




GPUコンピューティング No.1

導入

東京工業大学 学術国際情報センター

青木 尊之

GPUとは



GPGPU



GP GPU

(General-purpose computing on graphics processing units)

GPU を画像処理以外の一般的計算に使う

GPUの魅力

- 高性能: ハイエンド GPU はピーク **4 TFLOPS** 超
- 低価格: ハイエンドでもコンシューマタイプは **数万円**
- 手軽さ: 普通のPCにも装着できる
- プログラミング開発: 無償の開発環境

CPUと比較して単一
GPUは高消費電力



低消費電力: **FIOPS/W**

講義を受ける目的



GP GPU

- 既存のコードを GPU 化して高速に実行したい
- 新たに GPU プログラムを開発し、研究を促進したい
- これから主流となるであろう GPU のプログラミングをマスターしたい
- 超並列計算を習得したい
- 単位が欲しい




その、きっかけを得る

ショッキングなGPUの計算性能




レーリー-テラー不安定性成長



Video captured demonstration

$$\frac{\partial Q}{\partial t} + \frac{\partial E}{\partial x} + \frac{\partial F}{\partial y} = 0$$

$$Q = \begin{bmatrix} \rho \\ \rho u \\ \rho v \\ e \end{bmatrix} \quad E = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ eu + pu \end{bmatrix} \quad F = \begin{bmatrix} \rho v \\ \rho uv \\ \rho v^2 + p \\ ev + pv \end{bmatrix}$$



Core2 duo 1 core

GeForce GTX 260M



Y. Imai, T. Aoki and K. Takizawa, J. Comp. Phys., Vol. 227, Issue 4, 2263-2285 (2008)



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

TSUBAME 2.0

System (58 racks)

1442 nodes: 2952 CPU sockets,
4264 GPUs

Performance: 224.7 TFLOPS (CPU) ※ Turbo boost
2196 TFLOPS (GPU)

Total: **2420** TFLOPS

Memory: 103.9 TB



GP GPU

Rack (30 nodes)

Performance: 51.0 TFLOPS
Memory: 2.03 TB


Compute Node

(2 CPUs, 3 GPUs)

Performance: 1.7 TFLOPS
Memory: 58.0GB(CPU)
+9.7GB(GPU)



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




ORNL Jaguar vs Tsubame 2.0

Similar Peak Performance, 1/4 the Size and Power



GP GPU



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



科学と技術で未来を創造する

Supercomputer in the world



The Green500 list -- November 2010

Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)
1	1684.20	IBM Thomas J. Watson Research Center	NNSA/SC Blue Gene/Q Prototype	38.80
2+	1448.03	National Astronomical Observatory of Japan	GRAPE-DR accelerator Cluster, Infiniband	24.59
2	958.35	GSIC Center, Tokyo Institute of Technology	HP ProLiant SL390s G7 Xeon 6C X5670, Nvidia GPU, Linux/Windows	1243.80
3	933.06	NCSA	Hybrid Cluster Core i3 2.93Ghz Dual Core, NVIDIA C2050, Infiniband	36.00
4	828.67	RIKEN Advanced Institute for Computational Science	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	57.96
5	773.38	Universitaet Wuppertal	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54
5	773.38	Universitaet Regensburg	QPACE SFB TR Cluster, PowerXCell 8i, 3.2 GHz, 3D-Torus	57.54

Supercomputer in the world



2012 November

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	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
2	DOE/NNSA/LLNL United States	Sequoia - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1572864	16324.8	20132.7	7890
3	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 Vllifx 2.0GHz, Tofu interconnect Fujitsu	705024	10510.0	11280.4	12660
4	DOE/SC/Argonne National Laboratory United States	Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786432	8162.4	10066.3	3945
5	Forschungszentrum Juelich (FZJ) Germany	JUQUEEN - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	393216	4141.2	5033.2	1970
6	Leibniz Rechenzentrum Germany	SuperMUC - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR IBM	147456	2897.0	3185.1	3423
7	Texas Advanced Computing Center/Univ. of Texas United States	Stampede - PowerEdge C8220, Xeon E5- 2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi Dell	204900	2660.3	3959.0	
8	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

CPU/GPU Spec Sheet



		Intel Xeon X5670	Tesla C2050 /M2050	GeForce GTX Titan
GPU	Peak Performance [GFlops]	76.8*,153.6	515*,1030	1.3T*,4.5T
	Number of Processor	6	448	2688
	Core Clock [GHz]	2930	1150	837
Memory	Bandwidth[GB/s]	32.0	148.8	288.4
	Memory Interface [bit]	64	384	384
	Memory Clock [GHz]	1.333 (DDR3)	1.50 (GDDR5)	1.50 (GDDR5)
B _{peak} /F _{peak}	Bandwidth/Performance	0.416	0.289	0.221



Tesla M2050
Peak Power : 225W




Peak Power : 244W

GPUアーキテクチャの変更

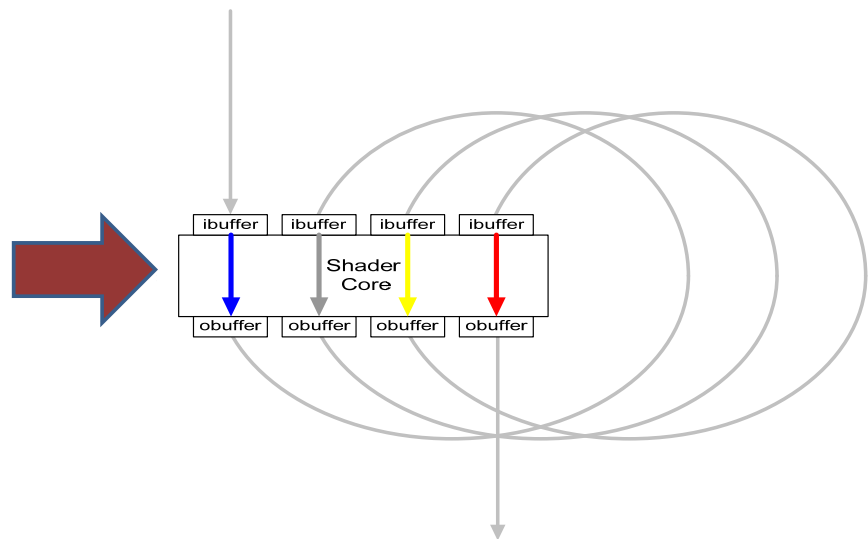


GP GPU

Graphics Pipeline



Unified Shader



Shader 言語



GP GPU

Unified Shader: プログラマブル・シェーダー

OpenGLやDirectX などのAPIに専用のプログラマブルなシェーディング機能

Open GL では version 1.5, DirectX では version 8 から

Shader プログラミング言語

OpenGL: GLSL 言語

DirectX: HLSL 言語

NVIDIA 独自の Cg (C for Graphics) 言語 (HLSL 似)

汎用計算を Graphics の機能に置き換えてプログラミング

TSUBAME に login



GP GPU

Windows 端末の Bash Shell から

```
$ ssh user_account@login-t2.g.gsic.titech.ac.jp
user_account@login-t2.g.gsic.titech.ac.jp's password:
```

インストールされているCUDA のバージョンの確認

```
/opt/cuda/3.0 3.1 3.2 4.0 4.1 5.0
が置いてある
```

現在のTSUBAMEには最新の **CUDA 5.0** がインストールされている。

CUDA 5.0



GP GPU

```
$ cd /opt/cuda/5.0
$ sh cuda.sh // 環境設定
```

CUDA コンパイラ `nvcc` のバージョンの確認

```
user_account@t2a006169:~> nvcc --version
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2012 NVIDIA Corporation
Built on Fri_Sep_21_17:28:58_PDT_2012
Cuda compilation tools, release 5.0, V0.2.1221
```


DeviceQuery



GP GPU

```
$ cd /opt/cuda/5.0/samples/1_Utillities/deviceQuery>  
$ ./deviceQuery
```

```
./deviceQuery Starting...  
  
CUDA Device Query (Runtime API) version (CUDA static linking)  
  
Detected 3 CUDA Capable device(s)  
  
Device 0: "Tesla M2050"  
  CUDA Driver Version / Runtime Version      5.0 / 5.0  
  CUDA Capability Major/Minor version number: 2.0  
  Total amount of global memory:            2687 MBytes (2817982464 bytes)  
  (14) Multiprocessors x ( 32) CUDA Cores/MP: 448 CUDA Cores  
  GPU Clock rate:                           1147 MHz (1.15 GHz)  
  Memory Clock rate:                        1566 Mhz  
  Memory Bus Width:                         384-bit  
  L2 Cache Size:                            786432 bytes  
  Max Texture Dimension Size (x,y,z)        1D=(65536), 2D=(65536,65535),  
                                              3D=(2048,2048,2048)  
  Max Layered Texture Size (dim) x layers    1D=(16384) x 2048,  
                                              2D=(16384,16384) x 2048  
  Total amount of constant memory:          65536 bytes
```

DeviceQuery



GP GPU

```
Total amount of shared memory per block:    49152 bytes  
Total number of registers available per block: 32768  
Warp size:                                   32  
Maximum number of threads per multiprocessor: 1536  
Maximum number of threads per block:        1024  
Maximum sizes of each dimension of a block:  1024 x 1024 x 64  
Maximum sizes of each dimension of a grid:   65535 x 65535 x  
                                              65535  
Maximum memory pitch:                        2147483647 bytes  
Texture alignment:                           512 bytes  
Concurrent copy and kernel execution:        Yes with 2 copy  
                                              engine(s)  
Run time limit on kernels:                   No  
Integrated GPU sharing Host Memory:          No  
Support host page-locked memory mapping:     Yes  
Alignment requirement for Surfaces:          Yes  
Device has ECC support:                      Enabled  
Device supports Unified Addressing (UVA):     Yes  
Device PCI Bus ID / PCI location ID:         6 / 0
```