SAND2017-10434D

Asynchronous Iterative Solvers: The ACHILES Library and Domain Decomposition Methods

Erik Boman, Siva Rajamanickam[1], Christian Glusa[1],
Edmond Chow[2], Paritosh Ramanan[1,2]
[1]*Sandia National Laboratories, Albuquerque, NM*
[2]*Georgia Tech, Atlanta, GA*

Abstract

We investigate the use of *asynchronous* parallel iterative methods to solve large linear systems on future extreme-scale computers. Our two contributions so far are:

1) ACHILES, a software library for AsynCHronous Iterative LinEar Solvers.
2) Experiments using the Restricted Additive Schwarz (RAS) [1] domain decomposition method (synchronous and asynchronous).

Motivation

Communication and synchronization costs are becoming increasingly important on extreme-scale parallel computers. Krylov methods often give good convergence for linear systems but they require global communication and synchronization for inner products (orthogonalization). Approaches such as pipelined and s-step (CA) methods can reduce such communication but not eliminate it. We investigate a quite different approach: Use a fixed-point (stationary) iterative method, without Krylov. Convergence may be slower but communication is reduced to nearest-neighbors and this is well suited for asynchronous methods. Chazan & Miranker [2] first suggested *chaotic relaxation* where processes can work at their own speed without global synchronization. Such methods are robust to load imbalance and system noise, and may thus be useful at extreme scale. We found that although there is some theoretical analysis for asynchronous iterative methods [4], very little work has been implemented and tested empirically on large parallel systems. Further, we wish to employ sophisticated iterative methods with better convergence properties.

Approach

The ACHILES approach is based on one-sided MPI calls using Remote Memory Access (RMA). We use MPI since it is standardized and widely available. It typically provides better performance than e.g. PGAS languages. Note we wish to support truly asynchronous methods, not just non-blocking overlap of computation and communication.
We investigate Schwarz-type domain decomposition as our parallel iterative method. Our initial experiments use 1-level RAS; we plan to extend this to 2-level methods and optimized Schwarz. The synchronous RAS method is:

$$x^{k+1} = x^k + M^{-1}(b - Ax^k)$$

Where

$$M = \sum_i R_i^T A_i^{-1} R_i$$

We use RAS as an iterative method, not as a preconditioner. We need to communicate either the solution or the residual at each update. The subdomain solve is assumed to be exact here, but could possibly be inexact (local iterative method).

Results

Our first test is sparse-matrix vector multiply (SpMV), a key kernel in scientific computing. This is not an iterative method, but mainly to verify that ACHILES performance is similar to that of current software, in particular Trilinos/Tpetra. These tests were run on a Intel Xeon Phi (KNL).

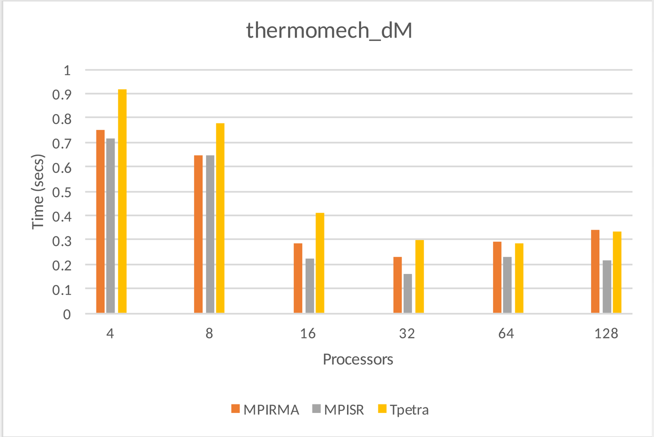

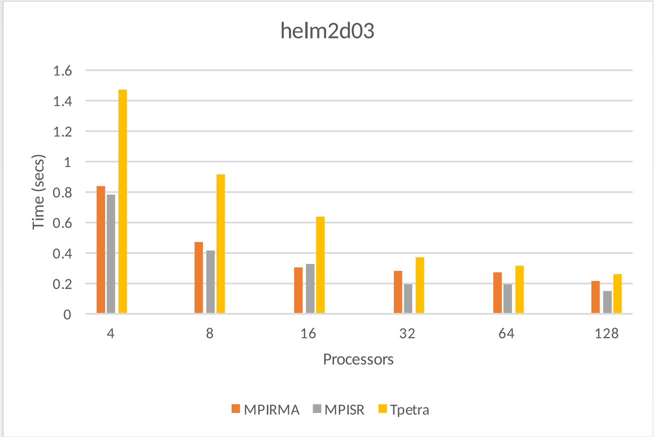

Our second test is a parallel implementation of the Restricted Additive Schwarz (RAS) method. We use it as a stationary iterative method, not as a preconditioner in a Krylov method. We assign one subdomain per MPI rank and typically one rank per core. In the asynchronous version, some processors may work faster than others (e.g. due to load imbalance or system noise). Each process uses the most recent information available from its neighbors. The concept of "iterations" is no longer meaningful so we need to look at times. We show results from a shared-memory workstation (16 cores), where communication is fast. We expect the results to change for distributed-memory parallel machines.

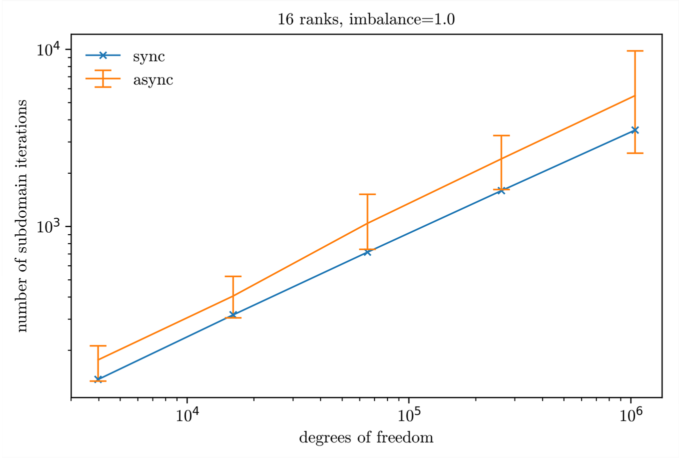

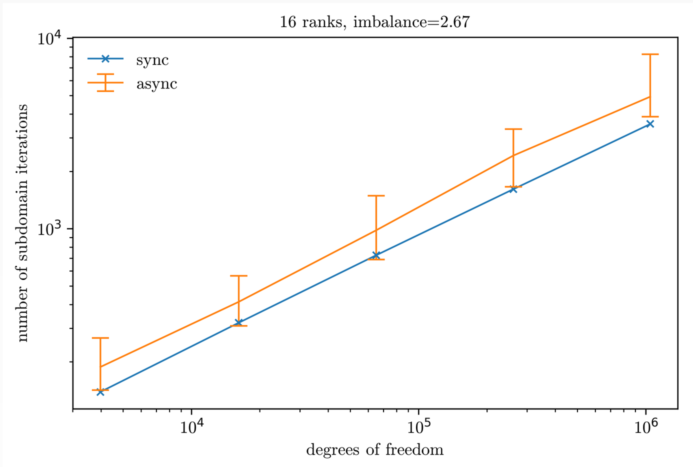

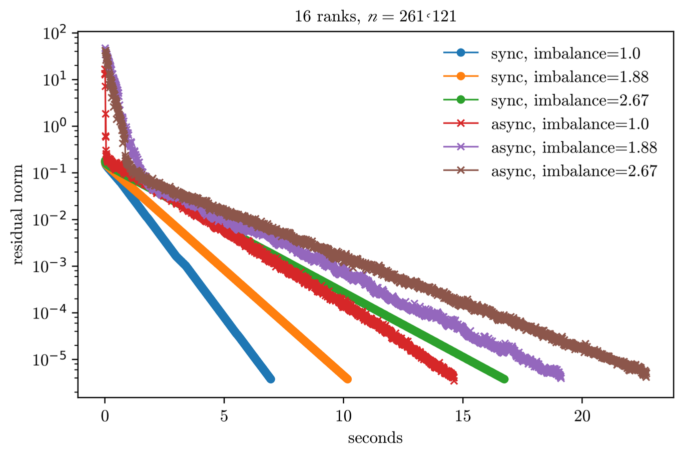

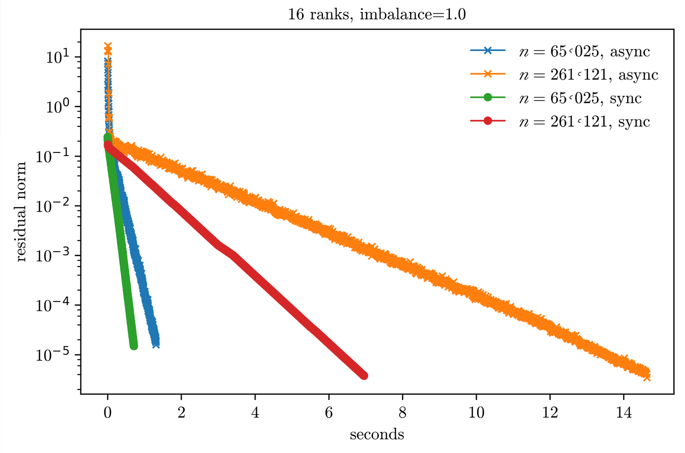

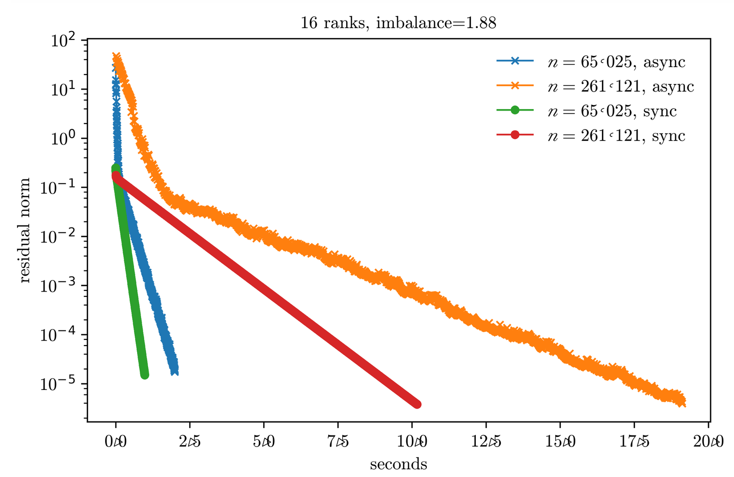

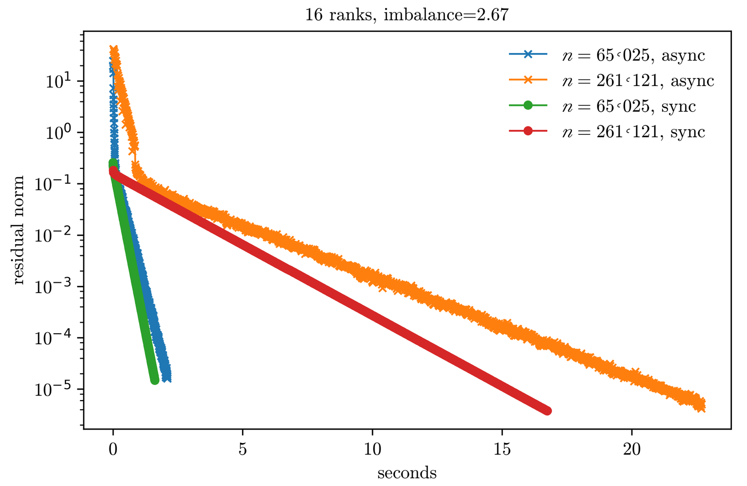

Conclusions and Future Work

We have presented preliminary work and results. We conclude that the ACHILES software library has high performance relative to current solver software, and is a useful foundation for testing asynchronous methods. RAS does not scale well with problem size and #processors so we plan to pursue two options: (a) Optimized Schwarz methods [3]. Optimizing the boundary conditions can greatly improve convergence. (b) Coarse-grid correction. Two-level methods are known to be highly scalable. The key challenge is how to handle the coarse grid in asynchronous methods. Asynchronous convergence detection (termination) is challenging future work.

This project is long-term fundamental research and we target mainly proof-of-concept work. Impact on DOE mission/applications is expected in the 10-year time frame, assuming continued funding.

Areas in which we can help

Our target is to help a wide variety of applications in computational science & engineering that need to solve large linear systems on extreme-scale computers. A project with similar goals that we view as a potential partner, is the ECP/PEEKS project. However, we have a longer time horizon and focus more on basic research.

Areas in which we need help

We could need help from computer science projects to optimize the 1-sided MPI communication on leadership class computers.
The distributed asynchronous termination detection problem is a challenging mathematics problem where help would be welcome.

References

[1] X.C. Cai and M. Sarkis, "A Restricted Additive Schwarz Preconditioner for General Sparse Linear Systems", SIAM J. Sci. Comp., 21(2), 1999.
[2] D. Chazan and W. Miranker, "Chaotic Relaxation", Lin. Alg. App.,2(2) 1969.
[3] M. J. Gander, "Optimized Schwarz Methods", SIAM J. Num. Anal. 44(2), 2006.
[4] Magoules, Szyld, Venet, "Asynchronous optimized Schwarz methods with and without overlap", Num. Math. 137.1 (2017): 199-227.

Georgia Tech

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.