

HY425 Lecture 07: Precise Interrupts, Implementing Speculation

Dimitrios S. Nikolopoulos

University of Crete and FORTH-ICS

October 31, 2011

Speculation basics

- ▶ Multiple-issue processors need more independent instructions
- ▶ Wide-issue processors may need to predict one or more branches per cycle
 - ▶ Hard to maintain high clock rate
- ▶ Speculation attempts to overcome problem
 - ▶ Execute program as if all branches were predicted correctly
 - ▶ Dynamic scheduling waits for branches instead
 - ▶ **Need mechanism to undo instructions from incorrect path**
- ▶ Three elements in speculation
 - ▶ Dynamic branch prediction
 - ▶ Ability to undo speculative instructions from wrong path
 - ▶ Dynamic scheduling

Superscalar execution

Characteristics

- ▶ Multiple instruction issue
- ▶ Dynamic (scoreboard, Tomasulo) or static (compiler) instruction scheduling
- ▶ Logic issues independent instructions, modulo constraints
 - ▶ Instructions dependent on pending instructions not issued
 - ▶ Instructions not issued due to hazards
 - ▶ Other constraints (e.g. on types of instructions issued per cycle)
- ▶ **Number of instructions issued per cycle variable**
- ▶ **CPI, IPC, limited by dependencies, hazards**

Maintaining precise exceptions

Reorder buffer (ROB)

- ▶ Instructions generating exceptions do not commit until they reach the head of the ROB
- ▶ In scoreboard, Tomasulo, instructions commit out of order:
 - ▶ **Out-of-order instructions modify state (register memory)**
 - ▶ **Exceptions from in-order instructions are imprecise**
- ▶ Recoverable memory exceptions
 - ▶ Program must resume execution
 - ▶ Example: page faults
 - ▶ No speculative state should exist when program resumes

Precise exceptions with speculation

Example

- ▶ Assume all instructions in the loop have issued twice
- ▶ Assume all instructions have completed execution
- ▶ LD, MULDD of first iteration have committed

```

Loop: LD  F0, 0(R1)
      MULDD F4, F0, F2
      SD  F4, 0(R1)
      SUBI R1, R1, 8
      BNE R1, R2, Loop
    
```

Precise exceptions with speculation (cont.)

ROB snapshot after two iterations

Reorder buffer					
Entry	Busy	Instruction	State	Destination	Value
1	no	LD F0, 0(R1)	Commit	F0	M[0+R[R1]]
2	no	MULD F4, F0, F2	Commit	F4	#1 × R[F2]
3	yes	SD F4, 0(R1)	Write result	0 + R[R1]	#2
4	yes	SUBI R1, R1, 8	Write result	R1	R[R1] - 8
5	yes	BNE R1, R2, Loop	Write result		
6	yes	LD F0, 0(R1)	Write result	F0	M[#4]
7	yes	MULD F4, F0, F2	Write result	F4	#6 × R[F2]
8	yes	SD F4, 0(R1)	Write result	0 + #4	#7
9	yes	SUBI R1, R1, 8	Write result	R1	#4 - 8
10	yes	BNE R1, R2, Loop	Write result		

Precise exceptions with speculation (cont.)

Assume first BNE not taken

Reorder buffer					
Entry	Busy	Instruction	State	Destination	Value
1	no	LD F0, 0(R1)	Commit	F0	M[0+R[R1]]
2	no	MULD F4, F0, F2	Commit	F4	#1 × R[F2]
3	yes	SD F4, 0(R1)	Commit	0 + R[R1]	#2
4	yes	SUBI R1, R1, 8	Commit	R1	R[R1] - 8
5	yes	BNE R1, R2, Loop	Write result		
6					
7					
8					
9					
10					

- ▶ **ROB after BNE cleared**
- ▶ **Fetch instructions from fall-through path**

Precise exceptions with speculation (cont.)

Using the ROB for exceptions

- ▶ Exception recorded in ROB
- ▶ Exception not thrown until instruction commits
- ▶ If exception happens in speculative instruction and processor mis-speculates:
 - ▶ Squash this and other speculative instructions
 - ▶ Optimization:
 - ▶ Handle exceptions as soon as they arise and earlier branches are resolved

Precise exceptions with speculation (cont.)

Using the ROB for exceptions

- ▶ Load and stores handled also in ROB
- ▶ Register/memory not updated until load/store reaches head of ROB
- ▶ WAW, WAR hazards in memory removed due to in-order commit
- ▶ RAW hazards handled by:
 - ▶ Preventing load from executing if prior store in ROB has same effective address
 - ▶ Preserve program order for computation of effective address of load with respect to earlier stores
- ▶ Forwarding possible from stores to later loads

Multiple issue with speculation

Mechanisms

- ▶ Similar to Tomasulo with multiple issue
- ▶ Must support multiple commits in one cycle

Example

```

Loop: LD   R2, 0 (R1)
      ADDI R2, R2, 1
      SD   R2, 0 (R1)
      ADDI R1, R1, 4
      BNE R2, R3, Loop
    
```

Multiple issue without speculation

Assume separate LD-ST and INT units

Iteration number	Instruction	Issues at clock cycle number	Executes at clock cycle number	Memory access at clock cycle number	Write CDB at clock cycle number	Comment
1	LD R2,0(R1)	1	2	3	4	First issue
1	ADDI R2,R2,1	1	5		6	Wait for LD
1	SD R2,0(R1)	2	3	7		Wait for ADDI
1	ADDI R1,R1,4	2	3		4	Execute directly
1	BNE R2,R3,Loop	3	7		Wait for ADDI	
2	LD R2,0(R1)	4	8	9	10	Wait for BNE
2	ADDI R2,R2,1	4	11		12	Wait for LD
2	SD R2,0(R1)	5	9	13		Wait for ADDI
2	ADDI R1,R1,4	5	8		9	Execute directly
2	BNE R2,R3,Loop	6	13		Wait for ADDI	
3	LD R2,0(R1)	7	14	15	16	Wait for BNE
3	ADDI R2,R2,1	7	17		18	Wait for LD
3	SD R2,0(R1)	8	15	19		Wait for ADDI
3	ADDI R1,R1,4	8	14		15	Execute directly
3	BNE R2,R3,Loop	9	19		Wait for ADDI	

Multiple issue with speculation

Assume separate LD-ST and INT units

Iteration number	Instruction	Issues at clock cycle number	Executes at clock cycle number	Read access at clock cycle number	Writes		Comment
					CDB at clock cycle number	Commits at clock cycle number	
1	LD R2,0(R1)	1	2	3	4	5	First issue
1	ADDI R2,R2,1	1	5		6	7	Wait for LD
1	SD R2,0(R1)	2	3		4	7	Wait for ADDI
1	ADDI R1,R1,4	2	3		4	8	Commit in order
1	BNE R2,R3,Loop	3	7		8	8	Wait for ADDI
2	LD R2,0(R1)	4	8		9	10	No execute delay
2	ADDI R2,R2,1	4	5	6	7	9	No Wait for LD
2	SD R2,0(R1)	5	6		7	10	Wait for ADDI
2	ADDI R1,R1,4	5	6		7	11	Commit in order
2	BNE R2,R3,Loop	6	10		11	11	Wait for ADDI
3	LD R2,0(R1)	7	8	9	10	12	Earliest possible
3	ADDI R2,R2,1	7	11		12	13	Wait for LD
3	SD R2,0(R1)	8	9		10	13	Wait for ADDI
3	ADDI R1,R1,4	8	9		10	14	Executes earlier
3	BNE R2,R3,Loop	9	13		14	14	Wait for ADDI

Renaming registers

Alternative renaming mechanisms

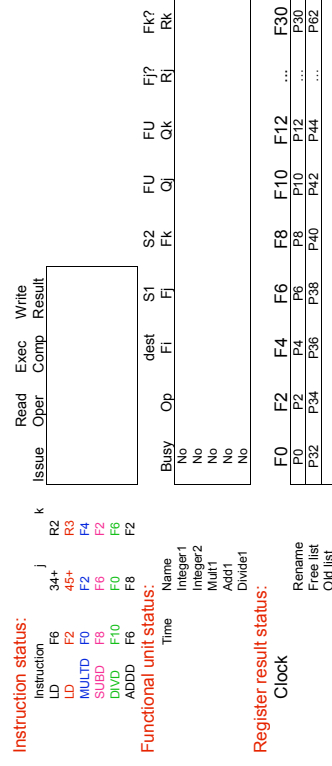
- ▶ Instruction target renaming done through reservation stations in Tomasulo
- ▶ Alternative implementations
 - ▶ Use ROB entries
 - ▶ Use a larger set of physical registers
 - ▶ Separate architecturally visible registers from physical registers
 - ▶ Architecture registers are registers visible to programmers
 - ▶ **Physical registers** > **Architecture registers** (e.g. 2x)
 - ▶ Physical registers used for renaming, if available

Renaming registers (cont.)

Implementation

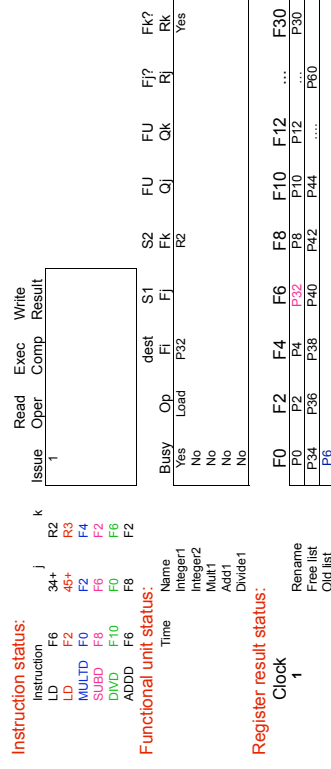
- ▶ Register renaming done through renaming map
 - ▶ Architecture register → physical register
 - ▶ Status of physical registers
 - ▶ **Free**: Register has not been assigned to instruction
 - ▶ **Invalid**: Register has been assigned to instruction for renaming, but value has not been computed yet
 - ▶ **Valid**: Register has been assigned to instruction for renaming, and value has been computed

Scoreboard with register renaming



Scoreboard with register renaming – Cycle 1

- ▶ P32 allocated from free list
- ▶ P6 (old physical register) moved to old list



Scoreboard with register renaming – Cycle 6

Instruction status:

Instruction	Issue	Read Op	Exec Comp	Write Result
LD F6	1	2	3	4
LD F2	2	3	4	5
MULTD F0	3	6		
SUBD F8	4	6		
DIVD F10	5			
ADD F6				

Functional unit status:

Time	Busy	Op	S1	S2	FU	Fk	Ok	Ri	Rk
Integer1	No		FI	FI					
Integer2	No	MULTD	P36	P34	P4			Yes	Yes
Add1	Yes	SUBD	P38	P32	P34			Yes	Yes
Divide1	Yes	DIVD	P40	P36	P32	Multi1		No	Yes

Register result status:

File-list	Rename	F0	F2	F4	F6	F8	F10	F12	...	F30
P36	P34	P4	P32	P38	P40	P36	P32	P34	P12	P30
P42	P44	P60	P6	P6	P2					
P0	P8	P10								

Instruction status:

Instruction	Issue	Read Op	Exec Comp	Write Result
LD F6	1	2	3	4
LD F2	2	3	4	5
MULTD F0	3	6		
SUBD F8	4	6		
DIVD F10	5			
ADD F6				

Functional unit status:

Time	Busy	Op	S1	S2	FU	Fk	Ok	Ri	Rk
Integer1	No		FI	FI					
Integer2	No	MULTD	P36	P34	P4			Yes	Yes
Add1	Yes	SUBD	P38	P32	P34			Yes	Yes
Divide1	Yes	DIVD	P40	P36	P32	Multi1		No	Yes

Register result status:

File-list	Rename	F0	F2	F4	F6	F8	F10	F12	...	F30
P36	P34	P4	P32	P38	P40	P36	P32	P34	P12	P30
P42	P44	P60	P6	P6	P2					
P0	P8	P10								

Dimitrios S. Nikolopoulos
 Recap
 Exceptions and rollbacks
Register renaming
 Branches and register renaming
 Summary

Scoreboard with register renaming – Cycle 7

Instruction status:

Instruction	Issue	Read Op	Exec Comp	Write Result
LD F6	1	2	3	4
LD F2	2	3	4	5
MULTD F0	3	6		
SUBD F8	4	6		
DIVD F10	5			
ADD F6				

Functional unit status:

Time	Busy	Op	S1	S2	FU	Fk	Ok	Ri	Rk
Integer1	No		FI	FI					
Integer2	No	MULTD	P36	P34	P4			Yes	Yes
Add1	Yes	SUBD	P38	P32	P34			Yes	Yes
Divide1	Yes	DIVD	P40	P36	P32	Multi1		No	Yes

Register result status:

File-list	Rename	F0	F2	F4	F6	F8	F10	F12	...	F30
P36	P34	P4	P32	P38	P40	P36	P32	P34	P12	P30
P42	P44	P60	P6	P6	P2					
P0	P8	P10								

Instruction status:

Instruction	Issue	Read Op	Exec Comp	Write Result
LD F6	1	2	3	4
LD F2	2	3	4	5
MULTD F0	3	6		
SUBD F8	4	6		
DIVD F10	5			
ADD F6				

Functional unit status:

Time	Busy	Op	S1	S2	FU	Fk	Ok	Ri	Rk
Integer1	No		FI	FI					
Integer2	No	MULTD	P36	P34	P4			Yes	Yes
Add1	Yes	SUBD	P38	P32	P34			Yes	Yes
Divide1	Yes	DIVD	P40	P36	P32	Multi1		No	Yes

Register result status:

File-list	Rename	F0	F2	F4	F6	F8	F10	F12	...	F30
P36	P34	P4	P32	P38	P40	P36	P32	P34	P12	P30
P42	P44	P60	P6	P6	P2					
P0	P8	P10								

Dimitrios S. Nikolopoulos
 Recap
 Exceptions and rollbacks
Register renaming
 Branches and register renaming
 Summary

Scoreboard with register renaming – Cycle 8

Instruction status:

Instruction	Issue	Read Op	Exec Comp	Write Result
LD F6	1	2	3	4
LD F2	2	3	4	5
MULTD F0	3	6		
SUBD F8	4	6		
DIVD F10	5			
ADD F6				

Functional unit status:

Time	Busy	Op	S1	S2	FU	Fk	Ok	Ri	Rk
Integer1	No		FI	FI					
Integer2	No	MULTD	P36	P34	P4			Yes	Yes
Add1	Yes	SUBD	P38	P32	P34			Yes	Yes
Divide1	Yes	DIVD	P40	P36	P32	Multi1		No	Yes

Register result status:

File-list	Rename	F0	F2	F4	F6	F8	F10	F12	...	F30
P36	P34	P4	P32	P38	P40	P36	P32	P34	P12	P30
P42	P44	P60	P6	P6	P2					
P0	P8	P10								

Instruction status:

Instruction	Issue	Read Op	Exec Comp	Write Result
LD F6	1	2	3	4
LD F2	2	3	4	5
MULTD F0	3	6		
SUBD F8	4	6		
DIVD F10	5			
ADD F6				

Functional unit status:

Time	Busy	Op	S1	S2	FU	Fk	Ok	Ri	Rk
Integer1	No		FI	FI					
Integer2	No	MULTD	P36	P34	P4			Yes	Yes
Add1	Yes	SUBD	P38	P32	P34			Yes	Yes
Divide1	Yes	DIVD	P40	P36	P32	Multi1		No	Yes

Register result status:

File-list	Rename	F0	F2	F4	F6	F8	F10	F12	...	F30
P36	P34	P4	P32	P38	P40	P36	P32	P34	P12	P30
P42	P44	P60	P6	P6	P2					
P0	P8	P10								

Dimitrios S. Nikolopoulos
 Recap
 Exceptions and rollbacks
Register renaming
 Branches and register renaming
 Summary

Scoreboard with register renaming – Cycle 9

Instruction status:

Instruction	Issue	Read Op	Exec Comp	Write Result
LD F6	1	2	3	4
LD F2	2	3	4	5
MULTD F0	3	6		
SUBD F8	4	6		
DIVD F10	5			
ADD F6				

Functional unit status:

Time	Busy	Op	S1	S2	FU	Fk	Ok	Ri	Rk
Integer1	No		FI	FI					
Integer2	No	MULTD	P36	P34	P4			Yes	Yes
Add1	Yes	SUBD	P38	P32	P34			Yes	Yes
Divide1	Yes	DIVD	P40	P36	P32	Multi1		No	Yes

Register result status:

File-list	Rename	F0	F2	F4	F6	F8	F10	F12	...	F30
P36	P34	P4	P32	P38	P40	P36	P32	P34	P12	P30
P42	P44	P60	P6	P6	P2					
P0	P8	P10								

Instruction status:

Instruction	Issue	Read Op	Exec Comp	Write Result
LD F6	1	2	3	4
LD F2	2	3	4	5
MULTD F0	3	6		
SUBD F8	4	6		
DIVD F10	5			
ADD F6				

Functional unit status:

Time	Busy	Op	S1	S2	FU	Fk	Ok	Ri	Rk
Integer1	No		FI	FI					
Integer2	No	MULTD	P36	P34	P4			Yes	Yes
Add1	Yes	SUBD	P38	P32	P34			Yes	Yes
Divide1	Yes	DIVD	P40	P36	P32	Multi1		No	Yes

Register result status:

File-list	Rename	F0	F2	F4	F6	F8	F10	F12	...	F30
P36	P34	P4	P32	P38	P40	P36	P32	P34	P12	P30
P42	P44	P60	P6	P6	P2					
P0	P8	P10								

Dimitrios S. Nikolopoulos
 Recap
 Exceptions and rollbacks
Register renaming
 Branches and register renaming
 Summary

Scoreboard with register renaming – Cycle 10

- ▶ P42 allocated from free list
- ▶ P32 pushed to old list, still in use by DIVD

Instruction status:

Instruction	j	k	R2	R3	R4	R5	R6	R7	R8	R9
LD	F6	F2	34+	45+	F2	F4	F6	F8	F10	F12
MULTD	F0	F2	F4	F6	F8	F10	F12	F14	F16	F18
SUBD	F8	F6	F4	F2	F0	F10	F12	F14	F16	F18
DIVD	F10	F8	F6	F4	F2	F0	F10	F12	F14	F16
ADD	F6	F8	F10	F12	F14	F16	F18	F20	F22	F24

Functional unit status:

Time	Name	Integer1	Integer2	5 Mult1	2 Add1	Divide1
10						

Register result status:

Free list	Old list
P36, P34, P44, P0	P42, P38, P40, P2, P10, P32

Clock 10

Scoreboard with register renaming – Cycle 12

Instruction status:

Instruction	j	k	R2	R3	R4	R5	R6	R7	R8	R9
LD	F6	F2	34+	45+	F2	F4	F6	F8	F10	F12
MULTD	F0	F2	F4	F6	F8	F10	F12	F14	F16	F18
SUBD	F8	F6	F4	F2	F0	F10	F12	F14	F16	F18
DIVD	F10	F8	F6	F4	F2	F0	F10	F12	F14	F16
ADD	F6	F8	F10	F12	F14	F16	F18	F20	F22	F24

Functional unit status:

Time	Name	Integer1	Integer2	4 Mult1	1 Add1	Divide1
12						

Register result status:

Free list	Old list
P36, P34, P44, P0	P42, P38, P40, P2, P10, P32

Clock 12

Scoreboard with register renaming – Cycle 11

Instruction status:

Instruction	j	k	R2	R3	R4	R5	R6	R7	R8	R9
LD	F6	F2	34+	45+	F2	F4	F6	F8	F10	F12
MULTD	F0	F2	F4	F6	F8	F10	F12	F14	F16	F18
SUBD	F8	F6	F4	F2	F0	F10	F12	F14	F16	F18
DIVD	F10	F8	F6	F4	F2	F0	F10	F12	F14	F16
ADD	F6	F8	F10	F12	F14	F16	F18	F20	F22	F24

Functional unit status:

Time	Name	Integer1	Integer2	5 Mult1	2 Add1	Divide1
11						

Register result status:

Free list	Old list
P36, P34, P44, P0	P42, P38, P40, P2, P10, P32

Clock 11

Scoreboard with register renaming – Cycle 13

Instruction status:

Instruction	j	k	R2	R3	R4	R5	R6	R7	R8	R9
LD	F6	F2	34+	45+	F2	F4	F6	F8	F10	F12
MULTD	F0	F2	F4	F6	F8	F10	F12	F14	F16	F18
SUBD	F8	F6	F4	F2	F0	F10	F12	F14	F16	F18
DIVD	F10	F8	F6	F4	F2	F0	F10	F12	F14	F16
ADD	F6	F8	F10	F12	F14	F16	F18	F20	F22	F24

Functional unit status:

Time	Name	Integer1	Integer2	3 Mult1	0 Add1	Divide1
13						

Register result status:

Free list	Old list
P36, P34, P44, P0	P42, P38, P40, P2, P10, P32

Clock 13

Scoreboard with register renaming – Cycle 18

- ▶ MULTD commits, recycle old register P0

Instruction status:

Instruction	F6	F8	F10	F12	F14
LD	34+	R3	R2	R3	R4
LD	F2	4E+	F2	F4	F5
MULTD	F0	F2	F4	F6	F7
SUBD	F8	F6	F2	F4	F9
DIVD	F10	F0	F6	F8	F9
ADD	F6	F8	F2	F4	F14

Functional unit status:

Time	Name	S1	S2	FU	FU?	FK?	RK
No	Integer1	Fi	Fi	Fi	Fi	Fi	Ri
No	Integer2	Fi	Fi	Fi	Fi	Fi	Ri
No	Mult1	Fi	Fi	Fi	Fi	Fi	Ri
Yes	Add1	Fi	Fi	Fi	Fi	Fi	Ri
40	Divide1	Fi	Fi	Fi	Fi	Fi	Ri

Register result status:

Free list	Old list
F0	F2
P36	F4
P44	P42
P8	P60
	P32

Branches and renaming

How do we undo speculative commits to registers?

- ▶ Commit stage and reorder buffer resolve branches in Tomasulo
 - ▶ Instructions commit in-order
 - ▶ Reorder buffer holds speculative results
 - ▶ Squashing instructions in ROB easy
- ▶ If speculative instructions modify physical registers
 - ▶ Physical registers hold speculative results
 - ▶ How do we undo these modifications?

Scoreboard with register renaming – Cycle 19

- ▶ SUBD commits, recycle old register P8

Instruction status:

Instruction	F6	F8	F10	F12	F14
LD	34+	R3	R2	R3	R4
LD	F2	4E+	F2	F4	F5
MULTD	F0	F2	F4	F6	F7
SUBD	F8	F6	F2	F4	F9
DIVD	F10	F0	F6	F8	F9
ADD	F6	F8	F2	F4	F14

Functional unit status:

Time	Name	S1	S2	FU	FU?	FK?	RK
No	Integer1	Fi	Fi	Fi	Fi	Fi	Ri
No	Integer2	Fi	Fi	Fi	Fi	Fi	Ri
No	Mult1	Fi	Fi	Fi	Fi	Fi	Ri
Yes	Add1	Fi	Fi	Fi	Fi	Fi	Ri
39	Divide1	Fi	Fi	Fi	Fi	Fi	Ri

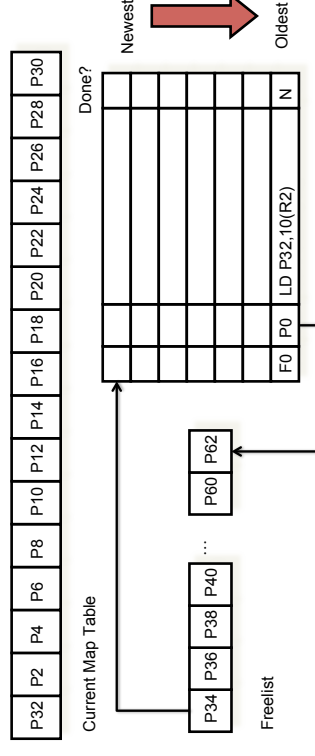
Register result status:

Free list	Old list
F0	F2
P36	F4
P44	P42
P8	P60
	P32

Register renaming with branches

MIPS R10000 solution

- ▶ Rename map stores instruction status
- ▶ Reorder buffer maintains instruction order



Concept of in-order instruction commit

Fixing speculation errors

- ▶ Out-of-order execution and speculation boost available ILP
- ▶ In-order instruction commit helps:
 - ▶ Resolve **WAW, WAR hazards**
 - ▶ Resolve memory hazards
 - ▶ Speculate past branches
 - ▶ Preserve precise exceptions
- ▶ Various implementations (ROB, register map, reservation stations)
- ▶ Used in most modern high-end processors

Implications of out-of-order execution

Complexity

- ▶ Expensive hardware
 - ▶ Large storage structures, associative searches
 - ▶ Complex issue logic
 - ▶ Complex branch prediction and speculation hardware
 - ▶ Hard to handle expensive exceptions
- ▶ **Diminishing performance returns**
 - ▶ Limits of ILP topic of next lecture
 - ▶ Finding more parallelism while reducing complexity topic of following lectures