrightarrow y$
	Blas_Index_Max(Vec)	Returns Maximum Vector Element
blas2++.h	Blas_Mat_Vec_Mult(GenMat,Vec)	Gen Mat * Vec
	Blas_Mat_Vec_Mult(BandMat,Vec)	Band Mat * Vec
	Blas_Mat_Vec_Mult(SymmMat,Vec)	Symm Mat * Vec
	Blas_Mat_Vec_Mult(SymmBandMat,Vec)	Symm Band Mat * Vec
	Blas_Mat_Vec_Mult(SpdMat,Vec)	SPD Mat * Vec
	Blas_Mat_Vec_Mult(LowerTriangMat,Vec)	Lower Triang Mat * Vec
	Blas_Mat_Vec_Mult(UpperTriangMat,Vec)	Upper Triang Mat * Vec
	Blas_Mat_Vec_Mult(UnitLowerTriangMat,Vec)	Unit Lower Triang Mat * Vec
	Blas_Mat_Vec_Mult(UnitUpperTriangMat,Vec)	Unit Upper Triang Mat * Vec
	Blas_Mat_Vec_Solve(LowerTriangMat,Vec)	Lower Triang Mat/Vec Solve
	Blas_Mat_Vec_Solve(UpperTriangMat,Vec)	Upper Triang Mat/Vec Solve
	Blas_Mat_Vec_Solve(UnitLowerTriangMat,Vec)	Unit Lower Triang Mat/Vec Solve
	Blas_Mat_Vec_Solve(UnitUpperTriangMat,Vec)	Unit Upper Triang Mat/Vec Solve
	Blas_R1_Update(GenMat,Vec,Vec)	$A \leftarrow \alpha x y^T + A$
	Blas_R1_Update(SymmMat,Vec,Vec)	$A \leftarrow \alpha x x^T + A$
	Blas_R1_Update(SpdMat,Vec,Vec)	$A \leftarrow \alpha x x^T + A$
	Blas_R2_Update(SymmMat,Vec,Vec)	$A \leftarrow \alpha x y^T + \alpha y x^T + A$
	Blas_R2_Update(SpdMat,Vec,Vec)	$A \leftarrow \alpha x y^T + \alpha y x^T + A$
blas3++.h	Blas_Mat_Mat_Mult(GenMat,GenMat)	Gen Mat * Gen Mat
	Blas_Mat_Trans_Mat_Mult(GenMatT,GenMat)	Gen Mat Transpose * Gen Mat
	Blas_Mat_Mat_Trans_Mult(GenMat,GenMatT)	Gen Mat * Gen Mat Transpose
	Blas_Mat_Mat_Mult(LowerTriangMat,GenMat)	Lower Triang Mat * Gen Mat
	Blas_Mat_Mat_Mult(UpperTriangMat,GenMat)	Upper Triang Mat * Gen Mat
	Blas_Mat_Mat_Mult(UnitLowerTriangMat,GenMat)	Unit Lower Triang Mat * Gen Mat
	Blas_Mat_Mat_Mult(UnitUpperTriangMat,GenMat)	Unit Upper Triang Mat * Gen Mat
	Blas_Mat_Mat_Mult(SymmMat,GenMat,GenMat)	Symm Mat * Gen Mat
	Blas_Mat_Mat_Solve(LowerTriangMat,GenMat)	Lower Triang Mat/Gen Mat Solve
	Blas_Mat_Mat_Solve(UpperTriangMat,GenMat)	Upper Triang Mat/Gen Mat Solve
	Blas_Mat_Mat_Solve(UnitLowerTriangMat,GenMat)	Unit Lower Triang Mat/Gen Mat Solve
	Blas_Mat_Mat_Solve(UnitUpperTriangMat,GenMat)	Unit Upper Triang Mat/Gen Mat Solve
	Blas_R1_Update(SymmMat,GenMat)	$C \leftarrow \alpha A A^T + \beta C$
	Blas_R2_Update(SymmMat,GenMat,GenMat)	$C \leftarrow \alpha A B^T + \alpha B A^T + \beta C$

Bibliography

[1] E. Anderson, Z. Bai, C. Bischof, J. W. Demmel, J. J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK: A portable linear algebra library for high-performance computers*, Computer Science Dept. Technical Report CS-90-105, University of Tennessee, Knoxville, 1990. (LAPACK Working Note 20).

[2] E. Anderson, J. J. Dongarra, S. Ostrouchov, *Installation Guide for LAPACK*, LAPACK Working Note 41, University of Tennessee, Computer Science Technical Report CS-92-151, Feb., 1992.

[3] J. Choi and J. J. Dongarra and D. W. Walker, *PB-BLAS : Parallel Block Basic Linear Algebra Subroutines on Distributed Memory Concurrent Computers*, Oak Ridge National Laboratory, Mathematical Sciences Section, in preparation, 1993.

[4] J. J. Dongarra, J. Du Croz, I. S. Duff, and S. Hammarling, *A set of Level 3 Basic Linear Algebra Subprograms*, ACM Trans. Math. Soft., 16 (1990), pp. 1–17.

[5] J. J. Dongarra, R. Pozo, D. W. Walker, *An Object Oriented Design for High Performance Linear Algebra on Distributed Memory Architectures*, Object Oriented Numerics Conference (OONSKI), Sunriver, Oregon, May 26-27, 1993.

[6] J. J. Dongarra, R. Pozo, D. W. Walker, *LAPACK++: A Design Overview of Object-Oriented Extensions for High Performance Linear Algebra*, Computer Science Technical Report, University of Tennessee, 1993.

[7] J. J. Dongarra, A. Lumsdaine, Xinhiu Niu, Roldan Pozo, Karin Remington, *A Sparse Matrix Library in C++ for High Performance Architectures*, Proceedings of the Object Oriented Numerics Conference, Sunriver, Oregon, April 1994.

[8] C. L. Lawson, R. J. Hanson, D. Kincaid, and F. T. Krogh, *Basic Linear Algebra Subprograms for FORTRAN usage*, ACM Trans. Math. Soft ., 5 (1979), pp. 308–323.

[9] B. T. Smith, J. M. Boyle, J. J. Dongarra, B. S. Garbow, Y. Ikebe, V. C. Klema, and C. B. Moler, *Matrix Eigensystem Routines – EISPACK Guide*, vol. 6 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, 2 ed., 1976.

Addison-Wesley, 1986.