Multigrid Solvers in Space and Time for Highly Concurrent Architectures

Robert Falgout

Center for Applied Scientific Computing (CASC)

Lawrence Livermore National Lab

Multigrid methods are important techniques for efficiently solving huge systems and they have already been shown to scale effectively on millions of cores. However, one of the major challenges facing computational science with future architectures is that faster compute speeds will be achieved through greater concurrency (more cores), since clock speeds are no longer increasing. Current petascale computers already have millions of cores, but future exascale machines are expected to have billions. This immense degree of parallelism requires a similar level of concurrency in the algorithms that run on them. One consequence of this is that time integration by traditional time marching will become a sequential bottleneck.

In this talk, we will discuss the multigrid method and its role in scientific computing. We will also discuss our efforts to develop multigrid algorithms for parallel time integration. The approach we use is based on multigrid reduction techniques and has the advantage of being easily integrated into existing codes because it builds directly on the original time-stepping scheme. Results for a variety of applications will be presented.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.