Title: Scalable Parallel-in-Time with Multigrid: Theory, Applications, and Recent Developments

Abstract: Recent growth in parallel-in-time methods is currently being driven by the rapidly changing nature of computer architectures. Future speedups will come through ever increasing numbers of processors (cores), not stagnant clock speeds. Previously, increasing clock-speeds compensated for traditional sequential time stepping when the problem size increased. However, this is no longer the case, leading to the sequential time integration bottleneck and subsequent need for parallelism in time. In this talk, we present the optimal-scaling multigrid reduction in time (MGRIT) method. MGRIT applies multigrid to the time dimension by solving the (non)linear systems that arise when solving for multiple time steps simultaneously. The result is a versatile approach that is nonintrusive and wraps existing time evolution codes. MGRIT allows for various time discretizations (e.g., Runge-Kutta and multistep) and for adaptive refinement/coarsening in time and space. Recent theoretical, methodological, and practical results for a variety of problems will be presented, e.g., nonlinear diffusion, powergrid systems, advection, and compressible Navier-Stokes. In particular, we will discuss recent advancements for MGRIT and hyperbolic problems, where semi-Lagrangian coarse time-grids are used for increased stability and accuracy, and also for MGRIT and chaotic problems, where an enhanced coarse-grid equation allows for meaningful speedups on truly chaotic problems. Lastly, a novel application of time-parallelism to deep learning is presented, where the layer dimension of deep neural networks is explicitly parallelized by an MGRIT method.