



CPU

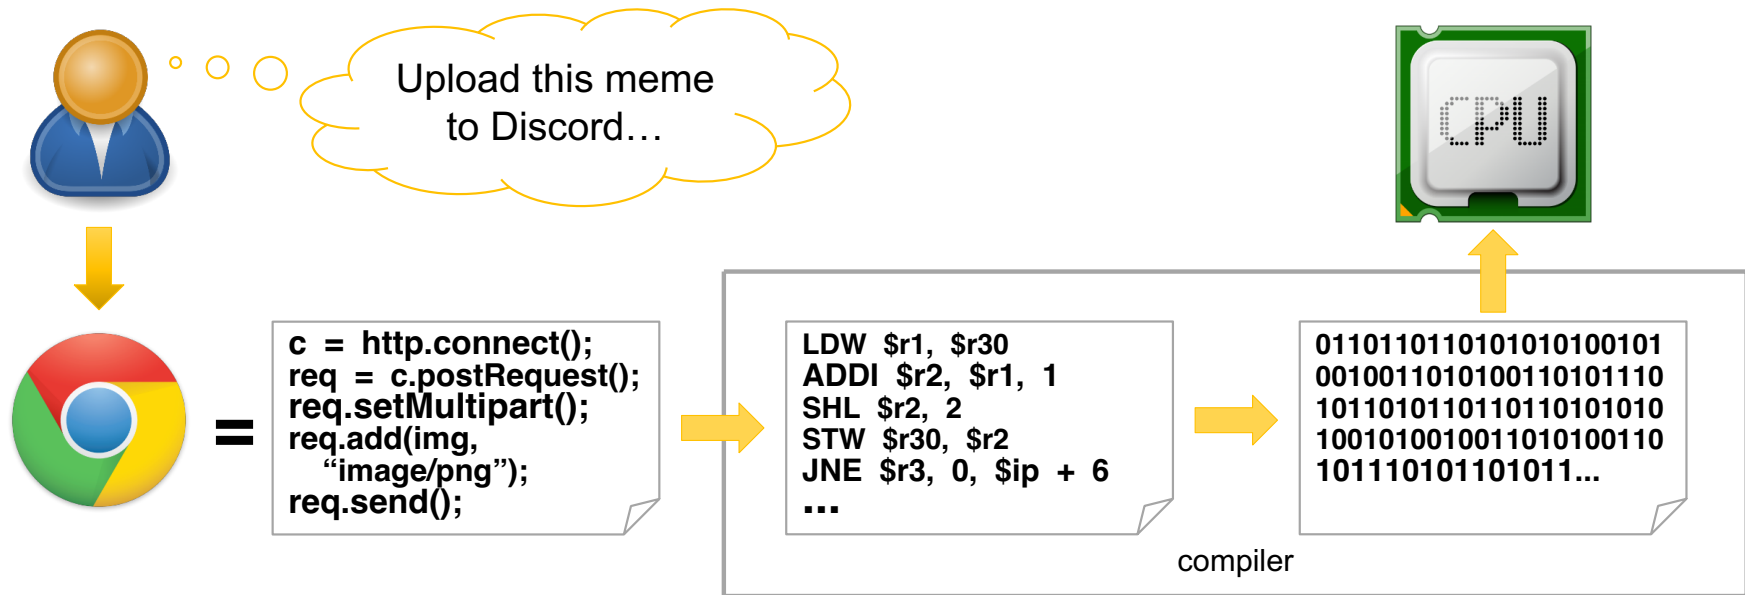
NSWI170 Computer Systems

Jakub Yaghob, Martin Kruliš



Semantic gap

- Gap between user intent and instructions executed by CPU



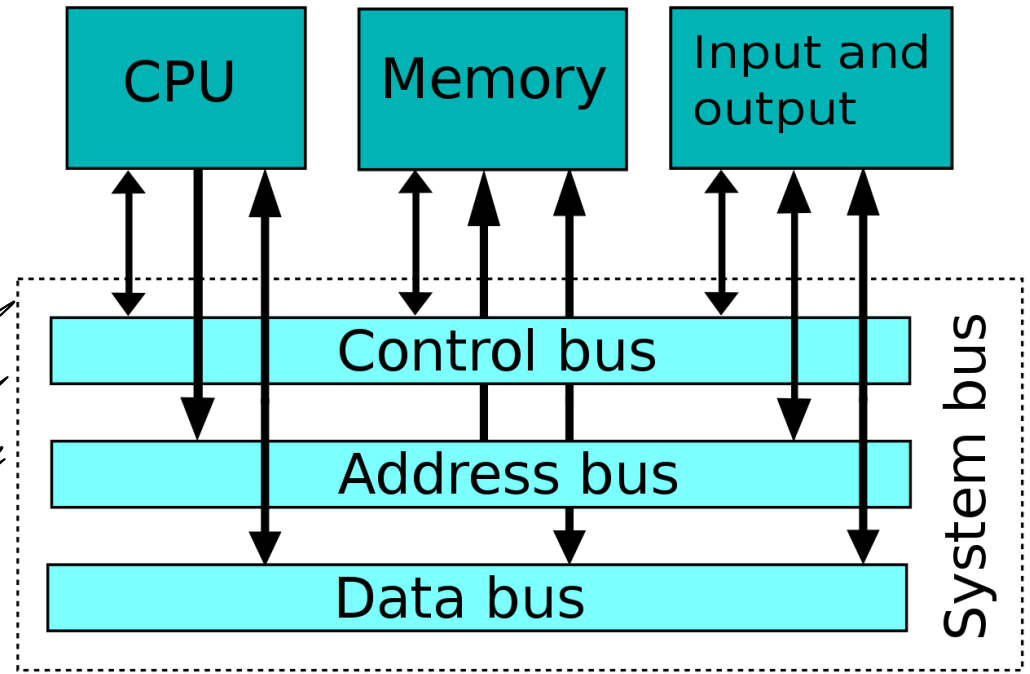


Von Neumann architecture

- System architecture
 - Program and data in the same memory
 - One shared bus
 - Evolution of the architecture

*Je všude HW jednoduchá
Ale pomalější, hlavně kvůli
multi-jádrům*


*Jedním společným
sdíleným sběrnicí
||
pomalejší,
protože je
jeden exkluzivní přenos*



Harvard architecture



→ Takle je uvěření pro kód



- Separate memory spaces
 - With separate buses
 - And separate addressing mechanism (instructions)
- Still used in some SoC microcontrollers

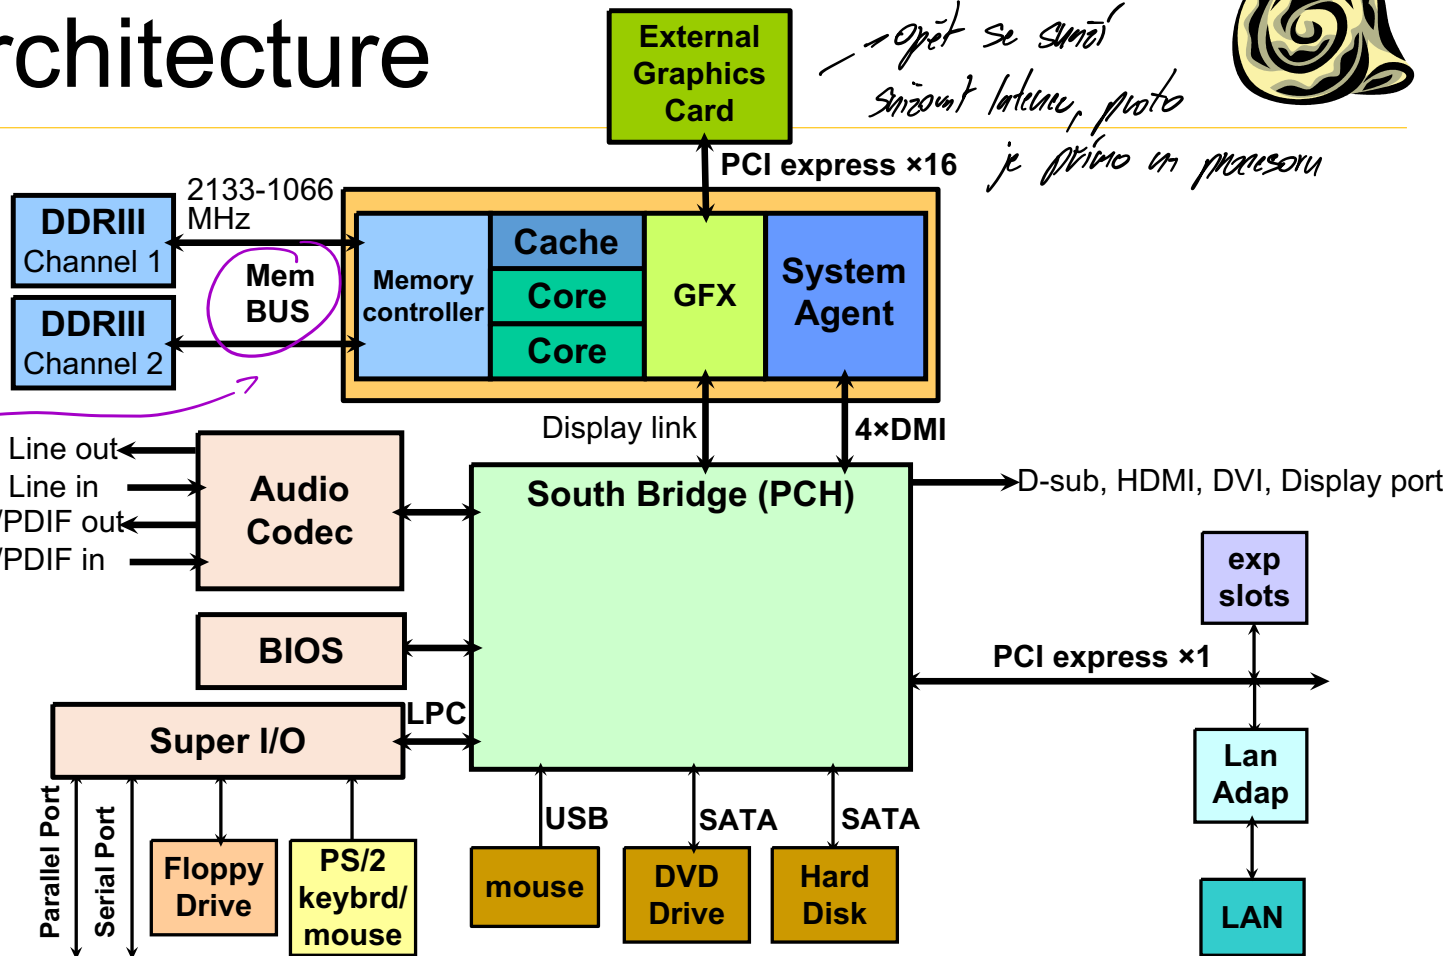
Většinou tam bývá ještě více datových míst

Real PC architecture



*opět se snaží
smažout latenci, proto
je přímo v procesoru*

- Sandy Bridge



*Vlastní paměť
sběrnice pro
vyšší rychlost*

*Funguje se in sériové
sběrnici, protože uvnitř
přesledky jiného vodiče*

CPU



- CPU Architecture

- Instruction Set Architecture (ISA)

- Abstract specification (contract between programmers and HW designers)

- Hardware architecture

- Actual implementation

} „Jak je to postavené“

*Popisuje logický chování procesoru,
více o detailech které ovlivňuje a co určuje*

- What is CPU?

- "Simple" machine that executes instructions

- Instruction – simple command, arithmetic operation, ...

- Registers, memory controllers, I/O

*x86 – nejpopulárnější
– dostupný*


*Na základě ISA se
dělá vlastní HW implementace.*



Instruction

- Simple command to the CPU
 - Binary encoding (CPU), assembler (programmers)

MIPS32 Add Immediate Instruction



Fixed vs. variable lengths

Equivalent mnemonic: **addi \$r1, \$r2, 350**

- Operands
 - Depending on ISA – max 1, 2, or 3 (some may be implicit)
 - Register, immediate value



Instruction cycle

- Sub-steps performed in every instruction
 - Load instruction from address stored in IP register
 - Decode the instruction
 - Load operands
 - Execute the operation
 - Store the result
 - Increment IP

Tchle jzan brzdy

Some steps may be skipped for some instructions



Instructions - motivation

- How can we execute the following code?

podmíněný
if (a < 3) b = 4; *nepodmíněný slok (abych přestočil ten else)*
else c = a << 2;

for (int i = 0; i < 5; ++i) a[i] = i;

```
int f(int p) { return p + 1; }  
void g() { auto r = f(42); }
```



Instruction classes

- Load instructions *→ Práceje s pamětí => jsou pomalí*
 - Memory -> register
 - Take a long time to execute, important to detect soon (fetch data ahead)
- Store instructions *→ Práceje s pamětí => je nejpomalejší*
 - Register/immediate -> memory
- Move instruction
 - Between registers
 - x86-64 also between registers and memory *→ To je prakticky load-store, move je ideální jen mezi registry*
 - Difficult to implement efficiently in HW



Instruction classes

- Arithmetic and logic instructions

- $+$, $-$, \ll , \gg , $\&$, $|$, \wedge , \sim
- $*$, $/$, $\%$

No funny stuff like `"str" * 4`

- Jumps

- Unconditional \times conditional
- Direct \times indirect \times relative

Tests required (`==`, `<`, ...)

tz musí mít

- Call, return

primo zadání adresy
adresy uložení v proměnné
přidání offsetu

- ...

softwarově

Stále musí být implementovaný způsob s úložitelnou adresou, kde jsem funkci volal.



Higher-level code structures

```
if (a < 3) b = 4; else c = a << 2;
```

```
...  
jge [a], 3  
store [b], 4  
jmp  
load r1, [a]  
shl r1, 2  
store [c], r1  
...
```

This is just a symbolic abstract
assembler
(not an actual one)



Higher-level code structures

```
for (int i = 0; i < 5; ++i) a[i] = i;
```

```
mov r1, 0  
jge r1, 5  
load r2, [a]  
add r2, r2, r1  
store [r2], r1  
add r1, r1, 1  
jmp  
...
```

Actually, the offset needs to be multiplied by `sizeof(a[0])`, but you got the point...

Registers *jednotky - stavby*



- Types

- General, integer, floating point, ^{bool}
- address, branch, flags, predicate,
- application, system,
- vector, ...

True: vybraná tato instrukce } nemusí se skáknout
False: nevybraná tato instrukce

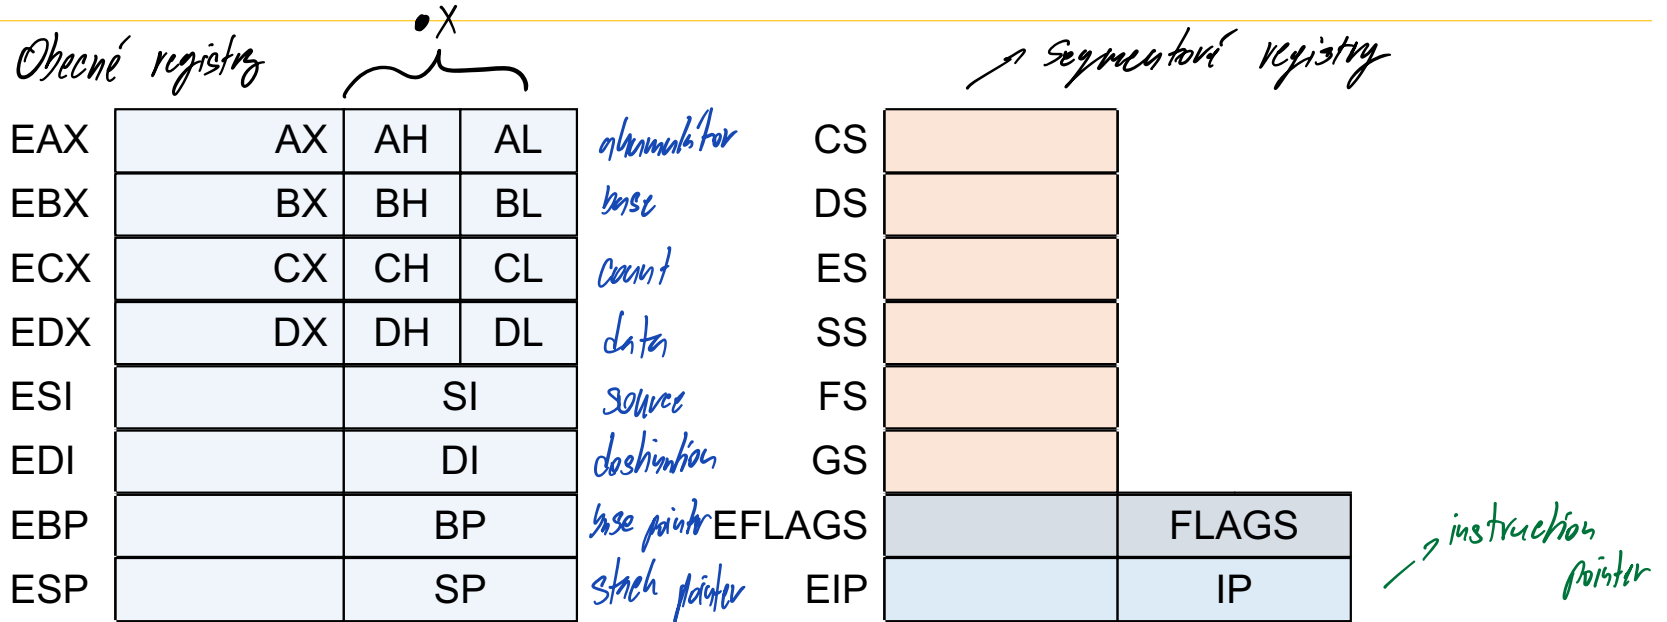
- Naming

- Direct (EAX, r01, ...) × stack (relative addressing to top of the stack)

- Aliasing



Registers – example 32-bit x86



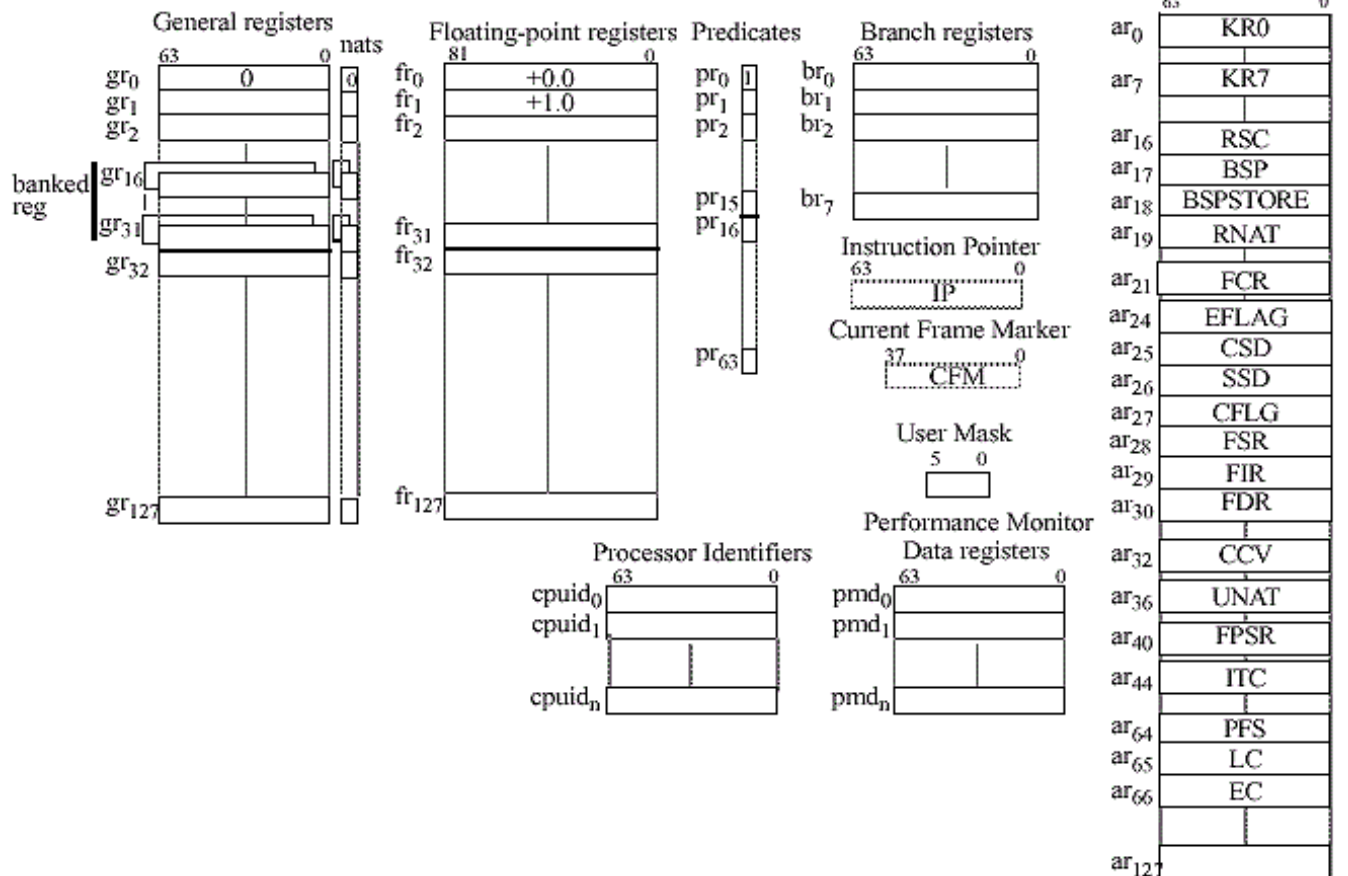
x86-64 extends registers to 64bits (RAX, RBX, ...)
+ adds general regs. R8 - R15

nejsem ortogonální -> registry jsou spojené s některými instrukcemi



Registers – example IA-64

APPLICATION REGISTER SET



MIPS – simple assembler



- Execution environment
 - 32-bit (almost general) registers r0-r31
 - r0 is always 0, writes are ignored → *běžná věc na sparcově architektuře*
 - r31 is a link register for the **jal** instruction
 - No stack
 - Realized by SW
 - No flags
 - Slightly different instruction set (especially tests and conditional jumps)
 - Program Counter (PC) register



Application Binary Interface

- Application Binary Interface (ABI) *— prescribes „job se to us about“*
 - Additional specification that accompanies ISA
 - Prescribes how the CPU should be used
 - Important for compiler designers
 - If application and a library are compiled separately, they need to be compatible
 - Binary interfaces between applications
 - Typical utilization (meaning) of registers
 - Stack implementation and layout
 - Function call conventions
 - ...

MIPS – register aliases *Takle specifikuje ABI*



Register	Name	Purpose	Preserve
\$r0	\$zero	0	N/A
\$r1	\$at	Assembler temporary	No
\$r2-\$r3	\$v0-\$v1	Return value	No
\$r4-\$r7	\$a0-\$a3	Function arguments	No
\$r8-\$r15	\$t0-\$t7	Temporaries	No
\$r16-\$r23	\$s0-\$s7	Saved temporaries – <i>garantovaná stabilita</i>	Yes
\$r24-\$r25	\$t8-\$t9	Temporaries	No
\$r26-\$r27	\$k0-\$k1	Kernel registers – DO NOT USE	N/A
\$r28	\$gp	Global pointer	Yes
\$r29	\$sp	Stack pointer	Yes
\$r30	\$fp	Frame pointer	Yes
\$r31	\$ra	Return address	Yes

*ostatní
funkce
to musí
uložit kdru
a pak zase
vratit*

Důležitost

MIPS – instructions



- Arithmetic

- **add \$rd,\$rs,\$rt**

- $R[rd] = R[rs] + R[rt]$

These are only symbolic placeholders, real example could look like **add \$r2,\$r4,\$r5**

- **addi \$rd,\$rs,imm16**

- $R[rd] = R[rs] + \text{signext}(\text{imm16})$

Immediate value, 16 bit number encoded directly in the instruction itself

- **sub \$rd,\$rs,\$rt**

Signed extension to 32bits (to match size of the registers)

- **subi \$rd,\$rs,imm16**



ISA comparison

MIPS

add \$t1,\$t1,\$t0

addi \$t1,\$t1,1

add \$t2,\$t0,\$t1

*↑
Třídresová instrukční sada je vždy vyžadována*

x86

→ tedy jde jen $eax += ebx$

add eax,ebx

add eax,1

or **inc eax**

mov eax,ebx

add eax,ecx



MIPS – instructions

- Logic operations

- **and** \$rd,\$rs,\$rt

- andi** \$rd,\$rs,imm16

- **or** \$rd,\$rs,\$rt

- ori** \$rd,\$rs,imm16

- **xor** \$rd,\$rs,\$rt

- xori** \$rd,\$rs,imm16

- **nor** \$rd,\$rs,\$rt

- R[rd] = R[rs] and/or/xor **zeroext**

- (imm16)

- No **not** instruction, use **nor** \$rd,\$rs,\$rs

- Shifts

- **sll** / **slr** \$rd,\$rs,shamt

- R[rd] = R[rs] << / >> shamt

Arithmetic shift (keeps the sign)

- **sra** \$rd,\$rs,shamt

ISA comparison



MIPS

nor \$t1,\$t2,\$t2

sll \$t1,\$t1,3

x86

mov eax,ebx

not eax

shl eax,3

MIPS – instructions

MIPS má řešení registry 32 bit only.
(Nemá aliasing)



• Memory access *kecama si 32bitu do registry*

„load word“

• **lw** \$rd,imm16(\$rs)

- $R[rd] = M[R[rs] + \text{signext32}(imm16)]$

„store word“

• **sw** \$rt,imm16(\$rs)

- $M[R[rs] + \text{signext32}(imm16)] = R[rt]$

„load byte“

• **lb** \$rd,imm16(\$rs)

- $R[rd] = \text{signext32}(M[R[rs] + \text{signext32}(imm16)])$

• **lbu** \$rd,imm16(\$rs)

- $R[rd] = \text{zeroext32}(M[R[rs] + \text{signext32}(imm16)])$

„store byte“

• **sb** \$rt,imm16(\$rs)

- $M[R[rs] + \text{signext32}(imm16)] = R[rt]$

• Moves *„load immediate“*

• **li** \$rd,imm32

- $R[rd] = imm32$
- Pseudo-instruction, translates to **lui** and **ori**

- Load upper immediate

• **move** \$rd,\$rs

- $R[rd] = R[rs]$



ISA comparison

MIPS

```
lw    $t1,1234($t0)
sw    $t1,1234($t0)
lb    $t1,1234($t0)
li    $t1,5678
move  $t1,$t0
```

x86

```
mov  eax,[ebx+1234]
mov  [ebx+1234],eax
mov  al,[ebx+1234]
mov  eax,5678
mov  eax,ebx
```



MIPS – instructions

- Jumps
 - **j addr** *→ přímý*
 - PC = addr
 - **jr \$rs** *→ nepřímý*
 - PC = R[rs]
 - **jal addr**
 - “Jump and link”
 - R[31] = PC+4
 - PC = addr
- jal instruction takes up 4 bytes

ISA comparison



MIPS

```
j    label
jr   $ra
```

```
label:
jal fnc
```

Placeholder that marks a place in the code, the actual address is computed by a compiler

x86

```
jmp label
jmp [ebx]
```



```
label:
call fnc
```

Return address stored to stack!

This is double indirection!!! Register **ebx** holds address where is the pointer that is loaded and used as target address for the jump!



MIPS – instructions

- Conditional jumps

- **beq \$rs,\$rt,addr** *→ rovnost*

- If $R[rs] == R[rt]$ then $PC = addr$ else $PC = PC+4$

- **bne \$rs,\$rt,addr**

- Analogical (not equal) *→ nerovnost*

Přestože nejsou vyjádřeny všechny podmínky, umíme je rovnou nebo druháky vyřešit.

- Testing

- **slt / sltu \$rd,\$rs,\$rt**

- If $R[rs] < R[rt]$ then $R[rd] = 1$ else $R[rd] = 0$

- **slti / sltiu \$rd,\$rs,imm16**

- If $R[rs] < \text{signext/zeroext}(imm16)$ then $R[rd] = 1$ else $R[rd] = 0$

sltu is unsigned version of **slt**

Only lesser-than test, greater-than is created by swapping operands

ISA comparison



MIPS

```
beq $t0,$t1,label
```

```
slt $t2,$t1,$t0  
bne $t2,$zero,label
```

```
slti $t2,$t1,5  
bne $t2,$zero,label
```

x86

```
cmp eax,ebx  
jz label
```

Universal compare
(result stored in flags)

Jump decision based on
flags

```
cmp eax,ebx  
jl label
```

```
cmp eax,5  
jl label
```

Code examples



- Simple for-loop

```
for (int i = 0; i < N; i++)  
{  
    A[i] = 42;  
}
```

```
addi $t0, $gp, 28    # $t0 <- address of A  
lw   $t1, 4($gp)    # fetch N  
move $t2, $zero     # i = 0  
ori  $t3, $zero, 42 # $t3 = 42  
j    cond           # go to condition  
  
body:  
sll  $t4, $t2, 2    # i*4 -> offset  
add  $t4, $t4, $t0  # A+i*4  
sw   $t3, 0($t4)    # A[i] = 42  
addi $t2, $t2, 1    # i = i+1  
  
cond:  
slt  $t4, $t2, $t1  # are we there yet?  
bne  $t4, $zero, body # no, we're NOT done
```



Code examples

$i = N*5+3;$

lw	\$t0, 4(\$gp)	# fetch N	lw	\$t0, 4(\$gp)	# fetch N
ori	\$t1, \$zero, 5	# 5	sll	\$t1, t0, 2	# N*4
mult	\$t0, \$t0, \$t1	# N*5(fake)	add	\$t1, \$t1, t0	# N*4+N
addi	\$t0, \$t0, 3	# N*5 + 3	addi	\$t1, \$t1, 3	# N+3
sw	\$t0, 0(\$gp)	# i = ...	sw	\$t1, 0(\$gp)	# i = ...


Code can be optimized...



Flags

- Only used by some ISA
- Control execution
- Check status of the last instruction
- Usual flags
 - Z – zero flag
 - S – sign flag
 - C – carry flag

x86 example



Reserved flags



System flags



Arithmetic flags



ISA consolidation

- Instruction set architecture
 - Abstract model of CPU
- Classification
 - CISC – Complex Instruction Set Computer
 - RISC – Reduced Instruction Set Computer
 - VLIW – Very Long Instruction Word
 - EPIC – Explicitly Parallel Instruction Computer
- Orthogonality
 - Accumulator
- Load-Execute-Store

ARM býval RISC, už je ale moc komplexní
x86 býval CISC, nyní se posouvá do RISC
- ty dvě věci jsou jen kvůli zpětné kompatibilitě

→ nem' žádný dekoder, aby to bylo rychlejší
↳ zjednodování, co lze dělat paralelně.

→ ARM, MIPS...
Jen x86 nemá dost registrů...

↑
Tohle ale nebrzdí execution, jen loading/store

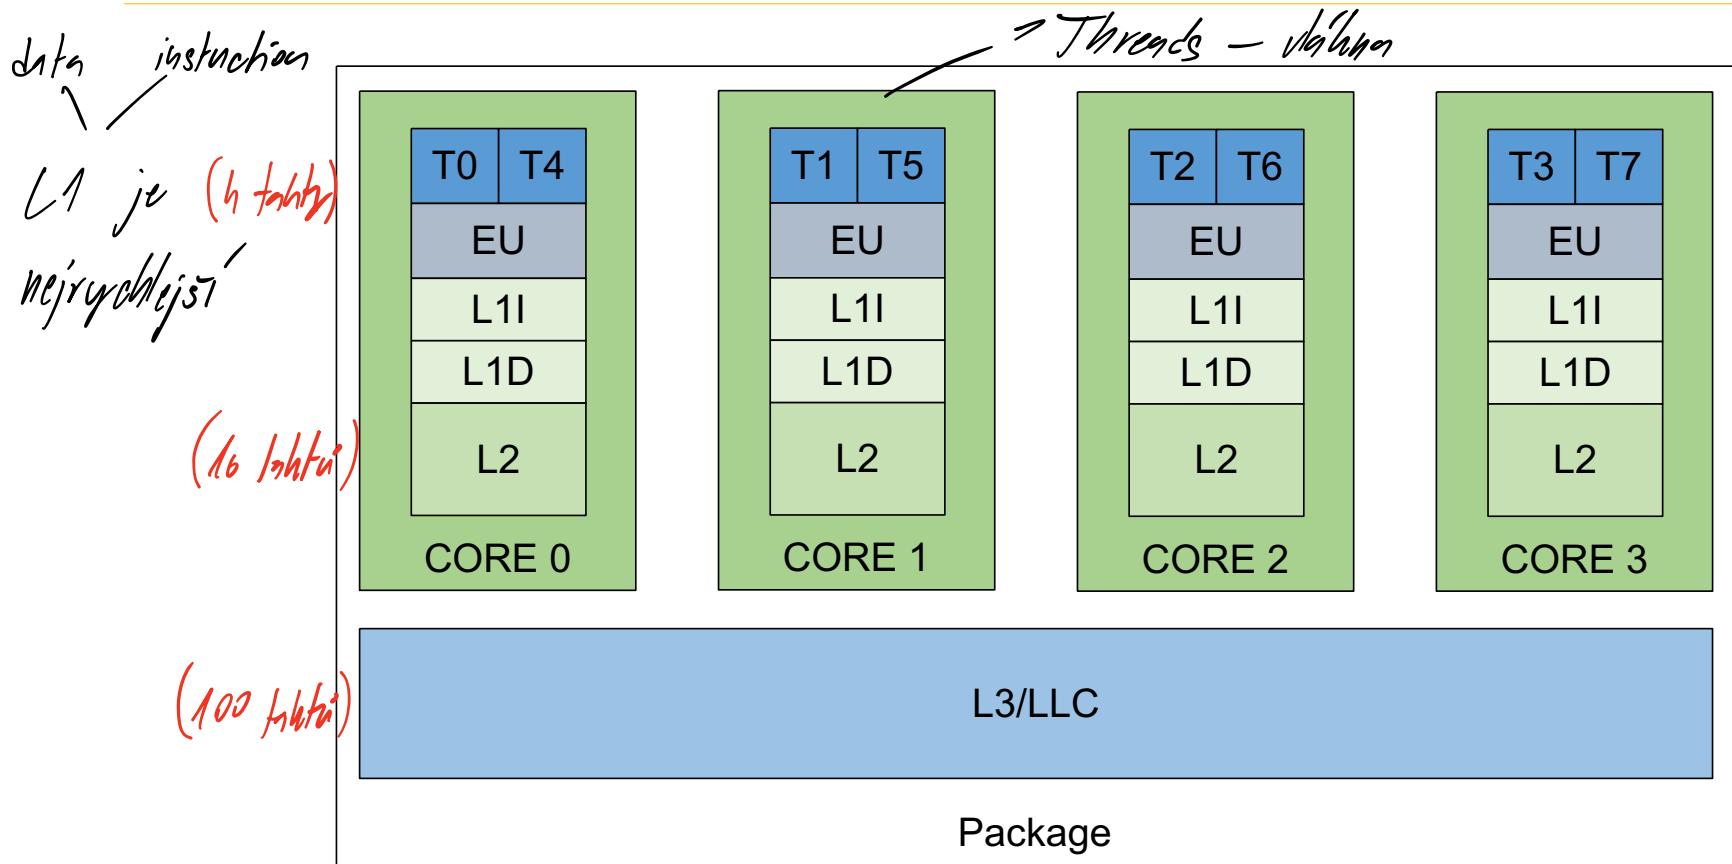


Hardware architecture

- CPU comprise...
 - Memory controller *— kváti vyčísí rychlosti paměti*
 - Cache hierarchy *— L1-3 ... různé rychlosti, velikosti*
 - Core
 - Registers
 - Types
 - Logical processor
 - Hyper threading



CPU – simplified scheme

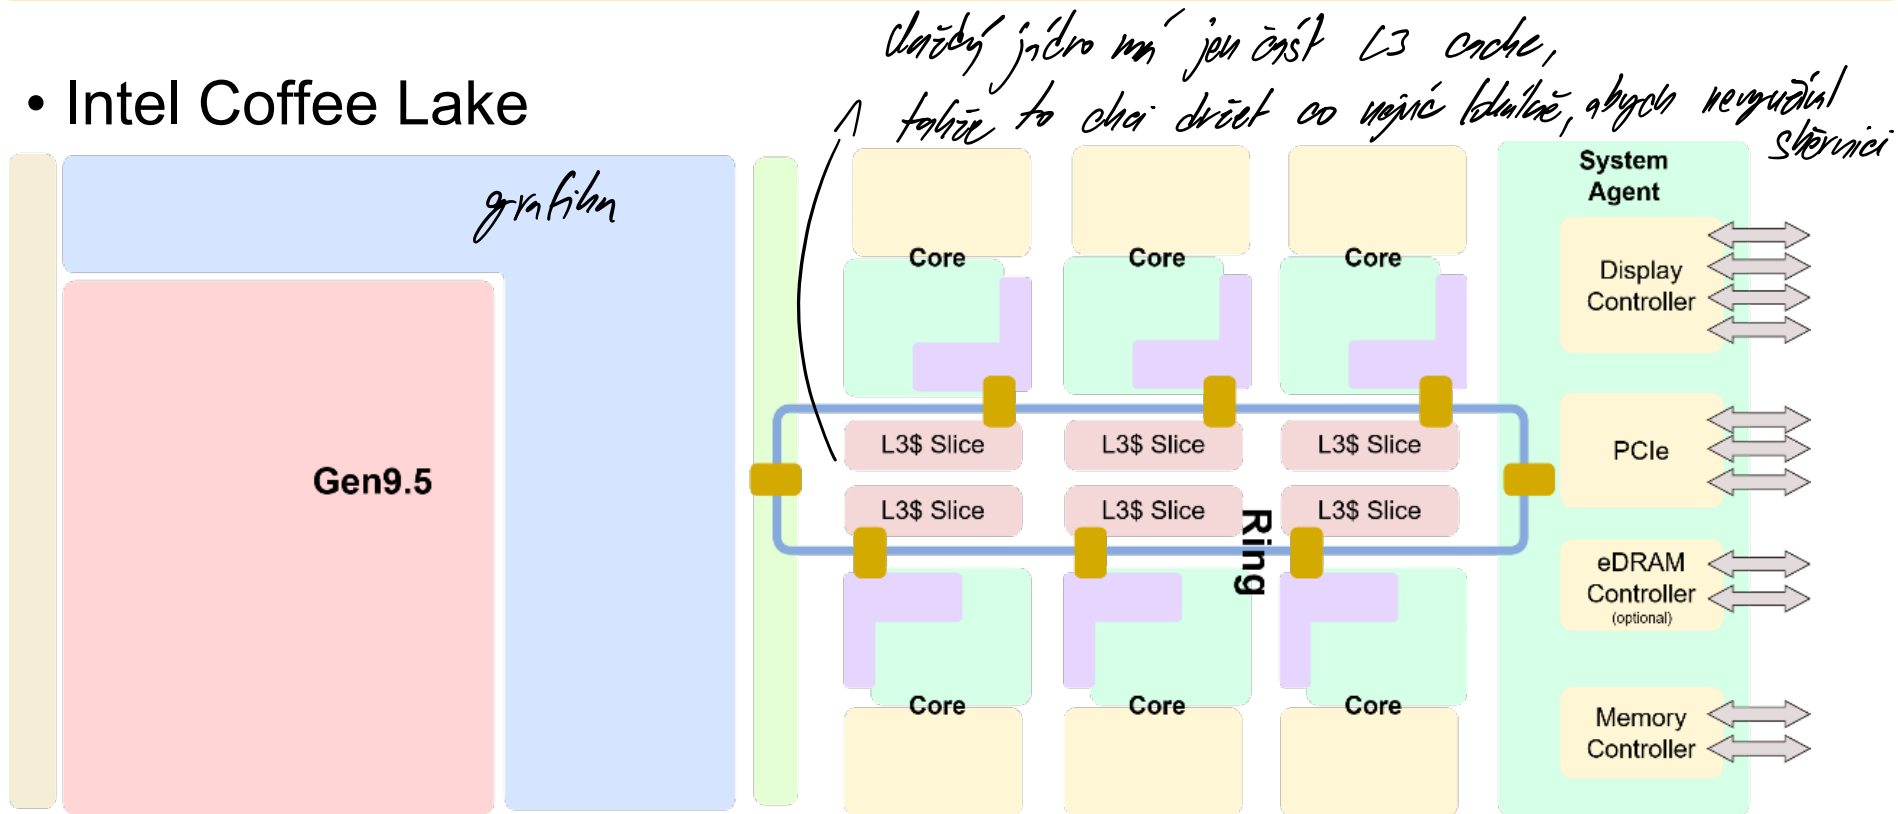


Jednotky nejsou totiž tolik vyčištěny, takže přichází ještě jedno vlákno. +20 až 50%
 (výpočetní jednotka je sdílená)



Scheme of Real CPU (Package)

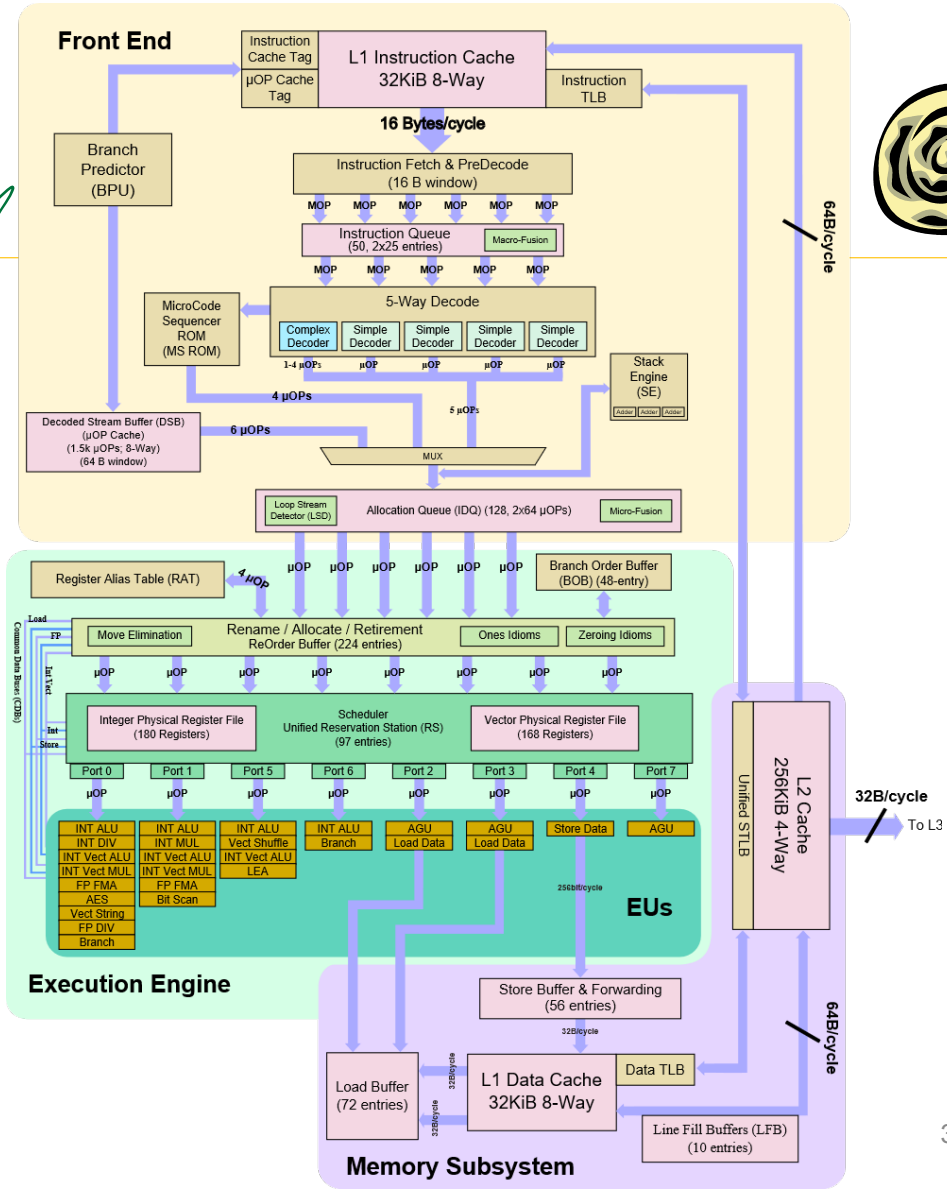
- Intel Coffee Lake




Real CPU

- Scheme of one core

*96%
использов.*



Real CPU Die



CPU architecture – pipeline



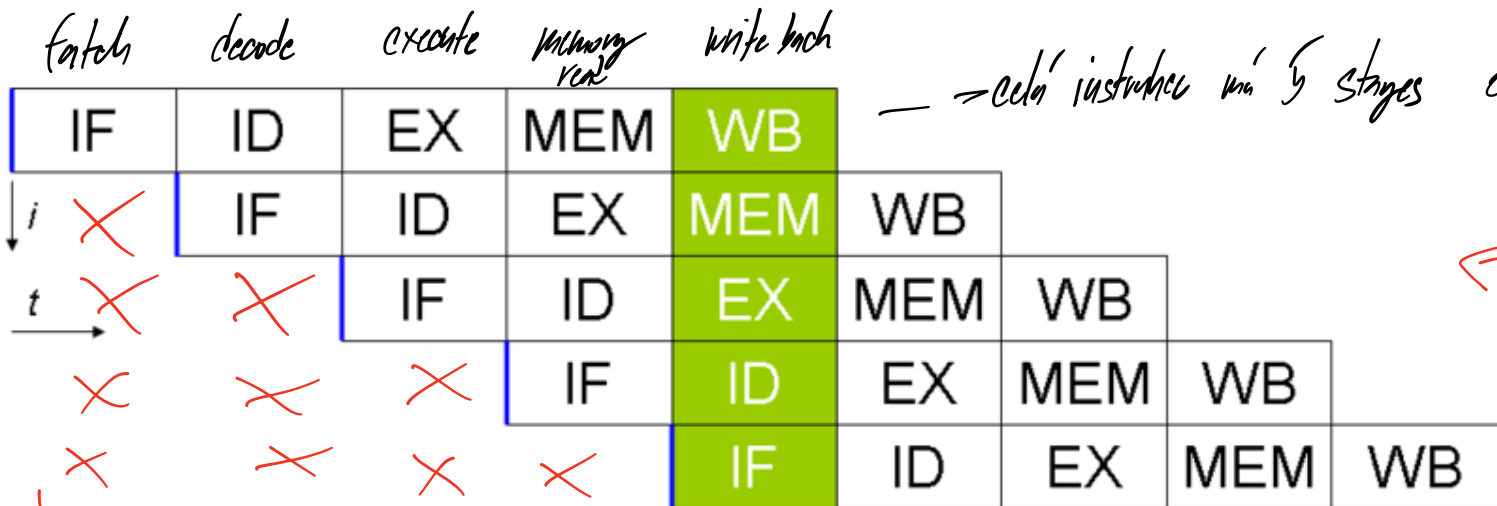
- Current CPU
 - 14-19 stages

– většinou jedem datum je jeden „fakt“

Může docházet ke kolísání

2a jednotku času jako základní úse

Nejhorší, co je, je při pipelingu, protože pak to musím spustit znovu



– celá instrukce má 5 stages ode

Náběh pipelingu, zejména

identický vzhled

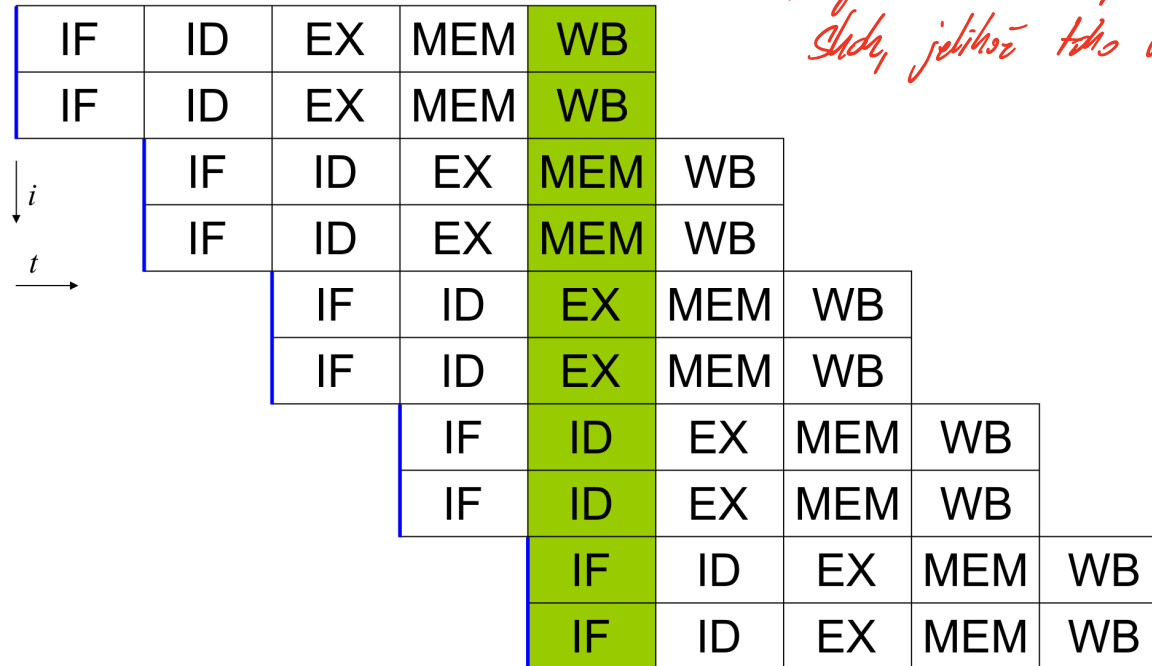
CPU architecture – superscalar processor



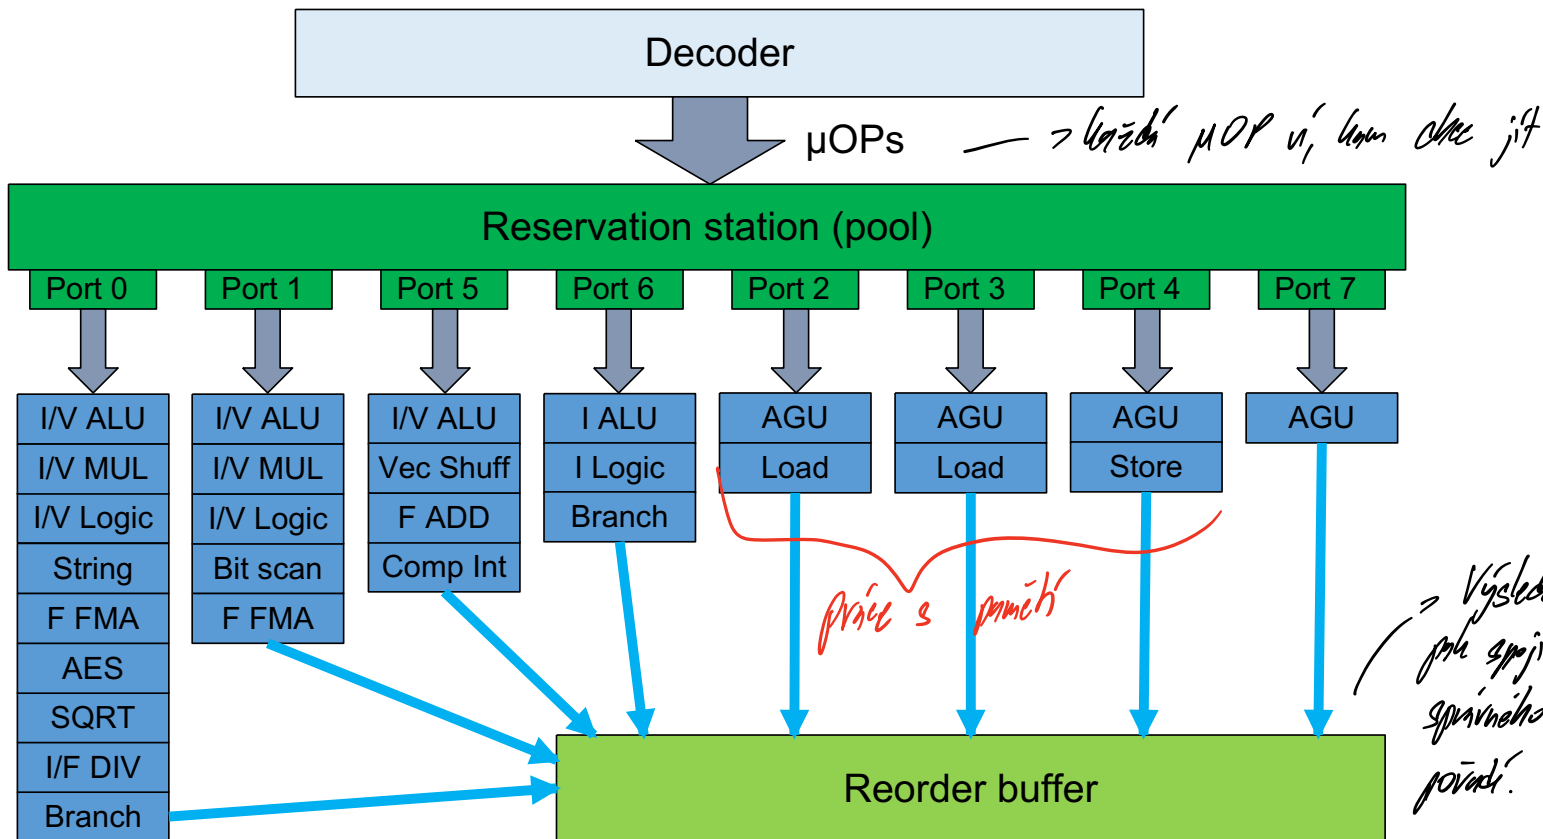
- Current CPU
 - 5-way, asymmetric

Můžeme mít více instrukcí mjedera

*Tady hodnotu kdi špatně odhadnuty
Shody, jelikož toho už hodnotě zabírají.*



CPU architecture – out-of-order execution



Discussion

