Multi-GPU 시스템을 이용한 입자계 기반 해석사례
Particle-based Simulations on multi-GPU Systems

2010-10

coahn@metariver.kr

Metariver Technology
www.metariver.kr

Metariver Technology

GPU & CUDA Technology

What is GPU?

Entertainment

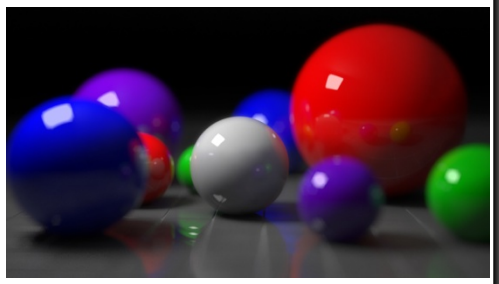

Games
Movie
CG
VR
:

Engineering
Medicine
Science
Finance
Biology
:

GPU

High Performance Computing

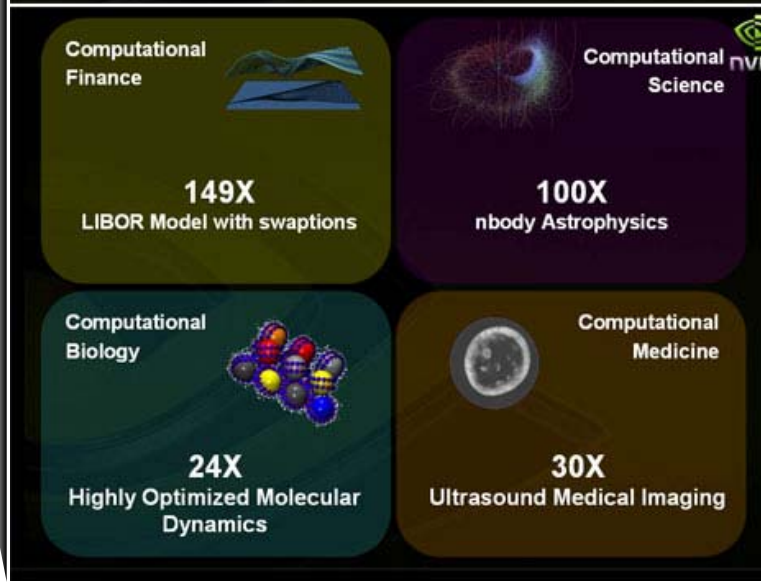

GPU & CUDA Technology

- GPU(Graphics Processing Unit)는 매우 빠른 시간에 복잡한 3차원 이미지를 화면에 출력하기 위해 사용되는 그래픽카드 전용 Processing Unit
- CUDA(Compute Unified Device Architecture)는 GPU를 이용하여 고속 연산이 가능하도록 하는 기술로서, 그래픽 작업을 처리하도록 작성된 고속 shader언어를 과학/공학 계산에 활용 가능한 형태로 사용할 수 있도록 작성된 형태
- CUDA는 GPU가 가지는 수십~수백 개의 core를 동시 활용함으로써 수많은 thread를 고속으로 처리할 수 있도록 하며, 이들 core들은 공통의 GPU 메모리를 사용가능
- 연산전용의 단일 GPU는 다수의 CPU를 고속의 network로 연결하여 대규모 계산을 공동 처리하도록 하는 기존의 supercomputing(병렬처리, Parallel Processing) 방법에 비해 보다 낮은 비용으로, 월등한 고속 연산 가능
- NVIDIA에서 shader언어를 다목적으로 사용할 수 있도록 한 GPGPU(General Purpose GPU) 기술을 더욱 확장한 것으로, 2008년 이후 미국 및 유럽 등 선진 연구자들에 의해 지속적으로 성공적인 활용사례가 보고되고 있음

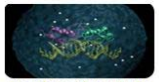
Bio-Informatics and Life Sciences


Computational Chemistry


Computational Electromagnetics and Electrodynamics

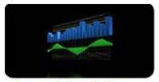
Computational Finance


Computational Fluid Dynamics


Data Mining, Analytics, and Databases

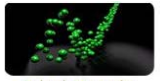
Molecular Dynamics


Weather, Atmospheric, Ocean Modeling, and Space Sciences

Parallel Computing (HPC)

Multiple Instruction Multiple Data

(conventional parallel processing)

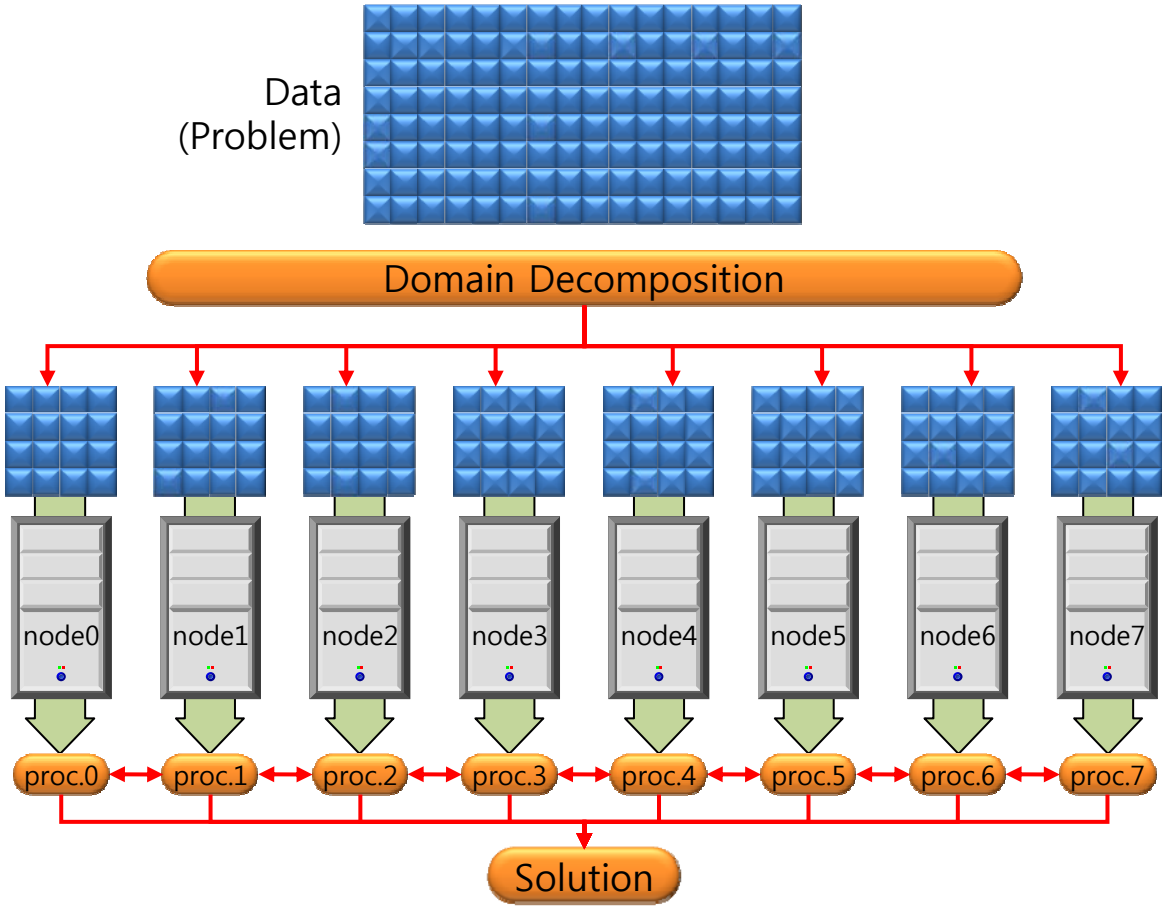

GPU Computing (CUDA)

Single Instruction Multiple Data

(CUDA technology)

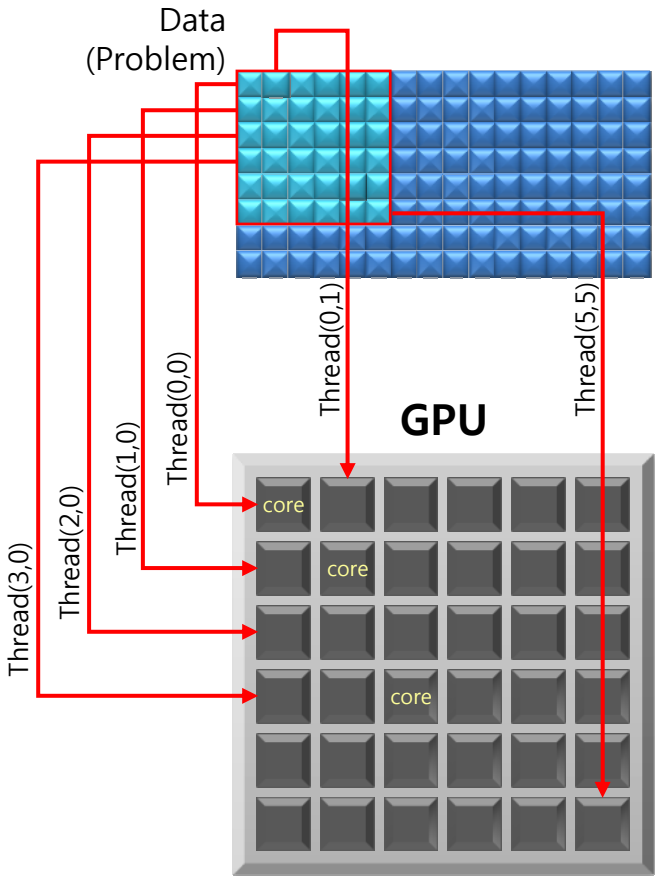

- GPU is a data-parallel processor
- Thousands of parallel threads / Thousands of data elements to process
- GeForce 8800 has 128 streaming processor cores and 512MB RAM
- Tesla C1060 has 240 streaming processor cores and 4GB RAM

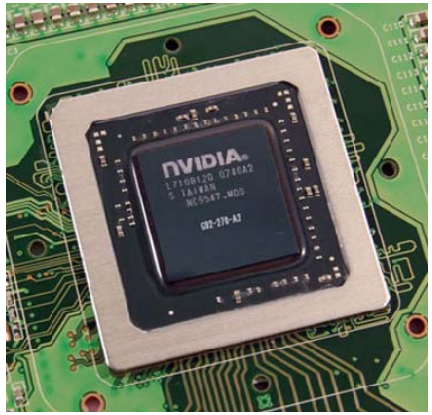

Metariver Technology

GPU Architecture

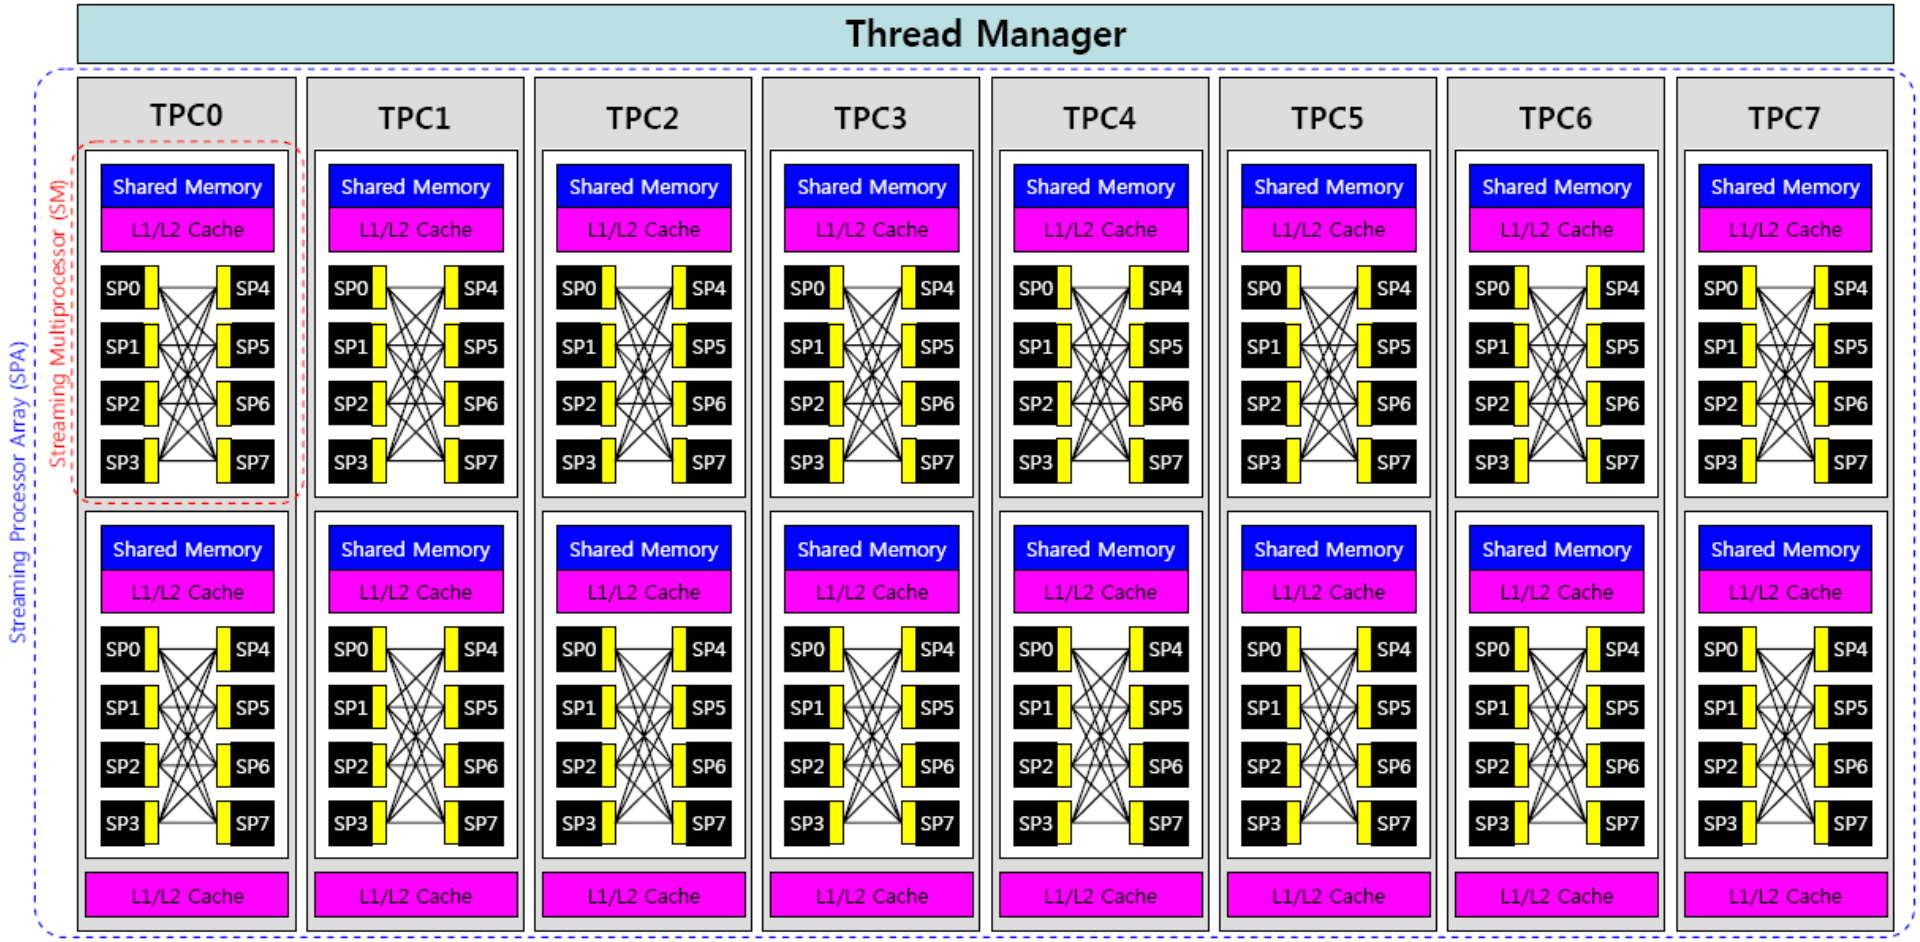

- SPA : Streaming Processor Array
- TPC : Texture Processor Cluster (2SM+TEX)
- SM : Streaming Multiprocessor (8SP) - Multi-threaded processor core, Fundamental processing unit for CUDA thread block
- SP : Streaming Processor - Scalar ALU for a single CUDA thread
- The size of the Shared Memory is 16KB.

Metariver Technology

Inside CUDA kernel

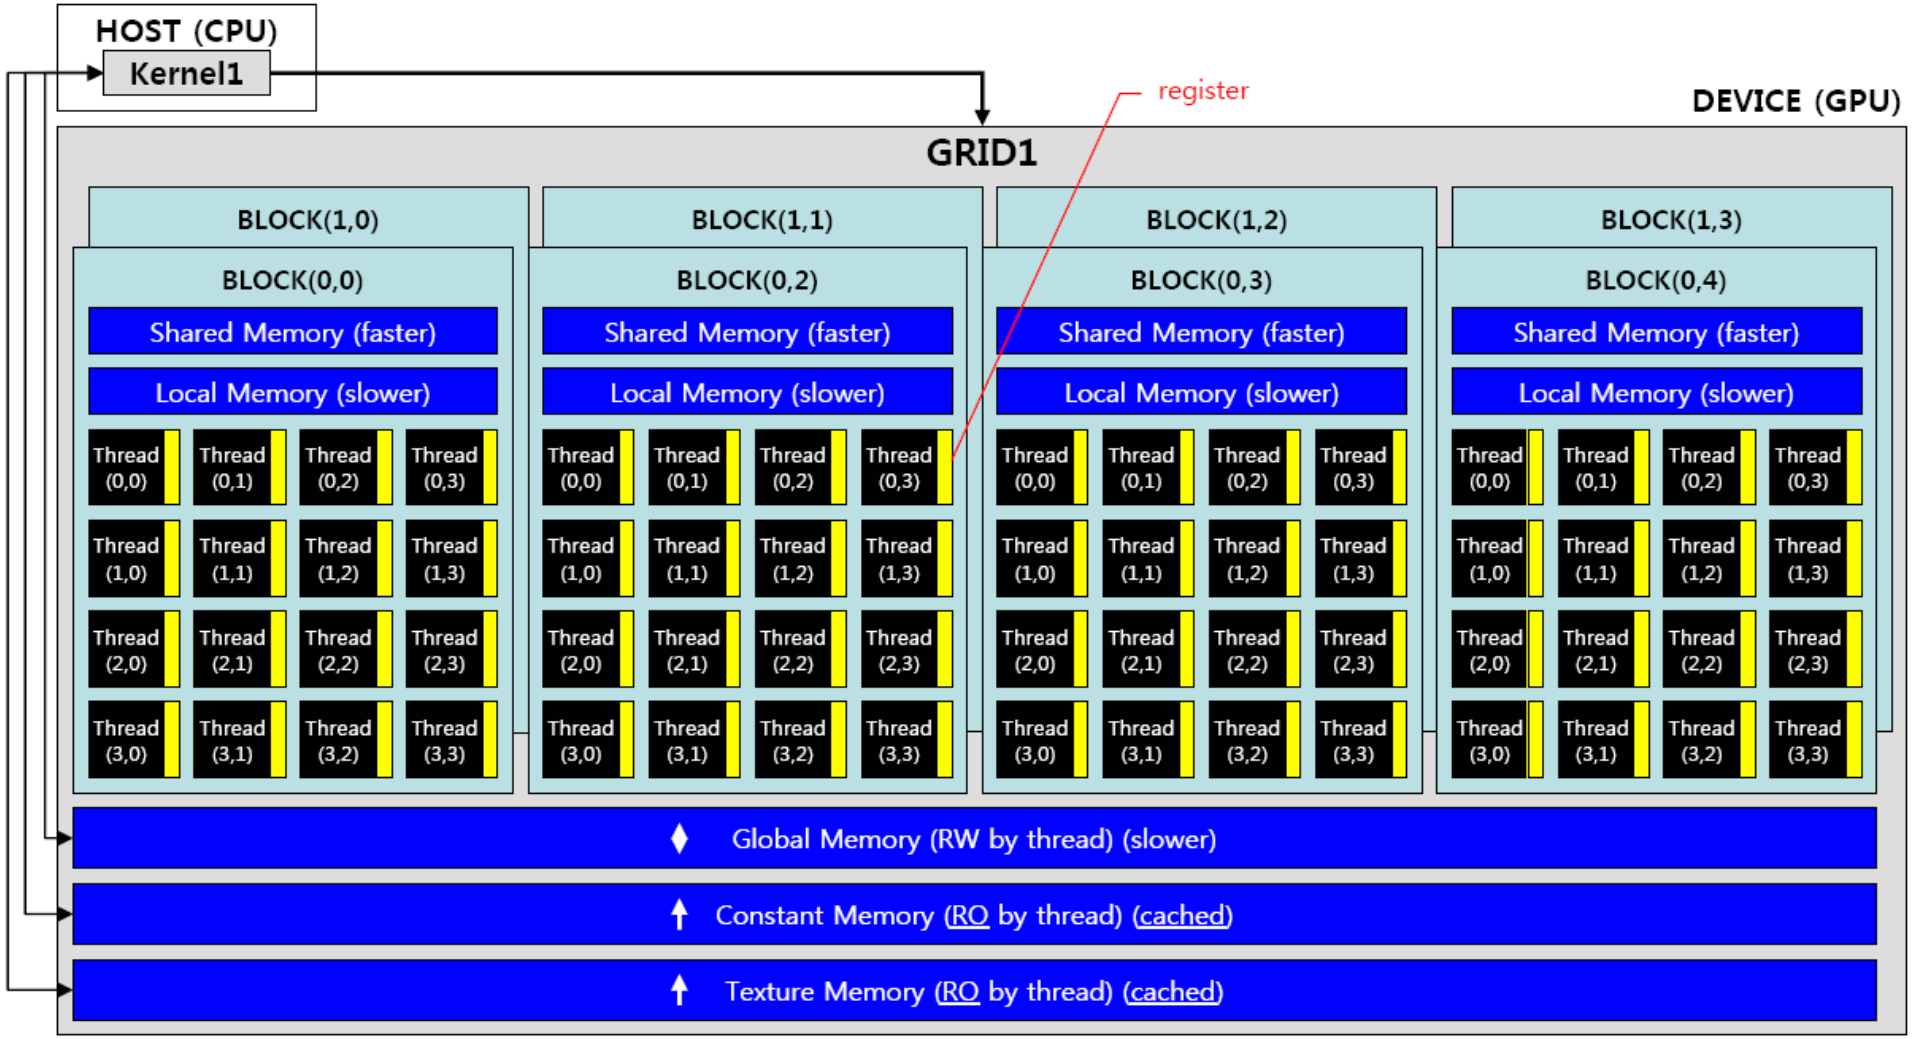

- Block ID : 1D/2D Thread ID : 1D/2D/3D
- Grid is launched on the SPA. Threads are assigned to SMs in block granularity. Up to 8 Blocks to each SM as resources allows.
- Host(CPU) can RW Global/Constant/Texture memories
- Two threads from different blocks CANNOT cooperate.

Metariver Technology

GPU vs. CPU

Wetting Simulation using SPH (Smoothed Particle Hydrodynamics)

- W,D : 50mm
- H : 100mm
- Dpore : 1.5mm
- N : 83,300 EA

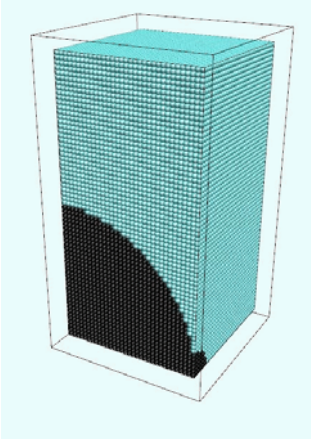

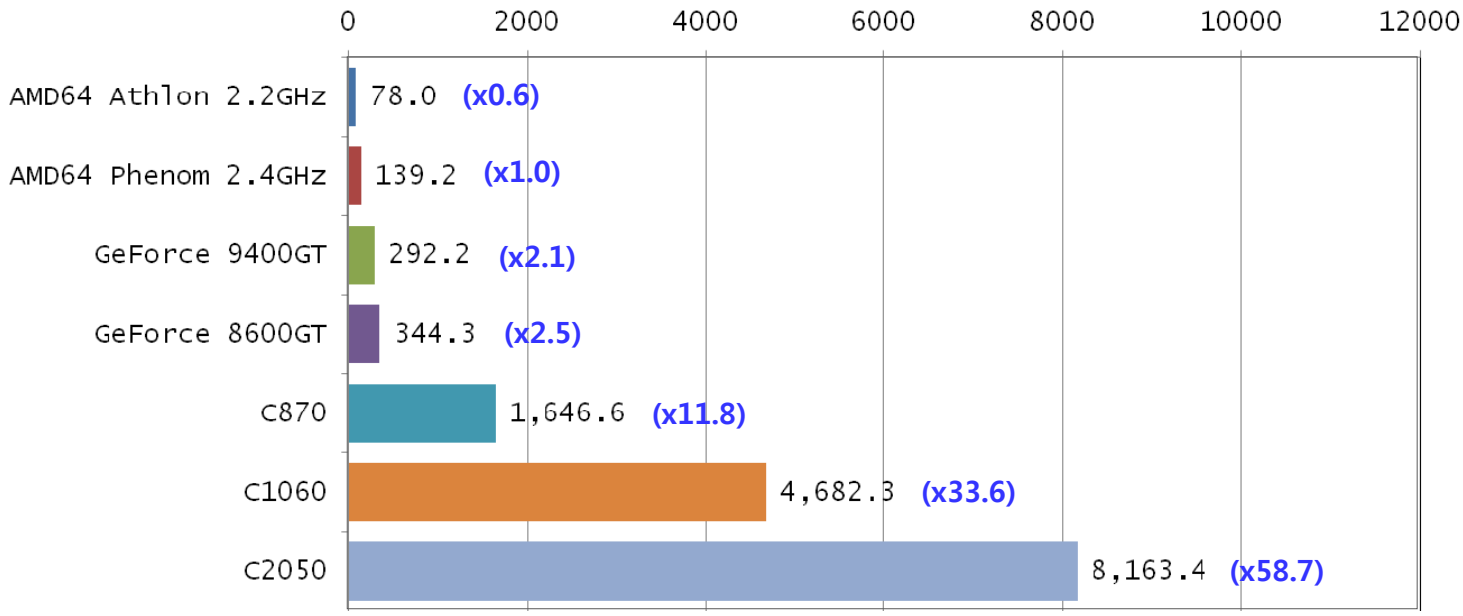

GPU vs. HPC

Lattice Boltzmann Method (CFD)

Lid-driven cavity flow

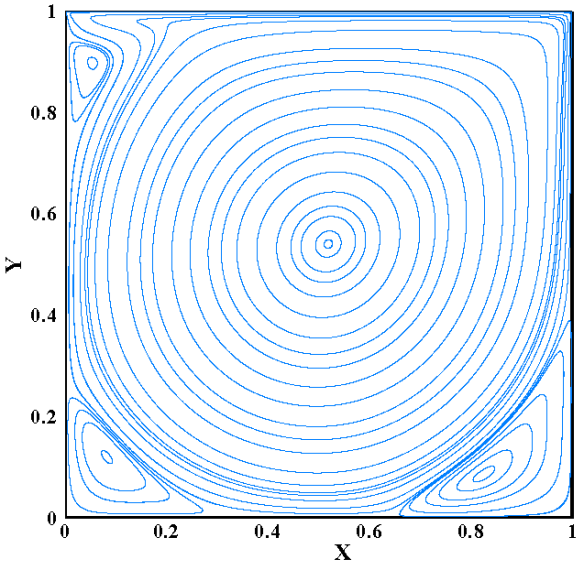

Speed-up :

$$s(n) = T_s / T_p$$

T_s : WCT of serial code [sec.]
T_p : WCT of parallel code [sec.]
n : the number of processors

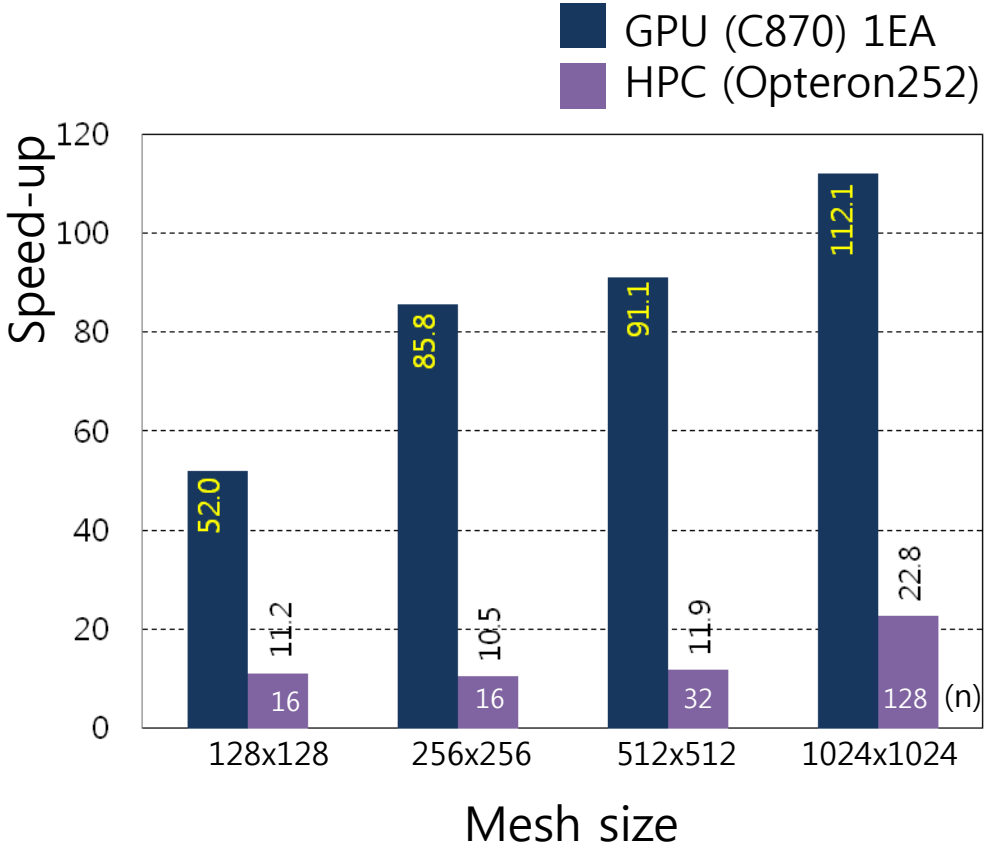

HPC vs. multi-GPU

HPC
(Parallel Processing)

multi-GPU

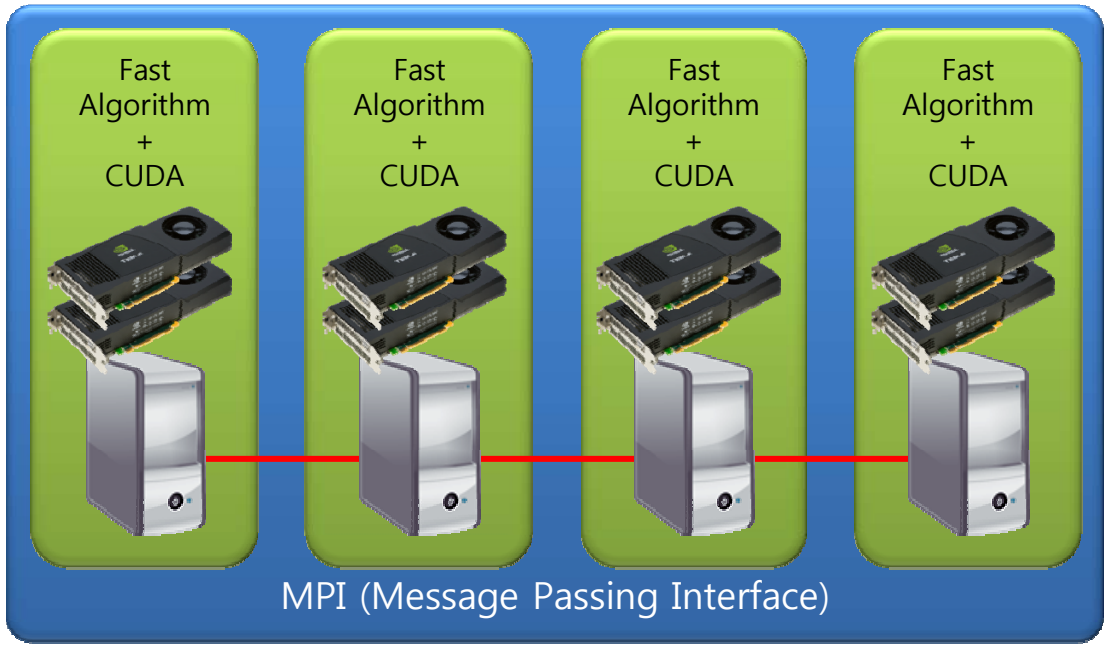

Inside multi-GPU system

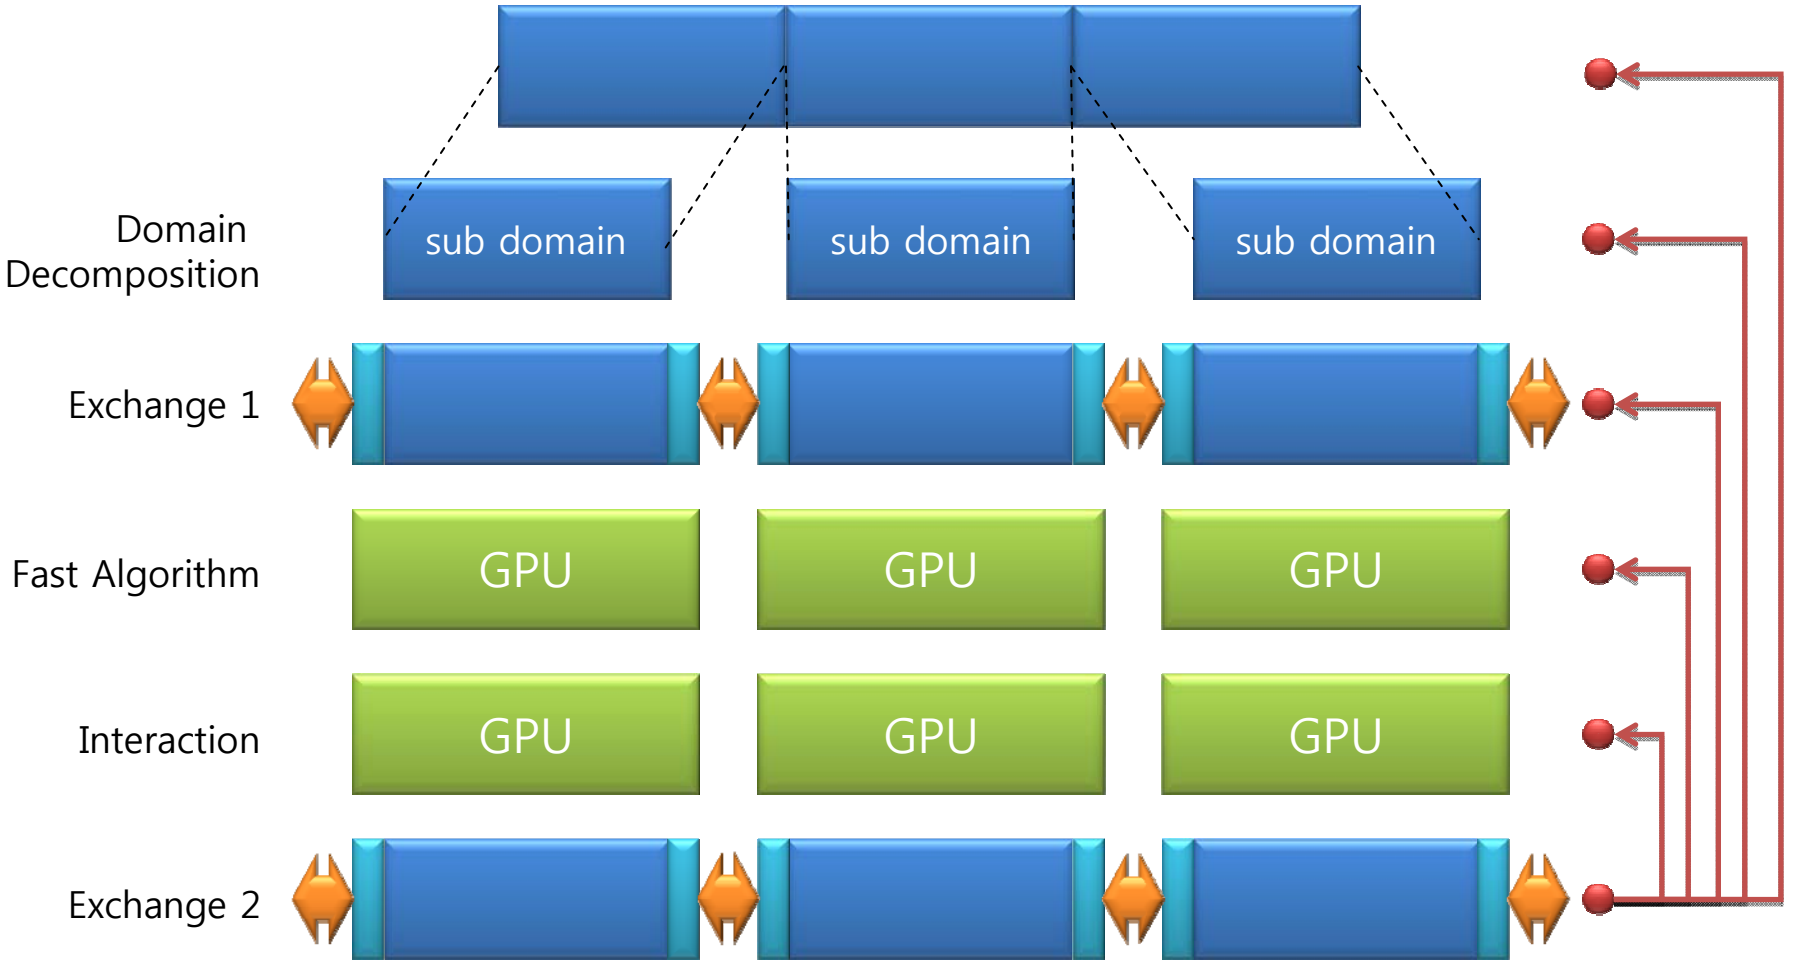

Metariver Technology

SAMADII

: Particle-based multi-physics solver

Metariver Technology

SAMADII

- Particle-based multi-physics solver
- H/W acceleration using multi-GPU (GPU cluster)
- S/W acceleration (Fast algorithm)
- Discrete Element Method (DEM)
- Magnetic particle, charged particle simulation
- Wetting simulation
- Smoothed-Particle Hydrodynamics (SPH) *
- Fluid-Solid Interaction (FSI) *
- Deformable body simulation *

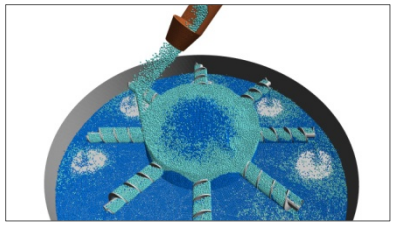
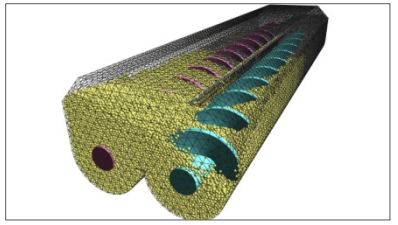
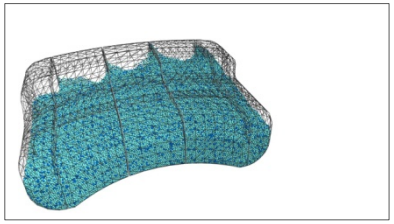
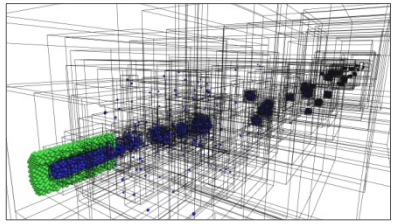
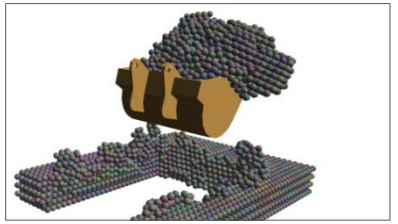
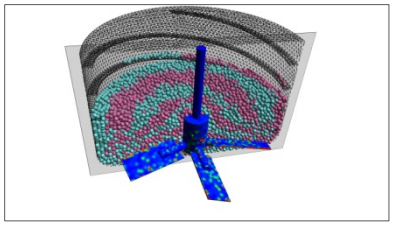
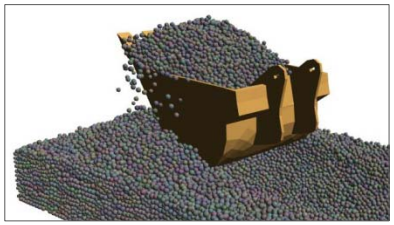
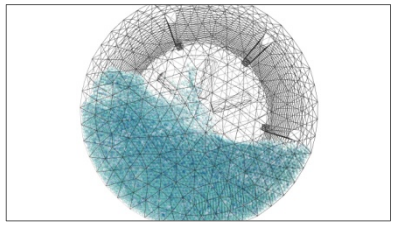
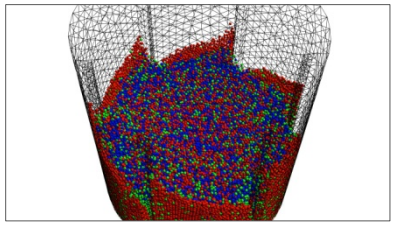

* under development

Metariver Technology

SAMADII

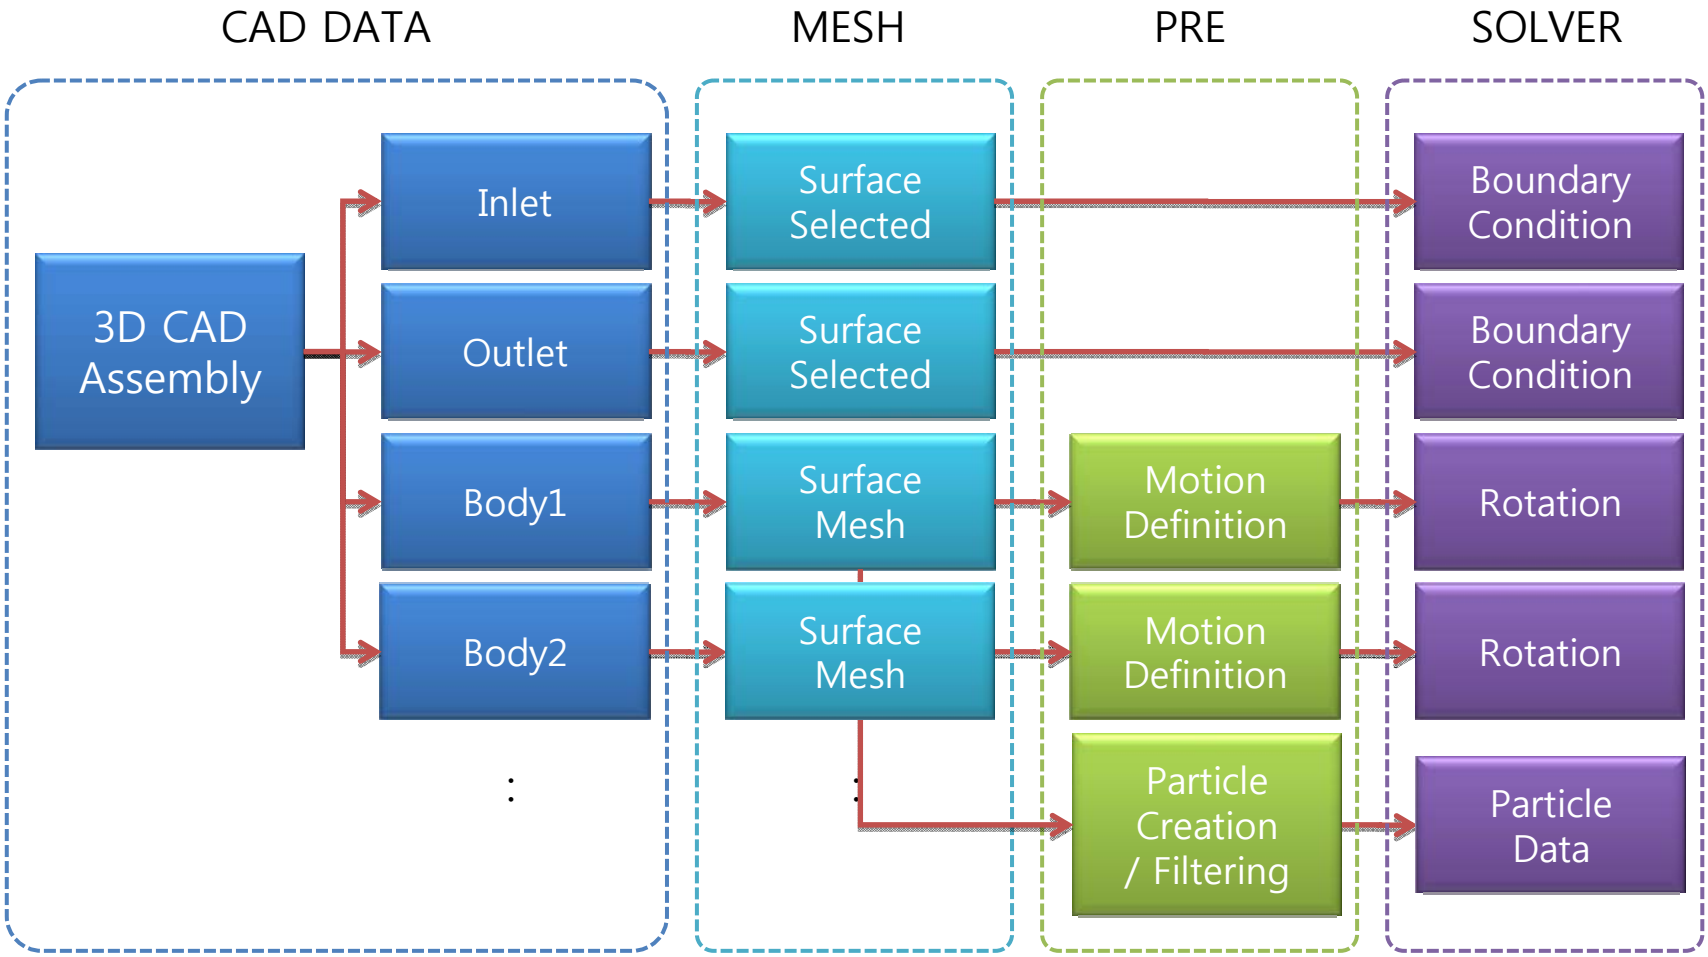

SAMADII - GUI

- Graphics Engine : **HyperCube4** (based on OpenGL & .NET, self-developed)
- Specially designed for high-speed rendering

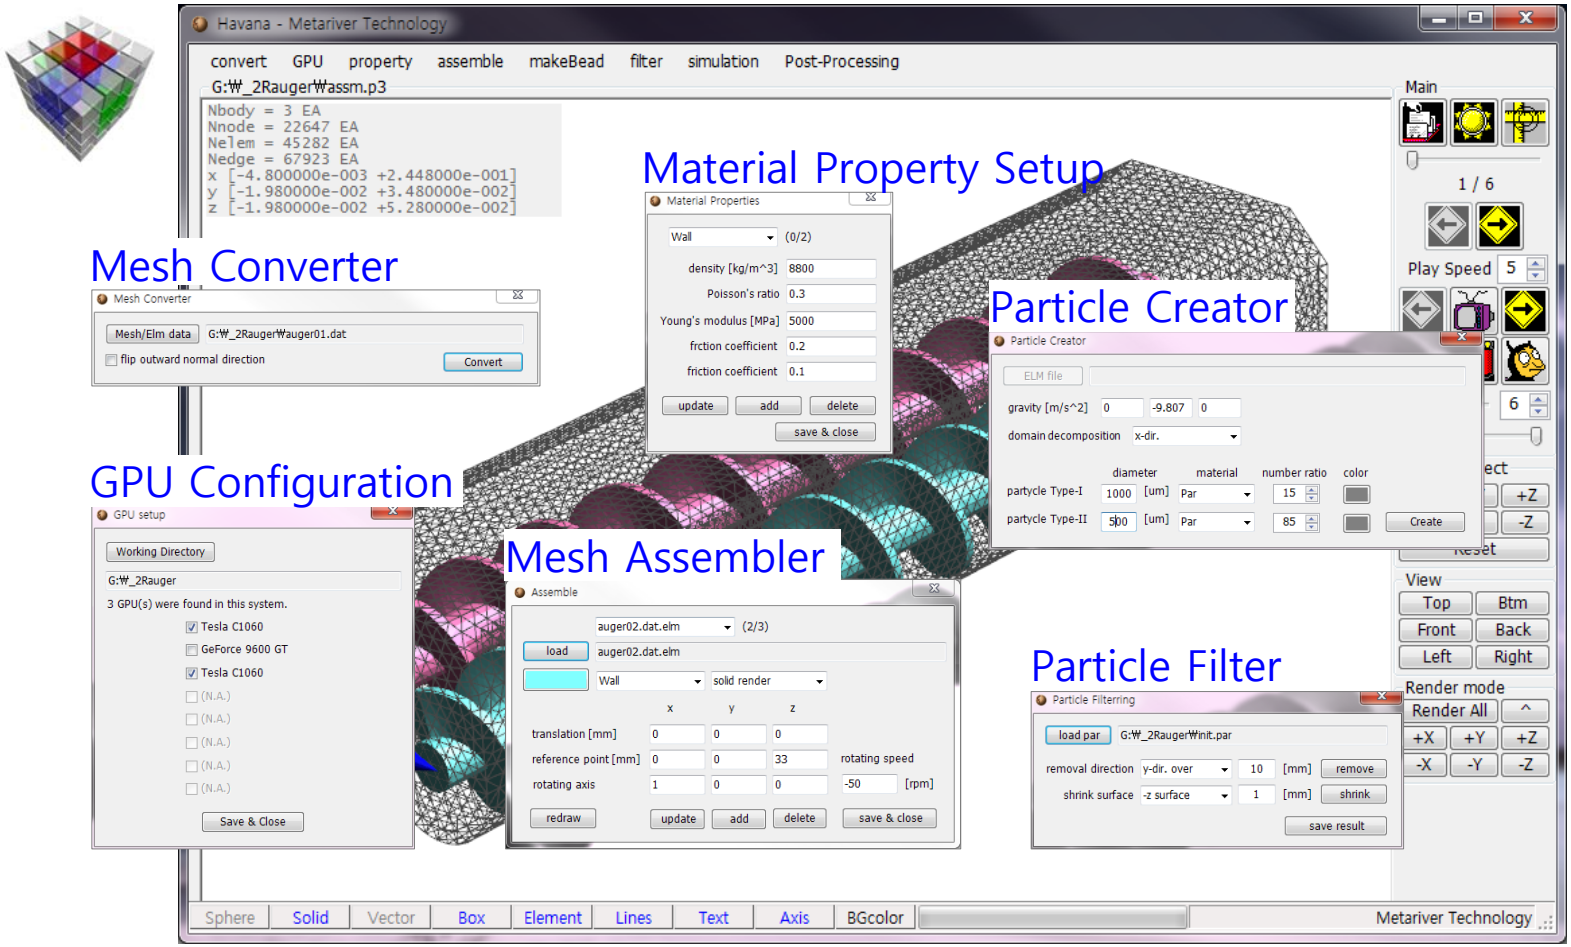

Example-1 : Excavator

N_p : 80,000 EA
N_e : 5,653 EA
500,000 steps
Multi-CPU (HPC 16 core)
2010-02

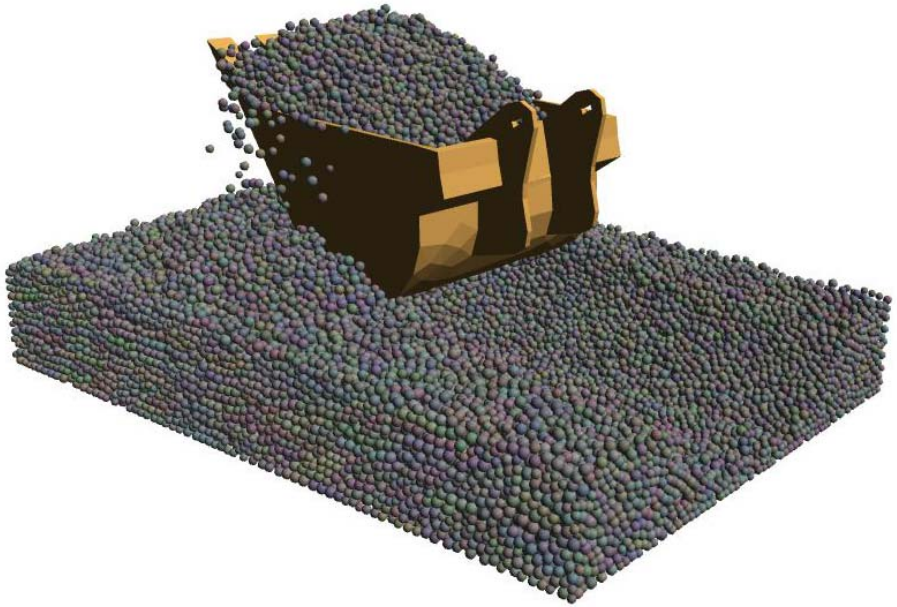

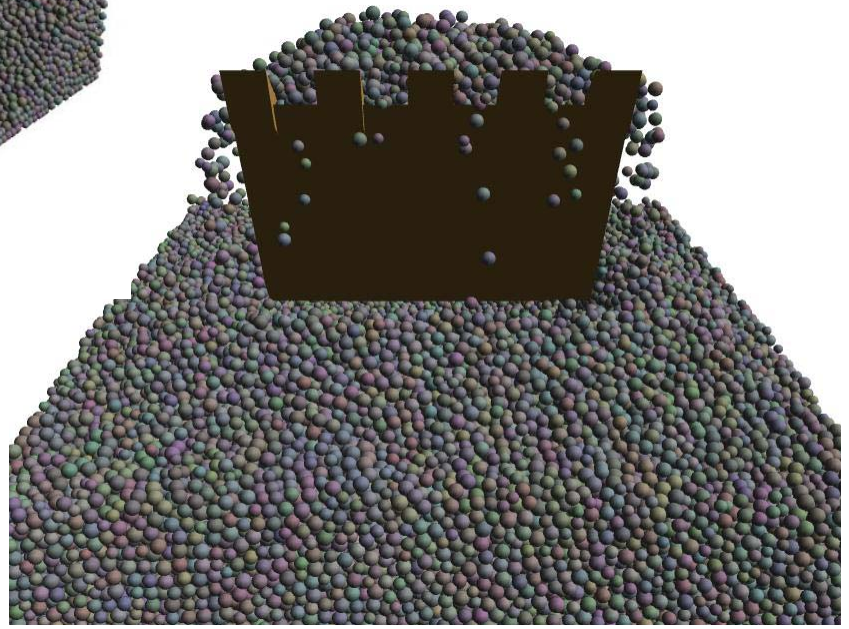

Metariver Technology

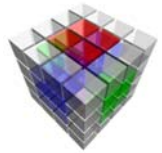

Metariver Technology

Example-1：Excavator

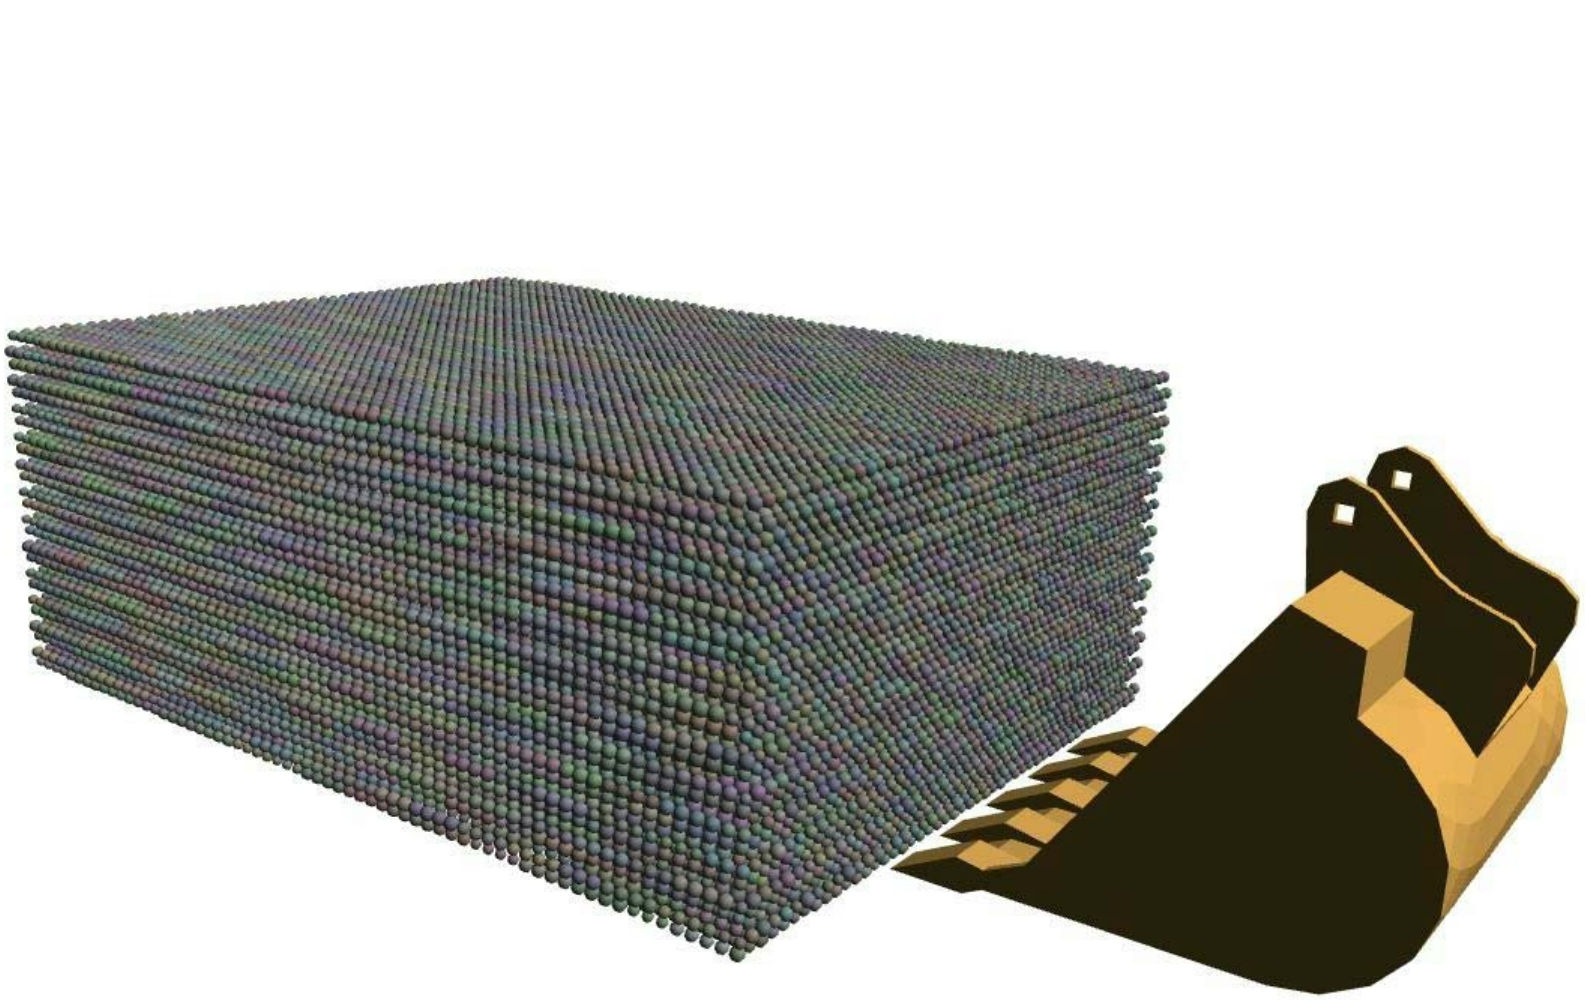

Metariver Technology

Example-2 : Agitator

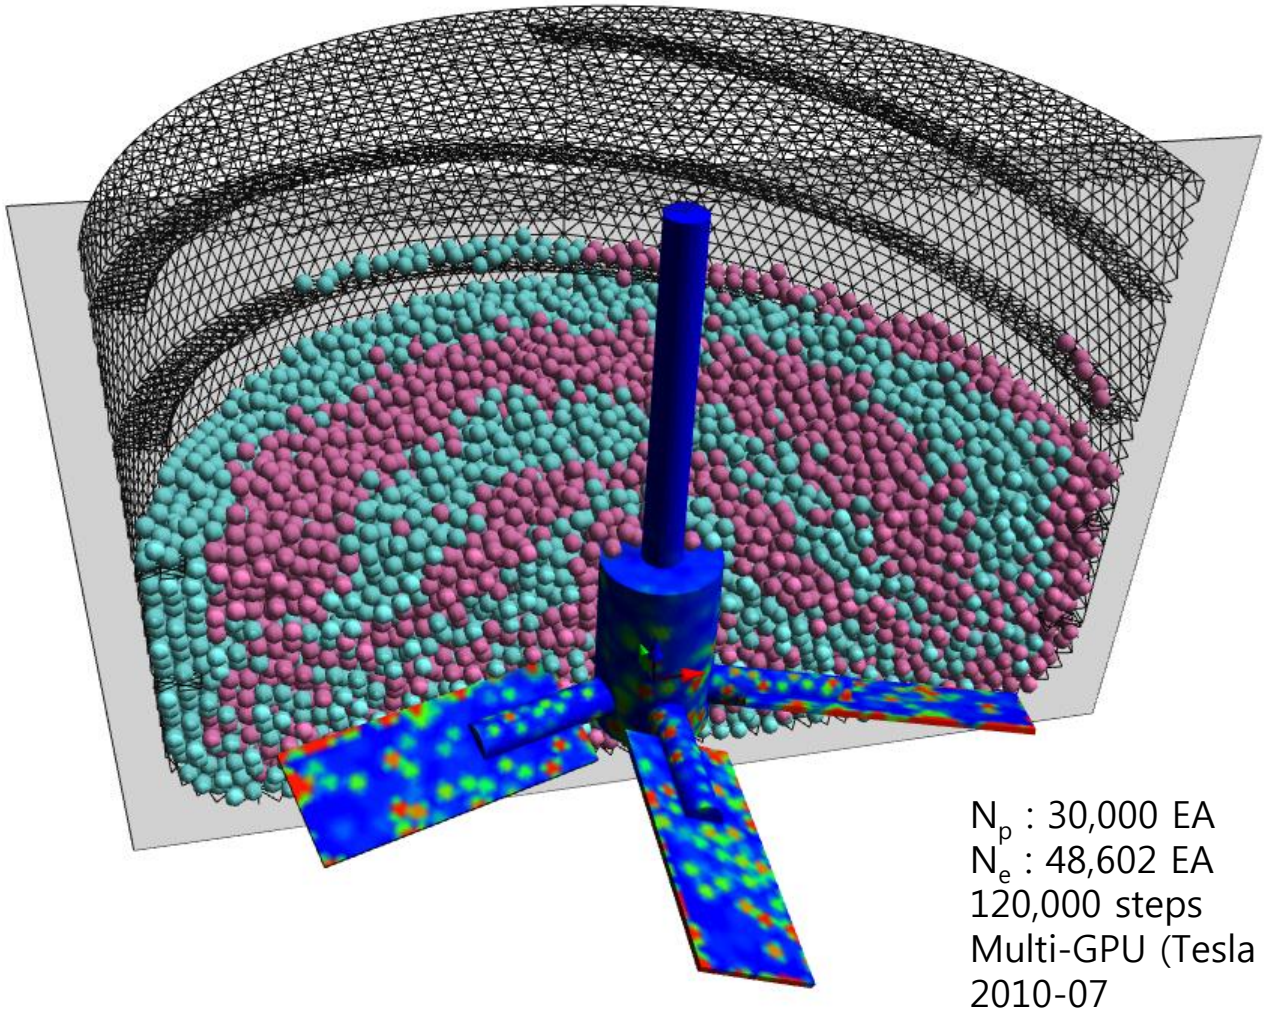

N_p : 30,000 EA
N_e : 48,602 EA
120,000 steps
Multi-GPU (Tesla C1060 x2)
2010-07

Example-2 : Agitator

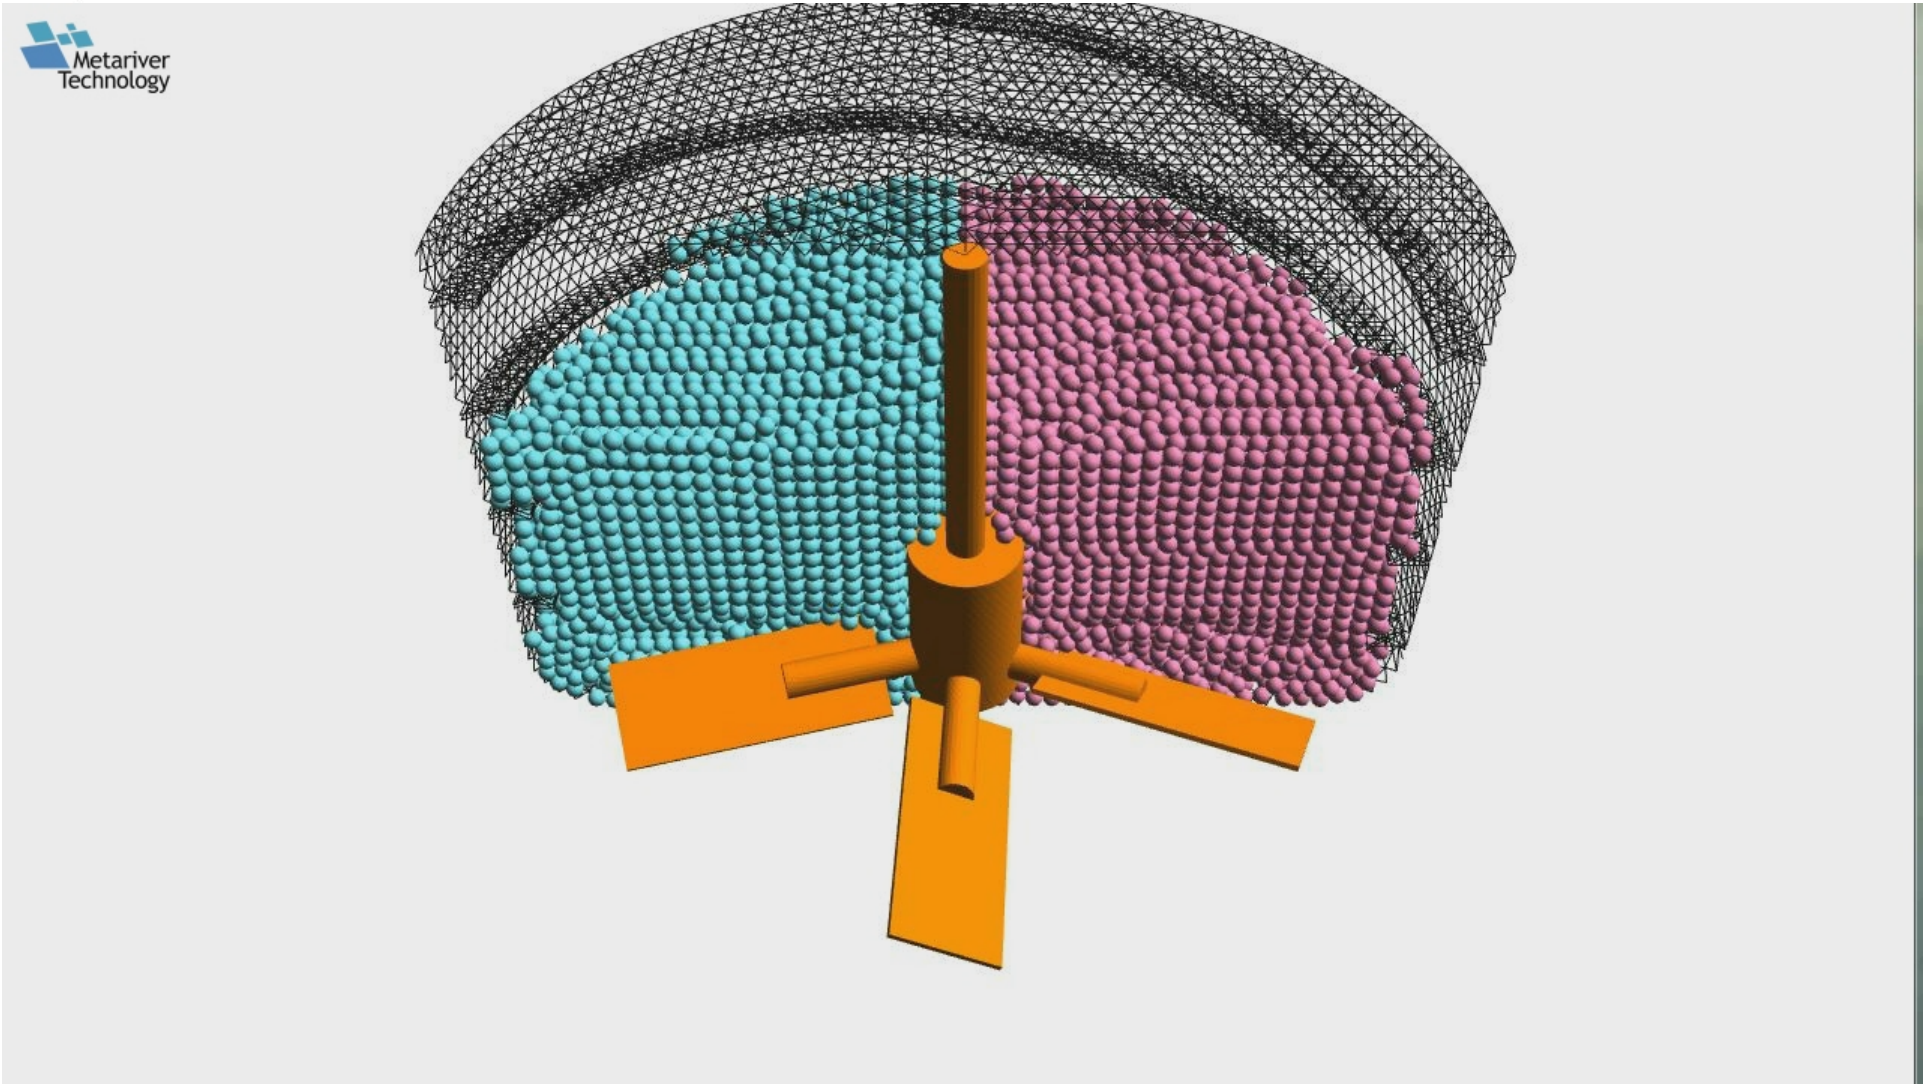

Example-3 : Blast Furnace

N_p : 640,000 EA
N_e : 64,314 EA
250,000 steps
Multi-GPU (Tesla C1060 x2)
2010-07

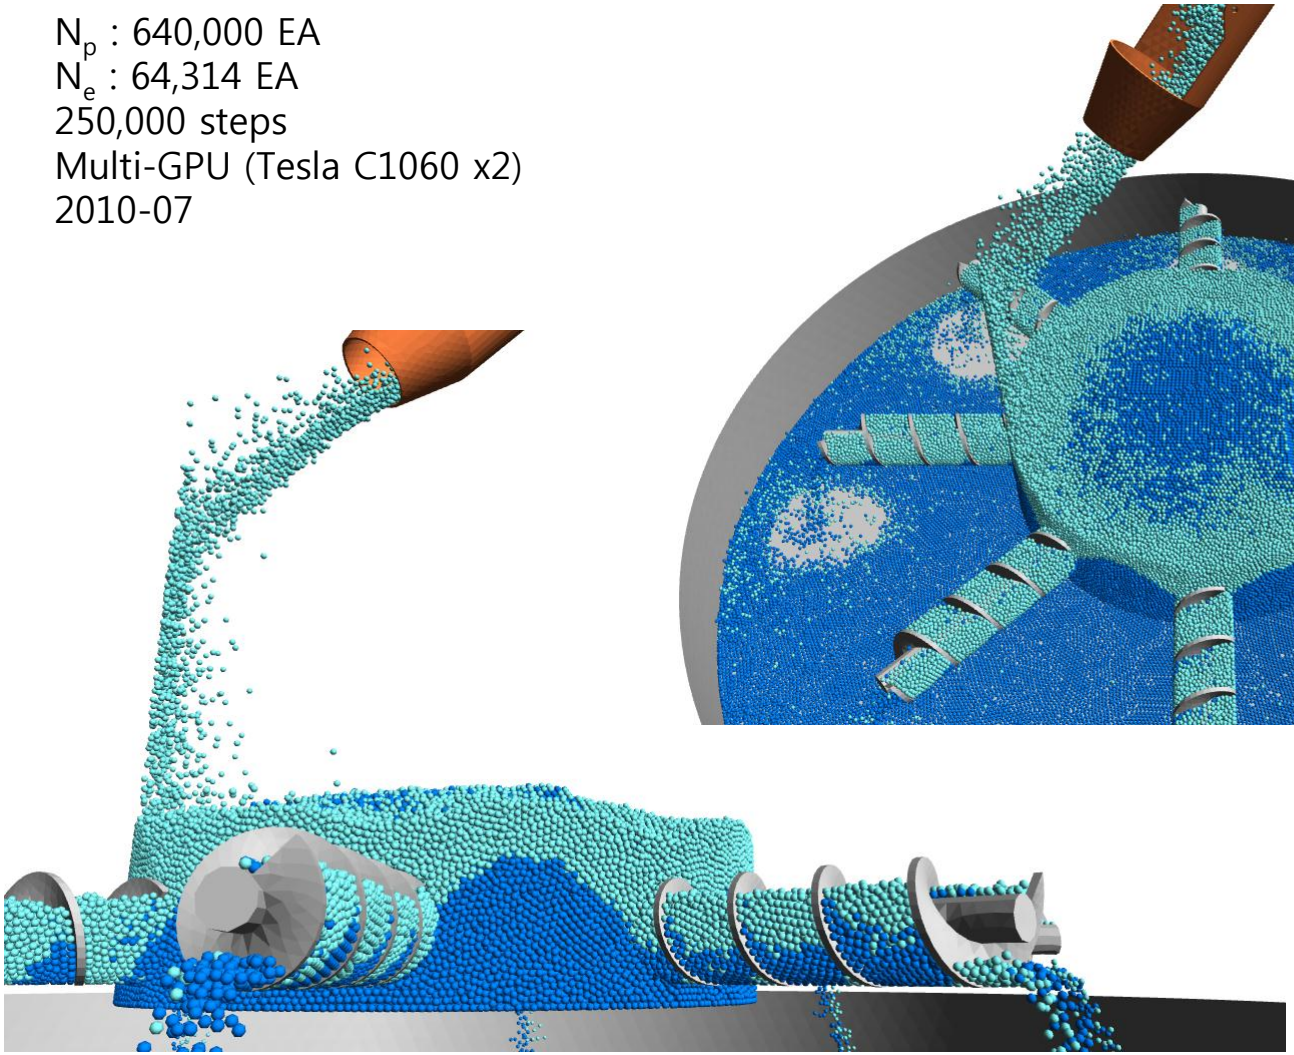

Metariver Technology

Example-3 : Blast Furnace

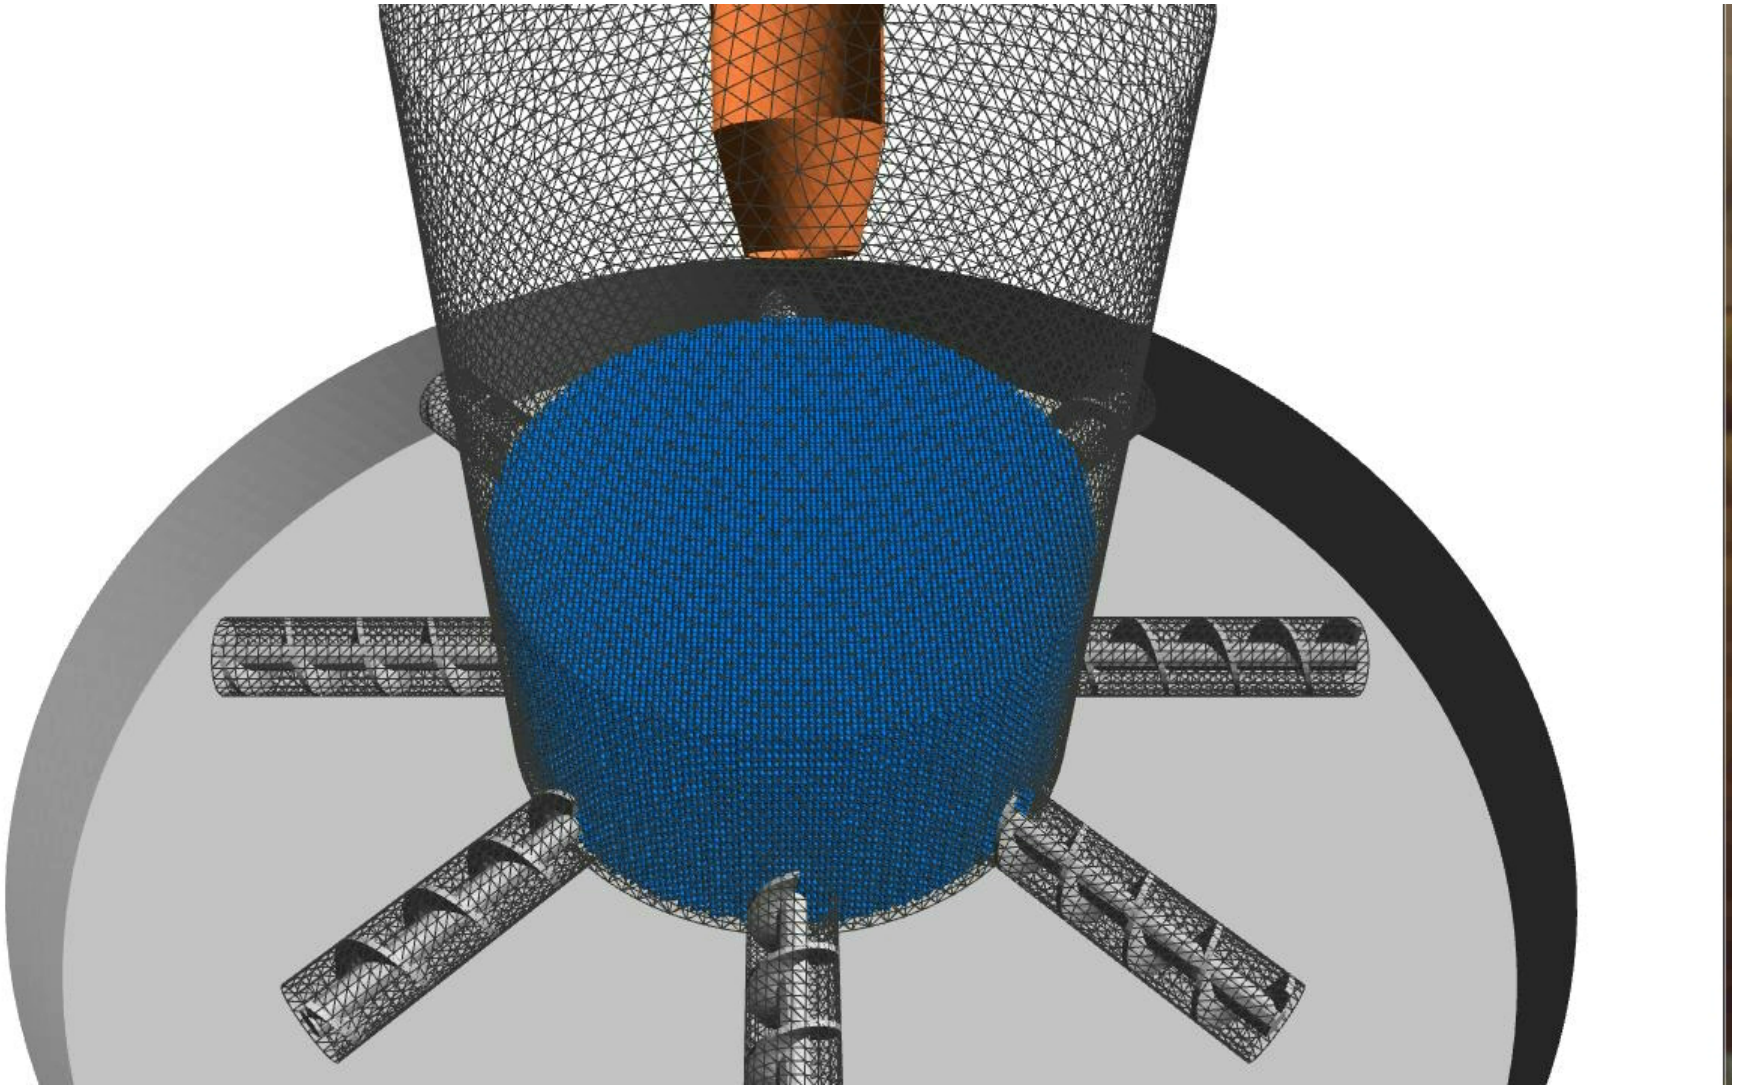

Metariver Technology

Visualization

: HyperCube4 & EnSight

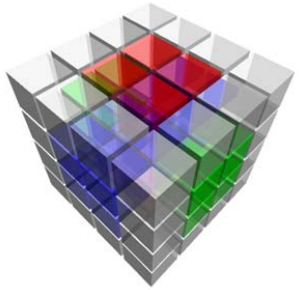

Sloshing

N_p : 248,845 EA
N_e : 12,580 EA
1,000,000 steps
Tesla C2050

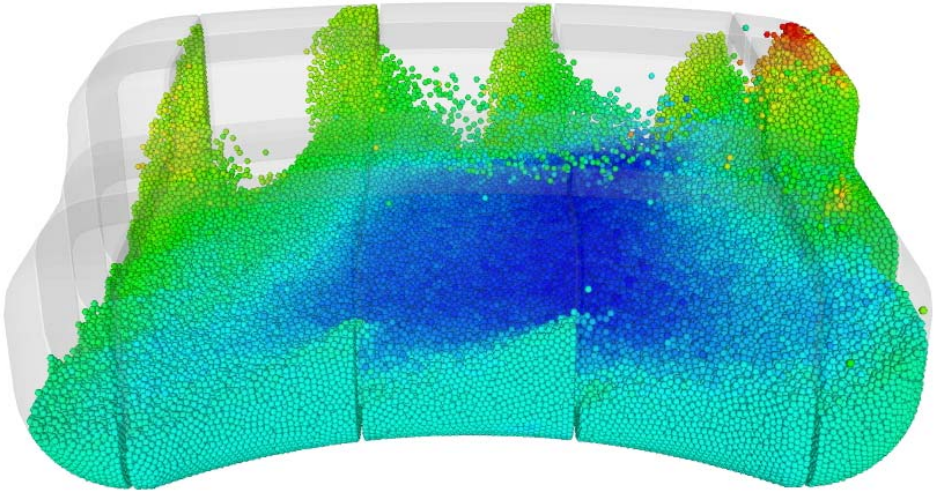

Sloshing

N_p : 248,845 EA
N_e : 12,580 EA
1,000,000 steps
Tesla C2050

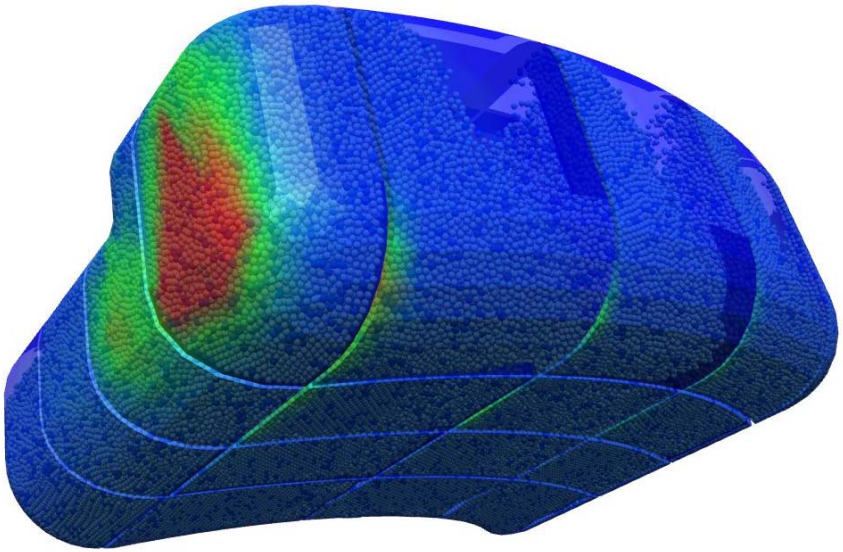

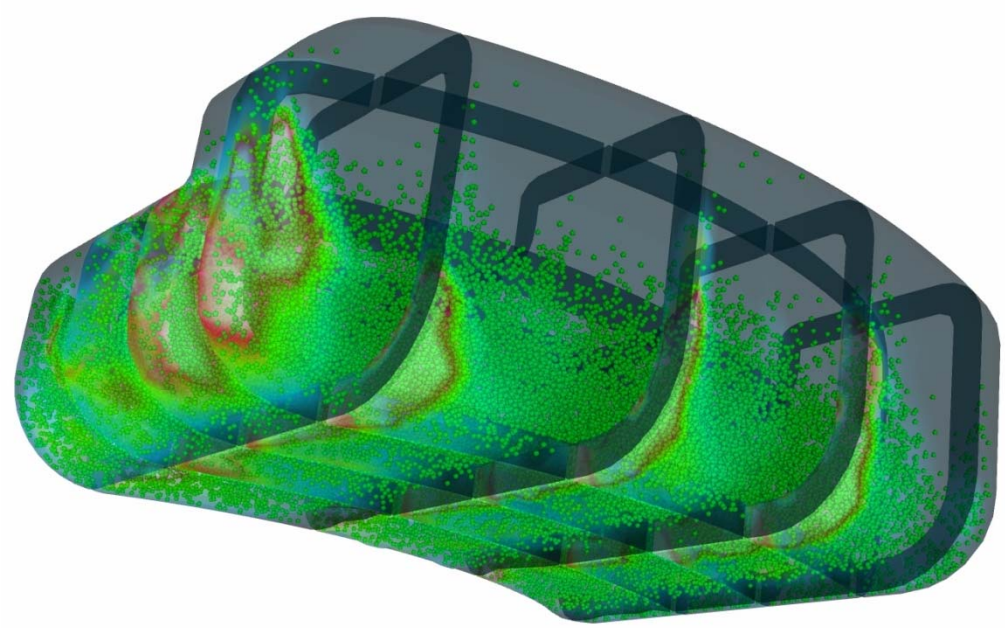

Sloshing

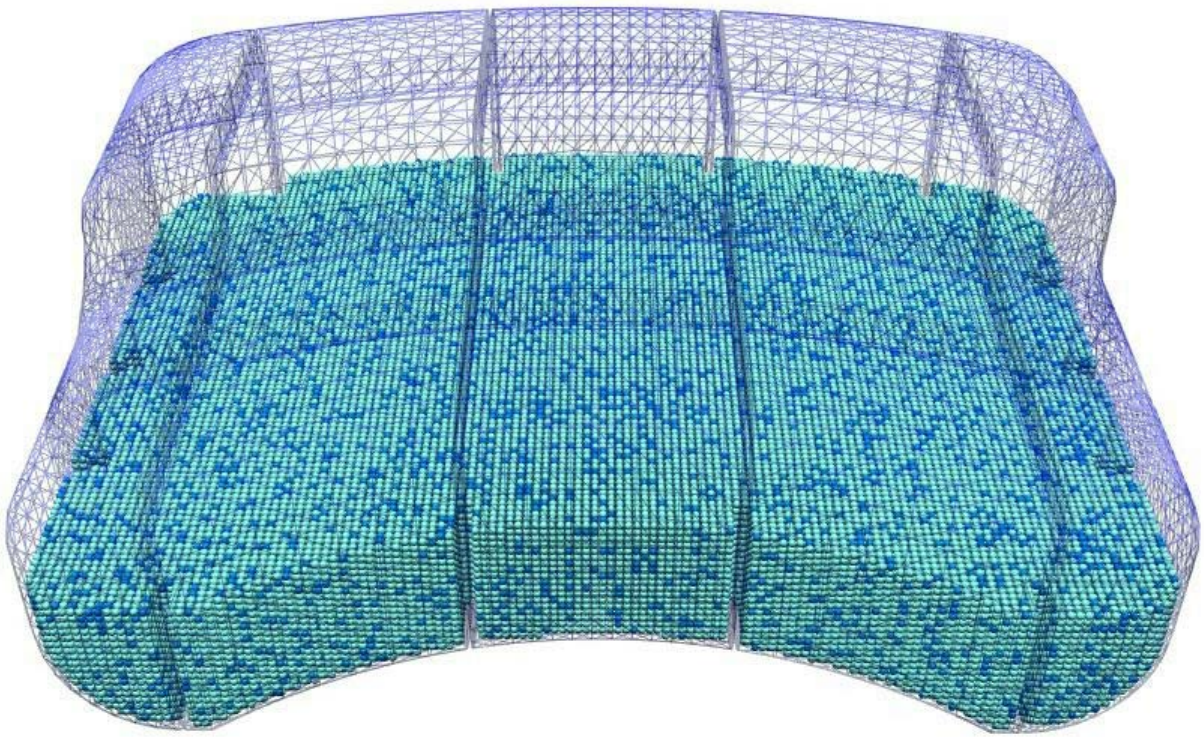

Sloshing

N_p : 248,845 EA
N_e : 12,580 EA
1,000,000 steps
Tesla C2050

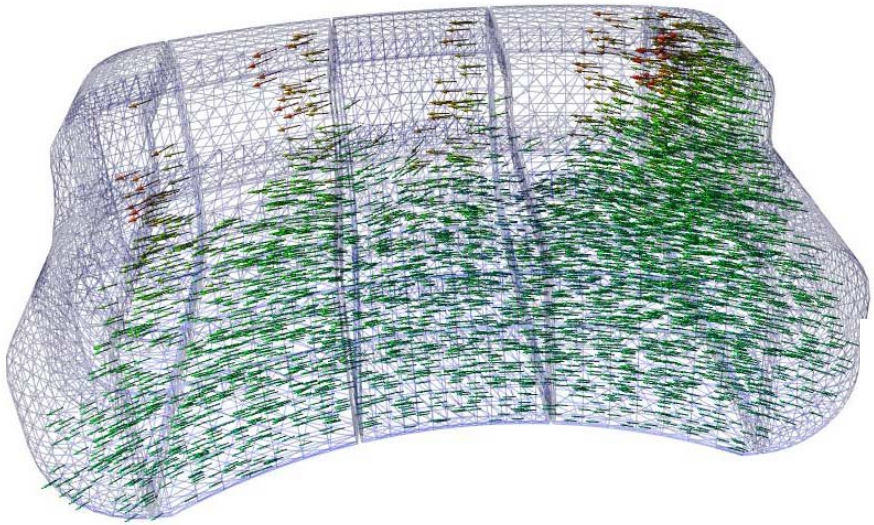

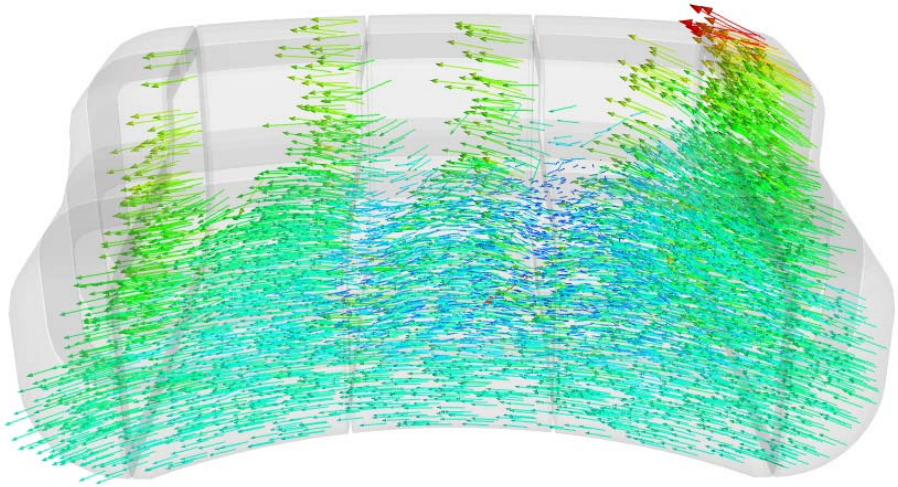

< Only Vector >

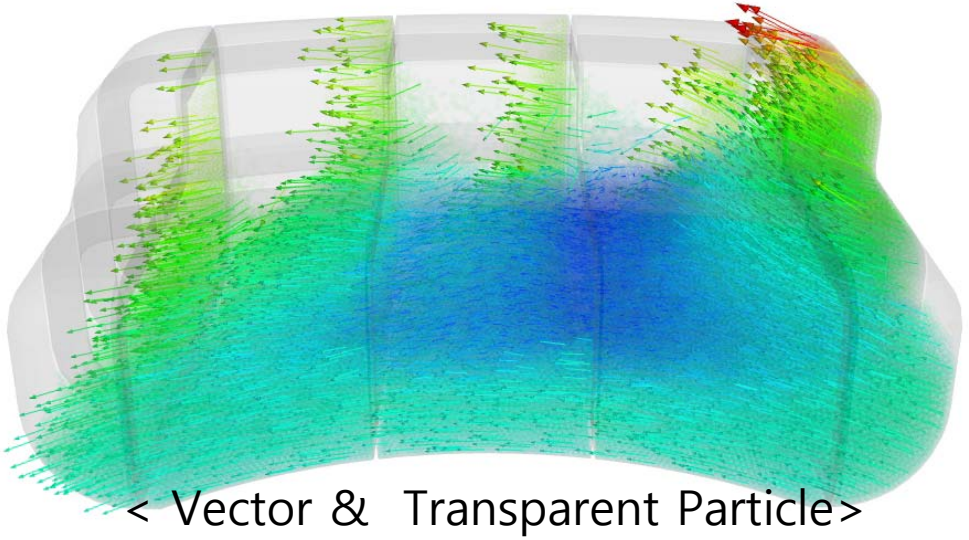

< Vector & Transparent Particle>

Sloshing

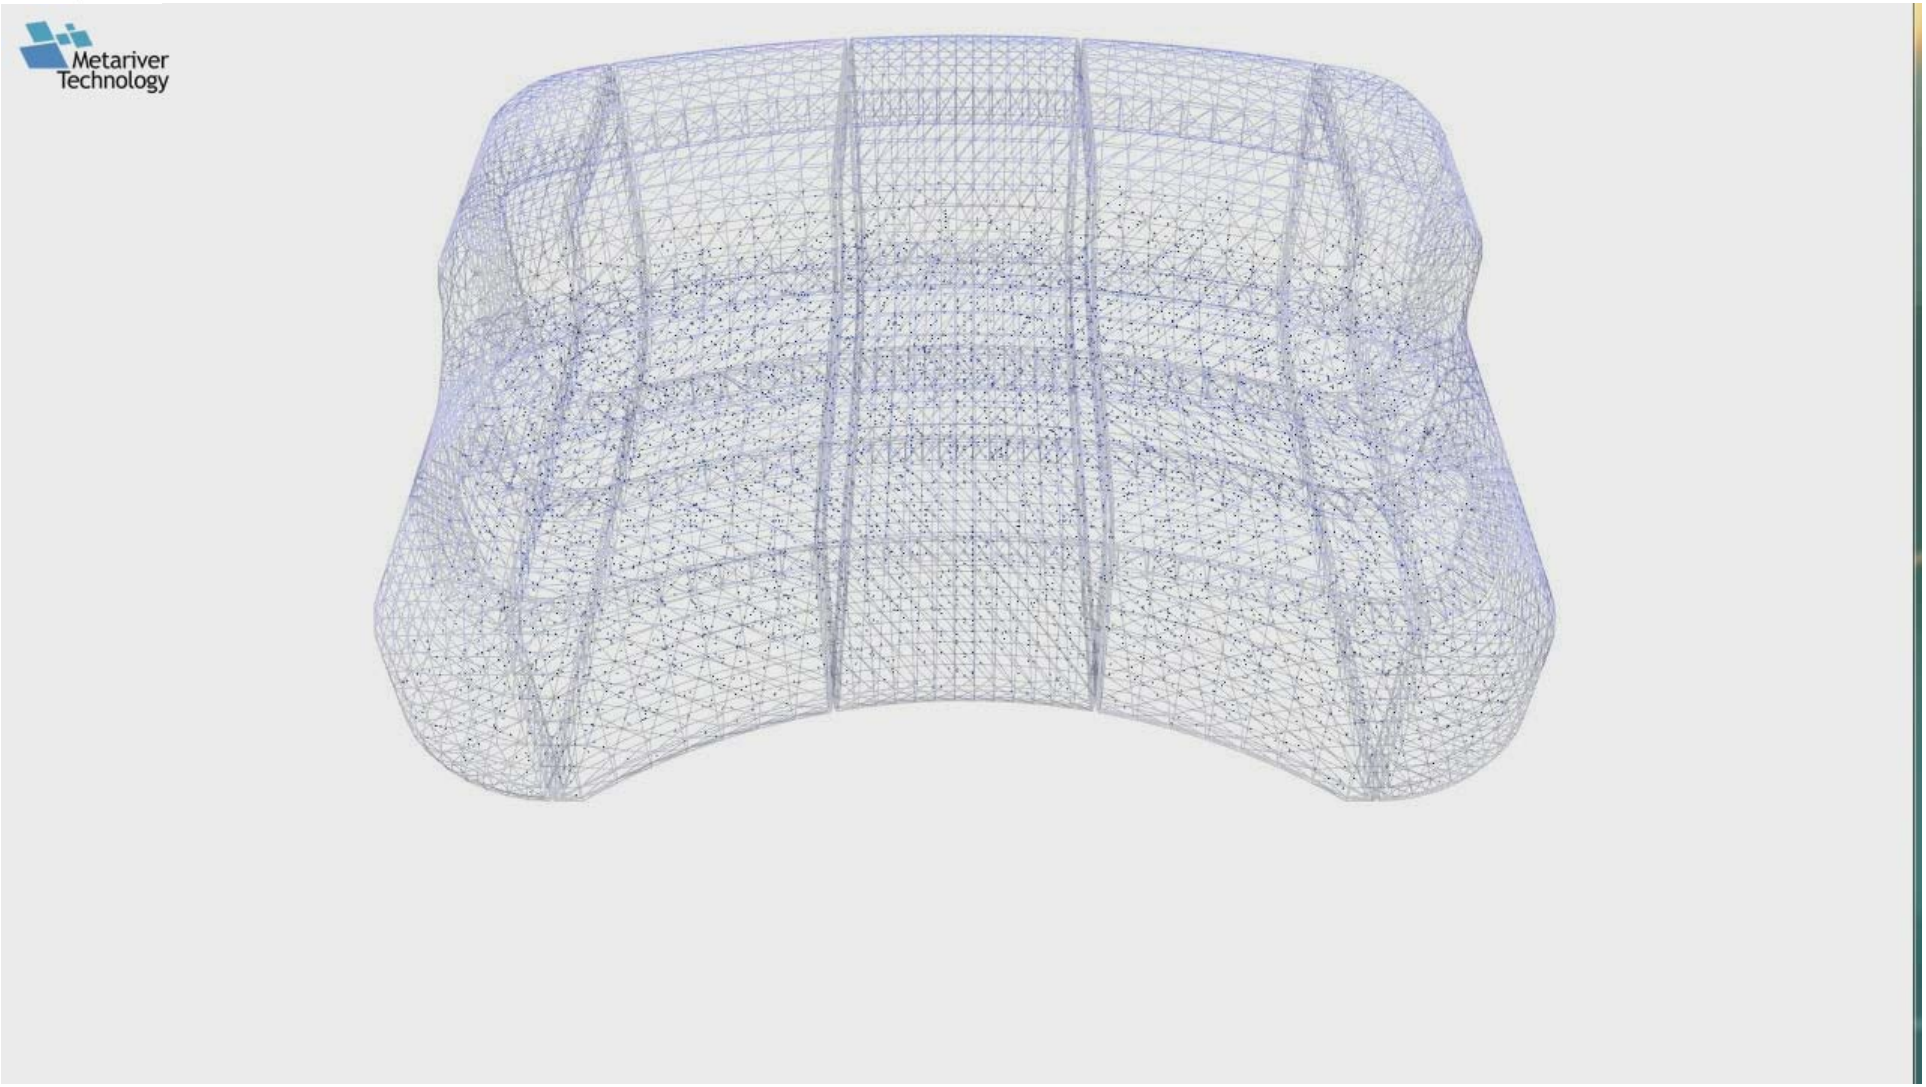

2-R Auger

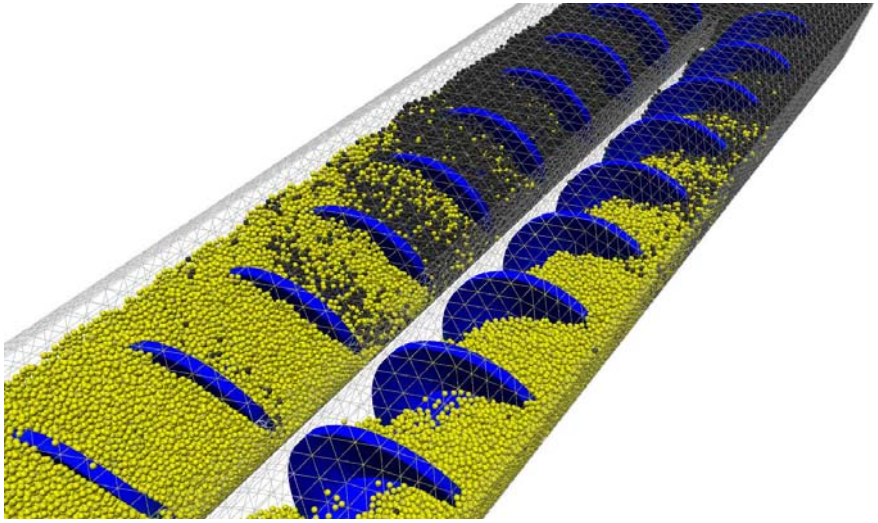

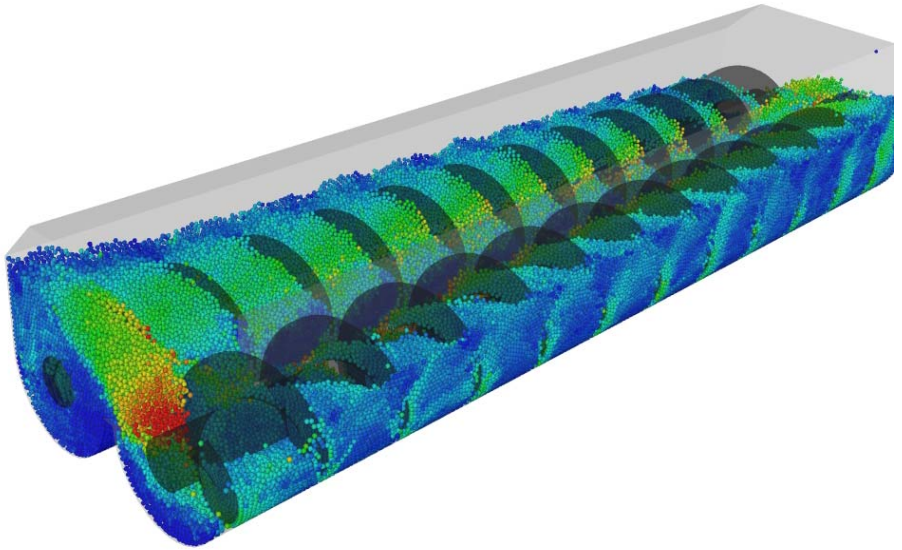

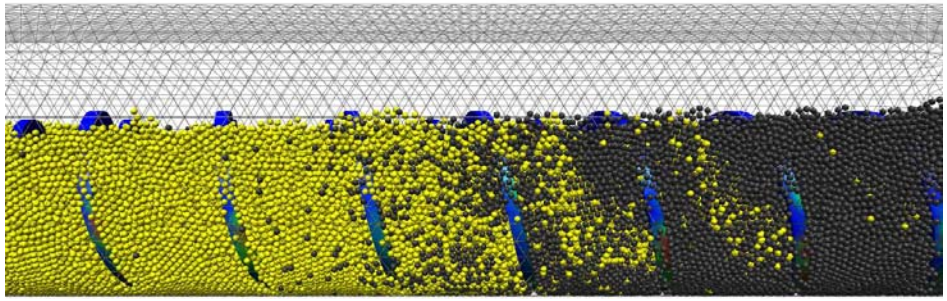

N_p : 219,391 EA
N_e : 45,282 EA
700,000 steps
Multi-GPU (Tesla C1060 x2)

2-R Auger

N_p : 219,391 EA
N_e : 45,282 EA
700,000 steps
Multi-GPU (Tesla C1060 x2)

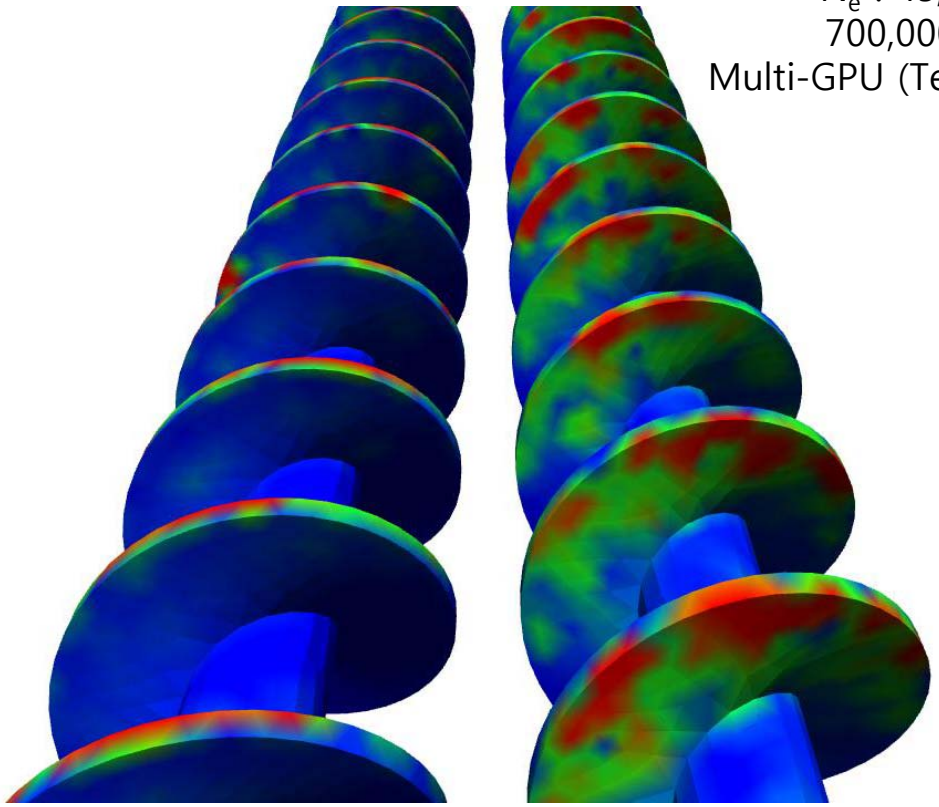

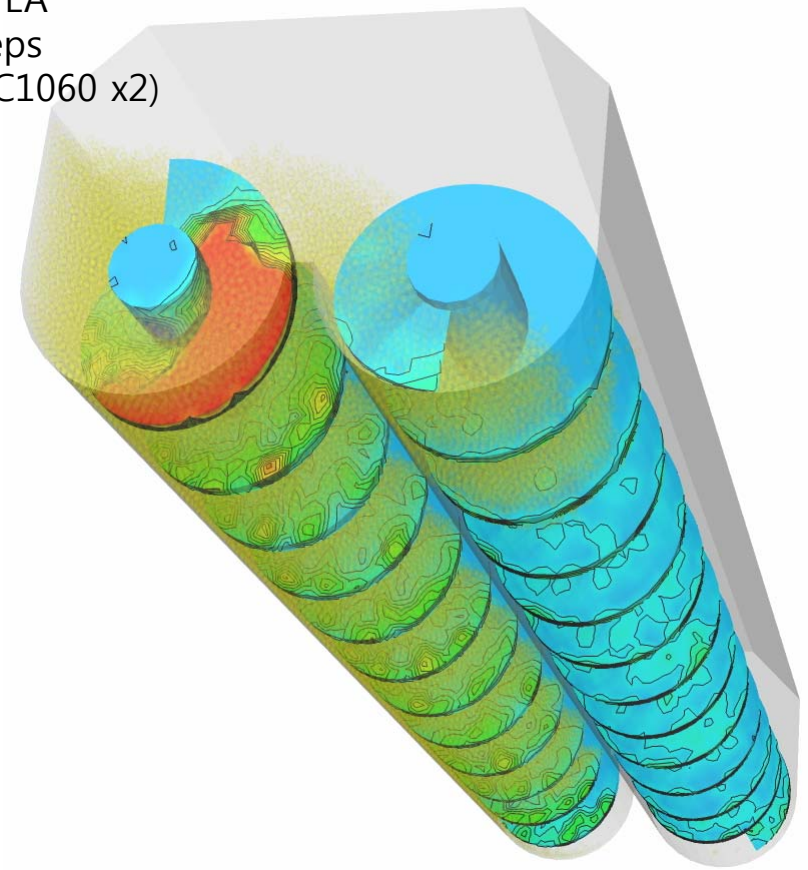

2-R Auger

N_p : 219,391 EA
N_e : 45,282 EA
700,000 steps
Multi-GPU (Tesla C1060 x2)

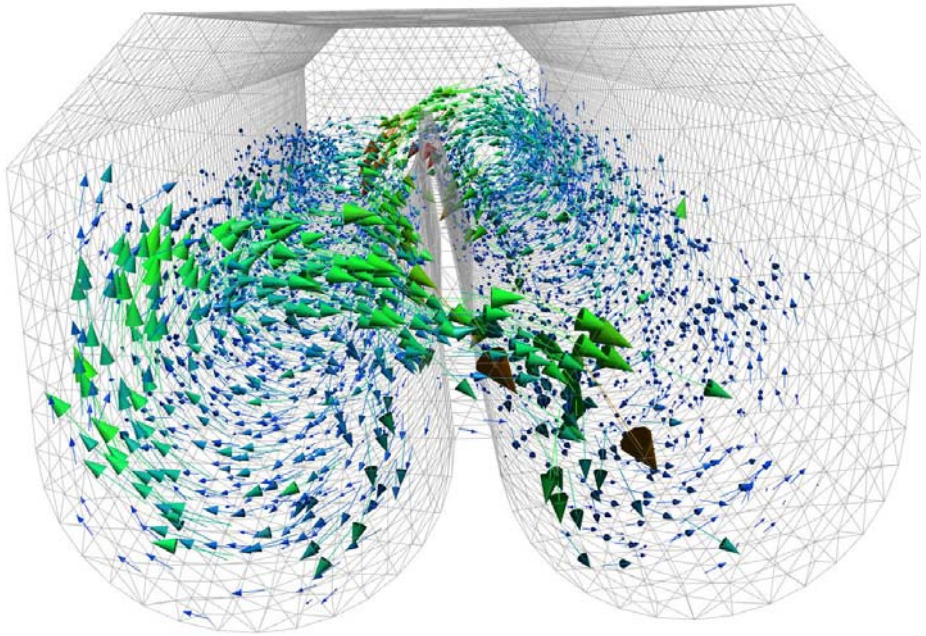

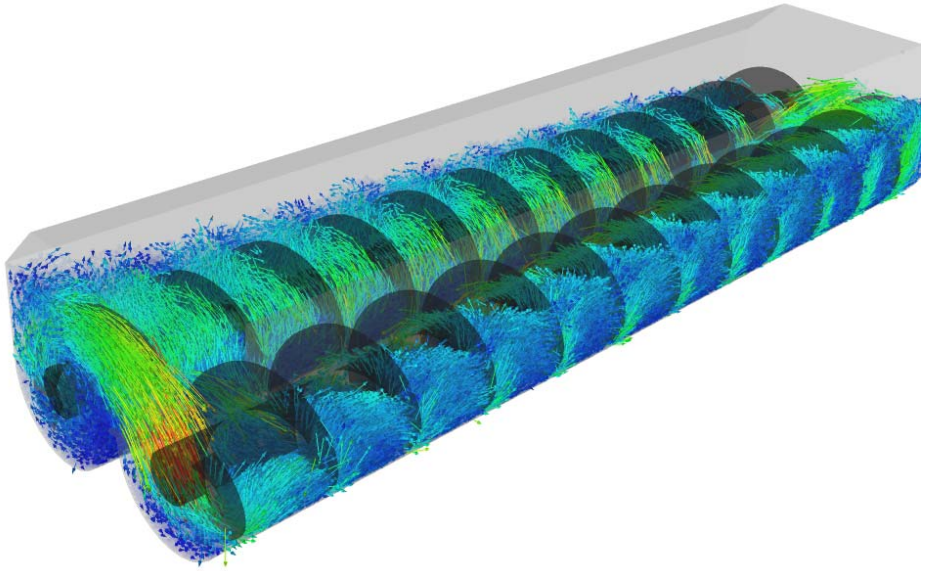

2-R Auger

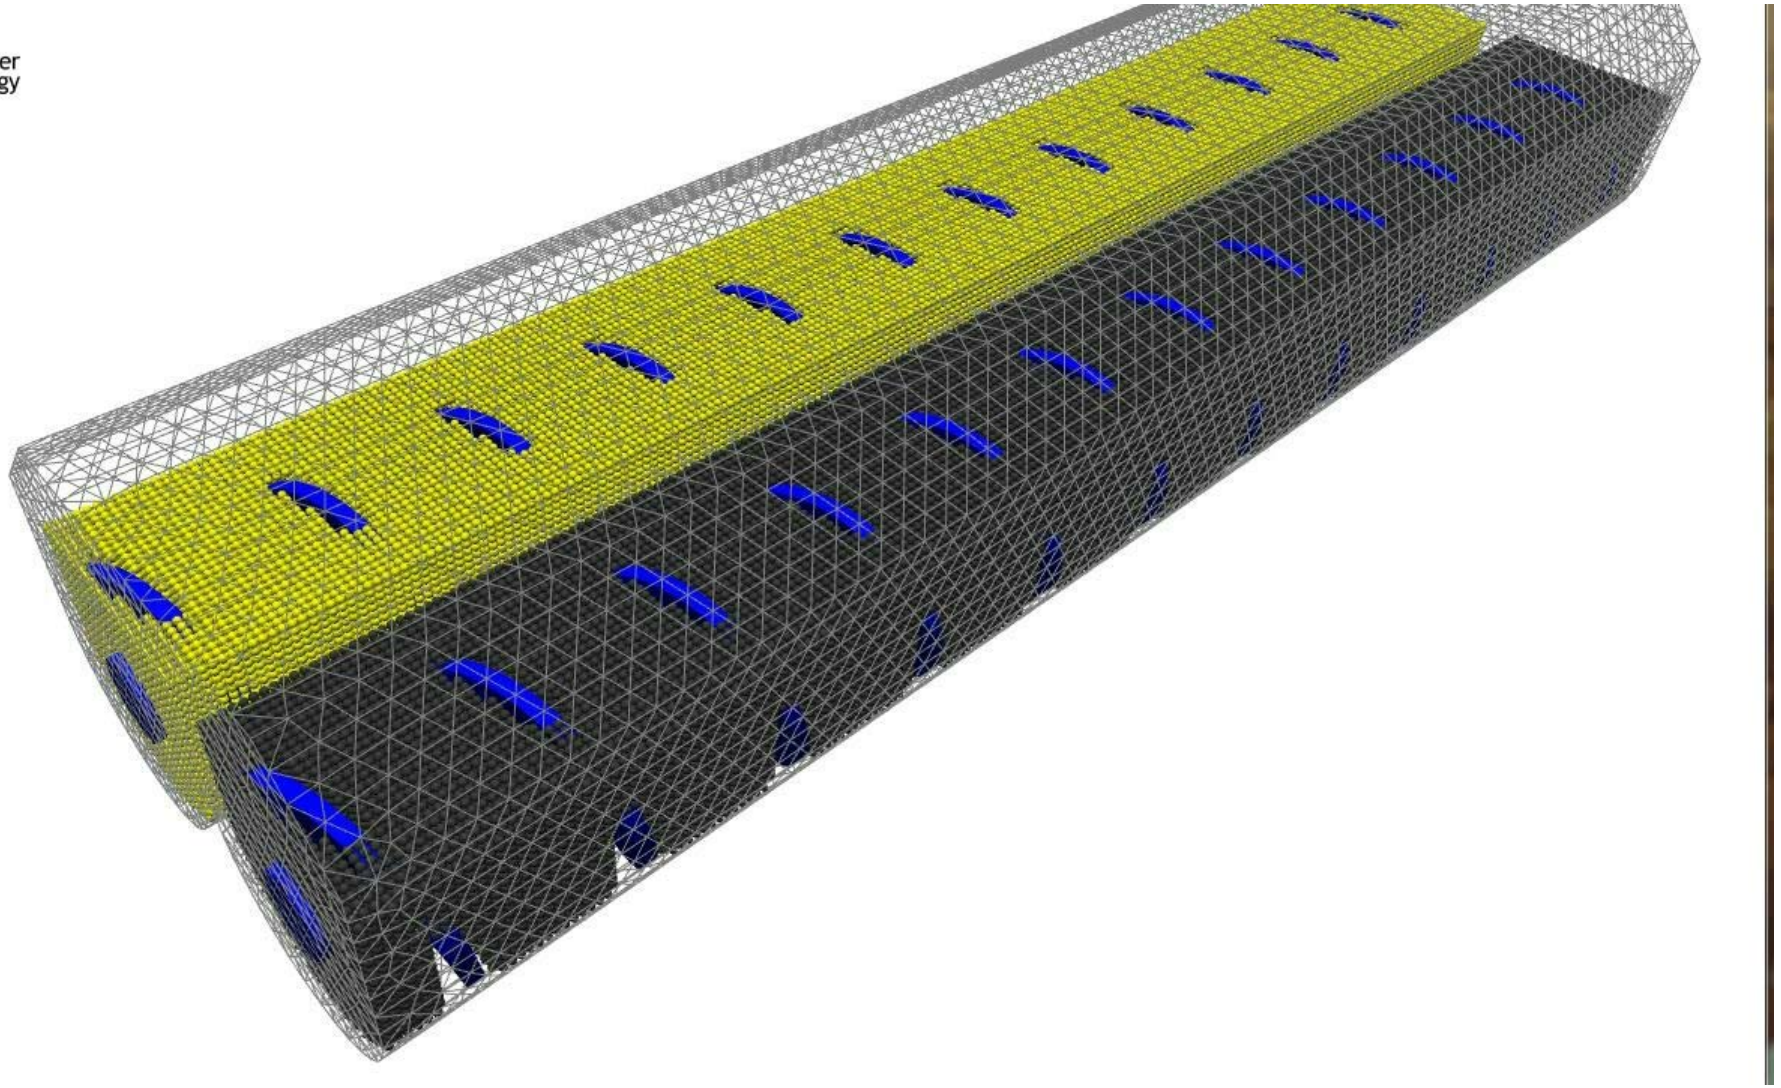

Rotating Drum

N_p : 124,266 EA
N_e : 6,048 EA
700,000 steps
Tesla C2050

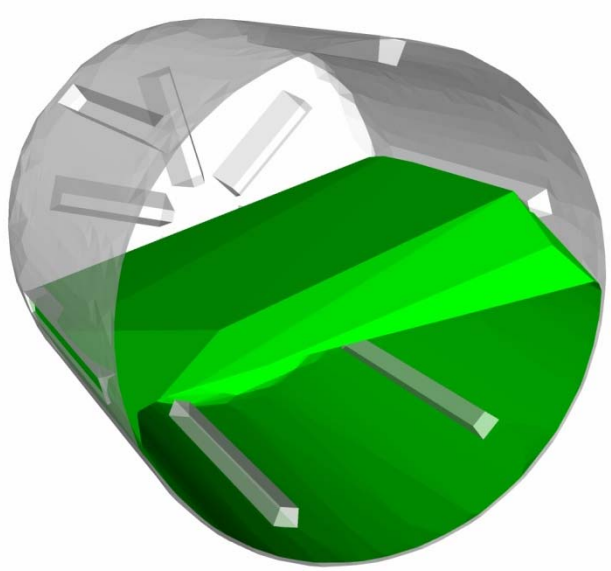

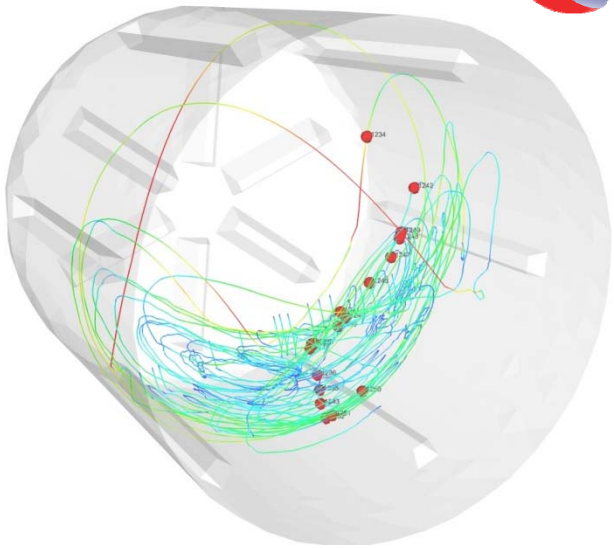

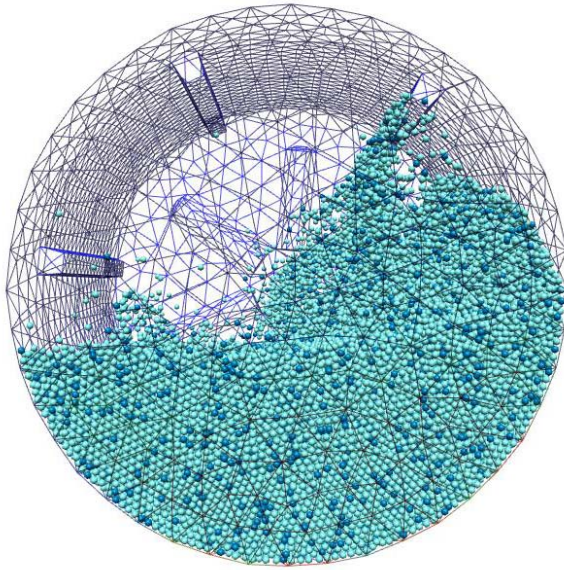

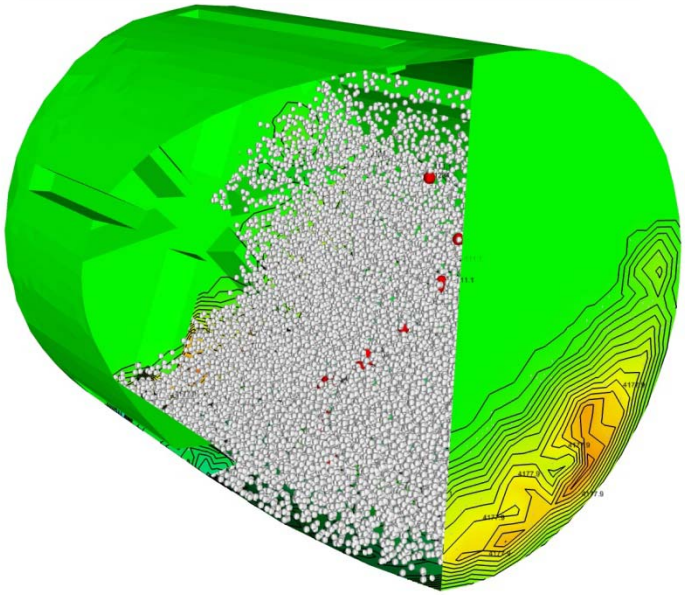

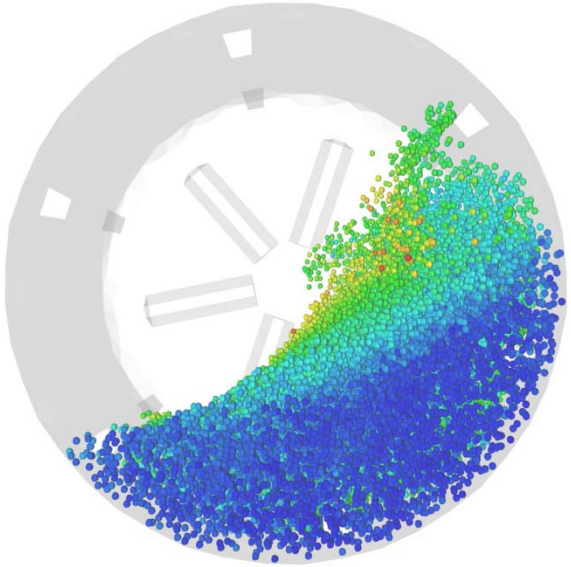

Metariver Technology

Rotating Drum

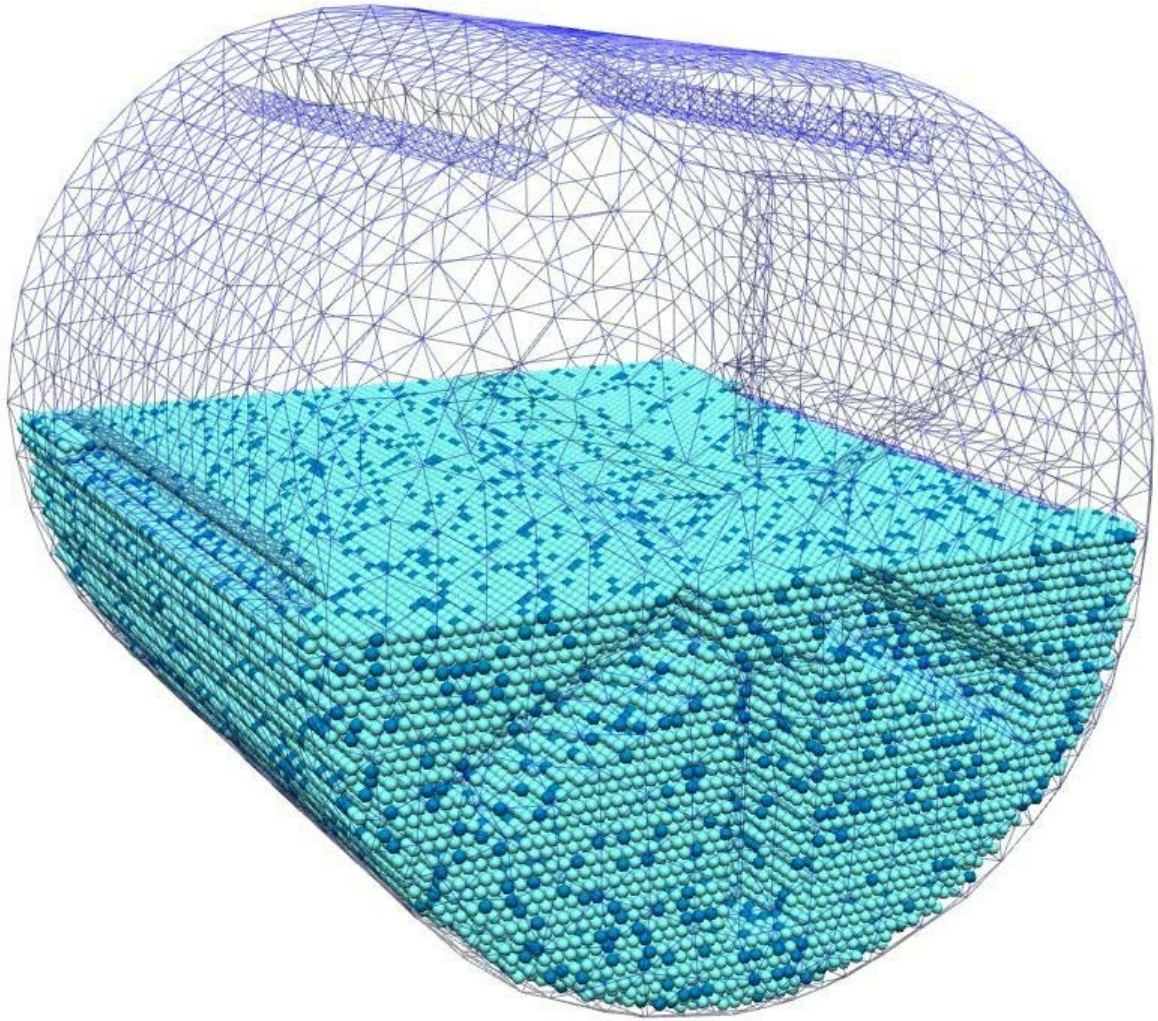

Metariver Technology

Rotating Drum

N_p : 124,266 EA
N_e : 6,048 EA
700,000 steps
Tesla C2050

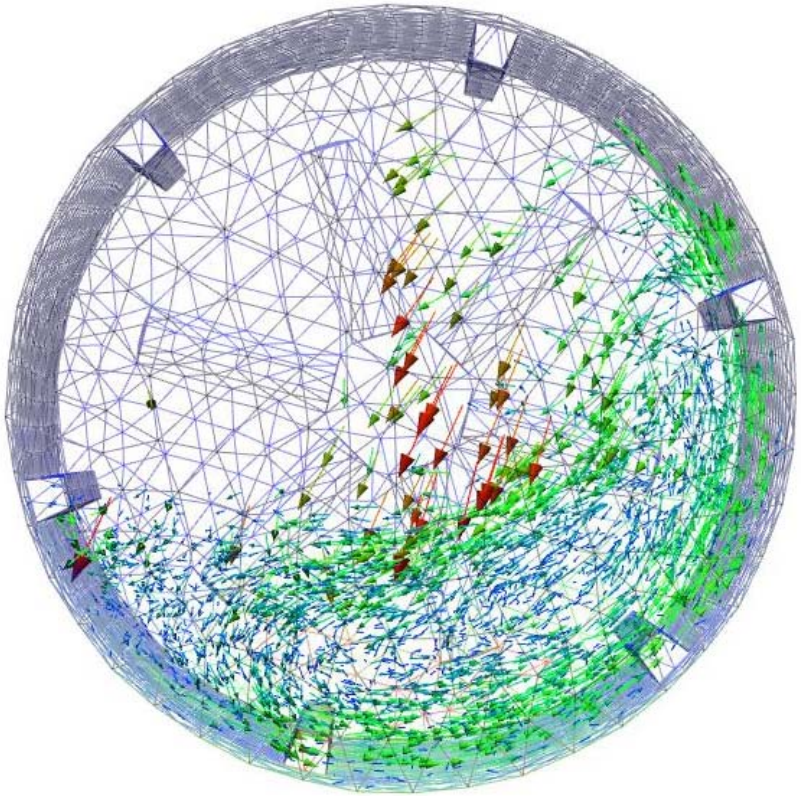

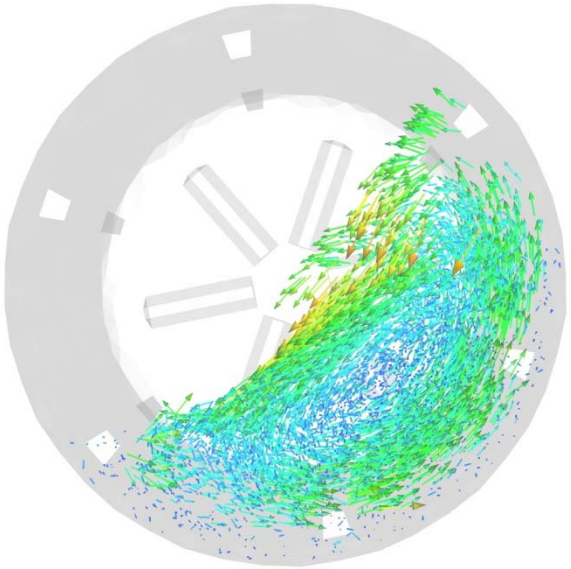

< Only Vector >

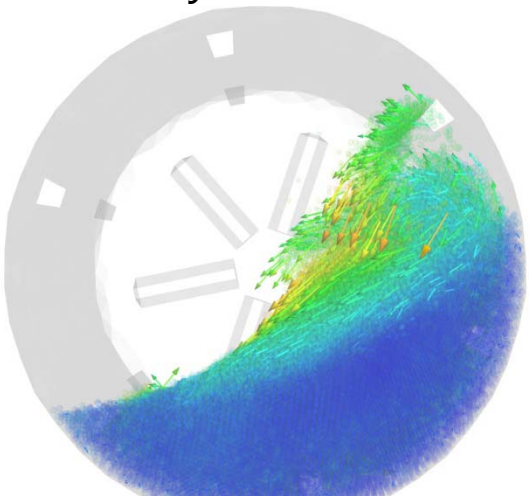

< Vector & Transparent Particle>

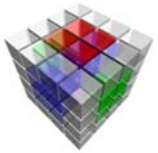

Metariver Technology

Metariver Technology

Rotating Drum