Domain Decomposition for Exascale Computing

Oliver Rheinbach[1]

Preconditioned Newton-Krylov algorithms using scalable multilevel preconditioners from domain decomposition (DD) or multigrid (MG) have been the workhorse for the parallel solution of nonlinear implicit finite element problems for several decades. In these methods, the problem is first linearized and then decomposed into parallel problems. On the contrary, in recent scalable nonlinear domain decomposition methods the nonlinear problem is directly decomposed into concurrent problems, i.e., before Newton linearization. This approach increases the concurrency of the algorithm, reduces the need for synchronization and can also reduce energy consumption.

We discuss recent nonlinear domain decomposition of the Nonlinear FETI-DP (Finite Element Tearing and Interconnecting) or Nonlinear BDDC (Balancing Domain Decomposition by Constraints) type to solve nonlinear hyperelasticity or plasticity problems and show parallel scalability to up to 800000 cores of the Mira BGQ supercomputer and 200000 cores of the Theta KNL supercomputer (both Argonne National Laboratory). We then apply these methods within a highly parallel two-scale numerical homogenization scheme using millions of MPI ranks for the simulation of micro-heterogeneous materials.

[1] Technische Universität Bergakademie Freiberg, Mathematics and Computer Sciences , Germany, oliver.rheinbach@math.tu-freibergt.de