GPU Hardware

Asrock X79 Extreme11

= 15 TFlops

for ~ 2500 EUR

+ 4x

AMD Radeon HD 7970

Parallel Hardware

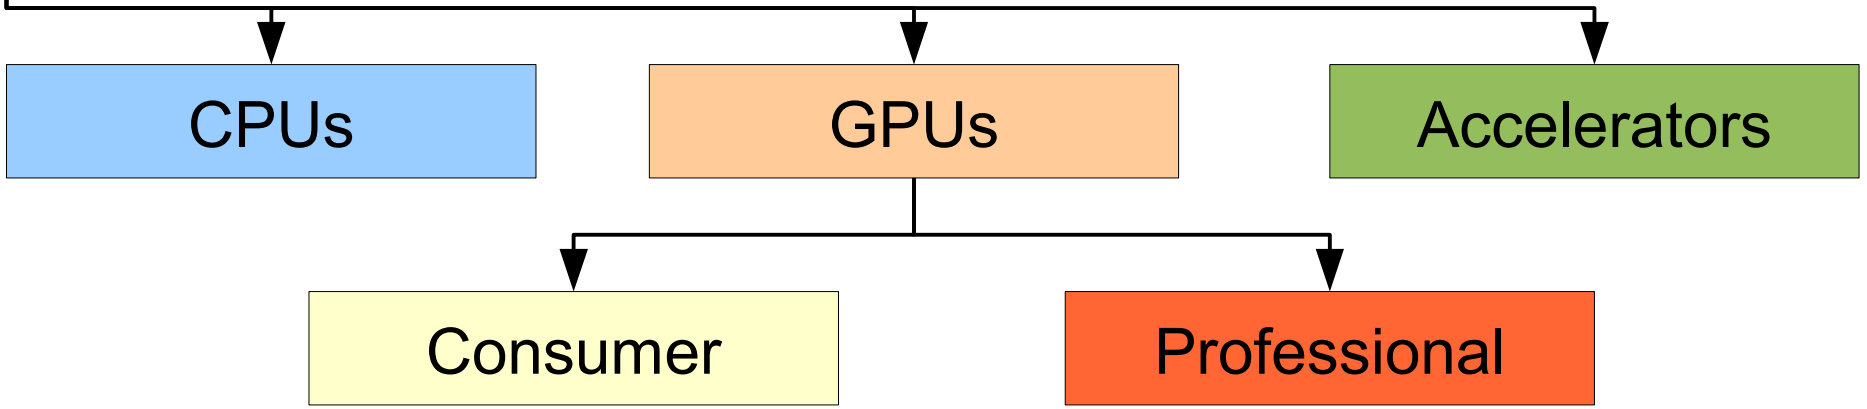

	E7-8870	Xeon/Phi	Tesla K20X	GeForce Titan	AMD HD7970	Power7+
SP	192	2020	3950	**4500**	3790	265
DP	96	1010	1310	**1300**	950	132
Mem	34.11	**320**	250	288	264	68
Texture			~ 180	**187.5**	118.4	
Max dev.	8	?	8	8	>=4	32
Price	$4,800	$2,800	$3,200	$1,000	**$400**	$$$$$$$$

Institute for Data Processing and Electronics
Karlsruhe Institute of Technology

GPUs vs. Accelerators

AMD & NVIDIA GPUs

- Mass-product: easy install, cheap prices, fast development, lots of 3rd party software
- May be installed in any desktop with x16 PCIe slot
- Have features of graphic processor: provides OpenGL functionality, texture engine may be used to accelerate some computations
- Only custom programming models are supported.
- Existing code can't be used directly

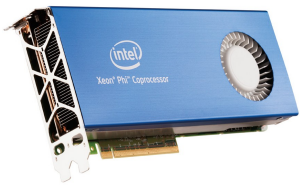

Intel Xeon Phi

- Professional-use only: only few Linux-distributions are supported, quite expensive
- Only server systems are supported, 64 bit PCI resource handling support is mandatory
- Computations only
- Standard programming models like OpenMP may be used.
- Some existing codes may be just recompiled to enjoy acceleration.

NVIDIA vs. AMD

- Supports both OpenCL and CUDA programming models and provides a number of highly optimized libraries
- Have advanced features like GPUDirect for cluster integration and real-time video processing.
- Large user base and lots of codes hand-tuned for latest architectures
- Good quality drivers and development tools
- No hardware specifications
- Most of consumer cards lack fast double-precision
- OpenCL support is limited to version 1.1 and sabotaged by NVIDIA lately
- Only GPU-computations are supported by OpenCL platform

- Only OpenCL is supported. Situation with libraries is improving, but performance is still inferior to CUDA libs.
- Does not include methods to bypass system memory
- Only few tuned codes are available. Besides, VLIW-based cards require very different type of tunning
- Very problematic drivers
- AMD publishes hardware specifications and there is hope for better quality open source drivers in future...
- Most of the cards have fast double-precision support
- OpenCL platform supports AMD GPUs and both Intel and AMD CPUs devices.
- Generally cheaper than NVIDIA

GeForce vs. Tesla

- Memory with error-correction (ECC)
- Generally have more global memory on board, though consumer GTX Titan has as much as Tesla K20
- Generally support faster double precision, though consumer GTX Titan have fast double precision mode as well
- GPUDirect for RDMA required to accelerate GPU-over-Infiniband in clusters is only available for Tesla K20
- Device sharing between multiple MPI processes (HyperQ) is only available for Kepler-based Tesla devices
- Monitoring and cluster management tools from NVIDIA are only available for Tesla devices
- Significantly more expensive and often slightly slower single-precision computations as compared to top consumer cards

Institute for Data Processing and Electronics
Karlsruhe Institute of Technology

GPU Platforms

SuperMicro 7047GR-TPRF (Intel C602 Chipset)
CPU: 2 x Xeon E5-2600 v2 family
Memory: 16 DDR3 sockets (512GB max)
PCIe: 4x x16 gen3, 2x x8 gen3, x4 gen2

Tyan FT77AB7059 (Intel C602 Chipset)
CPU: 2 x Xeon E5-2600 v2 family
Memory: 24 DDR3 sockets (768GB max)
PCIe: 8x x16 gen3

Asrock X79 Extreme11 (Intel X79 Chipset)
CPU: Core i7 (LGA 2011)
Memory: 8 DDR3 sockets (64GB max)
PCIe: 4x x16 gen3 (using PLX PEX8747 switch)

PCI Express Switches

Standard Configuration

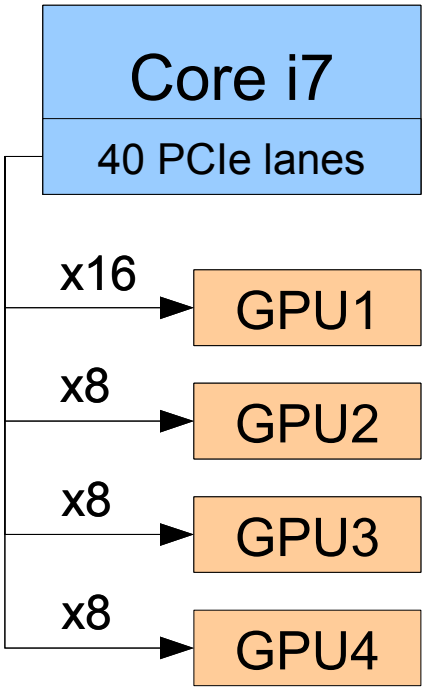

- Speed to most devices is always halved
- Transfer between devices is halved as well

Configuration with PCI express switch

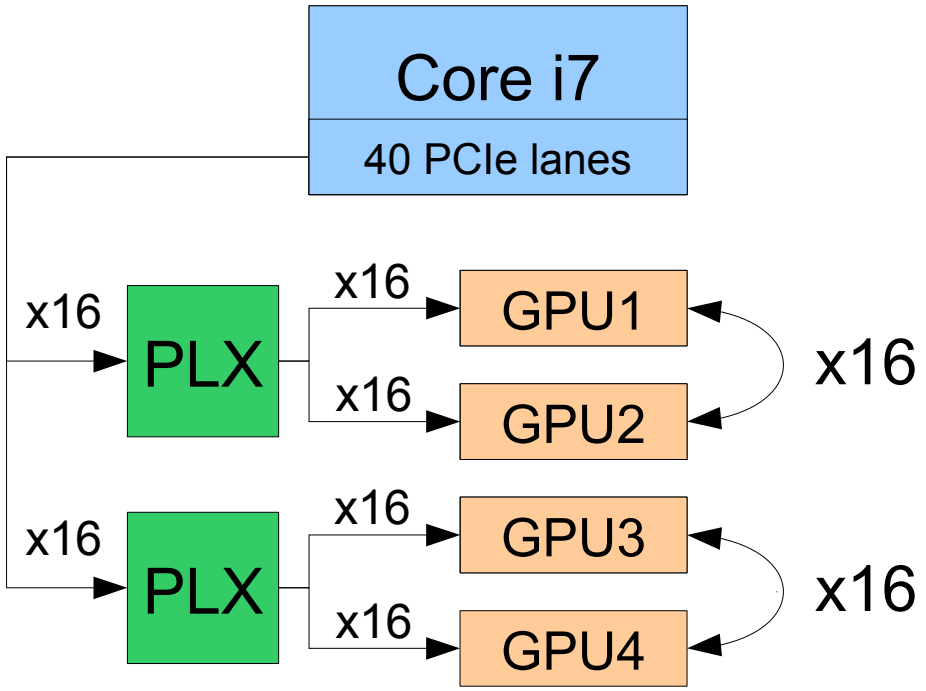

- Switched architecture allows full speed to devices if there is enough bandwidth
- In any case the direct transfer between GPUs
- According to NVIDIA, GPUDirect works better over switches when using chipset interconnect

Adding more GPUs

1 x PCIe x16 2.0
4 x GTX590
8 GPU cores

External GPU Enclosure
by One Stop Systems

Real-life test for tomography setup

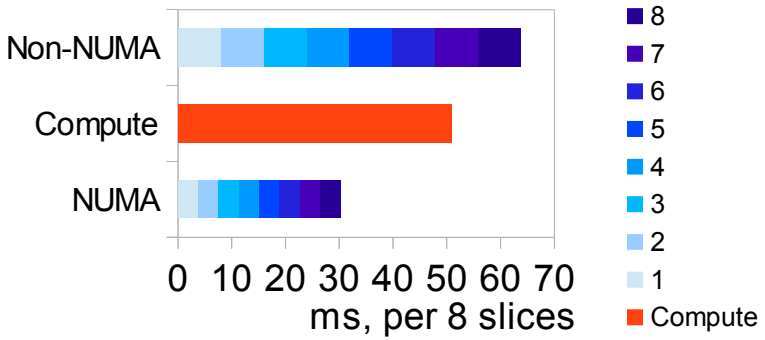

- Maximum 8 GPU cores (not cards) per system are working without major troubles. ASTRA Lab reported to run 13 GPU cores using specially modified BIOS
- To run more than 5 GPUs, NVIDIA driver have to be force to use MSI interrupts.
- Application startup times are linearly increased with number of GPU cores (with 8 cores it is about 3 seconds now)

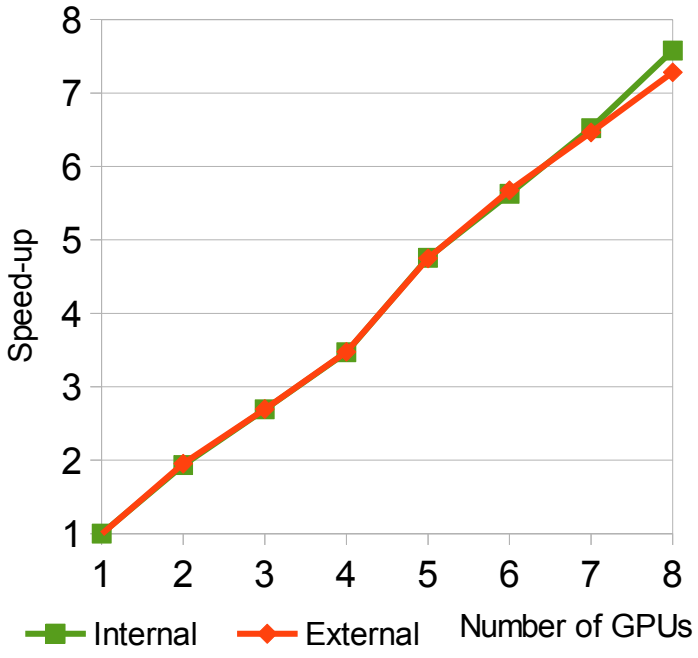

Institute for Data Processing and Electronics
Karlsruhe Institute of Technology

More things to consider

NUMA

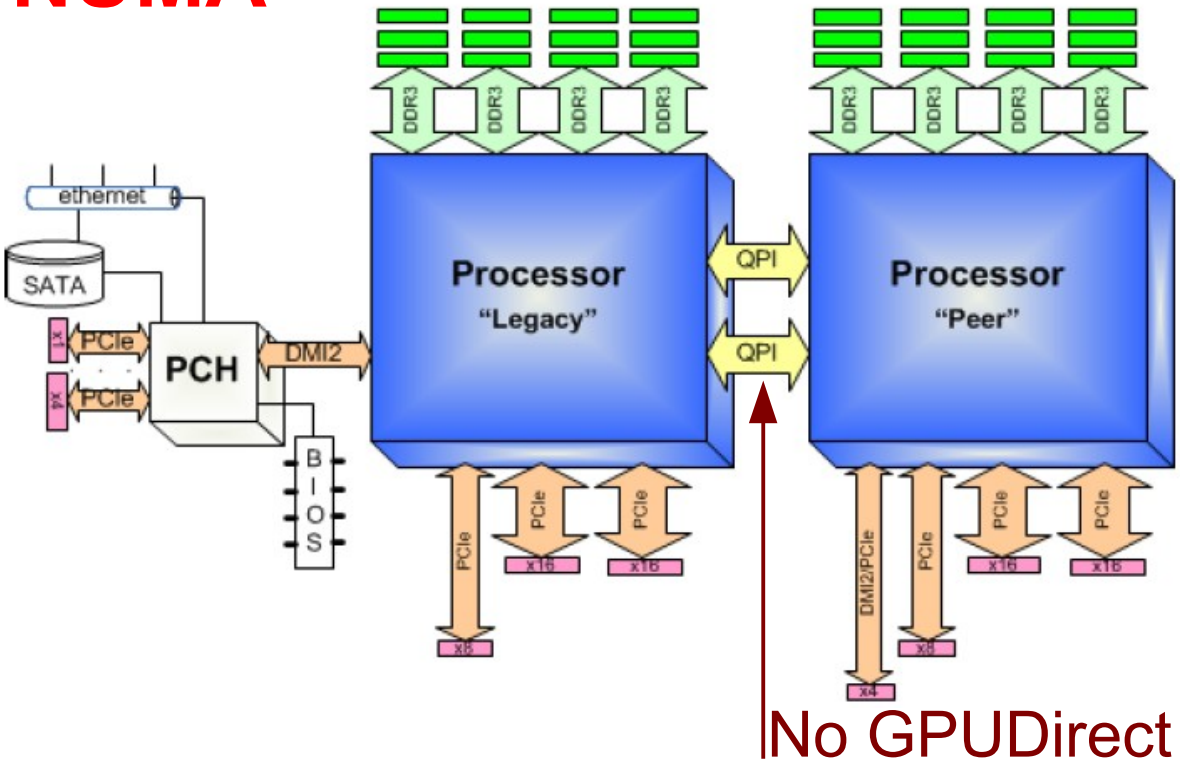

Xeon E5-2640

Bandwidth

QPI: 14.4 GB/s

DDR3: 10.6 GB/s (PC1333)

PCIe: 16 GB/s (gen3 x16)

QPI bus is even not enough to feed both GPU cards

No GPUDirect over QPI bus

I/O

- GPUDirect for RDMA will extremely improve cluster performance
- Storage I/O will definitively be an issue. SSD-based raid arrays may be a solution. However to go above GB/s speeds, a custom storage code have to be developed based on Kernel AIO, etc.