Parallel–Adaptive Implementation of Asynchronous Spacetime Discontinuous Galerkin Methods

Robert B. Haber[1], Amit Madhukar[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 implementation for asynchronous spacetime discontinuous Galerkin method for hyperbolic problems; see for example [1]. In order to circumvent synchronous bottlenecks that might prevent efficient use of next-generation exascale platforms, we set aside established techniques and abstractions of parallel computation, such as the domain decomposition method (DDM) and the bulk synchronous parallel (synchronous time marching) model, and replace them with a scalable, barrier-free asynchronous solution scheme and localized fine-grain spacetime adaptive meshing.

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 and naturally share the same granularity. Locality and shared granularity render most of the algorithm embarrassingly parallel.

Advancing a conforming space-like *front mesh* through the spacetime analysis domain is central to the Tent Pitcher algorithm. We maintain the front mesh as a global data structure to avoid DDM load balancing, a major synchronous bottleneck. Aside from querying and updating the front mesh, the algorithm is embarrassingly parallel. Adaptive meshing is implemented via modifications of the front mesh that govern refinement and coarsening of spacetime patches. However, because we must maintain the conforming property of the front mesh, these adaptive modifications are not naturally local. We present a *Lazy Refinement* method, first proposed but not implemented in [2], that localizes adaptive updates to the front mesh. Thus, adaptive front modifications, spacetime patch generation, and patch solutions share a common granularity and are embarrassingly parallel. We describe a software architecture that implements this scheme and present numerical results that demonstrate excellent scalability on a shared-memory host. Progress toward a distributed implementation and meshing in three spatial dimensions and time will also be discussed.

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).