DOE/ER/25095--1

DE92 017870

Progress Report
DOE Grant No. DE-FG02-91ER25095
June 1991 to Jan 1992

Sparse Matrix Algorithms on Distributed Memory Multiprocessors

Alex Pothen
Computer Science Department
The Pennsylvania State University
University Park, PA 16802

The past year has seen significant progress in algorithms and software for the solution of large-scale sparse systems of equations, least-squares problems, and optimization problems on advanced distributed-memory parallel machines. The progress made to date is described below; together with my students and colleagues, I am continuing to pursue several research issues on these topics.

1. Large-scale linear systems. In this area, we focused on three problems: the computation of good orderings for solving sparse systems of equations, algorithms and software for factoring sparse matrices on distributed-memory multiprocessors, and algorithms for solving sparse triangular systems on highly parallel machines.

1.1. Spectral nested dissection orderings. In joint work with my Ph.D. student Lie Wang and Horst Simon (NASA Ames) [13], we considered an algebraic approach to computing good parallel orderings for the factorization of large, sparse, symmetric positive definite matrices. In this approach, we use the adjacency graph of the matrix to form a matrix called the Laplacian matrix, and then use information about a particular eigenvector to compute a separator in the graph. This approach is then recursively employed to compute spectral nested dissection orderings. Our results on very large problems (with tens of thousands of unknowns) show that this approach is very successful in computing orderings that have better parallelism than the currently available methods such as minimum-degree and earlier variants of nested dissection. The new spectral orderings were used to compute the matrix factorizations on a Cray Y-MP/8 much faster than with minimum-degree and other orderings. Currently we are working on an efficient implementation of spectral nested dissection algorithm for the Cray; there have been several requests for this code from several groups of researchers, and we intend to make our software available to them. Lie presented this work at a 'Parallel Circus' organized by Professor Gene Golub (Stanford) and Dr. Esmond Ng at the Oak Ridge National Labs in Nov '91.

I spent the months of Oct and Nov '91 at the Institute of Mathematics and its Applications (IMA) at the University of Minnesota, at their invitation. While there, together with Professor Bojan Mohar of Ljubljana (formerly Yugoslavia), I used the spectral approach to design and analyze the performance of an algorithm for reducing the envelope size of a sparse matrix. This problem is important in several structural engineering codes, where envelope methods are used to solve large systems of equations. We showed that the Laplacian matrix could be used to greatly reduce the size of the envelope, and thereby the storage and arithmetic work required for the solution. This work [6] is being written up now. (The other work I performed while there will be described in the appropriate subsections below.)

1.2. Parallel Multifrontal factorization. The multifrontal method is known to be an efficient method for computing the Cholesky factorization of sparse matrices on vector and parallel computational environments. My Ph.D. student Chunguang Sun (now a postdoc at the Advanced

Computing Research Institute, Cornell University) investigated several issues in producing an efficient implementation of the multifrontal method on the iPSC/2 and iPSC/860 hypercubes. We used a data structure called the clique tree (which we had previously studied—see [5, 14]) to organize the computation using efficient dense matrix kernels, and designed a proportional mapping algorithm to map computational subtasks to the processors. We reported the first set of results on parallel execution times for irregular sparse systems for the hypercube machines, and efficiencies were comparable to the results obtained for the model regular grid problem. This work has been written up and submitted for publication [15, 16, 19]. We intend to make this software available for public use since we have received several requests for it. Chunguang described this work at the International Conference on Industrial and Applied Mathematics, Washington D. C. in July and at the SIAM linear algebra meeting at Minneapolis in Sep '91.

1.3. Highly parallel triangular solution. On massively parallel machines such as the Connection Machine, a bottleneck in the parallel solution of linear systems is the triangular solution part, since $O(n^2)$ floating point operations are performed on $O(n^2)$ elements. In the situation when the system involves multiple right-hand side vectors, a *partitioned inverse* approach can be used to significantly improve the parallelism by replacing triangular solutions by means of a sequence of matrix-vector multiplications. By minimizing the number of matrix-vector multiplications, we can obtain an algorithm for solving the triangular system efficiently in parallel on massively parallel machines. Together with Professor F. Alvarado (Wisconsin) we [9] designed a fast algorithm to reduce the number of matrix-vector multiplications in this approach when the input matrix is symmetric positive definite. This algorithm was faster by more than a hundred fold on a collection of problems over a previous algorithm [1]; it has an even greater edge in terms of auxiliary storage. This program is now being used in a software package called the Sparse Matrix Manipulation System (SMMS) for applications in Power Engineering.

Recently, with Barry Peyton (Oak Ridge National Labs) and a graduate student Xiaoqing Yuan [7], I was able to generalize the above problem to reduce the number of matrix-vector multiplications even further. This work has necessitated the development of the theory of a class of orderings of chordal graphs called *transitive elimination orderings.* Currently we are writing up this work; soon we expect to implement the new algorithm to study the gains it might bring.

I talked about this work at the SIAM Applied Linear Algebra meeting at Minneapolis in Sep '91 and at the IMA, University of Minnesota in Nov '91.

2. Least-squares problems. In this area, we worked on two problems: developing a parallel algorithm for solving the least-squares problem on a hypercube, and the correct structure prediction of the orthogonal factors of a sparse matrix.

2.1. Structure of orthogonal factors. A direct method for the solution of least-squares problems requires the computation of the orthogonal factors of the given sparse matrix. To do so efficiently, we require data structures that store only the nonzeros in the factors before the numerical factorization is computed. However, till last year, structure of sparse orthogonal factors could be correctly predicted only for a subset of matrices which possessed a property called the strong Hall property. Last year, Hare, Johnson, Olesky, and van den Driessche showed how the structures could be predicted in the absence of this property, but they could not show that the predicted structures were the best possible. In [8], I extended this work to show that the structures predicted were the best possible, and developed algorithms for efficiently computing the data structures of the factors. An important consequence of this work is that it makes it possible to design efficient algorithms for orthogonal factorization with pivoting for rank-deficient and ill-conditioned problems. We expect to soon provide an implementation of such an orthogonal factorization algorithm since rank-deficient problems form an important class of least-squares problems. I talked about this work at the IMA workshop on Sparse Matrix Algorithms in Minneapolis in Oct '91.

2.2. Parallel sparse orthogonal factorization. The development of parallel algorithms for computing the sparse orthogonal factorization has lagged behind the development of Cholesky factorization algorithms due to the greater difficulties associated with the former. My Ph.D. student, Padma Raghavan, (now a postdoc with Professor Mike Heath at the University of Illinois), and I developed a parallel sparse orthogonal factorization algorithm for the iPSC hypercube. Our algorithm uses the multifrontal idea for computing the factorization by means of a sequence of merges involving dense triangular matrices. The arithmetic cost of the algorithm is low because row-oriented Householder transformations are used to perform the numerical computations, and the communication costs represent a lower order term than the arithmetic cost. Our implementation of this algorithm performed well on the ipSC hypercube, and good speed-ups (comparable to Cholesky factorization) were ol ›rved.

Several areas for improvement of the algorithm remain. An important concern is reducing the number of messages sent during the factorization, since the high start-up costs of sending messages in currently available hypercubes demand that this number should be low. To achieve this, a hybrid of Givens and Householder transformations may be necessary.

Padma completed her thesis [18] last year, and this work is currently being written up. She talked about this work at the International Conference on Industrial and Applied Mathematics at Washington, D. C. in July '91.

3. Numerical Optimization. A central problem is the solution of large-scale numerical optimization problems is computing a sparse basis for the null space of a large, sparse, underdetermined matrix. A theoretical study of the sparse null space basis problem was made in [3, 4], and then algorithms for computing null space bases were designed and implemented in [17]. These algorithms compute a basis by successively computing null vectors which are linearly independent of previously computed null vectors. A fundamental problem associated with computing a sparse null space basis is identifying a condition on the zero-nonzero structure which would guarantee the linear independence of the computed null vectors. This was achieved in [4, 17] by restricting the basis to have an embedded identity or upper triangular matrix.

An important open question in this work was whether this restriction was unnecessarily strong. In recent work [2] with Professors Richard Brualdi (Wisconsin-Madison) and Shmuel Friedland (Illinois, Chicago), I have been able to characterize the structure of sparsest bases of dense underdetermined matrices in terms of a condition on the zero-nonzero structure of the basis. This work was also performed while I was visiting the IMA, University of Minneapolis. This problem turned out to be surprisingly difficult, and we had to employ techniques from algebraic geometry and multilinear algebra to solve the problem. We are currently trying to extend these results to the sparse case. Once this is done, I will work on the problem of computing sparse null space bases which are well-conditioned. This will make use of the recent work on orthogonal factorization with pivoting.

REFERENCES

[1] F. Alvarado and R. S. Schreiber. Optimal parallel solution of sparse triangular systems. SIAM J. Sci. Stat. Comput., to appear, 1992.

[2] R. A. Brualdi, S. Friedland, and A. Pothen. On the linear independence of elementary vectors. in preparation, Feb 1992.

[3] T. F. Coleman and A. Pothen. The null space problem I: Complexity. *SIAM J. Alg. & Disc. Meth.*, 7:527–537, 1986.

[4] T. F. Coleman and A. Pothen. The null space problem II: Algorithms. *SIAM J. Alg. & Disc. Meth.*, 8:544–563, 1987.

[5] J. G. Lewis, B. W. Peyton, and A. Pothen. A fast algorithm for reordering sparse matrices for parallel factorization. *SIAM J. Sci. Stat. Comput.*, 10:1146–1173, 1989. (Special section on Sparse Matrix Algorithms for Supercomputers).

[6] B. Mohar, A. Pothen, and H. D. Simon. A spectral approach to envelope reduction of sparse matrices. Work in preparation, 1992.

[7] B. W. Peyton, A. Pothen, and X. Yuan. Transitive orderings of chordal graphs and parallel sparse triangular solution. in preparation, Feb 1992.

[8] A. Pothen. Predicting the structure of sparse orthogonal factors. Technical Report CS-91-27, Computer Science, Penn State, Sep 1991. Submitted.

[9] A. Pothen and F. Alvarado. A fast reordering algorithm for parallel sparse triangular solution. *SIAM J. Sci. Stat. Comput.*, 13:645–653, March 1992.

[10] A. Pothen and C-J. Fan. Computing the block triangular form of a sparse matrix. *ACM Trans. on Math. Software*, 16:303–324, Dec 1990.

[11] A. Pothen and P. Raghavan. Parallel multifrontal orthogonal factorization for sparse matrices. Work in preparation, 1992.

[12] A. Pothen, H. D. Simon, and K-P. Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Mat. Anal. Appl.*, 11:430–452, 1990.

[13] A. Pothen, H. D. Simon, and L. Wang. Spectral nested dissection. Technical Report CS-92-01, Computer Science, Penn State, Jan 1992. Submitted.

[14] A. Pothen and C. Sun. Compact clique tree data structures in sparse matrix factorizations. In *Large-Scale Numerical Optimization, Proceedings in Applied Mathematics*, pages 180–204. Society for Industrial and Applied Mathematics, 1990.

[15] A. Pothen and C. Sun. A distributed multifrontal algorithm using clique trees. Technical Report CS-91-24, Computer Science, Penn State, Aug 1991. Submitted.

[16] A. Pothen and C. Sun. Distributed multifrontal factorization using clique trees. In *Proceedings of the Fifth SIAM Conference on Parallel Processing for Scientific Computing*. Society for Industrial and Applied Mathematics, 1991.

[17] Alex Pothen. *Sparse Null Bases and Marriage Theorems*. PhD thesis, Cornell University, Ithaca, New York, 1984.

[18] Padma Raghavan. *Distributed Sparse Matrix Factorization: QR and Cholesky Decompositions*. PhD thesis, The Pennsylvania State University, University Park, PA, 1991.

[19] Chunguang Sun. *Some Applications of Clique Trees to Sparse Cholesky Factorizations*. PhD thesis, The Pennsylvania State University, University Park, PA, 1991.

DATE
FILMED
8/18/92