

# A Revised Proposal for a Sparse BLAS Toolkit

SPARKER Working Note # 3

Sandra Carney\*    Michael A. Heroux<sup>†</sup>    Guangye Li<sup>†</sup>    Roldan Pozo<sup>‡</sup>  
Karin A. Remington<sup>‡</sup>    Kesheng Wu\*

January 1996

## Abstract

This paper describes a proposal for a “toolkit” of kernel routines for some of the basic operations in (iterative) sparse numerical methods. In particular, we describe an interface for routines which perform (i) sparse matrix times dense matrix product, (ii) the solution of a sparse triangular system with multiple right-hand-sides, (iii) the right permutation of a sparse matrix and (iv) a check for the integrity of a sparse matrix representation. The interfaces for these four operations are defined for a variety of common data structures and a set of guidelines is given to define interfaces for new data structures. The primary purpose of this toolkit is to provide a set of basic routines upon which the “User Level Sparse BLAS,” as described in [9], can be built. This paper is a revision of the original proposal found in [14].

## Keywords

Sparse matrices, sparse data structures, programming standards, sparse BLAS.

## 1 Introduction

Standard interfaces for numerical linear algebra software have been shown to be very useful if the interface is simple yet flexible enough for a wide variety of users. Examples of successful interfaces are the Level 1, Level 2 and Level 3 BLAS [16, 7, 6], LINPACK [5] and EISPACK [19], and more recently, LAPACK [1]. However, with the exception of perhaps the LINPACK banded storage scheme, widely adopted software standards for numerical linear algebra have been almost exclusively for dense linear systems. There are sparse extensions to the Level 1 BLAS (see [4]) which embody useful functionality. However, we need to consider higher level functionality to adequately address the needs of most iterative schemes.

Several efforts have been made at standards for sparse computations, but the task is

---

\*Department of Computer Science, University of Minnesota. Work supported in part by the Army Research Office and University of Minnesota AHPARC.

<sup>†</sup>Cray Research, Inc., 655 Lone Oak Dr. Eagan, MN 55121 USA.

<sup>‡</sup>National Institute of Standards and Technology, Gaithersburg, MD 20899

much more difficult. At the heart of this difficulty is the fact that there is no single widely accepted data structure for sparse matrices. In fact, there are many well-known data structures for sparse matrices and for each basic structure there are typically several minor variations. Any standard which would allow only a single data structure would severely limit the number of users that could effectively utilize the standard. Therefore, attempts at standards for sparse linear algebra software have tried to accommodate a variety of data structures.

One approach to standardized software for sparse linear systems is to provide a “toolkit”, or collection, of routines which perform common operations, such as sparse matrix-vector product, using a variety of common data structures. Some well-known examples of this are SPARSKIT by Saad [18] and the Iterative BLAS of Oppe and Kincaid [17]. The proposal described here is similar to the above two examples but is expressly developed to support the User Level Sparse BLAS proposed in [9].

Beyond this explicit relationship to the User Level Sparse BLAS, we want to emphasize that these kernels are generally useful for other high level interfaces. In particular, for “object-oriented” interfaces such as those described in [20, 12, 15], the sparse BLAS toolkit routines can serve as the primary computational kernels. Also, for iterative toolkits such as the “templates” described in [3], the sparse BLAS toolkit can provide a complimentary set of kernels.

## 1.1 Goals

The current goals of this proposal are to define an initial set of data-structure-dependent kernel routines for sparse iterative numerical methods and to establish a framework for extending this initial set as new data structures and operations are added. The purpose of these kernel routines is to provide the primary foundation for the User Level Sparse BLAS described in [9], i.e., to the extent possible, the User Level Sparse BLAS will be written using kernels from the sparse BLAS toolkit.

The ultimate goal is to eventually have the Sparse BLAS Toolkit become a *de facto* standard in much the same as the dense BLAS. We envisage that hardware vendors and software developers will provide optimized versions of (some) the kernel routines presented here as appropriate for the targeted scientific applications and hardware platforms. Application developers can then, via calls to the User Level Sparse BLAS, attain high performance, without code modifications, on a wide variety of architectures.

## 1.2 Deliverables

The final products of this proposal will be delivered in two phases. The first product will be a *basic* set of kernels (called the Basic Sparse BLAS Toolkit). The second product will be an *extended* set of kernels and a full test suite for all kernels (called the Extended Sparse BLAS Toolkit). The motivation for doing this is to allow delivery of the Basic Sparse BLAS Toolkit as soon as possible.

The Basic Sparse BLAS Toolkit includes kernels for computing for sparse matrix-dense matrix product and solution of sparse triangular systems for COO, CSC, CSR, BCO,

BSR, BSC and VBR data structures. The Extended Sparse BLAS Toolkit includes all other kernels described below. Specifically, it contains sparse matrix-dense matrix product and solution of sparse triangular systems for the rest of the supported data structures, he right permutation of a sparse matrix and the check for the integrity of a sparse matrix representation as described in this proposal.

All software is written in C in such a way that the routines are easily callable from Fortran, C or C++.

### 1.3 General Overview

In the remainder of this proposal we define the details of the Sparse BLAS Toolkit. In Section 2 we define the basic operations, independent of data structures. We also discuss our naming and argument conventions. Thus, this section defines the interfaces for the current set of routines in the Sparse BLAS Toolkit and also provides a framework for future extensions.

In Section 3 we define the supported data structures in detail. The precise definition of data structures is very important for this proposal since many varieties of the basic data structures exist. In effect, this proposal is attempting to standardize the definition of these data structures.

Section 4 discusses our rationale for the design. In this section we attempt to justify what is, and is not, included in our proposal. We also attempt to illustrate some of the intended uses for these kernels. Finally, in Section 5 we present some issues which are open for further discussion.

## 2 Scope of the Sparse BLAS Toolkit

The present scope of the Sparse BLAS Toolkit covers matrix-matrix product operations, the solution of triangular systems with multiple right-hand-sides, permutations of sparse matrices, and the integrity check of a sparse matrix representation. More precisely we have:

1. Matrix-matrix products:

$$C \leftarrow \alpha AB + \beta C$$

$$C \leftarrow \alpha A^T B + \beta C$$

2. Solution of triangular systems (supported for all but the COO and BCO data structures):

$$C \leftarrow \alpha DA^{-1}B + \beta C$$

$$C \leftarrow \alpha DA^{-T}B + \beta C$$

$$C \leftarrow \alpha A^{-1}DB + \beta C$$

$$C \leftarrow \alpha A^{-T}DB + \beta C$$

3. Right permutation of a sparse matrix (presently for JAD data structure only):

$$A \leftarrow AP$$

$$A \leftarrow AP^T$$

4. Integrity check of a sparse matrix representation:

$$(\text{NERR}) \leftarrow A$$

where  $\alpha$  and  $\beta$  are scalars,  $B$  and  $C$  are rectangular matrices,  $D$  is a (block) diagonal matrix, and  $A$  is sparse matrix. The sparse matrix  $A$  can be stored in one of a variety of sparse data structures as described in Section 3. NERR is the integer value which returns information about the integrity of a sparse matrix representation. Each of the above operations is available for single and double precision, real and complex data as appropriate.

## 2.1 Naming Conventions for Sparse BLAS Toolkit

Each Sparse BLAS Toolkit subroutine has a six-character name that is a function of the data type, the data structure of the sparse matrix and the type of operation. More precisely, each name is of the form  $XYYYZZ$  where

- X indicates the data type:
  - S - Single-precision real
  - C - Single-precision complex
  - D - Double-precision real
  - Z - Double-precision complex
- YYY indicates the data structure of the sparse matrix: (Note: This list may change. See Section 3.)
  - Point entry:
    - \* COO - Coordinate
    - \* CSC - Compressed sparse column
    - \* CSR - Compressed sparse row
    - \* DIA - Sparse diagonal
    - \* ELL - Ellpack/Itpack
    - \* JAD - Jagged diagonal
    - \* SKY - Skyline
  - Block entry:
    - \* BCO - Block coordinate
    - \* BSC - Block compressed sparse column
    - \* BSR - Block compressed sparse row
    - \* BDI - Block sparse diagonal

- \* BEL - Block Ellpack/Itpack
- \* VBR - Variable block compressed sparse row
- ZZ indicates the type of operation:
  - MM - Matrix-matrix product
  - SM - Solution of a triangular system with multiple right-hand-sides
  - RP - Right permutation of a sparse matrix (presently for JAD data structure only)
  - CK - Integrity check of a sparse matrix representation.

## 2.2 Argument Conventions for Sparse BLAS Toolkit

The argument conventions follow the spirit of the dense BLAS. The general order of arguments is:

1. Arguments specifying options.
2. Arguments specifying problem dimensions.
3. Input scalar associated with input matrices.
4. Description of sparse input matrices ( $args(A)$ , described below).
5. Description of dense input matrices.
6. Input scalar associated with input-output matrix.
7. Description of input-output matrix.
8. Error processing information.
9. Workspace.
10. Length of workspace.

In particular, we have

- Point entry matrix-matrix product:  
`XYYYMM( TRANSA, M, N, K, ALPHA, args(A), B, LDB, BETA, C, LDC, WORK, LWORK)`
- Block entry matrix-matrix product:  
`XYYYMM( TRANSA, MB, N, KB, ALPHA, args(A), B, BLDB, BETA, C, BLDC, WORK, LWORK)`
- Point entry solution of triangular system:  
`XYYYSM( TRANSA, M, N, UNITD, DV, ALPHA, args(A), B, LDB, BETA, C, LDC, WORK, LWORK)`

- Block entry solution of triangular system:  
`XYYYSM( TRANSA, MB, N, UNITD, DV, ALPHA, args(A), B, BLDB, BETA, C, BLDC, WORK, LWORK)`
- Right permutation of a sparse matrix:  
`XYYYRP( TRANSP, M, K, args(A), [ IPERM, ] WORK, LWORK)`
- Point entry integrity check of a sparse matrix representation:  
`XYYYCK( TRANSA, M, K, args(A), NERR, WORK, LWORK )`
- Block entry integrity check of a sparse matrix representation:  
`XYYYCK( TRANSA, MB, KB, args(A), NERR, WORK, LWORK )`

### 2.2.1 Common Arguments

The argument `TRANSA` is an integer argument. The possible values are

- `TRANSA` - Indicates how to operate with the sparse matrix.
  - 0 - Operate with the matrix (No-transpose).
  - 1 - Operate with the transpose of the matrix.
  - 2 - Operate with the conjugate transpose of the matrix.

2 is equivalent to 1 if the matrix is real.

The argument `M` is the number of rows in the matrix  $C$ , `N` is the number of columns in  $C$ , and `K` is the number of rows in  $B$ . `M` and `K` are the row and column dimensions of  $A$ , respectively, if `TRANSA` = 0 or the row and column dimensions of  $A^T$  if `TRANSA` = 1 or 2.

The argument `MB` is the number of block rows in the matrix  $C$ , and `KB` is the number of block rows in  $B$ . `MB` and `KB` are the block row and block column dimensions of  $A$ , respectively, if `TRANSA` = 0 or the block row and block column dimensions of  $A^T$  if `TRANSA` = 1 or 2. See Sections 3.4 and 3.5 for a full discussion of block entry matrices.

Negative or zero dimensions are allowed. However, if `M`, `N`, `K`, `MB`, or `KB` is less than or equal to zero, then no operations are performed.

The arguments `ALPHA` and `BETA` are both scalar inputs of the same data type as the matrices. These operands correspond to  $\alpha$  and  $\beta$ , respectively, as defined above.

The arguments `B` and `C` are rectangular arrays with first dimension `LDB` and `LDC`, respectively. If  $A$  is a constant block entry matrix, then `BLDB` and `BLDC` are the *block* first dimensions of  $B$  and  $C$ , respectively. This implies for a block entry dimension of `LB` that `LDB` = `BLDB`\*`LB` and `LDC` = `BLDC`\*`LB`. See Section 3.4 for details.

The argument `WORK` is an array of the same data type as  $A$  and of length `LWORK`. It is used as scratch space for optimizing performance. The minimal value for `LWORK` is  $\max(M, N)$  for point entry matrices and  $\max(MB * LB * LB, N)$  in the block entry case. The optimal value is returned in `WORK(1)` and can be obtained by reading `INT(WORK(1))`.

### 2.2.2 SM Arguments

The argument UNITD is an integer argument indicating whether or not the diagonal matrix  $D$  is unitary. The possible values for UNITD are

- 1 - Unitary diagonal matrix, i.e.,  $D$  is the identity. In this case the argument DV is ignored.
- 2 - Scale on the left (row scaling).
- 3 - Scale on the right (column scaling).

The argument DV is an array containing the diagonal entries of the (block) diagonal matrix  $D$ . It is described in more detail in Section 3.4.

### 2.2.3 RP Arguments

A permutation  $P$  is represented by an integer vector IPERM such that  $\text{IPERM}(i)$  is equal to the position of the only nonzero element in row  $i$  of  $P$ . For example, if

$$P = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix},$$

then  $\text{IPERM} = (3, 1, 2)$ . Note that, since IPERM is part of  $\text{args}(A)$  for the JAD data structure, there is no need to repeat it as an argument. However, for future extensions, we indicate where IPERM would be in the argument list.

The argument TRANSP is an integer argument. The possible values are

- TRANSP - Indicates how to operate with the permutation matrix.
  - 0 - Operate with the matrix (No-transpose).
  - 1 - Operate with the transpose of the matrix.

### 2.2.4 CK Arguments

The arguments NERR returns the number of errors detected in integrity check of a sparse matrix representation. Errors are processed by an error handling routine which prints the error information to standard error. The following checks are performed on the sparse matrix representation:

- Valid arguments.
  - Descriptors have valid values.
  - Pointers have non-decreasing order.
  - Indices are within problem dimensions.

- Floating-point values are valid, i.e., values represent finite floating-point number on the given architecture.
- Structural integrity.
  - Detects empty rows/columns or missing crucial entries, e.g., missing diagonal entry for a triangular matrix with non-unit diagonal.
  - Mismatch of descriptor and matrix graph, e.g., entries from the upper triangle are stored in for a matrix which is declared to be lower triangular.
- Numerical integrity.
  - Detects inappropriate zero values.
  - Detects extremely small values.
- Reducibility check. Checks to see if the matrix is structurally reducible, i.e., the matrix can be separated into two or more disjoint matrices. This step requires a workspace of size two times the minimal value for LWORK. If LWORK is too small, the reducibility check will not be performed.

### 2.2.5 Order of Arguments for $args(A)$

For all kernels, the data structure of  $A$  is determined by the `YYY` characters of the subroutine name. The argument  $args(A)$ , which is really a list of several arguments, will vary with the data structure. The general order of arguments is:

1. Descriptor array.
2. Array containing the non-zero values (entries) of input matrix.
3. First dimension of values array (if shape of this array differs from index arrays).
4. Array(s) containing indices corresponding to the entries of input matrix. If more than one index array:
  - (a) Arrays ordered in decreasing length.
  - (b) Arrays of same length are ordered with row indices first.
5. Length or first dimension of value/index arrays.
6. Array(s) containing pointers. If more than one pointer array:
  - (a) Arrays ordered in decreasing length.
  - (b) Arrays of same length are ordered with row pointers first, column pointers next and diagonal pointers last.
7. Length or first dimension of pointer arrays.
8. Array(s) representing left and/or right permutations of  $A$ . If more than one permutation array order left permutation array first.



<i>Data Structure</i>	<i>args(A)</i>
COO	DESCRA, VAL, INDX, JNDX, NNZ
CSC	DESCRA, VAL, INDX, PNTRB, PNTRE
CSR	DESCRA, VAL, INDX, PNTRB, PNTRE
DIA	DESCRA, VAL, LDA, IDIAG, NDIAG
ELL	DESCRA, VAL, INDX, LDA, MAXNZ
JAD	DESCRA, VAL, INDX, PNTR, MAXNZ, IPERM
SKY	DESCRA, VAL, PNTR
BCO	DESCRA, VAL, BINDX, BJNDX, BNNZ, LB
BSC	DESCRA, VAL, BINDX, BPNTRB, BPNTRE, LB
BSR	DESCRA, VAL, BINDX, BPNTRB, BPNTRE, LB
BDI	DESCRA, VAL, BLDA, IBDIAG, NBDIAG, LB
BEL	DESCRA, VAL, BINDX, BLDA, MAXBNZ, LB
VBR	DESCRA, VAL, INDX, BINDX, RPNTR, CPNTR, BPNTRB, BPNTRE

Table 1:  $args(A)$  for each data structure.

### 9. Block entry dimension.

The exact specification of  $args(A)$  for each supported data structure is in the following Section 3.

## 3 Storage Conventions for Sparse Matrices

In this section, we will describe the data structures and corresponding  $args(A)$ . Before continuing, one should note that the list of data structures for  $A$  presented here is not necessarily complete. We intend to add (or delete) items from this list if a justifiable argument can be presented. In particular, we will add a data structure if it can be shown that it offers a significant performance or algorithmic gains over the existing data structures for an important architecture or class of problems. We would rather not add data structures that are similar to the ones listed.

For example, there are several varieties of the compressed sparse row data structure (CSR) with little or no difference in performance between the variations and the CSR data structure presented here. Thus, we choose to avoid supporting all the variations.

Table 1 lists  $args(A)$  for all data structures. The following subsections describe the exact meaning of the arguments for each data structure.

### 3.1 Definitions and Notation

For the following discussion of data structures, we define some terms and notation for an  $m$  by  $k$  matrix  $A$ :

- Entry* - Any matrix coefficient which is handled explicitly (see [8]). There are two basic types of entries discussed in this proposal:
- *point* - Each entry is a single scalar value. Typically, the point entries of  $A$  are simply the non-zero elements of  $A$ .
  - *block* - Each entry is a dense rectangular matrix stored column by column. Block entries are derived from an assumed row and column partition of the matrix  $A$ .
- $NNZ(A)$  - The number of point entries of a matrix  $A$ .
- $NZE(A)$  - The set of point entries of a matrix  $A$ .
- $NZI(A)$  - The set of row indices corresponding to the point entries of a matrix  $A$ . The ordering of elements of  $NZI(A)$  should be the same as  $NZE(A)$ .
- $NZJ(A)$  - The set of column indices corresponding to the point entries of a matrix  $A$ . The ordering of elements of  $NZJ(A)$  should be the same as  $NZE(A)$ .
- $A_{i*}$  - The vector consisting of the elements from row  $i$  of the matrix  $A$ .
- $A_{*j}$  - The vector consisting of the elements from column  $j$  of the matrix  $A$ .
- $diag_i(A)$  - The vector consisting of the elements from a diagonal of  $A$ .
- $i < 0$  implies that  $diag_i(A)$  is a vector with the first  $|i|$  elements being unspecified or containing boundary information. The next  $q$  values of  $diag_i(A)$  are from the  $|i|^{th}$  diagonal below the main diagonal where  $q = \min(m - |i|, k)$
  - $i > 0$  implies that the first  $q$  elements,  $q = \min(m, k - i)$ , of  $diag_i(A)$  are from the  $i^{th}$  diagonal above the main diagonal of  $A$  and the remaining elements are unspecified or contain boundary information.
  - $i = 0$  implies  $diag_i(A)$  is the main diagonal of  $A$ .
- $PRF(A_{i*})$  - The set of elements from row  $i$  of  $A$  starting with the first point entry in row  $i$  up to and including the last point entry. Zero elements are included.
- $PRF(A_{*j})$  - The set of elements from column  $j$  of  $A$  starting with the first point entry in column  $j$  down to and including the last point entry. Zero elements are included.
- $\hat{A}$  - The matrix  $A$  with an assumed row and column partitioning.
- $BNNZ(\hat{A})$  - The number of block entries of a matrix  $A$  where  $\hat{A}$  denotes  $A$  with an assumed partition.
- $BNZE(\hat{A})$  - The set of block entries of a matrix  $A$ . Each block entry is a dense rectangular matrix stored column by column.

- $BNZJ(\hat{A})$  - The set of block column indices corresponding to the block entries of a matrix  $A$ . The ordering of elements of  $BNZJ(\hat{A})$  should be the same as  $BNZE(\hat{A})$ .
- $\hat{A}_{i*}$  - The block vector consisting of the elements from block row  $i$  of the partitioned matrix  $\hat{A}$ .
- $\hat{A}_{*j}$  - The block vector consisting of the elements from block column  $j$  of the partitioned matrix  $\hat{A}$ .
- $\bar{A}$  - The matrix  $PA$  where  $P$  is a permutation matrix.

- Note:**
1. The data structures presented here handle square or rectangular matrices, and allow for zero rows and columns.
  2. No order is assumed for storing the entries of the matrices except that for triangular matrices, we assume that the main diagonal elements are stored in proper relation to lower or upper triangle. For example, if a non-unit diagonal, upper triangular matrix is stored in CSR format, then we assume that the diagonal is the first entry in each row. This is necessary for efficient solution of triangular systems.

### 3.2 Descriptor Arguments

The first argument in  $args(A)$  is DESCRA. It is a nine-element integer array which describes the relevant characteristics of the matrix. The first three elements of DESCRA describe the matrix structure and type of main diagonal. The fourth through ninth elements of DESCRA can be used for data-structure-specific information. Presently DESCRA(5) indicates the ordering of block entry elements for constant block entry matrices (see Section 3.4) and DESCRA(6) is used in the case of the sparse diagonal data structures (DIA and BDI), where it indicates whether or not the matrix is a stencil-based operator. DESCRA(6) is also used by the COO, ELL, JAD, BCO and BEL data structures to assert that there are no repeated indices in the index arrays. The possible values of DESCRA(1)–DESCRA(3) are:

- DESCRA(1) - Matrix structure
  - 0 - General
  - 1 - Symmetric
  - 2 - Hermitian
  - 3 - Triangular
  - 4 - Anti-symmetric (Skew Symmetric/Hermitian)
  - 5 - Diagonal
- DESCRA(2) - Upper/Lower triangular indicator
  - 1 - Lower
  - 2 - Upper
- DESCRA(3) - Main diagonal type

- 0 - Non-unit
- 1 - Unit (diagonal elements are not stored)
- DESCRA(4) - Array base
  - 0 - Zero array base (typical for C/C++)
  - 1 - One array base (typical for Fortran)

**Note:**

1. For the solution of triangular systems, we only support DESCRA(1) = 3. We support all combinations of data and matrix structures for matrix multiplication.
2. The diagonal type for block entry matrices refers to the block diagonal. Currently we only support unit block diagonal for triangular solve routines. See Section 5.8 for a discussion of non-unit block diagonal entries.
3. If a matrix is declared to be anti-symmetric, we assume that if the diagonal entries are stored then they are explicitly zero.
4. The SKY data structure is not supported for a general matrix structure (DESCRA(1) = 0).
5. If DESCRA(4) = 1 then addresses based index vector values, e.g., INDX, would be assumed to start at one. Thus, a reference using  $INDX(i) = 1$  would point to the first element of an array. An array base of one is the default for Fortran programs. DESCRA(4) = 0 would indicate that a reference using  $INDX(i) = 1$  points to the second element. An array base of zero is assumed for C and C++ programs.

The remaining arguments in  $args(A)$  vary and are described below. For precision and conciseness, these arguments are described in terms of the notation defined above. Examples are given for each data structure to illustrate the definition. Section 3.3 describes each point entry data structure. Section 3.4 describes the BSR data structure and its relationship to CSR. Other constant block size data structures are not discussed since they are straight forward extensions of the point entry equivalents. In Section 3.5 we describe the VBR data structure.

### 3.3 Point Entry Data Structures

#### 3.3.1 COO - Coordinate

Three arrays are required for the COO format:

- VAL - A scalar array of length  $NNZ(A)$  consisting of  $NZE(A)$ , the entries of  $A$ , in any order.
- INDX - An integer array of length  $NNZ(A)$  consisting of  $NZI(A)$ , the corresponding row indices of the entries of  $A$ .

- JNDX - An integer array of length  $NNZ(A)$  consisting of  $NZJ(A)$ , the corresponding column indices of the entries of  $A$ .

We also need the argument  $NNZ = NNZ(A)$ .

For example, let

$$A = \begin{pmatrix} 11 & 0 & 13 & 14 & 0 \\ 0 & 0 & 23 & 24 & 0 \\ 31 & 32 & 33 & 34 & 0 \\ 0 & 42 & 0 & 44 & 0 \\ 51 & 52 & 0 & 0 & 55 \end{pmatrix}, \quad (1)$$

then one representation of  $A$  in COO format (assuming  $DESCRA(4) = 1$ ) would be as follows.

$$\begin{aligned} \text{VAL} &= ( 11 \ 51 \ 31 \ 32 \ 34 \ 52 \ 13 \ 23 \ 33 \ 14 \ 24 \ 42 \ 55 \ 44 ), \\ \text{INDX} &= ( 1 \ 5 \ 3 \ 3 \ 3 \ 5 \ 1 \ 2 \ 3 \ 1 \ 2 \ 4 \ 5 \ 4 ), \\ \text{JNDX} &= ( 1 \ 1 \ 1 \ 2 \ 4 \ 2 \ 3 \ 3 \ 3 \ 4 \ 4 \ 2 \ 5 \ 4 ). \end{aligned}$$

If  $A$  is symmetric (or Hermitian, triangular, or anti-symmetric) then we only need to store the lower (or upper) triangle. In this case we have (for the lower triangle)

$$\begin{aligned} \text{VAL} &= ( 11 \ 51 \ 31 \ 32 \ 52 \ 33 \ 42 \ 55 \ 44 ), \\ \text{INDX} &= ( 1 \ 5 \ 3 \ 3 \ 5 \ 3 \ 4 \ 5 \ 4 ), \\ \text{JNDX} &= ( 1 \ 1 \ 1 \ 2 \ 2 \ 3 \ 2 \ 5 \ 4 ). \end{aligned}$$

**Note:** If  $A$  is Hermitian and the diagonal of  $A$  is stored, then we will assume that the imaginary part of the diagonal is zero. Similarly, if  $A$  is anti-symmetric and the diagonal is stored, we will assume it is zero.

For the COO data structure,  $DESCRA(6)$  is used to indicate whether or not  $INDX$  has repeated values. In this case, if it is known that all columns of  $INDX$  have no repeated values, then there is no danger of overwriting values when executed in parallel. Performance improvements can be substantial if the user can guarantee there are no repeated indices and indicates so by setting  $DESCRA(6) = 1$ . The possible values for  $DESCRA(6)$  are:

- 0 - Unknown.
- 1 - No repeated indices.

### 3.3.2 CSC - Compressed Sparse Column

A matrix  $A$  is stored in the CSC format using four arrays.

- VAL - A scalar array of length  $NNZ(A)$  consisting of the entries of  $A$ :

$$\text{VAL} = (NZE(A_{*1}), NZE(A_{*2}), \dots, NZE(A_{*k})).$$

- **INDX** - An integer array of length  $NNZ(A)$  consisting of the row indices of entries of  $A$ :

$$\text{INDX} = (\text{NZI}(A_{*1}), \text{NZI}(A_{*2}), \dots, \text{NZI}(A_{*k})).$$

- **PNTRB** - An integer array of length  $k$  such that  $\text{PNTRB}(j) - \text{PNTRB}(1) + 1$  points to the location in **VAL** of the first element of  $\text{NZE}(A_{*j})$ . If  $\text{NZE}(A_{*j})$  is empty then set  $\text{PNTRB}(j) = \text{PNTRB}(j + 1)$ .
- **PNTRE** - An integer array of length  $k$  such that  $\text{PNTRE}(j) - \text{PNTRB}(1)$  points to the location in **VAL** of the last element of  $\text{NZE}(A_{*j})$ . If  $\text{NZE}(A_{*j})$  is empty then set  $\text{PNTRE}(j) = \text{PNTRB}(j)$ .

For example, the CSC representation of the matrix in Equation 1 would be:

$$\begin{aligned} \text{VAL} &= ( 11 \ 31 \ 51 \ 32 \ 42 \ 52 \ 13 \ 23 \ 33 \ 14 \ 24 \ 34 \ 44 \ 55 ), \\ \text{INDX} &= ( 1 \ 3 \ 5 \ 3 \ 4 \ 5 \ 1 \ 2 \ 3 \ 1 \ 2 \ 3 \ 4 \ 5 ), \\ \text{PNTRB} &= ( 1 \ 4 \ 7 \ 10 \ 14 ), \\ \text{PNTRE} &= ( 4 \ 7 \ 10 \ 14 \ 15 ), \end{aligned}$$

- Note:**
1. The actual values in **PNTRB** and **PNTRE** are not important, only their relative value from  $\text{PNTRB}(1)$ . This allows greater flexibility for the user. In particular, it is common to construct pointer arrays starting at 0. For Equation 1, we would then define **PNTRB** and **PNTRE** as

$$\begin{aligned} \text{PNTRB} &= ( 0 \ 3 \ 6 \ 9 \ 13 ), \\ \text{PNTRE} &= ( 3 \ 6 \ 9 \ 13 \ 14 ). \end{aligned}$$

For the rest of the examples in this proposal, we use  $\text{PNTRB}(1) = 1$ , but this is only for convenience.

2. One can represent **PNTRB** and **PNTRE** for this example by using a single array **PNTR** where

$$\text{PNTR} = ( 1 \ 4 \ 7 \ 10 \ 14 \ 15 ),$$

in which case  $\text{PNTRB}(1:5) = \text{PNTR}(1:5)$  and  $\text{PNTRE}(1:5) = \text{PNTR}(2:6)$ . In this way, the change from a single **PNTR** array in previous proposals to the two arrays **PNTRB** and **PNTRE** is compatible for user having only the **PNTR** array.

3. The two array approach to pointers offers much more flexibility than the previous one array approach. For example, if we define the array

$$\text{PNTRD} = ( 1 \ 4 \ 9 \ 12 \ 13 ),$$

to point to the diagonal elements, then letting  $\text{PNTRB} = \text{PNTRD}$  we can use just the lower triangular part of the general CSC representation without modifying or copying the other data structures.

4. A second example of increased flexibility is that there is no longer an implicit storage association between contiguous columns. The columns can be specified

in any order. In particular, for the example above it is possible to put column one as the last column stored:

$$\begin{aligned} \text{VAL} &= ( 32 \ 42 \ 52 \ 13 \ 23 \ 33 \ 14 \ 24 \ 34 \ 44 \ 55 \ 11 \ 31 \ 51 ), \\ \text{INDX} &= ( 3 \ 4 \ 5 \ 1 \ 2 \ 3 \ 1 \ 2 \ 3 \ 4 \ 5 \ 1 \ 3 \ 5 ), \\ \text{PNTRB} &= ( 12 \ 1 \ 4 \ 7 \ 11 ), \\ \text{PNTRE} &= ( 15 \ 4 \ 7 \ 11 \ 12 ), \end{aligned}$$

If  $A$  is symmetric then we only need to store the lower (or upper) triangle. In this case we have (for the lower triangle)

$$\begin{aligned} \text{VAL} &= ( 11 \ 31 \ 51 \ 32 \ 42 \ 52 \ 33 \ 44 \ 55 ), \\ \text{INDX} &= ( 1 \ 3 \ 5 \ 3 \ 4 \ 5 \ 3 \ 4 \ 5 ), \\ \text{PNTRB} &= ( 1 \ 4 \ 7 \ 8 \ 9 ). \\ \text{PNTRE} &= ( 4 \ 7 \ 8 \ 9 \ 10 ). \end{aligned}$$

### 3.3.3 CSR - Compressed Sparse Row (Point entry form)

A matrix  $A$  is stored in the CSR data structure using four arrays.

- VAL - A scalar array of length  $NNZ(A)$  consisting of the entries of  $A$ :

$$\text{VAL} = (NZE(A_{1*}), NZE(A_{2*}), \dots, NZE(A_{m*})).$$

- INDX - An integer array of length  $NNZ(A)$  consisting of the column indices of entries of  $A$ :

$$\text{INDX} = (NZJ(A_{1*}), NZJ(A_{2*}), \dots, NZJ(A_{m*})).$$

- PNTRB - An integer array of length  $m$  such that  $\text{PNTRB}(i) - \text{PNTRB}(1) + 1$  points to the location in VAL of the first element of  $NZE(A_{i*})$ . If  $NZE(A_{i*})$  is empty then set  $\text{PNTRB}(i) = \text{PNTRB}(i + 1)$ .
- PNTRE - An integer array of length  $m$  such that  $\text{PNTRE}(i) - \text{PNTRB}(1)$  points to the location in VAL of the last element of  $NZE(A_{i*})$ . If  $NZE(A_{i*})$  is empty then set  $\text{PNTRE}(i) = \text{PNTRB}(i)$ .

For example, the CSR representation of the matrix in Equation 1 would be:

$$\begin{aligned} \text{VAL} &= ( 11 \ 13 \ 14 \ 23 \ 24 \ 31 \ 32 \ 33 \ 34 \ 42 \ 44 \ 51 \ 52 \ 55 ), \\ \text{INDX} &= ( 1 \ 3 \ 4 \ 3 \ 4 \ 1 \ 2 \ 3 \ 4 \ 2 \ 4 \ 1 \ 2 \ 5 ), \\ \text{PNTRB} &= ( 1 \ 4 \ 6 \ 10 \ 12 ), \\ \text{PNTRE} &= ( 4 \ 6 \ 10 \ 12 \ 15 ). \end{aligned}$$

If  $A$  is symmetric then we only need to store the lower (or upper) triangle. Considering only the lower triangle of  $A$ , we have

$$\begin{aligned} \text{VAL} &= ( 11 \ 31 \ 32 \ 33 \ 42 \ 44 \ 51 \ 52 \ 55 ), \\ \text{INDX} &= ( 1 \ 1 \ 2 \ 3 \ 2 \ 4 \ 1 \ 2 \ 5 ), \\ \text{PNTRB} &= ( 1 \ 2 \ 2 \ 5 \ 7 ). \\ \text{PNTRE} &= ( 2 \ 2 \ 5 \ 7 \ 10 ). \end{aligned}$$

**Note:** See the discussion on the use of PNTRB and PNTRE in the previous discussion of the CSC data structure.

### 3.3.4 DIA - Sparse Diagonal (Point entry form)

Let  $LDA \geq \min(m, k)$ , and let NDIAG denote the number of non-zero diagonals of  $A$ . Two arrays are required for the DIA format:

- VAL - A two-dimensional LDA-by-NDIAG scalar array consisting of the NDIAG non-zero diagonals of  $A$  in any order.
- IDIAG - An integer array of length NDIAG consisting of the corresponding indices of the non-zero diagonals of  $A$  in VAL. Thus, if  $IDIAG(j) = i$  then the  $j^{th}$  column of VAL contains  $diag_i(A)$ .

For example, let

$$A = \begin{pmatrix} 11 & 0 & 13 & 0 & 0 \\ 21 & 0 & 0 & 24 & 0 \\ 31 & 32 & 33 & 0 & 35 \\ 0 & 42 & 0 & 44 & 0 \\ 0 & 0 & 53 & 0 & 55 \end{pmatrix},$$

then  $A$  would be stored in DIA format as follows.

$$VAL = \begin{pmatrix} * & * & 11 & 13 \\ * & 21 & 0 & 24 \\ 31 & 32 & 33 & 35 \\ 42 & 0 & 44 & * \\ 53 & 0 & 55 & * \end{pmatrix},$$

$$IDIAG = \begin{pmatrix} -2 & -1 & 0 & 2 \end{pmatrix}.$$

For the DIA data structure, DESCRA(6) is used to determine how we should treat the ‘\*’ elements. This is useful if the operator is truly grid-based and boundary values are stored in the matrix data structure. If  $DESCRA(6) = 1$ , then the ‘\*’ elements are accessed and assumed to hold boundary value information from a stencil operator. Note also that in this case elements of  $B$  are accessed which fall outside the normal domain. For instance, in the above example, the normal domain of each column,  $B(*, j)$ , of  $B$  is 1 through 5, but the first ‘\*’ element of VAL will be multiplied by  $B(-1, j)$  and added to  $C(1, j)$ . The last ‘\*’ element will be multiplied by  $B(7, j)$  and added to  $C(5, j)$ . If  $DESCRA(6) = 0$ , then ‘\*’ elements are unspecified and unused. If  $A$  is symmetric then we only need to store diagonals from the lower (or upper) triangle of  $A$ .

### 3.3.5 ELL - Ellpack/Itpack (Point entry form)

Let MAXNZ denote the maximum of the number of entries in each row of  $A$ , i.e.,

$$MAXNZ = MAXNZ(A) = \max_{1 \leq i \leq m} NNZ(A_{i*}),$$

and let  $LDA \geq m$ . A matrix  $A$  is stored in the ELL format using two arrays.



- VAL - A two-dimensional LDA-by-MAXNZ scalar array such that row  $i$  of VAL consists of  $NZE(A_{i*})$  padded by zero values if the length of  $NZE(A_{i*})$  is less than MAXNZ.
- INDX - A two-dimensional LDA-by-MAXNZ integer array such that row  $i$  of INDX consists of  $NZJ(A_{i*})$  padded by the integer value  $i$  if  $NNZ(A_{i*})$  is less than MAXNZ.

For example, the ELL representation of the matrix in Equation 1 would be:

$$\text{VAL} = \begin{pmatrix} 11 & 13 & 14 & 0 \\ 23 & 24 & 0 & 0 \\ 31 & 32 & 33 & 34 \\ 42 & 44 & 0 & 0 \\ 51 & 52 & 55 & 0 \end{pmatrix},$$

$$\text{INDX} = \begin{pmatrix} 1 & 3 & 4 & 1 \\ 3 & 4 & 2 & 2 \\ 1 & 2 & 3 & 4 \\ 2 & 4 & 4 & 4 \\ 1 & 2 & 5 & 5 \end{pmatrix}.$$

If  $A$  is symmetric then we only need to store entries from the lower (or upper) triangle of  $A$ .

**Note:** If  $A$  is lower triangular and non-unit, then all diagonal entries must be in the last column of value and all padding entries must precede the diagonal.

For the ELL data structure, DESCRA(6) is used to indicate whether or not columns of INDX have repeated values. This can be useful information when computing with the transpose or symmetric form. In this case, if it is known that all columns of INDX have no repeated values, then there is no danger of overwriting values when executed in parallel. Performance improvements can be substantial if the user can guarantee there are no repeated indices and indicates so by setting DESCRA(6) = 1. The possible values for DESCRA(6) are:

- 0 - Unknown.
- 1 - No repeated indices.

### 3.3.6 JAD - Jagged Diagonal (Point entry form)

The JAD format requires the specification of a (non-unique) permutation matrix  $P$  which permutes the rows of  $A$  in descending order of  $NNZ(A_{i*})$ . As mentioned in Section 2.2.3, a permutation is represented by an integer vector IPERM. Let  $\bar{A} = PA$  be the permuted matrix and let  $i'$  denote the row of  $\bar{A}$  corresponding to row  $i$  of  $A$ . Four arrays are needed:

- VAL - A scalar array of length  $NNZ(A)$  consisting of the entries of  $A$ . The first element of VAL is the first element of  $NZE(\bar{A}_{1*})$  followed by the first element of

$NZE(\bar{A}_{2*})$  and so on through the first value of  $NZE(\bar{A}_{m'*})$  where  $m'$  is the last row of  $\bar{A}$  that has non-zero entries (usually  $m' = m$ ). The  $m' + 1^{st}$  element of VAL is the second element of  $NZE(\bar{A}_{1*})$  and so on.

- **INDX** - An integer array of length  $NNZ(A)$  consisting of the corresponding column indices of the entries of  $A$ .
- **PNTR** - An integer array of length  $MAXNZ + 1$ , where  $MAXNZ$  is as defined in Section 3.3.5, such that  $PNTR(j) - PNTR(1) + 1$  points to the location in VAL of the  $j^{th}$  element of  $NZE(\bar{A}_{1*})$ .  $PNTR(MAXNZ + 1)$  is set to the value  $NNZ(A) + PNTR(1)$ .
- **IPERM** - An integer array of length  $m$  such that  $i = IPERM(i')$ . **IPERM** is used to determine the order in which rows of  $C$  are updated. If  $IPERM(1) = 0$ , then we assume by convention that  $P = I$ . One should note that **IPERM** represents  $P^T$ .

For example, an appropriate permutation (and its inverse) for the matrix in equation 1 would be:

$$P = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix}, P^T = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \end{pmatrix},$$

and  $\bar{A}$  would be:

$$\bar{A} = \begin{pmatrix} 31 & 32 & 33 & 34 & 0 \\ 11 & 0 & 13 & 14 & 0 \\ 51 & 52 & 0 & 0 & 55 \\ 0 & 0 & 23 & 24 & 0 \\ 0 & 42 & 0 & 44 & 0 \end{pmatrix}.$$

The elements of the JAD arrays would be:

$$\begin{aligned} \text{VAL} &= ( 31 \ 11 \ 51 \ 23 \ 42 \ 32 \ 13 \ 52 \ 24 \ 44 \ 33 \ 14 \ 55 \ 34 ), \\ \text{INDX} &= ( 1 \ 1 \ 1 \ 3 \ 2 \ 2 \ 3 \ 2 \ 4 \ 4 \ 3 \ 4 \ 5 \ 4 ), \\ \text{PNTR} &= ( 1 \ 6 \ 11 \ 14 \ 15 ), \\ \text{IPERM} &= ( 3 \ 1 \ 5 \ 2 \ 4 ). \end{aligned}$$

If  $A$  is symmetric then we only need to store the lower (or upper) triangle. In this case,  $P$  is computed for the triangular part and will typically be different than  $P$  in the general case.

For the JAD data structure, **DESCRA(6)** is used to indicate whether or not there are repeated values in **INDX** in the section of **INDX** starting at **PNTR(i)** and ending at **PNTR(i+1)-1**, for each  $i = 1, MAXNZ$ . This can be useful information when computing with the transpose or symmetric form. In this case, if it is known that each of these sections of **INDX** have no repeated values, then there is no danger of overwriting values when executed in parallel. Performance improvements can be substantial if the user can guarantee there are no repeated indices and indicates so by setting **DESCRA(6) = 1**. The possible values for **DESCRA(6)** are:

- 0 - Unknown.
- 1 - No repeated indices.

**Symmetrically Permuted JAD** To avoid permuting the rows of  $C$  each time a JAD kernel is called, one can explicitly permute the columns of  $\bar{A}$  by calling

DJADRP(0, 5, 5, DESCRA, VAL, INDX, PNTR, MAXNZ, IPERM, WORK, LWORK).

This call will compute

$$\bar{A}P^T = \begin{pmatrix} 33 & 31 & 0 & 32 & 34 \\ 13 & 11 & 0 & 0 & 14 \\ 0 & 51 & 55 & 52 & 0 \\ 23 & 0 & 0 & 0 & 24 \\ 0 & 0 & 0 & 42 & 44 \end{pmatrix}.$$

The elements of the JAD arrays for this fully permuted system would be:

$$\begin{aligned} \text{VAL} &= ( 33 \ 13 \ 51 \ 23 \ 42 \ 31 \ 11 \ 55 \ 24 \ 44 \ 32 \ 14 \ 52 \ 34 ), \\ \text{INDX} &= ( 1 \ 1 \ 2 \ 1 \ 4 \ 2 \ 2 \ 3 \ 5 \ 5 \ 4 \ 5 \ 4 \ 5 ), \\ \text{PNTR} &= ( 1 \ 6 \ 11 \ 14 \ 15 ). \end{aligned}$$

IPERM is no longer needed and an explicit zero value can be put in its place. Given the fully permuted matrix  $A$  and replacing  $C$  with  $PC$  and  $B$  with  $PB$ , we have an equivalent system to the original  $A$ ,  $B$  and  $C$  except in a permuted coordinate system.

### 3.3.7 SKY - Skyline (Point entry form)

A matrix  $A$  is stored in the SKY data structure using two arrays.

- If DESCRA(2) = 1, then we have:

- VAL - A scalar array consisting of:

$$\text{VAL} = (\text{PRF}(A_{1*}), \text{PRF}(A_{2*}), \dots, \text{PRF}(A_{m*}))$$

- PNTR - An integer array of length  $m + 1$  such that  $\text{PNTR}(i) - \text{PNTR}(1) + 1$  points to the location in VAL of the first element of  $\text{PRF}(A_{i*})$ .  $\text{PNTR}(m + 1)$  is set to the value  $\text{NNZ}(A) + \text{PNTR}(1)$ . If  $\text{PRF}(A_{i*})$  is empty, then set  $\text{PNTR}(i) = \text{PNTR}(i + 1)$ .

- If DESCRA(2) = 2, then:

- VAL - A scalar array consisting of:

$$\text{VAL} = (\text{PRF}(A_{*1}), \text{PRF}(A_{*2}), \dots, \text{PRF}(A_{*k}))$$

- PNTR - An integer array of length  $k + 1$  such that  $\text{PNTR}(j) - \text{PNTR}(1) + 1$  points to the location in VAL of the first element of  $\text{PRF}(A_{*j})$ .  $\text{PNTR}(k + 1)$  is set to the value  $\text{NNZ}(A) + \text{PNTR}(1)$ . If  $\text{PRF}(A_{*j})$  is empty, then set  $\text{PNTR}(j) = \text{PNTR}(j + 1)$ .

**Note:** This data structure is not supported for a general matrix structure (DESCRA(1) = 0).

For example, let

$$A = \begin{pmatrix} 11 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 31 & 32 & 33 & 0 & 0 \\ 0 & 42 & 0 & 44 & 0 \\ 51 & 52 & 0 & 0 & 55 \end{pmatrix},$$

then  $A$  would be stored in SKY format as follows (assuming DESCRA(2) = 1).

$$\begin{aligned} \text{VAL} &= ( 11 \ 31 \ 32 \ 33 \ 42 \ 0 \ 44 \ 51 \ 52 \ 0 \ 0 \ 55 ), \\ \text{PNTR} &= ( 1 \ 2 \ 2 \ 5 \ 8 \ 13 ). \end{aligned}$$

### 3.4 Constant Block Entry Data Structures

Each of the above data structures allows the matrix to have block entries where each entry of the matrix is an LB-by-LB dense block. Systems of this form typically arise when there are multiple unknowns per gridpoint of a discretized partial differential equation. By exploiting this property, we can reduce the amount of integer storage, increase the performance of the kernels and allow numerical algorithms to perform analysis on the block entries. Typically LB is a small number, less than ten, determined by the number of quantities measured at each gridpoint, e.g., velocity, pressure, viscosity, etc. However, there are cases where LB can be on the order of 100 as it is in the case of complex chemically reacting flows.

There are two common ways to order the elements of block entries in memory. One is to order the elements of each block entry contiguously, e.g., if  $LB = 2$  then, for each block entry, we store element (1,1) followed by element (2,1) followed by (1,2) and (2,2). Thus, each block entry is stored in  $LB^2$  contiguous memory locations. Another way is to store the (1,1) elements of all block entries contiguously followed by all the (2,1) elements and so on.

The first ordering is particularly suitable for cache-based microprocessors. It is also the natural data structure for performing analysis on the block entries since each entry is a dense LB-by-LB matrix which could be passed, for example, to an LAPACK routine. However, this ordering is typically very bad for vector computers, especially if LB is small, since it forces vectorization with non-unit stride memory access. The second ordering allows vectorization with unit-stride memory access, but loses the cache performance and the ability to pass block entries to other subroutines as standard dense arrays.

Since both orderings of block entry elements appear to be of importance, we provide a way for the user to choose the ordering. If DESCRA(5) = 0, then the sparse BLAS toolkit routines assume the first ordering, i.e., the elements of each block entry are stored contiguously. If DESCRA(5) = 1, then the second ordering is assumed.

Note that, for the sparse triangular solve routines, if the argument UNITD = 0 (non-unit diagonal), then the argument DV contains the block entries of the main diagonal of a block diagonal matrix. The dimensions of DV are determined by DESCRA(5). If

DESCRA(5) = 0, then the dimensions of DV are (LB,MB). Otherwise they are (MB,LB). Also, the dimensions of B and C are (LB,BLDB,NB) and (LB,BLDC,NB), respectively, if DESCRA(5) = 0, or (BLDB,LB,NB) and (BLDC,LB,NB) if DESCRA(5) = 1.

Although all of the point entry data structures discussed above have a constant block entry analogue, we only choose to support some of them. Below we present the BSR data structure which is the block entry form of the CSR data structure. The block entry forms of the other data structures are analogous.

**BSR - Constant Block Compressed Sparse Row Data Structure.** Let  $l = \text{LB}$  denote the dimension of the block entries of  $A$ ,  $\hat{A}$  denote  $A$  along with its row and column partition and  $m_b$  and  $k_b$  denote the block row and column dimensions, respectively.  $\hat{A}$  is stored in BSR format using three arrays. The order of elements depends on DESCRA(5).

- If DESCRA(5) = 0, we have:
  - VAL - An array of dimension  $(l, l, \text{BNNZ}(\hat{A}))$  consisting of the block entries of  $\hat{A}$ :
$$\text{VAL} = (\text{BNZE}(\hat{A}_{1*}), \text{BNZE}(\hat{A}_{2*}), \dots, \text{BNZE}(\hat{A}_{m_b*})).$$
  - BINDX - An integer array of length  $\text{BNNZ}(\hat{A})$  consisting of the block column indices of the block entries of  $\hat{A}$ :
$$\text{BINDX} = (\text{BNZJ}(\hat{A}_{1*}), \text{BNZJ}(\hat{A}_{2*}), \dots, \text{BNZJ}(\hat{A}_{m_b*})).$$
  - BPNTRB - An integer array of length  $m_b$  such that  $\text{BPNTRB}(i) - \text{BPNTRB}(1) + 1$  points to the location in BINDX of the the first block entry of  $\text{BNZE}(\hat{A}_{i*})$ . If  $\text{BNZE}(\hat{A}_{i*})$  is empty then set  $\text{BPNTRB}(i) = \text{BPNTRB}(i + 1)$ .
  - BPNTRE - An integer array of length  $m_b$  such that  $\text{BPNTRE}(i) - \text{BPNTRB}(1)$  points to the location in BINDX of the the last block entry of  $\text{BNZE}(\hat{A}_{i*})$ . If  $\text{BNZE}(\hat{A}_{i*})$  is empty then set  $\text{BPNTRE}(i) = \text{BPNTRB}(i)$ .

For example, let

$$\hat{A} = \left( \begin{array}{cc|cc|cc} 11 & 12 & 0 & 0 & 15 & 16 \\ 21 & 22 & 0 & 0 & 25 & 26 \\ \hline 0 & 0 & 33 & 0 & 35 & 36 \\ 0 & 0 & 43 & 44 & 45 & 46 \\ \hline 51 & 52 & 0 & 0 & 0 & 0 \\ 61 & 62 & 0 & 0 & 0 & 0 \end{array} \right), \quad (2)$$

then, since  $l = 2$ ,  $\hat{A}$  would be stored in BSR format as follows.

$$\text{VAL}(1 : 2, 1 : 2, 1) = \begin{pmatrix} 11 & 12 \\ 21 & 22 \end{pmatrix}, \quad \text{VAL}(1 : 2, 1 : 2, 2) = \begin{pmatrix} 15 & 16 \\ 25 & 26 \end{pmatrix},$$

$$\text{VAL}(1 : 2, 1 : 2, 3) = \begin{pmatrix} 33 & 0 \\ 43 & 44 \end{pmatrix}, \quad \text{VAL}(1 : 2, 1 : 2, 4) = \begin{pmatrix} 35 & 36 \\ 45 & 46 \end{pmatrix},$$

$$\text{VAL}(1 : 2, 1 : 2, 5) = \begin{pmatrix} 51 & 52 \\ 61 & 62 \end{pmatrix},$$

$$\begin{aligned}
\text{BINDX} &= ( 1 \ 3 \ 2 \ 3 \ 1 \ ), \\
\text{BPNTRB} &= ( 1 \ 3 \ 5 \ ). \\
\text{BPNTRE} &= ( 3 \ 5 \ 6 \ ).
\end{aligned}$$

The order of elements of VAL in memory would be:

$$(11, 21, 12, 22, 15, 25, 16, 26, 33, 43, 0, 44, 35, 45, 36, 46, 51, 61, 52, 62)$$

- If  $\text{DESCRA}(5) = 1$ , we have:

– VAL - An array of dimension  $(\text{BNNZ}(\hat{A}), l, l)$  consisting of the entries of  $\hat{A}$ :

$$\begin{aligned}
\text{VAL} &= ((\text{BNZE}(\hat{A}_{1*})(1, 1)), (\text{BNZE}(\hat{A}_{2*})(1, 1)), \dots, (\text{BNZE}(\hat{A}_{m_b*})(1, 1)), \\
&\quad (\text{BNZE}(\hat{A}_{1*})(2, 1)), (\text{BNZE}(\hat{A}_{2*})(2, 1)), \dots, (\text{BNZE}(\hat{A}_{m_b*})(2, 1)), \\
&\quad \vdots
\end{aligned}$$

$$(\text{BNZE}(\hat{A}_{1*})(l, l)), (\text{BNZE}(\hat{A}_{2*})(l, l)), \dots, (\text{BNZE}(\hat{A}_{m_b*})(l, l))),$$

where  $\text{BNZE}(\hat{A}_{i*})(i, j)$  denotes the  $(i, j)$  elements of the block entries.

– BINDX - An integer array of length  $\text{BNNZ}(\hat{A})$  consisting of the block column indices of the block entries of  $A$ :

$$\text{BINDX} = (\text{BNZJ}(\hat{A}_{1*}), \text{BNZJ}(\hat{A}_{2*}), \dots, \text{BNZJ}(\hat{A}_{m_b*})).$$

- BPNTRB - An integer array of length  $m_b$  such that  $\text{BPNTRB}(i) - \text{BPNTRB}(1) + 1$  points to the location in BINDX of the first block entry of  $\text{BNZE}(\hat{A}_{i*})$ . If  $\text{BNZE}(\hat{A}_{i*})$  is empty then set  $\text{BPNTRB}(i) = \text{BPNTRB}(i + 1)$ .
- BPNTRE - An integer array of length  $m_b$  such that  $\text{BPNTRE}(i) - \text{BPNTRB}(1)$  points to the location in BINDX of the last block entry of  $\text{BNZE}(\hat{A}_{i*})$ . If  $\text{BNZE}(\hat{A}_{i*})$  is empty then set  $\text{BPNTRE}(i) = \text{BPNTRB}(i)$ .

The matrix  $\hat{A}$  in Equation 2 would be stored in BSR format as follows.

$$\text{VAL}(1, 1 : 2, 1 : 2) = \begin{pmatrix} 11 & 12 \\ 21 & 22 \end{pmatrix}, \quad \text{VAL}(2, 1 : 2, 1 : 2) = \begin{pmatrix} 15 & 16 \\ 25 & 26 \end{pmatrix},$$

$$\text{VAL}(3, 1 : 2, 1 : 2) = \begin{pmatrix} 33 & 0 \\ 43 & 44 \end{pmatrix}, \quad \text{VAL}(4, 1 : 2, 1 : 2) = \begin{pmatrix} 35 & 36 \\ 45 & 46 \end{pmatrix},$$

$$\text{VAL}(5, 1 : 2, 1 : 2) = \begin{pmatrix} 51 & 52 \\ 61 & 62 \end{pmatrix},$$

$$\begin{aligned}
\text{BINDX} &= ( 1 \ 3 \ 2 \ 3 \ 1 \ ), \\
\text{BPNTRB} &= ( 1 \ 3 \ 5 \ ). \\
\text{BPNTRE} &= ( 3 \ 5 \ 6 \ ).
\end{aligned}$$

The order of elements of VAL in memory would be:

$$(11, 15, 33, 35, 51, 21, 25, 43, 45, 61, 12, 16, 0, 36, 52, 22, 26, 44, 46, 62)$$

If  $A$  is symmetric then we only need to store the lower (or upper) triangle. If the diagonal block entries are stored, we assume that the *entire* square block entry is stored.

### 3.5 Variable Block Entry Data Structures

It is often the case for problems with multiple unknowns per node, that the number of unknowns per node will vary, generating a matrix with a block structure where the size of the blocks varies correspondingly. If the variation in size is small, it may be advantageous to make all block rows the same size by adding identity equations to smaller blocks. However, often the the variation in size is large. For this case, we provide the variable block compressed sparse row (VBR) data structure.

**Variable Block Compressed Sparse Row Data Structure.** We could consider a variety of data structures to support problems with variable block matrices based on the point entry data structures discussed in Section 3.3. However, since to our knowledge there are few existing variable block data structures, we intend to restrict ourselves just in one data structure: the variable block compressed sparse row (VBR) which is a generalization of the supernodal data structures presented in [10, 11].

Formally we define the VBR data structure as follows. Consider an  $m$ -by- $k$  sparse matrix  $A$  along with a row partition  $P_r = \{i_1, i_2, \dots, i_{m_b+1}\}$  and column partition  $P_c = \{j_1, j_2, \dots, j_{k_b+1}\}$  such that  $i_1 = j_1 = 1$ ,  $i_{m_b+1} = m + 1$ ,  $j_{k_b+1} = k + 1$ ,  $i_p < i_{p+1}$ ,  $\forall p$ , and  $j_q < j_{q+1}$ ,  $\forall q$ . We denote  $A$ , along with its assumed row and column partitions  $P_r$  and  $P_c$ , by  $\hat{A}$ . Then we have the following equality

$$\hat{A} = \begin{pmatrix} \hat{A}_{11} & \hat{A}_{12} & \dots & \hat{A}_{1k_b} \\ \hat{A}_{21} & \hat{A}_{22} & & \vdots \\ \vdots & & \ddots & \vdots \\ \hat{A}_{m_b1} & \dots & \dots & \hat{A}_{m_bk_b} \end{pmatrix},$$

where  $\hat{A}_{pq}$  is of dimension  $(i_{p+1} - i_p)$ -by- $(j_{q+1} - j_q)$ .

If  $\hat{A}_{ij} \neq 0$  then  $\hat{A}_{ij}$  is a block entry of  $\hat{A}$ . In the VBR format, the block entries of  $\hat{A}$  are stored block row by block row. Each block entry is stored as a dense matrix in standard column major form. One floating point array and three to five integer arrays are used to store the matrix:

- VAL - A scalar array of length  $NNZ(A)$  consisting of the block entries of  $A$ :

$$\text{VAL} = (\text{BNZE}(\hat{A}_{1*}), \text{BNZE}(\hat{A}_{2*}), \dots, \text{BNZE}(\hat{A}_{m_b*})),$$

where each block entry is a dense rectangular matrix stored column by column.

- INDX - An integer array of length  $BNNZ(\hat{A}) + 1$  such that the  $i$ -th element of INDX points to the location in VAL of the the (1,1) element of the  $i$ -th block entry.  $\text{INDX}(BNNZ(\hat{A}) + 1)$  is set to the value  $NNZ(A) + 1$ , the ending number of nonzero point entries of  $A$  plus one.
- BINDX - An integer array of length  $BNNZ(\hat{A})$  consisting of the block column indices of entries of  $\hat{A}$ :

$$\text{BINDX} = (\text{BNZJ}(\hat{A}_{1*}), \text{BNZJ}(\hat{A}_{2*}), \dots, \text{BNZJ}(\hat{A}_{m_b*})).$$

- **RPNTR** - An integer array of length  $m_b + 1$  such that  $\text{RPNTR}(i) - \text{RPNTR}(1) + 1$  is the row index of the first point row in the  $i$ -th block row.  $\text{RPNTR}(m_b + 1)$  is set to  $m + \text{RPNTR}(1)$ . Thus, the number of point rows in the  $i$ -th block row is  $\text{RPNTR}(i + 1) - \text{RPNTR}(i)$ .
- **CPNTR** - An integer array of length  $k_b + 1$  such that  $\text{CPNTR}(j) - \text{CPNTR}(1) + 1$  is the column index of the first point column in the  $j$ -th block column.  $\text{CPNTR}(k_b + 1)$  is set to  $k + \text{CPNTR}(1)$ . Thus, the number of point columns in the  $j$ -th block column is  $\text{CPNTR}(j + 1) - \text{CPNTR}(j)$ .
- **BPNTRB** - An integer array of length  $m_b$  such that  $\text{BPNTRB}(i) - \text{BPNTRB}(1) + 1$  points to the location in **BINDX** of the first block entry of  $\text{BNZE}(\hat{A}_{i*})$ . If  $\text{BNZE}(\hat{A}_{i*})$  is empty then set  $\text{BPNTRB}(i) = \text{BPNTRB}(i + 1)$ .
- **BPNTRE** - An integer array of length  $m_b$  such that  $\text{BPNTRE}(i) - \text{BPNTRB}(1)$  points to the location in **BINDX** of the last block entry of  $\text{BNZE}(\hat{A}_{i*})$ . If  $\text{BNZE}(\hat{A}_{i*})$  is empty then set  $\text{BPNTRE}(i) = \text{BPNTRB}(i)$ .

- Note:**
1. For a general matrix ( $\text{DESCRA}(1) = 0$ ), **CPNTR** can be different from **RPNTR**. However, for a typical square matrix,  $\text{CPNTR} = \text{RPNTR}$  so only one copy needs to be kept. For all other matrix types, **RPNTR** must equal **CPNTR** and a single array can be passed for both arguments.
  2. The array **INDX** is not essential for reconstructing the matrix. However, it is essential for good performance of most matrix operations and should be computed one time and kept for later use.

To illustrate the VBR data structure, consider the following matrix with the indicated row and column partitioning:

$$\hat{A} = \left( \begin{array}{cc|ccc|c|ccc|cc} 4 & 2 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & -1 & 1 \\ 1 & 5 & 0 & 0 & 0 & 2 & 0 & 0 & 0 & 0 & -1 \\ \hline 0 & 0 & 6 & 1 & 2 & 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 7 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 2 & 9 & 3 & 0 & 0 & 0 & 0 & 0 \\ \hline 2 & 1 & 3 & 4 & 5 & 10 & 4 & 3 & 2 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 4 & 13 & 4 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 3 & 3 & 11 & 3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 7 & 0 & 0 \\ \hline 8 & 4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 25 & 3 \\ \hline -2 & 3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 8 & 12 \end{array} \right) .$$



The sparsity pattern of  $\hat{A}$  is

$$\hat{A} = \begin{pmatrix} * & 0 & * & 0 & * \\ 0 & * & * & 0 & 0 \\ * & * & * & * & 0 \\ 0 & 0 & * & * & 0 \\ * & 0 & 0 & 0 & * \end{pmatrix}.$$

It is stored in VBR format as follows:

$$\begin{aligned} \text{VAL} &= ( 4, 1, 2, 5; 1, 2; -1, 0, 1, -1; 6, 2, -1, 1, \\ & 7, 2, 2, 1, 9; 2, 0, 3; 2, 1; 3, 4, 5; 10; \\ & 4, 3, 2; 4, 3, 0; 13, 3, 2, 4, 11, 0, 2, 3, \\ & 7; 8, -2, 4, 3; 25, 8, 3, 12 ) , \end{aligned}$$

$$\text{INDX} = ( 1, 5, 7, 11, 20, 23, 25, 28, 29, 32, 35, 44, 48, 52 ) ,$$

$$\text{BINDX} = ( 1, 3, 5, 2, 3, 1, 2, 3, 4, 3, 4, 1, 5 ) ,$$

$$\text{RPNTR} = ( 1, 3, 6, 7, 10, 12 ) ,$$

$$\text{CPNTR} = ( 1, 3, 6, 7, 10, 12 ) ,$$

$$\text{BPNTRB} = ( 1, 4, 6, 10, 12, ) .$$

$$\text{BPNTRE} = ( 4, 6, 10, 12, 14 ) .$$

If  $A$  is symmetric (or Hermitian, triangular, or anti-symmetric), then we only need to store the block diagonal and block lower (or upper) triangle. If the diagonal block entries are stored, we assume that the *entire* square block entry is stored.

## 4 Rationale for Sparse BLAS Toolkit Design

In this section we discuss the reasons for our design. We believe there is little argument about the need for what we propose here. However, one may reasonably argue that there are other important kernels omitted. One may also argue with the details of our sparse matrix data structures and which data structures are supported or unsupported. We are aware of these issues and are willing to make changes if necessary. In the following Section 5 we list and discuss issues that we believe to be open for further review.

### 4.1 Emphasis on Higher-level Functionality

We have chosen to emphasize a very high level of functionality in this proposal, i.e., we have specified level-3 BLAS operations on block entry matrices. As we see it, the negative aspects of this are:

- For users who want to compute a simple sparse matrix vector product, these general interfaces might seem somewhat cumbersome.

- To provide a complete, optimal implementation of each of these kernels requires a substantial amount of effort.

The positive aspects are:

- There are fewer user-callable routines. By combining level-2 and level-3 operations we have reduced the number of routines by a factor of two.
- There is a greater opportunity for high performance. Users who can exploit level-3 operations and have block entry matrices can achieve significant performance gains.
- Users are encouraged to think in terms of higher-level operations.

## 4.2 Language Independent Specifications

We believe that the current proposal addresses all major issues related to ease-of-use from Fortran, C and C++. The latest changes of replacing PNTR (BPNTR) with PNTRB and PNTRE (BPNTRB and BPNTRE) and allowing for zero or one based arrays eliminated the know complaints.

## 4.3 Choice of Data Structures

Each of the data structures we propose is intended to support an important class of problems or algorithms. Below we list each point entry data structure and one or more of its intended uses. The intended uses for block entry data structures are analogous with added property that we exploit the property of multiple unknowns per gridpoint or related properties as mentioned in Section 3.4.

- COO - Coordinate: Most flexible data structure when constructing or modifying a sparse matrix.
- CSC - Compressed sparse column: Natural data structure for many common matrix operations including matrix multiplication and constructing or solving sparse triangular factors.
- CSR - Compressed sparse row: Natural data structure for many common matrix operations including matrix multiplication and constructing or solving sparse triangular factors.
- DIA - Sparse diagonal: Particularly useful for matrices coming from finite difference approximations to partial differential equations on uniform grids.
- ELL - Ellpack/Itpack: Appropriate for finite element or finite volume approximations to partial differential equations where elements are of the same type, but the gridding is irregular.
- JAD - Jagged diagonal: Appropriate for matrices which are highly irregular or for a general-purpose matrix multiplication where the properties of the matrix are not known *a priori*.

- SKY - Skyline: Particularly well suited for Cholesky or LU decomposition when no pivoting is required. In this case, all fill will occur within the existing non-zero structure.

## 4.4 Choice of Functionality

Below we discuss our motivation for each basic operation we chose. We realize that there are many other important operations and plan to implement more in the future. See Section 5.11 for specific future plans.

### 4.4.1 Matrix-Matrix Product and Solution of Sparse Triangular System

We believe there is little need to justify the first two operations which are supported in the Sparse BLAS Toolkit since sparse matrix-dense matrix product and the solution of a sparse triangular system are the primary computational kernels in many sparse linear equation and eigensystem solvers.

### 4.4.2 Right Permutation of a Sparse Matrix in JAD Format

The right permutation of a sparse matrix is necessary in cases where the rows of  $A$  were already permuted in the process of storing  $A$  in the sparse data structure. For example, if  $A$  is stored in JAD format, then the rows of  $A$  are ordered by decreasing number of entries. If the corresponding right permutation is not applied, then the result matrix,  $C$ , from a call to DJADMM will not be ordered consistently with the input matrix  $B$ . By permuting the columns of  $A$  and the rows of  $B$  *outside* the iteration loop, we can eliminate this inconsistency. For example, consider the simple iteration loop:

```
do i = 1, ...
  y ← Ax
  x ← x + ay
end do
```

If the matrix  $H$  is such that  $A = PH$  for some permutation matrix  $P$  then, in terms of the matrix  $H$ , the loop becomes:

```
do i = 1, ...
  y ← Hx
  y ← Py
  x ← x + ay
end do
```

The additional permutation  $Py$  in each iteration could significantly affect efficiency. By performing an explicit right permutation, we can eliminate permutations from the iteration loop as follows:

```

 $x \leftarrow P^T x$ 
 $H \leftarrow HP$ 
do  $i = 1, \dots$ 
   $y \leftarrow Hx$ 
   $x \leftarrow x + ay$ 
end do
 $x \leftarrow Px$ 

```

### 4.4.3 Check for Validity of Sparse Matrix Representation

Because of the potential large size of sparse problems and the complexity of sparse matrix data structures, it is very useful to have the capability to perform some kind of check on the sparse matrix representation. The checking routines presented here make no rigorous attempt to determine the numerical stability and, thus, are not fail-safe. However, they do improve the probability of finding and correcting most programming errors.

## 4.5 Use on a Distributed Memory Machine

Clearly the proposal presented here is intended for a shared memory programming environment. However, one can easily use these kernels on distributed memory machines in a natural way. Consider the MM operation and partition  $A$ ,  $B$ , and  $C$  so that each processor owns the same subset of rows of each array. Let  $A_i$ ,  $B_i^l$ , and  $C_i$  denote the portion of  $A$ ,  $B$ , and  $C$ , respectively, on the  $i^{\text{th}}$  processor. Let  $D_i$  denote the set of row indices that reside on processor  $i$ . Let  $B_i^r$  denote elements of  $B$  which are not on processor  $i$  but are reached to by  $A_i$ . Finally, further decompose each  $A_i$  into  $A_i^l$ , the portion of  $A_i$  that reaches to  $B_i^l$ , and  $A_i^r$ , the portion of  $A_i$  that reaches to  $B_i^r$ . Then a general distributed MM algorithm is as shown in Figure 1. Note that both computation steps can be performed using the MM kernels presented in this proposal. These kernels are also very useful for domain decomposition algorithms where all work in the inner iterations is done locally.

## 5 Open Issues

In this section we present a list of open issues. Certainly there are others which we have omitted, so this list will change with each revision.

### 5.1 Language Incompatibilities

**Fortran Character Variables** The original version of this proposal made extensive use of character variables (see [14]). This was motivated by the precedent established by the dense BLAS and because of the improved ability to read the Fortran calls to the toolkit kernels. However, in the interest of making the interface more accessible from other

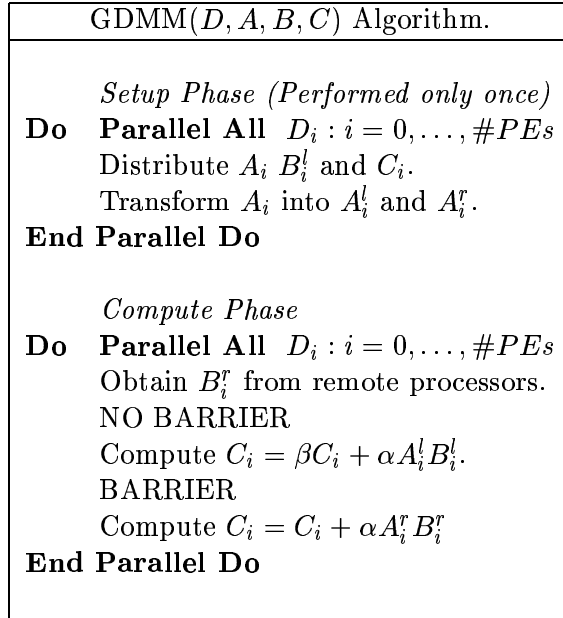


Figure 1: General Distributed Sparse Matrix-Dense Matrix Product Algorithm

languages, especially C, we have eliminated character variables and replaced them with less intuitive, but easier to use integer variables.

## 5.2 Omission of Rank Updates

The dense BLAS support rank updates of a matrix, e.g.,  $A \leftarrow A + \alpha xy^T$ . This operation is not useful if  $A$  is sparse since the update,  $xy^T$ , is dense. It would be possible to do an implicit rank update in the sense that we never really form the updated matrix  $A$ , but instead keep  $A$ ,  $x$  and  $y$ . However, we have not seen any reason for doing this since it is possible to use a combination of sparse and dense BLAS to perform this operation. For example, to compute  $w = (A + xy^T)z$  we first compute  $w = Az$  using a sparse BLAS kernel, then we compute  $\gamma = y^T z$  using the level-1 BLAS routine DDOT. Then we finally update  $w = w + \gamma x$  using the level-1 BLAS routine DAXPY.

## 5.3 Omission of SIDE

In the specification of the dense BLAS, the SIDE argument is used to indicate whether the operand matrix  $A$  is to the left or right of its associated operand. We assume that the sparse operand  $A$  always appears on the left of its associated operand. Thus, the SIDE argument is unnecessary.

## 5.4 Minimal support for pre-processing

In many cases, high-performance implementations of the Sparse BLAS Toolkit kernels will require some pre-processing that should be done only once. Repeated calls to a Sparse BLAS Toolkit kernel for a fixed matrix  $A$  would then use this pre-processed information to achieve better performance. In this proposal, there is minimal support for pre-processing. We include only one work array. However, because we believe that addressing these issues is not within the scope of this proposal, we hesitate to add much more support for pre-processing. Also, users should not need to be concerned about these implementation-specific details. The drawback is that some types of pre-processing may be difficult to implement.

## 5.5 Omission of MSR, MSC and Other Data Structures

The modified sparse row (MSR) and modified sparse column (MSC) data structures (see SPARSKIT User Guide [18]) are popular data structures and were originally supported in this proposal. Also, other data structures might be very useful. However, in the interest of having a small set of supported data structures, we have chosen to avoid data structures with similar performance characteristics.

**Note:** Matrix multiplication and the solution of triangular systems with MSR (and MSC) can be express using a combination of the DIA and with CSR (and CSC) with modification of the user's data structure.

## 5.6 Details of Constant Block Data Structures

The block data structures presented here are only one version and they differ in detail from those presented in [18]. Several questions arise:

1. Do we need to support both orderings of block entry elements? Presently we allow  $A$ ,  $B$  and  $C$  to have dimensions of  $A(\text{LB},\text{LB},\text{NNZ})$ ,  $B(\text{LB},\text{BLDB},*)$  and  $C(\text{LB},\text{BLDC},*)$  (contiguous) and  $A(\text{NNZ},\text{LB},\text{LB})$ ,  $B(\text{BLDB},\text{LB},*)$  and  $C(\text{BLDC},\text{LB},*)$  (non-contiguous). Can we get by with only one?
2. Presently the BINDX BPNTRB and BPNTRE arrays correspond to the *block* entries of the matrix. However, they could also correspond to the *scalar* elements by pointing to the (1,1) scalar element of each block. Compared with the present ordering scheme, this would not affect BPNTRB and bptre but, for the matrix in Equation 2 with  $\text{DESCRA}(5) = 0$ , BINDX as presented in Section 3.4 would become

$$\text{BINDX} = ( 1 \ 5 \ 3 \ 5 \ 1 ).$$

## 5.7 Variable Block Column Data Structure

Following the definition of the VBR structure, it is easy to define the variable compressed block column (VBC) data structure, i.e., blocks are stored block-column-by-block-column,

and the point entries are stored column-by-column within one block. One possible advantage of the VBC structure is that it could accelerate the matrix-vector (matrix-matrix) multiply operations on vector machines since the columns are longer in the VBC format. However, one needs to compute the leading dimensions of the columns. Our experience tells us that the drawback may overshadow the advantage in many cases. Therefore, we are not considering the VBC data structure in the current proposal, and will leave it as an open issue in the future.

## 5.8 Non-unit Block Diagonal Entries for Triangular Matrices

Currently we do not allow non-unit block diagonal entries for triangular matrices. This is because we are not sure how to invert these blocks when performing a block triangular solve. We see several possibilities: each diagonal block has the LU factorization of the block, or each has the explicit inverse, or each has some other factorization. Furthermore, we believe it is more convenient to keep the block diagonal separate, e.g., instead of using incomplete block LU we can use incomplete block LDU. Because of the possibilities and because of the support for block diagonal matrices we do not see the value in supporting non-unit block diagonal entries for triangular matrices.

## 5.9 Work Arrays and Dimensions

To facilitate optimization, we return in `WORK(1)` an integer in floating point form which is the optimal value for `LWORK`. Thus, setting `LWORK = INT(WORK(1))` will allow optimal performance from the kernel. This convention is used in LAPACK but we are open to other techniques. Our primary doubt about this approach is that the range of integers that we can express in floating point form is smaller than the standard range of integers if `WORK` and `LWORK` have the same element size.

## 5.10 Structurally-symmetric Nonsymmetric Matrices

Some nonsymmetric problems have a symmetric nonzero pattern which could be exploited in some of the routines presented here, e.g., matrix-matrix multiplication for the CSR data structure. In this case, it would be possible to store the structure for only the lower triangle of this type of matrix and the values for the entire matrix. It is possible to do this even under the constraints of the existing routines by making two calls, one for the lower triangle (in CSR format) and one for the upper (in CSC format) using the same integer vectors in both cases. There would be some improvement in performance by processing both the lower and upper triangle simultaneously, but it brings with it some complications to the user interface. At this time, we are not planning to support this capability.

## 5.11 Future Work

During the development of this proposal several other basic operations and data structures were considered for inclusion. In particular, there are three topics that are not discussed

here which we believe are important and should be presented in later working notes. They are as follows.

### 5.11.1 Sparse Matrix Update Kernel

A preliminary version of this proposal included a kernel for updating a sparse matrix. This operation is of the form

$$A \leftarrow A + \alpha A'$$

where  $A$  is a sparse matrix stored in any of the supported data structures and  $A'$  is another sparse matrix stored in one of a subset of the supported data structures. We consider this kernel to be very useful in the assembly of a sparse matrix and in converting from one data structure to another. However, at this point, we are not certain about the feasibility of the implementation of this kernel and must study the problem further before proposing it.

### 5.11.2 General Sparse Permutation Kernels

Presently we support only one sparse permutation kernel, the right permutation of the JAD data structure. This is necessary for the reasons presented in Section 4.4.2. However, left and right permutations of sparse matrices can be useful in a more general context (see [13]).

### 5.11.3 Sparse Matrix-Sparse Matrix Multiplication

Sparse matrix times sparse matrix multiplication is a basic kernel for constructing incomplete Cholesky and ILU preconditioners. We are considering kernels for this operation for a restricted set of data structures. See [2] for details.

## 6 Acknowledgements

Many people were of great help in the development of this proposal. In particular, we thank Ramesh Agarwal, Fernando Alvarado, Ed Anderson, Steve Ashby, Craig Douglas, Iain Duff, Fred Gustavson, Bill Harrod, David Kincaid, Michele Marrone, Giuseppe Radicati, Youcef Saad, Qasim Sheikh, Phuong Vu, Chao Yang, and Alex Yerebin.

## References

- [1] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK Users' Guide*. SIAM Pub., 1992.
- [2] Randolph E. Bank and Craig C. Douglas. Sparse matrix multiplication package (SMMP). *Advances in Computational Mathematics*, 1, 1993.



- [3] Richard Barrett, Michael Berry, Tony F. Cahn, James Demmel, June Donato, Jack Dongarra, Victor Eijkhout, Roldan Pozo, Rharles Romine, and Henk van der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. SIAM, 1993.
- [4] David S. Dodson, Roger G. Grimes, and John G. Lewis. Sparse extensions to the FORTRAN basic linear algebra subprograms. *ACM Transactions on Mathematical Software*, 17(2):253–263, June 1991.
- [5] J. J. Dongarra, J. Bunch, C. Moler, and G. Stewart. *LINPACK Users' Guide*. SIAM Pub., 1979.
- [6] Jack J. Dongarra, Jeremy Du Croz, Sven Hammarling, and Iain Duff. A set of level 3 basic linear algebra subprograms. *ACM Transactions on Mathematical Software*, 16(1):1–17, March 1990.
- [7] J.J. Dongarra, J. DuCroz, S. Hammarling, and R. Hanson. An extended set of fortran basic linear algebra subprograms. *ACM Transactions on Mathematical Software*, 14, 1988.
- [8] I. S. Duff, A. M. Erisman, and J. K. Reid. *Direct Methods for Sparse Matrices*. Monographs on Numerical Analysis. Oxford University Press, New York, 1986.
- [9] Iain Duff, Michele Marrone, and Giuseppe Radicati. A proposal for user level sparse BLAS. Technical Report TR/PA/92/85, CERFACS, December 1992.
- [10] Elegant Mathematics, Inc., Bothell, Washington. *The EM Symmetric Sparse Supernodal Format*, 1993.
- [11] Elegant Mathematics, Inc., Bothell, Washington. *The EM Unsymmetric Sparse Supernodal Format*, 1993.
- [12] John R. Gilbert, Cleve Moler, and Robert Schreiber. Sparse matrices in Matlab: Design and implementation. *SIAM J. Mat. Anal.*, 13(1):333–356, January 1992.
- [13] Fred G. Gustavson. Two fast algorithms for sparse matrices: Multiplication and permuted transposition. *ACM Transactions on Mathematical Software*, 4(3):250–269, 1978.
- [14] Michael A. Heroux. A proposal for a sparse BLAS toolkit. Technical Report TR/PA/92/90, CERFACS, December 1992.
- [15] Mark T. Jones and Paul E. Plassmann. BlockSolve v1.1: Scalable library software for the parallel solution of sparse linear systems. Technical Report ANL-92/46, Mathematics and Computer Science Division, Argonne National Laboratory, December 1992.
- [16] C. Lawson, R. Hanson, D. Kincaid, and F. Krogh. Basic linear algebra subprograms for fortran usage. *ACM Transactions on Mathematical Software*, 5, 1979.
- [17] Thomas C. Oppe and David R. Kincaid. Are there iterative BLAS? Technical report, Center for Numerical Analysis, The University of Texas at Austin, February 1990.

- [18] Youcef Saad. SPARSKIT: a basic tool kit for sparse matrix computations. Preliminary Version.
- [19] 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*, volume 6 of *Lecture Notes in Computer Science*. Springer–Verlag, New York, second edition, 1976.
- [20] Barry F. Smith and William D. Gropp. The design of data-structure-neutral libraries for the iterative solution of sparse linear systems, 1993. Submitted to SIAM J. Sci Comput.