Advanced Computer Architecture 27 January 2014

A) A processor has the following cache hierarchy: 32KB, 4-way L1 I-cache, 32 KB 8-way L1 D-cache, 32-byte block; 512 KB, 8-way associative, 64-byte block L2 cache; 8MB, 8-way associative, 64-byte block L3 cache. The latencies (disregarding virtual memory TLB) expressed in clock cycles are: 4 in L1, 10 in L2, 35 in L3. Only L1 is pipelined, L2 and L3 are not and only allow for two outstanding accesses to be performed, provided only one results in a miss (no miss-on-miss). Addressing uses 32 bits.

al) Assuming initially empty and invalidated cache lines throughout the hierarchy, show the content of each D-cache after the execution of the instruction LD $R1,(0000AFA0_{hex})$ (further assumption: the instruction is already in the I-cache).

a2) show a sequence of four memory access operations that follow the LD R1,(0000AFA0_{hex}), such that

- two ops hit in L1

one op hits in L2

one op misses throughout the cache hierarchy.

a3) discuss the possibility that a miss in L2 actually hits in L3.

B) The following pseudo-code fragment performs a circular shift among the elements of an array X[.] of 1024 4-byte integer entries:

for (i=0;i<1021;i++)</pre> X[i]=X[i+3]

b1) complete the code fragment, by adding a proper preamble;

b2) assuming empty and invalidated D-caches, compute the number of misses in all levels of the cache hierarchy, with a write-through policy.

b3) compute the number of blocks used in each level of the D-caches.

C) The processor has a superscalar, 2-way pipeline, that fetches, decodes issues and retires (commits) bundles containing each 2 instructions. The front-end in-order section (fetch and decode) consists of 2 stages. The issue logic takes 1 clock cycle, if the instructions in the bundle are independent, otherwise it takes 2 clock cycles. The architecture supports dynamic speculative execution, and control dependencies from branches are solved when the branch evaluates the condition, even if it is not at commit. The execution model obeys the attached state transition diagram.

There are 2 functional units (FUs) Int1-INT2 for integer arithmetics (arithmetic and local instructions, branches and jumps, no multiplication), 2 FUS FAdd1-Fadd2 for floating point addition/subtraction, a FU FMolt1 for floating point multiplication, and a FU for division, FDiv1. There are 12 integer (R0-R11) and 12 floating point (F0-F11) registers.

Speculation is handled through a 8-entry ROB, a pool of 4 Reservation Stations (RS) Rs1-4 shared among all FUs, 2 load buffers Load1-2, 2 store buffers Store1-2 (see the attached execution model): an instruction bundle is first placed in the ROB (if two entries are available), then each instruction is dispatched to one of the shared RS (if available) and then executed in the proper FU. FUs are *pipelined* (not the Fdiv one) and have the latencies quoted in the following table:

Int - 2	Fadd — 3
Fmolt — 5	Fdiv — 6

Further assumption • Data caches are described in point A) and are assumed empty and invalidated.

c1) show state transitions for the instructions of the first two iterations of the following code fragment (assume a conventional L3 miss time of 120 clock cycles), highlighting conflicts, if any:

PC01	MOVI	R1, 0x0000AFA0	 set base address of X[0]
PC02	MOVI	R5, 1021	 set loop terminating condition
PC03	MOVI	R2, 0	initialize loop controlling variable
PC04	LD	R4,12(R1)	load X[i+3]
PC05	ST	R4,0(R1)	store into X[i]
PC06	ADD	R2,R2,1	increase loop controlling variable
PC07	ADD	R1,R1,4	advance pointer into array by 1
PC08	BNEQ	R5,R2,PC04	testing for loop exit condition

c2) show ROB, RS and buffer status at the issue of the BNEQ of the first iteration; c3) estimate the CPI of the algorithm; c4) discuss the possible unrolling of the loop.

Dynamic speculative execution Decoupled ROB RS execution model

ISTRUCTION			INSTRUCTION STATE						
	n. ite	ROB pos	WO	RE	DI	EX	WB	RR	СО
PC01 MOVI R1, 0x0000AFA0									
PC02 MOVI R5, 1021									
PC03 MOVI R2, 0									
PC04 LD R4,12(R1)									
PC05 ST R4,0(R1)									
PC06 ADD R2,R2,1									
PC07 ADD R1,R1,4									
PC08 BNEQ R5,R2,PC04									

		Reservation station and load/store buffers								
	Busy	Ор	Vj	Vk	rob _j	ROB _k	ROB pos	Address		
Rs1										
Rs2										
Rs3										
Rs4										
Load1										
Load2										
Store1										
Store2										

${\rm ROB}_{\rm j}~{\rm ROB}_{\rm k}\texttt{:}$ sources not yet available ROB pos: ROB entry number where instruction is located

		Result Register status											
Integer	R0	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	R11	
ROB pos													
state													
Float.	FO	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	F11	
ROB pos													
state													

Reorder Buffer (ROB)							
ROB Entry#	Busy	Op	Status	Destination	Value		
1							
2							
3							
4							
5							
6							
7							
8							



Decoupled execution model for bundled (paired) instructions

The state diagram depicts the model for a dynamically scheduled, speculative execution microarchitecture equipped with a Reorder Buffer (ROB) and a set of Reservation Stations (RS). The RSs are allocated during the ISSUE phase, denoted as RAT (Register Alias Allocation Table) in INTEL microarchitectures, as follows: a bundle (2 instructions) if fetched from the QUEUE of decoded instructions and ISSUED if there is a free couple of consecutive entries in the ROB (head and tail of the ROB queue do not match); each *single* instruction is moved into a RS (if available) when all of its operands are available. Access memory instructions are allocated in the ROB and then moved to a load/store buffer (if available) when operands (address and data, if proper) are available .

States are labelled as follows:

- RE: Ready for Execution (all operands are available)
- DI: Dispatched (posted to a free RS or load/store buffer)

EX: Execution (moved to a load/store buffer or to a matching and free UF)

- WB: Write Back (result is ready and is returned to the Rob by using in exclusive mode the Common Data Bus CDB)
- RR: Ready to Retire (result available or STORE has completed)
- CO: Commit (result is copied to the final ISA register)

State transitions happen at the following events:

from QUEUE to WO:	ROB entry available, operand missing
from QUEUE to RE:	ROB entry available, all operands available
loop at WO:	waiting for operand(s)
from WO to RE:	all operands available
loop at RE:	waiting for a free RS or load/store buffer
from RE to DI:	RS or load/store buffer available
loop on DI:	waiting for a free UF
from DI to EX:	UF available
loop at EX:	multi-cycle execution in a UF, or waiting for CDB
from EX to WB:	result written to the ROB with exclusive use of CDB
from EX to RR:	STORE completed, branch evaluted
loop at RR:	instruction completed, not at the head of the ROB, or bundled with a not RR instruction
from RR to CO:	bundle of RR instructions at the head of the ROB, no exception raised

Resources

Register-to-Register instructions hold resources as follows:

ROB: from state WO (or RE) up to CO, inclusive;

RS: state DI

UF: EX and WB

Load/Store instructions hold resources as follows:

ROB: from state WO (or RE) up to CO, inclusive;

Load buffer: from state WO (or RE) up to WB

Store buffer: from state WO (or RE) up to EX (do not use WB)

Forwarding: a write on the CDB (WB) makes the operand available to the consumer in the same clock cycle. If the consumer is doing a state transition from QUEUE to WO or RE, that operand is made available; if the consumer is in WO, it goes to RE in the same clock cycle of WB for the producer.

Branches: they compute Next-PC and the branch condition in EX and optionally forward Next-PC to the "in-order" section of the pipeline (Fetch states) in the next clock cycle. They do not enter WB and go to RR instead.