

SPARSE TRIANGULAR SOLVE REVISITED: DATA LAYOUT CRUCIAL TO BETTER PERFORMANCE

BARRY SMITH* AND HONG ZHANG†

Abstract. A key to good processor utilization for sparse matrix computations is storing the data in the format that is most conducive to fast access by the memory system. In particular, for sparse matrix triangular solves the traditional compressed sparse matrix format is poor and minor adjustments to the data structure can increase the processor utilization dramatically. Such adjustments involve storing the L and U factors separately and storing the U rows “backwards” so that they are accessed in a simple streaming fashion during the triangular solves. Changes to the PETSc libraries to use this modified storage format resulted in over twice the floating point rate for some matrices. This improvement can be accounted for by a decrease in the cache misses and TLB (transaction lookaside buffer) misses in the modified code.

Key words. ILU-factorization, sparse triangular solve

AMS subject classifications.

1. Introduction. Many of the basic computational kernels in numeric software libraries were developed and implemented based on decades-old algorithms and techniques without serious consideration of computer architecture, e.g., the complex memory layout and data-fetching behavior. Traditionally, numerical algorithms and the associated programming subroutines are evaluated based on such factors as the mathematical error analysis, the rate of algorithmic convergence, and the flop counts. Hence many applications fail to achieve the anticipated speedup because of a mismatch between the data access patterns in the code and the data access patterns that are fastest on the given hardware.

As PETSc developers [2, 3], we have long been aware of the various memory bottlenecks in sparse matrix computation. We feel strongly that the *data access pattern* should become a standard in the evaluation of numerical algorithms and their implementations. Recently, we explored this concept on a computational kernel in PETSc: *the sparse triangular solve*. Through a simple reorganization of the data structure during the matrix factorization, we witnessed over 100 percent acceleration in the *sparse triangular solve* on a single core as a result of much better utilization of the memory subsystem. We are not changing how much data is accessed, nor are we changing the numerical algorithm. We are changing only the locations where the data is stored so that accessing it is as fast as possible. Essentially, we are decreasing the number of cache misses, fully utilizing each cache line (rather than having parts of cache lines not need in the next step of computation), and reducing the number of TLB (transaction lookaside buffer) misses.

The compressed sparse row (CSR) format is the most commonly used sparse matrix storage format. For sparse matrices with no additional structure (small dense blocks or values along certain diagonals, for example) the CSR format is appropriate for the sparse matrix-vector product kernel. Many sparse matrix software packages also provide LU and/or ILU factorization and triangular solver kernels. These are almost always implemented by using some simple variant of the CSR format. For example, the Euclid [14, 15], SPARSEKIT ILUT [16], and Aztec

*Mathematics and Computer Science Division, Argonne National Laboratory (bsmith@mcs.anl.gov).

†Computer Science Department, Illinois Institute of Technology (hzhang@mcs.anl.gov).

[13, 18] ILU implementations store the L and U factors interlaced by row, that is as $[L(1, :), U(1, :), L(2, :), U(2, :), \dots, L(n, :), U(n, :)]$. The Yale Sparse Matrix Package [6, 7] and hypre’s pilut [4, 8] store the L and U separately but still by row, starting with the first $[L(1, :), L(2, :), \dots, L(n, :)]$ and $[U(1, :), U(2, :), \dots, U(n, :)]$. All these variants are poor for system utilization because they result in slow memory access patterns on the back triangular solves. Hence the triangular solves as implemented have much lower floating-point rates than their corresponding matrix-vector products. By a simple change in the data layout we show that it is possible to bring the floating-point rate of the triangular solves up to almost as high as that for the sparse matrix-vector products.

In this paper we measure the efficiency of the triangular solves by their flop rates. High flop rates are not the ultimate goal; faster time to solution is the ultimate goal. But since for triangular solves the number of floating point operations remains the same, the flop rate is a good measure of the overall utilization of the compute node. Sparse matrix computations are always memory bandwidth limited. That is, some upper bound on the speed of computation is determined by the raw speed at which the memory can provide data for the process; see, for example, [1, 9, 10]. In this paper the focus is on how to get a particular computation closer to the memory bandwidth limit by taking into account other aspects of the memory system than simple raw bandwidth.

2. Modifications to the Data Structure. The *sparse triangular solve* and *matrix-vector product* are the dominating computational kernels in many large-scale iterative solvers. When an application uses Krylov subspace iteration with matrix-based preconditioners, for example, the incomplete LU (ILU) preconditioner, the *matrix-vector product* and *sparse triangular solve* are called repeatedly for generating Krylov subspaces; they often consume 70% or higher of the total execution time.

Assuming a sparse matrix A is stored in the conventional compressed row format, a practical way of implementing the ILU preconditioner is described in [17]:

Algorithm *sparse LU factor*:

Input: sparse matrix A

Output: sparse matrix factors L and U

For $i = 1, 2, \dots, n$ Do:

$w := A(i, :)$

For $k = 1, 2, \dots, i - 1$ Do:

$multiplier := A(i, k) / A(k, k)$

update $w := w - multiplier * U(k, :)$

EndDo

store w in $L(i, :)$ and $U(i, :)$

EndDo

For the ILU algorithm, the arithmetic operations described above occur only in the matrix entries that fall into a given nonzero sparse pattern. The conventional implementation of the *sparse LU factor* algorithm uses the standard compressed row format for the input matrix A . The entries of the output matrix factors L and U are naturally stored together in a single array with their rows interlaced:

$$(2.1) \quad factor_values = [L(1, :), U(1, :), L(2, :), U(2, :), \dots, L(n, :), U(n, :)].$$

The *sparse triangular solve* then follows the numerical factorization.

Algorithm *sparse triangular solve*:

Input: matrix factor L and U , right-hand side vector b

Output: solution vector x such that LUx approximates b

forward substitution

For $i = 1, 2, \dots, n$ Do

$$x(i) := b(i) - L(i, 1 : i - 1) * [x(1), \dots, x(i - 1)]^T$$

EndDo

backward substitution

For $i = n, n - 1, \dots, 1$ Do

$$x(i) := (x(i) - U(i, i + 1 : n) * [x(i + 1), \dots, x(n)]^T) / U(i, i)$$

EndDo

For each call of *sparse triangular solve*, the array *factor_values* in (2.1) is accessed twice: first, in the order of

$$L(1, :) \rightarrow L(2, :) \rightarrow \dots \rightarrow L(n, :)$$

during the forward substitution, then in the order of

$$U(n, :) \rightarrow U(n - 1, :) \rightarrow \dots \rightarrow U(1, :)$$

during the backward substitution. Here, for simplicity we present the solves without row and column permutations.

The flop counts for the *matrix-vector product*, are equal to the number of nonzero entries in A , which is the same as the flop counts for the *sparse triangular solve* when matrices L and U are obtained from the ILU(0) matrix factorization. However, because of the difference in the data layout of the matrix entries in the original matrix A and the matrix factors L and U , that is, contiguous array in A and noncontiguous array in L and U , the execution time of a *sparse triangular solve* usually takes twice as long as that of the *matrix-vector product*. The delay in the *sparse triangular solve* becomes more significant for larger matrices.

A careful examination of the *sparse triangular solve* algorithm reveals that the memory access can be improved through a simple reorganization of the matrix entries in the factored matrix L and U . Instead of storing the rows of L interlaced with U 's in the order being computed from the *sparse LU factor*, we arrange them in the **order of accessing** by the *sparse triangular solve*. Therefore the matrix entries are stored contiguously as

$$(2.2) \quad \text{factor_values} = [L(1, :), L(2, :), \dots, L(n, :), U(n, :), \dots, U(2, :), U(1, :)].$$

With this data layout, the *sparse triangular solve* reads the array *factor_values* (2.2) only once, from beginning to end, in comparison with two sweeps of (2.1) as in the previous implementation.

The idea of reorganizing matrix data is simple. Its implementation is simple, too: in the subroutine *sparse LU factor*, we store U entries from the end of the array *factor_values* instead of next to L 's, and we modify the values of the row and diagonal pointers accordingly. The existing subroutine *sparse triangular solve* requires only trivial editing that ensures the correct rows of L and U are accessed during the forward and backward substitution. Yet, the numerical experiments show an amazing acceleration: up to a 100 percent reduction in execution time.

Two important variants of the basic sparse triangular solves are as follows:

- Row and column permutations that are used to reduced fill in the factors or improve the convergence of the iterative method.
- Point-block storage of the factored matrix where the matrix has “natural” small blocks that it inherits from the continuous problem. For example, the full implicit discretization of the Navier-Stokes equations leads to sparse matrices with 5 by 5 dense blocks. In this case the CSR format is modified so only a single column index is needed to indicate the block column of the entire block. Computations using this block CSR format are faster because fewer loads of the column indices are needed.

Both these extensions can handle the modified storage proposed and benefit from it in the same manner.

3. Numerical Results. We tested the new data layout on two extreme cases – the traditional 7-point stencil on a unit cube and a matrix arising from the discretization of the compressible Euler equations – in order to get a measure of the effect of the modified data structure on both extremely sparse matrices and those arising in applications with several hundred nonzeros per row.

The experiments were conducted on a MacBook Pro with 2.8 GHz Intel Core 2 Duo and 1.067 GHz DDR3 memory using one core. ILU(0) with no row and column permutations was used for both CSR and block CSR formats. We are currently modifying all of the PETSc ILU solvers to use the new format and have obtained similar performance improvements with those as well. Our performance results were obtained by running the entire GMRES algorithm and profiling the relevant *matrix-vector products* and *sparse triangular solves*. We have done so because standalone benchmarking often produces unreasonably optimistic performance projections since much of the data is already in cache, whereas in the actual application it will not be. The numerical results were obtained by making multiple runs and using the average value; all runs had results within 5 percent of the average.

The floating-point rates were obtained by using the Intel ICC compiler version 10.1 with -O3 optimization. The cache misses and TLB misses were obtained by using the Apple profiling package Shark, Apple’s gcc version 4.0.1 with the -O3 option, and the Intel hardware counters L2_RQSTS and DTLB_MISSES.

TABLE 3.1
Performance Improvements for the Extremely Sparse Matrix

	Matrix-Vector Product	Triangular Solve	
		Noncontiguous Array	Contiguous Array
Flop rate (megaflops)	537	261 (49% of multiply)	447 (83% of multiply)
L2 cache misses (1000)	680	2,200	610
TLB misses (1000)	635	1,300	750

We considered two cases:

- **Extremely Sparse Matrix**, created with the 7-point stencil finite-difference scheme on a 65 by 65 by 65 cube. The performance results are given in Table 3.1.
- **Block Sparse Matrix**, the Jacobian matrix obtained from a fully implicit compressible Euler code on a mapped C-H mesh [12, 11]. It has a natural block size of 5. For this matrix we ran two studies. The first, where we use the traditional compressed sparse row format, is given in Table 3.2. The

TABLE 3.2
Performance Improvements for the Block Sparse Matrix

	Matrix-Vector Product	Triangular Solve	
		Noncontiguous Array	Contiguous Array
Flop rate (megaflops)	620	260 (42% of multiply)	660 (106% of multiply)
L2 cache misses (1000)	260	7,000	300
TLB misses (1000)	500	950	500

second, that uses the block compressed sparse row format, is given in Table 3.3.

TABLE 3.3
Performance Improvements for Block Sparse Matrix Using Block Compressed Sparse Row Storage

	Matrix-Vector Product	Triangular Solve	
		Noncontiguous Array	Contiguous Array
Flop rate (megaflops)	890	468 (53% of multiply)	758 (85% of multiply)
L2 cache misses (1000)	950	5,300	1,500
TLB misses (1000)	350	610	330

In Table 3.4 we provide the flop rates for the three studies in the previous tables for the IBM Blue Gene/P core which is a PowerPC 450 running at 850 MHz with memory running at 425 MHz.

TABLE 3.4
Flop Rate (in megaflops) Performance Improvements for the IBM Blue Gene/P Core

	Matrix-Vector Product	Triangular Solve	
		Noncontiguous Array	Contiguous Array
7-point stencil	98	46 (47% of multiply)	62 (63% of multiply)
Euler	148	99 (66% of multiply)	126 (85% of multiply)
Euler with block CSR	303	198 (65% of multiply)	260 86(% of multiply)

In all cases, using contiguous arrays in the *sparse triangular solve* accelerates execution time up to 100 percent in comparison with the traditional non-contiguous array storage of matrix factors. We also list the flop rate of the *matrix-vector product*. The newly proposed contiguous data layout makes the *sparse triangular solve* almost as fast as the *matrix-vector product*.

The time required for the sparse LU factorization does not change significantly depending on the storage of L and U , so this is not a matter of simply moving the computation time from one part of the code to another.

4. Conclusion. We have presented a case study in sparse matrix computations where a small change to the data structure for a sparse matrix results in a dramatic improvement in the performance of a computational kernel that uses the data structure. We note that for ILU factorizations, the factorization comes first, and traditionally that has dictated the data layout of the L and U factors. The factorization routine loads the factor values, in a natural way, from beginning to end. But this means that the solve routines access the values “backwards.” This emphasizes the importance of

picking a data structure, based not on how it is used first, but rather on how it will be accessed more often.

We plan to study data structures for other sparse matrix computations such as successive overrelaxation and Eisenstat's trick [5] to improve the performance of those computations as well.

REFERENCES

- [1] W. K. Anderson, W. D. Gropp, D. K. Kaushik, D. E. Keyes, and B. F. Smith. Achieving high sustained performance in an unstructured mesh CFD application. In *Proceedings of SC 99*, 1999. Winner of Gordon Bell Special Prize at SC1999.
- [2] Satish Balay, Kris Buschelman, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Barry F. Smith, and Hong Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.0.0, Argonne National Laboratory, 2009.
- [3] Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhauser Press, 1997.
- [4] E. Chow, A. Cleary, and R. Falgout. Design of the hypre preconditioner library. In *Proceedings of the SIAM Workshop on Object Oriented Methods for Inter-operable Scientific and Engineering Computing*. SIAM, 1999.
- [5] S. Eisenstat. Efficient implementation of a class of CG methods. *SIAM J. Sci. Stat. Comput.*, 2:1–4, 1981.
- [6] S. C. Eisenstat, H. C. Elman, M. H. Schultz, and A. H. Sherman. The (new) Yale Sparse Matrix Package. Technical Report YALE/DCS/RR-265, Department of Computer Science, Yale University, April 1983.
- [7] S. C. Eisenstat, H. C. Elman, M. H. Schultz, and A. H. Sherman. *Elliptic Problem Solvers II*, pages 45–52. Academic Press, 1984.
- [8] R. Falgout. hypre users manual. Technical Report Revision 2.0.0, Lawrence Livermore National Laboratory, 2006.
- [9] W. D. Gropp, D. K. Kaushik, D. E. Keyes, and B. F. Smith. High performance parallel implicit CFD. *Journal of Parallel Computing*, 27:337–362, 2001.
- [10] W. D. Gropp, D. K. Kaushik, D. E. Keyes, and B. F. Smith. Latency, bandwidth, and concurrent issue limitations in high-performance CFD. In *Proceedings of the First MIT Conference on Computational Fluid and Solid Mechanics*, June 2001.
- [11] William D. Gropp, David E. Keyes, Lois Curfman McInnes, and M. D. Tidriri. Parallel implicit PDE computations: Algorithms and software. In *Proceedings of Parallel CFD'97*, pages 333–344. Elsevier, 1998.
- [12] William D. Gropp, David E. Keyes, Lois Curfman McInnes, and M. D. Tidriri. Globalized Newton-Krylov-Schwarz algorithms and software for parallel implicit CFD. *Int. J. High Performance Computing Applications*, 14:102–136, 2000.
- [13] Scott A. Hutchinson, John N. Shadid, and Ray S. Tuminaro. Aztec user's guide version 1.1. Technical Report SAND95/1559, Sandia National Laboratories, October 1995.
- [14] D. Hysom and A. Pothen. A scalable parallel algorithm for incomplete factor preconditioning. *SIAM J. Sci. Comput.*, 22(6):2194–2215.
- [15] D. Hysom and A. Pothen. Euclid user manual (a scalable ILU preconditioning library for the parallel solution of sparse linear systems). Technical report, Old Dominion University, 2001.
- [16] Youcef Saad. SPARSKIT, a basic tool kit for sparse matrix computations. Technical Report 1029, Center for Supercomputing Research and Development, University of Illinois at Urbana-Champaign, 1990.
- [17] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. PWS Publishing Company, Boston, 1996.
- [18] Ray S. Tuminaro, Micheal Heroux, Scott A. Hutchinson, and John N. Shadid. Official Aztec user's guide version 2.1. Technical report, Sandia National Laboratories, 1999.