THE EFFECT OF HPC CLUSTER ARCHITECTURE ON THE SCALABILITY OF CAE SIMULATIONS

Pak Lui

HPC Advisory Council

Application Performance Manager

KEYWORDS

High-Performance Computing, HPC, system, CPU, network, CPU, architecture, Ethernet, InfiniBand, network, offloading, scalability, performance

ABSTRACT

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 with superior robustness, speed, and accuracy. Those simulations are designed to carry out on large-scale computational High-Performance Computing (HPC) systems effectively.

The breakthrough in HPC parallel computing that allows such complex analyses to be performed that generate the high-quality results, while reducing simulation time from days to just hours. Behind this type of computational improvement that makes the CFD software to perform, it involves complex calculations and data exchanges among computational systems in a HPC systems. The more complex simulations are being performed, the higher demands from the cluster performance are.

The recent trends in HPC cluster environments, such as the use of multi-core CPUs, hardware accelerators, and low-latency, high-speed Ethernet and InfiniBand cluster interconnect with offloading capabilities are changing the dynamics of clustered-based simulations. Software applications are being reshaped for higher parallelism and multi-

threads, and hardware configuration for solving the new emerging bottlenecks, in order to maintain high scalability and efficiency.

The HPC Advisory Council performed deep investigations on a few popular CFD software to evaluate its performance and scaling capabilities and to explore potential optimizations. The presentation will review the recent developments of high-performance clustering architectures, CPUs and interconnect solutions and how they can influence the runtime, scalability and performance of CAE simulations. This study will also present some of the optimization techniques and networking communication profiling to further understand some of the CFD software's dependencies on the communication networks and the underlying hardware. A deep analysis and communication profiling of a few well known CAE applications will be described and guidelines for optimized productivity will be explored. The research will review the effects by comparing various hardware using different simulation models.