
Grey Ballard
**Reducing Communication Costs for Sparse Matrix
Multiplication within Algebraic Multigrid**

P O Box 969
Mail Stop 9159
Livermore
CA 94550
gmballa@sandia.gov
Grey Ballard
Jonathan Hu
Christopher Siefert

We consider the sequence of sparse matrix-matrix multiplications performed during the setup phase of algebraic multigrid. In particular, we show that the most commonly used parallel algorithm is often not the most communication-efficient one for all of the matrix multiplications involved. By using an alternative algorithm, we show that the communication costs are reduced (in theory and practice), and we demonstrate the performance benefit for both model and real problems on large-scale distributed-memory parallel systems.

We consider the smoothed aggregation multigrid method, and for this talk we focus on a symmetric fine grid operator A . After fine-level rows of A are grouped into coarse-level aggregates, there are three sparse matrix multiplication operations performed to construct the coarse grid operator. The tentative prolongation matrix \hat{P} represents aggregate membership, and the final prolongation matrix P is computed by applying a step of damped Jacobi to \hat{P} , which is structurally equivalent to computing the product $A \cdot \hat{P}$. The coarse grid operator is given by the triple product $A_c = P^T A P$, which is typically performed with two more sparse matrix multiplications: $A \cdot P$ and $P^T \cdot (AP)$.

The standard approach for performing each of these parallel sparse matrix multiplications is to use a row-wise algorithm: for general $C = A \cdot B$, each processor owns a subset of the rows of A , a subset of the rows of B , and computes a subset of rows of C (which matches the distribution of A). The communication required is an exchange of rows of B among processors. We consider as an alternative the outer-product sparse matrix multiplication algorithm, where each processor owns a subset of the rows of A , the corresponding subset of columns of B , and the communication consists of exchanging partially unreduced rows of C among processors (assuming the desired output distribution is row-wise).

Because of its low arithmetic intensity, the performance of parallel algorithms for sparse matrix multiplication is typically bound by the cost of interprocessor communication. Based on communication cost analysis of model problems, we

conclude that the row-wise algorithm is the right choice for computing $A \cdot \hat{P}$ and $A \cdot P$ but the outer-product is the better choice for computing $P^T \cdot (AP)$. For a fine grid operator corresponding to a 3D 27-point stencil on a regular grid, for example, the row-wise algorithm for $P^T(AP)$ requires more than $5\times$ the communication of the outer-product algorithm. Furthermore, using the outer-product algorithm for $P^T \cdot (AP)$ avoids the explicit redistribution of P^T that is typically required in order to use the row-wise algorithm and involves as much interprocessor communication as a sparse matrix multiplication.

We have implemented the outer-product algorithm within the Trilinos software framework, and we compare communication costs and running times of the two approaches for representative matrices. To confirm the theoretical analysis, our experiments include regular-grid stencil matrices for 2D and 3D problems. We also perform tests on realistic problems, including a low Mach fluid dynamics application problem, to demonstrate the benefits of the new approach.