

CS152 Computer Architecture and Engineering
CS252 Graduate Computer Architecture

VLIW, Vector, and
Multithreaded Machines

Assigned 3/14/2018

Problem Set #4

Due 3/23/2018

<http://inst.eecs.berkeley.edu/~cs152/sp18>

The problem sets are intended to help you learn the material, and we encourage you to collaborate with other students and to ask questions in discussion sections and office hours to understand the problems. However, each student must turn in his own solution to the problems.

The problem sets also provide essential background material for the exams. The problem sets will be graded primarily on an effort basis, but if you do not work through the problem sets you are unlikely to succeed at the quizzes! Homework assignments are due at the beginning of class on the due date. Late homework will not be accepted.

Problem P4.1: Trace Scheduling

Trace scheduling is a compiler technique that increases ILP by removing control dependencies, allowing operations following branches to be moved up and speculatively executed in parallel with operations before the branch. It was originally developed for statically scheduled VLIW machines, but it is a general technique that can be used in different types of machines and in this question we apply it to a single-issue RISC-V processor.

Consider the following piece of C code (% is modulus) with basic blocks labeled:


```

A   if (data % 8 == 0)
B     X = V0 / V1;
    else
C     X = V2 / V3;
D   if (data % 4 == 0)
E     Y = V0 * V1;
    else
F     Y = V2 * V3;
G


```

Assume that **data** is a uniformly distributed integer random variable that is set sometime before executing this code.

The program's control flow graph is



The decision tree is



Path probabilities for 5.A:

A control flow graph and the decision tree both show the possible flow of execution through basic blocks. However, the control flow graph captures the static structure of the program, while the decision tree captures the dynamic execution (history) of the program.

Problem P4.1.A

On the decision tree, label each path with the probability of traversing that path. For example, the leftmost block will be labeled with the total probability of executing the path ABDEG. (Hint: you might want to write out the cases). Circle the path that is most likely to be executed.

Problem P4.1.B

This is the RISC-V code:

```
A:  lw   x1, data
    andi x2, x1, 7      # x2 <- x1 % 8
    beqz x2, C
B:  div  x3, x4, x5     # X <- V0 / V1
    j    D
C:  div  x3, x6, x7     # X <- V2 / V3
D:  andi x2, x1, 3     # x2 <- x1 % 4
    bnez x2, F
E:  mul  x8, x4, x5     # Y <- V0 * V1
    j    G
F:  mul  x8, x6, x7     # Y <- V2 * V3
G:
```

This code is to be executed on a single-issue processor **with perfect branch prediction**. Assume that the memory, divider, and the multiplier are all separate, long latency, **unpipelined** units that can be run in parallel.

Assume that the load takes x cycles, the divider takes y cycles, and the multiplier takes z cycles. Approximately how many cycles does this code take *in the best case*, *in the worst case*, and *on average*? (ignore the latency of ALU)

Problem P4.1.C

With trace scheduling, we can obtain the following code:

```
ACF: ld    x1, data
      div  x3, x6, x7    # X <- V2 / V3
      mul  x8, x6, x7    # Y <- V2 * V3
D:    andi x2, x1, 3     # x2 <- x1 % 4
      bnez x2, G
A:    andi x2, x1, 7     # x2 <- x1 % 8
      bnez x2, E
B:    div  x3, x4, x5    # X <- V0 / V1
E:    mul  x8, x4, x5    # Y <- V0 * V1
G:
```

We optimize only for the most common path, but the other paths are still correct. Approximately how many cycles does the new code take *in the best case, in the worst case* and *on average*? Is it faster *in the best case, in the worst case* and *on average* than the code in Problem P4.1.B?

Problem P4.2: VLIW machines

In this problem, we consider the execution of a code segment on a VLIW processor. The code we consider is the SAXPY kernel, which scales a vector X by a constant A, adding this quantity to a vector Y.

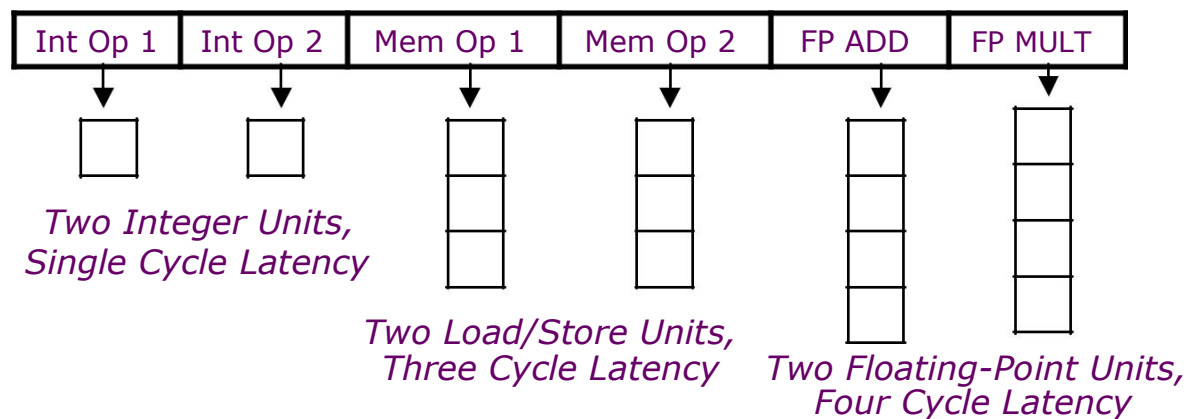
```
for(i = 0; i < N; i++) {
    Y[i] = Y[i] + A*X[i];
}
```

```
loop:  1.ld    f1, 0(x1)      # f1 = X[i]
        2.fmul  f2, f0, f1   # A * X[i]
        3.ld    f3, 0(x2)   # f3 = Y[i]
        4.fadd  f4, f2, f3   # f4 = Y[i] + A*X[i]
        5.sd    f4, 0(x2)   # Y[i] = f4
        6.addi  x1, x1, 4    # bump pointer
        7.addi  x2, x2, 4    # bump pointer
        8.bne  x1, x3, loop  # loop
```

Now we have a VLIW machine with six execution units:

- two ALU units, latency one cycle, also used for branch operations
- two memory units, latency three cycles, fully pipelined, each unit can perform either a store or a load
- two FPU units, latency four cycles, fully pipelined, one unit can perform **fadd** operations, the other **fmul** operations.

Below is a diagram of our VLIW machine:



Our machine has no interlocks. The result of an operation is written to the register file immediately after it has gone through the corresponding execution unit: one cycle after issue for ALU operations, three cycles for memory operations and four cycles for FPU operations. The old values can be read from the registers until they have been overwritten.

Problem P4.2.A: No Code Optimization

Schedule instructions for the VLIW machine in Table P4.2-1 without loop unrolling and software pipelining. What is the throughput of the loop in the code in floating point operations per cycle (FLOPS/cycle)?

Problem P4.2.B: Loop Unrolling

Schedule instructions for the VLIW machine in Table P4.2-2 only with loop unrolling. Write the assembly code by unrolling the loop once. What is the throughput of the loop in the code in floating point operations per cycle (FLOPS/cycle)? What is the speedup over Problem P4.2.A?

Problem P4.2.C: Software Pipelining

Schedule instructions for the VLIW machine in Table P4.2-3 only with software pipelining. Include the prologue and the epilogue in Table P4.2-3. What is the throughput of the loop in the code in floating point operations per cycle (FLOPS/cycle)? What is the speedup over Problem P4.2.A?

Problem P4.2.D: Loop Unrolling + Software Pipelining

Schedule instructions for the VLIW machine in Table P4.2-4 with both loop unrolling and software pipelining. Unroll the loop once as in Problem P4.2.B. Include the prologue and the epilogue in Table P4.2-3. What is the throughput of the loop in the code in floating point operations per cycle (FLOPS/cycle)? What is the speedup over Problem P4.2.A?

Problem P4.3: Vector Machines

In this problem, we analyze the performance of vector machines. We start with a baseline vector processor with the following features:

- **32 elements per vector register**
- **8 lanes**
- **One ALU per lane: 1 cycle latency**
- **One load/store unit per lane: 4 cycle latency, fully pipelined**
- **No dead time**
- **No support for chaining**
- **Scalar instructions execute on a separate 5-stage pipeline**

To simplify the analysis, we assume a magic memory system with no bank conflicts and no cache misses.

We consider execution of the following loop:

```
C code
for (i = 0; i < N; i++) {
    C[i] = A[i] + B[i] - 1;
}
```

```
loop:  1.LV    v1, (x1)      # load A
        2.LV    v2, (x2)      # load B
        3.ADDV  v3, v1, v2    # A + B
        4.SUBVS v4, v3, x4    # subtract x4 = 1
        5.SV    v4, (x3)      # store C
        6.ADDI  x1, x1, 128    # bump pointer
        7.ADDI  x2, x2, 128    # bump pointer
        8.ADDI  x3, x3, 128    # bump pointer
        9.SUBI  x5, x5, 32     # i++ (x5 = N)
       10.BNQZ  x5, loop      # loop
```

Problem P4.3.A: Simple Vector Processor

Complete the pipeline diagram in Table P4.4-1 of the baseline vector processor running the given code. The following **supplementary information** explains the diagram:

Scalar instructions execute in 5 cycles: fetch (**F**), decode (**D**), execute (**X**), memory (**M**), and writeback (**W**). A vector instruction is also fetched (**F**) and decoded (**D**). Then, it stalls (**—**) until its required vector functional unit is available. With no chaining, a dependent vector instruction stalls until the previous instruction finishes writing back all of its elements. A vector instruction is pipelined across all the lanes in parallel. For each element, the operands are read (**R**) from the vector register file, the operation executes on the load/store unit (**M**) or the ALU (**X**), and the result is written back (**W**) to the vector register file. A stalled vector instruction does not block a scalar instruction from executing.

Inst #	cycle																																									
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40		
1	F	D	R	M1	M2	M3	M4	W																																		
1				R	M1	M2	M3	M4	W																																	
1					R	M1	M2	M3	M4	W																																
1						R	M1	M2	M3	M4	W																															
2		F	D	—	—	—	R	M1	M2	M3	M4	W																														
2								R	M1	M2	M3	M4	W																													
2									R	M1	M2	M3	M4	W																												
2										R	M1	M2	M3	M4	W																											
3			F	D	—	—	—	—	—	—	—	—	—	—	—	R	X1	W																								
3																	R	X1	W																							
3																		R	X1	W																						
3																			R	X1	W																					
4			F	D	—																																					
4																																										
4																																										
4																																										
5				F	D	—																																				
5																																										
5																																										
5																																										
6					F	D	X	M	W																																	
7						F	D	X	M	W																																
8							F	D	X	M	W																															
9								F	D	X	M	W																														
10									F	D	X	M	W																													
1										F	D	—																														
1																																										
1																																										
1																																										

Table P4.3-1: Vector Pipeline Diagram (8 Lanes without Chaining)

Problem P4.3.B: Hardware Optimization (Chaining)

In this question, we analyze the performance benefits of chaining and additional lanes. Vector chaining is done through the register file and an element can be read (**R**) on the same cycle in which it is written back (**W**), or it can be read on any later cycle (the chaining is *flexible*). For this question, we always assume 32 elements per vector register, so there are 2 elements per lane with 8 lanes, and 1 element per lane with 32 lanes.

To analyze performance, we calculate the total number of cycles per vector loop iteration by summing the number of cycles between the issuing of successive vector instructions. For example, in Question P3.4.A, Inst #1(LV) begins execution in cycle 3, Inst #2(LV) in cycle 7 and Inst #3(ADDV) in cycle 16. Therefore, there are 4 cycles between Inst #1 and Inst #2 and 9 cycles between Inst #2 and Inst #3.

First, fill in Table P4.3-2 for 8 lanes with chaining, Table P4.3-3 for 16 lanes with chaining, and Table P4.3-4 for 32 lanes with chaining. Note that, with 8 lanes and chaining, Inst #4(SUBVS) cannot issue 2 cycles after Inst #3(ADDV) because there is only one ALU per lane.

Also, complete the following table. The first row corresponds to the baseline 8-lane vector processor with no chaining. The second row adds flexible chaining to the baseline processor, and the last two rows increase the number of lanes from 8 to 32.

Vector Processor Configuration	Number of cycles between successive vector instructions					Total cycles per vector loop iteration
	#1(LV) #2(LV)	#2(LV) #3(ADDV)	#3(ADDV) #4(SUBVS)	#4(SUBVS) #5(SV)	#5(SV) #1(LV)	
8 lanes, no chaining	4	9	6	6	4	29
8 lanes, chaining						
16 lanes, chaining						
32 lanes, chaining						

Inst #	cycle																																									
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40		
1	F	D	R	M1	M2	M3	M4	W																																		
1				R	M1	M2	M3	M4	W																																	
1					R	M1	M2	M3	M4	W																																
1						R	M1	M2	M3	M4	W																															
2		F	D	—	—	—	R	M1	M2	M3	M4	W																														
2								R	M1	M2	M3	M4	W																													
2									R	M1	M2	M3	M4	W																												
2										R	M1	M2	M3	M4	W																											
3			F	D	—	—	—	—	—	—	—	R	X1	W																												
3													R	X1	W																											
3														R	X1	W																										
3															R	X1	W																									
4				F	D	—																																				
4																																										
4																																										
4																																										
5					F	D	—																																			
5																																										
5																																										
5																																										
6						F	D	X	M	W																																
7							F	D	X	M	W																															
8								F	D	X	M	W																														
9									F	D	X	M	W																													
10										F	D	X	M	W																												
1											F	D	—																													
1																																										
1																																										
1																																										

Table P4.3-2: 8 Lanes with Chaining

Inst #	cycle																																											
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40				
1	F	D	R	M1	M2	M3	M4	W																																				
1				R	M1	M2	M3	M4	W																																			
2		F	D																																									
2																																												
3			F	D																																								
3																																												
4				F	D																																							
4																																												
5					F	D																																						
5																																												
6						F	D	X	M	W																																		
7							F	D	X	M	W																																	
8								F	D	X	M	W																																
9									F	D	X	M	W																															
10										F	D	X	M	W																														
1											F	D	—																															
1																																												

Table P4.3-3: 16 Lanes with Chaining

Inst #	cycle																																											
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40				
1	F	D	R	M1	M2	M3	M4	W																																				
2		F	D																																									
3			F	D																																								
4				F	D																																							
5					F	D																																						
6						F	D	X	M	W																																		
7							F	D	X	M	W																																	
8								F	D	X	M	W																																
9									F	D	X	M	W																															
10										F	D	X	M	W																														
1											F	D																																

Table P4.3-3: 32 Lanes with Chaining

Problem P4.4: Multithreading

This problem evaluates the effectiveness of multithreading using a simple database benchmark. The benchmark searches for an entry in a linked list built from the following structure, which contains a key, a pointer to the next node in the linked list, and a pointer to the data entry.

```
struct node { int
    key;
    struct node *next;
    struct data *ptr;
}
```

The following RISC-V code shows the core of the benchmark, which traverses the linked list and finds an entry with a particular key.

```
loop:  LW    x3, 0(x1)           # load a key
      LW    x4, 4(x1)           # load the next pointer
      SEQ   x3, x3, x2          # set x3 if x3 == x2
      BNEZ  x3, end             # found the entry
      ADD   x1, x0, x4          # check the next node
      BNEZ  x4, loop
end:
```

We run this benchmark on a single-issue in-order processor. The processor can fetch and issue (dispatch) one instruction per cycle. If an instruction cannot be issued due to a data dependency, the processor stalls. Integer instructions take one cycle to execute and the result can be used in the next cycle. For example, if `SEQ` is executed in cycle 1, `BNEZ` can be executed in cycle 2. We also assume that the processor has a perfect branch predictor with no penalty for both taken and not-taken branches.

Problem P4.4.A

Assume that our system does not have a cache. Each memory operation directly accesses main memory and takes 100 CPU cycles. The load/store unit is fully pipelined, and non-blocking. After the processor issues a memory operation, it can continue executing instructions until it reaches an instruction that is dependent on an outstanding memory operation. How many cycles does it take to execute one iteration of the loop in steady state?

Instruction	Start Cycle	End Cycle
LW x3, 0(x1)		
LW x4, 4(x1)		
SEQ x3, x3, x2		
BNEZ x3, End		
ADD x1, x0, x4		
BNEZ x1, Loop		

Problem P4.4.B

Now we add zero-overhead multithreading to our pipeline. A processor executes multiple threads, each of which performs an independent search. Hardware mechanisms schedule a thread to execute each cycle.

In our first implementation, the processor switches to a different thread every cycle using fixed round robin scheduling (similar to CDC 6600 PPUs). Each of the N threads executes one instruction every N cycles. What is the **minimum** number of threads that we need to fully utilize the processor, i.e., execute one instruction per cycle?

Problem P4.4.C

How does multithreading affect throughput (number of keys the processor can find within a given time) and latency (time processor takes to find an entry with a specific key)? Assume the processor switches to a different thread every cycle and is fully utilized. Check the correct boxes.

	Throughput	Latency
Better		
Same		
Worse		

Problem P4.4.D

We change the processor to only switch to a different thread when an instruction cannot execute due to data dependency. What is the minimum number of threads to fully utilize the processor now? Note that the processor issues instructions in-order in each thread.

Problem P4.5: Multithreading

Consider a single-issue in-order multithreading processor that is similar to the one described in Problem P4.4.

Each cycle, the processor can fetch and issue one instruction that performs any of the following operations:

- **load/store, 12-cycle latency (fully pipelined)**
- **integer add, 1-cycle latency**
- **floating-point add, 5-cycle latency (fully pipelined)**
- **branch, no delay slots, 1-cycle latency**

The processor **does not have a cache**. Each memory operation directly accesses main memory. If an instruction cannot be issued due to a data dependency, the processor stalls. We also assume that the processor has a perfect branch predictor with no penalty for both taken and not-taken branches.

Your job is to analyze the processor utilizations for the following two thread-switching implementations:

Fixed Switching: the processor switches to a different thread every cycle using fixed round robin scheduling. Each of the N threads executes an instruction every N cycles.

Data-dependent Switching: the processor only switches to a different thread when an instruction cannot execute due to a data dependency.

Each thread executes the following RISC-V code:

```
loop:  LD      f2, 0(x1)      # load data into f2
       ADDI   x1, x1, 4      # bump pointer
       FADD  f3, f3, f2      # f3 = f3 + f2
       BNE   f2, f4, loop    # continue if f2 != f4
```

Problem P4.5.A

What is the minimum number of threads that we need to fully utilize the processor for each implementation?

Fixed Switching: _____ **Thread(s)**

Data-dependent Switching: _____ **Thread(s)**

Problem P4.5.B

What is the minimum number of threads that we need to fully utilize the processor for each implementation if we change the **load/store latency to 1-cycle (but keep the 5-cycle floating-point add)**?

Fixed Switching: _____ **Thread(s)**

Data-dependent Switching: _____ **Thread(s)**

(CS252 only) Problem P4.5.C

Consider a **Simultaneous Multithreading (SMT)** machine with limited hardware resources. **Circle** the following hardware constraints that can limit the total number of threads that the machine can support. For the item(s) that you circle, **briefly describe** the minimum requirement to support **N** threads.

(A) Number of Functional Unit:

(B) Number of Physical Registers:

(C) Data Cache Size:

(D) Data Cache Associativity: