A Sparse Matrix Library in C++ for High Performance Architectures*

Jack Dongarra[§‡], Andrew Lumsdaine[◇], Xinhiu Niu[◇]
Roldan Pozo[‡], Karin Remington[§]

[§]Oak Ridge National Laboratory
Mathematical Sciences Section

[‡]University of Tennessee
Department of Computer Science

[◇]University of Notre Dame
Dept. of Computer Science & Engineering

Abstract

We describe an object oriented sparse matrix library in C++ built upon the Level 3 Sparse BLAS proposal [5] for portability and performance across a wide class of machine architectures. The C++ library includes algorithms for various iterative methods and supports the most common sparse data storage formats used in practice. Besides simplifying the subroutine interface, the object oriented design allows the same driving code to be used for various sparse matrix formats, thus addressing many of the difficulties encountered with the typical approach to sparse matrix libraries. We emphasize the fundamental design issues of the C++ sparse matrix classes and illustrate their usage with the preconditioned conjugate gradient (PCG) method as an example. We conclude by comparing performance results with equivalent Fortran 77 codes and discussing the impact of our design on elegance, performance, maintainability, portability, and robustness.

1 Introduction

Sparse matrices are pervasive in application codes which use finite difference, finite element or finite volume discretizations of PDEs for problems in computational fluid dynamics, structural mechanics, semiconductor simulation and other scientific and engineering applications. Nevertheless, comprehensive libraries for sparse matrix computations have not been developed and integrated to the same degree as those for dense matrices. Several factors contribute to the difficulty of designing such a comprehensive library. Different computer architectures, as well as different applications, call for different sparse matrix data formats in order to best exploit registers, data locality, pipelining, and parallel processing. Furthermore, code involving sparse matrices tends to be very complicated, and not easily portable, because the details of the underlying data formats are invariably entangled within the application code.

To address these difficulties, it is essential to develop codes which are as "data format free" as possible, thus providing the greatest flexibility for using the given algorithm (library routine) in various architecture/application combinations. In fact, the selection of an appropriate data structure can typically be deferred until link or run time. We describe an object oriented C++ library for sparse matrix computations which provides a unified interface for various iterative solution techniques across a variety of sparse data formats.

The design of the library is based on the following principles:

*This project was supported in part by the Defense Advanced Research Projects Agency under contract DAAL03-91-C-0047, administered by the Army Research Office, the Applied Mathematical Sciences subprogram of the Office of Energy Research, U.S. Department of Energy, under Contract DE-AC05-84OR21400, and by the National Science Foundation Science and Technology Center Cooperative Agreement No. CCR-8809615, and NSF grant No. CCR-9209815.

Clarity: Implementations of numerical algorithms should resemble the mathematical algorithms on which they are based. This is in contrast to Fortran 77, which can require complicated subroutine calls, often with parameter lists that stretch over several lines.

Reuse: A particular algorithm should only need to be coded once, with identical code used for all matrix representations.

Portability: Implementations of numerical algorithms should be directly portable across machine platforms.

High Performance: The object oriented library code should perform as well as optimized data-format-specific code written in C or Fortran.

To achieve these goals the sparse matrix library has to be more than just an unrelated collection of matrix objects. Adhering to true object oriented design philosophies, *inheritance* and *polymorphism* are used to build a matrix hierarchy in which the same codes can be used for various dense and sparse linear algebra computations across sequential and parallel architectures.

2 Iterative Solvers

The library provides algorithms for various iterative methods as described in Barrett *et. al.* [1], together with their preconditioned counterparts:

- Jacobi SOR (SOR)
- Conjugate Gradient (CG)
- Conjugate Gradient on Normal Equations (CGNE, CGNR)
- Generalized Minimal Residual (GMRES)
- Minimum Residual (MINRES)
- Quasi-Minimal Residual (QMR)
- Chebyshev Iteration (Cheb)
- Conjugate Gradient Squared (CGS)
- Biconjugate Gradient (BiCG)
- Biconjugate Gradient Stabilized (Bi-CGSTAB)

Although iterative methods have provided much of the motivation for this library, many of the same operations and design issues are addressed for direct methods as well. In particular, some of the most popular preconditioners, such as Incomplete LU Factorization (ILU) [1] have components quite similar to direct methods.

As an example of the C++ interface looks like, the core steps in the Preconditioned Conjugate Gradient method used to solve $Ax = b$ with preconditioner M, are presented below,

```
z = M.solve(r);
rho = r * z;
beta = rho /rho_1;
p = z + beta * p;
q = A * p;
alpha = rho / (p * q);
x += alpha * p;
r -= alpha * q;
rho_1 = rho;
```

Here the operators such as `*` and `+=` have been overloaded to work with matrix and vectors formats. This code fragment works for all of the supported sparse storage classes and makes use of the Level 3 Sparse BLAS in the matrix-vector multiply `A*p`.

3 Sparse Matrix types

We have concentrated on the most commonly used data structures which account for a large portion of application codes. The library can be arbitrarily extended to user-specific structures and will eventually grow. (We hope to incorporate user-contributed extensions in future versions of the software.) Matrix classes supported in the inital version of the library include:

`Sparse_Vector`: List of nonzero elements with their index locations. It assumes no particular ordering of elements.

`COOR_Matrix`: Coordinate Storage Matrix. List of nonzero elements with their respective row and column indices. This is the most general sparse matrix format, but it is not very space or computationally efficient. It assumes no ordering of nonzero matrix values.

`CRS_Matrix` : Compressed Row Storage Matrix. Subsequent nonzeros of the matrix rows are stored in contiguous memory locations and an additional integer arrays specifies where each row begins. It assumes no ordering among nonzero values within each row, but rows are stored in consecutive order.

`CCS_Matrix`: Compressed Column Storage (also commonly known as the *Harwell-Boeing* sparse matrix format [4]). This is a variant of CRS storage where columns, rather rows, are stored contiguously. Note that the CCS ordering of A is the same as the CRS of A^T.

`CDS_Matrix`: Compressed Diagonal Storage. Designed primarily for matrices with relatively constant bandwidth, the sub and super-diagonals are stored in contiguous memory locations.

`JDS_Matrix`: Jagged Diagonal Storage. Also know as `ITPACK` storage. More space efficient than CDS matrices at the cost of a gather/scatter operation.

`BCRS_Matrix`: Block Compressed Row Storage. Useful when the sparse matrix is comprised of square dense blocks of nonzeros in some regular pattern. The savings in storage and reduced indirect addressing over CRS can be significant for matrices with large block sizes.

`SKS_Matrix`: Skyline Storage. Also for variable band or profile matrices. Mainly used in direct solvers, but can also be used for handling the diagonal blocks in block matrix factorizations.

In addition, symmetric and Hermitian versions of most of these sparse formats will also be supported. In such cases only an upper (or lower) triangular portion of the matrix is stored. The trade-off is a more complicated algorithm with a somewhat different pattern of data access. Further details of each data storage format are given in [1] and [6].

4 Sparse Matrix Operations

Our library contains the common computational kernels required for solving linear systems by many direct and iterative methods. The internal data structures of these kernels are compatible with the proposed Level 3 Sparse BLAS, thus providing the user with large software base of Fortran 77 module and application libraries. Just as the dense Level 3 BLAS [3] have allowed for higher performance kernels on hierarchical memory architectures, the Sparse BLAS allow vendors to provide optimized routines taking advantage of indirect addressing hardware, registers, pipelining, caches, memory management, and parallelism on their

particular architecture. Standardizing the Sparse BLAS will not only provide efficient codes, but will also ensure portable computational kernels with a common interface.

There are two types of C++ interfaces to basic kernels. The first utilizes simple binary operators for multiplication and addition, and the second is a functional interfaces which can group triad and more complex operations. The binary operators provide for a simpler interface, e.g. $y = A*x$ denotes a sparse matrix-vector multiply, but may produce less efficient code. The computational kernels include:

- sparse matrix products, $C \leftarrow \alpha\ op(A)\ B + \beta C$.
- solution of triangular systems, $C \leftarrow \alpha D\ op(A)^{-1}\ B + \beta C$
- reordering of a sparse matrix (permutations), $A \leftarrow A\ op(P)$
- conversion of one data format to another, $A' \leftarrow A$,

where α and β are scalars, B and C are rectangular matrices, D is a (block) diagonal matrix, A and A' are sparse matrices, and $op(A)$ is either A or A^T.

In dealing with issues of I/O, the C++ library is presently designed to support reading and writing to Harwell-Boeing format sparse matrix files [4]. These files are inherently in compressed column storage; however, since sparse matrices in the library can be transformed between various data formats, this is not a severe limitation. File input is embedded as another form of a sparse matrix constructor; a file can be read and transformed into another format using conversions and the `iostream` operators. In the future, the library will also support other matrix file formats, such as a MATLAB™ compatible format, and IEEE binary formats. Sparse matrices can be also be initalized from conventional data and index vectors, thus allowing for a universal interface to import data from C or Fortran modules.

5 Conclusion

Using C++ for numerical computation can greatly enhance clarity, reuse, and portability. In the particular case of the C++ sparse matrix library described here, the details of the underlying sparse matrix data format are completely hidden at the algorithm level. Since the library is built upon the Level 3 Sparse BLAS, it should also provide performance comparable to that available with standard procedural languages. In the final paper we will include timing results between similar C++ and Fortran 77 codes.

References

[1] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, H. van der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, SIAM Press, 1994.

[2] J. Dongarra, R. Pozo, D. Walker, "LAPACK++: A Design Overview of Object-Oriented Extensions for High Performance Linear Algebra," Proceedings of Supercomputing '93, IEEE Press, 1993, pp. 162-171.

[3] J. Dongarra, J. Du Croz, I. S. Duff, S. Hammarling, "A set of level 3 Basic Linear Algebra Subprograms," *ACM Trans. Math. Soft.*, Vol. 16, 1990, pp. 1-17.

[4] I. Duff, R. Grimes, J. Lewis, "Sparse Matrix Test Problems," *ACM Trans. Math. Soft.*, Vol. 15, 1989, pp. 1-14.

[5] I. Duff, M. Marrone, G. Radicati, *A Proposal for User Level Sparse BLAS*, CERFACS Technical Report TR/PA/92/85, 1992.

[6] M. A. Heroux, *A Proposal for a Sparse BLAS Toolkit*, CERFACS Technical Report TR/PA/92/90, 1992.