Advanced Computer Architecture
Lab 1: MIPS-R10000

LAP

Outline

- MIPS R10000 is a superscalar processor from 1996
- Easy to understand:
  - Simple RISC ISA
  - Relatively old processor (many of the presented techniques have not changed)
- Great start before doing the homework!
  - Impossible to do the homework without understanding this lab very well
- Great opportunity to check your knowledge!
  - (Ask the TAs as many questions as you want)

Outline

- This processor view is great and high level but how things work in details?

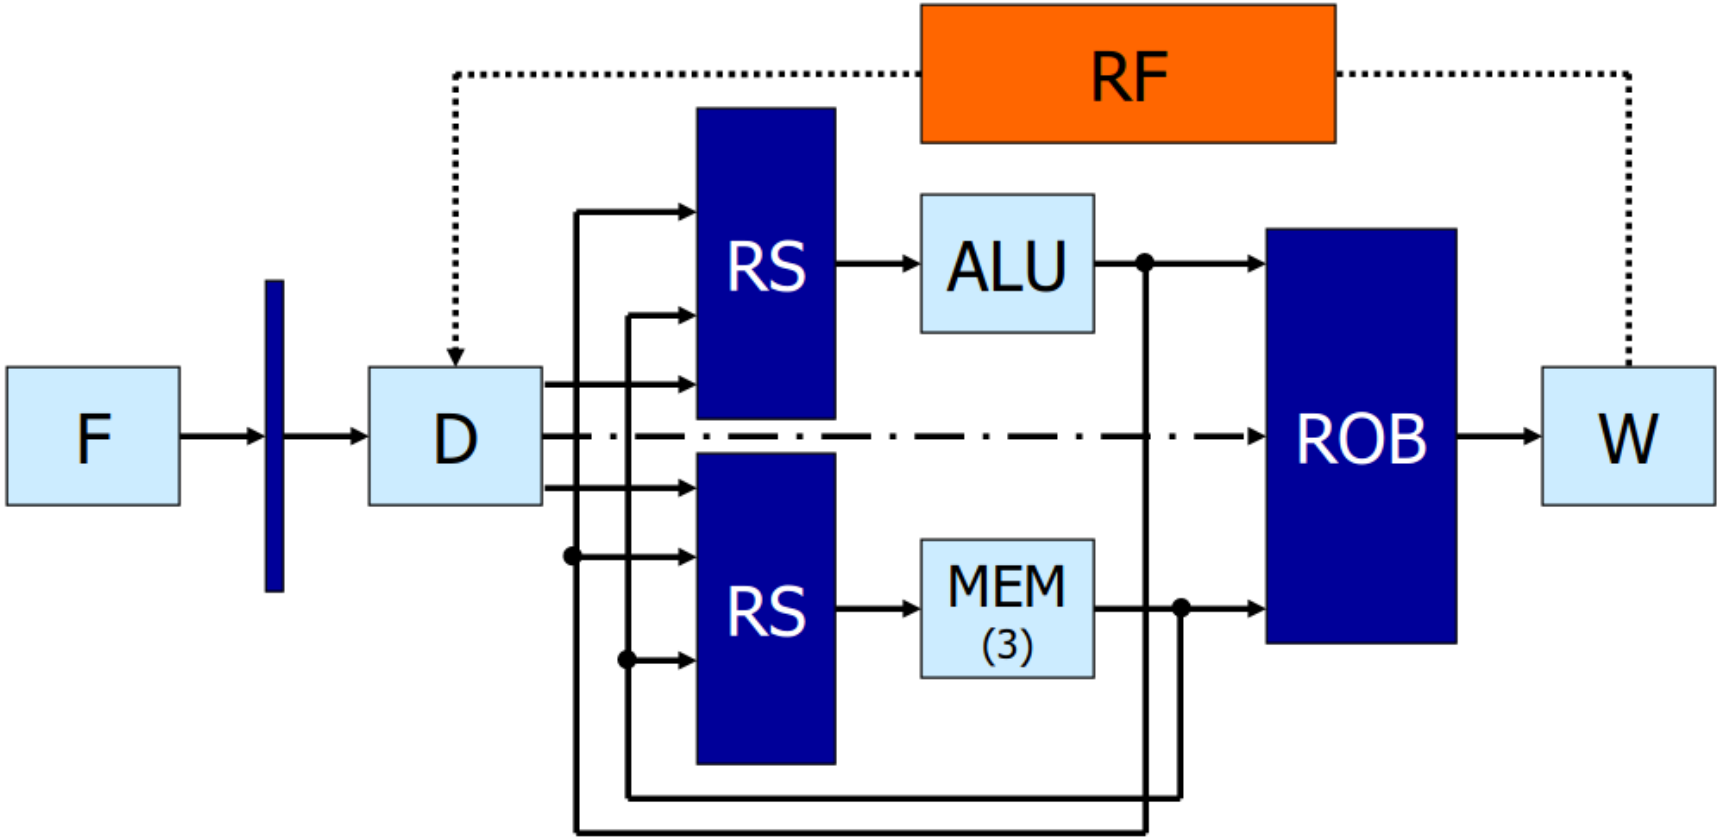

This lab

- First: Read the paper if not already done

THE MIPS R10000 SUPERSCALAR MICROPROCESSOR

Kenneth C. Yeager

Silicon Graphics, Inc.

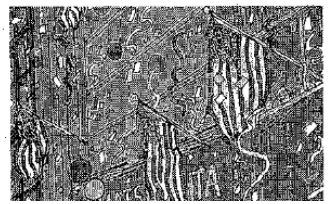

The Mips R10000 is a dynamic, superscalar microprocessor that implements the 64-bit Mips 4 instruction set architecture. It fetches and decodes four instructions per cycle and dynamically issues them to five fully-pipelined, low-latency execution units. Instructions can be fetched and executed speculatively beyond branches. Instructions graduate in order upon completion. Although execution is out of order, the processor still provides sequential memory

parallel while the processor executes other instructions. This type of cache design is called "nonblocking," because cache refills do not block subsequent accesses to other cache lines.

Processors rely on compiler support to optimize instruction sequencing. This technique is especially effective for data arrays, such as those used in many floating-point applications. For these arrays, a sophisticated compiler can optimize performance for a spe-

This lab

- Answer the questions
- Ask the TAs whenever you have any doubt on the paper
- We are here to help!

1 Pipeline and Register Renaming

1. Explain why in the 3rd stage of the pipeline (see Figure 2) the *register file* is read in the 2nd half of each cycle. What are the advantages and disadvantages of such an implementation?

2. Figure 5 describes register renaming. Why are the destination and source registers represented with 5 bits in the original instruction, and with 6 bits after the renaming?

3. Specify the function of the *queues*, *active lists*, and *map tables* in the R10000. Relate each of these to the structures described in the course (*reservation stations*, *ROB*, ...).

4. Describe the role of the components *Rdy*, *OpA*, *OpB*, *OpC*, *Dest*, *Old Dest*

This lab

- Fill the simulation spreadsheet
- Good starting point before implementing the simulator (HW1)
- You cannot implement the simulator without understanding this lab very well
- Use the TAs
- Simulation details on the pdf on Moodle

Program

```
LDC1 $F0, #0000($I1)
LDC1 $F1, #0000($I1)
LDC1 $F2, #0000($I1)
LDC1 $F3, #0000($I1)
MUL.S $F4, $F0, $F3
MUL.S $F5, $F1, $F2
SUB.S $F4, $F4, $F5
ADD.S $F5,$F0, $F1
MUL.S $F9, $F5, $F5
MUL.S $F10, $F4, $F0
SUB.S $F9, $F9, $F10
SQRT.S $F9, $F9
SDC1 F4, #0004($I1)
SDC1 F5, #0005($I1)
SDC1 F9, #0006($I1)
SDC1 $F10, #0007($I1)
```

Simulation

End of Cycle #0

Active List 1			
Adress	Old Dest	Log Dest	Done
x00			
x01			
x02			
x03			
x04			
x05			
x06			
x07			

Active List 2			
Adress	Old Dest	Log Dest	Done
x08			
x09			
x0A			
x0B			
x0C			
x0D			
x0E			
x0F			

Active List 3			
Adress	Old Dest	Log Dest	Done
x10			
x11			
x12			
x13			
x14			
x15			
x16			
x17			

Active List 4			
Adress	Old Dest	Log Dest	Done
x18			
x19			
x1A			
x1B			
x1C			
x1D			
x1E			
x1F			

FP Register Map Table	
Logical	Map
x00	x00
x01	x01
x02	x02
x03	x03
x04	x04
x05	x05
x06	x06
x07	x07
x08	x08
x09	x09
x0A	x0A
x0B	x0B
x0C	x0C
x0D	x0D
x0E	x0E
x0F	x0F
x10	x10
x11	x11
x12	x12
x13	x13
x14	x14
x15	x15
x16	x16
x17	x17
x18	x18
x19	x19
x1A	x1A
x1B	x1B
x1C	x1C
x1D	x1D
x1E	x1E
x1F	x1F

Busy Bit Table			
x00	0	x20	0
x01	0	x21	0
x02	0	x22	0
x03	0	x23	0
x04	0	x24	0
x05	0	x25	0
x06	0	x26	0
x07	0	x27	0
x08	0	x28	0
x09	0	x29	0
x0A	0	x2A	0
x0B	0	x2B	0
x0C	0	x2C	0
x0D	0	x2D	0
x0E	0	x2E	0
x0F	0	x2F	0
x10	0	x30	0
x11	0	x31	0
x12	0	x32	0
x13	0	x33	0
x14	0	x34	0
x15	0	x35	0
x16	0	x36	0
x17	0	x37	0
x18	0	x38	0
x19	0	x39	0
x1A	0	x3A	0
x1B	0	x3B	0
x1C	0	x3C	0
x1D	0	x3D	0
x1E	0	x3E	0
x1F	0	x3F	0

Free List								
FIFO1	x20	x24	x28	x2C	x30	x34	x38	x3C
FIFO2	x21	x25	x29	x2D	x31	x35	x39	x3D
FIFO3		x26	x2A	x2E	x32	x36	x3A	x3E
FIFO4		x27	x2B	x2F	x33	x37	x3B	x3F

FIFO Head

FIFO Tail

Address Queue			
Tag	Rdy	QpA	Dest

Pipeline	Stage 1	Stage 2	Stage 3	Stage 4	Stage 5	Stage 6	Stage 7
Instructions Positions	1 2 3 4						

FP Queue					
Tag	Rdy	QpA	QpB	QpC	Dest