
Amanda Bienz
**Reducing Communication Costs in Parallel Algebraic
Multigrid**

202 N Race Street
Apt 312
Urbana
IL 61801
`bienz2@illinois.edu`
Rob Falgout
William Gropp
Luke Olson
Jacob Schroder

Coarse-level fill-in as a result of a Galerkin product often yields increased complexity and large communication costs in algebraic multigrid (AMG) methods. Hence, the parallel efficiency of AMG may be improved by reducing this communication. The communication cost associated with a solver is dependent on the number of network links that are occupied at a given time. This link occupancy is equal to the product of the number of packets sent and the distance of each packet. Hence, the communication cost is composed of the number of messages sent, the size of messages sent, and the number of links traversed by each packet.

The number of messages in AMG, as well as their average size, can be minimized by eliminating unnecessary communication, for example through the use of non-Galerkin coarse grids. While this method can greatly reduce the number of packets sent, it does not take into account the distance these packets must travel. These methods are further improved by targeting the costly communication through the use of topology-aware methods. Removing small messages that traverse few links has little effect on per-iteration cost, yet can have a large impact on convergence. Furthermore, large messages that must travel a long distance can have a larger effect on communication cost than convergence rate, and their removal can be beneficial to the overall performance of AMG. In this talk we give an overview of the parallel efficiency in non-Galerkin methods and present directions that utilize topology to maximize efficiency.