Parallel-in-time integration based on spectral deferred corrections

Daniel Ruprecht, Technische Universität Hamburg

The still rapidly growing computing power of high-performance computing systems is fueled by an extremely rapid increase in available compute nodes, processors and cores. The currently fastest supercomputer, Frontier, at Oak Ridge National Laboratory, features more than 8.5 million cores and was the first computer to reach Exascale performance. However, translating this computing power into application performance is a huge challenge: among many other things, it requires numerical algorithms that are inherently designed to exploit concurrency. Essentially, algorithms need to be made up of lots of sub-tasks that can be solved independently.

For the numerical approximation of partial differential equations, the standard approach is parallelization by decomposition of the spatial computational domain. However, as more and more processes are used, the parts of the spatial domain handled by each process become smaller and smaller. Eventually, a code will spend most of its time transferring data between processes at which point adding more resources will no longer reduce solution times. These diminishing returns mean that parallelizing in space alone is not going to be enough to effectively use state-of-the-art supercomputers for models based on time-dependent partial differential equations. This fact is driving a growing interest in so-called parallel-in-time algorithms: integrators for initial value problems that can offer at least some degree of concurrency.

Parallel-in-time methods come in two flavours: parallel-across-the-steps algorithms like Parareal or MGRIT compute the solution on many time steps in parallel. They can scale to a large numbers of cores but can be challenging to use because they require the definition of a coarse level model, a process that is still more an art than a science. These algorithms also often have relatively poor parallel efficiency. By contrast, parallel-across-the-method approaches are "normal" time stepping schemes but designed in a way that some computations in each time step can be performed in parallel. The most obvious example would be a Runge-Kutta method with a diagonal Butcher table, where all stages can be computed at the same time. While they offer more limited parallelism, they are easier to use and can provide good parallel efficiency.

In this talk, I will describe a new approach to the design of parallel-across-the-method algorithms based on spectral deferred correction (SDC). SDC is an iterative solver for collocation methods/fully implicit Runge-Kutta methods that, in essence, uses a low-order method as preconditioner for a Picard iteration. Using a diagonal preconditioner allows to parallelize each iteration but requires carefully chosen coefficients to maintain good convergence and stability. I will present some theory on how to find good coefficients and show numerical examples illustrationg stability and efficiency of the approach. Finally, I will show performance results for an implementation of parallel SDC into the ICON-O ocean model run on a single node of the JUWELS cluster at Jülich Supercomputing Center.

This is joint work with T. Lunet, P. Freese, S. Götschel (TUHH), G. Čaklovic (KIT) and M. Schreiber (Grenoble).