Performance and Scalability of Parallel–Adaptive Asynchronous Spacetime Discontinuous Galerkin Methods

Amit Mudhakar[1], Robert Haber[1], Reza Abedi[2], Volodymyr Kindratenko[1]
[1]University of Illinois at Urbana-Champaign
[2]University of Tennessee Knoxville (UTK) / Space Institute (UTSI)

Abstract

The asynchronous Spacetime Discontinuous Galerkin (aSDG) method [1] is a powerful solution scheme for hyperbolic systems. It features unconditional stability, conservation over every spacetime cell, linear computational complexity, and support for arbitrarily high-order elements. However, its most powerful potential advantages are compatibility with an extremely powerful and dynamic form of adaptive spacetime meshing and a nearly perfect structure for parallel implementations. Ideally, these features would be exploited simultaneously. Unfortunately, it was not possible to evolve domain decompositions fast enough to maintain load balance in the face of dynamic adaptive meshing. Faced with a choice between serial adaptive solutions and parallel solutions on non-adaptive meshes, we chose the former more powerful option. Indeed, serial adaptive aSDG solvers typically outperformed conventional solvers running in parallel on large clusters, especially for multi-scale problems or problems with rapidly evolving domain geometry; cf. recent work on dynamic fracture [2].

This presentation describes recent work on a new approach to parallel–adaptive aSDG computations in which individual patch solutions, the basic unit of computation in aSDG solvers, replace subdomains as the means to define parallel solution tasks. A new 'lazy' adaptive meshing scheme parallelizes at the same patch-level granularity. The new parallel work unit comprises construction, solution, and local mesh adaptation for a single patch. A simple round-robin strategy preserves load balance by distributing these unit tasks across the available cores. We pin hyperthreaded meshing and solver threads to each core. The meshing thread handles patch construction and patch-local adaptive meshing; the solver thread carries out the more compute intensive patch solution. The software architecture operates asynchronously, at both the thread level within each core and the patch processing level across cores. We describe various optimizations for improved single-core performance and parallel scaling efficiency. For shared-memory parallel platforms we report parallel scaling efficiencies of 97% or better for non-adaptive runs and efficiencies above 90% for parallel–adaptive runs. We close with a discussion of continuing work, including a new implementation for large-scale distributed platforms.

References:

[1] R. Abedi, R. B. Haber, B. Petracovici. A spacetime discontinuous Galerkin method for elastodynamics with element-level balance of linear momentum, *Comput. Methods Appl. Mech. Eng.* 195 (2006) 3247–3273.
[2] R. Abedi, and R. Haber. Spacetime simulation of dynamic fracture with crack closure and frictional sliding, Advanced Modeling and Simulation in Engineering Sciences 5 (2018). doi:10.1186/s40323-018-0116-5