
Hormozd Gahvari
**Embedding a Performance Model into hypre's
BoomerAMG Solver**

Lawrence Livermore National Laboratory
P O Box 808
L-561
Livermore
CA 94551
`gahvari1@llnl.gov`
William Gropp
Kirk E. Jordan
Martin Schulz
Ulrike M. Yang

Agglomerating data on coarse grids is a commonly used technique in multigrid solvers to reduce communication and improve performance on parallel machines. In our past work, we showed that we can use a performance model at runtime to guide data agglomeration in algebraic multigrid to ensure that it is done in a way that improves performance. However, performance models require measurements of machine parameters and an idea of the factors that affect performance on each machine, which makes incorporating a performance model into production software a difficult task. In this talk, we discuss how we incorporated a performance model into the BoomerAMG solver in the hypre library to control data agglomeration on coarse grids. The model effectively makes decisions on when in the multigrid cycle to agglomerate data, and how much agglomeration to perform. It bases its decisions on runtime performance measurements, and conveniently allows for reuse of the measured parameters and decision information in future solves to save time. This results in improved performance and scalability while also leaving the user free of having to deal with all of the underlying machine details.