




ECE 498AL

Lecture 7: GPU as part of the PC Architecture

Objective

- To understand the major factors that dictate performance when using GPU as an compute accelerator for the CPU
 - The feeds and speeds of the traditional CPU world
 - The feeds and speeds when employing a GPU
 - To form a solid knowledge base for performance programming in modern GPU's
- Knowing yesterday, today, and tomorrow
 - The PC world is becoming flatter
 - Outsourcing of computation is becoming easier...

Review- Typical Structure of a CUDA Program

- Global variables declaration
- Function prototypes
 - `__global__ void kernelOne(...)`
- Main ()
 - allocate memory space on the device `cudaMalloc(&d_GblVarPtr, bytes)`
 - transfer data from host to device `cudaMemcpy(d_GblVarPtr, h_Gl...)`
 - execution configuration setup
 - kernel call `kernelOne<<<execution configuration>>>(args...);`  repeat as needed
 - transfer results from device to host `cudaMemcpy(h_GblVarPtr, ...)`
 - optional: compare against golden (host computed) solution
- Kernel `void kernelOne(type args,...)`
 - variables declaration - `__local__`, `__shared__`
 - automatic variables transparently assigned to registers or local memory
 - `syncthreads()...`

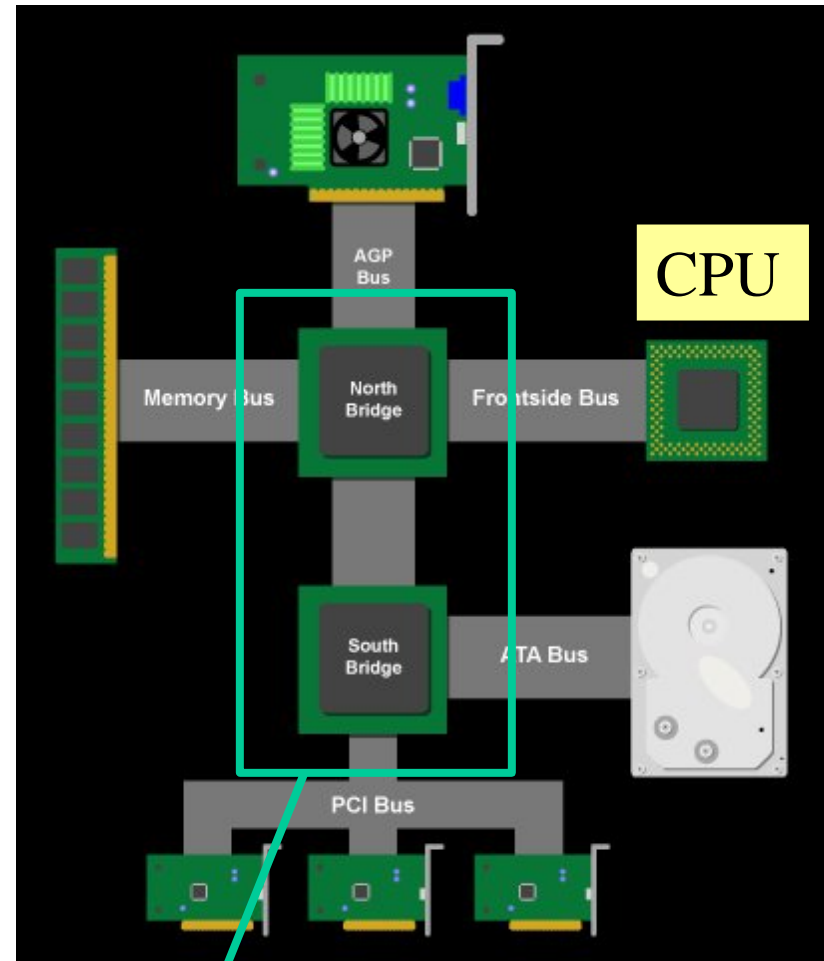
Bandwidth –

Gravity of Modern Computer Systems

- The Bandwidth between key components ultimately dictates system performance
 - Especially true for massively parallel systems processing massive amount of data
 - Tricks like buffering, reordering, caching can temporarily defy the rules in some cases
 - Ultimately, the performance goes falls back to what the “speeds and feeds” dictate


Classic PC architecture

- Northbridge connects 3 components that must be communicate at high speed
 - CPU, DRAM, video
 - Video also needs to have 1st-class access to DRAM
 - Previous NVIDIA cards are connected to AGP, up to 2 GB/s transfers
- Southbridge serves as a concentrator for slower I/O devices



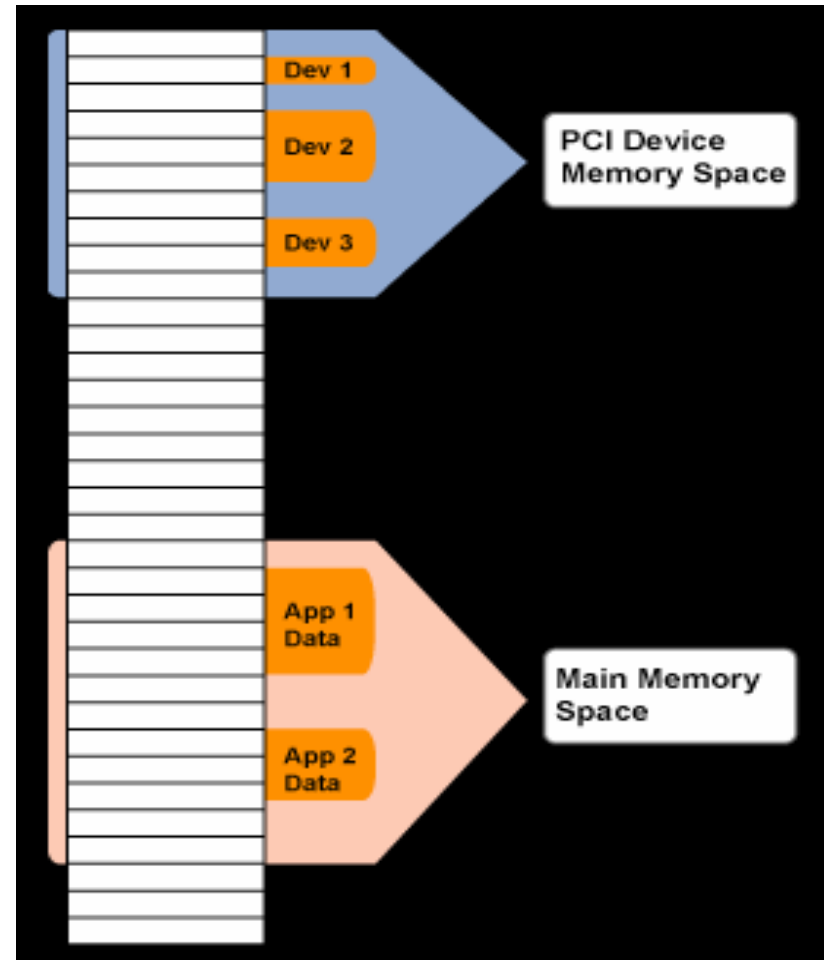
(Original) PCI Bus Specification

- Connected to the southBridge
 - Originally 33 MHz, 32-bit wide, 132 MB/second peak transfer rate
 - More recently 66 MHz, 64-bit, 512 MB/second peak
 - Upstream bandwidth remain slow for device (256MB/s peak)
 - Shared bus with arbitration
 - Winner of arbitration becomes bus master and can connect to CPU or DRAM through the southbridge and northbridge




PCI as Memory Mapped I/O

- PCI device registers are mapped into the CPU's physical address space
 - Accessed through loads/ stores (kernel mode)
- Addresses assigned to the PCI devices at boot time
 - All devices listen for their addresses




PCI Express (PCIe)

- Switched, point-to-point connection
 - Each card has a dedicated “link” to the central switch, no bus arbitration.
 - Packet switches messages form virtual channel
 - Prioritized packets for QoS
 - E.g., real-time video streaming



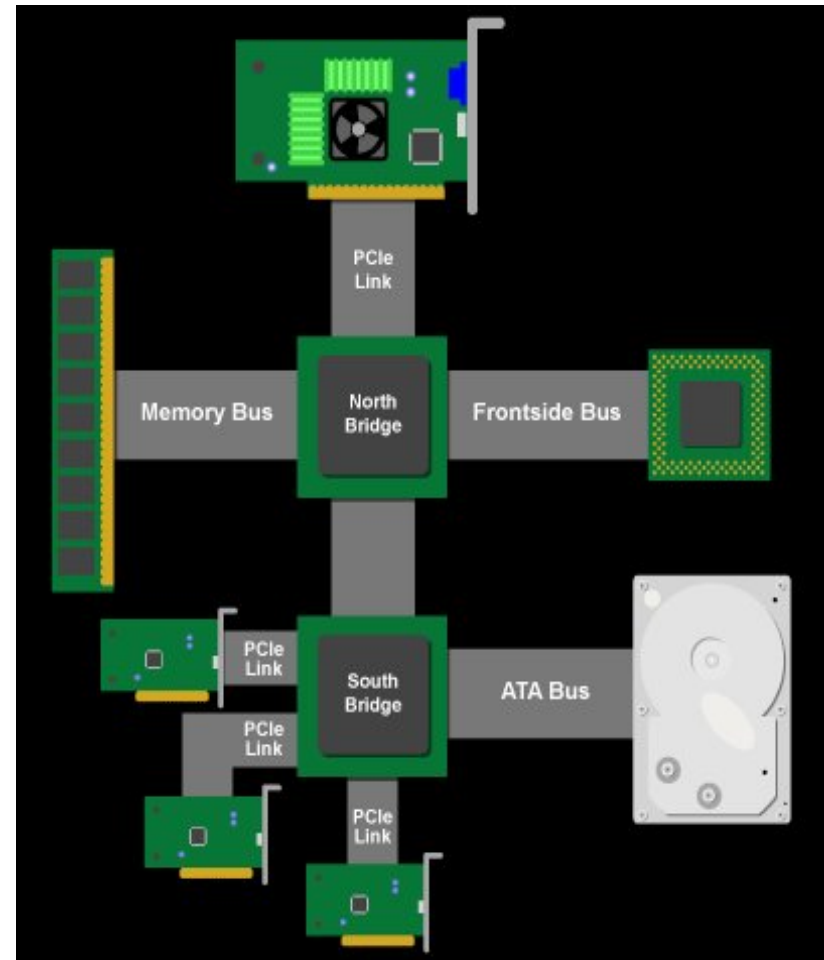
PCIe Links and Lanes

- Each link consists of one more lanes
 - Each lane is 1-bit wide (4 wires, each 2-wire pair can transmit 2.5Gb/s in one direction)
 - Upstream and downstream now simultaneous and symmetric
 - Each Link can combine 1, 2, 4, 8, 12, 16 lanes- x1, x2, etc.
 - Each byte data is **8b/10b** encoded into 10 bits with equal number of 1's and 0's; net data rate 2 Gb/s per lane each way.
 - Thus, the net data rates are 250 MB/s (x1) 500 MB/s (x2), 1GB/s (x4), 2 GB/s (x8), 4 GB/s (x16), each way



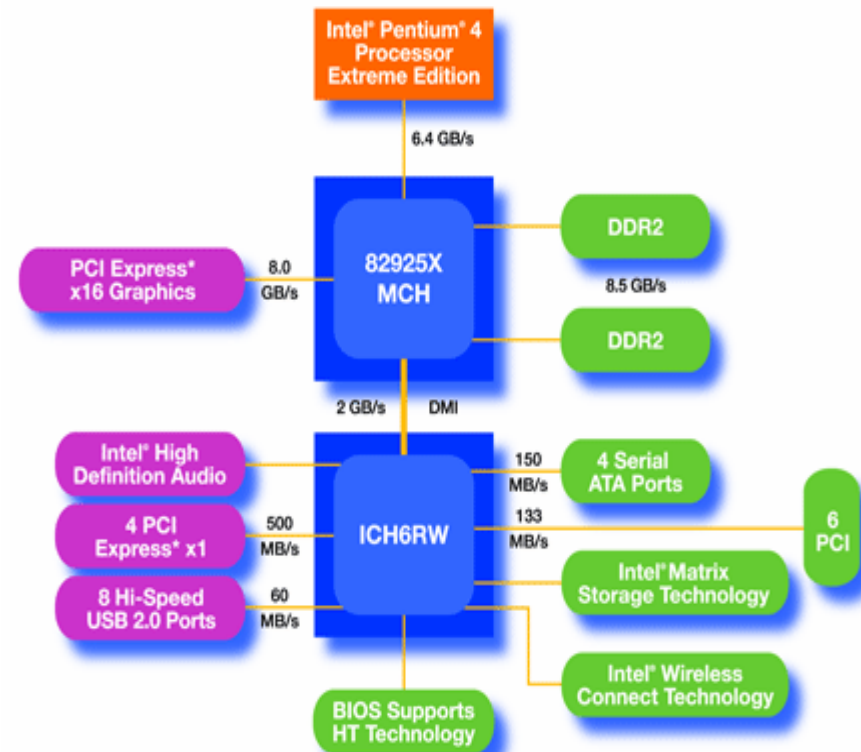
PCIe PC Architecture

- PCIe forms the interconnect backbone
 - Northbridge/Southbridge are both PCIe switches
 - Some Southbridge designs have built-in PCI-PCIe bridge to allow old PCI cards
 - Some PCIe cards are PCI cards with a PCI-PCIe bridge
- Source: Jon Stokes, PCI Express: An Overview
 - <http://arstechnica.com/article/s/paedia/hardware/pcie.ars>



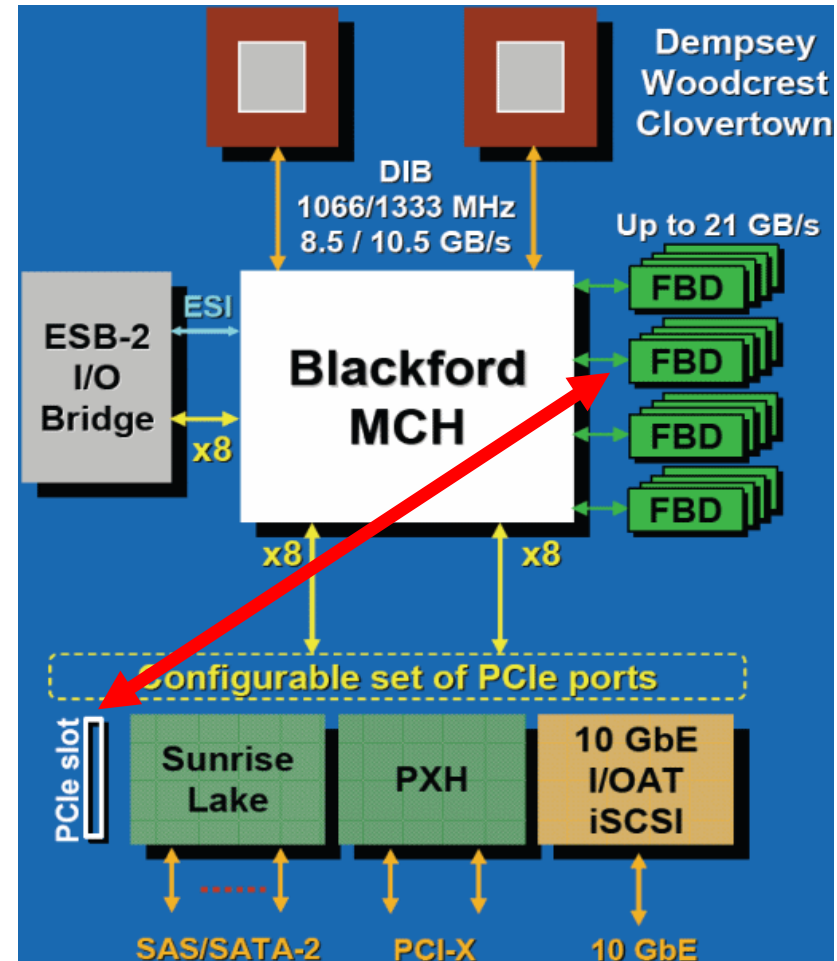
Today's Intel PC Architecture: Single Core System

- FSB connection between processor and Northbridge (82925X)
 - Memory Control Hub
- Northbridge handles “primary” PCIe to video/GPU and DRAM.
 - PCIe x16 bandwidth at 8 GB/s (4 GB each direction)
- Southbridge (ICH6RW) handles other peripherals



Today's Intel PC Architecture: Dual Core System

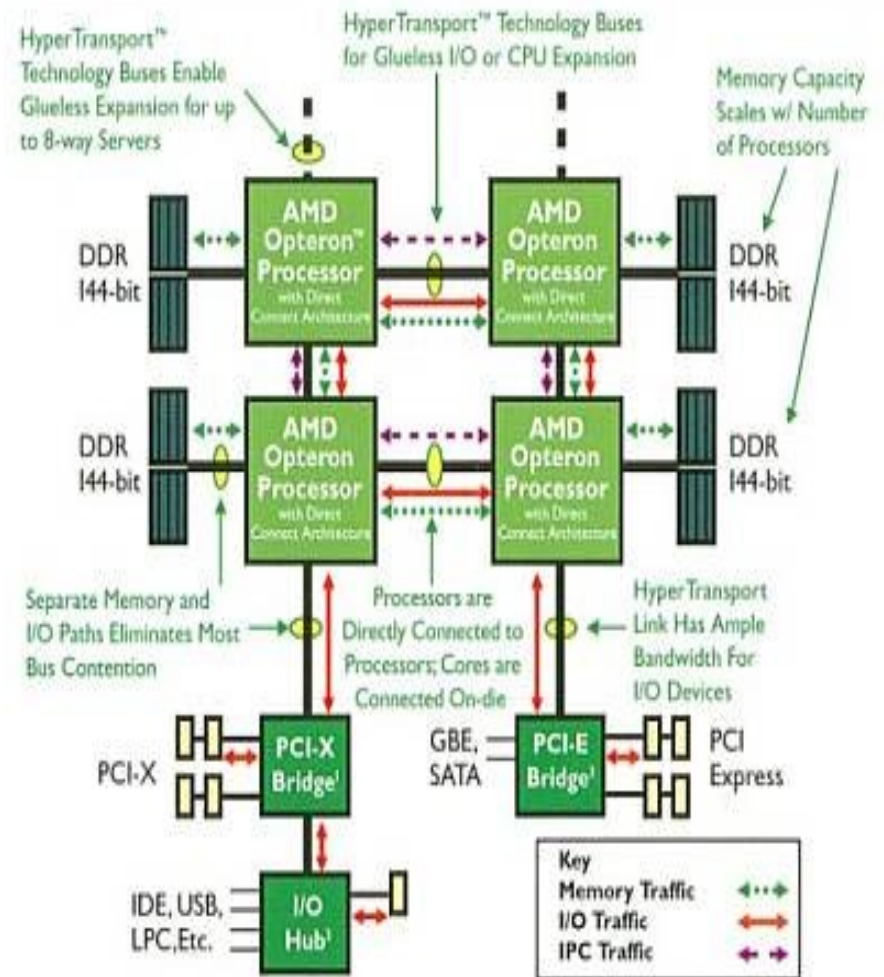
- Bensley platform
 - Blackford Memory Control Hub (MCH) is now a PCIe switch that integrates (NB/SB).
 - FBD (Fully Buffered DIMMs) allow simultaneous R/W transfers at 10.5 GB/s per DIMM
 - PCIe links form backbone
 - PCIe device upstream bandwidth now equal to down stream
 - Workstation version has x16 GPU link via the Greencreek MCH
- Two CPU sockets
 - Dual Independent Bus to CPUs, each is basically a FSB
 - CPU feeds at 8.5–10.5 GB/s per socket
 - Compared to current Front-Side Bus CPU feeds 6.4GB/s
- PCIe bridges to legacy I/O devices



Today's AMD PC Architecture


- AMD HyperTransport™ Technology bus replaces the Front-side Bus architecture
- HyperTransport™ similarities to PCIe:
 - Packet based, switching network
 - Dedicated links for both directions
- Shown in 4 socket configuraton, 8 GB/sec per link
- Northbridge/HyperTransport™ is on die
- Glueless logic
 - to DDR, DDR2 memory
 - PCI-X/PCIe bridges (usually implemented in Southbridge)

AMD Opteron™ Processor-based 4P Server




Today's AMD PC Architecture

- “Torrenza” technology
 - Allows licensing of coherent HyperTransport™ to 3rd party manufacturers to make socket-compatible accelerators/co-processors
 - Allows 3rd party PPU (Physics Processing Unit), GPUs, and co-processors to access main system memory directly and coherently
 - Could make accelerator programming model easier to use than say, the Cell processor, where each SPE cannot directly access main memory.




HyperTransport™ Feeds and Speeds

- Primarily a low latency direct chip-to-chip interconnect, supports mapping to board-to-board interconnect such as PCIe
- HyperTransport™ 1.0 Specification
 - 800 MHz max, 12.8 GB/s aggregate bandwidth (6.4 GB/s each way)
- HyperTransport™ 2.0 Specification
 - Added PCIe mapping
 - 1.0 - 1.4 GHz Clock, 22.4 GB/s aggregate bandwidth (11.2 GB/s each way)
- HyperTransport™ 3.0 Specification
 - 1.8 - 2.6 GHz Clock, 41.6 GB/s aggregate bandwidth (20.8 GB/s each way)
 - Added AC coupling to extend HyperTransport™ to long distance to system-to-system interconnect



Courtesy HyperTransport™ Consortium
Source: "White Paper: AMD HyperTransport Technology-Based System Architecture"

GeForce 7800 GTX Board Details



256MB/256-bit DDR3
600 MHz
8 pieces of 8Mx32