Front Side Buses SMP systems

COMP311 2008

Jamie Curtis

Traditional Intel Architecture

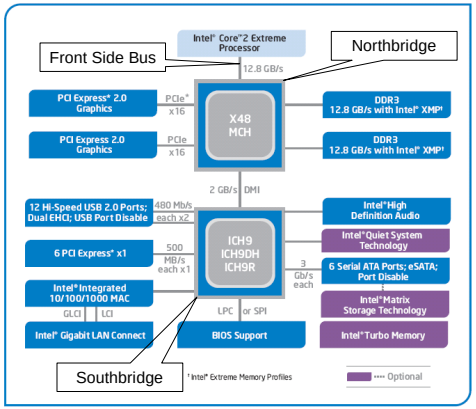

Clock vs. Data rates

- Clock rates no longer equal data rates
- Watch specifications as both can look identical
- e.g.
  - Intel's FSB is normally referred to as an 1600MHz bus
  - It is actually a 400MHz clock with a quad-pumped data rate
  - Should really be 1600MT/s

Frontside Bus & Dual Core

- 64 bit, quad-pumped 400MHz
  - 8 * 4 * 400 = 12.8GB/s
- Half Duplex
- Traditional P4 FSB has been point to point
  - Xeon's SMP requires chipsets support a proper multipoint bus for the FSB.
- How did a dual core Pentium D work ?

Intel Dual Core

- Slap two cores on the same die !
- How do they communicate ?
  - Over the FSB like a Xeon !
  - Requires a new chipset to support it.

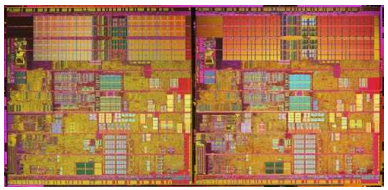

Caches

- There are two sets of L1 and L2 cache, one for each core.
- What happens if both cores are caching the same memory location ?
- Need a protocol to make sure this doesn't cause problems
  - Cache Coherency Protocol

Cache Coherency

- Intel use the MESI Protocol
  - ***M***odified
    - 1 cache contains a modified copy of the location
  - ***E***xclusive
    - 1 cache contains a un-modified copy of the location
  - ***S***hared
    - 2 or more caches contain un-modified copies of the location
  - ***I***nvalid
    - Another cache contains a modified copy of the location

MESI

- What happens if a core has an Invalid entry but tries to access it ?
  - The cache with the Modified entry needs to write the entry back to memory and become a Shared entry
- This means an Intel Pentium D has to involve the FSB in all of it's cache coherency updates even though they are on the same die

Intel Core

- Core architecture is designed for dual-core
  - Cache is now shared between cores
  - Cache coherency is now between each L1 and L2

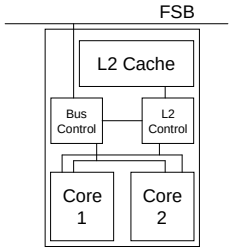

Intel Core 2 Quad

- Current quad core design adds two dual core designs together
- Cache coherency between dies again happens over the FSB

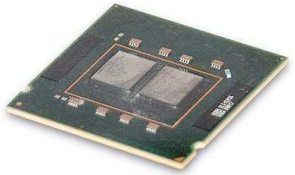

AMD K8 Architecture

- Integrated Memory Controller
  - Very low latency
  - CPU determines memory technology
    - Requires both CPU and Motherboard to be changed for a new type of memory
- Athlon 64, Socket 754
  - Single channel, 64bit DDR at 200MHz
    - 8 * 2 * 200 = 3.2GB/s

AMD K8 Memory Interfaces

- Athlon 64, Socket 939
  - Dual channel, 64bit DDR at 200MHz
    - 2 * 8 * 2 * 200 = 6.4 GB/s
- Athlon 64, Socket AM2
  - Dual channel, 64bit DDR2 at 400MHz
    - 2 * 8 * 2 * 400 = 12.8 GB/s
- Opteron, Socket 940
  - Dual channel, 64bit DDR at 200MHz
    - 2 * 8 * 2 * 200 = 6.4 GB/s

AMD K8 FSB

- AMD use the HyperTransport technology
  - Don't confuse with the very different Intel Hyper Threading Technology
  - Packet based, point to point link that provides the lowest possible latency
  - Full duplex bi-directional link
  - Available in 2, 4, 8, 16 or 32 bits wide
  - 50MHz – 1.4GHz clock rates
  - Clock rates and bit widths can be asymmetric
  - Double-pumped data rate
    - 1.4GHz * 2 * 4 = 11.2GB/s per direction

Athlon 64 FSB

- The Athlon 64 has a single HyperTransport link to connect to I/O subsystem
  - 16bit, bi-directional, DDR at 800MHz (754) or 1GHz (939)
    - 2 * 2 * 2 * 1000 = 8GB/s
- Can have either a single chip solution or stay with two chip, northbridge, southbridge combination

Opteron

- Uses HT for both the I/O interconnect and CPU interconnect
- CPU interconnect requires an additional Cache Coherency protocol addition to HT
- Come in three variants
  - 1xx – Memory controller and 3 HT links
  - 2xx – Memory controller and 3 HT links
  - 8xx – Memory controller and 3 HT links

Opteron

- What makes them different ?
  - Number of HT busses that support the CC protocol
    - 1xx – Zero
    - 2xx – One
    - 8xx – Three
- This allows you to scale to different numbers of processors
  - 1xx – 1 way, 2xx – 2way, 8xx – 4 or 8 way

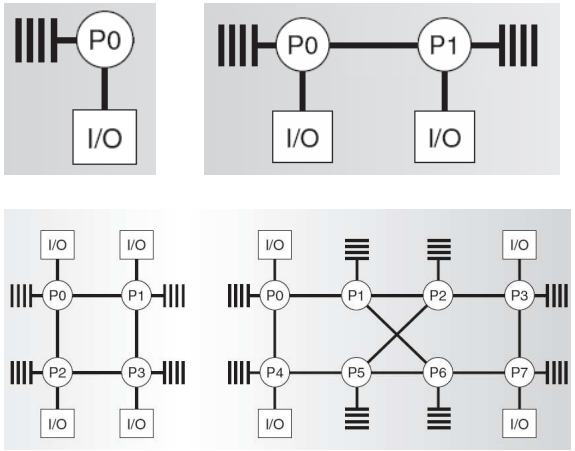

AMD MOESI

- Cache Coherency slightly different to Intel's
  - ***O***wner – This cache owns this memory location and it (not memory) services all requests for it from other caches
- The request goes across the high speed dedicated HT bus

AMD Dual Core

- In the dual core situation, it becomes even better:

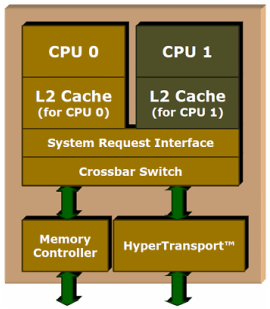

AMD Dual Core

- Requests now go across the System Request Interface
  - This runs at CPU core frequency
- System Request Interface also controls HT – Memory access as well as HT – HT communication in 2xx and 8xx Opterons

AMD Quad Core

- AMD again targeted native quad core instead of dual-die, dual core
  - Introduced shared L3 cache

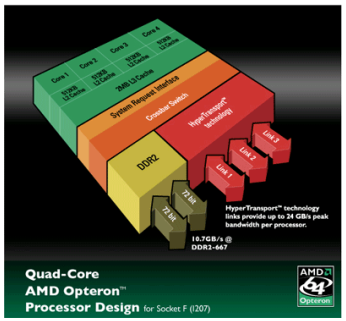

Intel Core i7

- Follows much of AMD's design decisions
  - On-die memory controller
  - High speed interconnect for both I/O and SMP communications
  - L1 and L2 per core, L3 shared per CPU
- New interconnect called QPI
  - QuickPath Interconnect

Quick Path

- Point to Point serial, full duplex
- Width and clock speed can be altered depending on the application
- Core i7 will release with 12.8GB/s each direction
  - 20bits ("lanes") wide (16data + 4 checksum)
  - 3.2GHz DDR clock
- Lanes can be reduced dynamically for power saving

Intel Core i7

- Initial releases indicate a single QPI interface for desktop processors and a dual QPI interface for Xeon servers
  - Allows for dual processor servers
- Potentially add extra QPI interfaces for higher processor counts
  - 4 QPI interfaces per processor for 4 way configurations