ECE 4750 Computer Architecture, Fall 2015

T10 Advanced Processors: Out-of-Order Execution

School of Electrical and Computer Engineering
Cornell University

revision: 2015-11-04-13-46

1. Incremental Approach to Exploring OOO Execution

- Gradually work through five different microarchitectures
- For each microarchitecture
  - overall pipeline structure
  - required hardware data-structures
  - example instruction sequence executing on microarchitecture
  - handling precise exceptions
- Several simplifications
  - all designs are single issue
  - assume code sequence never includes WAW or WAR dependencies
  - only support `addu`, `addiu`, `mul`

	Front-End or **Fetch/Decode**	**Issue**	**Writeback** or **Completion**	**Commit**	**Data Structures**
I3L	io	io	io	late	
I2OE	io	io	ooo	early	SB
I2OL	io	io	ooo	late	SB, ROB
IO2E	io	ooo	ooo	early	SB, IQ
IO2L	io	ooo	ooo	late	SB, IQ, ROB

```
a: mul    r1,   r2,   r3
b: addiu  r11,  r10,  1
c: mul    r5,   r1,   r4
d: mul    r7,   r5,   r6
e: addiu  r12,  r11,  1
f: addiu  r13,  r12,  1
g: addiu  r14,  r12,  2
```

2. IO Front-End/Issue/Completion, Late Commit

	Front-End or **Fetch/Decode**	**Issue**	**Writeback** or **Completion**	**Commit**	**Data Structures**
I3L	io	io	io	late	
I2OE	io	io	ooo	early	SB
I2OL	io	io	ooo	late	SB, ROB
IO2E	io	ooo	ooo	early	SB, IQ
IO2L	io	ooo	ooo	late	SB, IQ, ROB

The following is the basic in-order single-issue pipeline.

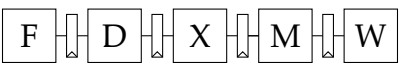

Split X/M stages into two functional units. Still single issue, so not strictly necessary but a nice incremental design step.

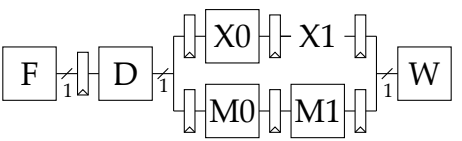

What if we want to incorporate a four-cycle pipelined integer multiplier?
Key Idea: Extend all pipelines to equal length.

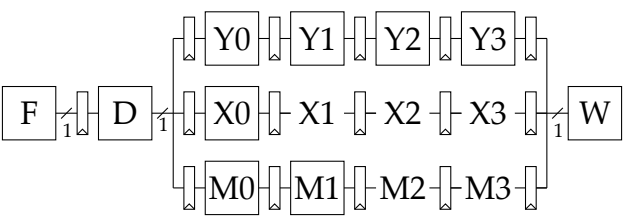

Cannonical I3L Pipeline

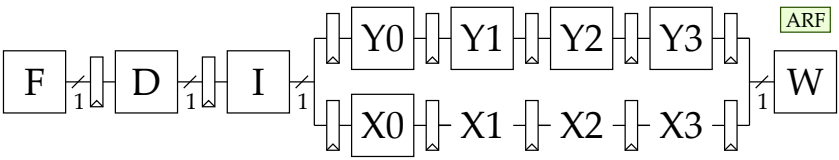

- To avoid increasing CPI, need full bypassing which can be expensive
- Add new issue stage which
  - reads architectural register file
  - performs hazard checking and includes bypass muxing
  - "issues" instruction to appropriate functional unit
- Include just X-pipe and Y-pipe since we are only focusing on `addu`, `addiu`, and `mul` instructions

Example Execution Diagrams

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				

3. IO Front-End/Issue, OOO Completion, Early Commit

	Front-End or Fetch/Decode	Issue	Writeback or Completion	Commit	Data Structures
I3L	io	io	io	late	
I2OE	io	io	ooo	early	SB
I2OL	io	io	ooo	late	SB, ROB
IO2E	io	ooo	ooo	early	SB, IQ
IO2L	io	ooo	ooo	late	SB, IQ, ROB

Cannonical I2OE Pipeline

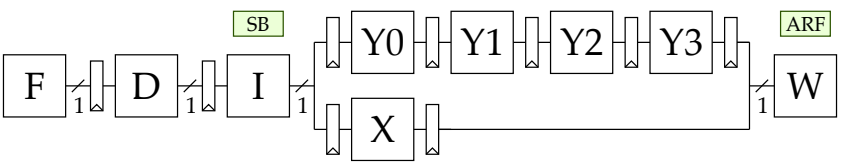

- Remove "dummy" pipeline stages
- Fewer bypass paths, significantly reduces hardware complexity
  - I3L has six bypass paths
  - I2OE has three bypass paths
  - Bypass from end of Y3, end of X, and W to end of I
- Scoreboard is used to centralize structural/data hazard detection
- WAW hazards are possible, which we ignore in this topic
- WAR hazards are not possible
- NOTE: Fewer stages does not necessarily mean better performance!

Data Structure: Scoreboard

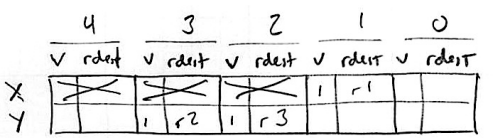

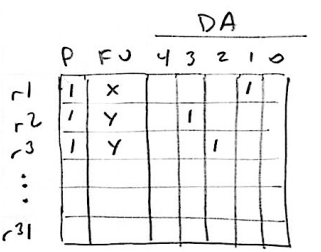

- Indexed by functional unit
  - **V**: valid bit
  - **rdest**: destination reg specifier
  - Entries shift to right every cycle
- Structural hazards: `addu` and `addiu` check col 2 valid bit to ensure no structural hazard on WB port
- RAW hazards: I stage compares current instruction source reg specifiers with every valid entry in SB
  - match in col 2–4 = stall I
  - match in col 0–1 = bypass into I
  - no match = read ARF
- Large number of comparisons make accessing SB expensive

- Indexed by reg specifier
  - **P**: pending bit
  - **FU**: functional unit
  - **WA**: when available?
  - WA bits shift to right every cycle
- Structural hazards: `addu` and `addiu` check no bits are set in col 2 to ensure no structural hazard on WB port
- I stage compares checks pending bit for each source register specifier
  - pending bit set = check WA to see if stall or bypass (FU says where to bypass from)
  - pending bit clear = read ARF
- Can use SB to stall to prevent WAW hazards

Example Execution Diagrams

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				

			WA Entry						
cycle	**D**	**I**	**r1**	**r5**	**r7**	**r11**	**r12**	**r13**	**r14**
0									
1									
2									
3									
4									
5									
6									
7									
8									
9									
10									
11									
12									
13									
14									
15									

Handling Precise Exceptions

Early commit requires the commit point to be in the decode stage. What if instruction d causes an exception?

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				

Not usually possible to detect all exceptions in the front-end, which motivates our interest in supporting late commit at the end of the pipeline.

4. IO Front-End/Issue, OOO Completion, Late Commit

	Front-End or Fetch/Decode	Issue	Writeback or Completion	Commit	Data Structures
I3L	io	io	io	late	
I2OE	io	io	ooo	early	SB
I2OL	io	io	ooo	late	SB, ROB
IO2E	io	ooo	ooo	early	SB, IQ
IO2L	io	ooo	ooo	late	SB, IQ, ROB

Cannonical I2OL Pipeline

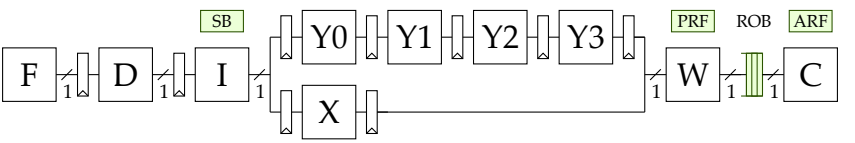

ARF					write
PRF		read		write	read
SB		read/write			
ROB	alloc			write	read/dealloc

- Add extra C stage for commit at end of pipeline
- Still use scoreboard to centeralize structural/data hazard detection
- Add physical regfile (PRF) and reorder buffer (ROB) between W/C
- PRF keeps uncommited results (a.k.a. future regfile, working regfile)
- Reorder buffer (ROB)
  - allocated in-order in D stage
  - updated out-of-order in W stage
  - deallocated in-order in C stage
- WAW hazards are possible, which we ignore in this topic
- WAR hazards are not possible

Data Structure: Reorder Buffer

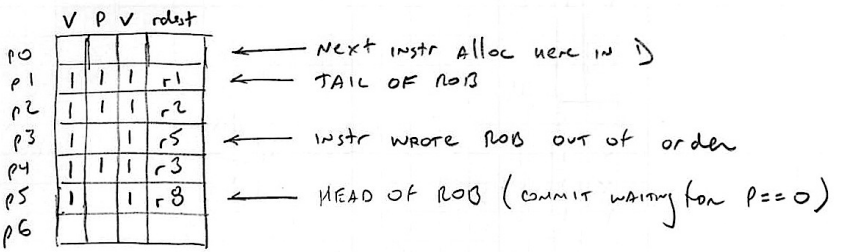

- ROB fields
  - **V**: valid bit (is this entry valid?)
  - **P**: pending bit (instruction in flight targeting this entry)
  - **V**: valid bit (is the dest reg specifier valid?)
  - **rdest**: destination reg specifier
- ROB managed like a queue, implemented with circular buffer
  - new instructions allocated ROB entries at tail
  - instructions update pending bit out-of-order
  - commit stage waits for pending bit of head to be clear

Example Execution Diagrams

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				

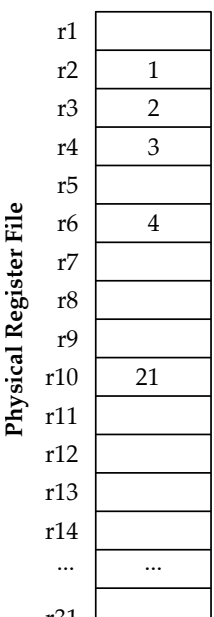

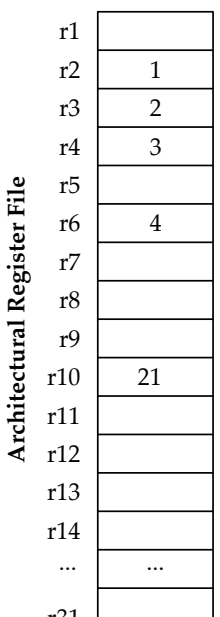

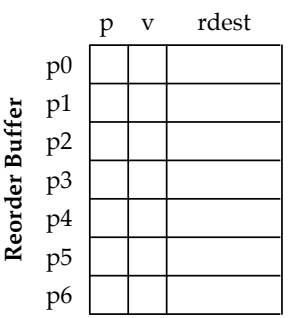

We can use a table to compactly illustrate how the ROB works.

			ROB Entry			
cycle	**D**	**I**	**0**	**1**	**2**	**3**
0						
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						
11						
12						
13						
14						
15						
16						
17						
18						
19						

Handling Precise Exceptions

Late commit means exceptions are handled in the C stage at the end of the pipeline. What if instruction a causes an exception?

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				

Need to copy values from ARF to PRF on an exception before redirecting the front of the pipeline to the exception handler. This copy may take multiple cycles. Also possible to include additional bits in I stage to indicate wether the most recent version of every given architectural register is in the ARF or PRF.

5. IO Front-End, OOO Issue/Completion, Early Commit

	Front-End or Fetch/Decode	Issue	Writeback or Completion	Commit	Data Structures
I3L	io	io	io	late	
I2OE	io	io	ooo	early	SB
I2OL	io	io	ooo	late	SB, ROB
IO2E	io	ooo	ooo	early	SB, IQ
IO2L	io	ooo	ooo	late	SB, IQ, ROB

Cannonical IO2E Pipeline

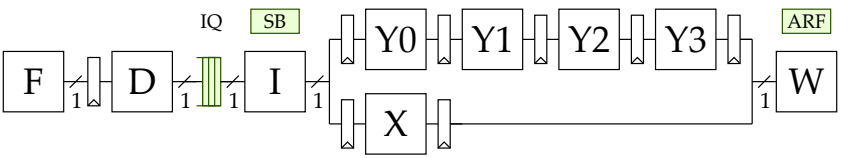

ARF		read	write
SB		read/write	
IQ	alloc	read/dealloc	write

- Still use scoreboard to centeralize structural/data hazard detection
- Add issue queue (IQ) between D and I stages
  - allocated in-order in D stage
  - updated out-of-order in W stage
  - deallocated out-of-order in I stage
- Do not necessarily want to wait for W stage to update IQ; we will need to assume *aggressive bypassing* which requires combinational communication between last stage of functional unit and I stage
- WAW hazards are possible, which we ignore in this topic
- WAR hazards are possible, which we ignore in this topic

Data Structure: Issue Queue

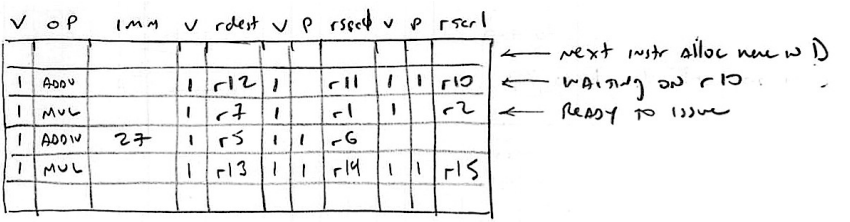

- IQ fields
  - **V**: valid bit (is this entry valid?)
  - **op**: instruction opcode
  - **imm** immediate value
  - **V**: valid bit (is the dest/src reg specifier valid?)
  - **P**: pending bit (is the src data ready?)
  - **rdest/rsrc**: destination/source reg specifiers
- IQ managed like a queue, implemented with circular buffer
  - new instructions allocated IQ entries at tail
  - instructions leave IQ out-of-order when ready
- Wakeup Logic: An instruction needs to update pending bits of dependent instructions when that instruction is in W stage (actually need to do this earlier to enable aggressive bypassing)
- Select Logic: Determine which instructions are ready to be issued, and then select which one to actually issue. Usually issue oldest ready instruction.

```
inst_ready = ( !val_src0 || !p_src0 )
          && ( !val_src1 || !p_src1 )
          && no structural hazards
```

Example Execution Diagrams

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				

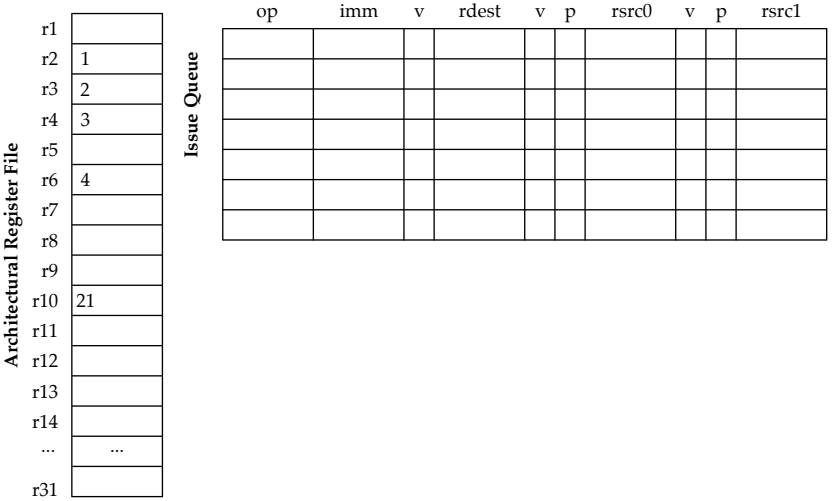

We can use a table to compactly illustrate how the IQ works.

			IQ Entry		
cycle	**D**	**I**	**0**	**1**	**2**
0					
1					
2					
3					
4					
5					
6					
7					
8					
9					
10					
11					
12					
13					
14					
15					
16					
17					
18					
19					

Handling Precise Exceptions

Early commit requires the commit point to be in the decode stage. What if instruction e causes an exception?

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				

Performance Benefit of OOO Execution

Does IO2E improve performance compared to I2OE? Let's assume all instructions are in issue queue.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				

Centeralized vs. Distributed IQs

IQs can either be centeralized or distributed across functional units. Distributed IQs are sometimes called reservation stations. This can naturally enable superscalar execution.

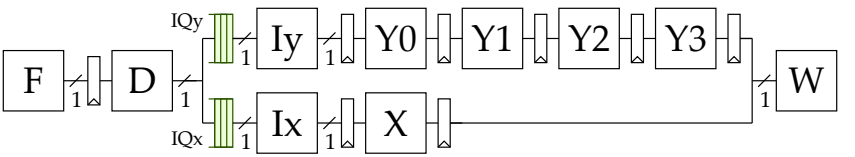

6. IO Front-End, OOO Issue/Completion, Late Commit

	Front-End or **Fetch/Decode**	**Issue**	**Writeback** or **Completion**	**Commit**	**Data Structures**
I3L	io	io	io	late	
I2OE	io	io	ooo	early	SB
I2OL	io	io	ooo	late	SB, ROB
IO2E	io	ooo	ooo	early	SB, IQ
IO2L	io	ooo	ooo	late	SB, IQ, ROB

Cannonical IO2L Pipeline

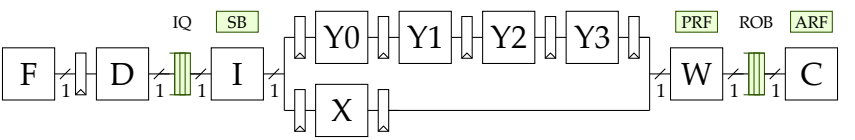

ARF					write
PRF		read		write	read
SB		read/write			
IQ	alloc	read/dealloc			
ROB	alloc			write	read/dealloc

- Use scoreboard to centeralize structural/data hazard detection
- Use IQ to enable out-of-order issue, ROB to enable late commit
- Overall organization:
  - In-order fetc/decode (front-end of pipeline)
  - Out-of-order issue/completion (middle of pipeline)
  - In-order commit (back-end of pipeline)
- WAW hazards are possible, which we ignore in this topic
- WAR hazards are possible, which we ignore in this topic

Example Execution Diagrams

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				

Handling Precise Exceptions

Late commit means exceptions are handled in the C stage at the end of the pipeline. What if instruction a causes an exception?

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				

Out-of-Order Dual-Issue Processor

Assume we can fetch, decode, issue, writeback, and commit two instructions per cycle.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
a: `mul r1, r2, r3`																				
b: `addiu r11, r10, 1`																				
c: `mul r5, r1, r4`																				
d: `mul r7, r5, r6`																				
e: `addiu r12, r11, 1`																				
f: `addiu r13, r12, 1`																				
g: `addiu r14, r12, 2`																				