
Paul Eller
**Non-blocking conjugate gradient methods for extreme
scale computing**

201 N Goodwin Ave
Urbana
IL 61801
`eller3@illinois.edu`
William Gropp

Many scientific and engineering applications use Krylov subspace methods to solve large systems of linear equations. For extreme scale parallel computing systems, the dot products in these methods (implemented using allreduce operations in MPI) can limit performance because they are a synchronization point or barrier. Therefore we seek to develop Krylov subspace methods that avoid blocking allreduce operations and provide greater parallel efficiency.

We present a rearranged preconditioned conjugate gradient method that overlaps communication in the dot products and computation to improve performance by hiding the cost of allreduces. These rearranged methods optimize memory access patterns to limit the cost of using additional vector operations that are introduced as part of the rearrangement of the algorithm. These methods provide the ability to hide message latency by overlapping communication with more time consuming computations so that messages are received prior to their use. This can allow the methods to avoid blocking due to waiting for messages and maintain better performance despite imbalances in communication time on different nodes caused by the network.

Performance models are developed to allow us to better understand the best case theoretical performance of these methods with different input parameters. We measured the performance for this method, a pipelined conjugate gradient method, a single allreduce conjugate gradient method, and the standard conjugate gradient method on Blue Waters in order to improve and validate both our performance models and our implementations.