Introductory presentation

Kevin Marquet
Citi lab / INSA Lyon / INRIA Socrate

Agenda

- Context: memristor, what we can do with it
- What we will *not* talk about
- What we will talk about today and tomorrow
  - Introduce each talk
- Points of view:
  - Hardware/software (scientific fields)
  - Application
  - Level of abstraction
- Present a starting INRIA project addressing together several fields

Memristor

Definition:
a dipole whose resistance depends on the charge that **had** flowed through it

Memristor

Definition:
a dipole whose resistance depends on the charge that **had** flowed through it

- Example: phase-change material

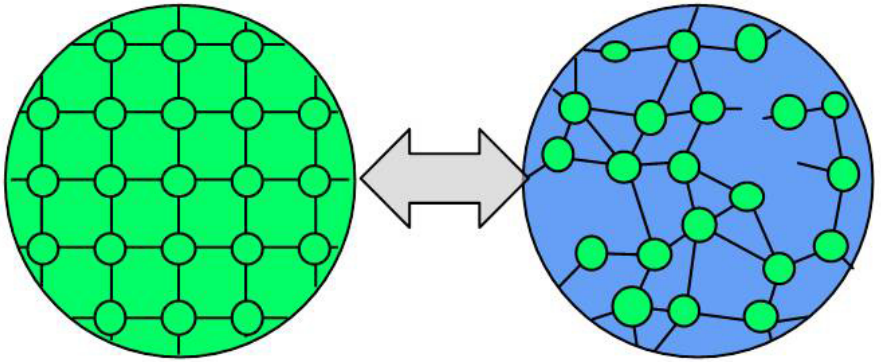

Memristor

Definition:
a dipole whose resistance depends on the charge that **had** flowed through it

- Example: phase-change material

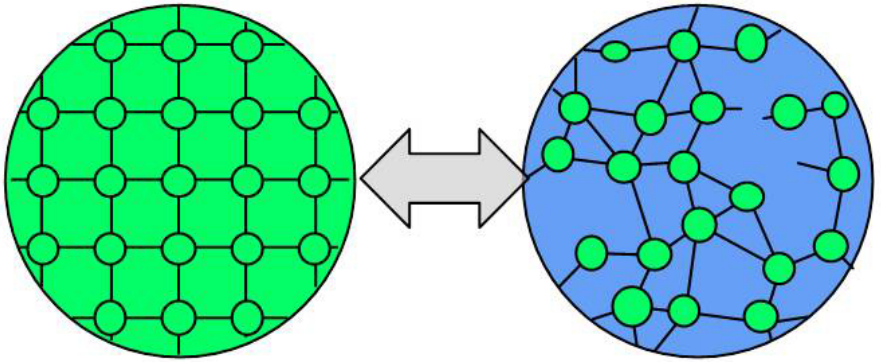

- Symbol:

Memory

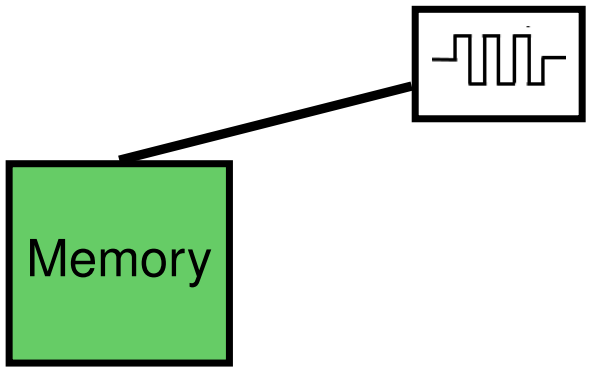

Digital computation

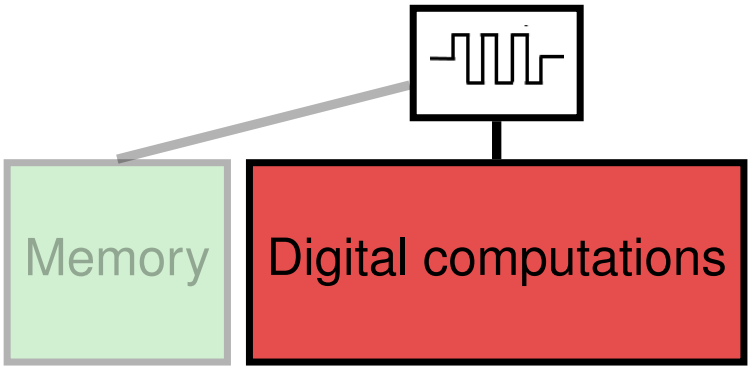

- MAGIC, IMPLY approaches

Digital computation

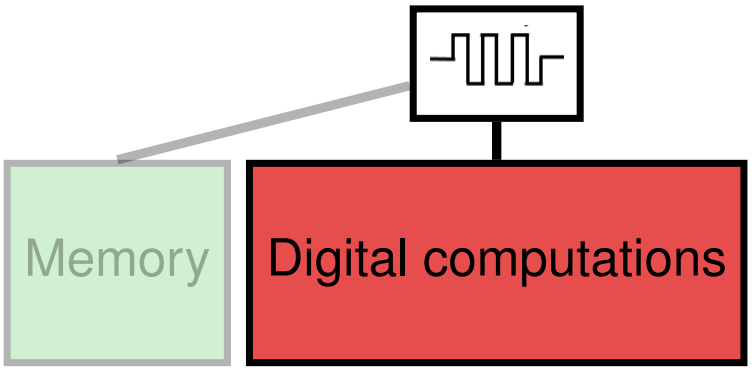

- MAGIC, IMPLY approaches

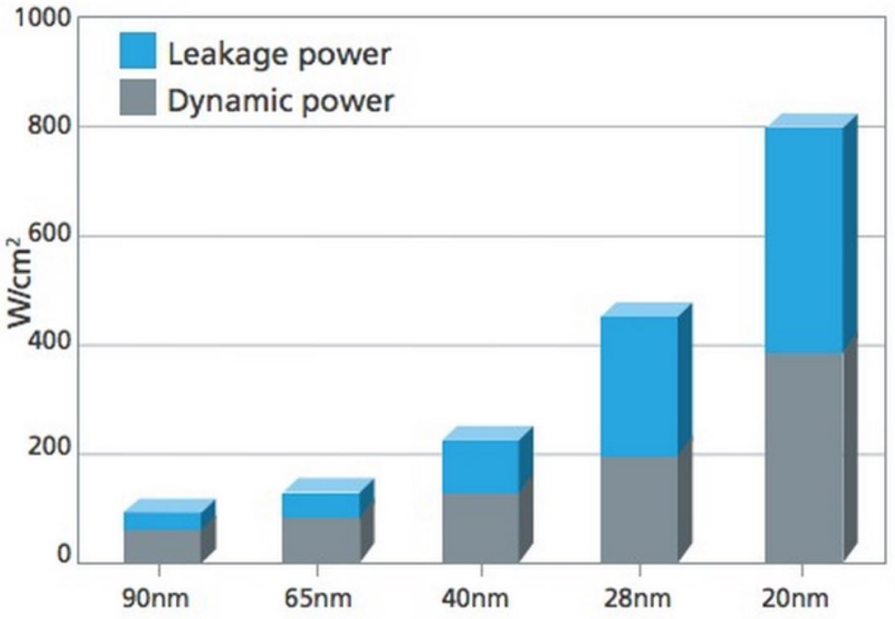

Digital computation

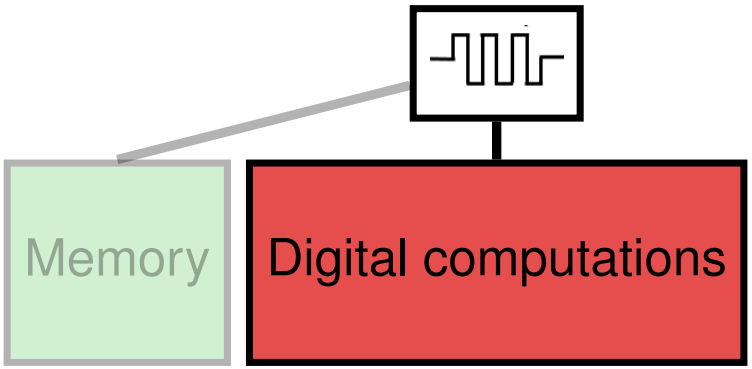

- MAGIC, IMPLY approaches
  - Leakage power of NoC : 43% at 45nm, to 54% at 32nm, to over 65% at 22nm. Citation : mp3

Digital computation

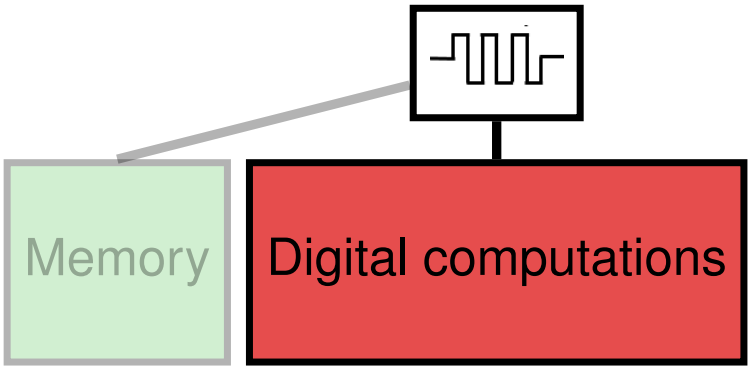

- MAGIC, IMPLY approaches
- fpga++ : use memristor as a reconfigurable switch
  - huge power saving
  - low-latency dynamic reconfiguration (today : about 10ms)
  - => reuse hardware

Analog circuit

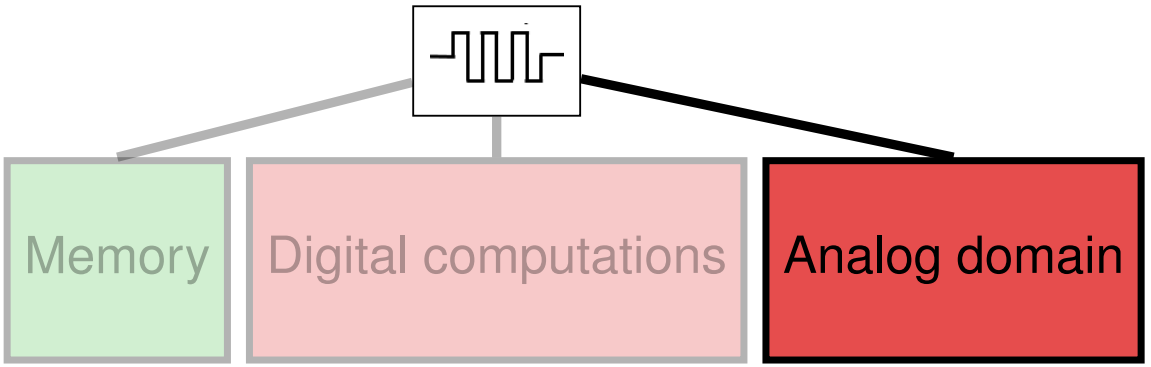

Analog circuit: circuit that transport a continuous signal

- Resistance of some memristors may vary continuously.
- Filters, amplifiers, oscillators, etc.
- Neuromorphic computing

Neuromorphic Computing

Neuromorphic computing: mimic neuro-biological architectures present in the nervous system

Neuromorphic Computing

Neuromorphic computing: mimic neuro-biological architectures present in the nervous system

- What does that got to do with the price of tea in China ?

Neuromorphic Computing

Neuromorphic computing: mimic neuro-biological architectures present in the nervous system

- What does that got to do with the price of tea in China ?
  - Learning is related to *Spike Timing Dependant Plasticity*

Neuromorphic Computing

Neuromorphic computing: mimic neuro-biological architectures present in the nervous system

- What does that got to do with the price of tea in China ?
  - Learning is related to *Spike Timing Dependant Plasticity*
- What does that got to do with the price of tea in China ?

Neuromorphic Computing

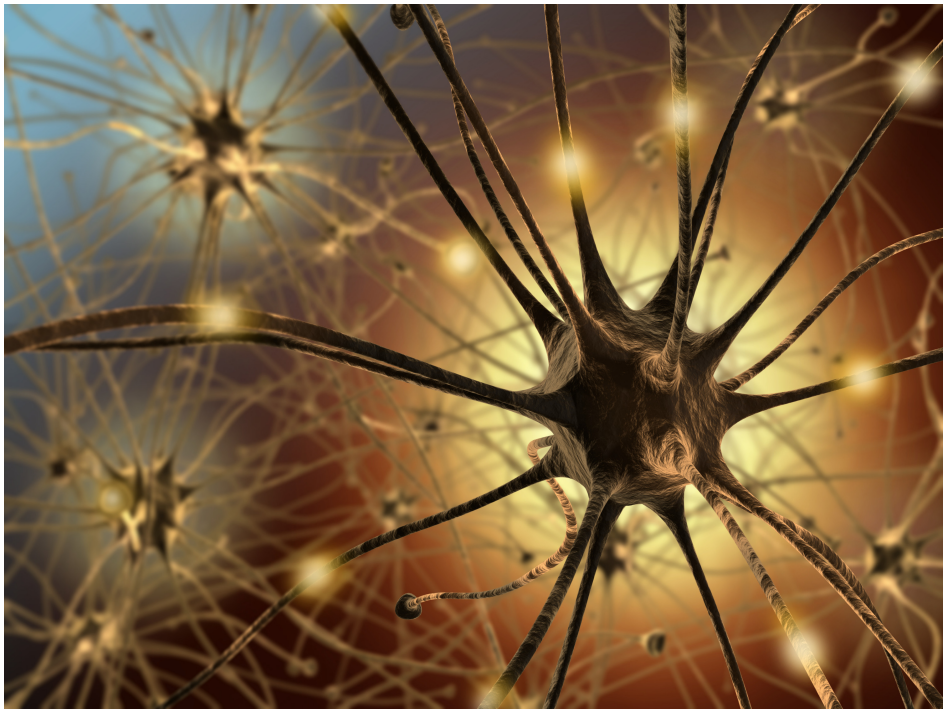

Neuromorphic Computing

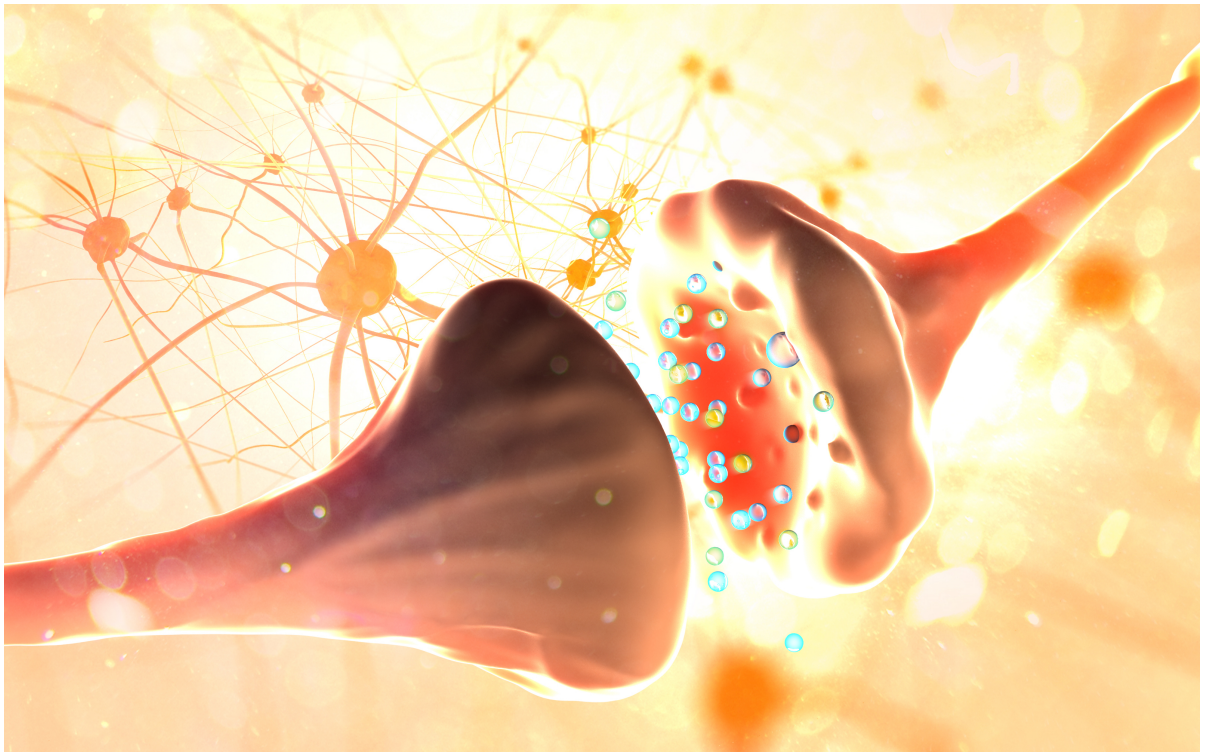

Learning is related to *Spike Timing Dependant Plasticity*

Neuromorphic Computing

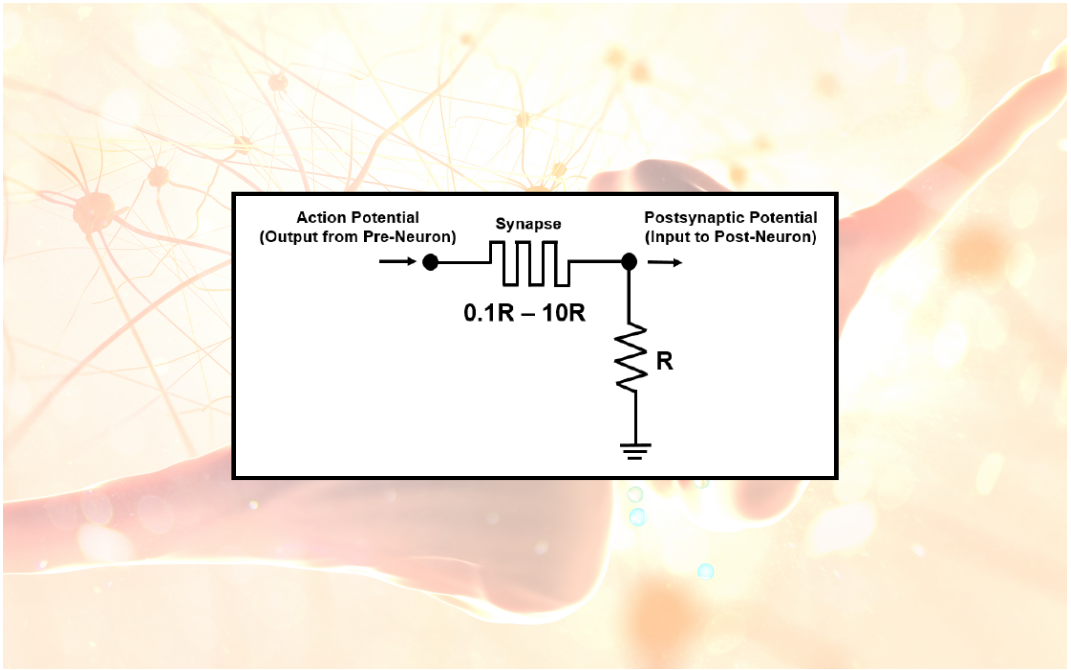

Learning is related to *Spike Timing Dependant Plasticity*

Neuromorphic Computing

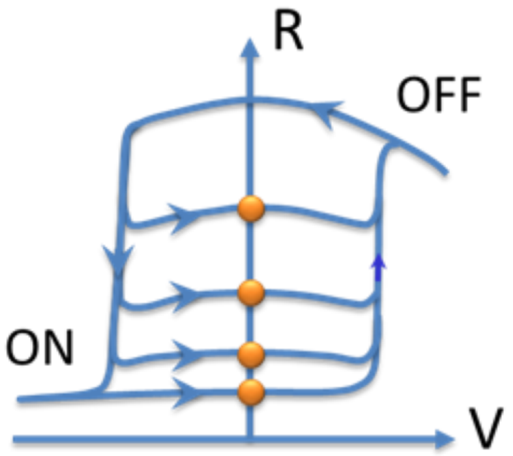

Neuromorphic Computing

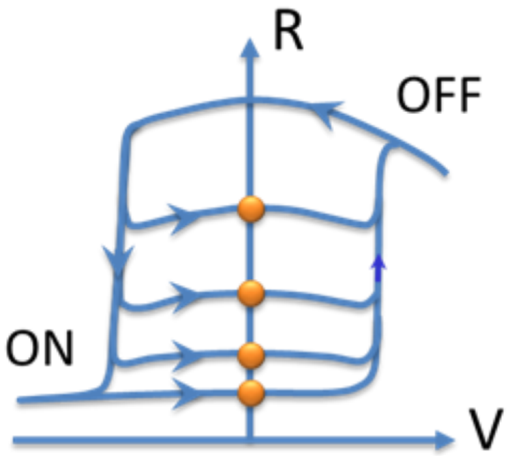

Theoretically, it can be done in pure software (neural network with dynamic adaptation of weights) but impossible in fact : for the same complexity size, 1500x slower than human brain, or Gigawatts are needed but heat dissipation is not possible.

Memory (NVRAM)

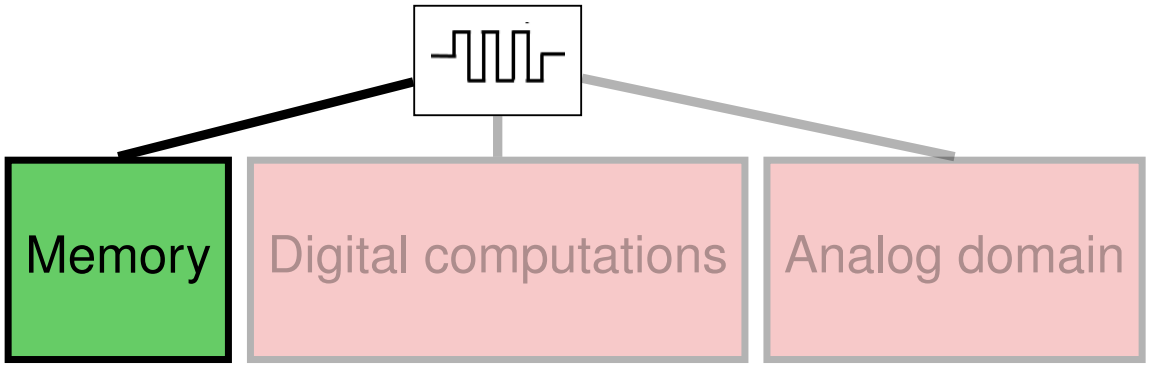

NVRAMs

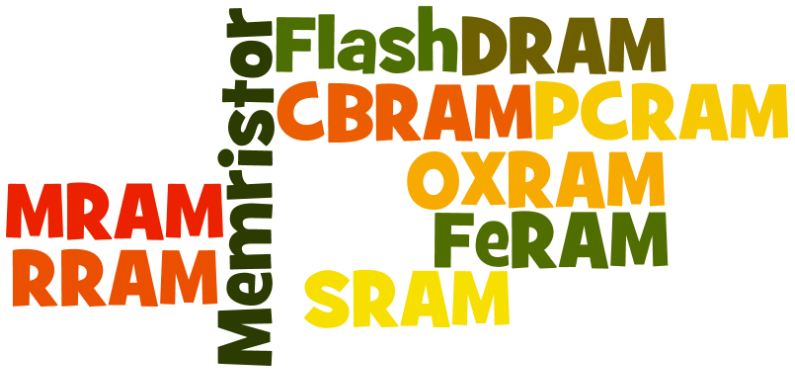

read/write times, scalability, endurance...

density, energy consumption

NVRAMs

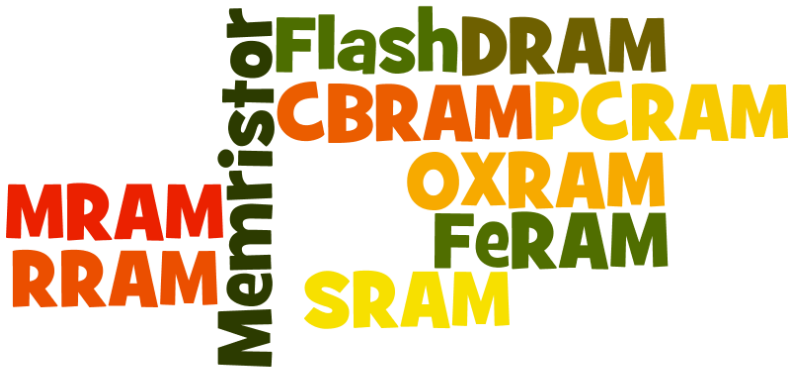

read/write times, scalability, endurance...

density, energy consumption

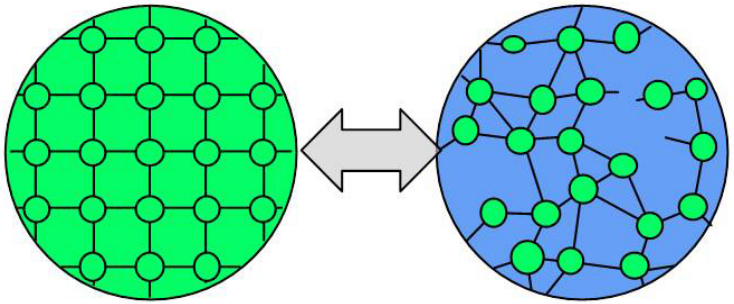

PCM

- Read time : 100ns
- Write time : 200ns

NVRAMs

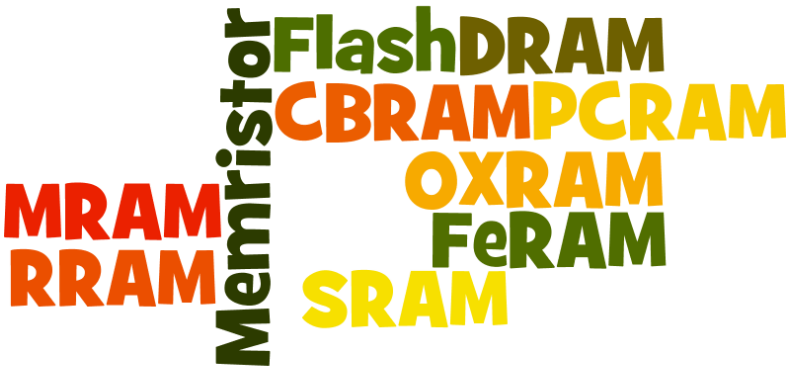

read/write times, scalability, endurance...
density, energy consumption

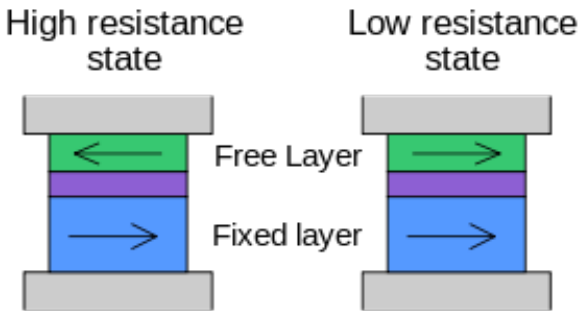

MRAM

- Read time : 10ns
- Write time : 20ns

NVRAMs

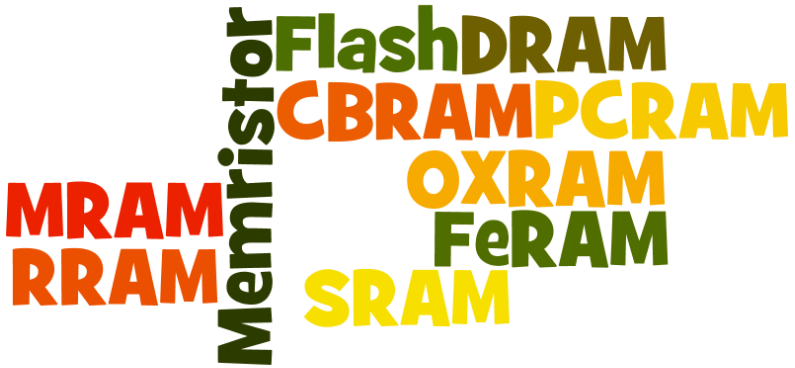

read/write times, scalability, endurance...
density, energy consumption

Tomorrow, 11h15: Breaking the Memory Bottleneck in Computing Applications with Emerging Memory Technologies: a System, Design, and Technology Perspective, **Michel Harrand, CEA**

Hardware approach ?

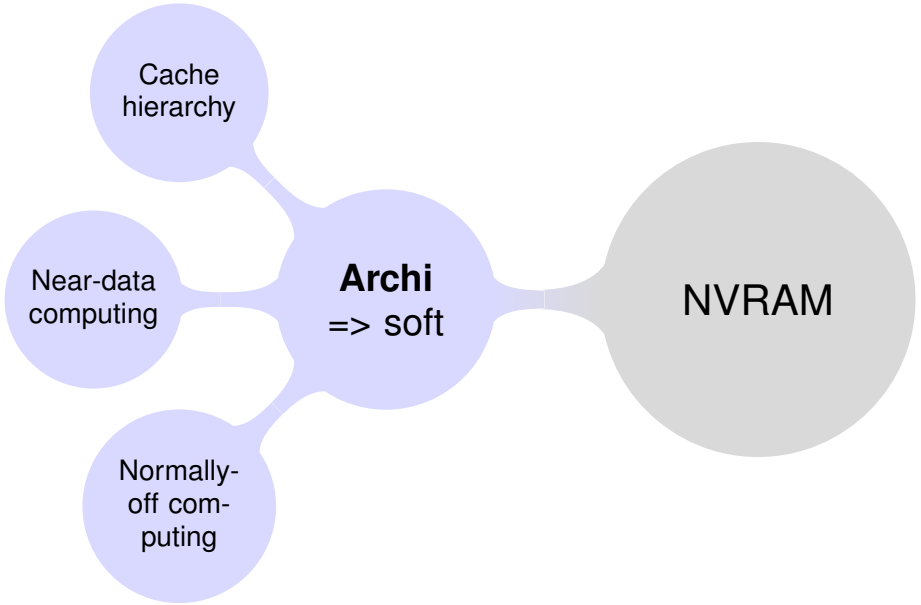

Hardware approach ?

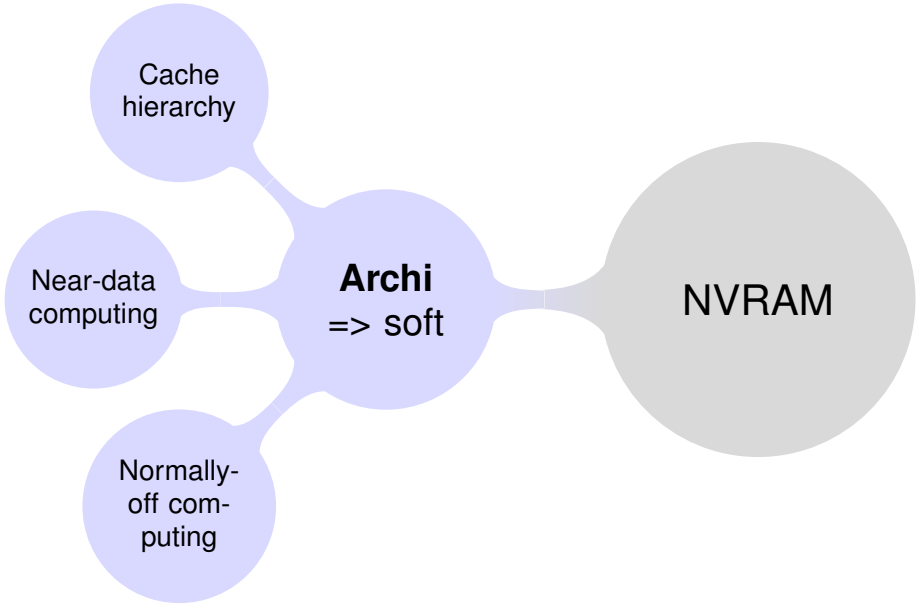

- cache: 20% of processor consumption
- ram: 25% of global consumption

Hardware approach ?

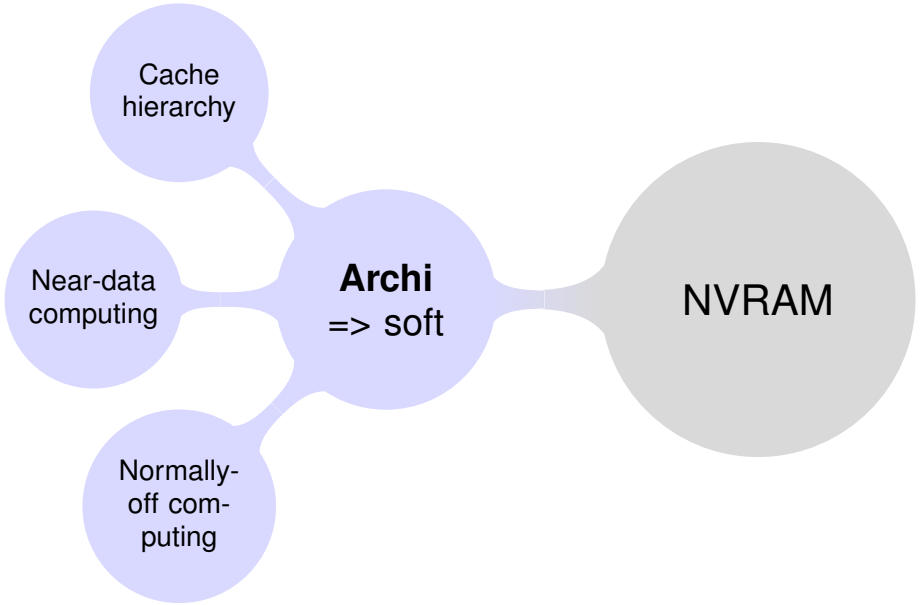

14h45: Maximize energy efficiency in normally-off system using NVRAM, **Stéphane Gros and Yeter Akgul, Evaderis**

- cache: 20% of processor consumption
- ram: 25% of global consumption

Software approach ?

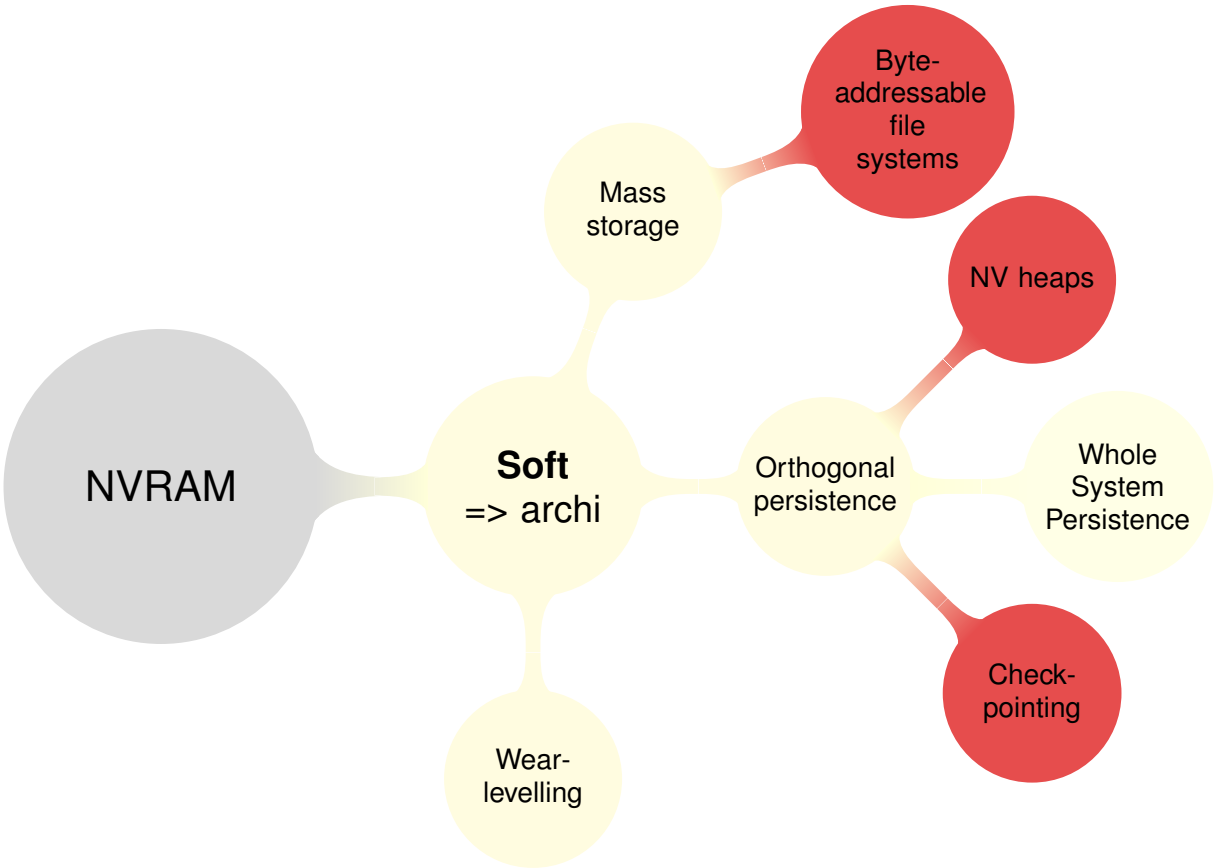

Software approach ?

mass storage:
can reach 80% of max power

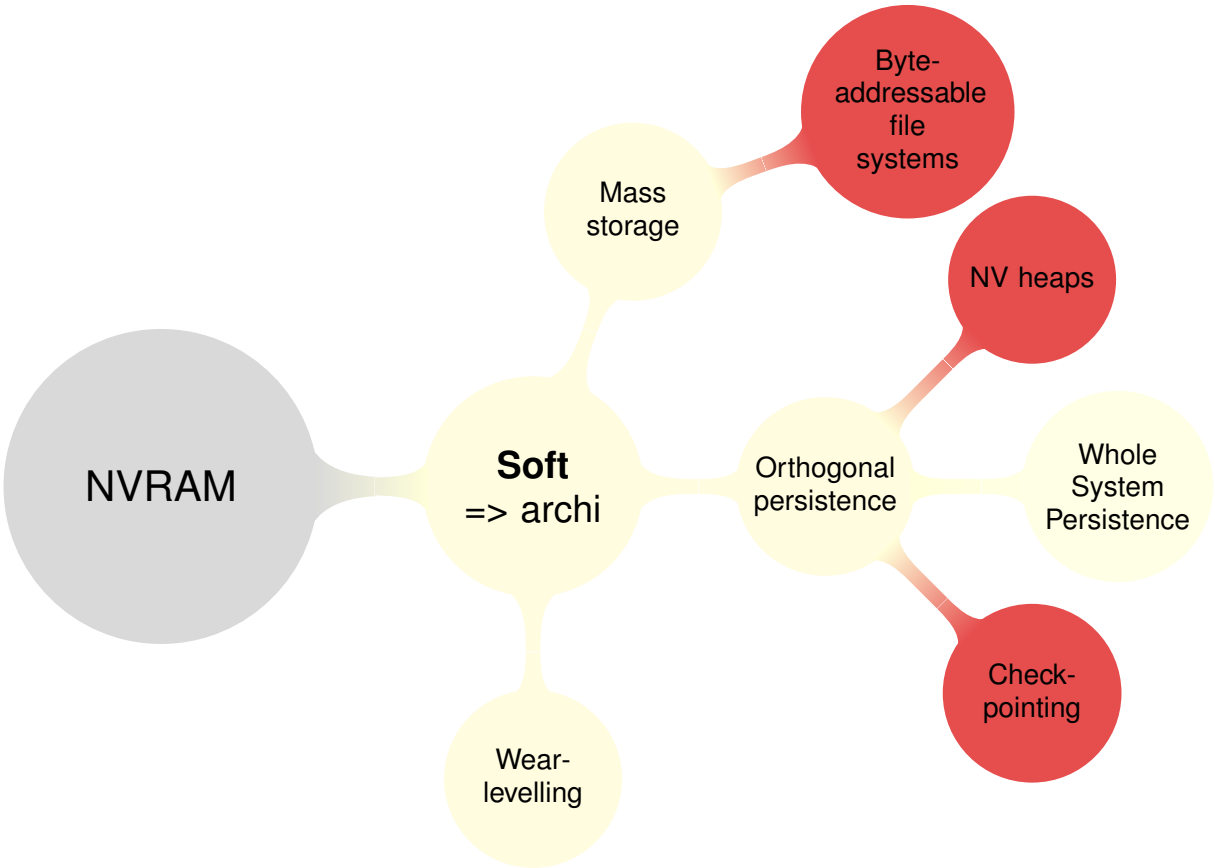

Software approach ?

Tomorrow, 9h45: Towards a single-level database architecture on byte-addressable non-volatile memory, **Ismail Oukid, SAP**

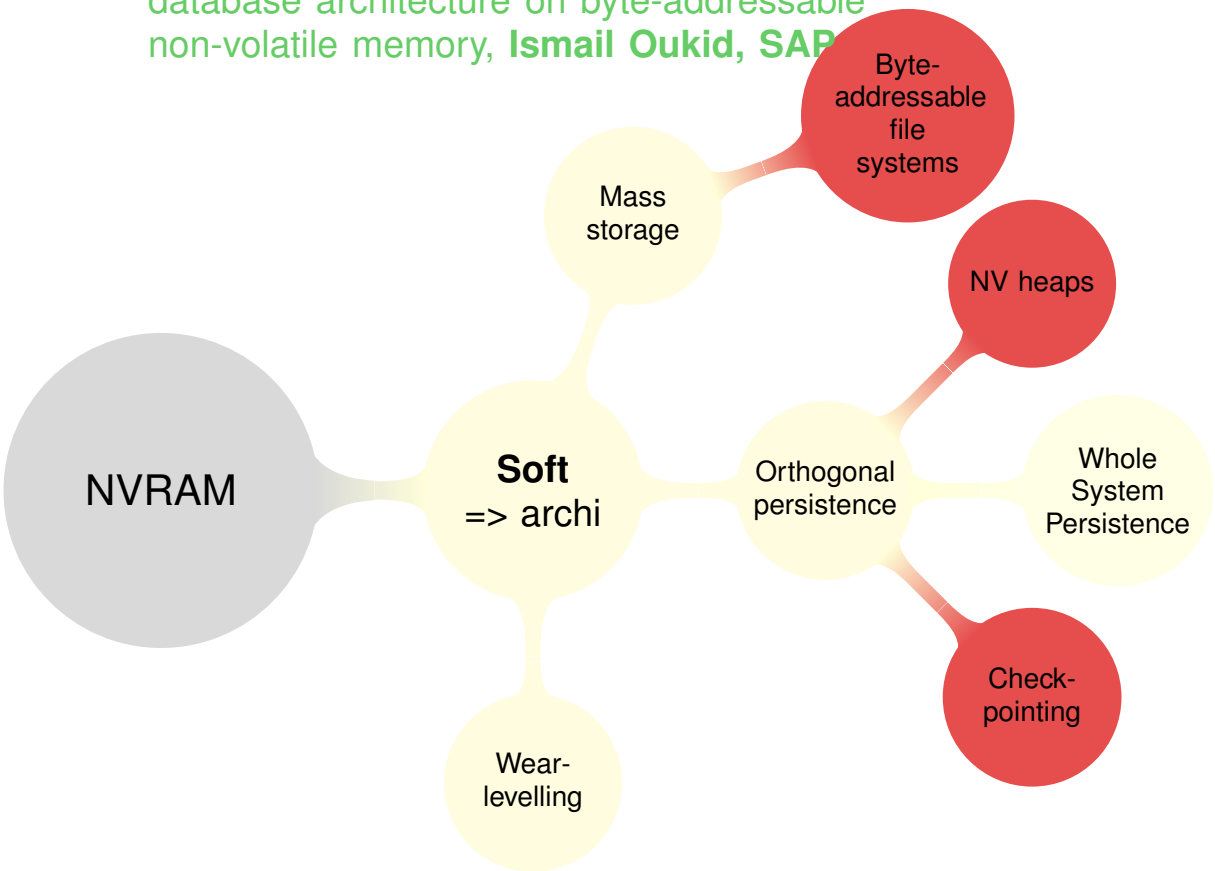

Software approach ?

transiently:
powered by energy harvesting (radio waves...)

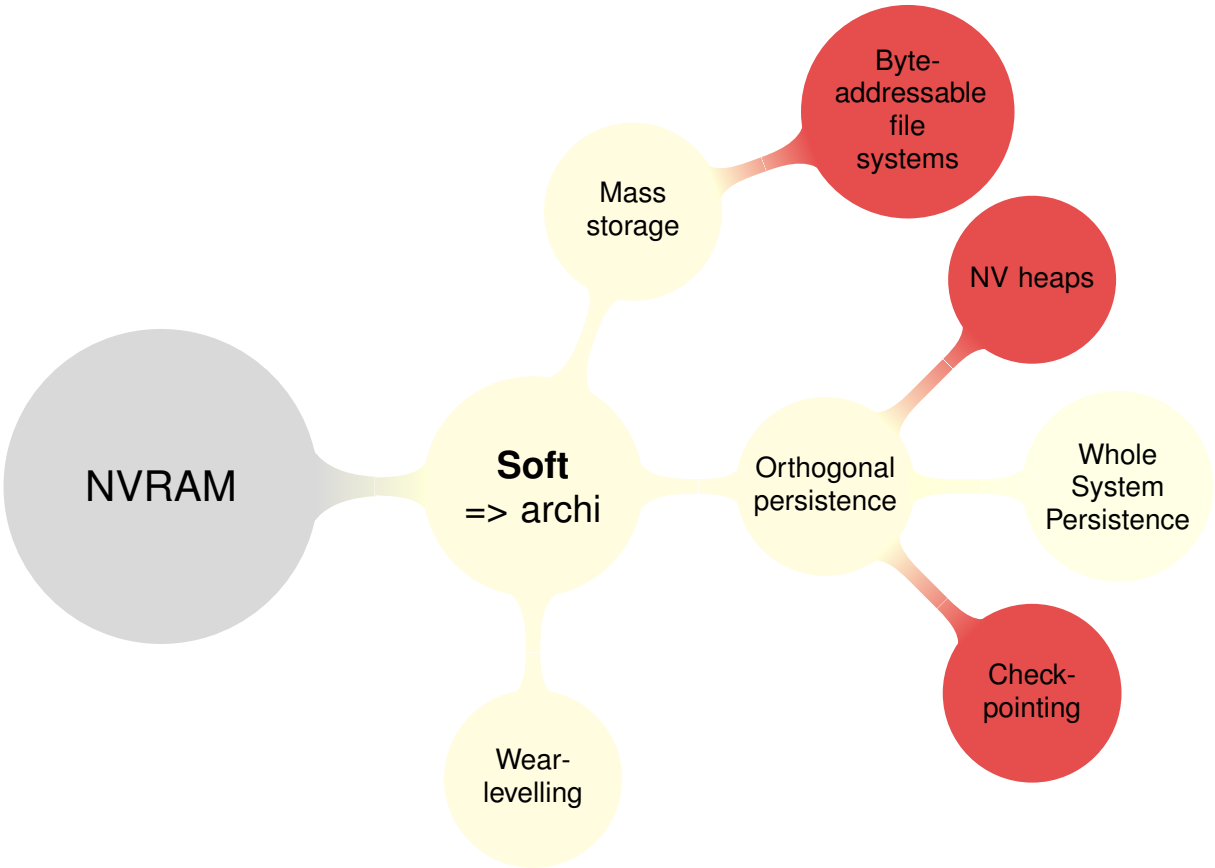

Software approach ?

transiently:
powered by energy harvesting (radio waves...)

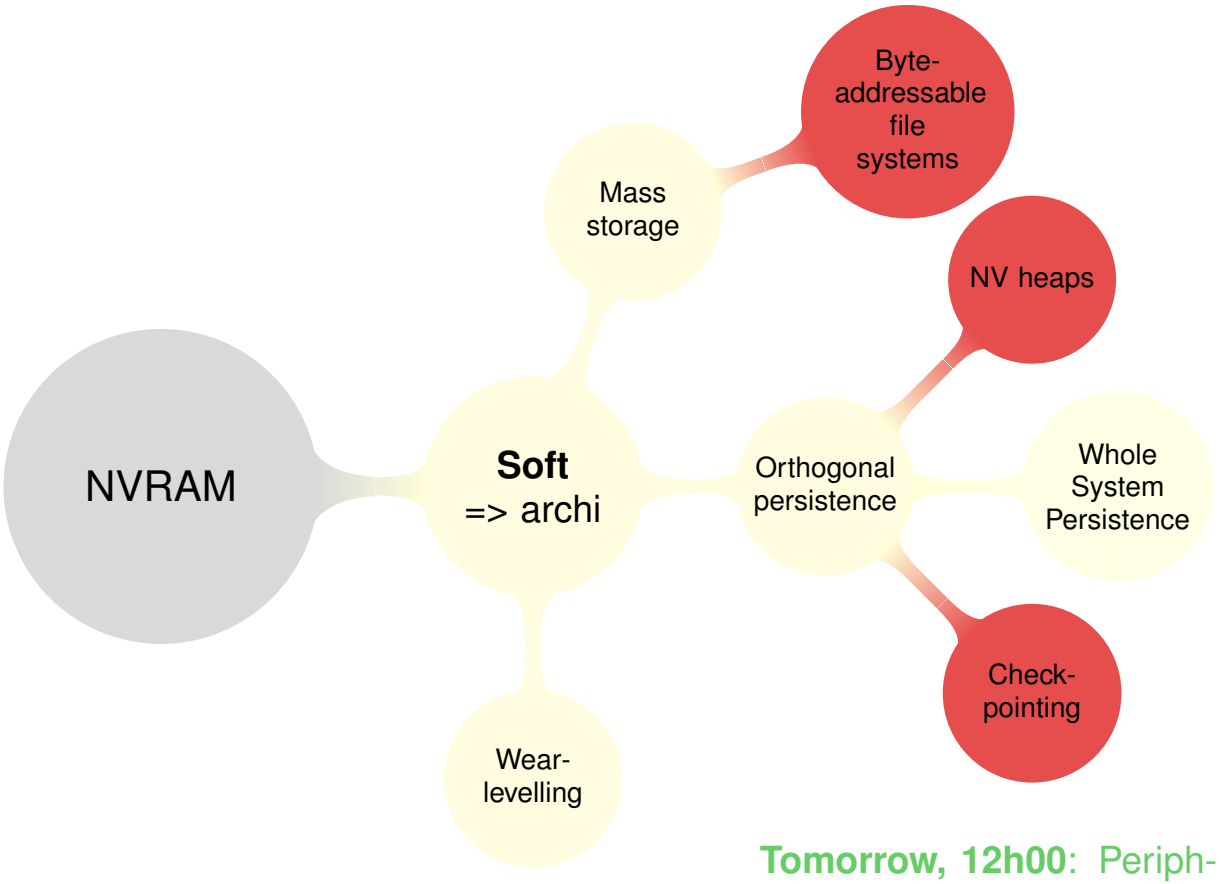

Tomorrow, 12h00: Peripheral State Persistence For Transiently Powered Systems, **Guillaume Salagnac, Inria**

Peripherals

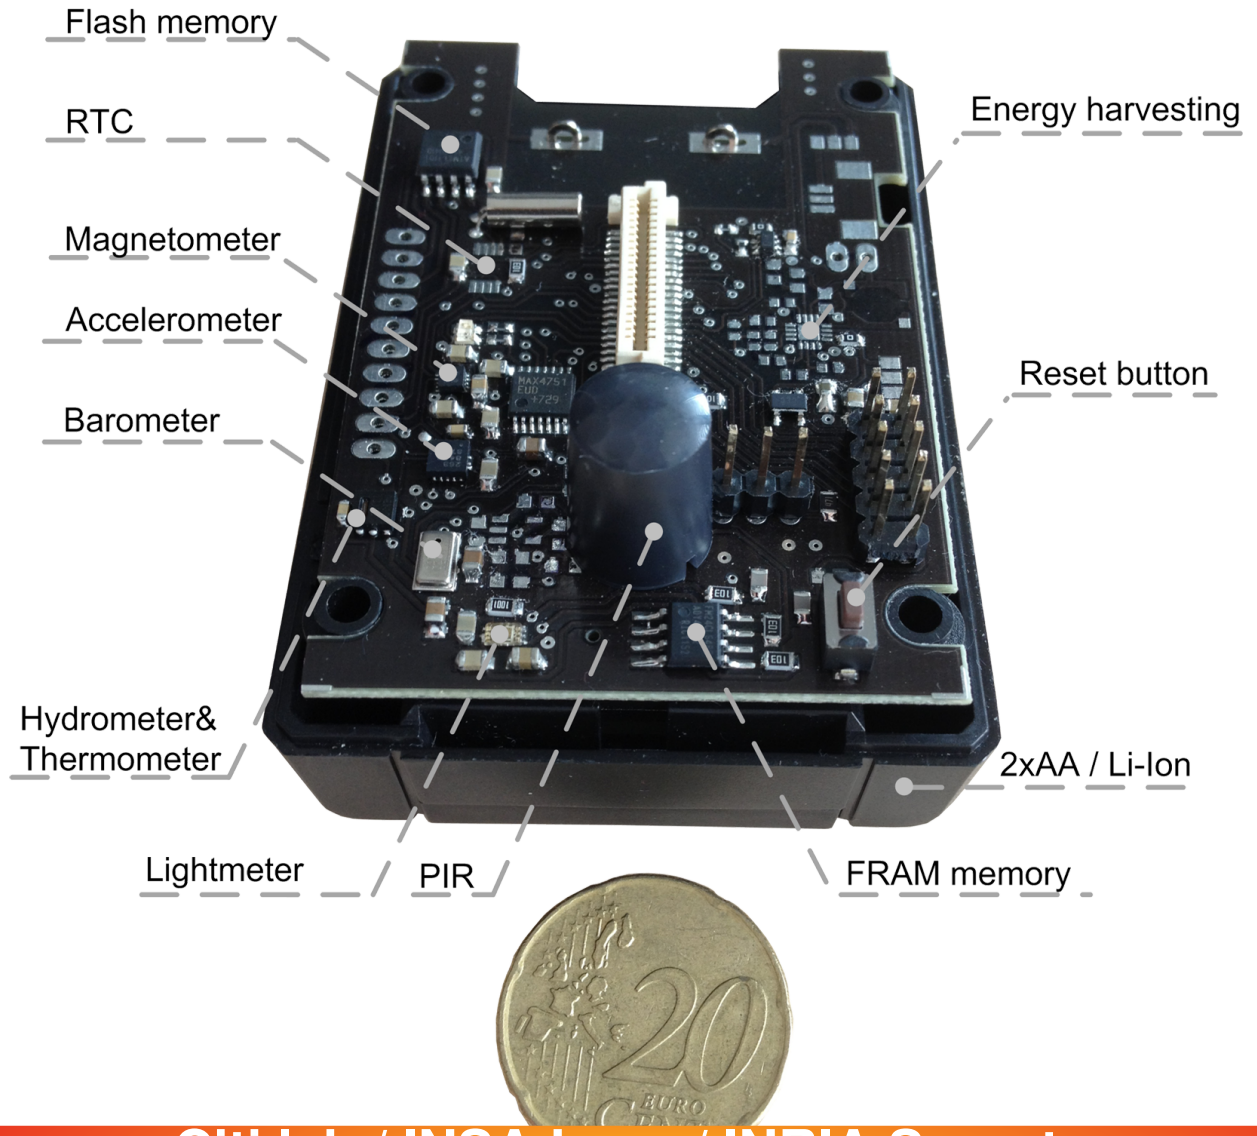

Hard-soft

- Must take into account hardware innovations at software-level of course
- Hardware benefits must be evaluated at software-level
- Pb: few platforms available

Hard-soft

- Must take into account hardware innovations at software-level of course
- Hardware benefits must be evaluated at software-level
- Pb: few platforms available
  - Extrapolate from dram accesses : No !
  - Fpga-based platforms

Hard-soft

- Must take into account hardware innovations at software-level of course
- Hardware benefits must be evaluated at software-level
- Pb: few platforms available
  - Extrapolate from dram accesses : No !
  - Fpga-based platforms
  - Simulation ;(

Hard-soft

- Must take into account hardware innovations at software-level of course
- Hardware benefits must be evaluated at software-level
- Pb: few platforms available
  - Extrapolate from dram accesses : No !
  - Fpga-based platforms
  - Simulation ;(

  - Explore memory hierarchies
  - Feedback to architects

14h00: From Embedded World to High Performance Computing using STT-MRAM, **Sohiane Senni, LIRMM**

Application fields

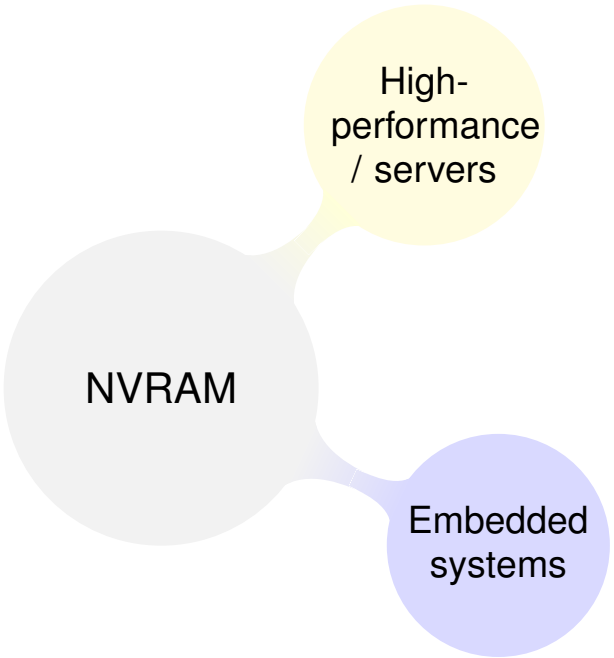

17h00: The role of NVM in the future of HPC workloads, **Patrick Demichel, HPE**

Tomorrow, 9h00: L-IOT : A flexible Platform for ultra-low power IoT, **Ivan Miro-Panades, CEA**

Compiler

- Because hardware changes
  - different time/space compromises
  - new mechanisms ? (*e.g.* checkpoint)
- Already in Intel's instruction set:
  - CLFLUSH, CLFLUSHOPT, CLWB, PCOMMIT..

16h15: NVRAM: New Opportunities for Compilers, **Erven Rohou, Inria**

Operating system

- Support for hardware specifics
- Data placement
- Paging vs. file system

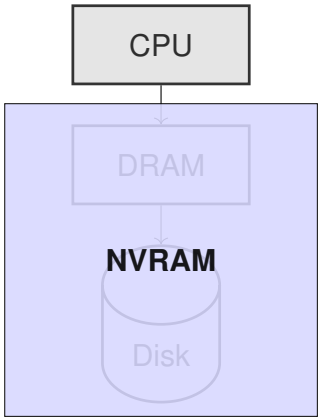

Operating system

- Support for hardware specifics
- Data placement
- Paging vs. file system
- Installation and launch

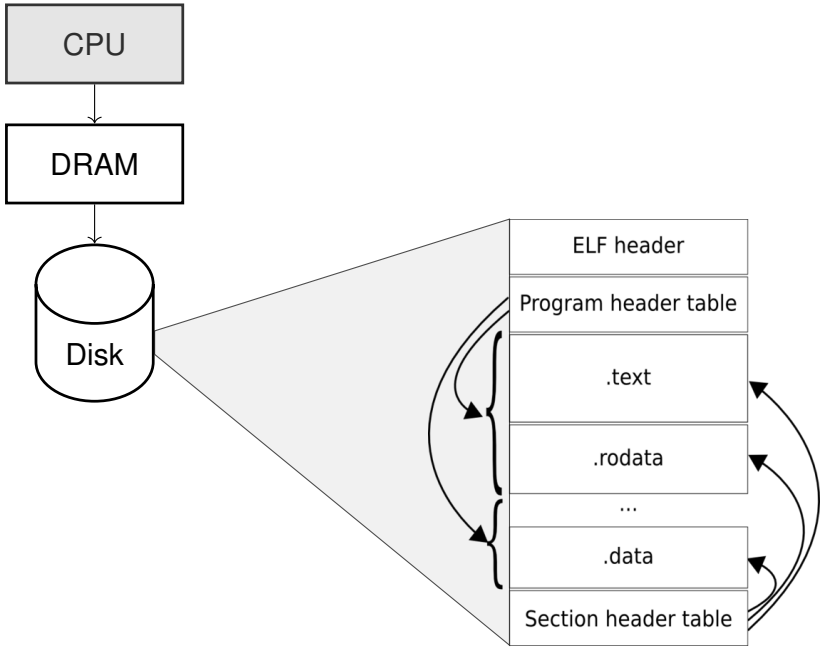

Operating system

- Support for hardware specifics
- Data placement
- Paging vs. file system
- Installation and launch

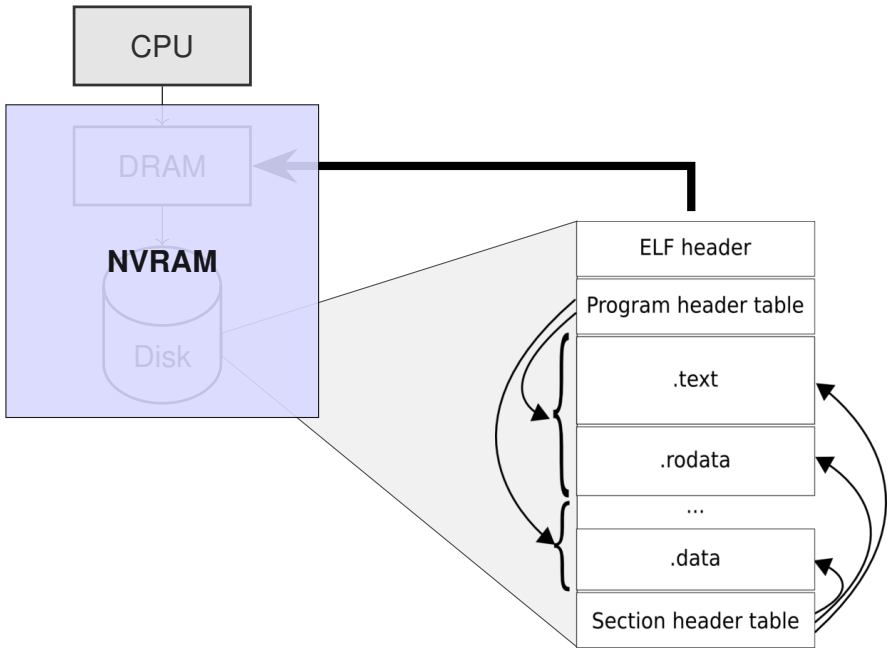

Operating system

- Support for hardware specifics
- Data placement
- Paging vs. file system
- Installation and launch

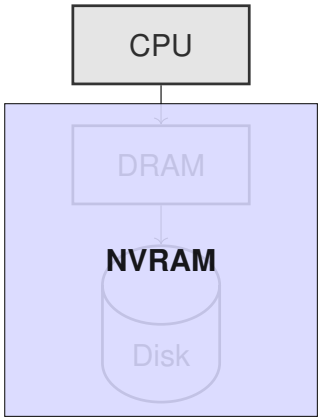

- No reboot...
  - bug ?
  - reset ? checkpoint ?

Programming model

- Persistent data
- Transiently-powered systems
- New programming paradigm needed
  - fpga++
  - neural computing

ZEP project

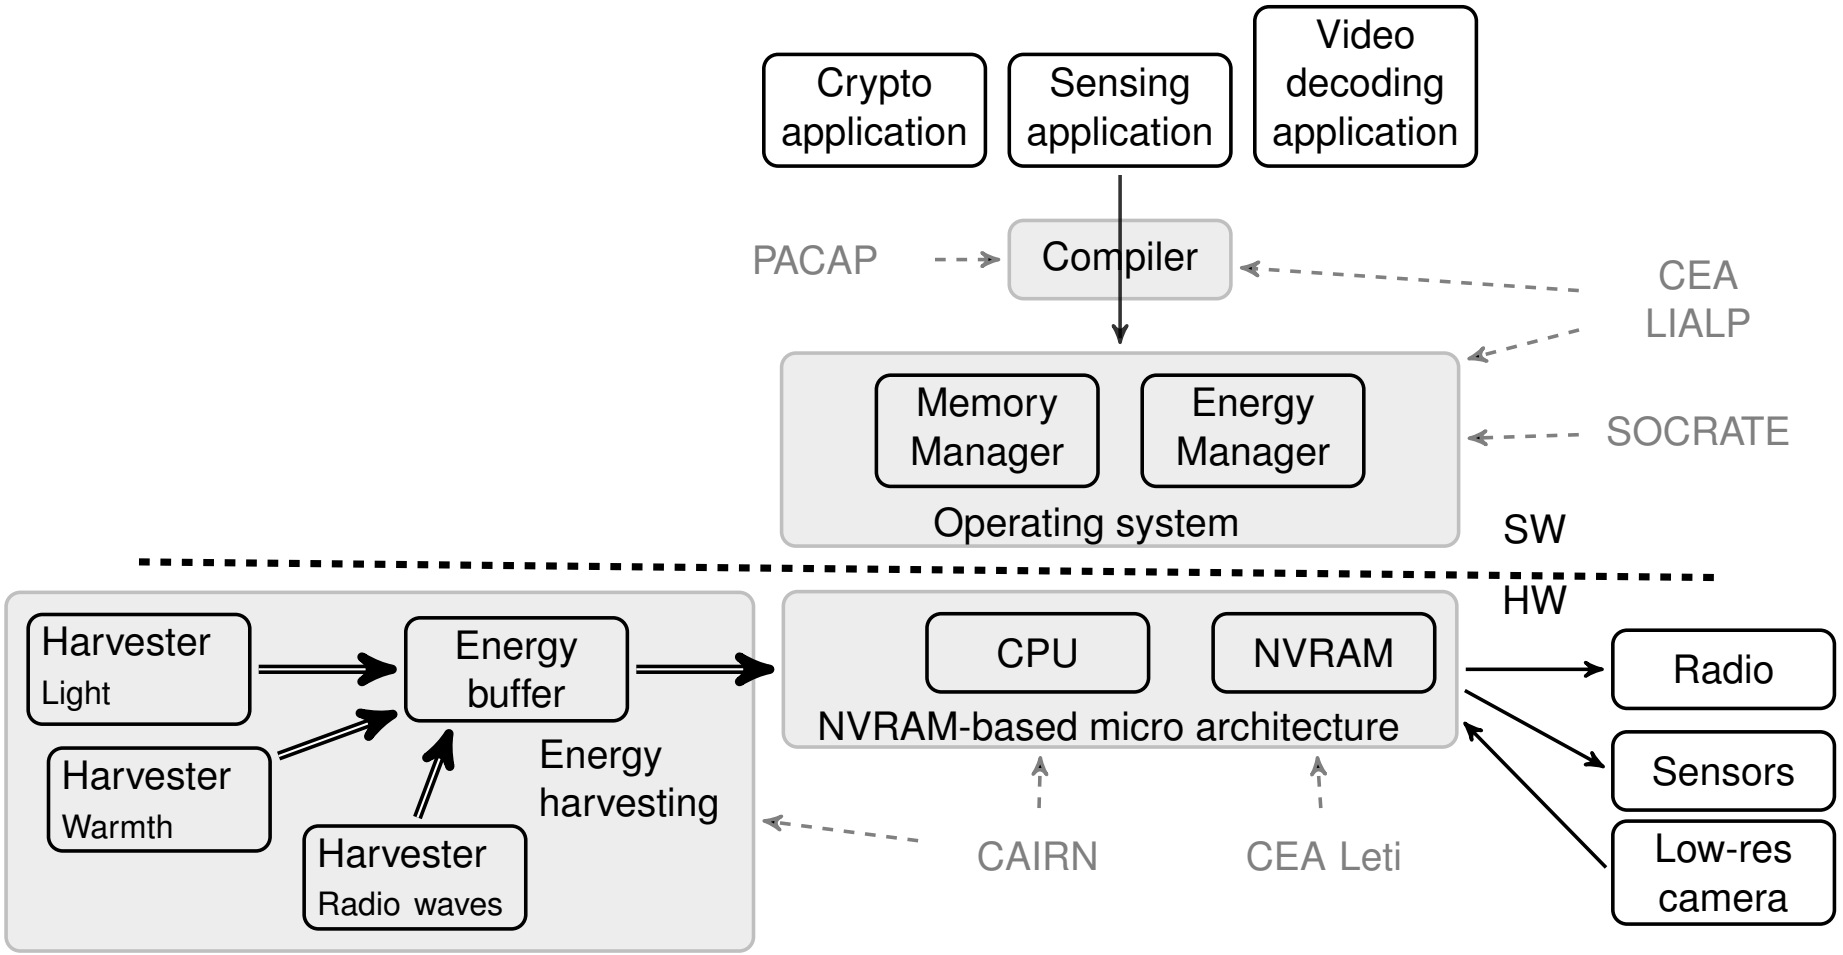
Crypto application
Sensing application
Video decoding application
PACAP
Compiler
CEA LIALP
Memory Manager
Energy Manager
SOCRATE
Operating system
SW
HW
Harvester Light
Harvester Warmth
Harvester Radio waves
Energy buffer
Energy harvesting
CPU
NVRAM
NVRAM-based micro architecture
CAIRN
CEA Leti
Radio
Sensors
Low-res camera

Program

https://project.inria.fr/iplzep/events/workshop-on-nvram-paris-may-29-30th/

Today

- **14:00** From Embedded World to High Performance Computing using STT-MRAM – Sophiane Senni, LIRMM
- **14:45** Maximize energy efficiency in normally-off system using NVRAM – Stéphane Gros and Yeter Akgul, Evaderis
- **15:30** === Pause ===
- **16:15** NVRAM: New Opportunities for Compilers – Erven Rohou, Inria
- **17:00** The role of NVM in the future of HPC workloads – Patrick Demichel HP
- **19:00** Diner

Tomorrow

- **08:45** Welcome
- **09:00** L-IOT: A flexible Platform for ultra-low power IoT – Ivan Miro-Panades, CEA
- **09:45** Towards a single-level database architecture on byte-addressable non-volatile memory – Ismail Oukid, SAP
- **10:30** === Pause ===
- **11:15** Breaking the Memory Bottleneck in Computing Applications with Emerging Memory Technologies: a System, Design, and Technology Perspective – Michel Harrand, CEA
- **12:00** Peripheral State Persistence For Transiently Powered Systems