

Cap6

Snoop-based Multiprocessor Design

Design Goals

Performance and cost depend on design and implementation too

Goals

- Correctness
- High Performance
- Minimal Hardware

Often at odds (riscos.....)

- High Performance => multiple outstanding low-level events
 - => more complex interactions
 - => more potential correctness bugs

We'll start simply and add concurrency to the design



6.1 Correctness Issues

Fulfill conditions for coherence and consistency

- Write propagation, serialization; for SC: completion, atomicity

Deadlock: all system activity ceases

- Cycle of resource dependences



Livelock: no processor makes forward progress although transactions are performed at hardware level

- e.g. simultaneous writes in invalidation-based protocol
 - each requests ownership, invalidating other, but loses it before winning arbitration for the bus

Starvation: one or more processors make no forward progress while others do.

- e.g. interleaved memory system with NACK on bank busy
- Often not completely eliminated (not likely, not catastrophic)

6.2 Base Cache Coherence Design

Até agora:

- Single-level write-back cache
- Invalidation protocol
- One outstanding memory request per processor
- Atomic memory bus transactions
 - For BusRd, BusRdX no intervening transactions allowed on bus between issuing address and receiving data
 - BusWB: address and data simultaneous and sinked by memory system before any new bus request
- Atomic operations within process
 - One finishes before next in program order starts

Examine write serialization, completion, atomicity

Then add more concurrency/complexity and examine again

Some Design Issues

Design of cache controller and tags

- Both processor and bus need to look up

How and when to present snoop results on bus

Dealing with write backs

Overall set of actions for memory operation not atomic

- Can introduce race conditions

New issues deadlock, livelock, starvation, serialization, etc.

Implementing atomic operations (e.g. read-modify-write)

Let's examine one by one ...

6.2.1 Cache Controller and Tags

Cache controller stages components of an operation

- Itself a finite state machine (but not same as protocol state machine)

Uniprocessor: On a miss:

- Assert request for bus
- Wait for bus grant
- Drive address and command lines
- Wait for command to be accepted by relevant device
- Transfer data

In snoop-based multiprocessor, cache controller must:

- Monitor bus and processor
 - Can view as two controllers: bus-side, and processor-side (ver fig 6.3)
 - With single-level cache: dual tags (not data) or dual-ported tag RAM
 - must reconcile when updated, but usually only looked up
- Respond to bus transactions when necessary (multiprocessor-ready)

6.2.2 Reporting Snoop Results: How?

Collective response from caches must appear on bus

Example: in MESI protocol, need to know

- Is block dirty; i.e. should memory respond or not?
- Is block shared; i.e. transition to E or S state on read miss?

Three wired-OR signals

- Shared: asserted if any cache has a copy
- Dirty: asserted if some cache has a dirty copy
 - needn't know which, since it will do what's necessary
- Snoop-valid: asserted when OK to check other two signals (equivalente a um strobe ou enable)
 - actually inhibit until OK to check

Illinois MESI requires priority scheme for cache-to-cache transfers

- Which cache should supply data when in shared state?
- Commercial implementations allow memory to provide data (ver Challenge e Enterprise)

Reporting Snoop Results: When?

Memory needs to know what, if anything, to do

1 Fixed number of clocks from address appearing on bus

- Dual tags required to reduce contention with processor (que tem prioridade)
- Still must be conservative (processor update both tags on write: E -> M; tags ficam ocupados)
- Pentium Pro, HP servers, Sun Enterprise

2 Variable delay

- Memory assumes cache will supply data till all say “sorry”
- Less conservative, more flexible, more complex
- Memory can fetch data and hold just in case (SGI Challenge)

3 Immediately: Bit-per-block in memory (existe bloco modificado em alguma cache?)

- Extra hardware complexity in commodity main memory system

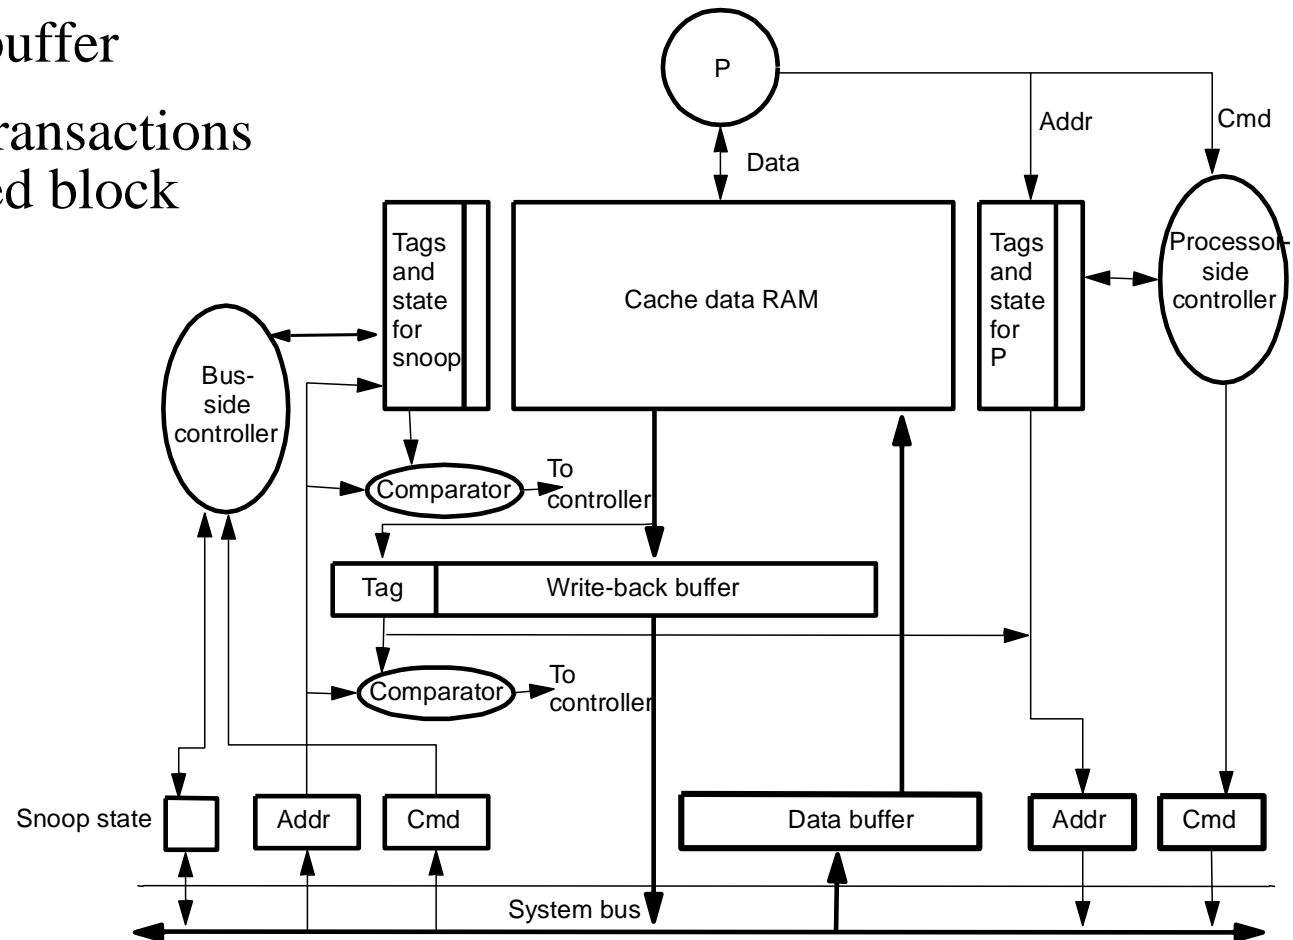
6.2.3 Writebacks

Duas transações: bloco buscado pelo miss e bloco enviado p/ mem(WB)

To allow processor to continue quickly, want to service miss first and then process the write back caused by the miss asynchronously

- Need write-back buffer
- Must handle bus transactions relevant to buffered block

- snoop the WB buffer
- comparador observa se alguém está precisando do bloco em WB, fornece o dado e cancela o pedido para acesso ao bus (alguém agora ficou com o dado)



6.2.5 Non-Atomic State Transitions

Nos diagramas (FSM) do Cap. 5, assumiu-se que as transições de estado eram instantâneas (ou atômicas)

Memory operation involves many actions by many entities, including bus transactions

- Look up cache tags, bus arbitration, actions by other controllers, (transferência de dados, finalização da transação)
- Even if bus is atomic, overall set of actions is not
- Can have race conditions among components of different operations

Expl 6.1: Suppose P1 and P2 attempt to write cached block A simultaneously (ambos estão no estado S)

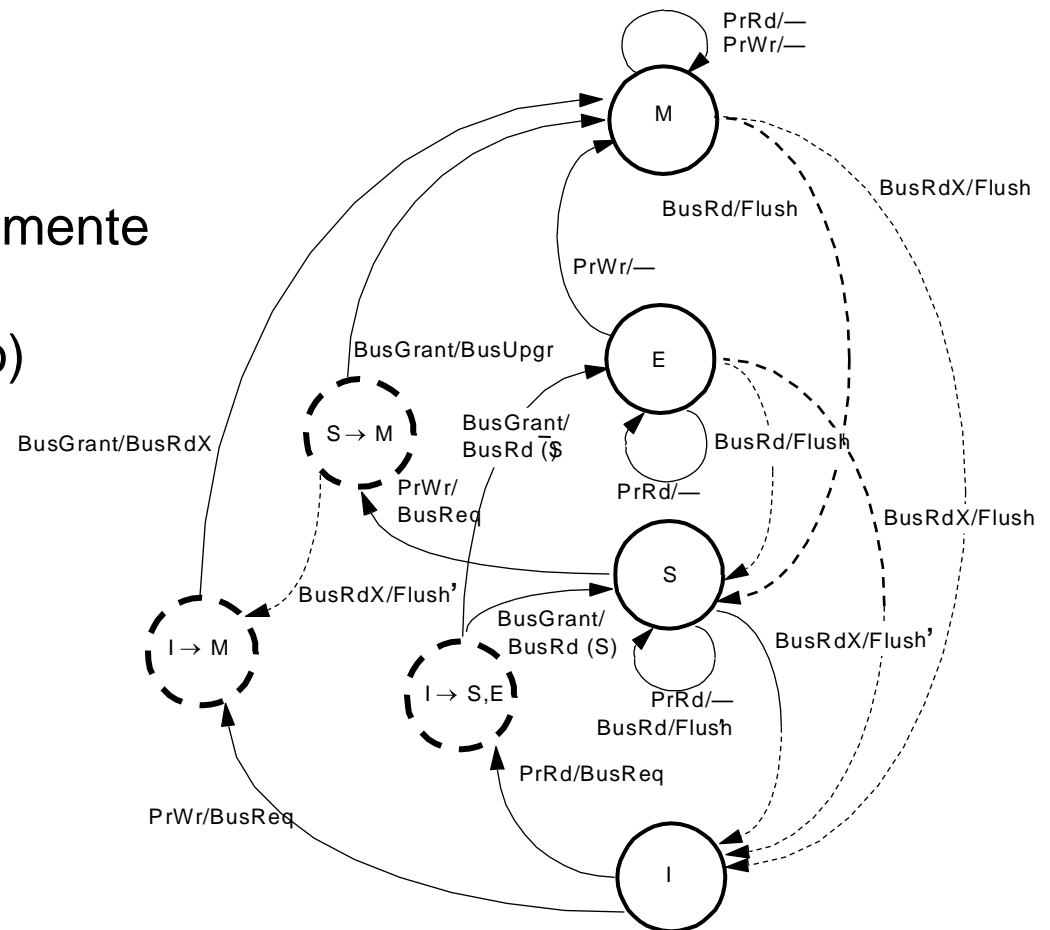
- Each decides to issue BusUpgr to allow S → M
 - Must handle requests for other blocks while waiting to acquire bus
 - Must handle requests for this block A
 - e.g. if P2 wins, P1 must invalidate copy and modify request to BusRdX

Handling Non-atomicity: Transient States

Two types of states

- Stable (e.g. MESI)
- Transient or Intermediate (introduzidos para eventualmente trocar o pedido em função da atividade no barramento)

- Normalmente, os estados instáveis não são codificados no estado de todos os blocos da cache (ficam no controlador)



- Increase complexity (mais difícil de garantir a corretude), so many seek to avoid

- e.g. don't use BusUpgr, rather other mechanisms to avoid data transfer (expl Sun Enterprise)(alguns problemas não aparecem com RdX)

6.2.6 Serialization

Processor-cache handshake must preserve serialization of bus order

- e.g. on write to block in S state, mustn't write data in block until ownership is acquired.
 - other transactions that get bus before this one may seem to appear later

Write completion for SC: needn't wait for inval to actually happen

- Just wait till it gets bus (here, will happen before next bus xaction) (não precisa aguardar a conclusão do RdX, simplesmente ter ganho o bus)
- *Commit* (ordem no bus está estabelecida) versus *complete*
- Don't know when inval actually inserted in destination process's local order, only that it's before next xaction and in same order for all procs
- Local write hits become visible not before next bus transaction
- Same argument will extend to more complex systems
- What matters is not when written data gets on the bus (write back), but when subsequent reads are guaranteed to see it

Write atomicity: if a read returns value of a write W, W has already gone to bus and therefore completed if it needed to

6.2.7, 6.2.8 Deadlock, Livelock, Starvation

Request-reply protocols can lead to protocol-level, *fetch deadlock*

- In addition to buffer deadlock discussed earlier
- When attempting to issue requests, must service incoming transactions
 - e.g. cache controller awaiting bus grant must snoop and even flush blocks
 - else may not respond to request that will release bus: deadlock

Livelock: many processors try to write same line. Each one:

- Obtains exclusive ownership via bus transaction (assume not in cache)
- Realizes block is in cache and tries to write it
- Livelock: I obtain ownership, but you steal it before I can write, etc.
- Solution: don't let exclusive ownership be taken away before write

Starvation: solve by using fair arbitration on bus and FIFO buffers

- May require too much buffering; if retries used, priorities as heuristics

6.2.9 Implementing Atomic Operations

Read-modify-write: read component and write component

- Cacheable variable, or perform read-modify-write at memory
 - cacheable has lower latency and bandwidth needs for self-reacquisition
 - also allows others spinning in cache without generating traffic while waiting
 - at-memory has lower lock transfer time
 - usually traffic and latency considerations dominate, so use cacheable
- Natural to implement with two bus transactions: read and write
 - can lock down bus (até completar a escrita): okay for atomic bus, but not for split-transaction
 - Mas existe solução melhor (better approach): get exclusive ownership, read-modify-write, only then allow others access (consegue acesso exclusivo e não libera até completar a escrita); melhor porque não bloqueia o barramento para operações em outros blocos
 - compare&swap more difficult in RISC machines (precisa operação: memória, registrador → memória) two registers+memory

Implementing LL-SC

HW Lock flag and lock address register at each processor

LL reads block, sets lock flag, puts block address in register

Incoming invalidations checked against address: if match, reset flag

- Also if block is replaced and at context switches

SC checks lock flag as indicator of intervening conflicting write

- If reset, fail; if not, succeed

Livelock considerations

- Don't allow replacement of lock variable between LL and SC
 - split (instruction and data cache) or set-assoc. (unified) cache,
 - or don't allow memory accesses between LL, SC
 - (also don't allow reordering of accesses across LL or SC) (porque isso poderia colocar outras instruções entre LL e SC)
- Don't allow failing SC to generate invalidations (not an ordinary write) (como aconteceria em um write comum)

Performance: both LL and SC can miss in cache

(2 misses no SharedState-SC x 1 miss no r-m-w)

- Prefetch block in exclusive state at LL (para evitar misses)
- But exclusive request reintroduces livelock possibility: use backoff

6.3 Multi-level Cache Hierarchies

How to snoop with multi-level caches? (mostrar fig 6.6)

- independent bus snooping at every level?
- muito caro e inadequado (\$1 on chip; precisaria pinos especiais para monitorar barramento; tag duplicada); saída → inclusion
- maintain cache inclusion (característica usual)

Requirements for *Inclusion*

- data in higher-level cache is subset of data in lower-level cache
- modified in higher-level (M em MESI ou Sm no Dragon) => marked modified in lower-level

Now only need to snoop lowest-level cache

- If L2 says not present (modified), then not so in L1 too
- If BusRd seen to block that is modified in L1, L2 itself knows this

Is inclusion automatically preserved?

- Replacements: all higher-level misses go to lower level
- Modifications (em estados, devem ser propagados cima/baixo)

Violations of Inclusion

The two caches (L1, L2) may choose to replace different block

- Differences in reference history
 - set-associative first-level cache with LRU replacement (história diferente da L2)
 - example: blocks m1, m2, m3 fall in same set of L1 cache... (ver texto)
- Split higher-level caches (instruções e dados)
 - instruction, data blocks go in different caches at L1, but may collide in L2
 - what if L2 is set-associative?
- Differences in block size

But a common case works automatically

- L1 direct-mapped, fewer sets than in L2, and block size same
 - (desde que bloco carregado em L1 também esteja presente em L2)

Preserving Inclusion Explicitly

Em algumas configurações não é possível garantir inclusão sem ação explícita

Propagate lower-level (L2) replacements to higher-level (L1)

- Invalidate or flush (if dirty) messages

Propagate bus transactions from L2 to L1

- Propagate all transactions (nem todas são relevantes para L1), or use inclusion bits (indica quais blocos estão em L1 → evita tráfego L1-L2)

Propagate modified state from L1 to L2 on writes?

- Write-through L1, or modified-but-stale bit (indica que o dado em L1 é que está atualizado) per block in L2 cache

Correctness issues altered?

- Not really, if all propagation occurs correctly and is waited for (up – down)
- Writes commit when they reach the bus, acknowledged immediately
- But performance problems, so want to not wait for propagation
- Discuss after split-transaction busses

Dual cache tags less important: each cache is filter for other (ver fig 6.7)

6.4 Split-Transaction Bus

Bus atômico desperdiça BW: fios do barramento ficam ociosos entre:
(endereço+com→bus) e (dado → bus)

Tipos de transação: BusRd (request / data),
BusUpgr (request / - / ack), BusRdX (request / data / ack)

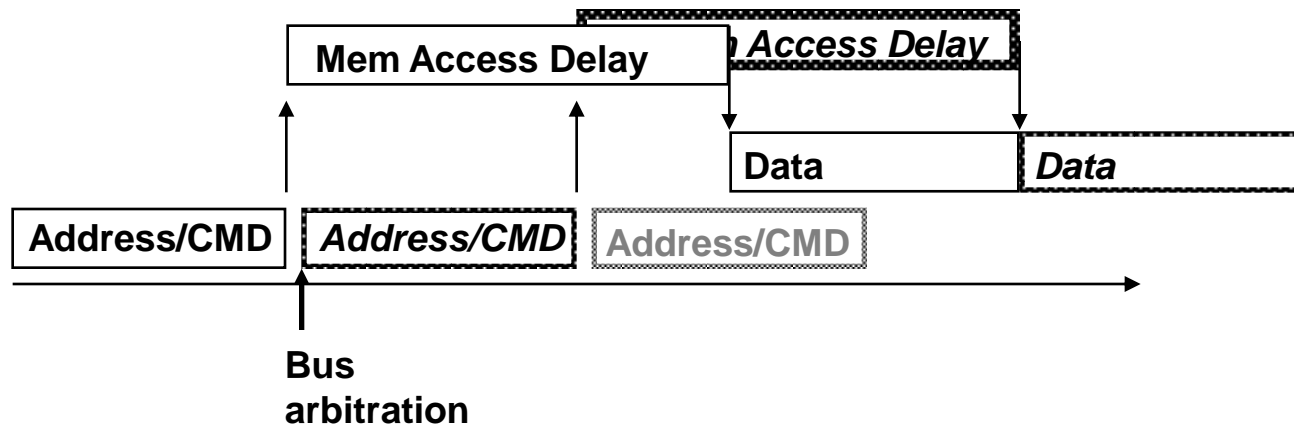
Split bus transaction into request and response sub-transactions

- Separate arbitration for each phase

Other transactions may intervene

- Improves bandwidth dramatically
- Response is matched to request
- Buffering between bus and cache controllers

Reduce serialization down to the actual bus arbitration



Complications

1- New request can appear on bus before previous one serviced

- Even before snoop result obtained
- Conflicting operations to same block may be outstanding on bus
- expl 6.2: P1, P2 write block in S state at same time
 - both get bus before either gets snoop result, so both think they've won
- Note: different from overall non-atomicity discussed earlier

2- Buffers are small, so may need *flow control* (evitar encher)

3- Buffering implies revisiting snoop issues

- When and how snoop results and data responses are provided
- In order w.r.t. requests? (PPro, DEC Turbolaser: yes; SGI, Sun: no)
- Snoop and data response together or separately?
 - SGI together, SUN separately

Large space, much industry innovation: let's look at one example first

6.4.1 Example (based on SGI Challenge)

No conflicting requests for same block allowed on bus

- 8 outstanding requests total, makes conflict detection tractable

Flow-control through *negative acknowledgement (NACK)*

- NACK as soon as request appears on bus, requestor retries
- Separate command (incl. NACK) + address and tag + data buses

Responses may be in different order than requests

- Order of transactions determined by requests
- Snoop results presented on bus with response

Look at (próximas transparências)

- Bus design, and how requests and responses are matched
- Snoop results and handling conflicting requests
- Flow control
- Path of a request through the system

6.4.2 Bus Design and Req-Resp Matching

Essentially two separate buses, arbitrated independently

- “Request” bus for command and address
- “Response” bus for data

Out-of-order responses imply need for matching req-response

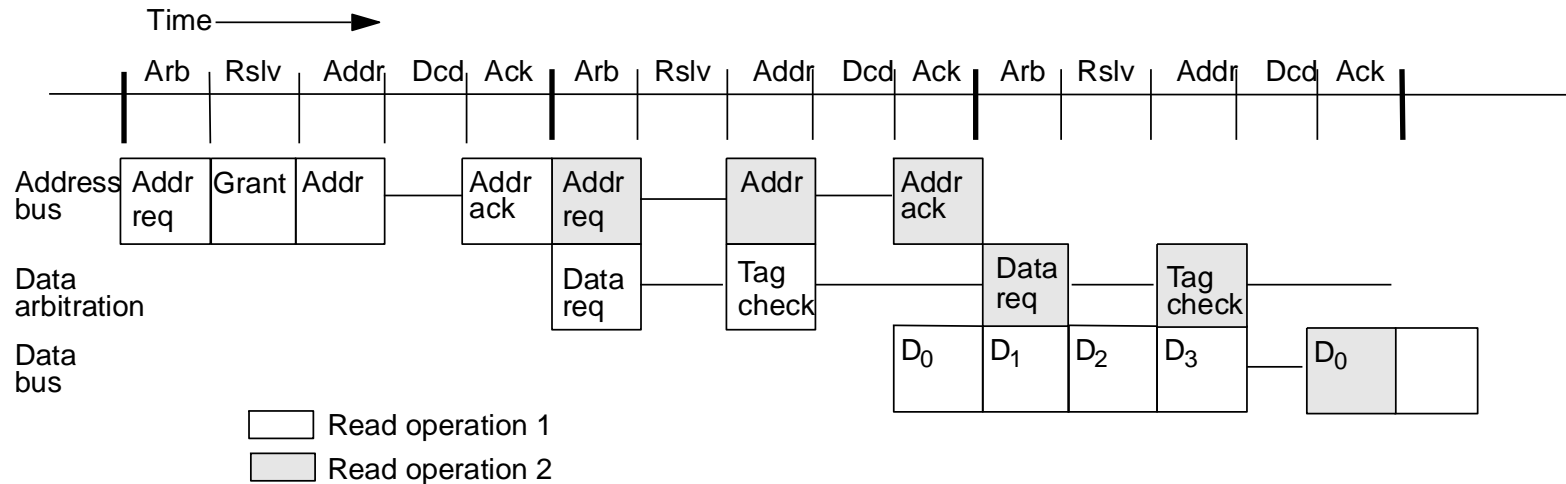
- Request gets 3-bit tag when wins arbitration (8 outstanding max)
- Response includes data as well as corresponding request tag
- Tags allow response to not use address bus, leaving it free

Separate bus lines for arbitration, and for snoop results

Bus Design (continued)

Each of request and response phase is 5 bus cycles (best case)

- Response: 4 cycles for data (bloco = 128 bytes = 1024 bits, 256-bit bus), 1 turnaround (para a resposta)
- Request phase (pipeline uniforme, também 5 ciclos): arbitration, resolution, address, decode, ack
- Request-response transaction takes 3 or more of these (address req, data req, data xfer = response)



Cache tags looked up in decode; extend ack cycle if not possible

- Determine who will respond, if any
- Actual response comes later, with re-arbitration

Write-backs have request phase only: arbitrate both data+addr buses (transmite dados junto com o request)

Upgrades have only request part; ack'ed by bus on grant (commit)

Bus Design (continued)

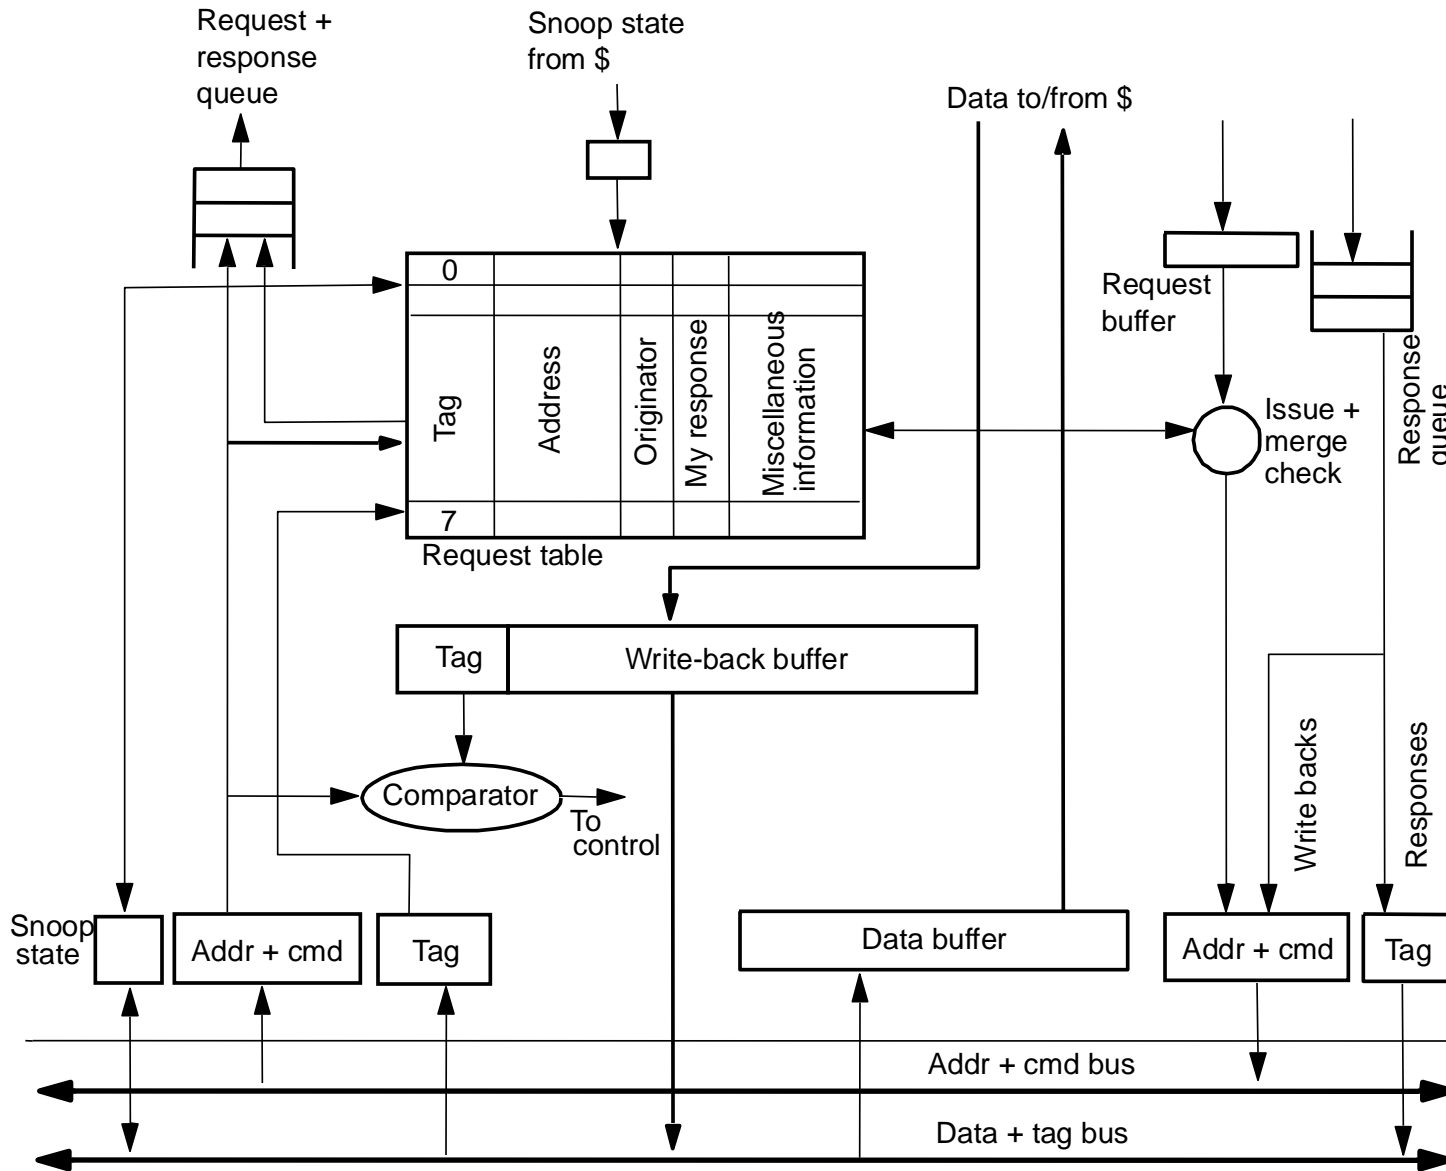
Tracking outstanding requests and matching responses

- Eight-entry “request table” in each cache controller
- New request on bus added to all at same index (3-bit tag), determined by tag
- Entry holds address, request type, state in that cache (if determined already), ...
- All entries checked on bus or processor accesses for match, so fully associative
- Entry freed when response appears, so tag can be reassigned by bus

Bus Interface with Request Table



confrontar



Adaptado dos slides da editora por Mario Côrtes – IC/Unicamp

6.4.3 Snoop Results and Conflicting Requests

Variable-delay snooping

Shared, dirty and inhibit (pode estender a duração da response phase) wired-OR lines, as before

Snoop results *presented* when response appears

- *Determined* earlier, in request phase, and kept in request table entry (naquela fase já se sabia quem forneceria o dado, mas pode levar tempo até o dado estar pronto)
- (Also determined who will respond)
- Writebacks and upgrades don't have data response or snoop result (ver início seção 6.4)

Avoiding conflicting requests on bus

- easy: don't issue request for conflicting request that is in request table

Recall (lembrar que) writes (foram) committed when request gets bus

6.4.4 Flow Control

Not just at incoming buffers from bus to cache controller (já visto)

Cache system's buffer for responses to its requests

- Controller limits number of outstanding requests, so easy (usar NACK)

Mainly needed (flow control) at main memory in this design

- Each of the 8 transactions can generate a writeback
- Can happen in quick succession (no response needed) → risco de buffer overflow
- SGI Challenge: separate NACK lines for address and data buses
 - Asserted before ack phase of request (response) cycle is done
 - Request (response) cancelled everywhere, and retries later
 - Backoff and priorities to reduce traffic and starvation
- SUN Enterprise: destination (em vez da origem iniciar) initiates retry when it has a free buffer
 - source keeps watch for this retry
 - guaranteed space will still be there, so only two “tries” needed at most

6.4.5 Handling a Read Miss

No caso de RD miss → Need to issue BusRd

First check request table. If hit:

- 1- If prior request exists for same block, want to grab data too!
 - “want to grab response” bit
 - “original requestor” bit
 - non-original grabber must assert sharing line so others will load in S rather than E state
- 2- If prior request incompatible with BusRd (e.g. BusRdX)
 - wait for it to complete and retry (processor-side controller)
- If no prior request (naquele instante), issue request and watch out for race conditions
 - conflicting request may win arbitration before this one, but this one receives bus grant before conflict is apparent
 - watch for conflicting request in slot before own (continuar olhando), degrade request to “no action” and withdraw till conflicting request satisfied

Upon Issuing the BusRd Request

All processors enter request into table, snoop for request in cache

Memory starts fetching block

Três possibilidades:

1. Cache with dirty block responds before memory ready

- Memory aborts on seeing response
- Waiters grab data
 - some may assert inhibit to extend response phase till done snooping
 - memory must accept response as WB (might even have to NACK se o seu buffer estiver cheio)

2. Memory responds before cache with dirty block

- Cache with dirty block asserts inhibit line till done with snoop
- When done, asserts dirty, causing memory to cancel response
- Later, cache with dirty issues response, arbitrating for bus

3. No dirty block: memory responds when inhibit line released

- Assume cache-to-cache sharing not used (for non-modified data)

Handling a Write Miss

Similar to read miss, except, (se não encontrar dado na cache):

- Generate BusRdX
- Main memory does not sink response since will be modified again
- No other processor can grab the data

If block present in shared state, issue BusUpgr instead

- No response needed
- If another processor was going to issue BusUpgr, changes to BusRdX as with atomic bus

6.4.6 Write Serialization

With split-transaction buses, usually bus order is determined by order of *requests* appearing on bus


- actually, the ack phase, since requests may be NACKed
- by end of this phase, they are committed for visibility in order

A write that follows a read transaction to the same location should not be able to affect the value returned by that read

- Easy in this case, since conflicting requests (para a mesma posição de memória) not allowed
- Read response precedes write request on bus

Similarly, a read that follows a write transaction won't return old value

Detecting Write Completion

Problem: invalidations don't happen as soon as request appears on bus (ver exemplo 6.3) 

- They're buffered between bus and cache
- Commitment does not imply performing or completion
- Need additional mechanisms

Key property to preserve: processor shouldn't see new value produced by a write before previous writes in bus order are visible to it

1. Don't let certain types of incoming transactions be reordered in buffers
 - in particular, data reply (para RD miss ou WR comittment ack) should not overtake invalidation request
 - okay for invalidations to be reordered: only reply actually brings data in (não re-ordenar reply; aplicar todas invalidações anteriores antes dos replys)
2. Allow reordering in buffers, but ensure important orders preserved at key points
 - e.g. flush incoming invalidations/updates from queues and apply before processor completes operation that may enable it to see a new value

Commitment of Writes (Operations)

More generally, distinguish between *performing* and *commitment* of a write w :

Performed w.r.t a processor: invalidation actually applied

Committed w.r.t a processor: guaranteed that once that processor sees the new value associated with W , any subsequent read by it will see new values of all writes that were committed w.r.t that processor before W .

Global bus serves as point of commitment, if buffers are FIFO

- benefit of a serializing broadcast medium for interconnect

Note: acks from bus to processor must logically come via same FIFO

- not via some special signal, since otherwise can violate ordering

Write Atomicity

Still provided naturally by broadcast nature of bus

Recall that bus implies:

- writes commit in same order w.r.t. all processors
- read cannot see value produced by write before write has committed on bus and hence w.r.t. all processors

Previous techniques allow substitution of “complete” for “commit” in above statements

- that’s write atomicity

Will discuss deadlock, livelock, starvation after multilevel caches plus split transaction bus

6.4.7 Alternatives: In-order Responses

FIFO request table suffices (ainda é necessário busca fully assoc. para bloquear pedidos conflitantes)

Dirty cache does not release inhibit line till it is ready to supply data

- No deadlock problem since does not rely on anyone else

But performance problems possible at interleaved memory

- Major motivation for allowing out-of-order responses

In-order responses allow conflicting requests more easily

- Two BusRdX requests one after the other on bus for same block
 - latter controller invalidates its block, as before
 - but earlier requestor sees later request before its own data response
 - with out-of-order response, not known which response will appear first
 - with in-order, known, and actually can use performance optimization
 - earlier controller responds to latter request by noting that latter is pending
 - when its response arrives, updates word, short-cuts block back on to bus, invalidates its copy (reduces ping-pong latency)

Other Alternatives

Fixed delay from request to snoop result also makes it easier

- Can have conflicting requests even if data responses not in order
- e.g. SUN Enterprise
 - 64-byte line and 256-bit bus => 2 cycle data transfer
 - so 2-cycle request phase used too, for uniform pipelines
 - too little time to snoop and extend request phase
 - snoop results presented 5 cycles after address (unless inhibited)
 - by later data response arrival, conflicting requestors know what to do

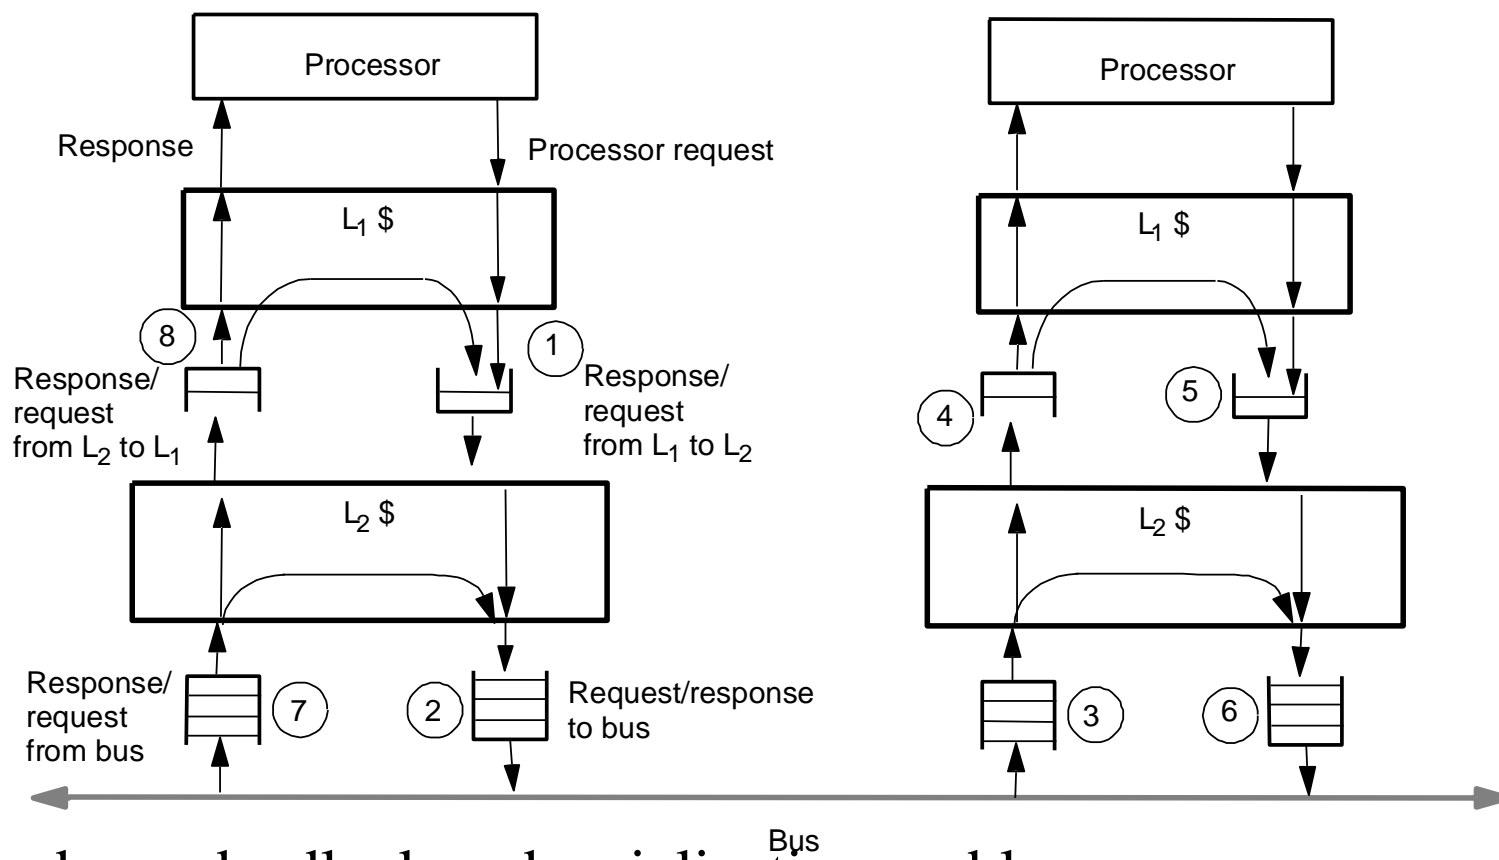
Don't even need request to go on same bus, as long as order is well-defined

- SUN SparcCenter2000 had 2 ST busses, Cray 6400 had 4 ST busses
- Multiple requests go on bus (um para cada bus) in same cycle
- Priority order established among them is logical order

Multi-Level Caches with ST Bus

Key new problem: many cycles to propagate through hierarchy

- Must let others propagate too for bandwidth, so queues between levels
- (ver sequencia de passos – RDMiss – a na legenda da figura 6.10)



Introduces deadlock and serialization problems

Deadlock Considerations (with multi-level caches + ST bus)

Fetch deadlock:

- Must buffer incoming requests/responses while request outstanding
- One outstanding request per processor (não precisa de buffer entre Proc e L1) => need space to hold p requests plus one reply (latter is essential)
- If smaller (or if multiple o/s requests), may need to NACK
- Then need priority mechanism in bus arbiter to ensure progress (evitar deadlock) (reservar no minimo um slot para resposta)

Buffer deadlock:

- L1 to L2 queue filled with read requests, waiting for response from L2
- L2 to L1 queue filled with bus requests waiting for response from L1
- Latter condition only when cache closer than lowest level is write back
- Could provide enough buffering, or general solutions discussed later

If max outstanding bus transactions smaller than total o/s cache misses, response from cache must get bus before new requests from it allowed

- Queues may need to support bypassing (passar na frente na fila)

Sequential Consistency (with multi-level caches + ST bus)

Separation of commitment from completion even greater now

- More performance-critical that commitment replace completion

Fortunately techniques for single-level cache and ST bus extend

- Just use them at each level
- i.e. either don't allow certain reorderings of transactions at any level
- Or don't let outgoing operation proceed past level before incoming invalidations/updates at that level are applied

6.4.9 Multiple Outstanding Processor Requests

So far assumed only 1 OS / processor: not true of modern processors

Danger: operations from same processor can complete out of order

- e.g. write buffer: until serialized by bus, should not be visible to others
- Uniprocessors use write buffer to insert multiple writes in succession
 - multiprocessors usually can't do this while ensuring consistent serialization
 - exception: writes are to same block, and no intervening ops in program order

Key question: who should wait to issue next op till previous completes

- Key to high performance: processor needn't do it (so can overlap)
- Queues/buffers/controllers can ensure writes not visible to external world and reads don't complete (even if back) until allowed (more later)

Other requirement: caches must be lockup free (ver texto) to be effective

- Merge operations to a block, so rest of system sees only one o/s to block

All needed mechanisms for correctness available (deeper queues for performance)

6.5 Case Studies of Bus-based Machines

SGI Challenge, with Powerpath bus

SUN Enterprise, with Gigaplane bus

- Take very different positions on the design issues discussed above

Overview

For each system:

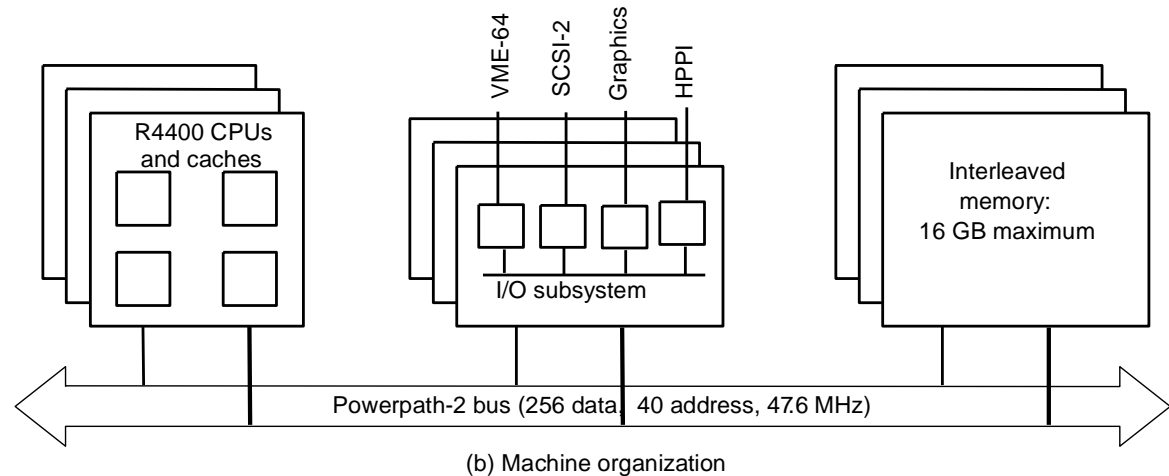
- Bus design
- Processor and Memory System
- Input/Output system
- Microbenchmark memory access results

Application performance and scaling (SGI Challenge)

SGI Challenge Overview



(a) A four-processor board



36 MIPS R4400 (peak 2.7 GFLOPS, 4 per board) or 18 MIPS R8000 (peak 5.4 GFLOPS, 2 per board)

8-way interleaved memory (up to 16 GB)

4 I/O busses of 320 MB/s each


1.2 GB/s Powerpath-2 bus @ 47.6 MHz, 16 slots, 329 signals

128 Bytes lines (1 + 4 cycles): $128B * 8bit = 1Kbit$

Split-transaction with up to 8 outstanding reads

- all transactions take five cycles

SUN Enterprise Overview



Up to 30 UltraSPARC processors (peak 9 GFLOPs)

Gigaplane™ bus has peak bw 2.67 GB/s; upto 30GB memory

16 bus slots, for processing or I/O boards

- 2 CPUs and 1GB memory per board
 - memory distributed, unlike Challenge, but protocol treats as centralized (acessada via barramento, portanto acesso uniforme)
- Each I/O board has 2 64-bit 25Mhz SBUSes

Bus Design Issues

Multiplexed versus non-multiplexed (separate addr and data lines)

Wide versus narrow data busses

Bus clock rate

- Affected by signaling technology, length, number of slots...

Split transaction versus atomic

Flow control strategy

6.5.1 SGI Powerpath-2 Bus

Non-multiplexed, 256-data/48-address (+cmd), 47.6 MHz, split transaction supporting 8 o/s requests

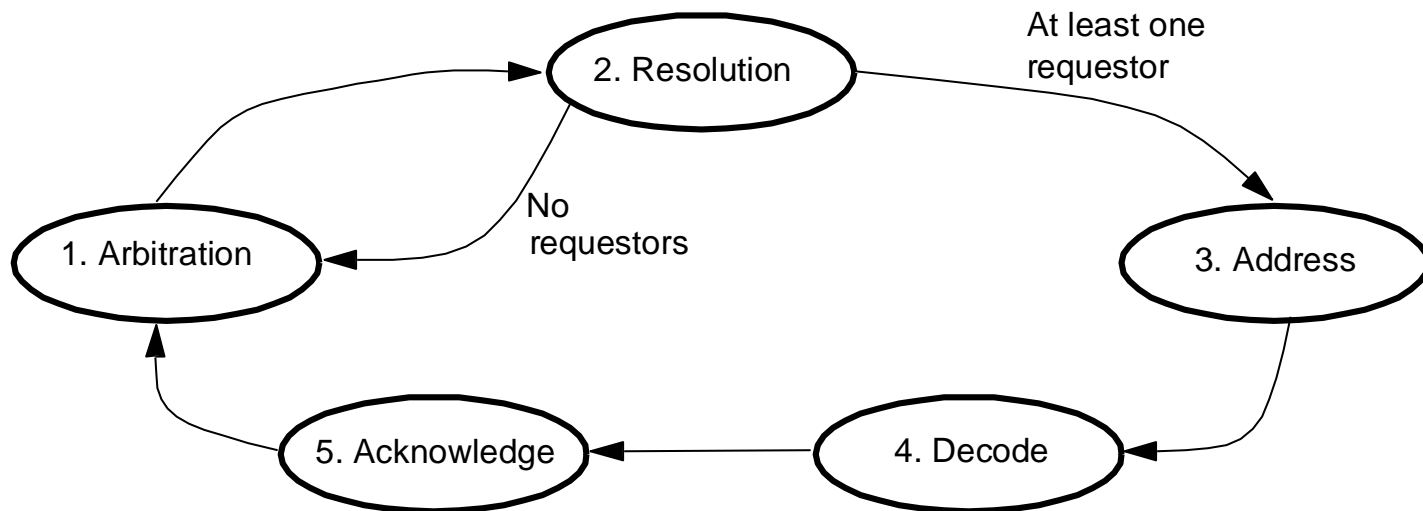
Wide => more interface chips so higher latency, but more bw at slower clock

Large block size also calls for wider bus

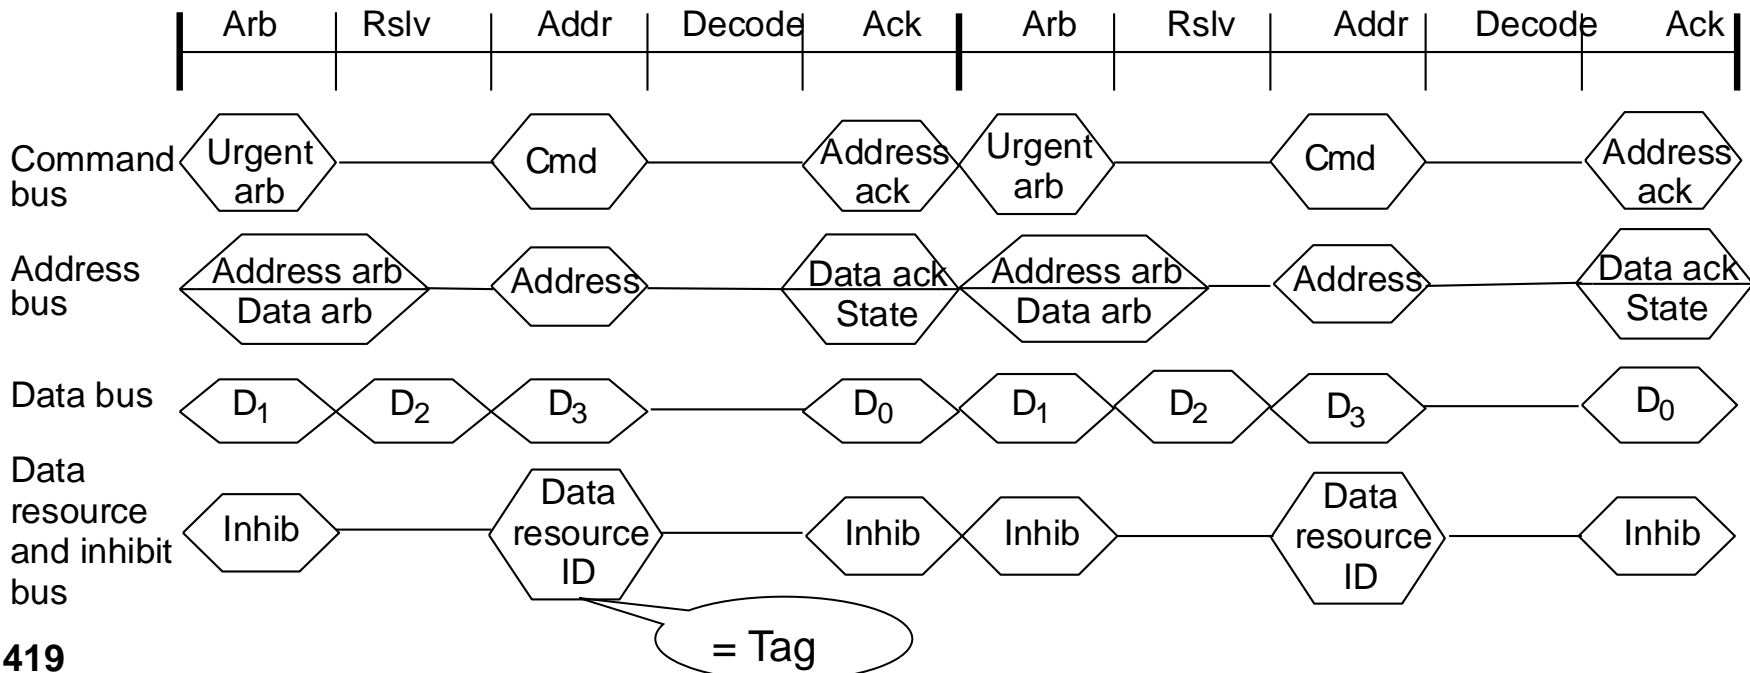
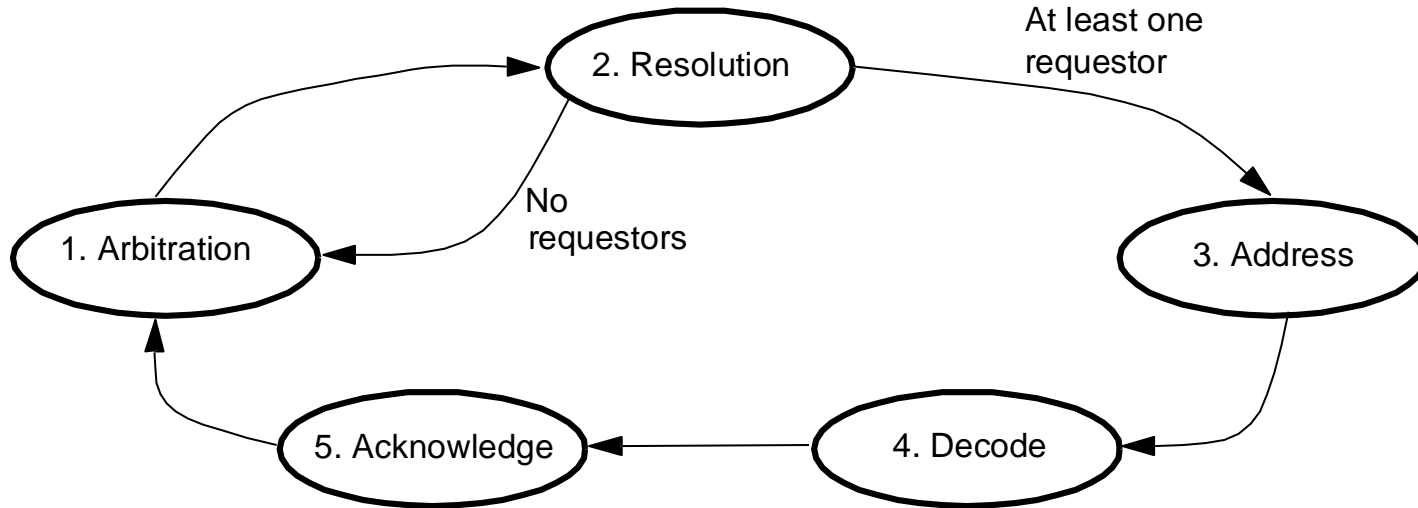
Uses Illinois MESI protocol (cache-to-cache sharing)

More detail in chapter (ver 16+16+16 bits e urgent bit lines para prevenir starvation)

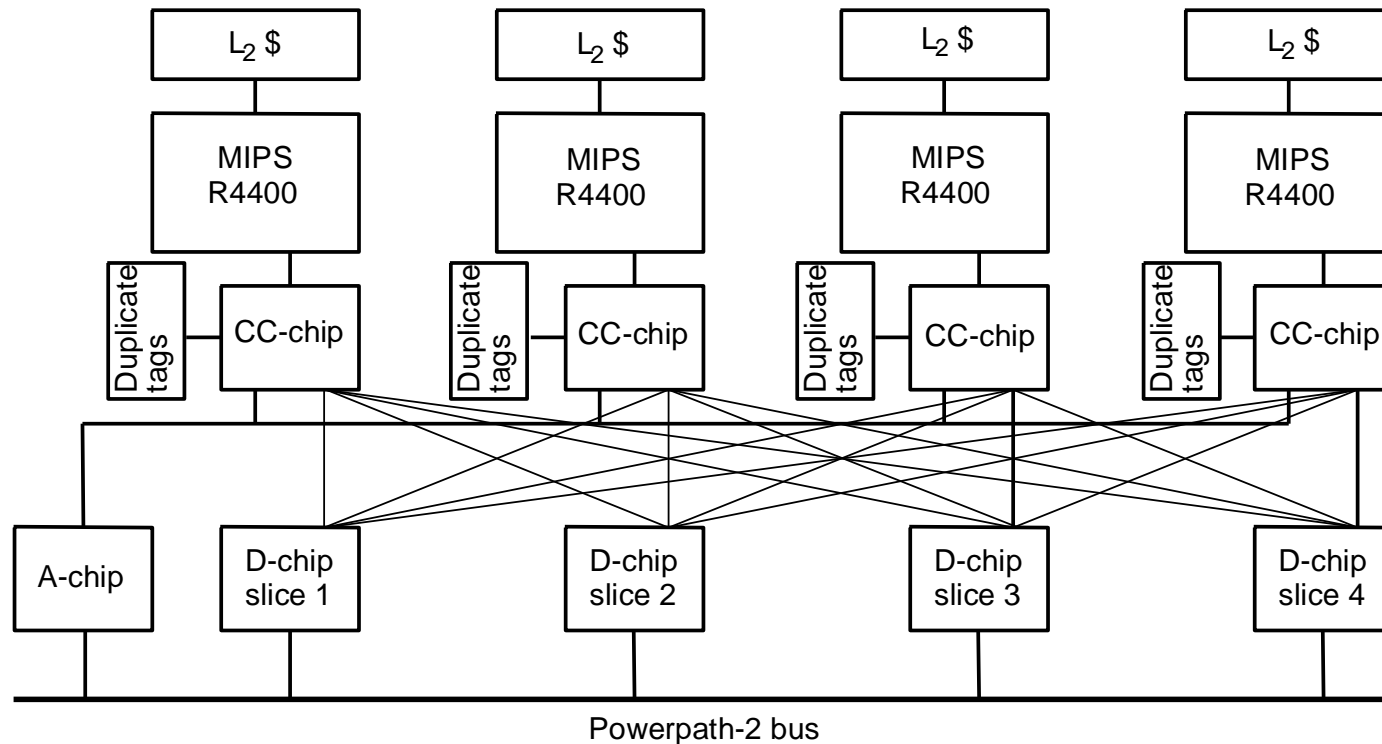
Na ausência de transações, máquina de espera (2 estados)



Bus Timing (detalhamento fig. 6.8)



6.5.2 Processor and Memory Systems



4 MIPS R4400 processors per board share A and D chips

A-chip has address bus interface, request table (8 linhas), control logic

CC (cache coherence) chip per processor has duplicate set of tags

Processor requests go from CC chip to A chip to bus

4 bit-sliced D chips interface CC chip to bus (4*64=256 bits); algum buffering

Memory Access Latency

Largura da memória 512 bits + ECC: linha da cache (1Kb)
carregada na memória em 2 ciclos

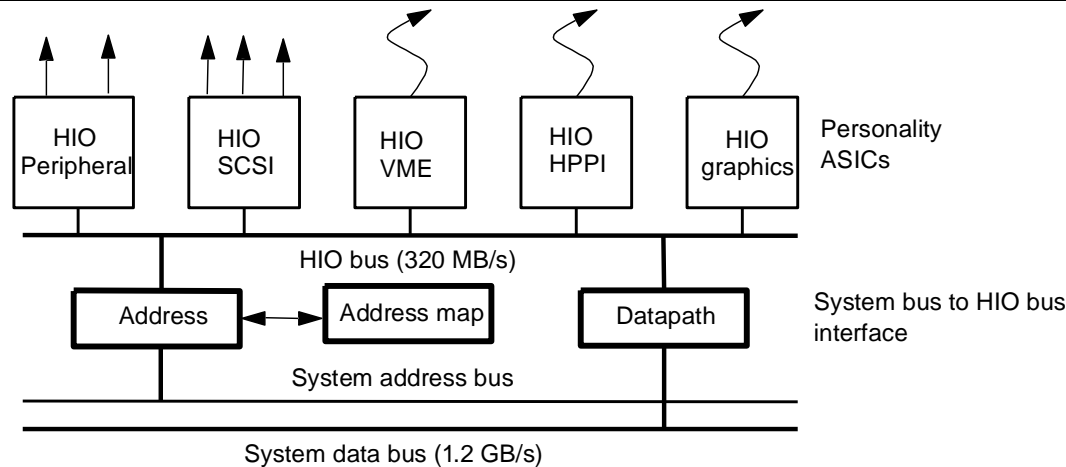
Memória é 2-way interleaved em cada placa (satura o bus)

250ns (12 ciclos) access time from address on bus to data on bus

But overall latency (L2 miss) seen by processor is 1000ns!

- 300 ns for request to get from processor to bus
 - down through cache hierarchy, CC chip and A chip
- 400ns later, data gets to D chips
 - 3 bus cycles to address phase of request transaction, 12 to access main memory, 5 to deliver data across bus to D chips
- 300ns more for data to get to processor chip
 - up through D chips, CC chip, and 64-bit wide interface to processor chip, load data into primary cache, restart pipeline

6.5.3 Challenge I/O Subsystem



Multiple I/O cards on system bus, each has 320MB/s HIO bus

- Personality ASICs connect these to devices (standard (Ethernet, SCSI, VME etc) and graphics)

Proprietary HIO bus


- 64-bit multiplexed address/data, same clock as system bus
- Split read transactions, up to 4 per device
- Pipelined, but centralized arbitration, with several transaction lengths
- Comunicação via DMA: address translation via mapping RAM in system bus interface

Why the decouplings? (Why not connect directly to system bus?) (HIO é menor 64 do que system bus 256)

I/O board acts like a processor to memory system

6.5.4 Challenge Memory System Performance

Read microbenchmark with various strides (afastamento nas leituras sucessivas ou amplitude do endereçamento) and array sizes



Ping-pong flag-spinning microbenchmark: round-trip time 6.2 μ s.

6.5.5 Sun Gigaplane Bus

Non-multiplexed, split-transaction, 256-data/41-address, 83.5 MHz

- Plus 32 ECC lines, 7 tag, 18 arbitration, etc. Total 388.

Cards plug in on both sides: 8 per side (total: 16 placas)

112 outstanding transactions, up to 7 from each board (7*16)

- Designed for multiple outstanding transactions per processor (lockup free caches)

Emphasis on reducing latency, unlike Challenge

- Speculative arbitration (collision based) if address bus not scheduled from prev. cycle
- Else regular 1-cycle arbitration, and 7-bit tag assigned in next cycle


Snoop result associated with request phase (5 cycles later)

Main memory can stake claim to data bus 3 cycles into this, and start memory access speculatively

- Two cycles later, asserts tag bus to inform others of coming transfer

MOESI protocol (owned state for cache-to-cache sharing)

Gigaplane Bus Timing



- **Duas operações BusRd (branco e cinza)**
- **Convenções: A D A D A D = address e dados; 1,2,3 etc= número da Board envolvida na transação; setas: caminho do 1o. BusRd**
- **Snoop signal nas linhas state: shared, owned, mapped, ignore.**
- **Board1 inicia com fast arbitration (já coloca o endereço); bem sucedida; Board2 responde (3) antes do resultado do snoop (5);**
- **Segundo BusRd: colisão entre 4 e 5; Board4 ganha; Board6 arbitra (7) pelo data bus e cancela em seguida (12) porque resultado snoop indica outra cache tem o dado; Board7 responde com o dado (11)**

6.5.6 Enterprise Processor and Memory System

2 procs per board, external L2 caches, 2 mem banks with x-bar

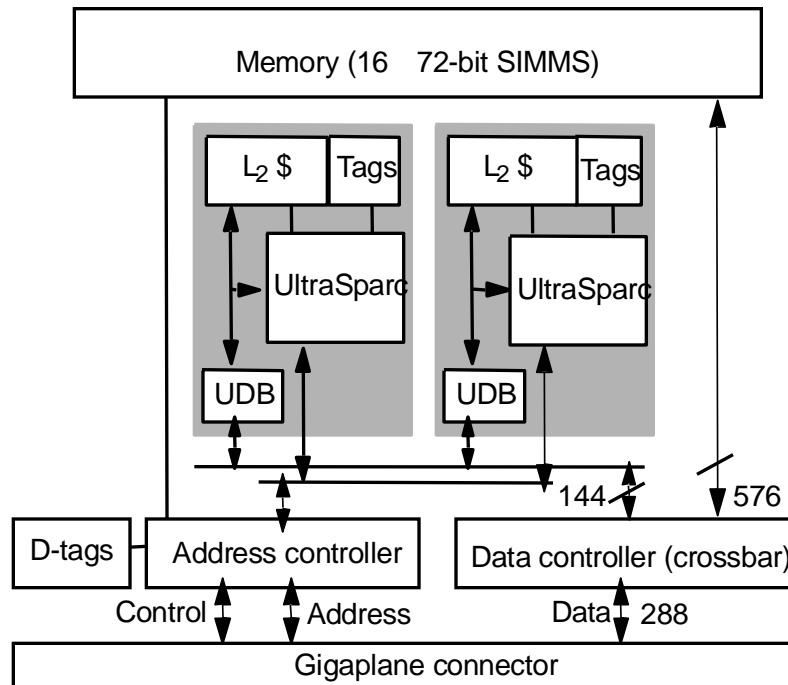
Data lines buffered through UDB to drive internal 1.3 GB/s UPA bus

Wide path to memory so full 64-byte (512b) line in 1 mem cycle (2 bus cyc)

Dtags = duplicate tags para a cache L2

Addr controller adapts proc and bus protocols, does cache coherence

- its tags keep a subset of states needed by bus (e.g. no M/E distinction)



6.5.7 Enterprise I/O System

I/O board has same bus interface
ASICs as processor boards

But internal bus half as wide, and no
memory path

Only cache block sized
transactions, like processing
boards

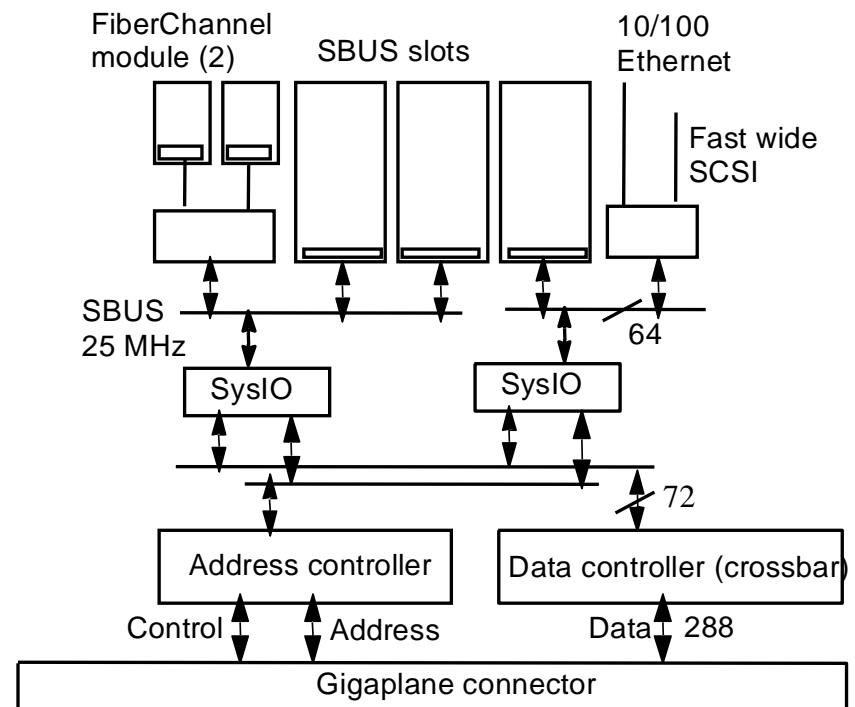
- Uniformity simplifies design
- ASICs implement single-block cache, follows coherence protocol

SysIO é percebido pelo barramento
como uma linha de cache

Performance and cost of I/O scale
with no. of I/O boards

Two independent 64-bit, 25 MHz Sbuses

- One for two dedicated FiberChannel modules connected to disk
- One for Ethernet and fast wide SCSI
- Can also support three SBUS interface cards for arbitrary peripherals




6.5.8 Memory Access Latency

300ns read miss latency

11 cycle min bus protocol at 83.5 Mhz is 130ns of this time

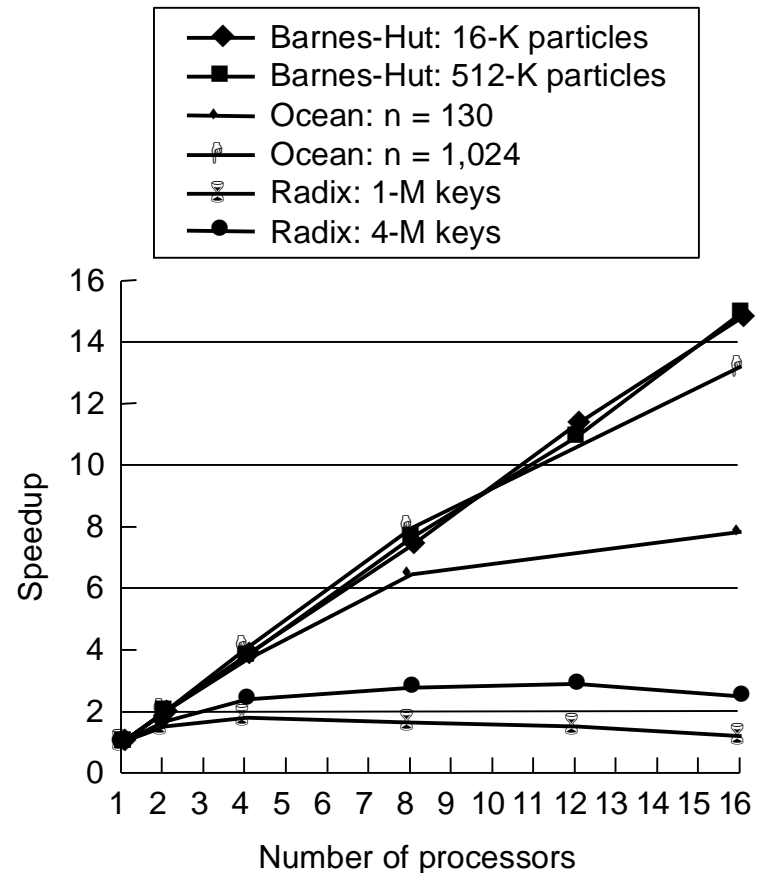
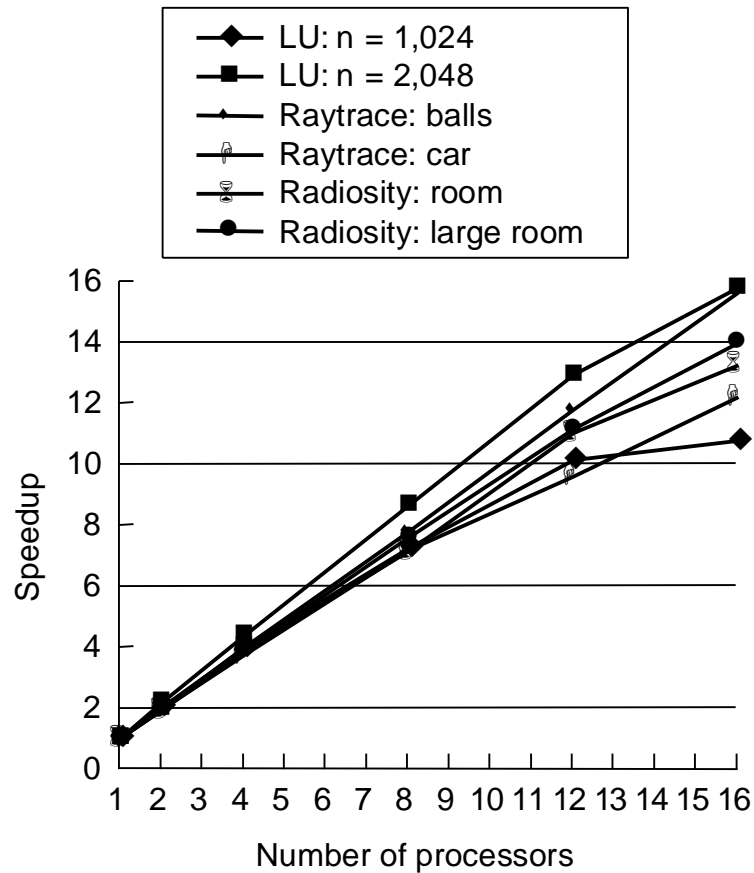
Rest is path through caches and the DRAM access

TLB misses add 340 ns



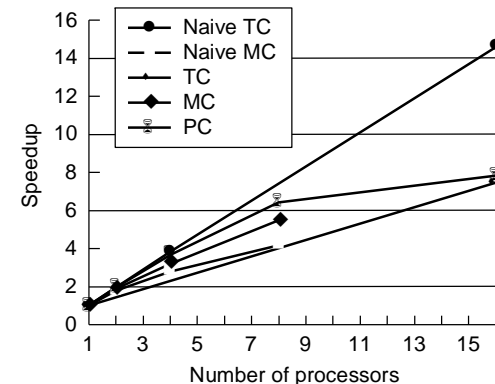
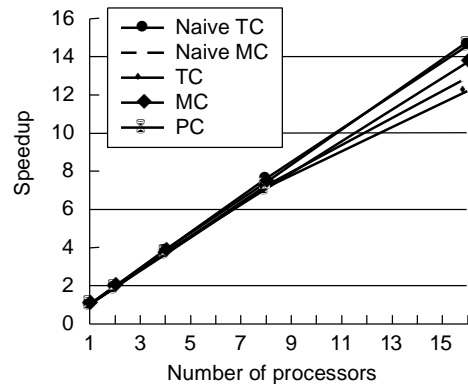
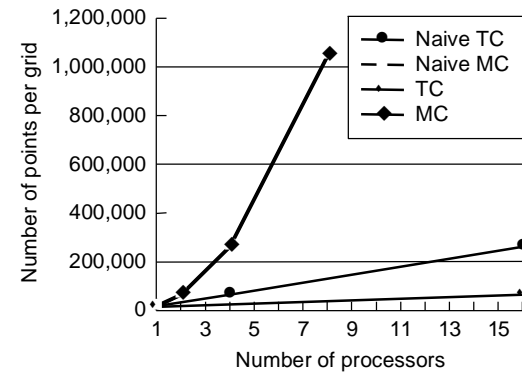
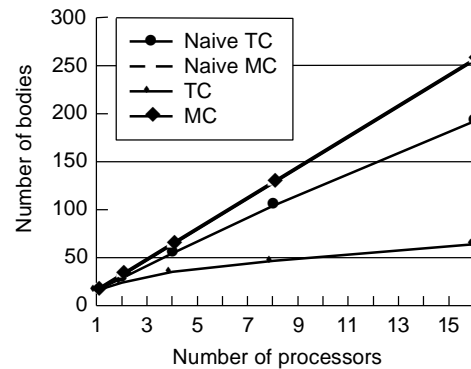
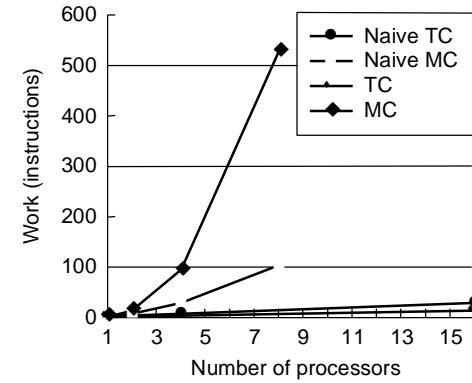

Ping-pong microbenchmark is 1.7 μ s round-trip (5 mem accesses)

6.5.9 Application Speedups (Challenge)



- Problem in Ocean with small problem: communication and barrier cost
- Problem in Radix: contention on bus due to very high traffic
 - also leads to high imbalances and barrier wait time

Application Scaling under Other Models



Barnes – Hut

Ocean

PC: problem-constrained
 TC: time-constrained
 MC: memory-constrained

Expl 6.3 pag 407



Cache 1 nível, split transaction (multiple O/S), reordenação de buffers processador/cache/bus é permitida, A e B = zero inicialmente. Que resultados são proibidos sob SC? Como podem acontecer?

P1	P2
A=1	rd B
B=1	rd A

P1	P2
A=1	B=1
rd B	rd A

- Proibido = AB = (0,1)
- P1 WR A commits (bus)
- P1 inicia WR em B (sem esperar, condições estendidas de SC)
- Invalidações em P2 são reordenadas (B antes de A)
- P2 tem Rd miss em B e lê novo valor B=1
- Quando P2 executa rd A, invalidação ainda no buffer
- P2 tem Rd Hit em A e lê A=0 (valor antigo)
- Proibido = AB = (0,0)
- P1 Wr A commits (bus)
- Continua e lê B=0 (velho, OK)
- P2 Wr B commits (bus)
- P2 prossegue e lê A
- Wr de B entra no bus depois do Wr de A (bus order) e P2 deveria ler o novo valor de A
- Mas a invalidação de P1 (Wr A) ainda está no buffer de entrada de \$2
- P2 tem Rd Hit em A e lê A=0 (velho)