


Tsunami Simulation on GPUs

Takayuki AOKI

*Global Scientific Information and Computing Center
Tokyo Institute of Technology*

Valdivia Earthquake, 1960




The Biggest earthquake:
several times M7~M8

142 died in Japan 1743 died in Chile




USGS ShakeMap : Concepcion, Chile
Sun May 22, 1960 19:11:17 GMT M 9.5 S38.23 W73.05 Depth: 35.0km ID:196005221911




Map Version 1.1 Processed Sun Nov 9, 2008 09:23:34 AM MST

PERCEIVED SHAKING	Not felt	Weak	Light	Moderate	Strong	Very strong	Severe	Violent	Extreme
POTENTIAL DAMAGE	none	none	none	Very light	Light	Moderate	Moderate/Heavy	Heavy	Very Heavy
PEAK ACC (%g)	<.17	.17-1.4	1.4-3.9	3.9-9.2	9.2-18	18-34	34-65	65-124	>124
PEAK VEL (cm/s)	<0.1	0.1-1.1	1.1-3.4	3.4-8.1	8.1-16	16-31	31-60	60-116	>116
INSTRUMENTAL INTENSITY	I	II-III	IV	V	VI	VII	VIII	IX	X+

TSUNAMI Disaster




Real-time TSUNAMI Simulation



ADPC : Asian Disaster Preparedness Center

Early Warning System:

Data Base



Real-time CFD

high accuracy



Shallow-Water Eq.

Conservative Form:

Assuming

hydrostatic balance

in the vertical direction,

3D → 2D equation

$$\frac{\partial h}{\partial t} + \frac{\partial hu}{\partial x} + \frac{\partial hv}{\partial y} = 0$$

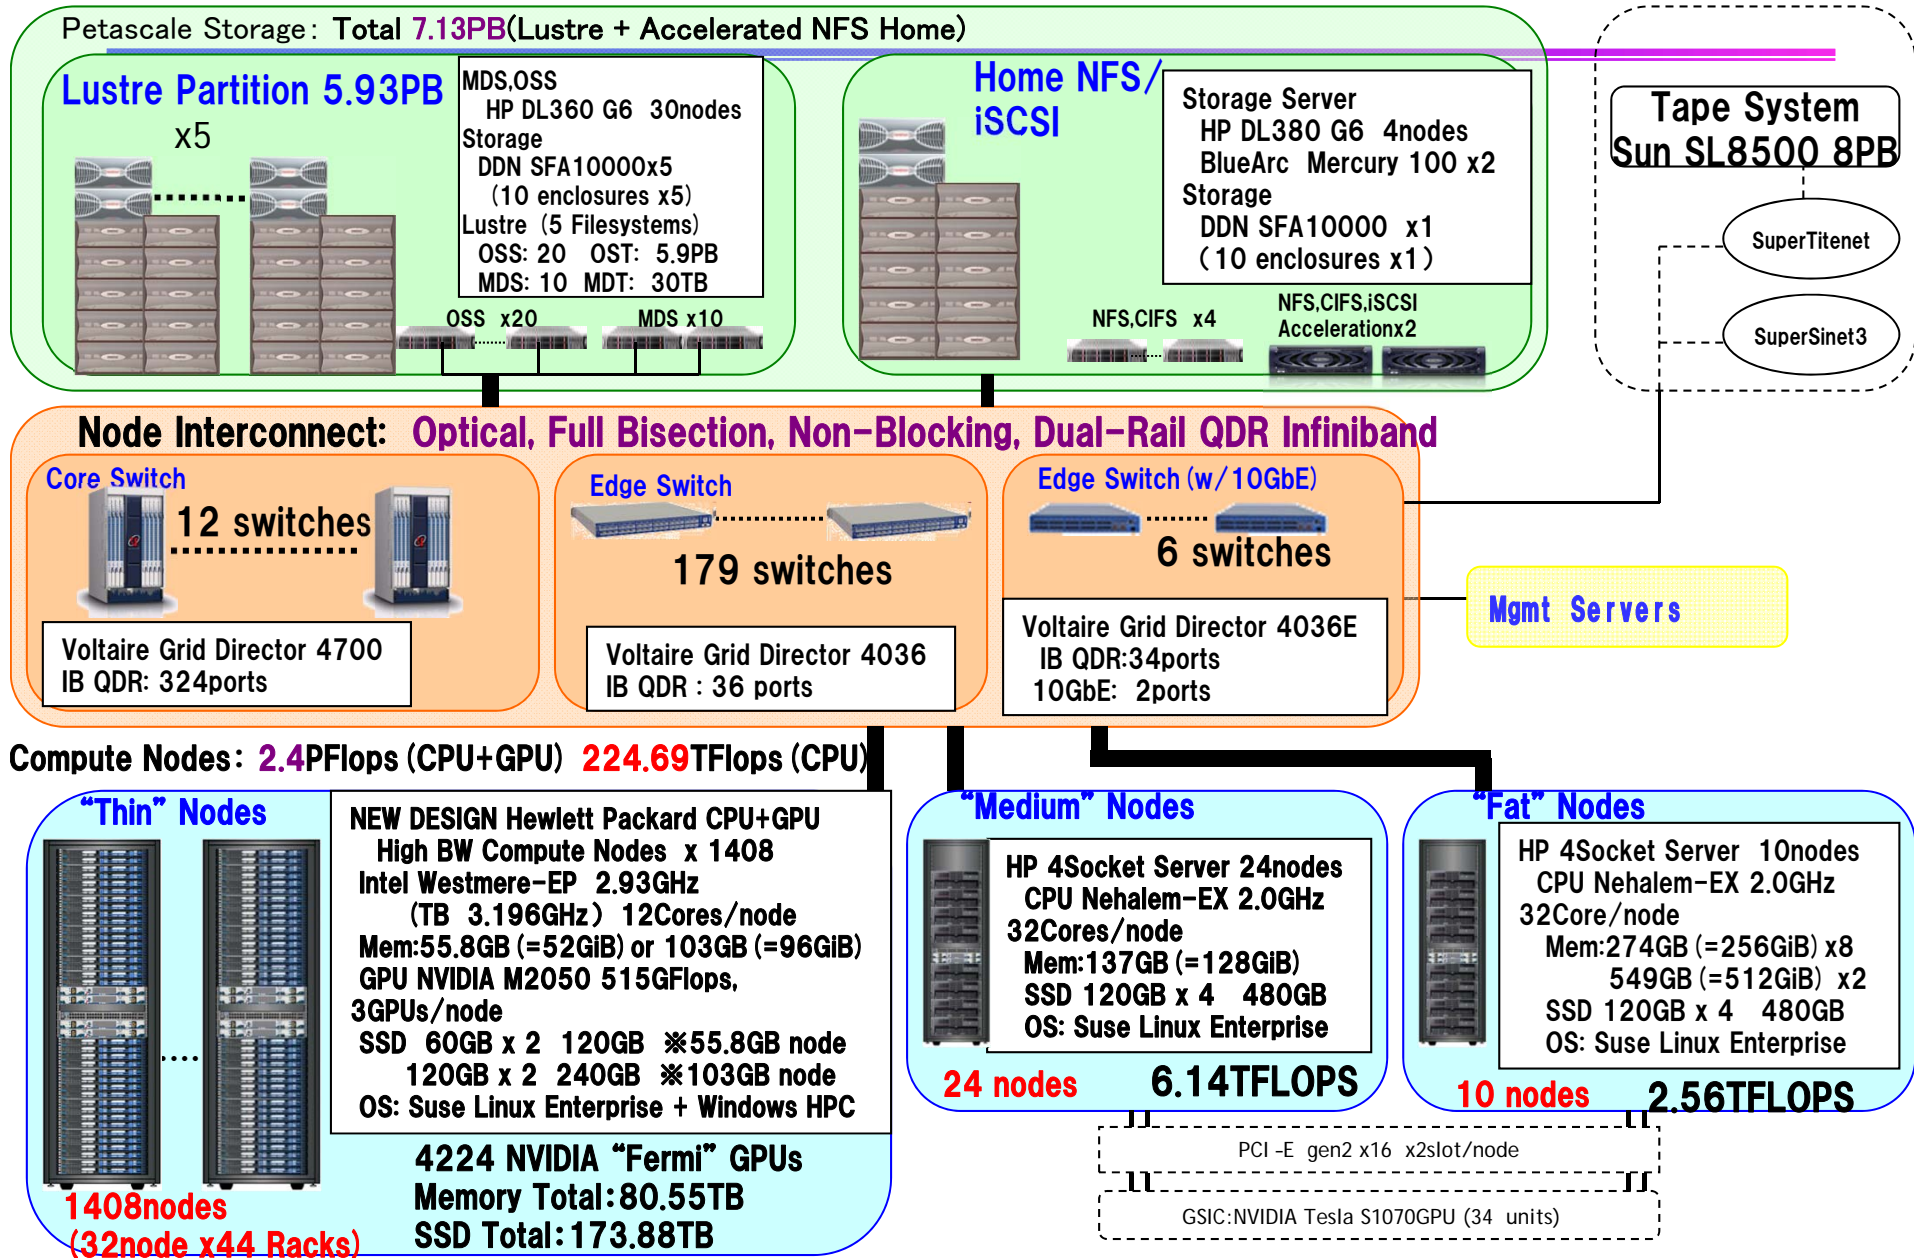
$$\frac{\partial hu}{\partial t} + \frac{\partial}{\partial x} \left(hu^2 + \frac{1}{2} gh^2 \right) + \frac{\partial huv}{\partial y} = -gh \frac{\partial z}{\partial x}$$

$$\frac{\partial hv}{\partial t} + \frac{\partial huv}{\partial x} + \frac{\partial}{\partial y} \left(hv^2 + \frac{1}{2} gh^2 \right) = -gh \frac{\partial z}{\partial y}$$

TSUBAME 2.0

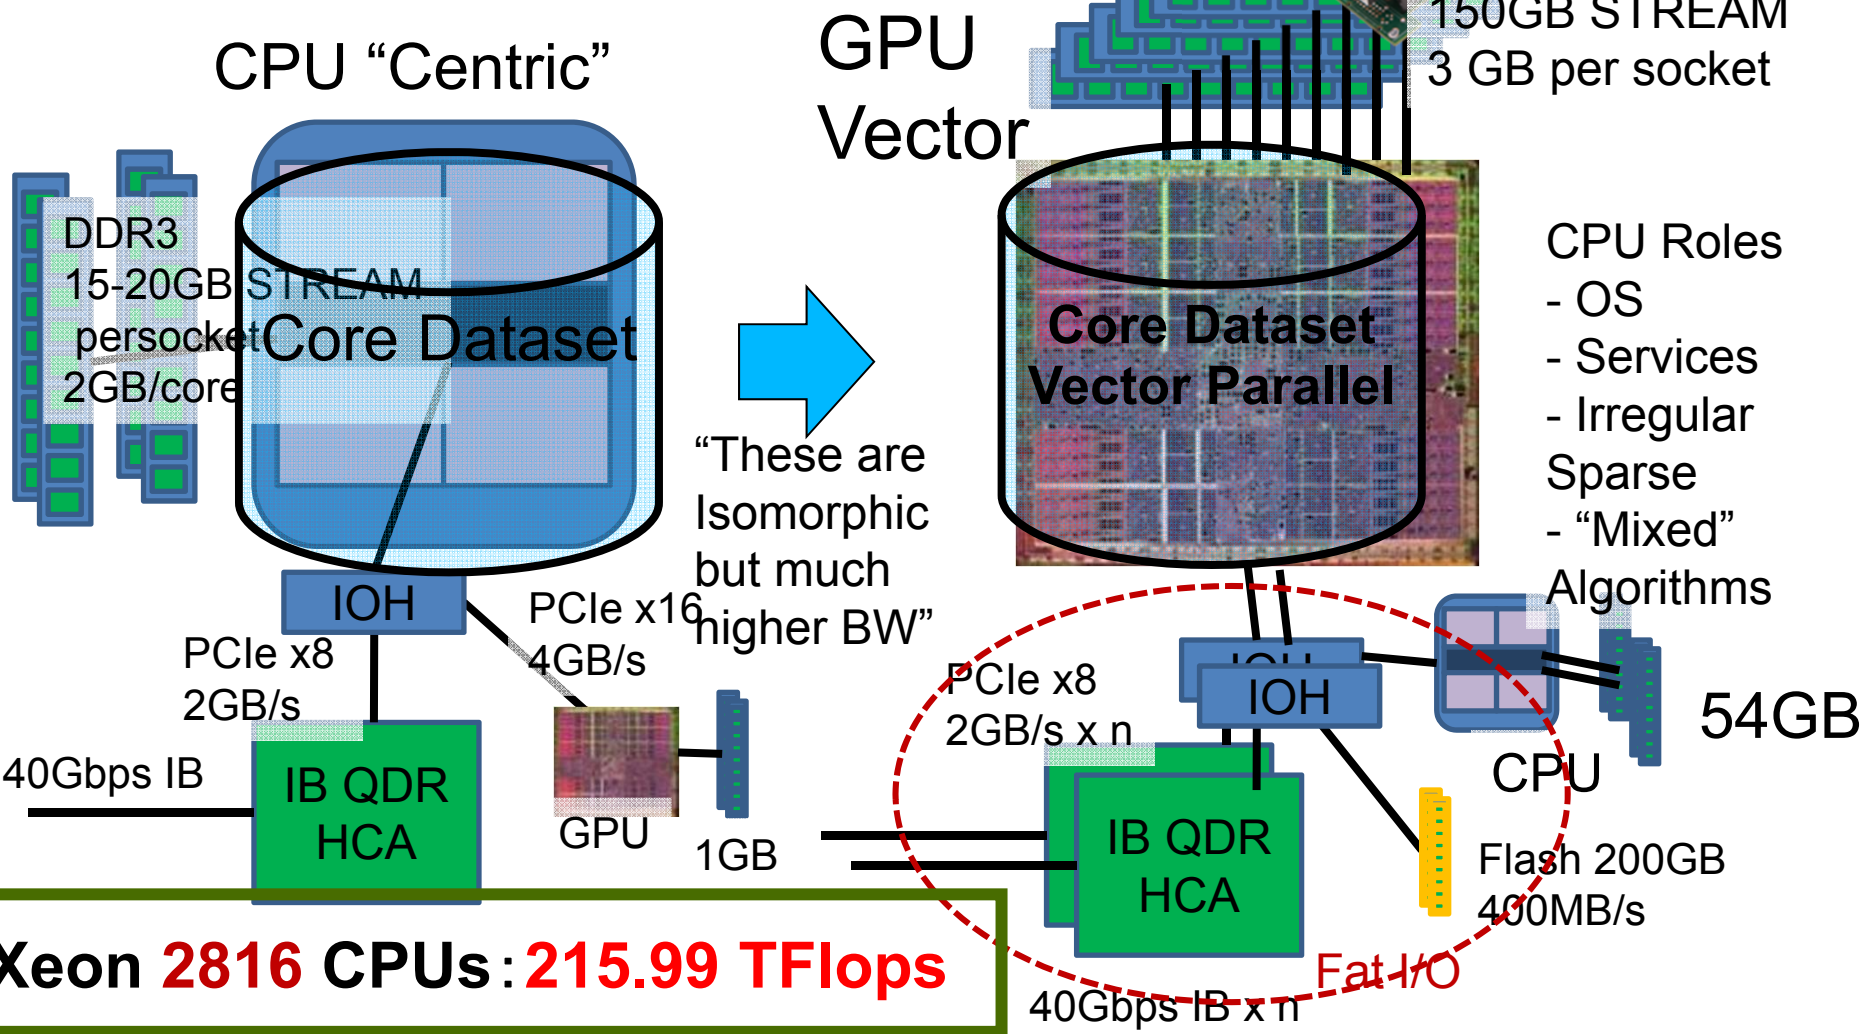


TSUBAME2.0 System Overview (2.4 Pflops/15PB)



TSUBAME 2.0: GPU Centric Nodes

Total M2050 **4224 GPU**
1408 nodes: **2175.36 TFlops**

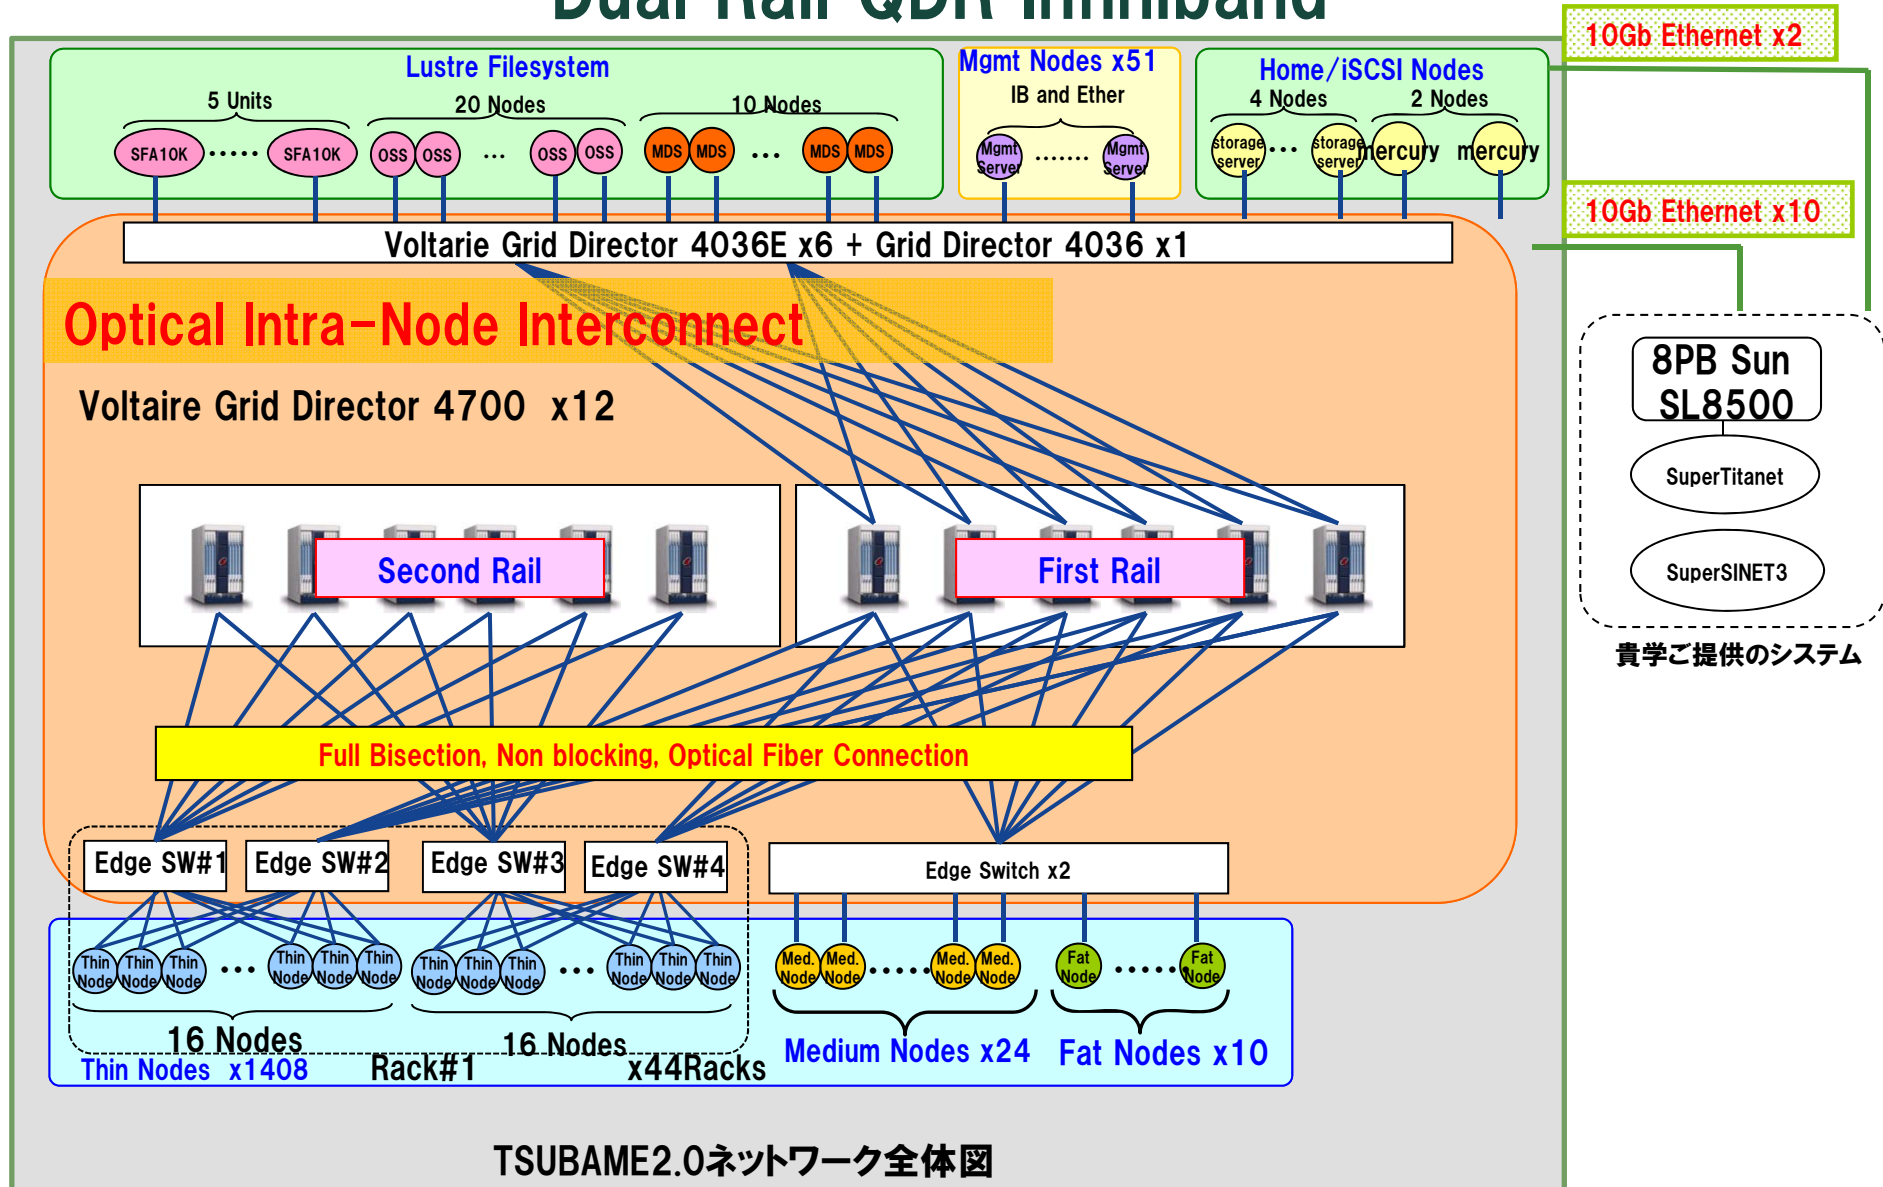



Xeon 2816 CPUs: 215.99 TFlops

GPU M2050



TSUBAME 2.0 Full Bisection Fat Tree, Optical, Dual Rail QDR Infiniband










STORAGETEK



TSUBAME2.0 Nov 1st, 2010

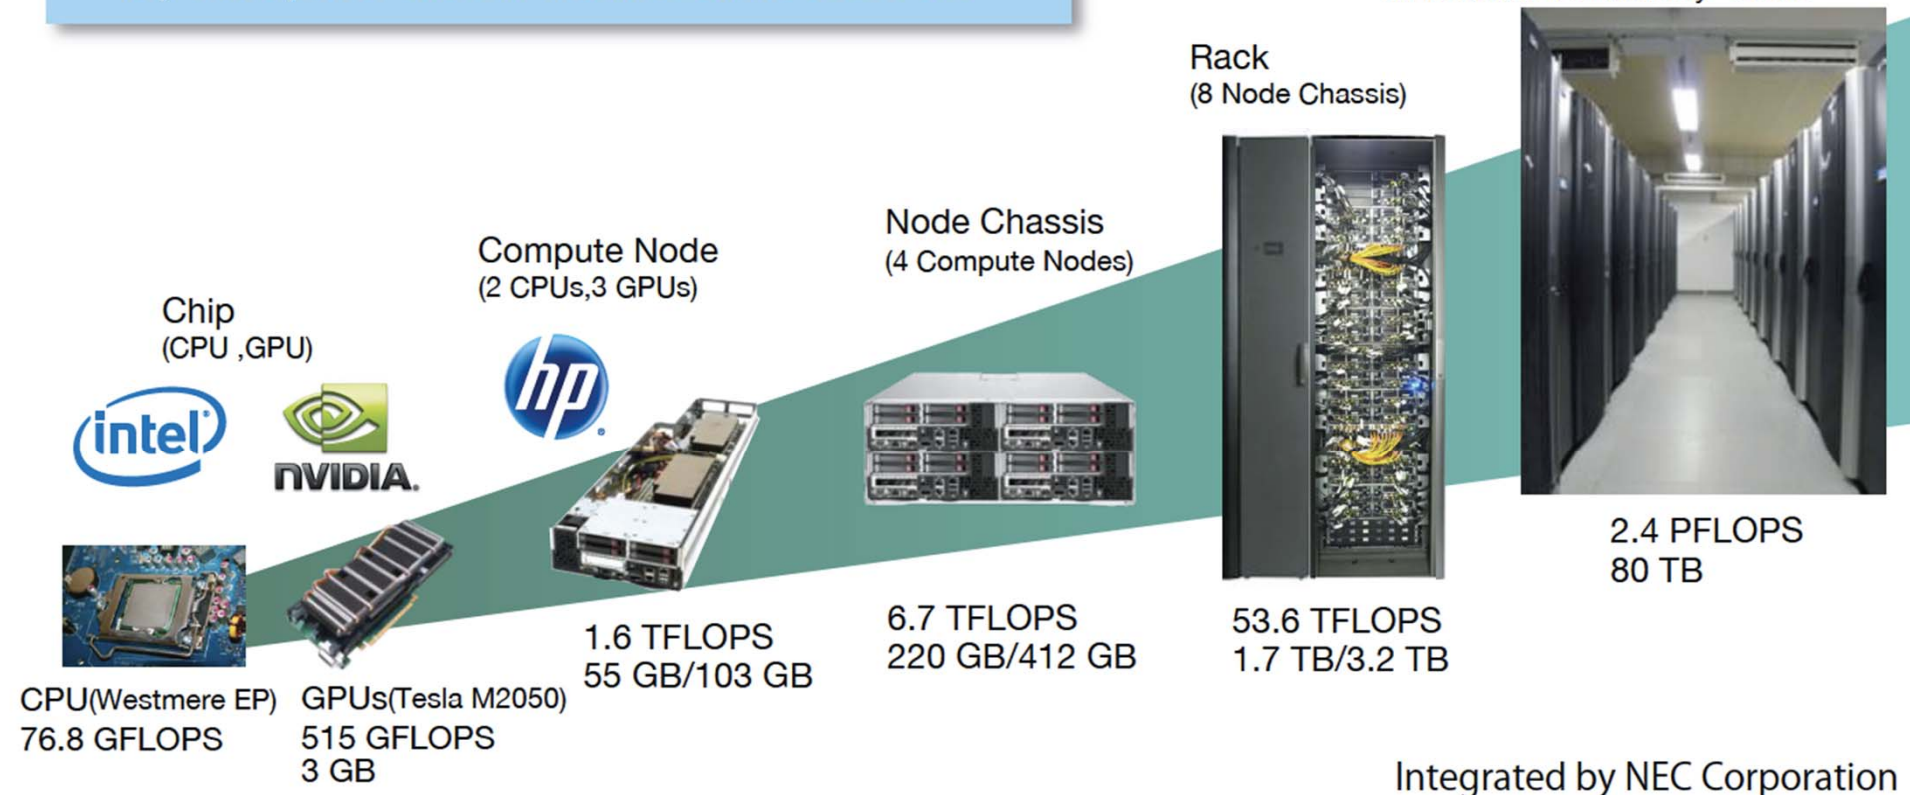


TSUBAME2.0: A GPU-centric Green 2.4 Petaflops Supercomputer

Tsubame 2.0: "Tiny" footprint, very power efficient

- Floorspace less than 200m² (2,100 ft²)
- Top-class power efficient machine on the Green 500

System
(42 Racks)
1408 GPU Compute Nodes,
34 Nehalem "Fat Memory" Nodes





科学と技術で未来を創造する

Supercomputer in the world





2010 November

Rank	Site	Computer/Year Vendor	Cores	R _{max}	R _{peak}	Power
1	National Supercomputing Center in Tianjin China	Tianhe-1A - NUDT YH Cluster, X5670 2.93Ghz 6C, NVIDIA GPU, FT-1000 8C / 2010 NUDT	186368	2566.00	4701.00	4040.00
2	DOE/SC/Oak Ridge National Laboratory United States	Jaguar - Cray XT5-HE Opteron 6-core 2.6 GHz / 2009 Cray Inc.	224162	1759.00	2331.00	6950.60
3	National Supercomputing Centre in Shenzhen (NSCS) China	Nebulae - Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU / 2010 Dawning	120640	1271.00	2984.30	2580.00
4	GSIC Center, Tokyo Institute of Technology Japan	TSUBAME 2.0 - HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows / 2010 NEC/HP	73278	1192.00	2287.63	1398.61
5	DOE/SC/LBNL/NERSC United States	Hopper - Cray XE6 12-core 2.1 GHz / 2010 Cray Inc.	153408	1054.00	1288.63	2910.00

ORNL Jaguar vs Tsubame 2.0

Similar Peak Performance, 1/4 the Size and Power




Supercomputer in the world



The Green500 list -- November 2010


Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)
<u>1</u>	1684.20	IBM Thomas J. Watson Research Center	NNSA/SC Blue Gene/Q Prototype	38.80
<u>2</u>	958.35	GSIC Center, Tokyo Institute of Technology	HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows	1243.80
<u>3</u>	933.06	NCSA	Hybrid Cluster Core i3 2.93Ghz Dual Core, NVIDIA C2050, Infiniband	36.00
<u>4</u>	828.67	RIKEN Advanced Institute for Computational Science	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	57.96
<u>5</u>	773.38	Forschungszentrum Juelich (FZJ)	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
<u>5</u>	773.38	Universitaet Regensburg	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
<u>5</u>	773.38	Universitaet Wuppertal	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
<u>8</u>	740.78	Universitaet Frankfurt	Supermicro Cluster, QC Opteron 2.1 GHz, ATI Radeon GPU, Infiniband	385.00

Power Efficiency



6600x
Faster

3x efficient



Laptop: SONY Vaio type Z (VPCZ1)
CPU: Intel Core i7 620M (2.66GHz)
MEMORY: DDR3-1066 4GBx2
OS: Microsoft Windows 7 Ultimate 64bit
HPL: Intel(R) Optimized LINPACK Benchmark for
Windows (10.2.6.015)
256GB HDD

18.1 Gflops


369 MFlops/Watt

Supercomputer: TSUBAME 2.0
CPU: 2714 Intel Westmere 2.93 Ghz
GPU: 4071 nVidia Fermi M2050
MEMORY: DDR3-1333 80TB + GDDR5 12TB
OS: SuSE Linux 11 + Windows HPC Server R2
HPL: Tokyo Tech Heterogeneous HPL
11PB Hierarchical Storage

1.192 Pflops

1037 MFlops/Watt

NVIDIA GPU



		Intel Core i7 Extreme	Tesla C2050 /M2050	GeForce GTX 580 Fermi
GPU	Peak Performance [GFlops]	51.2*, 102.4	515*, 1030	197*, 1576
	Number of Processor	4	448	512
	Core Clock [MHz]	3200	1476	1544
Memory	Bandwidth[GB/s]	32	148.8	192.1
	Memory Interface [bit]	64	384	384
	Memory Clock [GHz]	1.333 (DDR3)	1.55 (GDDR5)	2.00 (GDDR5)
B_{peak}/F_{peak}	Bandwidth/Performance	0.624	0.289	0.974









Tesla M2050
Peak Power : 225W




Peak Power : 244W

GPU Architecture

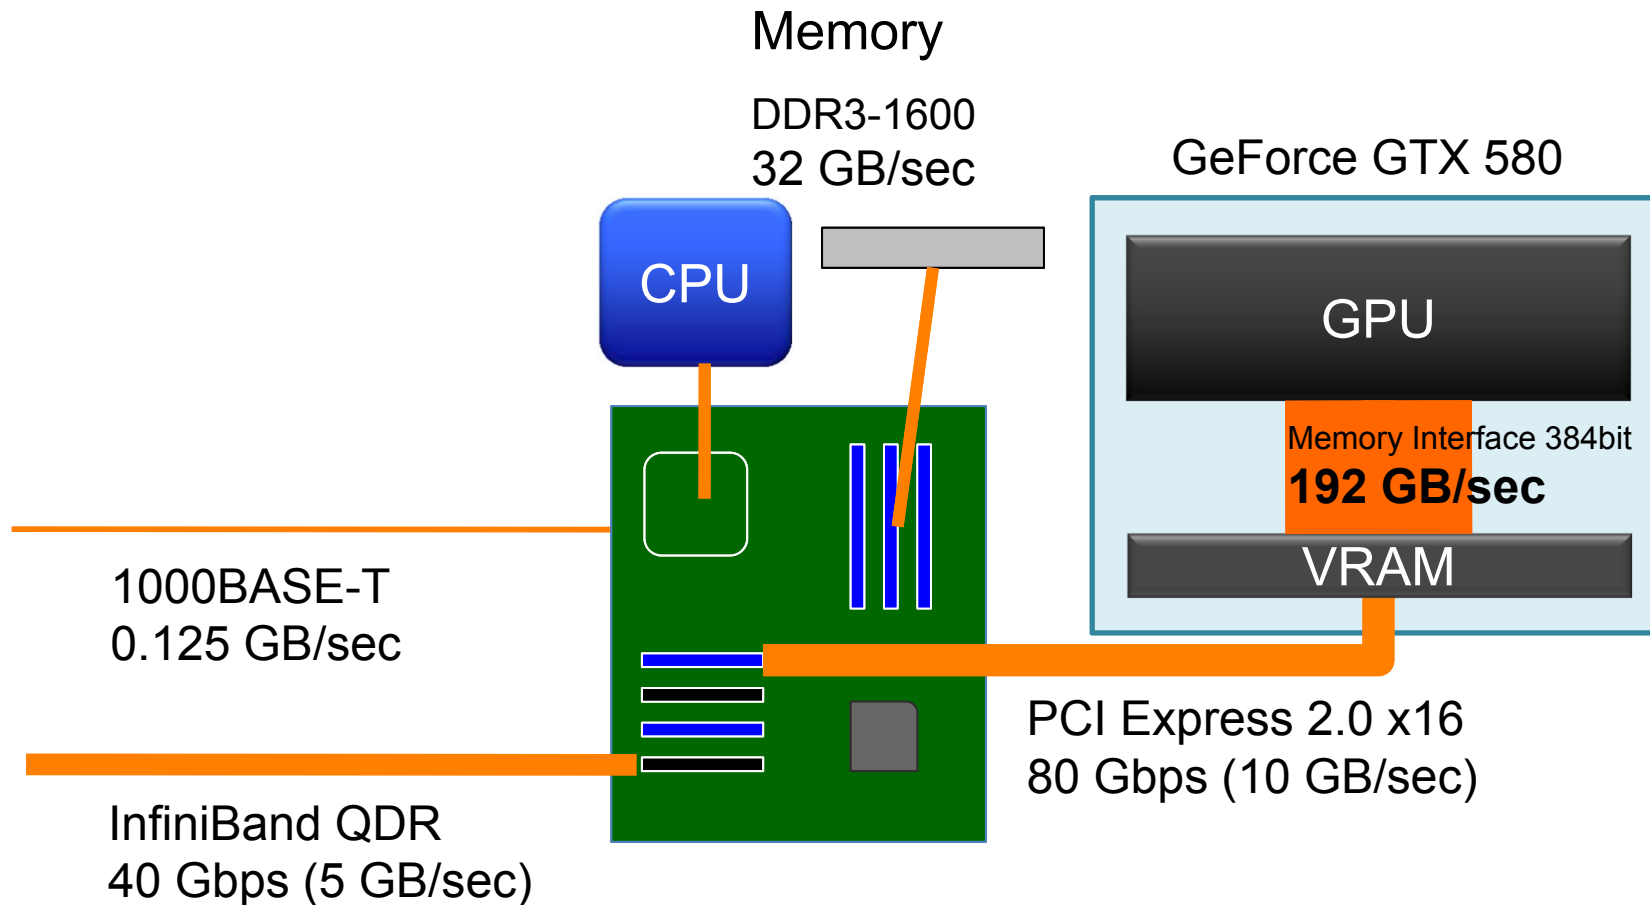


- | | | |
|---|---------------------------------|-------------------------|
|  | Global memory | ~6GB (VRAM) |
|  | Streaming Multiprocessor | ~16 (C2050 (GF100): 14) |
|  | Shared memory + L1 Cache | 64 Kbyte |
|  | Streaming Processor (CUDA core) | 8~48 per SM, total 512 |


Heterogeneous Computer



■ Several Bandwidth Bottle Necks



CFD Performances in GPU Computing




Partial GPU Implementation 30% up ~ × 3

- ⌘ Only Hot spot (Intensive part) : small cost
- ⌘ Overhead of host (CPU) memory device (GPU) memory communication

FULL GPU Implementation × 10 ~ × 100

- ⌘ Limitation of device (GPU) on board memory size


Real-time TSUNAMI Simulation



ADPC : Asian Disaster Preparedness Center

Early Warning System:

Data Base



Real-time CFD

high accuracy



Shallow-Water Eq.

Conservative Form:

$$\frac{\partial h}{\partial t} + \frac{\partial hu}{\partial x} + \frac{\partial hv}{\partial y} = 0$$

Assuming

hydrostatic balance


in the vertical direction,

$$\frac{\partial hu}{\partial t} + \frac{\partial}{\partial x} \left(hu^2 + \frac{1}{2} gh^2 \right) + \frac{\partial huv}{\partial y} = -gh \frac{\partial z}{\partial x}$$


3D → **2D equation**

$$\frac{\partial hv}{\partial t} + \frac{\partial huv}{\partial x} + \frac{\partial}{\partial y} \left(hv^2 + \frac{1}{2} gh^2 \right) = -gh \frac{\partial z}{\partial y}$$


Tsunami Modeling



free surface flow



Directional-Splitting Method



$$\frac{\partial U}{\partial t} + \frac{\partial F}{\partial x} + \frac{\partial G}{\partial y} = S \quad U = \begin{bmatrix} h \\ hu \\ hv \end{bmatrix}, \quad F = \begin{bmatrix} hu \\ hu^2 + \frac{1}{2}gh^2 \\ huv \end{bmatrix}, \quad G = \begin{bmatrix} hv \\ huv \\ hv^2 + \frac{1}{2}gh^2 \end{bmatrix}$$

First Step: x-directional computation

$$\frac{\partial h}{\partial t} + \frac{\partial hu}{\partial x} = 0 \quad \frac{\partial hv}{\partial t} + \frac{\partial uhv}{\partial x} = 0 \quad \frac{\partial hu}{\partial t} + \frac{\partial}{\partial x} \left(hu^2 + \frac{1}{2}gh^2 \right) = -gh \frac{\partial z}{\partial x}$$


Second Step: y-directional computation

$$\frac{\partial h}{\partial t} + \frac{\partial hv}{\partial y} = 0 \quad \frac{\partial hu}{\partial t} + \frac{\partial vhu}{\partial y} = 0 \quad \frac{\partial hv}{\partial t} + \frac{\partial}{\partial y} \left(hv^2 + \frac{1}{2}gh^2 \right) = -gh \frac{\partial z}{\partial y}$$

For Characteristics-based Method

For Conservative Semi-Lagrangian Method


CIP-CSL2 (Conservative Semi-Lagrangian)



R. Tanaka, T. Nakamura, and T. Yabe, Comp. Phys. Comm., 126, 232-243 (2000).

Continuum Eq.

$$\frac{\partial h}{\partial t} + \frac{\partial hu}{\partial x} = 0$$



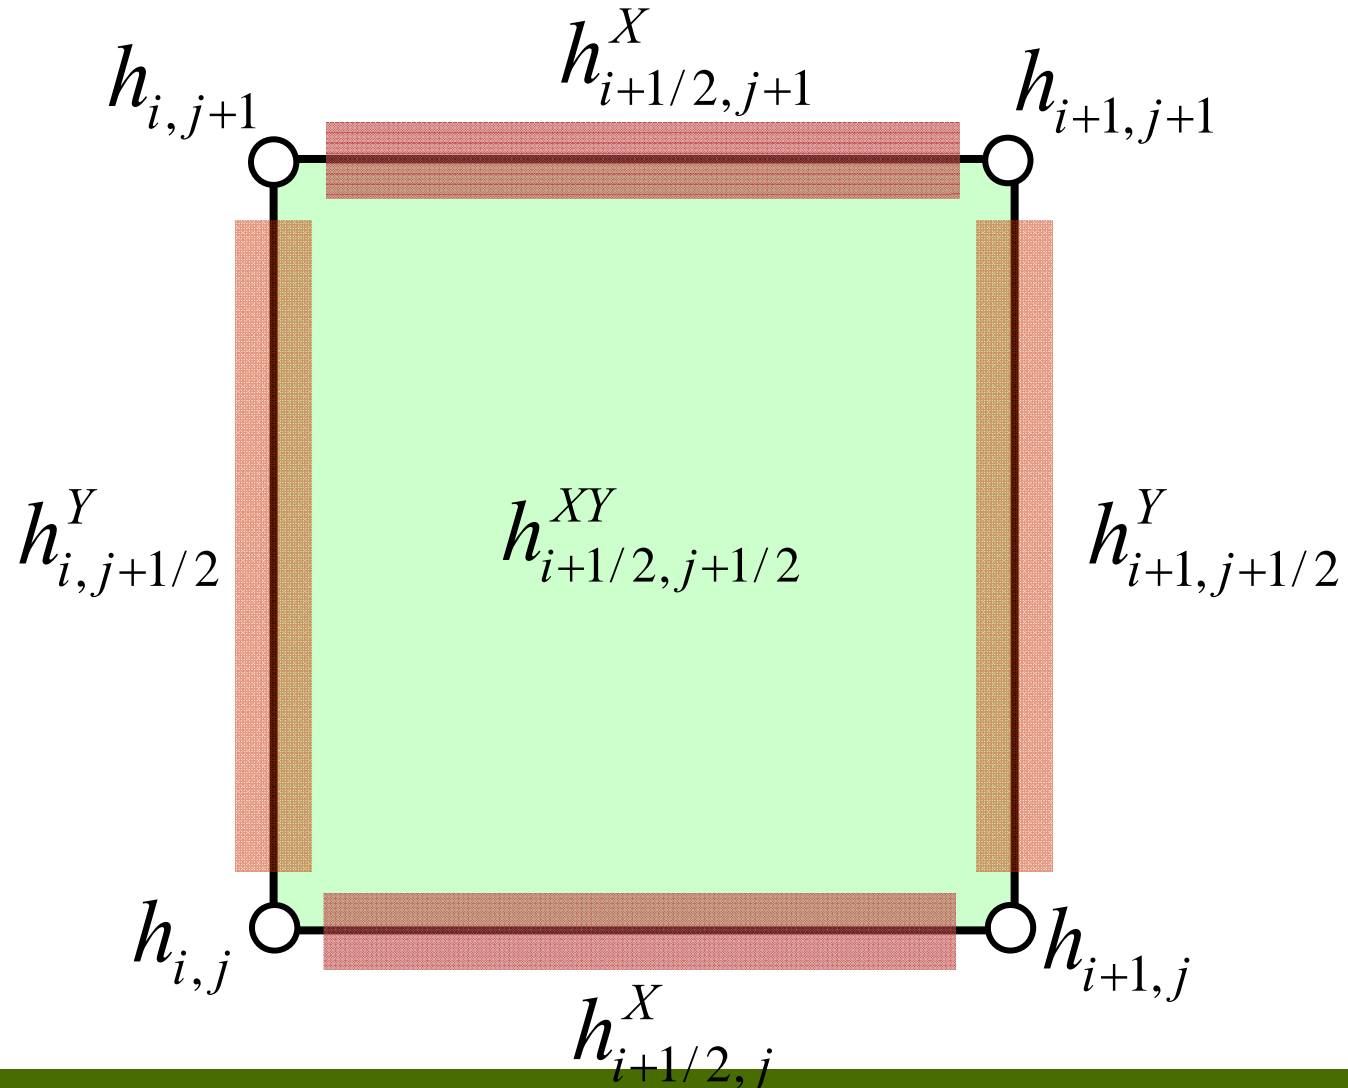

$$h_i(x) = a(x - x_i)^2 + b(x - x_i) + h_i \quad a = \frac{3h_{i+1} + 3h_i}{\Delta x^2} - \frac{6h_{i+1/2}^X}{\Delta x^3}, \quad b = \frac{6h_{i+1/2}^X}{\Delta x^2} - \frac{2h_{i+1} + 4h_i}{\Delta x}$$

$$h_{x,i} = \frac{6h_{i+1/2}^X}{\Delta x^2} - \frac{2h_{i+1} + 4h_i}{\Delta x} \quad h_i(x_i) = h_i^n \quad h_i(x_i + \Delta x) = h_{i+1}^n \quad \int_{x_i}^{x_i + \Delta x} h_i(x) dx = h_{i+1/2}^X$$


$$\Delta h_{i+1}^X = \int_{x_p}^{x_{i+1}} h_i^n(x) dx$$

$$h_i^{n+1} = h_j^n(x_i - u\Delta t) \quad h_{i+1/2}^{X,n+1} = h_{i+1/2}^{X,n} - \Delta h_{i+1}^n + \Delta h_i^n = -\left(\frac{a^n}{3} \xi^3 + \frac{b^n}{2} \xi^2 + h_i^n \xi \right)$$

2-dimensional Variable Configuration



Characteristics-Based Method




$$\frac{\partial U}{\partial t} + \frac{\partial F}{\partial x} = 0 \quad U = \begin{bmatrix} h \\ hu \end{bmatrix}, \quad F = \begin{bmatrix} hu \\ hu^2 + \frac{1}{2}gh^2 \end{bmatrix}$$

Factorizing:

$$\frac{\partial U}{\partial t} + A \frac{\partial U}{\partial x} = 0 \quad A = \frac{\partial F}{\partial U} = \begin{bmatrix} 0 & 1 \\ \Gamma^2 - u^2 & 2u \end{bmatrix}$$

$$\text{for } \Gamma = \sqrt{gh}$$



Riemann invariants :

$$\frac{\partial W}{\partial t} + \Lambda \frac{\partial W}{\partial x} = 0 \quad W = \begin{bmatrix} \Gamma + \frac{1}{2}u \\ \Gamma - \frac{1}{2}u \end{bmatrix}, \quad \Lambda = \begin{bmatrix} u + \Gamma & 0 \\ 0 & u - \Gamma \end{bmatrix}$$

$$\Gamma^\pm = \Gamma \pm \frac{1}{2}u$$


$$\Gamma^{n+1} = \frac{1}{2} \left\{ \Gamma^{+n+1} + \Gamma^{-n+1} + \frac{1}{2}(u^+ - u^-) \right\}$$

$$\frac{\partial \Gamma^\pm}{\partial t} + \lambda^\pm \frac{\partial \Gamma^\pm}{\partial x} = 0$$

$$u^{n+1} = \frac{1}{2} \left\{ u^+ - u^- + 2(\Gamma^+ - \Gamma^-) \right\}$$



$$\left. \begin{array}{l} \Gamma^{n+1} \\ u^{n+1} \end{array} \right\} \rightarrow \begin{array}{l} h^{n+1} \\ (hu)^{n+1} \end{array}$$

Hydrostatic Balance (1/2)



$$H = h + z \quad \Rightarrow \quad h = H - z$$

$$\frac{\partial hu}{\partial t} + \frac{\partial}{\partial x} \left(hu^2 + \frac{1}{2} gh^2 \right) = -gh \frac{\partial z}{\partial x}$$




$$\frac{\partial hu}{\partial t} + \frac{\partial}{\partial x} \left(hu^2 + \frac{1}{2} g(H-z)^2 \right) = -g(H-z) \frac{\partial z}{\partial x}$$

$$\frac{\partial hu}{\partial t} + \frac{\partial}{\partial x} \left(hu^2 + \frac{1}{2} gH^2 - ghHz + \frac{1}{2} gz^2 \right) = -gH \frac{\partial z}{\partial x} + gz \frac{\partial z}{\partial x}$$

$$\frac{\partial hu}{\partial t} + \frac{\partial}{\partial x} \left(hu^2 + \frac{1}{2} gH^2 \right) = gH \frac{\partial z}{\partial x} + gz \frac{\partial H}{\partial x} - gH \frac{\partial z}{\partial x}$$

$$\frac{\partial hu}{\partial t} + \frac{\partial}{\partial x} \left(hu^2 + \frac{1}{2} gH^2 \right) = gz \frac{\partial H}{\partial x}$$


Hydrostatic Balance (2/2)



For characteristics-based method,

$$\frac{\partial h}{\partial t} + \frac{\partial hu}{\partial x} = 0 \quad \frac{\partial hu}{\partial t} + \frac{\partial}{\partial x} \left(hu^2 + \frac{1}{2} gH^2 \right) = gz \frac{\partial H}{\partial x}$$

$$h = H - z$$




$$\frac{\partial H}{\partial t} + \frac{\partial Hu}{\partial x} = \frac{\partial zu}{\partial x} \quad \frac{\partial Hu}{\partial t} + \frac{\partial}{\partial x} \left(Hu^2 + \frac{1}{2} gH^2 \right) = u \frac{\partial zu}{\partial x}$$

Characteristics-based Method for H and Hu


+ Fractional Step :

$$\frac{\partial H}{\partial t} = \frac{\partial zu}{\partial x} \quad \frac{\partial Hu}{\partial t} = u \frac{\partial zu}{\partial x}$$


CUDA GPU Computing





x-directional Computation



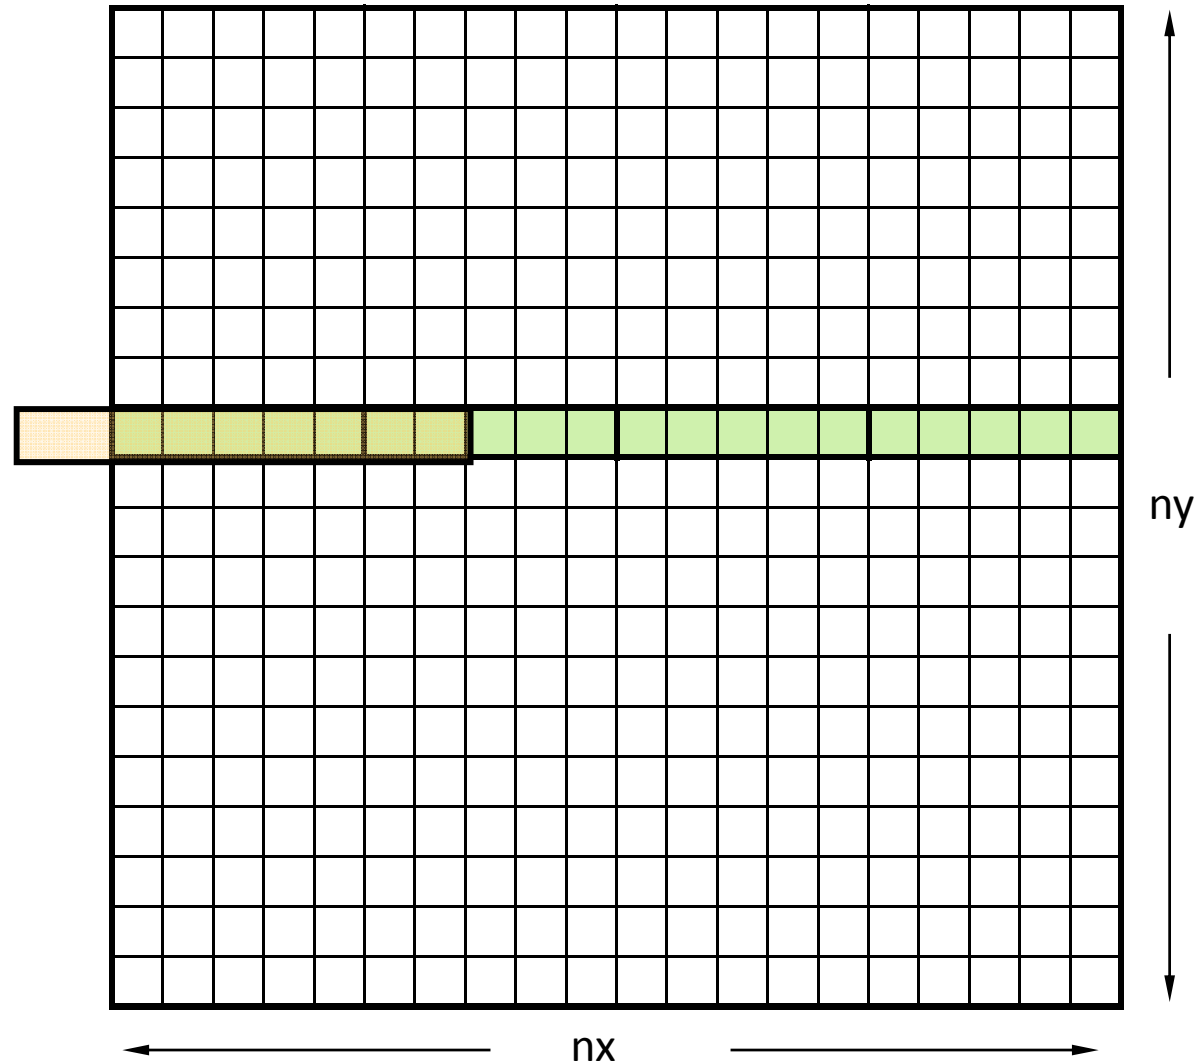

Data read n-step




Calculation



Data write to n+1 step





CUDA GPU Computing




y-directional Computation

blockDim.x=16




$2 * blockDim.y$



Data read & write



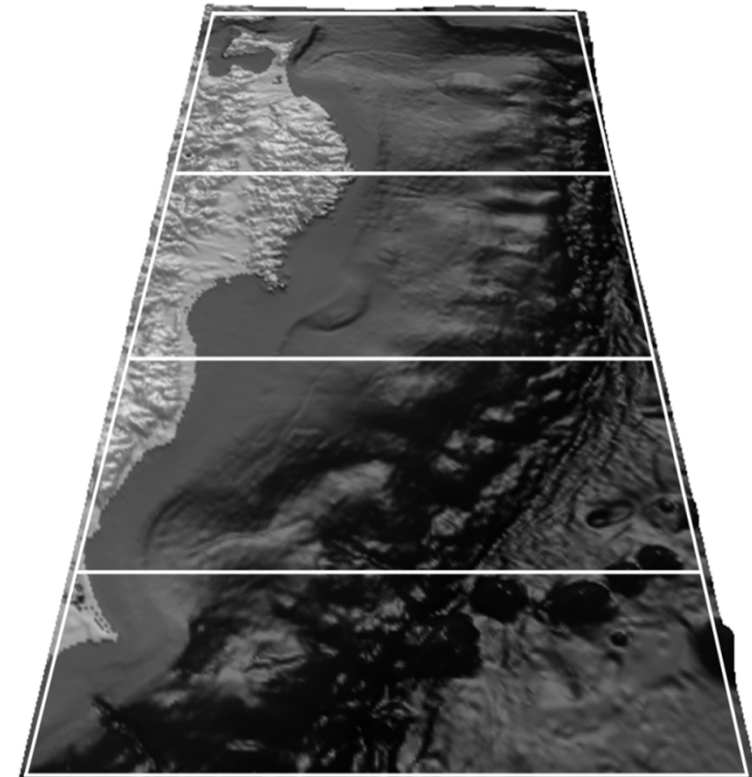


Calculation



SCREEN Capture





Large-scale Real-time Tsunami Simulator



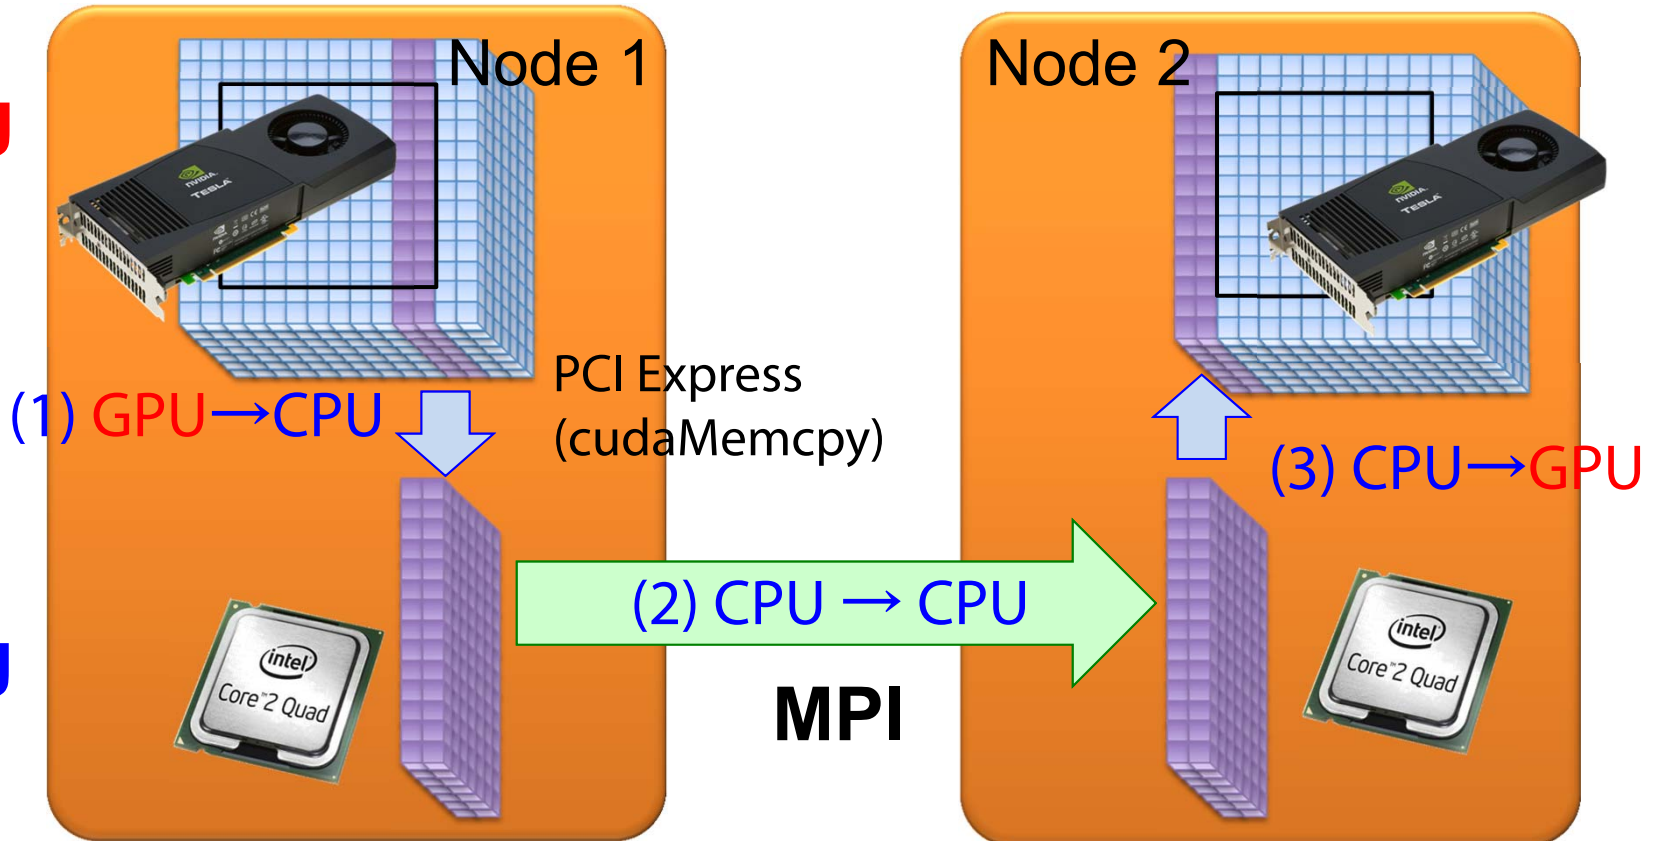
**8 GPU 400km × 800km
(100m mesh)**

within 3 min

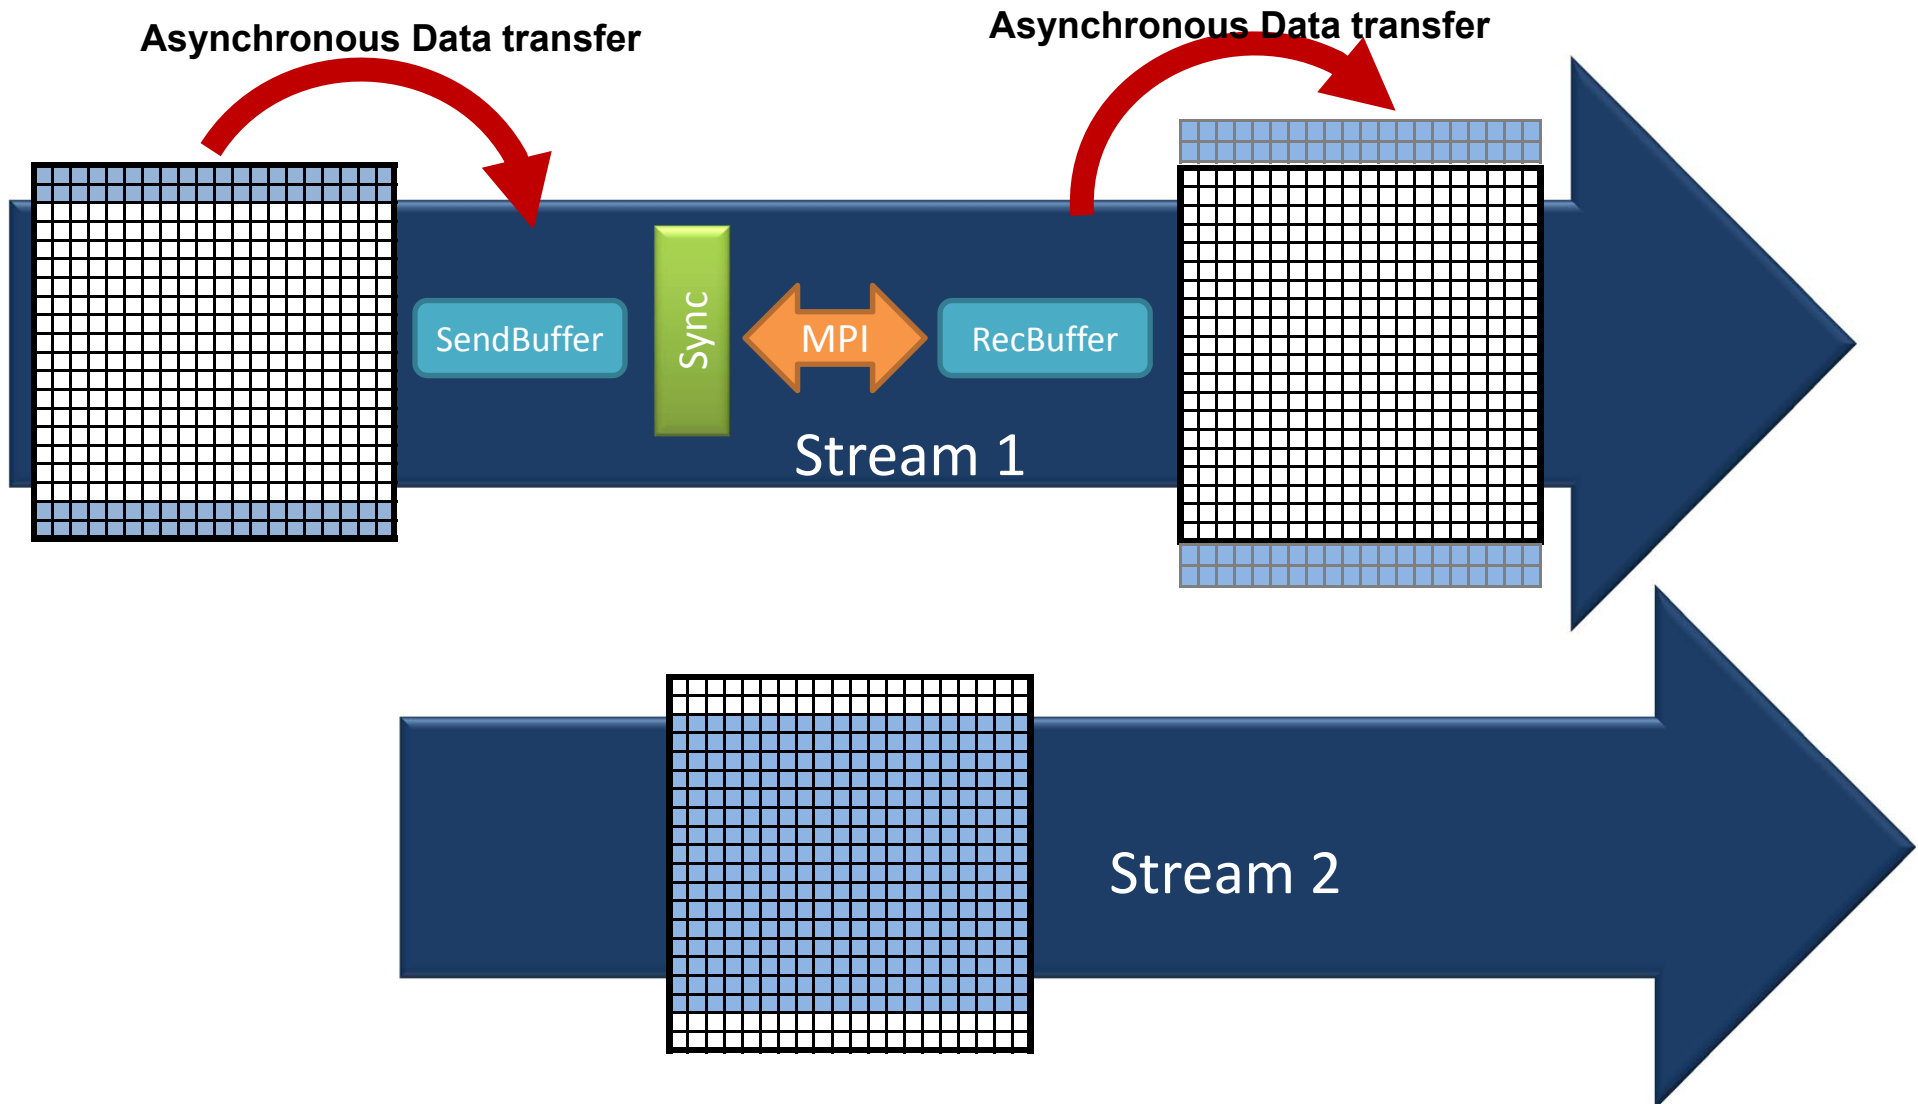

Multi-GPU : Domain decomposition







GPU



Overlapping between Computation and Communication




Large-scale Real-time Tsunami Simulator




**8 GPU 400km × 800km
(100m mesh)**

within 3 min

CPU-GPU Performance Comparison





(1 CPU Core based)




Results on Multi-node Computing

Tsubame 2.0




Multi-GPU Scalability




Tsubame 2.0


CPU-GPU Scalability Tsubame 2.0



Two-Phase Flow Simulation




Particle Method
ex. **SPH**



Mesh Method (Surface Capture)

- Navier-Stokes solver: Fractional Step
- Time integration: 3rd TVD Runge-Kutta
- Advection term: 5th WENO
- Diffusion term: 4th FD
- Poisson: MG-BiCGstab
- Surface tension: CSF model
- Surface capture: CLSVOF(THINC + Level-Set)


Low accuracy
< 10^6 particles




not splash

Numerical noise and unphysical oscillation

High accuracy > 10^8 mesh points



EQUATIONs for Two-Phase Flow



Time Integration : 3rd-order TVD Runge-Kutta

$$\nabla \cdot \mathbf{u} = 0$$

5th-upFD or 5th HJ-WENO

diffusion : 4th FD

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \mathbf{u} + \frac{1}{\rho} \mathbf{F}$$

Surface Tension CSF

5th HJ-WENO

$$\frac{\partial \phi}{\partial t} + (\mathbf{u} \cdot \nabla) \phi = 0 \quad \frac{\partial \phi}{\partial \tau} = \text{sgn}(\phi) (1 - |\nabla \phi|)$$


Re-initialization 5th HJ-WENO

Advection VOF THINC/WLIC


$$\frac{\partial \psi}{\partial t} + \nabla \cdot (\mathbf{u} \psi) = 0$$

$$\rho = \rho_g + (\rho_l - \rho_g) H(\phi)$$

$$\nu = \nu_g + (\nu_l - \nu_g) H(\phi)$$

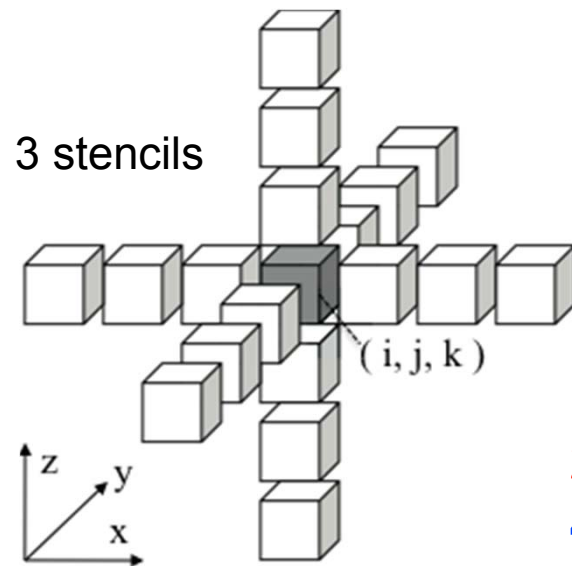


3D Advection Computation



Advection equation


$$\frac{\partial f}{\partial t} + (\mathbf{u} \cdot \nabla) f = 0$$




19 input values/cell
259 flop/cell (5th-WENO)
49 flop/cell (5th-up FD)

Discretization: Space : 5th-WENO
Time : 3rd TVD Runge-Kutta

312 GFlops (1GPU:GTX285)



Stencil Computation



■ Example: 2-dimensional diffusion Equation by FDM

$$\frac{f_{i,j}^{n+1} - f_{i,j}^n}{\Delta t} = \kappa \left(\frac{f_{i+1,j}^n - 2f_{i,j}^n + f_{i-1,j}^n}{\Delta x^2} + \frac{f_{i,j+1}^n - 2f_{i,j}^n + f_{i,j-1}^n}{\Delta y^2} \right)$$




$$f_{i,j}^{n+1} = c_0 f_{i,j}^n + c_1 f_{i+1,j}^n + c_2 f_{i-1,j}^n + c_3 f_{i,j+1}^n + c_4 f_{i,j-1}^n$$


FLOP = 9

Byte = 4*6 = 24 byte : read 5, write 1

FLOP/Byte = 9/24 = 0.375



Arithmetic INTENSITY: FLOP/Byte



FLOP = number of FP operation for applications


Byte = Byte number of memory access for applications

F = Peak Performance of floating point operation

B = Peak Memory Bandwidth

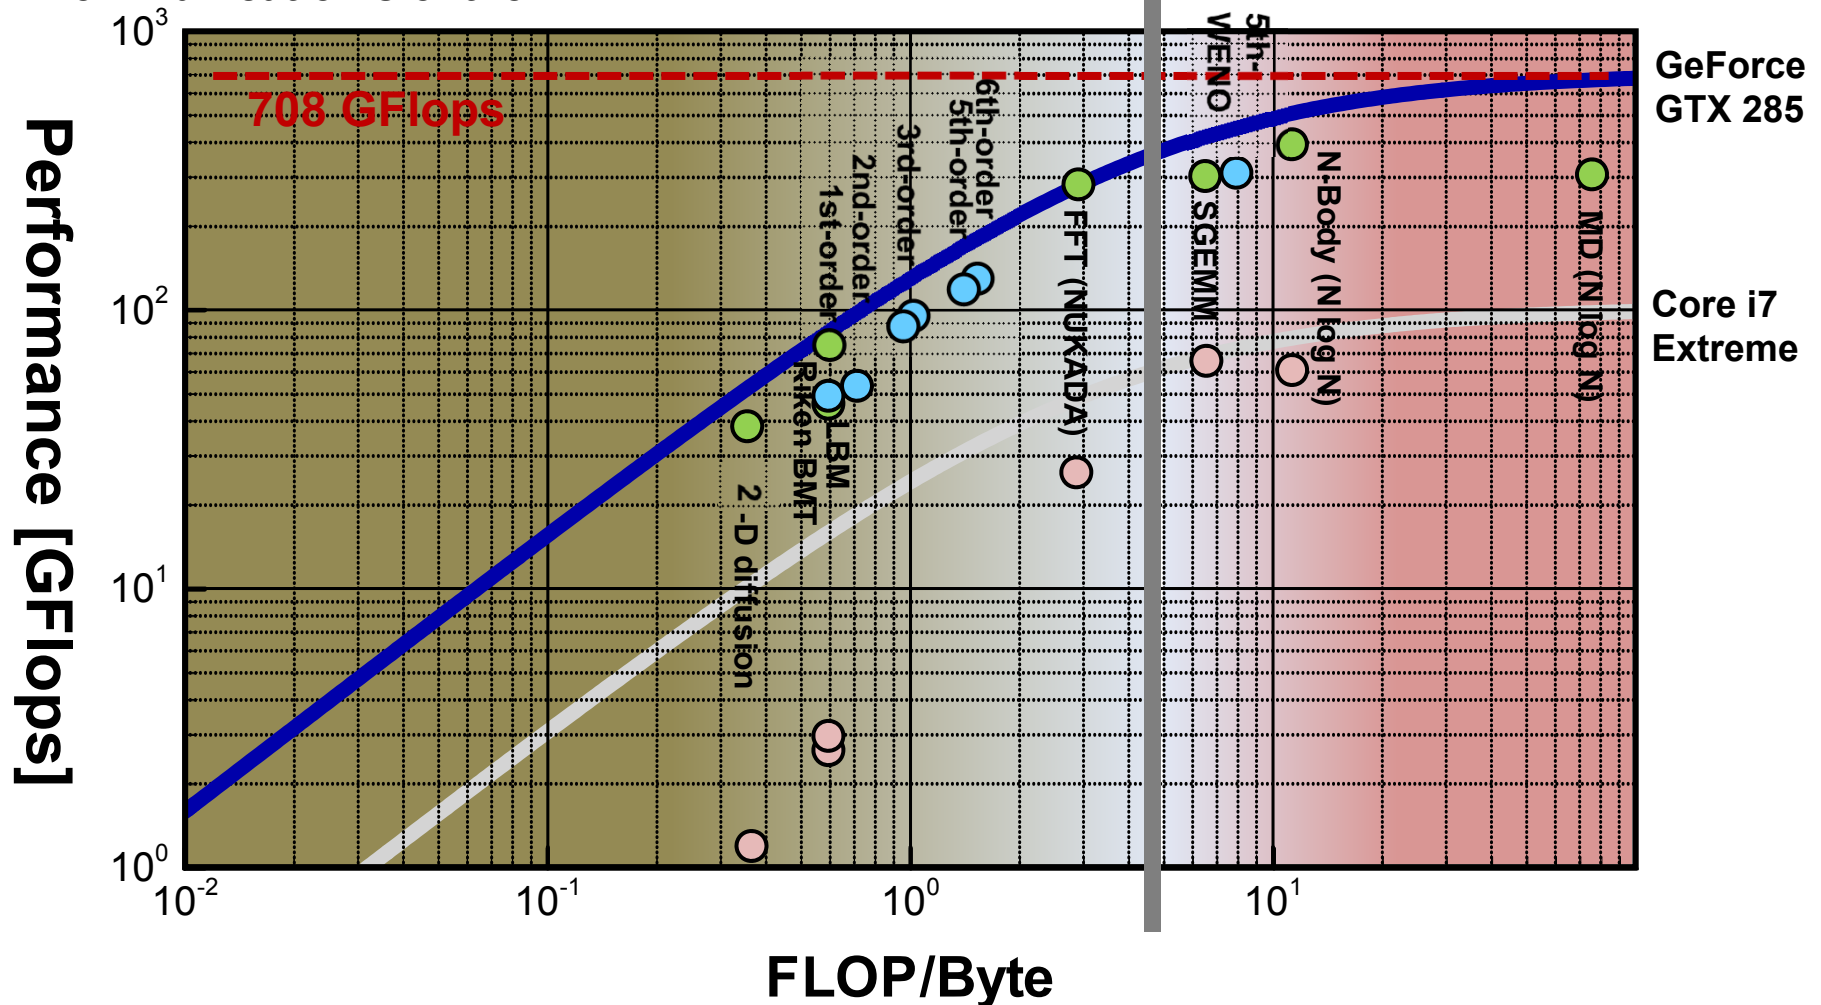
$$\begin{aligned} \text{Performance} &= \frac{\text{FLOP}}{\text{FLOP}/F + \text{Byte}/B + \alpha} \\ &= \frac{\text{FLOP/Byte}}{\text{FLOP/Byte} + F/B + \alpha} F \end{aligned}$$

Application Performances




Roofline model: Williams, Patterson 2008
 Communications of the ACM

$$\text{FLOP/Byte} = F/B$$

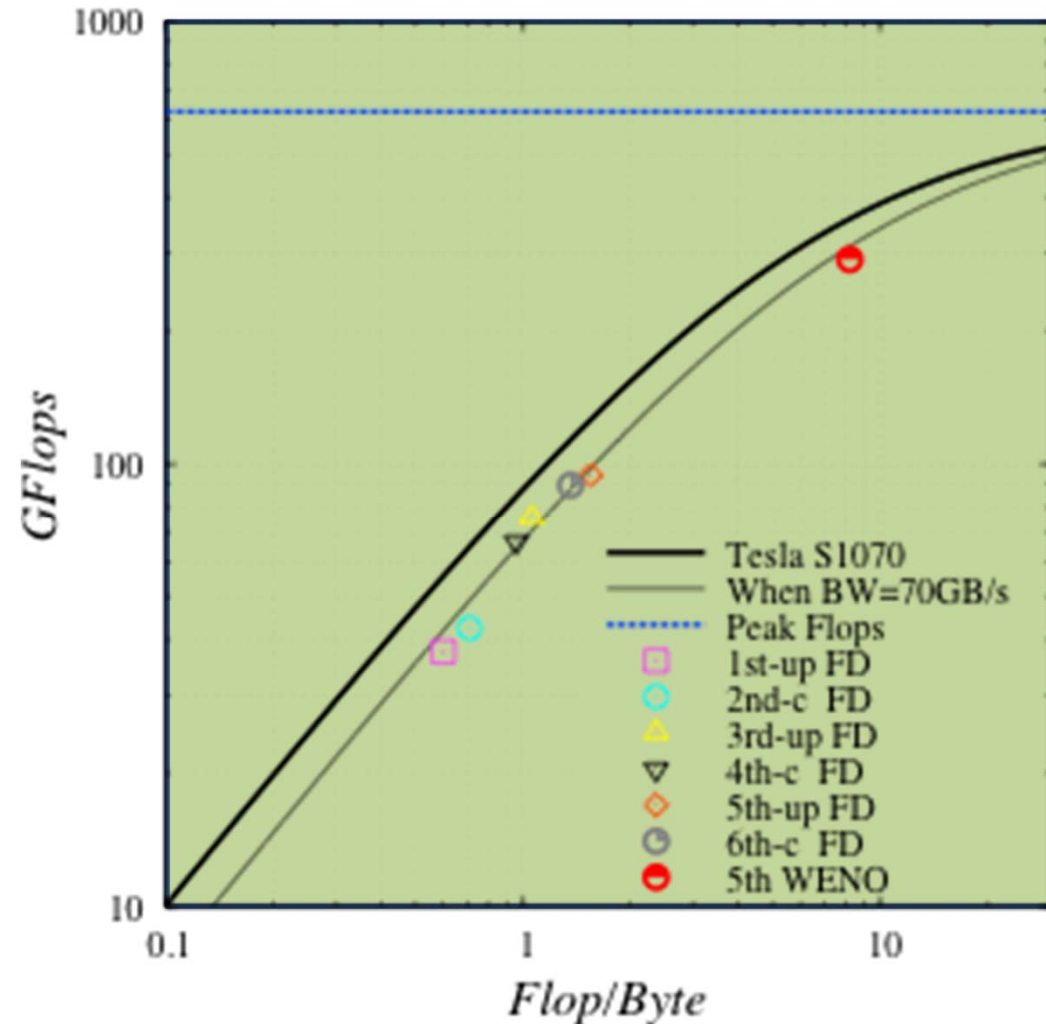


Performance of Advection Computation




Scheme	flop/byte		GFlops Tesla S1070
	no-SMem	SMem	
1st-up FD	0.29	0.60	37.61
2nd-c FD	0.34	0.71	42.46
3rd-up FD	0.38	1.06	75.49
4th-c FD	0.34	0.96	66.64
5th-up FD	0.45	1.55	94.58
6th-c FD	0.4	1.36	89.62
5th-WENO	2.40	8.22	289.66

3D advection 416x416x416 cells
Time integration: 3rd-order TVD Runge-Kutta



Level-Set method (LSM)



The Level-Set methods (LSM) use the signed distance function to capture the interface. The interface is represented by the zero-level set (zero-contour).

ϕ : **Level-Set function(distance function)**

H : Heaviside function

$$\begin{cases} H(\phi) = \frac{1}{2} & \phi > \varepsilon \\ H(\phi) = \frac{1}{2} \left(\frac{\phi}{\varepsilon} + \frac{1}{\pi} \sin \left(\frac{\pi\phi}{\varepsilon} \right) \right) & |\phi| \leq \varepsilon \\ H(\phi) = -\frac{1}{2} & \phi < -\varepsilon \end{cases}$$

Re-initialization for Level-Set function

$$\frac{\partial \phi}{\partial \tau} = \text{sgn}(\phi) (1 - |\nabla \phi|)$$

Advantage : Curvature calculation, Interface boundary

Drawback : Volume conservation

Contour of Level Set Function Heaviside Function





Fig. Takehiro Himeno, et. Al., JSME, 65-635,B(1999),pp2333-2340

Continuous Surface Force (CSF) model by Brackbill, Kothe and Zemach (1991)



GP GPU

The interfacial surface force is transformed to a volume force in the region near the interface via a delta function

Surface tension force

Curvature

↓


$$\mathbf{F}_S = \sigma \kappa \mathbf{n}$$

Normal vector

←

$$\kappa = -\nabla \cdot \mathbf{n} = -\nabla \cdot \frac{\nabla \phi}{|\nabla \phi|}$$

$$\mathbf{F}_S = \sigma \kappa \delta(\phi) \nabla \phi$$




Surface tension represented by volume force


Approximate delta function

$$\delta(\phi) = \frac{\partial H(\phi)}{\partial \phi} = \frac{1}{2} \left(\frac{1}{\varepsilon} + \frac{1}{\varepsilon} \cos \left(\frac{\pi \phi}{\varepsilon} \right) \right)$$

$$\int_{-\varepsilon}^{\varepsilon} \delta(\phi) d\phi = 1$$



Anti-diffusive Interface Capture



GP GPU

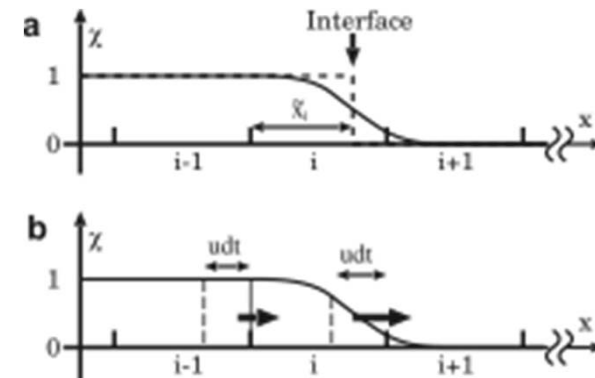
THINC (tangent of hyperbola for interface capturing) Scheme

[Xiao, etal, Int. J. Numer. Meth. Fluid. 48(2005)1023]

- VOF(volume of fluid) type interface capturing method
- Flux from tangent of hyperbola function
- Semi-Lagrangian time integration

$$F_i(x) = \frac{1}{2} \left(1 + \alpha \tanh \left(\beta \left(\frac{x - x_{i-1/2}}{\Delta x} - \tilde{x}_i \right) \right) \right)$$

$$\alpha = \begin{cases} 1 & (\text{if } n_x > 0) \\ -1 & (\text{if } n_x \leq 0) \end{cases}$$




- 1D implementation can be applied to 2D & 3D → Simple

$$Fl_{x,i+1/2} = - \int_{x_{i+1/2}}^{x_{i+1/2} - u_{i+1/2} \Delta t} F_{up}(x) dx \quad up = \begin{cases} i & (\text{if } u_{i+1/2} > 0) \\ i+1 & (\text{if } u_{i+1/2} \leq 0) \end{cases}$$

- Finite Volume like usage
 - * THINC is the method how to compute flux

→ 3 krenel (x, y, z) can be fused to 1 kernel. Merit in memory R/W

Sparse Matrix Solver




$$\mathbf{Ax} = \mathbf{b} \quad \text{for} \quad \nabla \cdot \left(\frac{1}{\rho} \nabla p \right) = \frac{\nabla \cdot \mathbf{u}}{\Delta t}$$


Krylov sub-space methods:
CG, BiCGStab, GMRes, , ,

Pre-conditioner:
Incomplete Cholesky,
ILU, MG, AMG,
Block Diagonal Jacobi

Non-zero Packing:
CRS → ELL, JDL



BiCGStab + MG



Collaboration with
Mizuho Information & Research Institute

Set $k = 0$ $\mathbf{r}_0 = \mathbf{p}_0 = \mathbf{M}^{-1}(\mathbf{b} - \mathbf{A}\mathbf{x}_0)$

for $k = 0; k < N; k++;$

$$\alpha_k = \frac{(\mathbf{r}_0, \mathbf{r}_k)}{(\mathbf{r}_0, \mathbf{M}^{-1}\mathbf{A}\mathbf{p}_k)} \quad \mathbf{q}_k = \mathbf{r}_k - \alpha_k \mathbf{M}^{-1}\mathbf{A}\mathbf{p}_k \quad \omega_k = \frac{(\mathbf{q}_k, \mathbf{M}^{-1}\mathbf{A}\mathbf{q}_k)}{(\mathbf{M}^{-1}\mathbf{A}\mathbf{q}_k, \mathbf{M}^{-1}\mathbf{A}\mathbf{q}_k)}$$

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{p}_k + \omega_k \mathbf{q}_k$$

$$\mathbf{r}_{k+1} = \mathbf{q}_k - \omega_k \mathbf{M}^{-1}\mathbf{A}\mathbf{q}_k$$


if $(\mathbf{r}_{k+1}, \mathbf{r}_{k+1}) < \varepsilon^2(\mathbf{b}, \mathbf{b})$ **exit;**

$$\beta_k = \frac{(\mathbf{r}_0, \mathbf{r}_{k+1})}{\omega_k (\mathbf{r}_0, \mathbf{M}^{-1}\mathbf{A}\mathbf{p}_k)}$$


$$\mathbf{p}_{k+1} = \mathbf{r}_{k+1} + \beta_k (\mathbf{p}_k - \omega_k \mathbf{M}^{-1}\mathbf{A}\mathbf{p}_k)$$

loop end

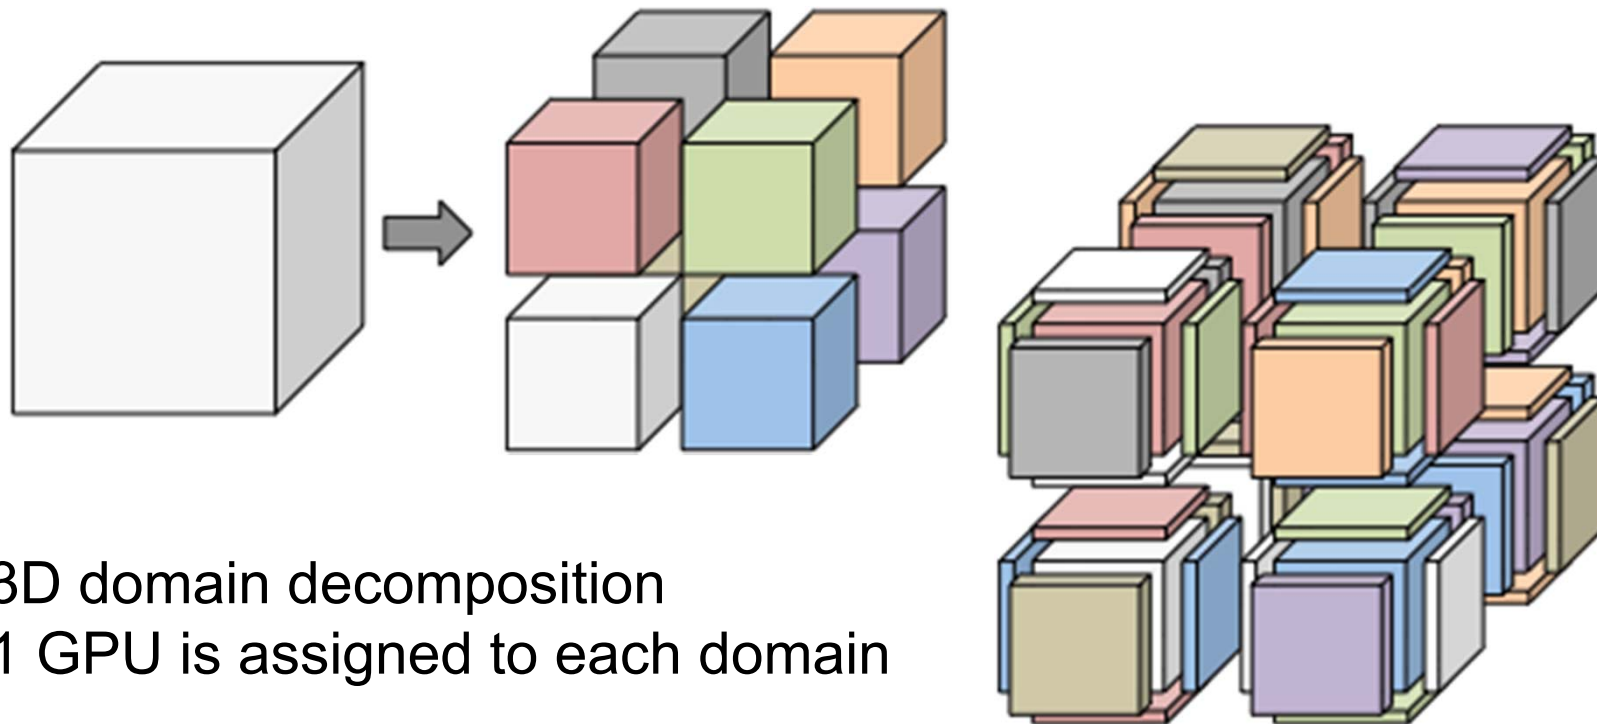

MG V-Cycle



GP GPU



Multi-Dimensional Domain Decomposition




- 3D domain decomposition
- 1 GPU is assigned to each domain


- Communication buffer for each face
- Host buffer & Device buffer

Milk Crown Formation


4.0 m/sec impact speed




Experiment




Simulation



Gunji, ishii, Saito, Sakai



Rayleigh-Taylor Instability with Surface Tension Force




When a heavy fluid is supported against gravity by a light fluid, a Rayleigh-Taylor instability develops in which perturbations of the interface grow exponentially in time as $\exp(nt)$ for small amplitudes.

$$n^2 = Kg \left[A - \frac{K^2 \sigma}{g(\rho_h - \rho_l)} \right]$$

$$A = \frac{\rho_h - \rho_l}{\rho_h + \rho_l}$$

$$\Phi = \frac{\sigma}{\sigma_c} = \frac{\sigma K^3}{g(\rho_h - \rho_l)}$$



$$\begin{aligned} \rho_l &= 0.25 \\ \rho_h &= 1.0 \\ K &= 1 \\ g &= 1 \\ \mu &= 0 \\ L_x &= 2\pi \\ L_y &= 6\pi \end{aligned}$$

Bellman, R., Pennington, R.H.: Effect of surface tension and viscosity on Taylor instability. *Q. Appl. Methods* **12**, 12, 151 (1954)


Drazin, P.G., Reid, W.H.: *Hydrodynamic Stability*. Cambridge University Press, Cambridge (1967)

Daly, B.J.: Numerical study of the effect of surface tension on interface instability. *Phys. Fluids* **12**, 1340 (1969)

Snapshots of the R-T Instability




GP GPU




1024x192 cells


Milk Crown



Drop on dry floor



Broken dam Problem



J.C.Martin and W.J. Moyce (1952)




FIGURE 1. Diagram of typical apparatus.




FIGURE 2. Two dimensional collapse of $n^2=1$ section.

Initial stages of dam-break flow

P.K.Stanby, A.Chegini and T.C.D.Barnes (1998)

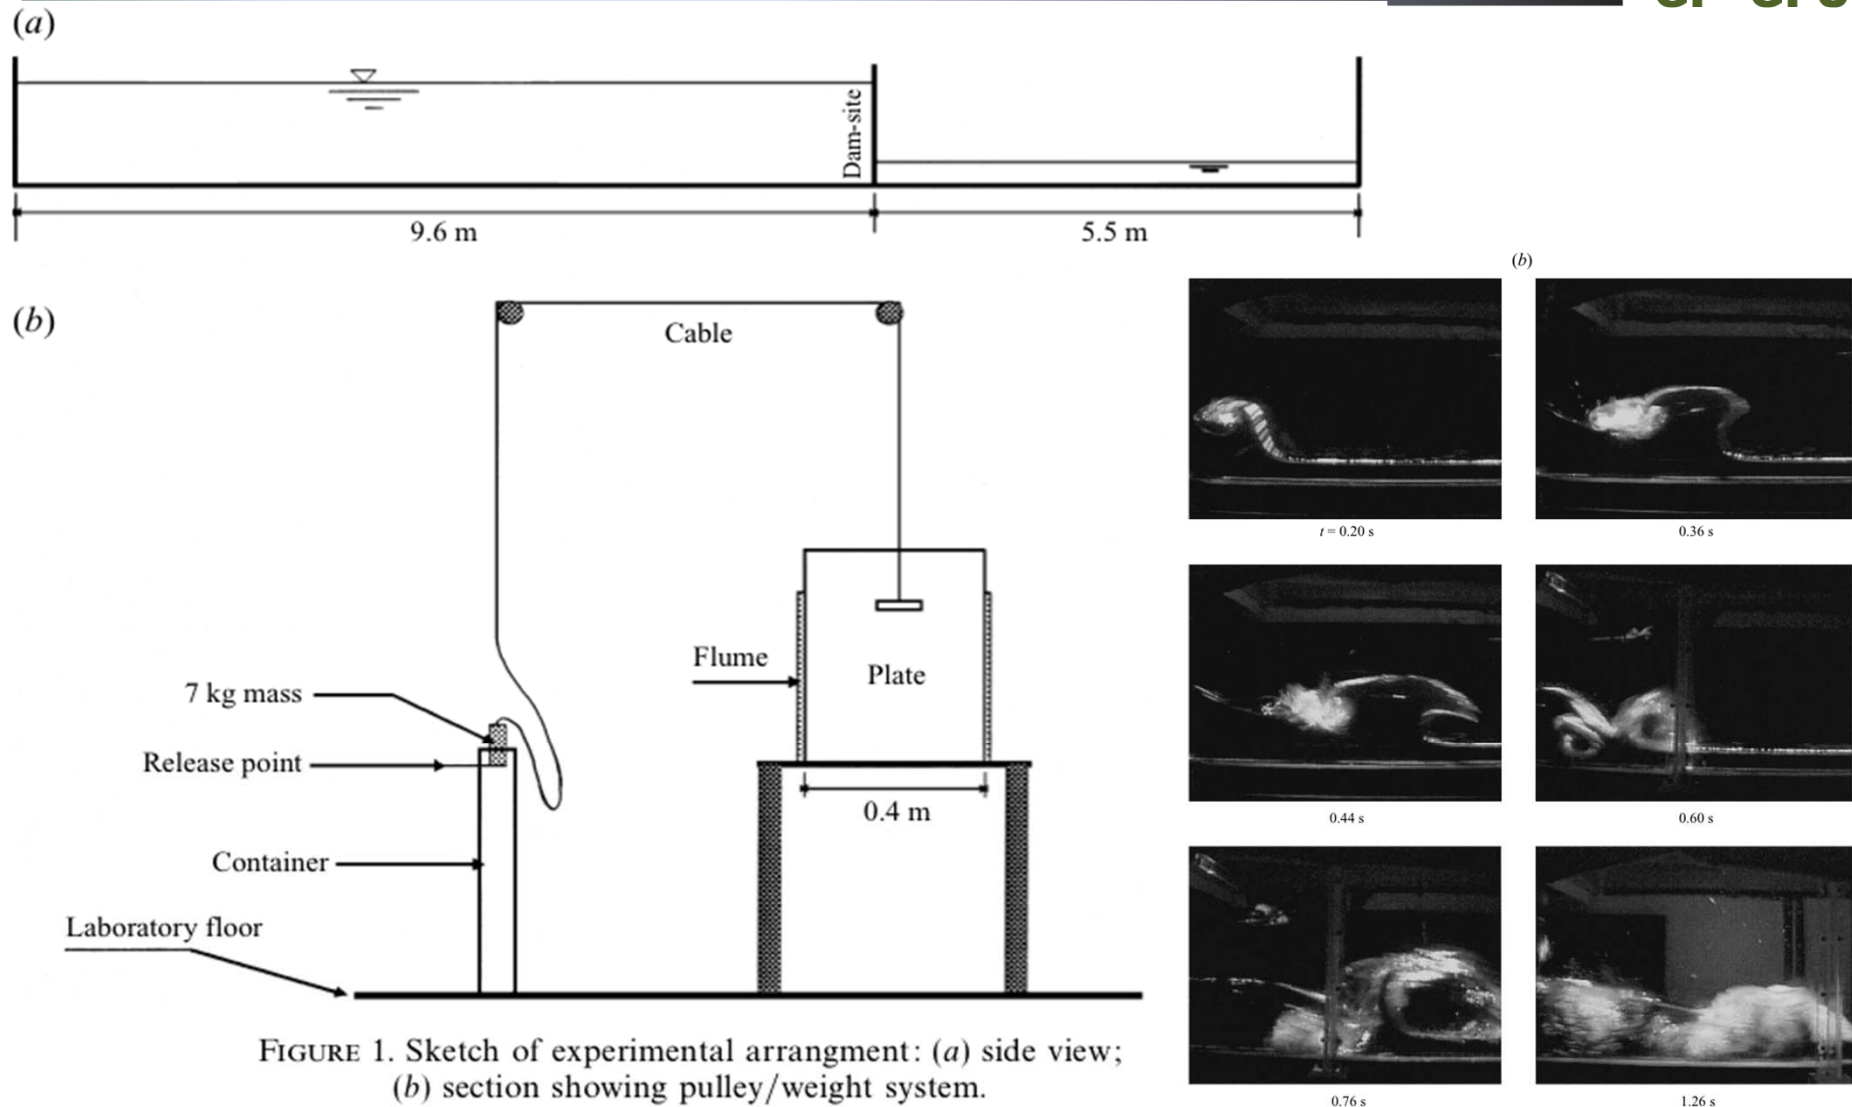










FIGURE 1. Sketch of experimental arrangement: (a) side view; (b) section showing pulley/weight system.







Experiment




Simulation







MULTI-GPU Performance



(TSUBAME 1.2)



Weather Prediction



Collaboration: Japan Meteorological Agency


Meso-scale Atmosphere Model:

Cloud Resolving Non-hydrostatic model

Compressible equation taking consideration of sound waves.



Atmosphere Model



Dynamical Process:

Full 3-D Navier-Stokes Equation


$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{1}{\rho} \nabla P - 2\boldsymbol{\Omega} \times \mathbf{u} - \boldsymbol{\Omega} \times (\boldsymbol{\Omega} \times \mathbf{r}) + \mathbf{g} + \mathbf{F}$$

Physical Process:

Cloud Physics, Moist, Solar Radiation, Condensation,
Latent heat release, Chemical Process, Boundary Layer

So called "Parameterization" including many empirical rules.

WRF GPU Computing



■ WRF (Weather Research and Forecast)

Community Code developed by NCAR, NCEP, OU, NOAA/FSL, AFWA

WSM5 (WRF Single Moment 5-tracer) Microphysics*

Represents condensation, precipitation and thermodynamic effects of latent heat release


1 % of lines of code, 25 % of elapsed time

⇒ 20 x boost in microphysics (1.2 - 1.3 x overall improvement)


WRF-Chem**

provides the capability to simulate chemistry and aerosols from cloud scales to regional

⇒ x 8.5 increase



Full GPU Implementation




■ ASUCA Production Code

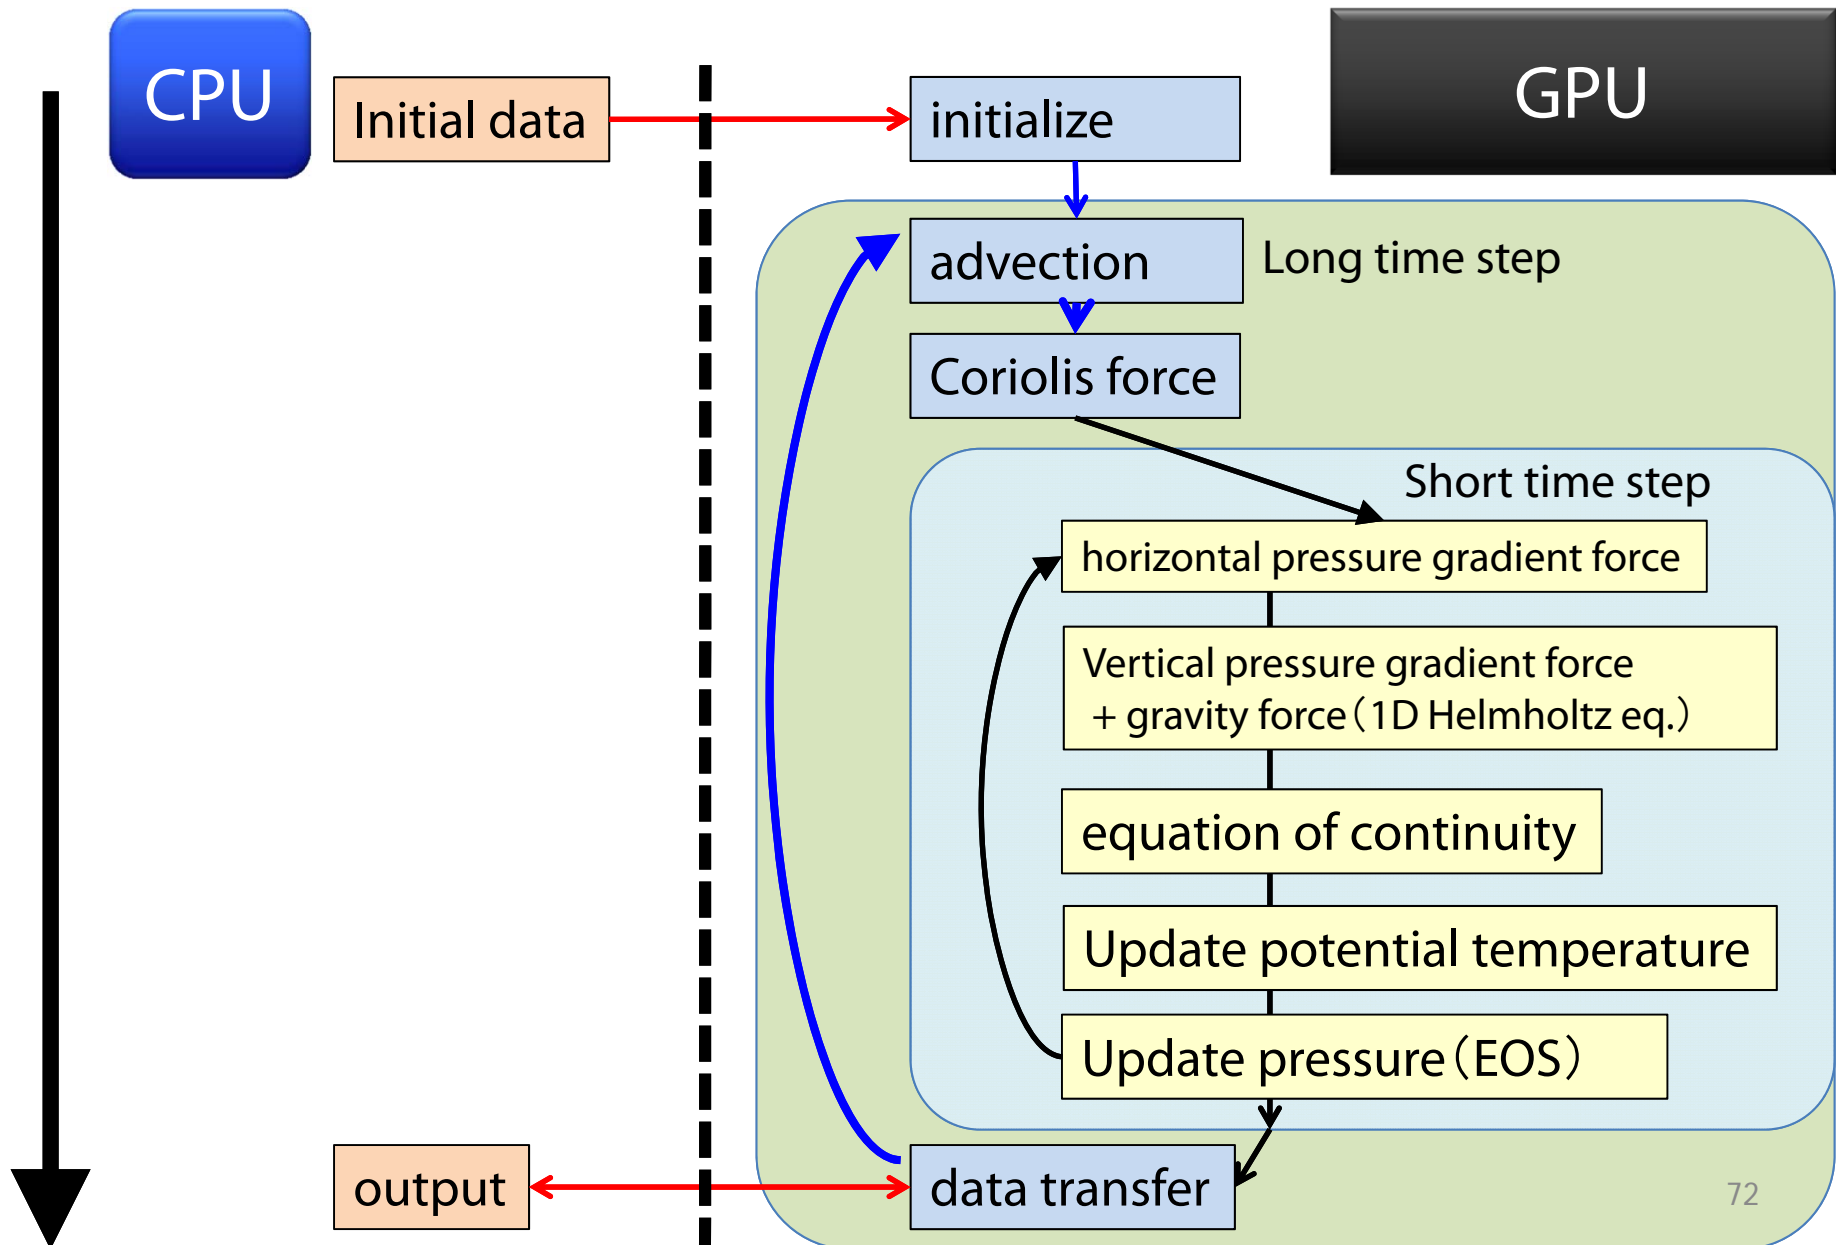
- ✓ A next-generation high resolution weather simulation code that is being developed by Japan Meteorological Agency (JMA)
- ✓ ASUCA succeeds the JMA-NHM as an operational non-hydrostatic regional model at JMA

■ Similar Structure as WRF


- ✓ HEVI (Horizontally explicit Vertical implicit) scheme
- ✓ Dynamical Core uses a numerical scheme with 3rd-order accuracy in time and space
 - Flux-form non-hydrostatic compressible equation
 - Generalized coordinate



Computational Flow of ASUCA



Entire Porting Fortran to CUDA



■ Rewrite from Scratch

Program init
implicit none

```
integer i  
integer a(10)  
do i = 1, 10  
  a(i) = i  
end do
```

✓ Original code
at JMA

```
#include <iostream>
```

```
int main()  
{  
  int i;  
  int a[10];  
  for(i=0;i<10;i++){  
    a[i] = i + 1;  
  }  
}
```

✓ Changing
array order

```
#include <cuda.h>  
__global__ void init(int  
*a){  
  a[threadIdx.x] =  
  threadIdx.x+1;  
}
```

CUDA

```
int main()  
{  
  int i;  
  int *a;  
  cudaMalloc(&a,sizeof(in  
t)*10);  
  init_1_10>>>(a);  
  cudaFree(a);  
}
```

✓ GPU code

z,x,y (k,i,j)-ordering


x,z,y (i,k,j)-ordering

x,z,y (i,k,j)-ordering

■ 1 Year

Introducing many optimizations, overlapping the computation with the communication, kernel fuse, reordering kernel execution


Implementation : Advection



Thread




Block




64 x 4 threads (2D) in a block

- Each thread specifies a (x, z) point, marching in y
 - ✓ Improve data transfer performance using domain decomposition



Using Shared Memory



- Shared Memory (SMem) = Software Managed Cache
- Read a 2D sub-domain from VRAM into SMem
- Advection : 12-point stencil
 - ✓ Store the xz-slice in $(64 + 3) \times (4 + 3)$ SMem

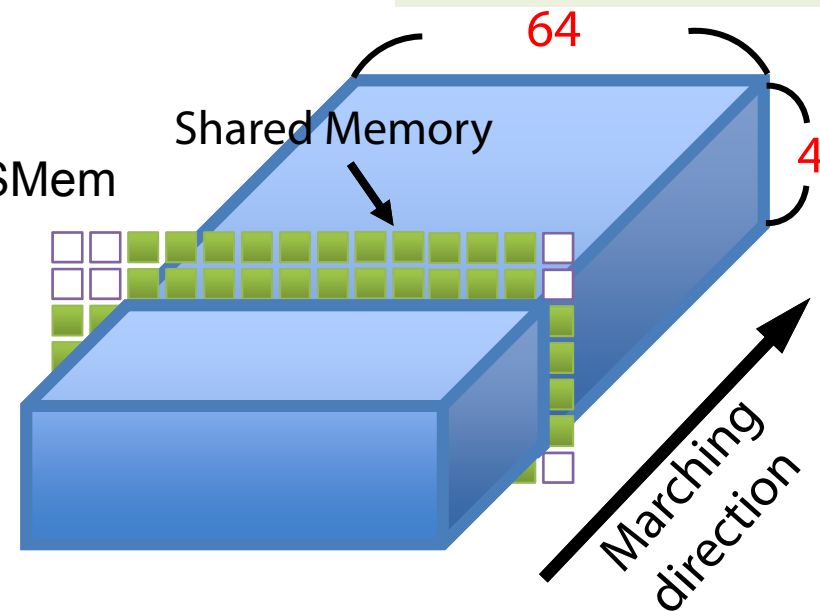
1 Block
= 64 x 4 threads

Access GMem directly : 4 + 4 read, 1 write



Using SMem : ~1


read, 1 write



- 2D sub-domain
- Halo
- Not in use

	Shared Memory	VRAM (Global Memory)
Access speed	~ 2 cycle	400-600 cycle
Capacity	16 kByte/Block	2 GByte (Total)

Using Registers in marching direction



■ Register

- ✓ Access speed : 1 cycle
- ✓ used for data not shared among threads

■ Advection : 12-point stencil


- ✓ Each thread keeps 4 y-elements in registers
- ✓ Elements are reuse

Access GMem directly : 4 + 4 + 4


read, 1write



Using SMem and Registers : ~1 read, 1write




Implementation : 1D Helmholtz equation




Thread 

Block

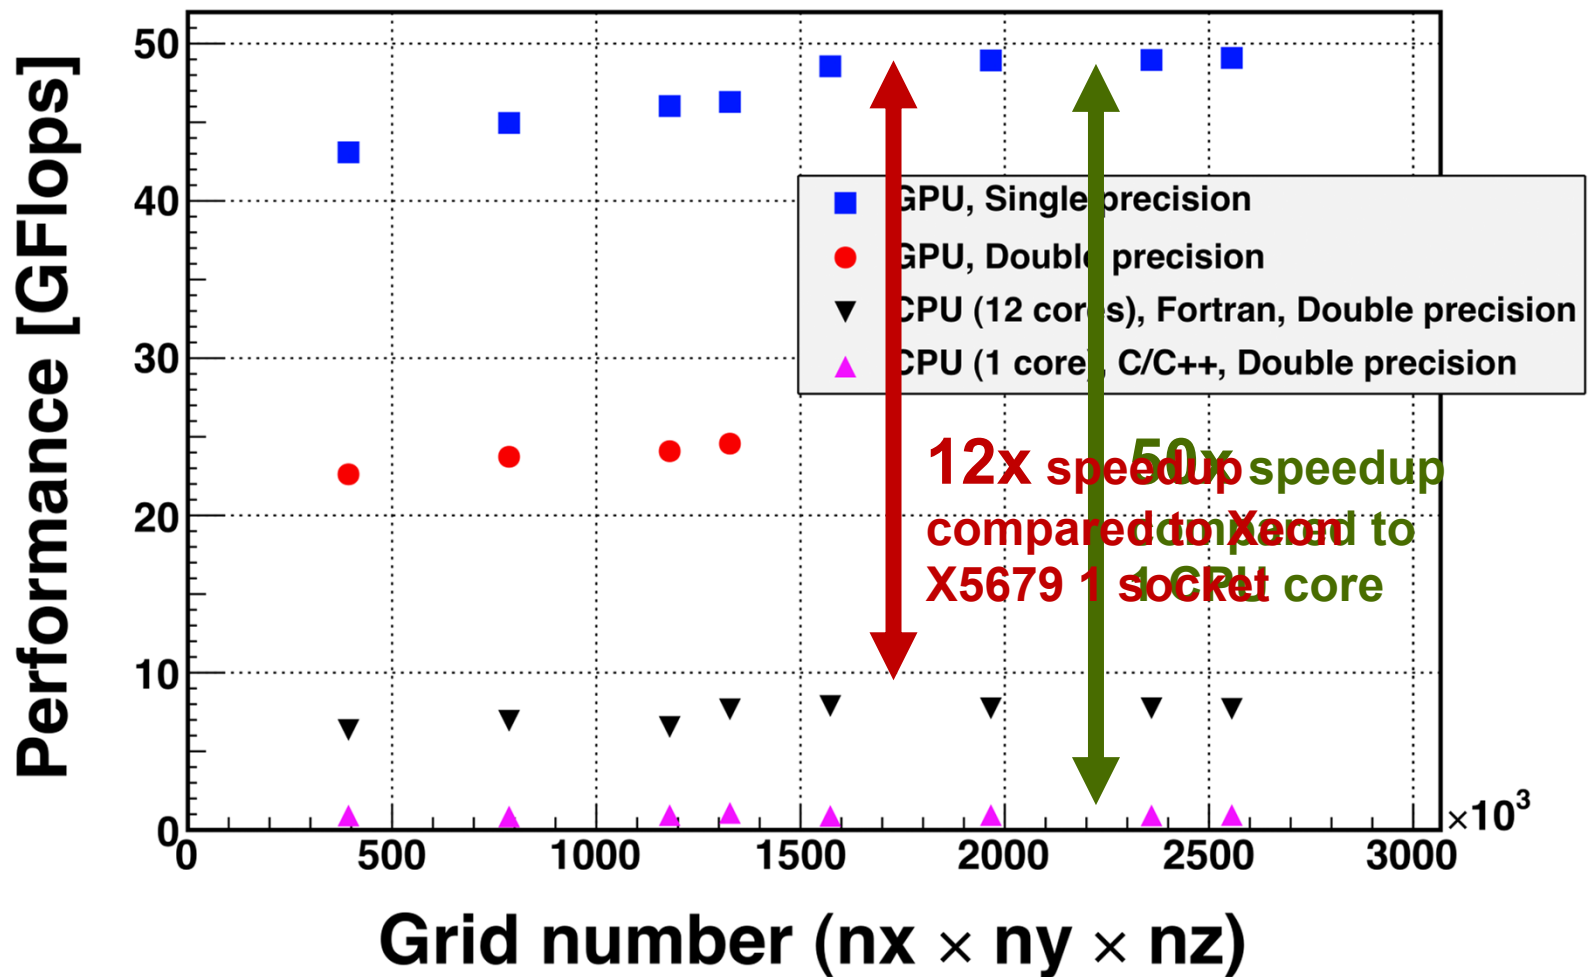



64 x 4 threads (2D) in a block

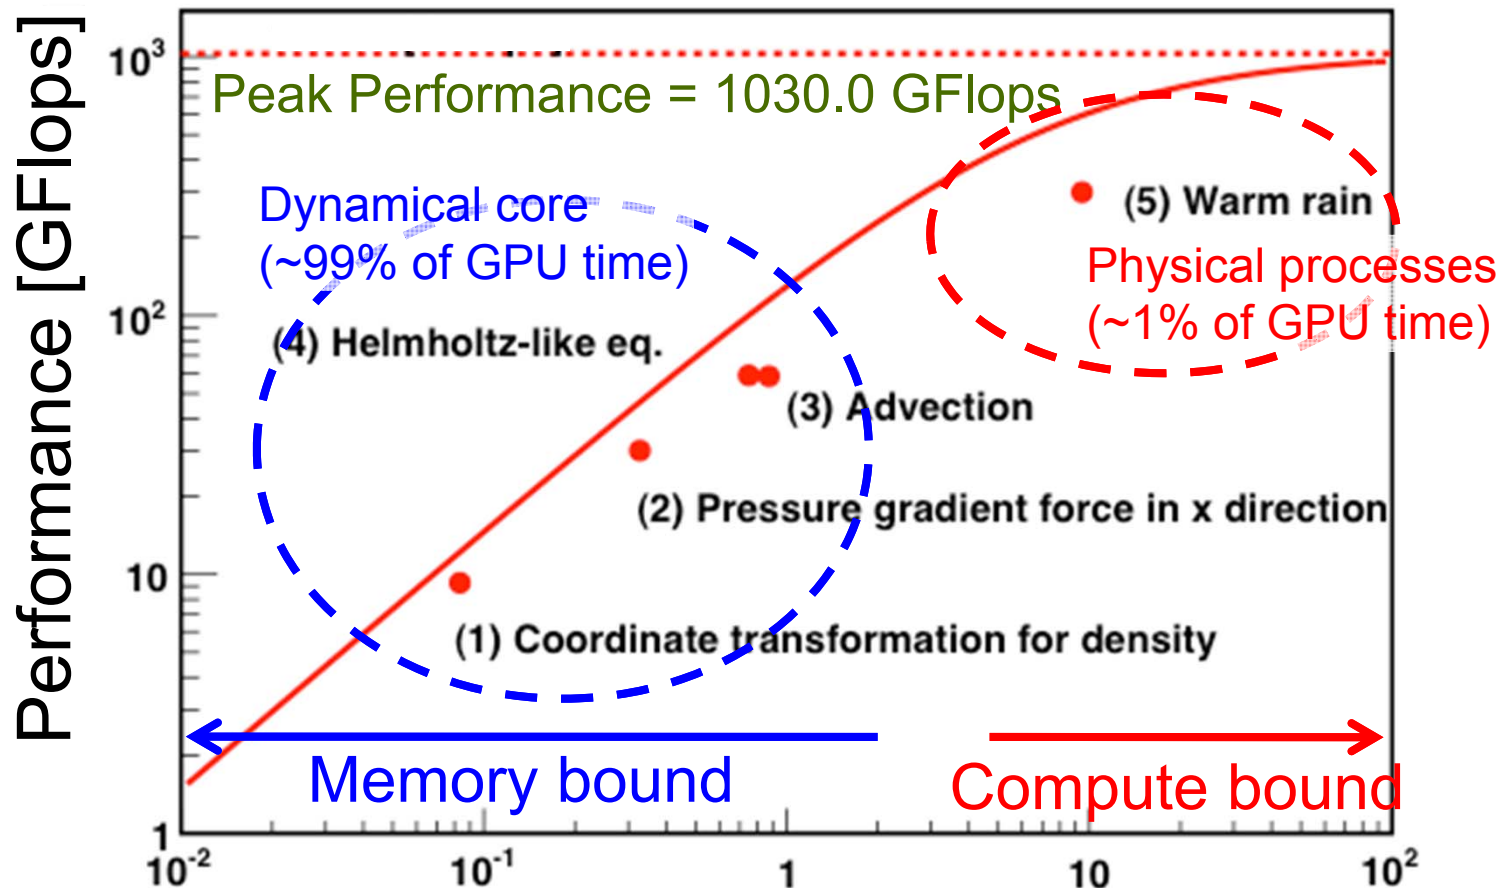

- 1D Helmholtz equation
 - ✓ Element in k depends on elements in $k \pm 1$
 - ⇒ marching in z direction



TSUBAME 2.0 (1 GPU)



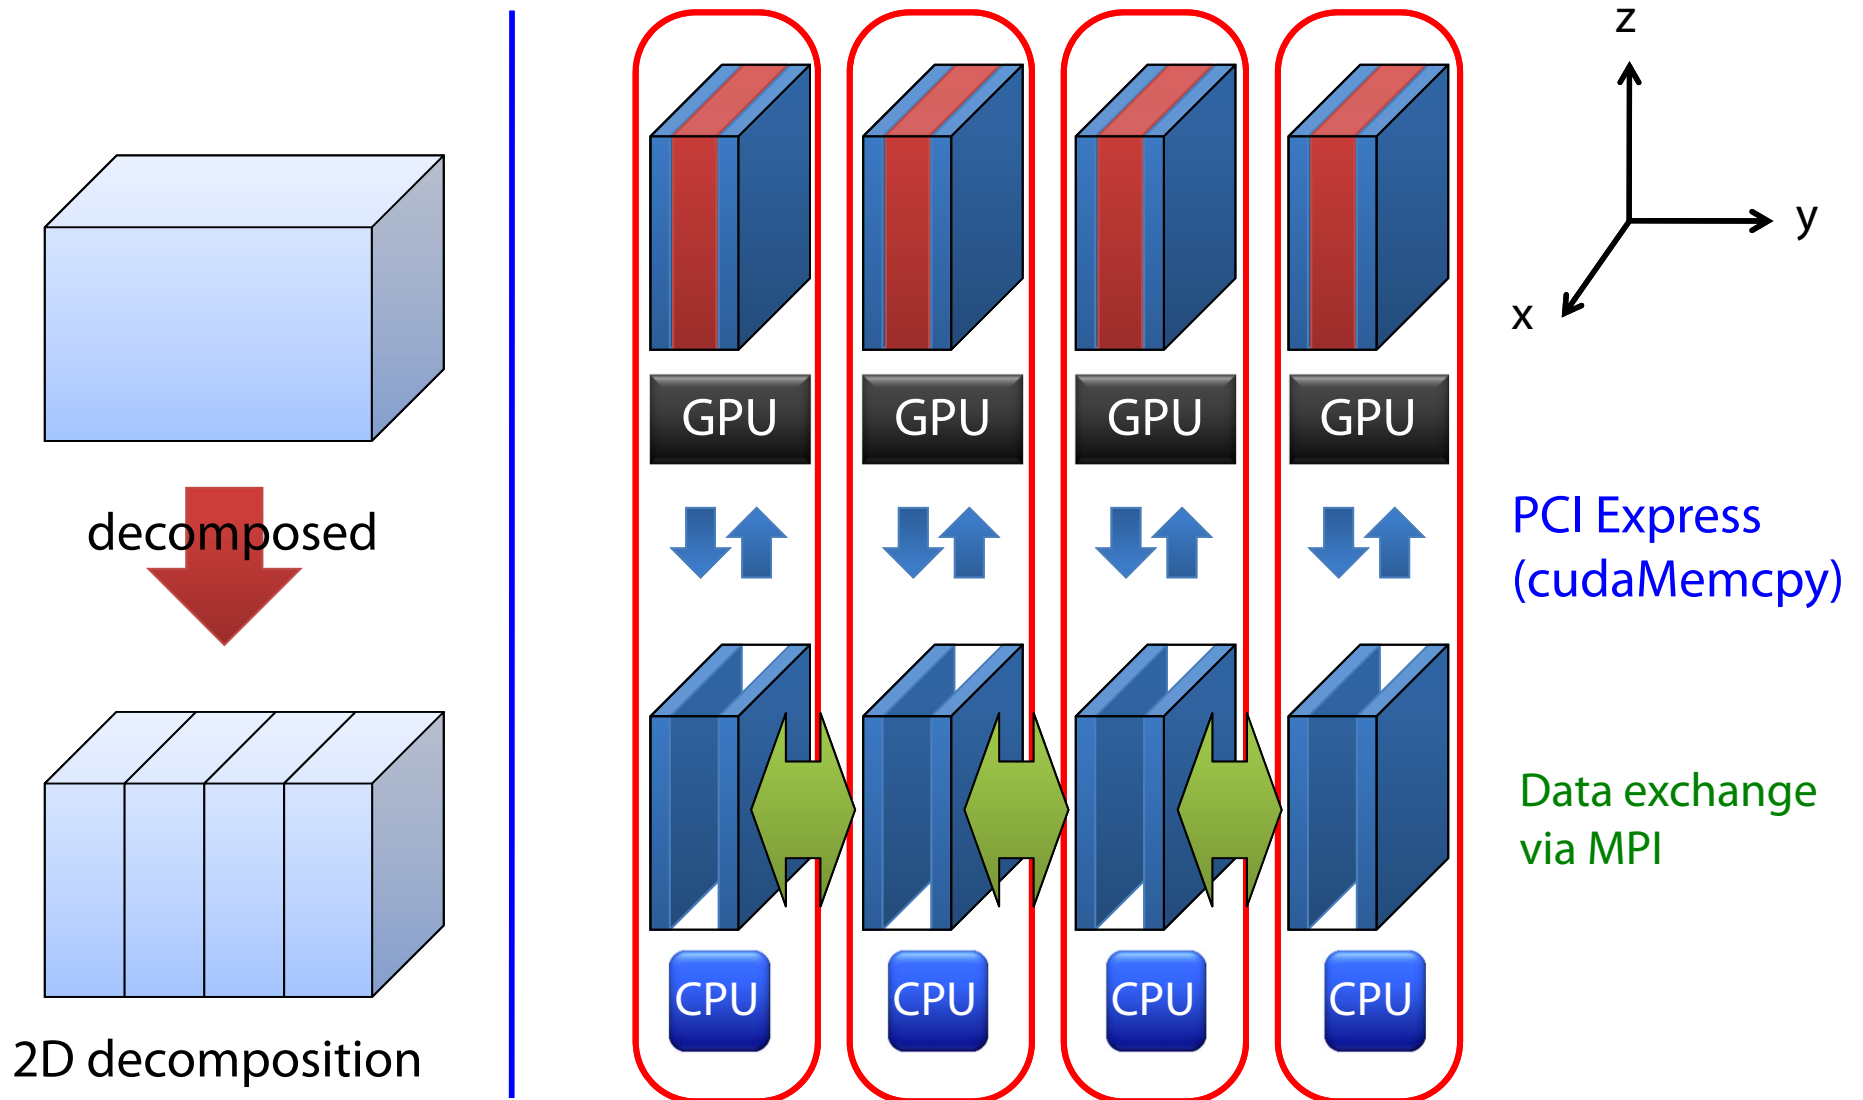

Performance of 5 kernels



$F_{peak} = 1030.0$ GFlops
 $B_{peak} = 148.0$ GByte/s


Arithmetic Intensity FLOP/Byte

Multi-GPU : Domain decomposition




ASUCA Typhoon Simulation

5km-horizontal resolution 479 × 466 × 48





ASUCA Typhoon Simulation

500m-horizontal resolution $4792 \times 4696 \times 48$



TSUBAME 2.0 Weak Scaling




145.0 Tflops
Single precision

76.1 Tflops
Double precision

Fermi core Tesla
M2050
3990 GPU

SUMMARY



FEATURES of GPU

High Performance and Low Power

■ Major differences from Previous Accelerators

ClearSpeed, Grape, , ,

High Memory Bandwidth

suitable for wide variety of applications

Consumer Product

inexpensive

Software Development Environment

CUDA, Open CL