

#### **Application Performances on Many-core Processors**

# Xeon Phi versus Kepler GPU





Takayuki Aoki

#### Tokyo Institute of Technology

Copyright © Takayuki Aoki , Global Scientific Information and Computing Center, Tokyo Institute of Technology

## **TOP 500 Ranking**

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)	
0	National University of Defense Technology China	Tianhe-2 (MilkyWay-2) - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P NUDT	3120000	33862.7	54902.4	17808	MIC Xeon F
0	DOE/SC/Oak Ridge National Laboratory United States	Titan - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560640	17590.0	27112.5	8209	GPU
0	DOE/NNSA/LLNL United States	Sequola - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1572864	17173.2	20132.7	7890	•
0	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIII6x 2.0GHz, Tofu interconnect Fujitsu	705024	10510.0	11280.4	12660	
0	DOE/SC/Argonne National Laboratory United States	Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786432	8586.6	10066.3	3945	
O	Texas Advanced Computing Center/Univ. of Texas United States	Stampede - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi SE10P Dell	462462	5168.1	8520.1	4510	MIC
0	Forschungszentrum Juelich (FZJ) Germany	JUQUEEN - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	458752	5008.9	5872.0	2301	
0	DOE/NNSA/LINL United States	Vulcan - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	393216	4293.3	5033.2	1972	
0	Leibniz Rechenzentrum Germany	SuperMUC - IDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR IBM	147456	2897.0	3185.1	3423	
10	National Supercomputing Center in Tianjin China	Tianhe-1A - NUDT YH MPP, Xeon X5670 6C 2.93 GHz, NVIDIA 2050 NUDT	186368	2566.0	4701.0	4040	GPU



MIC (eon Phi 31S1P GPU K20X

# **Processor SPECIFICATIONS**



	Xeon E5-	Xeon Phi	Tesla	Tesla	Tesla	GeForce
	2670	5110P	M2050	K20c	K20X	GTX Titan
Performance	333	2022	1030	3520 3.95		4500
(SP)	GFlops	GFlops	GFlops	GFlops GFlops		GFlops
Memory	51.2 GB/s	320 GB/s	148 GB/s	208 GB/s	250 GB/s	288 GB/s
Bandwidth		(ECC off)	(ECC off)	(ECC off)	(ECC off)	(nonECC)
Memory Size		8 GB	3 GB	5 GB	6 GB	6 GB
number of cores	8	60/61	448	2496	2688	2688
Clock Speed	2.6 GHz	1.053 GHz	1.15 GHz	0.706 GHz	0.732 GHz	0.837 GHz

Copyright © Takayuki Aoki , Global Scientific Information and Computing Center, Tokyo Institute of Technology





# **Memory Bandwidth**





## Performance of Diffusion Eq.



A primitive stencil application accessing neighbor 7 points using 2nd-order Finite Difference Method







# Xeon/Xeon Phi (1)



### Standard (simply directive) in the book (Intel)



Copyright © Takayuki Aoki , Global Scientific Information and Computing Center, Tokyo Institute of Technology





Peeling in the x-direction in the book (Intel)

#pragma omp parallel {
#pragma omp for collapse (2)
for (int jz = 0; jz < nz; jz++) {
for (int jy = 0; jy < ny; jy++) {
 jx = 0; . . .

#pragma simd
for (int jx = 1; jx < nx-1; jx++){
 j = nx\*ny\*jz + nx\*jy + jx;
 je = j+1; jw = j-1;
 jn = j+nx; js = j-nx;
 jt = j + nx\*ny; jb = j - nx\*ny;
 fn[j] = cc\*f[j]
 + ce\*f[je] + cw\*f[jw]
 + cn\*f[jn] + cs\*f[js]
 + ct\*f[jt] + cb\*f[jb];
 jx = nx - 1; . . .
 }
 }
 }
}</pre>

## Xeon/Xeon Phi (3)





Copyright © Takayuki Aoki , Global Scientific Information and Computing Center, Tokyo Institute of Technology

**GPU** code



### Standard

int jx = blockDim.x \* blockIdx.x + threadIdx.x; int jy = blockDim.y \* blockIdx.y + threadIdx.y; if (jx == nx-1) jxp = jx; if (jx == 0) jxm = jx; boundary
if (jy == ny-1) jyp = jy; if (jy == 0) jym = jy; condition #pragma unroll for (int jz=0; jz < nz; jz++) {
 jzm = (jz == 0) ? jz : jz-1;
 jzp = (jz == nz-1) ? jz : jz+1;
 boundary condition</pre> j = nx\*ny\*jz + nx\*jy + jx;j fn[j] = cc\*f[j]- cc\*f[je] + cw\*f[jw] + cn\*f[jn] + cs\*f[js] + ct\*f[jt] + cb\*f[jb]; S je ίw } in syncthreads(); jb

# **GPU tuning (1)**



### + Register reuse (reuse variable)





### Performances





Copyright © Takayuki Aoki , Global Scientific Information and Computing Center, Tokyo Institute of Technology

# **MIC Performances**





### **3-dimensional Euler Equation:**

$$\frac{\partial \rho}{\partial t} + \frac{\partial E}{\partial x} + \frac{\partial F}{\partial y} + \frac{\partial G}{\partial z} = 0$$

$$U = \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ e \end{pmatrix} \qquad E = \begin{pmatrix} \rho u \\ \rho uu + p \\ \rho vu \\ \rho wu \\ (e+p)u \end{pmatrix} \qquad F = \begin{pmatrix} \rho v \\ \rho uv \\ \rho uv \\ \rho vv + p \\ \rho wv \\ (e+p)v \end{pmatrix} \qquad G = \begin{pmatrix} \rho w \\ \rho uw \\ \rho uw \\ \rho vw \\ \rho ww + p \\ (e+p)w \end{pmatrix}$$

### **Conservative form**

Copyright © Takayuki Aoki , Global Scientific Information and Computing Center, Tokyo Institute of Technology

### **CPU, MIC, GPU Performances**



### **Kepler GPU Tuning**

- Shared Memory Use in the x-directional kernel.
- Super function unit
- Loop unrolling
- Variable reuse in the y- and z- loops
- Reduction of branch diverges



### MIC Performances of 3D Compressible Flow Computation



