OpenFOAM with GPU Solver Support

Matt Martineau[1], Simone Bnà[2], Stan Posey[3], Filippo Spiga[4]

[1]*NVIDIA Ltd., Bristol, UK, Developer Technology Group, mmartineau@nvidia.com*
[2]*CINECA SCAI, Bologna, IT, Numerical Software Development, s.bn@cineca.it*
[3]*NVIDIA Corporation, Santa Clara, CA, USA, CFD Program Mgmt., sposey@nvidia.com*
[4]*NVIDIA Ltd., Bristol, UK, HPC Developer Relations, fspiga@nvidia.com*

Current trends in high performance computing include the use of graphics processing units (GPUs) as massively parallel co-processors to CPUs that can accelerate numerical operations common to computational fluid dynamics (CFD) solvers. GPU-parallel CFD achieves speedups with CPUs from additional fine grain, or second-level parallelism under existing CPU-based distributed memory, or first-level scalable parallelism. For most CFD implementations, the GPU focus is on implicit sparse iterative solvers whereby linear algebra matrix operations that would be processed on the CPU are offloaded to the GPU for numerical acceleration, resulting in an overall simulation speedup.

During 2019, the ESI-OpenCFD OpenFOAM HPC Technical Committee introduced the PETSc4FOAM library that permits the plug-in of external solvers that conform to PETSc formats. NVIDIA developed an external solver for GPU offload of OpenFOAM pressure solves based on the AmgX solver library that was first introduced by NVIDIA in 2012. In the initial NVIDIA implementation, the OpenFOAM system matrix is copied from the CPU to the GPU, and the AmgX library applies an AMG preconditioner to a preconditioned conjugate gradient (PCG) linear solve for overall solver acceleration. Results are copied back to the CPU which completes the OpenFOAM job as the end-user normally observes.

This work will preview details of the AmgX development and its integration with OpenFOAM for multi-GPU and multi-node computations. Initial experiments with the standard benchmarks of 3d lid-driven cavity and motorbike cases demonstrate that AmgX can achieve as much as a 9x speedup of the pressure solver on an A100 GPU vs. an OpenFOAM GAMG-PCG solve on an x86 CPU server node with dual-socket 20 core “Broadwell” CPUs. With this community-supported solution, frequent software updates that are made to community-based libraries like PETSc, AmgX, and PETSc4FOAM will ensure that OpenFOAM users naturally benefit from the latest system software, compilers, system and processor hardware architectures, and OpenFOAM future releases.

Future work will be described which will include applications to more complex geometries and turbulence treatment such as LES, further optimizations for strong scaling across computational nodes of multiple GPUs, and the potential of more OpenFOAM coded applied to GPU acceleration such as matrix assembly procedures. The work is a collaboration among members of the HPC Technical Committee and current and future contributions will be described towards an execution plan for a community supported GPU-enabled OpenFOAM.