

Agenda

Evaluation of GPU platforms and components



Goal: Process up to 1GB/sec of image data (*1 Tflop/s required*)



In collaboration with ESRF:

European Synchrotron Radiation Facility

Polygone Scientifique Louis Néel, 6 rue Jules Horowitz, 38000 GRENOBLE


GPU Computing Hardware

Goal: Real-time data processing

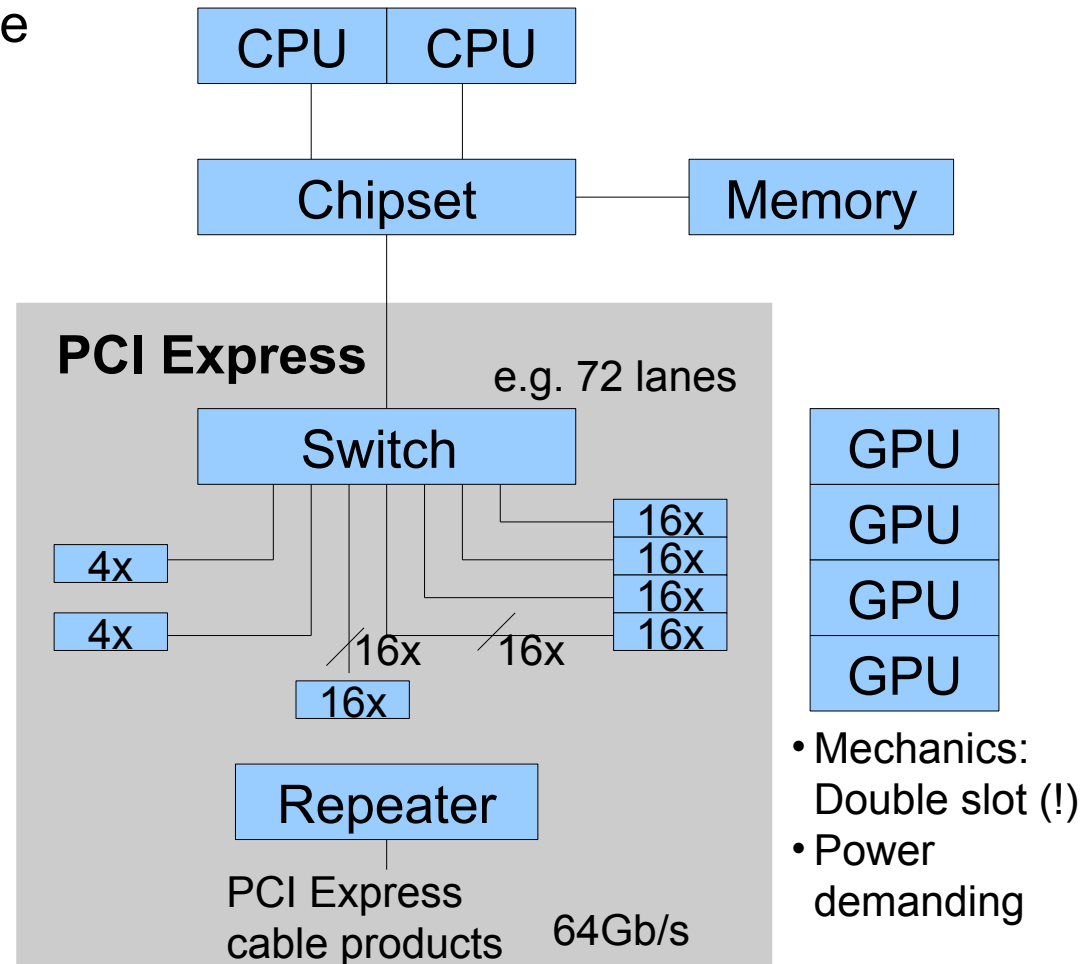
- PCI Express: 500MB/s / lane
 - Gen3 in 2011
- PCI Express external
- Mainboard types:
 - Server
 - Super computing
 - Gaming

Nvidia GPUs
CPUs
OpenCL
ATI GPUs
External PCIe
Storage

e.g. 6Gb/s



e.g. 10Gb/s



Sample Reconstruction Station



SuperMicro 7046GT-TRF (Dual Intel 5520 Chipset)
 CPU: 2 x Xeon X5650 (total 12 cores at 2.66 Ghz)
 GPUs: 2 x GTX 580 + 4 x GTX580 External
 Memory: 96 GB / 12 DDR3 slots (192GB max)

PCIe 2 x16 (8 GB/s):

PCIe 2 x16 (8 GB/s)

PCIe 2 x16 (8 GB/s):

PCIe 2 x16 (8 GB/s):

PCIe 2 x4 (2 GB/s):

PCIe 2 x4 (2 GB/s):

PCIe 1 x4 (1 GB/s):



2 x GTX 580

Measured bandwidth:
 ~ 5.7 GB/s to device
 ~ 6.3 GB/s from device



4 x GTX580



PCIe External



SAS Raid
 Areca ARC-1880x

2xSFF8088



16 x A7K2000
 ~ 1.6 GB/s



10 Gbit Net
 Intel 82598EB



SSD Raid
 4 x C300, 1420 MB/s



Frame Grabber
 Silicon Software, 850 MB/s

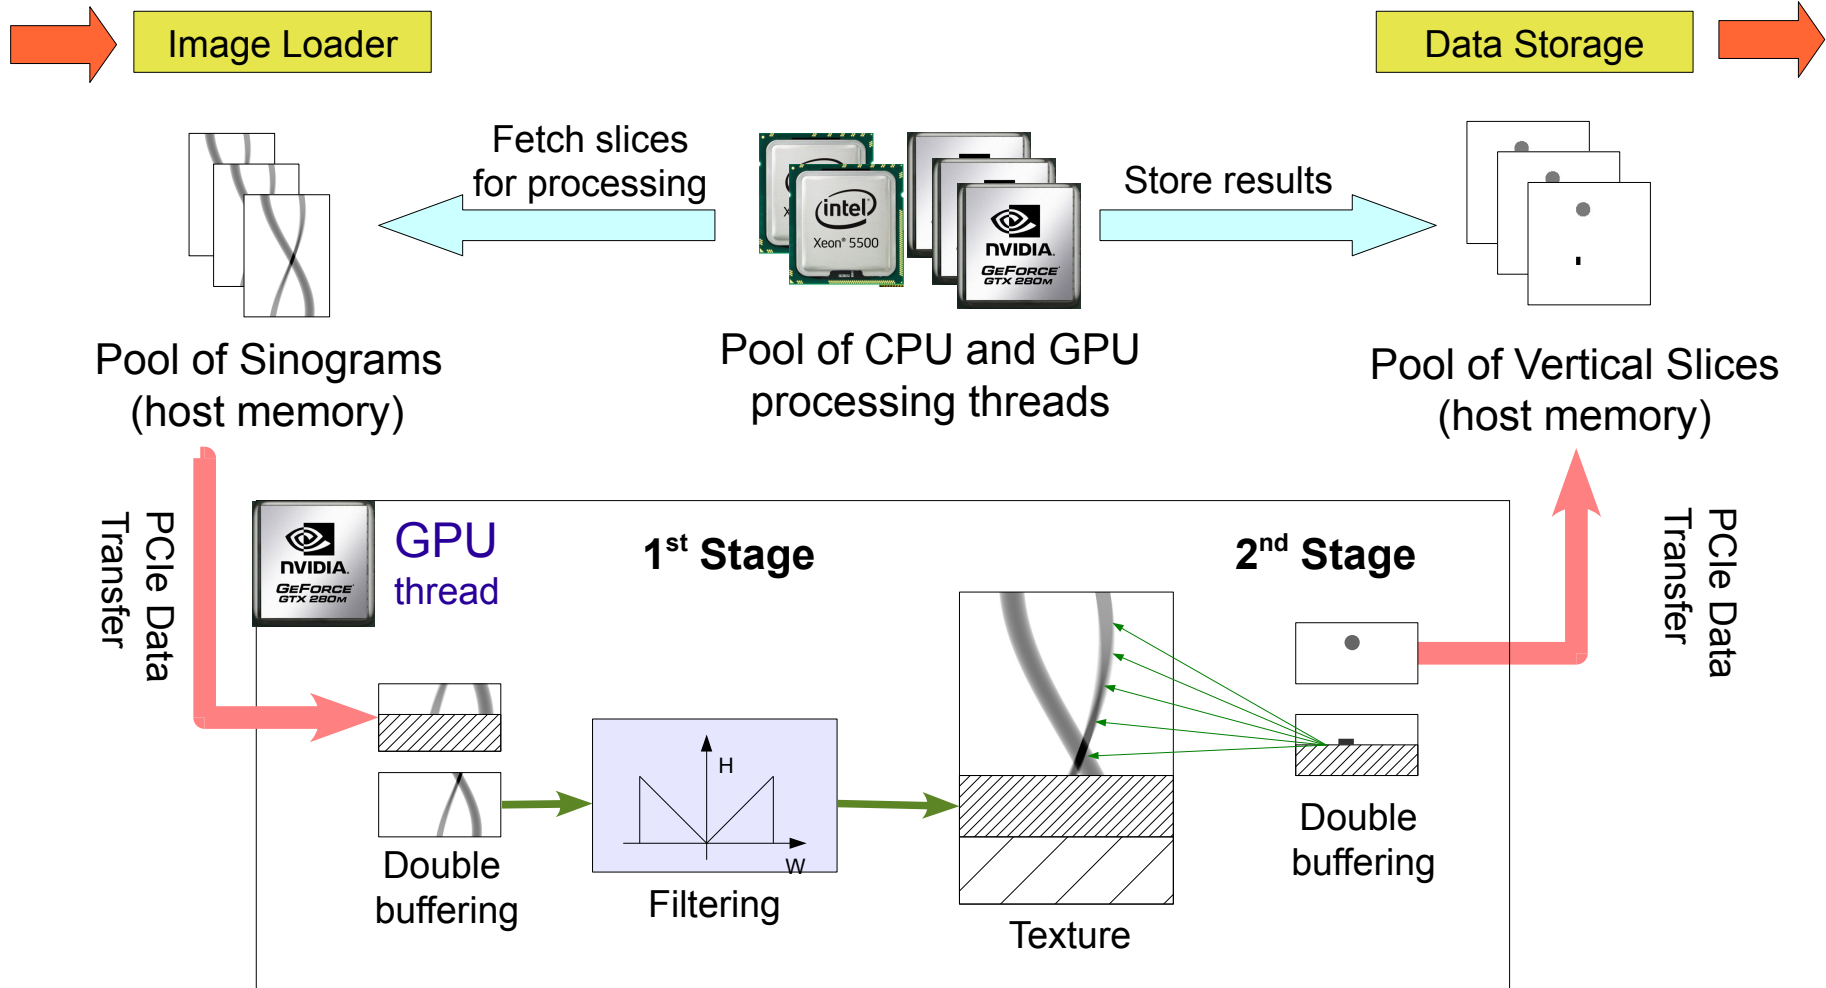
Camera
 Link



PCO
 edge

LSDF
 Large Scale Data Facility

PyHST: FBP Implementation



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




Performance: GPU vs. CPU

	Xeon Server	GPU Desktop	GPU Server
Type of Computation	CPU / Xeon X5650 12 cores, 2.66 GHz	GeForce GTX 280 1 core	2 x GTX295 + 2 x GTX580 6 cores
CPU	2 x Xeon E5650	Core2 E6300	2 x Xeon E5540
Memory	16GB DDR3	4GB DDR2	96GB DDR3
HDD/SSD	Hitachi A7K2000	2 x Intel X25-E	4 x Crucial RealSSD C300
Price	5500\$ (2000\$ CPUs)	1500\$ (400\$ GPU)	9000\$ (2000\$ GPU, 1200\$ SSD)
Software	SuSe 11.3, CUDA 3.2, MKL 10.2.1, gcc4.5 -O3 -march=nocona -mfpmath=sse		

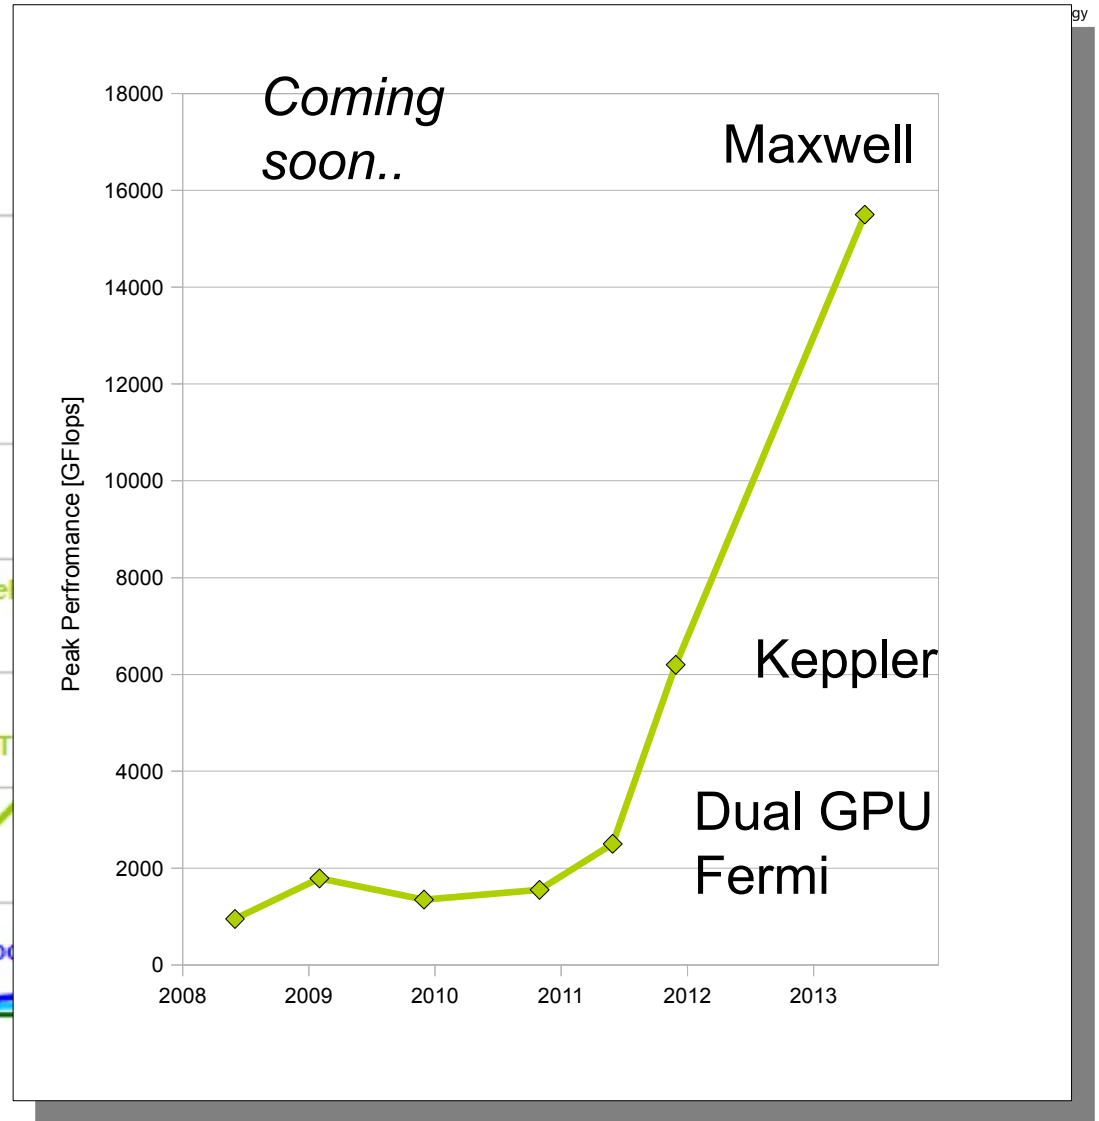



Nvidia GPU Family




Source: nvidia CUDA programmers guide

Nvidia GPU Family




Source: nvidia CUDA programmers guide

Performance: NVIDIA Desktop Products







Ratio of Operations




■ Transfer
■ Filtering
■ BP

	GTX280	GTX295	GTX480	GTX580
Architecture	GT200	GT200	Fermi	Fermi
Processors	1 x 240 @ 1.3 GHz	2 x 240 @ 1.3 GHz	1 x 480 @ 1.4 GHz	1 x 512 @ 1.54 GHz
Memory	1 GB @ 142 GB/s	900 MB @ 112 GB/s	1.5 GB @ 177 GB/s	1.5 GB @ 192 GB/s
Texture	1 x 48 GT/s	2 x 46 GT/s	1 x 42 GT/s	1 x 49.4 GT/s
Power Cons.	236 W	289 W	250 W	250 W
Price			\$400	\$500

Performance: NVIDIA Server Products





Ratio of Operations



■ Transfer
■ Filtering
■ BP


	GTX280	Tesla C1060	GTX580	Tesla C2070
Architecture	GT200	GT200	Fermi	Fermi
Processors	1 x 240 @ 1.3 GHz	1 x 240 @ 1.25GHz	1 x 512 @ 1.54 GHz	1 x 448 @ 1.15 GHz
Memory	1 GB @ 142 GB/s	4 GB @ 102 GB/s	1.5 GB @ 192 GB/s	6 GB @ 144 GB/s
Texture	48 GT/s	48 GT/s	49.4 GT/s	42 GT/s
Power Cons.	236 W	187.8 W	250 W	238 W
Price		\$2000	\$500	\$3500

Performance of OpenCL




* This is preliminary version, better performance is possible

Nvidia vs ATI: Peak Performance



	GTX280	GTX580	HD5970
Architecture	GT200	Fermi	Cypress
Processors	1 x 240 @ 1.3 GHz	1 x 512 @ 1.54 GHz	2 x 1600 @ 725 MHz
Memory	1 GB @ 142 GB/s	1.5 GB @ 192 GB/s	1.5 GB @ 153.6 GB/s
Texture	48 GT/s	49.4 GT/s	2 x 68 GT/s
Power Cons.	236 W	250 W	294 W
Price		\$400	\$650

Performance: Tomography




* Only single core of HD5970 is working with latest AMD SDK

** OpenCL support is experimental

	GTX280	GTX580	HD5970
Architecture	GT200	Fermi	Cypress
Processors	1 x 240 @ 1.3 GHz	1 x 512 @ 1.54 GHz	2 x 1600 @ 725 MHz
Memory	1 GB @ 142 GB/s	1.5 GB @ 192 GB/s	1.5 GB @ 153.6 GB/s
Texture	48 GT/s	49.4 GT/s	2 x 68 GT/s
Power Cons.	236 W	250 W	294 W
Price		\$400	\$650

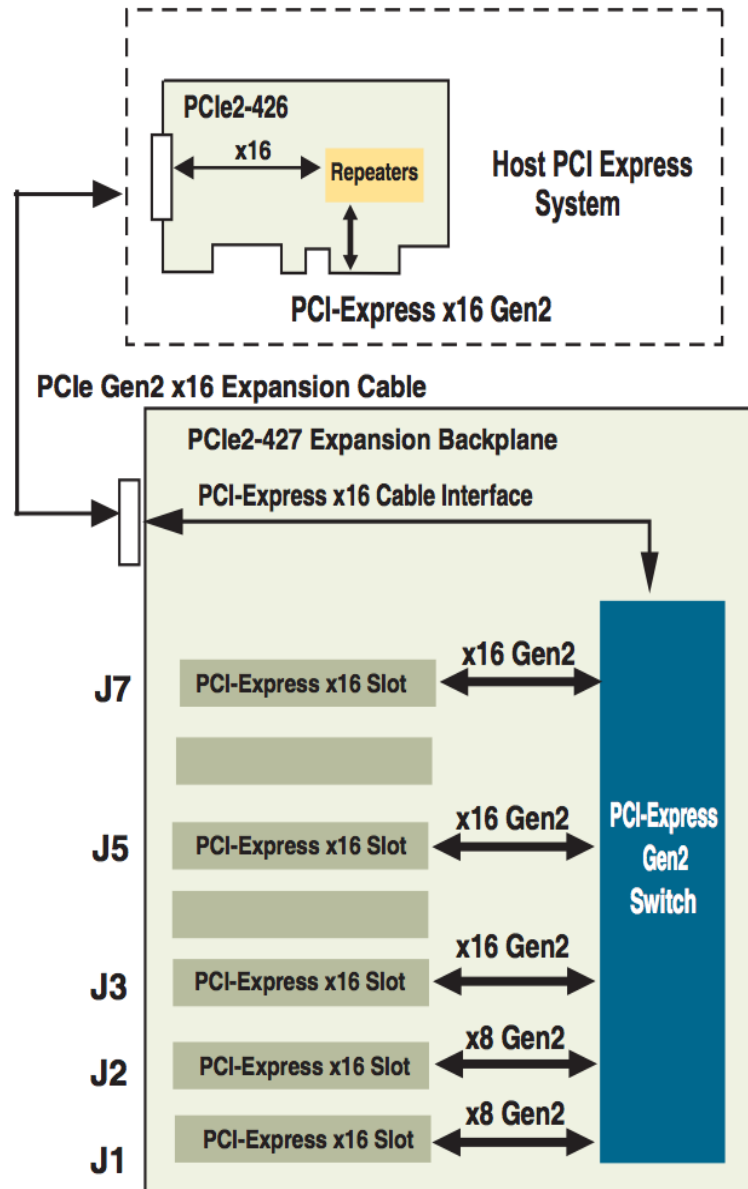
Utilization: Tomography



* Only single core of HD5970 is working with latest AMD SDK

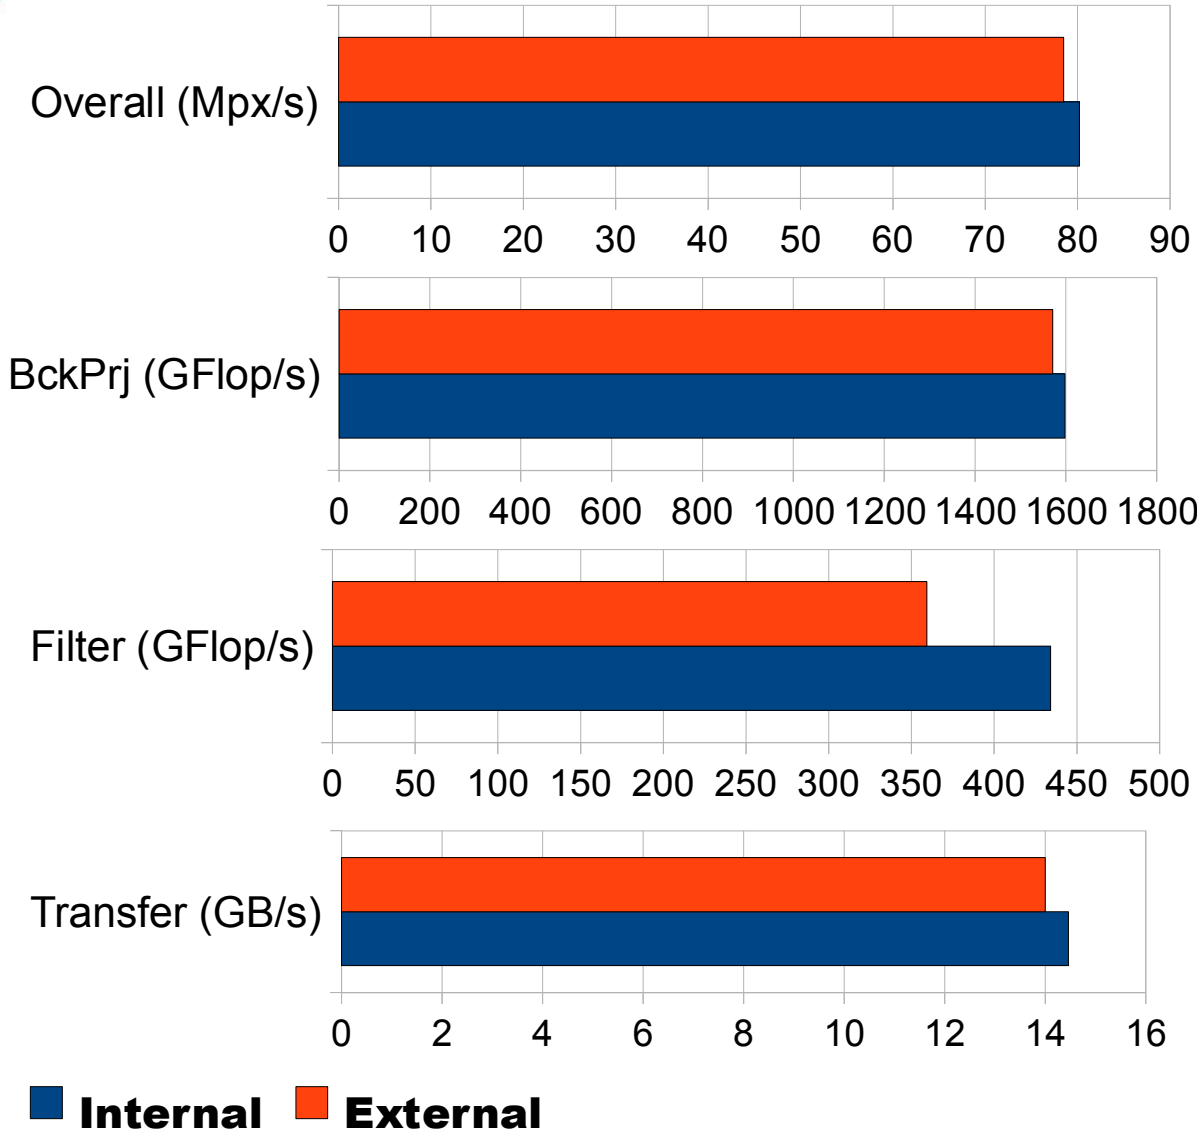
	GTX280	GTX580	HD5970
Architecture	GT200	Fermi	Cypress
Processors	1 x 240 @ 1.3 GHz	1 x 512 @ 1.54 GHz	2 x 1600 @ 725 MHz
Memory	1 GB @ 142 GB/s	1.5 GB @ 192 GB/s	1.5 GB @ 153.6 GB/s
Texture	48 GT/s	49.4 GT/s	2 x 68 GT/s
Power Cons.	236 W	250 W	294 W
Price		\$400	\$650

External GPU Boxes




Source:
Cyclone Microsystems

Performance: External GPU Boxes



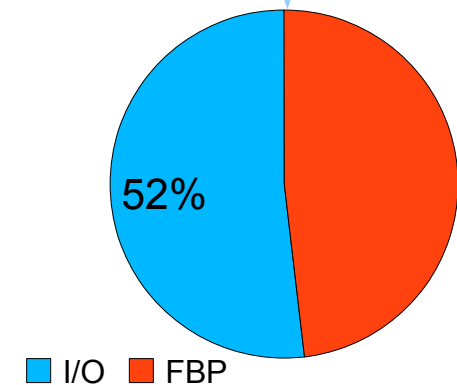


External GPU Box
PCIe Interface Card (16x)
4 External GPUs
4600 EUR

Performance: Storage Systems




Random read

- Hitachi A7K2000
- Areca 1680x
- Intel X25E
- 2 x Intel X25E Raid-0
- 2 x Crucial C300 Raid-0
- 4 x Crucial C300 Raid-0
- RamdDisk



Performance: RAID

12 Hitachi A7K200 in Raid6, OpenSuSe 11.3 Kernel 2.6.34



- Ext4 performance drops significantly if free space comes to the end. XFS on other hand have spurious reductions of speed on empty disk.
- Fragmentation reduces performance

- Excellent performance for tomography
- Increasing number of GPU platforms available
 - External boxes possible
- OpenCL
 - Good for nvidia
 - No universal code for all hardware (at least not now)
 - CPU cache not handled right?
 - ATI: Only one GPU, bad performance
- Storage
 - Different suitable solutions available:
RAM Disk, SSD, RAID Systems, ...
 - Strong file system dependency

What's next?

- Milestones
 - Documentation of the results?
 - PX?
 - Other (challenging) application?
 -