



Constructive Computer Architecture
Tutorial 8

Final Project Part 2: Coherence

Sizhuo Zhang
6.175 TA

Debugging Techniques

- ◆ Deficiency about \$display
 - Everything shows up together
- ◆ Distinct log file for each module: write to file
 - Also see src/unit_test/sc-test/Tb.bsv

```
Ehr#(2, File) file <- mkEhr(InvalidFile);
Reg#(Bool) opened <- mkReg(False);
rule doOpenFile(!opened);
  let f <- $fopen("a.txt", "w");
  if(f == InvalidFile) $finish;
  file[0] <= f; opened <= True;
endrule
rule doPrint;
  $fwrite(file[1], "Hello world\n");
endrule
```

Writing to
InvalidFile will
cause segfault.

Use EHR if the
logic will call
\$fwrite in the first
cycle

Debugging Techniques

- ◆ Deficiency about cycle counter
 - Rule for printing cycle may be scheduled before/after the rule we are interested in
 - Don't want to create a counter in each module
- ◆ Use simulation time
 - `$display("%t: evict cache line", $time);`
 - `$time` returns Bit#(64) representing time
 - In Scemi simulation, `$time` outputs: 10, 30, ...

Debugging Techniques

◆ Add sanity check

◆ Example 1

- Parent is handling upgrade request
- No other child has incompatible state
- Parent decides to send upgrade response
- Check: parent is not waiting for any child (waitc)

◆ Example 2

- D cache receives upgrade response from memory
- Check: must be in WaitFillResp state
- Process the upgrade response
- Check: if in I state, then data in response must be valid, otherwise data must be invalid (data field is Maybe type in the lab)

Coherence Protocol: Differences From Lecture

- ◆ In lecture: address type for byte address
 - Implementation: only uses cache line address
 - $\text{addr} \gg 6$ for 64B cache line
- ◆ In lecture: parent reads data using 0 cycle
 - Implementation: read from memory, long latency
- ◆ In lecture: voluntary downgrade rule
 - No need in implementation
- ◆ In lecture: Parent directory tracks states for all address
 - 32-bit address space \rightarrow huge directory
 - Implementation: usually parent is L2 cache, so only track address in L2 cache
 - We don't have L2 cache

Coherence Protocol: Differences From Lecture

- ◆ Work around for large directory
 - For each child, only tracks addresses in its L1 D cache

```
Vector#(CoreNum, Vector#(CacheRows, Reg#(CacheTag)))  
  tags <- replicateM(replicateM(mkRegU));  
Vector#(CoreNum, Vector#(CacheRows, Reg#(MSI)))  
  states <- replicateM(replicateM(mkReg(I)));
```

- To get MSI state for address a in core i

```
MSI s = tags[i][getIndex(a)] == getTag(a) ?  
        states[i][getIndex(a)] : I;
```

Load-Reserve (lr.w) and Store-Conditional (sc.w)

◆ New state in D cache

- `Reg#(Maybe#(CacheLineAddr)) la <- mkReg(Invalid);`
- Cache line address reserved by lr.w

◆ Load reserve: lr.w rd, (rs1)

- `rd <= mem[rs1]`
- Make reservation: `la <= Valid (getLineAddr(rs1));`

◆ Store conditional: sc.w rd, rs2, (rs1)

- Check la: la invalid or addresses don't match: `rd <= 1`
- Otherwise: get exclusive permission (upgrade to M)
 - ◆ Check la again
 - ◆ If address match: `mem[rs1] <= rs2; rd <= 0`
 - ◆ Otherwise: `rd <= 1`
 - ◆ If cache hit, no need to check again (address already match)
- Always clear reservation: `la <= Invalid`

Load-Reserve (lr.w) and Store-Conditional (sc.w)

◆ Cache line eviction

- Due to replacement, invalidation request ...
- May lose track of reserved cache line
 - ◆ Then clear reservation
- Compare evicted cache line with `la`
 - ◆ If match: `la <= invalid`
- This is how lr.w/sc.w pair ensures atomicity

Reference Memory Model

- ◆ Debug interface returned by reference model is passed into every D cache

```
interface RefDMem;  
    method Action issue(MemReq req);  
    method Action commit(MemReq req,  
        Maybe#(CacheLine) line, Maybe#(MemResp) resp);  
endinterface
```

```
module mkDCache#(CoreID id)(  
    MessageGet fromMem, MessagePut toMem,  
    RefDMem refDMem, DCache ifc);
```

- D cache calls the debug interface `refDMem`
- Reference model will check violation of coherence based on the calls

- ◆ Reference model: `src/ref`

Reference Memory Model

- ◆ `issue(MemReq req)`
 - Called when `req` issued to D cache
 - in `req` method of D cache
 - Give program order to reference model
- ◆ `commit(MemReq req, Maybe#(CacheLine) line, Maybe#(MemResp) resp);`
 - Called when `req` finishes processing (`commit`)
 - `line`: cache line accessed by `req`, set to Invalid if unknown
 - `resp`: response to the core, set to Invalid if no response
- ◆ Reference model checks when `commit` is called
 - `req` can be committed or not
 - `line` value is correct or not (not checked if Invalid)
 - `resp` is correct or not

Adding Store Queue

◆ New behavior for memory requests

- Ld: can start processing when store queue is not empty
- St: enqueue to store queue
- Lr, Sc: wait for store queue to be empty
- Fence: wait for all previous requests to commit (e.g. store queue must be empty)
 - ◆ Ordering memory accesses

◆ Issuing stores from store queue to process

- Only stall when there is a Ld request

Multicore Programs

- ◆ Run programs on 2-core system
- ◆ Single-thread programs
 - Programs/assembly, programs/benchmarks
 - core 1 starts looping forever at the very beginning
- ◆ Multithread programs
 - Programs/mc_bench
 - startup code (crt.S): allocate 128KB local stack for each core
 - main function: fork based on core id

```
int main() {
    int coreid = getCoreId();
    if(coreid == 0) { return core0(); }
    else { return core1(); }
}
```

Multicore Programs: mc_print

- ◆ Easiest one
- ◆ Two cores print "0" and "1" respectively
- ◆ Sample output:

```
---- ../..../programs/build/mc_bench/vmh/mc_print.riscv.vmh ----  
01  
PASSED  
  ■ (no cycle/inst count printed)
```

Multicore Programs: mc_hello

- ◆ Core 0 passes each character of a string to core 1
- ◆ Core 1 prints each character it receives
- ◆ Sample output:

```
---- ../../programs/build/mc_bench/vmh/mc_hello.riscv.vmh ----  
Hello World!  
This message has been written to a software FIFO by core 0 and  
read and printed by core 1.  
PASSED  
  ■ (no cycle/inst count printed)
```

Multicore Programs: mc_produce_consume

- ◆ Larger version of mc_hello
- ◆ Core 1 passes each element of an array to core 0
- ◆ Core 0 checks the data
- ◆ Sample output:

```
---- ../ ../programs/build/mc_bench/vmh/mc_produce_consume.riscv.vmh ----  
Benchmark mc_produce_consume  
Cycles (core 0) = xxx  
Insts (core 0) = xxx  
Cycles (core 1) = xxx  
Insts (core 1) = xxx  
Cycles (total) = xxx  
Insts (total) = xxx  
Return 0  
PASSED
```

Instruction counts may vary
due to variation in busy waiting
time, so IPC is not a good
performance metric.
Execute time is a better metric.

Multicore Programs: mc_median/vvadd/multiply

- ◆ Data parallel: fork-join style
- ◆ Core 0 calculates first half results
- ◆ Core 1 calculates second half results
- ◆ Sample output:

```
----- ../../programs/build/mc_bench/vmh/mc_median.riscv.vmh -----  
Benchmark mc_median  
Cycles (core 0) = xxx  
Insts (core 0) = xxx  
Cycles (core 1) = xxx  
Insts (core 1) = xxx  
Cycles (total) = xxx  
Insts (total) = xxx  
Return 0  
PASSED
```


Multicore Programs: mc_dekker

- ◆ Two cores contend for a mutex (Dekker's algo)
- ◆ After getting into critical section
 - increment/decrement shared counter, print core ID
- ◆ Sample output:

```
---- ../././programs/build/mc_bench/vmh/mc_dekker.riscv.vmh ----  
Benchmark mc_1dekker1  
100110...000  
Core 0 decrements counter by 600  
Core 1 increments counter by 900  
Final counter value = 300  
Cycles (core 0) = xxx  
Insts (core 0) = xxx  
Cycles (core 1) = xxx  
Insts (core 1) = xxx  
Cycles (total) = xxx  
Insts (total) = xxx  
Return 0  
PASSED
```

For implementation with
store queue, fence is
inserted in mc_dekker.

Multicore Programs: mc_spin_lock

◆ Similar to mc_dekker, but use spin lock implemented by lr.w/sc.w

◆ Sample output:

```
---- ../.. /programs/build/mc_bench/vmh/mc_spin_lock.riscv.vmh ----  
Benchmark mc1_spin_lock  
10101...000  
Core 0 increments counter by 300  
Core 1 increments counter by 600  
Final counter value = 900  
Cycles (core 0) = xxx  
Insts (core 0) = xxx  
Cycles (core 1) = xxx  
Insts (core 1) = xxx  
Cycles (total) = xxx  
Insts (total) = xxx  
Return 0  
PASSED
```

Multicore Programs: mc_incrementers

- ◆ Similar to mc_dekker, but use atomic fetch-and-add implemented by lr.w/sc.w
- ◆ Core ID is not printed
- ◆ Sample output:

```
---- ../..../programs/build/mc_bench/vmh/mc_incrementers.riscv.vmh ----  
Benchmark mc_incrementers  
  
core0 had 1000 successes out of xxx tries  
core1 had 1000 successes out of xxx tries  
shared_count = 2000  
Cycles (core 0) = xxx  
Insts (core 0) = xxx  
Cycles (core 1) = xxx  
Insts (core 1) = xxx  
Cycles (total) = xxx  
Insts (total) = xxx  
Return 0  
PASSED
```

Some Reminders

- ◆ Use CF regfile and scoreboard
 - Compiler creates a conflict in my implementation with bypass regfile and pipelined scoreboard
- ◆ Signup for project meeting
 - Half-page progress report
- ◆ Project deadline: 3:00pm Dec 9
- ◆ Final presentation (10min)