CFD WORKFLOW OPTIMIZATION ON LARGE SUPERCOMPUTERS

ROOH KHURRAM

Supercomputing Laboratory
King Abdullah University of Science and Technology (KAUST), Core Labs
Thuwal, 23955-6900, Saudi Arabia

e-mail: Rooh.Khurram@kaust.edu.edu
web page: https://www.researchgate.net/profile/Rooh_Khurram

Key words: High Performance Computing (HPC), High Performance Data Analytics (HPDA), Big Data, CFD applications.

Summary. Whether we identify some of our routine CFD simulations as big data or not, we are experiencing mixed HPC plus HPDA workflows on supercomputers. We will present case studies and share our experiences with CFD users of large supercomputing facilities.

1 INTRODUCTION

Big data is upon us. Whether we identify some of our routine CFD simulation as big data or not, we are experiencing mixed HPC + HPDA workflows on our flagship supercomputer, Shaheen II. In this presentation we will present various case studies that fall at the interface of HPC, BigData analytics and machine/deep learning. Traditionally, supercomputers are only used for purely compute intensive part of the full workflow. Rest of the steps are done on the smaller general-purpose clusters or even workstations. As the supercomputers are getting more powerful, the simulations are producing larger data sets that are not easy to copy from one system to another. Today's users need a supercomputer to perform all analysis. However, the supercomputer procuremts are typically done for maximizing the FLOPs for getting higher rank on the TOP500 list. The investment on compute typically comes at the expense of lopsided balance ratios – B/F etc. Local disks and memory are also limited, making it more difficult to optimize the full workflow. The parallel file system is also tuned for IO500 style IO pattern, and metadata servers get choked when non-standard IO is performed. This leaves the users with difficult options, either to move away from the supercomputers or reduce the problem size in terms of spatial and temporal resolutions. The challenge for supercomputer's support personal is to optimize workflows on supercomputers that are designed for LINPACK-type compute kernals. This is the dilemma of optimizing HPC and HPDA, simultaneously.

OpenFOAM based DNS project was one such research campaign that contained large scale turbulent simulation and high-speed big data analysis. OpenFOAM solver starts the simulation on a partitioned mesh that is stored in multiple folders which are equal to the number of partitions. The solution files are written in the folders for every prescribed time steps. The available resources on Shaheen II encourage researchers to weak-scale CFD simulations to tens or even hundreds of thousands of cores. Once the data is stored, the researchers analyze and visualize this *big data*. Consequently, this workflow was identified as "metadata killer" by our system administrators because it puts tremendous stress on metadata server of the lustre filesystem. In order to allow the researcher to continue using Shaheen II for the DNS simulations, we had to redesign the production workflow, with the aim of minimizing the size of the data and the speed of the data transfer. This effort resulted in a *press-of-a-button-ready* procedure for the production runs that are metadata friendly. RAMDISK was used for isolating metadata load on the lustre file system and partitioning of data is done in a way that each compute node's RAM acted like a local disk. The same technique was applied to CONVERGE based workflow as well.

In another workflow, we leveraged burst buffer, a Cray DataWarp technology, and customized a full production workflow to minimize data footprint. There is an increasing pressure on the HPC support scientists to learn standard tools of big data analytics to handle HPC workloads that require large data sets.

It is imperative to know what implications there are for supercomputing procurements, which are emanating from the convergence of HPC and HPDA. Only time will tell if the HPC community will continue its march towards designing and procuring compute intensive machines that rank high in the TOP500 list, or move towards more balanced systems that are more suitable for optimizing real life engineering workflows.

REFERENCES

[1] Geoffrey Fox, Judy Qiu, Shantenu Jha, Saliya Ekanayake, and Supun Kamburugamuve, "Big Data, Simulations and HPC Convergence", *in Workshop on Big Data Benchmarks. Springer, 2015, pp. 3–17.*

[2] Earl Joseph, "IDC Update on How Big Data Is Redefining High Performance Computing", *IEEE BigData, Washington DC, 2014*