


A GPU-based Architecture for Real-Time Data Assessment at Synchrotron Experiments

Agenda: Problem Description,
Performance Evaluation,
What do we do to run faster?
Hardware Evaluation



Goal: Process up to 30GB of image Data in a minute (*1 Tflop/s required*)




Mosquito head



In collaboration with ESRF:
European Synchrotron Radiation Facility
Polygone Scientifique Louis Néel, 6 rue Jules Horowitz, 38000 GRENOBLE

Solution: GPU Computing

- ♦ Highly parallel architecture: up to 480 parallel threads in the current NVIDIA architecture
- ♦ Up to 16 GPU cores in a single PC
- ♦ Rich set of libraries (FFT, BLAS, Lapack)
- ♦ OpenCL is an industry standard.
- ♦ Mass market: rapid development and cheap prices (cards are below 300 EUR)



AMD Opeteron 6176 (1500\$)

Number of cores: 12
 Frequency of cores: 2.3 Ghz
 Peak performance: 110.4/55.2 GFlop/s




NVIDIA GeForce GTX480 (500\$)

Parallel processors: 480
 Frequency of cores: 1.3 GHz
 Peak performance: 1345 / 168(672) GFlops




ATI Radeon HD 5970 (~700\$)

Parallel processors: 2 x 1600
 Frequency of cores: 725 MHz
 Peak performance: 1640 / 928 GFlops



GPU Servers



Asus Rampage III Extreme (1800\$ = Core i5 + 2 x GTX295 + 8 GB)

Chipset: x58, 36 PCIe 2.0 lanes; 6 DDR3 slots (48 GB max)

PCIe 2.0 x16: 4 (x8 if all 4 are used)

Max Peak Performance (ATI): 18.56 Tflops / 3.7 Tflops

Max Peak Performance (NVidia): 7.15 Tflops / 595 GFlops

Standard Desktop (800\$ = Core2 + GTX285 + 2 GB)

SuperMicro 7046GT-TRF (~8000\$ = 2 Xeon + 4 GPU + 96 GB)

Chipset: Dual Intel 5520, 72 PCIe 2.0 lanes, 12 DDR3 slots (192GB max)

PCIe 2.0 x16: 4 (full speed), x4: 2 (in x16 slots); PCIe 1.0 x4: (in x8 slot)

Max Peak Performance (ATI): 18.56 Tflops / 3.7 Tflops

Max Peak Performance (NVidia): 7.15 Tflops / 2.5 Tflops

NVIDIA Tesla S1070 (~8000\$ + 4000\$ host)

System: Requires separate host server

GPU Devices: 4 x Tesla C1060 (960 parallel processors at 1.44 GHz)

Peak performance: 4.14 Tflops / 345 Gflops

GPU Memory Size: 16 GB

Dual Xeon 5472 Server (5500\$ = 16 GB)

Max Peak Performance: 96 Gflops / 48 GFlops

Coolfax CXT8000 (36000\$ = 2 Xeon + 8 Tesla C2050 + 144GB)

Chipset: Dual Intel 5520, 72 PCIe 2.0 lanes

PCIe Switch: PLX PEX8647; PCIe 2.0 x16: 8 (full speed)

18 DDR3 Memory Slots: 288 GB max

Max Peak Performance: 10 Tflops / 5 Tflops

Back Projection Performance: 2336 Gflops (Estimated)

Sample Data-Set

Porose polyethylene grains in a conical plastic holder

Sample Data Set:

Number of Projections:	2000
Resolution of Projections:	1776x1707
Total Size:	24GB

Resulting 3D Image:



Resolution:	1691x1331x1311
Total Size:	11GB

Required Operations:


Data I/O:	35GB
Filtering:	~ 0.6 TFlop
Back Projection:	~ 53.0 TFlop

Maximum Performance Estimation:


Opteron 6176 (~ 110Gflops)	~ 10 minutes
HDD I/O (~ 50 MB/s)	~ 15 minutes



Test Setup and Results





	2 x Xeon E5472	GPU / CUDA	GPU Server / CUDA	Intel SSD X25E
Type of Computation	CPU / Xeon E5472 8 core, 3 GHz	NVIDIA GTX 280 1 core	GTX295, GTX480 6 cores	GPU / CUDA 240 cores
CPU	2 x Xeon E5472	Core2 E6300	2 x Xeon E5540	2 x Xeon E5540
Memory	16GB DDR3	4GB DDR2	96GB DDR3	96GB DDR3
HDD	WDC5000AACS	WDC5000AACS	WDC5000AACS	Intel X25-E
Price	5500\$ (2000\$ CPU)	1000\$ (400\$ GPU)	8000\$ (2000\$ GPU)	9000\$ (800\$ SSD)



Back Projection using Tesla S1070

Handling I/O Problem

Ratio between
I/O and Computations



Next steps:

- ▶ PCIe x16 card with SSD raid (up 1TB/s)
- ▶ Read from frame-grabber directly into the memory and process
- ▶ FPGA based frame reject and compression of image data



OCZ Z-Drive
Size: 1 TB, MLC
Read: 870MB/s
Write: 780MB/s


Conclusion

Goal: Process up to 30GB of image data in a minute

I/O

6 minutes

We will try disk-less in-memory processing to bring time under one minute.



Computations

40 seconds

Down from 30 minutes to less than a minute.
The computation problem is solved. ✓

Data Transfers is Dominating