Kenneth Jansen

Session 2

TITLE: Exascale Simulation of Turbulence

ABSTRACT:

For several decades, scale-resolving turbulence simulations have served as grand challenge problems capable of saturating the largest available computer's resources for as long (or longer) than the resource was available. Larger machines have allowed simulations at larger Reynolds numbers but that growth is slow due to the three-dimensional and unsteady nature of turbulence that drives the cost (degree-of-freedom count multiplied by time step count) to scale with the Reynolds number to powers ranging from roughly 1 to 4 depending upon the scale-resolving method chosen with higher exponents coming from more scales resolved vs. modeled.

While exponents on the Reynolds numbers are important, larger machines have also allowed expansion of geometric and flow complexity. In this talk, we will discuss the development of finite element methods for application to scale-resolving simulations on exascale computational resources. The methods developed provide very low dissipation and higher order accurate discretizations. Their ability to use unstructured grids to match grid resolution to the local needs of the scale resolving simulation makes them particularly attractive and efficient for complex geometry flows and more fundamental flows with a large spatial variation in the smallest required scale that must be resolved. The talk also provides a contrast of classical and newer approaches to equation formation, assembly, and solution that accounts for GPU hardware.

It is important to note that getting the solver to scale well and make efficient use of the emerging hardware is only one of the challenges of exascale computing. Other significant challenges addressed in this talk include preparing the inputs for simulations at this scale (pre-processing) and extracting meaningful insight from the massive spatial-temporal data stream that an exascale turbulence simulation produces. Insight extraction in the classical sense, writing data to files, becomes challenging due the relatively slow growth of I/O resources relative to computational resources which motivates development of in situ data analytics and compression techniques. This talk will discuss the general and specific challenges that have been addressed for unstructured grids in this area and will close with demonstrations.

If you are also looking for co-author/collaborator lists

AUTHORS:

Kenneth E Jansen, James R Wright, Jeff Hadley, Jed Brown, John A Evans, Alireza Doostan

(University of Colorado, Boulder)

Riccardo Balin, Umesh Unnikrishnan, Kris Rowe

(Argonne National Laboratory)

Mark Shephard

(Rensselaer Polytechnic Institute)