

Exascale Workflow for NekRS

We discuss two related tools for exascale workflows with NekRS. The first is parRSB, which is the parallel partitioner for Nek5000/RS that is based on recursive coordinate and recursive spectral bisection. The second is a new, communication-minimal, coarse-grid solve that is targeting exascale platforms such as Frontier/Aurora. The pressure solvers in Nek5000/RS rely on p-multigrid (pMG) with multilevel Schwarz smoothers for preconditioning an outer GMRES iteration. The coarsest level in pMG corresponds to the unstructured mesh comprising $p=1$ trilinear elements. For such meshes, the number of unknowns is approximately equal to the number of elements, E , and the number of unknowns per MPI rank is therefore approximately E/P . On GPU-based platforms, E/P is typically around 10,000, which is too small to fully occupy a GPU. Nek5000/RS mesh sizes, however, can be as large as $E = 1$ billion, so the over-all coarse problem is not small. To reduce communication in the coarse solve, we employ a two-level Schwarz method with a local solve that requires only a single pair of nearest-neighbor communications and with a global reduced-system solve that requires only $2 \log_2 P$ messages. On Frontier, the communication for the reduced-system solve is on par with a single MPI-allreduce.

Relevance for Nek [100 words max]

Paul

Primary author: RATHNAYAKE, Thilina

Co-authors: OLSON, Luke; FISCHER, Paul

Presenter: FISCHER, Paul