# High Performance Computing at Intel Xeon Phi Knights Landing Cluster with OpenFOAM

Zhi Shang, Honggao Liu
High Performance Research Computing, Texas A&M University, TX 77843, USA

**Description:**

In the oil and gas industry, gas, oil, water and sand particles all normally exist in a reservoir. Oil and gas companies need accurate information on the location of oil within porous media (such as rock and sand) to help them determine where to drill and to evaluate existing oil reservoirs. It is also important to know where extensive sand deposits exist within porous media. Therefore there are lots of data processing and physics computing in analyzing the situation of the reservoir. With new methods available to address complex physical phenomena, and advances in powerful computing platforms, the ability to model fundamental flow physics at high resolution becomes both essential and possible. It offers the promise of developing reduced order models for large scale processes that retain the fundamental physics in a rational manner [1]. Fig. 1 shows the geometry of a 3D digital porous media sample.

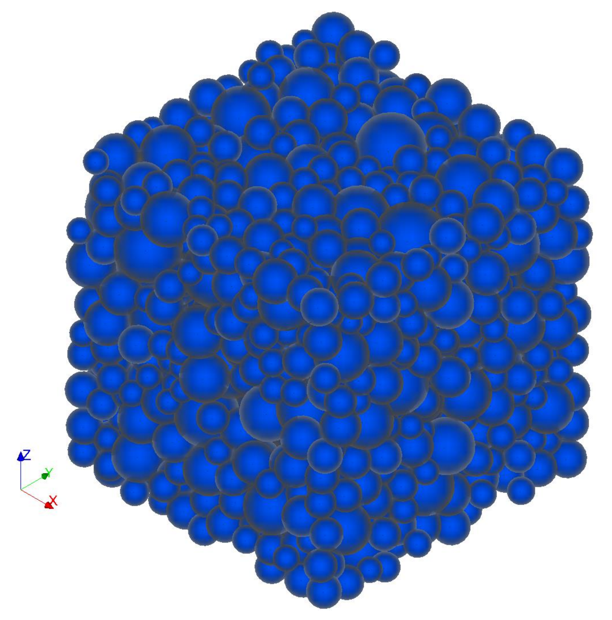

Fig. 1 Structure of micro-pores in oil reservoir

For simulating complex multiphase flows in oil reservoirs, a mathematical model was developed to describe the multiphase flow in the oil reservoir based on OpenFOAM [2]. In the multiphase flow model used in this study, every single fluid phase was presented by Navier-Stokes fluid flow governing equations and solid particles were presented directly by the Lagrangian motion equation using discrete particle model (DPM). Therefore the mathematical model for the multiphase flow is a coupled Eulerian-Lagrangian model. Furthermore, in the multiphase flow model, the VOF (void of fluid) method was introduced to capture the interface between fluid phases. Due to the Eulerian-Lagrangian coupling, the coupling between discrete particle model (DPM) and volume of fluid (VOF) method was also developed for this multiphase flow model. Fig. 2 shows the validating simulation of solid particles passing through the oil layer and water layer with gas bubbles.

The solid particles with gas were injected into the inlet hole at the bottom. First, the gas-particles pass through the water layer in the inner riser, then enter the oil layer, and are finally thrown out off the free surface at the top of oil surface. During the processing of solid particles thrown out off the free surfaces, some of the water and oil droplets may also be carried out into the air. Due to gravity, these heavy particles with liquid droplets fall down into the oil layer and subsequently into the water layer again outside of the inner riser. This processing is a gas-oil-water-particle 4-phase flow which was simulated by the multiphase flow model with coupled DPM and VOF developed for this study.

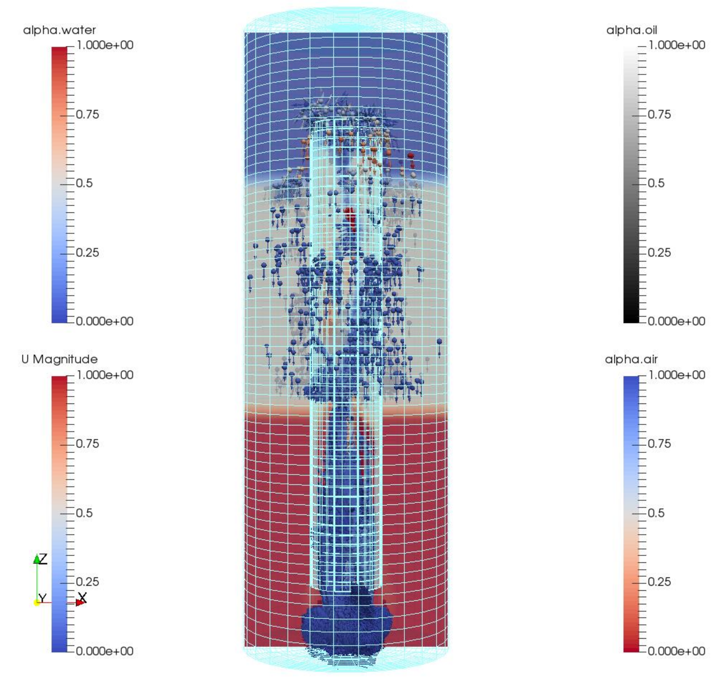

Fig. 2 Multiphase flow simulation by coupled DPM and VOF

In the oil reservoir the flow normally is a multiphase flow including gas, oil, water and solid particles. Therefore the simulation should be physically realistic for gas-oil-water-particle 4-phase flow, i.e. the 4 phases flowing together pass through the porous media with a micro-pore structure. Because the flow process is very complicated and the structure is also complex, the simulation is very time-consuming at large scale. It usually requires a powerful supercomputer to perform these numerical simulations. In this study, the numerical simulations were performed on Intel's Endeavor cluster [3] with Intel Xeon Phi Knight's Landing cards (KNL) at the same time with the optional optimizations for HPC.

Fig. 3 shows the model's computational performance with vectorization optimization (green) and without vectorization optimization (red) for simulating the gas-oil-water-particle 4-phase flow passing through the porous media. From the comparisons shown in Fig. 3, it can be seen that KNL is able to perform and optimize computations for such complex CFD simulations.

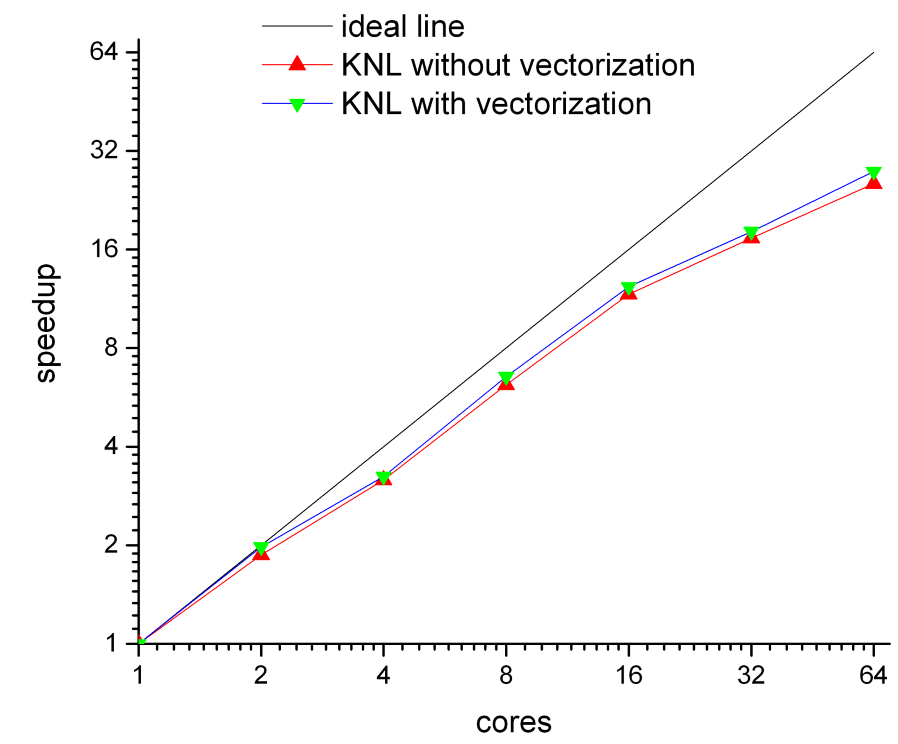


Fig. 3 Speed up of Intel Xeon Phi Knights Landing card (KNL)

**Acknowledgements:**

The authors would like to thank and acknowledge support by the Intel Parallel Computing Center at Texas A&M University, College Station, TX 77843, USA, and Intel. Portions of this research were conducted with high performance computing resources provided by the Intel® Parallel Computing Center.

**References:**

[1] PoreSim Research Consortium, http://www.poresim.org, 2016
[2] OpenFOAM, http://www.openfoam.com, 2016.
[3] Endeavor - Intel Cluster, https://www.top500.org/system/176908, 2016.

**Related Websites:**

IPCC@TAMU (http://hprc.tamu.edu/ipcc/index.php)
Intel PCC (https://software.intel.com/en-us/ipcc)

# Clusters at High Performance Research Computing (HPRC)

## Terra: 307-node Lenovo x86 Cluster

Terra is a 8,512-core Lenovo commodity cluster with

- 304 compute nodes based on Intel's 64-bit 14-core *Broadwell* processors and 64 GB memory.
- Among these nodes, 48 nodes have one *Tesla K80* GPU each and 128 GB memory.
- In addition, there are 3 login nodes (one with GPU), each with 128 GB of memory.
- The interconnecting fabric is based on the *Intel Omni-Path* in a two-level fat-tree topology.
- 3 PB (raw) IBM Spectrum Scale (GPFS) file system via GSS24/26 appliance.
- Slurm as scheduler

### Terra Schematic:

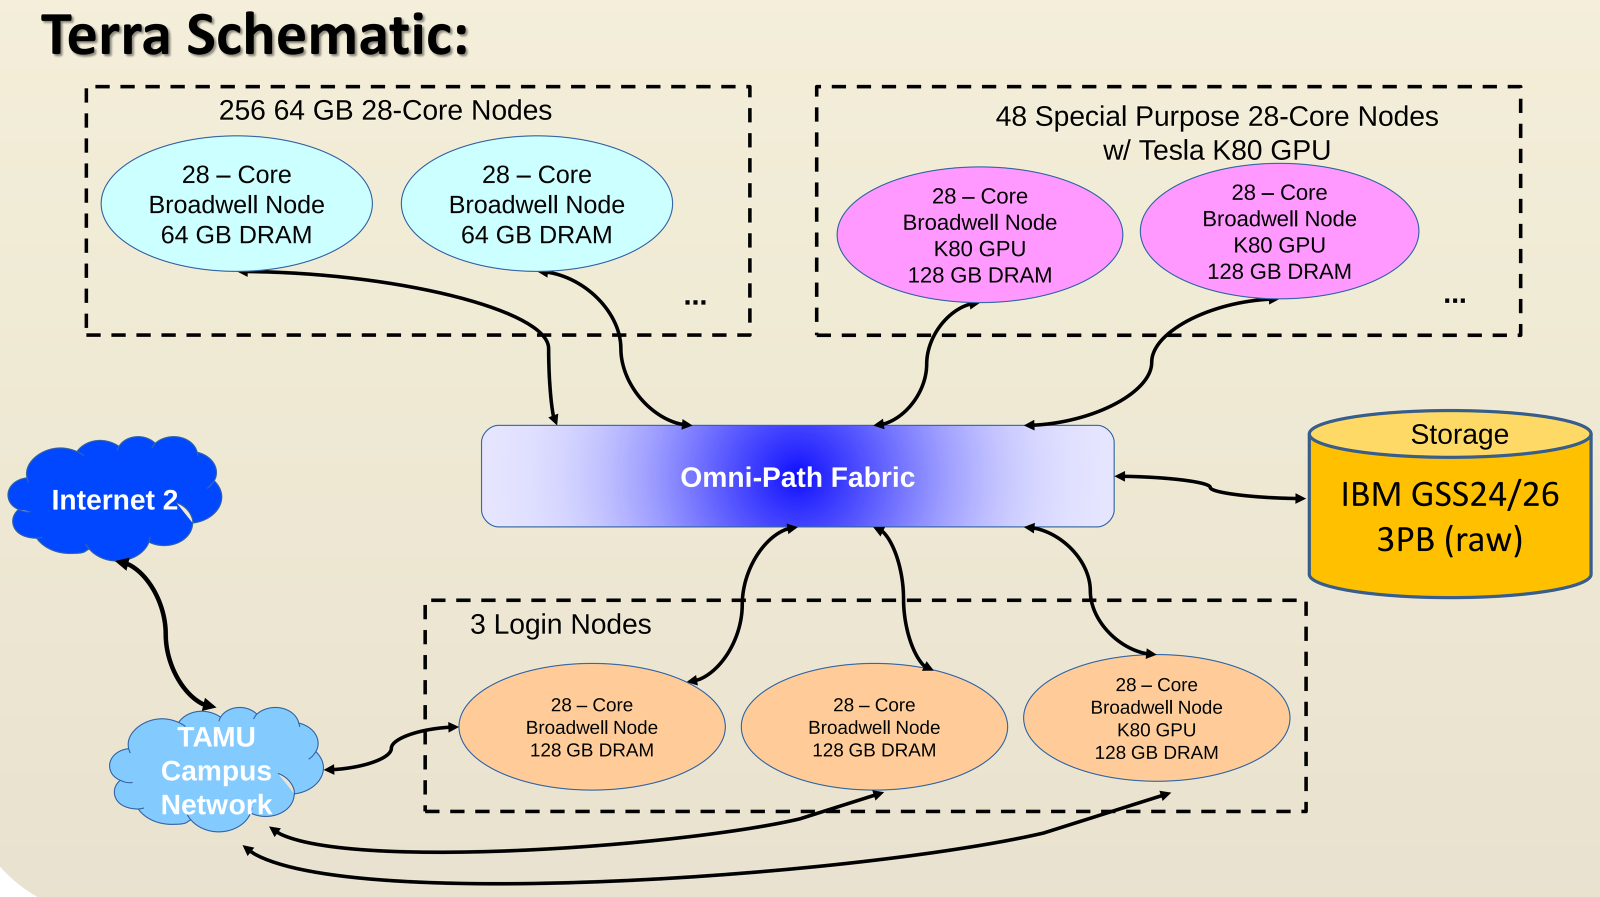


## Ada: 860-node IBM x86 Cluster

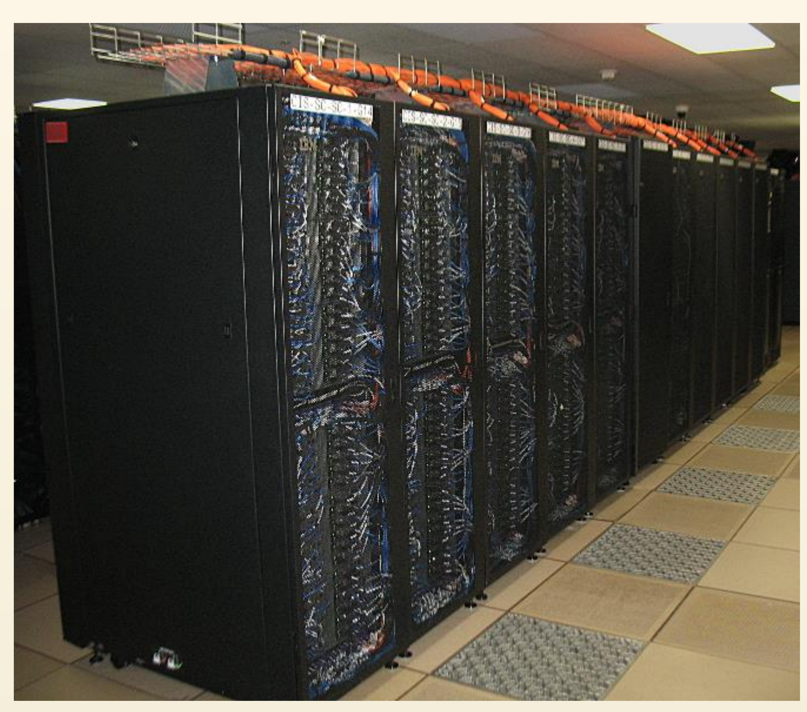

Ada is a 17,500-core, 860-node cluster with:

- 837 20-core compute nodes with two Intel 10-core 2.5GHz *IvyBridge* processors.
- Among these nodes, 30 nodes have 2 GPUs (*K20*) each and 9 nodes have 2 *Phi* coprocessors.
- 15 compute nodes are 1TB and 2TB memory, 4-processor SMPs with the Intel 10-core 2.26GHz *Westmere* processor.
- 8 20-core login nodes with two Intel 10-core 2.5GHz *IvyBridge* processors and 1 GPU, 2 GPUs, or 2 Phi coprocessors
- Nodes are interconnected with *FDR-10 InfiniBand* fabric in a two-level fat-tree topology.
- 4 PB (raw) IBM Spectrum Scale (GPFS) file system via GSS26 appliance.
- Platform LSF as scheduler.
- 2 nodes with 40 Gbps link to Internet2 (Science DMZ)

### Ada Schematic:

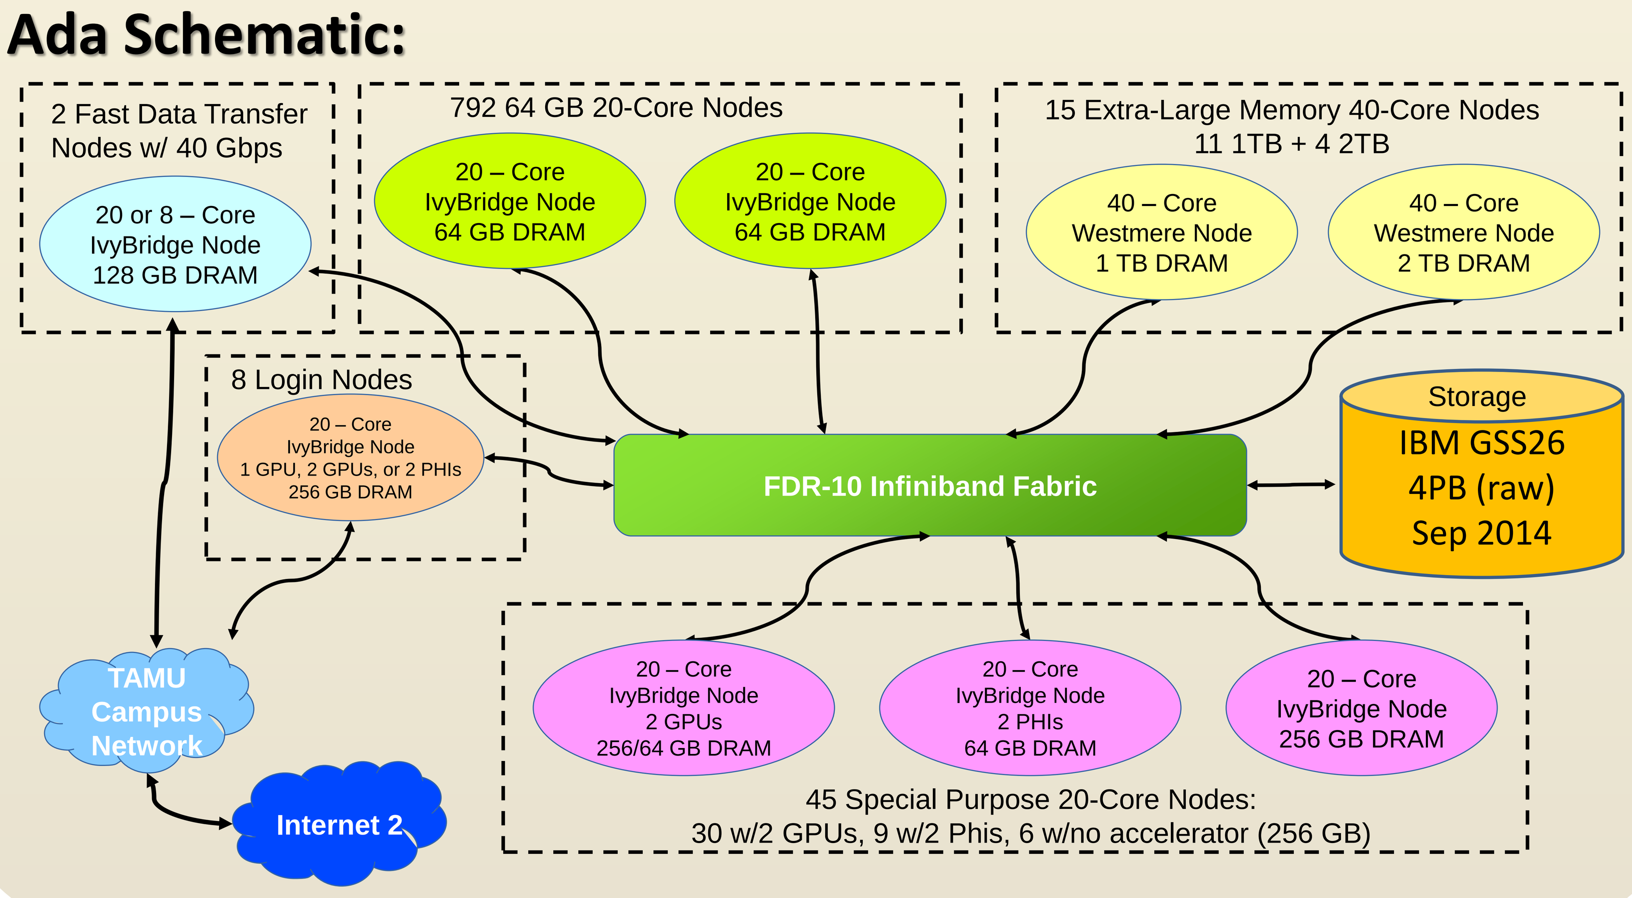


## Curie: 50-node POWER7+ cluster

Picture shown mix of Curie and Crick nodes.

- Curie is an 768-core IBM Power7+ cluster with nodes based on IBM's 64-bit 16-core Power7+ processors.
- In addition to the 48 nodes are 2 login nodes with 256GB of memory per node.
- Curie's file system (GPFS) and batch scheduler (LSF) are shared with Ada cluster.

## Crick: 25-node POWER7+ cluster

Picture shown mix of Curie and Crick nodes.

- Crick is a 368-core IBM Power7+ BigData cluster with nodes based on IBM's 64-bit 16-core Power7+ processors.
- Included in the 23 nodes are 1 BigSQL node with 256GB of memory per node and 14TB (raw) of storage and 22 data nodes with 14TB (raw) storage for GPFS-FPO and local caching.
- In addition to these nodes are 2 login nodes with 128GB of memory per node,
- Crick is primarily used for big data analytics (IBM InfoSphere BigInsights, Data Explorer Resource; Query Routing for InfoSphere Data Explorer).

http://hprc.tamu.edu