

HW3 Solutions

1. (P11.1) Using the syntax in Figure 11-2, show how to use the load-linked/store conditional primitives to synthesize a compare-and-swap operation.

```
/* r1 contains compare value, r2 contains swap value */
cmpswap: ll r0, A
          cmp r0,r1
          bne fail
          stc r2, A
          bfail cmpswap
fail:     ...
```

2. (P11.8) Real coherence controllers include numerous transient states in addition to the ones shown in Figure to support split-transaction buses. For example, when a processor issues a bus read for an invalid line (I), the line is placed in a IS transient state until the processor has received a valid data response that then causes the line to transition into shared state (S). Given a split-transaction bus that separates each bus command (bus read, bus write, and bus upgrade) into a request and response, augment the state table and state transition diagram of Figure to incorporate all necessary transient states and bus responses. For simplicity, assume that any bus command for a line in a transient state gets a negative acknowledge (NAK) response that forces it to be retried after some delay.

Note: for writebacks, we assume that once the data shows up on the bus as a BD command, the processor issuing the writeback also sees the BD and can then transition to I or S. Similarly, we assume that any subsequent bus read/write will then be satisfied by memory (this is sometimes called the writeback race).

3. (P11.10) Assuming a processor frequency of 1 GHz, a target CPI of 2, a per-instruction level-2 cache miss rate of 1% per instruction, a snoop-based cache coherent system with 32 processors, and 8-byte address messages (including command and snoop addresses), compute the inbound and outbound snoop bandwidth required at each processor node.

Outbound snoop rate = $.01 \text{ miss/inst} \times 1 \text{ inst/2 cyc} \times 1 \text{ cyc/ns} \times 8 \text{ bytes/miss} = .04 \text{ b/ns} = 40 \text{ million bytes per second}$

Inbound snoop rate = $31 \times 40 = 1240 \text{ million bytes per second} = 1182 \text{ MB/sec.}$

4. Cacti problem -- solution not provided
5. Niagara problem. (a) solution not provided, but dedicated core per thread should be fastest, followed by 2 threads per core, followed by 4 threads per core (b) most likely cause is destructive interference due to sharing of L1 data cache. (c) open-ended problem.

MESI cache coherence protocol for Problem 8.

Current State s	Event and Local Coherence Controller Responses and Actions (s' refers to next state)						
	Local Read (LR)	Local Write (LW)	Local Eviction (EV)	Bus Read (BR)	Bus Write (BW)	Bus Upgrade (BU)	Bus Data (BD)
Invalid (I)	Issue bus read if no sharers then s' = IE else s' = IS	Issue bus write s' = IM	s' = I	Do nothing	Do nothing	Do nothing	Error
ItoS (IS)	Stall	Stall	Stall	NAK	NAK	NAK	s' = S
ItoE (IE)	Stall	Stall	Stall	NAK	NAK	NAK	s' = E
ItoM (IM)	Stall	Stall	Stall	NAK	NAK	NAK	s' = M
Shared (S)	Do nothing	Issue bus upgrade s' = M	s' = I	Respond shared	s' = I	s' = I	Error
Exclusive (E)	Do nothing	s' = M	s' = I	Respond shared s' = S	s' = I	Error	Error
Modified (M)	Do nothing	Do nothing	Write data back; s' = I	Respond dirty; Write data back; s' = MS	Respond dirty; Write data back; s' = MI	Error	Error
MtoI (MI)	Do nothing	Stall	Stall	NAK	NAK	NAK	s' = I
MtoS (MS)	Do nothing	Stall	Stall	NAK	NAK	NAK	s' = S