

# Evaluating asynchronous Schwarz solvers for Exascale.

*Wednesday, 6 November 2019 11:15 (30 minutes)*

With the commencement of the exascale computing era, we realize that the majority of the leadership supercomputers are heterogeneous and massively parallel even on a single node with multiple co-processors such as GPU's and multiple cores on each node. For example, ORNL's Summit accumulates six NVIDIA Tesla V100's and 42 core IBM Power9's on each node.

At this scale of parallelism, the traditional bulk-synchronous programming model will not be able to leverage the compute power of the hardware efficiently. Hence, it is necessary to develop and study asynchronous algorithms that circumvent this issue. The Schwarz methods are a class of Domain decomposition solvers which allow for decomposing the main domain into smaller subdomains, solve the sub-domain problems in parallel and exchange the information between iterations.

In this study, we examine the asynchronous version of one of these (Restricted Additive Schwarz) solvers where we do not explicitly synchronize, but allow for communication of the data between the subdomains to be completely asynchronous. Thereby, we remove the bulk synchronous nature of the algorithm. We accomplish this using the RDMA functions of MPI. We study the benefits of using such an asynchronous solver over its synchronous counterpart on both multi-core and on multiple GPU's. Detecting convergence in these asynchronous solvers is tricky, and careful consideration of hardware is necessary to make efficient use of the hardware. We show that this concept can nevertheless render attractive runtime benefits over the synchronous counterparts.

**Primary authors:** NAYAK, Pratik (Karlsruhe Institute of Technology); ANZT, Hartwig (Karlsruhe Institute of Technology)

**Presenter:** NAYAK, Pratik (Karlsruhe Institute of Technology)

**Session Classification:** Day II