Shared–Memory Parallel Implementation of High-Order Asynchronous Spacetime Discontinuous Galerkin Methods

Amit Madhukar[1], Robert B. Haber[1], Volodymyr Kindratenko[1], and Reza Abedi[2]

[1]University of Illinois at Urbana-Champaign
[2]University of Tennessee Knoxville (UTK) / Space Institute (UTSI)

We present a new parallel-adaptive shared-memory implementation of high-order asynchronous spacetime discontinuous Galerkin methods for hyperbolic problems; see for example [1] for a serial implementation. We use the *Tent Pitcher* algorithm [2] to generate fully unstructured spacetime meshes that satisfy a *causality constraint* to enable locally implicit aSDG solutions. These involve local Galerkin projections on a sequence of spacetime patches (small clusters of spacetime finite elements) that inherit the stability of implicit solvers while the overall solution exhibits the linear computational complexity reminiscent of explicit methods. The duration of each patch is determined independently and is not restricted by the order of the local basis. The processes of constructing and solving patches are interleaved, asynchronous and share the same granularity, so most of the algorithm is embarrassingly parallel.

Advancing a conforming space-like *front mesh* through the spacetime analysis domain subject to the causality constraint is the heart of the Tent Pitcher algorithm. This front mesh is the only global data structure. We copy fragments of the front mesh into private data structures called footprints that render the generation and solution of each new patch embarrassingly parallel. Gather and scatter operations between the front and footprint meshes are not embarrassingly parallel, but they represent a tiny fraction of the overall computational expense.

In this presentation, we focus on the architectural details of our parallel implementation. These include coarse-grained, patch-level parallelization across multiple cores; strict separation of shared front-mesh and private patch-level data; task queues and use of hardware threads for asynchronous execution of major operations, such as footprint construction, patch generation, patch solution, and front updates; and NUMA-aware data storage. We present numerical performance results to demonstrate near perfect scaling for high-order models of varying order and the effects of various software optimizations.

References:

[1] R. Abedi, B. Petracovici, and R. B. Haber. "A spacetime discontinuous Galerkin method for linearized elastodynamics with element-wise momentum balance. *Comp. Methods Appl. Mechs. Engnrg.* 195(25-28), 3247–3273 (2006).

[2] Shripad Thite. *Spacetime Meshing for Discontinuous Galerkin Methods*. Ph.D. thesis, Dept. Computer Science, Univ. Illinois Urbana-Champaign, August (2005).