


Lecture 6: Snooping Protocol Design Issues

- Topics: split transaction buses, case studies

Split Transaction Bus

- What would it take to implement the protocol correctly while assuming a split transaction bus?
- Split transaction bus: a cache puts out a request, releases the bus (so others can use the bus), receives its response much later
- Assumptions:
 - only one request per block can be outstanding
 - separate lines for addr (request) and data (response)

Split Transaction Bus



Design Issues

- When does the snoop complete? What if the snoop takes a long time?
- What if the buffer in a processor/memory is full? When does the buffer release an entry? Are the buffers identical?
- How does each processor ensure that a block does not have multiple outstanding requests?
- What determines the write order – requests or responses?

Design Issues II

- What happens if a processor is arbitrating for the bus and witnesses another bus transaction for the same address?
- If the processor issues a read miss and there is already a matching read in the request table, can we reduce bus traffic?

Shared Cache Designs

- There are benefits to sharing the first level cache among many processors (for example, in a CMP):
 - no coherence protocol
 - low cost communication between processors
 - better prefetching by processors
 - working set overlap allows shared cache size to be smaller than combined size of private caches
 - improves utilization
- Disadvantages:
 - high contention for ports
 - longer hit latency (size and proximity)
 - more conflict misses

TLBs

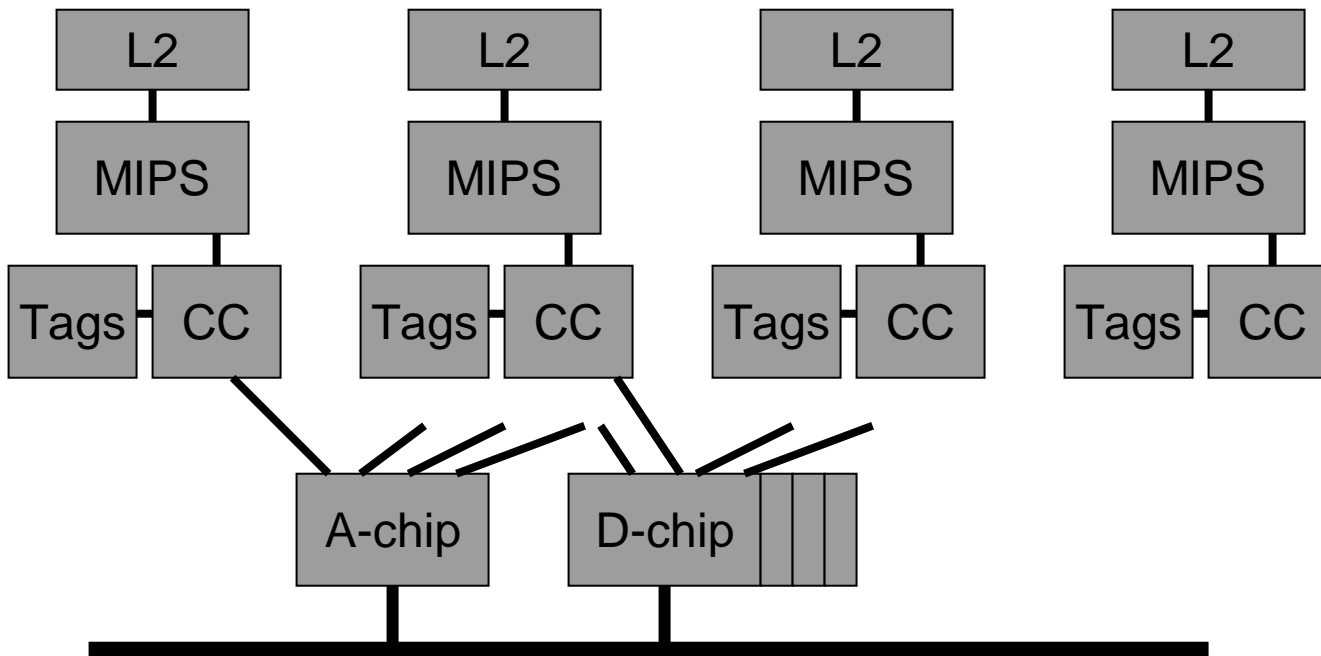
- Recall that a TLB caches virtual to physical page translations
- While swapping a page out, can we have a problem in a multiprocessor system?
- All matching entries in every processor's TLB must be removed
- TLB shutdown: the initiating processor sends a special instruction to other TLBs asking them to invalidate a page table entry

Case Study: SGI Challenge

- Supports 18 or 36 MIPS processors
- Employs a 1.2 GB/s 47.6 MHz system bus (Powerpath-2)
- The bus has 256-bit-wide data, 40-bit-wide address, plus 33 other signals (non multiplexed)
- Split transaction, supporting eight outstanding requests
- Employs the MESI protocol by default – also supports update transactions

Processor Board

- Each board has four processors (to reduce the number of slots on the bus from 36 to 9)
- A-chip has request tables, arbitration logic, etc.



Latencies

- 75ns for an L2 cache hit
- 300ns for a cache miss to percolate down to the A-chip
- Additional 400ns for the data to be delivered to the D-chips across the bus (includes 250ns memory latency)
- Another 300ns for the data to reach the processor
- Note that the system bus can accommodate 256 bits of data, while the CC-chip to processor interface can handle 64 bits at a time

Sun Enterprise 6000

- Supports 30 UltraSparcs
- 2.67 GB/s 83.5 MHz Gigaplane system bus
- Non multiplexed bus with 256 bits of data, 41 bits of address, and 91 bits of control/error correction, etc.
- Split transaction bus with up to 112 outstanding requests
- Each node speculatively drives the bus (in parallel with arbitration)

Latencies

- L2 hits take 40ns
- Memory access takes a total of 300ns, including 130ns for the bus transfer

Title

- Bullet