Combinatorial Algorithms for Parallel Sparse Matrix Distributions

Erik G. Boman *

Abstract

Combinatorial algorithms have long played a crucial enabling role in parallel computing. An important problem in sparse matrix computations is how to distribute the sparse matrix and vectors among processors. We review graph, bipartite graph, and hypergraph models for both 1d (row or column) distributions and 2d distributions. A valuable tool is hypergraph partitioning. We present results using the parallel partitioner from the Zoltan toolkit, which is capable of efficiently partitioning very large data sets (matrices).

An important problem in parallel scientific computing is how to distribute work (load) among processors. We wish to minimize the communication cost while maintaining approximate load balance. This problem is often modeled as graph (or hypergraph) partitioning, where the vertices are computational tasks and edges represent data dependencies.

We focus on a common kernel in many numerical algorithms: multiplication of a sparse matrix by a vector. For example, this operation is often the most computationally expensive part of iterative methods for linear systems and eigensystems. More generally, many data dependencies in scientific computation can be modeled as graphs or hypergraphs, which again can be represented as (usually sparse) matrices. The question is how to distribute the nonzero matrix entries (and the vector elements) among processors. The sparse case is quite complicated and is a rich source of combinatorial problems. This problem has been studied in detail in [1] and in [2, Ch.4].

The most common matrix distribution is a one-dimensional (1d) decomposition of either matrix rows or columns. The communication needed for matrix-vector multiplication with a 1d distribution is shown in Figure 1. This structure can be modelled as a graph, a bipartite graph, or a hypergraph, a generalization of a graph. It has been shown [1] that the hypergraph model accurately represents communication volume and is therefore superior to the standard graph model. However, the bipartite graph model can also be used. The bipartite and hypergraph models have another advantage in that they work for nonsymmetric matrices while the graph model assumes symmetry. The communication volume can often be further reduced by using more complex 2d distributions, which again lead to new graph and hypergraph models. In the full paper we will show a new representation of fine-grain 2d matrix distribution using the bipartite graph model.

*Discrete Algorithms and Math Department, Sandia National Laboratories, Albuquerque, NM 87185, USA. egboman@sandia.gov

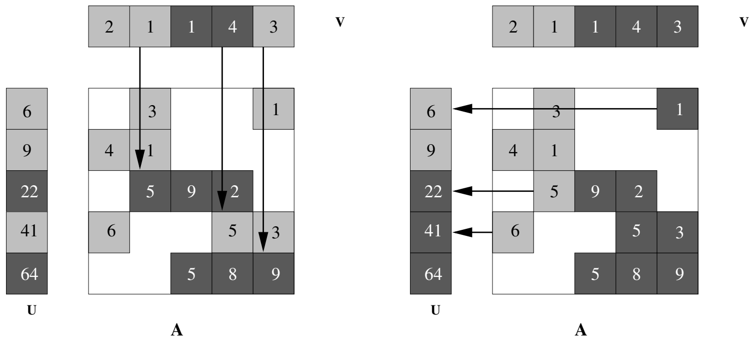

Figure 1: Row (left) and column (right) distribution of a sparse matrix for multiplication $u = Av$. There are only two processors, indicated by dark and light shading, and communication between them is shown with arrows. In this example, the communication volume is three words (integers) in both cases.

The Zoltan toolkit was developed at Sandia National Labs and contains a wide range of partitioning and load-balancing methods, including a parallel hypergraph partitioner [3]. We show results from three different application areas: the Tramonto DFT code for nanoscale fluid simulation, circuit simulation, and DNA electrophoreses. All produce structurally symmetric sparse matrices. We computed row partitions using graph partitioning (ParMetis) and hypergraph partitioning (Zoltan). The reduction in communication volume (cuts) depends on the matrix structure; up to a factor three for circuit problems. The actual run-time in applications is also reduced, but less dramatically than the communication volume indicates. The measured performance increase for the matrix-vector kernel in the Tramonto and Xyce examples was 13% and 18%, respectively, compared to graph partitioning on a 64-processor cluster.

Acknowledgements

The Zoltan hypergraph partitioner was developed with Umit Catalyurek, Karen Devine, and Bob Heaphy. We thank Rob Bisseling and Bruce Hendrickson for many fruitful discussions.

Sandia is a multiprogram laboratory operated by Sandia Corporation, a LockheedMartin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

References

[1] Ü. Çatalyürek and C. Aykanat. Hypergraph-partitioning-based decomposition for parallel sparse-matrix vector multiplication. *IEEE Trans. Parallel Dist. Systems*, 10(7):673–693, 1999.

[2] R. H. Bisseling. *Parallel Scientific Computing: A structured approach using BSP and MPI.* Oxford University Press, 2004.

[3] K.D. Devine, E.G. Boman, R.T. Heaphy, R.H. Bisseling, and U.V. Catalyurek. Parallel hypergraph partitioning for scientific computing. In *Proc. of IPDPS'06*. IEEE, 2006. To appear.