Development and application of a Peta-scale GPU cluster for multi-scale discrete simulation — *Mole-8.5*

Xiaowei Wang, Wei Ge, Xianfeng He, Feiguo Chen, Li Guo, Jinghai Li
Institute of Process Engineering, Chinese Academy of Sciences, Beijing 100190

Introduction

Mole-8.5 is the first GPGPU supercomputer (Rpeak of about 1100 Tflops) using NVIDIA Tesla C2050 in the world, designed and established in April 2010 by Institute of Process Engineering (IPE), Chinese Academy of Sciences, one of the NVIDIA CCOEs. It holds the No. 19 spot on the June 2010 TOP500 list of worldwide supercomputers, ranked No. 3 in China. In the latest release of Green500 list, Mole-8.5 debuts at No. 8 for the most energy-efficient supercomputers in the world.

Development of Mole-8.5

Mole-8.5 consists of 372 nodes, including displaying and computing nodes. It is the successor of the first supercomputer with 1.0 Petaflops peak performance in single precision in China, which was a hybrid system including four units integrating NVIDIA and AMD GPUs announced on April 20, 2009. A designing philosophy utilizing the similarity between hardware, software and the problems to be solved is embodied, based on the multi-scale method and discrete simulation approaches developed at IPE. The whole system is connected with Gigabit Ethernet and QDR Infiniband network. *Mole-8.5* has some unique advantages over the HPC system of same performance based on CPU, for example the high performance/price ratio, the area occupied by it is only about 150 M^2. The linpack result of 320 nodes of Mole-8.5 is 2.073e+05 Gflops with a power consumption of about 480 KWatt, therefore the average power efficiency is 431 Mflop/s/Watt, manifesting an energy efficient supercomputer.

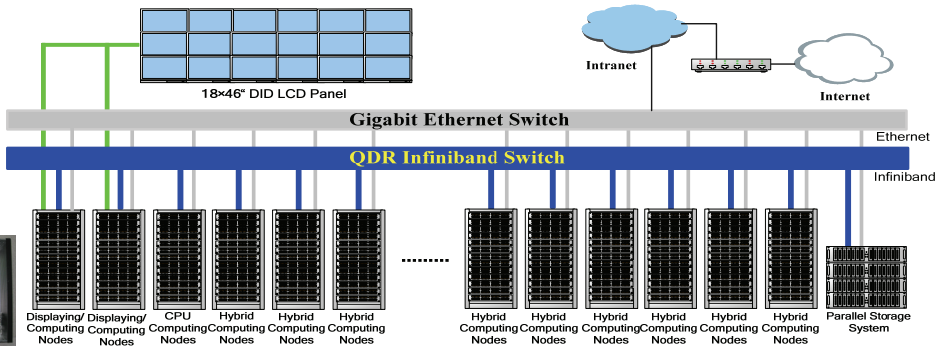

Application of Mole-8.5

With the multi-scale discrete software developed by IPE, the *Mole-8.5* system has already carried out large-scale simulations of high scientific significance covering areas such as chemical engineering, material science, biochemistry, oil exploitation and recovery, metallurgy. The relative efficiency of the system for many of these practical applications is more favorable than of other CPU and GPU-CPU hybrid systems, justifying its GPU-dominated architecture for multi-scale discrete simulation.

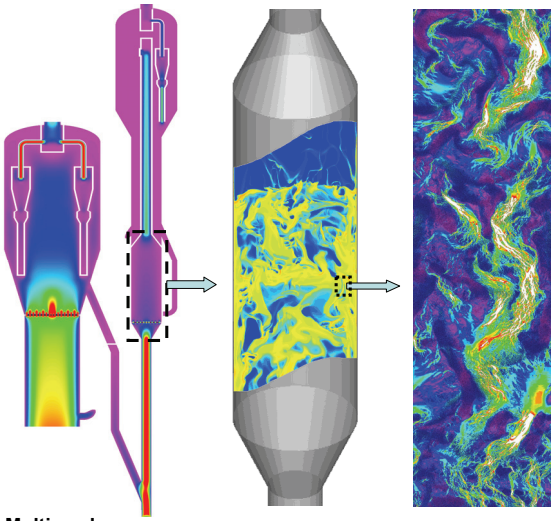

Multi-scale simulation of gas-solid flow

2D simulation of a section (~1.2M cells, 60GPUs)

DNS simulation in a cell (100K~1M solids, 1G grids, 600GPUs)

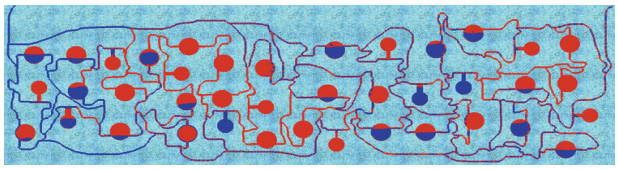

Multi-scale simulation of water flooding in fracture-cave type reservoir (area: 500×150 m, 40GPUs)

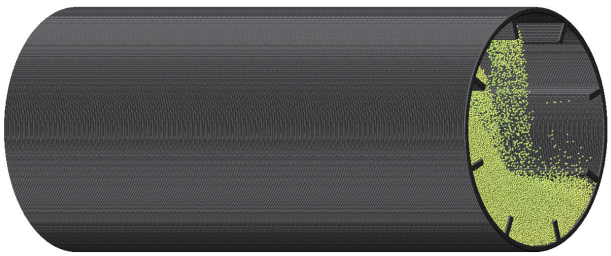

Quasi-real time simulation of a industrial rotating drum (~9.6M particles, 1/13 physical speed, 270GPUs)

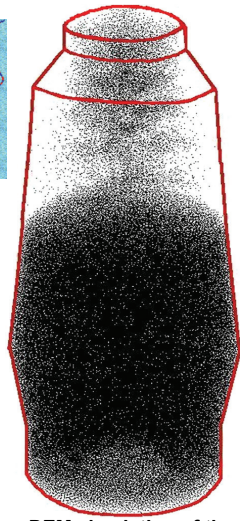

DEM simulation of the bedding of feed-stocks in a metallurgical furnace

Conclusion and perspective

- The development and application of Mole-8.5 demonstrate the supercomputer as a paradigm of green computation with innovative and effective architecture for multi-scale discrete simulation.
- The consistency among the system to be simulated , the physical model, the numerical algorithm and the computing hardware is critical to high-efficiency and low-cost supercomputing.