The Effect of InfiniBand In-Network Computing on OpenFoam Simulations

Ophir Maor,Yong Qin, David Cho, Gerardo Cisneros-Stoianowski

HPC-AI Advisory Council

From concept to engineering, and from design to test and manufacturing, engineers from wide ranges of industries face ever increasing needs for complex, realistic models to analyze the most challenging industrial problems; Finite Element Analysis is performed in an effort to secure quality and speed up the development process. Powerful virtual development software is developed to tackle these needs for the finite element-based Computational Fluid Dynamics (CFD) simulations such as OpenFoam with superior robustness, speed, and accuracy. Those simulations are designed to carry out on large-scale computational High-Performance Computing (HPC) systems effectively.

The latest revolution in HPC is the effort around the co-design approach, a collaborative effort to reach Exascale performance by taking a holistic system-level approach to fundamental performance improvements, is In-Network Computing. The CPU-centric approach has reached the limits of its scalability in several aspects, and In-Network Computing acting as “distributed co-processor” can handle and accelerates performance of various data algorithms, such as reductions and more.

The past focus for smart interconnects development was to offload the network functions from the CPU to the network. With the new efforts in the co-design approach, the new generation of smart interconnects will also offload data algorithms that will be managed within the network, allowing users to run these algorithms as the data being transferred within the system interconnect, rather than waiting for the data to reach the CPU. This technology is being referred to as In-Network Computing, which is the leading approach to achieve performance and scalability for Exascale systems. In-Network Computing transforms the data center interconnect to become a “distributed CPU”, and “distributed memory”, enables to overcome performance walls and to enable faster and more scalable data analysis.

The new generation of HDR 200G InfiniBand In-Network Computing technology includes several elements - Scalable Hierarchical Aggregation and Reduction Protocol (SHARP), a technology that enables to execute data reduction algorithm on the network devices instead of the host based processor. Other elements include smart Tag Matching and rendezvoused protocol, and more. These technologies are in use at some of the recent large scale supercomputers around the world, including the top TOP500 platforms. HPC-AI Advisory Council performed deep investigations on OpenFoam to evaluate its performance and scaling capabilities when using InfiniBand interconnect. The study reveals the influence of the applications on runtime, scalability and performance of the simulations.