

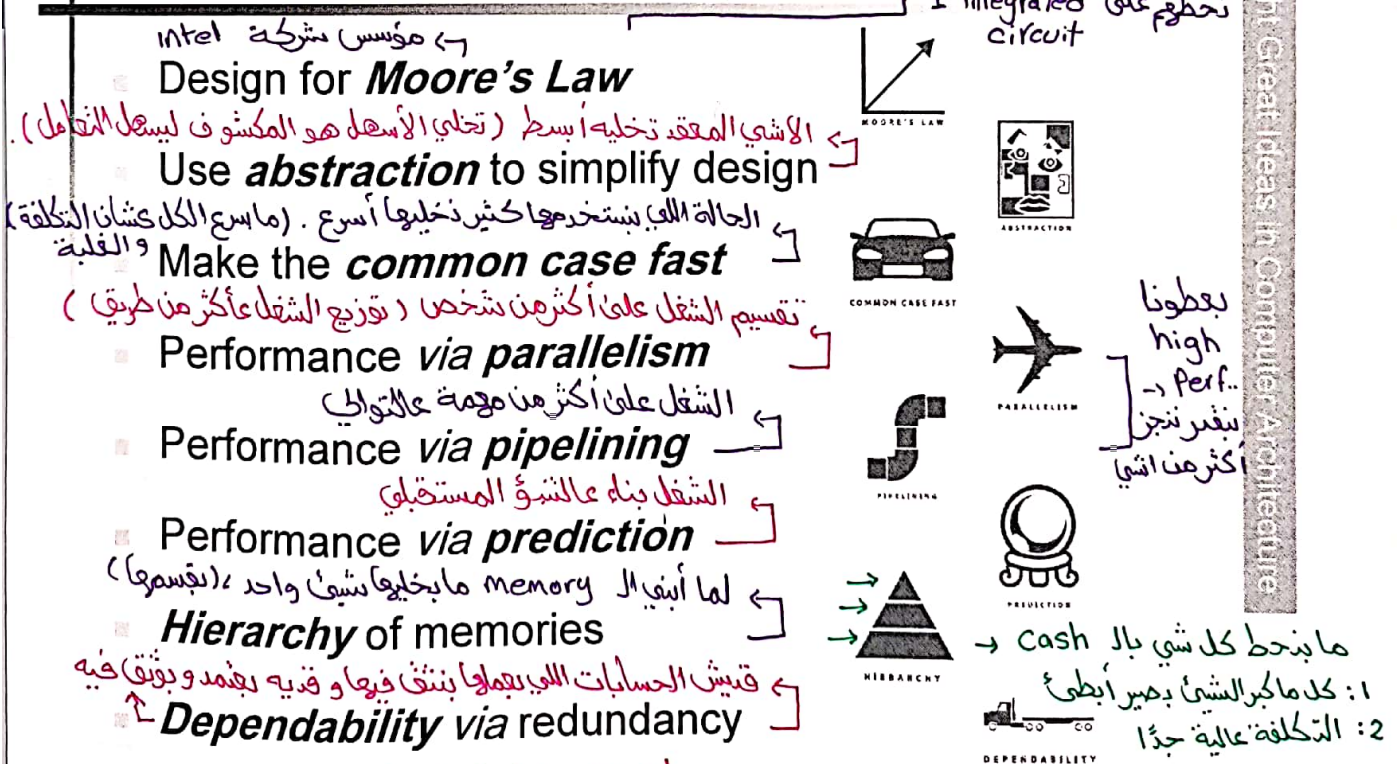
# Content

- 1.2 Eight Great Ideas in Computer Architecture (Review)
- 1.5 Technologies for Building Processors and Memory
- 1.6 Performance (Review)
- 1.7 The Power Wall
- 1.8 The Sea Change: The Switch from Uniprocessors to Multiprocessors
- 1.9 Real Stuff: Benchmarking the Intel Core i7
- 1.10 Fallacies and Pitfalls
- 1.11 Concluding Remarks



## Eight Great Ideas

ملاحظة تبيّن انه كل سنة أو سنة ونصف  
 بنضاع عدد transistor الذي يقدر  
 نخطم على integrated circuit



§1.2 Eight Great Ideas in Computer Architecture



# Content

- 1.2 Eight Great Ideas in Computer Architecture  
(Review)
- 1.5 Technologies for Building Processors and Memory
- 1.6 Performance (Review)
- 1.7 The Power Wall
- 1.8 The Sea Change: The Switch from Uniprocessors to Multiprocessors
- 1.9 Real Stuff: Benchmarking the Intel Core i7
- 1.10 Fallacies and Pitfalls
- 1.11 Concluding Remarks




بحتاج لتحقيق higher performance : وجود تكلفة و دوائر لعمل المعجات  
و بعالم الحواسيب هاي الدوائر لازم تكون أكبر و بتكون أغلى وأسرع.

## Technology Trends

\* vacuum tube و transistor واحد بينما ال chip فيها أكثر من transistor \*

Electronics technology continues to evolve

- Increased capacity and performance
- Reduced cost



DRAM capacity

Year	Technology	Relative performance/cost
1951	Vacuum tube	1
1965	Transistor	35
1975	Integrated circuit (IC)	900
1995	Very large scale IC (VLSI)	2,400,000
2013	Ultra large scale IC	250,000,000,000





# Semiconductor Technology

المادة الرئيسية ووظيفتها

Silicon: semiconductor

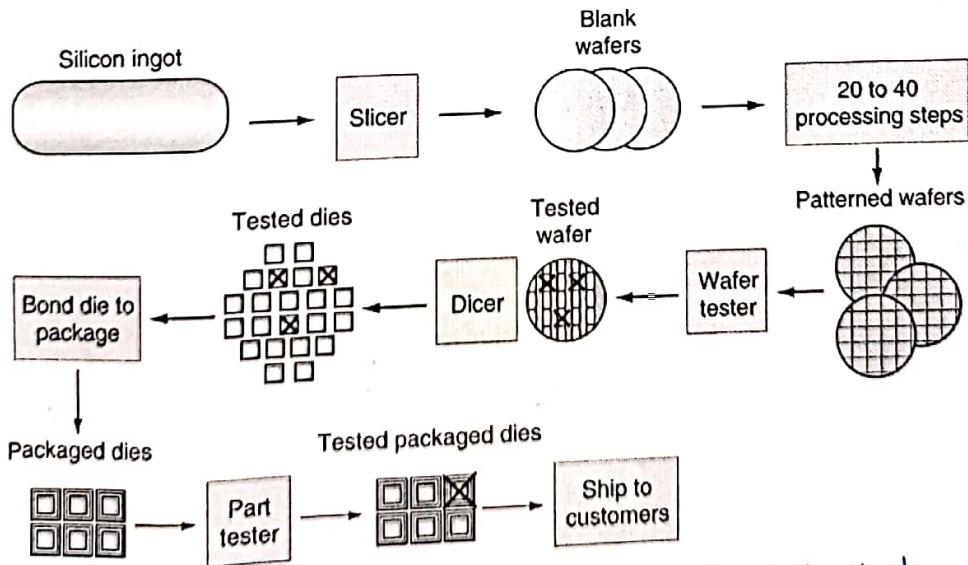
Add materials to transform properties:

- ↳ Conductors موصل
- ↳ Insulators عازل
- Switch مفتاح

## Manufacturing ICs

كيف نبتعمل

« تكلفة عشان هيك حاليا شركات معينة اللي بتعملها بس »




\* كل ما يكون عندي خبرة عدد القطع الخرابية بتقل فال Yield بيمس أعلى فالنكفة بتقل و مالك ذلك.

Yield: <sup>نسبة</sup> proportion of working dies per wafer  
 عدد ال dies الصالجات تقسيم العدد الكلي  
 لإنه بال wafer ممكن يكون فيه dies معطلة.

# Intel Core i7 Wafer

كل ما زادت مساحة ال die ينهس تكلفتها أعلى



مشكلتهم ممكن يكونوا حلالين

مصنوع عليها هالقد قطر

ينحدد قديه ال transistor صفار أو كبار

باستخدام هاي

300mm wafer, 280 chips, 32nm technology

Each chip is 20.7 x 10.5 mm

كلما كان الرقم أصغر ال transistor يكونوا أقل



أبعادها

(  $cost \propto area$  ← ولكن هيش linear ، ممكن يكون تربيعي )

## Integrated Circuit Cost

العلاقة بين ال cost

$$Cost \text{ per die} = \frac{Cost \text{ per wafer}}{Dies \text{ per wafer} \times Yield}$$

عدد ال Chips بال wafer

$$Dies \text{ per wafer} \approx \frac{Wafer \text{ area}}{Die \text{ area}}$$

بتطلع 280 تقريبا Chips

$$Yield = \frac{1}{(1 + (Defects \text{ per area} \times Die \text{ area}/2))^2}$$

له بالأصل أكبر من 280

كلما كانت الكثافة أقل كلما كان ال Yield أفضل

### Nonlinear relation to area and defect rate

بس لإنه اللي جايبين بأطراف الدائرة هيش كاملين فبنلغيهم

- Wafer cost and area are fixed
- Defect rate determined by manufacturing process
- Die area determined by architecture and circuit design

$$* Defect \text{ per area} = \frac{\# \text{ of defects}}{wafer \text{ area}}$$





# Content

- 1.2 Eight Great Ideas in Computer Architecture  
(*Review*)
- 1.5 Technologies for Building Processors and Memory
- 1.6 Performance (*Review*)
- 1.7 The Power Wall
- 1.8 The Sea Change: The Switch from Uniprocessors to Multiprocessors
- 1.9 Real Stuff: Benchmarking the Intel Core i7
- 1.10 Fallacies and Pitfalls
- 1.11 Concluding Remarks



## Response Time and Throughput

- Response time
  - How long it takes to do a task → قديه وقت لأنفذ برنامج أو instruction
- Throughput
  - Total work done per unit time → قديه بنفصر نعمل نعمل بوحدة الزمن
  - e.g., tasks/transactions/... per hour استجابة
- How are response time and throughput affected by
  - Replacing the processor with a faster version?
  - Adding more processors?
- We'll focus on response time for now...



# Relative Performance

Define Performance =  $1/\text{Execution Time}$

"X is  $n$  time faster than Y" « Performance =  $\frac{1}{T}$  »

لعمري يكون أكبر ما يمكن

لعمري يكون أقل ما يمكن

Parallel units عدد ال

$$\text{Performance}_X / \text{Performance}_Y$$

ومش واحد لا P

$$= \text{Execution time}_Y / \text{Execution time}_X = n$$

$$\text{Throughput} = \frac{P}{T}$$

لما اشتغل عال Parallelism أو pipelining مش دقيقة فبمير

Time العملية الواحدة

Example: time taken to run a program

- 10s on A, 15s on B
- Execution Time<sub>B</sub> / Execution Time<sub>A</sub> = 15s / 10s = 1.5
- So A is 1.5 times faster than B



# Measuring Execution Time

النوع Elapsed time

- Total response time, including all aspects
  - Processing, I/O, OS overhead, idle time
- Determines system performance

النواحي

CPU time

الوقت المستغرق في معالجة وطريقة هيته


- Time spent processing a given job
  - Discounts I/O time, other jobs' shares
- Comprises user CPU time and system CPU time
- Different programs are affected differently by CPU and system performance





# CPU Clocking

- Operation of digital hardware governed by a constant-rate clock <sup>يتحكم</sup>



- Clock period: duration of a clock cycle
  - e.g., 250ps = 0.25ns =  $250 \times 10^{-12}$ s
- Clock frequency (rate): cycles per second
  - e.g., 4.0GHz = 4000MHz =  $4.0 \times 10^9$ Hz



# CPU Time

$$\begin{aligned} \text{CPU Time} &= \text{CPU Clock Cycles} \times \text{Clock Cycle Time} \\ &= \frac{\text{CPU Clock Cycles}}{\text{Clock Rate}} \end{aligned}$$

- Performance improved by
  - Reducing number of clock cycles
  - Increasing clock rate
  - Hardware designer must often trade off clock rate against cycle count



# Instruction Count and CPI

$$\text{Clock Cycles} = \text{Instruction Count} \times \text{Cycles per Instruction}$$

$$\text{CPU Time} = \text{Instruction Count} \times \text{CPI} \times \text{Clock Cycle Time}$$

$$f = \frac{\text{Instruction Count} \times \text{CPI}}{T} \leftarrow \text{Clock Rate (f)}$$

أهم معادلات

- Instruction Count for a program
  - Determined by program, ISA and compiler
- Average cycles per instruction
  - Determined by CPU hardware
  - If different instructions have different CPI
    - Average CPI affected by instruction mix



وقت انجاز task

Chapter 1 — Computer Abstractions and Technology

Ex 8 Time = 0.1 seconds / CPI = 0.5 / IC = 400 million Instructions

بقدر ينفذ 2 instr. في ال cycle الواحدة

( f = ?? )

Sol :  $0.1 = 400 \times 10^6 \times 0.5 \times T$   
 $\rightarrow T = 5 \times 10^{-10}$   
 $f = \frac{1}{T} \rightarrow f = 2 \times 10^9$   
 $= 2 \text{ GHz}$

## CPI in More Detail

- If different instruction classes take different numbers of cycles

$$\text{Clock Cycles} = \sum_{i=1}^n (\text{CPI}_i \times \text{Instruction Count}_i)$$

- Weighted average CPI

$$\text{CPI} = \frac{\text{Clock Cycles}}{\text{Instruction Count}} = \sum_{i=1}^n \left( \text{CPI}_i \times \frac{\text{Instruction Count}_i}{\text{Instruction Count}} \right)$$

Relative frequency





# Performance Summary

\* كلما قل الـ time زاد الـ performance

## The BIG Picture

$$\text{CPU Time} = \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Clock cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Clock cycle}}$$

IC كـ في program في البرنامج  
 CPI ↑  
 T ↑  
 الوحدة التي بتطلع CPU time →

- Performance depends on
    - Algorithm: affects IC, possibly CPI
    - Programming language: affects IC, CPI
    - Compiler: affects IC, CPI
    - Instruction set architecture: affects IC, CPI,  $T_c$
- = قديه محتاج Seconds لينفذ program معين



كيفية تقليل 1 -  $\frac{\text{Instructions}}{\text{Program}}$  يعتمد على الخوارزميات و كيف استخدموا للتقليل من الـ Instructions بالعملية الواحدة.

Content 2 -  $\frac{\text{clock cycles}}{\text{Instructions}}$  تخلفوا أكثر من Instruc. بنفس الـ cycle يعني يعمل Pipelining و هكذا.

1.2 Eight Great Ideas in Computer Architecture 3 -  $\frac{\text{Seconds}}{\text{clock cycle}}$  لو بدى أقل الـ Period لازم اشغل على high frequency (Review) عن طريق استخدام smaller transistors لإنتاج بصيروا أسرع

1.5 Technologies for Building Processors and Memory

1.6 Performance (Review)

1.7 The Power Wall

1.8 The Sea Change: The Switch from Uniprocessors to Multiprocessors

1.9 Real Stuff: Benchmarking the Intel Core i7


1.10 Fallacies and Pitfalls

1.11 Concluding Remarks

# Power Trends

حنكبي عن استهلاك الطاقة الكوبانية

Processors \*  
يستهلوكوا طاقة  
كوبانية وهي  
الطاقة مع الزمن  
بتزداد \*



بتشحن وبتفرغ

## In CMOS IC technology

ال Power تتناسب طرديا معها

transistors ما في تيار بعد فيهم فيننظر الهم انهم capacitors

$$\text{Power} = \frac{1}{2} \times \text{Capacitive load} \times \text{Voltage}^2 \times \text{Frequency}$$

Dynamic Power

(يعتمد على : حجم العدد ل transistors)

كلما قل بقول ال Power حجم ال transistor

5V → 1V

×1000

يعني كل ما زاد ال transistor



له من 1.5 ولها أقل من 1 بشوي وما بتل أكثر من هيك

leakage current \* supply voltage = \* ال static power ال chip يستهلكها

## Reducing Power

- Suppose a new CPU has
  - 85% of capacitive load of old CPU
  - 15% voltage and 15% frequency reduction

$$\frac{P_{new}}{P_{old}} = \frac{C_{old} \times 0.85 \times (V_{old} \times 0.85)^2 \times F_{old} \times 0.85}{C_{old} \times V_{old}^2 \times F_{old}} = 0.85^4 = 0.52$$

### The power wall

- We can't reduce voltage further
- We can't remove more heat
- How else can we improve performance?



له بزيوا ال Performance بالوقت الحاضر عن طريق





# Multiprocessors

one chip MicroProcessor : برغمة الوحدة الحاسبة بال Computer الكن شتر انما تكون على

## Multicore microprocessors

More than one processor per chip

Requires explicitly parallel programming  
← لازم نلجا اليه لنتفيد من Multicore

Compare with instruction level parallelism

← كل في عبء

عالمبروجين انهم يعيدوا كتابة برامج حتى تصير Parallel Program حتى تصير تشغيل Multiple cores

Hardware executes multiple instructions at once

Hidden from the programmer

Hard to do

Programming for performance

لانهم برلمجك لازم تكون مع لازم تطيل Performance اعلى

Load balancing

Optimizing communication and synchronization

← نخليه قليل وسريع

لانهم يحتاجوا وقت ليتفاهموا مع بعض

بديك تقسم العمل على اجزاء متساوية بحيث ال Multiple cores

كل core ياخذ تشغل قد التالي لانه لو واحد أخذ تشغل أكثر من غيره رح ياخذ كل البرنامج ولازم يشغله لحتى يخلص

## Content

1.2 Eight Great Ideas in Computer Architecture (Review)

1.5 Technologies for Building Processors and Memory

1.6 Performance (Review)

1.7 The Power Wall

1.8 The Sea Change: The Switch from Uniprocessors to Multiprocessors

1.9 Real Stuff: Benchmarking the Intel Core i7

1.10 Fallacies and Pitfalls

1.11 Concluding Remarks







# SPEC Power Benchmark

Power consumption of server at different workload levels

Power مهمة لإنفا تبستهلك مال وكل ما زاد  
 ال Power بنزيد فانورة الكوربا  
 ال Power وللحفاظ على الطبيعة  
 \* وانشان البطارية تعيش

- Performance: ssj\_ops/sec
- Power: Watts (Joules/sec)

Power benchmark ← مجموع

$$\text{Overall ssj\_ops per Watt} = \left( \sum_{i=0}^{10} \text{ssj\_ops}_i \right) / \left( \sum_{i=0}^{10} \text{power}_i \right)$$

مجموع

بعمني لو بيدي  
 أشوق الجهاز  
 قديه بيأخذ وقت  
 لمهمة  
 وحدة  
 بعضي لما بيدي  
 أشغل أكثر من  
 مهمة  
 بنفس  
 الوقت

The SPEC cpu 2017 benchmark package contains 43 benchmarks, organised into four suites:

- 1- SPEC speed integer and SPEC speed floating ] are used for comparing time for a computer to complete single tasks
- 2- SPEC rate integer and SPEC rate floating ] measure the throughput or work per unit of time



نستخدم بـ general app

Chapter 1 — Computer Abstractions and Technology — 2

مثل Compiler أو Sorting أو Excel  
 وهذا . يستخدم للبرمجيات  
 التي بتعملنا  
 أصور  
 مختلفة

غالبية الحسابات العلمية والهندسية  
 تستخدمها.

## SPECpower\_ssj2008 for Xeon X5650

operation per second ← وحرته

Target Load %	Performance (ssj_ops)	Average Power (Watts)
100%	865,618	258
90%	786,688	242
80%	698,051	224
70%	607,826	204
60%	521,391	185
50%	436,757	170
40%	345,919	157
30%	262,071	146
20%	176,061	135
10%	86,784	121
0%	0	80
Overall Sum	4,787,166	1,922
Σssj_ops/Σpower =		2,490

← تشغيل  
 ال Processor  
 على  
 different target  
 load

لـ كلما كان أكبر كلما كان أفضل



بصير في تشغيل وكل ما زاد  
 بزيد الاستهلاك  
 للطاقة .

← كلما زاد ال target load بنزيد ال average power  
 Chapter 1 — Computer Abstractions and Technology — 2



# Content

- 1.2 Eight Great Ideas in Computer Architecture  
(Review)
- 1.5 Technologies for Building Processors and Memory
- 1.6 Performance (Review)
- 1.7 The Power Wall
- 1.8 The Sea Change: The Switch from Uniprocessors to Multiprocessors
- 1.9 Real Stuff: Benchmarking the Intel Core i7
- 1.10 Fallacies and Pitfalls
- 1.11 Concluding Remarks



## Pitfall: Amdahl's Law

- Improving an aspect of a computer and expecting a proportional improvement in overall performance

$$T_{\text{improved}} = \frac{T_{\text{affected}}}{\text{improvement factor}} + T_{\text{unaffected}}$$

*Handwritten notes:*   
 ← الزمن الجديد  $T_{\text{improved}}$    
 ← العالم Computer eng قديم  $T_{\text{affected}}$    
 ← ينقل عليه  $T_{\text{unaffected}}$    
 ← تقسم ال exec. time  $T_{\text{affected}}$    
 ← ال  $T_{\text{unaffected}}$    
 ←  $T_{\text{unaffected}}$    
 ← ما ينقل عليه  $T_{\text{unaffected}}$

- Example: multiply accounts for 80s/100s → *بيأخو وقت 80 من أصل 100*
- How much improvement in multiply performance to get 5x overall? *لما 5 =  $\frac{T_{\text{old}}}{T_{\text{improved}}}$*

$$20 = \frac{80}{n} + 20$$

Can't be done!

$$5 = \frac{T_{\text{old}}}{T_{\text{improved}}}$$

$$5 = \frac{100}{T_i}$$

$$\Rightarrow T_i = 20$$

*Handwritten notes:*   
 ← speedup  $5 = \frac{100}{T_i}$    
 →  $T_i = 20$

- Corollary: make the common case fast



# Fallacy: Low Power at Idle

- Look back at i7 power benchmark
  - At 100% load: 258W
  - At 50% load: 170W (66%)
  - At 10% load: 121W (47%)
- Google data center
  - Mostly operates at 10% – 50% load
  - At 100% load less than 1% of the time
- Consider designing processors to make power proportional to load



# Pitfall: MIPS as a Performance Metric

- MIPS: Millions of Instructions Per Second
  - Doesn't account for
    - Differences in ISAs between computers
    - Differences in complexity between instructions

$$\begin{aligned} \text{MIPS} &= \frac{\text{Instruction count}}{\text{Execution time} \times 10^6} \\ &= \frac{\text{Instruction count}}{\frac{\text{Instruction count} \times \text{CPI}}{\text{Clock rate}} \times 10^6} = \frac{\text{Clock rate}}{\text{CPI} \times 10^6} \end{aligned}$$

- CPI varies between programs on a given CPU





# Content

- 1.2 Eight Great Ideas in Computer Architecture  
(*Review*)
- 1.5 Technologies for Building Processors and  
Memory
- 1.6 Performance (*Review*)
- 1.7 The Power Wall
- 1.8 The Sea Change: The Switch from  
Uniprocessors to Multiprocessors
- 1.9 Real Stuff: Benchmarking the Intel Core i7
- 1.10 Fallacies and Pitfalls
- 1.11 Concluding Remarks



## Concluding Remarks

- Cost/performance is improving
  - Due to underlying technology development
- Execution time: the best performance measure
- Power is a limiting factor
  - Use parallelism to improve performance



# Chapter 4

## The Processor

*Adapted by Prof. Gheith Abandah*

## Contents

- 4.6 Pipelined Datapath and Control (Review)
- 4.7 Data Hazards: Forwarding versus Stalling
- 4.8 Control Hazards
- 4.9 Exceptions
- 4.10 Parallelism via Instructions
- 4.11 Real Stuff: The ARM Cortex-A53 and Intel  
Core i7 Pipelines
- 4.14 Fallacies and Pitfalls
- 4.15 Concluding Remarks



# Contents

## 4.6 Pipelined Datapath and Control (Review)

Five-Stage Pipeline

Pipeline Control

Pipeline Hazards



## Five-Stage Pipeline

**F:** Fetch instruction from the instruction memory → بتجيب ال inst من inst. memory

**D:** Decode instruction and read operands → تحليل ال inst وقرائتها

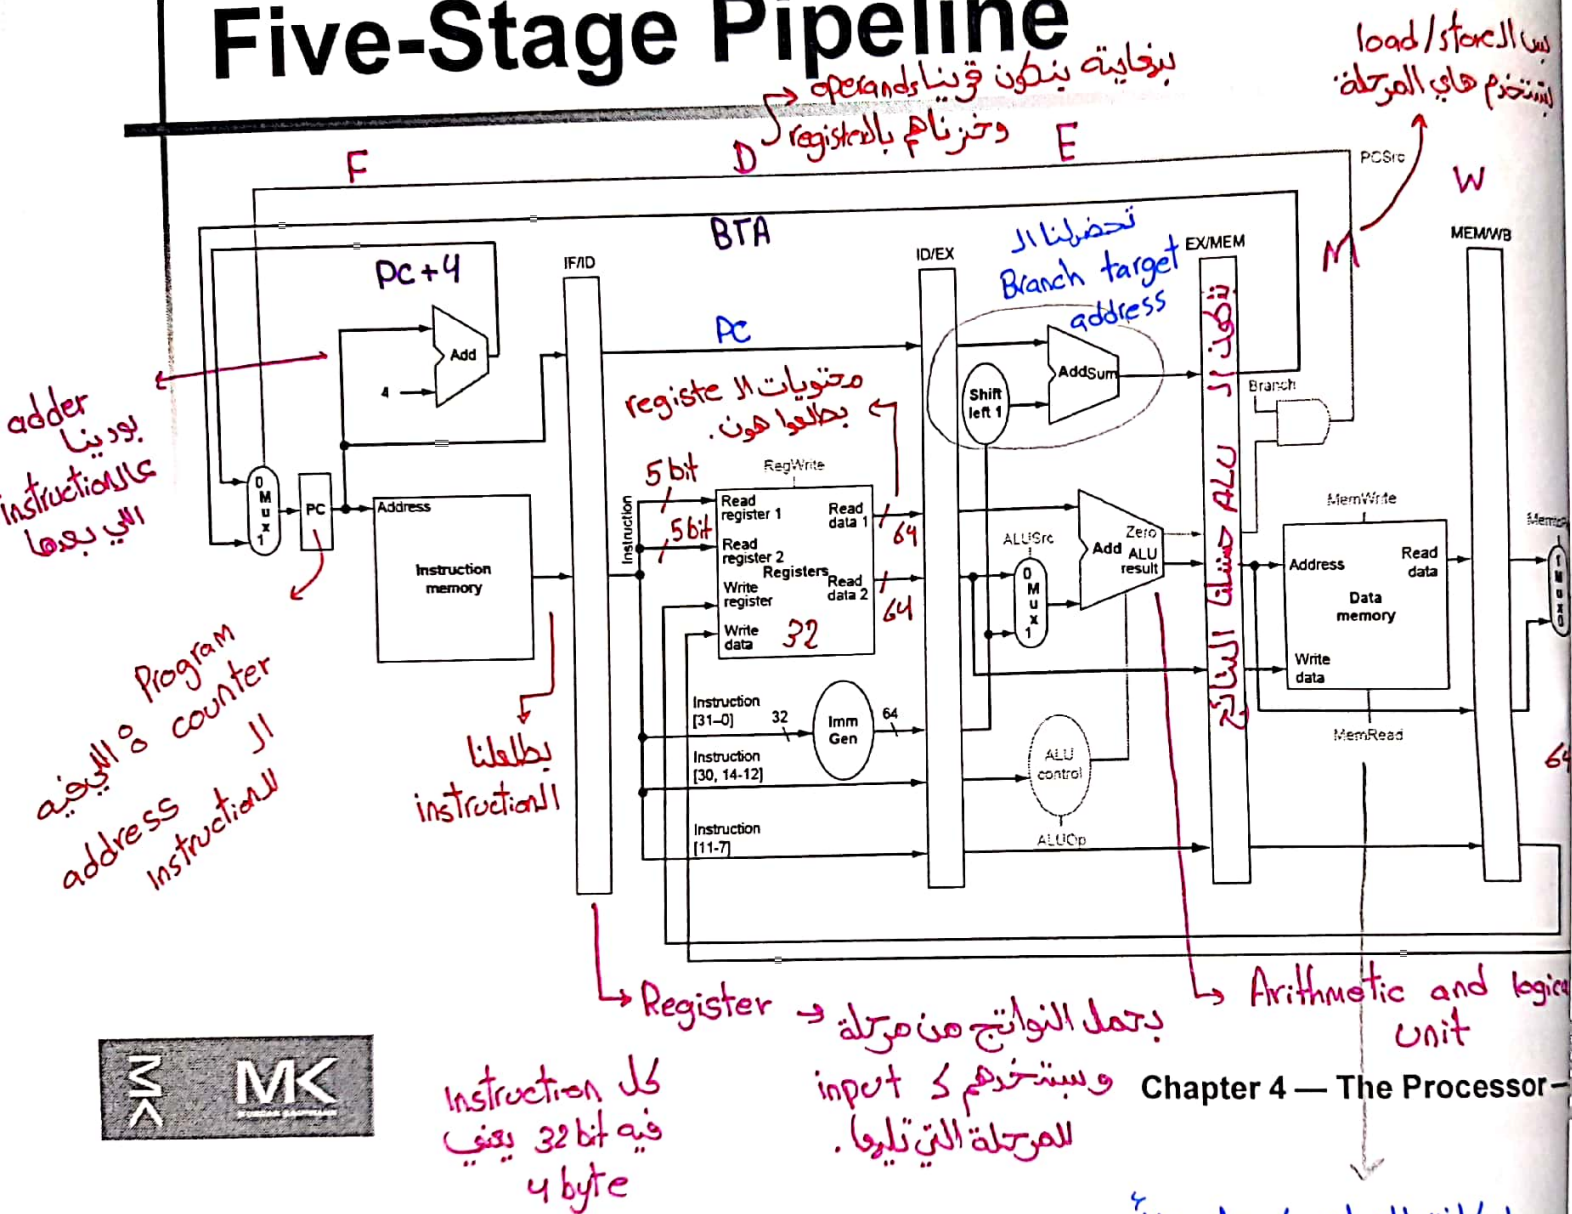
**E:** Execute operation or calculate address → الحساب

**M:** Memory access → اذا كانت Load / Store

**W:** Write result to the register



# Five-Stage Pipeline



## Pipelined Control

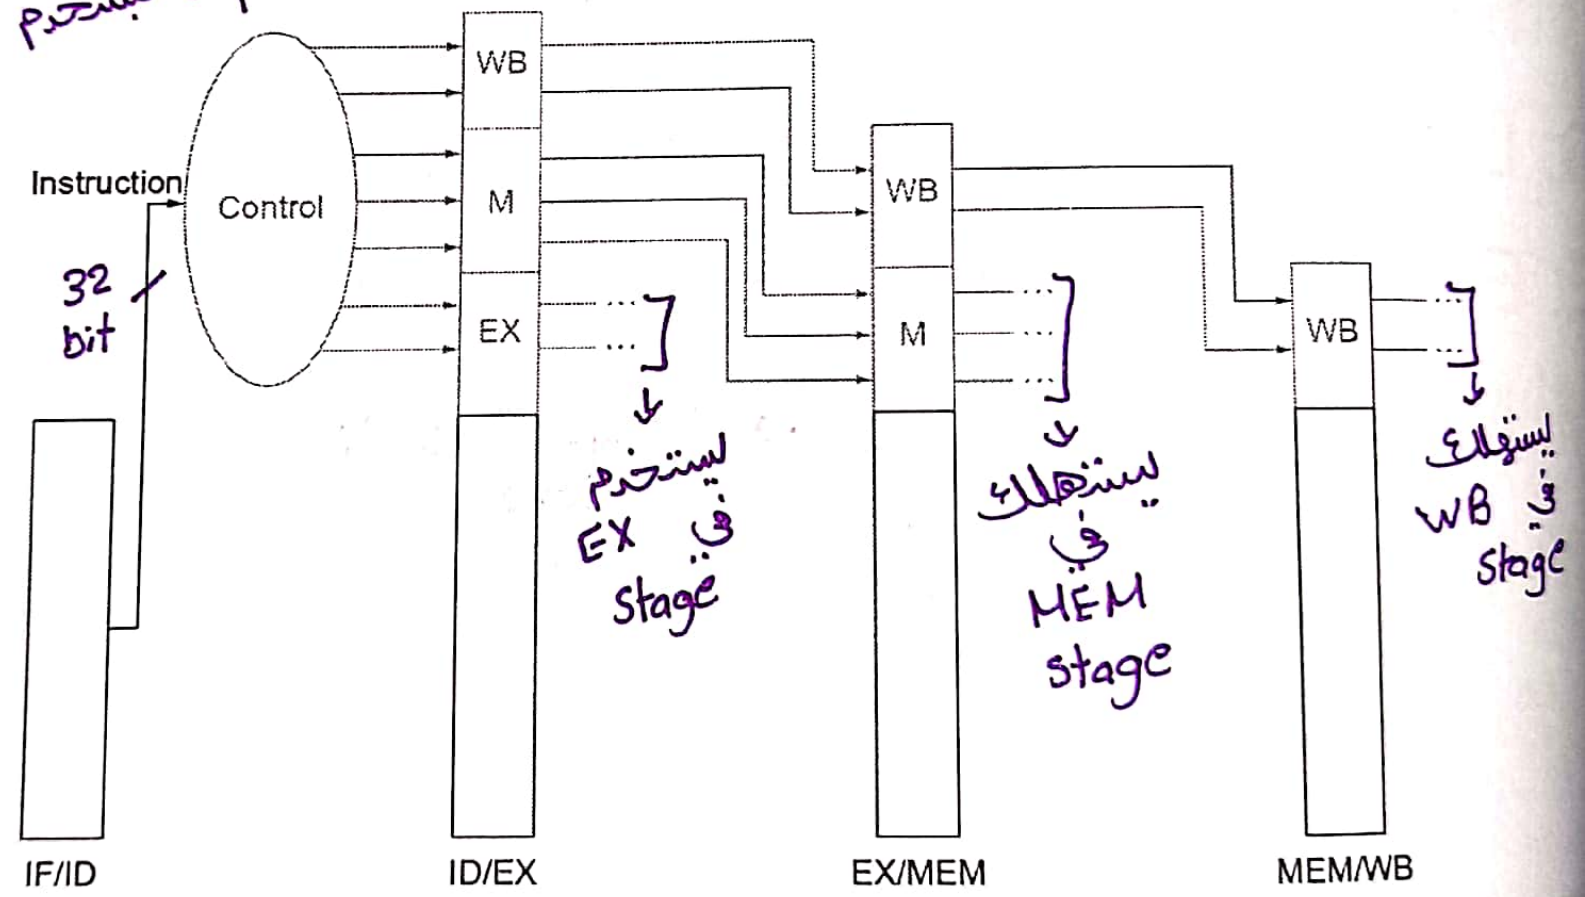
Control signals derived from instruction



# Control signals derived from instruction

- As in single-cycle implementation

استخدم 7-bit للبيتم opcode  
لنظارة control signals



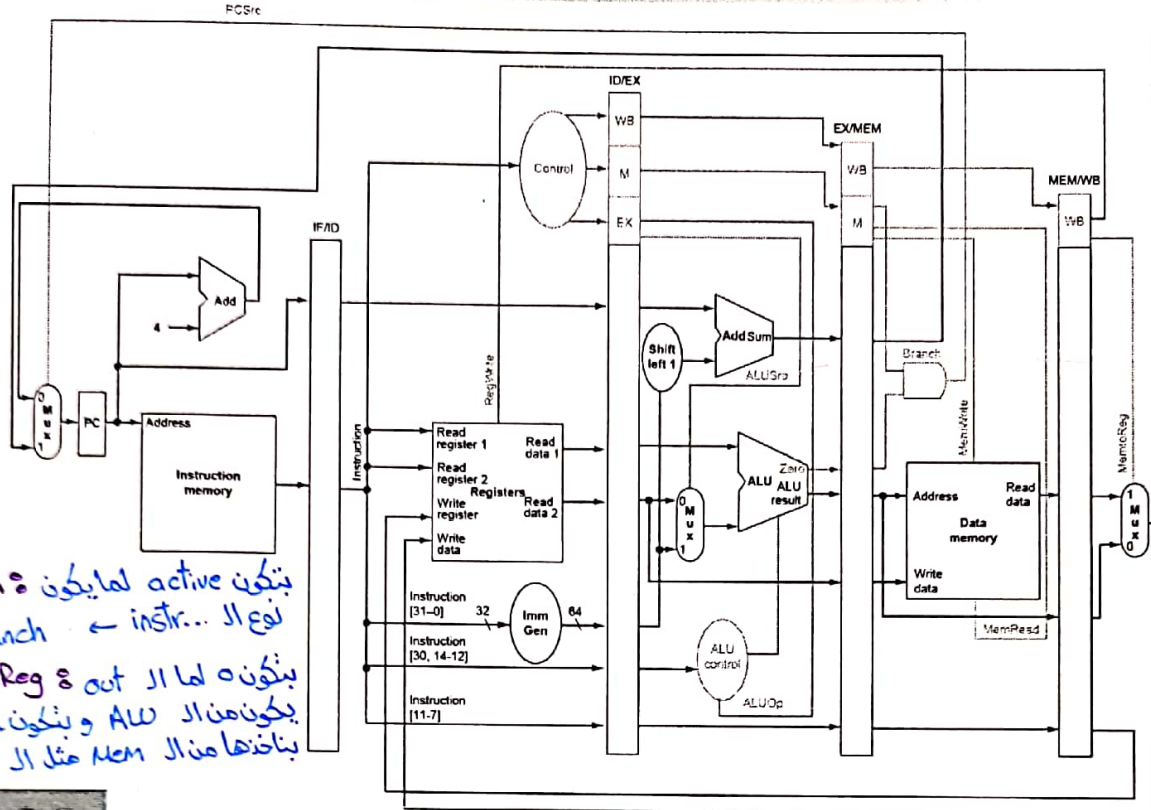
# Pipelined Control

\* Rewrite

يكون 1 عشان بنينا نكتب

register و لما بنينا نكتب مثل : ال Branch ما بنينا ال reg ومثل ال store كمان

load مثل ال add/sub



\* Branch : يكون active لما يكون نوع ال instr... ← branch  
 \* MemWrite : يكون لما ال out من ال ALU ويكون 1 لما ال data بنانها من ال Mem مثل ال load



\* ALUSrc : لو كانت 0 فال ALU بجربا ال data من ال register  
 لو كانت 1 يدخل ال imm instruction

\* ALUOp : عبارة عن 2 bits بتختار العملية اللي تنعملها ال ALU

## Hazards

\* Mem Write : store : يكون 1 لما بنينا نكتب ال Mem مثل ال instruction  
 \* Mem read : load : يكون 1 لما بنينا نقرأ من ال Mem مثل ال instruction

Situations that prevent starting the next instruction in the next cycle

\* الفائدة الأساسية من ال Pipelining : لتحسين ال performance

- 1 Structure hazards → في مشكلة بلا organisation ← توقف خط العمل
  - A required resource is busy → بنزيد عدد ميمنا التحسينات
- 2 Data hazard → تأخذ اللي بتقرأ ال data بينما ما تنظروا اللي بنكتب
  - Need to wait for previous instruction to complete its data read/write
- 3 Control hazard → العلاقة بـ jump وال Branch
  - Deciding on control action depends on previous instruction





# Contents

- 4.6 Pipelined Datapath and Control (Review)
- 4.7 Data Hazards: Forwarding versus Stalling
- 4.8 Control Hazards
- 4.9 Exceptions
- 4.10 Parallelism via Instructions
- 4.11 Real Stuff: The ARM Cortex-A53 and Intel Core i7 Pipelines
- 4.14 Fallacies and Pitfalls
- 4.15 Concluding Remarks



Chapter 4 — The Processor

# Contents

- 4.7 Data Hazards: Forwarding versus Stalling
  - Hardware Solutions [ Data Hazards in ALU Instructions
  - Load-Use Data Hazard
  - Software Solutions [ Code Scheduling

# Data Hazards in ALU Instructions

Consider this sequence: *register instructions* *كلم*  
*store* *ال*

sub x2, x1, x3

and x12, x2, x5

or x13, x6, x2

add x14, x2, x2

sd x15, 100(x2)

*read from register and* *store \**

*write to memory*

*read from memory and write* *load \**  
*to register*

*reg محتويات* *كـ* *تنقر من ال* *offset* *base register*

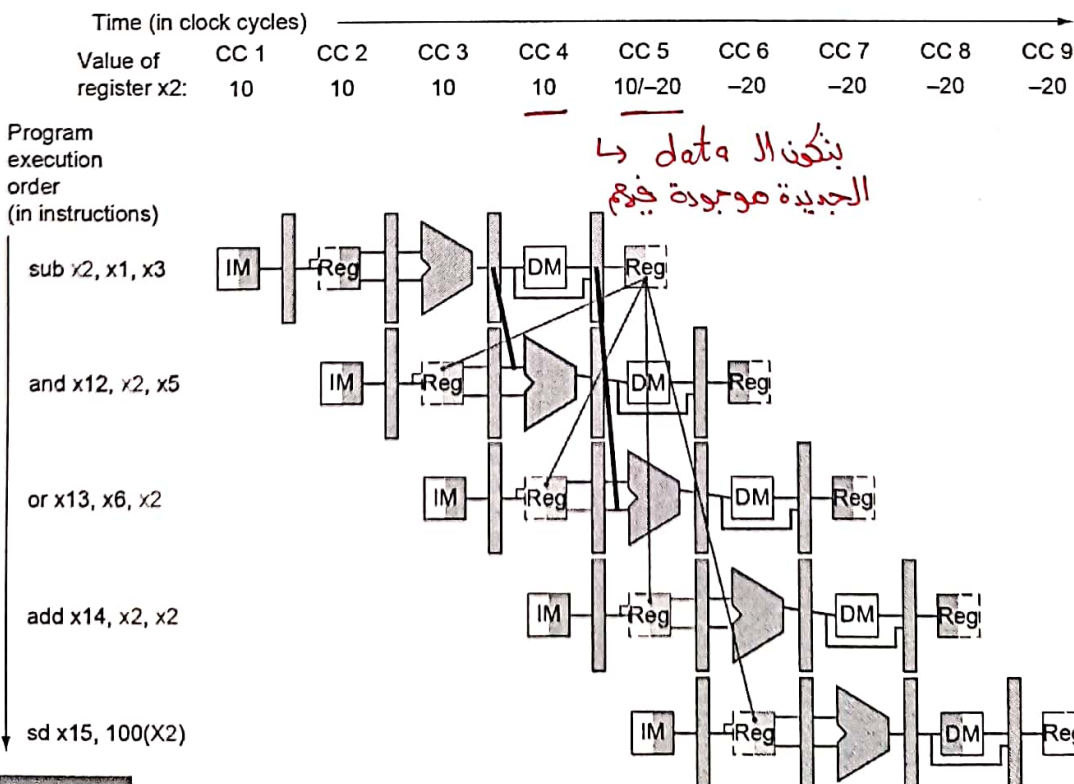
There are multiple true data dependencies read-after-write (RAW), on register x2.

We can resolve hazards with stalls or forwarding.

*x15*  
*وبينا*  
*نخزنوا*  
*بال Mem*  
*وال Mem*  
*مكانه =*  
*محتويات + 100*  
*x2*

*حلال Hazards بطريقتين* *1- Stalls*  
*2- forward*

## Dependencies & Forwarding

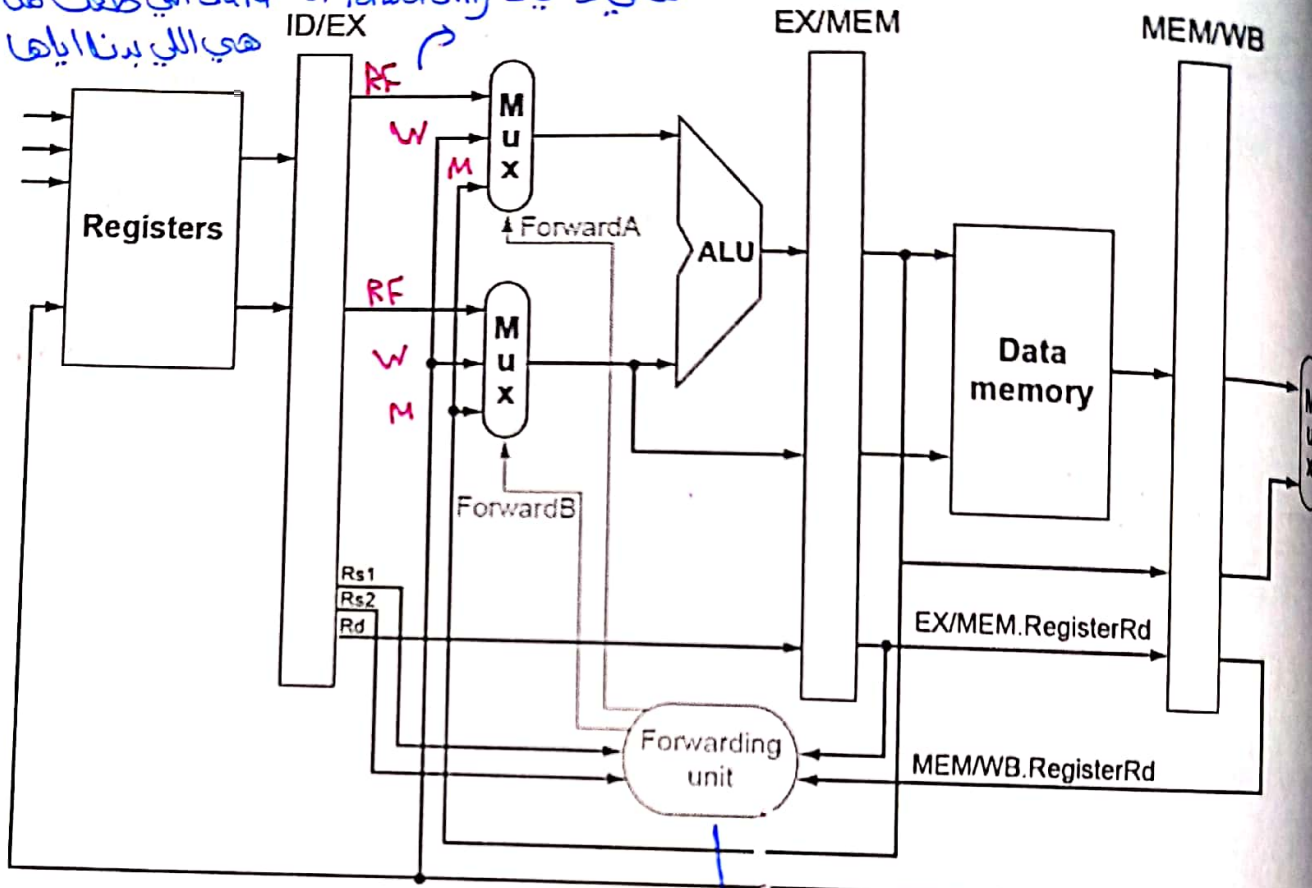




# Forwarding Paths

reg file

ما في داعي ل forwarding ال data اللي طلعت من ال هي اللي بدنا اياها



بها تكشف هل في data hazards  
 وهل في حاجة ل forwarding ولا  
 في حاجة بيه تعطي control signal  
 ملائمة ل Mux لحتى ناخذ ال data من المكان الصحيح



## Load-Use Data Hazard

- Can't always avoid stalls by forwarding
  - If value not computed when needed
  - Can't forward backward in time!

Program execution

200 400 600 800 1000 1200 1400

Slide 8 :

عدد ال cycles

	1	2	3	4	5	6	7	8	9	10	
I1	F	D	E	M	W						* تسمى : "Multicycles Diagram" لأنه احنا بنتابع بأكثر من cycle بنوع يخبر ال instructions
I2		F	D	E	M	W					
I3			F	D	E	M	W				
I4				F	D	E	M	W			
I5					F	D	E	M	W		

↳ 5 Instructions

\* لو مافي Stalls بنحتاج ال 9 cycles لتنفيذ ال instructions =  $5 + 4 = 9$

↳ مجموع عدد ال cycles الزايدة بكل instruction

↳ عدد ال cycles اللي احتجهم لتنفيذ I1

\* لو ما كان عندي pipeline كان احتجت  $25 = 5 \times 5$  cycles لتنفيذ ال instructions

Slide 11 : Assume no forwarding (except through the Register file) and hazards are solved by stalls

	1	2	3	4	5	6	7	8	9	10	
Sub $X_2, X_1, X_3$	F	D	E	M	W						* بنا 2 cycles 1 stall بين ما حصلنا القيمة الجيبية بنظربا F بين ما نطلب مرحلة ال D اللي قبلها
and $X_{12}, X_2, X_5$		F	D	<u>D</u>	<u>D</u>	E	M	W			
or $X_3, X_6, X_2$			F	<u>F</u>	<u>E</u>	D	E	M	W		
add $X_{14}, X_2, X_7$						F	D	E	M	W	
sd $X_5, \text{lo}(X_2)$							F	D	E	M	

\* بنا 11 stalls احتجت 11 cycles =  $4 + 10 + \text{stalls}$

كل instruction بيها 1 cycle لتصل ال W / عدد ال cycles الزيادة

( ثابت ) 4 cycles لاحتق نغبي ال pipeline

\* هون في 2 stalls بس لأنه ال stall اللي عند ال F لسببه نفس سبب ال Stall اللي  
دار عند ال D

\* هاد الحل مو كويس لأنه نزل من ال performance



Slide 12 : forwarding solution :

بنهاية ال E رح يكون الناتج الجديد موجود

	1	2	3	4	5	6	7	8	9	10
sub $X_2, X_1, X_3$	F	D	E	M	W					
and $X_2, X_2, X_5$		F	D	E	M	W				
or $X_3, X_6, X_2$			F	D	E	M	W			
add $X_4, X_2, X_2$				F	D	E	M	W		
sd $X_5, 100(X_2)$					F	D	E	M	W	

\* خلاصه 9 cycles \*

## Homework 1

A multicore process has a clock rate of 3.0 GHz and voltage of 1.0 V. Assume that, on average, it consumes 75 W of static power and 30 W of dynamic power when all its six cores are active. It executes a program consisting of  $10^{10}$  instructions on a single core with the following instruction mix.

Instruction type	CPI	Frequency
Arithmetic	1.0	50%
Load/store → أبطأ استبي	2.0	40%
Branch instructions	1.5	10%

(a) Find the total execution time for this program on one core?

$$\text{CPU time} = IC \times \sum(\text{CPI} \times \text{frequency}) \div \text{clock rate} = 10^{10} \times (1.0 \times 0.5 + 2.0 \times 0.4 + 1.5 \times 0.1) \div 3 \times 10^9 = 4.833 \text{ Second}$$

(b) When this program is parallelized and run on the six cores, the number of Arithmetic and Load/store instructions per core is divided by 4, but the number of branches remains the same? What is the execution time of the parallel program.

$$10^{10} \times (1.0 \times 0.5 \div 4 + 2.0 \times 0.4 \div 4 + 1.5 \times 0.1) \div 3 \times 10^9 = 1.583 \text{ Sec}$$

(c) What is the speedup of the parallel program?  $\text{Speedup} = 4.833 \div 1.583 = 3.05$  ممكن يوميل أحسنم  
من هيك يعني  
6 مثلاً

(d) Assume that the power consumed is proportional to the number of active cores. What are the energies consumed by the serial program on the single core and the parallel program on the six cores?  $\text{Energy of the serial program} = \text{exec. time} \times \text{power} = 4.833 \times (75 + 30) \div 6 = 85 \text{ Joules}$

$\text{Energy of the parallel program} = \text{exec. time} \times \text{power} = 1.583 \times (75 + 30) = 166 \text{ Joules}$

(e) What is the average capacitive load of this processor when the six cores are active?

$$\text{Power} = \frac{1}{2} \times \text{capacitive load} \times \text{voltage}^2 \times \text{frequency}$$

$$30 = \frac{1}{2} \times \text{capacitive load} \times 1^2 \times 3 \times 10^9 \rightarrow \text{capacitive load} = 20 \text{ nano farads}$$

(f) What is the average current drawn by this processor when the six cores are active?

$$\text{Power} = \text{current} \times \text{voltage} \Rightarrow 75 + 30 = \text{current} \times 1 \rightarrow \text{current} = 105 \text{ Ampere}$$

كبير يعتبر →  
لم يدخلوها عن طريق عشرات ال

لو هسي higher performance رج استهلاك طاقة أكبر \* (c)



# Load-Use Hazard Detection

- Check when using instruction is decoded in ID stage
- ALU operand register numbers in ID stage are given by
  - IF/ID.RegisterRs1, IF/ID.RegisterRs2
- Load-use hazard when *شروط الاستخدام* *load-use*
  - ID/EX.MemRead and *5 bits*  
*5 bits*  $((\text{ID/EX.RegisterRd} = \text{IF/ID.RegisterRs1}) \text{ or } (\text{ID/EX.RegisterRd} = \text{IF/ID.RegisterRs2}))$
- If detected, stall and insert bubble



*D. Rs1 = E. Rd*  
or  
*D. Rs2 = E. Rd*  
*E. MemRead*

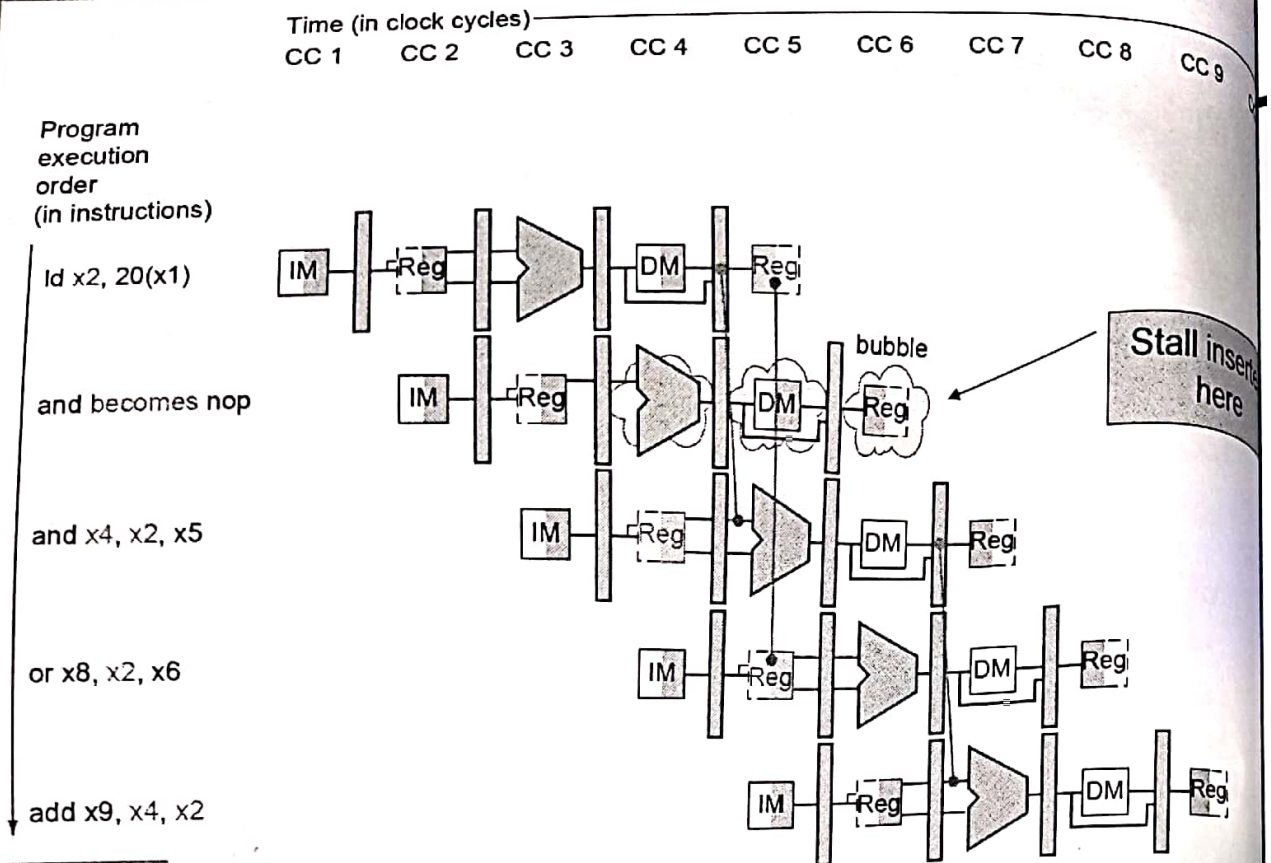


Chapter 4 — The Processor — 15

## How to Stall the Pipeline

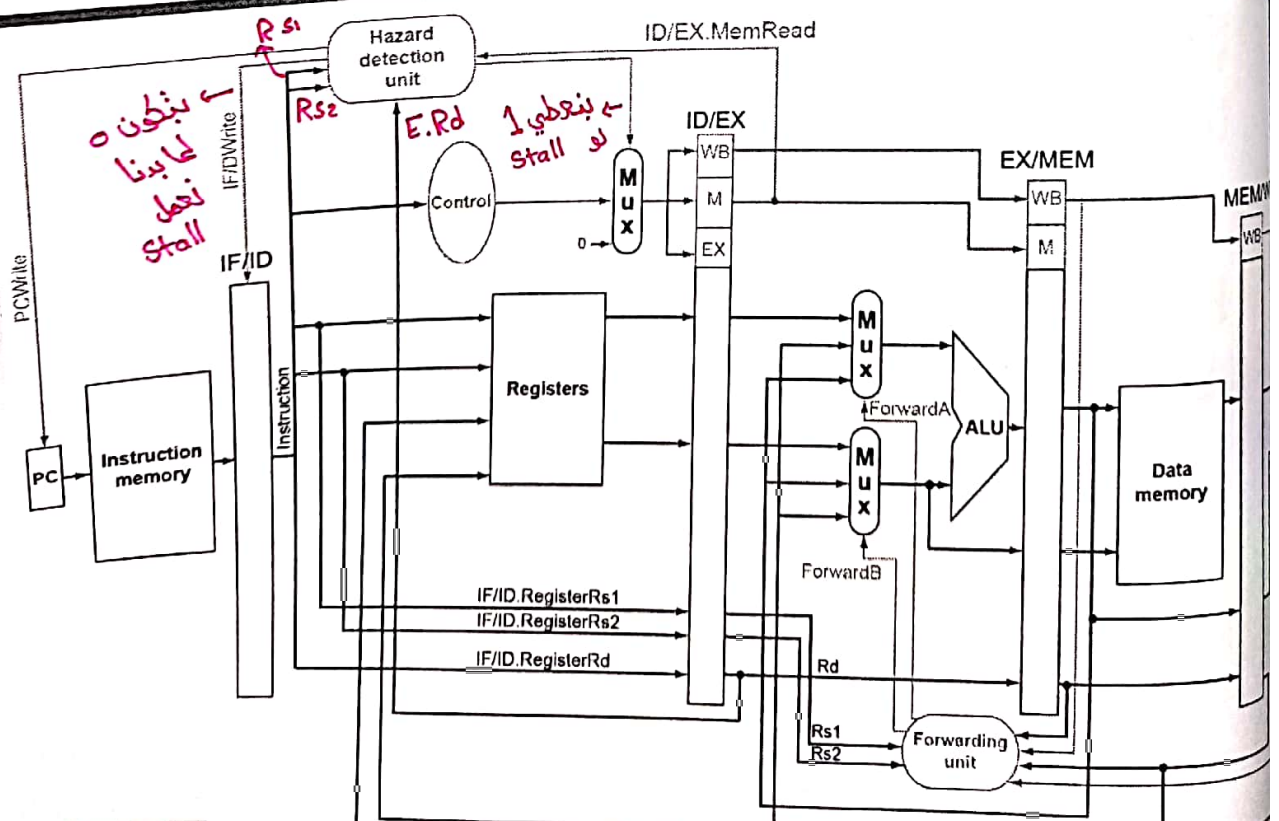
- Force control values in ID/EX register to 0
  - EX, MEM and WB do nop (no-operation)
- Prevent update of PC and IF/ID register
  - Using instruction is decoded again
  - Following instruction is fetched again
  - 1-cycle stall allows MEM to read data for 1d
    - Can subsequently forward to EX stage

# Load-Use Data Hazard



Chapter 4 — The Processor

# Datapath with Hazard Detection





# Stalls and Performance

## The BIG Picture

- Stalls reduce performance
  - But are required to get correct results «*بئجناجه بحالات*»
- Compiler can arrange code to avoid hazards and stalls
  - Requires knowledge of the pipeline structure

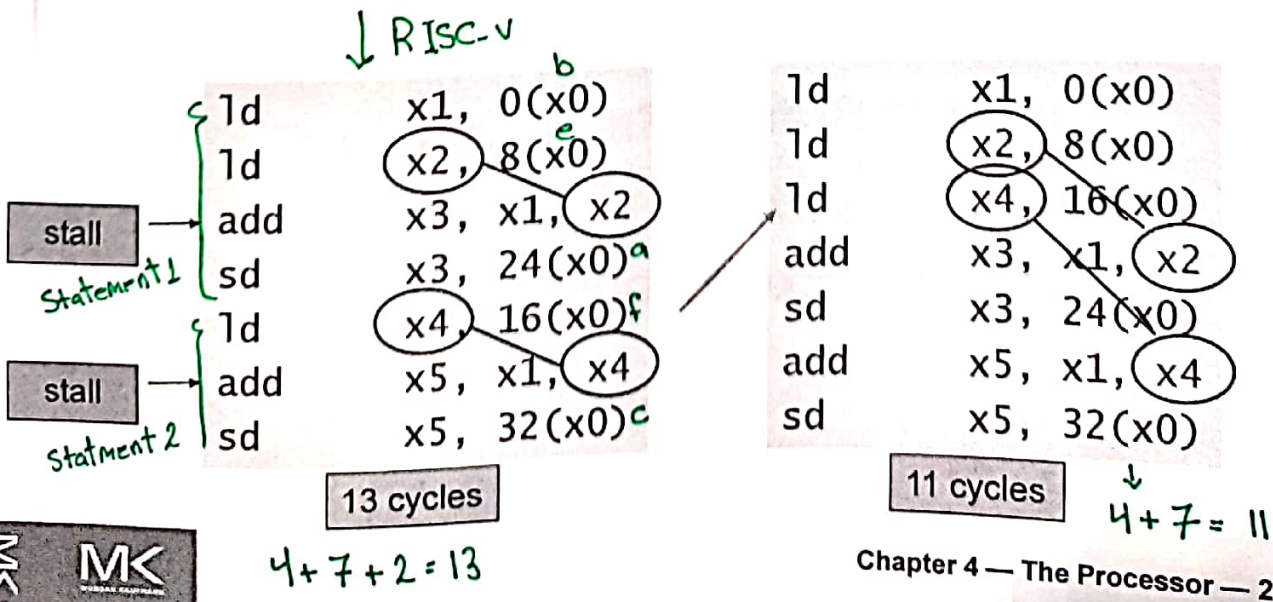


→ slide 16  
Chapter 4 — The Processor — 19

\* شو بئجناجه نفل باء Stall cycle  
 -1 Signal ترح لا نفل Pc disable  
 -2 disable IR → Register after Pc  
 -3 No-operation (o) to Execution

## Code Scheduling to Avoid Stalls

- Reorder code to avoid use of load result in the next instruction
- C code for  $a = b + e; c = b + f;$



# Contents

- 4.6 Pipelined Datapath and Control (Review)
- 4.7 Data Hazards: Forwarding versus Stalling
- 4.8 Control Hazards
- 4.9 Exceptions
- 4.10 Parallelism via Instructions
- 4.11 Real Stuff: The ARM Cortex-A53 and Intel Core i7 Pipelines
- 4.14 Fallacies and Pitfalls
- 4.15 Concluding Remarks



# Contents

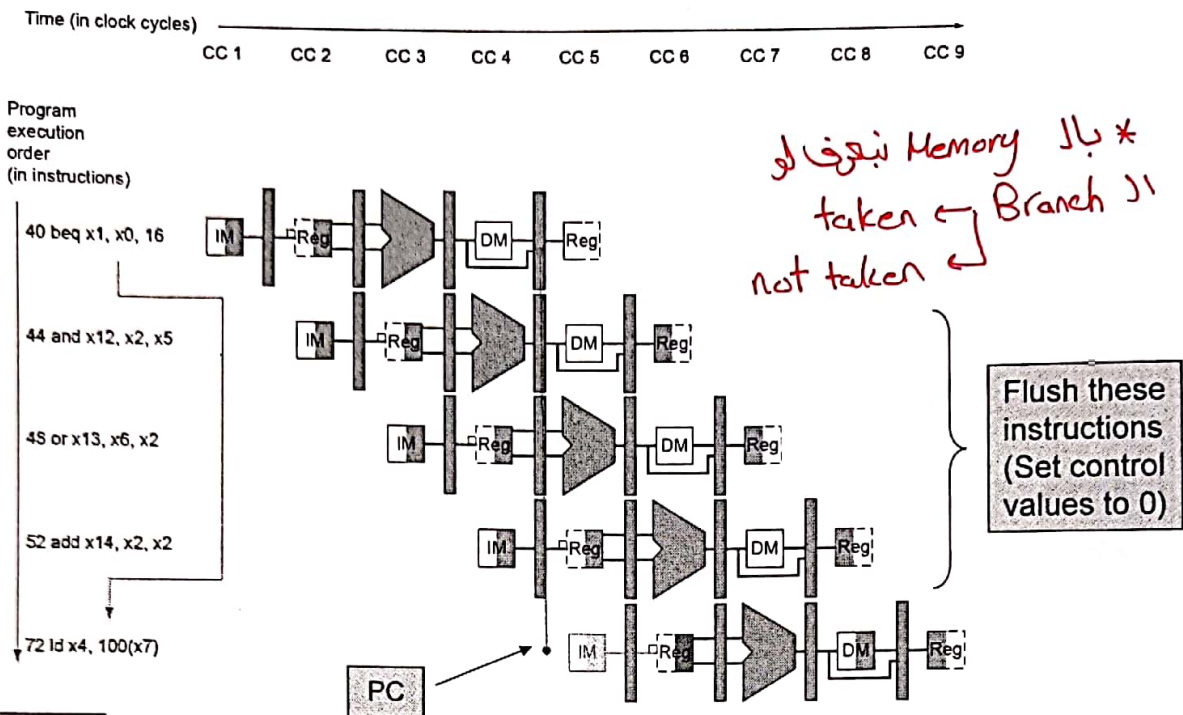
- 4.8 Control Hazards
  - Branch Hazards
  - Reducing Branch Delay
  - Branch Prediction
  - Dynamic Branch Prediction
  - Calculating Branch Target
  - Imprecise Exceptions

لعملوا delay و Stalls في كيف  
تقللوا عن طريق



# Branch Hazards

If branch outcome determined in MEM




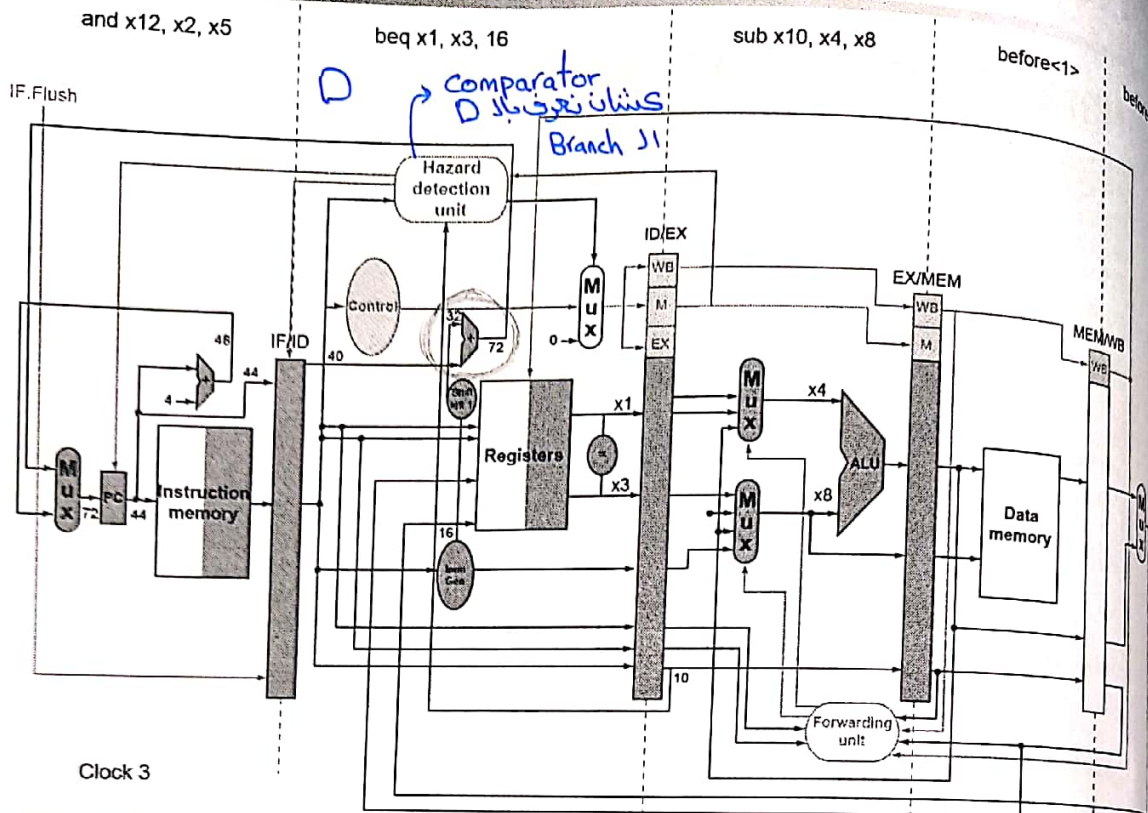
## Reducing Branch Delay

- Move hardware to determine outcome to ID stage
  - Target address adder
  - Register comparator
- Example: branch taken
 

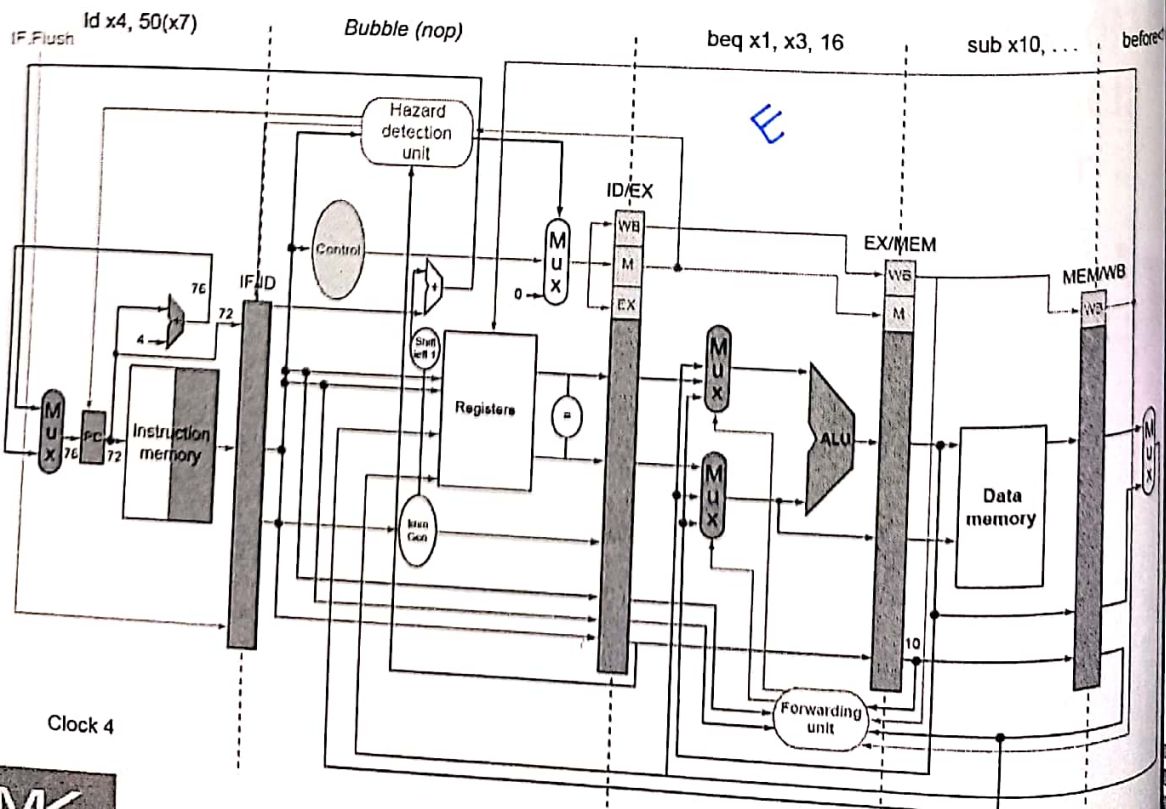
```

36: sub x10, x4, x8
40: beq x1, x3, 16 // PC-relative branch
                        // to 40+16*2=72
44: and x12, x2, x5
48: orr x13, x2, x6
52: add x14, x4, x2
56: sub x15, x6, x7
...
72: ld x4, 50(x7)
            
```

# Example: Branch Taken



# Example: Branch Taken







# Dynamic Branch Prediction

- In deeper and superscalar pipelines, branch penalty is more significant
- Use dynamic prediction
  - Branch prediction buffer (aka branch history table)
  - Indexed by recent branch instruction addresses
  - Stores outcome (taken/not taken)
  - To execute a branch
    - Check table, expect the same outcome
    - Start fetching from fall-through or target
    - If wrong, flush pipeline and flip prediction



## Branch History Table (BHT)

NT: 0  
T: 1

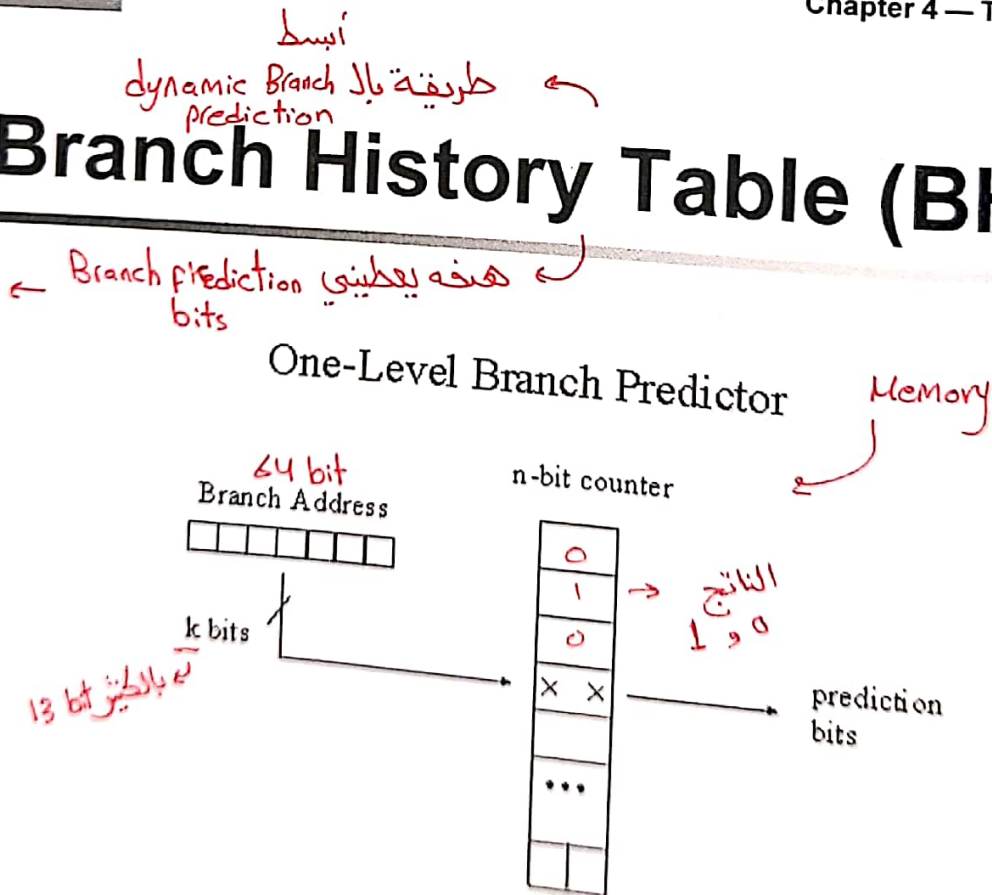


Table size =  $n \times 2^k$  bits


كمثال  $2 \times 2^{10} = 2^k$  bits





# 1-Bit Predictor: Shortcoming

- Inner loop branches mispredicted twice!

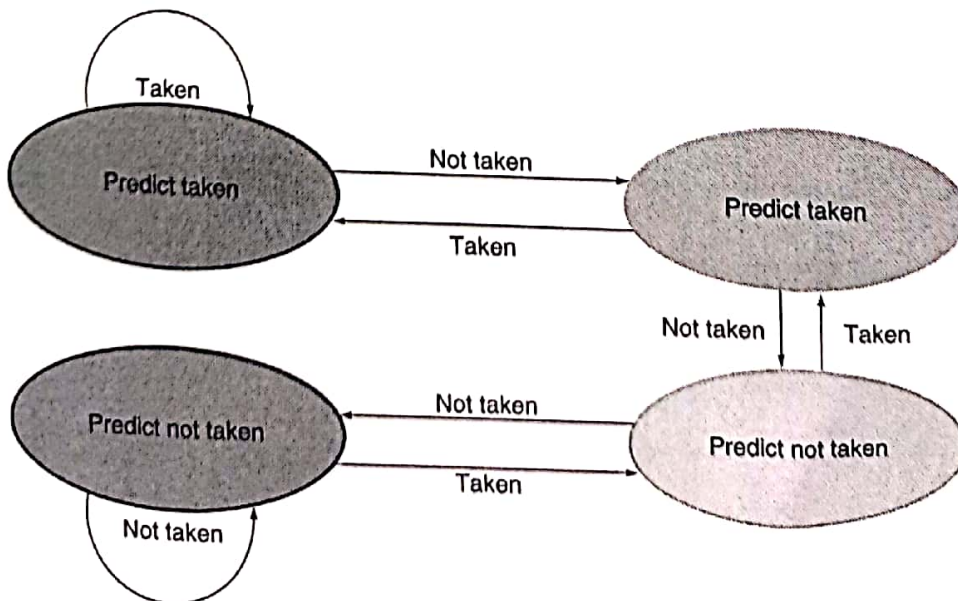


- Mispredict as taken on last iteration of inner loop
- Then mispredict as not taken on first iteration of inner loop next time around

MK

## 2-Bit Predictor

- Only change prediction on two successive mispredictions



MK

# Calculating the Branch Target

- Even with predictor, still need to calculate the target address
  - 1-cycle penalty for a taken branch
- Branch target buffer
  - Cache of target addresses
  - Indexed by PC when instruction fetched
    - If hit and instruction is branch predicted taken, can fetch target immediately

كشأن بالمستقبل لو مرت علي نفس ال branch وكان ال address تابعاً نفس  
الذي مخزن عندي فيودينياً ال BTA عن طريق ال MUX

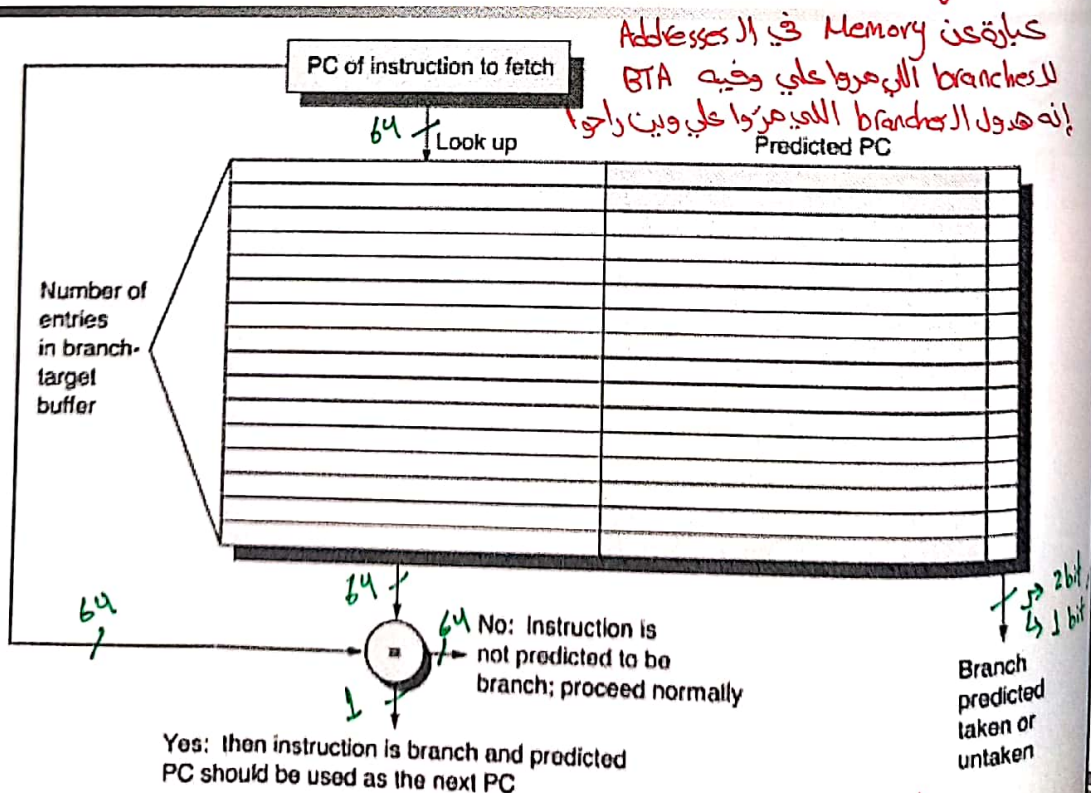


Chapter 4 — The Processor

كما يمر علي Branch بخزن ال address تبعها ولوين بشرح

من ال PC في Address هاد ال Address يستخدمه لنوجد ال BTB

## Branch Target Buffer (BTB)



له ممكن يخزنوا معاه ال BHT بي يعني الأوقات ما يكون موصولين  
سوا غا بالامفصولين





# Contents

---

- 4.6 Pipelined Datapath and Control (Review)
- 4.7 Data Hazards: Forwarding versus Stalling
- 4.8 Control Hazards
- 4.9 Exceptions → *بخلينا نوقف تشغيل البرنامج ونروح نحلله*
- 4.10 Parallelism via Instructions
- 4.11 Real Stuff: The ARM Cortex-A53 and Intel Core i7 Pipelines
- 4.14 Fallacies and Pitfalls
- 4.15 Concluding Remarks



# Contents

---

- 4.9 Exceptions
  - Exceptions and Interrupts
  - Handling Exceptions
  - Exceptions in a Pipeline
  - Exception Example
  - Multiple Exceptions

# Exceptions and Interrupts

- “Unexpected” events requiring change in flow of control
  - Different ISAs use the terms differently
- 1) ■ Exception
  - Arises within the CPU *يعني داخل ال pipeline وهو بيتنجز ممكن يبيس error يحتاج معالجة*
    - e.g., undefined opcode, syscall, ... *و division by zero*
- Interrupt
  - From an external I/O controller *بيجونا من الخارج*
- Dealing with them without sacrificing performance is hard



## Handling Exceptions

- Save PC of offending (or interrupted) instruction *ال instruction التي عملت Exception*
  - In RISC-V: Supervisor Exception Program Counter (SEPC)
- Save indication of the problem
  - In RISC-V: Supervisor Exception Cause Register (SCAUSE)
  - 64 bits, but most bits unused
  - Exception code field: 2 for undefined opcode, 12 for hardware malfunction, ... *هنا كلهم مستخدمات*
- Jump to handler
  - Assume at 0000 0000 1C09 0000<sub>hex</sub>



# An Alternate Mechanism

- Vectored Interrupts
  - Handler address determined by the cause
- Exception vector address to be added to a vector table base register:
  - Undefined opcode      00 0100 0000<sub>two</sub>
  - Hardware malfunction: 01 1000 0000<sub>two</sub>
  - ....:                      ...
- Instructions either
  - Deal with the interrupt, or
  - Jump to real handler

MK

Chapter 4 — The Processor — 39

## Handler Actions

- Read cause, and transfer to relevant handler
- Determine action required
- If restartable
  - Take corrective action
  - use SEPC to return to program
- Otherwise
  - Terminate program
  - Report error using SEPC, SCAUSE, ...


MK

Slide 14.8 Load-use data hazard

	1	2	3	4	5	6	7	8	9	10	ثابت ↑ → 4+2+1
ld $x_1, 0(x_2)$	F	D	E	M	W						↑ Instruction ↑ stalls
sub $x_4, x_1, x_5$		F	D	D	E	M	W				= 7 cycles

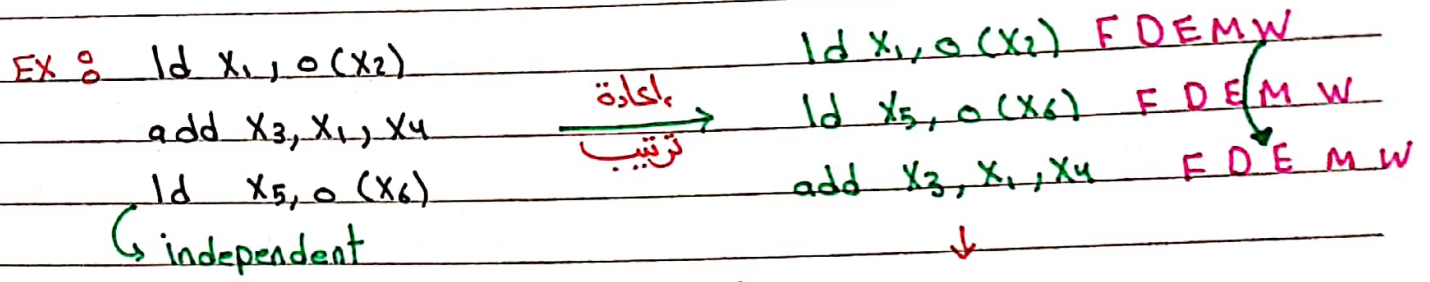
\* ال Sub وال add بنجزو ال data بلا Execute stage  
 \* ال load بنجزو ال data بلا Memory بنواينها  
 ال stall كلاس ال ستان ال data بلا load  
 بنجزو نواية ال Memory بعين كلاس  
 Forwarding ستان نحد القيم الجديدة .

\* ال Load-use ال يحتاج نحد stall و Forwarding  
 \* الدائرة اللي بنحتاجها لتنفيذ  
 ال  $D.RS_2 = E.Rd$  او ال  $D.RS_1 = E.Rd$   
 بنطوي 1 او  $A = B$





Slide 20: Code Scheduling to Avoid Stalls:



هون بهيا ترتيب ما يحتاج اول  
Stall بس Forwarding

Slide 23: Branch Hazards

Solving branches in the Memory stage

\* بنحلوم ل null instruction

لانه ما بدنا نينفذوا

		1	2	3	4	5	6	7	8	9	10	
40	beq X1, X0, 16	F	D	E	M	W						* بال Memory بنعرف اذا ال Branch not taken او taken
44	and X12, X2, X5		F	D	E	n	n					بنغايتمها
48	or X13, X6, X2			F	D	n	n	n				* افترضنا انه
52	add X4, X2, X2				F	n	n	n	n			ال Branch is taken
72	ld X4, 100(X7)					F	D	E	M	W		

target address = 40 + 2 \* 16 = 72

لانه بنواليا ال cycle 4 وبداية cycle 5 سوف

لانه half bytes لازم نضرب في 2

انه ال Branch is taken في ال F

\* هون ال performance يتقل لانه كانه عندي 3 stall cycles

Branch delay = 3

\* و ال تقليل ال Branch delay ← هناك حل بانو بيد ال Branch ال يجوز بال Memory stage

بغلا يجوز بال Decode stage

Slide 24 8

Solving branches in the Decode stage → Branch نعرف نتيجة ال

		1	2	3	4	5	6	7	8	9	10	
40	beq X <sub>1</sub> , X <sub>0</sub> , 16	F	D	E	M	W						* افترضنا انه Branch is taken
44	and X <sub>12</sub> , X <sub>2</sub> , X <sub>5</sub>		F	n	n	n	n					
48	or X <sub>13</sub> , X <sub>6</sub> , X <sub>2</sub>											
52	add X <sub>14</sub> , X <sub>2</sub> , X <sub>2</sub>											
72	ld X <sub>4</sub> , 100(X <sub>7</sub> )			F	D	E	M	W				

Branch delay =  $\frac{1}{1}$  \*

↙ Branch prediction ← التحسين الجيد هو ان تقوم بعمل  
بخلي ال Hardware يتنبأ بالنتيجة

Slide 28 8 Predict Not taken

Solving branches in the Decode stage

① Assume branch is not taken

		1	2	3	4	5	6	7	8	9	10	
	beq X <sub>1</sub> , X <sub>0</sub> , L	F	D	E	M	W						بغيا نوقفنا وكان نوقفنا صح وما ضيعنا اي cycle
	I <sub>2</sub>		F	D	E	M	W					→ Branch Delay
	L I <sub>T</sub>											= 0

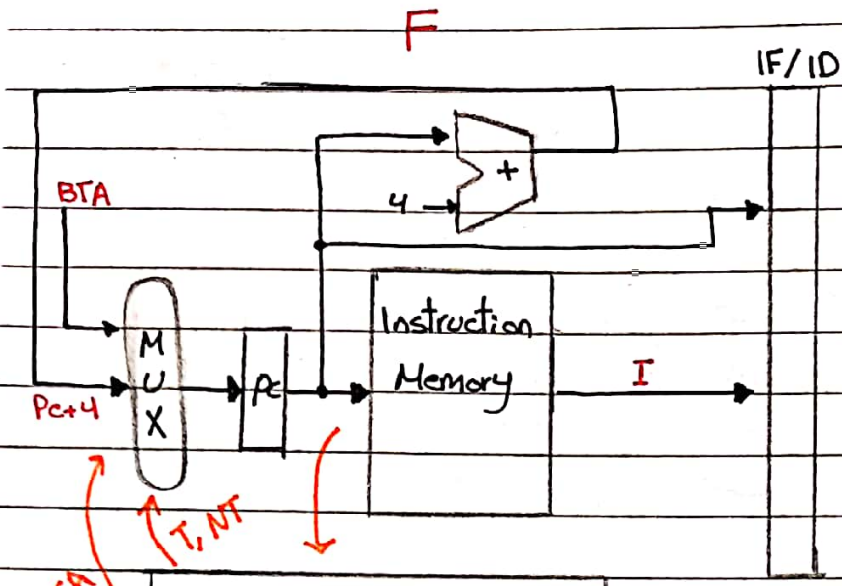
② Assume branch is taken

avg Branch Delay =  $\frac{1}{2}$  cycle

		1	2	3	4	5	6	7	8	9	10	
	beq X <sub>1</sub> , X <sub>0</sub> , L	F	D	E	M	W						→ Branch Delay = 1
	I <sub>2</sub>		F	n	n	n	n					
	L I <sub>T</sub>			F	D	E	M	W				




Slide 28 Branch Prediction Unit



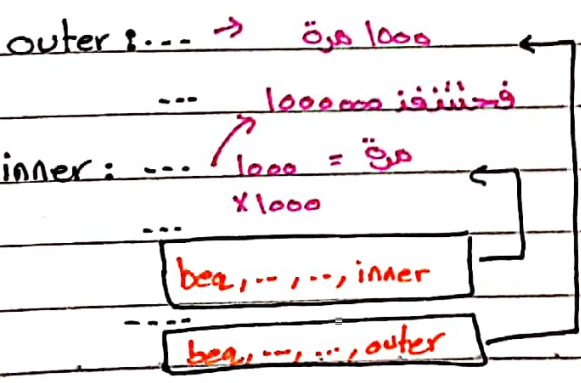
Branch prediction unit → Address للتقسيم وشغلته ← من ال instruction التي قاسمين بنجبوا بيتنا اذا taken or not taken ← Branch ال وأيها مسؤولة عن اخطاء ال BTA

Slide 30 BHT 1-bit predictor



لو افترضنا ال Prediction taken وطقنا فاس taken بنظنا اذا و في اي حالات الباقية

Assume 1000 iterations for the inner loop and for the outer loop



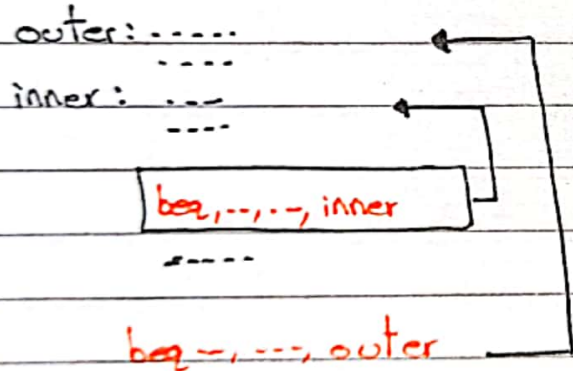
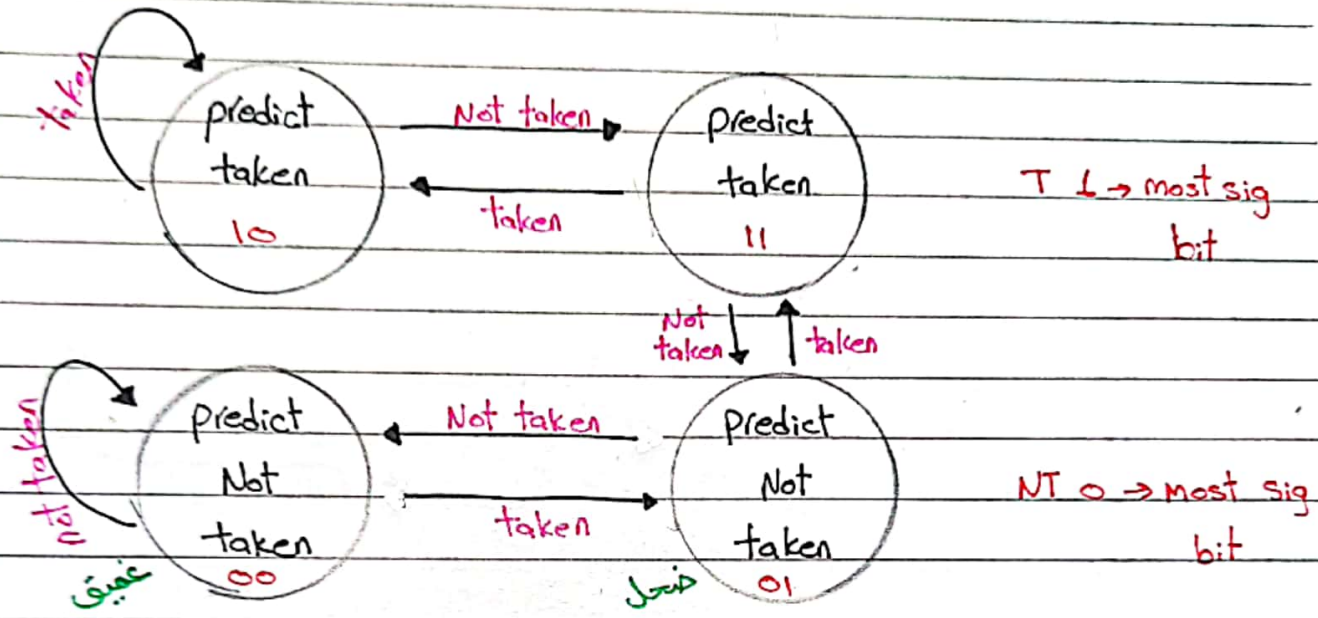
Iteration	997	998	999	0	1
Prediction Result	T/C	T/C	T/I	NT/I	T/C

\* T: taken / NT: Not taken  
C: correct prediction / I: incorrect prediction

\* ال one bit بتغير رأيه بس مرة فانتجونا اظ ال 2 bit بتغير رأيه بطريقة أبطأ.

Slide 32

2-bit predictor



Iteration	997	998	999	0	1
Prediction result	T/c	T/c	T/l	T/c	T/c

1 miss Prediction



# Slide 38 & Handling Exceptions

\* ميزة هاي الطريقة انما Simple hardware

سبب Software أسرع

11 \* لتفرض حمار Exception في 13


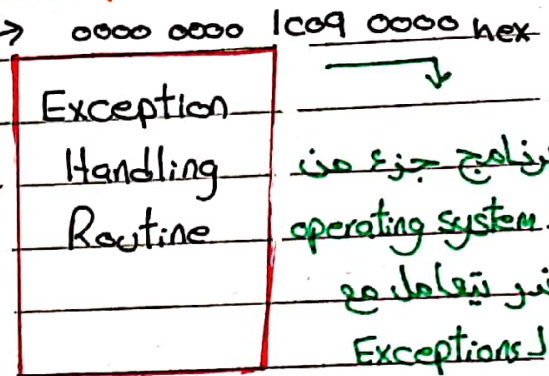
12

13 \* المفروض ما نكمل 14 و 15

14 بل نروح على Address الي

15 موجود عنده ال Exception handling

وبال RISC-V هاد ال Address المخصص



Additional Register لازم نضيف

اللي هم SCAUSE, SEPC

هدول بس يميز في Exception لازم نخزن ال Address ال Instruction الي صار عليها ال Exception وبنفس الوقت بي اخزن سبب ونوع ال Exception

\* Exception Handling Routine : رح يستفيد من سبب ال Exception الي خزناه بال SCAUSE ويعالجه ال Exception ويعين من ال SEPC بيقرر يرجع ال Instruction الي توقفنا عندها ويتابع تنفيذ البرنامج التالي

\* ال Exception Handling Routine: مش دايم بقدر يرجع ال Instruction الي توقفنا عندها

انما قدر يحد المشكلة وبنقدر نكمل البرنامج بيرجع، وبعين المرات بوجح مكان تافي بال

operating system لو كان ما بقدر يحد المشكلة مثل مشكلة: illegal code

وساعتها ال OS بطلع ال user رسالة بتقله انو عندك مشكلة ما بنقدر نكمل برنامجك

## Slide 40 :


11 undefined opcode 00 0100 0000-two

12 Hardware malfunction: 01 1000 0000-two

13 \* هاي الطريقة حسب نوع ال Exception بنروح ال Address مختلف

14 \* ميزة هاي الطريقة بتخليه أسرع

15 \* ال Software أسول



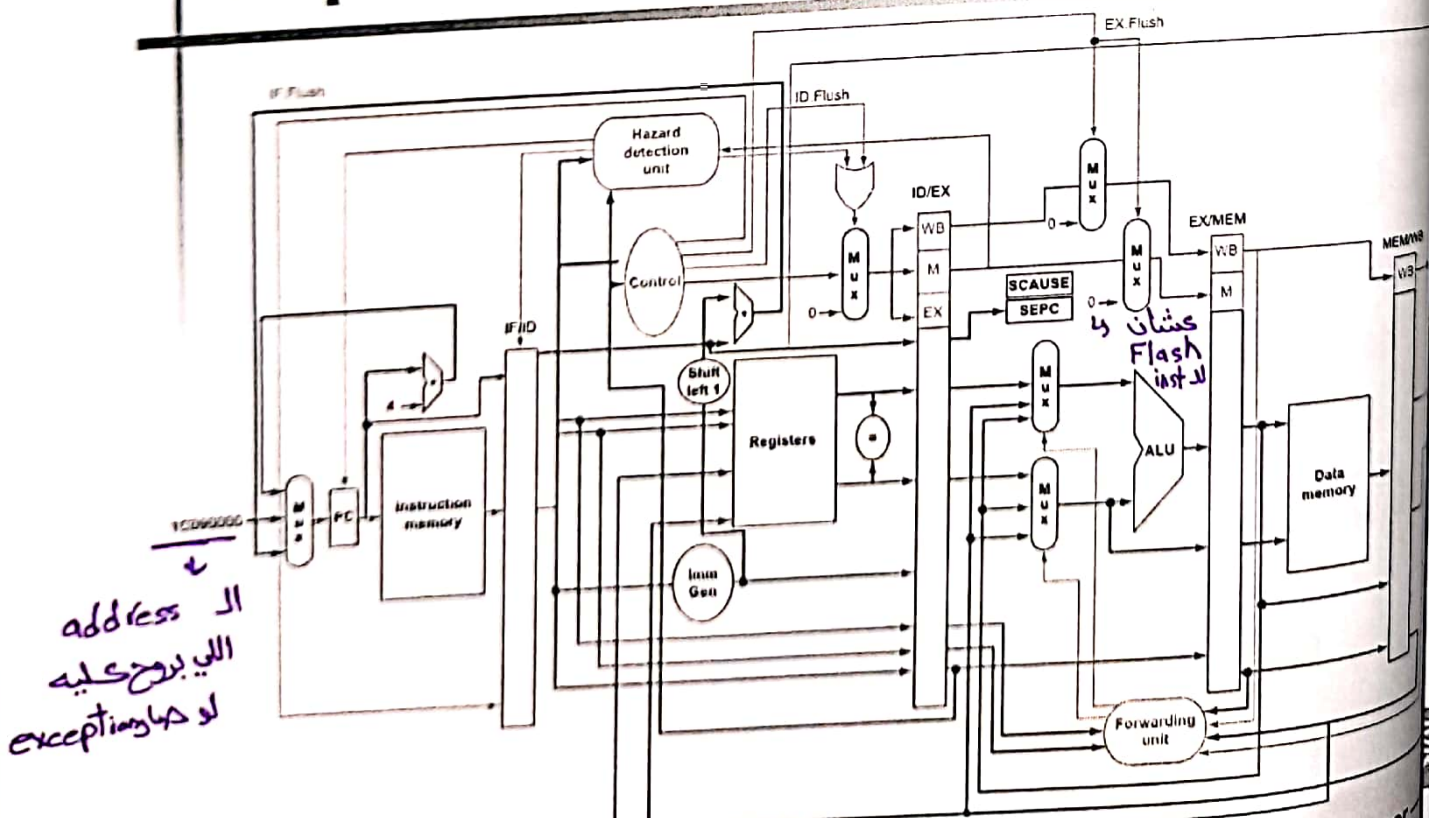
# Exceptions in a Pipeline

التعامل مع  
الاستثناءات

- Another form of control hazard
- Consider malfunction on add in EX stage
  - add x1, x2, x1
  - Prevent x1 from being clobbered
  - Complete previous instructions
  - Flush add and subsequent instructions
  - Set SEPC and SCAUSE register values
  - Transfer control to handler
- Similar to mispredicted branch
  - Use much of the same hardware



## Pipeline with Exceptions





# Exception Properties

- Restartable exceptions → يعني بنقدر نرجع للinstruction التي مار عليها  
 Exception بعد ما نحلها
- Pipeline can flush the instruction
- Handler executes, then returns to the instruction  
 ↓  
 كسوم ال  
 unrestartable
- Refetched and executed from scratch
- PC saved in SEPC register
- Identifies causing instruction  
 ↓  
 ال instruction  
 ال ديمار عندها exception



## Exception Example

### Exception on add in

40	sub	x11, x2, x4	} هسهول لازم يشغفوا ويخلصوا
44	and	x12, x2, x5	
48	orr	x13, x2, x6	
4c	add	x1, x2, x1	← "هوه ديمار فيه "exception"
50	sub	x15, x6, x7	} هسهول بنعملهم Flash
54	ld	x16, 100(x7)	

...

### Handler

1c090000	sd	x26, 1000(x10)	} بالنهاية رج يعيل ld x27 ld x26
1c090004	sd	x27, 1008(x10)	

بالعادة ال Exception handling subroutine بيلشوا ب sd : لإنه حتى يشغف بده reg

يقرا ويكتب مفهوم وينفس الوقت مابده يخرب ال data اللي بالبرنامج الأظلي وما ينفها

Save معلوم







# Multiple Exceptions

- Pipelining overlaps multiple instructions
  - Could have multiple exceptions at once
- Simple approach: deal with exception from earliest instruction
  - Flush subsequent instructions
  - "Precise" exceptions → زمان ما كانوا بتقروا يعملوا هيلع
- In complex pipelines
  - Multiple instructions issued per cycle
  - Out-of-order completion
  - Maintaining precise exceptions is difficult!

Chapter 4 — The Processor — 47

## Imprecise Exceptions

- Just stop pipeline and save state
    - Including exception cause(s)
  - Let the handler work out
    - Which instruction(s) had exceptions
    - Which to complete or flush
      - May require "manual" completion
  - Simplifies hardware, but more complex handler software
  - Not feasible for complex multiple-issue out-of-order pipelines
- Processors كانت لما يحسبوا exception ما بتعرف مين خلص ومين ما خلص وهو عملية معقدة و بعض المرات لو خلص exception ما بتقدر ترجع للبرنامج الأصلي لا زه ما بتعرف لو كل ال instruction اللي قبل ال exception خلصوا ولا لا و لا اللي بعدها

Chapter 4 — The Processor — 48

# Contents

- 4.6 Pipelined Datapath and Control (Review)
- 4.7 Data Hazards: Forwarding versus Stalling
- 4.8 Control Hazards
- 4.9 Exceptions
- 4.10 Parallelism via Instructions
- 4.11 Real Stuff: The ARM Cortex-A53 and Intel Core i7 Pipelines
- 4.14 Fallacies and Pitfalls
- 4.15 Concluding Remarks

Multiple instructions every cycle وهذا مهم بال modern Processors



# Contents

- 4.10 Parallelism via Instructions
  - Instruction-Level Parallelism (ILP)
  - Multiple Issue
    - Static Multiple Issue
    - VLIW
    - Scheduling Static Multiple Issue
    - Loop Unrolling
    - Dynamic Multiple Issue
  - Register Renaming
  - Speculation
  - Why Do Dynamic Scheduling

الجهاز يتنفيذ أوامر كثيرة







# Static Multiple Issue

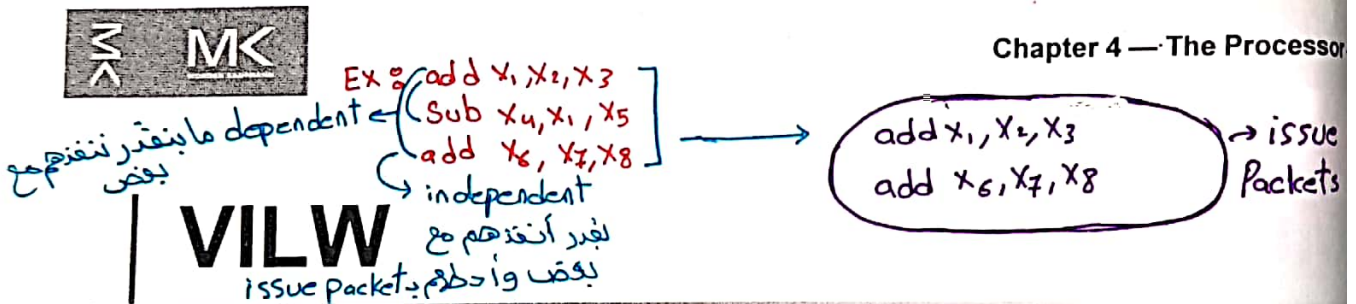
- Compiler groups instructions into "issue packets" → *issue packets → independent instructions* يجمع
- Group of instructions that can be issued on a single cycle
- Determined by pipeline resources required
- Think of an issue packet as a very long instruction
- Specifies multiple concurrent operations
- ⇒ Very Long Instruction Word (VLIW)

Static Multiple issue

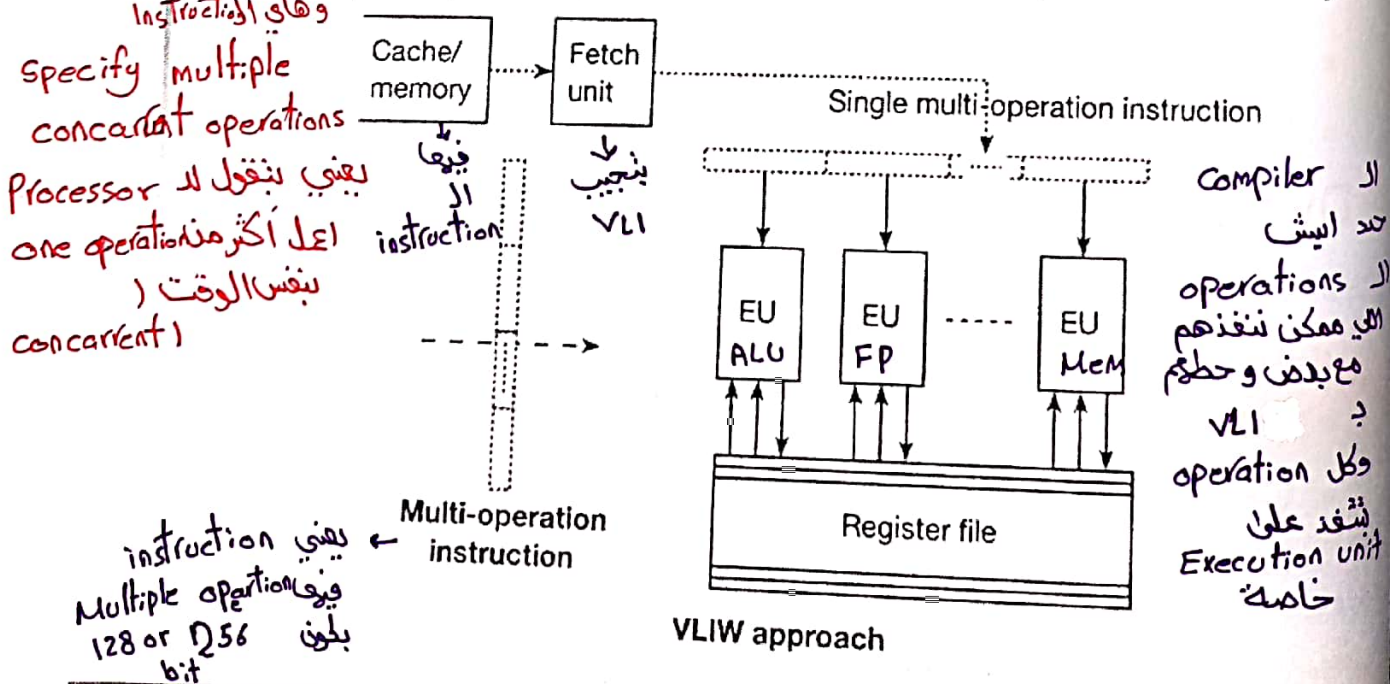
والله يشبه الـ

مثال Itanium

Chapter 4 — The Processor



يمكن تغييره نوع من الـ VLIW → Static Multiple issue الـ instruction الوحدة أكبر بكثير من 32 bit (very long instruction word, 1024 bits!) وهي الوحدتين





يمكن نلاحظه لحد الآن

# Scheduling Static Multiple Issue

Excellant performance ~~هو~~ good performance يعطى و يوفر كوربا و يعطى

Compiler must remove some/all hazards

- Reorder instructions into issue packets
- No dependencies within a packet
- Possibly some dependencies between packets

→ instruction Packet يكون فيه Packet وال Packet و

Varies between ISAs; compiler must know!

Pad with nop if necessary → independent instruction

كافية يجمعهم به Packet واحد يعنى باقي ال packet بـ no operations

يعني عبي الحالات الفاضية



# RISC-V with Static Dual Issue

Two-issue packets

- One ALU/branch instruction
- One load/store instruction (NOP)
- 64-bit aligned

يعني ال issue packet يكون فيه 2 instructions

وحدة يكون ALU/وحدة load/store branch

ولو ما قدر يلقى بضيف (NOP)

ALU/branch, then load/store

Pad an unused instruction with nop

صفت 4 زي ال

مقرض انه مافي stalls ولا في hazards

Address	Instruction type	Pipeline Stages						
		IF	ID	EX	MEM	WB		
n	ALU/branch	IF	ID	EX	MEM	WB		
n + 4	Load/store	IF	ID	EX	MEM	WB		
n + 8	ALU/branch		IF	ID	EX	MEM	WB	
n + 12	Load/store		IF	ID	EX	MEM	WB	
n + 16	ALU/branch			IF	ID	EX	MEM	WB
n + 20	Load/store			IF	ID	EX	MEM	WB

issue Packet  
issue Packet  
issue Packet

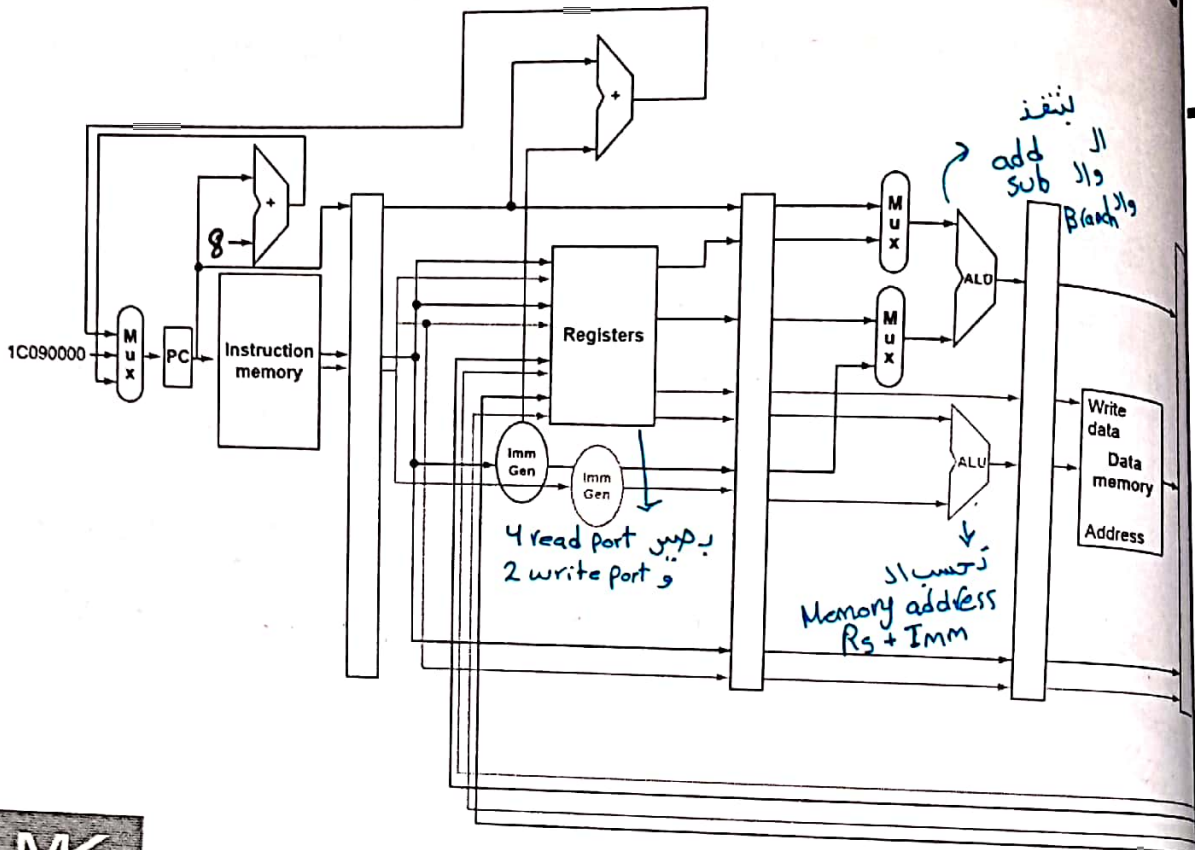
ثعبير ال throughput افضل لانه بتنفذ 2 في cycle

2 IPC → instruction per cycle

1/2 CPI → clock per instruction



# RISC-V with Static Dual Issue



## Hazards in the Dual-Issue RISC-V

- More instructions executing in parallel
- EX data hazard → بين ال issue وال issue اللي بعده ممكن يكون في dependent ونحتاج forwarding
- Forwarding avoided stalls with single-issue
- Now can't use ALU result in load/store in same packet
- Split into two packets, effectively a stall
- Load-use hazard
  - Still one cycle use latency, but now two instructions
- More aggressive scheduling required

← مايقدر نحلطهم مع بعض  
لانه ينفذوا على بعض

```
add x10, x0, x1
ld x2, 0(x10)
```





# Forwarding in Dual-Issue RISC-V

In addition to forwarding from M and W to E, there are additional forwarding paths among the two pipelines, e.g.:

- From W in memory pipeline to E in ALU pipeline

```
ld x31, 0(x20)
add x31, x31, x21
```

↓  
أخذه قبل

- From M in ALU pipeline to M in memory pipeline

```
add x31, x31, x21
sd x31, 0(x20)
```



## Scheduling Example

\* حصلنا على  $IPC = 1.25$  كشان ال nop  
فلو حطينا بدل ال nop instructions  
بزيد ال IPC فلذلك بنستخدم حد احسن والي هو  
loop unrolling

Schedule this for dual-issue RISC-V

```
Loop: ld x31, 0(x20) // x31=array element
      add x31, x31, x21 // add scalar in x21
      sd x31, 0(x20) // store result
      addi x20, x20, -8 // decrement pointer
      blt x22, x20, Loop // branch if x22 < x20
```

2 issue Packets

	ALU/branch	Load/store	cycle
Loop:	nop	ld x31, 0(x20)	1
	addi x20, x20, -8	nop	2
	add x31, x31, x21	nop	3
	blt x22, x20, Loop	sd x31, 8(x20)	4

حطيناها اول وحدة لانها اول عمليه  
ممكن نخط ب 1 برضه

ما قدرنا نخطه  
بالتاليه  
هذه  
load use hazard

IPC = 5/4 = 1.25 (c.f. peak IPC = 2)  
لاننا بنيجي بعد ال add  
بس ممكن احطها ب 3 برضه  
ال add قبله  
فلانم انليها 8 مو 5 كشان  
نعكس اثر ال add  
بنخط زهد  
دينا  
instruction e  
issue Packet



# Loop Unrolling

فيها بتزيد عدد ال instructions الموجودة كنا عشان نزيد ال IPC

- Replicate loop body to expose more parallelism

- فوائده
  - Reduces loop-control overhead

- Use different registers per replication

- Called "register renaming"
  - بخلينا نتغلب على أنواع لظافية من ال dependence

صنجان لتعمل loop unrolling →

- Avoid loop-carried "anti-dependencies"
  - Store followed by a load of the same register
  - Aka "name dependence", write-after-read
  - Or "output dependence", write-after-write
    - Reuse of a register name



```
if (i <= 10)
    a[i] = a[i] + 3;
```

مثال Chapter 4 — The Processor

```
if (i <= 10) {
    a[i] = a[i] + 3;
    a[i+1] = a[i+1] + 3;
}
```

instructions بعمله بطريقة احسن وبكبر ال حصا وبتزيد

## Unrolling Steps

بعملها البرمج وهاي الانيام بعملها ال compiler

انت بتختار

1. Replicate the loop instructions n times
2. Remove unneeded loop overhead
3. Modify instructions
4. Rename registers
5. Schedule instructions





# Loop Unrolling Example

	ALU/branch	Load/store	cycle
Loop:	addi x20, x20, -32	ld x28, 0(x20)	1
	nop	ld x29, 24(x20)	2
	add x28, x28, x21	ld x30, 16(x20)	3
	add x29, x29, x21	ld x31, 8(x20)	4
	add x30, x30, x21	sd x28, 32(x20)	5
	add x31, x31, x21	sd x29, 24(x20)	6
	nop	sd x30, 16(x20)	7
	blt x22, x20, Loop	sd x31, 8(x20)	8

32 general

32 Floating

بس عادي لاننا كثر register  
 $IPC = 14/8 = 1.75$

ما وصلنا 2 عثمان لسا في nop

بس هاي طريقة احسن

بانه قبل كنا نستخدم بس x31, x20... بس هالا حرمنا نستخدم x28, x29, x30

Closer to 2, but at cost of registers and code size

فلو بومني ال Performance ما بومني لو حجم البرنامج كبر نشوي

هو بحسن ال Performance بس على حساب ال registers وال code size يعني استعملنا Memory usage ولكن هاي المشاكل مقبولة وال Mem هاي الايام رخيصة وكيرة

Chapter 4 — The Processor — 63

32 bit

\* ال code اللي قبل :  $5 \times 4 = 20$  byte  
 ↓  
 instructions

كبر حجمه اكثر من 3 اضعاف

\* ال code الجديد :  $8 \times 2 \times 4 = 64$  bytes  
 ↓  
 instruction per issue packets

issue packets



Slide 41 : Exceptions in a pipeline :

بي أفترق - إنه بين اعلاوا E في مشكلة Exception ابي جمع int + int كلو كبير ما بقدر  
أخزنه بـ X1

		1	2	3	4	5	6	7	8	9	10	11	12	13
I1		F	D	E	M	W								
add X1, X2, X1		F	D	E*	M	W								
I3				F	D	E	W							
I4					F	D	W							
I5						F	W							
IHS							F	D	E	M	W			

بثكون عن Many

Instructions

ممكن ننتقل ال Mem بس كما توصل ال W

ما بديس تعملوا وبتحول ال Null وبتروح ال IHS

وهيك ختمنا إنه ال Instruction اللي قبلوا خلدوا لإنه ما حلرهم مشكلة

عملنا ال Flush لإنه ما بديس يكملوا

بعد ما حلينا ال exception باستخدام register SEPC بقدر إنه يرجع ال I3 مش

ال add بوزا المثال خصيصا لإنه عملية ال add ما بتزبط لإنه الناتج كبير ما بتخزن.

\* فإحنا ممكن نرجع لنفس ال instruction أو اللي بعدها وهذا بعتمد على الحالة.

Slide 47 : Multiple exceptions :

حضرنا إنه illegal opcode ما بتقدر تعملوا E

		1	2	3	4	5	6	7	8	9	10	11	12	13
I1		F	D	E	M	W								
add X1, X2, X1		F	D	E*	M*	W								
I3 (bad)				F	D*	E*	W			F	D*	E*	M*	W
I4					F	D	W							
I5						F	W							
IHS							F	D	E	M	W			
							F	D	E	M	W			

بالأول بتحل مشكلة ال add قبل ال I3 وبتروح وبيجي دور I3 بعدين



Slide 58 : Hazards and forwarding in Dual issue static processor :

Dependent load ↓	1	2	3	4	5	6	7	8	9	10
add $X_{10}, X_0, X_1$	F	D	E	M	W					
ld $X_2, 0(X_{10})$		F	D	E	M	W				

Forwarding

→ لا يجب ان  
issue packet  
و حصة لا  
يتم العمل بها  
فلازم بها ان

ان لا يتم الـ add الى الـ ld متاخره في issue packet

Slide 58 : Load use hazard :

	1	2	3	4	5	6	7	8	9	10
ld $X_1, 0(X_2)$	F	D	E	M	W					
sub $X_4, X_1, X_5$			F	D	E	M	W			

→ مايقدر نرسم  
بعض الـ issue packet  
فلازم بها الـ sub  
issue packet  
ثانية بس يكون  
فيه بعد الـ

الـ issue packet كلاس الـ

Slide 59 : From W in Memory pipeline to E in ALU Pipeline :

	1	2	3	4	5	6	7	8	9	10
ld $X_{31}, 0(X_{20})$	F	D	E	M	W					
sub $X_{31}, X_{31}, X_{21}$			F	D	E	M	W			

Slide 59 : From M in ALU pipeline to M in Memory pipeline :

	1	2	3	4	5	6	7	8	9	10	
add X31, X31, X21	F	D	E	M	W						
sd X31, 0(X20)	F	D	E	M	W						→ *السؤال load use hazard لأنه لا لل sd يتحسب من X20 وهو ما يعنى

علا instruction التي فوقها وال X31 يتحسب ابتداءً من cycle 3 وال sd بدأ تأخذ ال X31 وتبقى  
على Memory وتبقى على M ال cycle 4 ال M  
فيغير أحكام بـ issue packet بشرط نعمل forwarding .

Slide 62 : Example of loop unrolling :

Loop :

```
ld X31, 0(X20)
add X31, X31, X21
sd X31, 0(X20)
addi X20, X20, -8
blt X22, X20, loop
```

① Replicate the loop instructions  $n$  times → 4

Loop :

```
ld X31, 0(X20)      ld X31, 0(X20)
add X31, X31, X21   add X31, X31, X21
sd X31, 0(X20)      sd X31, 0(X20)
addi X20, X20, -8   addi X20, X20, -8
x [ blt X22, X20, loop ] x [ blt X22, X20, loop ]
```

```
ld X31, 0(X20)      ld X31, 0(X20)
add X31, X31, X21   add X31, X31, X21
sd X31, 0(X20)      sd X31, 0(X20)
addi X20, X20, -8   addi X20, X20, -8
x [ blt X22, X20, loop ] x [ blt X22, X20, loop ]
```



## ② Remove unneeded loop overhead

Loop:

ld X31, 0(X20)

ld X31, <sup>-16</sup>0(X20)

add X31, X31, X21

add X31, X31, X21

sd X31, 0(X20)

sd X31, <sup>-16</sup>0(X20)

ld X31, <sup>-8</sup>0(X20)

ld X31, <sup>-24</sup>0(X20)

add X31, X31, X21

add X31, X31, X21

sd X31, <sup>-8</sup>0(X20)

sd X31, <sup>-24</sup>0(X20)

addi X20, X20, <sup>-32</sup>-8

blt X22, X20, loop

## ③ Modify instructions

Loop:

ld X31, 0(X20)

ld X31, -16(X20)

add X31, X31, X21

add X31, X31, X21

sd X31, 0(X20)

sd X31, -16(X20)

ld X31, -8(X20)

ld X31, -24(X20)

add X31, X31, X21

add X31, X31, X21

sd X31, -8(X20)

sd X31, -24(X20)

addi X20, X20, -32

blt X22, X20, loop

## ④ Rename registers

Loop:

ld X28, 0(X20)

ld X30, -16(X20)

add X28, X28, X21

add X30, X30, X21

sd X28, 0(X20)

sd X30, -16(X20)

ld X29, -8(X20)

ld X31, -24(X20)

add X29, X29, X21

add X31, X31, X21

sd X29, -8(X20)

sd X31, -24(X20)

addi X20, X20, -32

blt X22, X20, loop

\* ال Renaming رح يساعد في ازالة بين ال instructions ولنا جاز Schedule حتى اقل  
من المشاكل التي ممكن تطلع عندي (anti-dependence, output dependence, WAW, WAR)

## ⑤ Schedule instructions

	ALU / branch	Load / store	cycle	
Loop:	addi X20, X20, -32	ld X28, 0(X20)	1	ld X28, 0(X20) add X28, X28, X21 sd X28, 0(X20)
	nop	ld X29, 24(X20)	2	ld X29, -8(X20) add X29, X29, X21 sd X29, -8(X20)
	add X28, X28, X21	ld X30, 16(X20)	3	ld X30, -16(X20) add X30, X30, X21 sd X30, -16(X20)
	add X29, X29, X21	ld X31, 8(X20)	4	ld X31, -24(X20) add X31, X31, X21 sd X31, -24(X20)
	add X30, X30, X21	sd X28, 32(X20)	5	sd X28, 32(X20) ld X31, -24(X20) add X31, X31, X21 sd X31, -24(X20)
	add X31, X31, X21	sd X29, 24(X20)	6	sd X29, 24(X20) addi X20, X20, -32 blt X22, X20, loop
	nop	sd X30, 16(X20)	7	sd X30, 16(X20) blt X22, X20, loop
	blt X22, X20, loop	sd X31, 8(X20)	8	sd X31, 8(X20)

\* نخلص تعبئة ال loads ال يفرق تنفيذها لسبب ما يمكن لا في ال instructions ال data التي بنحتاجها

ال load بنجربها

$$-8 = X - 32 / -16 = X - 32 / -24 = X - 32$$

\* بغير ال offset شي بنجرب مع ال addi :



# Dynamic Multiple Issue

أحدث وأقوى  
ويعرفوا كغيرها أكثر

- “Superscalar” processors → dynamic Multiple issue الاسم الثاني لا
- CPU decides whether to issue 0, 1, 2, ...  
each cycle  
مثلا ال ALU ما بيسر البعتوا مع 2 instruction مع ابعت
- Avoiding structural and data hazards
- Avoids the need for compiler scheduling
  - Though it may still help
  - Code semantics ensured by the CPU

ما بحتاجه  
ولكن لو عمل مثلا  
Loop unrolling  
يحسن أكثر

\* درجة ال Code مسؤولية ال CPU مش مسؤولية ال Compiler



# Dynamic Pipeline Scheduling

Allow the CPU to execute instructions out of order to avoid stalls → انه اذافيه instructions قادرين يستوا بعض بطولم ينتسوا لكن لو في بعدهم instructions آخرين بقدر اخدمهم فيخدمهم

But commit result to registers in order

Example

ld x31, 20(x21)

add x1, x31, x2

sub x23, x23, x3

andi x5, x23, 20

Can start sub while add is waiting for ld

حتوا حافظ على ترتيب البرنامج ببعث النتيجة  
الرفائية لا register على الترتيب  
الأحلي للبرنامج.

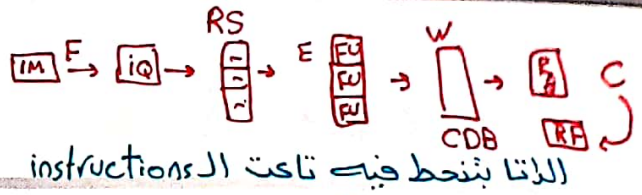
اهم اشياء  
يعملها







# Pipeline Stages



**F:** Fetch from instr. memory (IM) to instr. queue (IQ).

**I:** Issue from IQ to reservation stations (RS), reading ready operands from register file (RF).

**E:** Execute when functional unit (FU) is free and instr. In RS has ready operands. → متغيراتها جاهزين

**W:** Write result from FU through common data bus (CDB) to reorder buffer (ROB) and RS.

**C:** Commit results in order from ROB to RF and memory → بال Store بنكتب فيها


Loads have **FIAMWC**, stores have **FIAC**. **A:** Address calculation



بالدال  
تسمى  
Address calculation

بالا بنكتب بلا  
Memory

# Register Renaming



Reservation stations and reorder buffer effectively provide register renaming

On instruction issue to reservation station

If operand is available in register file or reorder buffer

- Copied to reservation station

- No longer required in the register; can be overwritten

If operand is not yet available

- It will be provided to the reservation station by a function unit

- Register update may not be required





# Examples

غير ال RAW في أمور ثانية

- Assume superscalar processor of degree 4
- Name dependence (WAR)

```
mul x1, x2, x3
add x4, x1, x5
ld x5, 16(x21)
```

بال 5 stage Pipeline  
لما كنا نتقدم بالترتيب ما كان في مشكلة بنوم

له ممكن نحلها بتغيير اسم ال Register  
الخوف لما يغير كرم ترتيب في الشيفر وتعمل ال ld

قبل ال add  
وهيك بغير  
خطأ

## Output dependence (WAW)

```
mul x1, x2, x3
add x4, x1, x5
ld x1, 16(x21)
```

الخوف مرة ثانية لأنه هاي ال instruction  
اللي هي ال ld تتبقى ال mul  
وهيك خطأ.



ممكن نحلها بإنه نغير  
اسم ال Register

# Speculation

مبنى على ال Dynamic Exec  
وتزيد عليه ال Speculation

“Guess” what to do with an instruction

- Start operation as soon as possible
- Check whether guess was right
  - If so, complete the operation
  - If not, roll-back and do the right thing

Common to static and dynamic multiple issue

## Examples

- Speculate on branch outcome
  - Roll back if path taken is different
- Speculate on load
  - Roll back if location is updated

ال هاردوير  
بيعمل تنبؤ  
زي ال Branch  
متأكد بتوقع لو ال  
Branch بتوقع بيشي  
بنفذ ال instruction  
اللي بتوقع  
وهيك ممكن  
تتوقع ال Branch  
وتتوقع ال load  
وتتوقع ال location  
وتتوقع ال path  
وتتوقع ال path  
وتتوقع ال path



# Compiler/Hardware Speculation

- Compiler can reorder instructions
  - e.g., move load before branch
  - Can include “fix-up” instructions to recover from incorrect guess
- Hardware can look ahead for instructions to execute
  - Buffer results until it determines they are actually needed
  - Flush buffers on incorrect speculation



جدا مفيد لأنه التنبؤ حاليا يكون صحيح بنسبة 95% تقريبا

## Branch Speculation

- Predict branch and continue issuing

- Don't commit until branch outcome determined → ما يفعل commit الا بعد

نعمل لا Branch عشوائيا نتأكد من التنبؤ

- **Example:** Assume a superscalar processor of degree 2 and the branch prediction is not taken.

بغدر يجيب 2 instructions  
Every cycle


```
I0    ld    x1, 0(x20)
I1    beq   x1, x2, skip
I2
I3
I4
```





# Load Speculation

- Avoid load and cache miss delay
  - Load before completing outstanding stores
  - Predict the effective address or loaded value
  - Bypass stored values to load unit
- Don't commit load until speculation cleared
- Example: Superscalar of degree 3.



independent  
ما يعتمدوا ببعض  
على الأكلب

ما يعتمد على اللي  
فوق



في حالات نادرة يكونوا  
dependent لو قتلنا ال

Address اللي حطيناه بـ x1 بالصيغة كان يساوي  
ال Address اللي موجود بـ x21 وهاي حالة نادرة

لو بدون Speculation  
حتميطر نشأخر لنعرف  
سبب بالـ Speculation لا

# Speculation and Exceptions

- What if exception occurs on a speculatively executed instruction?
  - e.g., speculative load before null-pointer check
- Static speculation
  - Can add ISA support for deferring exceptions
- Dynamic speculation
  - Can buffer exceptions until instruction completion (which may not occur)

# Exceptions Examples

- Assume superscalar processor of degree 3 with 2 address calculation units
- E1: Predict branch as not take, but resolve to taken. The 1d has exception in M.

```
beq x1, x2, L1
1d x5, 16(x21)
```

- E2: Assume first sd has exemption in C.

```
1d x1, 0(x20)
sd x1, 0(x21)
sd x2, 16(x21)
```



Chapter 4 — The Processor — 75  
منهوا إرفا بتحل ال WAW والشائج لبتخزن حسب الترتيب الأصلي تاخوم. وجود ROB يمكننا إنه تعمل speculation وأيضا يمكننا نتعامل مع ال Exception بطريقة سليمة

## Why Do Dynamic Scheduling?

- Why not just let the compiler schedule code?
- Not all stalls are predicabile
  - e.g., cache misses
- Can't always schedule around branches
  - Branch outcome is dynamically determined
- Different implementations of an ISA have different latencies and hazards





# Does Multiple Issue Work?

## The BIG Picture

- Yes, but not as much as we'd like
- Programs have real dependencies that limit ILP
- Some dependencies are hard to eliminate
  - e.g., pointer aliasing
- Some parallelism is hard to expose
  - Limited window size during instruction issue
- Memory delays and limited bandwidth
  - Hard to keep pipelines full
- Speculation can help if done well



## Power Efficiency

Complexity of dynamic scheduling and speculations requires power  
 Multiple simpler cores may be better

more transistors, more circuits  
 much more complex than single pipeline  
 ما كان بحاجة لحظ هروحة  
 تغير خلال 40 سنة تقريبا

Microprocessor	Year	Clock Rate	Pipeline Stages	Issue Width	Out-of-Order/Speculation	Cores/Chip
Intel 486	1989	25 MHz	5	1	No	1
Intel Pentium	1993	66 MHz	5	2	No	1
Intel Pentium Pro	1997	200 MHz	10	3	Yes	1
Intel Pentium 4 Willamette	2001	2000 MHz	22	3	Yes	1
Intel Pentium 4 Prescott	2004	3600 MHz	31	3	Yes	1
Intel Core	2006	2930 MHz	14	4	Yes	2
Intel Core i5 Nehalem	2010	3300 MHz	14	4	Yes	2-4
Intel Core i5 Ivy Bridge	2012	3400 MHz	14	4	Yes	8

مختلف عن الـ simple pipeline الذي درسته  
 له لما ال Pipeline stages بتزيد الاشياء الي  
 بدي لعله بـ stage ورتبة بغير اقل  
 والدايرة بـ stage ورتبة بتشير الي  
 وبالتالي بتقدر تشتغل على shorter clock period يعني اقل perform...  
 وهذا كان يزيد ال Power كثير  
 بس سيطروا عليه آخر السنوات  
 وقلوا عدد ال Pipeline stages  
 Chapter 4 — The Processor





# Contents

- 4.6 Pipelined Datapath and Control (Review)
- 4.7 Data Hazards: Forwarding versus Stalling
- 4.8 Control Hazards
- 4.9 Exceptions
- 4.10 Parallelism via Instructions
- 4.11 Real Stuff: The ARM Cortex-A53 and Intel Core i7 Pipelines
- 4.14 Fallacies and Pitfalls
- 4.15 Concluding Remarks

مثالين مهمات  
modern Process

## Cortex A53 and Intel i7

استخدم من شركة Apple في بعض هواتفها والسبب، انه مصمم بصرف للهواتف، أقل

برضه يستخدم في ال high end Processor ال computer زي ال PC

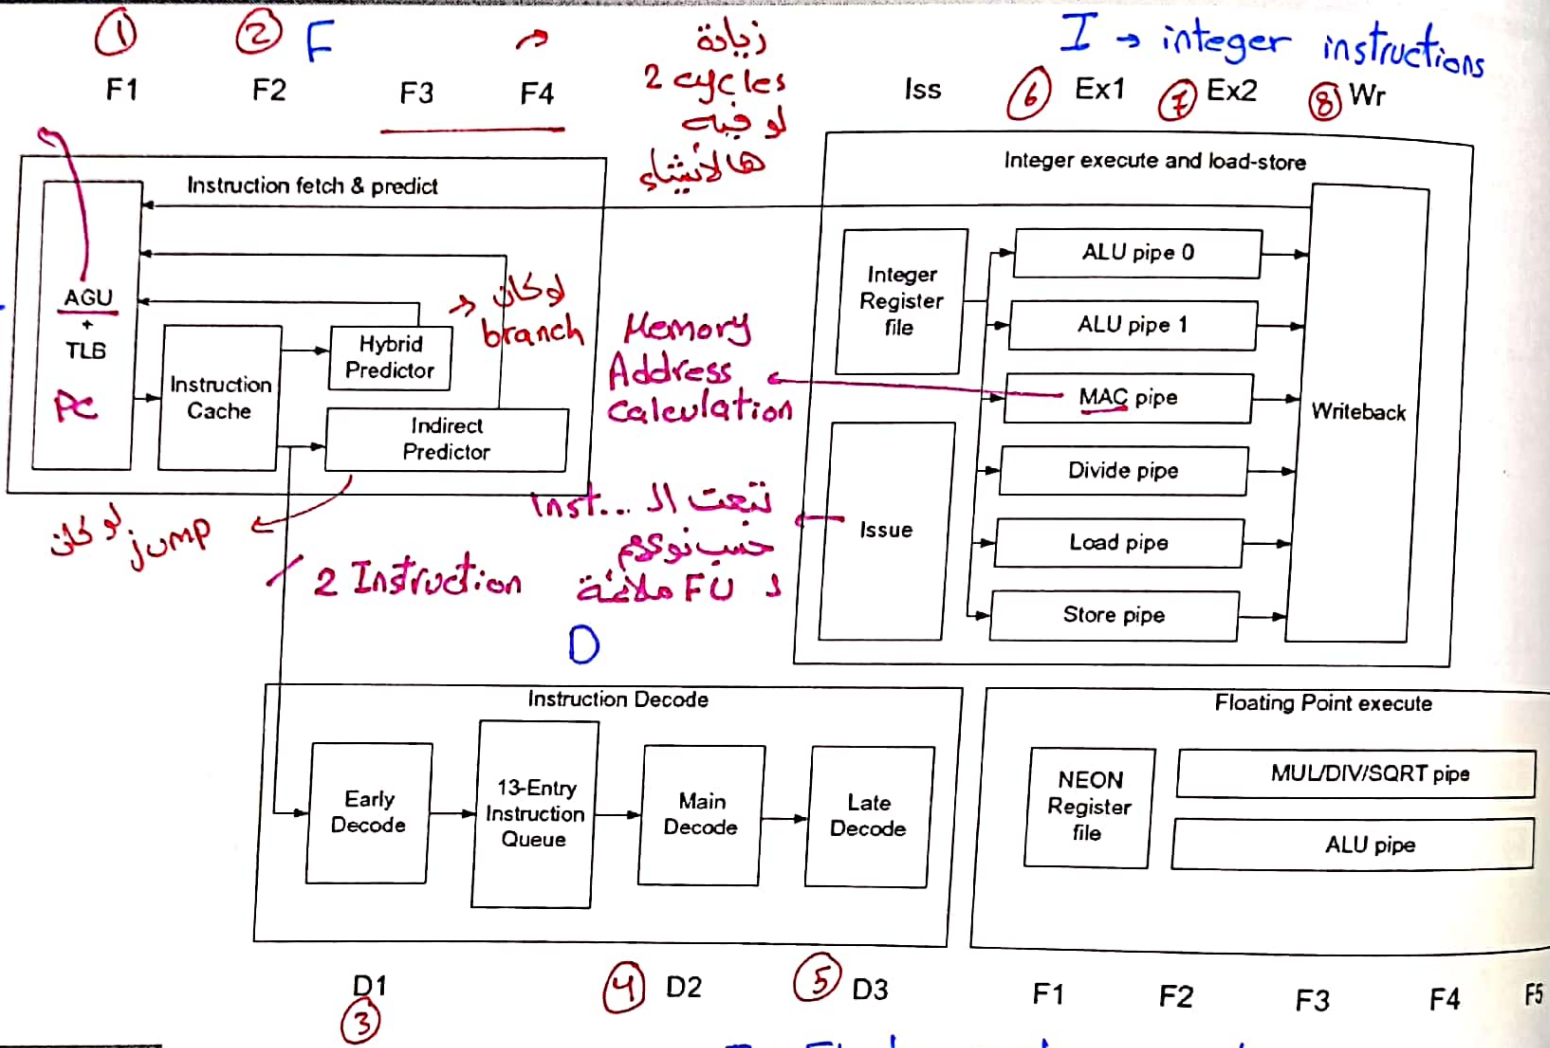
Processor	ARM A53	Intel Core i7 920
Market	Personal Mobile Device	Server, cloud
Thermal design power	100 milliWatts (1 core @ 1 GHz)	130 Watts
Clock rate	1.5 GHz	2.66 GHz
Cores/Chip	4 (configurable)	4 → ممكن هلاقي 8 و 6
Floating point?	Yes	Yes
Multiple issue?	Dynamic	Dynamic
Peak instructions/clock cycle	2 → 2 instruction every cycle	4 → 4 instruction every cycle
Pipeline stages	8 ← أكثر لانه بيشغل Frequency أقل	14
Pipeline schedule	Static in-order	Dynamic out-of-order with speculation
Branch prediction	Hybrid	2-level
1 <sup>st</sup> level caches/core	16-64 KiB I, 16-64 KiB D	32 KiB I, 32 KiB D
2 <sup>nd</sup> level caches/core	128-2048 KiB	256 KiB (per core)
3 <sup>rd</sup> level caches (shared)	(platform dependent)	2-8 MB

\$4.11 Real Stuff: The ARM Cortex-A53 and Intel Core i7 Pipelines



# ARM Cortex-A53 Pipeline

في 8 stages بأصعب الأحوال Address generation unit ال أهم اشيا فيه هو ال Instruction cach اللي فيه ال Instructions اللي يستخدم



زيادة 2 cycles لو فيه هالاشيا

Memory Address calculation

تبع ال Inst... حسب نوع ال FU ملأنة 2 Instruction

I → Floating point instructions


لو كان Floating في 10 cycles or stages



# ARM Cortex-A53 Performance

بالدراسة  
هاي  
قلسوا قديه  
بيان  
time  
Sec  
وقديه بنسخه  
Instructions

Sec x f → cycle  
cycle  
IC  
قسسوا عدد العنصر  
اللي بنحتاجهم لتنفيذ  
كل ال Instructions  
لهذا البرنمج  
بطلع معي هذا  
ال ratio



Memory hirany  
Stalls

المشوق  
الها بسبب: يا انه  
ال Instructions في بينا اعتماد بعض  
او انه ال Processor يحاول يطلب ال data من ال cache وهيك بغير

Chapter 4 — The Processor

بحتاج فتره طويله  
ليجي ال data من ال Mem

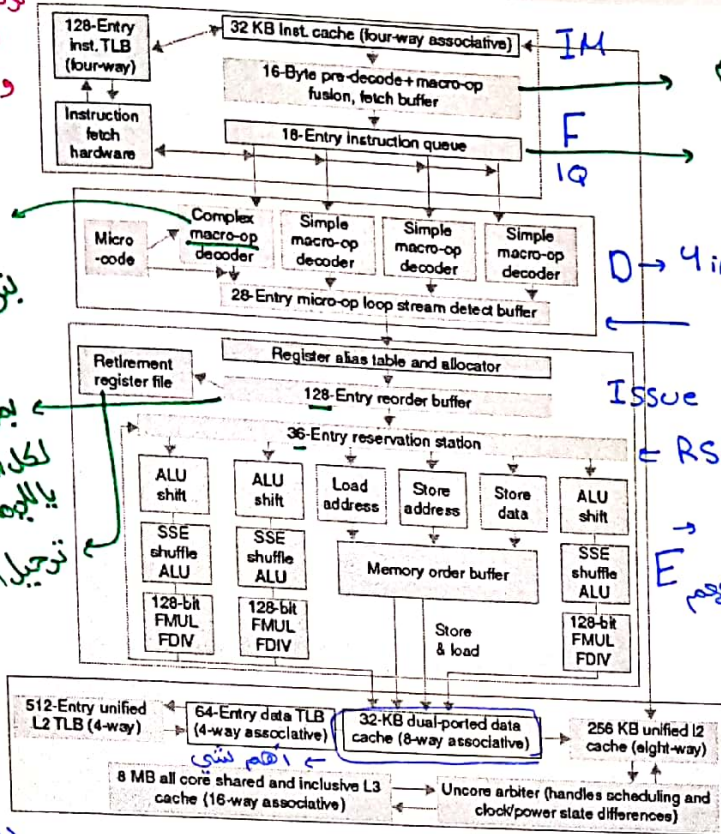


# Core i7 Pipeline

→ Processor much more complex  
Superscalar Processor with degree 4

وقت بخلاف من وقت  
وقت اجرای FU و ROB و RS

ما نعرف كم  
Instructions  
بمع 18 مجموعة  
→ 18 \* 16 byte



Complex instructions  
تترجمها ل Simple

بمجرد Issue ينجز  
لكل Instruction بلا ROB  
يالمصنوع بغير ال Commit  
ترتيب النتائج الى حسب ترتيبها  
الاصلي

4 instruction  
ال decode instructions  
ينحطوا فيه



تنفيذ ال Instructions ب (FU)  
Functional unit حسب نوعهم  
Floating + بقروا يعملوا ل  
Point

مش  
خطي  
نوع  
بالوقت  
الحاضر

# Core i7 Performance

تقلنا ال Performance

اسوا ما يمكن  
100%



اسوا وا

3 مرات أسرع من ال ARM  
لانه يشتغل على high frequency  
في dynamic exc... ولو مع بعضنا  
بخطي نتائج افضل  
لكن على حساب ال Power

# Contents

- 4.6 Pipelined Datapath and Control (Review)
- 4.7 Data Hazards: Forwarding versus Stalling
- 4.8 Control Hazards
- 4.9 Exceptions
- 4.10 Parallelism via Instructions
- 4.11 Real Stuff: The ARM Cortex-A53 and Intel Core i7 Pipelines
- 4.14 Fallacies and Pitfalls
- 4.15 Concluding Remarks



Chapter 4 — The Processor

## Fallacies

- Pipelining is easy (!)
  - The basic idea is easy
  - The devil is in the details
    - e.g., detecting data hazards
- Pipelining is independent of technology
  - So why haven't we always done pipelining?
  - More transistors make more advanced techniques feasible
  - Pipeline-related ISA design needs to take account of technology trends
    - e.g., predicated instructions





# Pitfalls

- Poor ISA design can make pipelining harder
  - e.g., complex instruction sets (VAX, IA-32)
    - Significant overhead to make pipelining work
    - IA-32 micro-op approach
  - e.g., complex addressing modes
    - Register update side effects, memory indirection
  - e.g., delayed branches
    - Advanced pipelines have long delay slots



## Contents

- 4.6 Pipelined Datapath and Control (Review)
- 4.7 Data Hazards: Forwarding versus Stalling
- 4.8 Control Hazards
- 4.9 Exceptions
- 4.10 Parallelism via Instructions
- 4.11 Real Stuff: The ARM Cortex-A53 and Intel Core i7 Pipelines
- 4.14 Fallacies and Pitfalls
- 4.15 Concluding Remarks

# Concluding Remarks

- Pipelining improves instruction throughput using parallelism
  - More instructions completed per second
  - Latency for each instruction not reduced
- Hazards: structural, data, control
- Multiple issue and dynamic scheduling (ILP)
  - Dependencies limit achievable parallelism
  - Complexity leads to the power wall





Slide 67 : Dynamic Scheduling :

Single issue  $\rightarrow$  1 instruction يعني يجيب بس

	1	2	3	4	5	6	7	8	9	10
ld X31, 20(X21)	F	I	A	M	W*	C				
add X1, X31, X2		F	I	-	-	E	W	C		
sub X23, X23, X3			F	I	E	W*	-	-	C	
andi X5, X23, 20				F	I	-	E	W	-	C

\* بنعلم I بالترتيب الأصلي تبعهم وكذلك C ولكن E مش بالترتيب الجاهز بنفذ

Slide 69 :

$\rightarrow$  3 things يعني أي Stage بتقدر تعمل

\* Triple Issue : Name dependence (WAR)

Ex تنفيذ إن شاء الله mul تحتاج 3 cycles بال

	1	2	3	4	5	6	7	8	9	10
mul X1, X2, X3	F	I	E	E	E	W	C			
add X4, X1, X5	F	I	-	-	-	-	E	W	C	
ld X5, 16(X21)	F	I	A	M	W	-	-	-	C	

↓  
نجمع محتويات X21 مع 16

\* الكتابة على ال RF بتتم بال C مش ال W

\* Triple Issue : output dependence (waw)

Ex تنفيذ إن شاء الله mul تحتاج 3 cycles بال

كبت ب ROB مش بال RF

	1	2	3	4	5	6	7	8	9	10
mul X1, X2, X3	F	I	E	E	E	W	C			
add X4, X1, X5	F	I	-	-	-	-	E	W	C	
ld X1, 16(X21)	F	I	A	M	W	-	-	-	C	

هون بكتب X1

هون بكتب X1

RF مش بال reorder buffer له هاي كبت بال

وهون X بتخزن بال RF لاني نتيجة X الثانية

Slide 72 :

Example: Assume a superscalar processor of degree 2 and the branch prediction is not taken. (Correct Prediction)

	1	2	3	4	5	6	7	8	9	10	توضيح I3 C
Id $x_1, 0(x_{20})$	F	I	A	M	W	C					في نفس ال cycle عادي
bez $x_1, x_2, skip$	F	I				E	W	C			لما كذا C كذا
I3		F	I	E	W			C			انه تنبؤ صح
I4		F	I	E		W			C		بعلم F يا
...											cycle 2
Skip:											انه Not taken

3 write في ال W انه ما يقرب

of degree 2 بنفس ال cycle ال البرنامج

C/W/M/I/F in 2 فكل ما يقرب

Functional unit بس ال E يعني قويه سني



Slide 72 :

Example: Assume a superscalar processor of degree 2 and the Branch prediction is Not taken. (Incorrect Prediction)

	1	2	3	4	5	6	7	8	9	10	توضيح E
Id $x_1, 0(x_{20})$	F	I	A	M	W	C					بس E كذا
bez $x_1, x_2, skip$	F	I				E	W	C			انه تنبؤ صح
I3		F	I	E	W			n			ع
I4		F	I	E		W		n			لما بعلم C ال
...											Branch ال
Skip:									F		C كذا



Slide 73 :

Example : Load speculation. Assume a superscalar processor of degree 3.

Predict the second load does not depend on the store.

(Correct prediction)

	1	2	3	4	5	6	7	8	9	10	
ld $x_1, o(x_{20})$	F	I	A	M	W	C					
sd $x_2, o(x_1)$	F	I				A	C				بنكتب على Mem
ld $x_3, o(x_{21})$	F	I	A	M	W		C				بال Commit

Slide 73 :

Example : Load speculation. Assume a superscalar processor of degree 3.

Predict the second load does not depend on the store. (Incorrect prediction)

	1	2	3	4	5	6	7	8	9	10	
ld $x_1, o(x_{20})$	F	I	A	M	W	C					
sd $x_2, o(x_1)$	F	I				A	C				] → RAW
ld $x_3, o(x_{21})$	F	I	A	M	W		R	F			

اكتشفنا ان  $x_1 = x_{21}$

علفوف انه  $x_1 \neq x_{21}$

لاننا اخذت قيمة خاطئة

بنعدها مرة ثانية لانها كانت خطأ بالبارية وهيك بنمشي صح.

Slide 75 : Speculation and Exceptions

Example 1: Assume superscalar processor of degree 3 with 2 address calculations units.

Predict branch as not taken, but resolve to taken. The ld has exception in M.

هونينكشف بانوا Taken

	1	2	3	4	5	6	7	8	9	10
beq X1, X2, L1	F	I	E	W	C					
ld X5, 16(X21)	F	I	A	M	N					

↳ جيناها بناء على الافتراض اننا ال beq  
 Not taken ←  
 هونا اكتشفنا انه فيه Exception

Exception Handling Subroutine ← EHS

بس بالتصميم الحديثة، انا فيه Exception احنا بنعالجها، الا بال Commit بس عشان  
 احنا بننفذ ال ld عالفاضي لاننا ال beq ← taken فبيسا طة بنعالجها null  
 وال Exception ولا كانه صار

Example 2 : Assume superscalar processor of degree 3 with 2 address calculation units.

Assume First sd has exception in C.

	1	2	3	4	5	6	7	8	9	10
ld X1, 0(X20)	F	I	A	M	W	C				
sd X1, 0(X21)	F	I	A	-	-	N				
sd X2, 16(X21)	F	I	-	A	-	N				
EHS :							F			

عرفنا انه فيه exception  
 فنبحول من C الى null

وبيروح ال EHS وبس اخلص والامور تكون ملائمة يرجع بعمل F ال instruction اللي  
 عملناوم null ويرجع انفسهم مرة ثانية



# Chapter 5

## Large and Fast: Exploiting Memory Hierarchy

↓  
التكبير على

Memory

الهدف، اني اضع Memory حديثة تكون كبيرة وبتقضي الوقت

سرعة وهذا مش موضوع سرعة

Adapted by Prof. Gheith Abandah

# Contents

- 5.1 Introduction
- 5.2 Memory Technologies
- 5.3 The Basics of Caches
- 5.4 Measuring and Improving Cache Performance
- 5.5 Dependable Memory Hierarchy
- 5.11 Redundant Arrays of Inexpensive Disks
- 5.6 Virtual Machines
- 5.7 Virtual Memory
- 5.8 A Common Framework for Memory Hierarchy
- 5.9 Using a Finite-State Machine to Control a Simple Cache
- 5.10 Cache Coherence
- 5.13 The ARM Cortex-A53 and Intel Core i7 Memory Hierarchies
- 5.16 Fallacies and Pitfalls
- 5.17 Concluding Remarks

→ حذرنا من ناحية ال Memory

Chapter 5 — Large and Fast: Exploiting Memory Hierarchy — 2

صفة من صفات البرامج لما نتعامل مع ال Memory  
ولنفس الوقت موجودة بحياتنا?

## Principle of Locality

Programs access a small proportion of their address space at any time → مثلا برنامج معين ما يشتغل بكل امكانياته يشتغل على شغلات معينة فيه بس.

### Temporal locality

- Items accessed recently are likely to be accessed again soon → اذا اشيت طلبته هسا في احتمال كبير انك تطلبه بعد شوي.
- e.g., instructions in a loop, induction variables

### Spatial locality

- Items near those accessed recently are likely to be accessed soon → اذا في اشيت عملته access احتمال اعلى access لا locations اللي قريبين منك
- E.g., sequential instruction access, array data

↳ Ex: for (i=0 ... )

Chapter 5 — Large and Fast: Exploiting Memory Hierarchy — 3



# Taking Advantage of Locality

المغزى من ال disk ولما نسخر الجواز محتوياتها  
 ما يشرح  
 هذا البرامج الي  
 مخزنة  
 Smart Phones

انه ال Mem تكون اثير  
 هنا نوع واثير  
 من طبيعة  
 كبير وبطيئ  
 هنية صف سريعة  
 كبير ورخيصة  
 وهذا المستوى الثاني  
 بعد ال disk / Flash  
 الشغلات الي  
 تنقل اليها والقرينة  
 وهاي سريعة  
 لكن صغيرة

Memory hierarchy


Store everything on disk or Flash

Copy recently accessed (and nearby) items from disk to smaller DRAM memory

■ Main memory

Copy more recently accessed (and nearby) items from DRAM to smaller SRAM memory

■ Cache memory attached to CPU



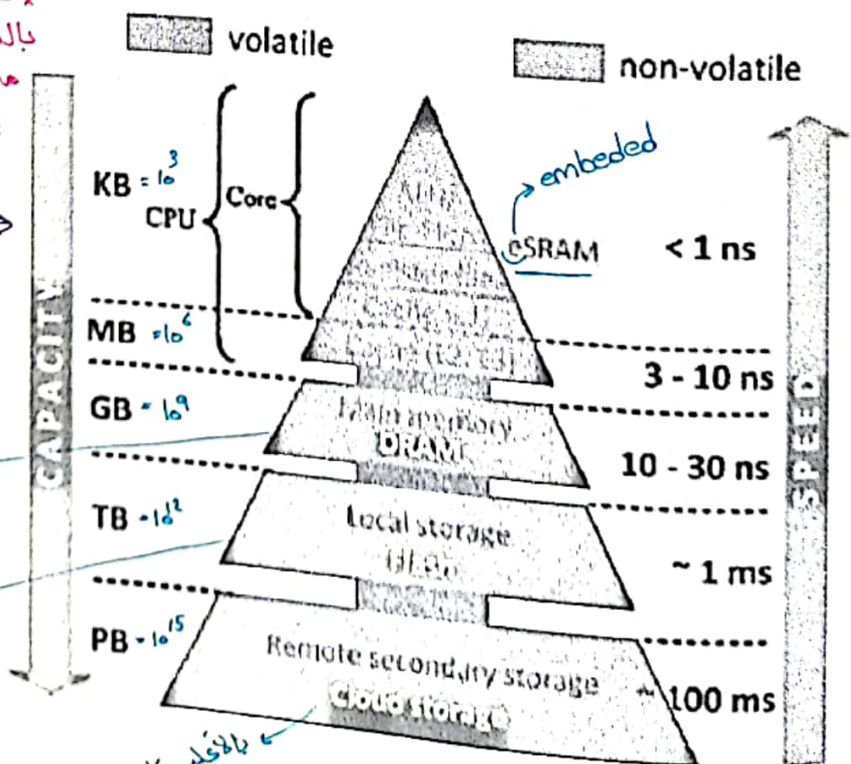
# Memory Hierarchy

\* تصميم مثل هذا التصميم  
 بالنسبة لا Processor  
 ما رح يدخل مع بعض  
 Fast and large

Fast → هي الي نستخدمها  
 أكثر

Large →

دايما DRAM  
 بالاطيب  
 disk /  
 Flash



بالاطيب disk

# Memory Hierarchy Levels

Memory يتقسم لـ blocks  
يعني لما نجيب data من تحت لفة  
بنجيبها blocks مش مفردة

اسم ثاني إله

Block (aka line): unit of copying

May be multiple words 64 bytes → block لا الواحد

If accessed data is present in

upper level → Cache طلب شغل من ال Processor  
بنقول من نوع هاي ال access انجا Hit

Hit: access satisfied by upper level

Hit ratio: hits/accesses → 
$$= \frac{\text{hits}}{\text{hits} + \text{misses}}$$

If accessed data is absent


Miss: block copied from lower level

Time taken: miss penalty

Miss ratio: misses/accesses

$$= 1 - \text{hit ratio}$$

Then accessed data supplied from upper level



لغينا اللي بدنا اياه بال Cache

بنأخذ ال block اللي فيه ال word المطلوبة وبننسخه ال Upper Memory Level

الداتا اللي بدنا ياهو مش موجودة بال Cache موجودة بال lower Memory level

الزمن الإضافي اللي يحتاجه كين ما أحيب ال data

Chapter 5 — Large and Fast: Exploiting Memory Hierarchy — 6

Ex: hit ratio = 95% → دايتا عالي  
miss ratio = 100% - 95% = 5%

## Contents

5.1 Introduction

5.2 Memory Technologies → أنواع ال Memory Technologies

Introduction

SRAM

DRAM

Flash

Disk Storage



# Memory Technology (2012)

- Static RAM (SRAM) → الأسرع ولكن الأغلى
  - 0.5ns – 2.5ns, \$2000 – \$1000 per GB
- Dynamic RAM (DRAM)
  - 50ns – 70ns, \$10 – \$20 per GB
- Flash memory
  - 5,000ns – 50,000ns, \$0.75 – \$1.00 per GB
- Magnetic disk
  - 5ms – 20ms, \$0.05 – \$0.10 per GB
- Ideal memory
  - Access time of SRAM
  - Capacity and cost/GB of disk

لا جيجا بايت  
الواحدة




Chapter 5 — Large and Fast: Exploiting Memory Hierarchy-

1970 → Static Random Access Memory

## SRAM Technology

"SRAMs lose data content on power loss"

- Static RAM
- 6-8 transistors per bit
- Fast but not dense → كثيف
- Often has standby mode → وضع الاستعداد



IDT61078A/A  
CMOS Static RAM (16K / 16K x 1-Bit)

### Pin Configurations

A0	1	20	VCC
A1	2	19	A13
A2	3	18	A12
A3	4	17	A11
A4	5	16	A10
A5	6	15	A9
A6	7	14	A8
DOU1	8	13	A7
WE	9	12	DN
GND	10	11	CS

$2^4 \times 2^4 \times 2^{10} = 16 \text{ k bits}$





Chapter 5 — Large and Fast: Exploiting Memory Hierarchy-

# DRAM Technology

بعضه انه تخزين المعلومات  
 Capacitor و يحتاج  
 one transistor

- Data stored as a charge in a capacitor
  - Single transistor used to access the charge
  - Must periodically be refreshed
    - Read contents and write back
    - Performed on a DRAM "row"



\* 1 transistor  
 بتخزن 1 bit

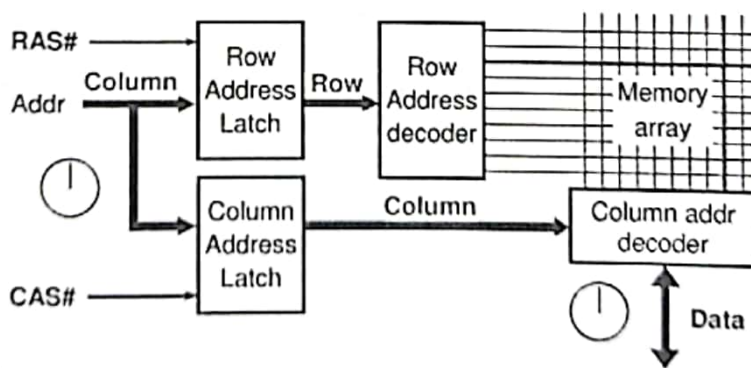


\* DRAM Refresh Each memory cell 16 times Per Second (or more) or risk losing content.

## Classic DRAM



### Basic DRAM chip




• DRAM access sequence

- Put Row on addr. bus
- Assert RAS# (Row Addr. Strobe) to latch Row
- Put Column on addr. bus
- Wait RAS# to CAS# delay and assert CAS# (Column Addr. Strobe) to latch Col
- Get data on address bus after CL (CAS latency)





# Classic DRAM



Every access - individual

2009-2013

©S.Maciulevičius

- Low bandwidth



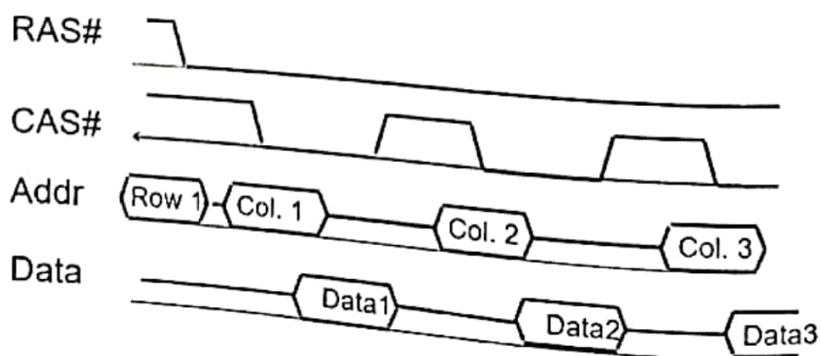
Chapter 5 — Large and Fast: Exploiting Memory Hierarchy-1

التعليقات

# Advanced DRAM Organization

- Access an entire row and save it in a row buffer.
- **Fast page mode:** supply successive words from the row buffer with reduced latency

وقت الاستجابة

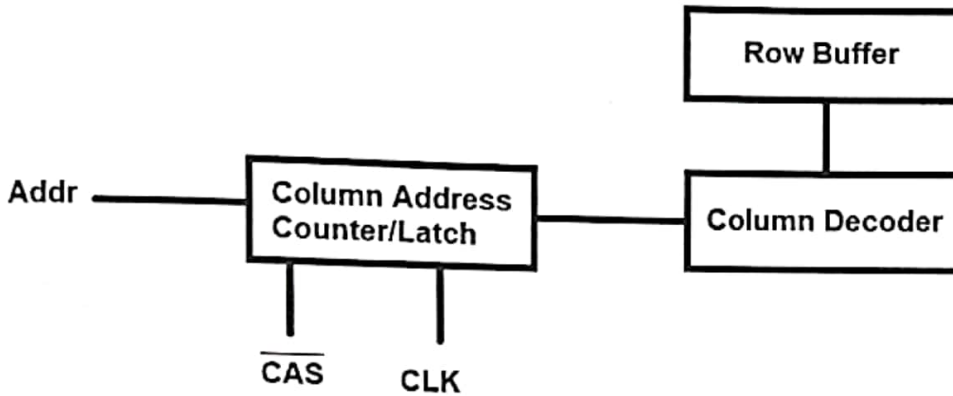


Chapter 5 — Large and Fast

Scanned with CamScanner

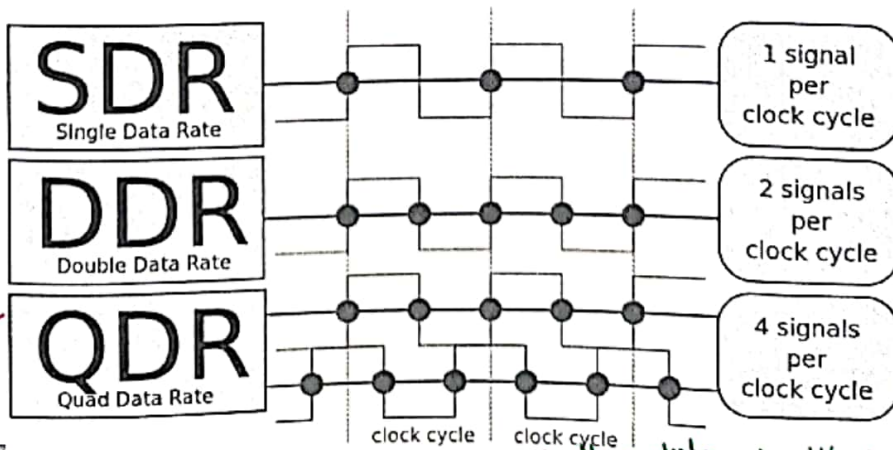
# Advanced DRAM Organization

**Synchronous DRAM (SDRAM)** has a counter that increments the column address using a clock signal.



# Advanced DRAM Organization

- **Double data rate (DDR) SDRAM**
  - Transfer on rising and falling clock edges
- **Quad data rate (QDR) SDRAM**
  - Separate DDR inputs and outputs



فكرته انه في clock والدا تا ينتقل مرتين كل cycle عند ال Raising Edge وعند ال falling edge فبالا cycle الواحدة بيتنقل 2 data transfers وهذا النوع من ال Mem ال bandwidth تاسه على

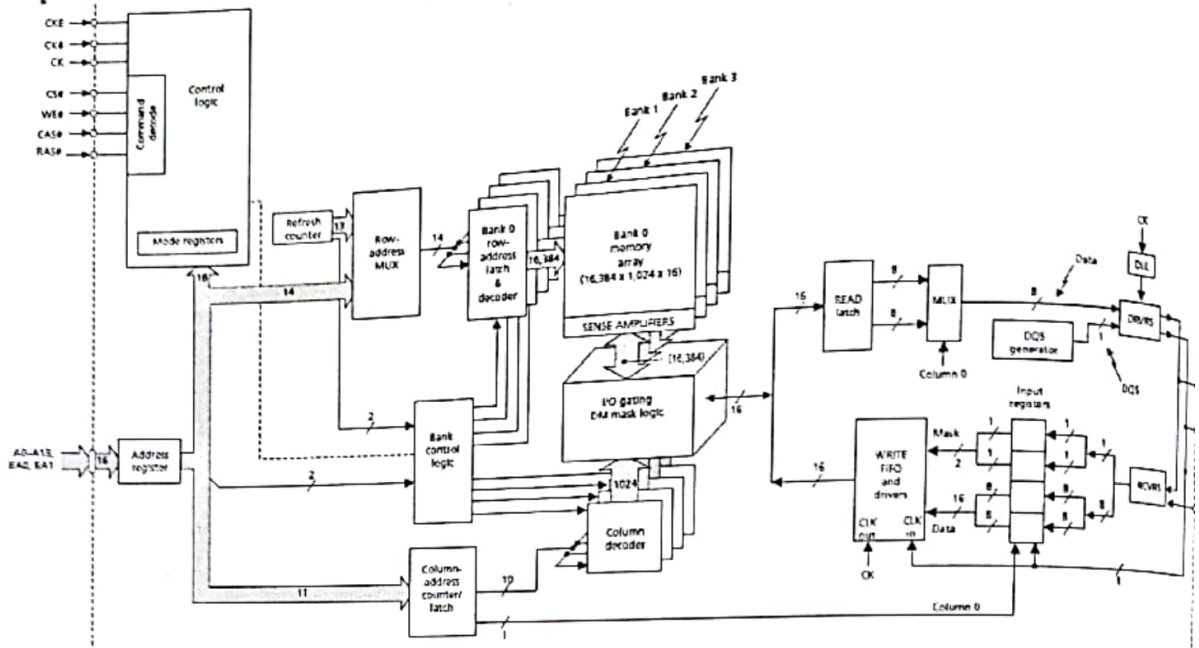
معظم الذاكرة تستخدم 1 Port فيه 2 Ports



# Micron 1Gb DDR-SDRAM

شرح

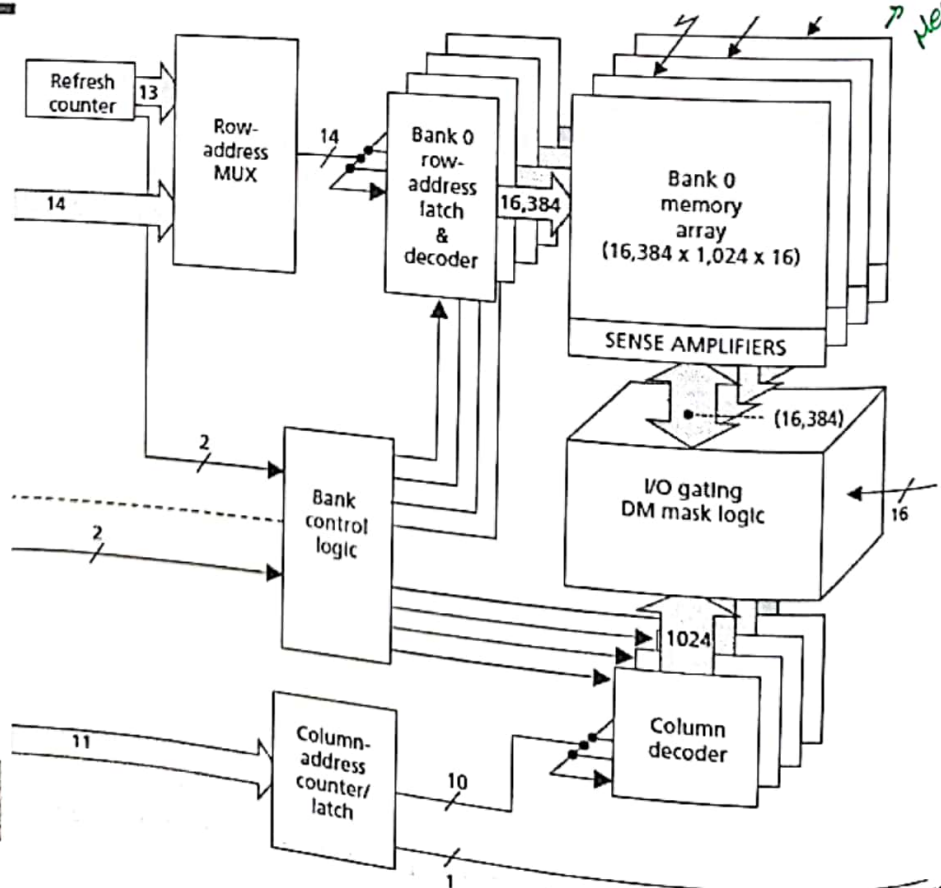
MT46V128M8 - 32 Meg X 8 X 4 Banks, Datasheet



Chapter 5 - Large and Fast: Exploiting Memory Hierarchy-

# Micron 1Gb DDR-SDRAM

داخل  
البنية من  
Mem array



# DRAM Generations

Year introduced	Chip size	\$ per Gb	Total access time to a new row/column	Average column access time to existing row
1980	64 Kibibit	\$1,500,000	250 ns	150 ns
1983	256 Kibibit	\$500,000	185 ns	100 ns
1985	1 Megibit	\$200,000	135 ns	40 ns
1989	4 Megibit	\$50,000	110 ns	40 ns
1992	16 Megibit	\$15,000	90 ns	30 ns
1996	64 Megibit	\$10,000	60 ns	12 ns
1998	128 Megibit	\$4,000	60 ns	10 ns
2000	256 Megibit	\$1,000	55 ns	7 ns
2004	512 Megibit	\$250	50 ns	5 ns
2007	1 Gibibit	\$50	45 ns	1.25 ns
2010	2 Gibibit	\$30	40 ns	1 ns
2012	4 Gibibit	\$1	35 ns	0.8 ns


بفضل التسعير  
الزمن

يتم تحديث  
مع الزمن



# DRAM Generations

Year	Capacity	\$/GB
1980	64Kbit	\$1500000
1983	256Kbit	\$500000
1985	1Mbit	\$200000
1989	4Mbit	\$50000
1992	16Mbit	\$15000
1996	64Mbit	\$10000
1998	128Mbit	\$4000
2000	256Mbit	\$1000
2004	512Mbit	\$250
2007	1Gbit	\$50





# DRAM Performance Factors

- Row buffer
  - Allows several words to be read and refreshed in parallel
- Synchronous DRAM
  - Allows for consecutive accesses in bursts without needing to send each address
  - Improves bandwidth
- DRAM banking
  - Allows simultaneous access to multiple DRAMs
  - Improves bandwidth




Chapter 5 — Large and Fast: Exploiting Memory Hierarchy-


## Increasing Memory Bandwidth

تصاميم مختلفة للذاكرة


تحتاج وقت طويل



a. One-word-wide memory organization



b. Wider memory organization



c. Interleaved memory organization

لا يأخذ وقت طويل ال transfer سريع

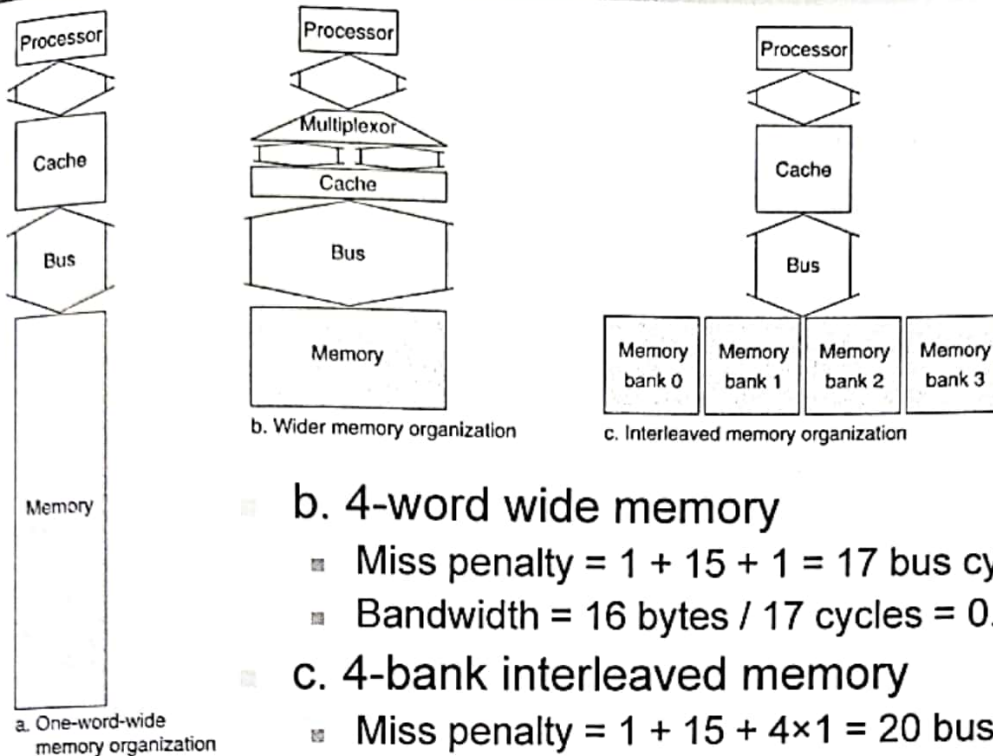
To get 16-byte block:

- a. One-word wide memory
  - Miss penalty =  $4 \times (1 + 15 + 1) = 68$  bus cycles
  - Bandwidth =  $16 \text{ bytes} / 68 \text{ cycles} = 0.24 \text{ B/cycle}$



Chapter 5 — Large and Fast: Exploiting Memory Hierarchy-


# Increasing Memory Bandwidth



- b. 4-word wide memory
  - Miss penalty =  $1 + 15 + 1 = 17$  bus cycles
  - Bandwidth =  $16 \text{ bytes} / 17 \text{ cycles} = 0.94 \text{ B/cycle}$
- c. 4-bank interleaved memory
  - Miss penalty =  $1 + 15 + 4 \times 1 = 20$  bus cycles
  - Bandwidth =  $16 \text{ bytes} / 20 \text{ cycles} = 0.8 \text{ B/cycle}$

# Increasing Memory Bandwidth

- d. DDR-SDRAM
  - Miss penalty =  $1 + 15 + 4 \times 0.5 = 18$  bus cycles
  - Bandwidth =  $16 \text{ bytes} / 18 \text{ cycles} = 0.89 \text{ B/cycle}$







# Flash Storage

سريع  
بسبب التخزين  
مع الدقة

- Nonvolatile semiconductor storage
  - 100× – 1000× faster than disk
  - Smaller, lower power, more robust
  - But more \$/GB (between disk and DRAM)





Chapter 6 — Storage and Other I/O Topics — 2

## Flash Types

- NOR flash: bit cell like a NOR gate
  - Random read/write access
  - Used for instruction memory in embedded systems
- NAND flash: bit cell like a NAND gate
  - Denser (bits/area), but block-at-a-time access
  - Cheaper per GB
  - Used for USB keys, media storage, ...
- Flash bits wears out after 1000's of accesses
  - Not suitable for direct RAM or disk replacement
  - Wear leveling: remap data to less used blocks

# Disk Storage

- Nonvolatile, rotating magnetic storage



## Disk Sectors and Access

- Each sector records
  - Sector ID
  - Data (512 bytes, 4096 bytes proposed)
  - Error correcting code (ECC)
    - Used to hide defects and recording errors
  - Synchronization fields and gaps
- Access to a sector involves
  - Queuing delay if other accesses are pending
  - Seek: move the heads
  - Rotational latency
  - Data transfer
  - Controller overhead





# Disk Access Example

- Given
  - 512B sector, 15,000rpm, 4ms average seek time, 100MB/s transfer rate, 0.2ms controller overhead, idle disk
- Average read time
  - 4ms seek time
  - +  $\frac{1}{2} / (15,000/60) = 2\text{ms}$  rotational latency
  - +  $512 / 100\text{MB/s} = 0.005\text{ms}$  transfer time
  - + 0.2ms controller delay
  - = 6.2ms
- If actual average seek time is 1ms
  - Average read time = 3.2ms




# Disk Access Example 2

- Given
  - 15,000rpm, 2MB/cylinder
- Sustainable peak transfer rate?

# Disk Performance Issues

- Manufacturers quote average seek time
  - Based on all possible seeks
  - Locality and OS scheduling lead to smaller actual average seek times
- Smart disk controller allocate physical sectors on disk
  - Present logical sector interface to host
  - SCSI, ATA, SATA
- Disk drives include caches
  - Prefetch sectors in anticipation of access
  - Avoid seek and rotational delay



## Contents

5.1 Introduction

5.2 Memory Technologies

5.3 The Basics of Caches

مبداً مستخدم بشكل واسع برقيات  
وبسول علينا.

Direct Mapped Cache

Cache Example

} Reading

Larger Block Sizes

Writing to the Cache

Example: Intrinsicity FastMATH



# Cache Memory


## Cache memory

ال Mem تتكون من طبقات قريبة من ال CPU


The level of the memory hierarchy closest to the CPU

Given accesses  $X_1, \dots, X_{n-1}, X_n$

أيسر حالة Cache واحد



a. Before the reference to  $X_n$



b. After the reference to  $X_n$

How do we know if the data is present?

Where do we look?

## Direct Mapped Cache

أيسر أنواع ال Cache

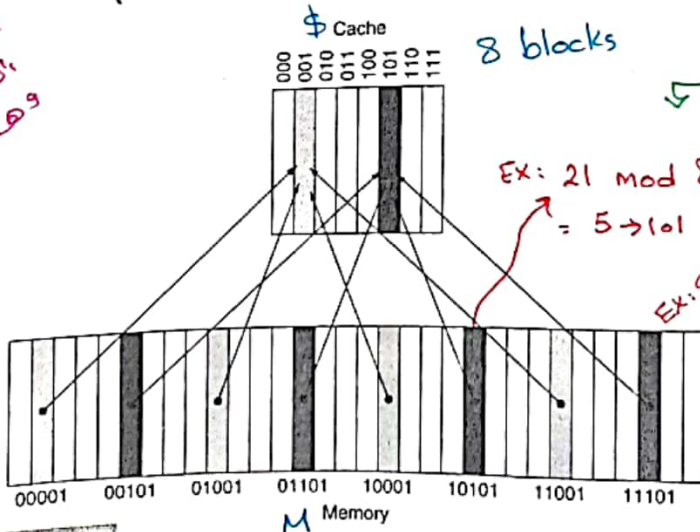
Location determined by address

هو بحد طمان ال block في ال Cache

Direct mapped: only one choice

(Block address) modulo (#Blocks in cache)

أي اشي بدنا بجيبه من ال Memory وهذا يعتمد على فقط address تبعه



block Address يتخذ ال blocks وينقسم على عدد ال Cache ويتخذ الباقي

#Blocks is a power of 2

Use low-order address bits

Least Significant bits

# Tags and Valid Bits

How do we know which particular block is stored in a cache location?

- Store block address as well as the data
- Actually, only need the high-order bits


Called the tag  
 high Address bit  
 يعني يحتاج اجيب ال bits المناقيلت بس لانه آخر 3 انا اخذتوم اصلا

لما نجيب block من ال Memory مش بس بنخزن ال data تاخده منخزن كمان ال Address تاخده

What if there is no data in a location?

- Valid bit: 1 = present, 0 = not present
- Initially 0

لما نخزن tag ولما ال Processor يطالع address ال Cache بنخص ال Address اللي طلبه مع ال tag الموجود بال Cache انا نتساوا معناته ال data اللي بيدنا ياها موجودة بالكاش انا ما نتساوا معناته في هنا Miss



## Cache Example

- 8-blocks, 1 word/block, direct mapped
- Initial state

< 32 bit > or < 1 w >

	Index	V	Tag	Data
0	000	N = 0		
1	001	N = 0		
2	010	N = 0		
3	011	N = 0		
4	100	N = 0		
5	101	N = 0		
6	110	N = 0		
7	111	N = 0		



# Cache Example

Word addr	Binary addr	Hit/miss	Cache block
22	10 110	Miss	110

بالعنوان

لأنه ما في باره cache block  
أشبه

Index	V	Tag	Data
000	N		
001	N		
010	N		
011	N		
100	N		
101	N		
110	Y	10	Mem[10110]
111	N		



data و وطبقها باره data  
Part Chapter 5 — Large and Fast: Exploiting Memory Hierarchy

و بعد من ال Most Significant bit  
في هون cache tag باره عنوانه بال مستقبل تعرفه ش

# Cache Example

Word addr	Binary addr	Hit/miss	Cache block
26	11 010	Miss	010

Index	V	Tag	Data
000	N		
001	N		
010	Y=1	11	
011	N		Mem[11010]
100	N		
101	N		
110	Y		
111	N	10	Mem[10110]



# Cache Example

Word addr	Binary addr	Hit/miss	Cache block
22	10 110	Hit	110
26	11 010	Hit	010

Index	V	Tag	Data
000	N		
001	N		
010	Y	11	Mem[11010]
011	N		
100	N		
101	N		
110	Y	10	Mem[10110]
111	N		

نفس ال  
طوله اللي  
بدنا يحاها

# Cache Example

Word addr	Binary addr	Hit/miss	Cache block
16	10 000	Miss	000
3	00 011	Miss	011
16	10 000	Hit	000

Index	V	Tag	Data
000	Y	10	Mem[10000]
001	N		
010	Y	11	Mem[11010]
011	Y	00	Mem[00011]
100	N		
101	N		
110	Y	10	Mem[10110]
111	N		

ال  
tag اللي فيه  
نفس اللي انا طالبي




# Cache Example

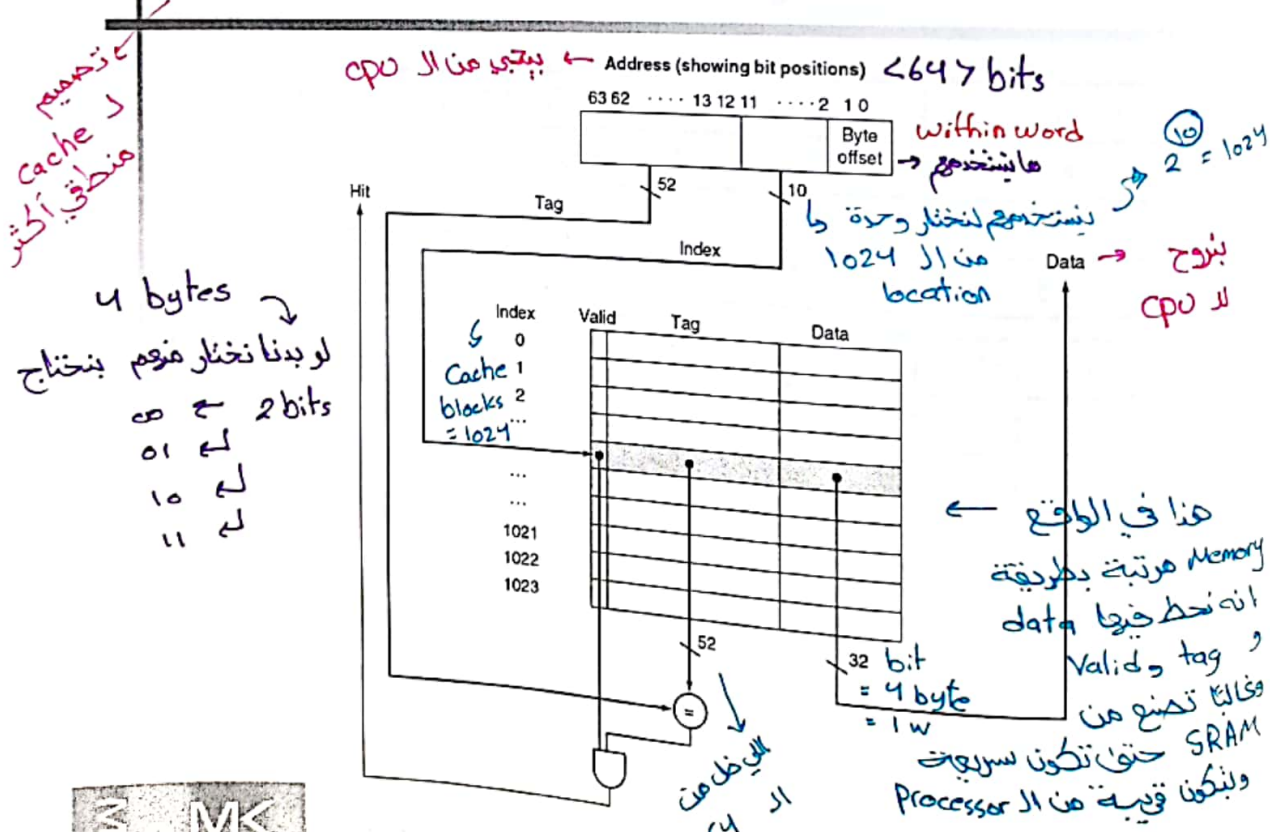
Word addr	Binary addr	Hit/miss	Cache block
18	10 010	Miss	010

لإظهار ما الذي موجود فيه مثل نفس الذي أنا طالبه

Index	V	Tag	Data
000	Y	10	Mem[10000]
001	N		
010	Y	10	Mem[10010]
011	Y	00	Mem[00011]
100	N		
101	N		
110	Y	10	Mem[10110]
111	N		



# Address Subdivision



# Example: Larger Block Size

أكبر من 1w

64 blocks, 16 bytes/block


To what block number does address 1200 map?

هذا الـ address كامل 1200 فيدي الأقي الـ block فيقسم على 16

Block address =  $\lfloor 1200/16 \rfloor = 75$

Block number =  $75 \text{ modulo } 64 = 11$

Location 11  
بالـ cache



54

< 64 bit >  
↓  
2<sup>6</sup> = 64

يمكنونا نختار

1 من 16

$$2^4 \cdot 16 = 16 \cdot 16 = 16_2 \cdot 2^4 = 4$$

## Block Size Considerations

لأنه لما يصير عندك Miss لـ 1w

Larger blocks should reduce miss rate

بنجيب block

Due to spatial locality

بسنتفيد من الـ spatial locality  
ورانه الـ Processor عالب حيطب الـ near by words

But in a fixed-sized cache

اللي قبلها أو بعدها فلما أجيب الـ block  
كامل بيخى عجلي Some misses بالمستقبل القريب

Larger blocks  $\Rightarrow$  fewer of them

More competition  $\Rightarrow$  increased miss rate

كلما أكبر الـ block  
الولد يصير عندي عدد blocks أقل

Larger blocks  $\Rightarrow$  pollution

يعني بتزيد عالـ cache بـ data  
ما بنحتاجها

Larger miss penalty

Can override benefit of reduced miss rate

Early restart and critical-word-first can help

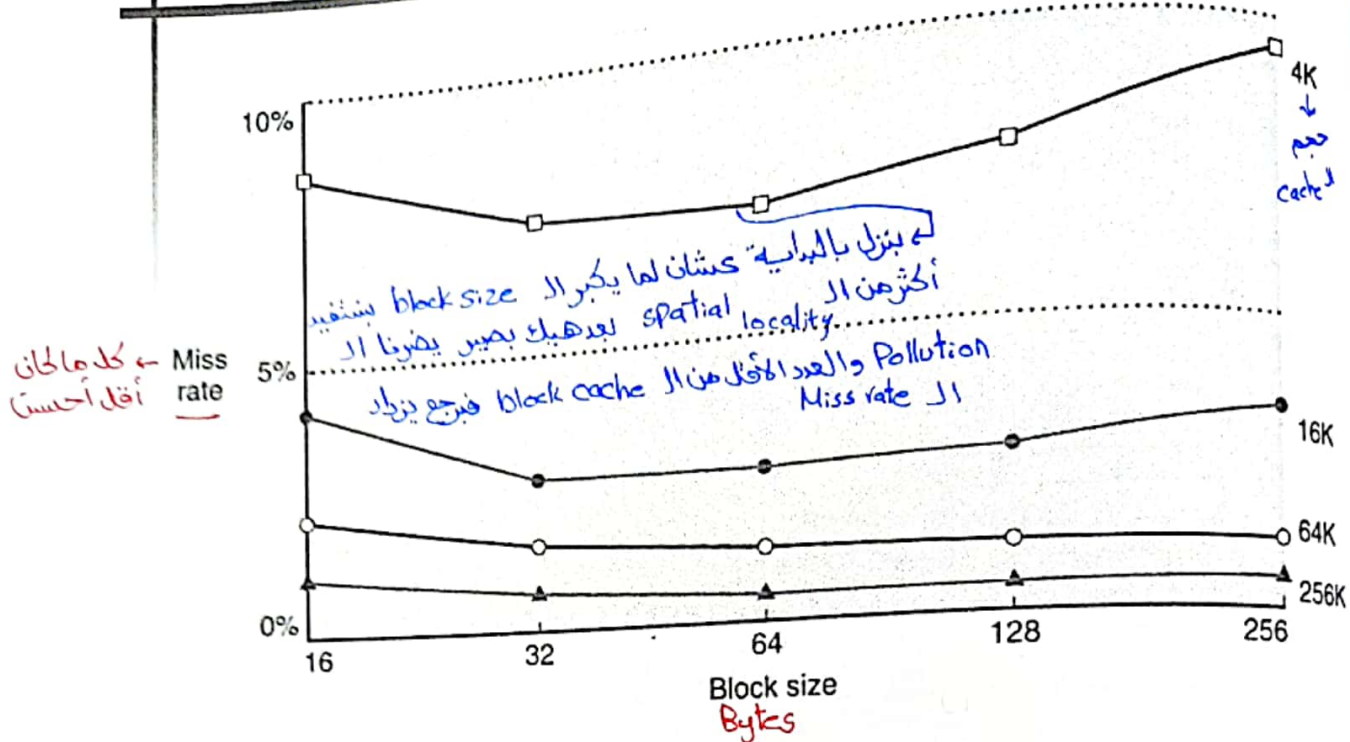
بمجرد ما توصل الـ critical word  
بمجرد ما توصل الـ critical word  
بمجرد ما توصل الـ critical word  
بمجرد ما توصل الـ critical word  
بمجرد ما توصل الـ critical word

حسب ترتيبها بالمثال بنسبت الـ address  
لما عول اول اشي بعدين بنطلب باقي الـ block



# Block Size Considerations

العلاقة بين حجم الcache وال block size وال Miss rate



كلما كان Miss rate أقل أحسن

له ينزل بالبرايه كشان لما يكبر ال block size بتفيد أكثر من ال spatial locality بعد هيك بصير يضربنا ال

ال Miss rate ال Pollution والعدد الأقل من ال block cache فيجوز زياد



## Cache Misses

- On cache hit, CPU proceeds normally
- On cache miss →
  - Stall the CPU pipeline
  - Fetch block from next level of hierarchy
  - Instruction cache miss
    - Restart instruction fetch
  - Data cache miss
    - Complete data access

بده يعطل ال Instruction او ال Store / load و data صا يلفظوا بتكون ال Miss

لما يكون في cache miss ال data ما بتكون موجوده بال cache وبتنا زمن أطول لتجيبوا فيعملوا

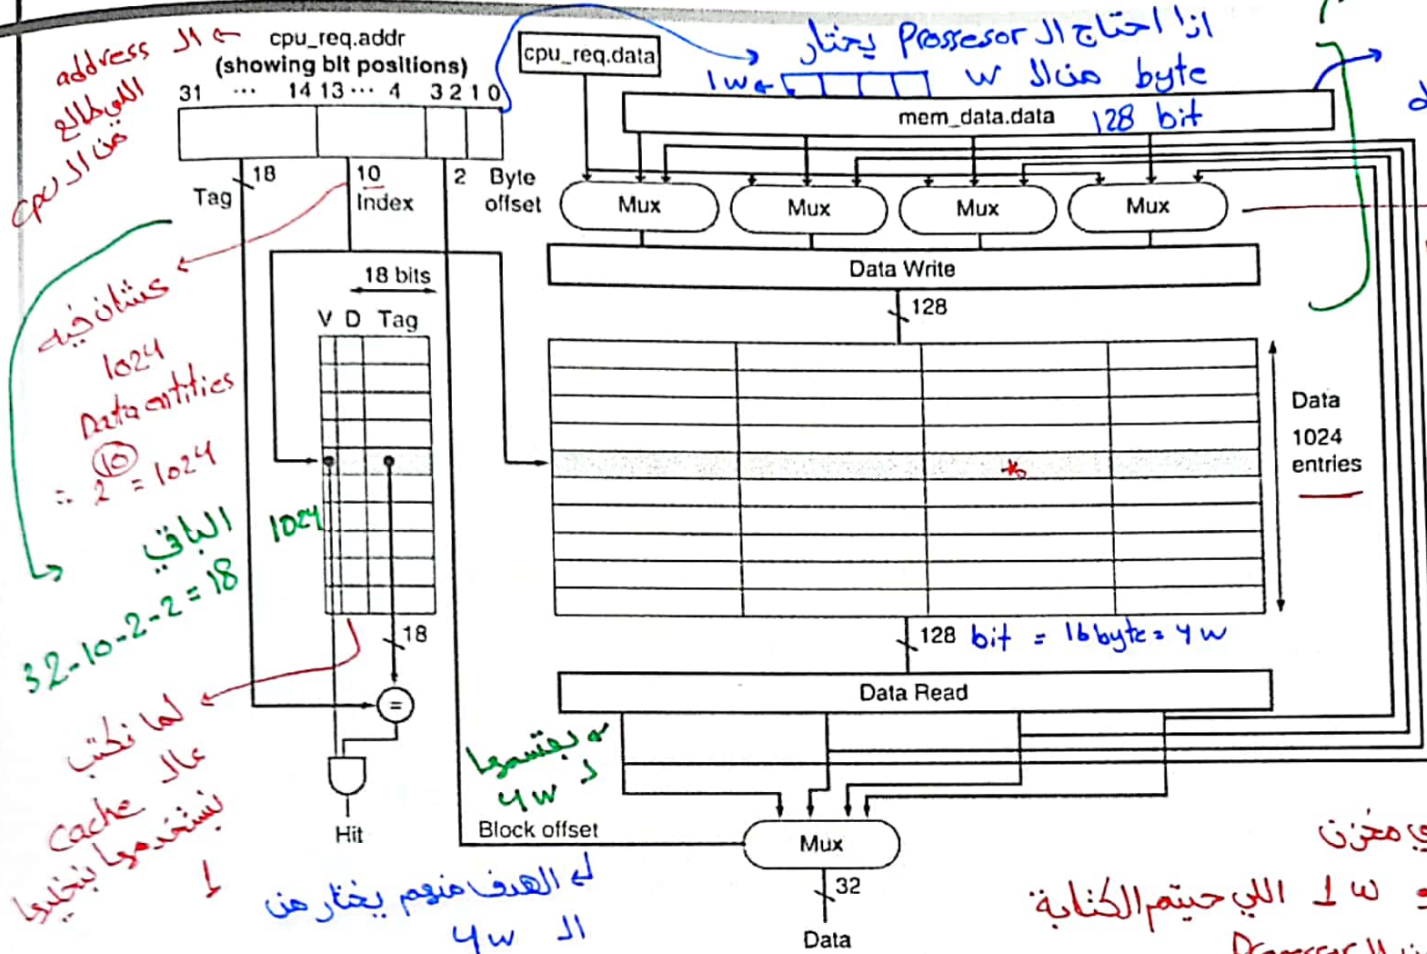
بطلب من اللى زحتمه العيت بطلب من اللى تحت تحت هيك

بطلب ال Instruction مرة تانيه



# تصميم Cache بنظرنا ونقرأ عليه

## Writing to the Cache



address ← ال  
المعالج  
من ال CPU

عشان خليه  
1024  
Data entities  
= 2<sup>10</sup> = 1024  
الباقي  
32-10-2-2=18  
لما نكتب  
عال Cache  
بنستخدموا بنخيلوا

تصميم Cache بنظرنا ونقرأ عليه

اذا احتاج ال Processor يختار  
1w byte  
منا ال Processor يختار

لما نكتب  
عال Cache  
بنستخدموا بنخيلوا

لما نكتب  
عال Cache  
بنستخدموا بنخيلوا

لما نكتب  
عال Cache  
بنستخدموا بنخيلوا

to write to the  
Cache  
بنستخدمها لما يغير في  
data فنجيب ال  
Memory Miss

عشان لما بننا  
نكتب حنكتب  
عجزة من ال block  
مش كله  
فبداية بقرا ال block  
كامل و يدخله  
عال MUX  
واللي ما تغيروا  
لدخلوا زي ما هم  
وال MUX المسوول  
عن ال word اللي  
لدها تغير

ف ش... باخدم من اللي مخزن  
بال Cache و ال اللي حيتم الكتابة  
كليوا بيجي من ال Processor

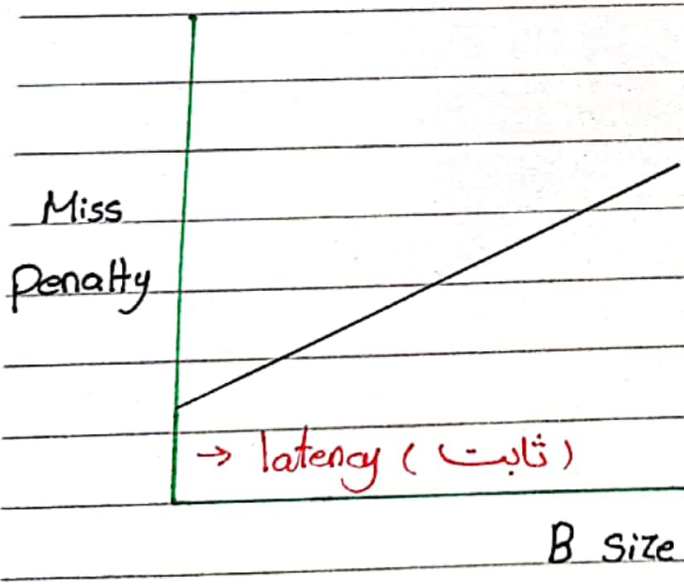

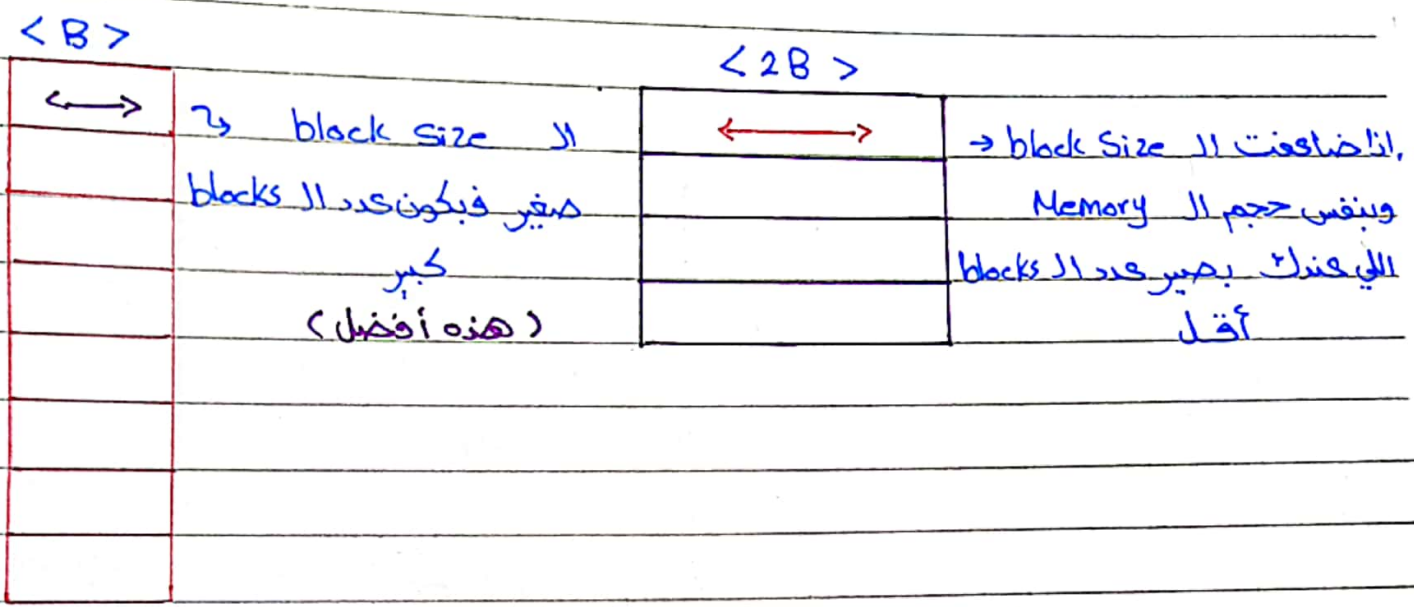
- \* الكتابة بال Cache ال 3 محاور 1- من ال Processor
- 2- من ال Memory
- 3- من ال Cache للاشي اللي ما به يتغير

## Write-Through

( في أكثر من طريقة للتخلد مع ال write )



Slide 43 : Block Size Considerations



# Write-Through

( في اكثر من طريقة للتغاط مع ال write )



هون بقالك يكتب عال cache  
بس . بس ما بياخذ بعين  
لا اعتبار انه ممكن  
تفقد ال data  
وتخبر معلومات  
ال \$ مختلفة  
عن معلومات  
ال M

Processors بده يكتب شي واد  
address اللي بده  
يطلب عليه  
موجود بالcache  
→ اول اصل

On data-write hit, could just update the block in cache

But then cache and memory would be inconsistent

Write through: also update memory

But makes writes take longer

- e.g., if base CPI = 1, 10% of instructions are stores, write to memory takes 100 cycles
  - Effective CPI =  $1 + 0.1 \times 100 = 11$

Solution: write buffer

- Holds data waiting to be written to memory
- CPU continues immediately
  - Only stalls on write if write buffer is already full

دخله يكتب عال cache  
وال write buffer لحد  
ما يتعطل اذا مالحق ال  
cache

write buffer على ال CPU وتبقى خيلا ما يترجم ال Memory





# Write-Back

مستخدمة أكثر الآن  
وهو التصميم الذي يريح  
ال Memory

write through

Alternative: On data-write hit, just update the block in cache → فقط بنكتب على block التي بال Cache

- Keep track of whether each block is dirty  $0=1$

When a dirty block is replaced

- Write it back to memory
- Can use a write buffer to allow replacing block to be read first

متيح بس الحساب

كشأن أمل مشكله البطء



# Write Allocation

data مش موجودة بال Cache

What should happen on a write miss?

Alternatives for write-through

- Allocate on miss: fetch the block
- Write around: don't fetch the block (No Allocate)
  - Since programs often write a whole block before reading it (e.g., initialization)

For write-back

- Usually fetch the block

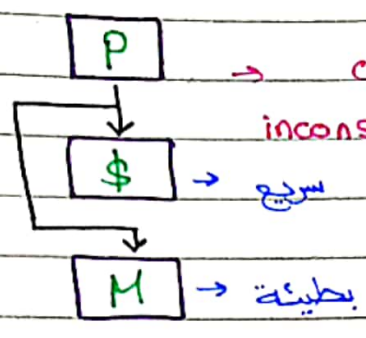
بموجود بس مش  
write Allocate







Slide 47 : Write through :

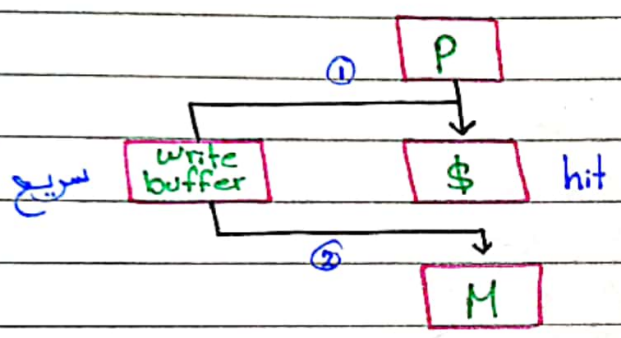


لما ال Processor يكتب يكتب على ال cache  
ويكتب على Memory فيكون ال cache inconsistent

→ سريع

→ بطيئة

\* Write buffer :



\* لما يكتب على ال M مباشرة، أنا بيأول  
ال data ال Memory وبديش يالها تضعي علي  
ال data اللي بيدي يالها تروح ال Memory بعت  
نسخة من الي حكتك بال Cache ال write buffer  
وهي عبارة عن Memory سريعة في حد ذاتها


ال Processor بكون نشغله و ال write buffer كايحتاجه يكتب بال Memory و هاد بخلنا  
مشكلة ال Performance

\* ال write buffer لازم يكون كبير كفاية إذا ال Processor بعت Many store instructions  
وذا بعض يتخزنوا بال write buffer ومع الوقت يترحلوا ال Memory

\* بترجع ال Memory لأنه ال Processor بيشتغل كثير وينفذ Many instructions وكل  
store instructions ياللي به ينفذهم راح يوحلوا ال Memory فترجع ال Memory بال  
accesses و فينروح ال Write-Back

Slide 48 : write-Back :

on hit → Processor write just on cache ( سريع )

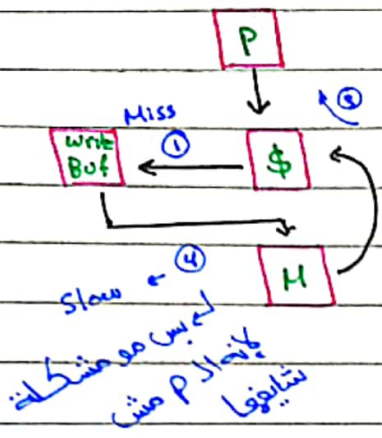


\* لما يصير في Miss و به يعامل replacement : يعني به الاشئ اللي يجيبه  
من ال Mem بتبدل ال block موجود بال cache وفيه data  
فمجانا ما تضعي ال Data قبلما جاس ال replacement أنقل  
ال block اللي بال cache و عليه ال data جديدة replacement  
ال Memory

\* بال Miss تكون بطيئة لأنها بتعمل replacement و ال M بطيئة جدًا

Slide 48 :

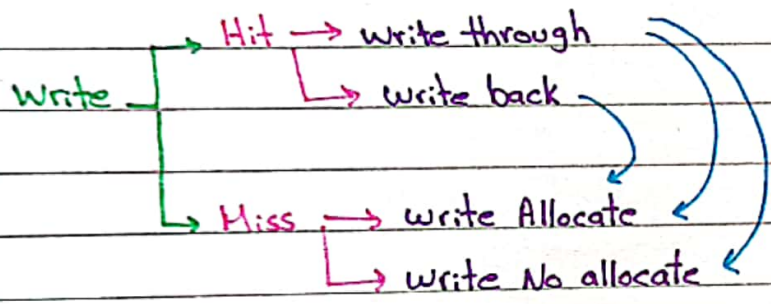
\* لأحد مشكلة البطء في ال Write-Back :



① لو كان hit بدل يقرأ ويكتب من ال cache بسرعة ما يروح ال M  
② إن صار Miss كخطوة أول بتنقل ال block الي بيينا نعمل replacement  
ال ال write buffer من ال cache وهاي عملية سريعة  
بجدين بتعمل replacement بعد هيك ال P خلص بياني  
اللي بيها . مع بطيئة لكن هو مشكلة  
بجدين علفنا راحتنا بتعمل write back

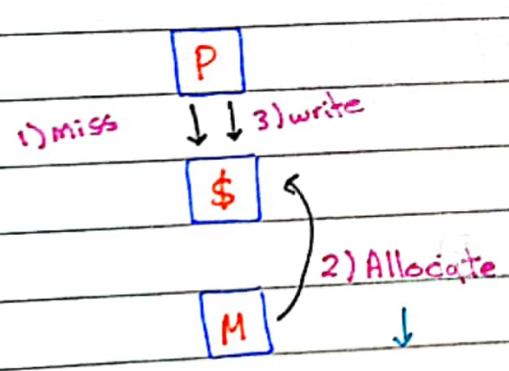
Slide 49 : write Allocate:

read-miss- Allocate

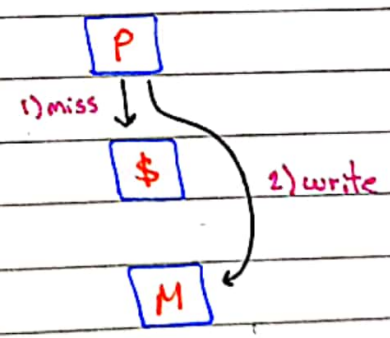


Allocate

No Allocate



لما يصير في Miss بتجيب ال data من ال Mem بجدين يتكتب عليها



لما يصير في Miss ال data هو موجودة بال M والمقر النهائي تبعها هو ال Memory فيكتب مباشرة بال Memory

\* الأفضل ال Allocate لأنه بتلغيني أستفيد من ال spatial locality  
\* لمرحلة الأوط بتفكر ال No Allocate لأنه أسرع بالبداية . بس ال Allocate عنده بعد نظر



# Contents

5.1 Introduction

5.2 Memory Technologies

5.3 The Basics of Caches

5.4 Measuring and Improving Cache Performance

← كيف نحسب ← Measuring Cache Performance

↳ Memory Average Access Time

Associative Caches

Multi-level Caches

Interactions with Advanced CPUs

Interactions with Software



## Measuring Cache Performance

### Components of CPU time

سوء طانت  
miss  
hit si  
F I A M W C

① Program execution cycles → الزمن اللي يقضوه الinstructions داخل ال Pipeline (F I E W C)

Includes cache hit time

② Memory stall cycles

Mainly from cache misses → نسيبم Missis

With simplifying assumptions:

كل ما كان اولي كلها ال Stall اعلى

Memory stall cycles

Memory accesses

$$= \frac{\text{Memory stall cycles}}{\text{Memory accesses}} \times \text{Miss rate} \times \text{Miss penalty} \quad (1)$$

Missis  
hits + Missis

كل ما كانت اعلى كل ما ال Stall اعلى

$$= \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Misses}}{\text{Instruction}} \times \text{Miss penalty} \quad (2)$$

له من بين كلا اللي نوزناه قويه في Missis

حسب الموطيات بختار اي معادلة استخدم



# Cache Performance Example

## Given

- I-cache miss rate = 2%
- D-cache miss rate = 4%
- Miss penalty = 100 cycles → *انا كان في Miss بلخظت 100*
- Base CPI (ideal cache) = 2
- Load & stores are 36% of instructions

## Miss cycles per instruction

- I-cache:  $0.02 \times 100 = 2$
- D-cache:  $0.36 \times 0.04 \times 100 = 1.44$

Actual CPI =  $2 + 2 + 1.44 = 5.44$

- Ideal CPU is  $5.44/2 = 2.72$  times faster



## Average Access Time

- <sup>→ fast</sup> Hit time is also important for performance
- Average memory access time (AMAT)

- $AMAT = \text{Hit time} + \text{Miss rate} \times \text{Miss penalty}$

## Example

- CPU with 1ns clock, hit time = 1 cycle, miss penalty = 20 cycles, I-cache miss rate = 5%

- $AMAT = 1 + 0.05 \times 20 = 2\text{ns}$

- 2 cycles per instruction

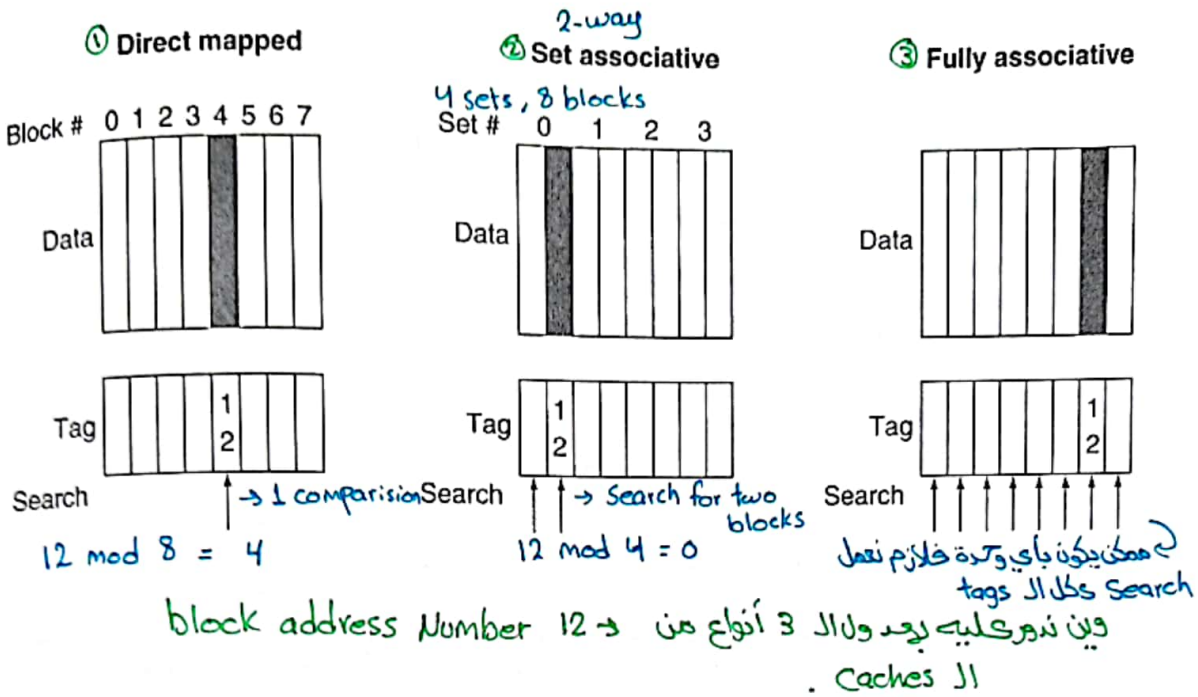
*→  $f = \frac{1}{T} = 1\text{ GHz}$*

*miss time = hit time + miss penalty  
= 1 + 20 = 21*





# Associative Cache Example



## Spectrum of Associativity

→ Associativity يعني قابلية ممكن تزيد

For a cache with 8 entries = 8 blocks

One-way set associative (direct mapped)

Block	Tag	Data
0		
1		
2		
3		
4		
5		
6		
7		

Two-way set associative

Set	Tag	Data	Tag	Data
0				
1				
2				
3				

$8/2 = 4 \text{ sets}$

Four-way set associative

Set	Tag	Data	Tag	Data	Tag	Data	Tag	Data
0								
1								

→  $8/4 = 2 \text{ sets}$

Eight-way set associative (fully associative)

Tag	Data	Tag	Data	Tag	Data	Tag	Data	Tag	Data	Tag	Data	Tag	Data

$8/8 = 1 \text{ sets} = \text{Fully Associative}$





# Associativity Example

- Compare 4-block caches
  - Direct mapped, 2-way set associative, fully associative
  - Block access sequence: 0, 8, 0, 6, 8 cpu → طالعزم جاهزات
- Direct mapped

Block address	Cache index	Hit/miss	Cache content after access			
			0	1	2	3
0	0	miss	Mem[0]			
8	0	miss	Mem[8]			
0	0	miss	Mem[0]			
6	2	miss	Mem[0]		Mem[6]	
8	0	miss	Mem[8]		Mem[6]	

↓  
 Block Address mod 4 = بطاقت من  
 وبتعديب لانه في 5 accesses  
 5 missis



التصنيف  
 Performance

# Associativity Example

- 2-way set associative

Block address	Cache index	Hit/miss	Cache content after access	
			Set 0	Set 1
0	0	miss	Mem[0]	
8	0	miss	Mem[0]	Mem[8]
0	0	hit	Mem[0]	Mem[8]
6	0	miss	Mem[0]	Mem[6]
8	0	miss	Mem[8]	Mem[6]

↓  
 Block Address mod 2 = بطاقت من

↳ 4 missis

- Fully associative

Block address	Hit/miss	Cache content after access	
		one set	
0	miss	Mem[0]	
8	miss	Mem[0]	Mem[8]
0	hit	Mem[0]	Mem[8]
6	miss	Mem[0]	Mem[6]
8	hit	Mem[0]	Mem[8]

لم كل ال cache مفتوح ابى فيغير استخدم ابى index

↳ 3 missis



# How Much Associativity

Increased associativity decreases miss rate → Performance بالتالي التحسن بغيره

But with diminishing returns

Simulation of a system with 64KB D-cache, 16-word blocks, SPEC2000

- Direct Mapped cache ←
- 1-way: 10.3%
  - 2-way: 8.6% ) 1.7%
  - 4-way: 8.3% ) 0.3%
  - 8-way: 8.1% ) 0.2%
- Miss Rate


16x4 = 64 bytes  
 cache blocks =  $\frac{64 \text{ KB}}{64} = 1 \text{ K} = 1024 \text{ block}$

\* مش دائما بنسفير لما نروح ل higher Associativity



15 = 4 blocks → 256 sets / (4-way) → 4 comparators لأنه في

## Set Associative Cache Organization



\* لو صار Miss فنروح نتجيب ال data من ال Memory وبتخطوا بلا block

مستحيل يجينا 1 على الأربعة بس يكون 1 على واحد منهم أو كلهم أحضار

256 x 4 = 1024 block  
 موزين على 256 set

بتطلع تسمى اي block hit مار عندها hit





# Replacement Policy

- Direct mapped: no choice
- Set associative
  - Prefer non-valid entry, if there is one
  - Otherwise, choose among entries in the set
- Least-recently used (LRU) الاشي التي من زمان ما استخدمت  
في الذاكرة بنقلها
  - Choose the one unused for the longest time
  - Simple for 2-way, manageable for 4-way, too hard beyond that له أو أكثر بصيرال hardware انجب
- Random غير منظمة
  - Gives approximately the same performance as LRU for high associativity



## Multilevel Caches

- Primary cache attached to CPU = L1
  - Small, but fast
- Level-2 cache services misses from primary cache = L2 = Secondary
  - Larger, slower, but still faster than main memory
- Main memory services L-2 cache misses
- Some high-end systems include L-3 cache Secondary



# Multilevel Cache Example

## Given

- CPU base CPI = 1, clock rate = 4GHz  $\rightarrow \frac{1}{f} = 0.25 \text{ ns}$
- Miss rate/instruction = 2%
- Main memory access time = 100ns = 400 cycles

## With just primary cache

- Miss penalty =  $100\text{ns}/0.25\text{ns} = 400 \text{ cycles}$
- Effective CPI =  $1 + 0.02 \times 400 = 9$



## Example (cont.)

- Now add L-2 cache
  - Access time = 5ns
  - Global miss rate to main memory = 0.5%
- Primary miss with L-2 hit
  - Penalty =  $5\text{ns}/0.25\text{ns} = 20 \text{ cycles}$
- Primary miss with L-2 miss
  - Extra penalty = 500 cycles  $\rightarrow$  400 cycles
- $\text{CPI} = 1 + 0.02 \times 20 + 0.005 \times 500 = 3.9$
- Performance ratio =  $9/3.9 = 2.3$

↳ level of cache<sup>1</sup>      ↳ 2 level of cache





# Multilevel Cache Considerations

## Primary cache

- Focus on minimal hit time

## L-2 cache

- Focus on low miss rate to avoid main memory access
- Hit time has less overall impact

## Results

- L-1 cache usually smaller than a single cache
- L-1 block size smaller than L-2 block size

Miss Rate  
↓  
Memory



Slide 55 : AMAT:

Find the average Memory Access time of a Memory hierarchy with the following specifications.

Memory Level	Hit Time	Miss Rate
L1 cache	1 cycle	6.0%
L2 cache	10 cycles	4.0%
L3 cache	20 cycles	2.0%
Main Memory	250 cycles	0%

```

    graph TD
      P[P] --> L1[L1]
      L1 --> L2[L2]
      L2 --> L3[L3]
      L3 --> M[M]
    
```

$AMAT = Hit\ time + Miss\ rate \times Miss\ Penalty$

هو عبارة عن وقت الوصول للبيانات في الـ Main Memory  
 ما عليه Miss فال average = 250 فال  
 بنا نروح من L3 الـ Memory و نحتاج 250  
 فال Miss penalty لـ L3 = 250 ، فال السؤال جيبين من الـ أسئلة الأخرى .

$AMAT_{L3} = 20 + 0.02 \times 250 = 25\ cycles$

$AMAT_{L2} = 10 + 0.04 \times 25 = 11\ cycles$

$AMAT_{L1} = 1 + 0.06 \times 11 = 1.66\ cycles$

∴ AMAT = 1.66 → processor التي بتستوفها الـ

Slide 64 : LRU Replacement Policy :

4-way associative cache, initially empty

Address	1	2	3	4	1	5	3	6
MRU	1	2	3	4	①	5	③	6
		1	2	3	4	1	5	3
			1	2	3	4	1	5
LRU				1	2	3	4	1

Information maintained for LRU replacement for each set

له يعني بتحتاج 4 register فيهم بين  
 الـ MRU بين الـ LRU

\* MRU = Most recently used

\* LRU = Least Recently used



# Interactions with Advanced CPUs

*dynamic execution*  
Out-of-order CPUs can execute instructions during cache miss →

- Pending store stays in load/store unit
- Dependent instructions wait in reservation stations
  - Independent instructions continue

Effect of miss depends on program data flow →

■ Much harder to analyse

■ Use system simulation →

الاعتماد عليه بالوقت الحاضر ليصرف سريع ولا لأنه بالمعادلات صعب نطلع النتيجة

هدول  
مار فيوم  
Store  
add  
sub

له هاي ما بتعتمد عليهم  
فبنفذوها ما بتستنى

نتيجه





# Interactions with Software

مختلفة Algorithms 2

Misses depend on memory access patterns

- Algorithm behavior
- Compiler optimization for memory access



For large data Radix ال بنين افضل ولكن لما نحسب ال clock cycle ال بطول ال Quick Sort احسن وفعلاً هو الامن (نعم انه يحتاج عدد Instructions اكثر)

ال Radix ال كثير cache miss فيستخدم ال large Memory عشان هيك بطيئ

يستخدم حجم اقل من ال Mem وال cache miss اقل

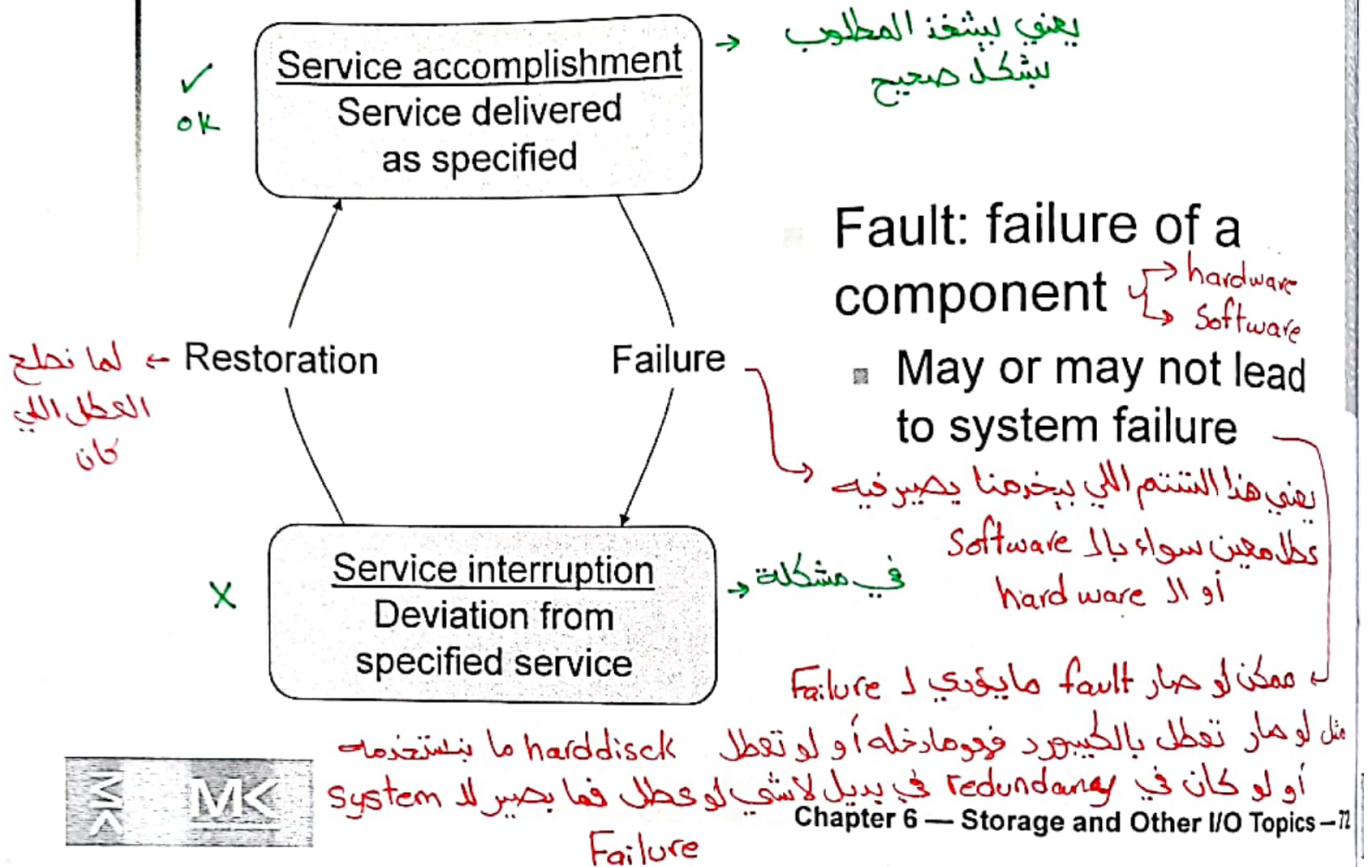
ال Radix ال large Memory وبعني فيه ال data الي بييجي بعدين لما يجي به يقرأها بقراها بشكل متعجب من مكان امكن ولانه ال cache ما بوسعك كل حالاتها بحير عنده cache miss. بينما ال quick فيه شغل على منطقة منطقة at a time وفيه منه better Memory access

## Contents

- 5.1 Introduction
- 5.2 Memory Technologies
- 5.3 The Basics of Caches
- 5.4 Measuring and Improving Cache Performance
- 5.5 Dependable Memory Hierarchy → يعني الي بنحطه بال Mem نرد نقدر نقرأه
  - Dependability
  - Error Correction Codes



# Dependability → نفي هذا المستمر تشغيل ولا خراب



## Dependability Measures

- معتمد عليها ←
- Reliability: mean time to failure (MTTF)
  - Service interruption: mean time to repair (MTTR)
  - Mean time between failures
    - MTBF = MTTF + MTTR
  - Availability =  $\frac{MTTF}{MTTF + MTTR}$ 

تضمنت availability تكون واحد = 100% ليس أحسن لتستم يكون إليه 99.999...%
  - Improving Availability
    - Increase MTTF: fault avoidance, fault tolerance, fault forecasting
    - Reduce MTTR: improved tools and processes for diagnosis and repair
- منتج ←

# The Hamming SEC Code → تصحيح الأخطاء

Hamming distance → بقدر أمثل الأرقام

Number of bits that are different between two bit patterns

تبتك  
شكوك  
صعينة

0 0 0 0  
1 0 0 1 → different = 1 bit  
2 0 1 0 → " = 2 bit

Minimum distance = 2 provides single bit error detection → بقدر أمثل الأرقام واكتشف الخطأ

E.g. parity code

Minimum distance = 3 provides single error correction, 2 bit error detection

بقدر أمثل الأرقام واكتشف الخطأ واحده



## Encoding SEC

To calculate Hamming code:

- Number bits from 1 on the left
- All bit positions that are a power 2 are parity bits
- Each parity bit checks certain data bits:

Bit position	1	2	3	4	5	6	7	8	9	10	11	12
Encoded data bits	p1	p2	d1	p4	d2	d3	d4	p8	d5	d6	d7	d8
odd → p1	X		X		X		X		X		X	
Parity bit → p2		X	X			X	X			X	X	
coverate → p4				X	X	X	X					X
p8								X	X	X	X	X

ببعضها في شكل  
لوح





# Decoding SEC

→ Single Error Correction

Value of parity bits indicates which bits are in error

- Use numbering from encoding procedure
- E.g.
  - Parity bits = 0000 indicates no error
  - Parity bits = 1010 indicates bit 10 was flipped

## Example:

- What will be stored for 1001 1010?  
 (Note: 9 and 10 are written above the underlined bits)
- If you read 0111 0010 1110, is there error? Correct it.



# SEC/DED Code

→ Single Error Correction / Double Error Detection

Add an additional parity bit for the whole word ( $p_n$ ) →

Parity Bit

Make Hamming distance = 4

Decoding: Single Errors يكونوا Single Errors الـ بقدر انصحاليه  
 Double errors يكونوا Double errors الـ بقدر انصحاليه

Let  $H = \text{SEC parity bits}$

- $H = 0, p_n$  even, no error
- $H \neq 0, p_n$  odd, correctable single bit error
- $H = 0, p_n$  odd, error in  $p_n$  bit
- $H \neq 0, p_n$  even, double error occurred

Correctable 1 error ←  
 لأنه H باش مكانه  
 ما بقدرنا P ما بقدرنا  
 بومنا الـ data تكون صحه  
 ما بقدرنا نصلح الـ طابوع  
 سب بقدرنا نصح

8 Parity Bit  
 $7 = H$   
 $1 = P_n$

ECC DRAM uses SEC/DED with 8 bits protecting each 64 bits

→ Error correction code



# Contents

- 5.1 Introduction
- 5.2 Memory Technologies
- 5.3 The Basics of Caches
- 5.4 Measuring and Improving Cache Performance
- 5.5 Dependable Memory Hierarchy
- 5.11 Redundant Arrays of Inexpensive Disks = RAID
- 5.6 Virtual Machines
- 5.7 Virtual Memory
- 5.8 A Common Framework for Memory Hierarchy
- 5.9 Using a Finite-State Machine to Control a Simple Cache
- 5.10 Cache Coherence
- 5.13 The ARM Cortex-A53 and Intel Core i7 Memory Hierarchies
- 5.16 Fallacies and Pitfalls
- 5.17 Concluding Remarks



## RAID

- Redundant Array of Inexpensive (Independent) Disks
  - Use multiple smaller disks (c.f. one large disk)
  - Parallelism improves performance → كشان ال data تكون موزعة بأكثر من disk
  - Plus extra disk(s) for redundant data storage
- Provides fault tolerant storage system → شتمل الأنظمة وما يتأثر فيها
  - Especially if failed disks can be “hot swapped” ↓ وهو ال Server شغال بتقدر إنك تشيل ال component المعطلة وتحط واحد جديد
- ① RAID 0
  - No redundancy (“AID”?)
    - Just stripe data over multiple disks
  - But it does improve performance





# RAID 1 & 2

- RAID 1: Mirroring
  - N + N disks, replicate data
    - Write data to both data disk and mirror disk
    - On disk failure, read from mirror
- RAID 2: Error correcting code (ECC)
  - N + E disks (e.g., 10 + 4)
  - Split data at bit level across N disks
  - Generate E-bit ECC → P, P<sub>2</sub>, P<sub>4</sub>, P<sub>8</sub>
  - Too complex, not used in practice



# RAID 3: Bit-Interleaved Parity

- N + 1 disks
  - Data striped across N disks at byte level
  - Redundant disk stores parity
  - Read access
    - Read all disks → البيانات موزعة فلازم نقرأ من كل الـ disks
  - Write access
    - Generate new parity and update all disks
  - On failure → لو تعطل واحد يجيبه من الـ Parity  
البيانات اللي
  - Use parity to reconstruct missing data
- Not widely used



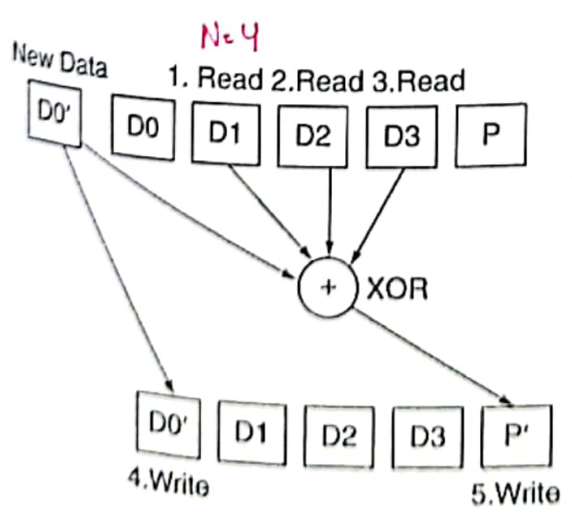
# RAID 4: Block-Interleaved Parity

- N + 1 disks
  - Data striped across N disks at block level
  - Redundant disk stores parity for a group of blocks
  - Read access
    - Read only the disk holding the required block
  - Write access
    - Just read disk containing modified block, and parity disk
    - Calculate new parity, update data disk and parity disk
  - On failure
    - Use parity to reconstruct missing data
- Not widely used

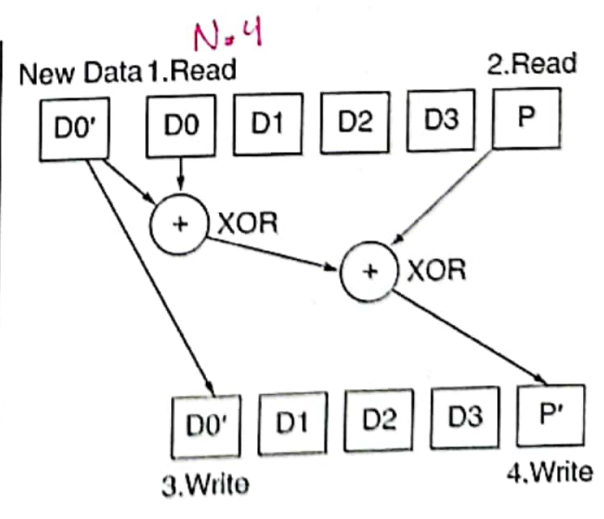
\* الأوفر RAID 3 / RAID 4 من ال RAID 2 لإنه أرخص

## RAID 3 vs RAID 4

مستخدم أكثر من RAID 3



N-1 reads  
2 writes

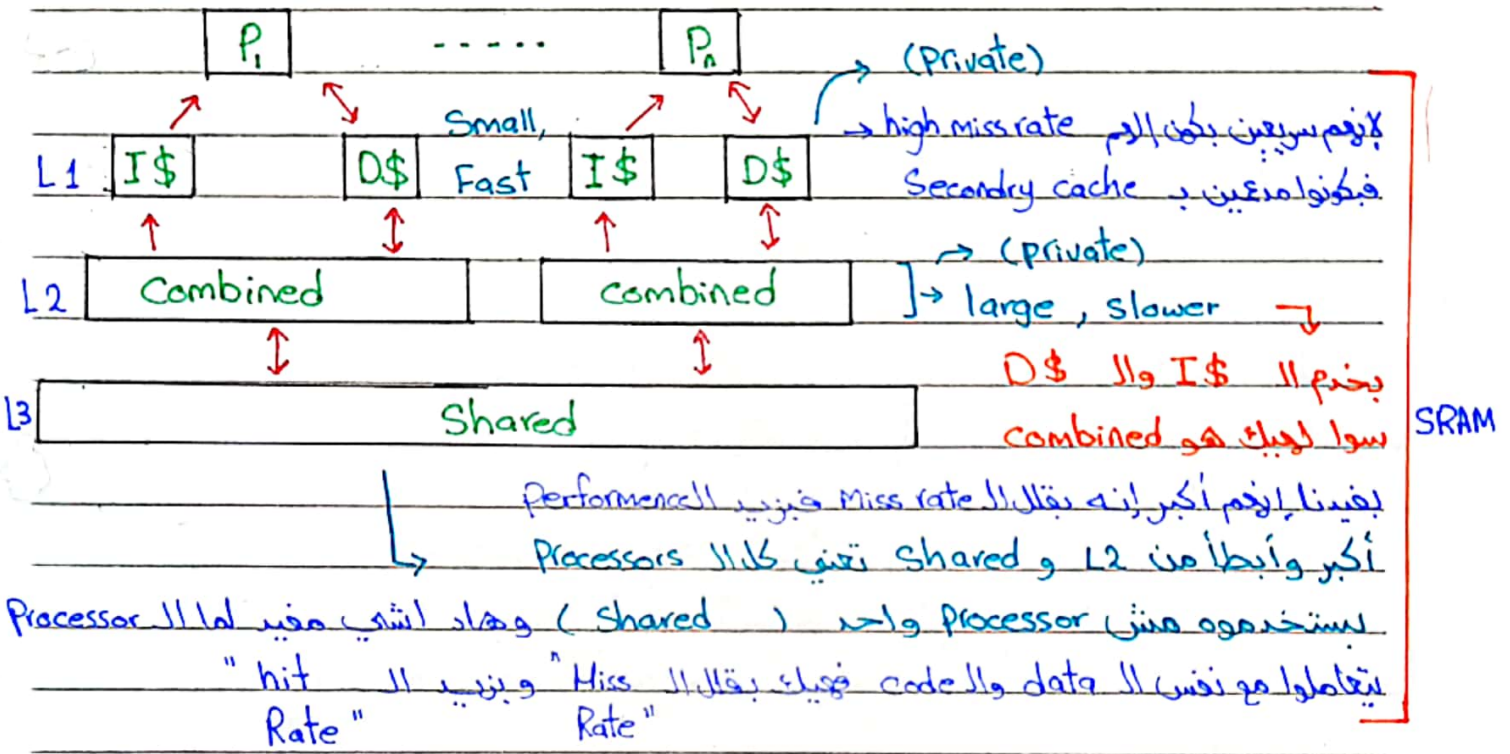


2 reads  
2 writes

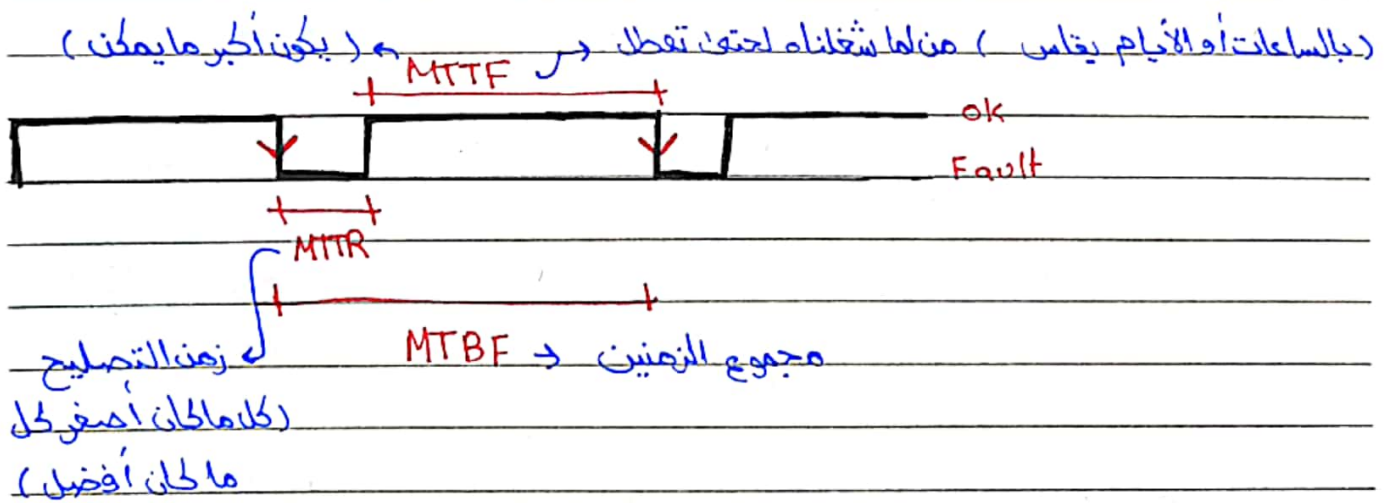


### Slide 68 : Multilevel Caches

\* microprocessor with multiple cores.



### Slide 73 : Dependability Measures :



### slide 73 8 Improving Availability.

#### ① Increase MTTF:

- **Fault Avoidance:** من البداية أتجنب الخطأ بإزالة software و hardware يكونوا أفضل ما يمكن ومن الصعب أحصل عليها لأننا غالباً جداً بنكون

- **Fault tolerance:** يعني لو تعطل عندي اشي يكون في الـ backup عشان ما نوقف الشغل مثل اننا يكون فيه هروبين تبريد مثلاً .

- **Fault Forecasting:** أنتبأ بالأعطال اللي ممكن تصير وبعمل علفا تبديل القطعة لها أريف، إنه حيلنا عمرهم أو حيلنا .

#### ② Reduce MTTR:

أحسن الـ tools عشان تنبهني بالأخطاء وتشخصها وبتساعدني: 1- يكون عندي مكان أدخل عليه العطل وأحده بيخبر ما أثره الباقي 2- نزيد أعداد المصالحين 3- نستعمل harddisk ثاني فلانم يكون عندي قطع غير متأول احتياطاً .

### Slide 74 8 Hamming Distance :

Distance = 1


#### Example Binary

0 00

1 01

2 10

3 11



Parity

Ex: 011 -> 0 ✓

010 -> 1 ✗

Distance = 2

#### Example Binary with Parity

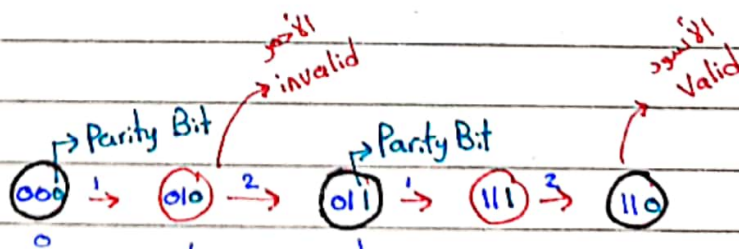
0 00 0

1 01 1

2 10 1

3 11 0

even ←



وهون نفس المبدأ

من 0 لـ 1 بنحتاج نعمل تغييرين هي

أول تغيير وثاني تغير تغير الـ Parity

من 0 ← 1

لو عدد الـ زوجي بنحط

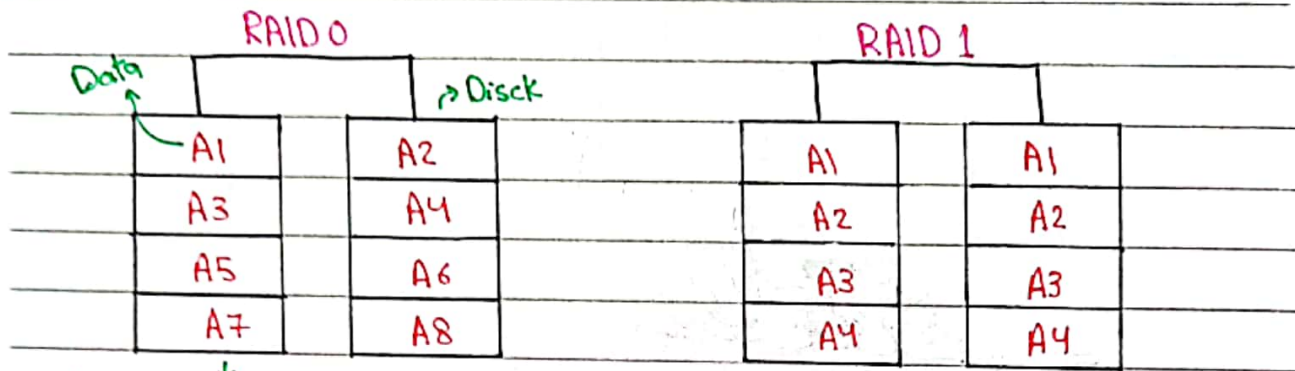
0 ولو زوجي بنحط 1





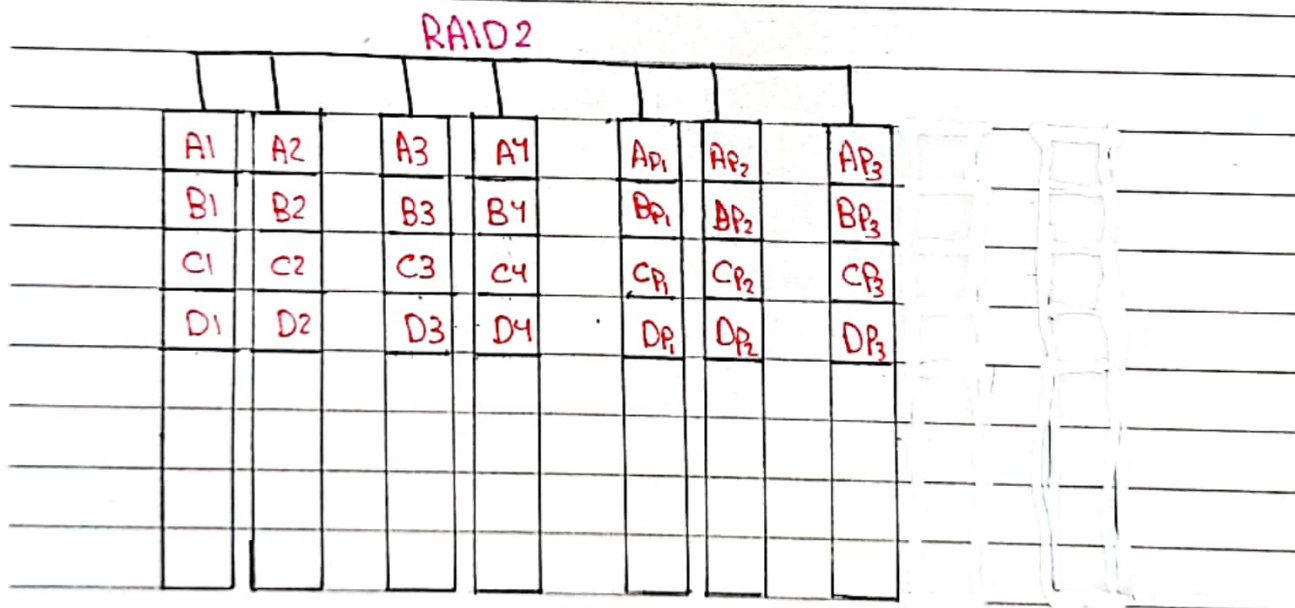


Slide 79 & RAID



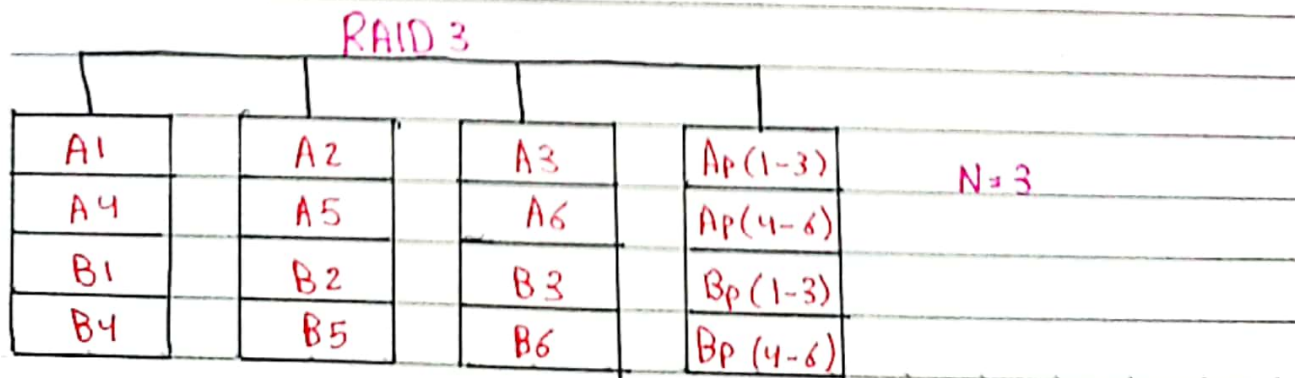
لو هاد تابل فيتكون  
Data مكتوب على  
البي عليه

\* لازم يكون عدد زوجي من ال disk  
\* لو تطلب واحد بقا من ال Mirror ناسه  
\* ال بيتطلب بسبب انه و ينقل عليه ال data



Disck1 Disck2 Disck3 Disck4 Disck5 Disck6

\* غالي و صعب ال hardware



## RAID 4

	A	B	XOR
A1	0	0	0
B1	0	1	1
C1	1	0	1
D1	1	1	0

إذا في تغير الح زكس ال Parity

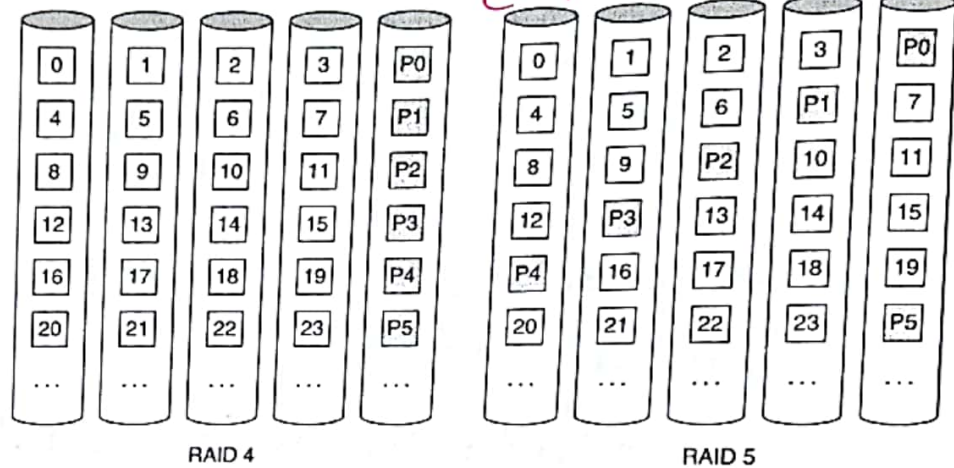
و بنستخدم ناتج ال XOR مع ال Parity القيمة ونحسم XOR بطل ال Parity الجديدة.

# RAID 5: Distributed Parity

( أفضل واحد في الوقت الحاضر )

- N + 1 disks
  - Like RAID 4, but parity blocks distributed across disks
    - Avoids parity disk being a bottleneck
- Widely used


مباروا كلهم زي بعض والباريتي موزع



# RAID 6: P + Q Redundancy

كشان لو بالصفحة تعطل disks 2 مانزوح علينا وبمناقلنا ييجو الذااتا

- N + 2 disks
  - Like RAID 5, but two lots of parity
  - Greater fault tolerance through more redundancy
- Multiple RAID → له مؤسسات الكبيرة اللي الوا اكثر منافع
- More advanced systems give similar fault tolerance with better performance
- Example RAID 51





# RAID Summary

RAID can improve performance and availability

- High availability requires hot swapping → *لأنه بدل ما يكون عندك Disk واحد يكون أكثر من Disk*

Assumes independent disk failures

- Too bad if the building burns down!

*إذا تعطلت اشياء ما بوقف الجواز وبتبدل المعطل بجديد  
إذا تعطل واحد من البقية ما بتعطلوا  
إذا احترقت البناية واح كمدال RAID والاحل تروح لل Multiple RAID مثل 51 RAID*

## Contents

- 5.1 Introduction
- 5.2 Memory Technologies
- 5.3 The Basics of Caches
- 5.4 Measuring and Improving Cache Performance
- 5.5 Dependable Memory Hierarchy
- 5.11 Redundant Arrays of Inexpensive Disks
- 5.6 Virtual Machines → *العلاقة ب cloud*
- 5.7 Virtual Memory
- 5.8 A Common Framework for Memory Hierarchy
- 5.9 Using a Finite-State Machine to Control a Simple Cache
- 5.10 Cache Coherence
- 5.13 The ARM Cortex-A53 and Intel Core i7 Memory Hierarchies
- 5.16 Fallacies and Pitfalls
- 5.17 Concluding Remarks

يستخدم مثلا لو بي ايشغل برنامج ال operating system ال مختلف واللي عندي

# Virtual Machines

ليقلد او يجالي

Host computer emulates guest operating system and machine resources

ما يشغل كل ال computer

بسا ال cores

اللي حطت

عليهم

Improved isolation of multiple guests

Avoids security and reliability problems

Aids sharing of resources → بنظم كل App قديه ياخذ من ال resources

Virtualization has some performance impact

Feasible with modern high-performance computers

Examples

1. IBM VM/370 (1970s technology!)

2. VMWare من شركة oracle

3. Microsoft Virtual PC

كل VM ما يتشوف الا

ملفاتنا ولو

توصلت

ما بنحرب

غيرها



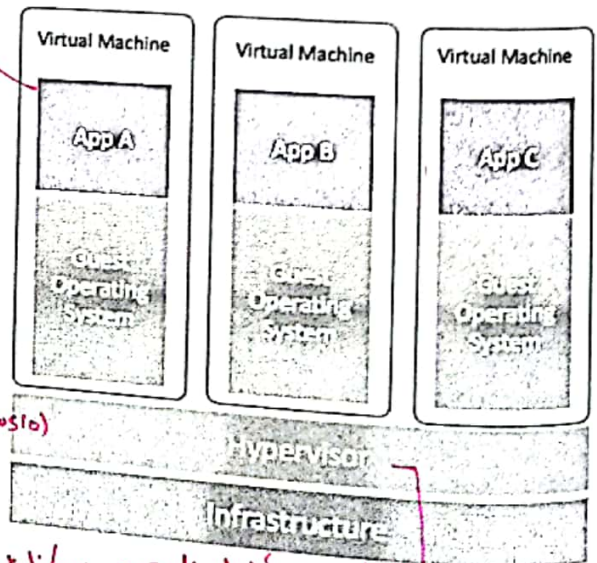
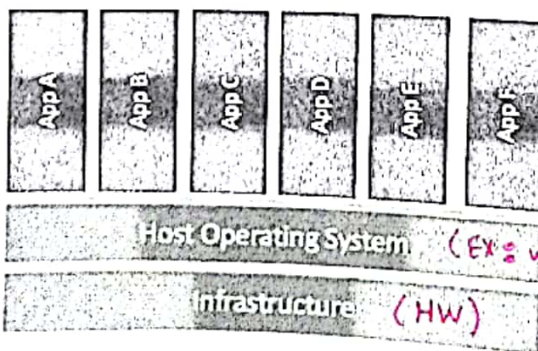
# Virtual Machines

cloud



لو خرب واحد من ال App ما يتاثروا الباقى

عن ايجواز



نوع من انواع ال ال بيملك تعمل اكثر من VM



# Virtual Machine Monitor

- Maps virtual resources to physical resources
  - Memory, I/O devices, CPUs
- Guest code runs on native machine in user mode → *User mode* *بشغل بال*
- Traps to VMM on privileged instructions and access to protected resources
- Guest OS may be different from host OS
- VMM handles real I/O devices
  - Emulates generic virtual I/O devices for guest



## Instruction Set Support

- User and System modes
- Privileged instructions only available in system mode
  - Trap to system if executed in user mode
- All physical resources only accessible using privileged instructions
  - Including page tables, interrupt controls, I/O registers



# Contents

- 5.5 Dependable Memory Hierarchy
- 5.11 Redundant Arrays of Inexpensive Disks
- 5.6 Virtual Machines
- 5.7 Virtual Memory

## Introduction

← باستخدامها بغير ال access بطيئاً

Fast Translation Using a TLB

Memory Protection → ما في Process بتغير تداخل

Virtual Memory → code أو data ل Process تاني وهذا فائدة ال



# Virtual Memory

زاحة افتراضية


الموجودة بكل أنظمة الحواسيب ،  
يعني البرامج بتفترض شي والحقيقة  
شيء آخر

Use main memory as a “cache” for secondary (disk) storage

← يعني ال hardware  
Data بتخل ال  
Memory من ال

Managed jointly by CPU hardware and the operating system (OS)

Programs share main memory →



لا Cache  
لكن نقل ال Data  
بين ال Memory  
وال disk بين  
يتعاون  
Software +  
hardware

Each gets a private virtual address space holding its frequently used code and data

Protected from other programs

CPU and OS translate virtual addresses to physical addresses

VM “block” is called a page

VM translation “miss” is called a page fault



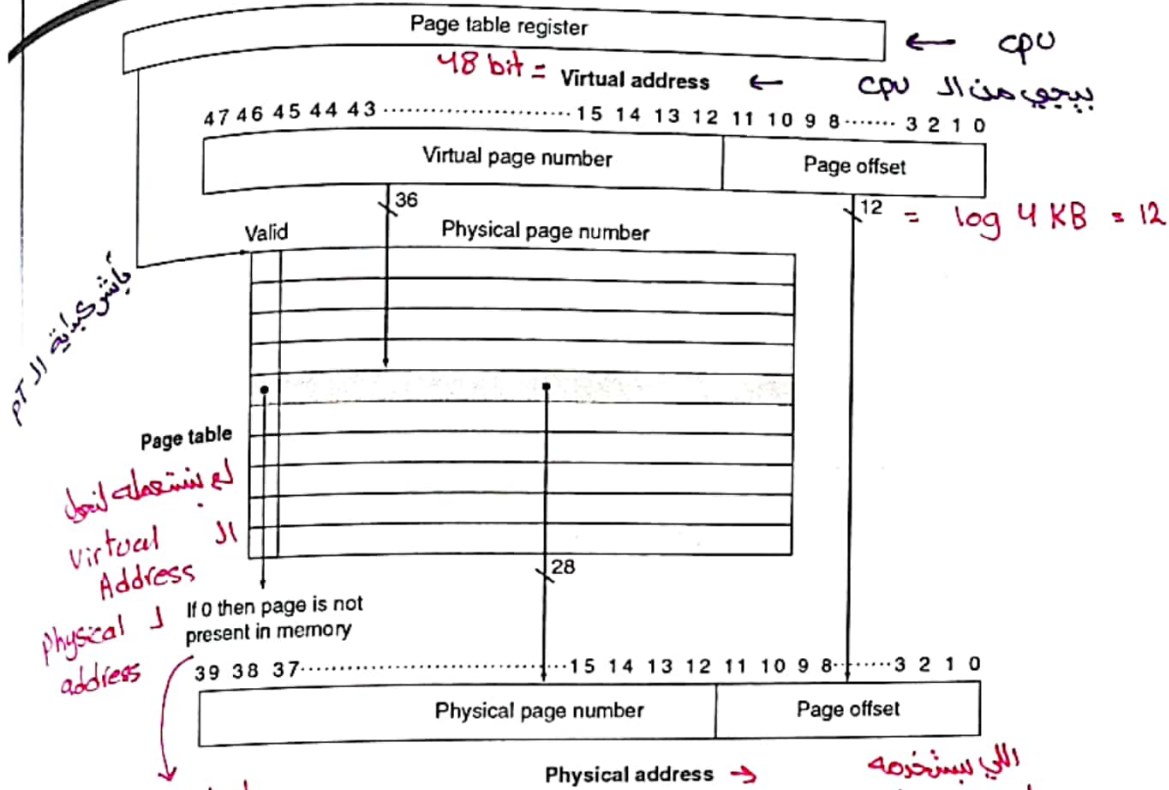








# Translation Using a Page Table



بیشتر عیب از ال PT

لکه نیست صحت لینوی ال

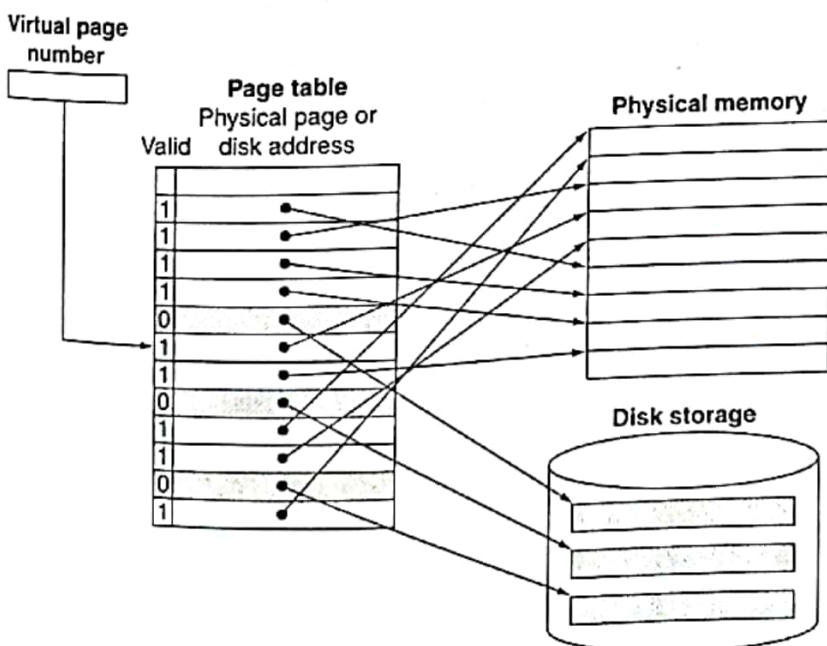
Virtual Address  
Physical address

لو كانت  
ا يكون  
Page fault  
ويكون موجود بال disk

اللي بيستخدمه  
Mem Access ↓



# Mapping Pages to Storage



# Replacement and Writes


- To reduce page fault rate, prefer least-recently used (LRU) replacement
  - Reference bit (aka use bit) in PTE set to 1 on access to page
    - Periodically cleared to 0 by OS
    - A page with reference bit = 0 has not been used recently
- Disk writes take millions of cycles
  - Block at once, not individual locations
  - Write through is impractical
  - Use write-back
    - Dirty bit in PTE set when page is written

كل لم ثانية ال OS  
يصغر ال Use bit

Mem ينطلب بال  
ما ينعتجنا ال disk

مش عملي

هو العملي



طرق تسرع ال translation

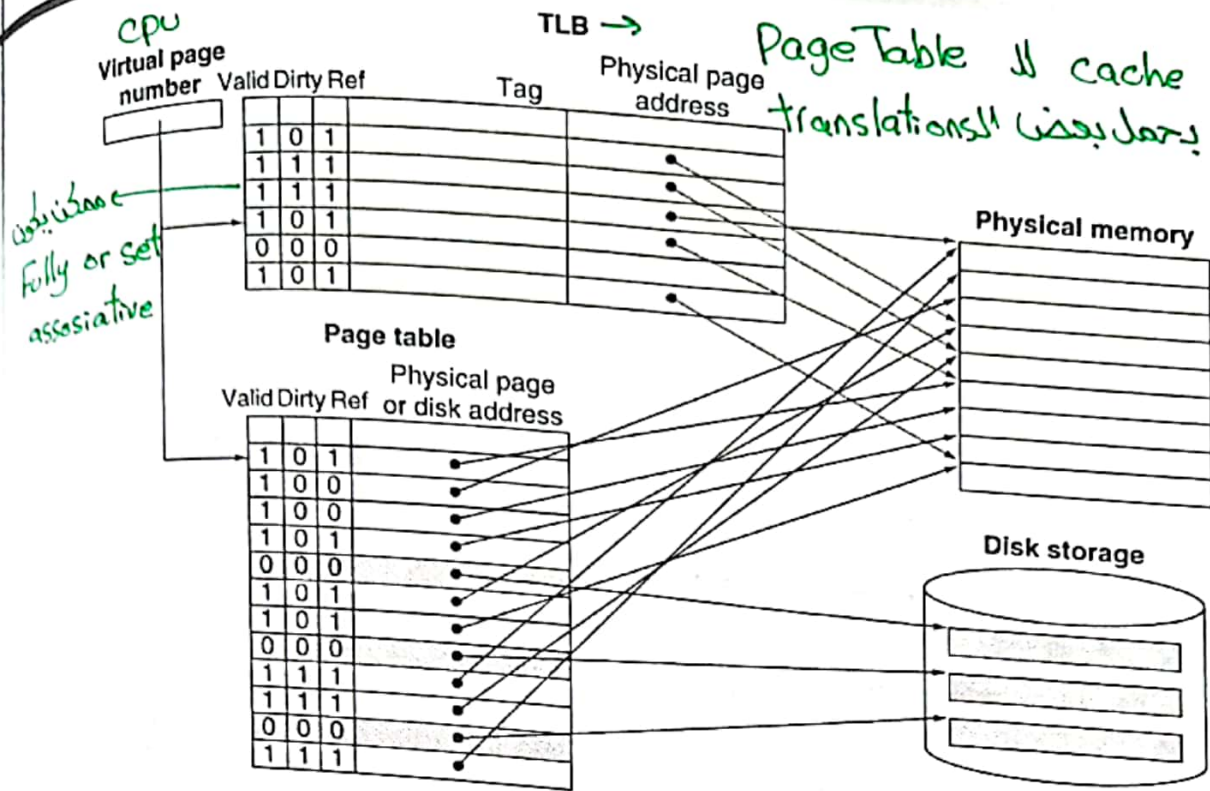
## Fast Translation Using a TLB

- Address translation would appear to require extra memory references  $VA \rightarrow PT \rightarrow PA \rightarrow M$ 
  - One to access the PTE
  - Then the actual memory access
- But access to page tables has good locality
  - So use a fast cache of PTEs within the CPU
  - Called a Translation Look-aside Buffer (TLB)
  - Typical: 16–512 PTEs, 0.5–1 cycle for hit, 10–100 cycles for miss, 0.01%–1% miss rate
  - Misses could be handled by hardware or software

أبداً جدامن ال Process  
ال Page الوحدة بنستخدمها مرات عديدة  
good locality  
له ال CPU يتروح ال TLB وهو مافيه كل شي في جزء من ال translations ولو ما لقيت جزء بلك تروح ال PT وحتي تروح ال PT وتجييب المطلوب وتخطي



# Fast Translation Using a TLB




## TLB Misses

- ① If page is in memory → *ال Page موجودة في ال Memory*
  - Load the PTE from memory and retry *بما ان ال Page موجودة في ال TLB*
  - Could be handled in hardware → *بعض قبيح ال PT Simple فيستخدم ال hardware*
  - Can get complex for more complicated page table structures
  - Or in software → *في ال instructions محسوبة بدقة*
    - Raise a special exception, with optimized handler
- ② If page is not in memory (page fault) → *بكل ال system يتصرف ال OS*
  - OS handles fetching the page and updating the page table
  - Then restart the faulting instruction



# TLB Miss Handler


- TLB miss indicates
  - ① Page present, but PTE not in TLB هاي انبيط
  - ② Page not present (page fault)
- Must recognize TLB miss before destination register overwritten مثلا لو عندي x<sub>1</sub> دا لو عملت TLB Miss عننا لازم نخليها تكمل ولا تكتب x<sub>1</sub> لان ال translation ما صار لسا.
- Raise exception Software لا يتصرف
- Handler copies PTE from memory to TLB
- Then restarts instruction
- If page not present, page fault will occur



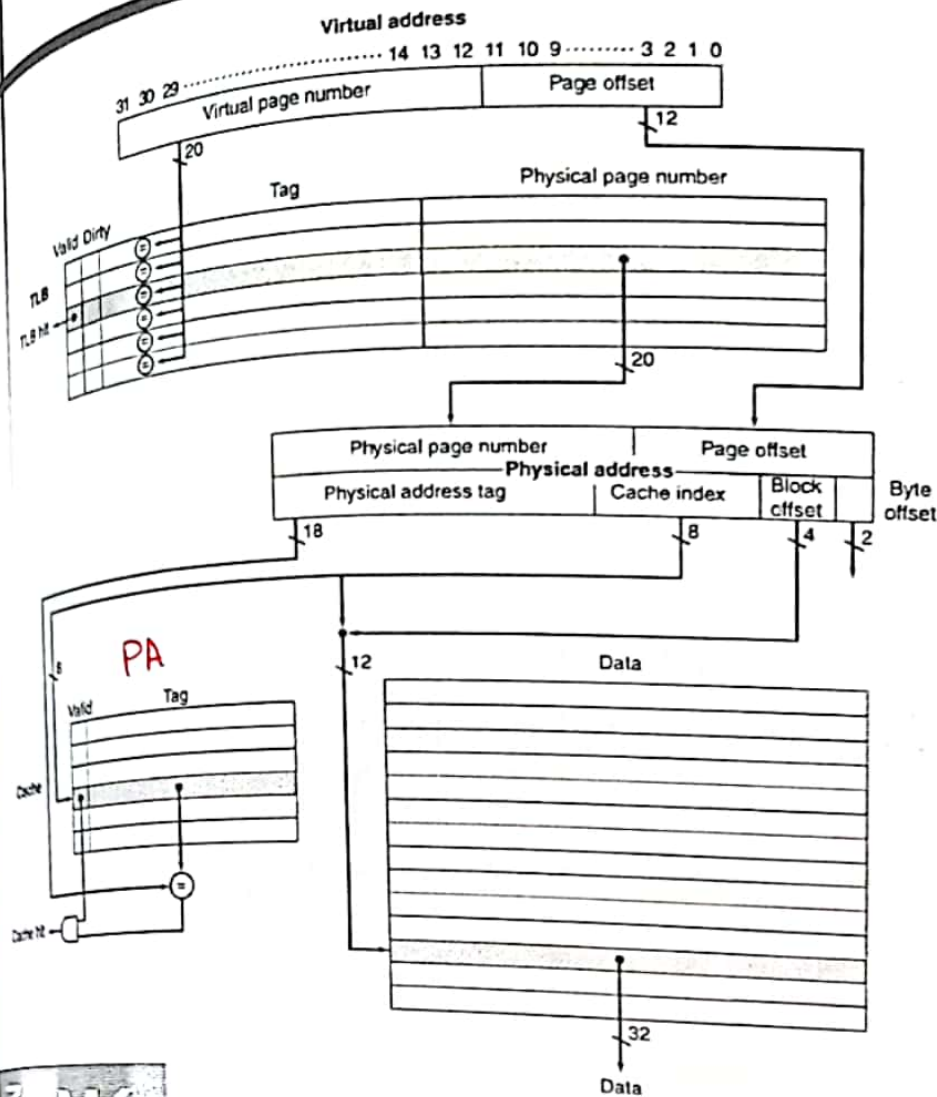
# Page Fault Handler

جزء من ال OS  
بشغل لما يجرى  
عنا page fault

- Use faulting virtual address to find PTE
- Locate page on disk
- Choose page to replace
  - If dirty, write to disk first
- Read page into memory and update page table
- Make process runnable again
  - Restart from faulting instruction



# TLB and Cache Interaction



If cache tag uses physical address

- Need to translate before cache lookup

Alternative: use virtual address tag

- Complications due to aliasing

Different virtual addresses for shared physical address

لے یعنی ممکن ہو سکتا ہے  
 2 P → Shared وہی one page  
 یطووعا different Addresses

## RAID 5

A <sub>1</sub>	A <sub>2</sub>	A <sub>3</sub>	A <sub>p</sub>
B <sub>1</sub>	B <sub>2</sub>	B <sub>p</sub>	B <sub>3</sub>
C <sub>1</sub>	C <sub>p</sub>	C <sub>2</sub>	C <sub>3</sub>
D <sub>p</sub>	D <sub>1</sub>	D <sub>2</sub>	D <sub>3</sub>


## RAID 6

A <sub>1</sub>	A <sub>2</sub>	A <sub>3</sub>	A <sub>p</sub>	A <sub>q</sub>
B <sub>1</sub>	B <sub>2</sub>	B <sub>p</sub>	B <sub>q</sub>	B <sub>3</sub>
C <sub>1</sub>	C <sub>p</sub>	C <sub>q</sub>	C <sub>2</sub>	C <sub>3</sub>
D <sub>p</sub>	D <sub>q</sub>	D <sub>1</sub>	D <sub>2</sub>	D <sub>3</sub>
E <sub>q</sub>	E <sub>1</sub>	E <sub>2</sub>	E <sub>3</sub>	E <sub>p</sub>

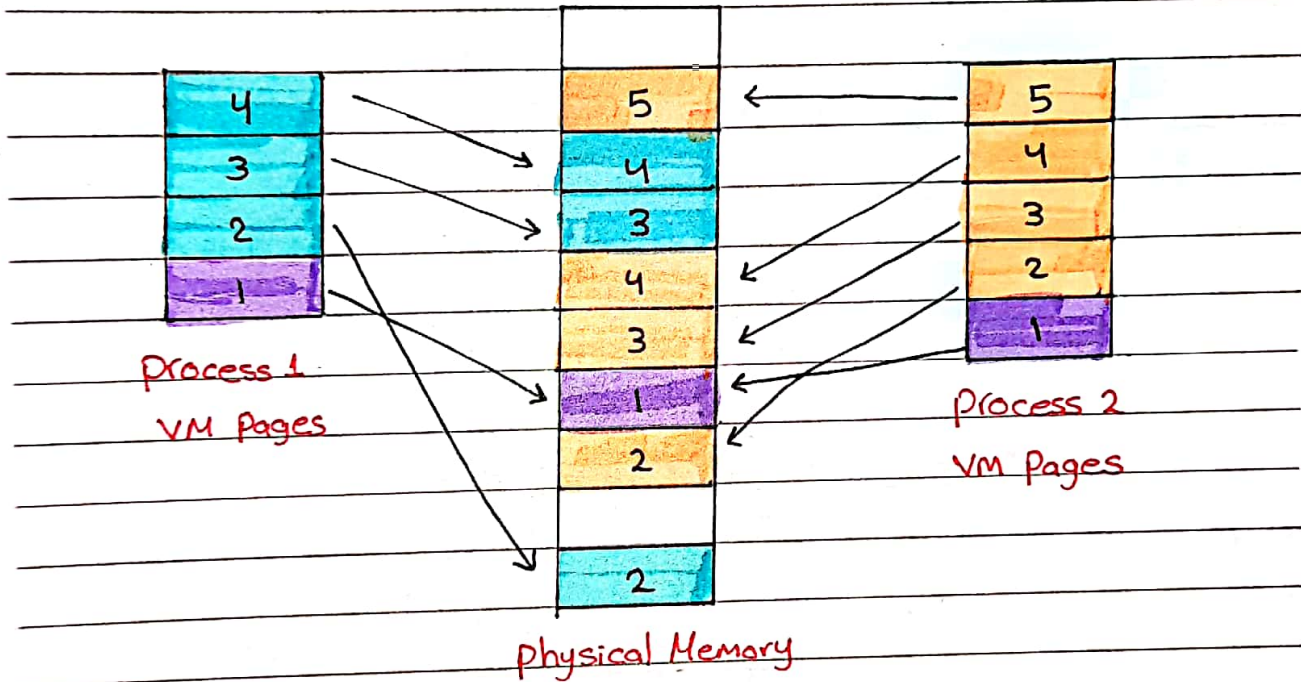
\* ال P يحسب سطرًا أفقيًا .

\* ال Q تحسب بطريقة أخرى كسطر ماثل مثل أفقي ( قطري diagonal )





Slide 94 : Sharing the physical Memory :




- \* مقسمة لـ Pages و الـ page الوحدة الكثر شي = 4KB
- \* بالأمثلة كل Process لها Pages خاصة فيها، إلا إذا بنا نعمل Sharing زي رقم 1
- \* الأبيش يعبر عن pages فاضية
- \* المبرمج يفترض انه بيغير يستخدم أي Address من أكبر رقم لأقل رقم
- \* مسؤولية الـ OS + CPU hardware بتوزعها الـ Virtual spaces بالـ Physical Memory (Pages)

## Slide 106 : TLB - Cache Interaction :


### \* Physical Address Tag :

\* يعني ال cache يكون فيه physical Address مثلا ال Memory  
فلازم اول يكون ال TLB و هيك يكون ال cache ال بيتا .

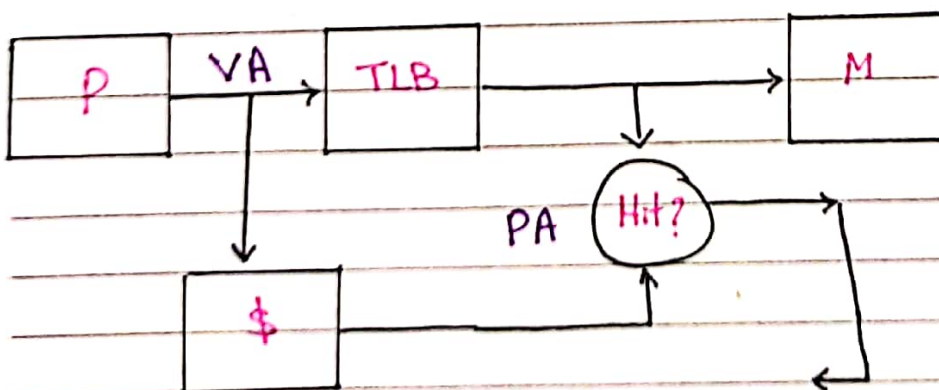


### \* Virtual Address Tag :

\* يعني ال tags ال في ال cache يكون ال Virtual Address مثلا في حاجة  
في Translation cache ال لو صار Miss .



\* Physical Address Tag :  $cache\ size = Page\ Size * associativity$   
في بيتا الحالة الخالصة .






\* ما في Page في Process بتقدر تشوف Page ثانية في Process ثاني وبالتالي هنا Protection . الا انا كانا

# Memory Protection

Process 1



Process 2  
 Parallel Program وطلبوا من ال System ان يعلو Shared فيقدرنا يشوفوا بعض

Different tasks can share parts of their virtual address spaces (read) (write) (Fetch) r w F


But need to protect against errant access access في حبيجة

Requires OS assistance

هو اللي يكتشف الاخطاء

هو اللي ينسق ويحدد

## Hardware support for OS protection



Privileged supervisor mode (aka kernel mode)

Privileged instructions

Page tables and other state information only accessible in supervisor mode → لبنودنيا لل operating system

System call exception (e.g., ecall in RISC-V)





# Contents

- 5.1 Introduction
- 5.2 Memory Technologies
- 5.3 The Basics of Caches
- 5.4 Measuring and Improving Cache Performance
- 5.5 Dependable Memory Hierarchy
- 5.11 Redundant Arrays of Inexpensive Disks
- 5.6 Virtual Machines
- 5.7 Virtual Memory → Fundamental (أساسي)
- 5.8 A Common Framework for Memory Hierarchy
- 5.9 Using a Finite-State Machine to Control a Simple Cache
- 5.10 Cache Coherence
- 5.13 The ARM Cortex-A53 and Intel Core i7 Memory Hierarchies
- 5.16 Fallacies and Pitfalls
- 5.17 Concluding Remarks



## The Memory Hierarchy

### The BIG Picture

- Common principles apply at all levels of the memory hierarchy  
\* كل اشئ هو cache اللي تحتها موجود فيه جزء من المعلومات اللي تحتها.
- Based on notions of caching
- At each level in the hierarchy
  - Block placement → اذا جينا اشئ من تحت ل فوق وين نخطه .
  - Finding a block → وين هذا ال address اذا موجود بال level اللي اخنا فيه .
  - Replacement on a miss → اذا كل الاماكن مليانة وين نخط .
  - Write policy → لما بدنا نكتب كيف نكتب



# Block Placement

- Determined by associativity
  - Direct mapped (1-way associative)
    - One choice for placement *كل Mem block له مكان واحد بتحدد بال index*
  - n-way set associative (2/4/8/16) *مش أكثر من هيك*
    - n choices within a set
  - Fully associative
    - Any location *مانا index كل الأماكن بال cache محتاجت*
- Higher associativity reduces miss rate
  - Increases complexity, cost, and access time

*مشكلة  
higher associativity*



## Finding a Block

Associativity	Location method	Tag comparisons
Direct mapped	Index <i>بتحدد عدد ال block بال cache</i>	1 $\oplus$
n-way set associative	Set index, then search entries within the set	n
Fully associative	Search all entries $\$$	#entries
	Full lookup table <i>VM</i>	0

- Hardware caches
  - Reduce comparisons to reduce cost
- Virtual memory
  - Full table lookup makes full associativity feasible
  - Benefit in reduced miss rate

*Page table*

*VA  $\rightarrow$  [PT]  $\rightarrow$  PA*



# Replacement

choices ما في direct map بال  
high/fully associative بال

- Choice of entry to replace on a miss
  - Least recently used (LRU)
    - Complex and costly hardware for high associativity
  - Random
    - Close to LRU, easier to implement
- Virtual memory
  - LRU approximation with hardware support

recently used : يقسم ال Page لجزئين  
و Not recently used و بينم ال Replacement  
من ال not recently used  
وال bit ال ال VM ال ال يستخدمها ال OS



هي ( 2 bit ) ← يعني ال OS كل  
فتره وفترة بحد

فيها ال Pages و بحد فيها ال بيدينا  
ال ال Page ال ال بحيرلوا access بنحول ال و هيك بنكون

أشترنا عال  
Pages  
ال ال

# Write Policy



Recently  
used  
وال ال خيلوا  
مفر  
Not recently  
Used

- Write-through
  - Update both upper and lower levels
  - Simplifies replacement, but may require write buffer
- Write-back
  - Update upper level only
  - Update lower level when block is replaced
  - Need to keep more state
- Virtual memory
  - Only write-back is feasible, given disk write latency





# Sources of Misses

تصنيف ال cache Misses حسب سببهم .

## Compulsory misses (aka cold start misses)

- First access to a block

لما أول مرة برنا نعمل ونشغل البرنامج

## Capacity misses

- Due to finite cache size
- A replaced block is later accessed again

اجباري تكون Miss لأنها مش موجودة .

لحجم ما يكون واسع في ال cache فلو علنا access معلومة ما لقينها لأنه اذختم كارنا اشي ثاني فبغير في Miss

## Conflict misses (aka collision misses)

- In a non-fully associative cache
- Due to competition for entries in a set
- Would not occur in a fully associative cache of the same total size

إذا بدى أقل ال capacity Miss ال تزيد حجم ال cache

يعني يعرف بال cache مكان معين، ال index معين، انه ال P بطلب هذا ال M عدد blocks



Chapter 5 — Large and Fast: Exploiting Memory Hierarchy — 114

نفس ال index بال cache فلو طلب الأول بيحي عادي لو طلب الثاني بده يجيب ال cache بس يطرد الأول فبغير الأول والثاني يخلوهم يطردوا بعض .

# Cache Design Trade-offs




مع انه ال cache في أماكن خاضعة لسبب أشوا عتس ال index

$$* AMAT = HT + miss\ rate * Miss\ Penalty$$

Design change	Effect on miss rate	Negative performance effect
Increase cache size	Decrease capacity misses	May increase access time
Increase associativity	Decrease conflict misses	May increase access time
Increase block size	Decrease compulsory misses	Increases miss penalty. For very large block size, may increase miss rate due to pollution.

لما ال block size يكثر يجيب اجيب ب Miss وحدة more data فلو بختا بكون وقتنا عوالي Misses بالمستقبل

# Data Cache Miss Rate



تشانة كل الأنوع  
Miss rate



\* العلاقة بين ال Miss rate و ال Associativity و بتقل ال Miss rate كل ما يزيد ال Associativity السبب إنه ال conflict Missis بتقل

## Contents

\* لما يكبر ال cache بتقل ال Miss rate لأنه ال Miss capacity  
\* لما نعمل ل 128 KiB ولو زدنا ال Associativity بتقل ال Miss rate و السبب وجود ال compulsory Missis

فلا يمكن  
ينقص  
أكثر حد  
هيك

- 5.1 Introduction
- 5.2 Memory Technologies
- 5.3 The Basics of Caches
- 5.4 Measuring and Improving Cache Performance
- 5.5 Dependable Memory Hierarchy
- 5.11 Redundant Arrays of Inexpensive Disks
- 5.6 Virtual Machines
- 5.7 Virtual Memory
- 5.8 A Common Framework for Memory Hierarchy
- 5.9 Using a Finite-State Machine to Control a Simple Cache
- 5.10 Cache Coherence
- 5.13 The ARM Cortex-A53 and Intel Core i7 Memory Hierarchies
- 5.16 Fallacies and Pitfalls
- 5.17 Concluding Remarks






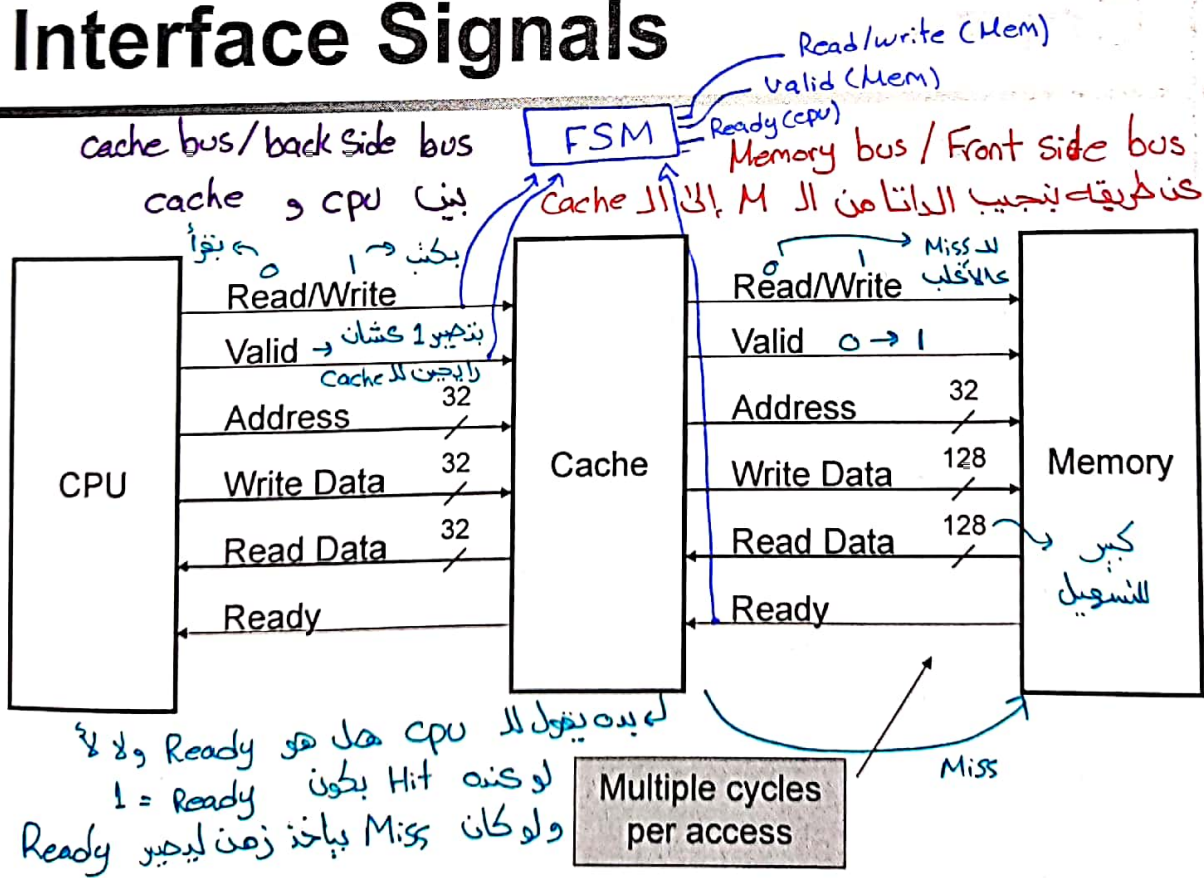
# Cache Control

## Example cache characteristics

- Direct-mapped, write-back, write allocate
- Block size: 4 words (16 bytes)
- Cache size: 16 KB (1024 blocks)
- 32-bit byte addresses
- Valid bit and dirty bit per block
- Blocking cache → write-back لأن CPU waits until access is complete



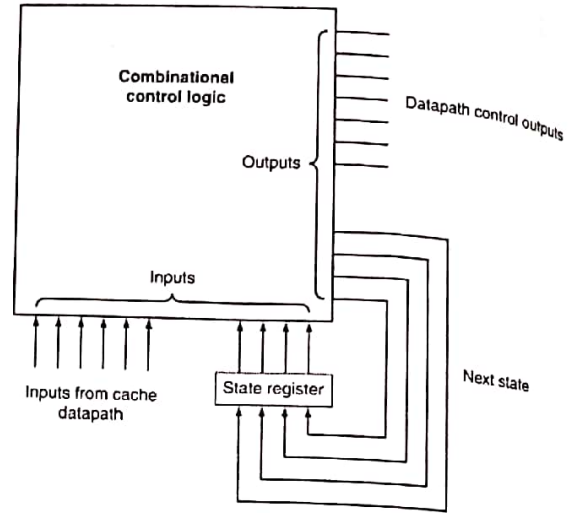
## Interface Signals





# Finite State Machines (Sequential)

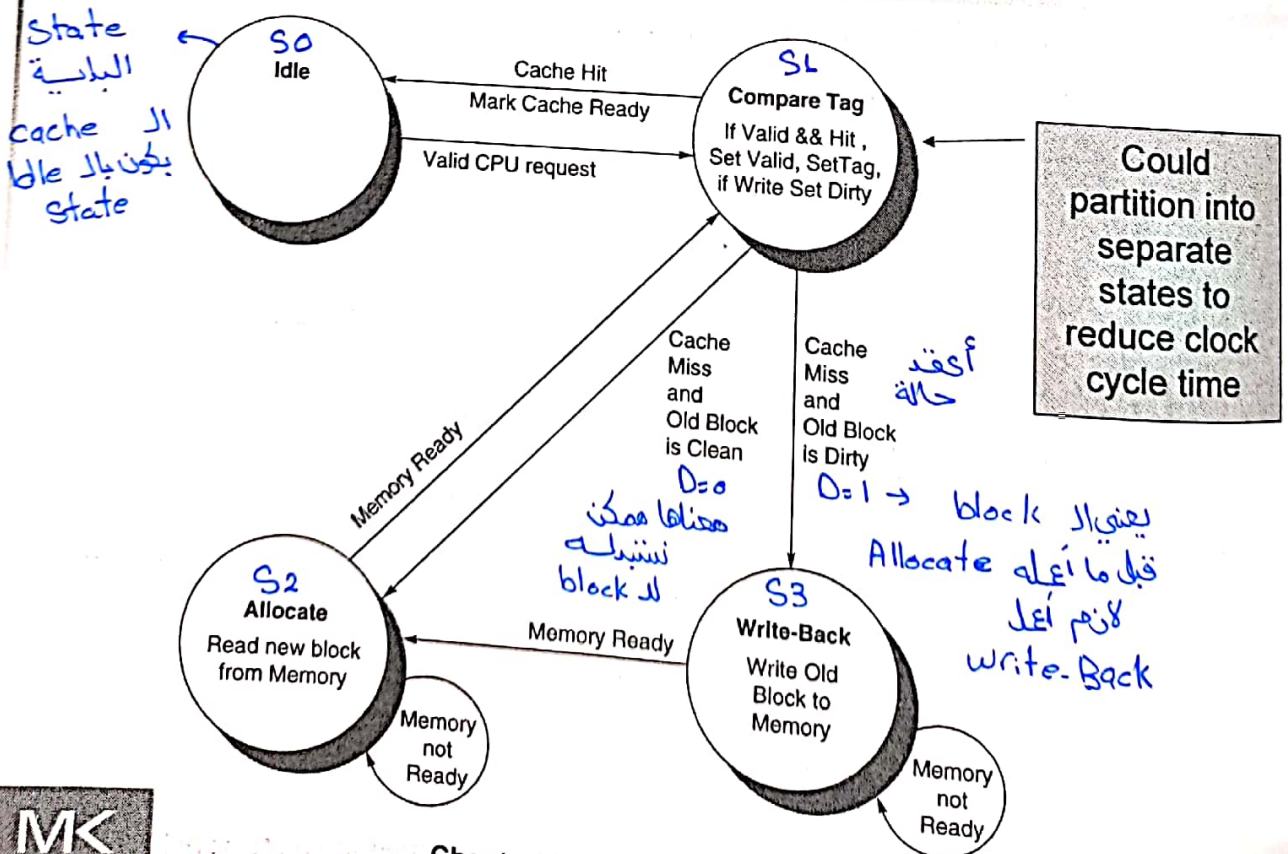
- Use an FSM to sequence control steps
- Set of states, transition on each clock edge
  - State values are binary encoded
  - Current state stored in a register
  - Next state =  $f_n$  (current state, current inputs)
- Control output signals =  $f_o$  (current state)



Chapter 5 — Large and Fast: Exploiting Memory Hierarchy — 120

Sequential circuits → input out بعينها و ال State و ال  
 Combinational circuits → ما فيو ذاكرة و ال out بعينها عا ال input

## Cache Controller FSM



Chapter 5 — Large and Fast: Exploiting Memory Hierarchy — 121

# Contents

- 5.1 Introduction
- 5.2 Memory Technologies
- 5.3 The Basics of Caches
- 5.4 Measuring and Improving Cache Performance
- 5.5 Dependable Memory Hierarchy
- 5.11 Redundant Arrays of Inexpensive Disks
- 5.6 Virtual Machines
- 5.7 Virtual Memory
- 5.8 A Common Framework for Memory Hierarchy
- 5.9 Using a Finite-State Machine to Control a Simple Cache
- 5.10 Cache Coherence *Multiple cpus*
- 5.13 The ARM Cortex-A53 and Intel Core i7 Memory Hierarchies
- 5.16 Fallacies and Pitfalls
- 5.17 Concluding Remarks



*متراپطة ومتعاضة ومتزامنة*

## Cache Coherence Problem

Suppose two CPU cores share a physical address space

Write-through caches → *هاي المشكلة بنحسب بال write-back ولا*

Time step	Event	CPU A's cache	CPU B's cache	Memory
0				$x = 0$
1	CPU A reads X	0		0
2	CPU B reads X	0	0	0
3	CPU A writes 1 to X	1	0	1

*رح يكتب عال cache وعال Memory L =  
وهون في مشكلة لانها ال cache حار في اي وقت نسخ منها X وواحد بدل والثاني لا  
فلو CPU B قرا X من قبلها همتب ا وهي مشكلة والحد في الحل في الحلف*

§5.10 Parallelism and Memory Hierarchies: Cache Coherence





# Coherence Defined

تعريف cache coherence

Informally: Reads return most recently written value → يعني لما ال CPU تعمل read لازم تقرا اجدد قيمت

Formally:

التعريف المطلوب

- ① P writes X; P reads X (no intervening writes)  
⇒ read returns written value
- ② P<sub>1</sub> writes X; P<sub>2</sub> reads X (sufficiently later)  
⇒ read returns written value  
بعد فترة زمنية مقبولة
- ③ P<sub>1</sub> writes X, P<sub>2</sub> writes X  
⇒ all processors see writes in the same order  
End up with the same final value for X

ننشوف بأي ترتيب كتبوا و بناخذ آخر كتابية كتبوها أو انكتب



## Cache Coherence Protocols

Operations performed by caches in multiprocessors to ensure coherence

- Migration of data to local caches
  - Reduces bandwidth for shared memory
- Replication of read-shared data
  - Reduces contention for access

أنواع ال cache coherence

① Snooping protocols

- Each cache monitors bus reads/writes

② Directory-based protocols

- Caches and memory record sharing status of blocks in a directory

دليل






Modified ← Shared → Invalid  
 (MSI) ← كسنان في اكثر من نوع وحناخذ

# Invalidating Snooping Protocols

- Cache gets exclusive access to a block when it is to be written
  - Broadcasts an invalidate message on the bus
  - Subsequent read in another cache misses
    - Owning cache supplies updated value

CPU activity	Bus activity	CPU A's cache	CPU B's cache	Memory
		I	I	X = 0
CPU A reads X	Cache miss for X	0	I	0
CPU B reads X	Cache miss for X	S 0	S 0	0
CPU A writes 1 to X	Invalidate for X	M 1	I	0
CPU B read X	Cache miss for X	S 1	S 1	1



# Memory Consistency

لما P واحد بيكتب الثانيين حتى يشوفوا الكتابة.

- When are writes seen by other processors
  - "Seen" means a read returns the written value
  - Can't be instantaneously → بسبب انه في تباعد بين Processors
- Assumptions
  - A write completes only when all processors have seen it
  - A processor does not reorder writes with other accesses ← sequential consistency هيا الايام لانه بظلي البرامج بطيئة وازا موجود لازم يكون اسمه
- Consequence
  - P writes X then writes Y
    - ⇒ all processors that see new Y also see new X
  - Processors can reorder reads, but not writes



# Contents

- 5.1 Introduction
- 5.2 Memory Technologies
- 5.3 The Basics of Caches
- 5.4 Measuring and Improving Cache Performance
- 5.5 Dependable Memory Hierarchy
- 5.11 Redundant Arrays of Inexpensive Disks
- 5.6 Virtual Machines
- 5.7 Virtual Memory
- 5.8 A Common Framework for Memory Hierarchy
- 5.9 Using a Finite-State Machine to Control a Simple Cache
- 5.10 Cache Coherence
- 5.13 The ARM Cortex-A53 and Intel Core i7 Memory Hierarchies
- 5.16 Fallacies and Pitfalls
- 5.17 Concluding Remarks



## Multilevel On-Chip Caches

Characteristic	ARM Cortex-A53	Intel Core i7
L1 cache organization	Split instruction and data caches	Split instruction and data caches
L1 cache size	Configurable 16 to 64 KiB each for instructions/data	32 KiB each for instructions/data per core
L1 cache associativity	Two-way (I), four-way (D) set associative	Four-way (I), eight-way (D) set associative
L1 replacement	Random	Approximated LRU
L1 block size	64 bytes	64 bytes
L1 write policy	Write-back, variable allocation policies (default is Write-allocate)	Write-back, No-write-allocate
L1 hit time (load-use)	Two clock cycles	Four clock cycles, pipelined
L2 cache organization	Unified (instruction and data)	Unified (instruction and data) per core
L2 cache size	128 KiB to 2 MiB	256 KiB (0.25 MiB)
L2 cache associativity	16-way set associative	8-way set associative
L2 replacement	Approximated LRU	Approximated LRU
L2 block size	64 bytes	64 bytes
L2 write policy	Write-back, Write-allocate	Write-back, Write-allocate
L2 hit time	12 clock cycles	10 clock cycles
L3 cache organization	-	Unified (instruction and data)
L3 cache size	-	8 MiB, shared
L3 cache associativity	-	16-way set associative
L3 replacement	-	Approximated LRU
L3 block size	-	64 bytes
L3 write policy	-	Write-back, Write-allocate
L3 hit time	-	35 clock cycles

← الصفحة

1 GHz

3-4 GHz

↑ P ↓ D  
I 32k D 32k

80w each block



# 2-Level TLB Organization

Characteristic	ARM Cortex-A53	Intel Core i7
Virtual address	48 bits	48 bits
Physical address	40 bits	44 bits
Page size	Variable: 4, 16, 64 KiB, 1, 2 MiB, 1 GiB	Variable: 4 KiB, 2/4 MiB
TLB organization	1 TLB for instructions and 1 TLB for data per core	1 TLB for instructions and 1 TLB for data per core
Cache	Both micro TLBs are fully associative, with 10 entries, round robin replacement	Both L1 TLBs are four-way set associative, LRU replacement
PTE ↓	64-entry, four-way set-associative TLBs	L1 I-TLB has 128 entries for small pages, seven per thread for large pages
VA بچوں	TLB misses handled in hardware	L1 D-TLB has 64 entries for small pages, 32 for large pages
PA 3!	The micro TLBs are L1 TLBs that are backed by two 64-entry L2 TLBs.	The L2 TLB is four-way set associative, LRU replacement
		The L2 TLB has 512 entries
		TLB misses handled in hardware

→  $2^{44} = 16 \text{ TB}$

→ L1 TLB I 128+7 TLB D 64+32  
L2 TLB 512

MK

## Supporting Multiple Issue


Both have multi-banked caches that allow multiple accesses per cycle assuming no bank conflicts

Other optimizations

- Return requested word first
- Non-blocking cache
  - Hit under miss
  - Miss under miss



### Slide 123 : Cache Coherence.




يخزنه القيمة القيمة لـ x الذي  
Bus

مفرد  
فذلك بطل في تعاون بين A و B  
وهذا هو cache coherence problem  
أي شيء يروح عال M بين عال Bus

Time step	Event	cpu A's cache	cpu B's cache	Memory
0				0
1	cpu A reads x	0		0
2	cpu B reads x	0	0	0
3	cpu A writes 1 to x	1	0	1

### Slide 126 : Snooping Protocols.



لما يراقب ال cache ال Bus ويشوف  
مار فيه Exec. copy بخلي  
invalid يعني بشيل المفرد

وال M خذ فيها القيمة القيمة واما رجينا ليا  
read B يكون Miss فيطلع عال Bus ← Read Miss  
و يياخذ ال 1 من ال A

cpu activity	Bus Activity	cpu A's cache	cpu B's cache	Memory
				0
cpu A reads x	Cache Miss for x	0		0
cpu B reads x	Cache Miss for x	0	0	0
cpu A writes 1 to x	invalidate for x	1		0
cpu B reads x	Cache Miss for x	1	1	1

## Slide 127 : Memory Consistency.

Processor A	Processor B
Flag = False ;	
X = Compute();	while (flag == false);
Flag = True ;	Print (X);

2 Processors يعملوا مع بعض : A يجعل حصة معينة ويكتب قيمة جديدة لـ X  
وبمجرد flag = true يعني إنه خلص .

B به يطبع القيمة الجديدة لـ X بس قبل ما يقرأها ويطلعها يكون براقب ال flag لو كان  
false بزمال while او حلات ال Flag = true بطلع من ال loop ويقدر إنه ينفذ

ال Print . ولو فيه Sequential Consistency فالتا تطلع ال flag = t معنا كتب  
قيمة X الجديدة فبشغل البرنامج صح .

هاي الأيام ما في Sequential Consistency وال Processors يكتب وفي out of order

execution فممكن انه قيمة ال Flag تظهر لـ B قبل ما تطوره القيمة الجديدة لـ X

وبالتالي البرنامج ما يشغل صح فالحل انه البرمجين هالأيام لما يكتبوا Parallel Programs

يعتصوا على تقنيات خاصة ( For synchronisation ) حتى تضمنوا انه الشغل

صح .

# Supporting Multiple Issue

(سبب احنا جندي بشكل عام)

→ ARM and Intel

- Both have multi-banked caches that allow multiple accesses per cycle assuming no bank conflicts (الدفتر بشرحها)

## ■ Other optimizations

- ① ■ Return requested word first
- ② ■ Non-blocking cache
  - Hit under miss
  - Miss under miss
- ③ ■ Data prefetching

سبب احنا نفلق عندنا  
Capacity  
Missis





# Contents

- 5.1 Introduction
- 5.2 Memory Technologies
- 5.3 The Basics of Caches
- 5.4 Measuring and Improving Cache Performance
- 5.5 Dependable Memory Hierarchy
- 5.11 Redundant Arrays of Inexpensive Disks
- 5.6 Virtual Machines
- 5.7 Virtual Memory
- 5.8 A Common Framework for Memory Hierarchy
- 5.9 Using a Finite-State Machine to Control a Simple Cache
- 5.10 Cache Coherence
- 5.13 The ARM Cortex-A53 and Intel Core i7 Memory Hierarchies
- 5.16 Fallacies and Pitfalls
- 5.17 Concluding Remarks



## Pitfalls

- Byte vs. word addressing
  - Example: 32-byte direct-mapped cache, 4-byte blocks
    - Byte 36 maps to block 1
    - Word 36 maps to block 4
- Ignoring memory system effects when writing or generating code
  - Example: iterating over rows vs. columns of arrays
  - Large strides result in poor locality



# Pitfalls

- In multiprocessor with shared L2 or L3 cache
  - Less associativity than cores results in conflict misses
  - More cores  $\Rightarrow$  need to increase associativity
- Using AMAT to evaluate performance of out-of-order processors
  - Ignores effect of non-blocked accesses
  - Instead, evaluate performance by simulation

MK

Chapter 5 — Large and Fast: Exploiting Memory Hierarchy — 134

# Pitfalls

- Extending address range using segments
  - E.g., Intel 80286
  - But a segment is not always big enough
  - Makes address arithmetic complicated
- Implementing a VMM on an ISA not designed for virtualization
  - E.g., non-privileged instructions accessing hardware resources
  - Either extend ISA, or require guest OS not to use problematic instructions

MK

Chapter 5 — Large and Fast: Exploiting Memory Hierarchy — 135


# Concluding Remarks

- Fast memories are small, large memories are slow
  - We really want fast, large memories ☹️
  - Caching gives this illusion 😊
- Principle of locality
  - Programs use a small part of their memory space frequently
- Memory hierarchy
  - L1 cache ↔ L2 cache ↔ ... ↔ DRAM memory ↔ disk
- Memory system design is critical for multiprocessors





# Slide 131 : Supporting Multiple Issue.



« هذا ما توضحه cache مقسوم لـ 2 Banks »  
فيخرج الـ Set اللى الـ Address تاوم زوجي  
بـ Bank 0 والـ Set اللى الـ Address تاوم فردى  
بالـ Bank الآخر. ولما يكون في 2 Bank معناها  
فيه 2 Parts فممكن الـ P access 2 بنفس الوقت  
وهيك بنكون في الـ Performance



# Slide 131 : Data Prefetching

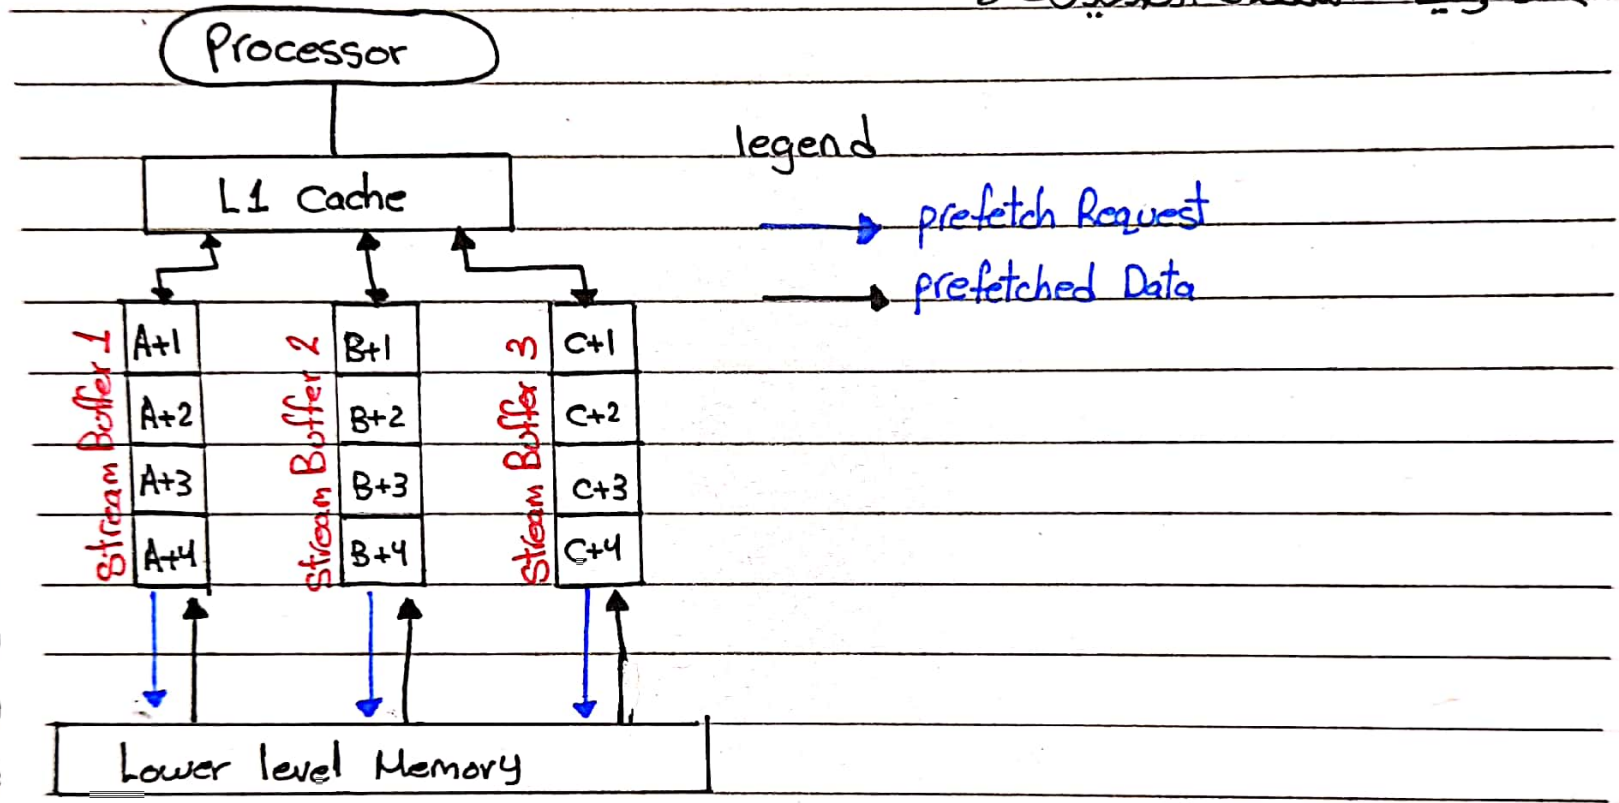
Access Block A

→ Processor يطلب A ال prefetch circuit ينسقب الأبحاث ويتجهل Next sequential Block ال prefetching

Prefetch Block A+1

ممتاز وينوف على حالنا Miss

← الطريقة المربعة المربعة





# Chapter 6

## Parallel Processors from Client to Cloud

لـ مثل العواتف  
والنوت بوك

*Adapted by Prof. Gheith Abandah*

# Contents

- 6.1 Introduction
- 6.2 The Difficulty of Creating Parallel Programs
- 6.3 SISD, MIMD, SIMD, SPMD, and Vector
- 6.4 Hardware Multithreading
- 6.5 Shared Memory Multiprocessors
- 6.6 Introduction to Graphics Processing Units
- 6.7 Clusters and Message-Passing Multiprocessors
- 6.8 Introduction to Multiprocessor Network Topologies
- 6.10 Multiprocessor Benchmarks and Performance Models
- 6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU
- 6.12 Multiple Processors and Matrix Multiply
- 6.13 Fallacies and Pitfalls
- 6.14 Concluding Remarks



## Introduction

§6.1 Introduction

■ Goal: connecting multiple computers to get higher performance

- Multiprocessors (بجيب هيك اسم الجواز )
- Scalability, availability, power efficiency (بصرف طاقة أقل )

multiple processors are more power efficiency than single processors

■ Task-level (process-level) parallelism

■ High throughput for independent jobs

■ Parallel processing program

■ Single program run on multiple processors

multiple jobs ال  
ينفذوا على Multiple cores

■ Multicore microprocessors

■ Chips with multiple processors (cores)

Ex. SPMD  
+ data





# Hardware and Software

## Hardware

- ① Serial: e.g., Pentium 4
- ② Parallel: e.g., quad-core Xeon e5345

## Software

- ① Sequential: e.g., matrix multiplication
- ② Concurrent: e.g., operating system

Sequential/concurrent software can run on serial/parallel hardware

- Challenge: making effective use of parallel hardware

كان موجود قبل 20 سنة

ينفذ 4 program بنفس الوقت

لو تنفيذون بيأخذ core واحد ويترك الباقي

ممكن تنفيذ

مفقد قسمه

لأجزاء كل جزء إلى مجموعة خاصة فيه ويتعاون مع بعض

لو تنفيذ

هون رح ينفذ

بزنه أقل

وبنفذوا وحدة

ورا الثانية



## What We've Already Covered

- §2.11: Parallelism and Instructions
  - Synchronization
- §3.6: Parallelism and Computer Arithmetic
  - Subword Parallelism → Performance أعلى
- §4.10: Parallelism and Advanced Instruction-Level Parallelism
- §5.10: Parallelism and Memory Hierarchies
  - Cache Coherence

بلاصغ





# Contents

- 6.1 Introduction
- 6.2 The Difficulty of Creating Parallel Programs → البرمجة صعبة
- 6.3 SISD, MIMD, SIMD, SPMD, and Vector
- 6.4 Hardware Multithreading
- 6.5 Shared Memory Multiprocessors
- 6.6 Introduction to Graphics Processing Units
- 6.7 Clusters and Message-Passing Multiprocessors
- 6.8 Introduction to Multiprocessor Network Topologies
- 6.10 Multiprocessor Benchmarks and Performance Models
- 6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU
- 6.12 Multiple Processors and Matrix Multiply
- 6.13 Fallacies and Pitfalls
- 6.14 Concluding Remarks



# Contents

- 6.2 The Difficulty of Creating Parallel Programs
  - Parallel Programming
  - Amdahl's Law
  - Scaling
  - Strong and Weak Scaling

# Parallel Programming

- Parallel software is the problem
- Need to get significant performance improvement
  - البرنامج هشا بس يشتغل مع لا يشتغل مع → ويكون الperformance عالي و التحسين جدا عالي يكون
  - Otherwise, just use a faster uniprocessor, since it's easier!
    - لم لو كان صعب عليك تقمل البروجت بتحسين عالي فبستخدم ال faster uniprocessor.

## Difficulties

- Partitioning → تجزء الشغل اللي بديك تقمله لأجزاء
- Coordination → خلال التنفيذ لازم يسنقرو ال cores كيف ييلشوا ويخلصوا مع بعض و يشتغلوا مزبوط
- Communications overhead

لما نكتب Parallel Program لازم نتنبه لهاي الصعوبات

لما يهيبوا ال processors يتعاونوا مع بعضا يحسبوا يحتاجوا

Data من بعض فبحسب في زمن كبير فيرنا لما ال processor يشتغل



لحاله بدون ما يشارف

Processors آخرين.

## Amdahl's Law

بكيلنا عن ال speed up ممكن نحصله

Sequential part (1-f) Speed up ممكن نحصله بعينه

- Sequential part can limit speedup
- Example: 100 processors, 90x speedup?
  - $T_{new} = T_{parallelizable}/100 + T_{sequential}$
  - $Speedup = \frac{1}{(1 - F_{parallelizable}) + F_{parallelizable}/100} = 90$
  - Solving:  $F_{parallelizable} = 0.999$
  - Need sequential part to be 0.1% of original time
    - وهذا قولنا صعب فوهنا واحد من الصعوبات اللي بنواجهها



# Scaling Example

أمثلة بتفريغنا العلاقة بين حجم المسألة وعدد ال Processors

نعنيها Serial

Workload: sum of 10 scalars, and  $10 \times 10$  matrix sum

لـ بقدر أجزءه

Speed up from 10 to 100 processors → مقارنة نسبة Processor

Single processor: Time =  $(10 + 100) \times t_{add}$  110 . لح يحتاج عملية جمع

10 processors  
 Time =  $10 \times t_{add} + 100/10 \times t_{add} = 20 \times t_{add}$   
 Speedup =  $110/20 = 5.5$  (55% of potential)

بدلها يكون 110 يكون 20

100 processors  
 Time =  $10 \times t_{add} + 100/100 \times t_{add} = 11 \times t_{add}$   
 Speedup =  $110/11 = 10$  (10% of potential)

Eff =  $\frac{S}{P} \times 100\%$

Processors عدد ال

Assumes load can be balanced across processors

بخط سببي



## Scaling Example (cont)

What if matrix size is  $100 \times 100$ ?

Single processor: Time =  $(10 + 10000) \times t_{add}$

10 processors

Time =  $10 \times t_{add} + 10000/10 \times t_{add} = 1010 \times t_{add}$   
 Speedup =  $10010/1010 = 9.9$  (99% of potential)

Eff (وهذا جيد)

100 processors

Time =  $10 \times t_{add} + 10000/100 \times t_{add} = 110 \times t_{add}$   
 Speedup =  $10010/110 = 91$  (91% of potential)

Eff

Assuming load balanced





# Strong vs Weak Scaling

- Strong scaling: problem size fixed
  - As in example
- Weak scaling: problem size proportional to number of processors
  - 10 processors,  $10 \times 10$  matrix
    - Time =  $20 \times t_{\text{add}}$
  - 100 processors,  $32 \times 32$  matrix
    - Time =  $10 \times t_{\text{add}} + 1000/100 \times t_{\text{add}} = 20 \times t_{\text{add}}$
  - Constant performance in this example

لو كبرنا المسألة  
لما نستخدم عدد  
أكبر من ال  
Processors  
بمبير  
Weak  
Scaling



## Contents


- 6.1 Introduction
- 6.2 The Difficulty of Creating Parallel Programs
- 6.3 SISD, MIMD, SIMD, SPMD, and Vector
- 6.4 Hardware Multithreading
- 6.5 Shared Memory Multiprocessors
- 6.6 Introduction to Graphics Processing Units
- 6.7 Clusters and Message-Passing Multiprocessors
- 6.8 Introduction to Multiprocessor Network Topologies
- 6.10 Multiprocessor Benchmarks and Performance Models
- 6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU
- 6.12 Multiple Processors and Matrix Multiply
- 6.13 Fallacies and Pitfalls
- 6.14 Concluding Remarks

# Contents

- 6.3 SISD, MIMD, SIMD, SPMD, and Vector Flynn's Classification
- Vector Processors
- SIMD Instruction Extensions

## Instruction and Data Streams

An alternate classification



		Data Streams	
		Single	Multiple
Instruction Streams	Single	<b>SISD:</b> → أبسط انسي Intel Pentium 4	<b>SIMD:</b> SSE (vector) instructions of x86
	Multiple	<b>MISD:</b> Data بترتيل Instruction بترتيل No examples today	<b>MIMD:</b> intel core i7 Intel Xeon e5345

SPMD: Single Program Multiple Data

- A parallel program on a MIMD computer
- Conditional code for different processors

Processor بترتيل تنفيذ مختلفة على كل Processor الآخر



# Vector Processors

بتطلع  
Data ال  
Streams  
L value

- Highly pipelined function units
- Stream data from/to vector registers to units

- Data collected from memory into registers
- Results stored from registers to memory

## Example: Vector extension to RISC-V

- v0 to v31: 32 × 64-element registers, (64-bit elements)
- Vector instructions
  - fld.v, fsd.v: load/store vector (Floating Point)
  - fadd.d.v: add vectors of double
  - fadd.d.vs: add scalar to each element of vector of double

Vector الريب ممكن ←

Scalar  
مثلا →

- Significantly reduces instruction-fetch bandwidth  
(element 64) بتغذ instruction ل ←



## Example: DAXPY ( $Y = a \times X + Y$ )

### ① Conventional RISC-V code:

```

fld    f0, a(x3) // load scalar a
addi   x5, x19, 512 // end of array x
loop: fld    f1, 0(x19) // load x[i]
      fmul.d f1, f1, f0 // a * x[i]
      fld    f2, 0(x20) // load y[i]
      fadd.d f2, f2, f1 // a * x[i] + y[i]
      fsd    f2, 0(x20) // store y[i]
      addi  x19, x19, 8 // increment index to x
      addi  x20, x20, 8 // increment index to y
      bltu  x19, x5, loop // repeat if not done
    
```

بأشركا الخباية  
ل بتغذوا 64 مرة  
 $2 + 8 * 64 = 514$  instruction

بأشركا ال Array


for(i=0, i<64, i++)  
y[i] = a \* x[i] + y[i];

### ② Vector RISC-V code:

```

fld    f0, a(x3) // load scalar a
fld.v  v0, 0(x19) // load vector x
fmul.d.vs v0, v0, f0 // vector-scalar multiply
fld.v  v1, 0(x20) // load vector y
fadd.d.v v1, v1, v0 // vector-vector add
fsd.v  v1, 0(x20) // store vector y
    
```

هذا البرنامج أفضل بس 6 instruction





# Vector vs. Scalar

- Vector architectures and compilers
  - Simplify data-parallel programming
  - Explicit statement of absence of loop-carried dependences
    - Reduced checking in hardware
  - Regular access patterns benefit from interleaved and burst memory
  - Avoid control hazards by avoiding loops
- More general than ad-hoc media extensions (such as MMX, SSE)
  - Better match with compiler technology



## SIMD (Single Instruction Multiple Data)

- Operate elementwise on vectors of data
  - E.g., MMX and SSE instructions in x86
    - Multiple data elements in 128-bit wide registers
- All processors execute the same instruction at the same time
  - Each with different data address, etc.
- Simplifies synchronization
- Reduced instruction control hardware
- Works best for highly data-parallel applications

طالو بفریم  
AVX 256 bit  
AVX 512 bit  
↳ بقدر تقسیمه  
double ← 8 DP ↓  
single ← 16 SP



# Vector vs. Multimedia Extensions

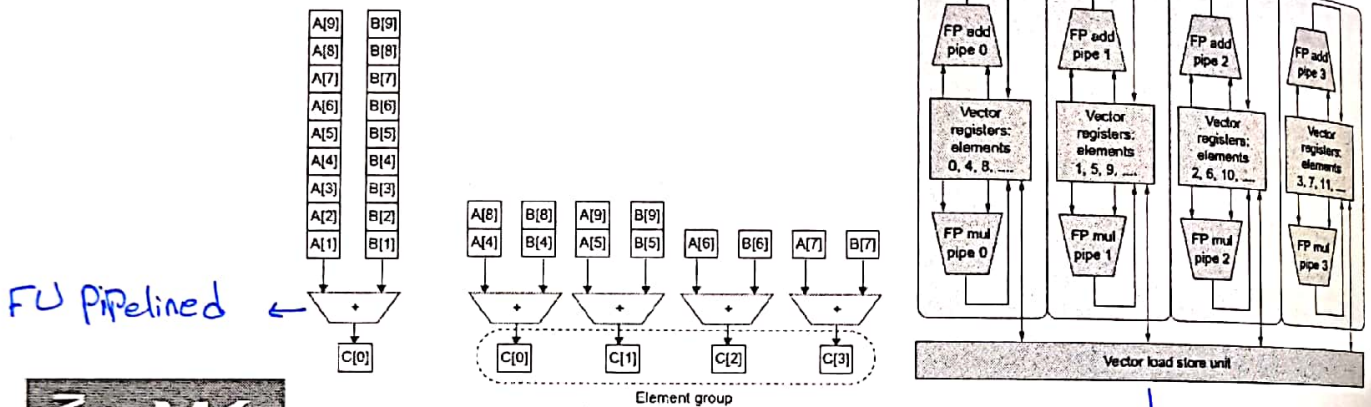
مثل SIMD

- Vector instructions have a variable vector width, multimedia extensions have a fixed width
- Vector instructions support strided access, multimedia extensions do not
- Vector units can be combination of pipelined and arrayed functional units:

انست  
تحدد طول  
Vector ال

لازم تشغيل  
Data ع  
حجب بعضها

يكون في قفزة بين Data والثابت



FU Pipelined



4 FU حارة

بذل ما كل الشغل ينعمل ع

تقسوا ال register ل 4 اجزاء

## Contents

بسول وبفلا عدد ال input/output لل register الواحد

- 6.1 Introduction
- 6.2 The Difficulty of Creating Parallel Programs
- 6.3 SISD, MIMD, SIMD, SPMD, and Vector
- 6.4 Hardware Multithreading
- 6.5 Shared Memory Multiprocessors
- 6.6 Introduction to Graphics Processing Units
- 6.7 Clusters and Message-Passing Multiprocessors
- 6.8 Introduction to Multiprocessor Network Topologies
- 6.10 Multiprocessor Benchmarks and Performance Models
- 6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU
- 6.12 Multiple Processors and Matrix Multiply
- 6.13 Fallacies and Pitfalls
- 6.14 Concluding Remarks



# Multithreading

Sol 4 Hardware Multithreading

Performing multiple threads of execution in parallel → *Multiple process* *واحد بنفذ Core*

- Replicate registers, PC, etc.
- Fast switching between threads

## ① Fine-grain multithreading

- Switch threads after each cycle
- Interleave instruction execution
- If one thread stalls, others are executed

## ② Coarse-grain multithreading

- Only switch on long stall (e.g., L2-cache miss)
- Simplifies hardware, but doesn't hide short stalls (eg, data hazards)



# Simultaneous Multithreading

③ In multiple-issue dynamically scheduled processor (*SMT*)

- Schedule instructions from multiple threads
- Instructions from independent threads execute when function units are available
- Within threads, dependencies handled by scheduling and register renaming

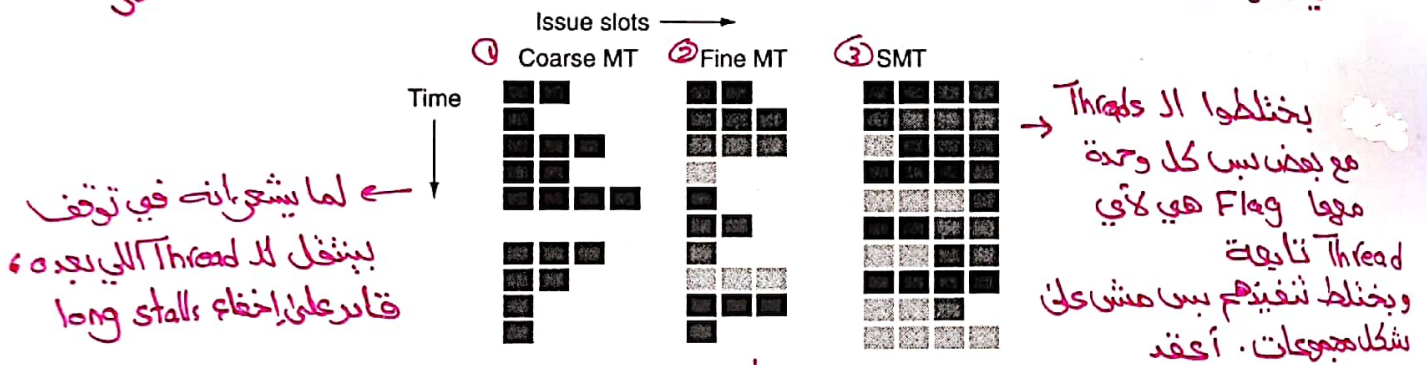
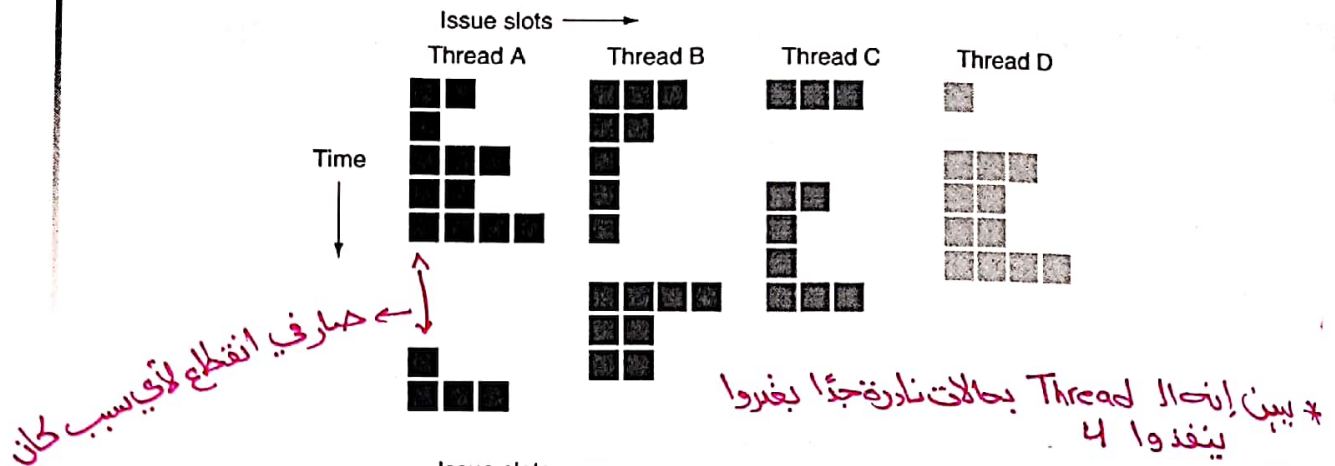
Example: Intel Pentium-4 HT → *Hyper Threading*

- Two threads: duplicated registers, shared function units and caches





# Multithreading Example



له كل cycle يجيب من Thread

وهاد كويس لأنه بياعد ال Instruction بـ thread واحد عن بعض وهاد بحسن لو كان في dependence، قادر على إخفاء

Short/small  
• Stalls  
أبسط

## Future of Multithreading

- Will it survive? In what form?
- Power considerations ⇒ simplified microarchitectures
  - Simpler forms of multithreading
- Tolerating cache-miss latency
  - Thread switch may be most effective
- Multiple simple cores might share resources more effectively

# Contents

- 6.1 Introduction
- 6.2 The Difficulty of Creating Parallel Programs
- 6.3 SISD, MIMD, SIMD, SPMD, and Vector
- 6.4 Hardware Multithreading
- 6.5 Shared Memory Multiprocessors
- 6.6 Introduction to Graphics Processing Units
- 6.7 Clusters and Message-Passing Multiprocessors
- 6.8 Introduction to Multiprocessor Network Topologies
- 6.10 Multiprocessor Benchmarks and Performance Models
- 6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU
- 6.12 Multiple Processors and Matrix Multiply
- 6.13 Fallacies and Pitfalls
- 6.14 Concluding Remarks

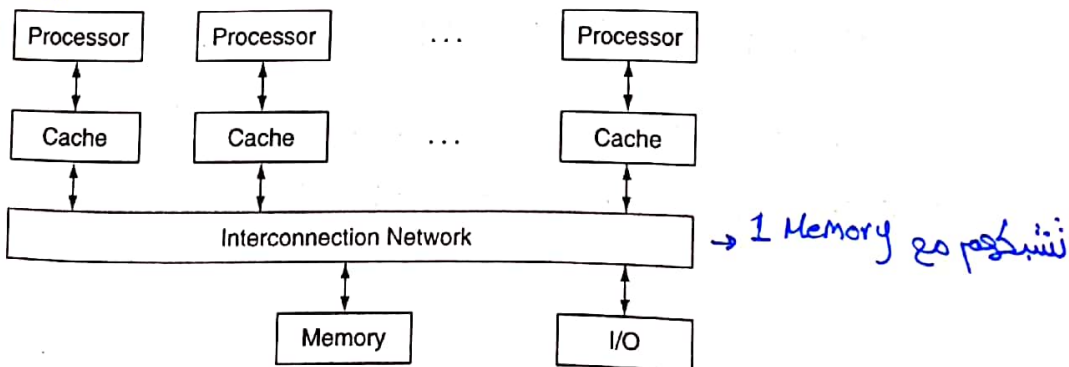


## Shared Memory

- SMP: shared memory multiprocessor
  - Hardware provides single physical address space for all processors
  - Synchronize shared variables using locks
  - Memory access time

① UMA (uniform) vs. NUMA (nonuniform) → أنواع الـ shared memory

له تعيين نفس البعد  
في الـ Mem




§6.5 Multicore and Other Shared Memory Multiprocessors



# Example: Sum Reduction


- Sum 64,000 numbers on 64 processor UMA
  - Each processor has ID:  $0 \leq P_n \leq 63$  كشان عنا 64 processor
  - Partition 1000 numbers per processor
  - Initial summation on each processor
 

```
sum[Pn] = 0;
for (i = 1000*Pn;
    i < 1000*(Pn+1); i += 1)
    sum[Pn] += A[i];
```
- Now need to add these partial sums
  - Reduction: divide and conquer
  - Half the processors add pairs, then quarter, ...
  - Need to synchronize between reduction steps



# Example: Sum Reduction

(التصنيف ويكون مسؤولية النص التمثالي انه ينقل قيمة الفوقاني)



```
half = 64;
do
```

```
    synch();
```

```
    if (half%2 != 0 && Pn == 0)
```

```
        sum[0] += sum[half-1];
```

```
        /* conditional sum needed when half is odd;
        processor0 gets missing element */
```

```
        half = half/2; /* dividing line on who sums */
```

```
        if (Pn < half) sum[Pn] += sum[Pn+half];
```

```
    while (half > 1);
```

في الهدف منه  
اذا كان الرقم  
فزي كشان نعالج  
المشكلة





# Contents

- 6.1 Introduction
- 6.2 The Difficulty of Creating Parallel Programs
- 6.3 SISD, MIMD, SIMD, SPMD, and Vector
- 6.4 Hardware Multithreading
- 6.5 Shared Memory Multiprocessors
- 6.6 Introduction to Graphics Processing Units
- 6.7 Clusters and Message-Passing Multiprocessors
- 6.8 Introduction to Multiprocessor Network Topologies
- 6.10 Multiprocessor Benchmarks and Performance Models
- 6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU
- 6.12 Multiple Processors and Matrix Multiply
- 6.13 Fallacies and Pitfalls
- 6.14 Concluding Remarks



## History of GPUs

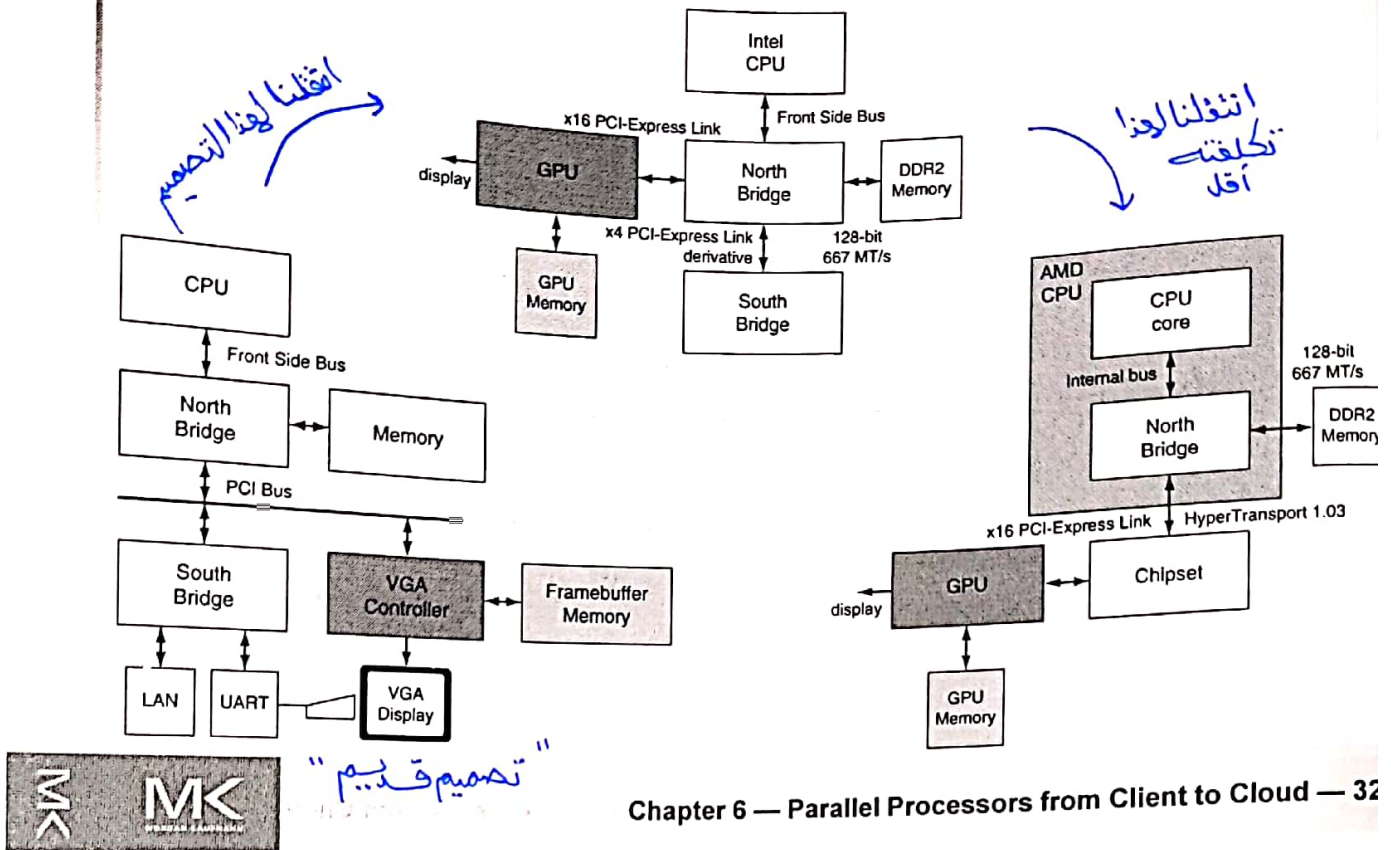
- Early video cards → بالثانينات
  - Frame buffer memory with address generation for video output
- 3D graphics processing → بنهاية الثمانينات
  - Originally high-end computers (e.g., SGI)
  - Moore's Law ⇒ lower cost, higher density
  - 3D graphics cards for PCs and game consoles → هذا البلايستيشن
- Graphics Processing Units GPUs
  - Processors oriented to 3D graphics tasks
  - Vertex/pixel processing, shading, texture mapping, rasterization  
تحويل الأنماط المضادة

6.6 Introduction to Graphics Processing Units



تحويل الفيديو لـ Signals

# Graphics in the System (Computer)



## GPU Architectures

- Processing is highly data-parallel → Abot of data
- GPUs are highly multithreaded بشكل عظيم
- Use thread switching to hide memory latency نفس الحسابات
- Less reliance on multi-level caches
- Graphics memory is wide and high-bandwidth متنوعة التطبيقات
- Trend toward general purpose GPUs
- Heterogeneous CPU/GPU systems أفضل من ال CPU
- CPU for sequential code, GPU for parallel code أفضل من ال CPU
- Programming languages/APIs أفضل من ال CPU
- DirectX, OpenGL
- C for Graphics (Cg), High Level Shader Language (HLSL)
- Compute Unified Device Architecture (CUDA) Nvidia



# Modern Computer

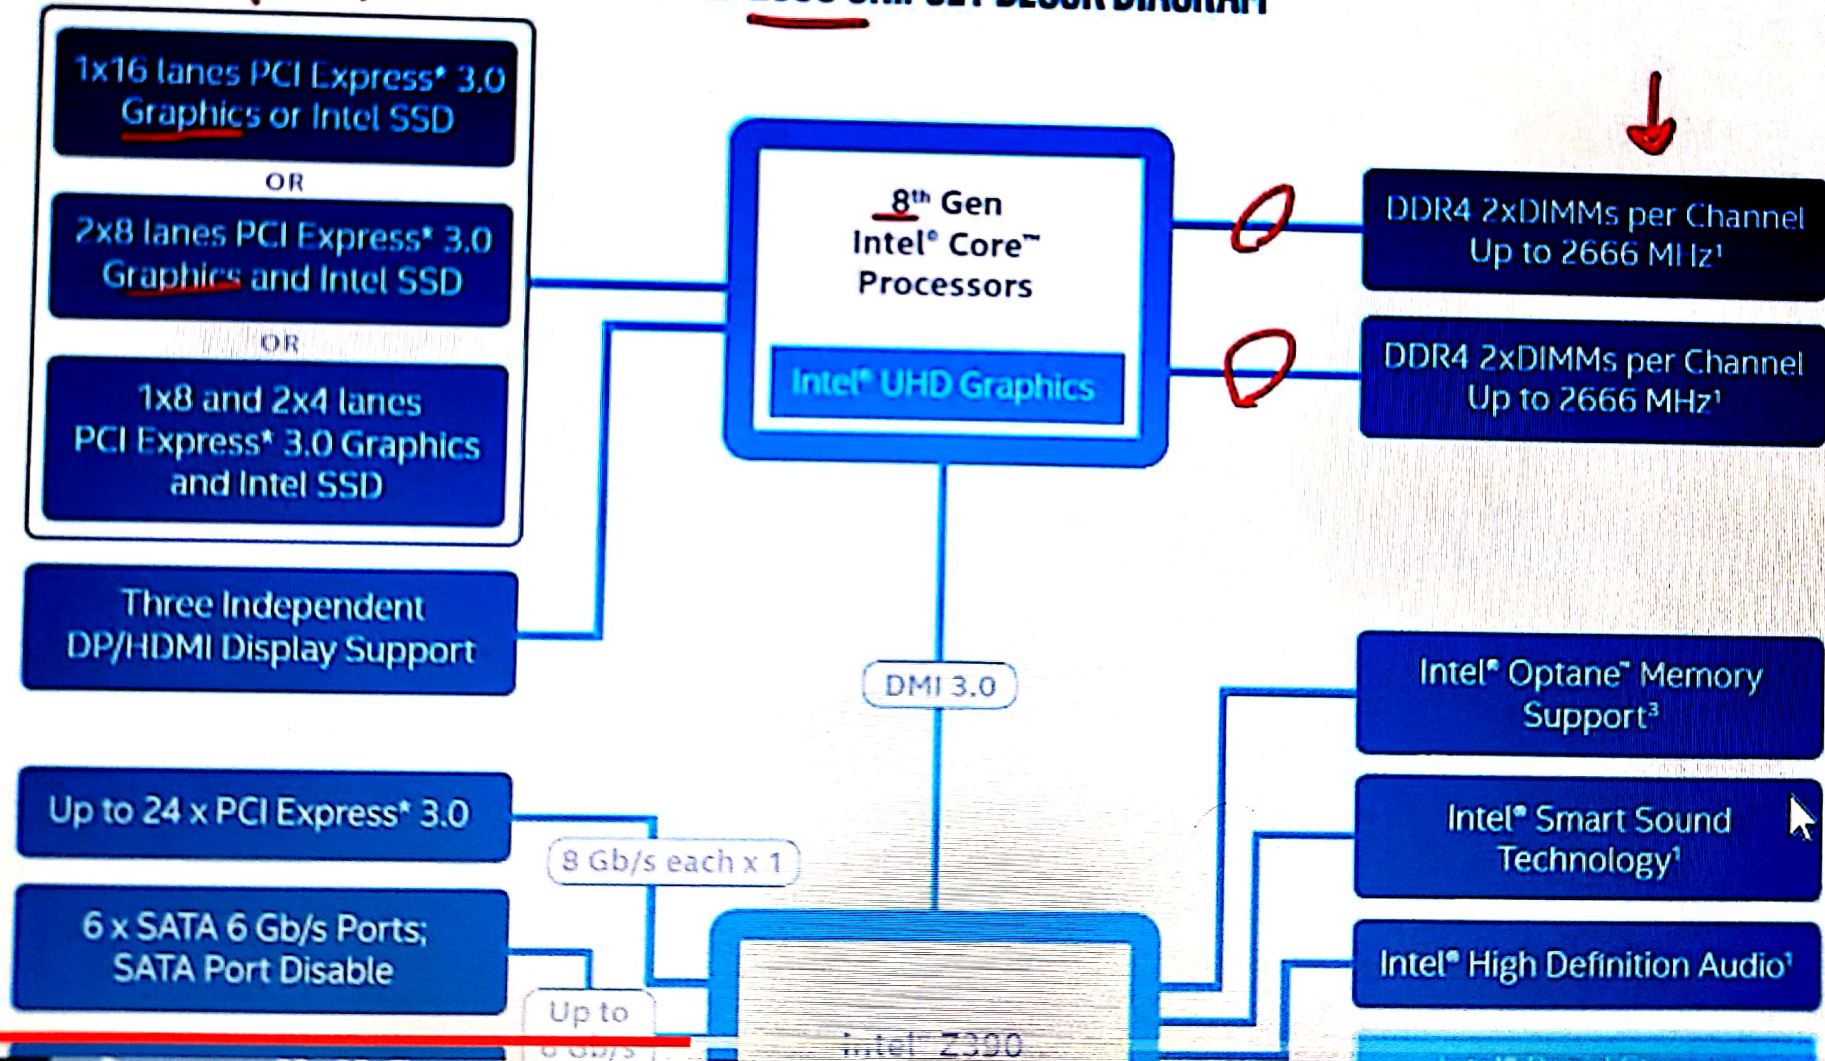
Thursday, April 23, 2020 9:45 AM

i3

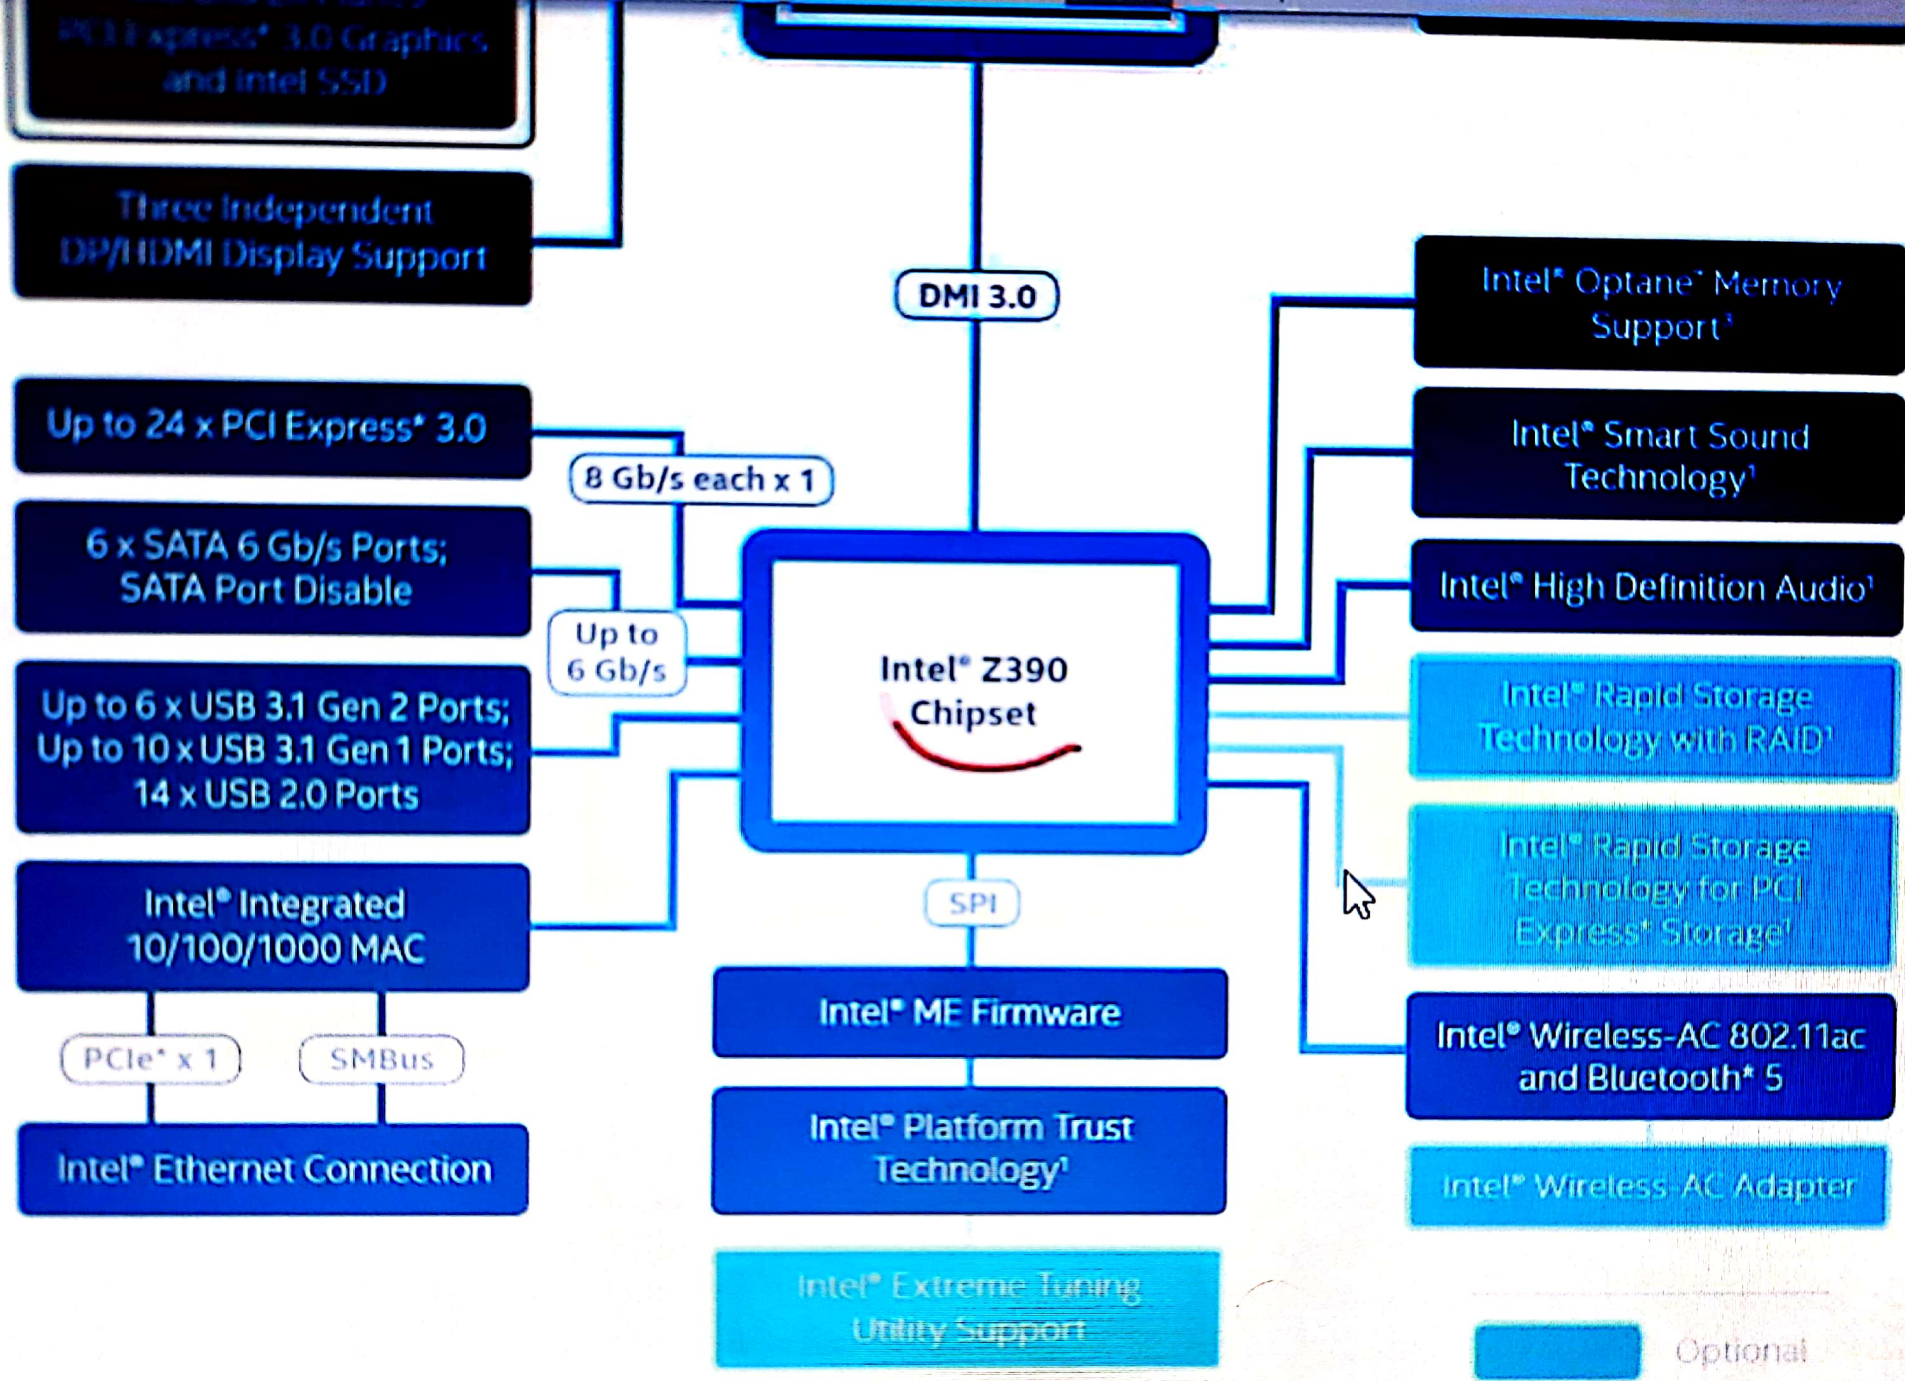
i7

Nvidia

## INTEL® Z390 CHIPSET BLOCK DIAGRAM

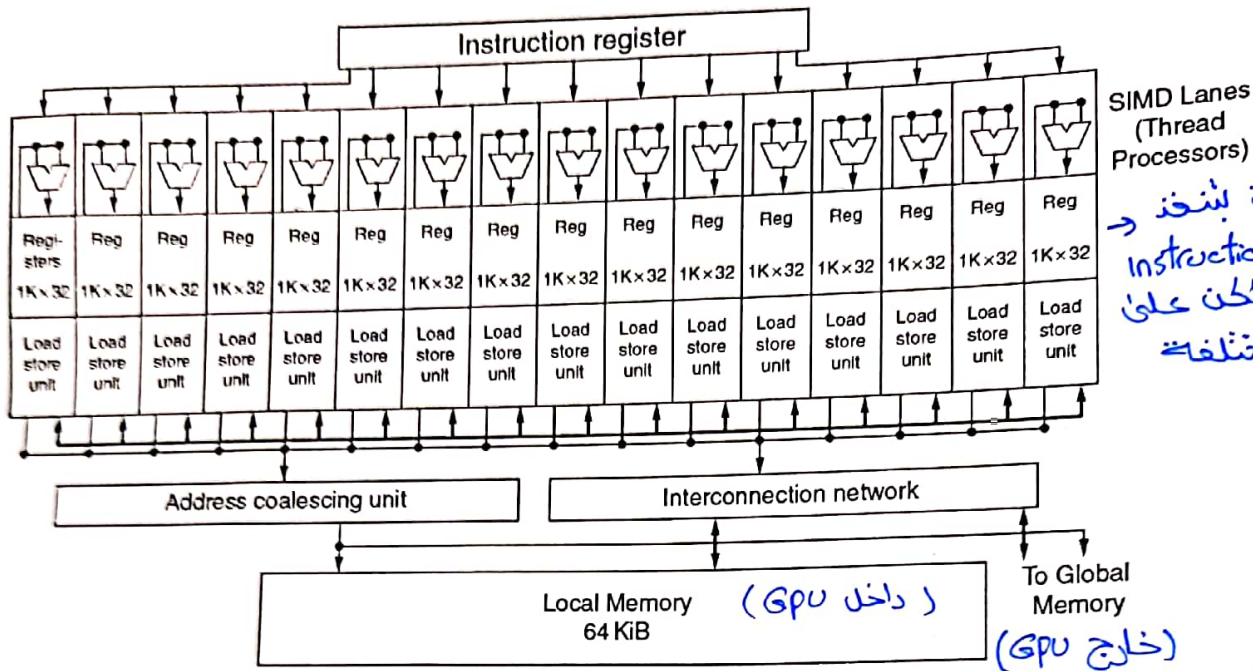






# Example: NVIDIA Fermi

Multiple SIMD processors, each as shown:

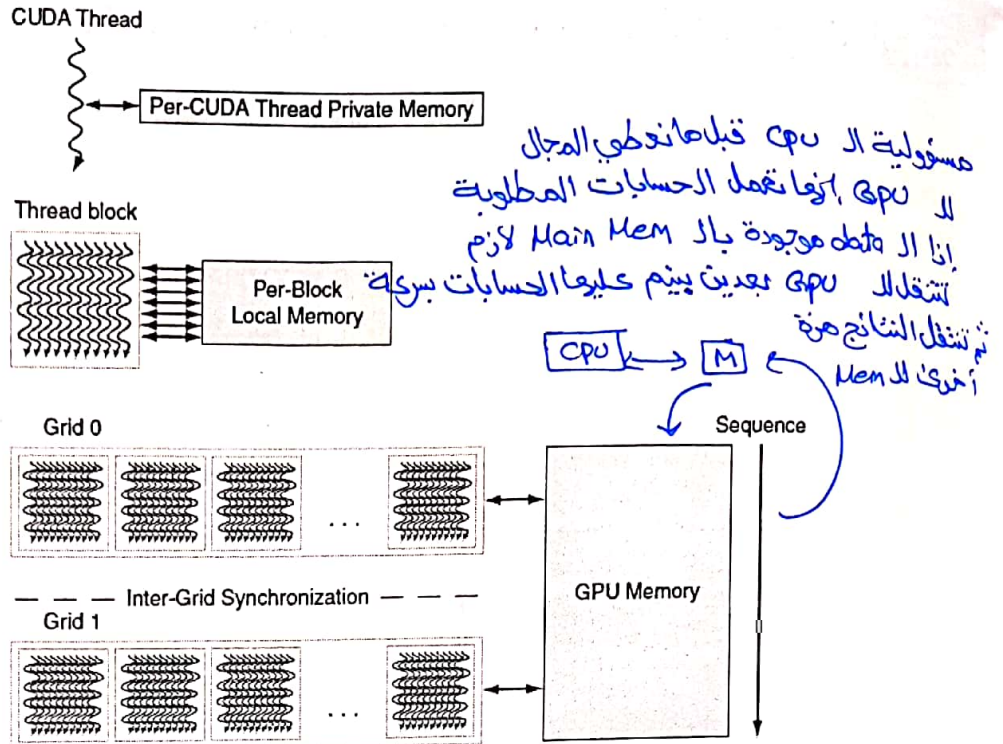


# Example: NVIDIA Fermi

- SIMD Processor: 16 SIMD lanes
- SIMD instruction
  - Operates on 32 element wide threads
  - Dynamically scheduled on 16-wide processor over 2 cycles
- 32K x 32-bit registers spread across lanes
  - 64 registers per thread context



# GPU Memory Structures



## Classifying GPUs

- Don't fit nicely into SIMD/MIMD model
  - Conditional execution in a thread allows an illusion of MIMD
    - But with performance degradation
    - Need to write general purpose code with care

	Static: Discovered at Compile Time	Dynamic: Discovered at Runtime
Instruction-Level Parallelism	VLIW	Superscalar
Data-Level Parallelism	SIMD or Vector	<u>Tesla Multiprocessor</u> GPUs



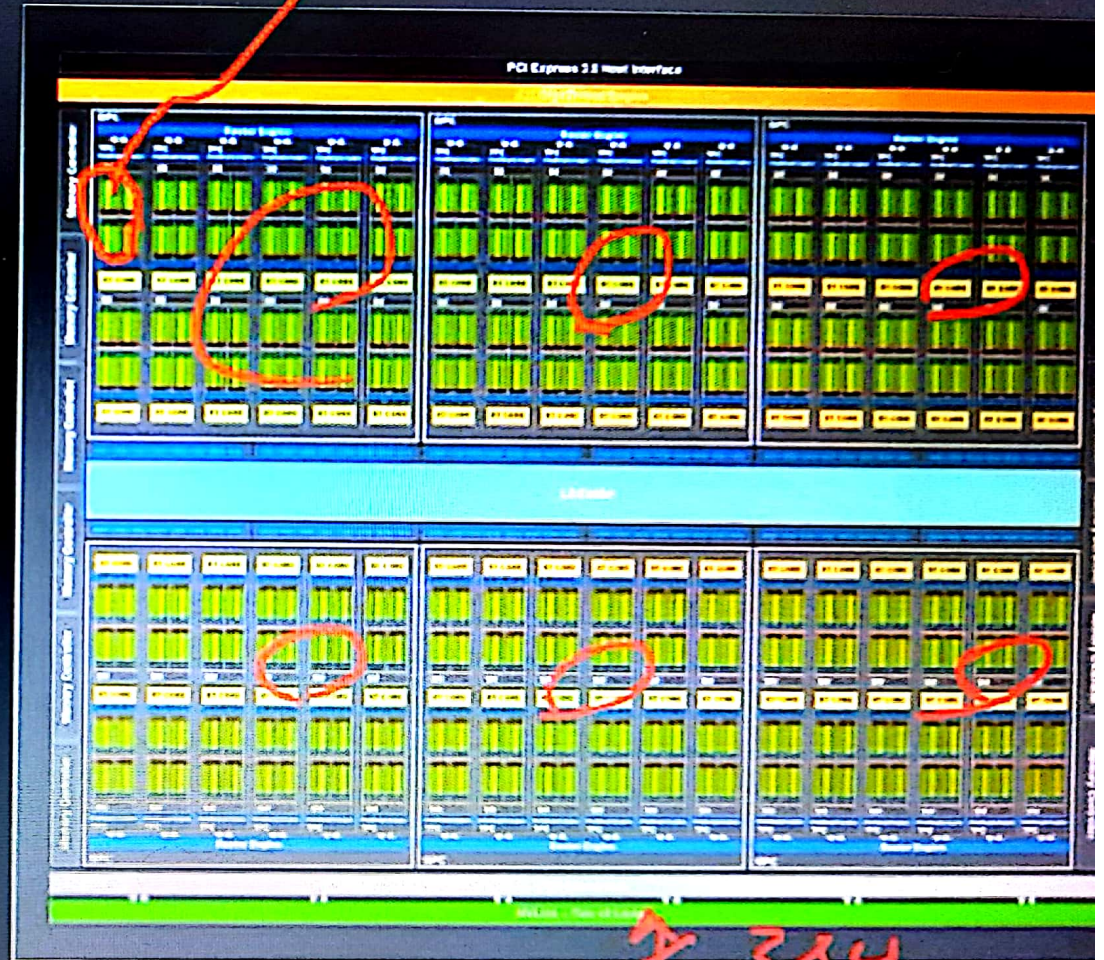


# INTRODUCING TURING

## TU102 - FULL CONFIG

18.6 BILLION TRANSISTORS

SM	72
<u>CUDA CORES</u>	4608
TENSOR CORES	576
RT CORES	72
GEOMETRY UNITS	36
TEXTURE UNITS	288
ROP UNITS	96
MEMORY	<u>384-bit 7 GHz GDDR6</u>
NVLINK CHANNELS	2



384  
GM 8GB





# RTX 2080 Ti

Thursday, April 23, 2020 9:53 AM

## INTRODUCING TURING

**TU102 - FULL CONFIG**  
18.6 BILLION TRANSISTORS

SM	72
CUDA CORES	4608
TENSOR CORES	576
RT CORES	72
GEOMETRY UNITS	36





# Putting GPUs into Perspective

CPU

Feature	Multicore with SIMD	GPU
SIMD processors	4 to 8	8 to 16
SIMD lanes/processor	2 to 4	8 to 16 <u>or 32 or 64</u>
Multithreading hardware support for SIMD threads	2 to 4	16 to 32
Typical ratio of single precision to double-precision performance	2:1	2:1
Largest cache size	8 MB	0.75 MB
Size of memory address	64-bit	64-bit
Size of main memory	8 GB to 256 GB	4 GB to 6 GB
Memory protection at level of page	Yes	Yes
Demand paging	Yes	No
Integrated scalar processor/SIMD processor	Yes	No
Cache coherent	Yes	No



## Guide to GPU Terms

Type	More descriptive name	Closest old term outside of GPUs	Official CUDA/NVIDIA GPU term	Book definition
Program abstractions	Vectorizable Loop	Vectorizable Loop	Grid	A vectorizable loop, executed on the GPU, made up of one or more Thread Blocks (bodies of vectorized loop) that can execute in parallel.
	Body of Vectorized Loop	Body of a (Strip-Mined) Vectorized Loop	Thread Block	A vectorized loop executed on a multithreaded SIMD Processor, made up of one or more threads of SIMD instructions. They can communicate via Local Memory.
	Sequence of SIMD Lane Operations	One iteration of a Scalar Loop	CUDA Thread	A vertical cut of a thread of SIMD instructions corresponding to one element executed by one SIMD Lane. Result is stored depending on mask and predicate register.
Machine object	A Thread of SIMD Instructions	Thread of Vector Instructions	Warp	A traditional thread, but it contains just SIMD instructions that are executed on a multithreaded SIMD Processor. Results stored depending on a per-element mask.
	SIMD Instruction	Vector Instruction	PTX Instruction	A single SIMD instruction executed across SIMD Lanes.
Processing hardware	Multithreaded SIMD Processor	(Multithreaded) Vector Processor	Streaming Multiprocessor	A multithreaded SIMD Processor executes threads of SIMD instructions, independent of other SIMD Processors.
	Thread Block Scheduler	Scalar Processor	Giga Thread Engine	Assigns multiple Thread Blocks (bodies of vectorized loop) to multithreaded SIMD Processors.
	SIMD Thread Scheduler	Thread scheduler in a Multithreaded CPU	Warp Scheduler	Hardware unit that schedules and issues threads of SIMD instructions when they are ready to execute; includes a scoreboard to track SIMD Thread execution.
	SIMD Lane	Vector lane	Thread Processor	A SIMD Lane executes the operations in a thread of SIMD instructions on a single element. Results stored depending on mask.
Memory hardware	GPU Memory	Main Memory	Global Memory	DRAM memory accessible by all multithreaded SIMD Processors in a GPU.
	Local Memory	Local Memory	Shared Memory	Fast local SRAM for one multithreaded SIMD Processor, unavailable to other SIMD Processors.
	SIMD Lane Registers	Vector Lane Registers	Thread Processor Registers	Registers in a single SIMD Lane allocated across a full thread block (body of vectorized loop).



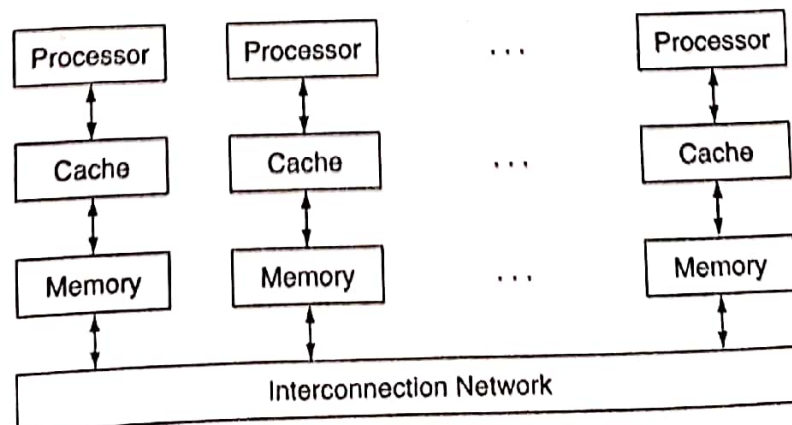
# Contents

- 6.1 Introduction
- 6.2 The Difficulty of Creating Parallel Programs
- 6.3 SISD, MIMD, SIMD, SPMD, and Vector
- 6.4 Hardware Multithreading
- 6.5 Shared Memory Multiprocessors
- 6.6 Introduction to Graphics Processing Units
- 6.7 Clusters and Message-Passing Multiprocessors
- 6.8 Introduction to Multiprocessor Network Topologies
- 6.10 Multiprocessor Benchmarks and Performance Models
- 6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU
- 6.12 Multiple Processors and Matrix Multiply
- 6.13 Fallacies and Pitfalls
- 6.14 Concluding Remarks



## Message Passing

- Each processor has private physical address space
- Hardware sends/receives messages between processors





ما يتعاونوا عن طريق الذاكرة

# Loosely Coupled Clusters

Network of independent computers

- Each has private memory and OS
- Connected using I/O system  
E.g., Ethernet/switch, Internet → يطبق

Suitable for applications with independent tasks

- Web servers, databases, simulations, ...

High availability, scalable, affordable

Problems

- Administration cost (prefer virtual machines) ← انا واحد توصل البيئية بكملوا ← يستخدمه معروا عنشان يسول
- Low interconnect bandwidth ← كسب  
c.f. processor/memory bandwidth on an SMP  
↓  
هدول أفضل



## Sum Reduction (Again)

Sum 64,000 on 64 processors →

64000 موزعة  
64 P كل P يتاخذ 1000

First distribute 1000 numbers to each

- The do partial sums

```
sum = 0;  
for (i = 0; i < 1000; i += 1)  
    sum += AN[i];
```

Reduction

- Half the processors send, other half receive and add
- The quarter send, quarter receive and add, ...



# Sum Reduction (Again)

- Given send() and receive() operations

```
limit = 64; half = 64; /* 64 processors */
do
    half = (half+1)/2; /* send vs. receive
                        dividing line */
    if (Pn >= half && Pn < limit)
        send(Pn - half, sum);
    if (Pn < (limit/2))
        sum += receive();
    limit = half; /* upper limit of senders */
while (half > 1); /* exit with final sum */
```

هنا: الجديد المدمج

- Send/receive also provide synchronization
- Assumes send/receive take similar time to addition



## Grid Computing

- Separate computers interconnected by long-haul networks
  - E.g., Internet connections
  - Work units farmed out, results sent back
- Can make use of idle time on PCs
  - E.g., SETI@home, World Community Grid

\* يعني لو الجهاز مو بيستخدمه بقدر اتيجه عالانترنت لحد به بيستخدمه \*



# Contents

- 6.1 Introduction
- 6.2 The Difficulty of Creating Parallel Programs
- 6.3 SISD, MIMD, SIMD, SPMD, and Vector
- 6.4 Hardware Multithreading
- 6.5 Shared Memory Multiprocessors
- 6.6 Introduction to Graphics Processing Units
- 6.7 Clusters and Message-Passing Multiprocessors
- 6.8 Introduction to Multiprocessor Network Topologies
- 6.10 Multiprocessor Benchmarks and Performance Models
- 6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU
- 6.12 Multiple Processors and Matrix Multiply
- 6.13 Fallacies and Pitfalls
- 6.14 Concluding Remarks

MK

## Interconnection Networks

تشبيك ال Processors أو ال Computers مع بعض


مخانة ال Torus 2D عن ال Mesh التي بالصورة

انه يسرع عملية التواصل لانه تشبيك الاطراف مع بعض  
وخلال ال latency .

### Network topologies


- Arrangements of processors, switches, and links

①




Bus

②




Ring

③




2D Mesh  
(Torus)

④



N-cube (N = 3)  
dimension =  $\log N$   
latency = dimension

⑤



Fully connected latency = 1  
كل نقطة مشبوكة مع الكل

$\frac{N(N-1)}{2}$


مشكلته مكلف جدا وعدد الاسلاك كبير  $\uparrow$

6.8 Introduction to Multiprocessor Network Topologies

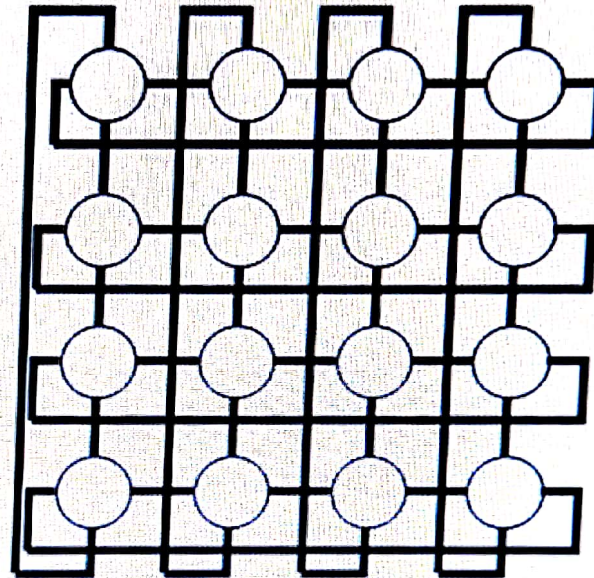
MK

# Interconnection Networks


Sunday, April 26, 2020 10:02 AM



Mesh Topology



2D Torus

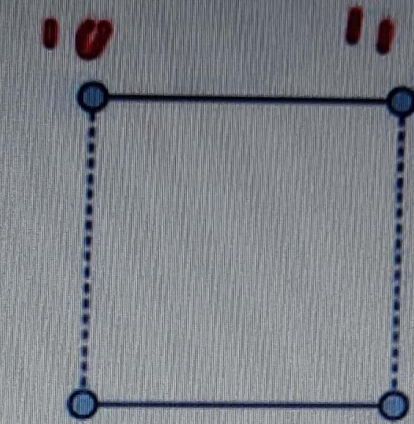


3D Torus

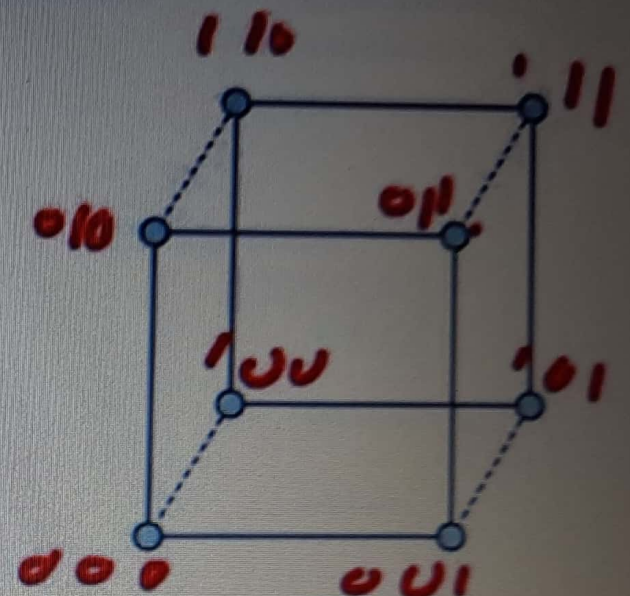




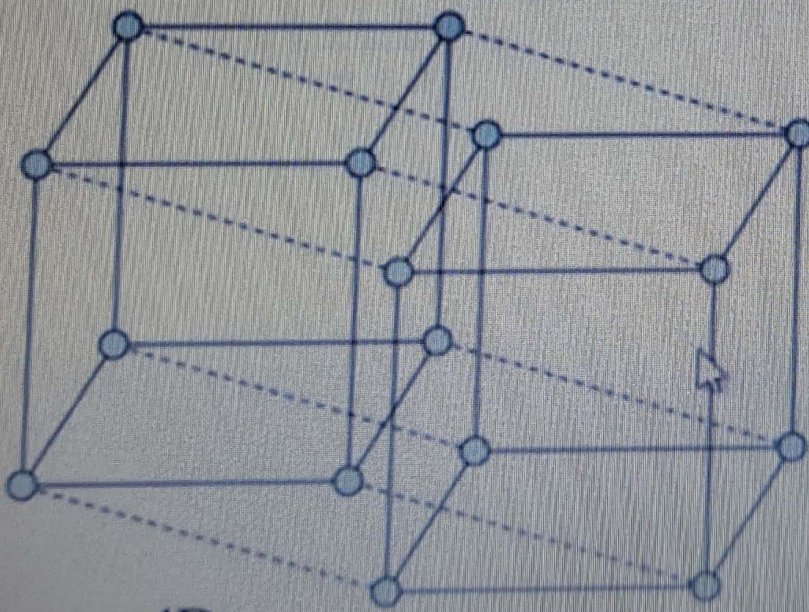
1D



2D



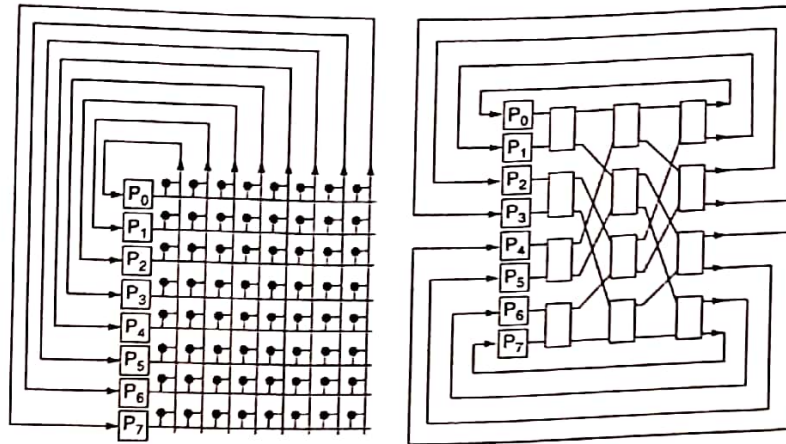
3D



4D



# Multistage Networks




a. Crossbar

b. Omega network

لنقلوا من التكلفة

مكلف جدًا ال switch غالي



c. Omega network switch box

ليست ختموا ال switch box



## Network Characteristics

### Performance

- Latency per message (unloaded network) أقل ما يمكن
- Throughput بمعدل حساب ال cost
- Link bandwidth → يعتمد على عدد الأسلاك وال frequency
- Total network bandwidth
- Bisection bandwidth → لما أقسم ال network لـ 2 link يضل عندي
- Congestion delays (depending on traffic)

بفضل ال

أعلن ما يمكن ليس مش حساب ال cost


### Cost ( معقولة )

### Power

### Routability in silicon







$$d = \log_2 64 = 6$$

$$\frac{64 \times 63}{2}$$

Evaluation category	Bus	Ring	2D mesh	2D torus	Hypercube	Fat tree	Fully connected
<b>Performance</b>							
BW <sub>Bisection</sub> in # links	<u>1</u>	<u>2</u>	<u>8</u>	<u>16</u>	<u>32</u>	32	1024
Max (ave.) hop count	1 (1)	32 (16)	14 (7)	8 (4)	6 (3)	11 (9)	1 (1)
<b>Cost</b>							
I/O ports per switch	NA	3	5	5	7	4	64
Number of switches	NA	64	64	64	64	192	64
Number of net. links	1	64	112	128	192	320	2016
Total number of links	1	128	176	192	256	384	2080

**Figure F.15** Performance and cost of several network topologies for 64 nodes. The bus is the standard reference at unit network link cost and bisection bandwidth. Values are given in terms of bidirectional links and ports. Hop count includes a switch and its output link, but not the injection link at end nodes. Except for the bus, values are given for the number of network links and total number of links, including injection/reception links between end node devices and the network.



Slide 9 : Amdahl's law :

$$T_{old} = T_{old} * (1 - f + f)$$


$$T_{new} = T_{old} * (1 - f + f/p)$$

$$Speed\ up = T_{old} / T_{new} = 1 / (1 - f + f/p)$$

عدد ال Processors

Serial → Told → Parallelizable

← زمن قديم وكبير



← بي اوله الزمن جديد عن طريق ال

Parallel Program

Tnew

\* Serial : لا بد انك تنفذها بتقدير واحد

\* parallelizable : ممكن انجزها لانه بيغذي اكثر من P

$$\therefore \lim_{P \rightarrow \infty} speed\ up = \frac{1}{1 - f} \rightarrow \text{سلي Speedup ممكن انحصار}$$

Slide 11 : Strong Scalling ( تحلي المسألة ثابتة ونزيد عدد ال Processors )

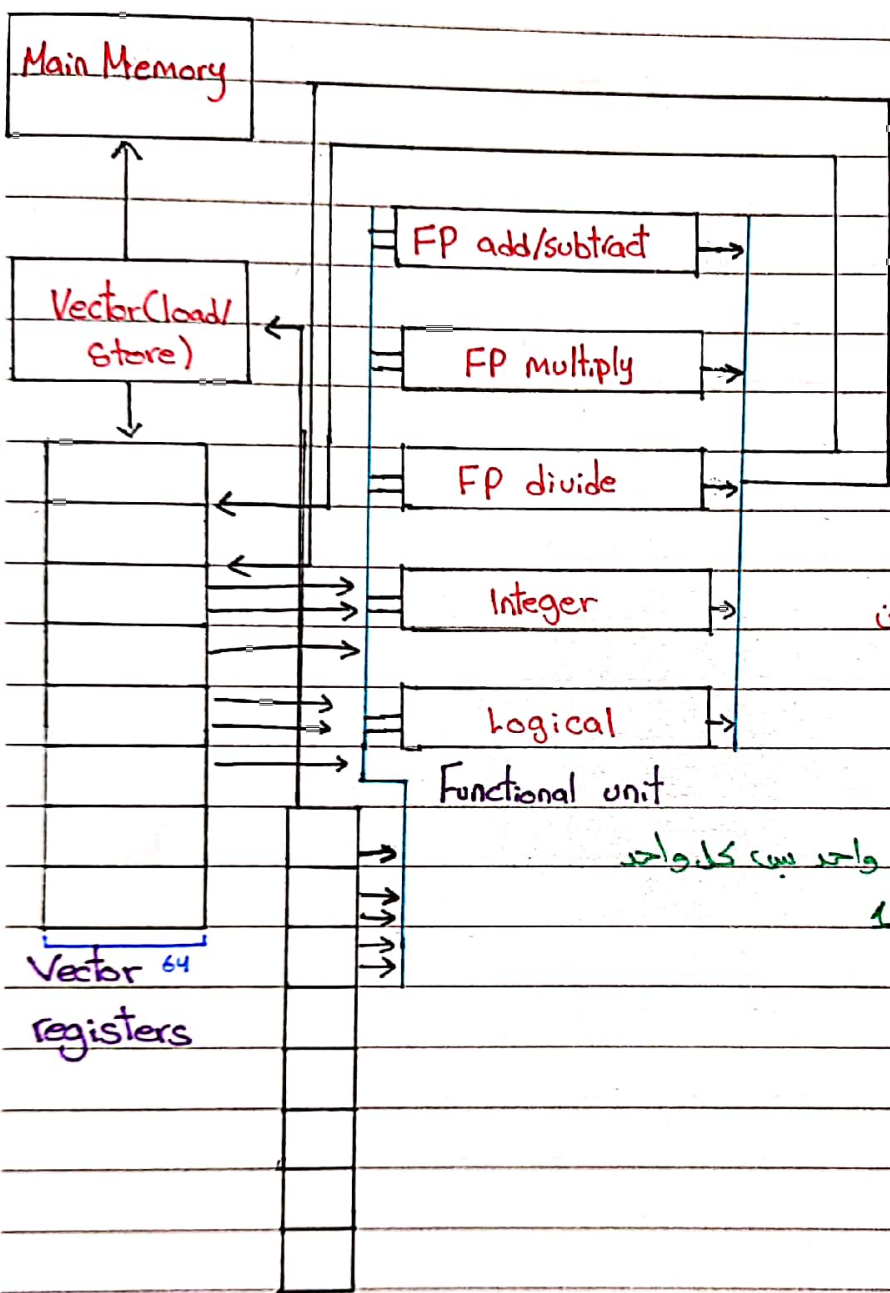
Size			P=1		P=10		P=100	
10 + 100	Speedup	Efficiency	1	100%	5.5	55%	10	10%
10 + 1000	Speedup	Efficiency	1	100%	9.2	92%	51	51%
10 + 10000	Speedup	Efficiency	1	100%	9.9	99%	91	91%

→ اضافة من ال ركوز

\* بالتحليل كل ما زاد عدد ال Processors ال Efficiency ينقل لى المسألة ثابتة .



Slide 16 : Vector Processor (رسمية توضيحية)

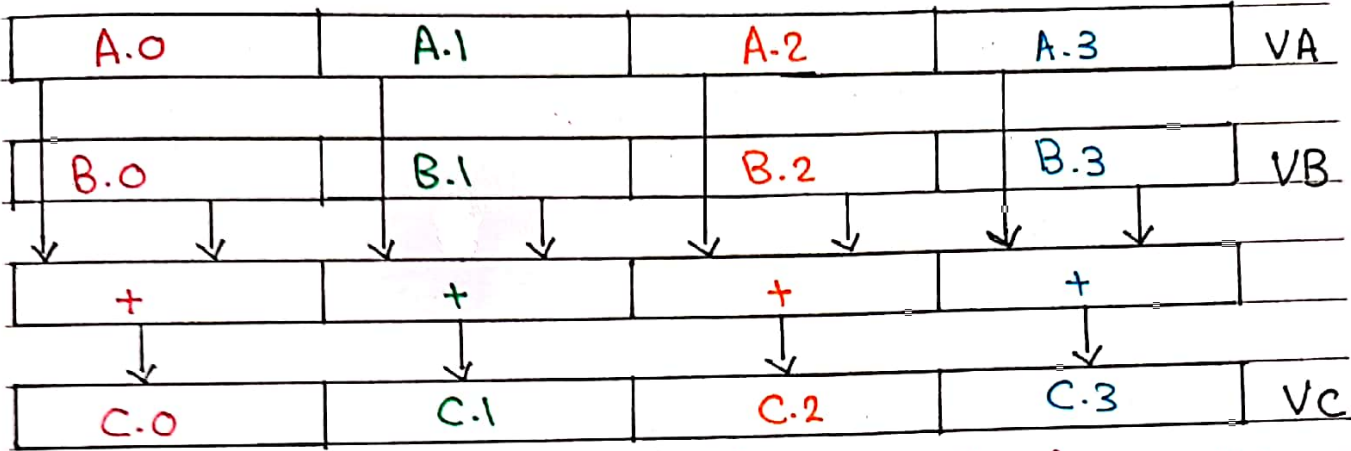


\* المبرمج هنا، اولى داخل  
 Vector computer  
 في Vector registers  
 وهو يعمل على Vectors  
 يكون فيه 32 الـ 32  
 واحد في الـ  
 Vector register  
 value 1  
 Vector  
 من الـ Vector register  
 يكون الـ 64 (64 register)  
 $\Rightarrow 32 VR \times 64 = 2K reg$   
 (مستطيقه هون)

\* Scalar  
 Register  
 32  
 value 1


Scalar registers

Slide 19 & SIMD (add VC, VA, VB)

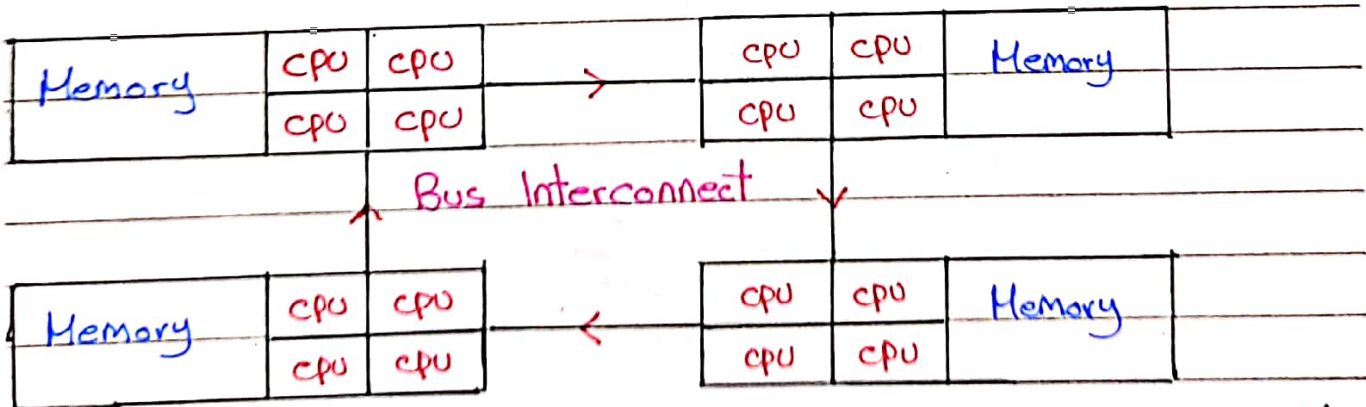


4 additions done by 1 instruction, 4x performance increase

Slide 22 & Multithreading



Slide 27 & NUMA



\* كل مجموعة لها جزء خاص من ال Memory وال 4 مشتركين وأي CPU يشوف الباقي لكن ال Memory اللي في CPU قريبة أكثر من باقي ال Memory ويوصلوا بشكل أسرع.



## Slide 28 : Sum Reduction

## Serial :

$$\text{Sum} = 0;$$

$$\text{for } (i=0; i < 64000; i+=1)$$

$$\text{SUM} += A[i];$$
1000 element P<sub>63</sub>

63999

→ Sum[63]

Array  
A

## Parallel :

$$\text{Sum}[P_n] = 0;$$

$$\text{for } (i=1000 * P_n; i < 1000 * (P_n+1); i+=1)$$

$$\text{Sum}[P_n] += A[i];$$
1000 P<sub>i</sub>  
1000 P<sub>0</sub>

→ Sum[63]

0

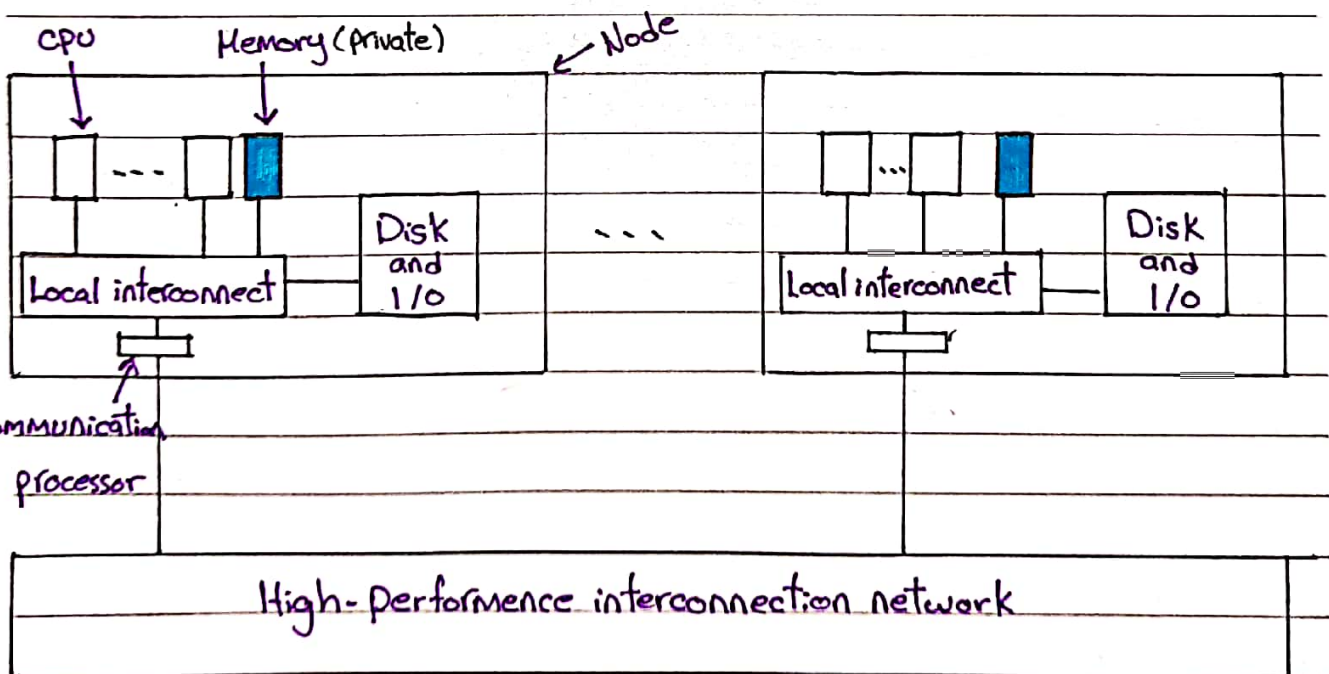
→ مثالين ختم الوقت، ينقسم كل الشغل

على ( 64 Processor ) . حتملى اقل بكثير من ال Serial

بطلعلي 64 جواب وبرد بجمعهم كلهم برقم واحد

وبدل ما اعد loop مثالنا جمع ال 64 بدل اللى يسلايد 29

## Slide 41 : Message-passing Multi-computers



# Parallel Benchmarks

نستعمل عدة تطبيقات مختلفة تقاينا  
بين الآتين

- Linpack: matrix linear algebra
- SPECrate: parallel run of SPEC CPU programs
  - Job-level parallelism
- SPLASH: Stanford Parallel Applications for Shared Memory
  - Mix of kernels and applications, strong scaling
- NAS (NASA Advanced Supercomputing) suite
  - computational fluid dynamics kernels
- PARSEC (Princeton Application Repository for Shared Memory Computers) suite
  - Multithreaded applications using Pthreads and OpenMP

MK

تغذت  
بجودة



# Code or Applications?

- Traditional benchmarks
  - Fixed code and data sets
- Parallel programming is evolving
  - Should algorithms, programming languages, and tools be part of the system?
  - Compare systems, provided they implement a given application
  - E.g., Linpack, Berkeley Design Patterns
- Would foster innovation in approaches to parallelism

بقرى الابنطار بالطرق  
الى الواءلاقة بار  
Parallelism



## الهدف نفوم الperformance ونعرفه هل هو كويس ولا محتاج الحتب بونامج احسنه

# Modeling Performance


- Assume performance metric of interest is achievable GFLOPs/sec  
قيده الcomputer بقدر يعمل  
Floating point operations وهش دا ايما هو المعيار الوحيد بين هو الهم هون
- Measured using computational kernels from Berkeley Design Patterns
- Arithmetic intensity of a kernel
  - FLOPs per byte of memory accessed
- For a given computer, determine
  - Peak GFLOPS (from data sheet)
  - Peak memory bytes/sec (using Stream benchmark)

موجعات للتعيين  
بين ال kernel  
المختلفات

لو تشغلنا كل ال ALUs  
بوعنا ال computer  
قيدينا ممكن نعمل  
GFLOPs



# Roofline Diagram



Attainable GPLOPs/sec  
 = Max ( Peak Memory BW × Arithmetic Intensity, Peak FP Performance )



## Comparing Systems

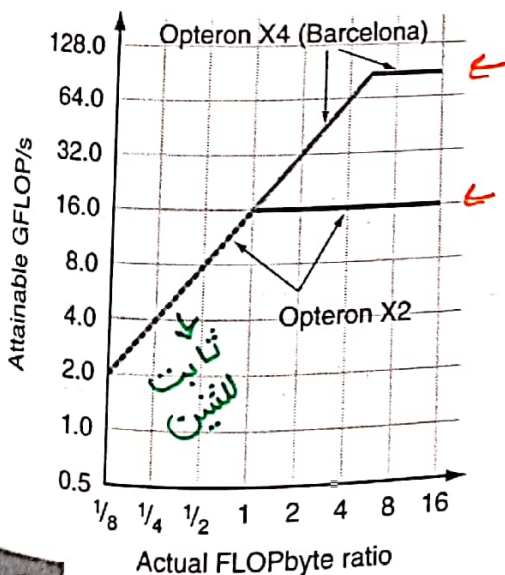
### Example: Opteron X2 vs. Opteron X4

- 2-core vs. 4-core, 2× FP performance/core, 2.2GHz vs. 2.3GHz, 1 × 2 SIMD vs. 2 × 2 SIMD

لم كل cycle تنفيذ فيها 1 SIMD ابعائتين

بعد 2 SIMD op كل cycle  
 كل وحدة فيها عملتين

- Same memory system



To get higher performance on X4 than X2

- Need high arithmetic intensity
- Or working set must fit in X4's 2MB L-3 cache

opteron x2: Peak FP Performance = 2 cores/chip \*  
 2 SIMD/core \* 2 FP/SIMD \* 2.2 Gcycle/s = 2\*2\*2\*2.2 =  
 17.6 GFLOPs/s

opteron x4: Peak FP Performance = 4 cores/chip \* 4 SIMD/  
 core \* 2 FP/SIMD \* 2.3 Gcycles/s = 4\*4\*2\*2.3 =  
 73.6 GFLOPs/s



# Optimizing Performance

تحول لنوع  
العمل  
Performance  
النقل

## Optimize FP performance


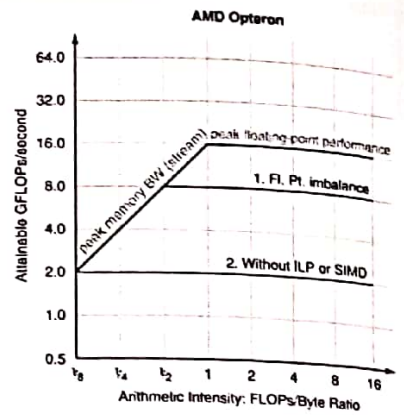
- Balance adds & multiplies
- Improve superscalar ILP and use of SIMD instructions

Instruction level Parallelism

## Optimize memory usage


- Software prefetch
  - Avoid load stalls
- Memory affinity
  - Avoid non-local data accesses

الحسنة تأخذ  
تصنيفها لسطر كامل  
قبل ما تنقل  
للعمود  
اللي بعده



# Optimizing Performance

Choice of optimization depends on arithmetic intensity of code



Arithmetic intensity is not always fixed

- May scale with problem size
- Caching reduces memory accesses
  - Increases arithmetic intensity



# Contents

- 6.1 Introduction
- 6.2 The Difficulty of Creating Parallel Programs
- 6.3 SISD, MIMD, SIMD, SPMD, and Vector
- 6.4 Hardware Multithreading
- 6.5 Shared Memory Multiprocessors
- 6.6 Introduction to Graphics Processing Units
- 6.7 Clusters and Message-Passing Multiprocessors
- 6.8 Introduction to Multiprocessor Network Topologies
- 6.10 Multiprocessor Benchmarks and Performance Models
- 6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU
- 6.12 Multiple Processors and Matrix Multiply
- 6.13 Fallacies and Pitfalls
- 6.14 Concluding Remarks



## i7-960 vs. NVIDIA Tesla 280/480

	<sup>CPU</sup> Core i7-960	<sup>GPU</sup> GTX 280	<sup>GPU</sup> GTX 480	Ratio 280/17	Ratio 480/17
Number of processing elements (cores or SMs)	4	30	15	7.5	3.8
Clock frequency (GHz)	3.2	1.3	1.4	0.41	0.44
Die size	263	576	520	2.2	2.0
Technology	Intel 45 nm	TCMS 65 nm	TCMS 40 nm	1.6	1.0
Power (chip, not module)	130	130	167	1.0	1.3
Transistors	700 M	1400 M	3100 M	2.0	4.4
Memory bandwidth (GBytes/sec)	32	<sup>أربع</sup> 141	177	4.4	5.5
Single precision SIMD width	4	8	32	2.0	8.0
Double precision SIMD width	<sup>أحسن</sup> 2	<sup>أسوأ</sup> 1	16	0.5	8.0
Peak Single precision scalar FLOPS (GFLOP/sec)	26	117	63	4.6	2.5
Peak Single precision s SIMD FLOPS (GFLOP/Sec)	102	311 to 933	515 to 1344	3.0-9.1	6.6-13.1
(SP 1 add or multiply)	N.A.	(311)	(515)	(3.0)	(6.6)
(SP 1 instruction fused)	N.A.	(622)	(1344)	(6.1)	(13.1)
(face SP dual issue fused)	N.A.	(933)	N.A.	(9.1)	-
Peak double precision SIMD FLOPS (GFLOP/sec)	51	78	515	1.5	10.1





86.11 Real-World Benchmarking and Profiles: i7 vs. Tesla





# Rooflines

high performance جاسی



# Benchmarks

Kernel	Units	Core i7-960	GTX 280	GTX 280/ i7-960
SGEMM	GFLOP/sec	94	364	3.9
MC	Billion paths/sec	0.8	1.4	1.8
Conv	Million pixels/sec	1250	3500	2.8
FFT	GFLOP/sec	71.4	213	3.0
SAXPY	GBytes/sec	16.8	88.8	5.3
LBM	Million lookups/sec	85	426	5.0
Solv	Frames/sec	103	52	0.5
SpMV	GFLOP/sec	4.9	9.1	1.9
GJK	Frames/sec	67	1020	15.2
Sort	Million elements/sec	250	198	0.8
RC	Frames/sec	5	8.1	1.6
Search	Million queries/sec	50	90	1.8
Hist	Million pixels/sec	1517	2583	1.7
Bilat	Million pixels/sec	83	475	5.7



# Performance Summary

GPU (480) has 4.4 X the memory bandwidth

- Benefits memory bound kernels

GPU has 13.1 X the single precision throughput, 2.5 X the double precision throughput

- Benefits FP compute bound kernels

CPU cache prevents some kernels from becoming memory bound when they otherwise would on GPU

GPUs offer scatter-gather, which assists with kernels with strided data

Lack of synchronization and memory consistency support on GPU limits performance for some kernels

MK

Chapter 6 — Parallel Processors from Client to Cloud — 62

## Contents

6.1 Introduction

6.2 The Difficulty of Creating Parallel Programs

6.3 SISD, MIMD, SIMD, SPMD, and Vector

6.4 Hardware Multithreading

6.5 Shared Memory Multiprocessors

6.6 Introduction to Graphics Processing Units

6.7 Clusters and Message-Passing Multiprocessors

6.8 Introduction to Multiprocessor Network Topologies

6.10 Multiprocessor Benchmarks and Performance Models

6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU

\*6.12 Multiple Processors and Matrix Multiply *منه مطلوب*

6.13 Fallacies and Pitfalls

6.14 Concluding Remarks

MK

Chapter 6 — Parallel Processors from Client to Cloud — 63

Scanned with CamScanner



# Fallacies

- Amdahl's Law doesn't apply to parallel computers
  - Since we can achieve linear speedup
  - But only on applications with weak scaling
- Peak performance tracks observed performance
  - Marketers like this approach!
  - But compare Xeon with others in example
  - Need to be aware of bottlenecks

6.13 Fallacies and Pitfalls

6.1  
6.2  
6.3  
6.4  
6.5  
6.6  
6.7  
6.8  
6.1  
6.1  
6.1  
6.1  
6.1



# Pitfalls

- Not developing the software to take account of a multiprocessor architecture
  - Example: using a single lock for a shared composite resource
    - Serializes accesses, even if they could be done in parallel
    - Use finer-granularity locking

C  
P  
S  
P  
S

# Contents

- 6.1 Introduction
- 6.2 The Difficulty of Creating Parallel Programs
- 6.3 SISD, MIMD, SIMD, SPMD, and Vector
- 6.4 Hardware Multithreading
- 6.5 Shared Memory Multiprocessors
- 6.6 Introduction to Graphics Processing Units
- 6.7 Clusters and Message-Passing Multiprocessors
- 6.8 Introduction to Multiprocessor Network Topologies
- 6.10 Multiprocessor Benchmarks and Performance Models
- 6.11 Benchmarking Intel Core i7 960 and NVIDIA Tesla GPU
- 6.12 Multiple Processors and Matrix Multiply
- 6.13 Fallacies and Pitfalls
- 6.14 Concluding Remarks



## Concluding Remarks


- Goal: higher performance by using multiple processors
- Difficulties
  - Developing parallel software
  - Devising appropriate architectures
- SaaS importance is growing and clusters are a good match
- Performance per dollar and performance per Joule drive both mobile and WSC

6.14 Concluding Remarks



# Concluding Remarks (con't)

- SIMD and vector operations match multimedia applications and are easy to program
- Adding 2 cores/chip every 2 years.
- Doubling SIMD operations every 4 years.



## Slide 53 of Roofline Model

DAXPY:

for (i=0; i&lt;N; i+=1)

 $Y[i] += S * X[i] + Y[i];$ 

↓            ↑    ↑    ↑    ↑

∴ Arithmetic Intensity =  $\frac{2 \text{ Flops}}{(2 \text{ loads} + 1 \text{ store})}$

$$= \frac{2}{(3 * 8)} = 0.0833 \text{ FLOP/byte}$$

8 byte loads

(Memory intensive ليل فوالب)