



Multi-Scale Simulation of Multi-Phase Systems Using GPUs

Wei Ge
Institute of Process Engineering
Chinese Academy of Sciences

SC2010, Nov. 13-19, New Orleans

Process engineering and supercomputing

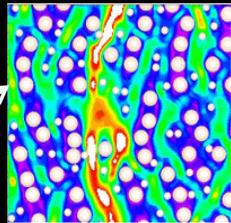
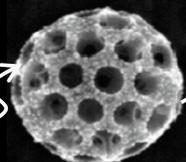
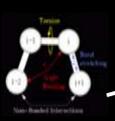
New paradigm in HPC - **CPU+GPU**
best for multi-scale discrete simulation

Simulation challenges:
speed, accuracy, efficiency

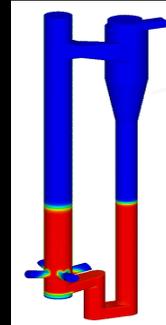
huge scale gap

Products

$10^{-10} \sim 10^{-6} \text{m}$, $10^{-15} \sim 10^{-9} \text{s}$



molecule cluster particle emulsion



reactor



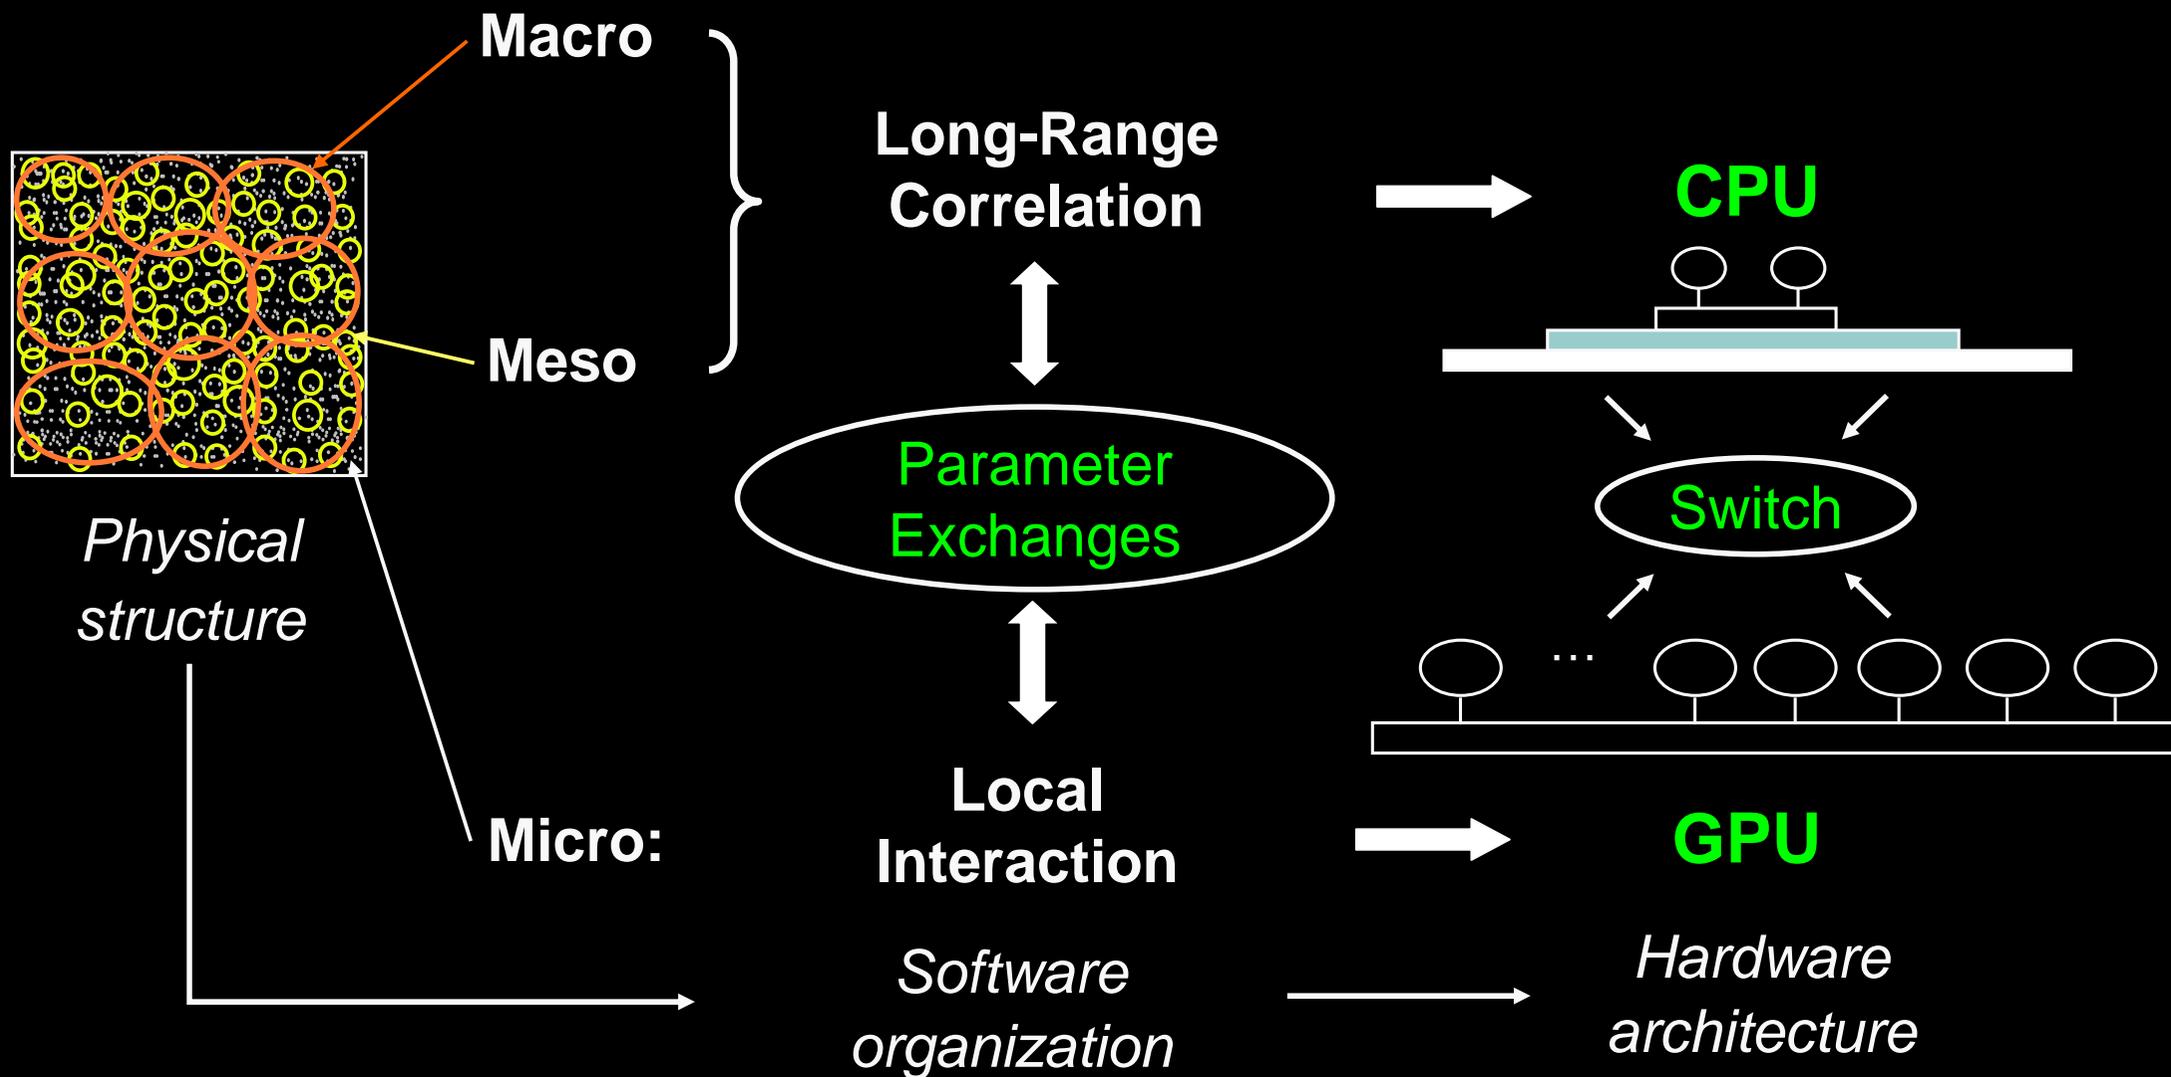
process



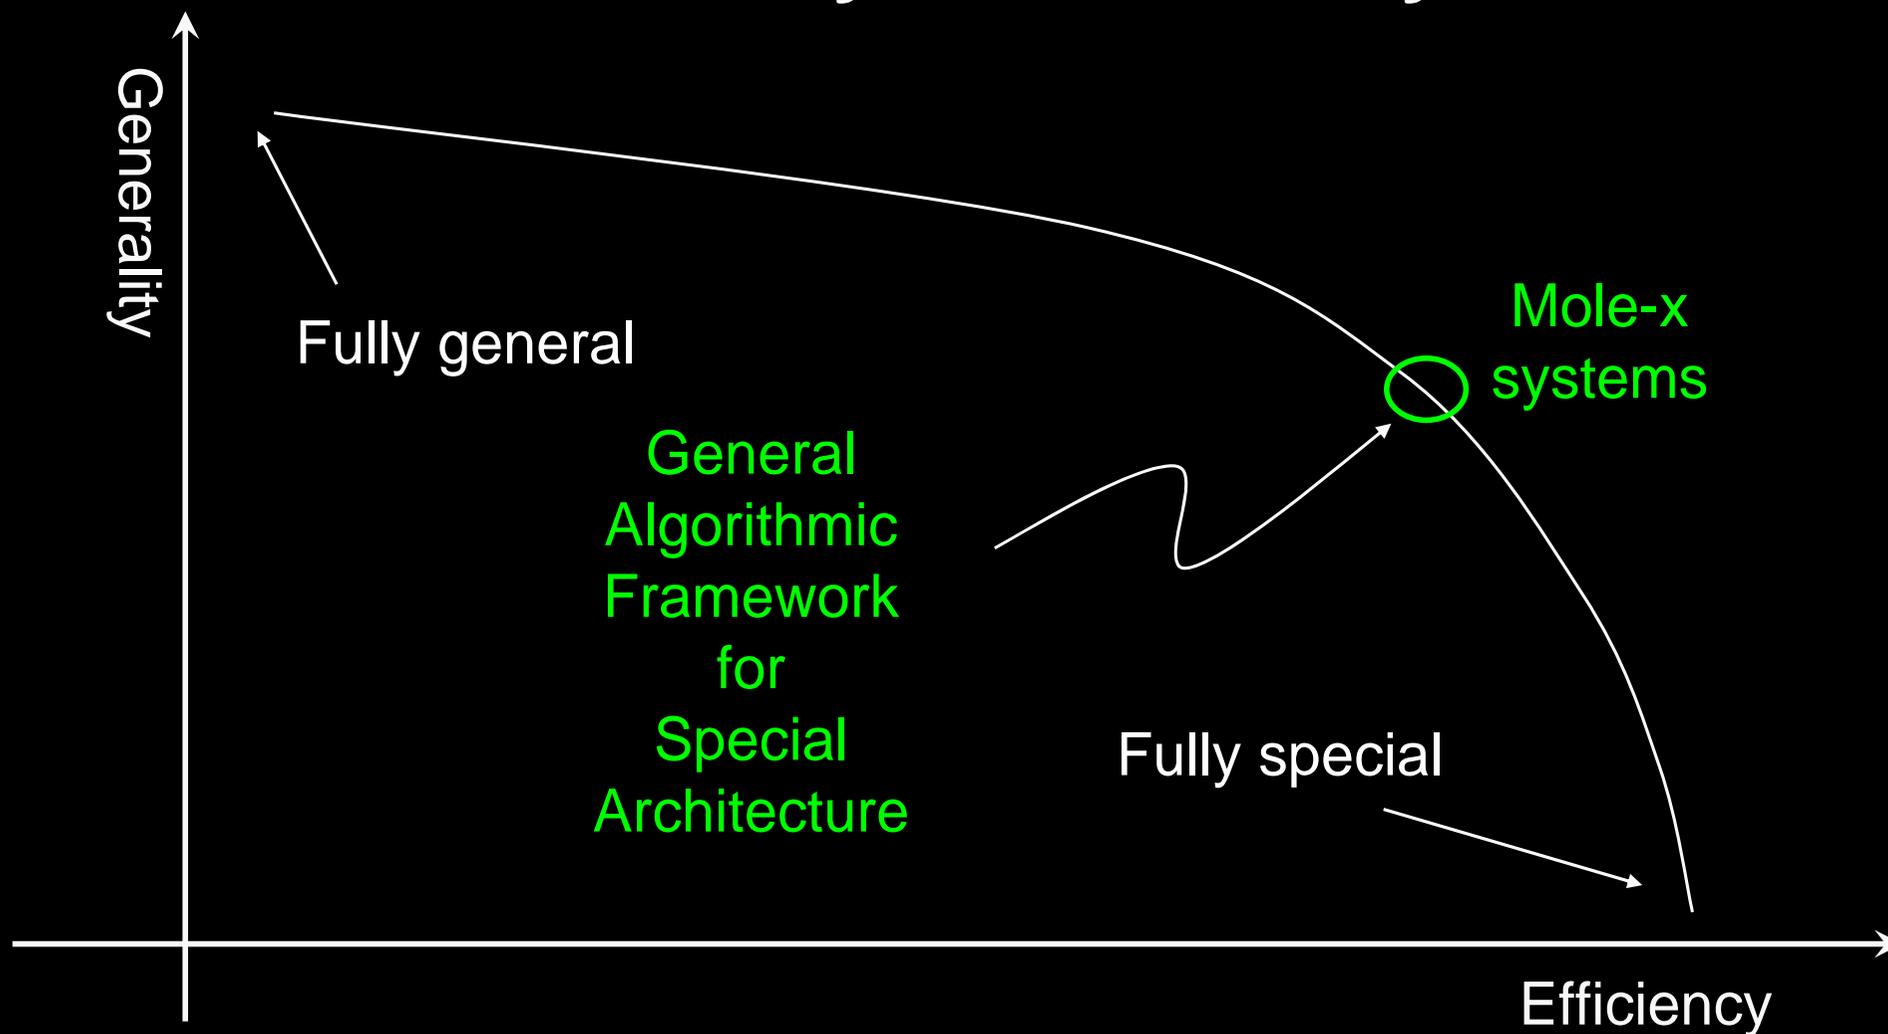
environment

Products
 $1 \sim 10^3 \text{m}$, $1 \sim 10^6 \text{s}$

Consistency: Physics, Algorithm, Architecture



Generality vs Efficiency



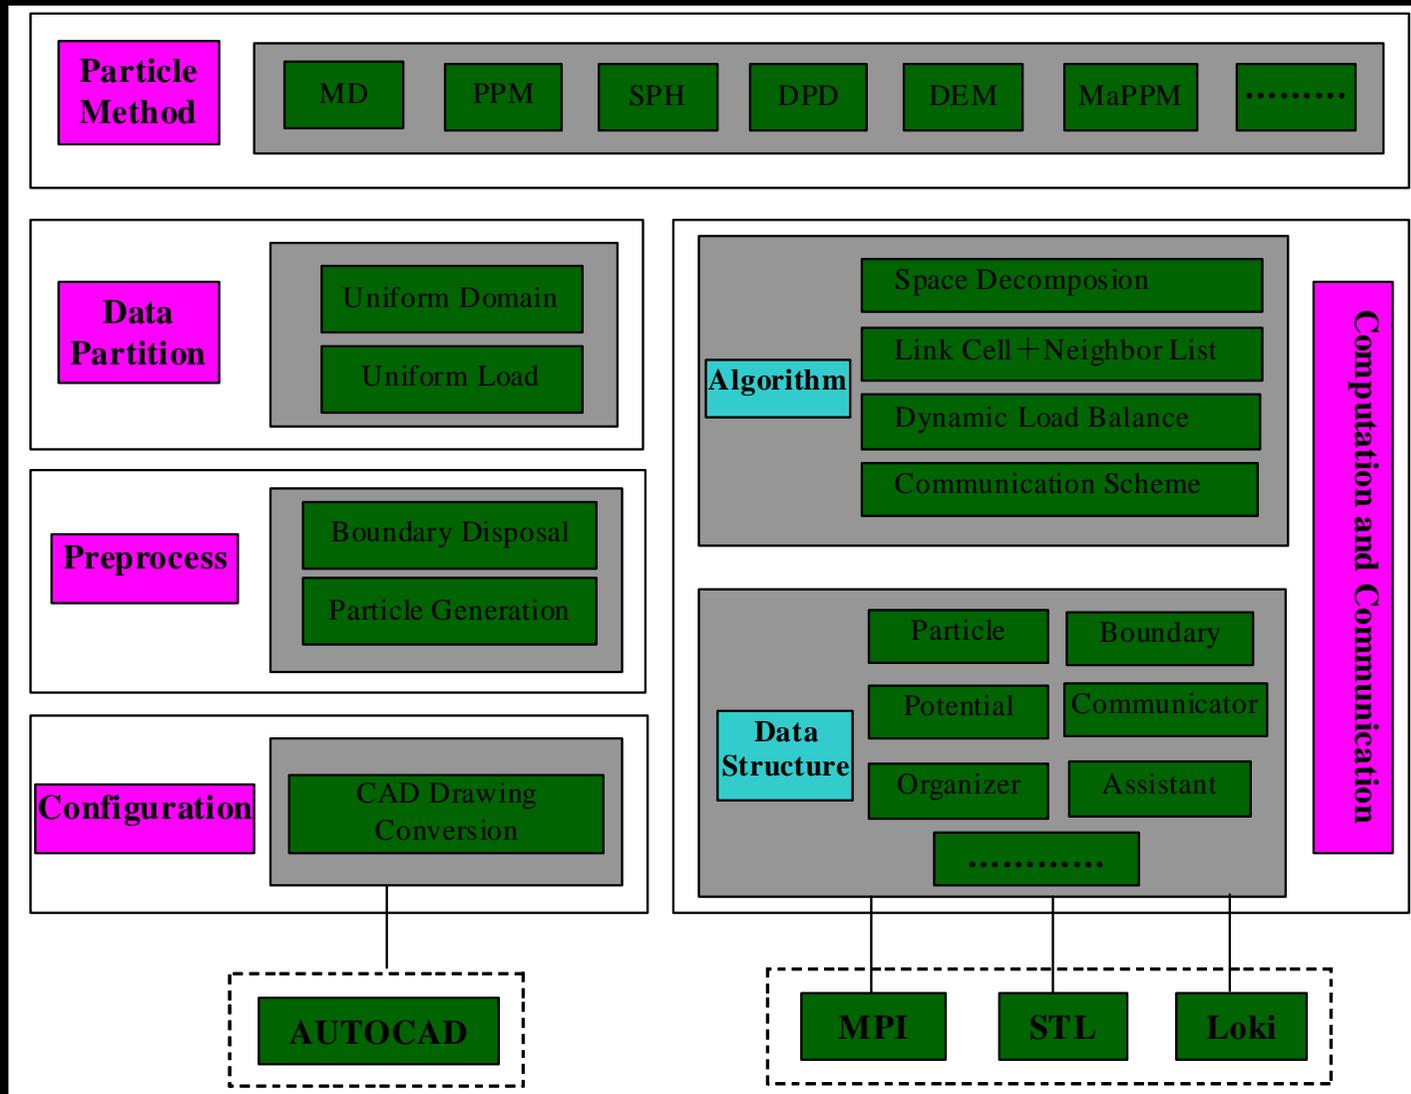
Multi-scale discrete approaches

Micro-scale: fluctuating, conservative
MD, DSMC, LGA, **PPM**, ...

Meso-scale: fluctuating, dissipative
DPD, FPM, DSPH, LBM, ...

Macro-scale: smooth, dissipative
SPH, MPS, DEM, **MaPPM**, ...

General platform for discrete simulation



*Ge et al.,
Chin. Sci. Bull.,
47:1172, 2002;*

*Tang, 2005
Doctor Thesis*

Example: Direct simulation of gas-solid flow

Particle

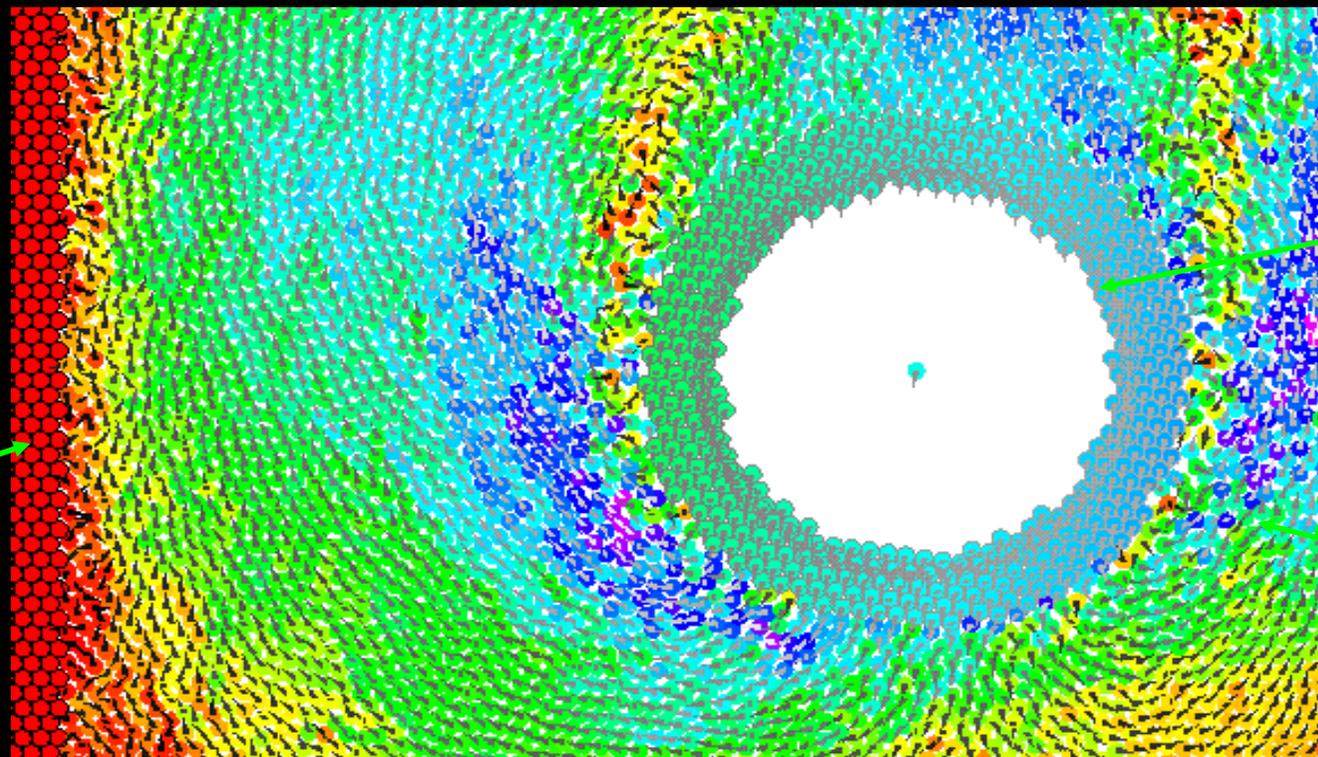
Continuum

Velocity difference \rightarrow tangential stress

viscosity

Density difference \rightarrow normal stress

pressure

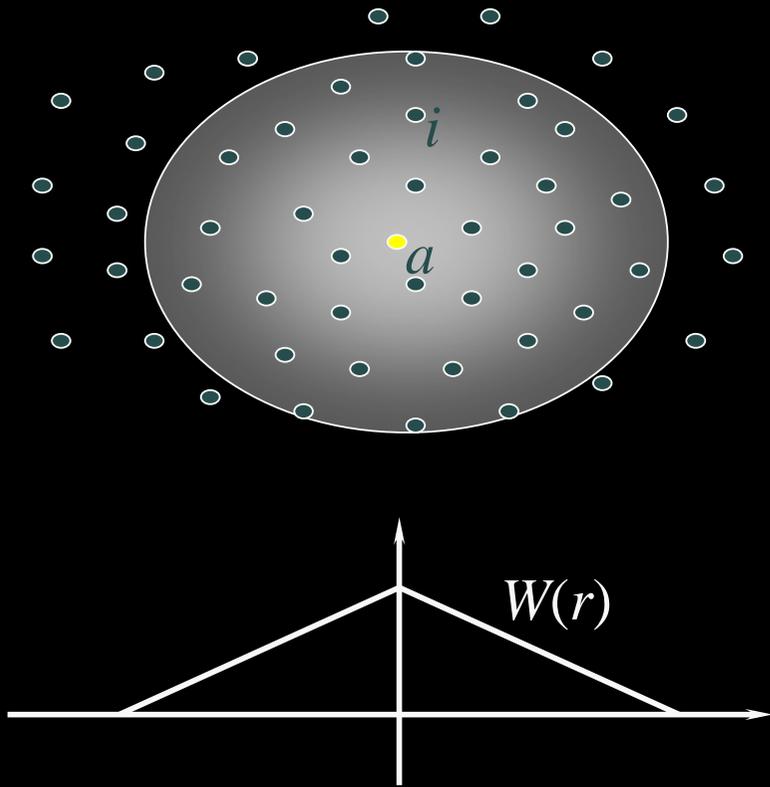


Fixed particle
(boundary)

Bundled particle
(solids)

Free particle
(fluid)

A straightforward formulation



$$\nabla f \Big|_a = D \sum_i \frac{f_{ia}}{r_{ai}^2} \mathbf{r}_{ai} \frac{m_i}{\rho_i} W_{ai}$$

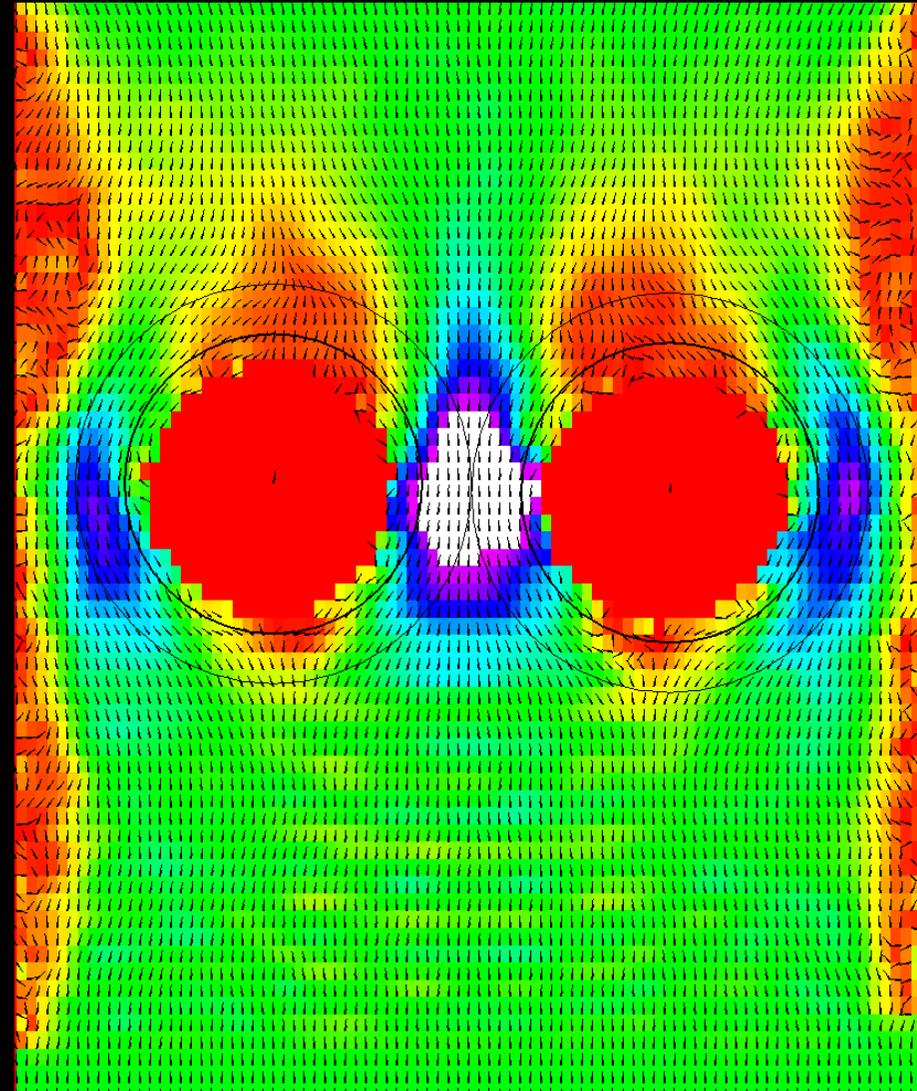
$$\Delta f \Big|_a = 2D \sum_i \frac{f_{ia}}{r_{ai}^2} \frac{m_i}{\rho_i} W_{ai}$$

2000



SGI Indigo2
800Mflops

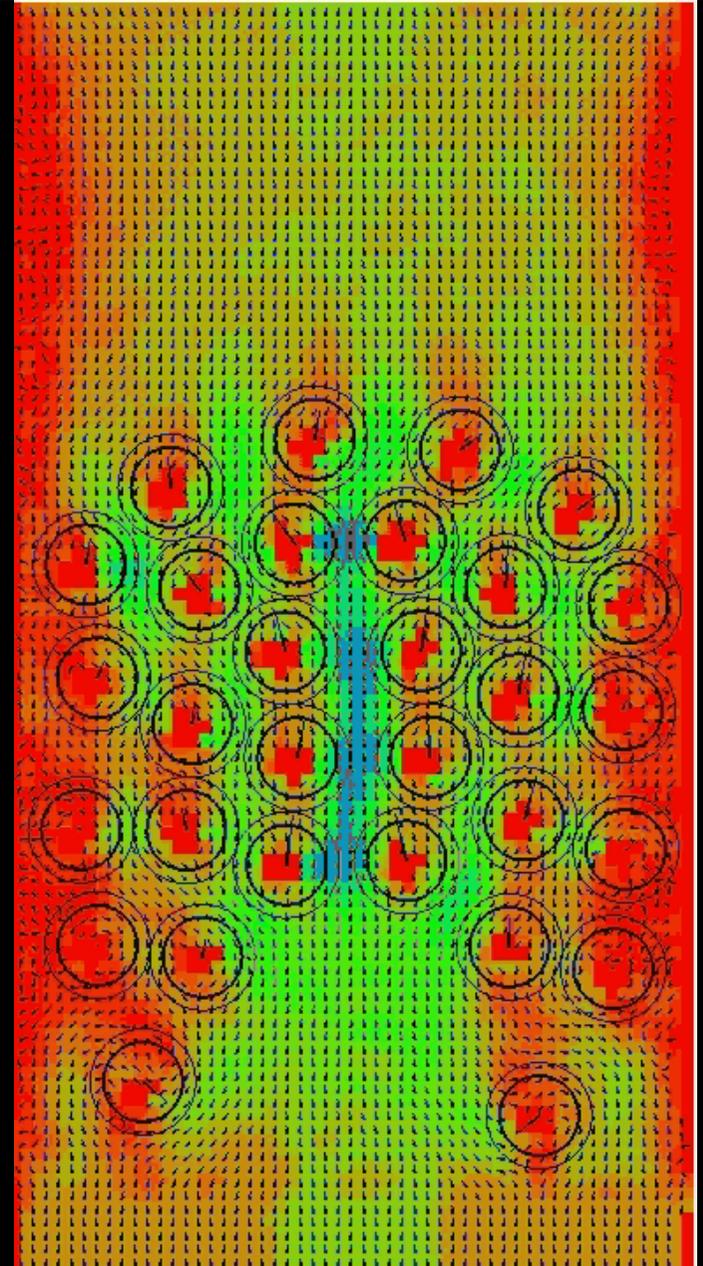
A fall of
two particles
20K fluid particles



2003

Fluidization of
30 particles
100k fluid particles

Compaq DL360 Cluster
100Gflops Peak
(~30th, Top100 China, 2002)



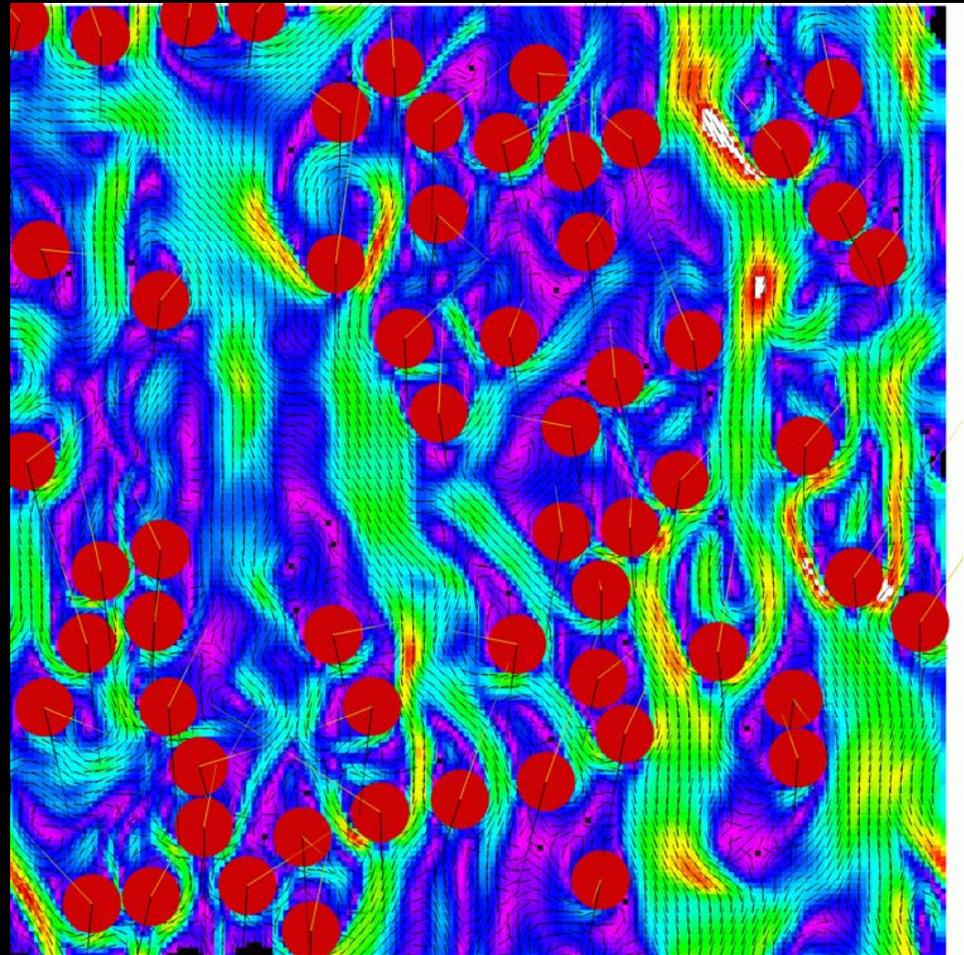
2005

DeepComp 6800
5.3Tflops (14th, Top500, 2003.11)

lenovo



1024 solid particles
5M fluid particles



Ma et al., Chem. Eng. Sci., 61:6878, 2006

2008



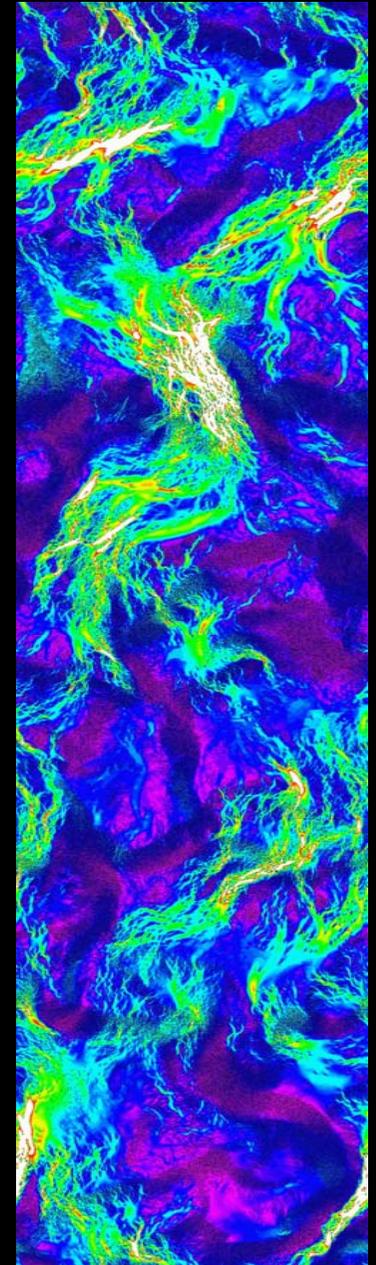
Host: 126 × HP8600
CPU: 252 × Intel E5430 2.66GHz
GPU: 200 × NV Tesla C870
+20 × NV GeForce 9800 × 2
Net.: Gigabit Ethernet (mesh+tree)
Sw.: H3C 7506R
OS: RedHat Linux 5.2
Rreal: 20 ~ 40 Tflops (SP)

Mole-9.7
127 Tflops (SP)

6.023e(23-9.7) flops

118K solid particles
400M fluid particles

Xiong et al., Chem. Eng. Sci., 65:5356, 2010



2010

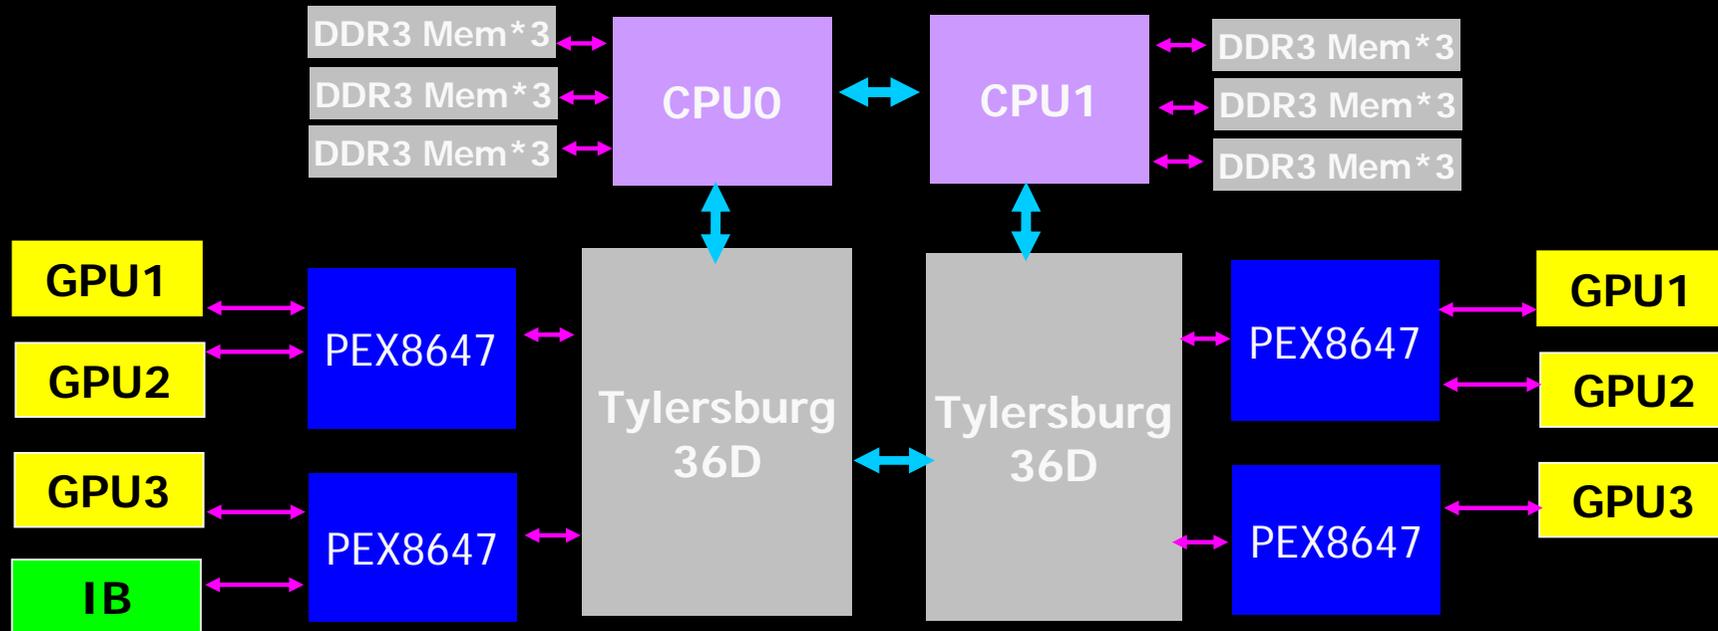
Mole-8.5

(born on April 24)

Rpeak SP: 2Petaflops
Rpeak DP: 1Petaflops
Linpack: 207.3Tflops (19th, Top500)
Mflops/Watt: 431 (8th, Green500)
Memory: 17.2TB(RAM), 6.6TB(VRAM)
Storage: 76TB(Nastron)+320TB(HP)
Data Comm.: Mellanox QDR InfiniBand
Inst. Comm.: H3C Gigabit Ethernet
Occupied area: 150m² (with internal cooling)
Weight: 12.6T (with internal cooling)
Max Power: 600kW+200kW(cooling)
System: CentOS 5.4, PBS
Monitor: Ganglia, GPU monitor
Languages: C, C++, CUDA, OpenCL



Node layout of Mole-8.5



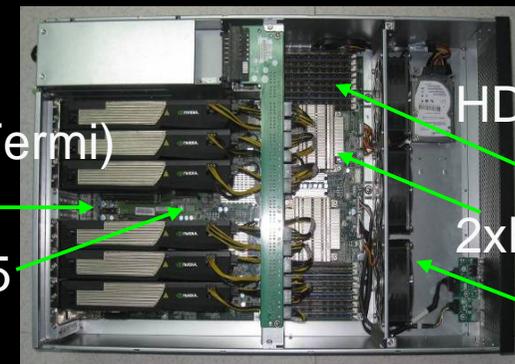
Bottleneck:

DeMem → PCIE → IB

6xC2050 (Fermi)

QDR IB

Tyan S7015



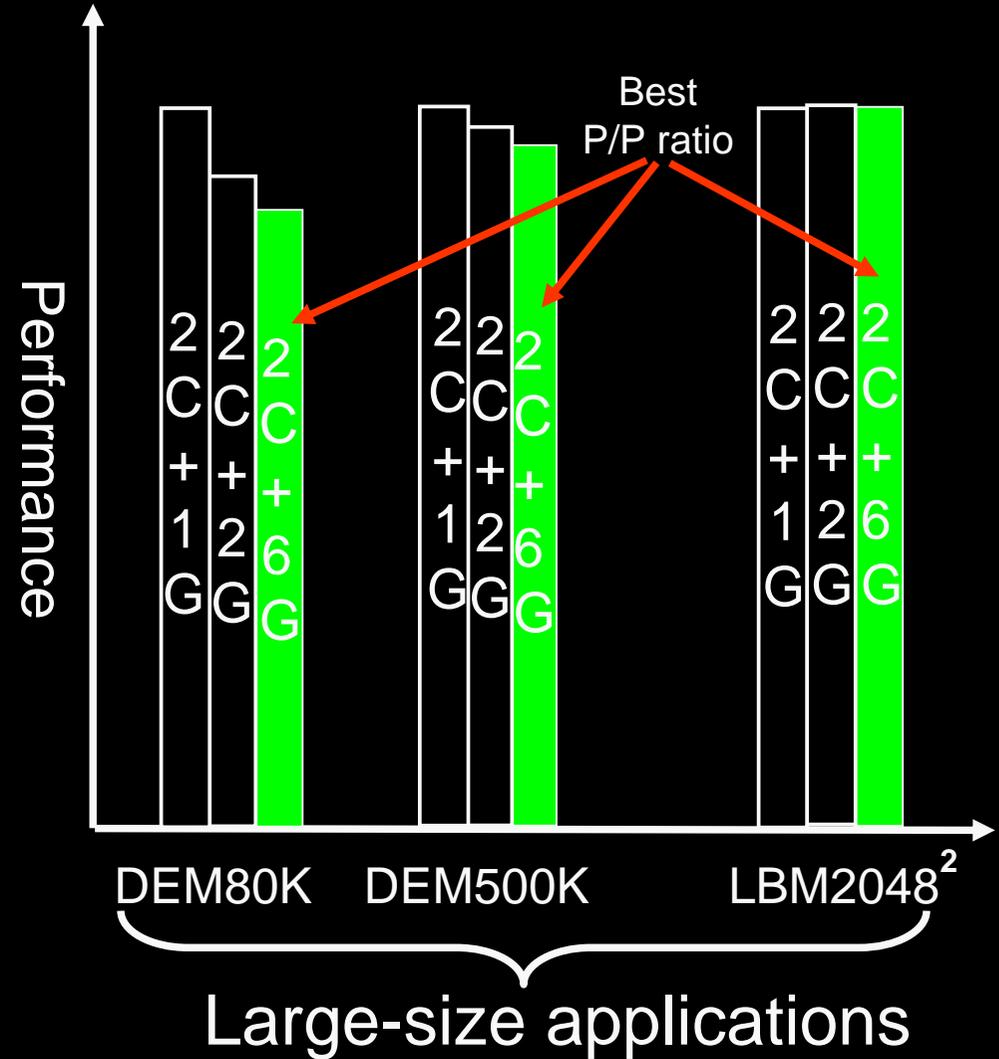
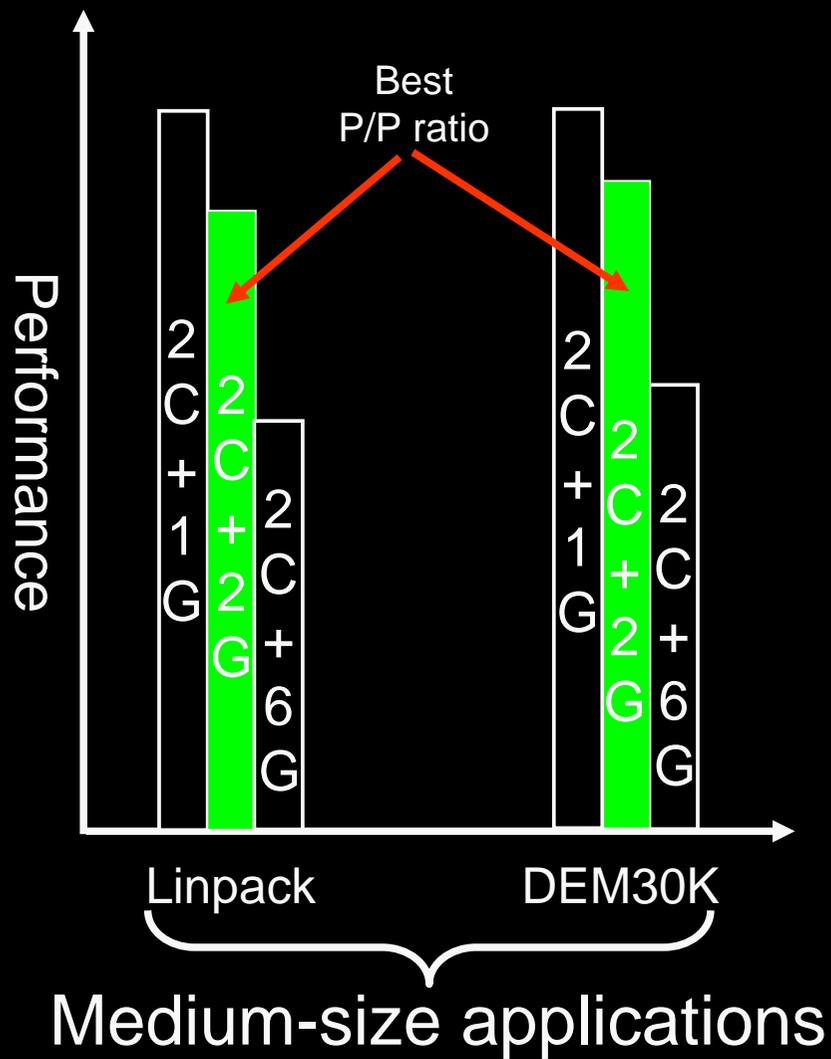
HD

Mem

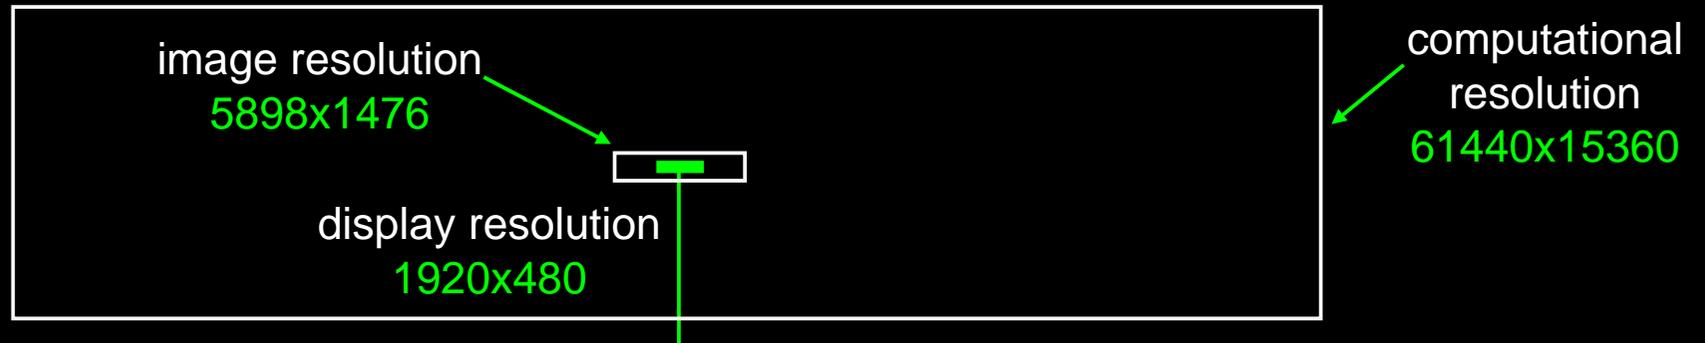
2xE5520/70

Fan

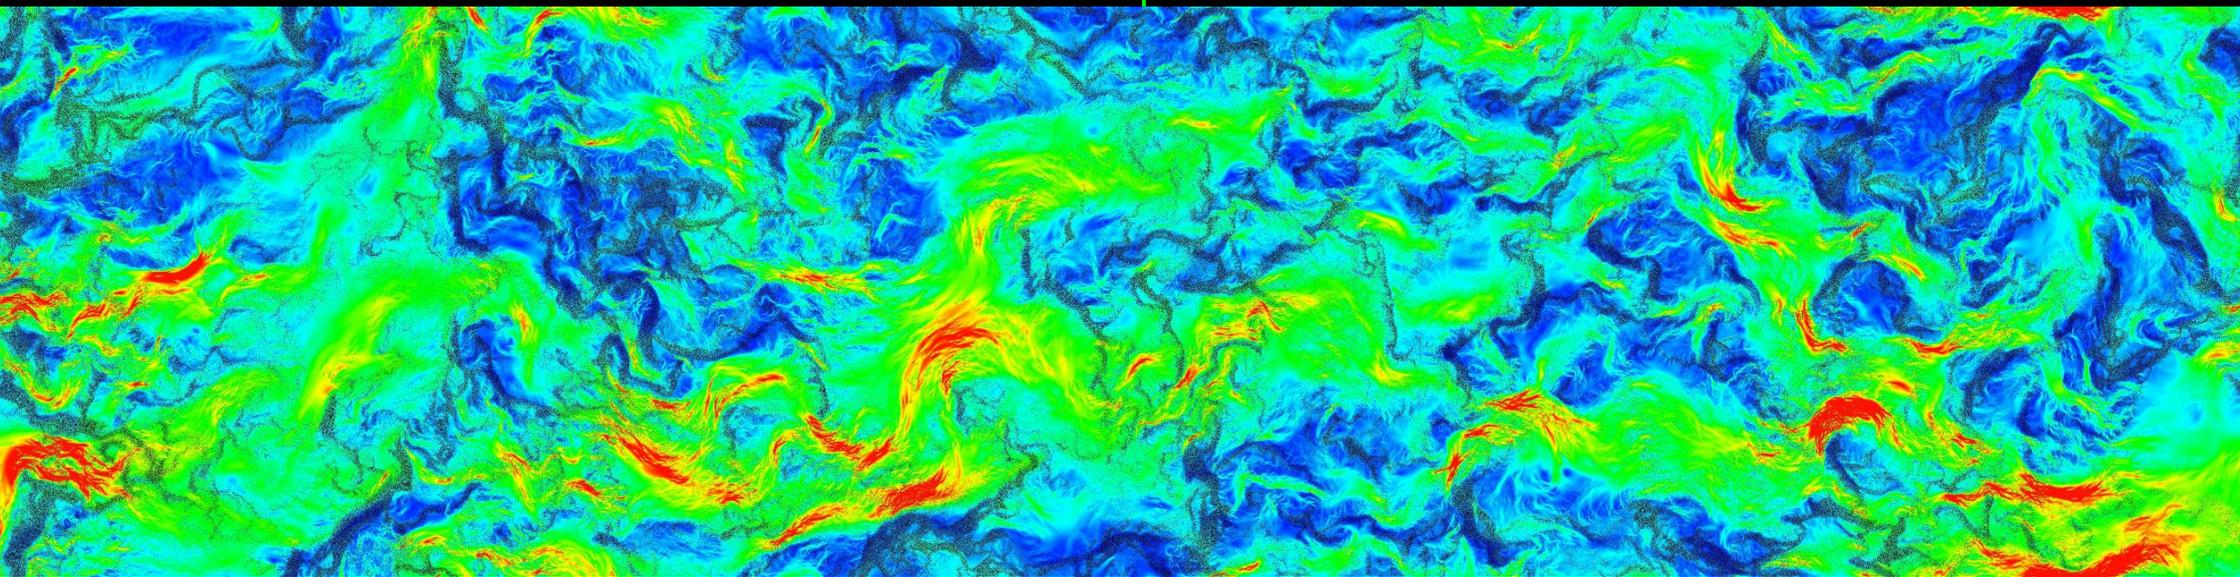
Effect of the CPU-GPU ratio on performance



1M solids & 1G fluid Ps @ 576 GPUs

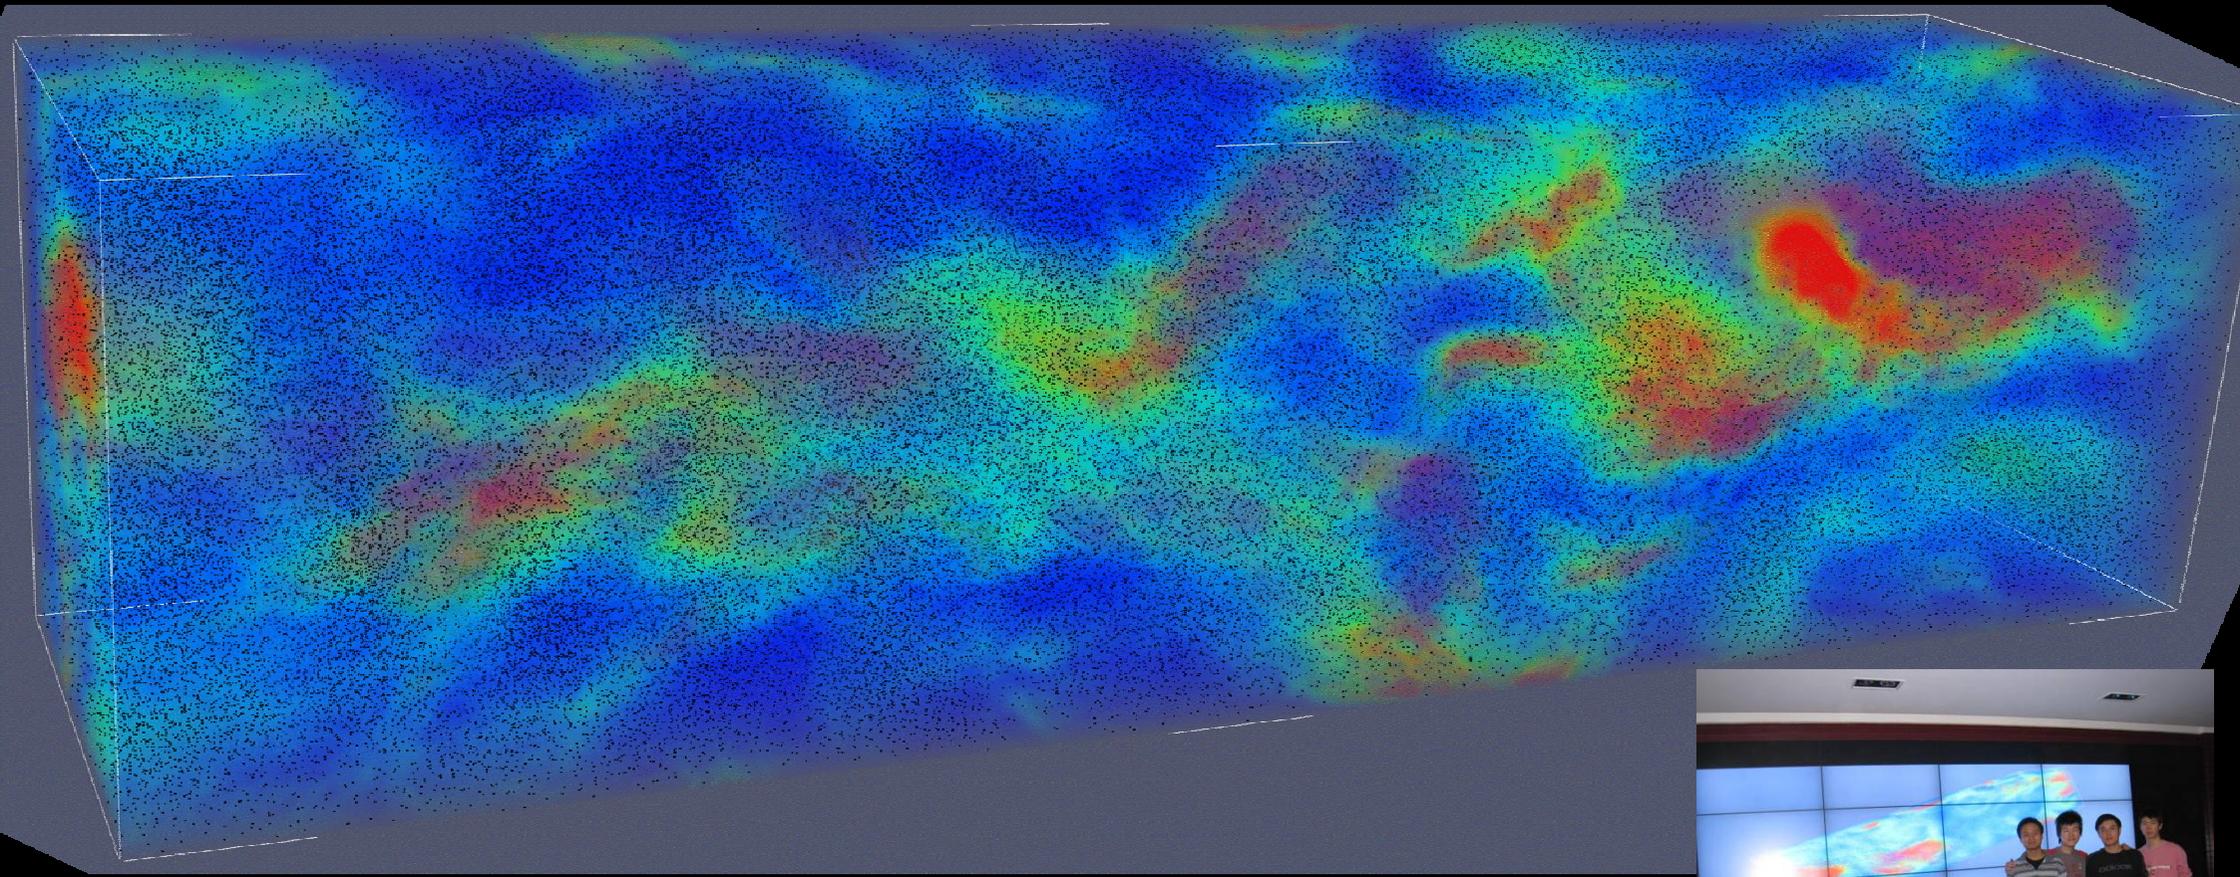


Xiong et al., 2010



Can you paint with all the *colors of the wind*?

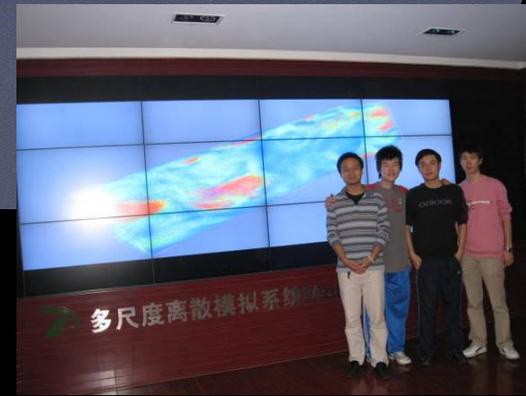
100K solids & 1G fluid Ps in 3D @ >1K GPUs



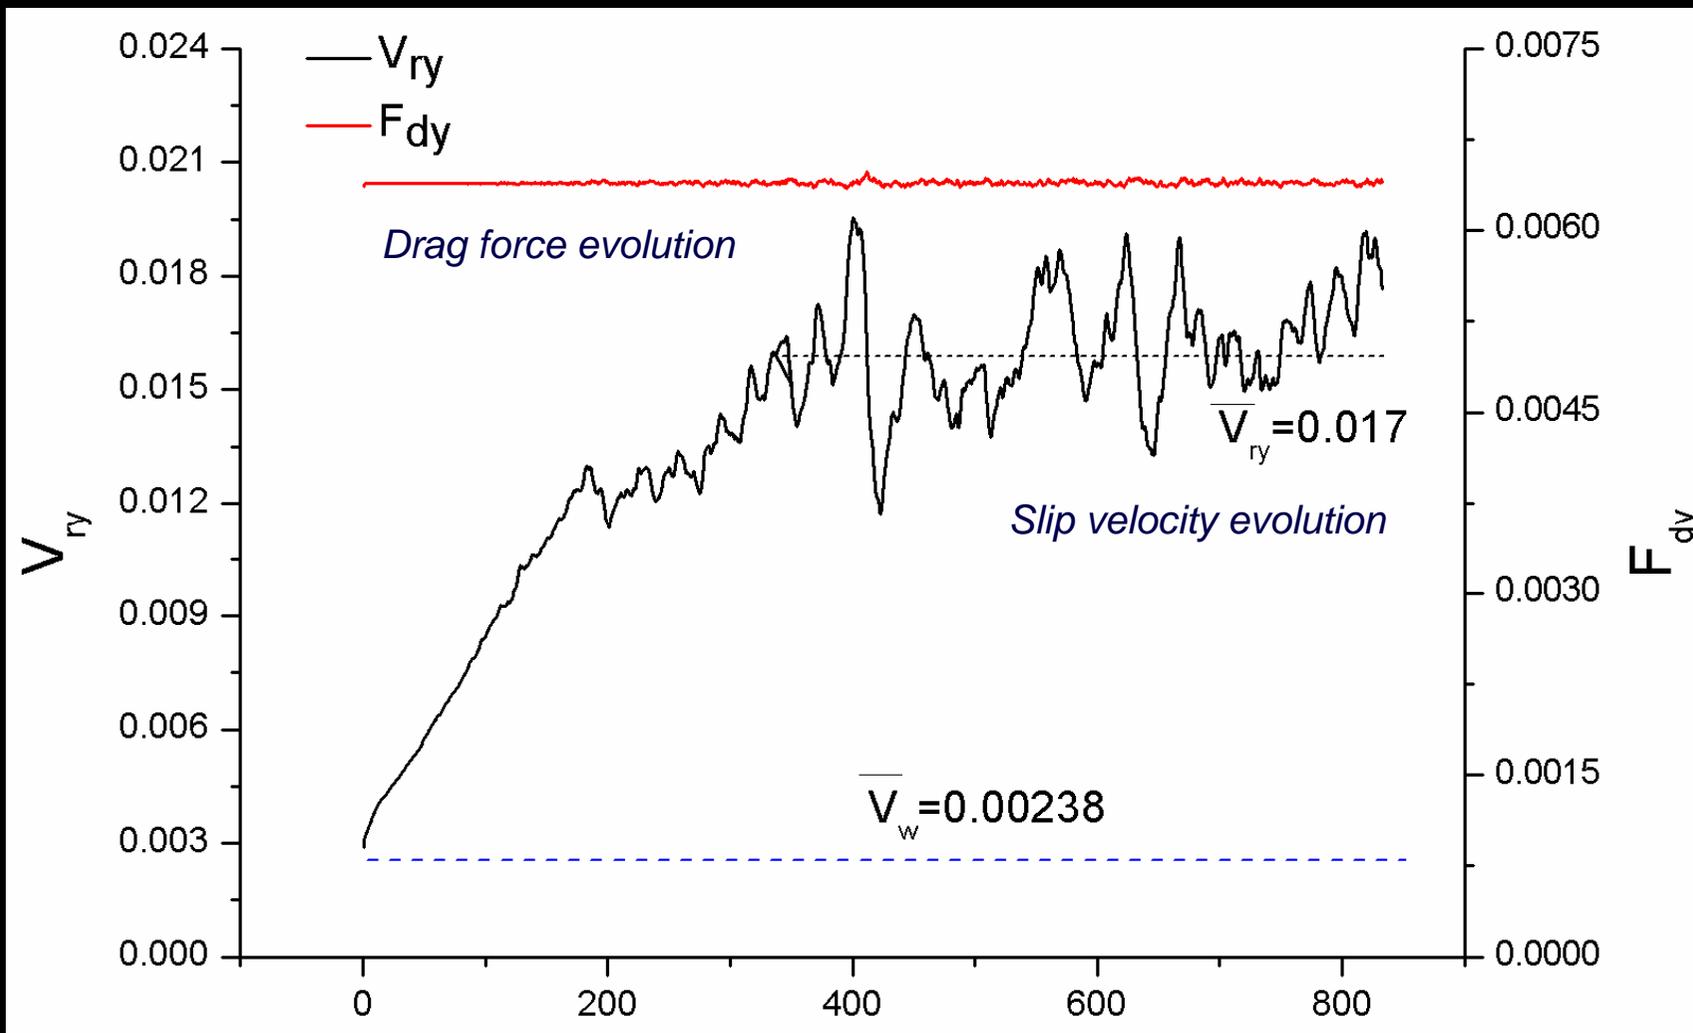
Xiong et al., 2010

ICT-IPE visualization team against the display wall

Sept. 18, 2010, photo by Xiaowei Wang



New constitutive laws for continuum models



Ma et al.,
Chem. Eng. Sci.,
61:6878, 2006

Xiong et al.,
2010

Distributed GPU-supercomputing in China

Collective capacity:
4.907Petaflops SP
1.300Petaflops DP

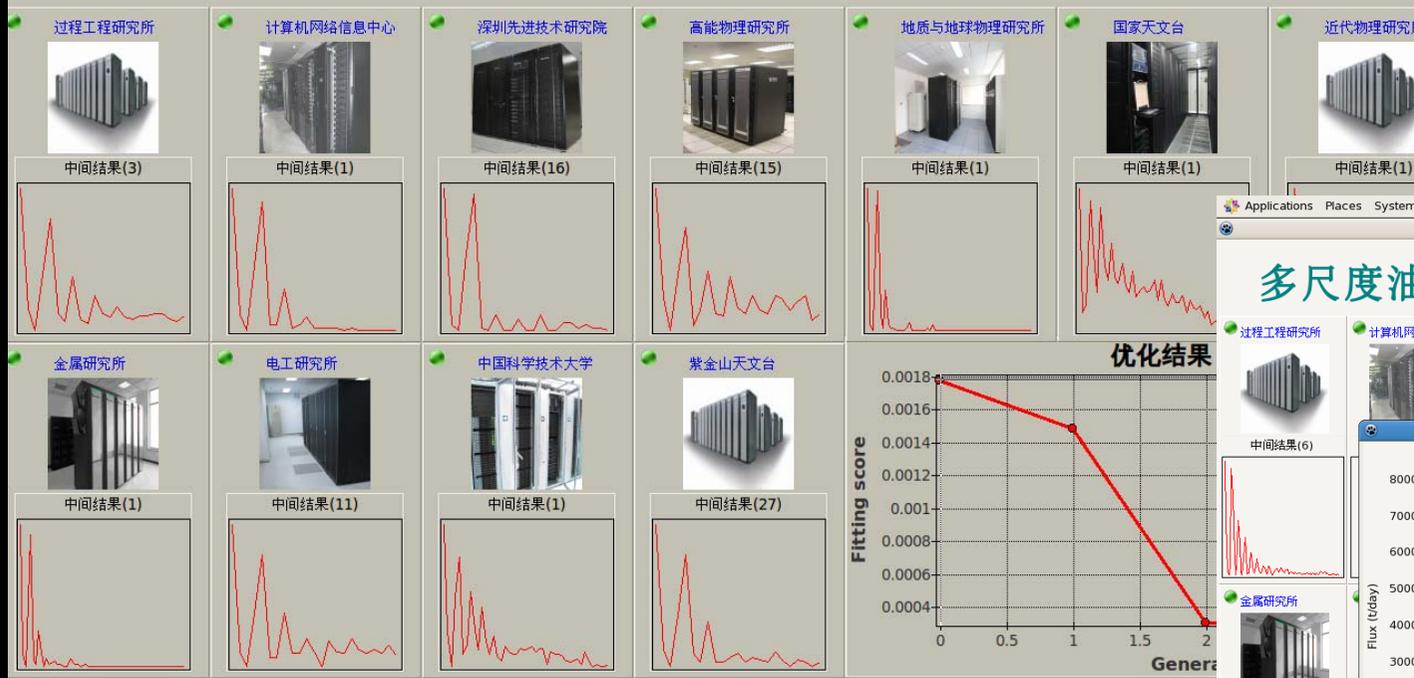
IPE Mole-8.5
3Petaflops SP, 1Petaflops DP



过程所 NV+AMD	高能所 AMD	金属所 NV	深圳先进院 NV	国家天文台 NV	地质地球所 NV+AMD	网络中心 NV+AMD	近代物理所 NV	紫金山台 NV	电工所 NV	中科大 NV
1P SP	200T SP	183T SP	173T SP	160T SP	200T SP	300T SP	202T SP	183T SP	101T SP	205T SP
100T DP	40T DP	15T DP	14T DP	13T DP	17T DP	39T DP	17T DP	15T DP	8T DP	17T DP

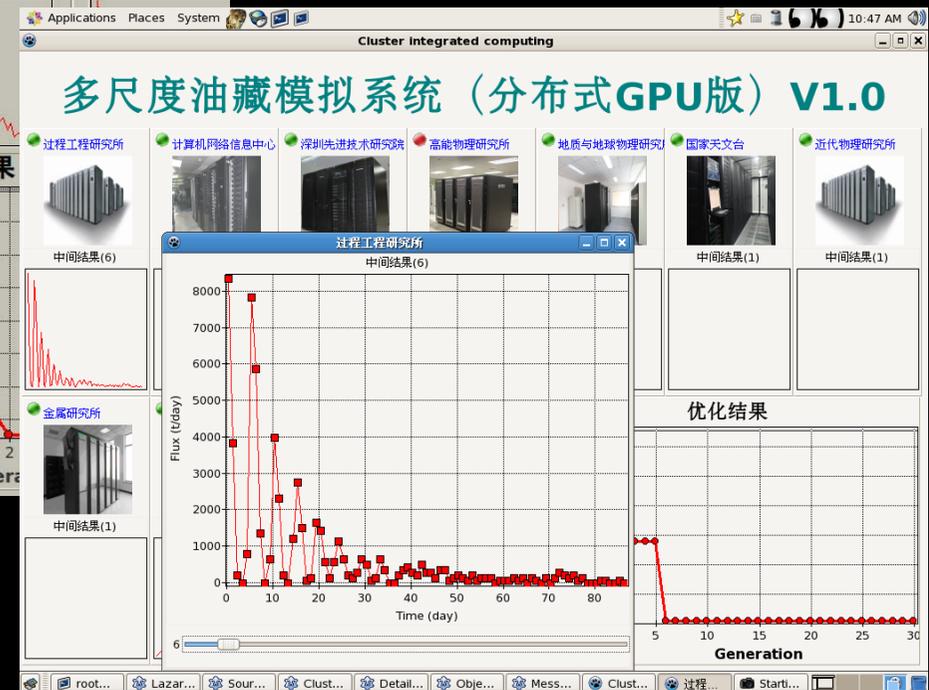
Distributed GPU computing for oil recovery

多尺度油藏模拟系统（分布式GPU版）V1.0



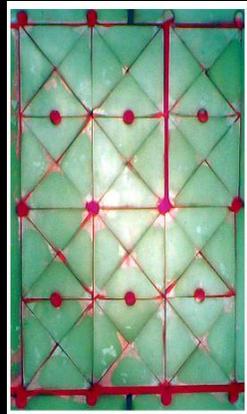
Distributed clusters:
**Simulation of individual
oil wells**

Li & Guo et al., 2010



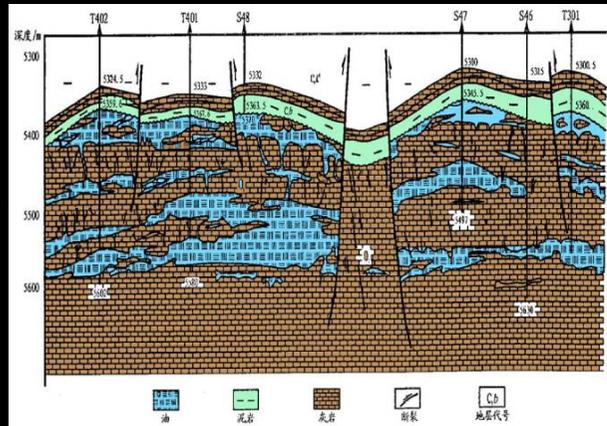
Main cluster:
Overall optimization and control

Oil recovery: fracture-cave type oil fields



Physical Experiment

0.5x0.8m days

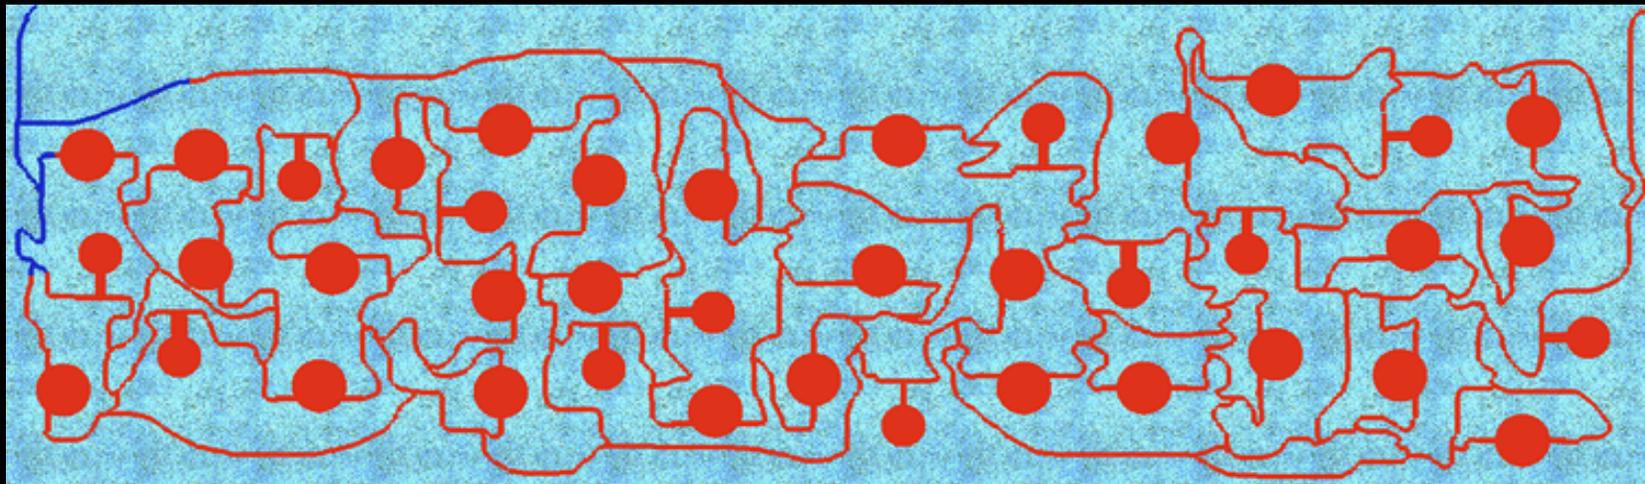


Geological Structure

20x20km years

Simulation
80 GPUs
500x150m
month
(in hours)

Li et al., 2009

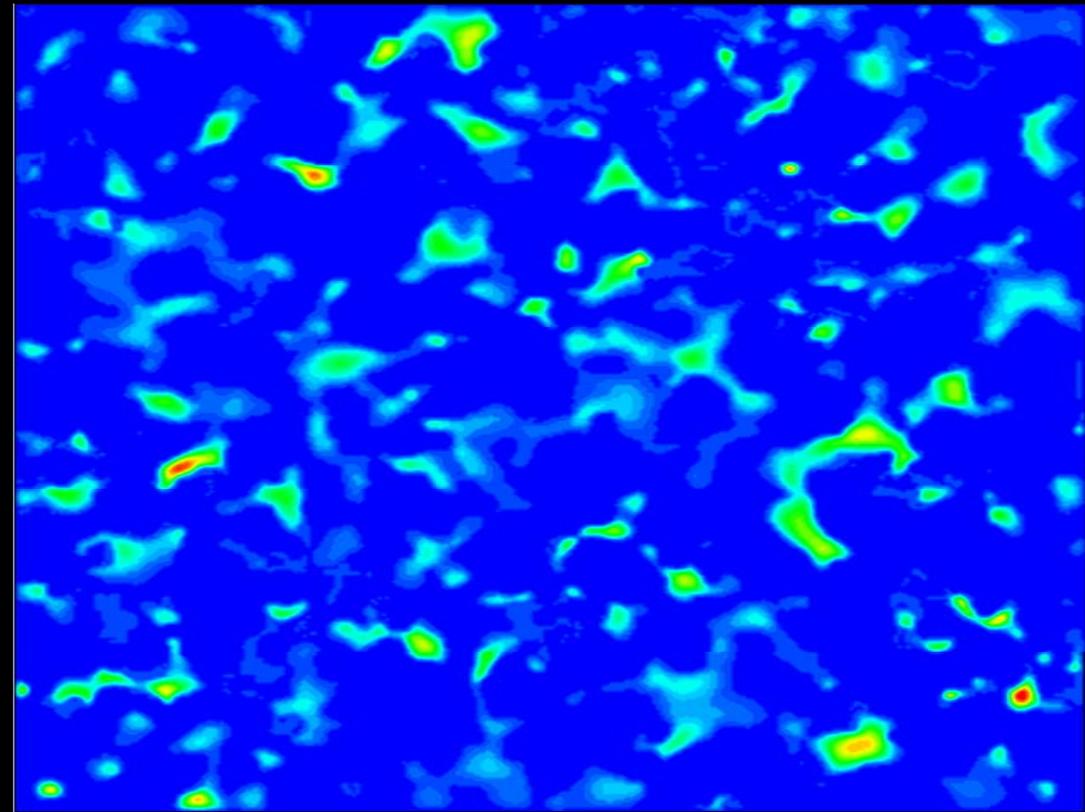
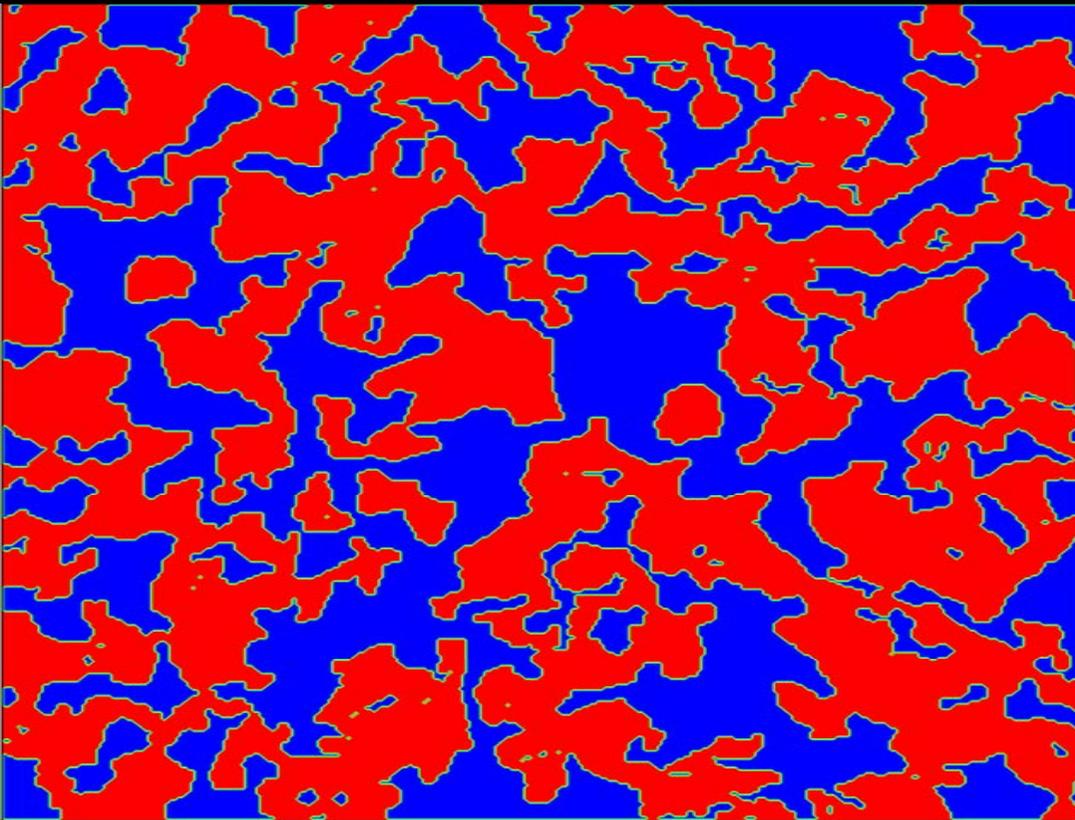


Integrated system delivered to industry

Permeability of rock samples

Experiment: 7220 mD 2 weeks
Tube network model: 8076 mD

Multi-scale model: 7322 mD <1 hour



Diversity of GPU applications in CAS

Astrophysics (2):

Biology & biochemistry:

Electrical engineering:

Geophysics:

High energy physics (2):

Material science:

Mechanics:

Medical instruments:

N-body, Image processing

Genomics, Protein structure

IC design evaluation

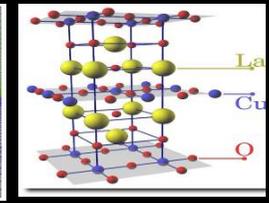
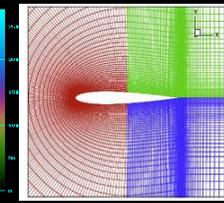
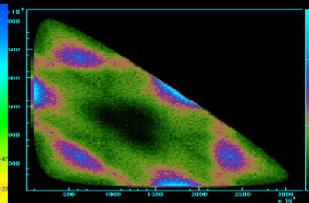
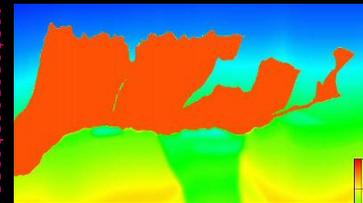
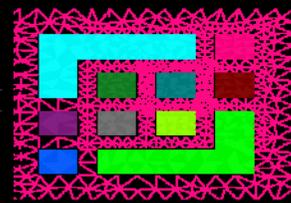
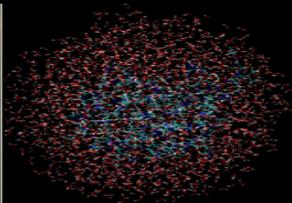
Seismic data analysis

Partial wave analysis, MD

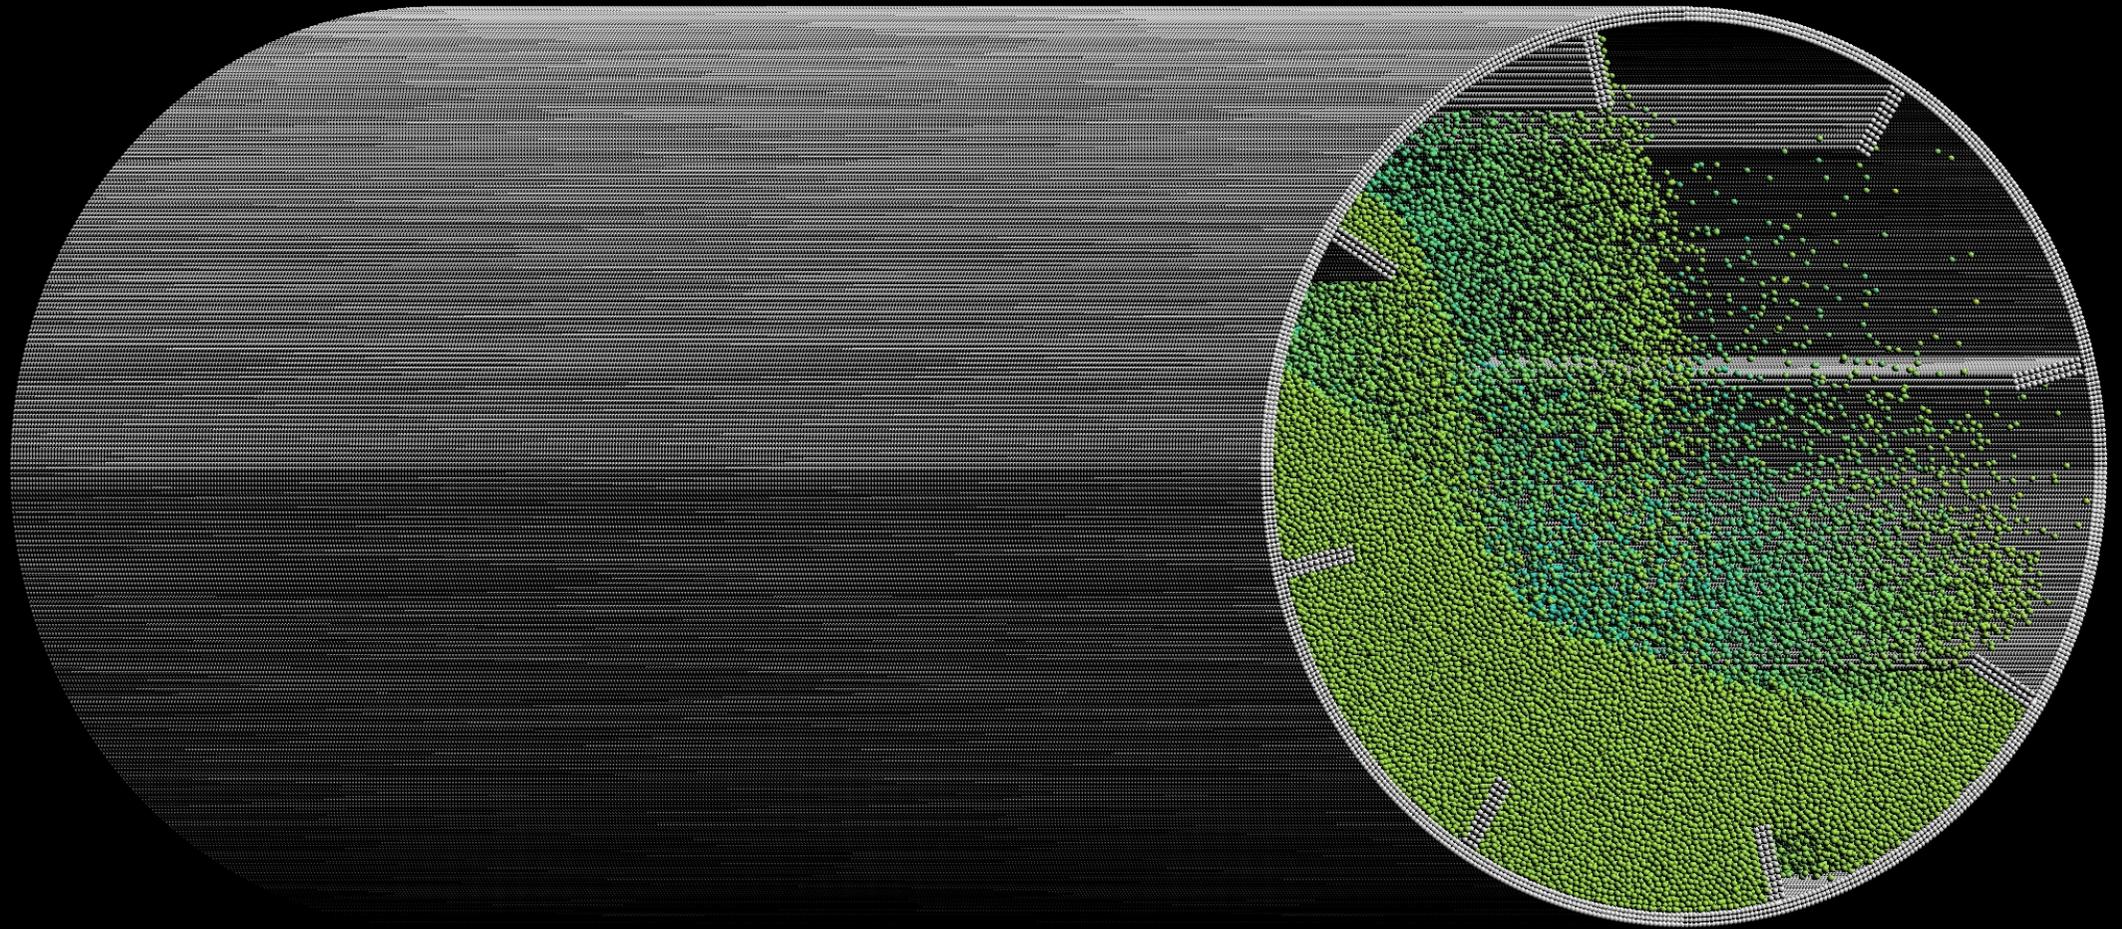
Lattice distortion and defect

CFD, multi-phase flow

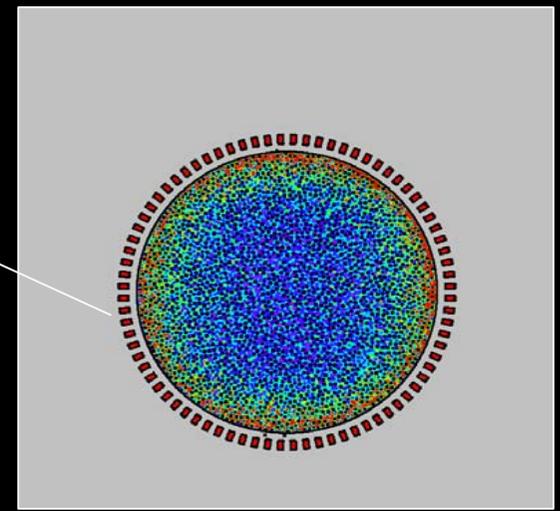
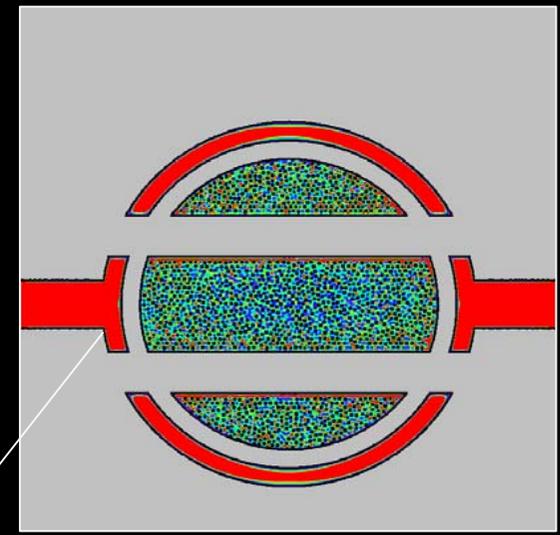
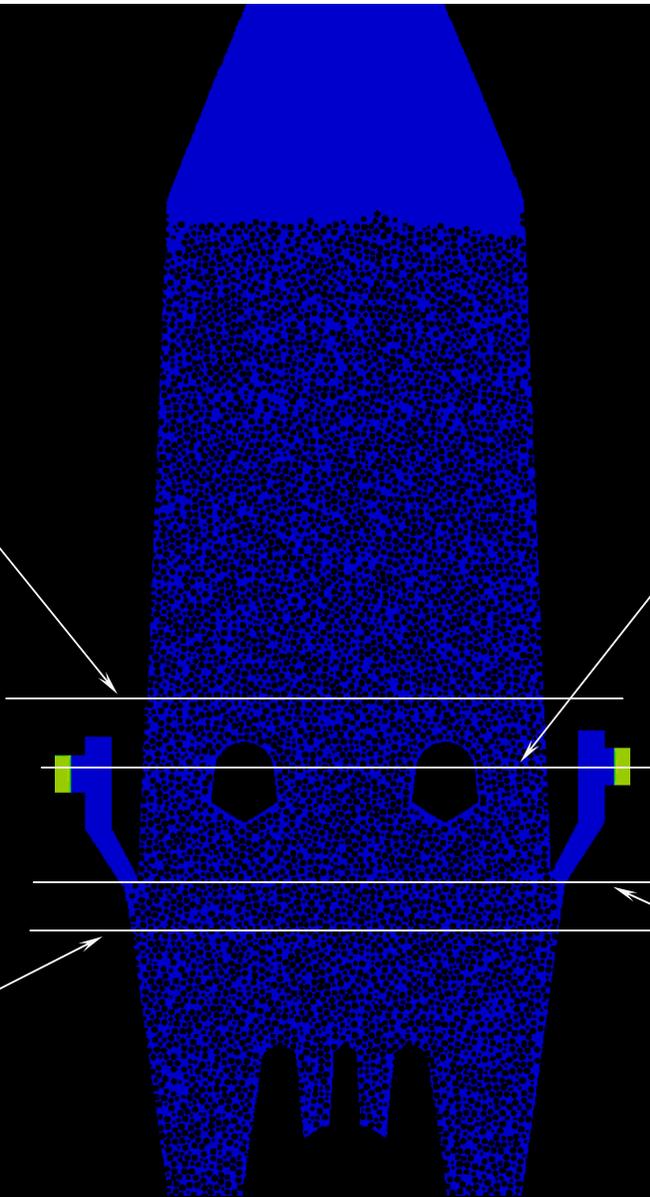
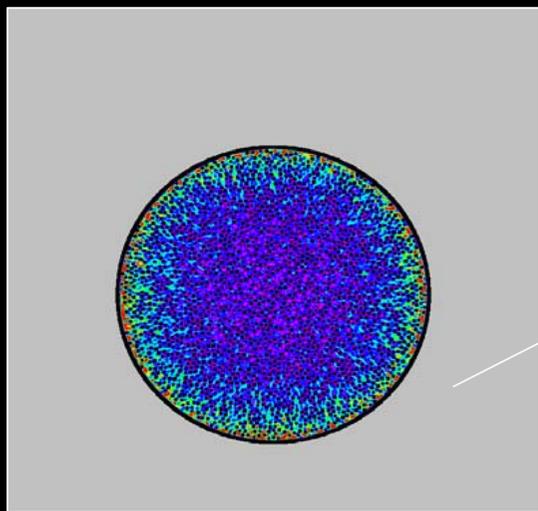
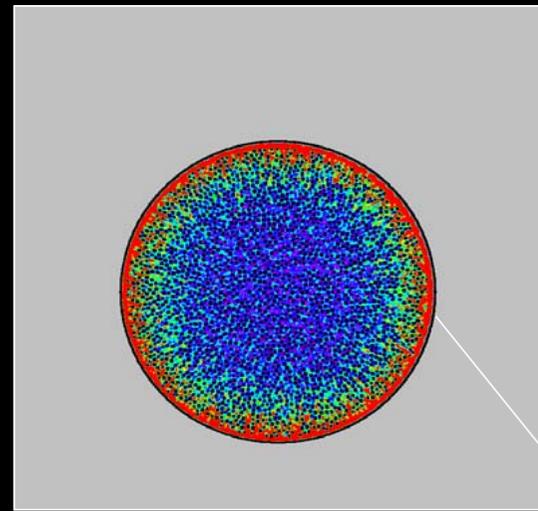
CT image reconstruction



Rotating drum: 9.6M solids, 270GPUs, 13.5*1.5m, 1/9 realtime



Xu et al., submitted to Particuology, 2010



Velocity: Minimum

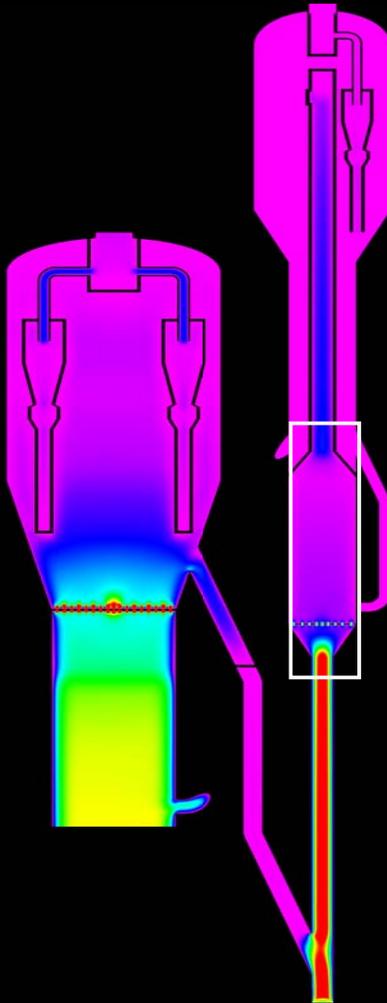


Maximum

Wang et al., 2009

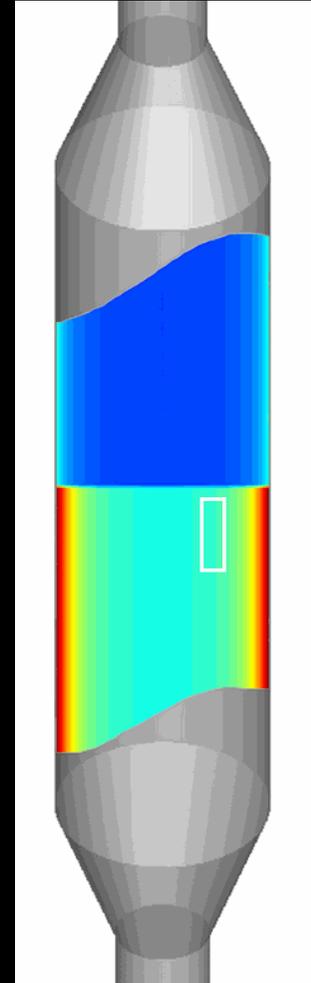
Simulation of gas solid flow on multi-scales

Reactor:
9*40m
3D
EMMS



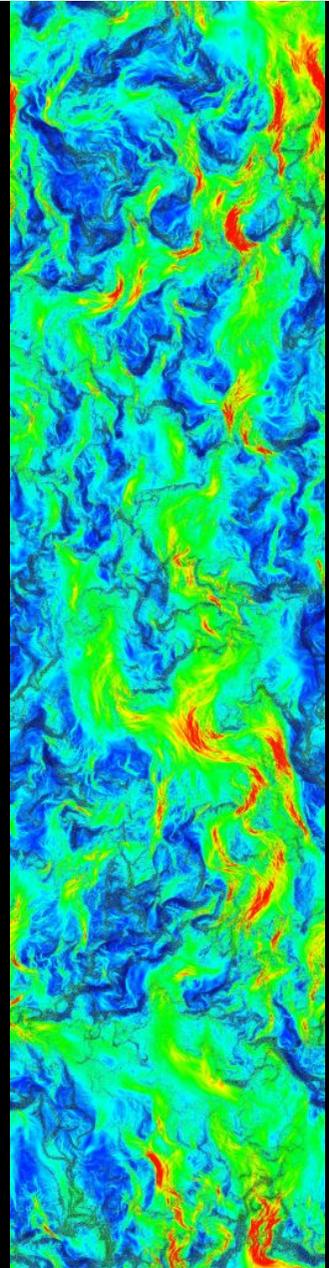
100M grids
432 GPUs
~3s
~100x*
speedup

Section:
3*10m
2D
CFD+
EMMS



1.2M cells
96 GPUs
Quasi-
realtime
~50x
speedup

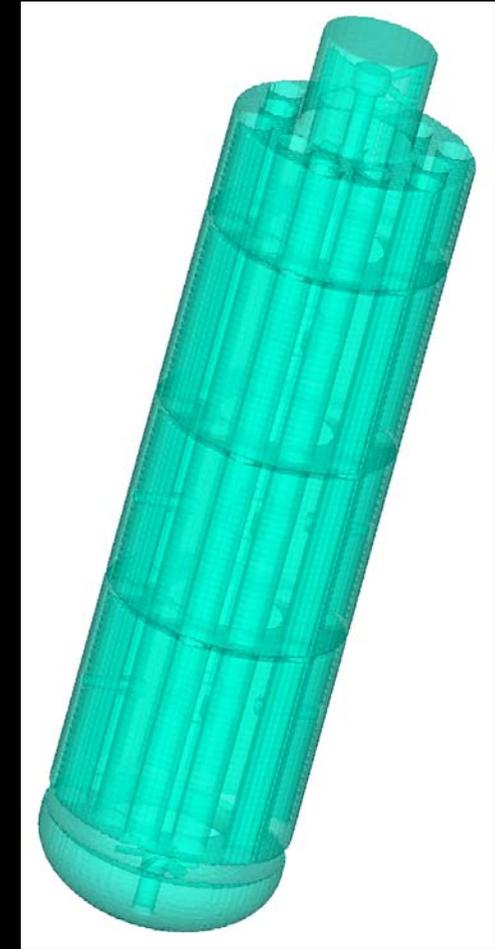
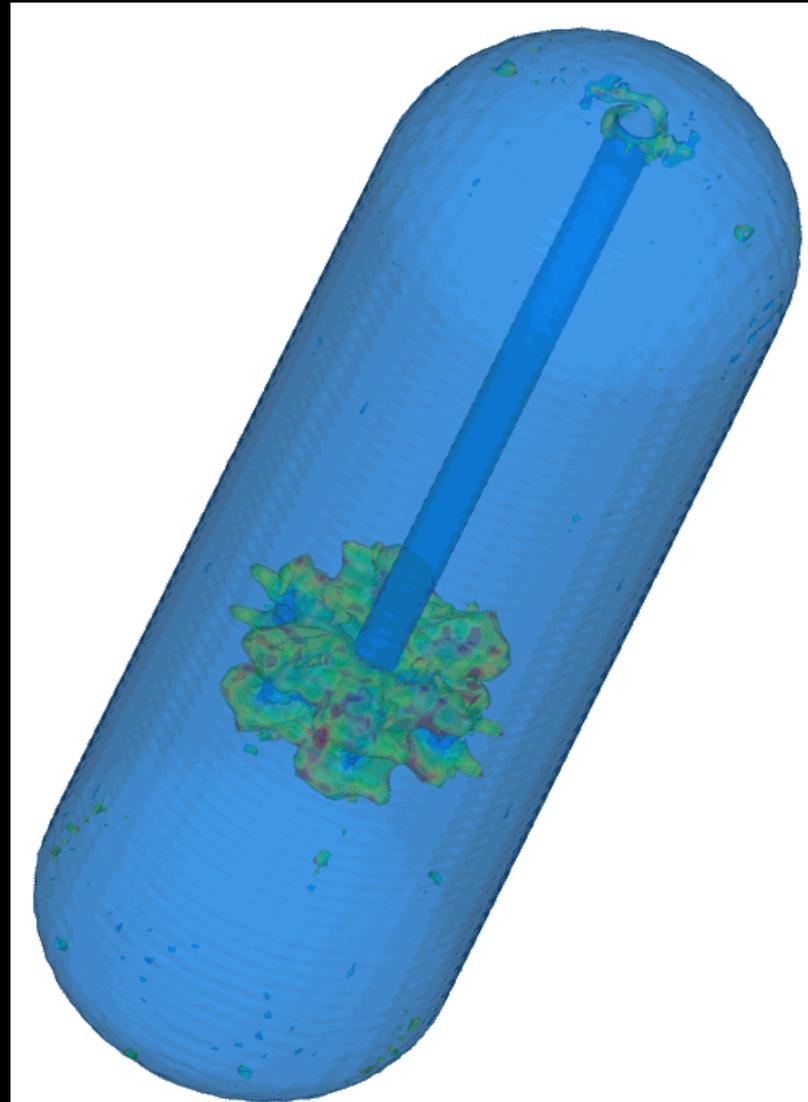
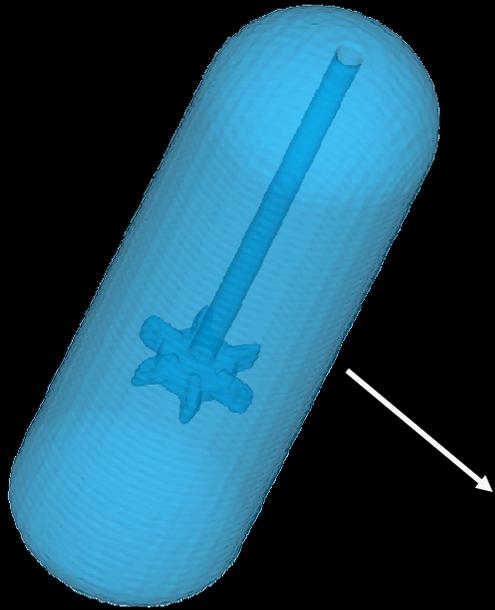
Cell:
10*48cm
2D
DNS



1M solids
~ 1G fluids
576 GPUs
30~50x
speedup

* one C2050 as compared with one core of Intel E5430 at 2.66GHz, both in single precision

Industrial three-phase systems

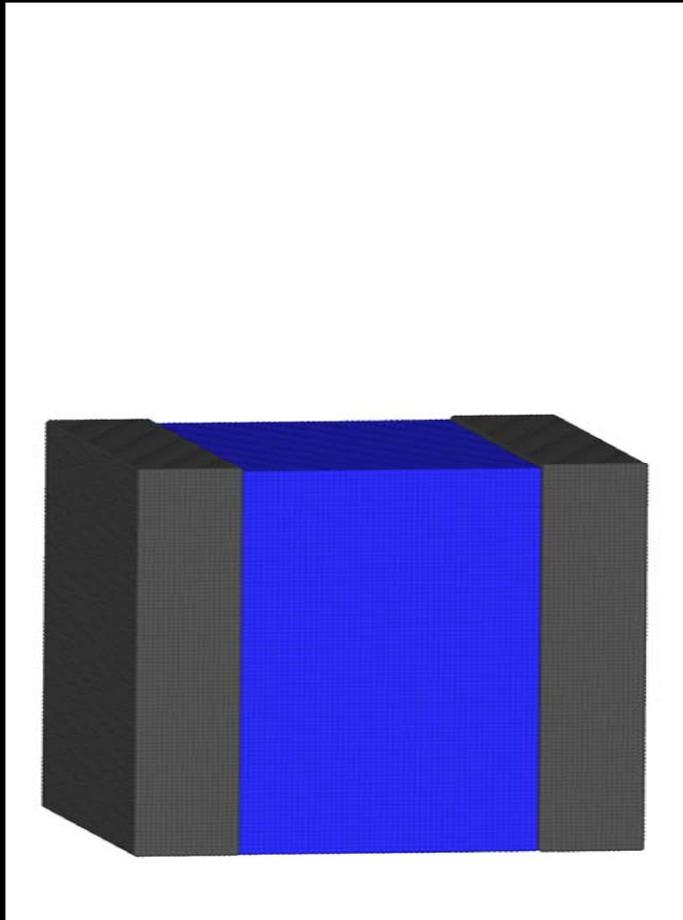


Max. scale: 1600 GPUs

Micro-/Nano-fluidics:

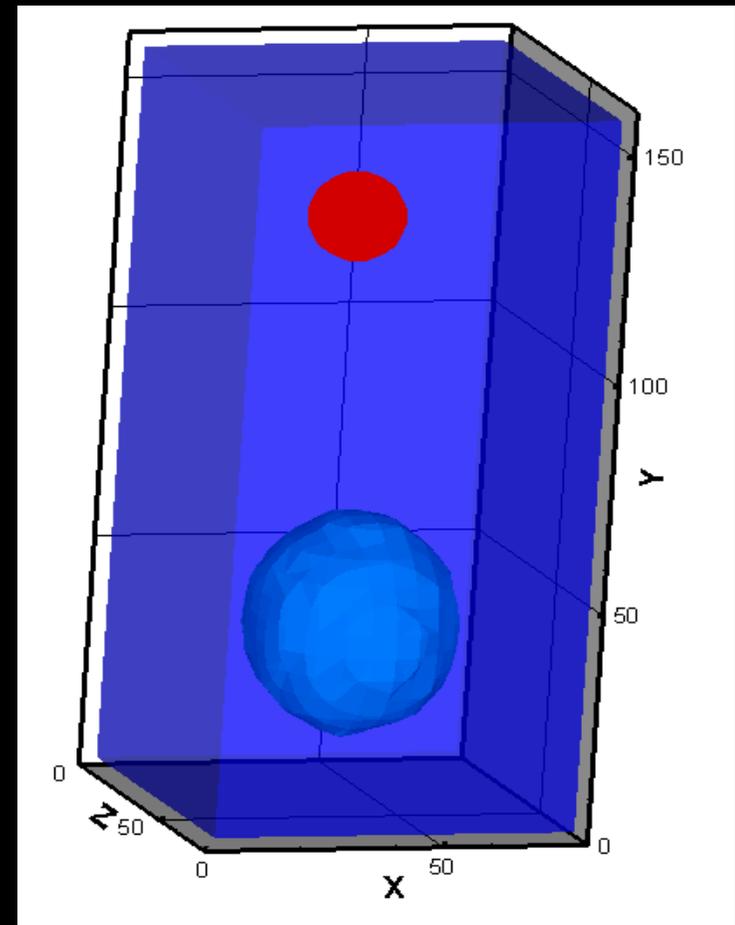
3D simulation: bubble-particle in liquid
0.1*0.1*0.15 μm , bubble mean velocity 3m/s
LJ/PP fluid at 60K, NVT ensemble

7M particles, 2 GPUs
15~50x speedup

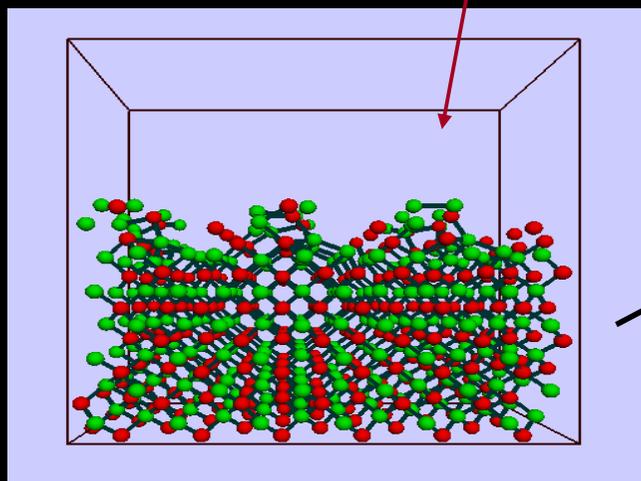
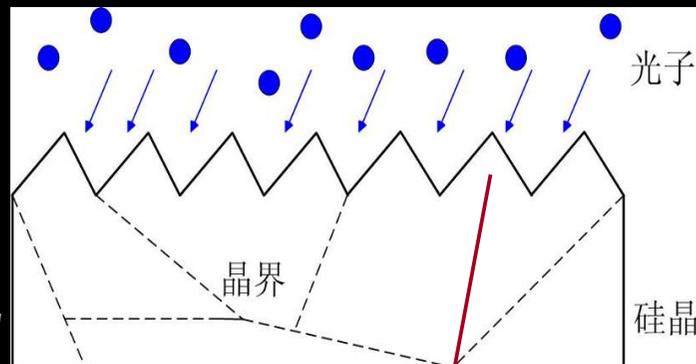


*Chen et al.,
NMF, 2008;
CSB, 2008*

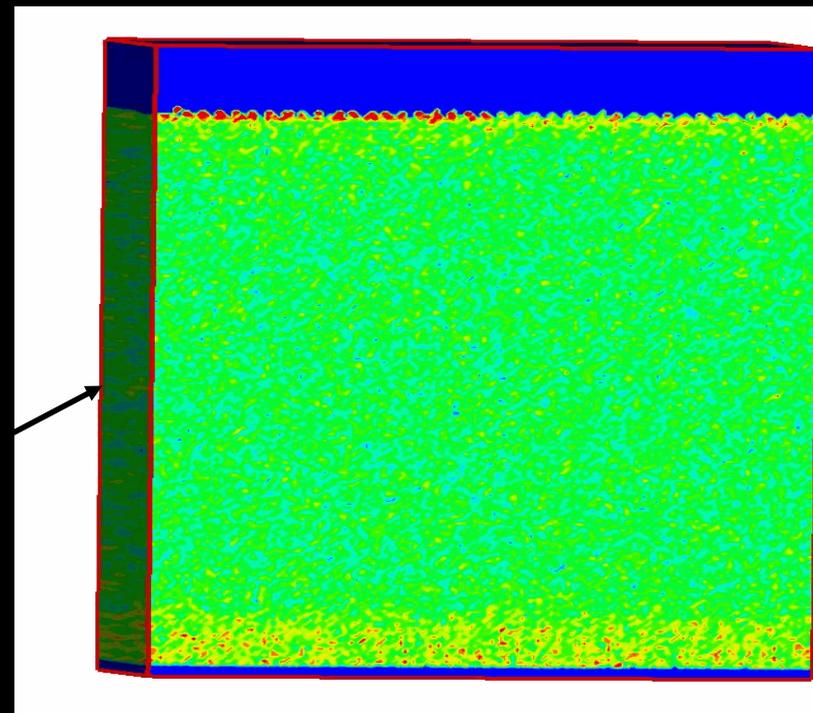
3D simulation: gas-liquid phase transition
0.1*0.05*0.1 μm ,
LJ/PP fluid at 60K, NVT ensemble
1M particles, 2GPU



Material: multi-scale structure in solar cells

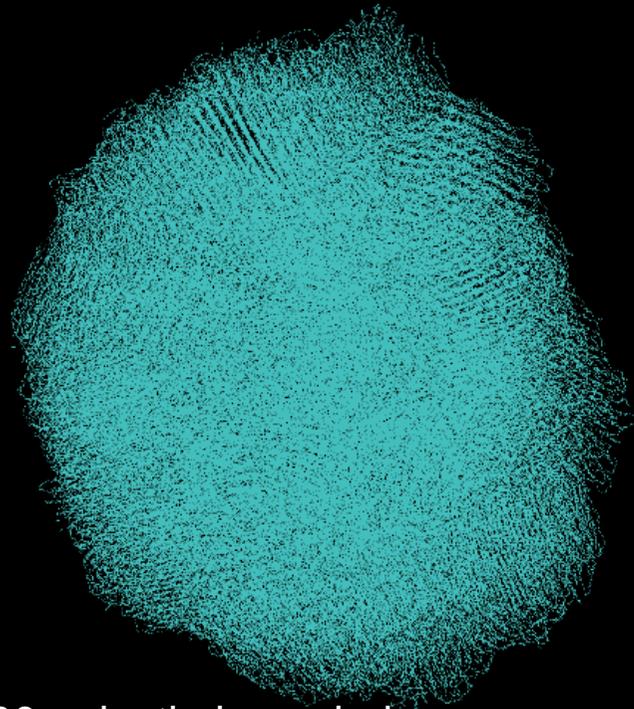


From $10E-10m$ to $10E-3m$
>20 billion atoms
1500GPU, 20x speedup
16% peak performance



Chemistry: complex and macro molecules

Xu et al., Mol. Simu., 2010



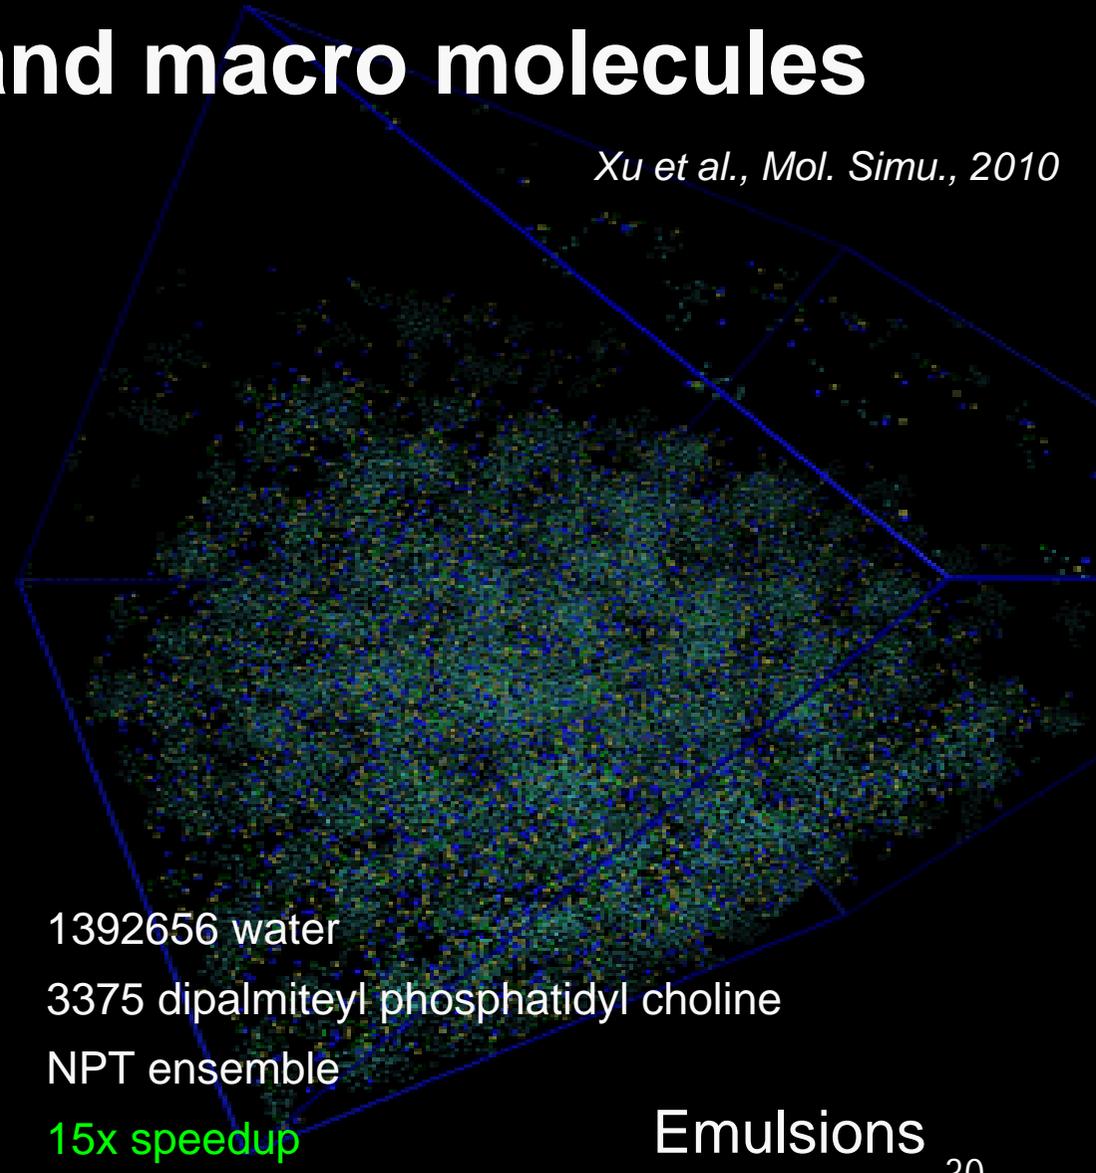
1200 polyethylene chains

Chain length: 300 CH₂

NVT ensemble

40x speedup

Polymers



1392656 water

3375 dipalmitoyl phosphatidyl choline

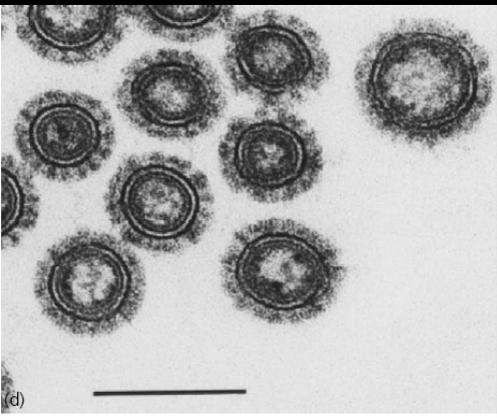
NPT ensemble

15x speedup

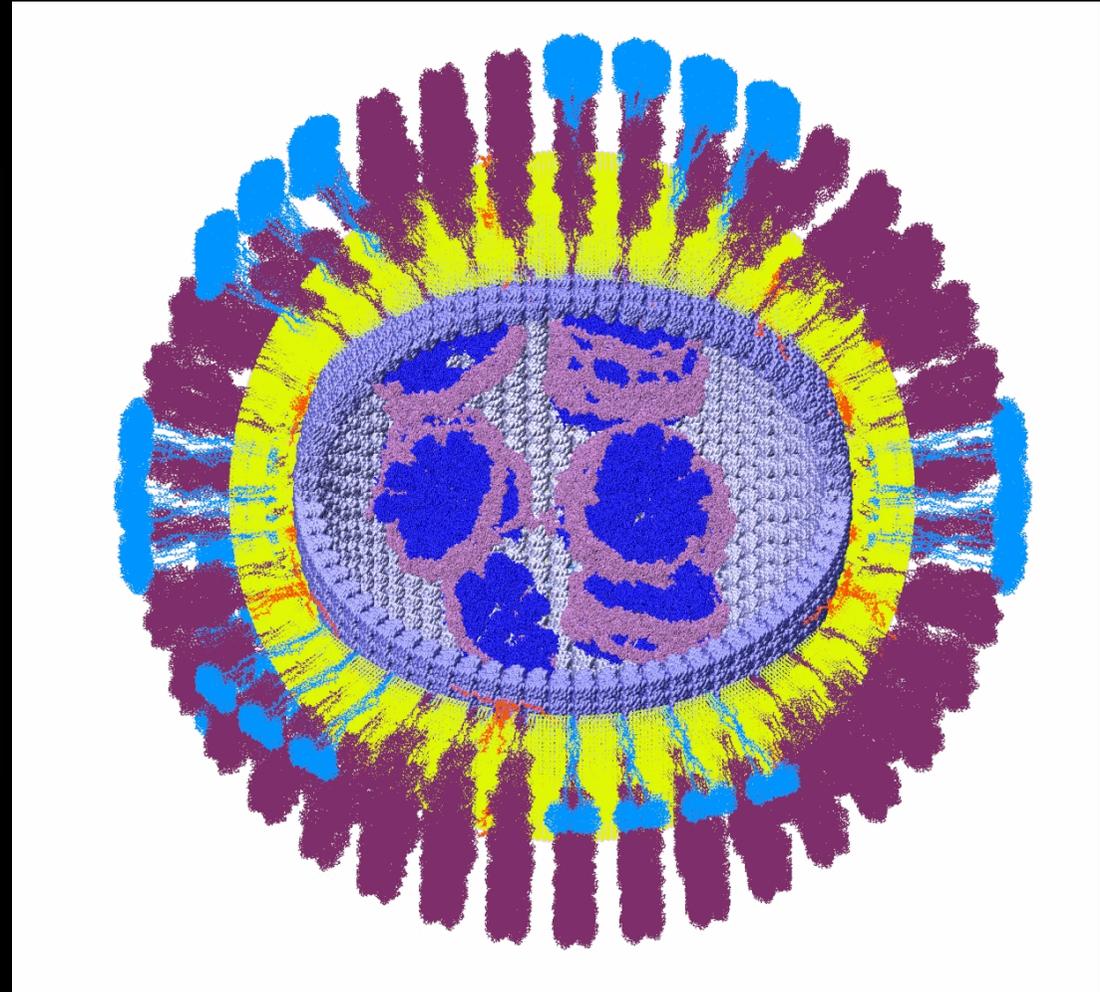
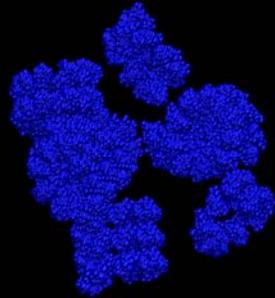
Emulsions

Molecular dynamics simulation of *Swine flu*

H1N1



100 nm



300M atom/radicals, >100ns, 1728GPUs

Xu & Ren et al., 2010

Prospect of virtual process engineering

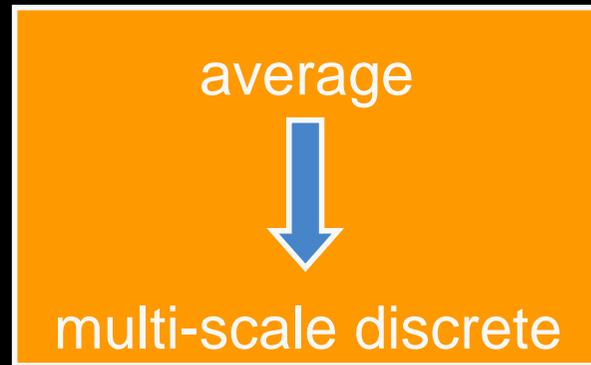
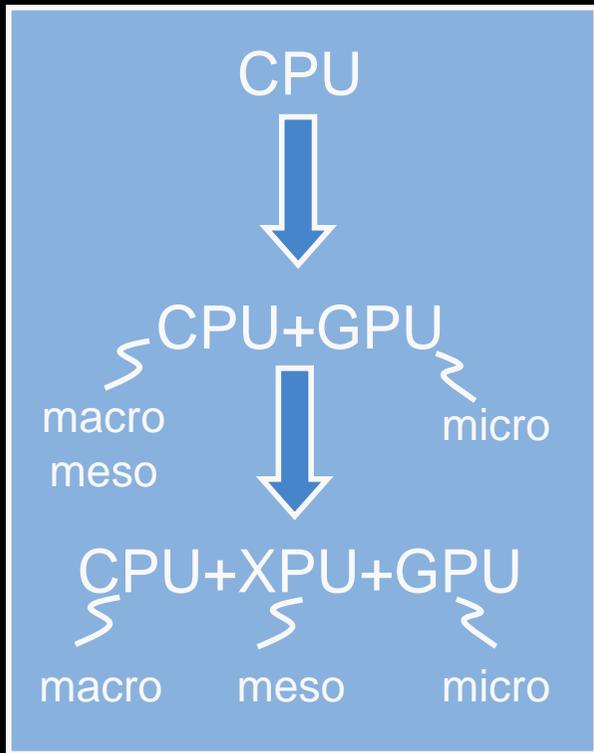
hardware

+

software



application

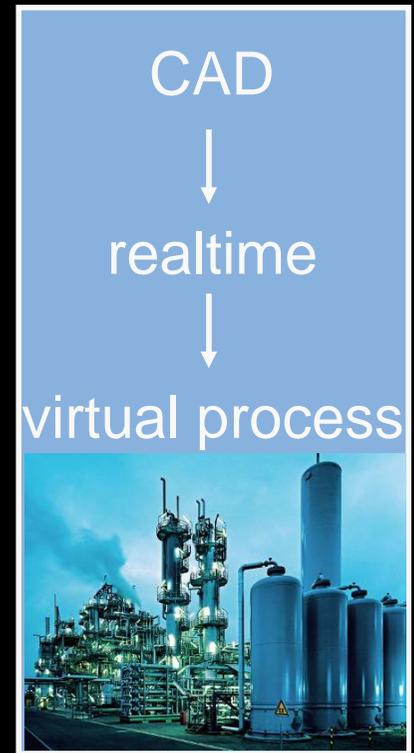


simulator



pre-/post-processing

engineering design



This work is supported by:

NSFC MOF MOST CAS

SinoPec PetroChina BaoSteel YeTech

BHPb Alstom nVidia GE Unilever Shell

www.multiscalesci.org

www.ipe.cas.cn/emms (under construction)



GroupEx2010, Shanxi, Photo by Xianfeng He

Thanks!