The 18th International Conference on Finite Elements in Flow Problems (FEF2015)
16-18 March 2015, Taipei, Taiwan

Comparison of Various Explicit CFD schemes applied to GPU Computing with Unstructured Tetrahedral Grids

Josh (Chieh-Ming) Chang[1] and Matthew Ross Smith[2]

[1] Department of Mechanical Engineering,
National Cheng Kung University, Tainan, Taiwan
leader0319@hotmail.com
[2] Department of Mechanical Engineering,
National Cheng Kung University, Tainan, Taiwan.
msmith@mail.ncku.edu.tw

Key Words: *Graphics Processing Units, GPU, Parallel Computing, Unstructured Grid, Finite Volume Method, Computational Fluid Dynamics, CFD, Riemann solver.*

The use of Graphics Processing Units (GPUs) to accelerate explicit Finite Volume Method (FVM) based schemes has been well established over the previous 5 years. Conventional wisdom and research results have demonstrated that a high degree of locality in the numerical scheme translates to a high parallel efficiency when applied to GPU (or indeed, any vector) computation. While this approach is generally tested using Cartesian or structured grids, application of high-locality schemes to GPU computing using unstructured grids is still lacking. Here we compare and contrast several FVM CFD schemes – both local and non-local – to computation using GPU devices. The unstructured grid consists of tetrahedrons, and double precision computation is required to ensure correct local to global coordinate transformation for flux computations. Specifically, we investigate three schemes – (i) the approximate Riemann solver proposed by Jacobs in 1999, which is highly computationally intensive and non-local, (ii) the conventional Harten, Lax and van Leer (HLL) method, which is computationally inexpensive and non-local, and (iii) the Split HLL (SHLL) which is a computational inexpensive and local variant of the HLL method. These three techniques, while all falling under the traditional family of Godunov type approaches for CFD, vary greatly in their data access patterns the the ratio of on-device memory access to computation ratio. Hence, we discuss the performance characteristics of these approaches when applied to the same computational (unstructured) grid and strategies for increasing the performance of each. The solvers are finally applied to an industrial application and results compared.