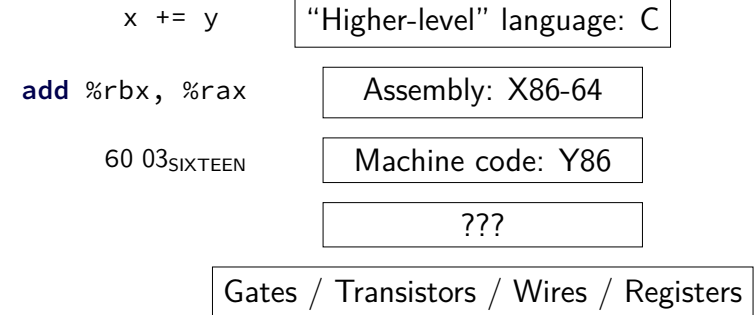


## CS3330 — overview

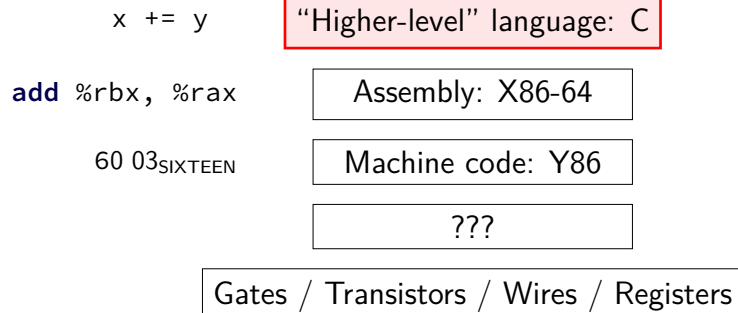
1

## layers of abstraction



2

## layers of abstraction



3

## Why C?

*almost* a subset of C++  
notably removes classes, new/delete, iostreams  
other changes, too, so C code often not valid C++ code

**direct correspondence** to assembly

4

## Why C?

*almost* a subset of C++

notably removes classes, new/delete, iostreams  
other changes, too, so C code often not valid C++ code

**direct correspondence** to assembly

Should help you understand machine!  
Manual translation to assembly

4

## Why C?

*almost* a subset of C++

notably removes classes, new/delete, iostreams  
other changes, too, so C code often not valid C++ code

**direct correspondence** to assembly

But “clever” (optimizing) compiler  
might be confusingly indirect instead

4

## homework: C environment

get a C compiler

options:

lab accounts + SSH  
Linux (native or VM)  
online IDE (e.g. Cloud9, Koding)

5

## assignment compatibility

supported platform: lab machines

many use laptops

trouble? we'll say to use lab machines

most assignments: C and Unix-like environment

also: tool written in Rust — but we'll provide binaries  
previously written in D + needed D compiler

6

## layers of abstraction

`x += y`

“Higher-level” language: C

`add %rbx, %rax`

Assembly: X86-64

`60 03SIXTEEN`

Machine code: Y86

???

Gates / Transistors / Wires / Registers

7

## X86-64 assembly

in theory, you know this (CS 2150)

in reality, ...

8

## 32 versus 64-bit note

some of you may have learned 32-bit in 2150  
(the course has changed)

differences mostly: more, bigger registers

9

## layers of abstraction

`x += y`

“Higher-level” language: C

`add %rbx, %rax`

Assembly: X86-64

`60 03SIXTEEN`

Machine code: Y86

???

Gates / Transistors / Wires / Registers

10

## Y86-64??

Y86: our textbook's X86-64 subset

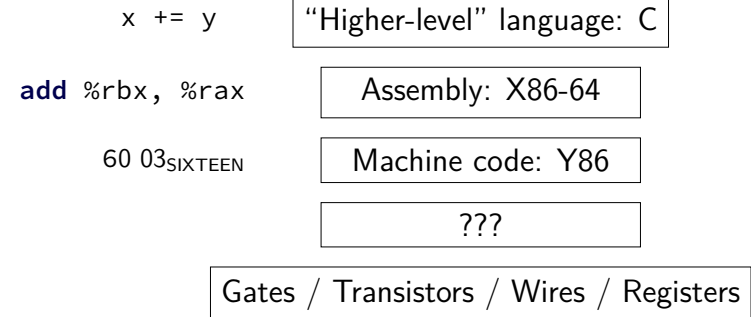
much simpler than real X86-64 encoding  
(which we will not cover)

not as simple as 2150's ICBM

- variable-length encoding
- mostly full register set
- full conditional jumps
- stack-manipulation instructions

11

## layers of abstraction



12

## hardware

most of the semester

13

## goals/other topics

understand how hardware works for...

program performance

what compilers are/do

weird program behaviors (segfaults, etc.)

14

## goals/other topics

understand how hardware works for...

### program performance

what compilers are/do

weird program behaviors (segfaults, etc.)

15

## program performance

naive model:

one instruction = one time unit

number of instructions matters, but ...

16

## program performance: issues

### parallelism

fast hardware is parallel  
needs multiple things to do

### caching

accessing things recently accessed is faster  
need reuse of data/code

(more in other classes: **algorithmic** efficiency)

17

## goals/other topics

understand how hardware works for...

program performance

### what compilers are/do

weird program behaviors (segfaults, etc.)

18

## what compilers are/do

understanding weird compiler/linker rrors

if you want to make compilers

debugging applications

19

## goals/other topics

understand how hardware works for...

program performance

what compilers are/do

weird program behaviors (segfaults, etc.)

20

## weird program behaviors

what is a segmentation fault really?

how does the operating system interact with programs?

if you want to handle them — writing OSs

21

## interlude: powers of two

	...		...
$2^0$	1	$2^{11}$	2 048
$2^1$	2	$2^{12}$	4 096
$2^2$	4	$2^{13}$	8 192
$2^3$	8	$2^{14}$	16 384
$2^4$	16	$2^{15}$	32 768
$2^5$	32	$2^{16}$	65 536
$2^6$	64	...	...
$2^7$	128	$2^{20}$	1 048 576 <b>M</b> (or Mi)
$2^8$	256	...	...
$2^9$	512	$2^{30}$	1 073 741 824 <b>G</b> (or Gi)
$2^{10}$	1 024 <b>K</b> (or Ki)	$2^{31}$	2 147 483 648
		$2^{32}$	4 294 967 296
		...	...

22

## powers of two: forward

$$2^{35}$$

$$2^{21}$$

$$2^9$$

$$2^{14}$$

## powers of two: forward

$$2^{35} = 2^5 \cdot 2^{30} = 32G \text{ (30 = G)}$$

$$2^{21}$$

$$2^9$$

$$2^{14}$$

## powers of two: forward

$$2^{35} = 2^5 \cdot 2^{30} = 32G \text{ (30 = G)}$$

$$2^{21}$$

$$2^9$$

$$2^{14}$$

## powers of two: forward

$$2^{35} = 2^5 \cdot 2^{30} = 32G \text{ (30 = G)}$$

$$2^{21} = 2^1 \cdot 2^{20} = 2M \text{ (20 = M)}$$

$$2^9$$

$$2^{14}$$

## powers of two: forward

$$2^{35} = 2^5 \cdot 2^{30} = 32G \text{ (30 = G)}$$

$$2^{21} = 2^1 \cdot 2^{20} = 2M \text{ (20 = M)}$$

$$2^9 = 512$$

$$2^{14}$$

23

## powers of two: forward

$$2^{35} = 2^5 \cdot 2^{30} = 32G \text{ (30 = G)}$$

$$2^{21} = 2^1 \cdot 2^{20} = 2M \text{ (20 = M)}$$

$$2^9 = 512$$

$$2^{14} = 2^4 \cdot 2^{10} = 16K$$

23

## powers of two: backward

16G

128K

4M

256T

24

## powers of two: backward

$$16G = 16 \cdot 2^{30} = 2^{30+4} = 2^{34}$$

128K

4M

256T

24



## powers of two: backward

$$16\text{G} = 16 \cdot 2^{30} = 2^{30+4} = 2^{34}$$

$$128\text{K} = 128 \cdot 2^{10} = 2^{10+7} = 2^{17}$$

4M

256T

24

## powers of two: backward

$$16\text{G} = 16 \cdot 2^{30} = 2^{30+4} = 2^{34}$$

$$128\text{K} = 128 \cdot 2^{10} = 2^{10+7} = 2^{17}$$

$$4\text{M} = 4 \cdot 2^{20} = 2^{20+2} = 2^{22}$$

$$256\text{T} = 256 \cdot 2^{40} = 2^{40+8} = 2^{48}$$

24

## rest of today/tomorrow

brief preview of circuits, CPUs

assembly and linking

selected things about C

25

## layers of abstraction

`x += y`

“Higher-level” language: C

`add %rbx, %rax`

Assembly: X86-64

`60 03SIXTEEN`

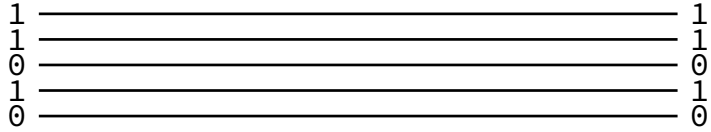
Machine code: Y86

???

Gates / Transistors / Wires / Registers

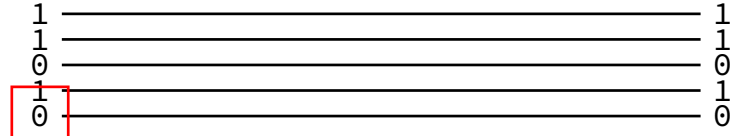
26

## Circuits: Wires



27

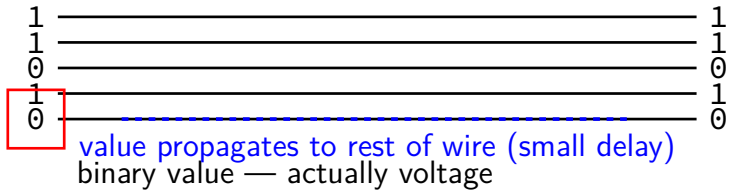
## Circuits: Wires



binary value — actually voltage

27

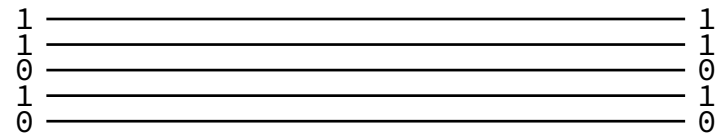
## Circuits: Wires



value propagates to rest of wire (small delay)  
binary value — actually voltage

27

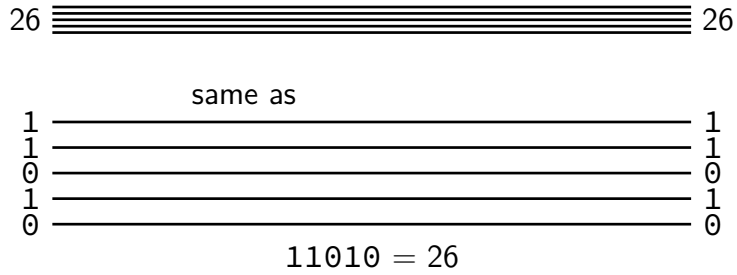
## Circuits: Wire Bundles



11010 = 26

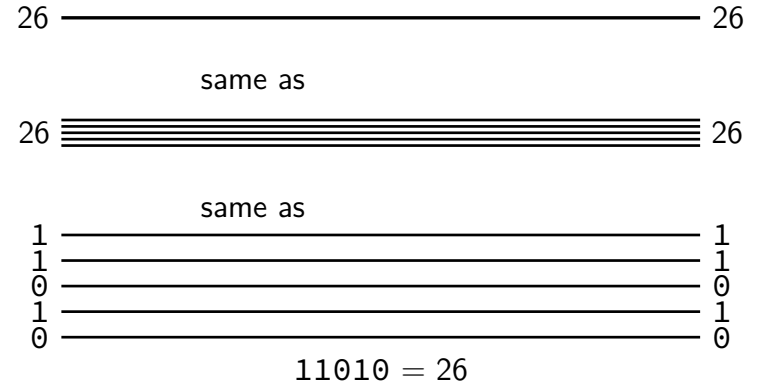
28

## Circuits: Wire Bundles



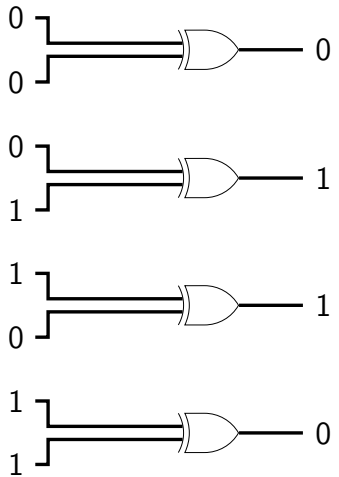
28

## Circuits: Wire Bundles



28

## Circuits: Gates

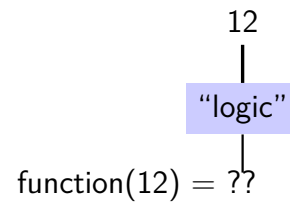


29

## Circuits: Logic

want to do calculations?

generalize gates:



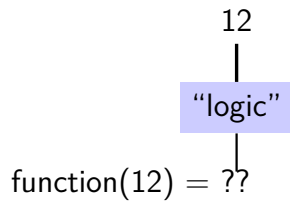
30

## Circuits: Logic

want to do calculations?

generalize gates:

output wires contain result of function on input  
changes as input changes (with delay)



30

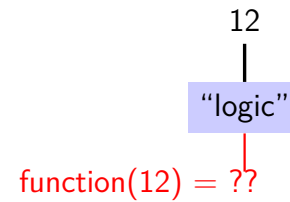
## Circuits: Logic

want to do calculations?

generalize gates:

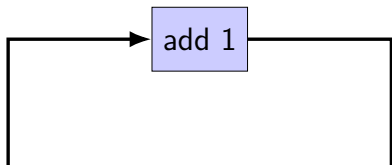
output wires contain result of function on input  
changes as input changes (with delay)

need not be same width as output



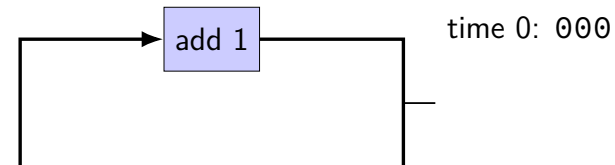
30

## example: (broken) counter circuit



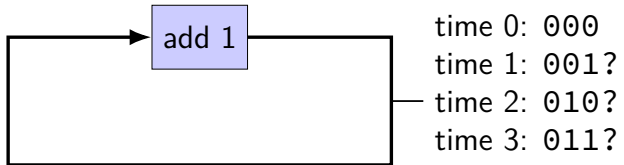
31

## example: (broken) counter circuit



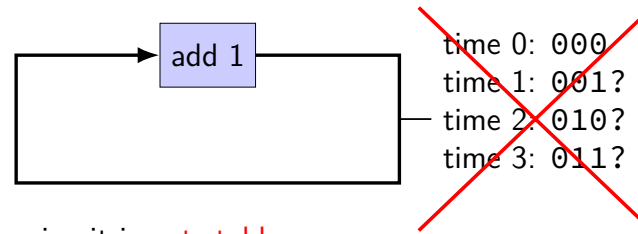
31

## example: (broken) counter circuit



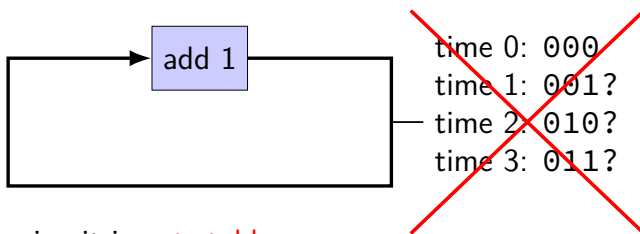
31

## example: (broken) counter circuit



31

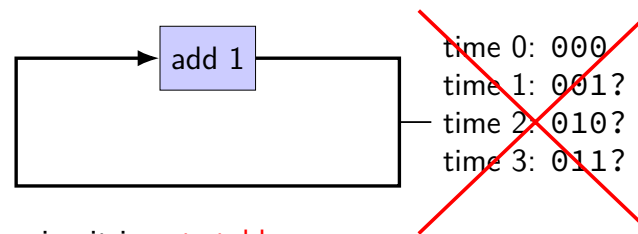
## example: (broken) counter circuit



circuit is **not stable**  
**transient values** during changes  
can't transition from 001 to 010  
without 011 or 000

31

## example: (broken) counter circuit



circuit is **not stable**  
**transient values** during changes  
can't transition from 001 to 010  
without 011 or 000  
halfway voltages — hard to predict behavior

31

## circuits: state

logic performs calculations all the time

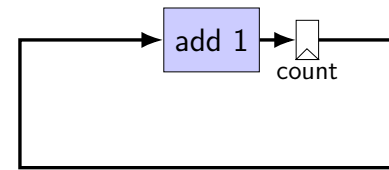
never stores values!

need **extra elements** to store values  
registers, memory

more on these later in the course

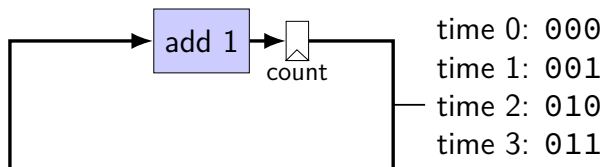
32

## example: counter circuit (corrected)



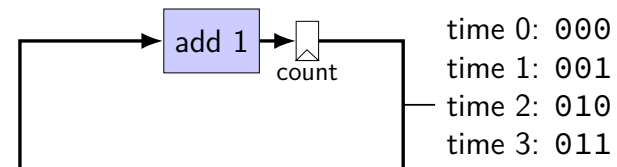
33

## example: counter circuit (corrected)



33

## example: counter circuit (corrected)

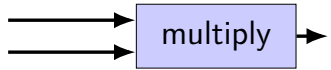
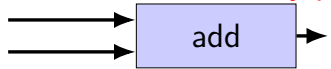


add **register** to store current count  
updates based on "clock signal" (not shown)  
avoids intermediate updates  
much more on this later in the semester

33

## parallel hardware

hardware is **inherently parallel**



most hardware design: making it **sequential**



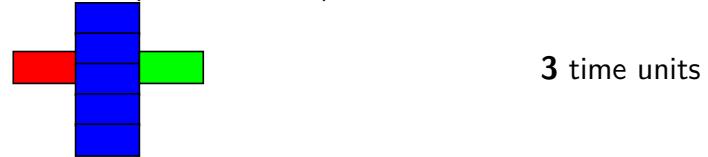
34

## parallelism and bottlenecks

Serial:



Parallel (blue 5x faster):



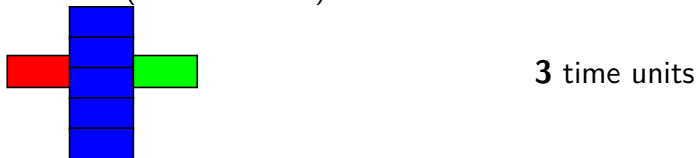
35

## parallelism and bottlenecks

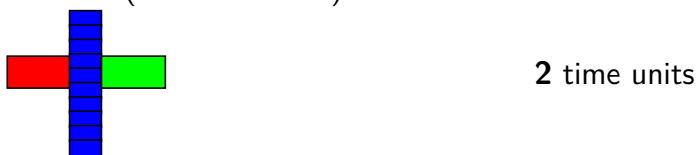
Serial:



Parallel (blue 5x faster):



Parallel (blue 10x faster):



35

## Amdahl's Law

formula in textbook

benefits of speedup limited by **non-sped-up parts**

parallelism:

anything not parallelized will be significant

or in math:

$$\text{time} = \text{serial part} + \text{parallel part} \div \text{parallelism}$$

36

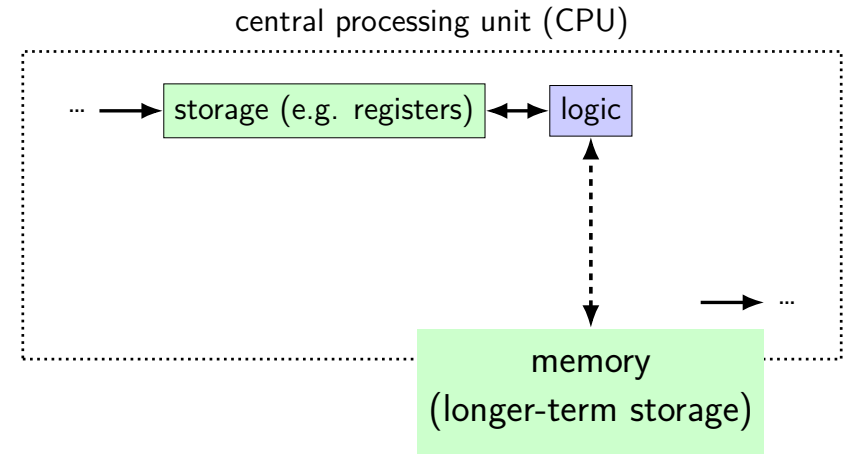
## not just parallelism

time = serial part + parallel part  $\div$  parallelism

time = unoptimized part + optimized part  $\div$  speedup

37

## constructing a computer



38

## layers of abstraction

`x += y`

“Higher-level” language: C

`add %rbx, %rax`

Assembly: X86-64

`60 03SIXTEEN`

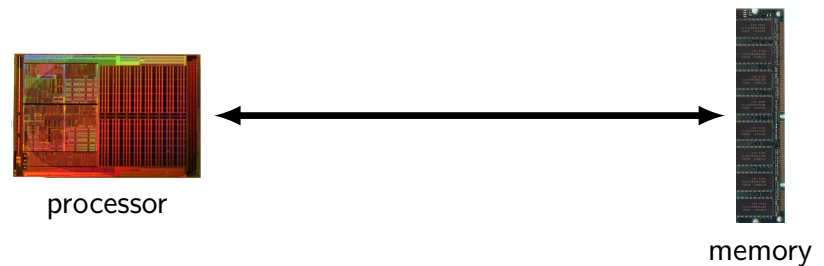
Machine code: Y86

???

Gates / Transistors / Wires / Registers

39

## processors and memory

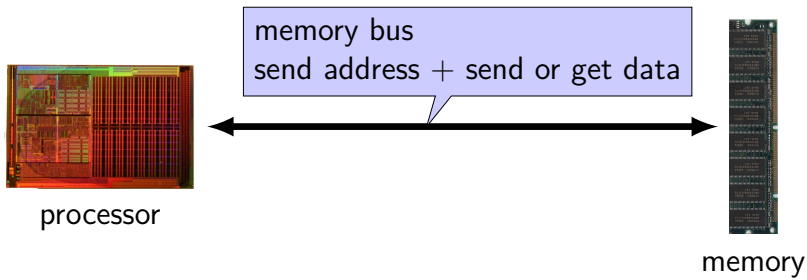


Images:  
Single core Opteron 8xx die: Dg2fer at the German language Wikipedia, via Wikimedia Commons  
SDRAM by Arnaud 25, via Wikimedia Commons

40



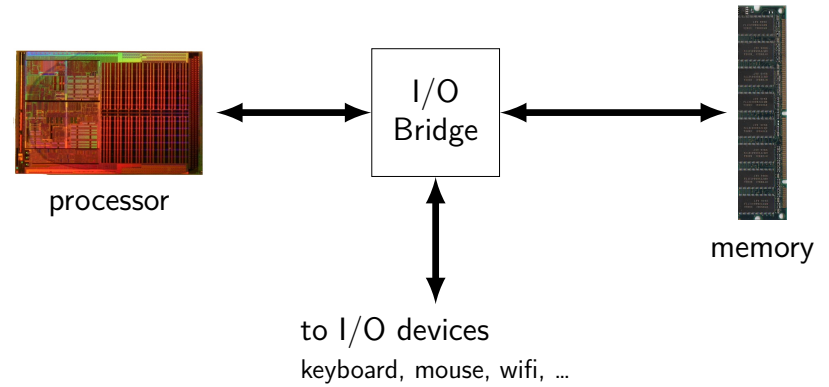
## processors and memory



Images:  
Single core Opteron 8xx die: Dg2fer at the German language Wikipedia, via Wikimedia Commons  
SDRAM by Arnaud 25, via Wikimedia Commons

40

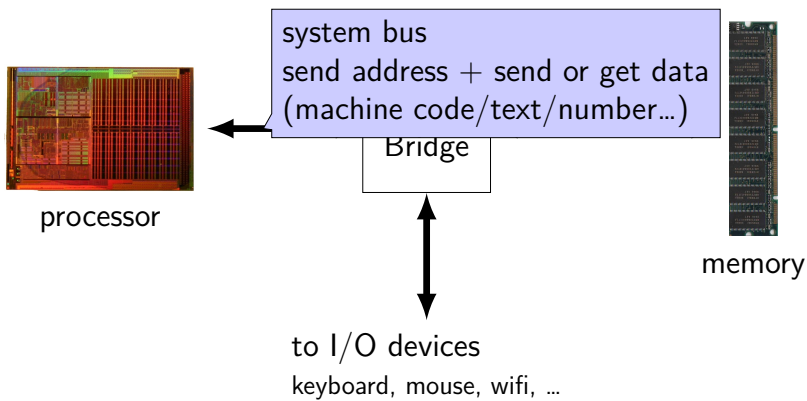
## processors and memory



Images:  
Single core Opteron 8xx die: Dg2fer at the German language Wikipedia, via Wikimedia Commons  
SDRAM by Arnaud 25, via Wikimedia Commons

40

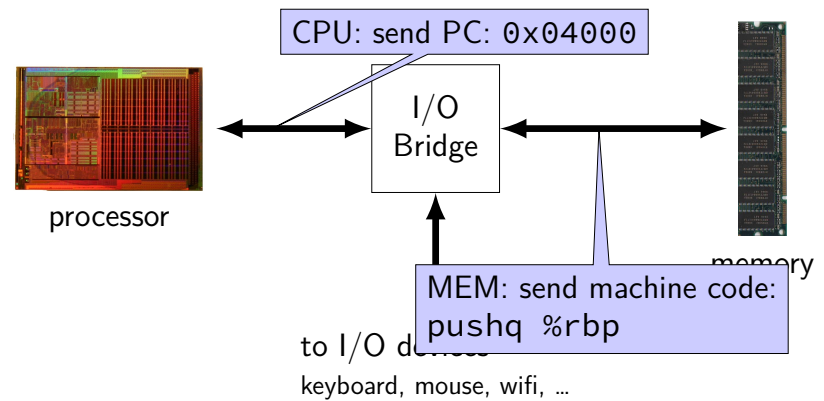
## processors and memory



Images:  
Single core Opteron 8xx die: Dg2fer at the German language Wikipedia, via Wikimedia Commons  
SDRAM by Arnaud 25, via Wikimedia Commons

40

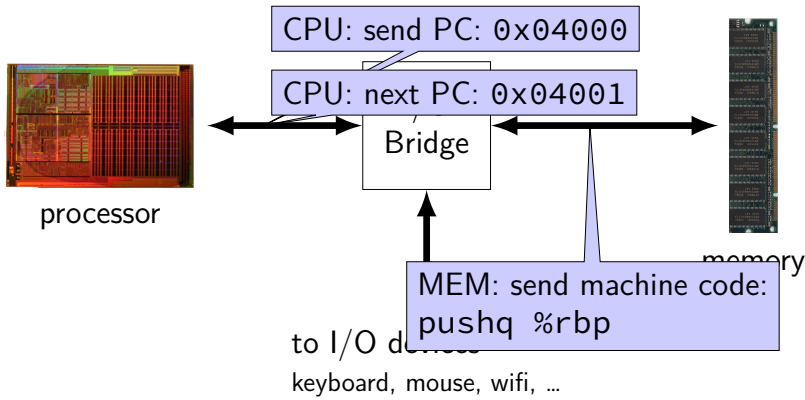
## processors and memory



Images:  
Single core Opteron 8xx die: Dg2fer at the German language Wikipedia, via Wikimedia Commons  
SDRAM by Arnaud 25, via Wikimedia Commons

40

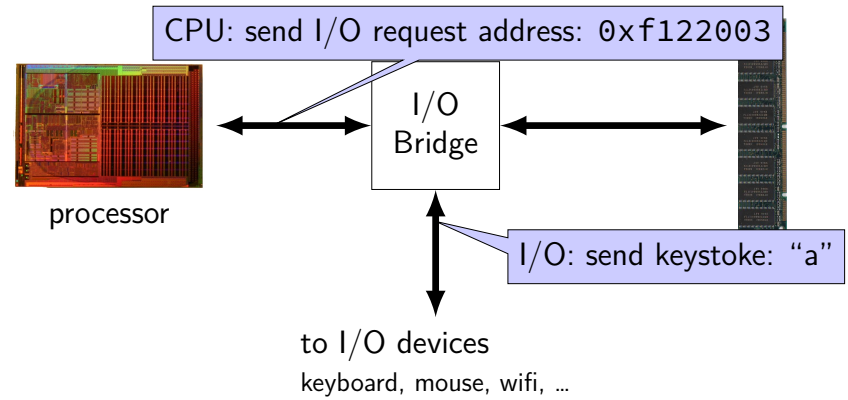
## processors and memory



Images:  
Single core Opteron 8xx die: Dg2fer at the German language Wikipedia, via Wikimedia Commons  
SDRAM by Arnaud 25, via Wikimedia Commons

40

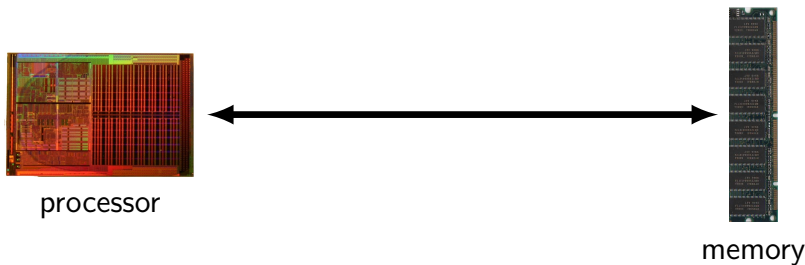
## processors and memory



Images:  
Single core Opteron 8xx die: Dg2fer at the German language Wikipedia, via Wikimedia Commons  
SDRAM by Arnaud 25, via Wikimedia Commons

40

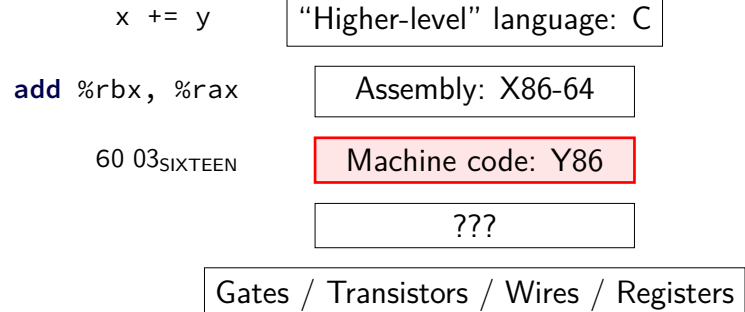
## processors and memory



Images:  
Single core Opteron 8xx die: Dg2fer at the German language Wikipedia, via Wikimedia Commons  
SDRAM by Arnaud 25, via Wikimedia Commons

41

## layers of abstraction



42

## memory

address	value
0xFFFFFFFF	0x14
0xFFFFFFF0	0x45
0xFFFFFFF4	0xDE
...	...
0x00042006	0x06
0x00042005	0x05
0x00042004	0x04
0x00042003	0x03
0x00042002	0x02
0x00042001	0x01
0x00042000	0x00
0x00041FFF	0x03
0x00041FFE	0x60
...	...
0x00000002	0xFE
0x00000001	0xE0
0x00000000	0xA0

43

## memory

address	value
0xFFFFFFFF	0x14
0xFFFFFFF0	0x45
0xFFFFFFF4	0xDE
...	...
0x00042006	0x06
0x00042005	0x05
0x00042004	0x04
0x00042003	0x03
0x00042002	0x02
0x00042001	0x01
0x00042000	0x00
0x00041FFF	0x03
0x00041FFE	0x60
...	...
0x00000002	0xFE
0x00000001	0xE0
0x00000000	0xA0

array of bytes (byte = 8 bits)  
CPU interprets based on how accessed

43

## memory

address	value
0xFFFFFFFF	0x14
0xFFFFFFF0	0x45
0xFFFFFFF4	0xDE
...	...
0x00042006	0x06
0x00042005	0x05
0x00042004	0x04
0x00042003	0x03
0x00042002	0x02
0x00042001	0x01
0x00042000	0x00
0x00041FFF	0x03
0x00041FFE	0x60
...	...
0x00000002	0xFE
0x00000001	0xE0
0x00000000	0xA0

address	value
0x00000000	0xA0
0x00000001	0xE0
0x00000002	0xFE
...	...
0x00041FFE	0x60
0x00041FFF	0x03
0x00042000	0x00
0x00042001	0x01
0x00042002	0x02
0x00042003	0x03
0x00042004	0x04
0x00042005	0x05
0x00042006	0x06
...	...
0xFFFFFFF4	0xDE
0xFFFFFFF0	0x45
0xFFFFFFFF	0x14

43

## endianness

address	value
0xFFFFFFFF	0x14
0xFFFFFFF0	0x45
0xFFFFFFF4	0xDE
...	...
0x00042006	0x06
0x00042005	0x05
0x00042004	0x04
0x00042003	0x03
0x00042002	0x02
0x00042001	0x01
0x00042000	0x00
0x00041FFF	0x03
0x00041FFE	0x60
...	...
0x00000002	0xFE
0x00000001	0xE0

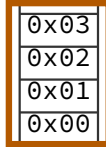
```
int *x = (int*)0x42000;  
cout << *x << endl;
```

44

# endianness

address	value
0xFFFFFFFF	0x14
0xFFFFFFFFE	0x45
0xFFFFFFFFD	0xDE
...	...
0x00042006	0x06
0x00042005	0x05
0x00042004	0x04
0x00042003	0x03
0x00042002	0x02
0x00042001	0x01
0x00042000	0x00
0x00041FFF	0x03
0x00041FFE	0x60
...	...
0x00000002	0xFE
0x00000001	0xE0

```
int *x = (int*)0x42000;  
cout << *x << endl;
```



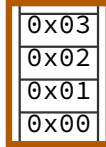
# endianness

address	value
0xFFFFFFFF	0x14
0xFFFFFFFFE	0x45
0xFFFFFFFFD	0xDE
...	...
0x00042006	0x06
0x00042005	0x05
0x00042004	0x04
0x00042003	0x03
0x00042002	0x02
0x00042001	0x01
0x00042000	0x00
0x00041FFF	0x03
0x00041FFE	0x60
...	...
0x00000002	0xFE
0x00000001	0xE0

```
int *x = (int*)0x42000;  
cout << *x << endl;
```

0x03020100 = 50462976

0x00010203 = 66051



# endianness

address	value
0xFFFFFFFF	0x14
0xFFFFFFFFE	0x45
0xFFFFFFFFD	0xDE
...	...
0x00042006	0x06
0x00042005	0x05
0x00042004	0x04
0x00042003	0x03
0x00042002	0x02
0x00042001	0x01
0x00042000	0x00
0x00041FFF	0x03
0x00041FFE	0x60
...	...
0x00000002	0xFE
0x00000001	0xE0

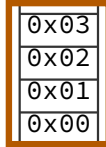
```
int *x = (int*)0x42000;  
cout << *x << endl;
```

0x03020100 = 50462976

little endian  
(least significant byte has lowest address)

0x00010203 = 66051

big endian  
(most significant byte has lowest address)



# endianness

address	value
0xFFFFFFFF	0x14
0xFFFFFFFFE	0x45
0xFFFFFFFFD	0xDE
...	...
0x00042006	0x06
0x00042005	0x05
0x00042004	0x04
0x00042003	0x03
0x00042002	0x02
0x00042001	0x01
0x00042000	0x00
0x00041FFF	0x03
0x00041FFE	0x60
...	...
0x00000002	0xFE
0x00000001	0xE0

```
int *x = (int*)0x42000;  
cout << *x << endl;
```

0x03020100 = 50462976

little endian  
(least significant byte has lowest address)

0x00010203 = 66051

big endian  
(most significant byte has lowest address)

