New developments in sparse matrix partitioning for parallel computations

Rob Bisseling Daan Pelt

Sparse matrix partitioning is at the core of many parallel computations in the field of scientific computing. Partitioning data well leads to good load balance and possibly tremendous savings in communication time. Optimal partitioning of small sparse matrices can give insights into how larger matrices should be partitioned. We present a branch-and-bound method for optimal bipartitioning of small matrices, and demonstrate its use in developing a data base of optimal communication volumes for about 150 out of the 200 matrices of smallest size in the Florida sparse matrix collection.

Furthermore, we derive lessons for heuristic partitioning, for both structured and unstructured problems and we present the new heuristic overpainting method, and compare it to partitioning using the Mondriaan package.