ARC: Computer Architecture

tanguy.risset@insa-lyon.fr
Lab CITI, INSA de Lyon
Version du March 27, 2025

Tanguy Risset

March 27, 2025

Table of Contents

ARC course presentation

- Schedule:
  - Course 4h
  - labs (TP) 20h
  - Evaluation (In french): un seul devoir papier en fin de cours
- skills and knowledge learned in ARC cours:
  - Bolean logic, arithmetics
  - combinatorial and sequential logic circuits, automata.
  - Processor architecture, datapath, compilation process, RISC architecture
  - Assembly code, link with high level programming languages
  - Simple processor design, simple assembly program analysis.
  - Link with compilation, operating systems and programming
- Moddle (open): frames, labs, various document
- Course based on the two IF architecture course: AC and AO (open courses on Moodle).

From electron to Von-Newman CPU

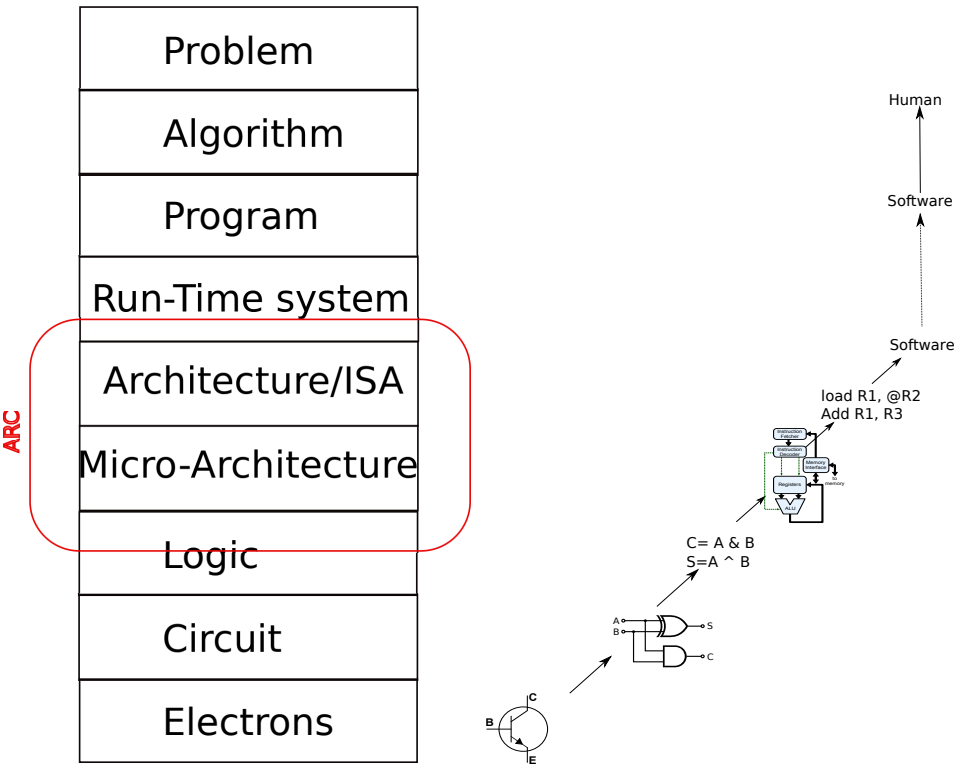

Computer architecture usefulness

- How to solve a problem with electrons:
- ARC is useful
  - For general knowledge of a computer scientist
  - To understand pro/cons of modern complex architectures
  - For embedded system programming

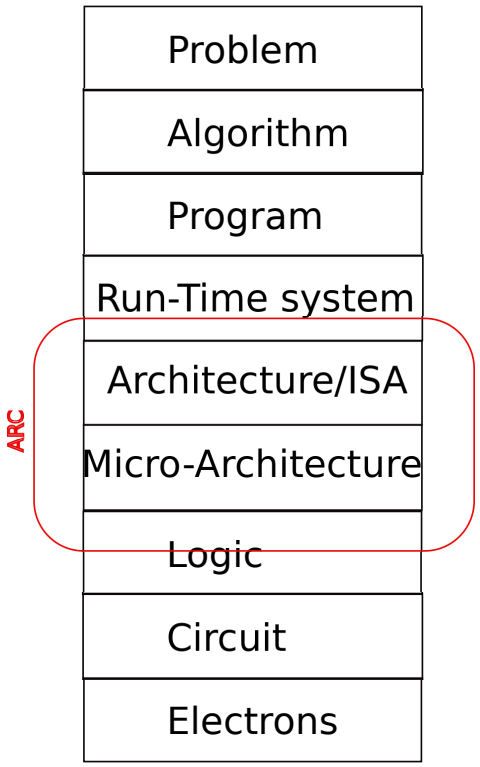

Table of Contents

History of computing

- Ancient time: various arithmetics systems
- 17th century (Pascal and Leibniz): notion of mechanical calculator
- 1822 Charles Babbage Difference engine (tabulate polynomial functions)
- 1854 Georges Boole proposes the so-called Boolean logic.
- (More details on the poly or on Internet)

from Yale Babylonian Collection, $\simeq$ 1600 BC

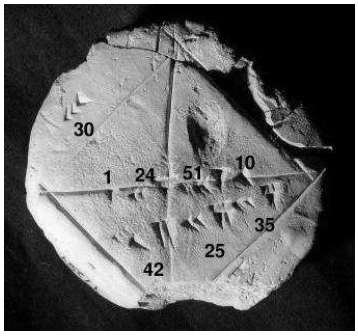

http://www.math.ubc.ca/~cass/Euclid/ybc/ybc.html

Difference Machine close-up

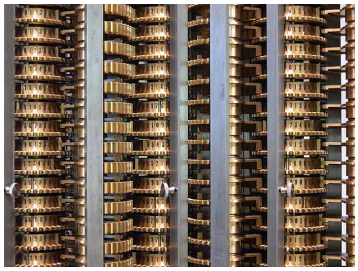

By By Carsten Ullrich - Own work, CC BY-SA 2.5

History of computers

- 1936: Alan Turing's PhD on a universal abstract machine
- 1941: Konrad Suze builds the Z3 first programmable computer (electro-mechanic)
- 1946: ENIAC is the first electronic calculator
- 1949: Turing and Von Neumann build the first universal electronic computer: the Manchester Mark 1
- (More details on the poly or on Internet)

Alan Turing

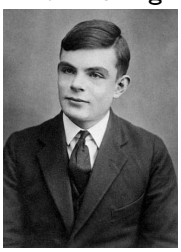

Z3 computer at Deutches Museum, Munich

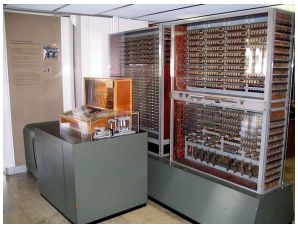

By Venusianer, CC BY-SA 3.0

Manchester Mark 1 1948

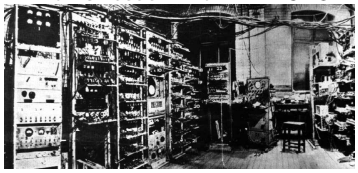

Table of Contents

Transistor

- Discovered in 1947 at Bell Labs: (**trans**fer res**istor**)
- Could replace the thermionic triode (vacuum tube) that allow radio and telephone technologies.
- Principle: flow or Interrupt current between Source and Drain, depending on Gate status

- Can be seen as a switch
- Wildly used after Integrated Circuit invention (1958)

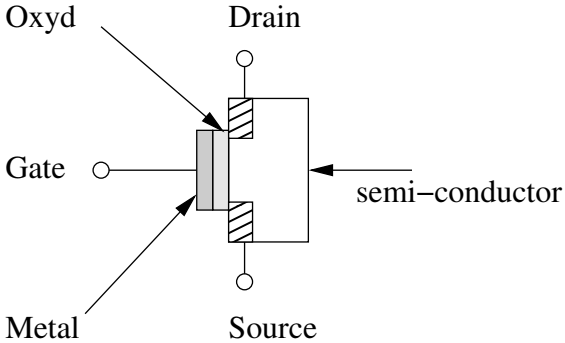

Mosfet technology

Popular Transistor technology: CMOS

- CMOS: Complementary Metal Oxide Semiconductor
- Two logical levels : 0 = 0V and 1 = 3V
- Two types of transistors
  - nMOS : current flows if gate is 1
  - pMOS : current flows if gate is 0
- Mainly used to realize basic logical gates (NOT, NAND, NOR, etc.)

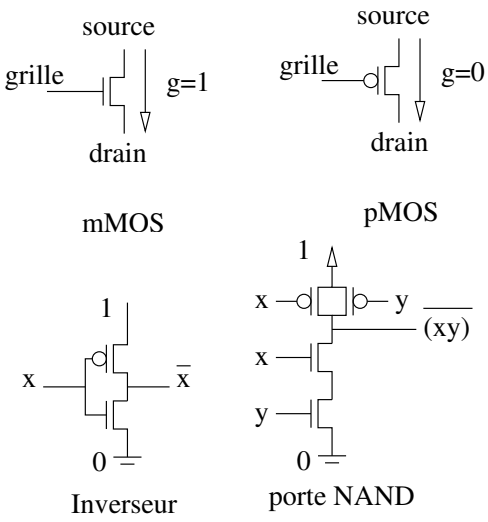

Moore's low

- Gordon Moore, co-founder of Fairchild Semiconductor and Intel, predicted in "a doubling every two year in the number of components per integrated circuit"
- Contributed to world economic growth
- Slow down in 2015 and is ended now.

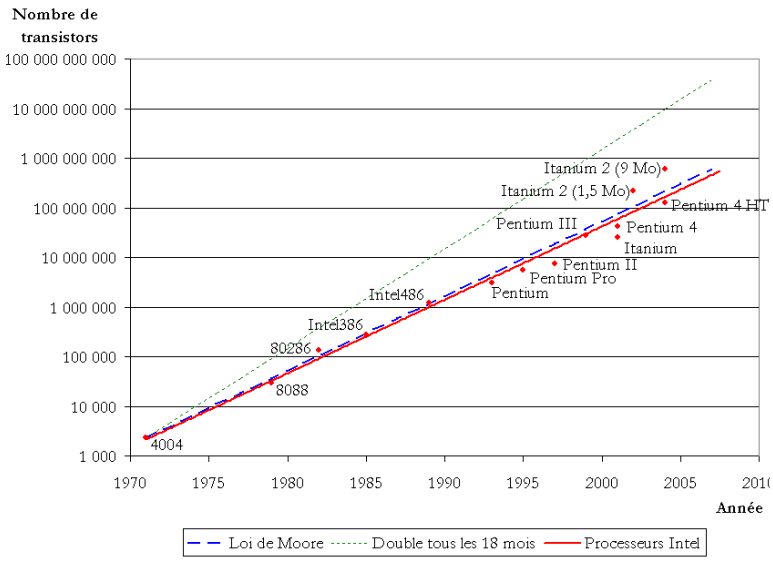

Boolean functions

Boole Algebra is equipped with three operations

- a unary operation, **negation**, noted `NOT`;
- two binary commutative, associative operations:
  - **conjunction** — `AND`, with 1 as neutral element;
  - **disjunction** — `OR`, with 0 as neutral element;
- `AND` is distributive over `OR`

If a and b are 2 boolean variables, we write:

$$\text{NOT}(a) = \overline{a}, \quad \text{AND}(a, b) = ab = a.b, \quad \text{OR}(a, b) = a + b$$

Boolean Cheat Sheet

- neutral elements: $a + 0 = a, \quad a \cdot 1 = a$
- absorbing elements: $a + 1 = 1, \quad a \cdot 0 = 0$
- idempotence: $a + a = a, \quad a \cdot a = a$
- tautology/antilogy: $a + \overline{a} = 1, \quad a \cdot \overline{a} = 0$
- commutativity: $a + b = b + a, \quad ab = ba$
- distributivity: $a + (bc) = (a + b)(a + c), \quad a(b + c) = ab + ac$
- associativity: $a + (b + c) = (a + b) + c = a + b + c,$
  $a(bc) = (ab)c = abc$
- De Morgan's law: $\overline{ab} = \overline{a} + \overline{b},$
  $\overline{a + b} = \overline{a} \cdot \overline{b}$
- others: $a + (ab) = a, \quad a + (\overline{a}b) = a + b,$
  $a(a + b) = a, \quad (a + b)(a + \overline{b}) = a$

Elementary logical gates

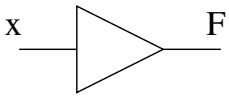

Amplifier:
$F = x$

x	F
0	0
1	1

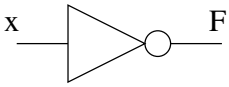

NOT: $F = \overline{x}$

x	F
0	1
1	0

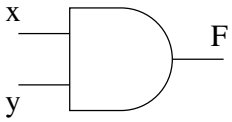

AND: $F = x\ y$

x	y	F
0	0	0
0	1	0
1	0	0
1	1	1

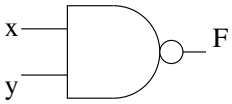

NAND:
$F = \overline{(x\ y)}$

x	y	F
0	0	1
0	1	1
1	0	1
1	1	0

Elementary logical gates

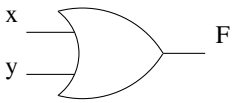

OR:
$F = x + y$

x	y	F
0	0	0
0	1	1
1	0	1
1	1	1

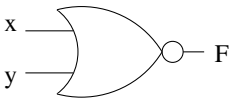

NOR:
$F = \overline{(x + y)}$

x	y	F
0	0	1
0	1	0
1	0	0
1	1	0

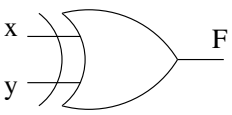

XOR:
$F = x \oplus y$

x	y	F
0	0	0
0	1	1
1	0	1
1	1	0

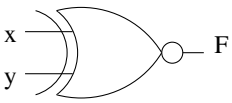

XNOR:
$F = x \odot y$

x	y	F
0	0	1
0	1	0
1	0	0
1	1	1

Combinatorical circuit Design

1. Boolean description of the problem:
   - Compute y and z from a, b and c
   - y is 1 if a is 1 or b and c are 1.
   - z is 1 if b or c is 1 (but not both) or if a, b et c are 1.
2. Truth table
3. Logic equation
   - $y = \overline{a}bc + a\overline{b}\overline{c} + a\overline{b}c + ab\overline{c} + abc$
   - $z = \overline{a}\overline{b}c + \overline{a}b\overline{c} + a\overline{b}c + ab\overline{c} + abc$
4. Optimized logic equations
   - $y = a + bc$
   - $z = ab + \overline{b}c + b\overline{c}$
5. logic gates

input			output	
a	b	c	y	z
0	0	0	0	0
0	0	1	0	1
0	1	0	0	1
0	1	1	1	0
1	0	0	1	0
1	0	1	1	1
1	1	0	1	1
1	1	1	1	1

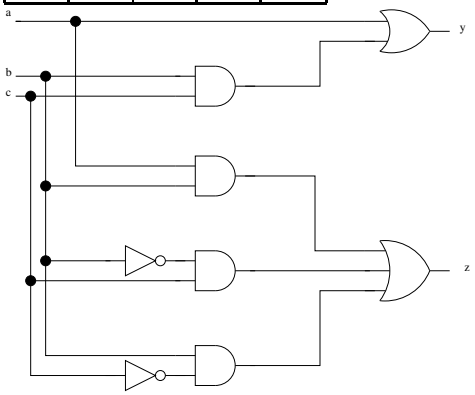

Disjunctive Normal Form (DNF)

- In Boolean logic, a logical formula in Disjunctive Normal Form (*Forme normale disjonctive* in French) if:
  - It is a disjunction of one or more clauses
  - where the clauses are conjunction of literals
  - a literal is a variable, a constant or 'not' a variable
- Otherwise put, it is an OR of ANDs.
- Example of DNF:
  - $x.\bar{y}.\bar{z} + \bar{t}.u.v$
  - $(a \wedge b) \vee \neg c$
- Example not in DNF:
  - $\overline{(x + y)}$
  - $a \vee (b \wedge (c \vee d))$

Conjunctive Normal Form (CNF)

- In Boolean logic, a formula is in conjunctive normal form (*forme normale conjonctive* in French) if:
  - it is a conjunction of one or more clauses,
  - where a clause is a disjunction of literals;
  - a literal is a variable, a constant or 'not' a variable
- Otherwise put, it is an AND of ORs.
- Example of CNF:
  - $(x + y + \bar{z})(\bar{x} + z)$
  - $(a + \bar{b} + \bar{c})(\bar{d} + \bar{a})$
  - $x + y$
- Example not in CNF
  - $\overline{(x + y)}$
  - $x(y + (z.t))$

From Truth table to DNF

- Back to previous example (z is 1 if b or c is 1 (but not both) or if a, b et c are 1.)
- Truth table on the right, z is 1 if and only if one of the five condition identified occurs.
- It is easy to find a conjunction that is valid in a unique case: example: $\bar{a}.\bar{b}.c$ is 1 if and only if: $a = 0$, $b = 0$ and $c = 1$ (double arrow on the right)
- by adding all the conjunction valid only on each of the five cases identified on the right, we get a DNF formulae that has exactly that truth table.

input				
a	b	c	z	
0	0	0	0	
0	0	1	1	⇐
0	1	0	1	←
0	1	1	0	
1	0	0	0	
1	0	1	1	←
1	1	0	1	←
1	1	1	1	←

Hence the possible formulae for z: $z = \overline{a}\overline{b}c + \overline{a}b\overline{c} + a\overline{b}c + ab\overline{c} + abc$

How can it be simplified?

Simple Boolean optimization: Karnaugh Table (1)

- Karnaugh map (*tables de Karnaugh*) use a "visual" reprentation of a simple property:
  $(a.\bar{b}) + (a.b) = a.(\bar{b} + b) = a$
- The first step in the method is to transform the truth table (3 or 4 input variables) of the function in a two-dimensional array (split into two parts of the set of variables)
- Rows and columns are indexed by the valuations of the corresponding variables in such a way that between two rows (or columns) only one boolean value changes.
- In our example (3 variables):

a b c	0 0	0 1	1 1	1 0
0	0	1	1	0
1	1	0	1	1

Simple Boolean optimization: Karnaugh Table (2)

- Then, we try to cover all '1' of the table by forming groups.
  - each group contains only adjacent '1'
  - must form a rectangle
  - the number of elements of a group must be a power of two.
- each group correspond to a possible optimization of the DNF
- In our example:

a b c	0 0	0 1	1 1	1 0
0	0	1	1	0
1	1	0	1	1

- example : Three groups:
  - $\bar{a}.b.\bar{c} + a.b.\bar{c}$ simplifies to $b.\bar{c}$
  - $a.b.\bar{c} + a.b.c$ simplifies to $a.b$
  - $a.\bar{b}.c + \bar{a}.\bar{b}.c$ simplifies to $\bar{b}.c$
- hence $z = \overline{a}\overline{b}c + \overline{a}b\overline{c} + a\overline{b}c + ab\overline{c} + abc$ simplifies to $z = a.b + \bar{b}.c + b.\bar{c}$

Well formed cicruits

As far as combinatorial circuits are concerned, a "Well formed" circuit is:

- A logic gate
- A wire
- Two well formed circuits next to each other
- Two well formed circuits, the outputs of one being the inputs of the other
- Two well formed circuits sharing a common input

It can be shown that it correspond to an acyclic graph of logic gates.

- No cycles, no ouptuts conected together

Usefull combinatorics logic components

- n input multiplexer
- decoder $log(n) \rightarrow n$
- n bits adder
- n bits comparator
- n bits ALU
- etc.

Memorizing: latches and Flip-Flops

- Set-Reset Latch (SR latch, *Bascule RS*): When R and S are reset, Q and $\overline{Q}$ keep their previous value.

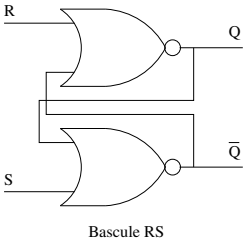

Bascule RS

S	R	Q	$\overline{Q}$
0	1	0	1
1	1	forbidden	forbidden
1	0	1	0
0	0	Q_{n-1}	$\overline{Q_{n-1}}$

- Gated D latch (Flip-flop, register, *Bascule D*): sample input data on clock rising edge and keeps the value when clock is 0.

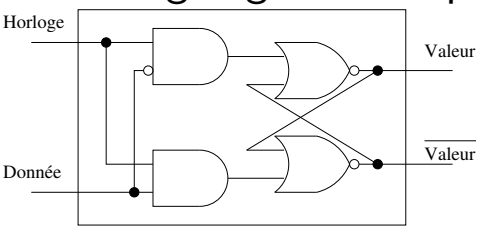

latches and Flip-Flops: other common representation

- Latch (*verrou*)

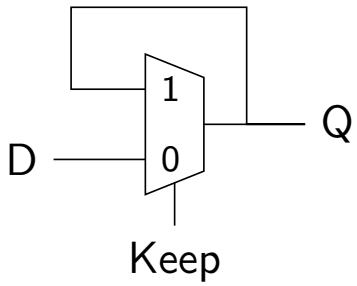

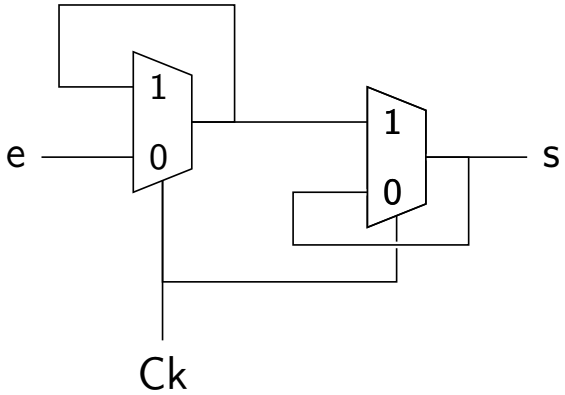

- Flip-Flop (register)

Sequential logic

Sequential logic combines logic function and memorizing, it opens the way to synchronous circuits, automata, programs, algorithms....

- n bits register
- n bits counter
- state machine
- CPU
- Computer

Sequential circuit design

- Extremely complex in general.
- Many computation models:
  - Sequential
    - State machine
    - control + data-path
  - task parallelism (communicating tasks)
  - Data parallelism (data-flow)
  - Asynchronous circuits
- Important notion use every where: finite state machine (*automate*)

Logic in ARC: Digital software

In ARC: use of Digital software (`https://github.com/hneemann/Digital`)

- Design basic logic components (TD1)
- Design of a memory (sequential component, TD2)
- Design of dedicated circuit: integer division (TD3).
- Study of a Von Neumann 8-bit processor (TD4)

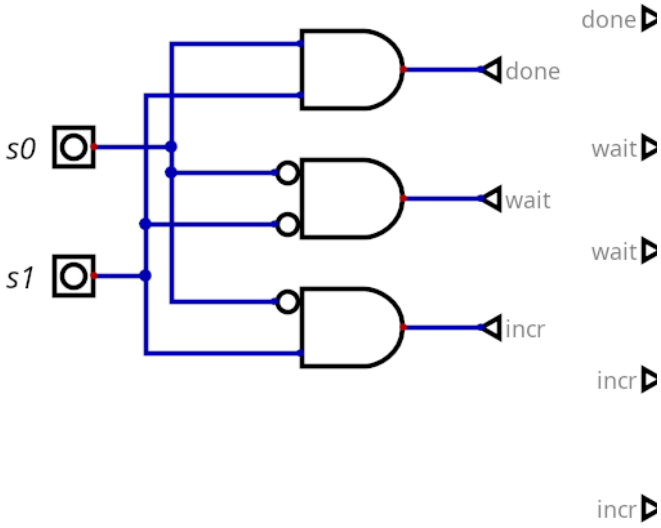

Table of Contents

What is a Von Neumann machine?

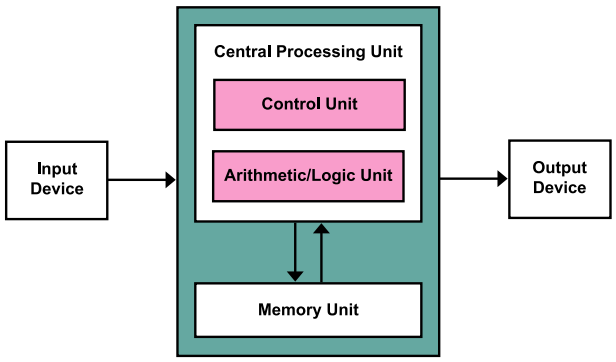

- Computer architecture Model (also called *Princeton* architecture) proposed after J. Von Neumann report: "First Draft of a Report on the EDVAC".
- Usually abstracted as a processor connected to a memory
- The memory is accessed (*randomly*) with an address (i.e. unlike a Turing machine)
- The memory contains both data and program (unlike a Harvard machine).

How does it work?

Compilation, Assembly code and binary code

High Level Language ⇒	Assembly code ⇒	Binary code ⇒
`int a,b,c;`	`load R0, @b`	`01001011...10101`
`a = b + c;`	`load R1, @c`	`01001010...10001`
	`add R3,R0,R1`	`...`
	`store R3, @a`	`10010011...00011`

Fast compilation thanks to Donald Knuth (and others..)

- The programmer:
  - Write a program (say a C program: `ex.c`)
  - Compiles it to an object program `ex.o`
  - links it to obtain an executable `ex`

content of `ex.c`

```
#include <stdio.h>

int main()
{
  printf("hello World\n");

  return(0);
}
```

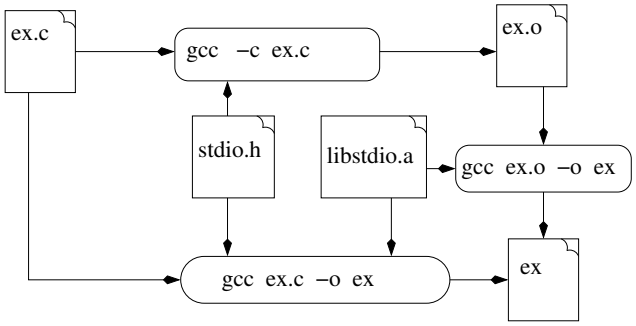

Program execution on a Processor (8 general purpose registers)

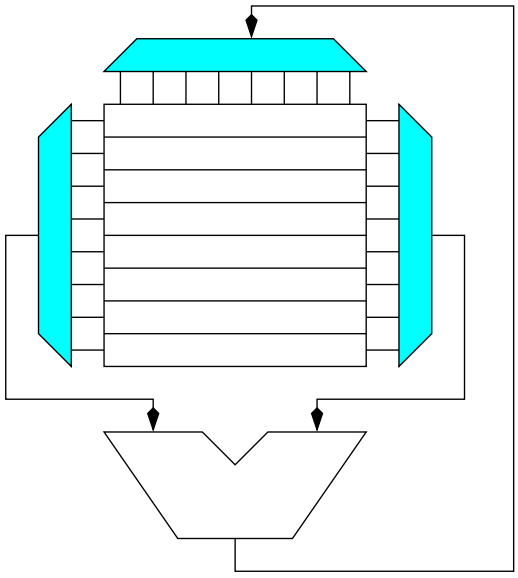

Program execution on a Processor (8 general purpose registers)

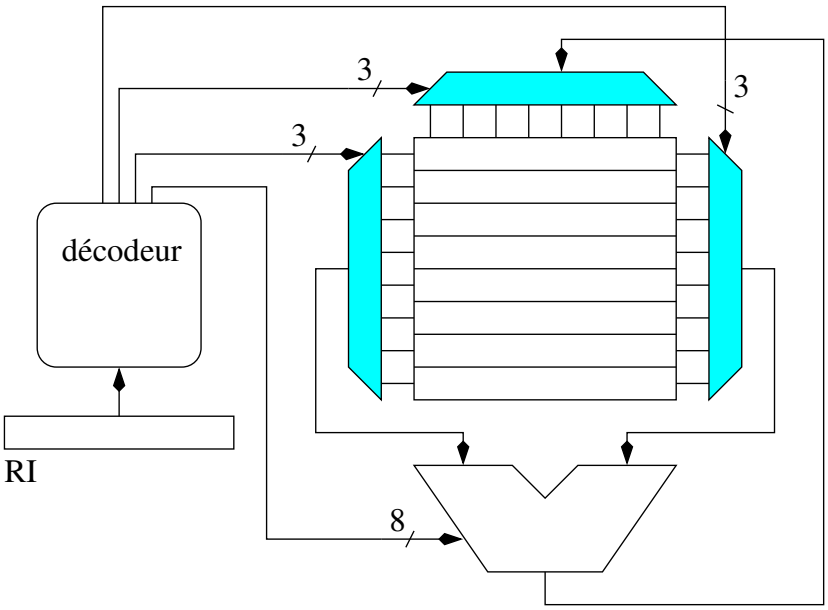

Program execution on a Processor (8 general purpose registers)

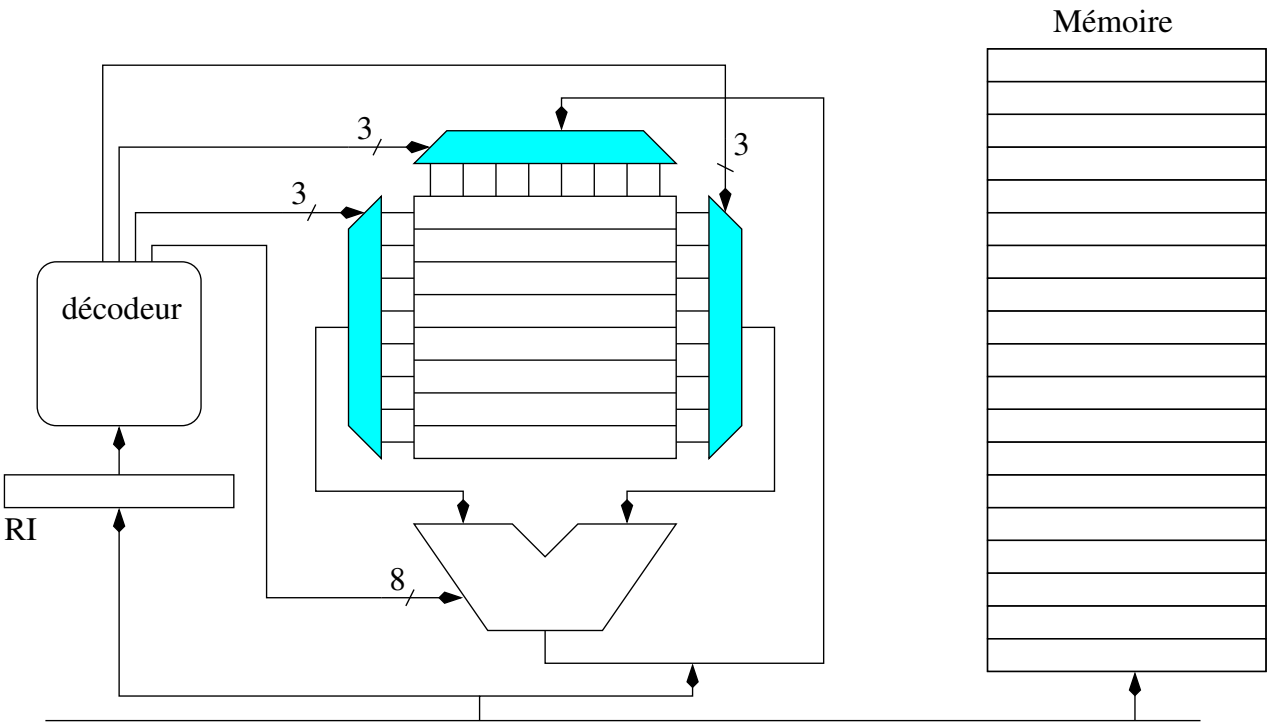

Program execution on a Processor (8 general purpose registers)

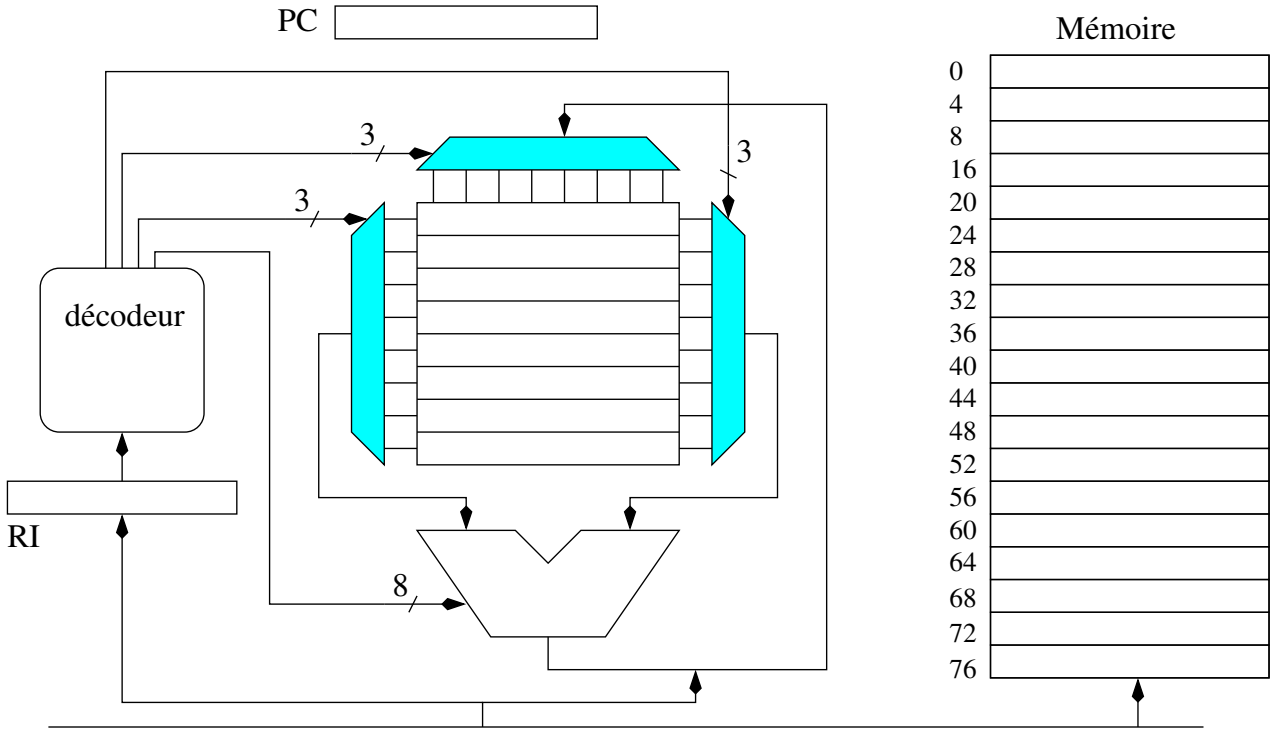

Program execution on a Processor (8 general purpose registers)

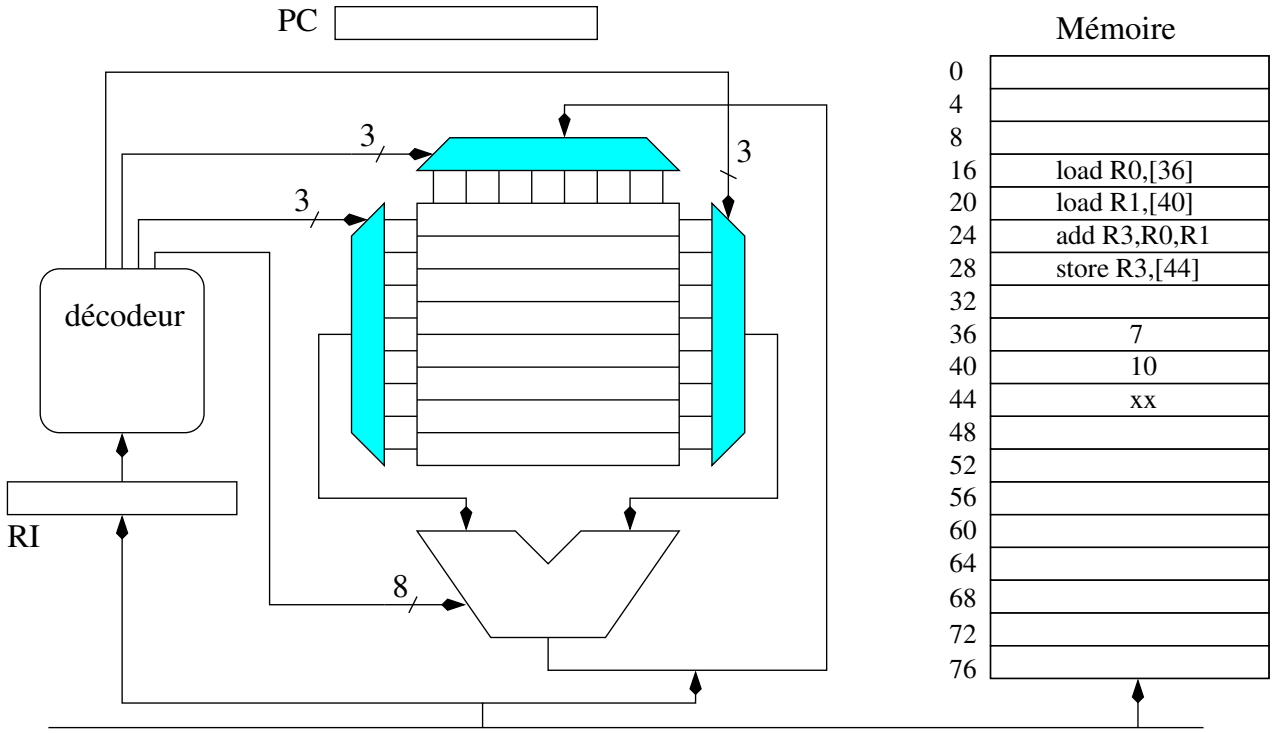

Program execution on a Processor (8 general purpose registers)

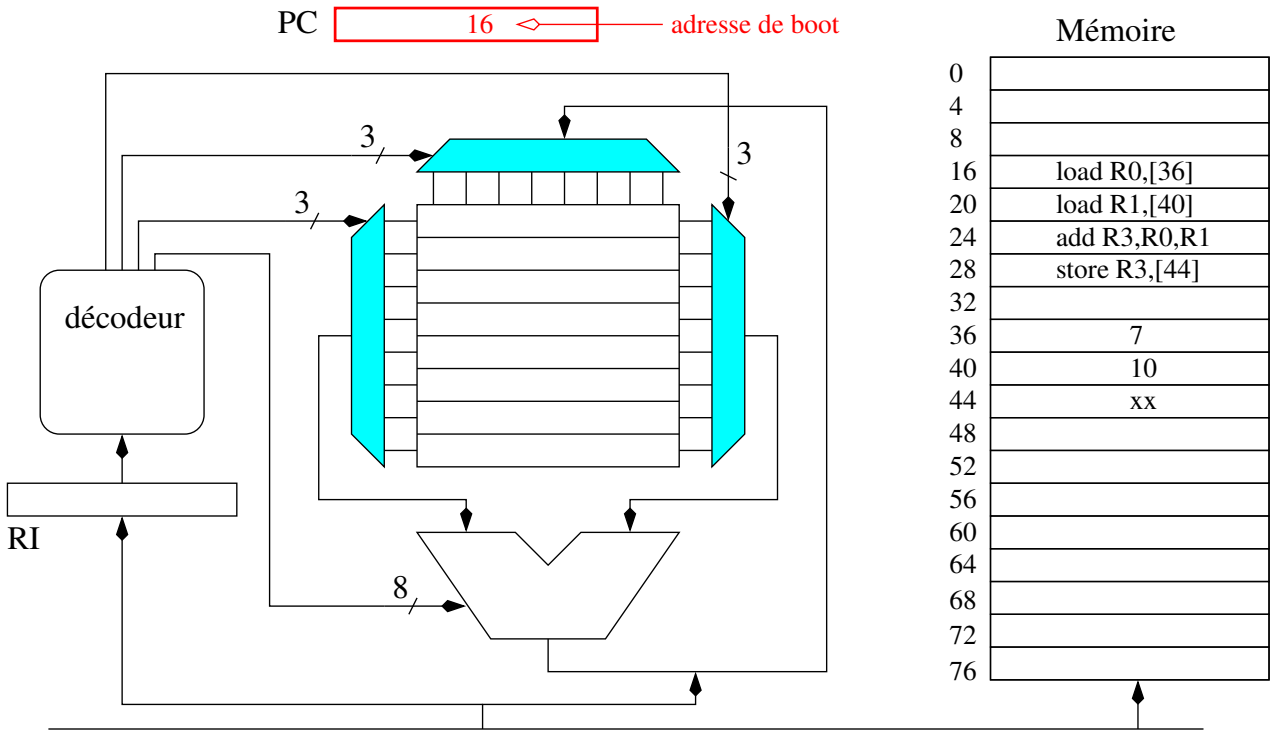

Program execution on a Processor (8 general purpose registers)

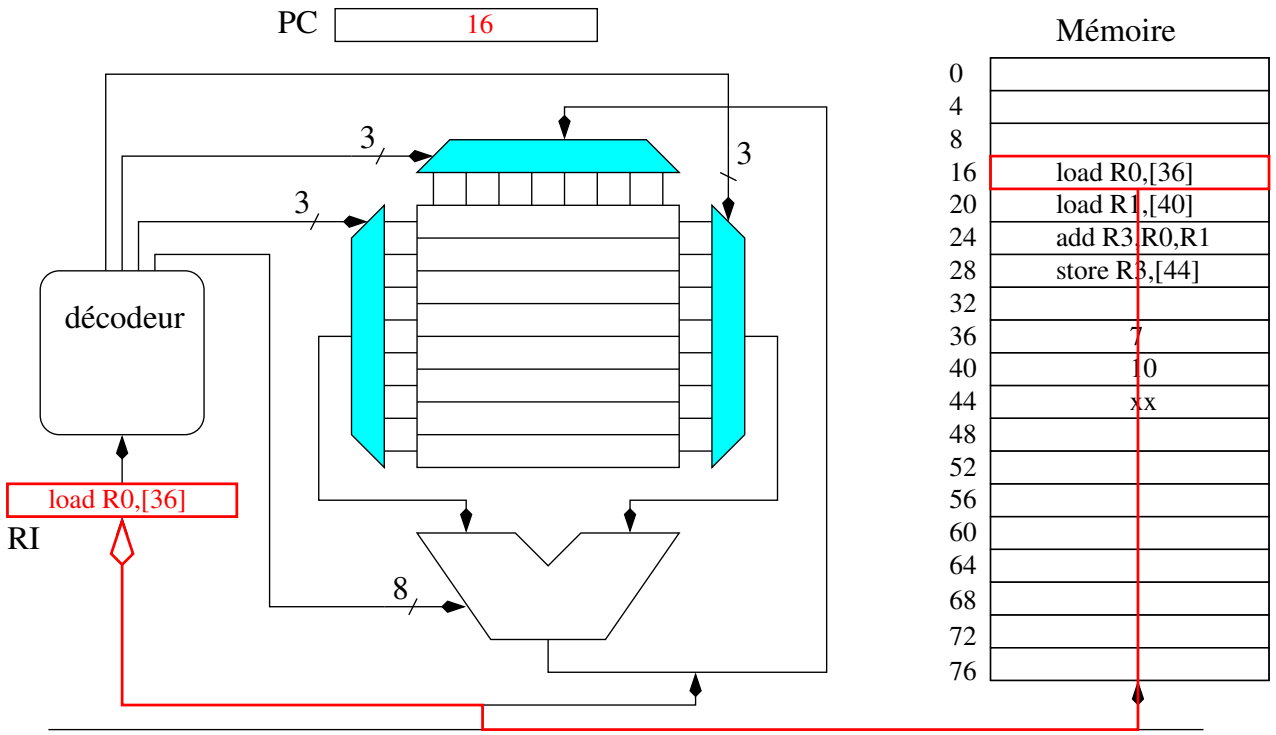

Program execution on a Processor (8 general purpose registers)

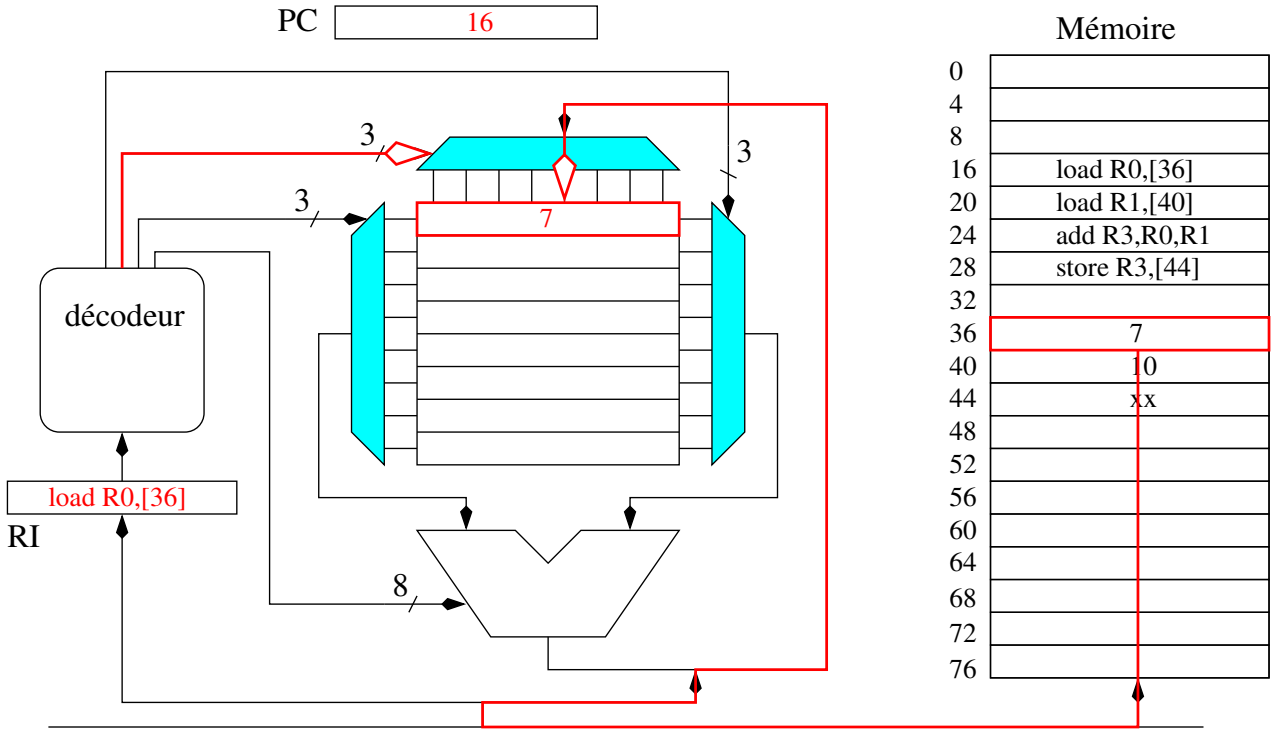

Program execution on a Processor (8 general purpose registers)

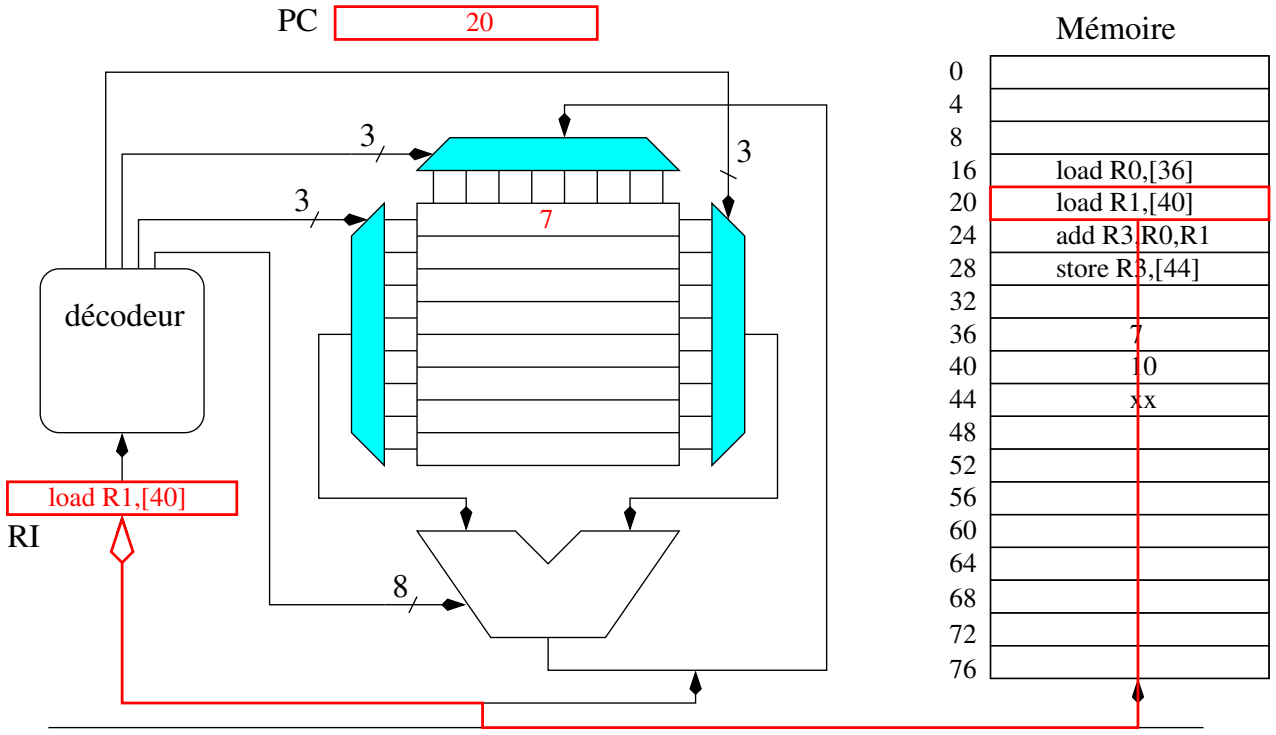

Program execution on a Processor (8 general purpose registers)

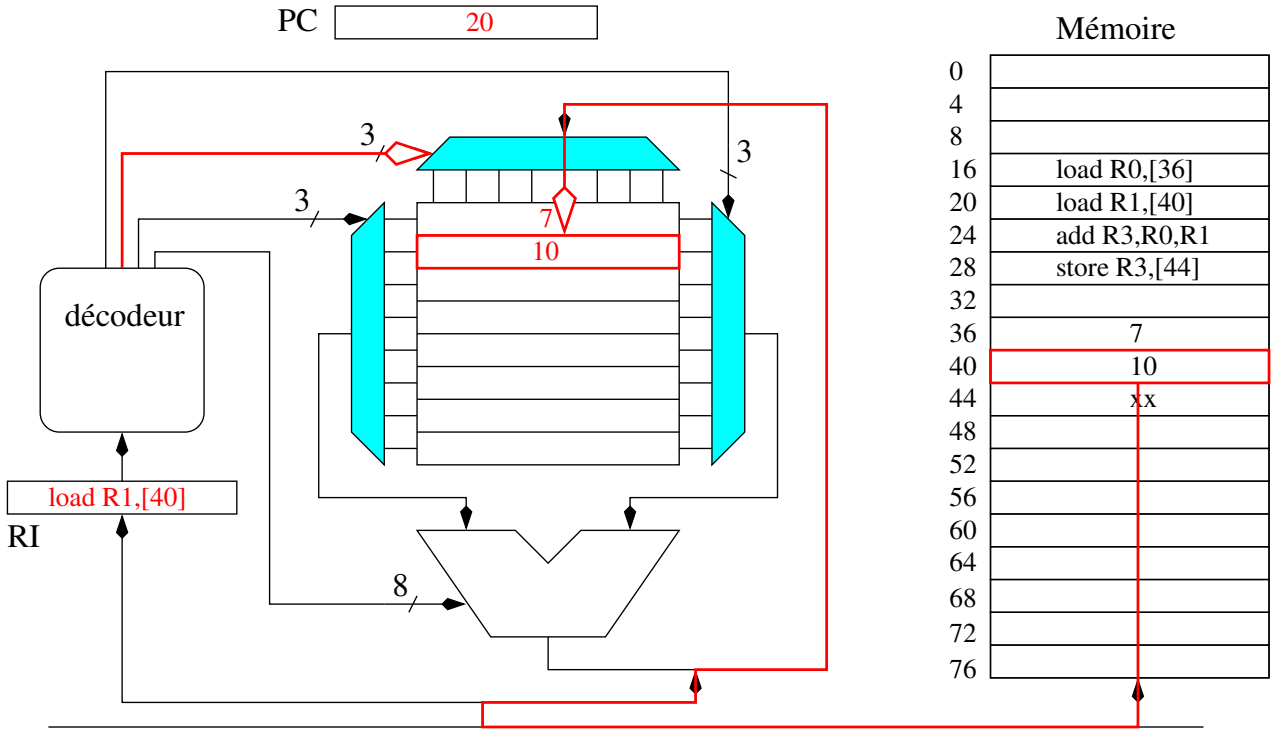

Program execution on a Processor (8 general purpose registers)

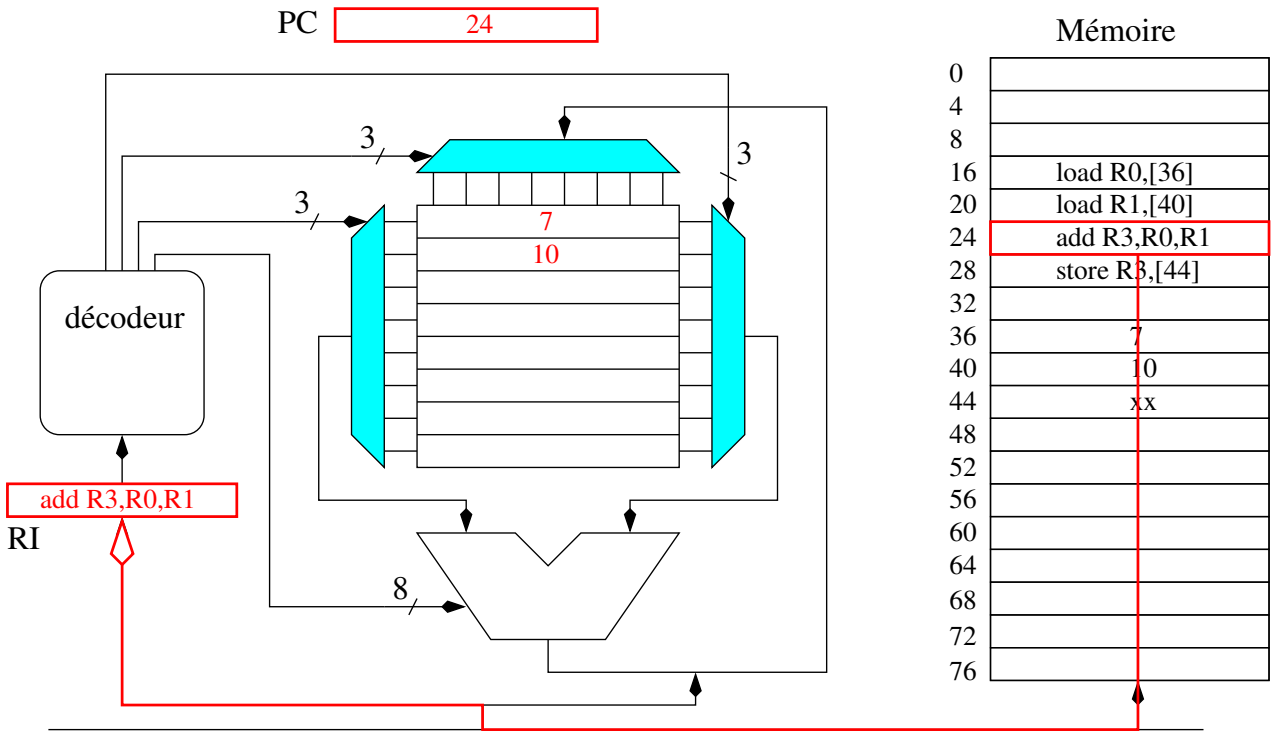

Program execution on a Processor (8 general purpose registers)

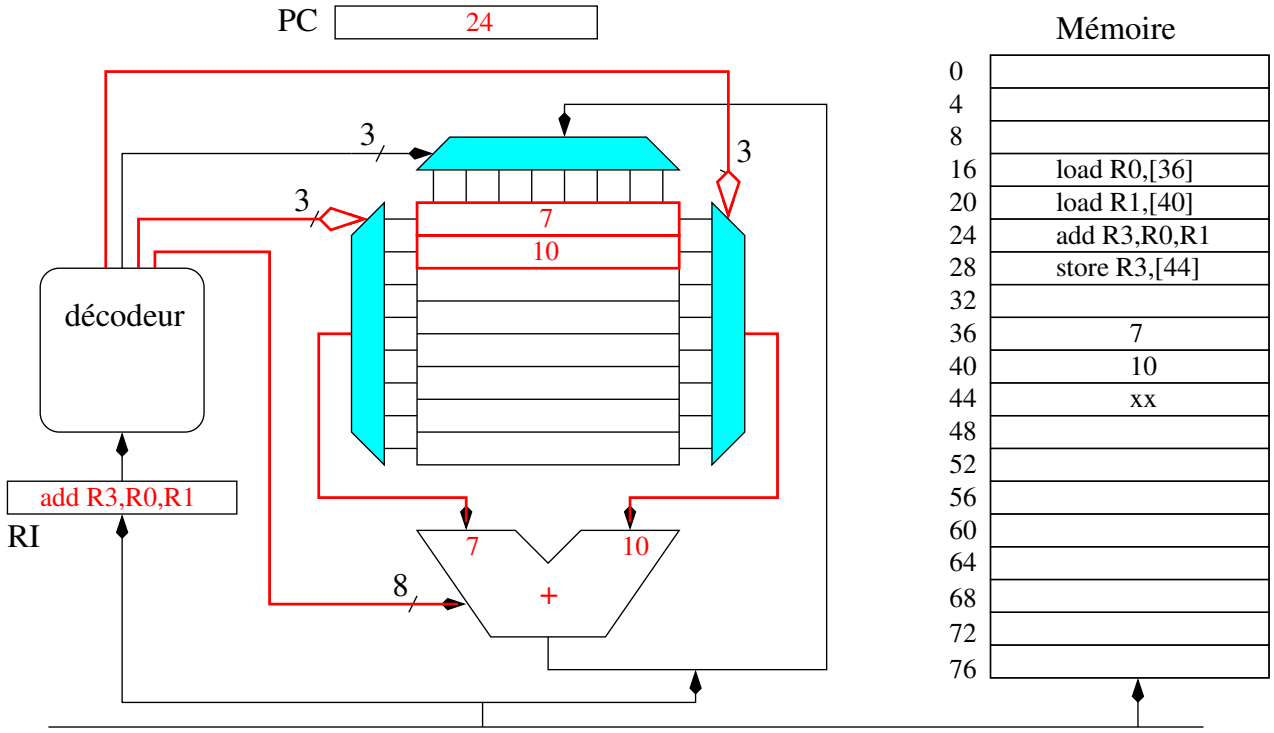

Program execution on a Processor (8 general purpose registers)

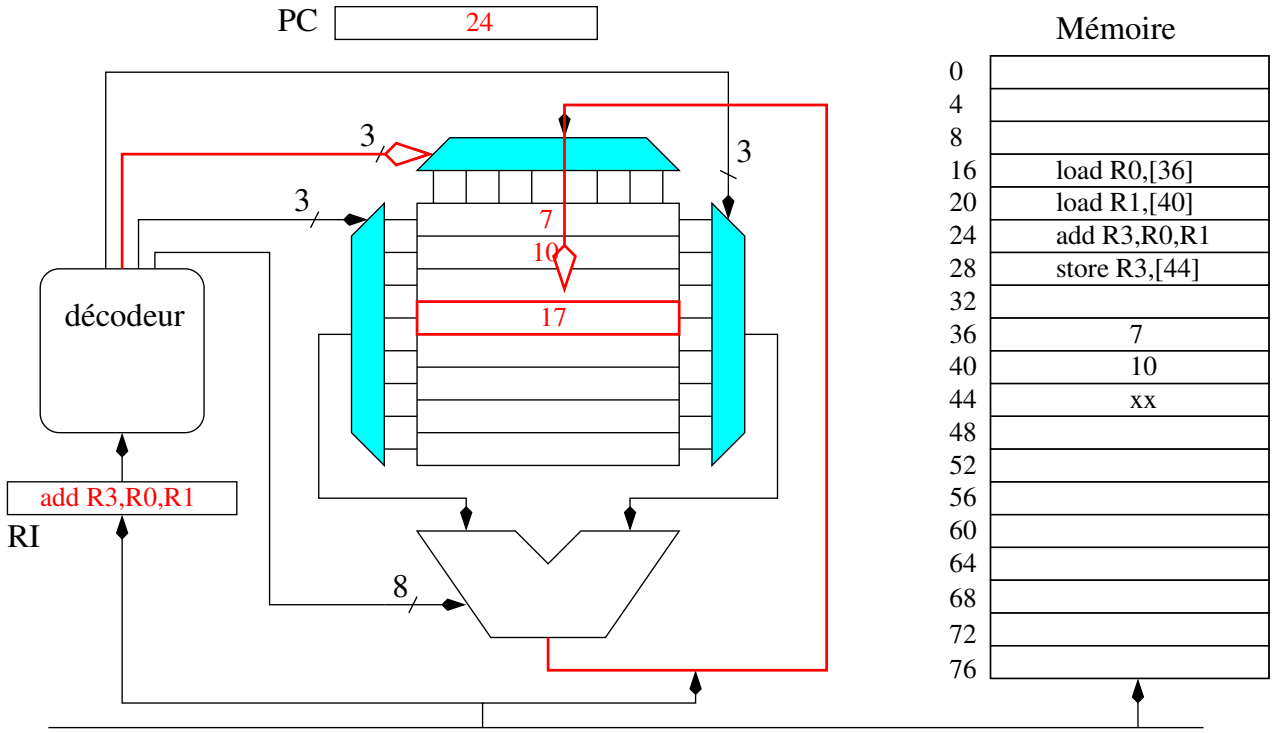

Program execution on a Processor (8 general purpose registers)

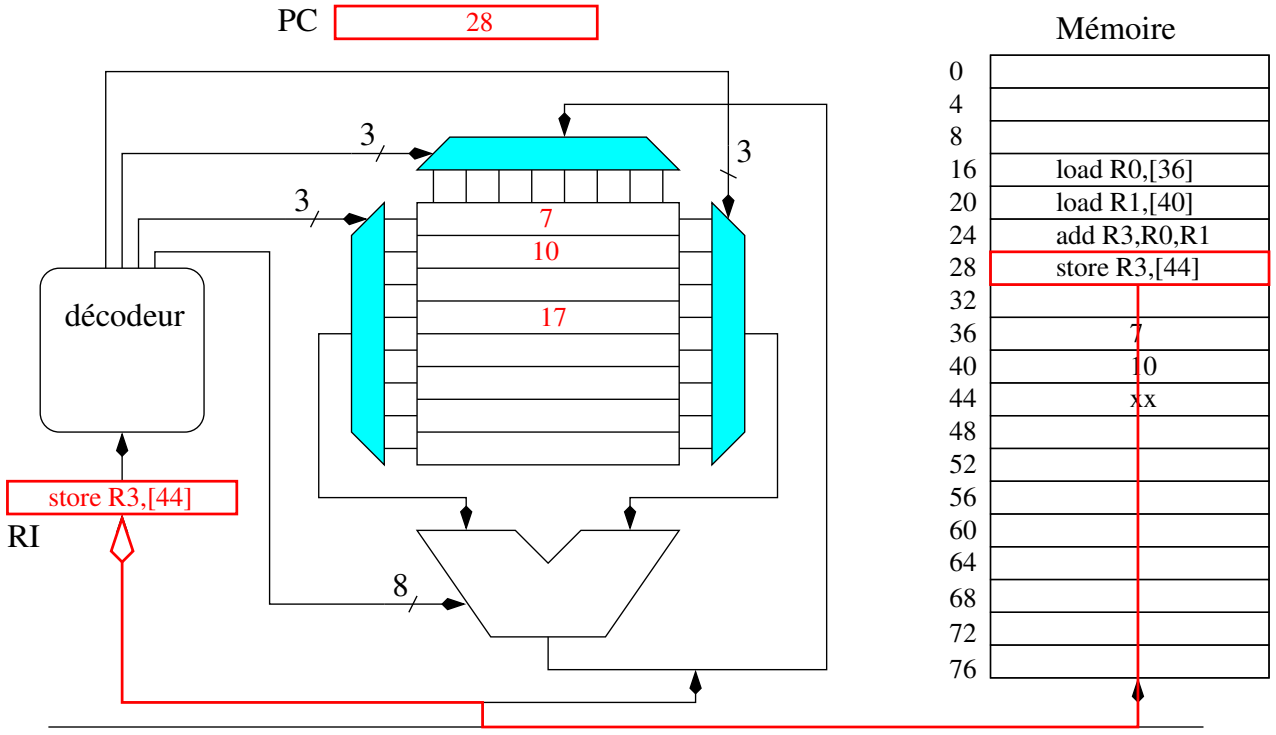

Program execution on a Processor (8 general purpose registers)

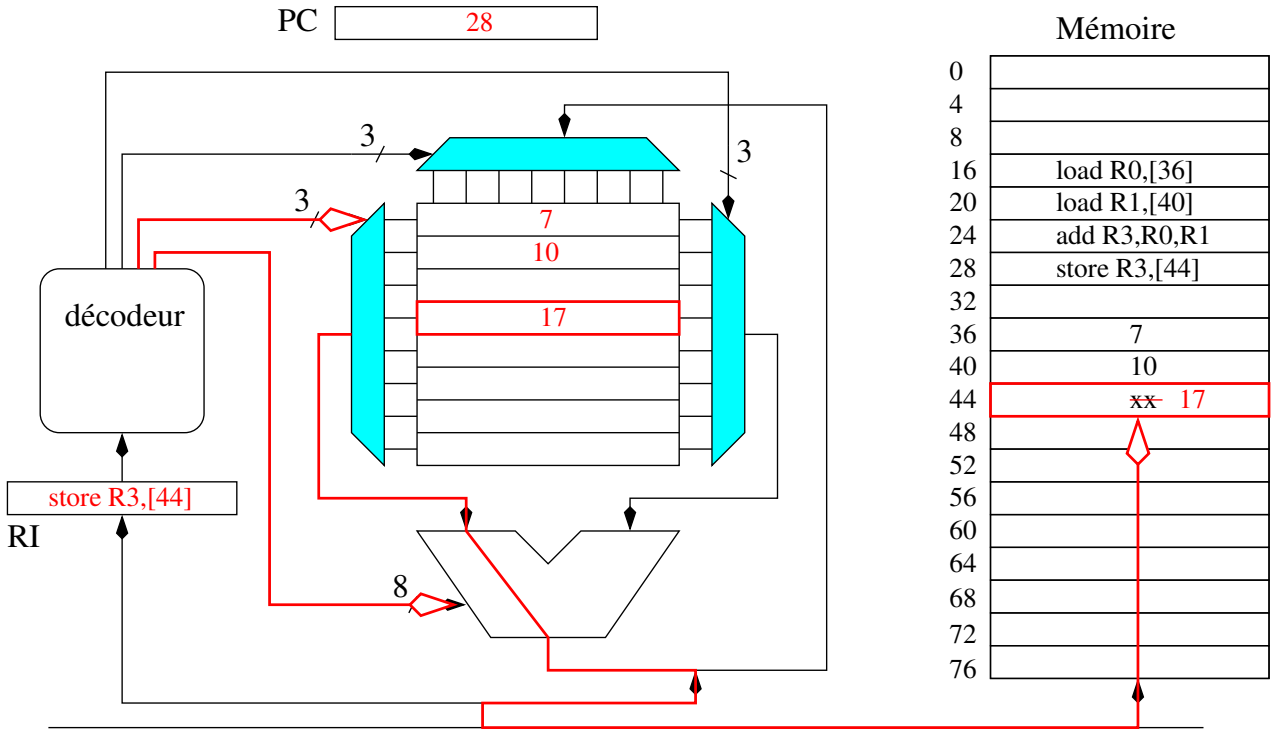

Computer Architecture in ARC

- Design of a simple dedicated circuit in logisim
- Study of a simple processor in logisim
- Overview of assembly code principles
- Compilation basics
- embedded system case study

Add on: two's complement representation

- Two's complement (*complément à deux*) is the most common representation for negative integers
- For a number on N bits:
  - Positive integers from 0 to $2^{N-1}-1$ are represented with usual binary encoding
  - Negative integer x from -2^{N-1} to -1 are represented by coding in binary the positive number $2^N - |x|$
  - Hence Negative integers always have the last (i.e. most significant) bit at 1, and positive always have the last bit at 0
- Example with $N = 3$
  - Integers between -4_{10} and 3_{10} can be represented
  - -1_{10} is represented as 111_2 $(2^3 - 1 = 7)$
  - -2_{10} is represented as 110_2 $(2^3 - 2 = 6)$
  - -4_{10} is represented as 100_2 $(2^3 - 4 = 4)$

Add on: two's complement representation (2)

- Two's complement have an important property: Addition "classical" algorithm works (except that the overflow should be ignored).
- Example:
  - $-1_{10} + (-2_{10}) = 111_2 + 110_2 = 1101_2 =$(ignoring the carry/overflow)$101_2 = -3$
  - $-1_{10} + 2_{10} = 111_2 + 010_2 = 1001_2 =$(ignoring the carry/overflow)$001_2 = 1$
- For $x > 0$, $x \leq 2^{N-1}$, The representation of $-x$ on N bit two's complement can be obtained by:
  - Complementing each bits of x
  - adding 1 to the resulting integer
- Example:
  - with $N = 3$ and $x = 3_{10} = 011_2$, complement of x is 100_2 adding 1 gives $101_2 = -3_{10}$
  - With N=8 and $x = 96_{10} = 01100000_2$ complement of x is 10011111, adding one is $-96_{10} = 10100000_2$, indeed $256 - 96 = 160 = 10100000_2$

Table of Contents

Automata

- Definition (Wikipedia): An automaton is a self-operating machine, or a machine or control mechanism designed to automatically follow a predetermined sequence of operations, or respond to predetermined instructions.
- In computer science:
  - Used in language theory to build compilers
  - Used in any technical domain: to describe predetermined behaviour.
  - Used in computer architecture: to design dedicated circuit.
  - A computer is a specific automaton.

Notion d'automate

- Un automate est une collection de K états numérotés de 0 à K-1, ainsi qu'une collection de transitions
- Un état particulier est l'état initial.
- Tous les états sont soit des états d'acceptation et soit des états de refus
- Les transitions, sont étiquetées
  1. soit par des actions (par exemple, je lis la lettre `x`)
  2. soit par des condition (par exemple, la lettre `x` est présente)
- le triplets (état 1, lettre `x`, état 2) signifie: si je suis dans l'état 1 et que je lis la lettre `x`, alors je vais dans l'état 2.

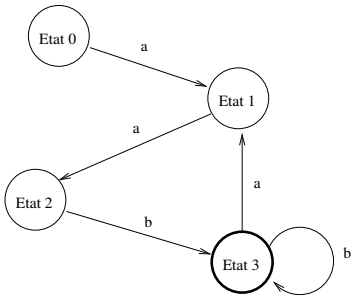

Notion d'automate

- Fonctionnement d'un automate
  - Initialisation de l'automate dans l'état
  - il lit les lettres du mot une par une
    - s'il trouve une transition possible, il l'exécute,
    - sinon il répond «le mot n'appartient pas au langage»;
  - si l'automate arrive à effectuer des transitions jusqu'à la dernière lettre du mot, il regarde alors dans quel état il termine:
    - si l'état appartient à la classe d'acceptation, l'automate répond «le mot appartient au » (on dit que le mot est reconnu),
    - sinon, il répond «le mot n'appartiennent pas au langage».

Notion de mot reconnu

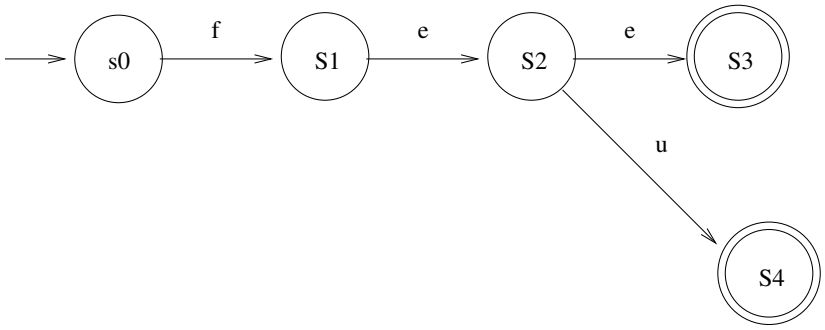

- `fee` → reconnu
- `feu` → reconnu
- `fei` → non reconnu (impossible de lire `'i'`)
- `fe` → non reconnu (arrêt dans un état non final)

Link with architecture: Computers are automata

- Every computing machine is an automata
- Computer are *universal* in the sense that the program gives much flexibility in the action performed.
- In fact the basic action of a computer is very repetitive:
  - Read the instruction at $PC in memory
  - decode the instruction
  - send the decoding to the ALU (or to memory if it is a load)
  - increment $PC
- Dedicated circuits (ASICs) are automata designed for specific tasks.

Table of Contents

Example from the poly

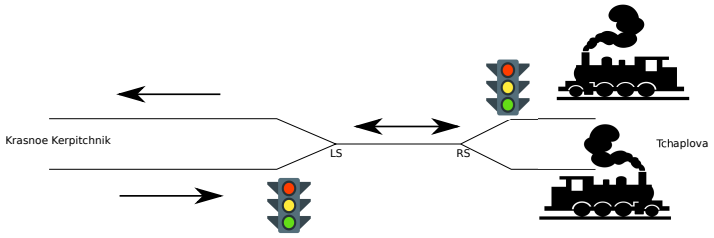

- A piece of unique train track for both train directions between the cities `T.` et `K.`
- Sensors triggered by train weight on rallways will command red lights when the track is used by a train.
- Modeling:
  - A booleen `A` (for 'Ampoule') indicating the state of the red light
  - Two booleans (LS for Left Sensor and RS for Rigth sensor) indicating the states of the sensors
  - An automaton to command the red lights

The Russian train automaton

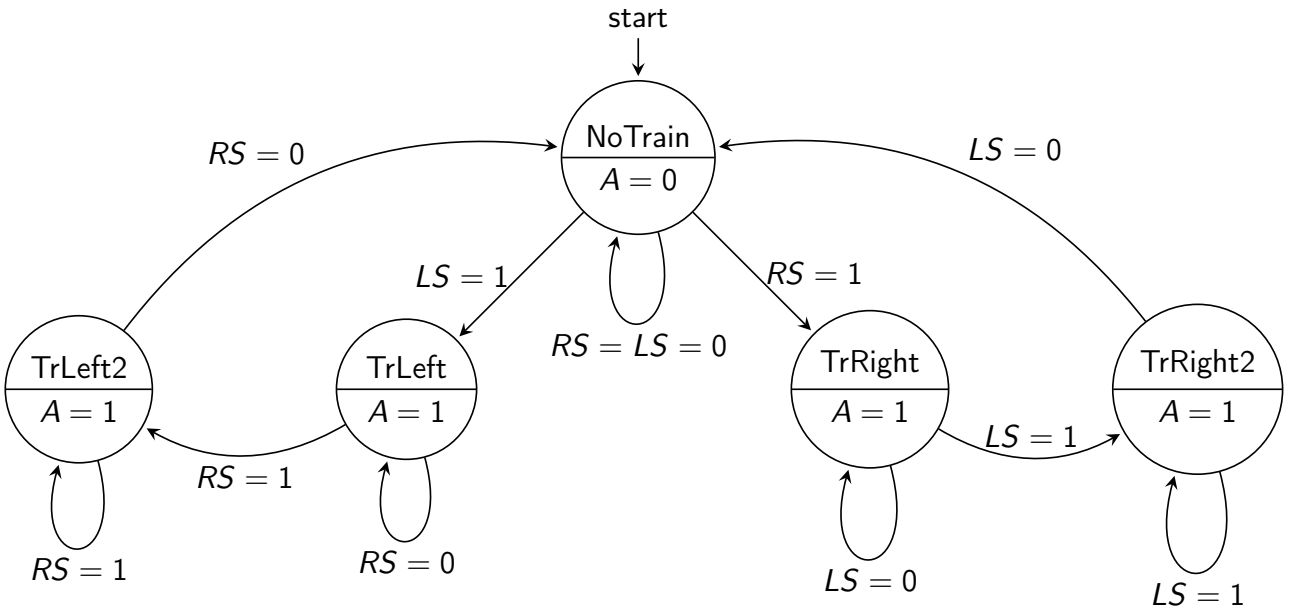

The Russian train automaton

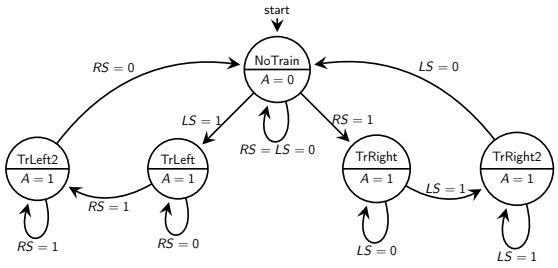

- Circles are *states* of the automaton (e.g. `NoTrain` state models the cases where no train stand on the track).
- States specifies output Values (here only one: `A`)
- Arrows are *transitions*, labeled by event that triggered them.

Back to the Russian train example

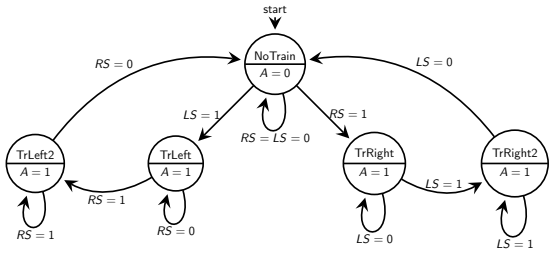

- The Inputs are `RS` and `LS` sensors Boolean values
- The Output is the value of Boolean `A`
- The functions (Transition and Output) can be defined by tables ⇒
- `X` means 'don't care'

s	x=(LS, RS)	s'=T(s,x)
NoTrain	00	NoTrain
NoTrain	01	TrRight
NoTrain	10	TrLeft
NoTrain	11	XXX
TrRight	0X	TrRight
TrRight	1X	TrRight2
TrRight2	1X	TrRight2
TrRight2	0X	NoTrain

s	y=F(s)
NoTrain	0
TrRight	1
TrRight2	1

Implementation of a synchronous automaton as a circuit

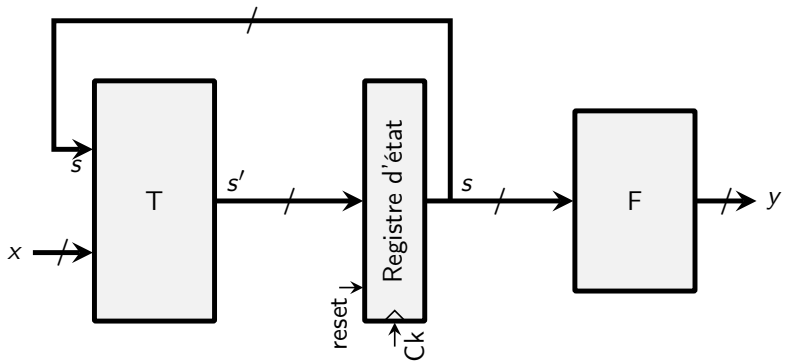

- `s` is current state, `s'` is next state, `x` are input bits, `y` are output bits.
- `Ck` and `reset` are not considered as inputs
- State change will occur on each rising edge of the Clock.

Implementation in Logisim

- We need to store 5 States, hence we need at least 3 bits:

value (binary)	state
100	NoTrain
000	TrRight1
001	TrRight2
010	TrLeft
011	TrLeft2

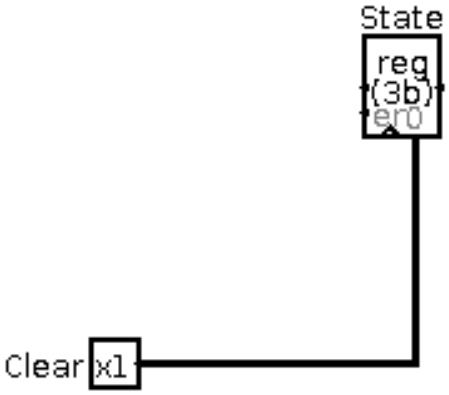

Russian train output function

- The output function is easy: A is on iff state is ''NoTrain'

s	y=F(s)
NoTrain	0
TrRight	1
TrRight2	1

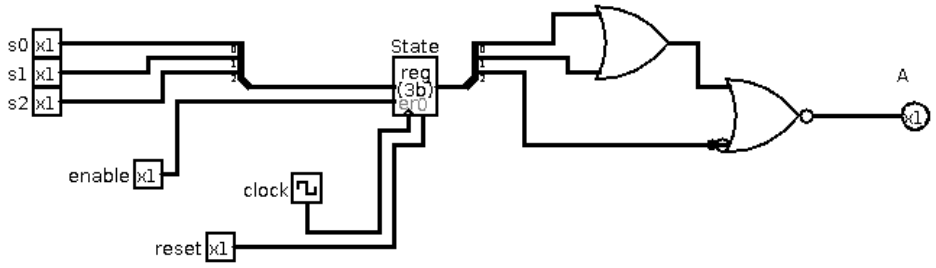

Russian train Transition function: more complicater

s	x=(LS, RS)	s'=T(s,x)
100 (NoTrain)	00	NoTrain
100 (NoTrain)	01	TrRight
100 (NoTrain)	10	TrLeft
100 (NoTrain)	11	XXX
000 (TrRight)	0X	TrRight
000 (TrRight)	1X	TrRight2
001 (TrRight2)	1X	TrRight2
001 (TrRight2)	0X	NoTrain
010 (TrLeft)	X0	TrLeft
010 (TrLeft)	X1	TrLeft2
011 (TrLeft2)	X1	TrLeft2
011 (TrLeft2)	X0	NoTrain

Table of Contents

Comming back to automata

- Automata are very widely used in computer science in different domains.
- In ARC we use them to *control the execution of dedicated synchronous circuits*
- As soon as a dedicated circuit is designed, there is an automaton designed.

Mealy and Moore automata

- We have seen a *Moore automaton*: output only depend on the state (not on the input), usually simpler to handle.
- The most general form of an automaton has a moore output and a mealy output

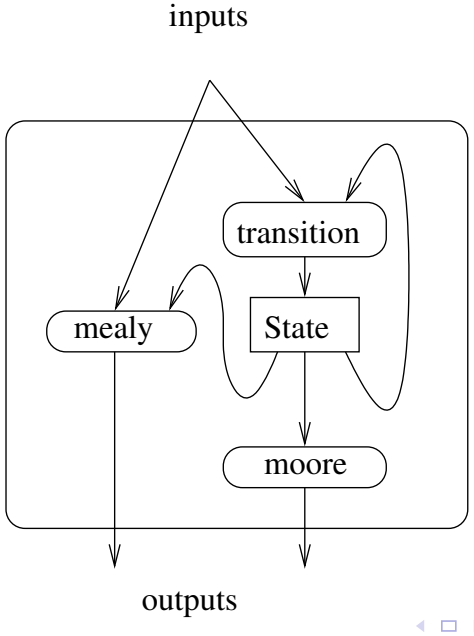

Summery: from Algorithm to Circuit

- From algorithm to automata (states and input/output)
- From automata to synchronous automata
- From synchronous automata to digital design

Lab topic: circuit for integer division

```
n := entrée N
p := entrée P
x := 0
q := 0
tant que x+p ⩽ n
    x := x+p
    q := q+1
fin tant que
sortie Q := q
```


Lab topic: proposed circuit to realize it

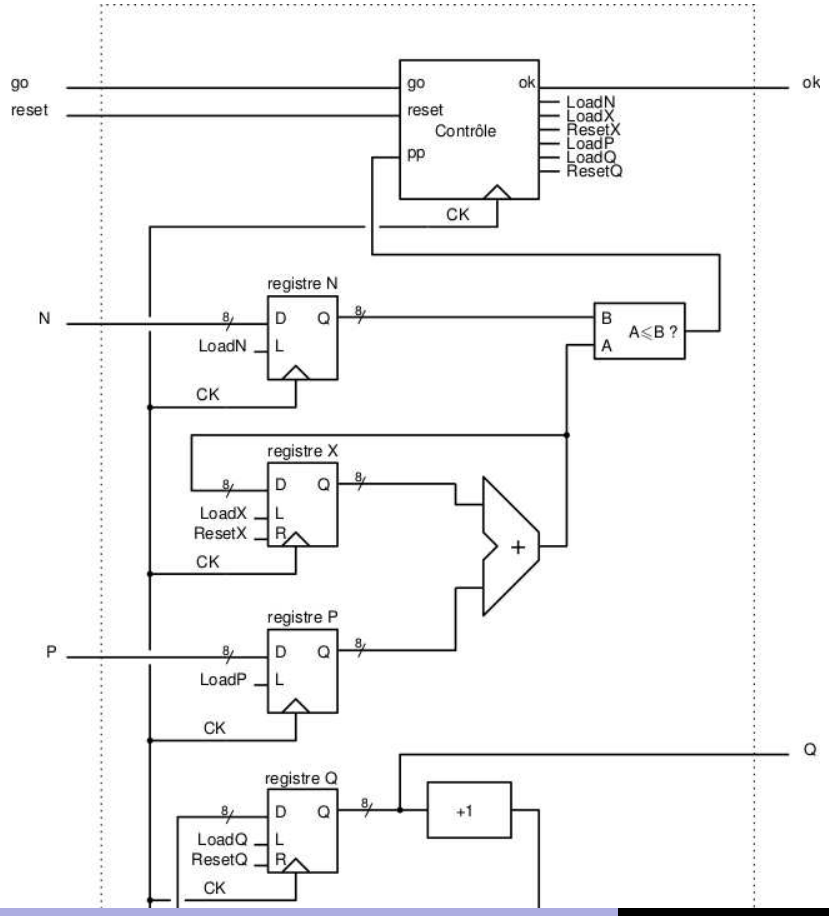

Table of Contents

Set Architecture Statement (ISA)

- The *instruction set* (Set Architecture statement: ISA) is of paramount importance
  - It determines the basic instructions executed by the CPU.
  - It's a balance between the hardware complexity of the CPU and the ability to express the required actions
  - It is represented in a symbolic way: the assembly code/language (ex: ADD R1,R2)
  - The tool that translates symbolic assembly code in binary code (i.e. machine code) is also called the **assembler**
- Two types of ISA:
  - CISC: Complex Instruction Set Computer
  - RISC: Reduce Instruction Set Computer

CISC: Complex Instruction Set Computer

- An instruction can code several elementary operations
  - Ex: a load, an add and a store (in memory operations)
  - Ex: computer a linear interpolation of several values in memory
- Need a mode complex hardware (specifically hardware accelerators)
- High variability in size and execution time for different instructions
- Produce a more compact code but more complex to generate
- Vax, Motorola 68000, Intel x86/Pentium

Example: instructions ISA of Pentium

JE EIP + displacement

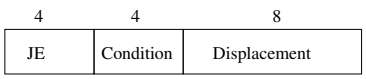

Call

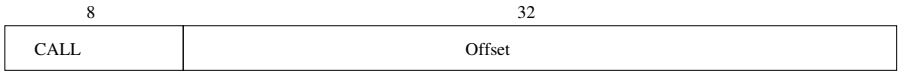

Mov $EBX, [EDI+displacement]

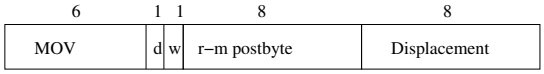

Push ESI

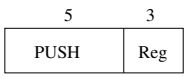

Add $EAX, Immediate

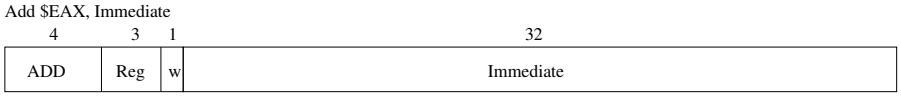

Test $EDX, Immediate

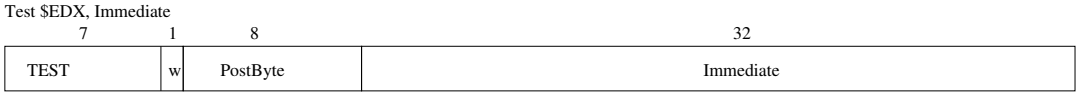

RISC: Reduced Instruction Set Computer

- Small simple instructions, all having the same size, and (almost) the same execution time.
- no complex instruction
- Clock speed increase with pipelining (between 3 and 7 pipeline stages)
- Code simpler to generate but less compact
- Every modern processor use this paradigm: SPARC, MIPS, ARM, PowerPC, etc.

Example: instructions of MSP430 ISA

1 operand instruction

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	1	0	0	opcode			B/W	Ad		Dest reg.			

relative Jumps

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	1	condition			PC offset (10 bits)									

2 operands instruction

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
opcode			Dest reg.				Ad	B/W	As			Dest reg.			

Examples:

- `PUSB.B R4`
- `JNE -56`
- `ADD.W R4,R4`

Exemple of Pentium ISA

- Write a simple C program `toto.c`
- Type `gcc -S toto.c` and get the `toto.s` file
- you can also use the *compiler explorer*: `https://gcc.godbolt.org/`
-

```
main() {
  int i=17;
  i=i+42;
  printf("%d\n", i);
}
```

⇒

```
(... instructions ...)
movl $17, -4(%rbp)
addl $42, -4(%rbp)
(...  printf params... )
call printf
(... instrucitons ...)
```


Disassemby

- compile the assembly code: `gcc toto.s -o toto`
- disassemble with `objdump`: `objdump -d toto`

```
Adresses  Instructions binaires    Assembleur
(...)
 40052c: 55                        push   %rbp
 40052d: 48 89 e5                  mov    %rsp,%rbp
 400530: 48 83 ec 10               sub    $0x10,%rsp
 400534: c7 45 fc 11 00 00 00      movl   $0x11,-0x4(%rbp)
 40053b: 83 45 fc 2a               addl   $0x2a,-0x4(%rbp)
 40053f: 8b 45 fc                  mov    -0x4(%rbp),%eax
 400542: 89 c6                     mov    %eax,%esi
% (...)
```

common properties of ISA

- An ISA first defines the types of data on which the processors can compute (32 bit memory addresses, integer of various sizes, etc.)
- Then it contains various types of instructions:
  - Computation instructions (add, sub, or, and, . . .), with various number of operands
  - Memory addessing instructions (load, store)
  - stack management instructions (push, pop)
  - Flow control instructions (jumps)
  - subroutine calls

Table of Contents

RISC-V history

- The RISC paradigm was invented 1980 David Patterson (UC Berkeley) and John Hennessy (Stanford U.).
- They described the MIPS architecture in their books "Computer Organization and Design" and "Computer Architecture: A Quantitative Approach."
- The MIPS was built by a commercial compagny (MIPS was in Nintendo 64, Sony PlayStation, PlayStation 2) and use in many architecture courses (including 3TC-ARC!).
- Hennessy and Patterson received the ACM A.M. Turing Award in 2017 (`https://amturing.acm.org/byyear.cfm`).

From 1980 to 2010, the development of the fifth generation of the RISC research project started and,led to the RISC-V (pronounced "risk-five").

RISC-V project

- RISC-V is an open instruction set architecture (ISA), this is fairly new!
- RISC-V International is a global nonprofit organization that owns and maintains the RISC-V ISA intellectual property.
- Its members range from individuals to organizations like Google, Intel, and Nvidia.

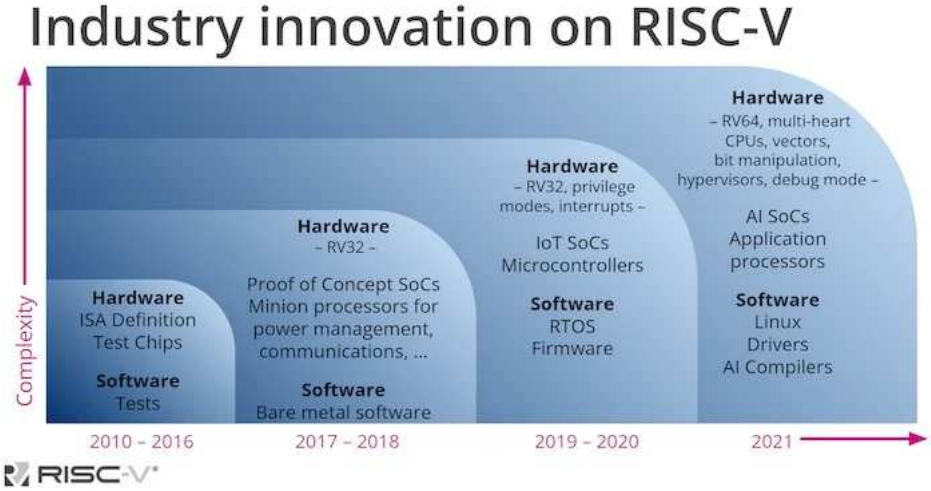

(from `https://www.allaboutcircuits.com/`)

RISC-V Processor

- RISC-V is specified by its ISA (RV32I, for integer 32 bits for instance).
- Many extension of the ISA are specified (32 bits, 64 bits, 128 bits, Atomic Instructions, Compressed Instructions, etc.)
- The architecture can be pipelined or not, it can target small embedded systems or large powerfull machines.
- Each processor compagny can build it own RISC-V implementation as long as it respect the ISA specification.
- We will study the RV32I base integer ISA that implements the necessary operations to achieve basic functionality with 32-bit integers.

RISC-V ISA basics (RV32)

- a *register-to-register* (or *load/store*) architecture
- RISC-V use 3-adress instructions (destination is the first operand)
- 32 32-bits registers (x0-x31) plus a A program counter (pc) of
- register can be name x0, x1, etc. or with their more explicit ABI names: x0 is "zero", x1 is ra (return adress), etc (`https://en.wikichip.org/wiki/risc-v/registers`)
- x0 is hardwired to value 0
- x1 is the return adress (ra)
- x2 is the stack pointer (sp)
- 32-bit address space: Addressable memory of 2^{32} bytes
  - $\Leftrightarrow 2^{30}$ words of 4 bytes

understanding RISC assembly

- From C to assembly:

```
riscv64-linux-gnu-gcc -S NtimesN.c -o NtimesN.S
```

NtimesN.c

```
[...]
scanf("%d",&N);
i = N*N + 3*N;
printf("i=%d\n",i);
[...]
```

NtimesN.s

```
[...]
call  __isoc99_scanf@plt #call to scanf
lw    a5,4(sp)           #N is now in a5
mulw  a2,a5,a5           #a2 <- N*N
sliiw a4,a5,1            #a4 <- 2*N (shift
addw  a5,a4,a5           #a4 <- 2*N+N
addw  a2,a2,a5           #a5 <- N*N + 3*N
lla   a1,.LC1            #printf args (To b
li    a0,1               #.. explained late
call  __printf_chk@plt   #call to printf
[...]
```


RISC-V register

- 32 registers in the *register file*
- Named
  - by their number: `x0 x1 ...x31`
  - or by their name `zero ra sp fp a0 a1 ...a7 s1 s2 ...`
- x0 (zero) contains value 0
- a0 ...a7 are used to pass arguments of a function call
- a0 a1 are used to transmit functions result
- t0 ...t6 and s0 ...s11 are working registers, used for CPU computations
- sp is the stack pointer
- fp is the frame pointer (explained later)
- ra contains the return address (after the end of current function)
- gp is a pointer to global area
- tp is the thread pointer

Risc-V assembly addressing mode

- The addressing mode defines how the operands of each instruction are interpreted.
- RISC-V has four addressing modes:
  - Immediate addressing: the operand is a constant within the instruction
  - Register addressing: where the operand represents a register.
  - Base addressing: the operand is an address which is the sum of a register and a constant (sometimes called indirect addressing)
  - PC-relative addressing: the operand is an address which is the sum of PC and a constant

Example of RISC-V adressing mode

- Register addressing
  - `add x1, x2, x3`
  - puts in `x1` the value of `x2` plus the value of `x3`.
  - `x1=x2+x3`
- Immediate addressing
  - `addi x1, x2, 0x0f`
  - `addi x1, x2, 15`
  - puts in `x1` the value of `x2` plus 15.
  - `x1=x2+15`
- Base addressing
  - `lw x1, 10(x3)`
  - puts in `x1` the value situated in memory at the address obtained by adding 10 to the content of `x3`.
  - `x1=Memory[x3+10]`
- `bne a1, a2, label`
  - branch to address of `label` if values in `a1` and `a2` are different.
  - `if (a1 != a2) then $PC=label`

Format of Risc-V instructions

- 3 types of format: R-Type, I-Types and B-Types
- R-types:

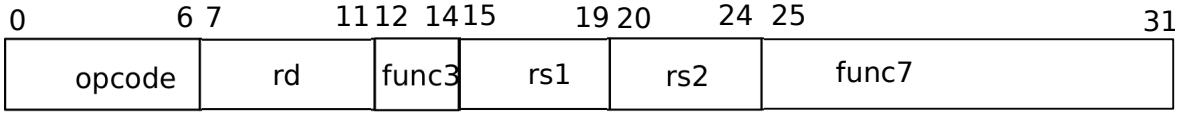

  - Used for 3-registers instructions
  - `opcode` is the operation code that specifies the operation
  - `rs1` and `rs2` are the first and second source register
  - `rd` is the destination register
  - `func3` and `func7` are used with `op` to select arithmetic operation (additionnal opcode fields)

I-Types instruction

- I-Types instruction are used for load, store, branch and immediate instruction.

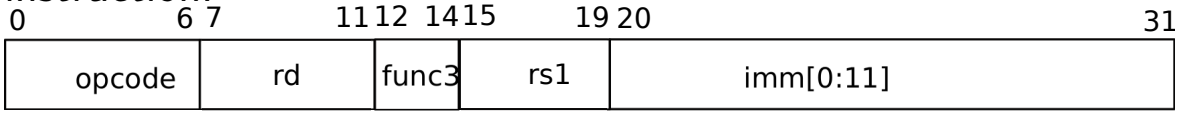

  - `rs1` is a source register
  - `rd` is a destination register
  - `func3` is additionnal opcode field
  - The `imm` field is a 16 bit's integer in two's-complement code , ranging from -32 768 to 32 767 (remind that this is a problem in many cases)

B-Types instruction

- B-Types instruction are used for Branch instructions

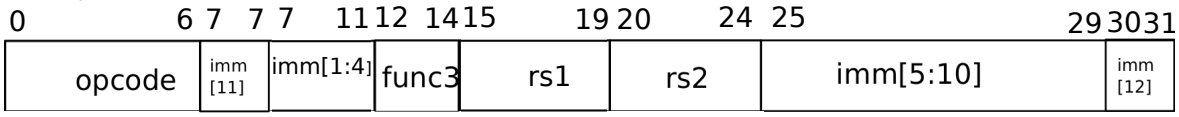

  - The `imm` split-field is a 13 bit's integer containing an address (always even hence bit 0 is implicitely 0).
  - can jump from address 0 to 2^{13}=1MB from $PC.
  - For longer jump, on can use others instrction (PC absolute), barely used.

Basic arithmetic and logic instruction

- R-Types instructions: add, sub and, or, xor
  - `add rd, rs1, rs2 // rd = rs1 + rs2`
  - `xor rd, rs1, rs2 // rd = rs1 ˆrs2`
- I-types for immediate operand operation:
  - `addi rd, rs1, 4 // rd = rs1 + 4`
  - `li rd, 4 // rd = 4`, pseudo (`addi rd, zero, 4`)

Load and store

- load and store operation use *indexed addressing*
  - the address operand specifies a signed constant and a register
  - These values are added to generate effective address
- byte instruction: `lb` and `sb` transfer one byte
  - `lb rd, 20(rs1)  // rd=Memory[rs1+20][0:7]`
  - `lw rd, 20(rs1)  // rd=Memory[rs1+20]` (i.e.[0:31]
  - `sb rd, 20(rs1)  // Memory[rd+20][0:7]=rs1`
  - `sb` stores only the lowest byte of operand register
- Word instruction: `lw` and `sw` operates on word (i.e. 32 bits)
- Remind that address have to be aligned to 32 bit world, hence must be multiple of 4.

Branches

- Conditional branch
  - `bne rs1, rs2, Label`
  - if `rs1` and `rs2` have different values, the next instruction to execute is at address Label (i.e. `pc = Label`
  - `beq rs1, rs2, Label  // same thing if rs1=rs2`
- Unconditionnal branch
  - `j offset // next instruction executed is at address PC+offset`
  - `jr rs1 // next instruction executed is at address contained in rs1`
- These are the only way of implementing loops in assembly:

```
[...]
        li s2, 1
while:  beq s1, zero, done
        sub s1, s1, s2
        j while
done:
```

```
t2=1
while (t1 != 0) {
   t1 = t1 - t2
}
```

Function control flow in RISC-V

- RISC-V uses the *jump-and-link* (`jal`) instruction to call functions
  - Example:

    ```
    jal ra, fact
    ```

  - saves the return address (i.e. the address of the following instruction) in the `ra` register and jumpt to the label `label` (code of `fact` function)
- At the end of the execution of `fact`, the instruction `ret` jumps back to the address stored in `ra` (pseudo: `jalr x0, ra, 0`)
- Arguments transmited to `Fact` are stored in registers `a0 ...a7`
- Return values of `Fact` are stored in registers `a0, a1`

Who save the register during Function call?

- When a function call occurs: `jal ra, fact`, who save the register?
  - The Caller (who knows which register he will use after the call)?
  - Or the callee (who knows which register he will use during its execution)?
- This convention is part of the *calling convetion* or ABI *application binary interface*.
- For MIPS:
  - `ra, t0 ...t6, a0 ...a7,` are caller saved
  - `fp, s1 ...s11` are callee saved

Function call example with MIPS

- Let says: function B calls function C
- Function B wants to save `t0`, `t1` and `a0` because it will need them after the return of C.
- this is done using the stack via the stack pointer `sp`

The Stack

- The stack is use to store all *local* information (in the sense local to the current function)
- That includes (say for function C):
  - local variable
  - Callee saved register if needed
  - Return address (i.e. the instruction following the `jal C` instruction).
  - (sometimes) the parameters passed to C
  - (sometimes) the result of C
  - In many ISA, the parameters and the results are passed through dedicated registers
- All these data constitute the frame of the fonction instance.
- the frame pointeur points to the frame of the current function
- For MIPS, the frame pointer is `fp`

Function B calls C

```
B   ...                 beguinning of  B
    ...
    sw t0,0(sp)      saving  t0 in stack
    sw t1,-4(sp)     saving  t1 in stack
    sw a0,-8(sp)     saving  a0 in stack
    sub sp,sp,12     correct stack pointer
    jal ra, C        call to C function
    lw a0,4(sp)      restoring return addresse of B from stack
    lw t1,8(sp)      restoring s1 from stack
    sw t0,12(sp)     restoring  s0
    add sp,sp,12     adjusst  stack pointeur value
    ...
    ret             end of B
    ...
```


Sketching code of C function

```
C:
    addi    sp,sp,-40     # C need 40 Bytes for its frame
    sw      ra,32(sp)    # store return address (inst. in B)
    sw      fp,28(sp)    # store frame pointer
    sw      s0,24(sp)    # store s0 (because C uses it)
    move    fp,sp        # fp <- sp: frame pointer of C set
       ....
      ....
    lw      ra,32(sp)    # ra <- return address (in B)
    lw      fp,28(sp)    # fp  <- frame pointeur of B
    lw      s0,24(sp)    # restore s0
    addi    sp,sp,40     # sp <- sp+40, restore B stack pointe
    ret                  # return to ra (B function)
```

Table of Contents

Procedure abstraction

- Let's pause a while to come back to high level langage
- What is a function (or a procedure)?
- How its isolation mecanisme (local variable) is implemented?
- This is implemented with a very fundamental mecanism: the Stack and the Activation Record (or Frame) of each procedure.

Notion of procedure

- Procedures (or functions) are the basic units for compilers
- Three important abstraction:
  - Control abstraction: parameter passing and result transmission
  - Memory abstraction: variable lifetime (local variables)
  - Interface: procedure's signature

Procedure Control Transfer

- Transfer mechanism of control between procedures:
  - when calling a procedure, the control is given to the procedure called;
  - when this called procedure ends, the control is returned to the calling procedure.
  - Two calls to the same procedure create two em independent instances (or invocations).
- two useful graphic representations:
  - The call graph: represents the information written in the program.
  - The call tree: represents a particular execution.

Call Graph

```
procedure calc;
begin { calc}
    ...
end;
procedure call1;
    var y...
    procedure call2
        var z: ...
        procedure call3;
            var y....
            begin { call3}
                x:=...
                calc;
            end;
        begin { call2}
            z:=1;
            calc;
            call3;
        end;
    begin { call1}
        call2;
        ...
    end;
```

Call Graph:

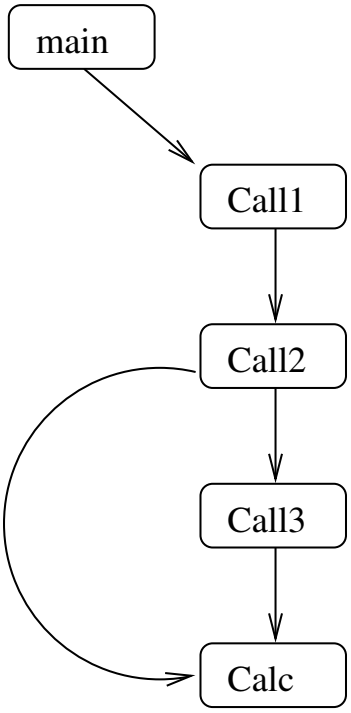

Call Tree

```
procedure calc;
begin { calc}
    ...
end;
procedure call1;
    var y...
    procedure call2
        var z: ...
        procedure call3;
            var y....
            begin { call3}
                x:=...
                calc;
            end;
        begin { call2}
            z:=1;
            calc;
            call3;
        end;
    begin { call1}
        call2;
    end;
```

Call tree for one particular execution:

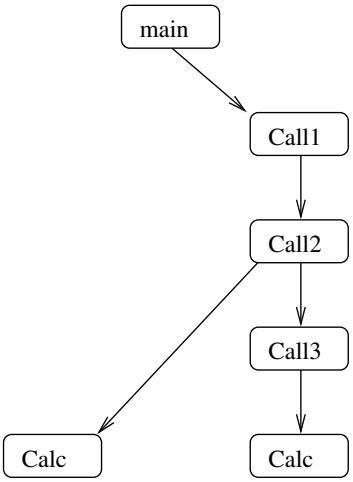

main calls $call_1$
$call_1$ calls $call_2$
$call_2$ calls *calc*
calc returns to $call_2$
$call_2$ calls $call_3$
$call_3$ calls*calc*
calc returns to $call_3$
$call_3$ returns to $call_2$
$call_2$ returns to $call_1$
$call_1$ returns to *main*

Execution Stack

- The transfer of control mechanism between procedures is implemented thanks to the *execution stack*.
- The programmer has this vision of virtual memory:

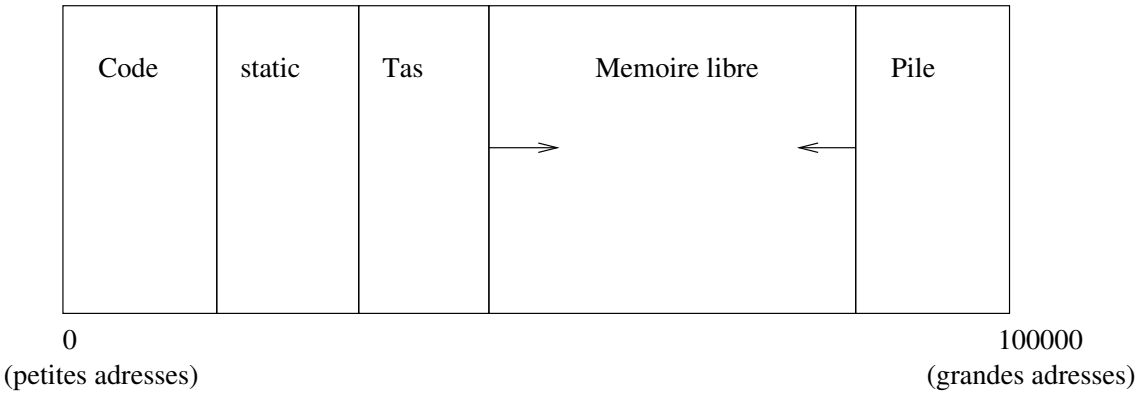

- The *heap* is used for dynamic allocation.
- The *stack* is used for the management of contexts of procedures (local variable, etc.)

Function call: status of the stack

Before the call
(AR=Activation Record)

after the call

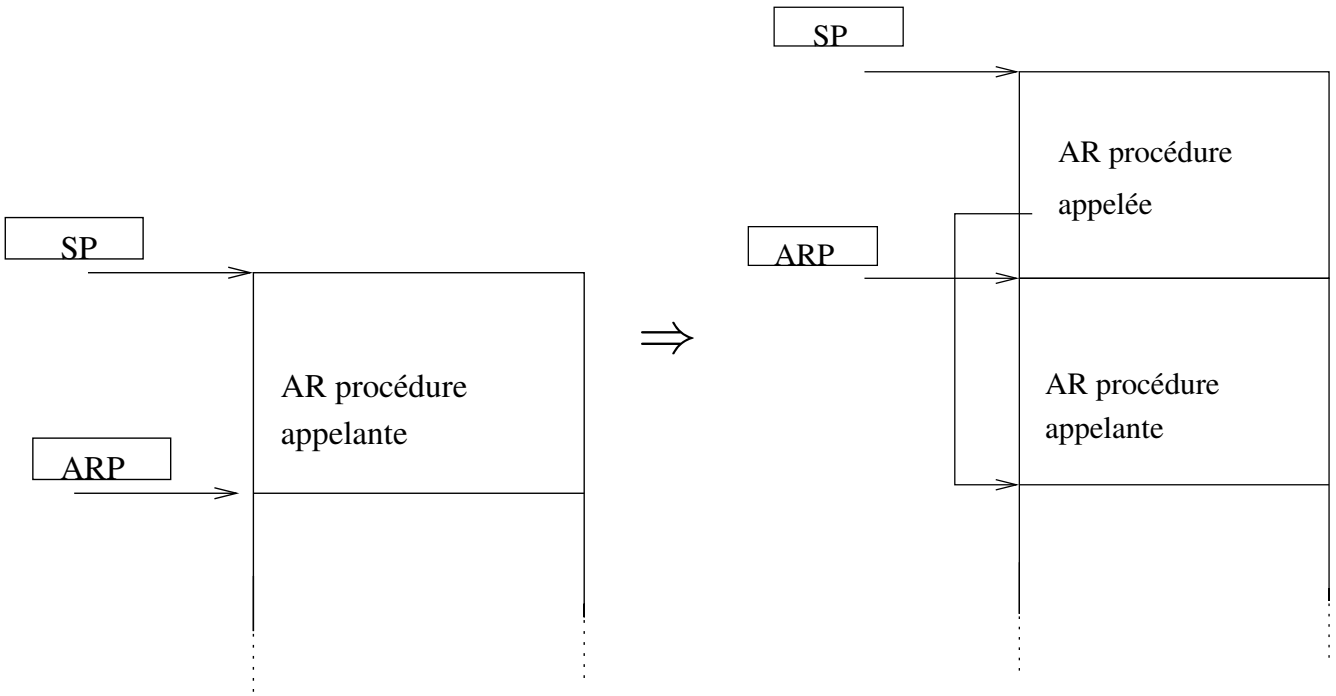

Activation record

- Calling a procedure: Stacking the *activation record* (or *frame*).
- Need of a dedicated pointer for that: the *activation record pointer* (ARP) or *frame pointeur* (`fp`))
- The frame allows to set up the *context* of the procedure.
- This frame contains
  - The space for local variables declared in the procedure
  - Information for restoring the context of the calling procedure:
    - Pointer to the frame of the calling procedure (ARP or FP for em frame pointer).
    - Address of the return instruction (statement following the call of the appellant proceedings).
    - Eventually save the state of the registers at the time of the call.

Content of the Frame

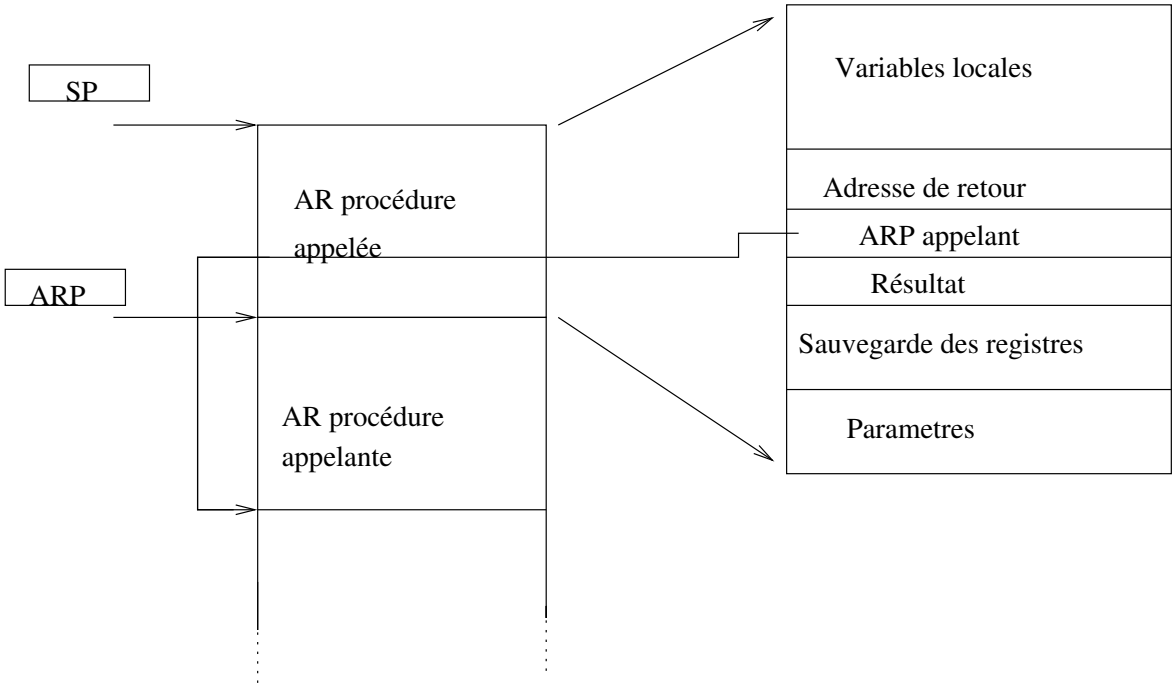

Return to calling function

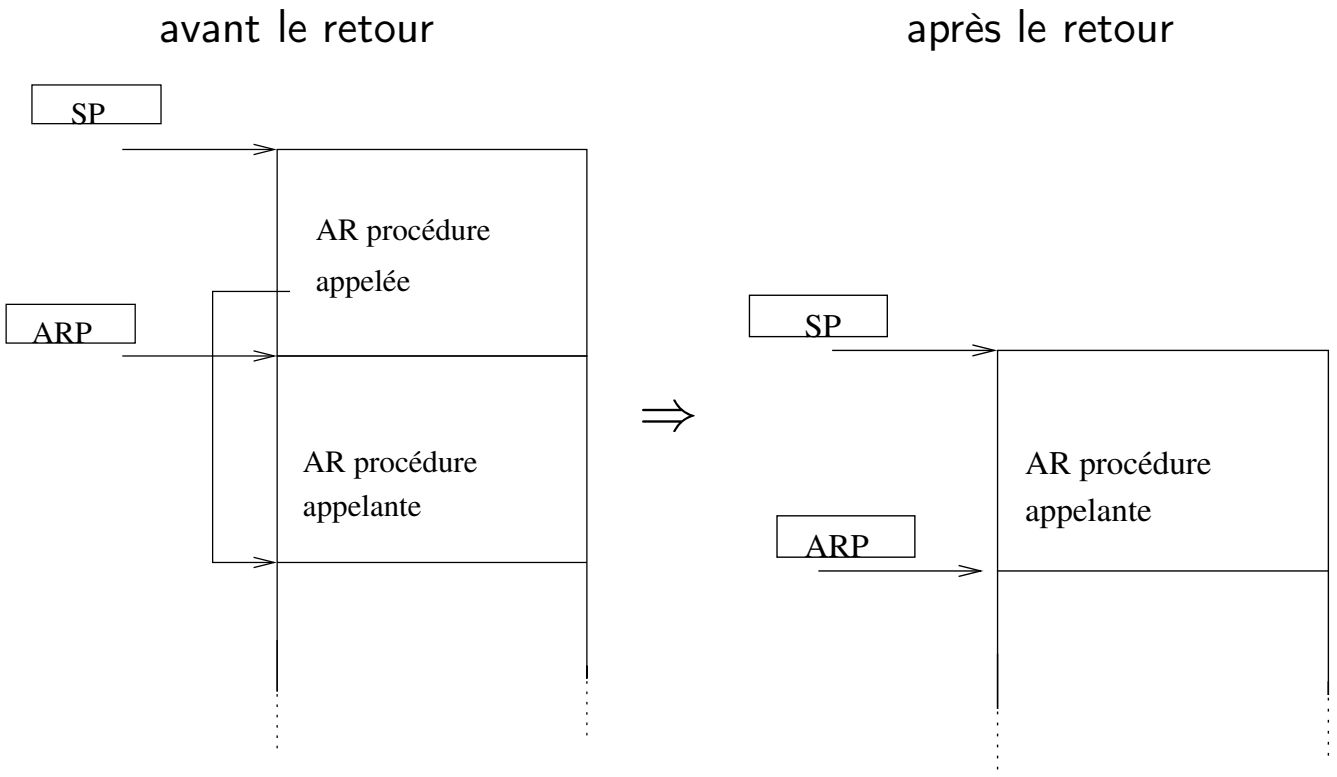

Table of Contents

Coming back to previous call example with B and C

- Let says: function B calls function C
- Function B wants to save `t0`, `t1` and `a0` because it will need them after the return of C.
- this is done using the stack via the stack pointer `sp`

The Stack

- The stack is use to store all *local* information (in the sense local to the current function)
- That includes (say for function C):
  - local variable
  - Callee saved register if needed
  - (sometimes) Return address (i.e. the instruction following the `jal ra, C` instruction).
  - (sometimes) the parameters passed to C
  - (sometimes) the result of C
  - In many ISA, the parameters and the results are passed through dedicated registers
- All these data constitute the frame of the fonction instance.
- the frame pointeur points to the frame of the current function
- For RISC-V, the frame pointer is `fp`

Function B calls C

```
B   ...                 beguinning of  B
    ...
    sw t0,0(sp)      saving  t0 in stack
    sw t1,-4(sp)     saving  t1 in stack
    sw a0,-8(sp)     saving  a0 in stack
    sub sp,sp,12     correct stack pointer
    jal ra, C        call to C function
    lw a0,4(sp)      restoring return addresse of B from stack
    lw t1,8(sp)      restoring s1 from stack
    sw t0,12(sp)     restoring  s0
    add sp,sp,12     adjusst  stack pointeur value
    ...
    ret             end of B
    ...
```


Sketching code of C function

```
C:
    addi    sp,sp,-40     # C need 40 Bytes for its frame
    sw      ra,32(sp)    # store return address (inst. in B)
    sw      fp,28(sp)    # store frame pointer
    sw      s0,24(sp)    # store s0 (because C uses it)
    move    fp,sp        # fp <- sp: frame pointer of C set
       ....
      ....
    lw      ra,32(sp)    # ra <- return address (in B)
    lw      fp,28(sp)    # fp  <- frame pointeur of B
    lw      s0,24(sp)    # restore s0
    addi    sp,sp,40     # sp <- sp+40, restore B stack pointe
    ret                  # return to ra (B function)
```

RISC-V Assembly for programme fib

Fibbonacci suite program:

```
int fib (int i)
{
  if (i<=1) return(1);
  else return(fib(i-1)+fib(i-2));
}

int main (int argc, char *argv[])
{
  fib(2);
}
```


Assembleur RISC-V pour programme fib

```
fib:
    addi    sp,sp,-48  # SP <- SP-48 :AR de 48 octet (12 mots)
    sd    ra,40(sp)    # stocke adresse retour (64 bits) a SP+40
    sd    s0,32(sp)    # sauvegarde registre s0
    sd    s1,24(sp)    # sauvegarde registre s1
    addi    s0,sp,48   # s0=ARP/FP <- SP
    mv    a5,a0        # a5 <- arg1 (N)
    sw    a5,-36(s0)   # stock arg1 (N) dans la pile (SP-12)
    lw    a5,-36(s0)   # instruction inutile (supprimée si optimisation)
    sext.w    a4,a5    # a4 <- sign extension a5(32)
    li    a5,1         # a5 <- 1
    bgt    a4,a5,.L2   # if (a4 > 1)  sauter a .L2)
    li    a5,1         # ici on a arg1=N<=1 donc a5 <- res=1
    j    .L3           # sauter à .L3
.L2:
    lw      a5,-36(s0)   # ici N>1, a5 <- N
    addiw   a5,a5,-1     # a5 <- a5 - 1
    sext.w  a5,a5        # sign extension
    mv      a0,a5        # a0 <- a5 (set arg in a0 for recursive call)
    call    fib          # recursive call
    mv      a5,a0        # a5 <- result from recursive call
    mv      s1,a5        # s1 <- a5
    lw      a5,-36(s0)   # a5 <- N
    addiw   a5,a5,-2     # a5 <- a5 -2
    sext.w  a5,a5        # sign extension
    mv      a0,a5        # a0 <- a5 (set arg in a0 for recursive call)
    call    fib          # recursive call
    mv      a5,a0        # a5 <- result from recursive call
    addw    a5,s1,a5     # a5 <- fib(N-1)+fib(N-1)
    sext.w  a5,a5        # sign extension
```

Assembleur RISC-V pour programme `fib`

```
.L3:
    mv        a0,a5         #a0 <- a5 (set result in a0)
    ld        ra,40(sp)     # restaure ra
    ld        s0,32(sp)     # restaure s0
    ld        s1,24(sp)     # resaure s1
    addi      sp,sp,48      # restaure sp
    jr        ra            # return
    .size     fib, .-fib
    .section  .rodata
    .align    3
.LC0:
    .string    "le resultat est %d "
    .text
    .align    2
    .globl    main
    .type     main, @function
main:
    addi      sp,sp,-32   # set AR for main
    sd        ra,24(sp)
    sd        s0,16(sp)
    addi      s0,sp,32
    mv        a5,a0       #store arg og main
    sd        a1,-32(s0)
    sw        a5,-20(s0)
    li        a0,2        # we call fig(2)
    call      fib
    mv        a5,a0       # get fib result
    mv        a1,a5       #set args for printf
    lla       a0,.LC0
    call      printf@plt
    li        a5,0
```


example 1, if-then

```
        bne s0, s1, Test
        add s2, s0, s1
Test:
```

example 2, if-then-else

```
      beq s4, s5, Lab1
      add s6, s4, s5
      j Lab2
Lab1: sub s6, s4, s5
Lab2:
```


example 3, looping

```
      li t2, 0
      li t3, 1
while: beq t1, zero, done
      add t2, t1, t2
      sub t1, t1, t3
      j while
done:
```

example 4, static variable

```
.globl main
.type main, @function
        .data
var1:   .word  23          # declare storage for var1; initial
                           # value is 23

        .text
main:
        lw t0, var1          # load contents of RAM location int
                             # register $t0:  t0 = [var1] ( = 23
        li t1, 5             #  $t1 = 5   ("load immediate")
        la t2, var1          # load address of var1
        sw t1, (t2)          # store contents of register t1
                             #into RAM:  [var1] = t1 = 5
done:
        jr   ra
```


example 5, array accesses

```
        .globl main
        .type main, @function
        .data
array1: .space 12          #  declare 12 bytes of storage to
                           # hold array of 3 integers
        .text
main:   la t0, array1      #  load base address of array into
                           #register $t0
        li t1, 5           #t1 = 5   ("load immediate")
        sw t1, (t0)        #first array element set to 5;
                           #indirect addressing
        li t1, 13          #t1 = 13
        sw t1, 4(t0)       #second array element set to 13
        li t1, -7          #t1 = -7
        sw t1, 8(t0)       #third array element set to -7
        jr   ra
```

Documentation on RISV-V assembly

- The RISC-V Instruction Set Manual Volume I: User-Level ISA

  https://github.com/riscv/riscv-isa-manual/releases/download/Ratified-IMAFDQC/riscv-spec-20191213.pdf
- Risc-V assembly manual on github

  https://github.com/riscv-non-isa/riscv-asm-manual/blob/master/riscv-asm.md
- CheatSheet on ARC Moodle site.

Table of Contents

Program execution on a Processor (8 general purpose registers)

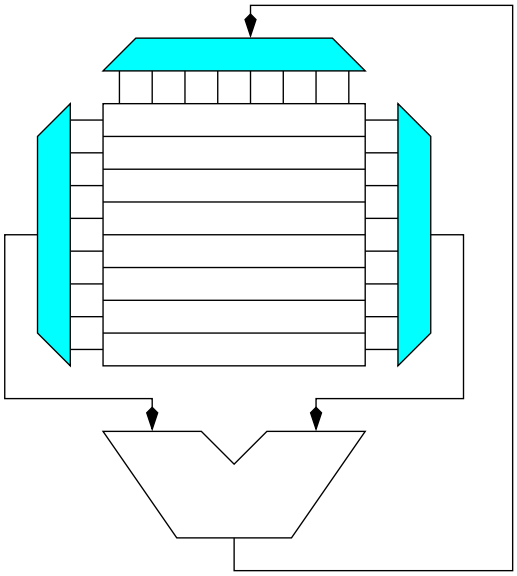

Program execution on a Processor (8 general purpose registers)

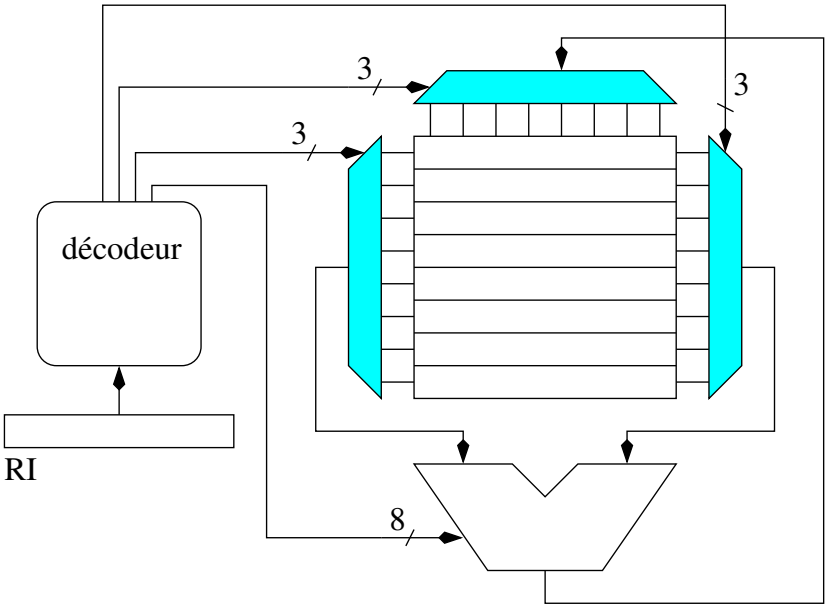

Program execution on a Processor (8 general purpose registers)

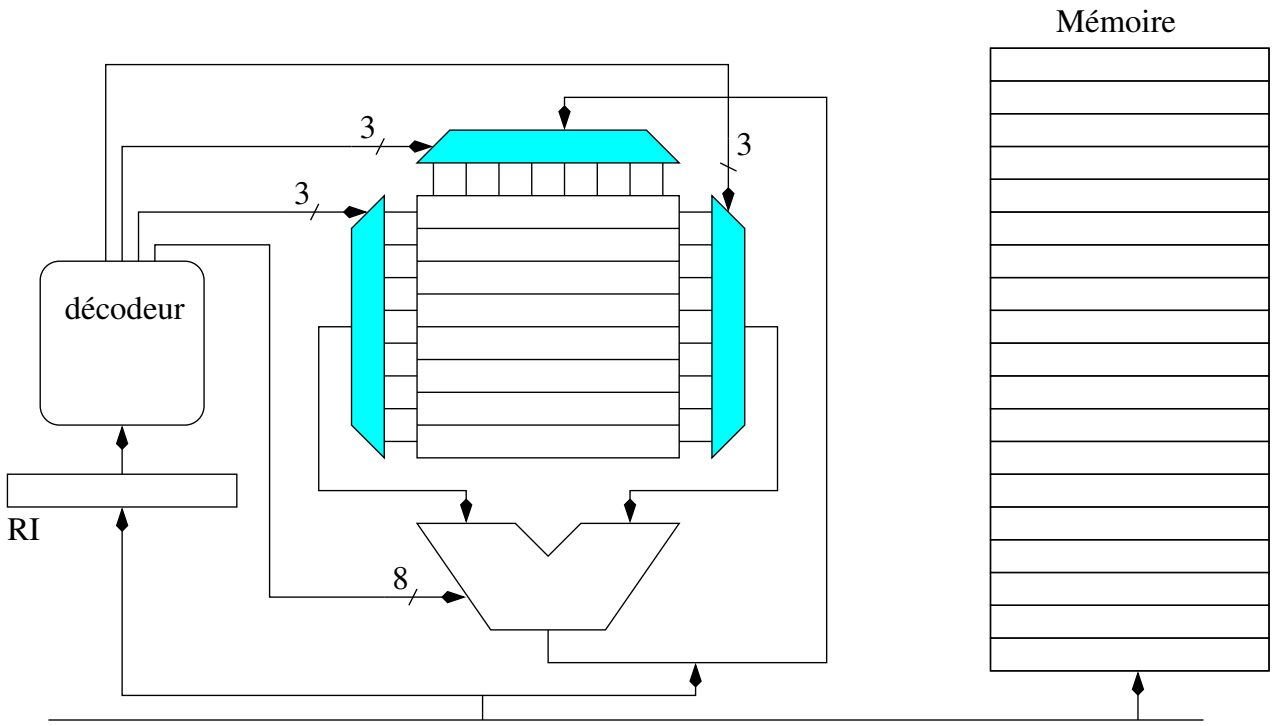

Program execution on a Processor (8 general purpose registers)

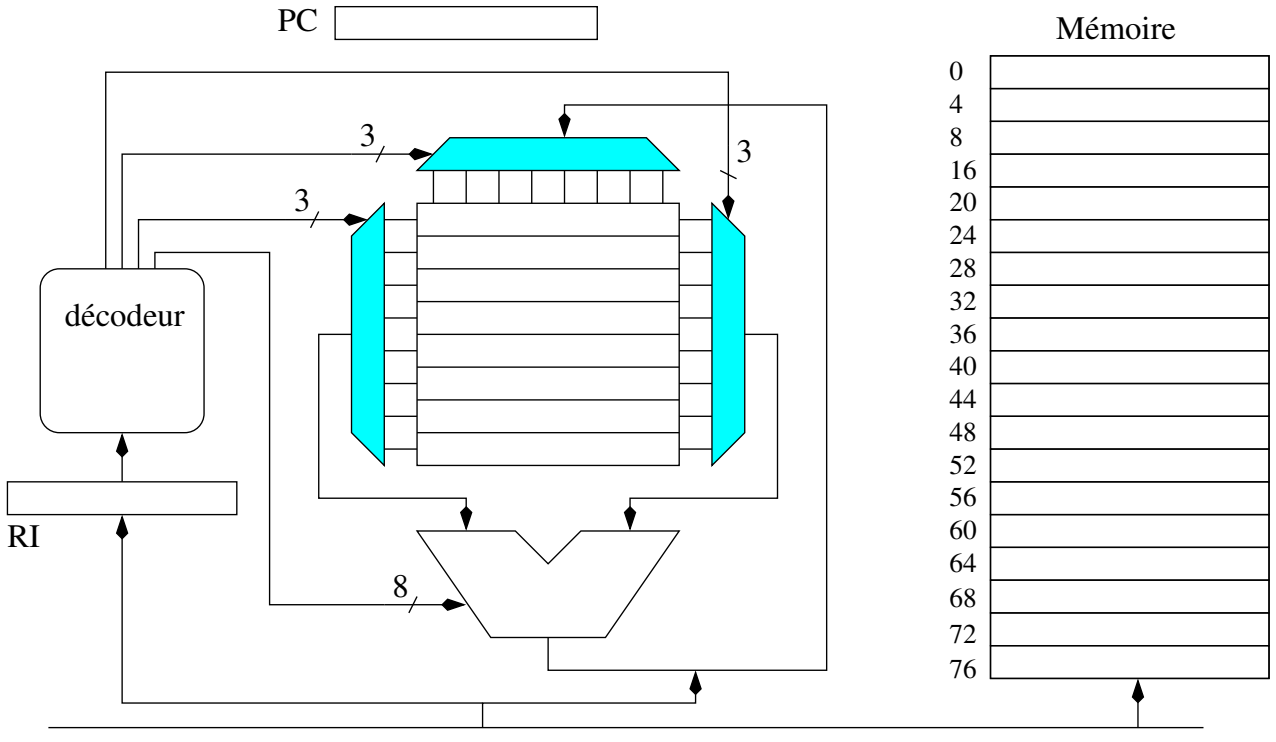

Program execution on a Processor (8 general purpose registers)

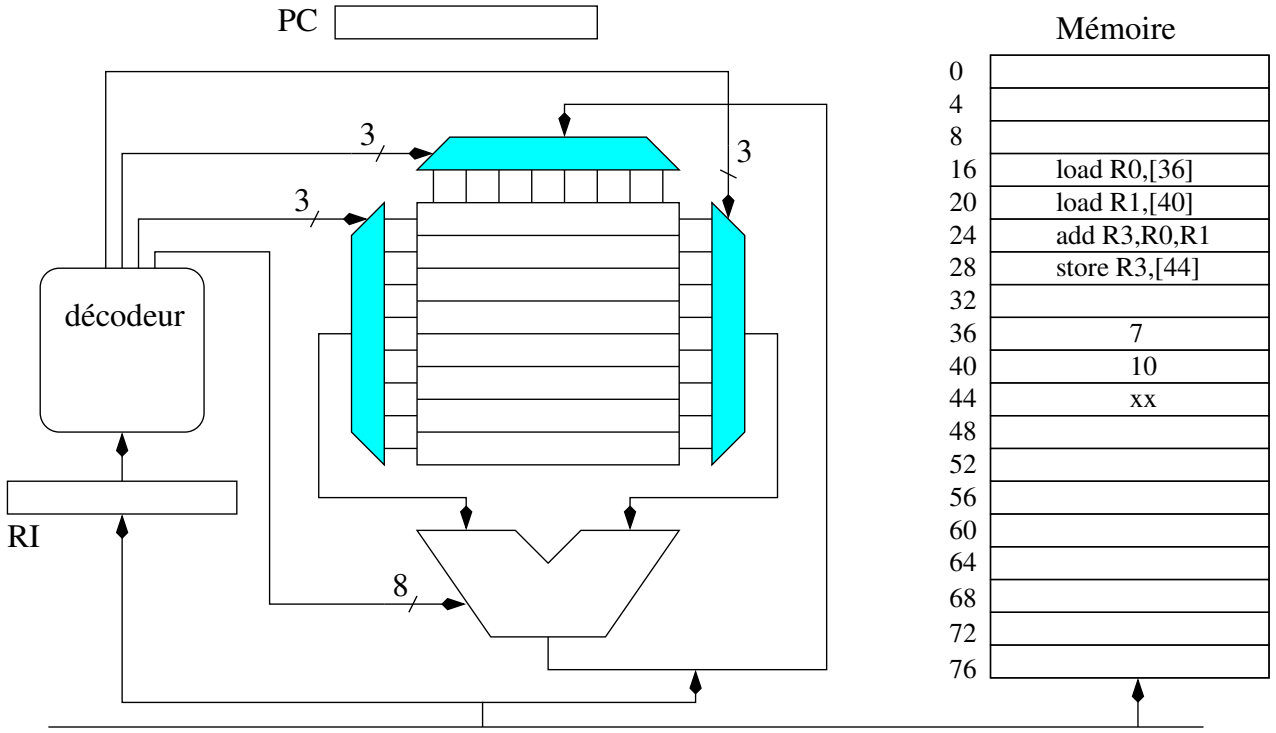

Program execution on a Processor (8 general purpose registers)

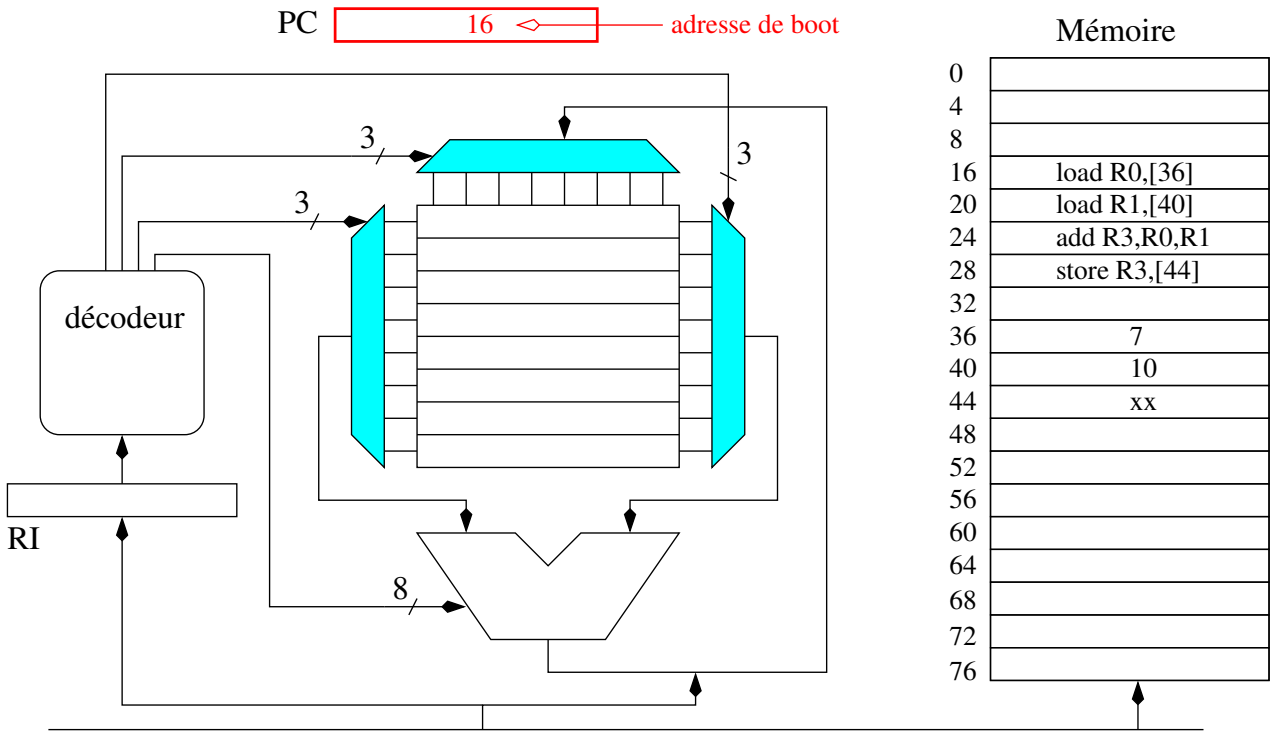

Program execution on a Processor (8 general purpose registers)

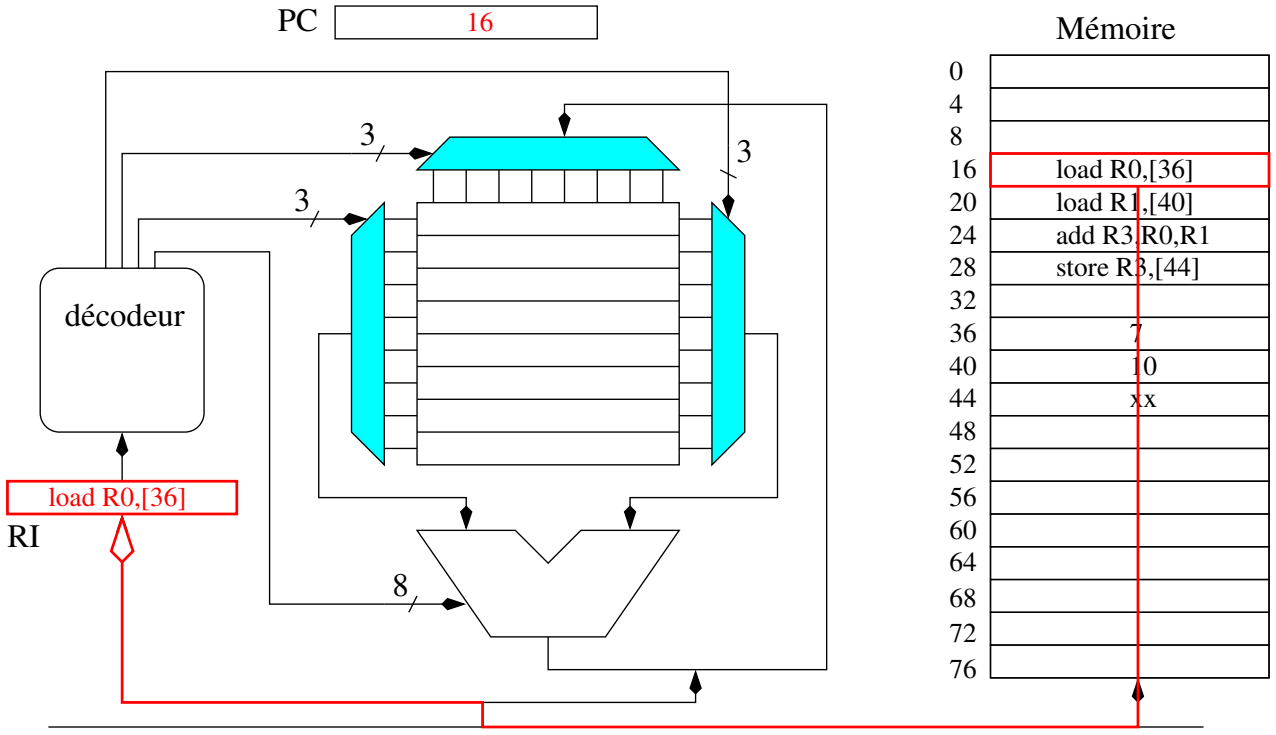

Program execution on a Processor (8 general purpose registers)

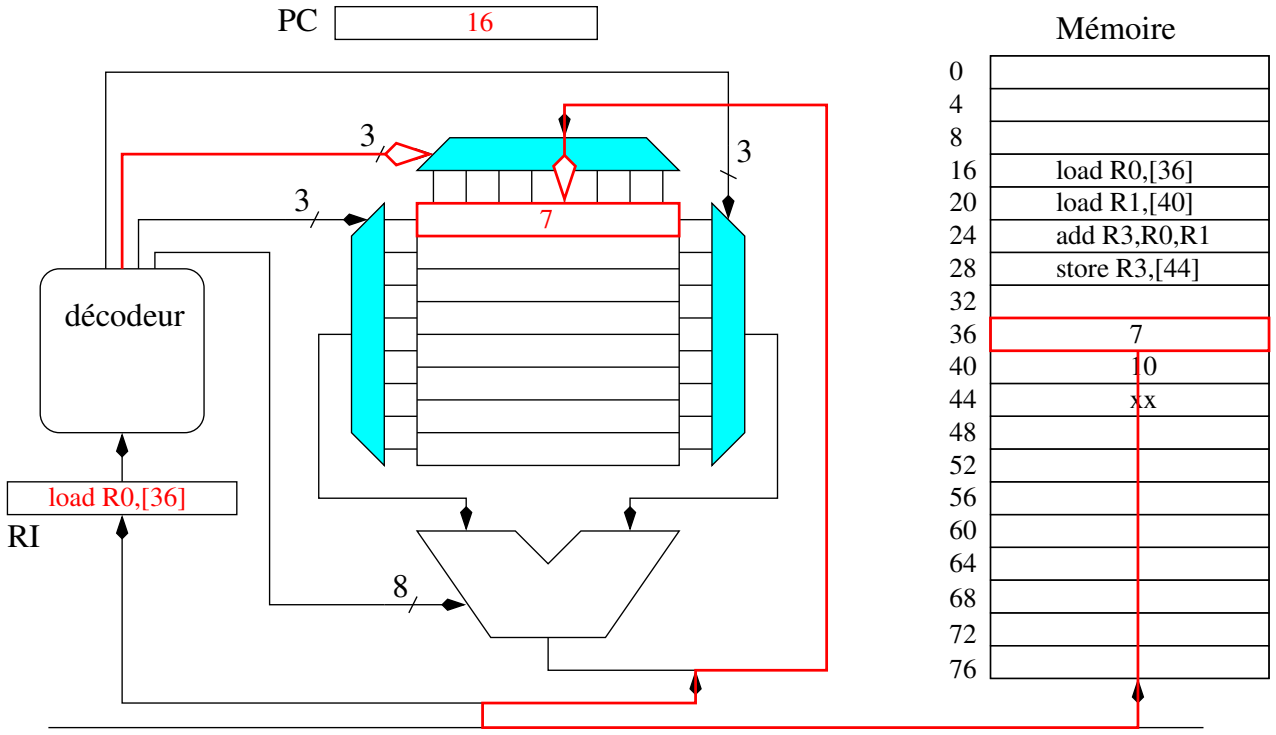

Program execution on a Processor (8 general purpose registers)

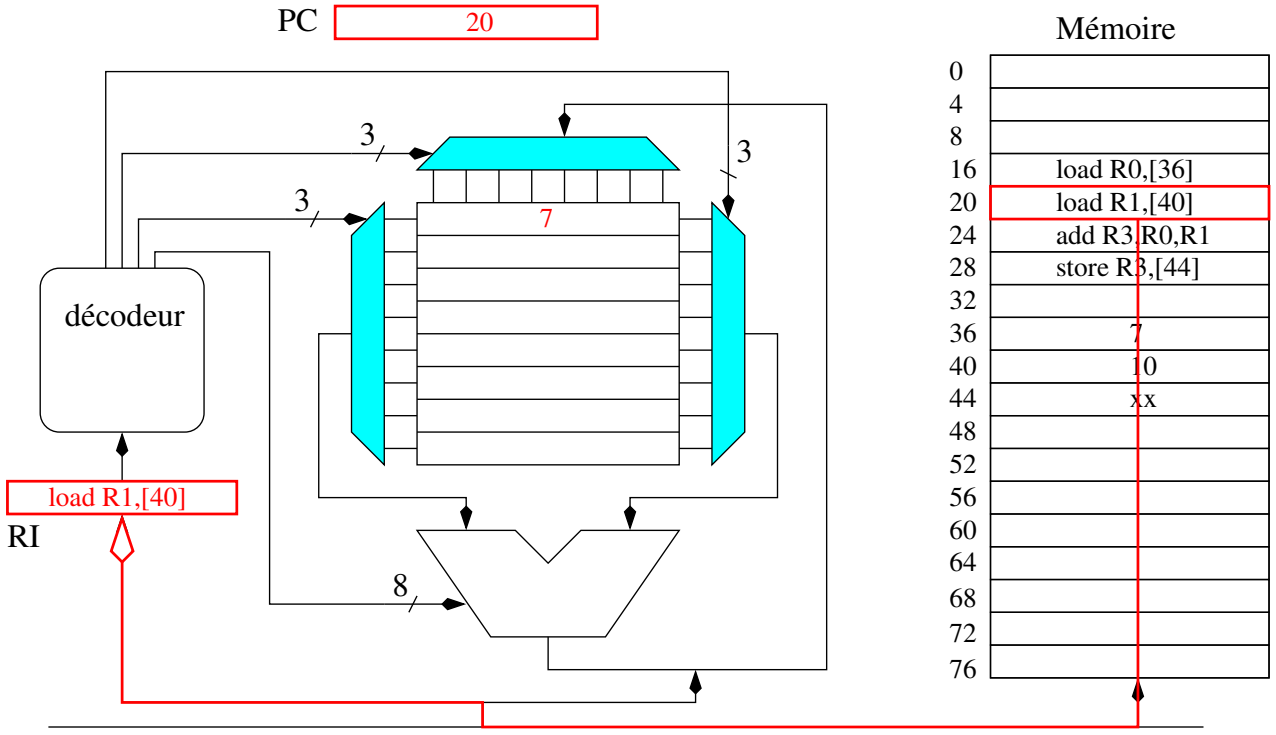

Program execution on a Processor (8 general purpose registers)

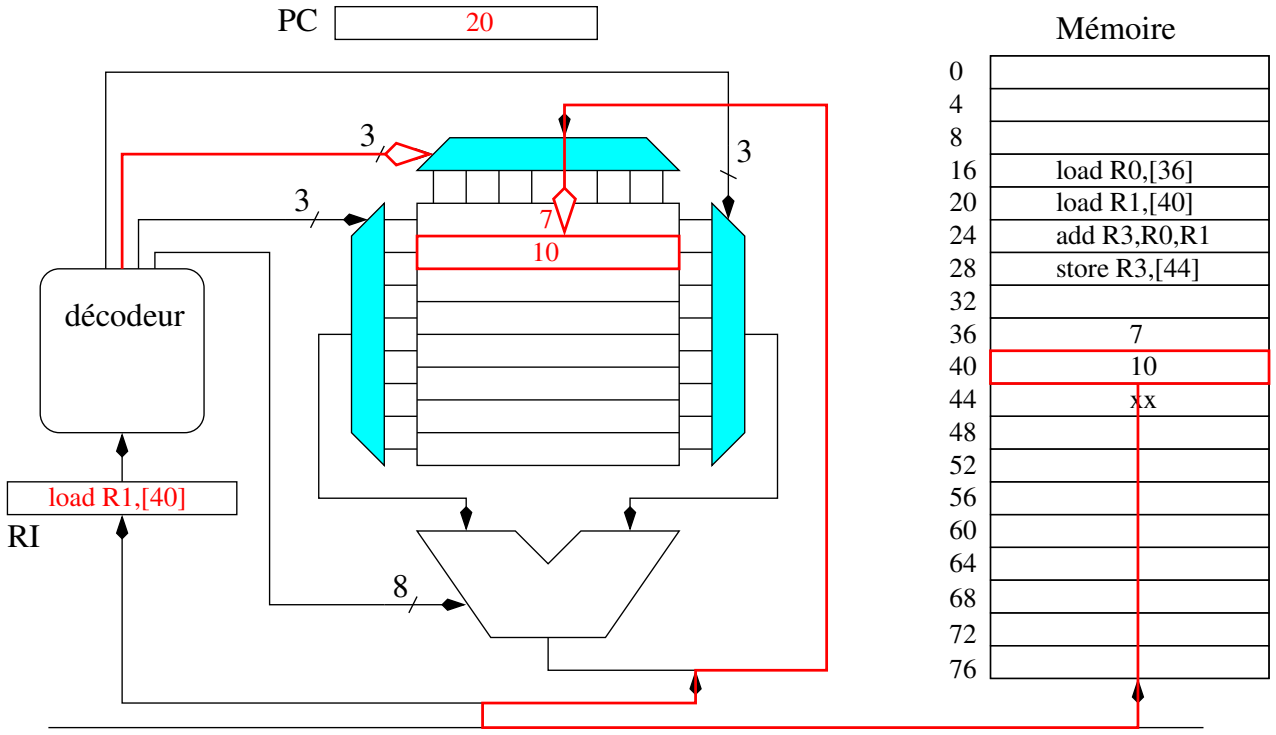

Program execution on a Processor (8 general purpose registers)

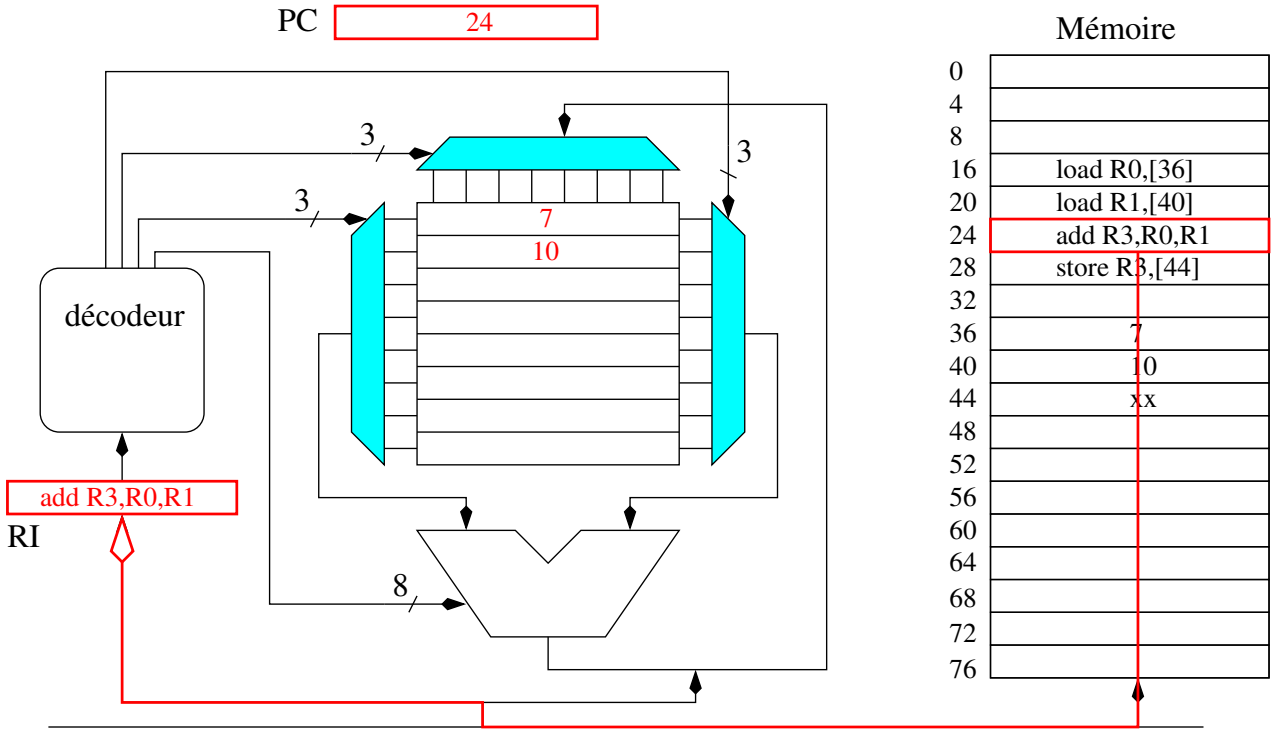

Program execution on a Processor (8 general purpose registers)

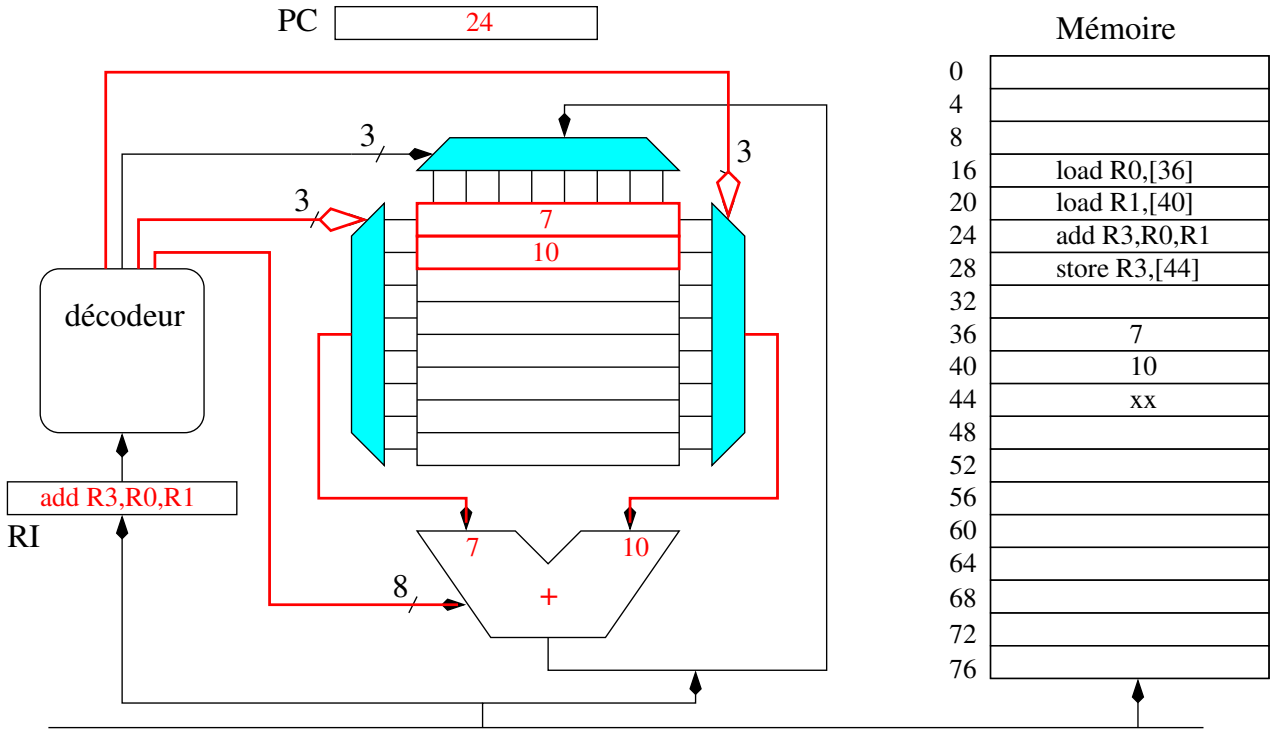

Program execution on a Processor (8 general purpose registers)

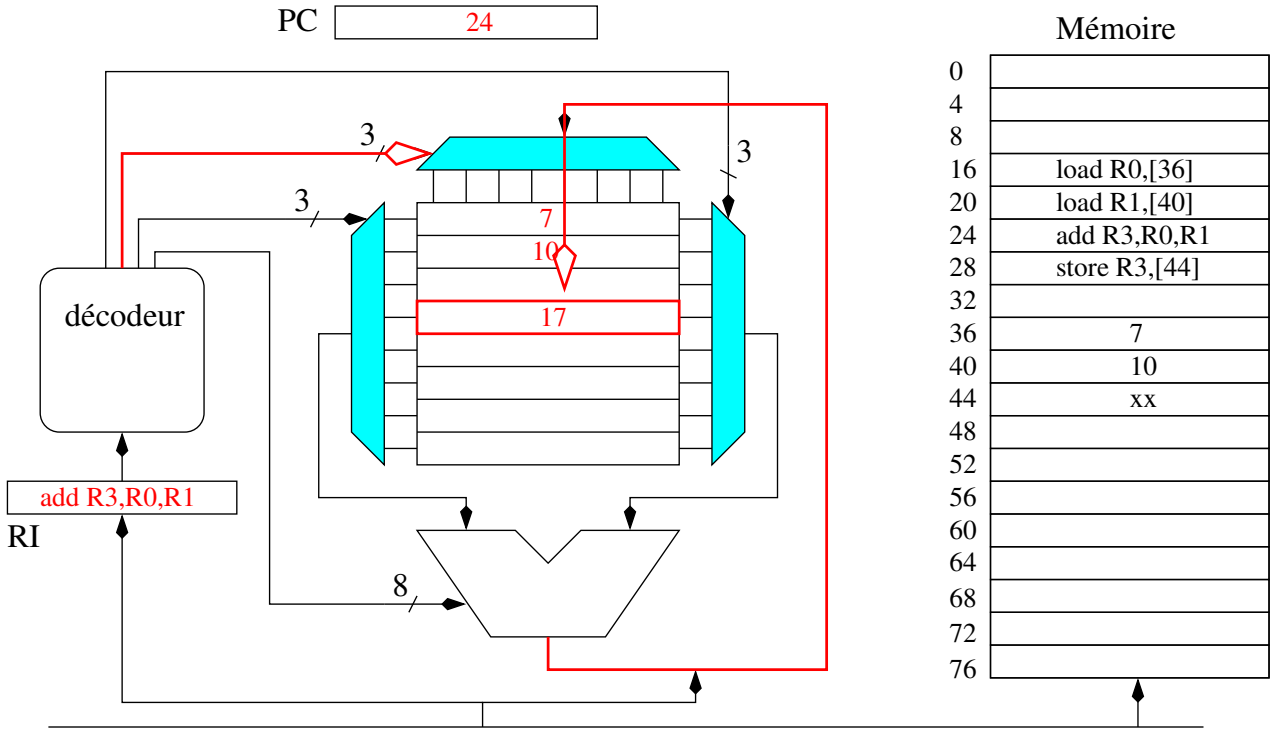

Program execution on a Processor (8 general purpose registers)

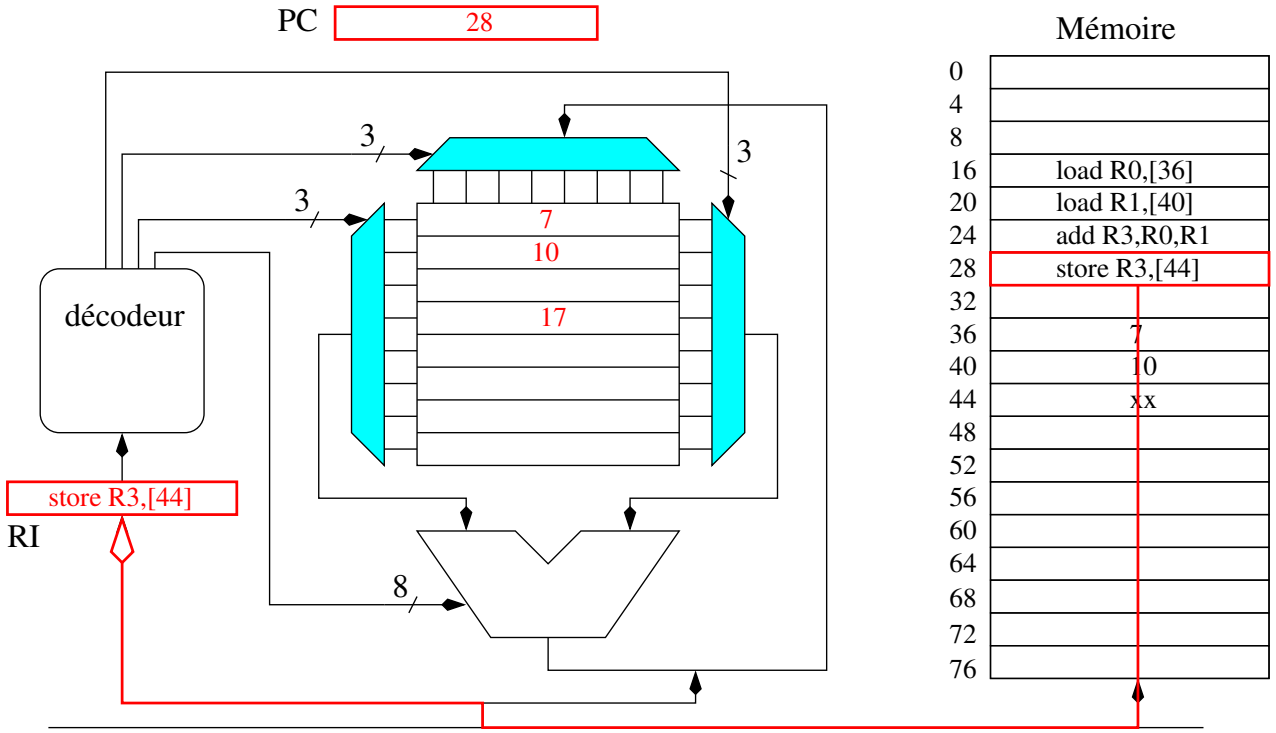

Program execution on a Processor (8 general purpose registers)

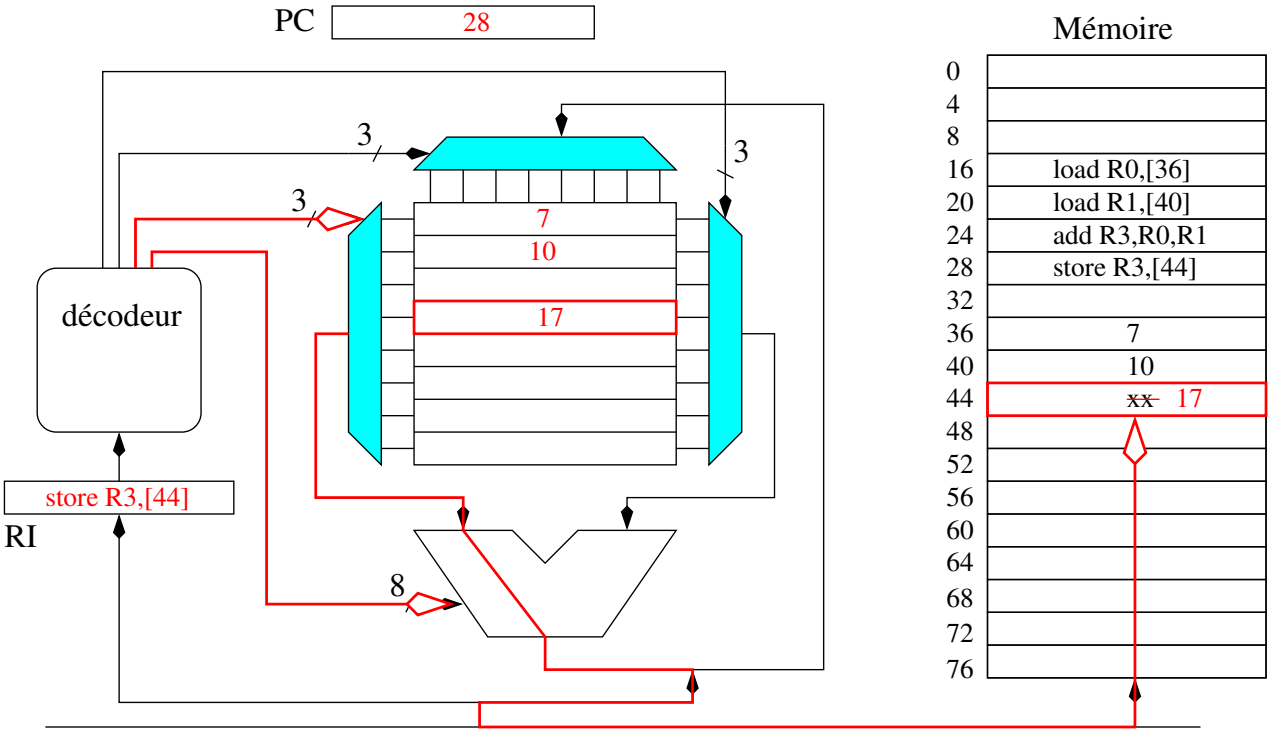

The "Von Neumann cycle"

- The so-called Von Neumann cycle is simply the decomposition of the execution of an instruction in several independent stages.
- The number of stages depend on the processor, usually 5 stages are commonly used as example:
  - **Instruction Fetch** (IF)
    - Reads the instruction from memory (at address $PC) and write it in $IR.
  - **Instruction Decode** (ID)
    - computes what needs to be computed before execution: jump address destination, access to register, etc.
  - **Execute** (EX)
    - executes the instruction: ALU computation if needed
  - **Memory Access** (MEM)
    - Loads (or stores) data from memory if needed
  - **Write Back** (WB)
    - Writes the result into the register file if needed

The MIPS example

- The RISC paradigm was invented by Berkeley and popularized by Henessy and Patterson in the book on MIPS
- MIPS stands for *Microprocessor without Interlocked Pipeline Stages* and propose and architecture to execute each stage independently

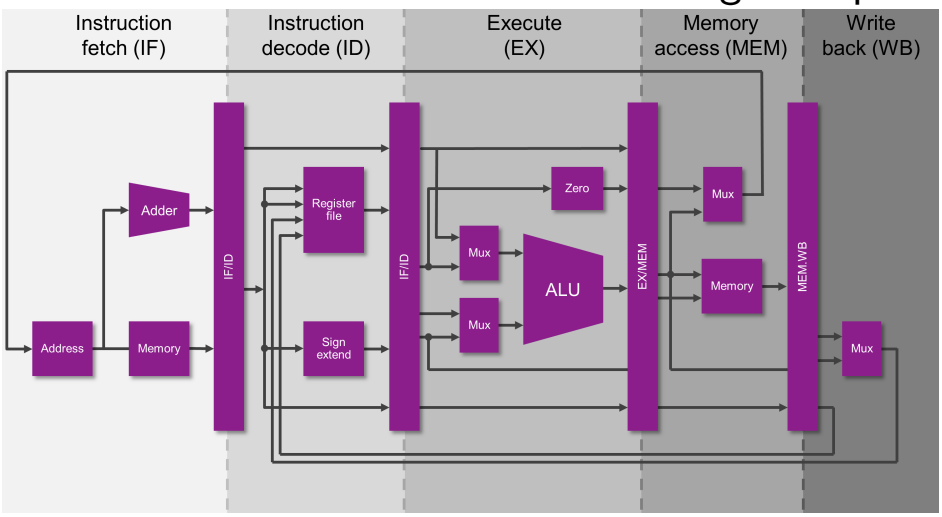

from MIPS website `https://www.mips.com/`

Christian Wolf's slides

- Use Christian Wolf slides for explaining MIPS instruction pipeline
- Here

example of MIPS pipeline CPU architecture

- Taken from Henessy/patterson book

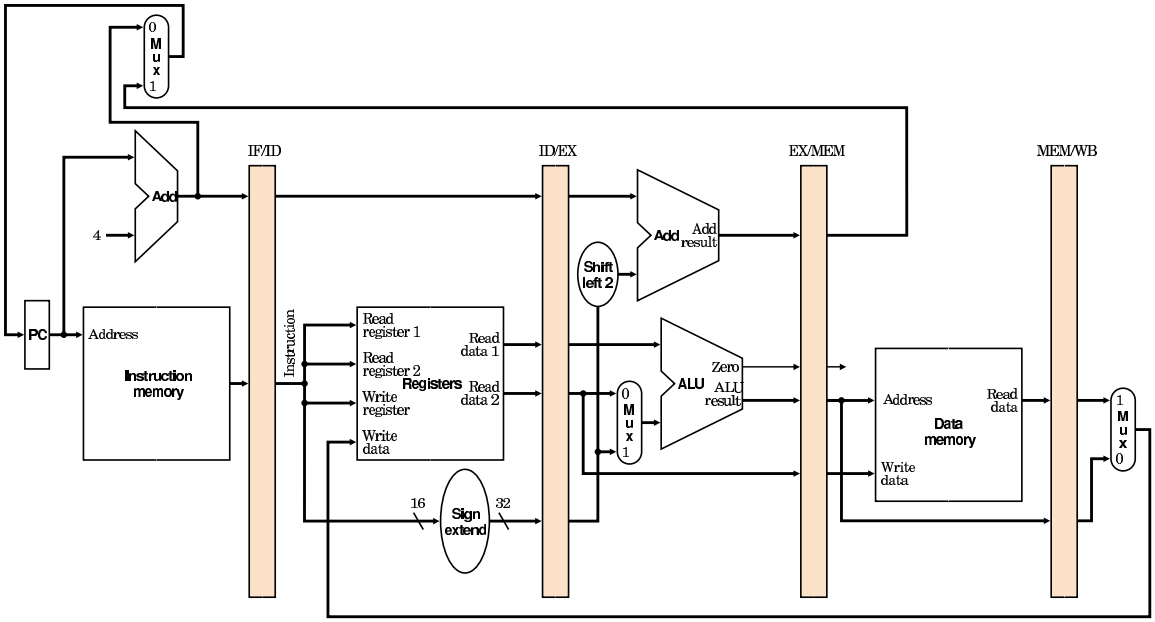

Illustration of bubble on MIPS

- When next instruction cannot be fetched directly (because it need the result of previous instruction for instance) it creates a "bubble"
- For instance: an addition using a register that was just loaded
- The value of the register will be available after the MEM stage of first instruction, hence we can delay on only on cycle, provided there is a *shortcut*.

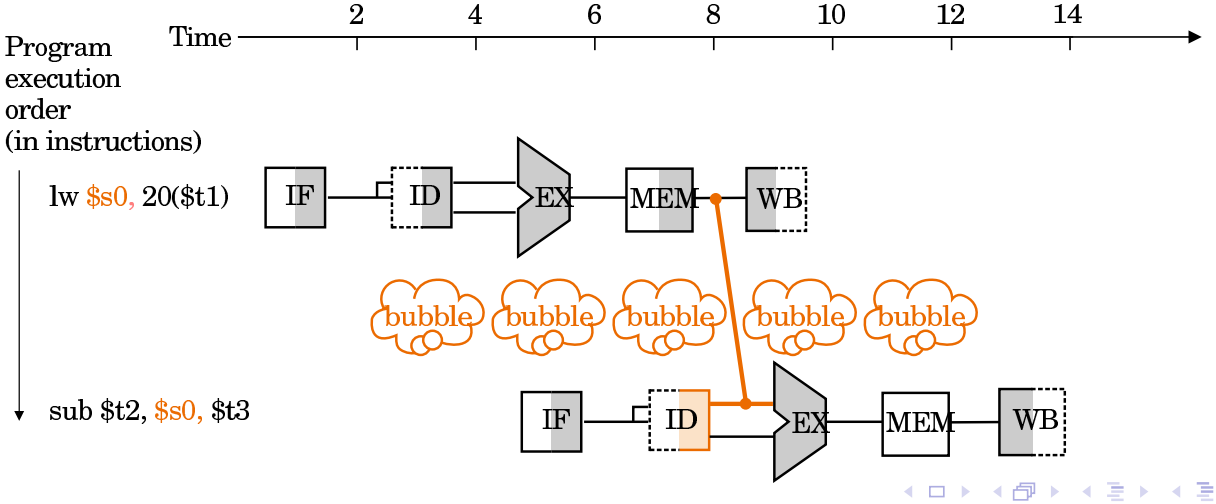

Another illustration of instruction pipeline

- Go back to our previous representation of the processor and memory:
  - Von Neumann computer= Memory + CPU
  - CPU= = control Unit + Datapath
  - Datapath= ALU + Register file

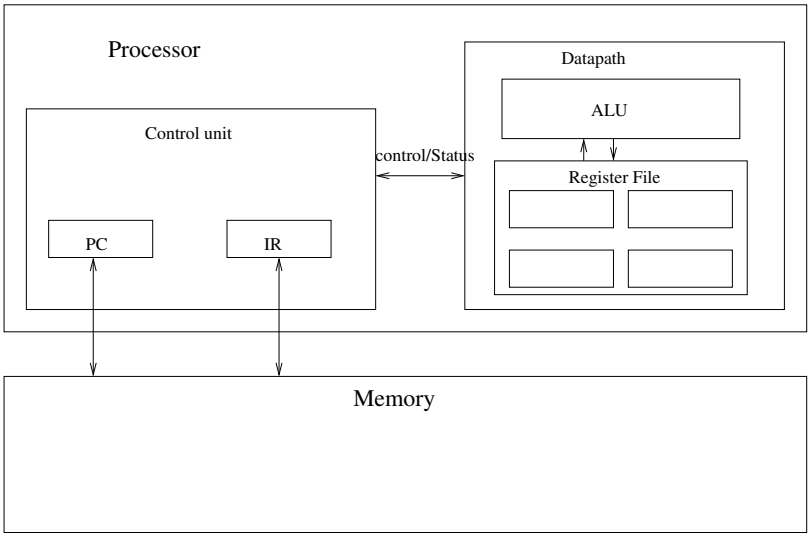

A pipeline example from MIPS

- Execute the sequence of assemby instruction:
  - load value at address 500 in register R0
  - Add 1 to R0 and put result in R1
  - store value of Register R1 at address 500
- (Think of `i=i+1`)
- Code:

```
la R0,500
add R1, R0, 1
sw  R1,500
```

First possible execution: without pipeline

- Before execution starts, $PC contains the address of the first instruction: 100

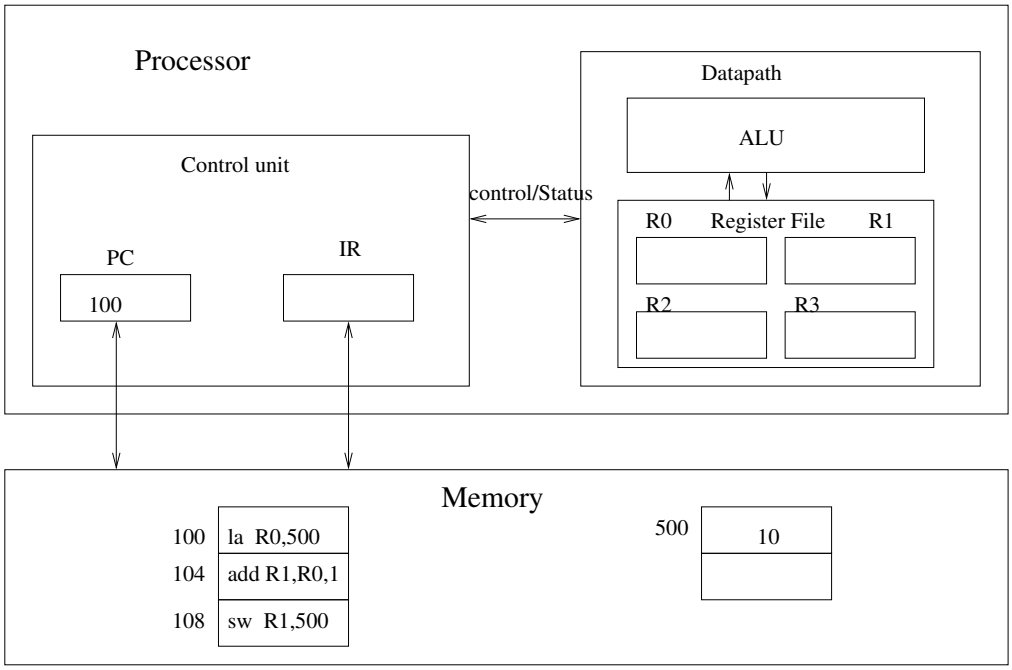

cycle 1

- Instruction Fetch

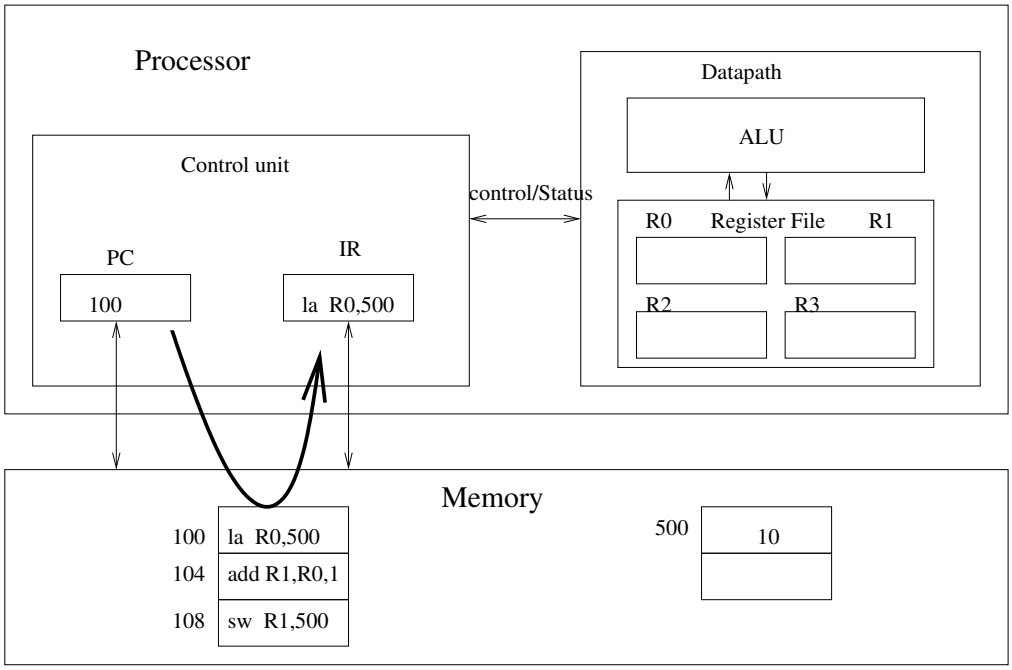

cycle 2

- Instruction Decode

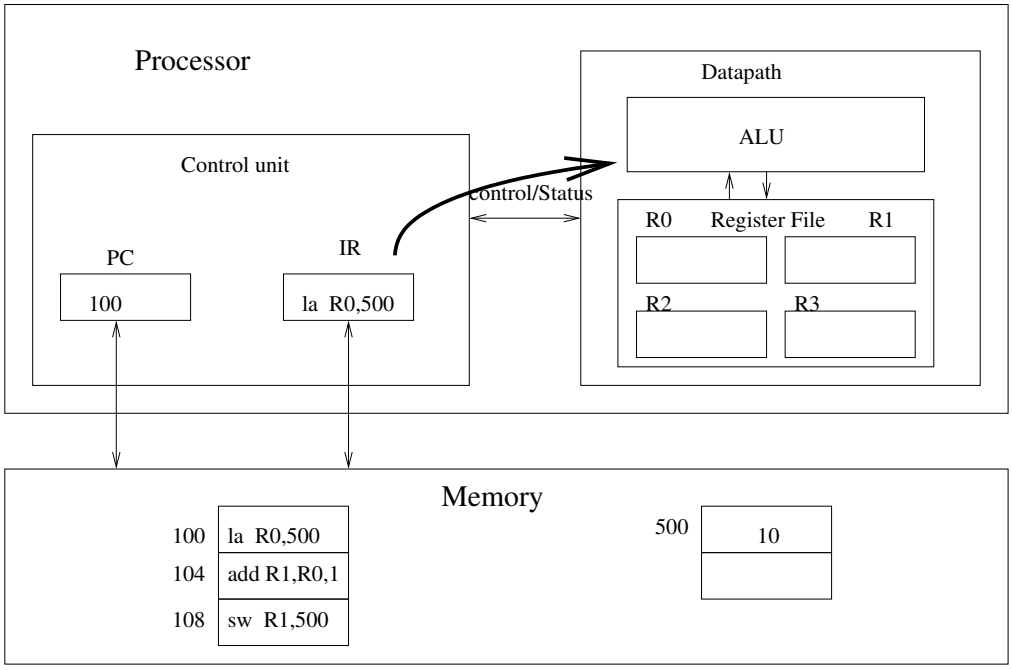

cycle 3

- Execute (nothing for load)

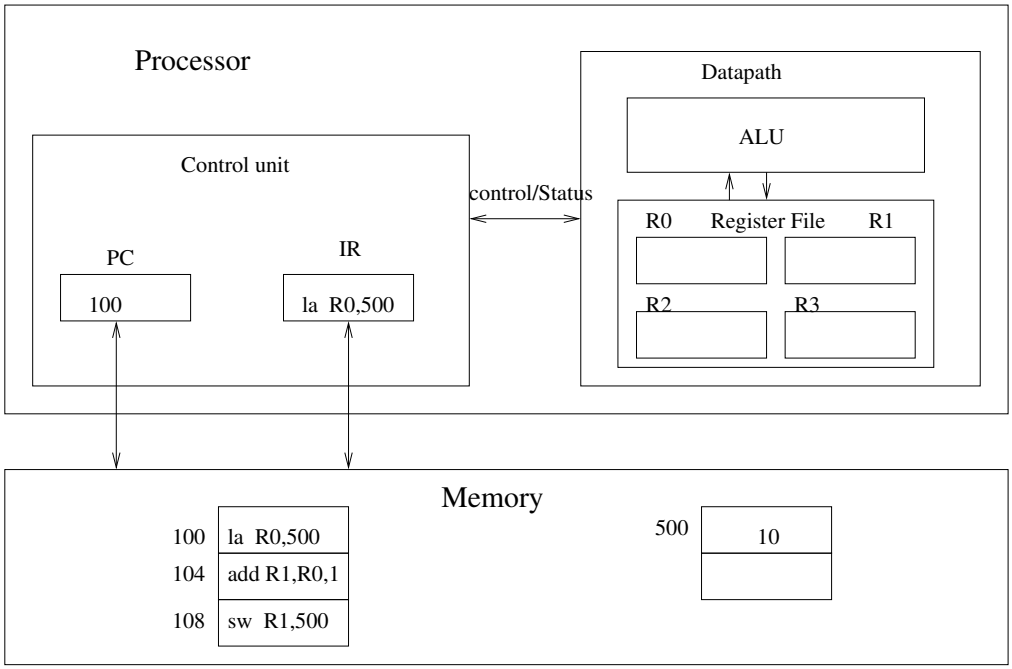

cycle 4

- Memory access

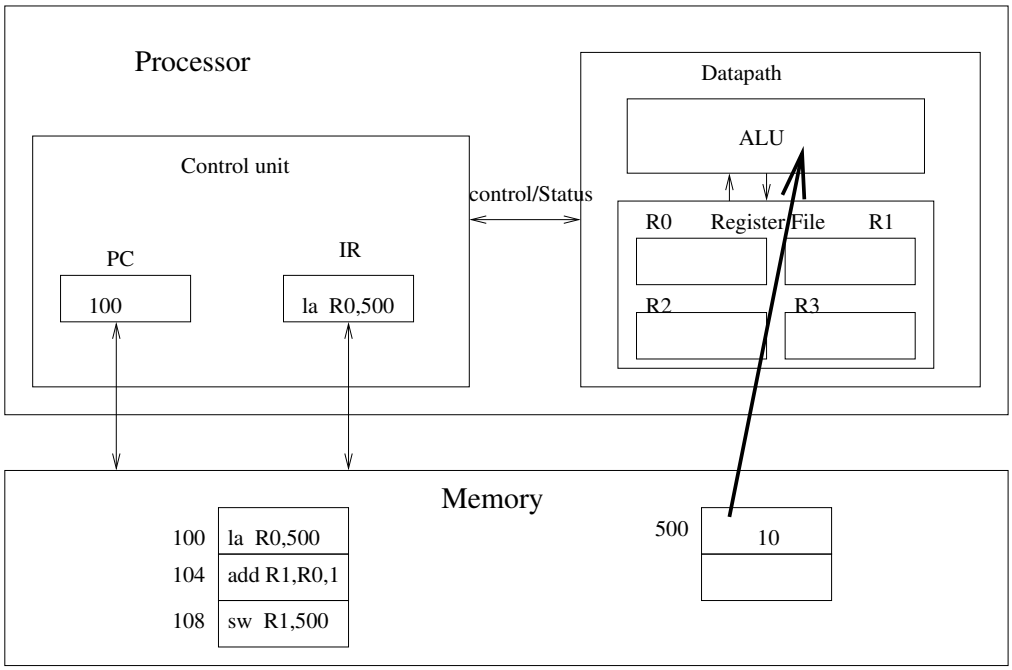

cycle 5

- Write Back

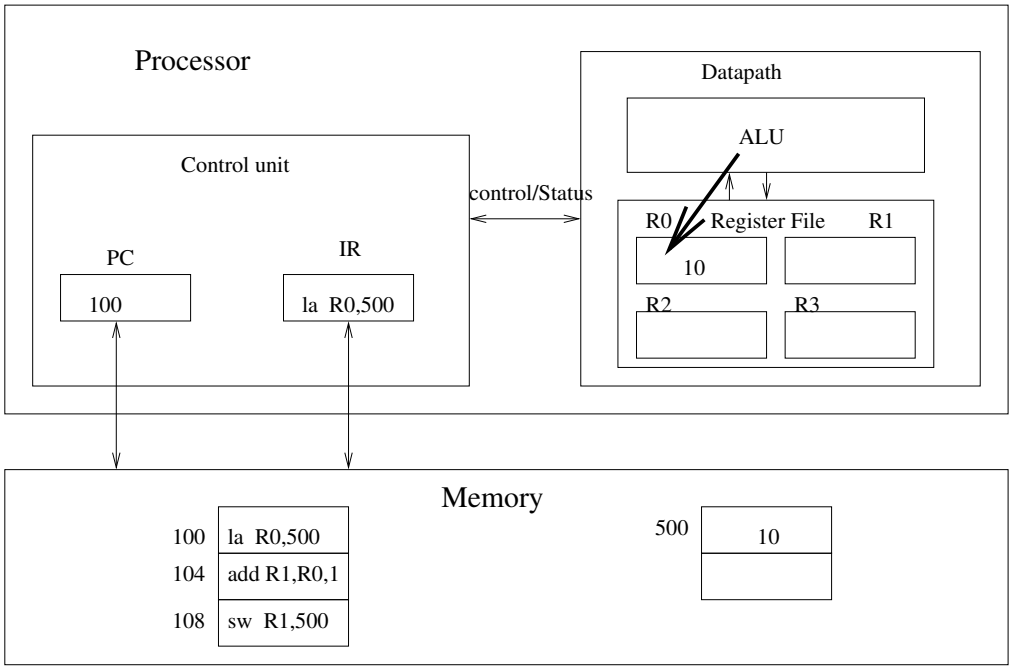

cycle 6

- increment $PC
- Fetch next instruction
- etc. etc.

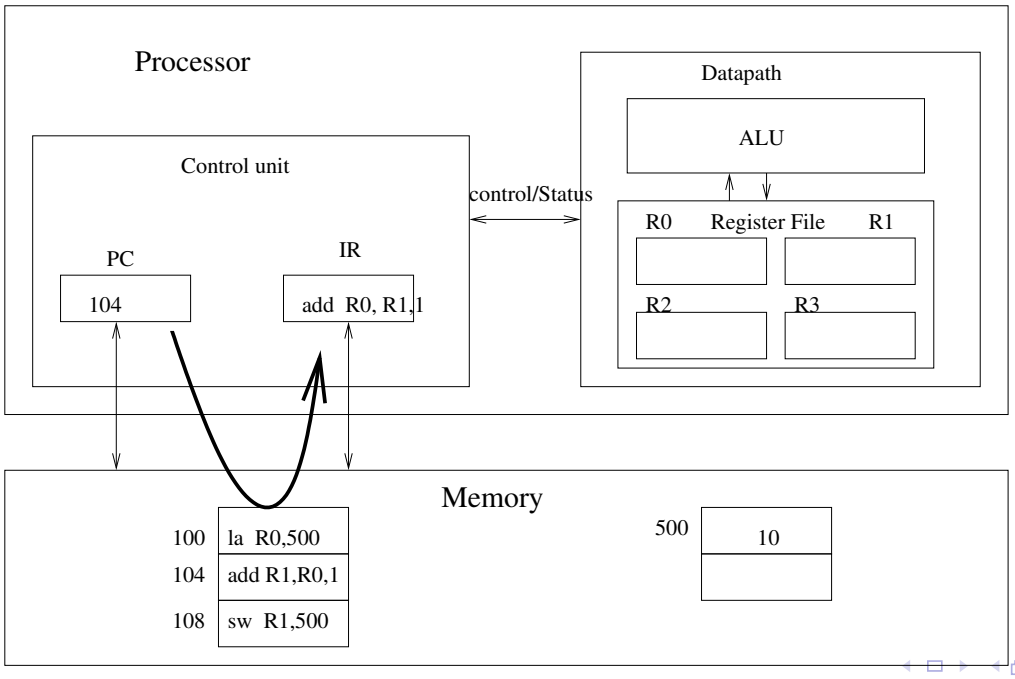

Counting CPI for non-pipelined architecture

- CPI= Cycle per instruction
- 5 cycles for executing on instruction
- $\Rightarrow$ 15 cycles for 3 instructions.

Example of pipelined execution

- Instruction Fetch (for 'load' instruction)

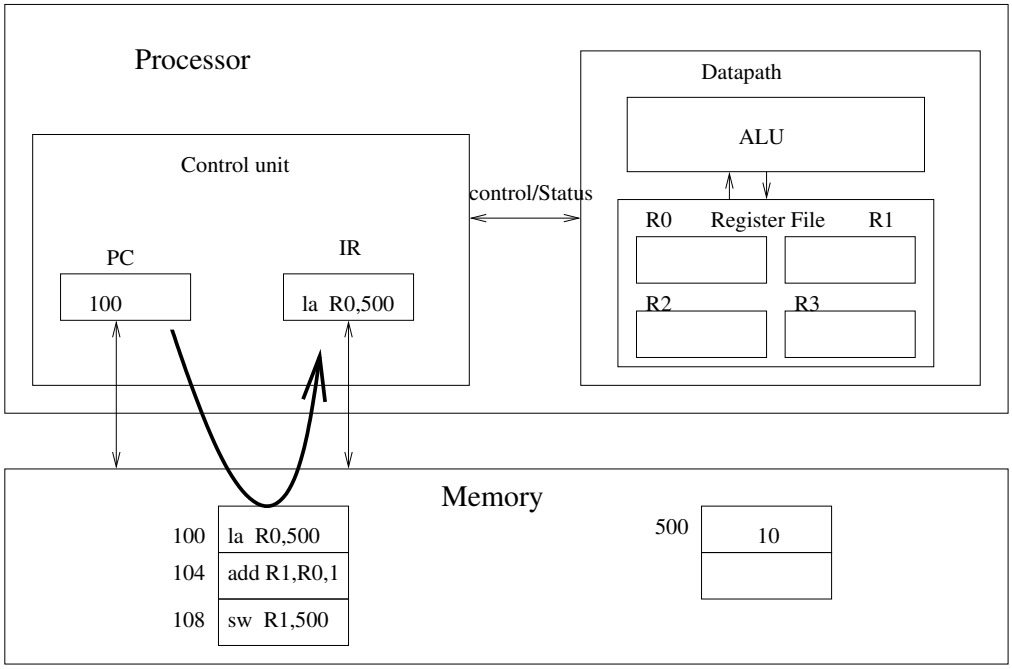

cycle 2

- Instruction Decode (for load)
- Instruction Fetch (for 'nothing' because of a bubble: instruction 'add' delayed)

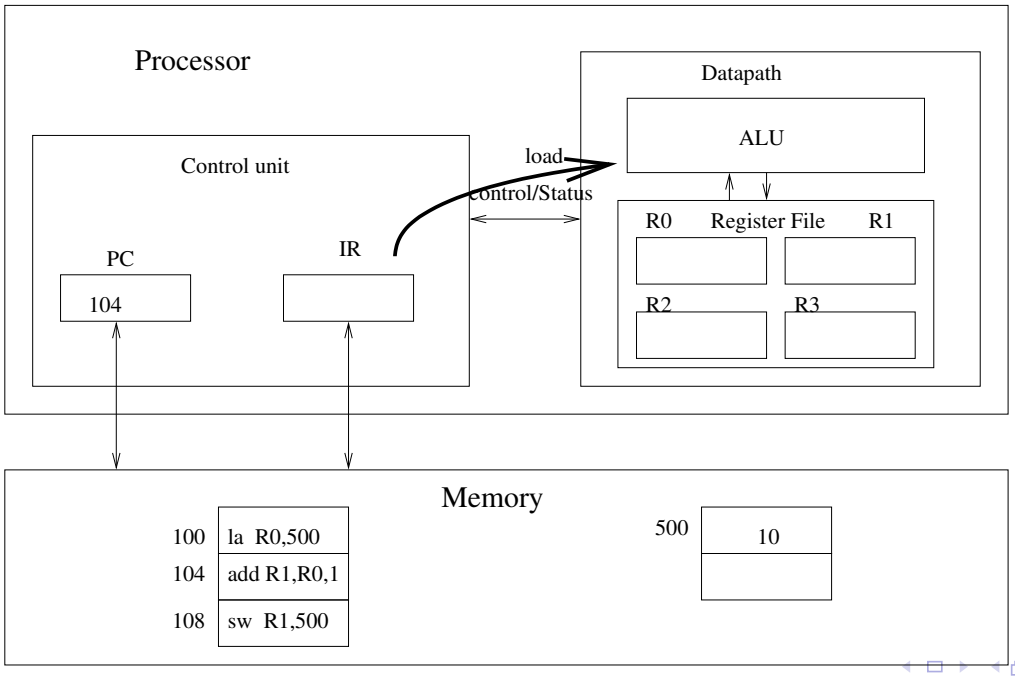

cycle 3

- Execute (for load: nothing to do)
- Instruction Decode (for 'nothing')
- Instruction fetch (for 'add')

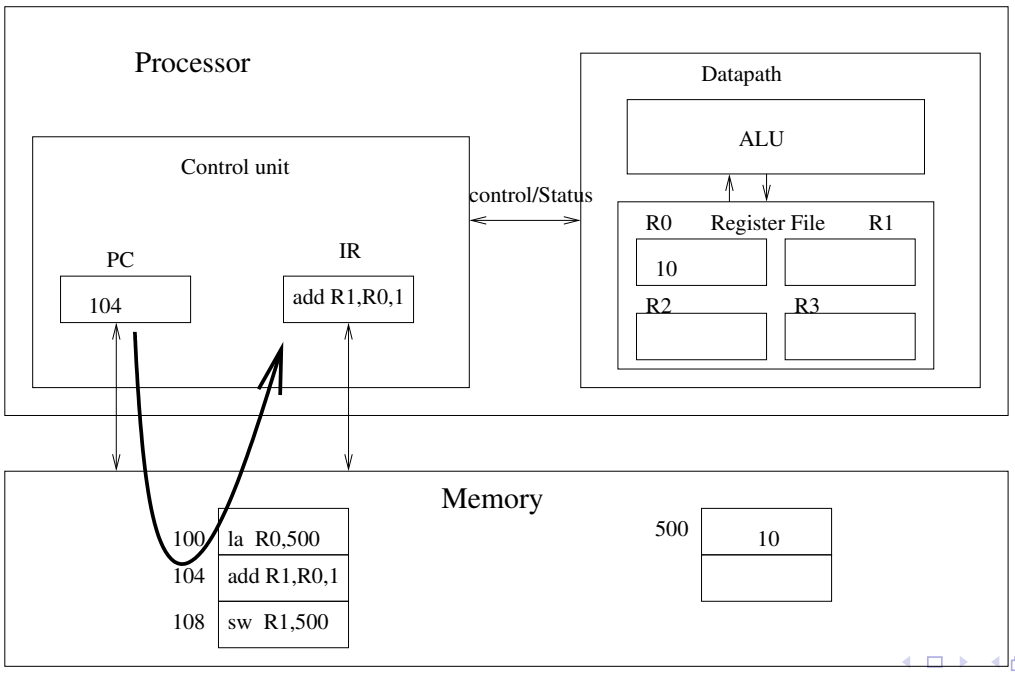

cycle 4

- Memory access (for load)
- Execute (for 'nothing')
- Instruction Decode (for add)
- Instruction fetch (for store)

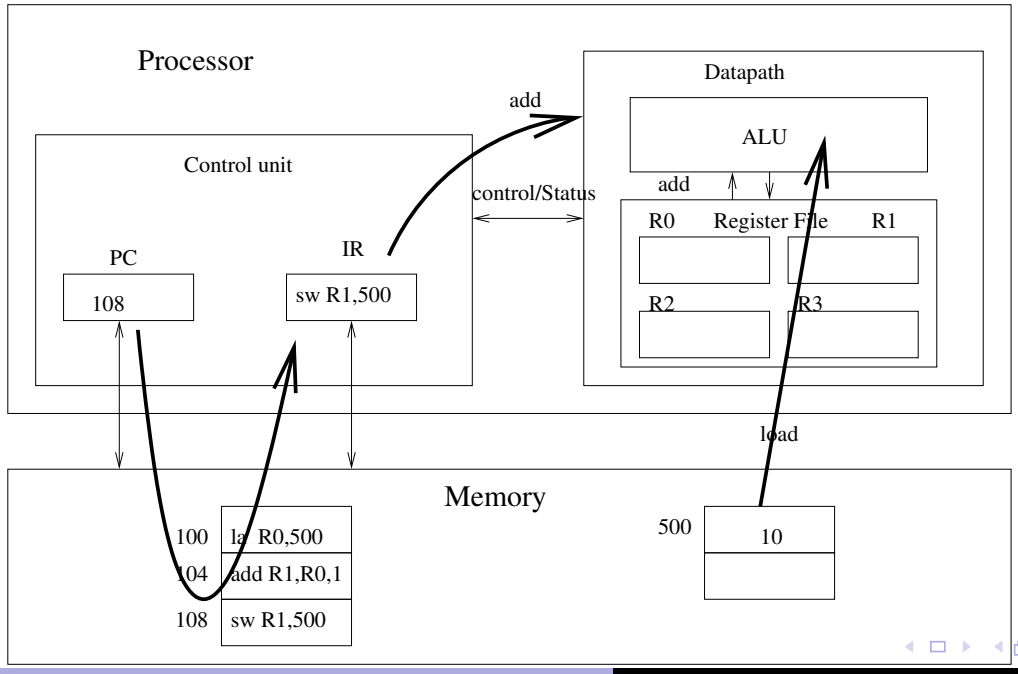

cycle 5

- Write Back (instruction load)
- Memory access (for 'nothing')
- Execute (instruction add: bypass)
- Instruction Decode store

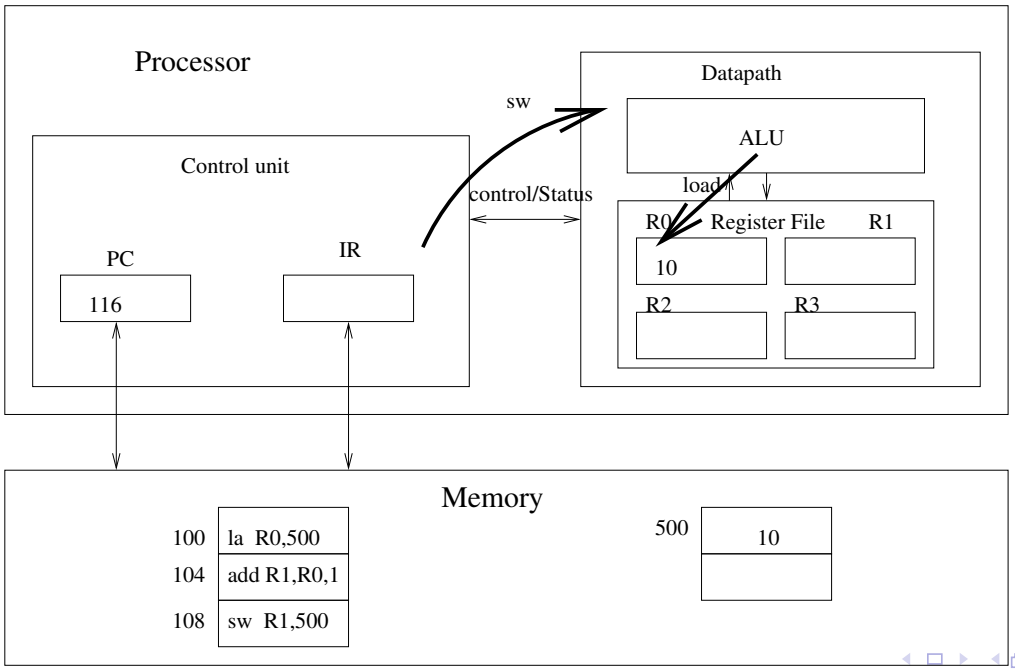

cycle 6

- Write Back (for 'nothing')
- Memory access (instruction add, nothing to do)
- Execute (instruction store: nothing to do)

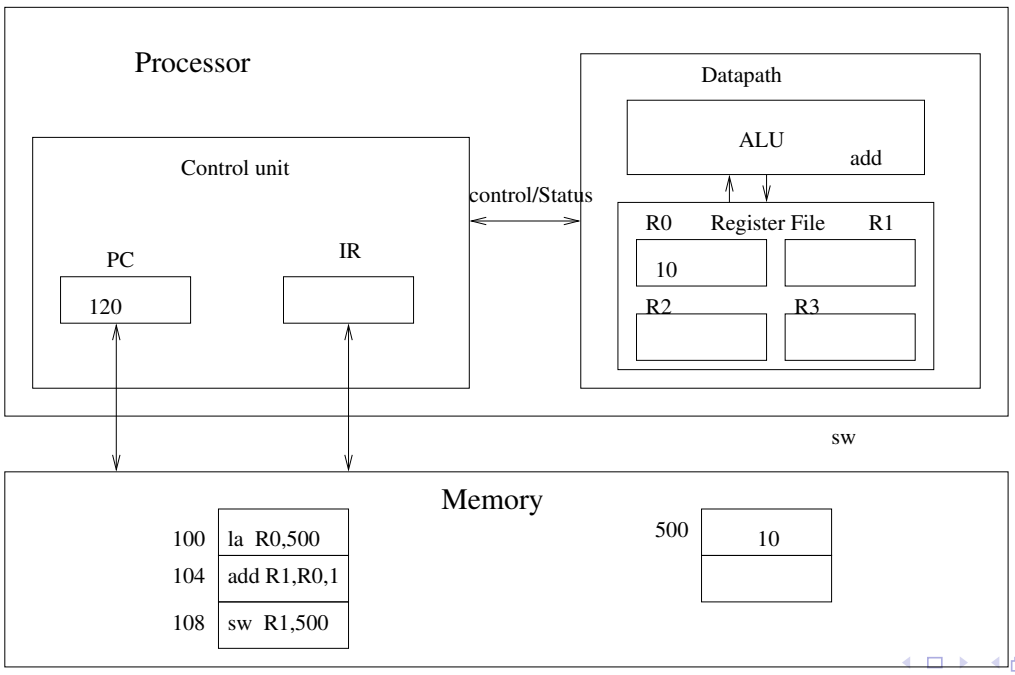

cycle 7

- Write Back (instruction add)
- Memory access (instruction store: bypass)

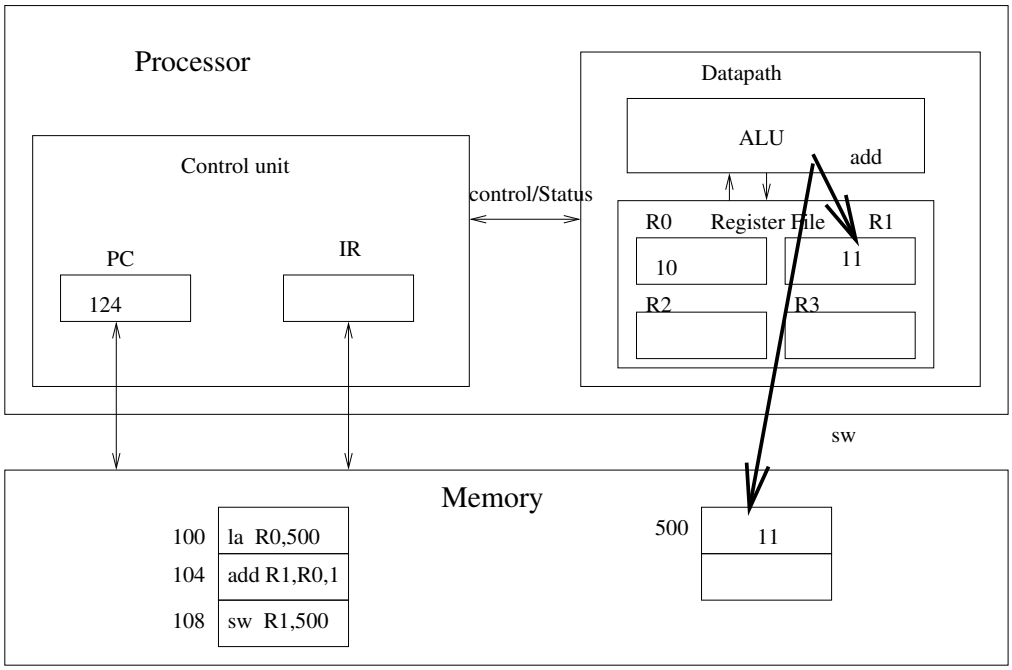

Counting CPI for both architectures

- Non-pipelined architecture:
  - 5 cycles for one instruction
  - $\Rightarrow$ 15 cycles for 3 instructions.
- Pipelined architecture:
  - 5 cycles for one instruction
  - 8 cycles for 3 instructions.
  - $\Rightarrow$ without bubbles, one instruction per cycle
  - A 'jump' instruction interrupt the pipeline (need to wait for the address decoding to fetch next instruction) $\Rightarrow$ *pipeline stall*
  - Some ISA allow to use these *delay slots*: one or two instruction *after* the jump are executed before the jump occurs.

Du langage à l'exécution

Rappels d'architecture

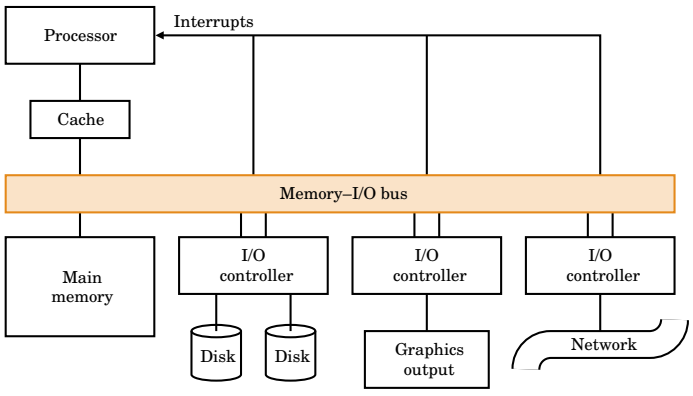

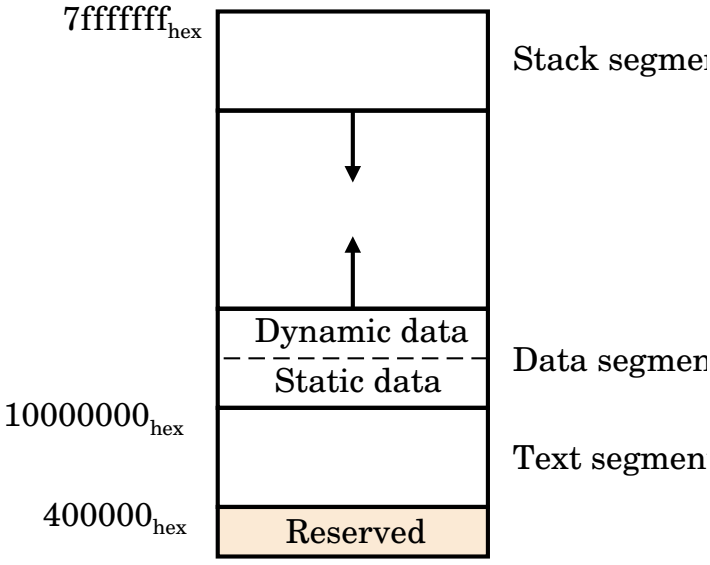

Architecture view from the programmer

- Modern systems allow
  - To run multiple independent programs in parallel (process)
  - To access memory space larger than physical memory available (virtual memory)
- For the programmer: all this is transparent
  - Only one program runs with very large memory available
- The processor view memory contains:
  - The code to execute
  - Static data (size known at compile time)
  - Dynamic data (size known at runtime: the heap, and the space needed for the execution itself: the battery)
- The programmer sees only the data (static and dynamic)

compilation process

- the complete process will translate a C program into code executable (loading and execution will take place later).

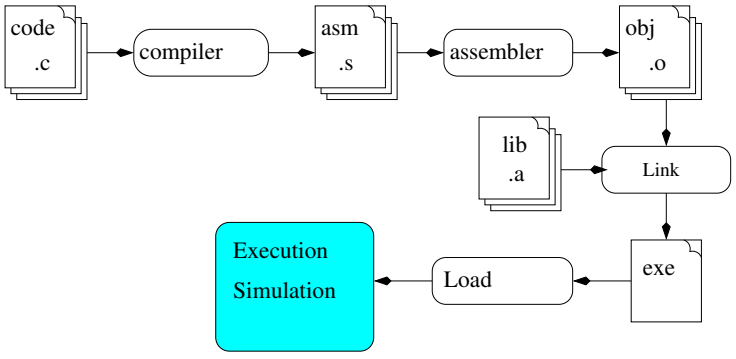

- We often call *compilation* the set compiler + assembler
- The gcc compiler also includes an assembler and linking process (accessible by options)

Your compilation process

- The programmer:
  - Write a program (say a C program: `ex.c`)
  - Compiles it to an object program `ex.o`
  - links it to obtain an executable `ex`

content of `ex.c`

```
#include <stdio.h>

int main()
{
  printf("hello World\n");

  return(0);
}
```

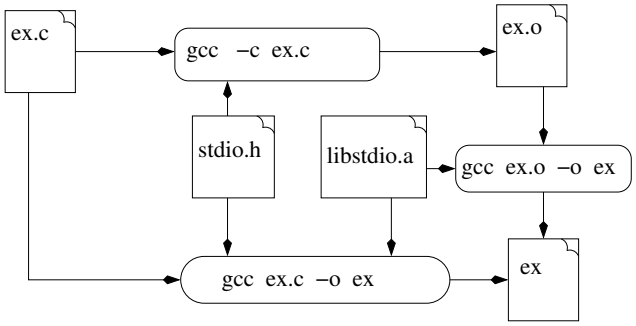

Zooming on "compilation"

- The compilation process is divided in 3 phases:

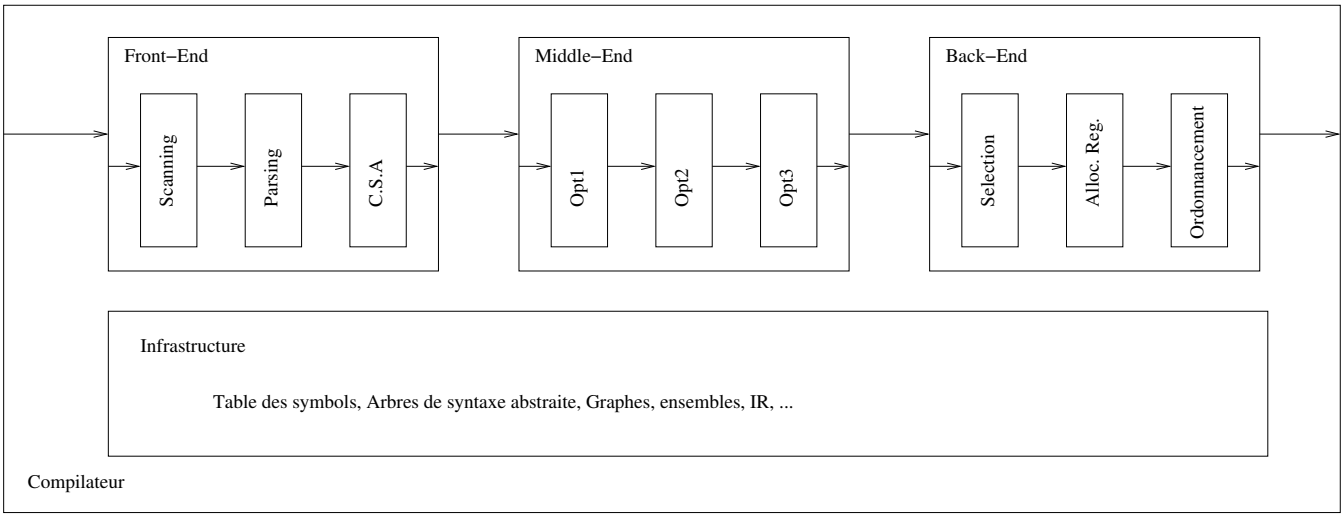

Compilation: the front-end

- The front end of an embedded code compiler uses the same techniques as traditional compilers (we can want to include assembler parts directly)
- Parsing LR(1): the parser is usually generated with dedicated *metacompilation* tools such as `Flex` et `bison` for GNU

Compilation: The middle-end

- Some phases of optimizations are added, they can be very calculative
- Some example of optimisation independent of the target machine architecturre
  - Elimination of redundant expressions
  - dead code elimination
  - constant propagation
- Warning: optimization can hinder the understanding of the assembler (use the -O0 options with tt gcc)

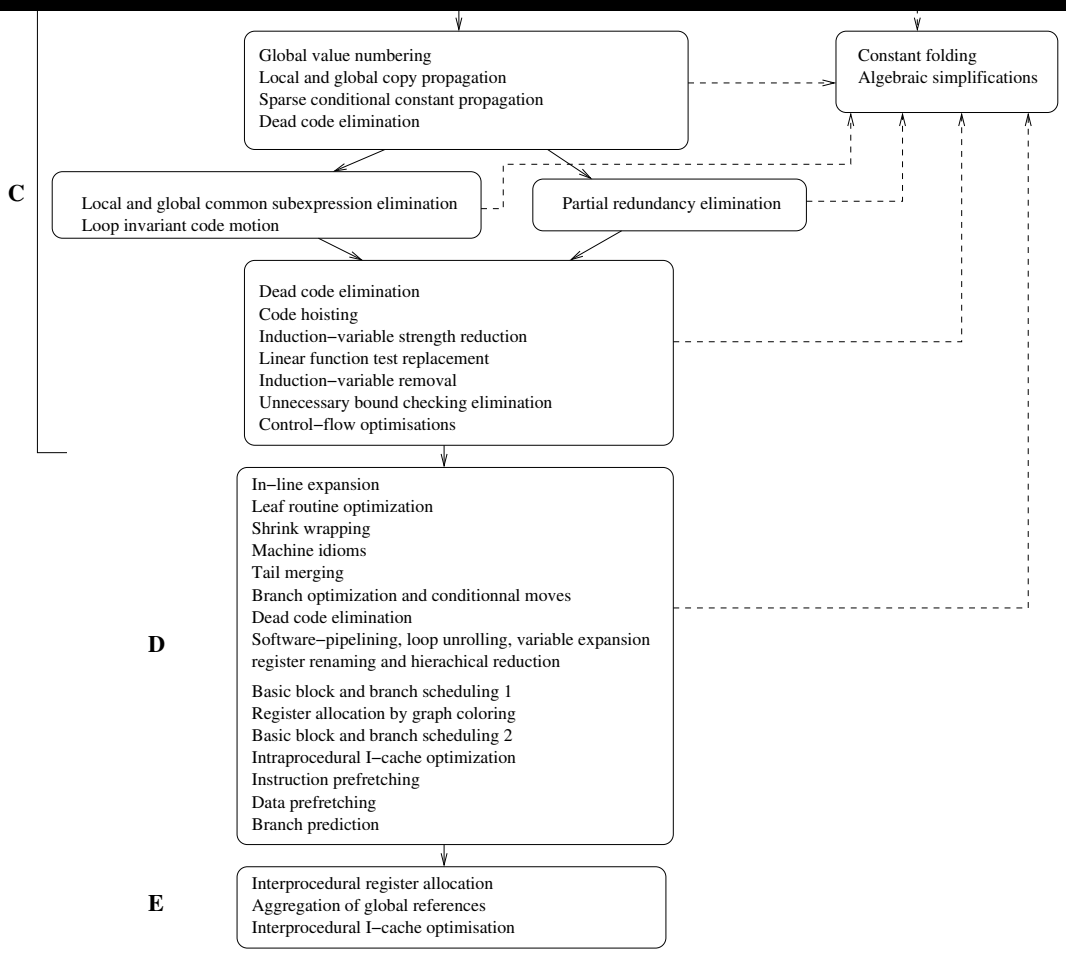

Compilation: The back-end

- The code generation phase is dedicated to architecture target. Retargetable compilation techniques are used for architectural families.
- The most important steps important are:
  - Code selection
  - Register allocation
  - instruction scheduling

GCC

- The `gcc` command runs several program depending on the options
  - The pre-processer `cpp`
  - The compiler `cc1`
  - The assembleur `gas`
  - The Linker `ld`
- `gcc -v` allow to visualize the different programs called by `gcc`

The pre-processer cpp or gcc -E

- the task of the pre-processor are :
  - elimination of comments,
  - inclusion of source files
  - macro substitution (#define)
  - conditionnal compilation.
- Example:

ex1.c

```
#define MAX(a, b) ((a) > (b) ? (a) : (b))
...
f=MAX(3,b);
```

ex1.i

```
#define MAX(a, b) ((a) > (b) ? (a) : (b))
...
f=((3) > (b) ? (3) : (b));
```

The compiler cc1 or gcc -S

- generate assembly code
- `gcc -S main.c -o main.S`
- Exemple :

```
void main()
{ int i;
  i=0;

  while (1)
   {
     i++;
     nop();
   }
}
```


Assembly code generated (for MSP430)

```
mov    #2558,  SP          ; stack initialization de la pile
mov    r1,  r4             ; r4 <- SP
mov    #0,  0(r4)          ; i initialization
inc    0(r4)               ; i++
nop                        ; nop();
jmp    $-6                 ; unnconditionnal jump (PC-6):
incd    SP                 ;
br    #0x1158              ;
```

Assembly code produce by mspgcc -S

```
        .text
        .p2align 1,0
.global         main
        .type   main,@function
main:
/* prologue: frame size = 2 */
.L__FrameSize_main=0x2
.L__FrameOffset_main=0x6
        mov     #(__stack-2), r1
        mov     r1,r4
/* prologue end (size=3) */
        mov     #llo(0), @r4
.L2:
        add     #llo(1), @r4
        nop
        jmp     .L2
/* epilogue: frame size=2 */
        add     #2, r1
        br      #__stop_progExec__
/* epilogue end (size=3) */
/* function main size 14 (8) */
```


Assembler as ou gas

- transform an assembly code into object code (binaire representation of symbolic assembly code)
- Option `-c` of `gcc` allow to conbine compilation et assembly
  `gcc -c main.c -o main.o`

Linking: ld

- Produce the executable (`a.out` by default) from object code of programs and library used
- There are two ways to use libraries in a program
  - Dynamic or shared libraries (default option): the code of the library is not included in the executable, the system dynamically loads the code of the library in memory when calling the program. We need than only *one* version of the library in memory even if several programs use the same library. The library must be em installed on the machine, before running the code.
  - Static libraries: the code of the library is included in the executable. The executable file is bigger but you can run it on a machine on which the library is not installed.

Binary file manipulation

Some usefull command:

- `nm`
  Allow to know symboles (i.e. label: function names) used in an object file or executable

```
trisset@hom\$ nm fib.elf | grep main
000040c8 T main
```

- `objdump` allow to anlayze a binary file. For instance it can get correspondance between binary representation and assembly code

```
trisset@hom$ objdump -f fib

fib:     file format elf32-msp430
architecture: msp:43, flags 0x00000112:
EXEC_P, HAS_SYMS, D_PAGED
start address 0x00001100
```