INF2270 — Spring 2010

Philipp Häfliger

Lecture 9: Course Summary/Repetition (2/2)

UNIVERSITETET I OSLO

content

Lecture Summary and Brief Repetition: von Neumann Architecture

- Overview
- CU
- EU, ALU
- Memory
- I/O

Optimizing Hardware Performance

- Memory Hierarchy
- Pipelining
- Superscalar CPU

content

Lecture Summary and Brief Repetition: von Neumann Architecture

- Overview
- CU
- EU, ALU
- Memory
- I/O

Optimizing Hardware Performance

- Memory Hierarchy
- Pipelining
- Superscalar CPU

content

Von Neumann Overview

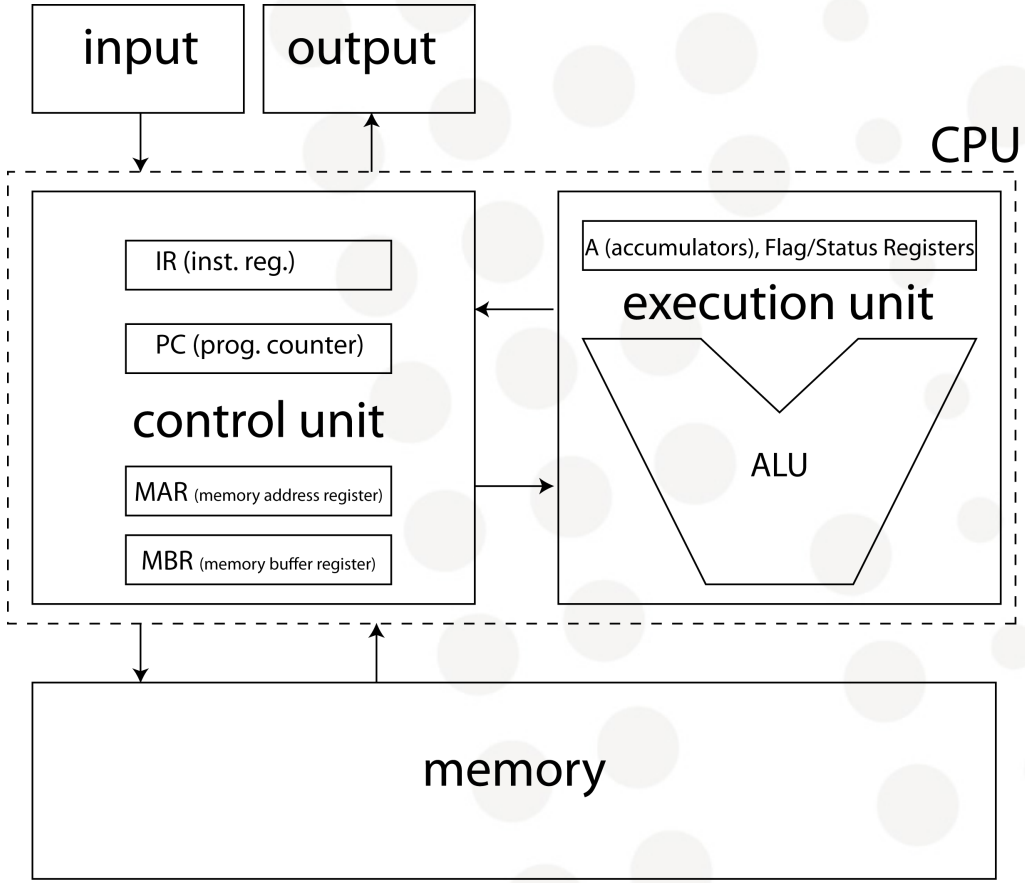

content

Lecture Summary and Brief Repetition: von Neumann Architecture

- Overview
- CU
- EU, ALU
- Memory
- I/O

Optimizing Hardware Performance

- Memory Hierarchy
- Pipelining
- Superscalar CPU

CU Registers

PC: (program counter, also called instruction pointer (IP)) the register holding the memory address of the next machine code instruction.

IR: (instruction register) the register holding the machine code of the instruction that is executed.

MAR: (memory address register) half of the CPU-memory interface, holding the memory address to be read or written to.

MBR: (memory buffer register) the other half of the CPU-memory interface, holding the data just read from or to be written to the memory. Typically connectable as input to the ALU.

CU as FSM

The CU is a FSM with its registers defining its state and the memory provideds its input, i.e. the program and the data. However, it is not usually decribed with a complete characteristic table and often designed as several interlinked FSMs.

Register Transfer Language (RTL)

expression	meaning
X	register X or unit X
[X]	the content of X
←	replace/insert or execute code
M()	memory M
[M([X])]	memory content at address [X]

An example:

[IR] ← [MBR] transfer the content of the MBR to the IR

Microarchitecture

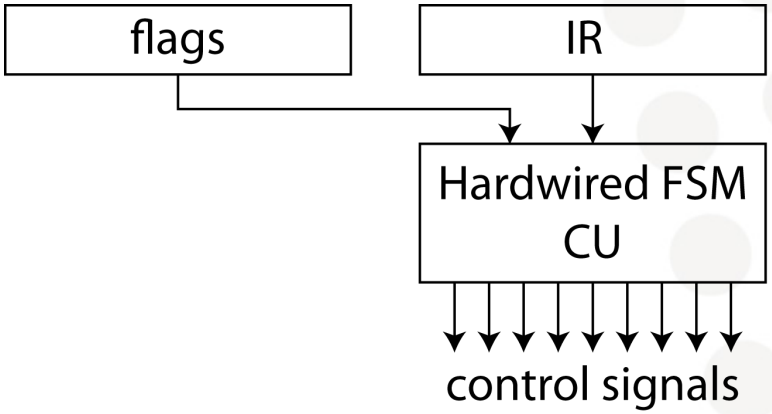

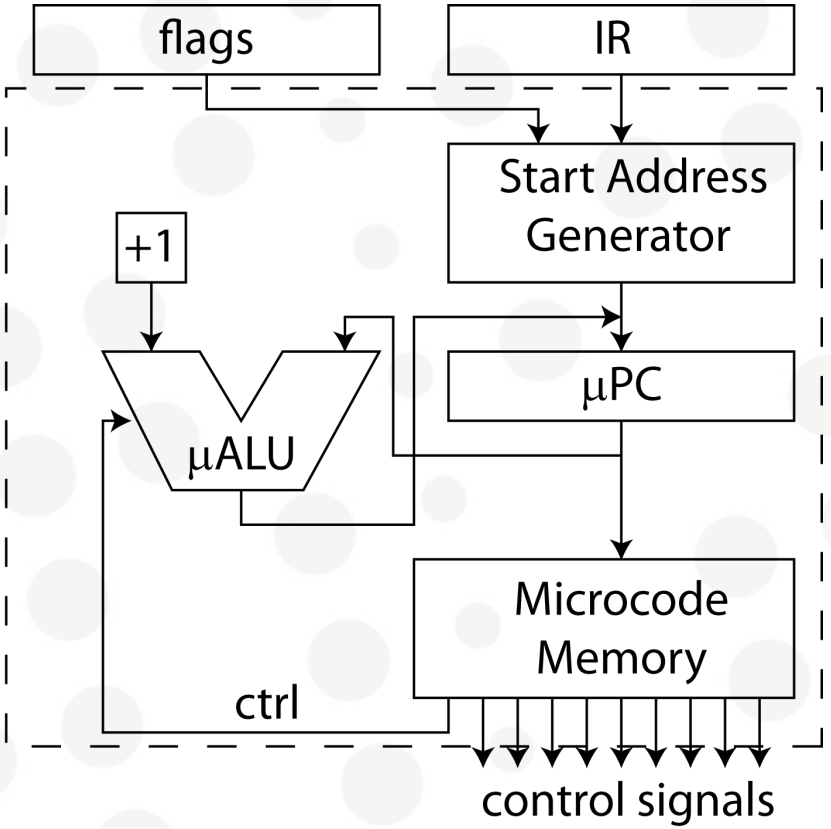

content

Lecture Summary and Brief Repetition: von Neumann Architecture

- Overview
- CU
- EU, ALU
- Memory
- I/O

Optimizing Hardware Performance

- Memory Hierarchy
- Pipelining
- Superscalar CPU

ALU Registers

accumulator: a dedicated register that stores one operand and the result of the ALU. Several accumulators (or general purpose registers in the CPU) allow for storing of intermediate results, avoiding (costly) memory accesses.

flag/status register: a register where each single bit stands for a specific property of the result from (or the input to) the ALU, like carry in/out, equal to zero, even/uneven ...

1-bit ALU

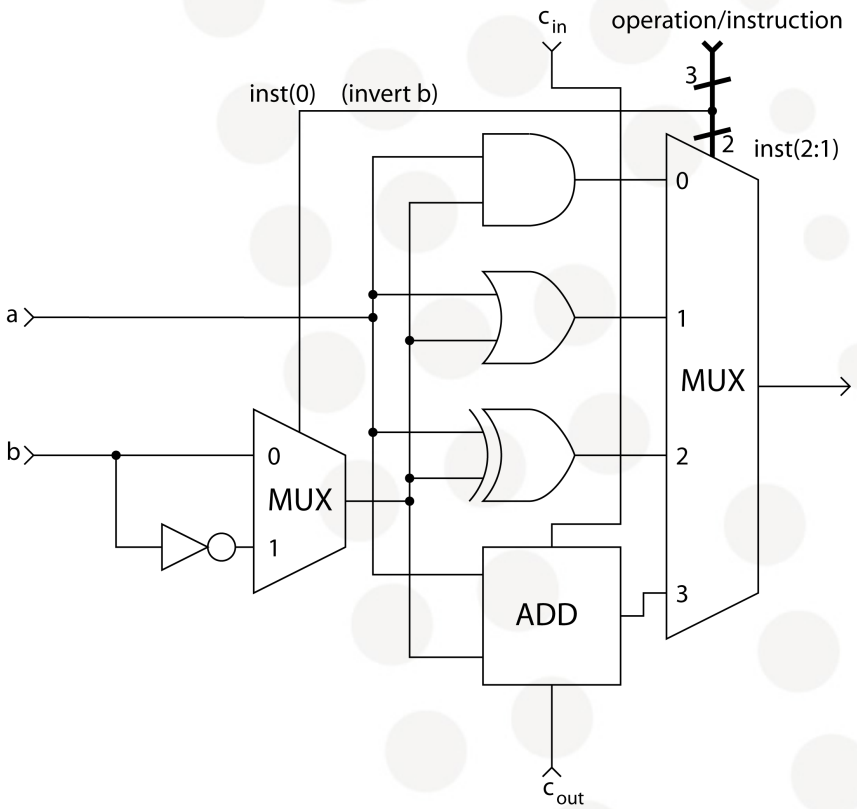

content

Lecture Summary and Brief Repetition: von Neumann Architecture

- Overview
- CU
- EU, ALU
- Memory
- I/O

Optimizing Hardware Performance

- Memory Hierarchy
- Pipelining
- Superscalar CPU

random access memory (RAM)

The name refers to a memory where one can address individual words at any time, as opposed to memory content that can only be accessed in sequence (e.g. read out from a camera frame, or a hard drive that has to rotate to the right place first).

SRAM

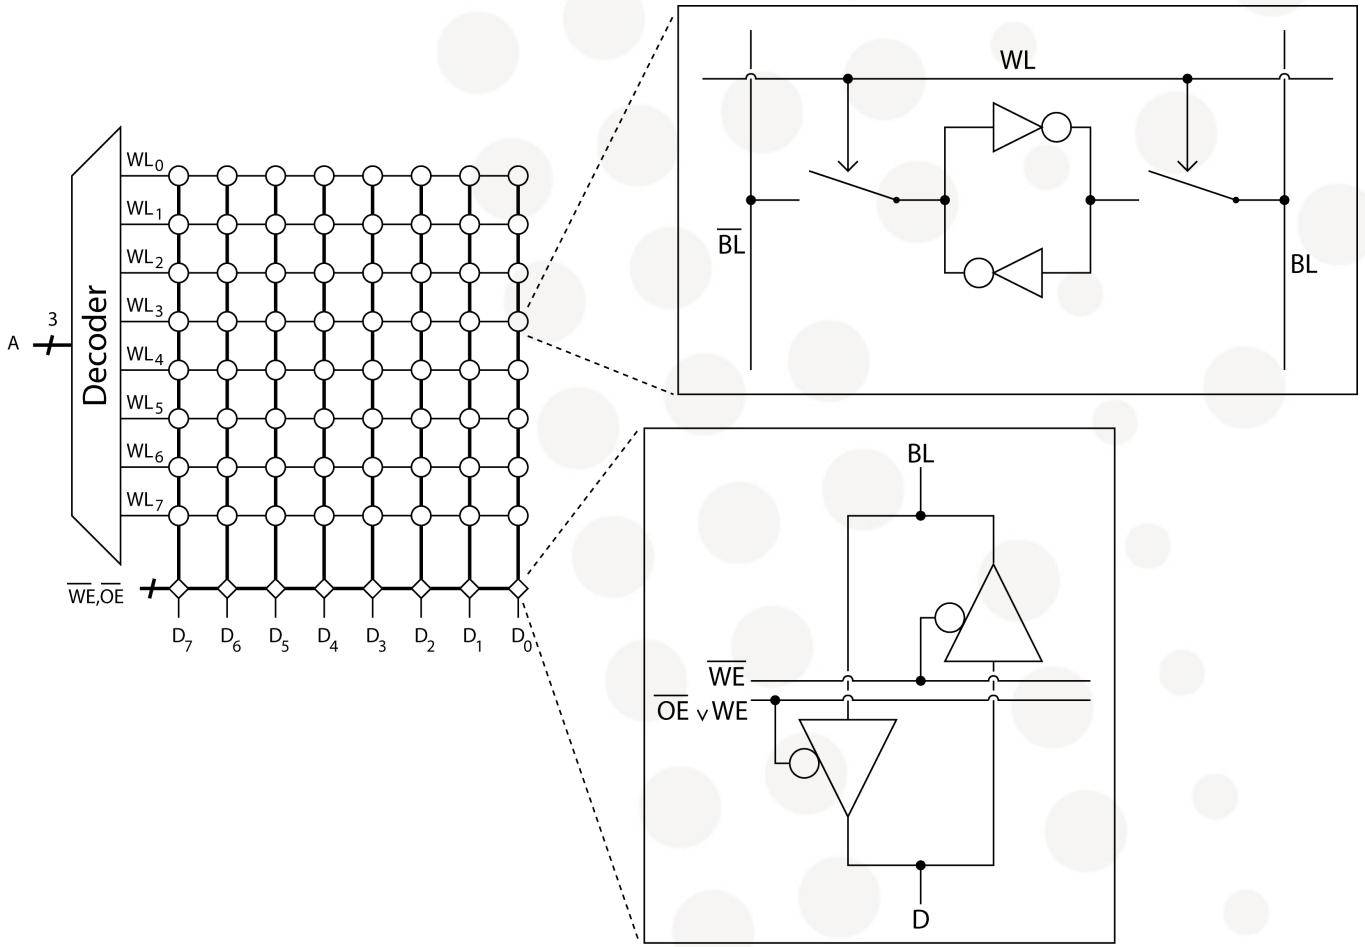

DRAM

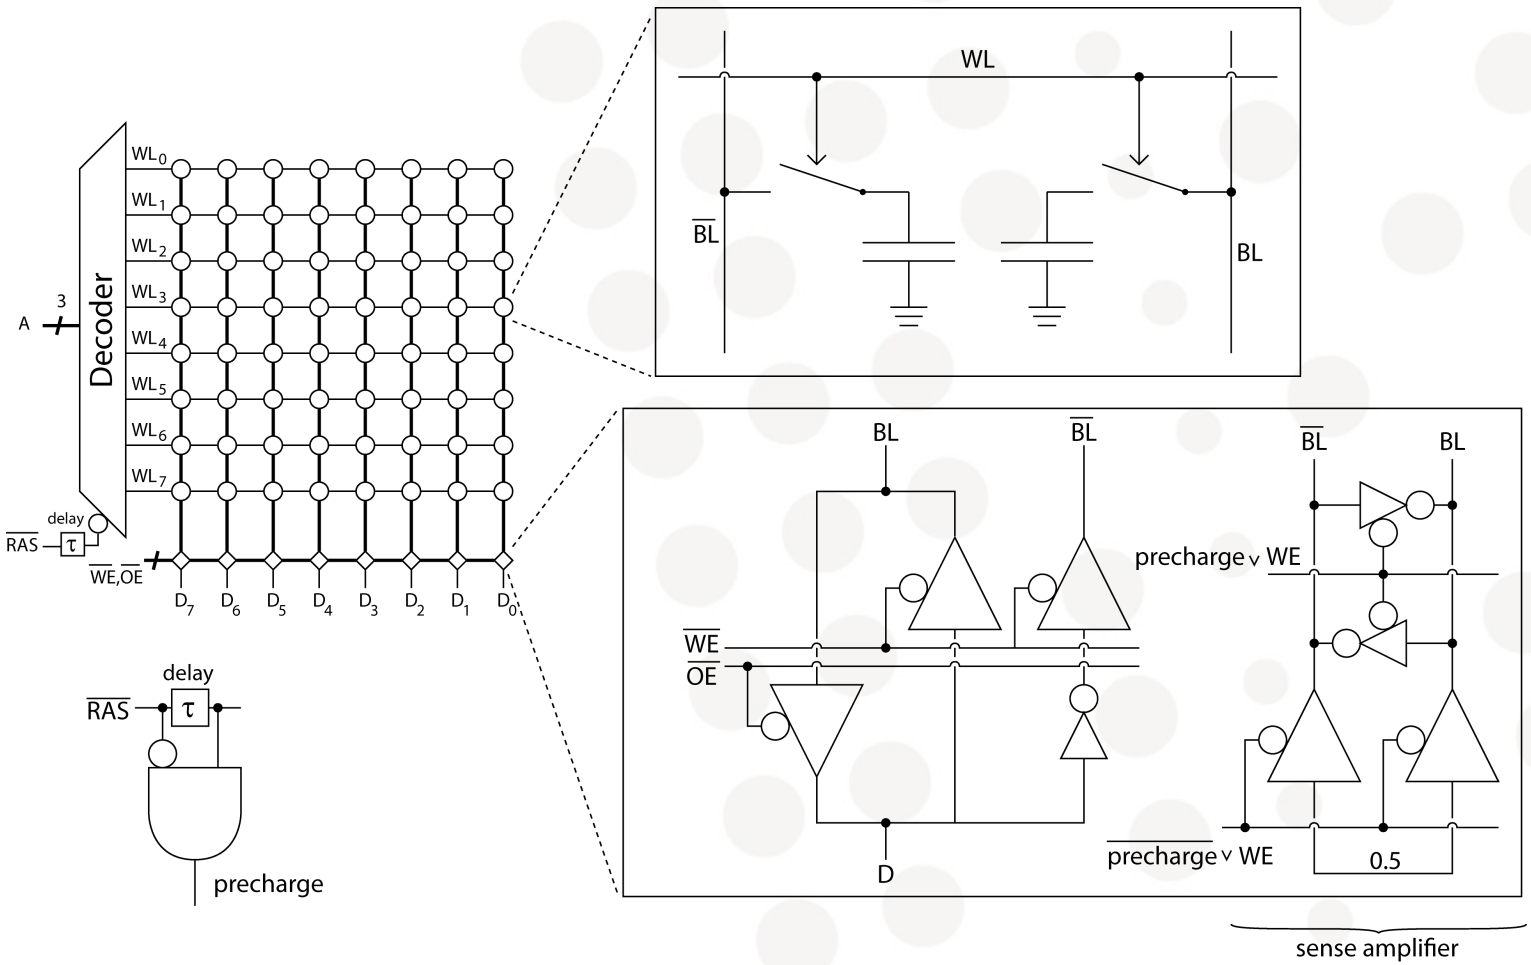
WL
BL
BL
Decoder
A
3
WL0
WL1
WL2
WL3
WL4
WL5
WL6
WL7
delay
RAS
WE,OE
D7
D6
D5
D4
D3
D2
D1
D0
BL
BL
WE
OE
D
precharge ∨ WE
BL
BL
precharge ∨ WE
0.5
sense amplifier
delay
RAS
precharge

SRAM vs. DRAM

	SRAM	DRAM
access speed	+	-
memory density	-	+
no refresh needed	+	-
simple internal control	+	-
price per bit	-	+

content

Lecture Summary and Brief Repetition: von Neumann Architecture

- Overview
- CU
- EU, ALU
- Memory
- I/O

Optimizing Hardware Performance

- Memory Hierarchy
- Pipelining
- Superscalar CPU

Interrupt Controlled I/O

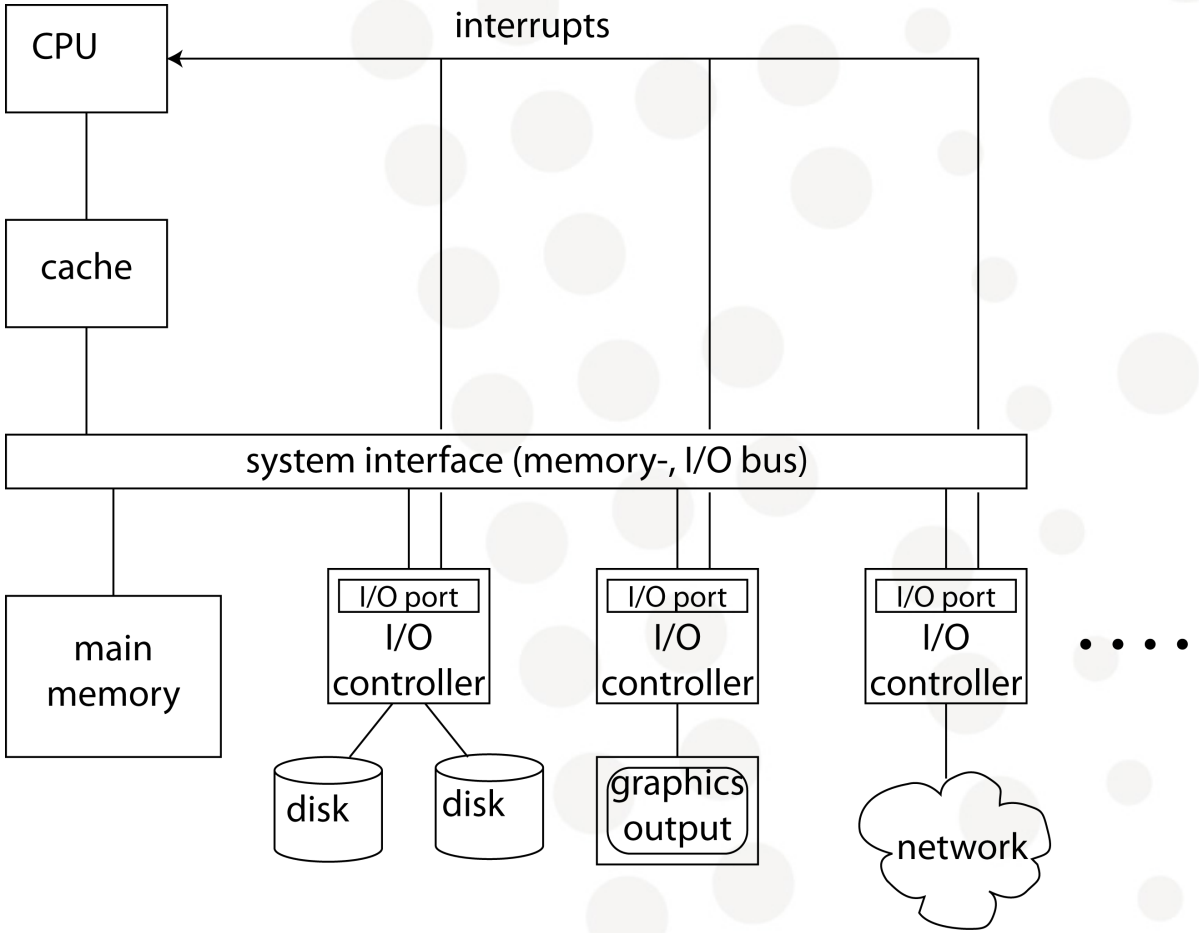

content

Lecture Summary and Brief Repetition: von Neumann Architecture
- Overview
- CU
- EU, ALU
- Memory
- I/O

Optimizing Hardware Performance
- Memory Hierarchy
- Pipelining
- Superscalar CPU

content

Memory Hierarchy Motivation

A hierarchical memory structure is employed to ameliorate the von Neumann bottle neck by trying to keep frequently accessed data in the hierarchy's fastest level, trading in speed vs. cost and speed vs. size.

Memory Hierarchy Structure

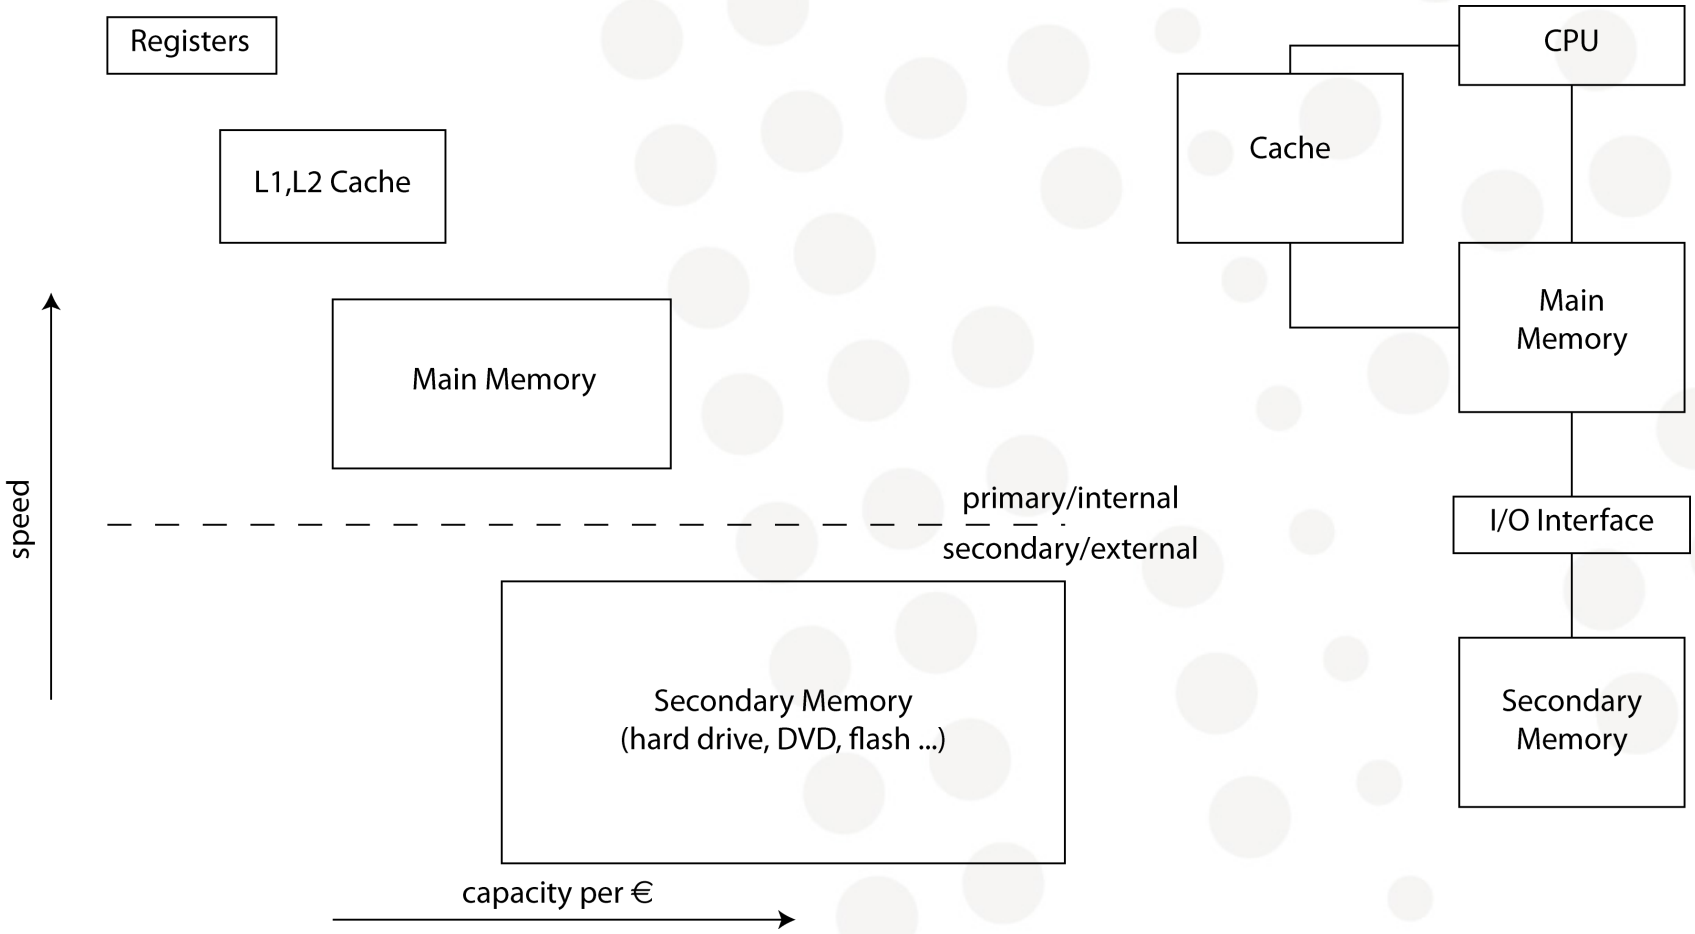

Access Times and Size

registers	~ 1ns	~ 100B
L1 (on CPU) cache	$\sim \geq$ 1ns	~ 10kB
L2,L3 (off CPU) cache	~ 2ns-10ns	~ 1MB
main memory (DRAM)	~ 20ns-100ns	~ 1GB
SSD/flash	~ 100ns-1μs	~ 10-100GB
hard disc	~ 1ms	~ 0.1-1TB

content

Lecture Summary and Brief Repetition: von Neumann Architecture

- Overview
- CU
- EU, ALU
- Memory
- I/O

Optimizing Hardware Performance

- Memory Hierarchy
- Pipelining
- Superscalar CPU

Pipelined Instruction Execution

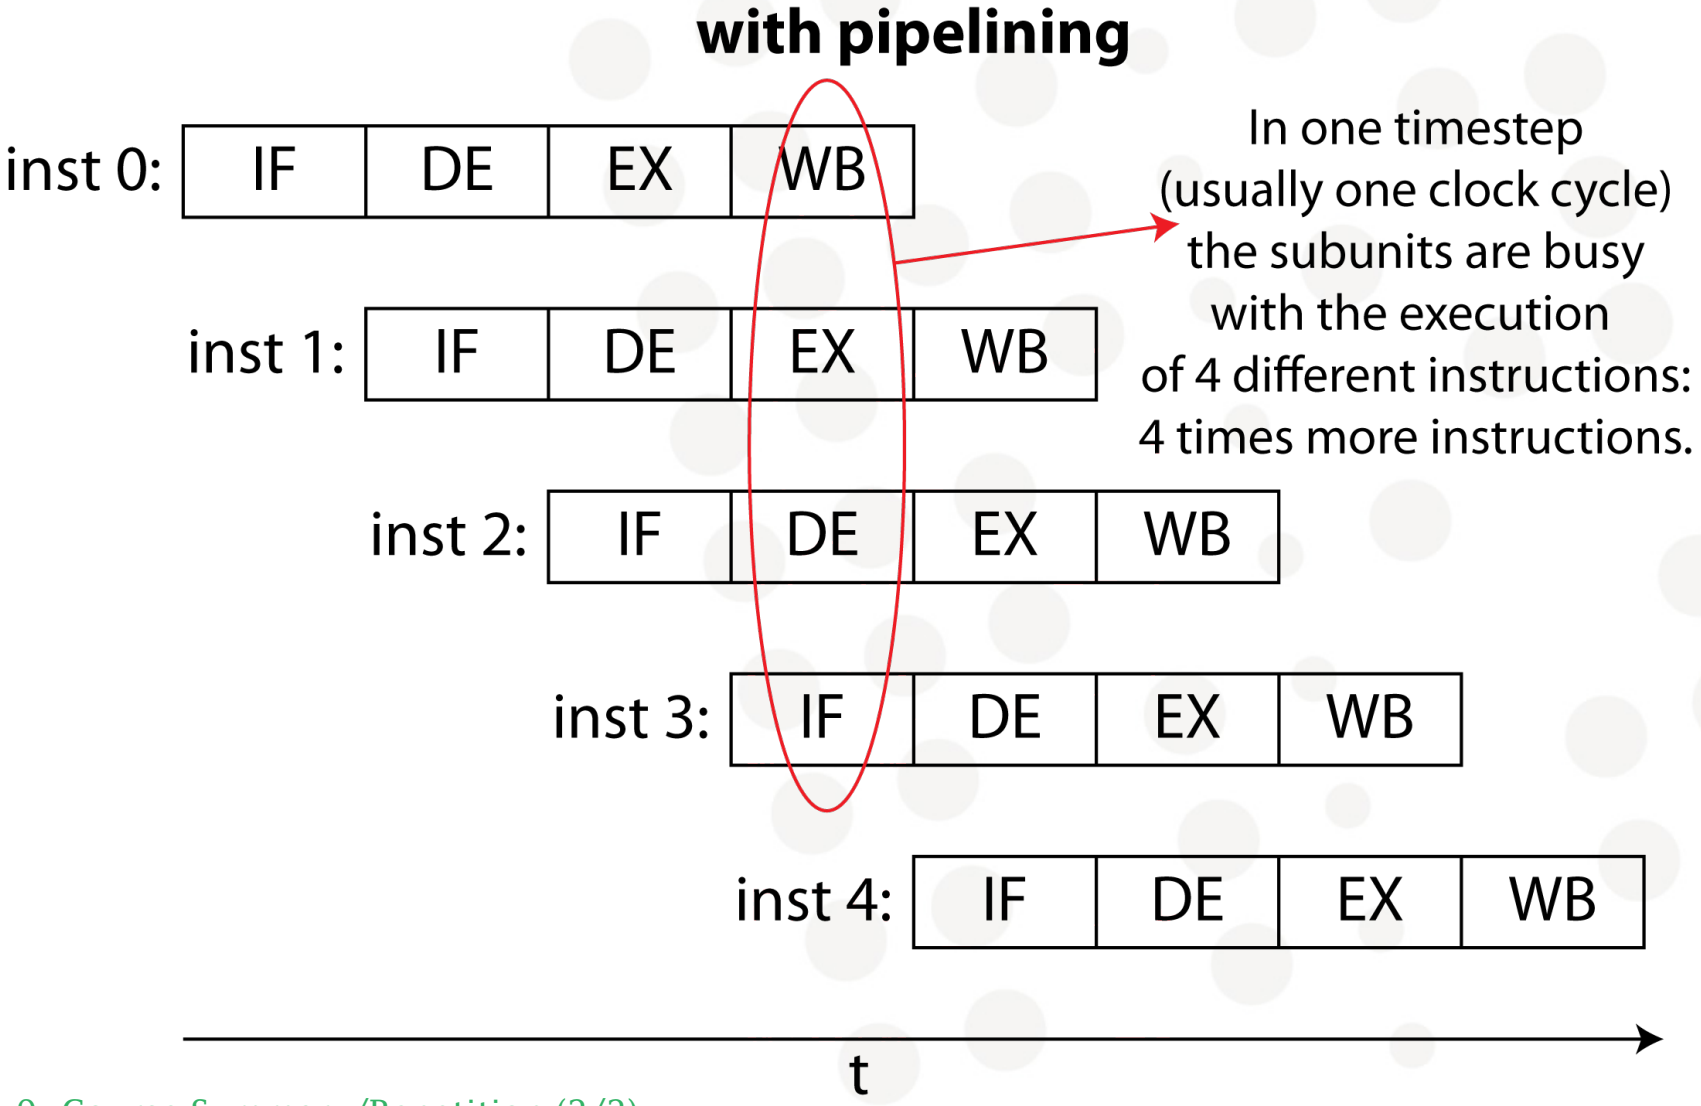

content

Lecture Summary and Brief Repetition: von Neumann Architecture

- Overview
- CU
- EU, ALU
- Memory
- I/O

Optimizing Hardware Performance

- Memory Hierarchy
- Pipelining
- Superscalar CPU

Superscalar Execution