
Christopher M. Siefert
**Won't you be my neighbor? A look at neighbor discovery
for algebraic multigrid**

Sandia National Laboratories
P O Box 5800
MS 1322
Albuquerque
NM 87185-1322
`csiefer@sandia.gov`

At the core of many MPI-based codes, including almost all MPI-based PDE codes, is the idea of a “neighbor” — the list processors to whom point-to-point messages get sent and from whom they are received. Thus, one of the very first bits of parallel code an implementer must develop (or borrow from elsewhere) is code that does arbitrary neighbor discovery. More specifically, this code answers the question, “Who owns a particular degree of freedom?” This code is necessary for things as mundane as assembling sparse matrices and performing matrix-vector products. Thus a non-trivial literature has developed on performing arbitrary neighbor discovery with minimum storage and message cost. With such a tool in the proverbial toolbox, it is not uncommon for the arbitrary neighbor discovery code to be used again and again — even for problems where additional information is available to make neighbor discovery *not* so arbitrary.

In this talk we will focus neighbor discovery in two (related) computational kernels involved in the setup phase of algebraic multigrid — migration of sparse matrices between processors and sparse matrix-matrix multiplication. We will show that for these kernels, passing additional information during point-to-point communication can allow one to perform neighbor discovery more efficiently. For these kernels, it can be performed using a number of messages and amount of storage independent of the number of processors (assuming that the number of neighbors is bounded independently).