
Pieter Ghysels
Hiding global communication in the GMRES algorithm

Middelheimcampus M G 113
Middelheimlaan 1
2020 Antwerp
Belgium
`pieter.ghysels@ua.ac.be`
Tom Ashby
Karl Meerbergen
Wim I. Vanroose

In many Krylov solvers running on current large scale compute clusters, global reductions for both inner products and norms is starting to become a bottleneck; a problem which will only get worse as we approach the era of exascale machines. Although flop rates keep increasing through additional parallelism, link latencies hit their physical limits. As compute nodes are often distributed over large datacenters, these latencies start to play an important role in algorithm design. Furthermore, load imbalance, system noise and processor variability (due to extreme reduction in scale and voltage) will make global communication an even more costly global *synchronization* step [1].

Recently there has been much interest in so called s -step Krylov methods [2] that expand the Krylov base with s new vectors at once before performing an orthogonalization step. This reduces the number of global communications by a factor s .

In this work, we propose the use of non-blocking or asynchronous global collective operations to hide the latency of global communication in the GMRES algorithm. A standard GMRES iteration [3], based on classical (iterated) Gram-Schmidt orthogonalization, requires at least two global reductions, one for orthogonalization and one for normalization. First of all, we show that this can be reduced to just a single reduction. To avoid stability problems, the matrix vector product $z_{i+1} = Av_i$ is modified with a shift σ_i as $z_{i+1} = (A - \sigma_i I) v_i$. Proper choices of the shifts are discussed. Furthermore, we present a class of pipelined GMRES algorithms that allow overlapping of the inner product latency with other computations and communications. Depending on the total reduction latency, a larger pipelining depth l can be used. The resulting algorithm keeps two sets of Krylov base vectors

$$V_{i-l+1} = [v_0, v_1, \dots, v_{i-l}] \quad \text{and} \quad Z_{i+1} = [z_0, z_1, \dots, z_{i-l}, z_{i-l+1}, \dots, z_i],$$

with $V^T V = I$. Iteration i expands the V_{i-l+1} base using only results of reductions started l iterations ago while the Z_{i+1} base is expanded by a vector created

by the (shifted) matrix vector product. To prevent a blow-up of the condition number of the Z_i base, a *correction* is applied to Z_{i+1} , which introduces some redundant calculations compared to the original GMRES algorithm.

The numerical stability of the presented GMRES variation is examined using a collection of test matrices from applications. To assess the scalability and overall performance of the pipelined algorithms, we developed an analytical performance model, which allows us to study scaling aspects and determine bottlenecks on large numbers of nodes. As non-blocking collective operations are proposed for the upcoming MPI-3 standard, they will likely become available in most MPI libraries soon.

We compare our approach with s -step or *communication-avoiding* GMRES, as presented in [2]. We present some preliminary results on pipelining applied to the CG algorithm, for which the results look promising since the short recurrence keeps the amount of redundant calculations small.

In the pipelined GMRES algorithms, data dependencies between subsequent steps in the algorithm have been relaxed. As a result, the network hardware constraints can be relaxed. Apart from hiding global reduction latencies on large distributed clusters, this could also allow more efficient fine grained work scheduling on current multicore and heterogeneous (CPU + accelerator) architectures.

Bibliography

- [1] T. Hoefer, T. Schneider, A. Lumsdaine, *Characterizing the Influence of System Noise on Large-Scale Applications by Simulation*. International Conference for High Performance Computing, Networking, Storage and Analysis (SC'10), Nov, 2010.
- [2] M. Hoemmen, *Communication-avoiding Krylov subspace methods*. PhD, University of California, 2010.
- [3] Y. Saad, M.H. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*. SIAM J. Sci. Stat. Comput., 7:3, 856-869, 1986.