Ecole IN2P3: Electronique pour la physique

Référence CENSEA17100052-01

Présenté le : 1/04/2019

Gestion des versions

Indice	Evolution	Date
00	Création	15/09/2018
01	Modification pour l'ecole des technique de base des deteceturs	01/04/2019

Validation du Document

Action	Nom	Fonction/Entité	Date	Visa
Rédigé par	F. Druillole	Responsable CENBG/SEA	14/09/2018	
Vérifié par				
Vérifié par				
Approuvé par				
Approuvé par				

Documents de référence Projet

Document	Référence	Emetteur	Date
Bruit de fond et mesures – Aspects Théoriques	R310V2	Techniques de l'ingénieur	18/12/2006
Bruit de fond et mesures – Mesures et applications en conception	R311V1	Technique de l'ingénieur	18/12/2006
Op Amp Noise Theory ans Applications	SLOA082	Texas Istruments	
NOISE ANALYSIS OF FET TRANSIMPEDANCE AMPLIFIERS	SBOA060	BurrBrown	02/1994
Presentation electronique multi-detecteur		Laurent Leterrier LPCCaen	

ECOLE IN2P3 : Techniques de base des acquisitions multi-détecteurs

l'électronique du détecteur à la mesure
III – Prospectives & Synthèse

CARGESES, du 1er au 5 Avril 2019

Un Système d'acquisition : Dominant Design

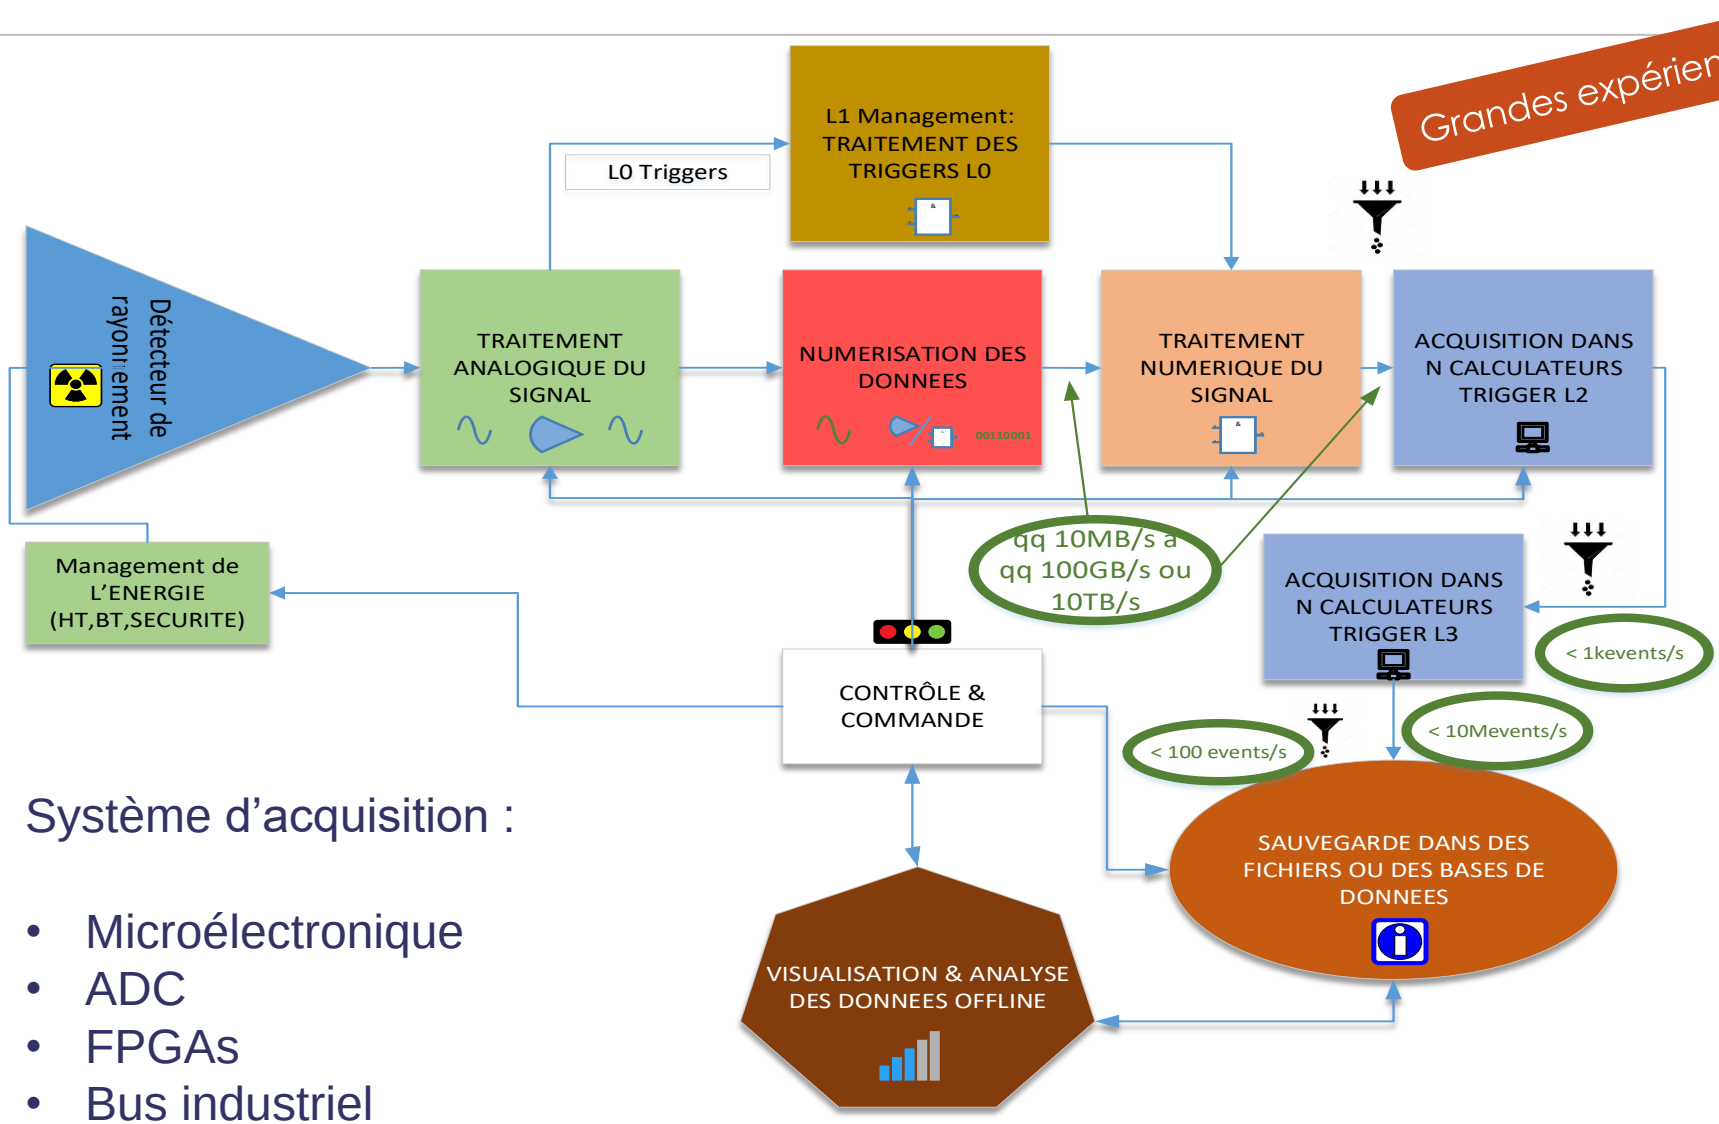

Système d'acquisition :

- Microélectronique
- ADC
- FPGAs
- Bus industriel
- Ferme de PC

Les Technologies émergentes pour les DAQ

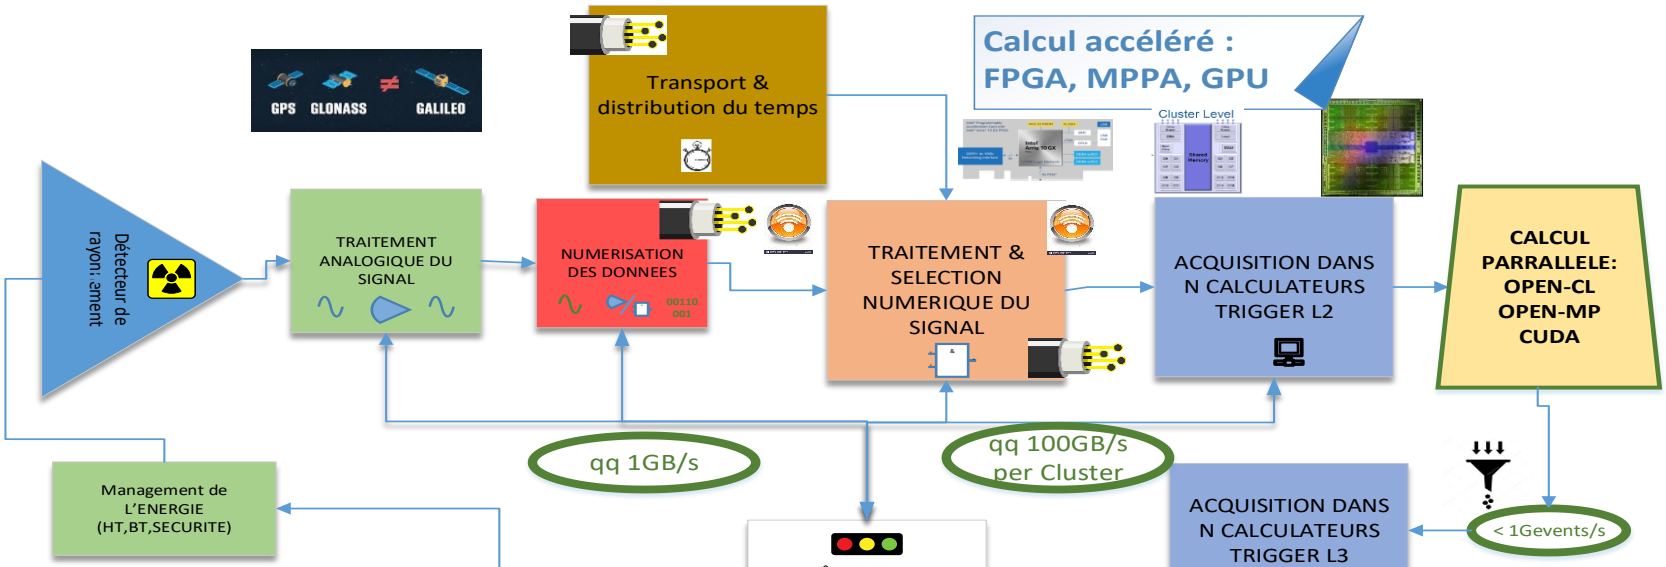

Les points de rupture :

- Transmission Optique et RF des FE
- Distribution d'horloges précises (qq ps)
- Calcul accéléré et massivement parallèle
- Algorithmes IA (*embarqués*)

Un ASIC FE Système : CATIROC

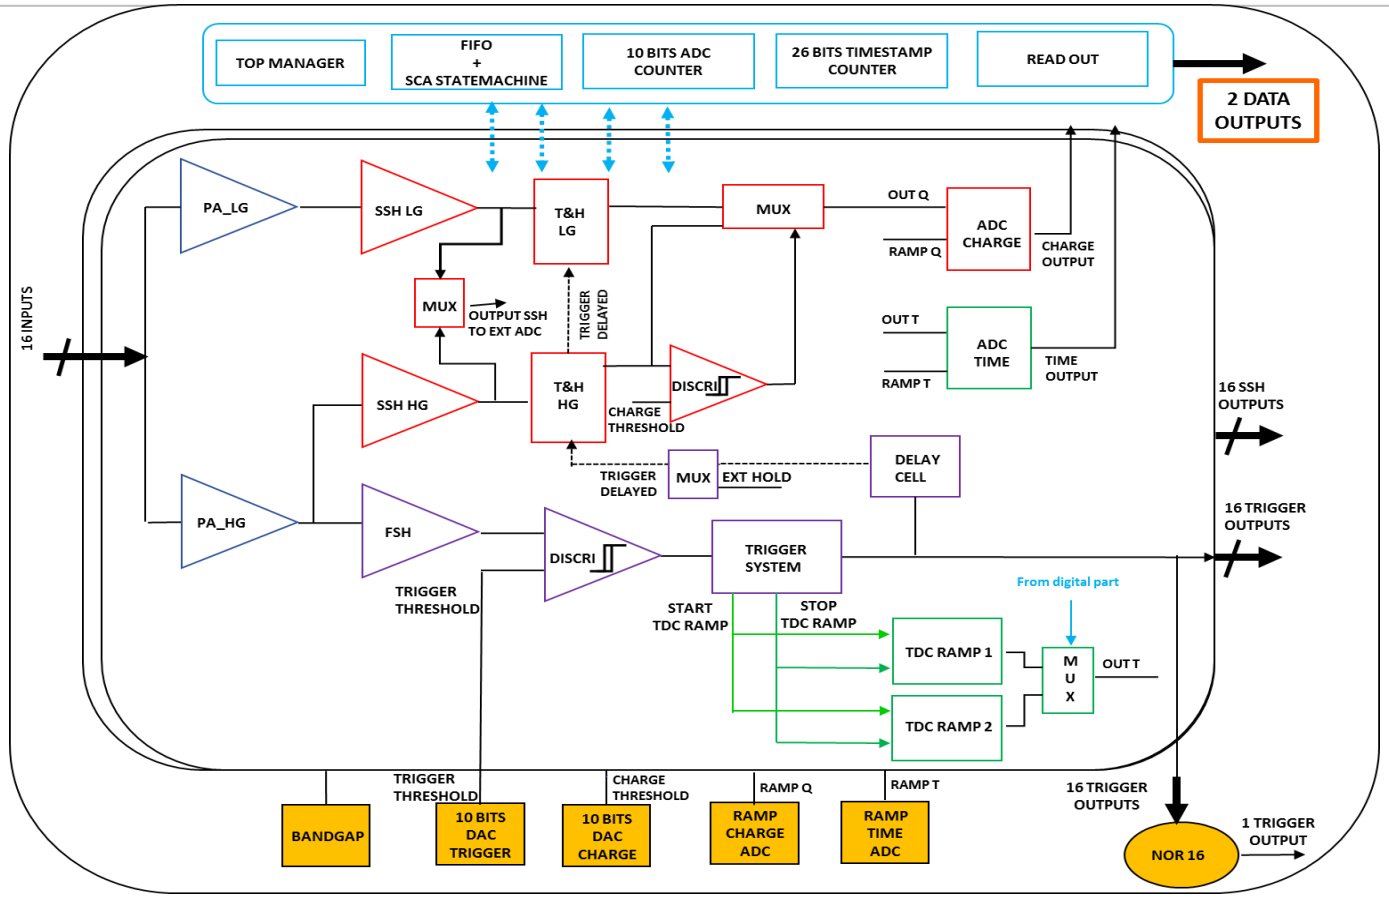

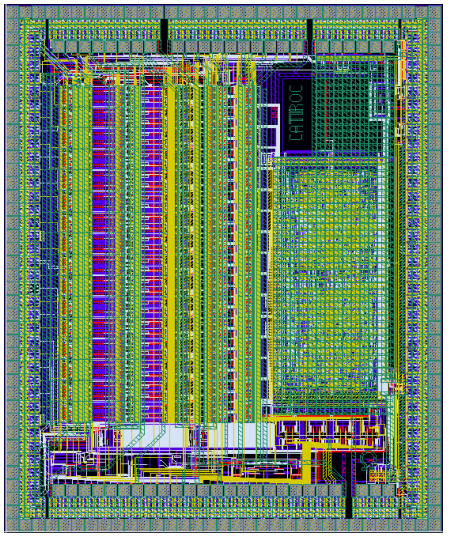

- LG/HG
- SHAPER
- DISCRILINATEUR
- TAC Fine Time
- Compteur coardse Time
- ADC Wilkinson (double rampe)

➔ ADC intégré dans l'ASIC

Transmission RF des front-end

A 3D VIEW OF RF COMMUNICATION STANDARDS

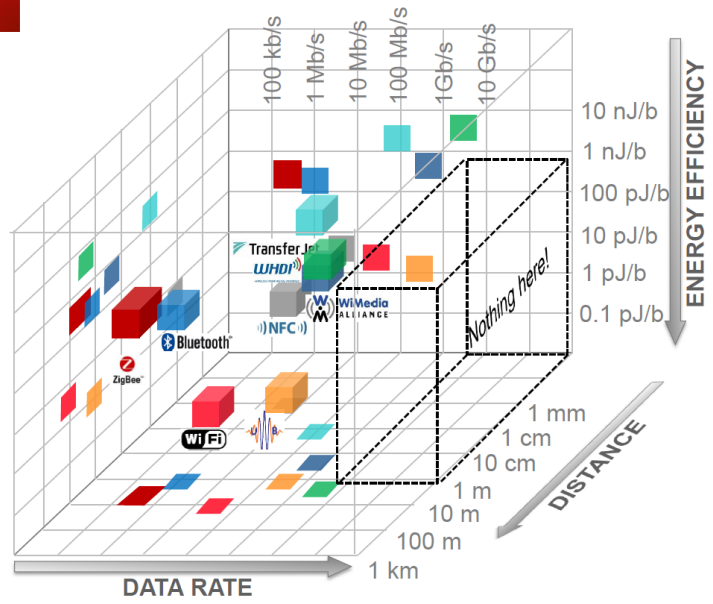

RF COMMUNICATIONS FUNDAMENTALS

AN OVERVIEW OF MODULATION

Digital modulations:

1. **Baseband digital message signal**: $m(t)$
2. **Analog sinusoidal carrier signal**:
   A. Carrier signal: $A_c \cos(2\pi f_c t + \phi_c)$
3. **ASK: Amplitude Shift Keying**.
   A. Message signal changes the carrier's **amplitude** : $A_i(t)$.
4. **FSK: Frequency Shift Keying**.
   A. Message signal changes the carrier's **frequency** : $f_i(t)$.
5. **PSK: Phase Shift Keying**.
   A. Message signal changes the carrier's **phase** : $\phi_i(t)$.

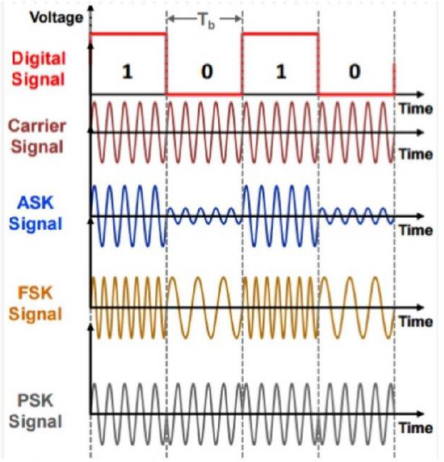

Une place à prendre :

Transmission inter-carte pour remplacer les câbles

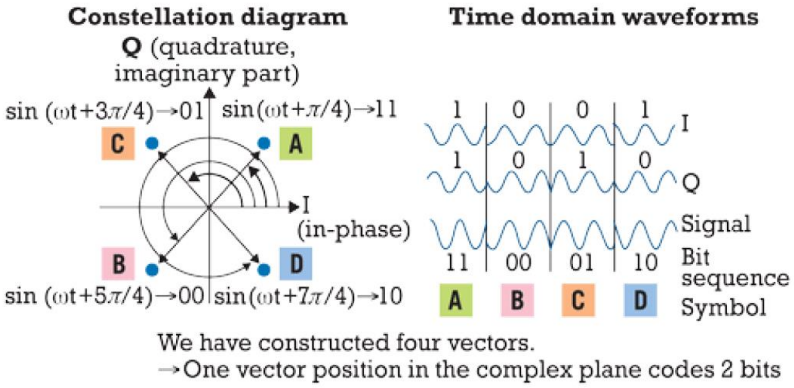

RF COMMUNICATIONS FUNDAMENTALS

AN OVERVIEW OF MODULATION

Complex modulation schemes

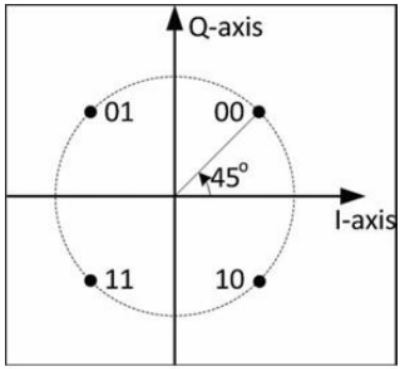

Figure 1. Constellation View of QPSK

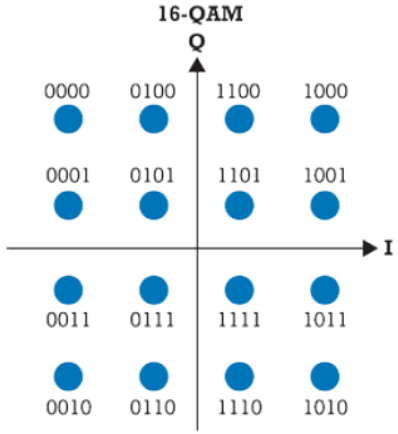

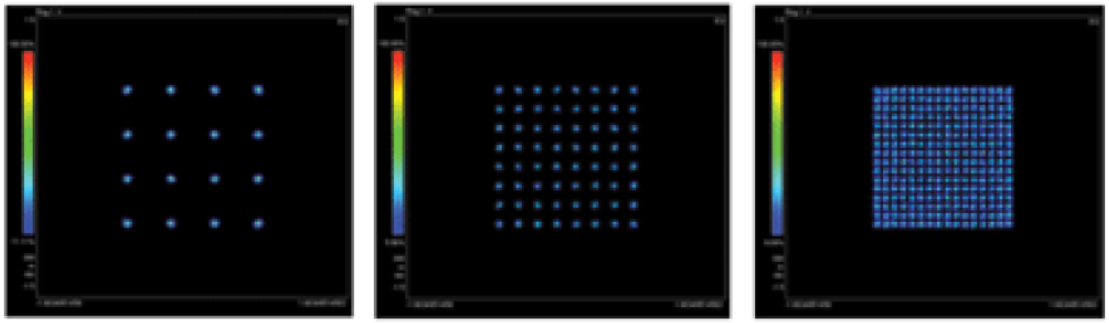

16-QAM
4 bits/symbol

64-QAM
6 bits/symbol

256-QAM
8 bits/symbol

Transmission RF des front-end

FROM RESEARCH TO INDUSTRY
cea tech

www.cea.fr

leti & list

Multi-Gigabit Wireless Data Transfer for High Energy Physics Applications

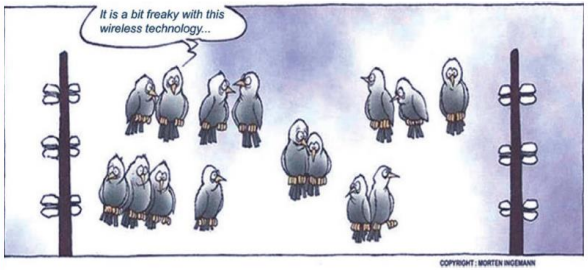

cedric.dehos@cea.fr

WADAPT Collaboration

- Wadapt: "**Wireless Allowing Data And Power Transmission**"
- Objectives:
  - Definition of the needs of data connectivity for particle-physics detectors
  - Evaluation of the wireless technologies for data and power transfer
  - Hardening, specific design and prototyping
- Consortium
  - **CERN**, European Organisation for Nuclear Research, Geneva, Switzerland
  - **CEA/DSM/IRFU**, Gif-sur-Yvette, France
  - **CEA/LETI**, Grenoble, France
  - **University of Heidelberg**, Germany
  - **University of Uppsala**, Sweden
  - **University of Bergen**, Norway
  - **Argonne National Laboratory**, Argonne, USA
  - **Gangneung National University**, Korea

Context: massive cable plant

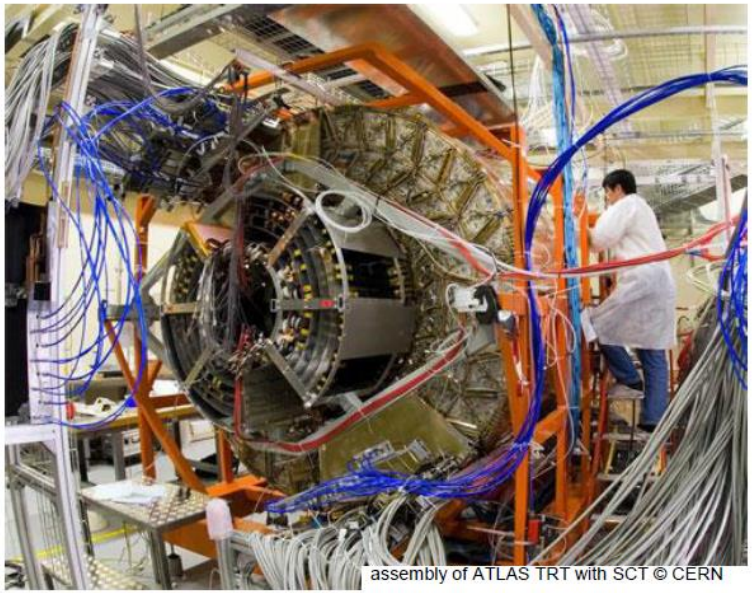

assembly of ATLAS TRT with SCT © CERN

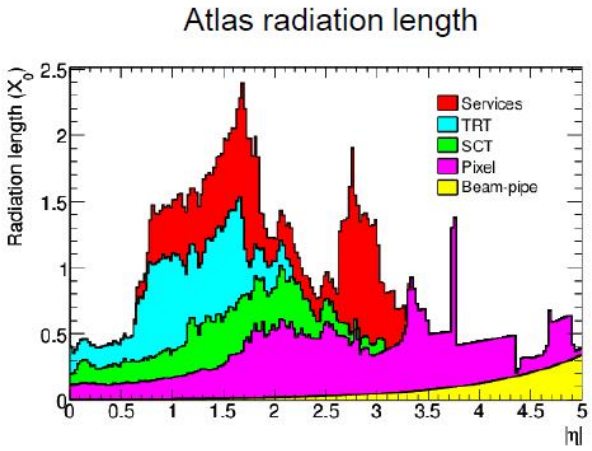

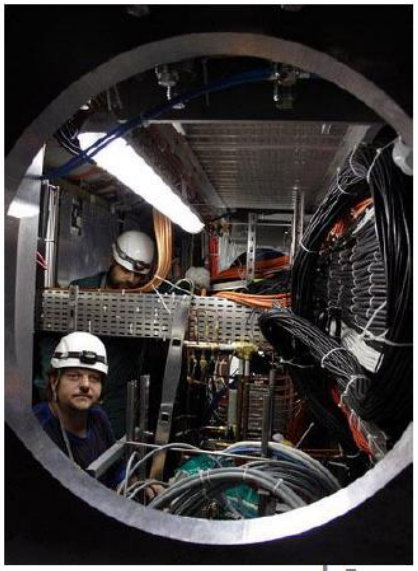

- **Impact on the measurements**
  - Multiple scattering and nuclear interactions
  - Dead-zone areas
- **Impact on the installation and the operation**
  - Cables and connectors are fragile
  - Cable path is not so flexible
  - Design constraints

Why Wireless ?

- Minimize material budget of cables/connectors
  Limited radiation length because of massive services in region between Barrel and Disks
- Axial readout induce important latencies
  Direct communication between layers (radial readout)
- More flexible transceiver placement
- Point-to-Multipoint links, interlayer intelligence
- Data follows event topology enabling fast triggering

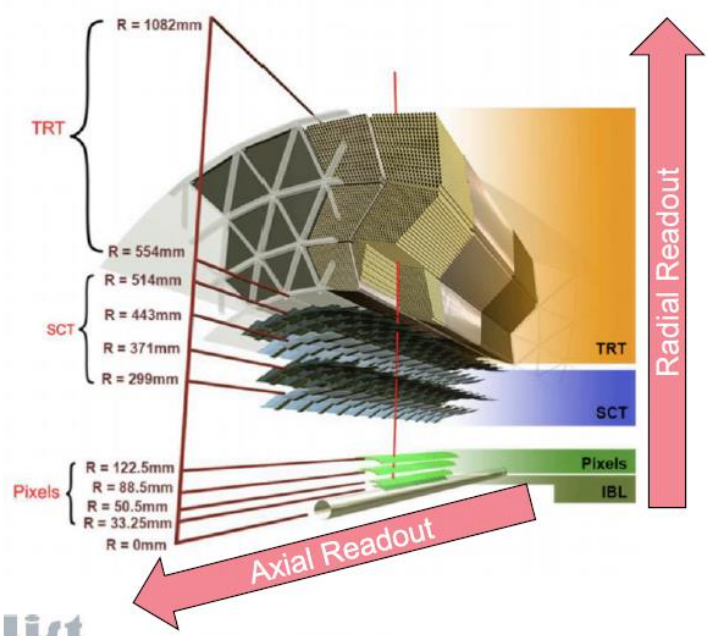

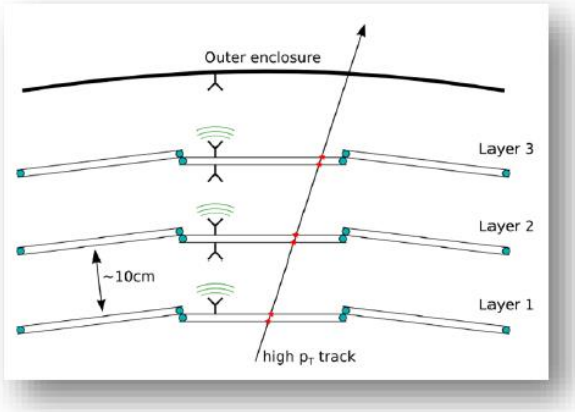

Wireless readout principle
(R. Brenner, Uppsala Univ)

cea tech

Feasibility studies, CEA Leti

- PRBS 8b/10b
- 60GHz TRX package on test board
- 9dB horn antennas
- 3cm range
- Oscilloscope eye and jitter analysis

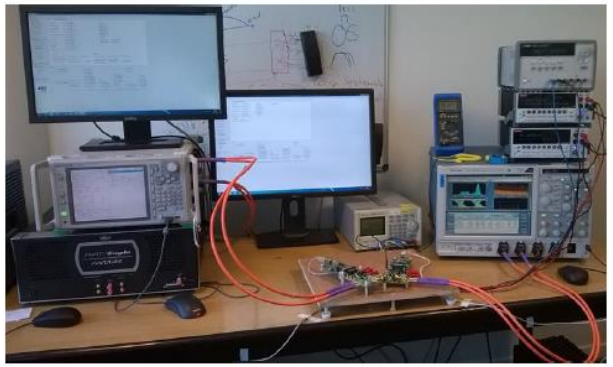

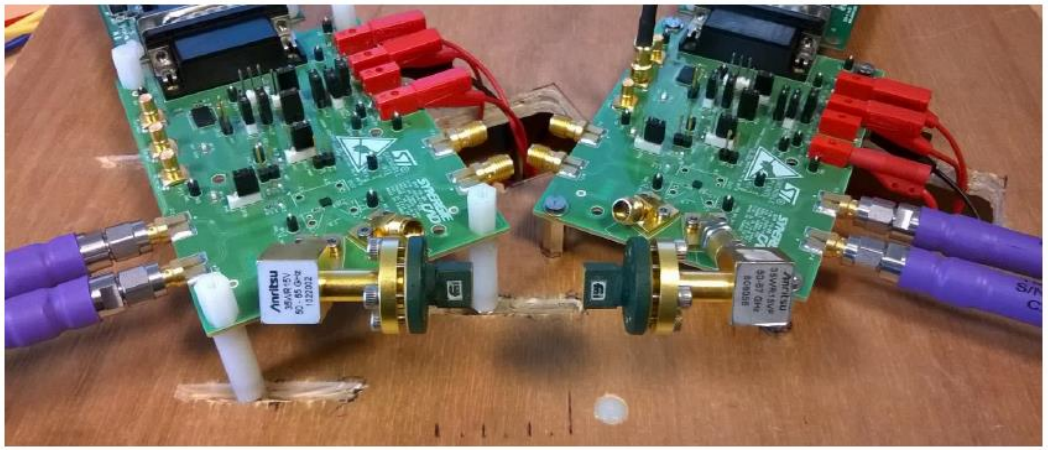

Context: … but who is feeding the network ?

In 2016:
- 1'200'000 petabytes have been exchanged worldwide
- 49 petabytes have been produced by LHC experiments

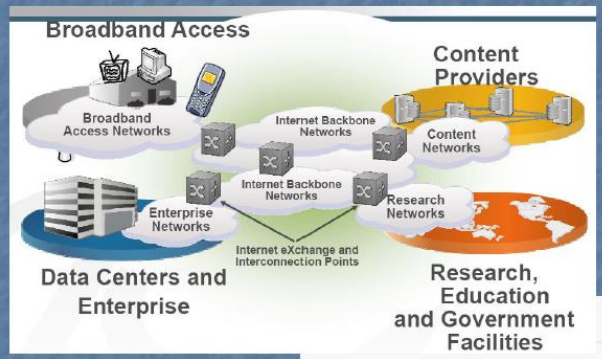

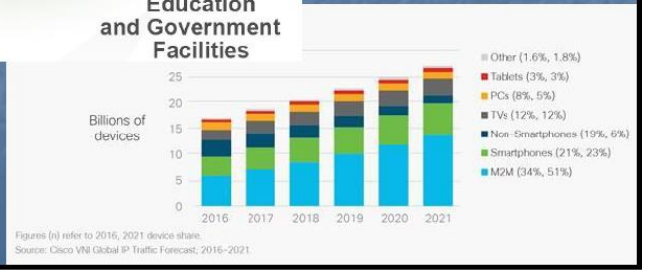

12 Nov 2018

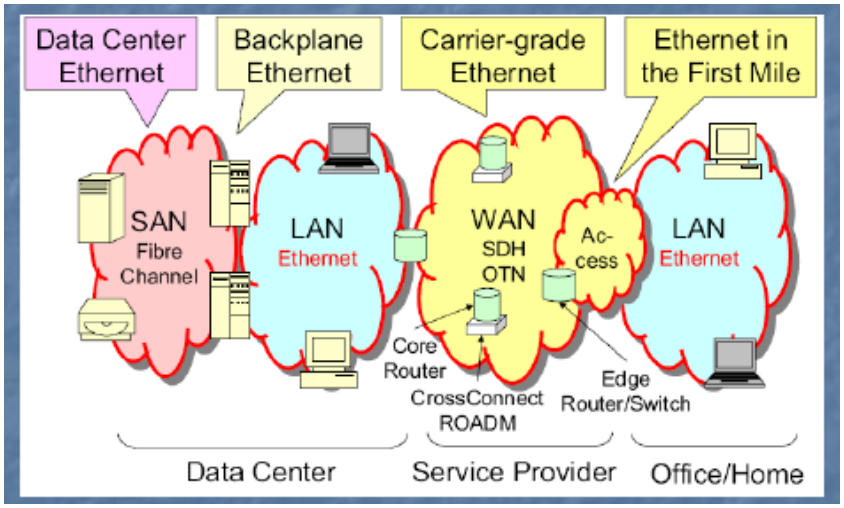

Bit Rates vs Standards

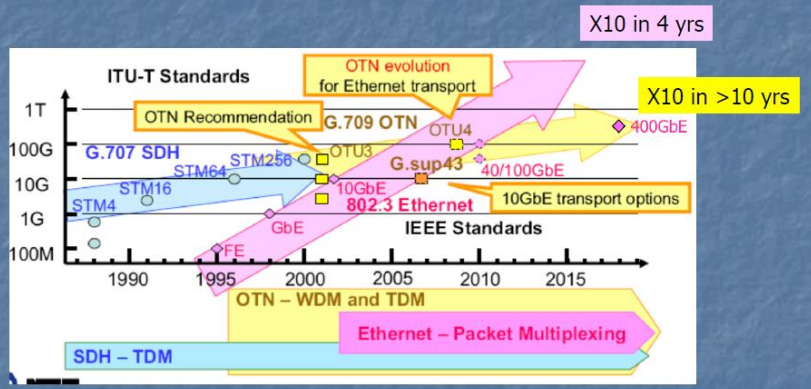

Transmission Optique

IEEE802.3ba Ethernet Roadmap

https://ethernetalliance.org/the-2018-ethernet-roadmap/

Ethernet Interfaces and Nomenclature

	Electrical Interface	Backplane	Twinax Cable	Twisted Pair (1 Pair)	Twisted Pair (4 Pair)	MMF	500m PSM4	2km SMF	10km SMF	40km SMF	80km SMF
10BASE-		T1S?		T1S/T1L							
100BASE-				T1							
1000BASE-				T1	T						
2.5GBASE-		KX		T1S?	T						
5GBASE-		KR		T1S?	T						
10GBASE-				T1S?	T						
25GBASE-	25GAUI	KR	CR/CR-S		T	SR			LR	ER	
40GBASE-	XLAUI	KR4	CR4		T	SR4/eSR4	PSM4	FR	LR4	ER4	
50GBASE-	LAUI-2/50GAUI-2 50GAUI-1	KR	CR			SR		FR	LR	ER	
100GBASE-	CAUI/10 CAUI-4/100GAUI-4 100GAUI-2 100GAUI-1	KR4 KR2 KR1	CR10 CR4 CR2 CR1			SR10 SR4 SR2	PSM4 DR	10X10 CWDM4 CLR4 100G-FR	LR4 4WDM-10 100G-LR	ER4 4WDM-40 ?	?
200GBASE-	200GAUI-4 200GAUI-2	KR4 KR2	CR4 CR2			SR4	DR4	FR4	LR4	?	?
400GBASE-	400GAUI-16 400GAUI-8 400GAUI-4	KR4	CR4			SR16	DR4	FR8 400G-FR4	LR8 ?	?	?

Physical layer	40 Gigabit Ethernet	100 Gigabit Ethernet
Backplane	n.a.	100GBASE-KP4
Improved Backplane	40GBASE-KR4	100GBASE-KR4
7 m over twinax copper cable	40GBASE-CR4	100GBASE-CR10 100GBASE-CR4
30 m over "Cat.8" twisted pair	40GBASE-T	
100 m over OM3 MMF	40GBASE-SR4	100GBASE-SR10 100GBASE-SR4
125 m over OM4 MMF[34]		
2 km over SMF, serial	40GBASE-FR	100GBASE-CWDM4[97]
10 km over SMF	40GBASE-LR4	100GBASE-LR4
40 km over SMF	40GBASE-ER4	100GBASE-ER4

Gray Text = IEEE Standard Red Text = In Standardization Green Text = In Study Group
Blue Text = Non-IEEE standard but complies to IEEE electrical interfaces

Outlook

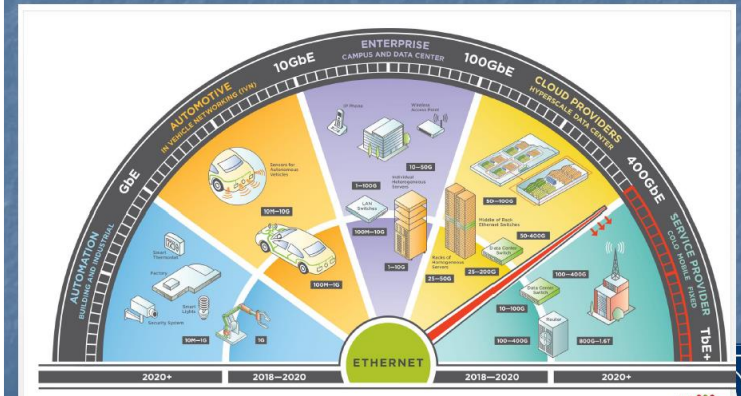

Ethernet Interfaces and Nomenclature

	Electrical Interface	Backplane	Twinax Cable	Twisted Pair (1 Pair)	Twisted Pair (4 Pair)	MMF	500m PSM4	2km SMF	10km SMF	40km SMF	80km SMF
10BASE-		T1S?		T1S/T1L							
100BASE-				T1							
1000BASE-				T1	T						
2.5GBASE-		KX		T1S?	T						
5GBASE-		KR		T1S?	T						
10GBASE-				T1S?	T						
25GBASE-	25GAUI	KR	CR/CR-S		T	SR			LR	ER	
40GBASE-	XLAUI	KR4	CR4		T	SR4/eSR4	PSM4	FR	LR4	ER4	
50GBASE-	LAUI-2/50GAUI-2 50GAUI-1	KR	CR			SR		FR	LR	ER	
100GBASE-	CAUI/10 CAUI-4/100GAUI-4 100GAUI-2 100GAUI-1	KR4 KR2 KR1	CR10 CR4 CR2 CR1			SR10 SR4 SR2	PSM4 DR	10X10 CWDM4 CLR4 100G-FR	LR4 4WDM-10 100G-LR	ER4 4WDM-40 ?	?
200GBASE-	200GAUI-4 200GAUI-2	KR4 KR2	CR4 CR2			SR4	DR4	FR4	LR4	?	?
400GBASE-	400GAUI-16 400GAUI-8 400GAUI-4	KR4	CR4			SR16	DR4	FR8 400G-FR4	LR8 ?	?	?

Physical layer	40 Gigabit Ethernet	100 Gigabit Ethernet
Backplane	n.a.	100GBASE-KP4
Improved Backplane	40GBASE-KR4	100GBASE-KR4
7 m over twinax copper cable	40GBASE-CR4	100GBASE-CR10 100GBASE-CR4
30 m over "Cat.8" twisted pair	40GBASE-T	
100 m over OM3 MMF 125 m over OM4 MMF[34]	40GBASE-SR4	100GBASE-SR10 100GBASE-SR4
2 km over SMF, serial	40GBASE-FR	100GBASE-CWDM4[37]
10 km over SMF	40GBASE-LR4	100GBASE-LR4
40 km over SMF	40GBASE-ER4	100GBASE-ER4

Gray Text = IEEE Standard Red Text = In Standardization Green Text = In Study Group

Blue Text = Non-IEEE standard but complies to IEEE electrical interfaces

ethernet alliance

12 Nov 2018

49

State of the Art

- 15000 optical fibers for CMS tracker, ≤ 5 Gb/s per fiber
- Only a fraction of data is read out
- Lasers located on detector

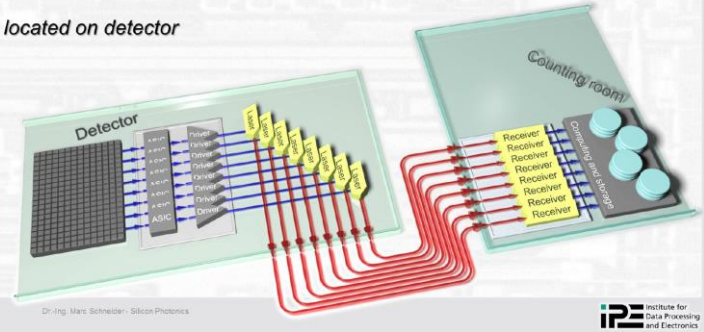

Read out ALL data?

- Efficient and radiation-hard modulators
- Lasers located off detector
- Wavelength division multiplexing (WDM) ⇒ fewer fibers

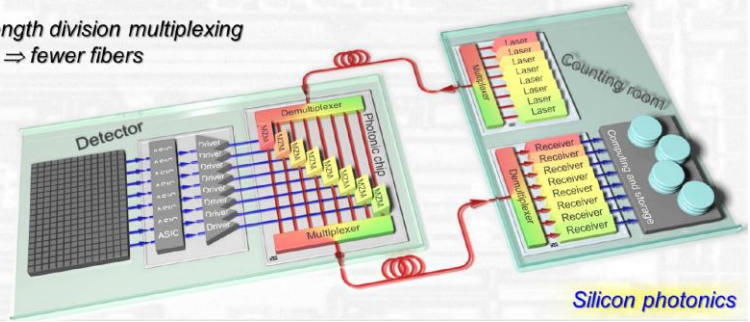

Silicon photonics

Silicon-on-Insulator: Waveguides

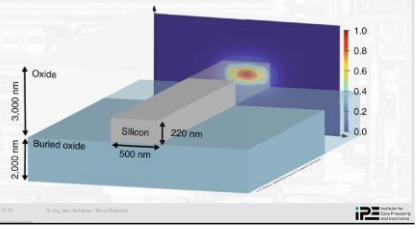

Integrated photodiodes

Example: Evanescent-coupled waveguide photodiode

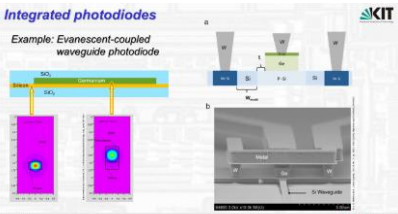

Silicon photonic modulators

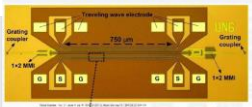

Xi Xiao, Hao Xu, Xianyao Li, Zhiyong Li, Tao Chu, Yude Yu, and Jinzhong Yu,

High-speed, low-loss silicon Mach–Zehnder modulators with doping optimization

Optics Express, Vol. 21, Issue 4, pp. 4116-4125 (2013), https://doi.org/10.1364/OE.21.004116

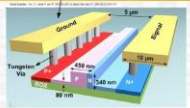

Silicon photonic modulators

Ring modulators

- Resonant and slightly absorbing ring, coupled to a bus waveguide
- At the resonance wavelength, light is coupled out of the bus waveguide and damped away

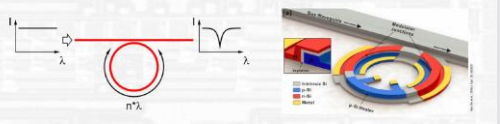

(De-)Multiplexers

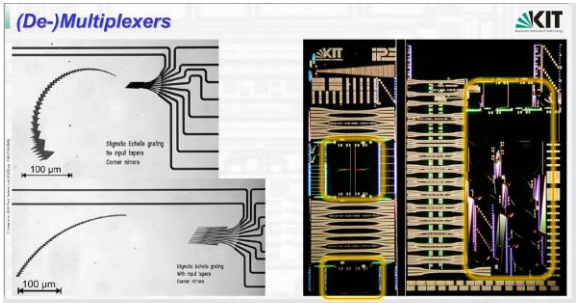

In-plane fiber-chip coupling

- Fiber axis aligned in parallel to chip surface
- SMF is polished at an angle θ smaller than 45 °
- Total internal reflection: Optical field couples radially from SMF to grating coupler

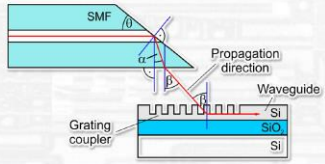

Optical probing of silicon photonic components

Mismatch between waveguide and fiber modes

- Mode size conversion required
- Edge-coupled inverted tapers
- Grating couplers for surface coupling

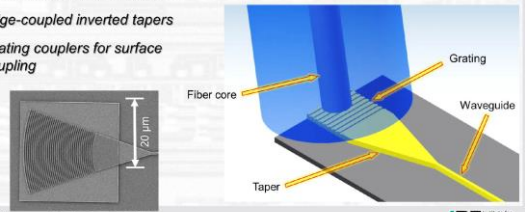

Synchronisation des canaux (transmission du temps)

Pictures of frequency standards

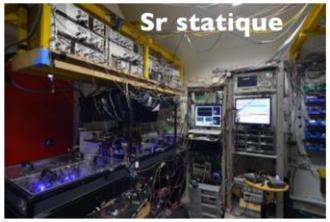

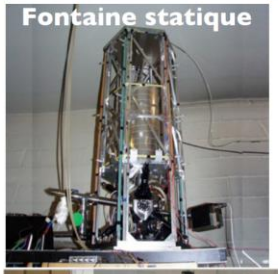

Time/Frequency dissemination nowadays

ntp-p1.obspm.fr

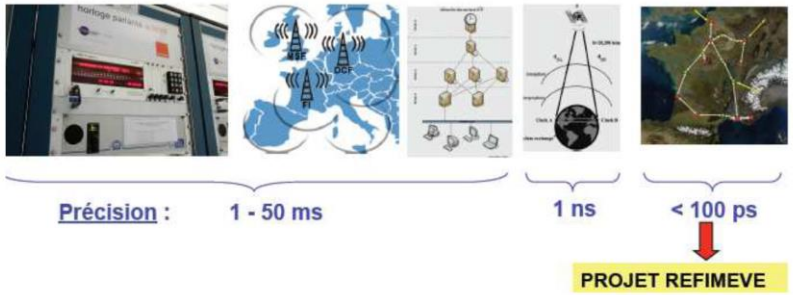

Courtesy M.-C. Angonin

Means to compare/disseminate clocks

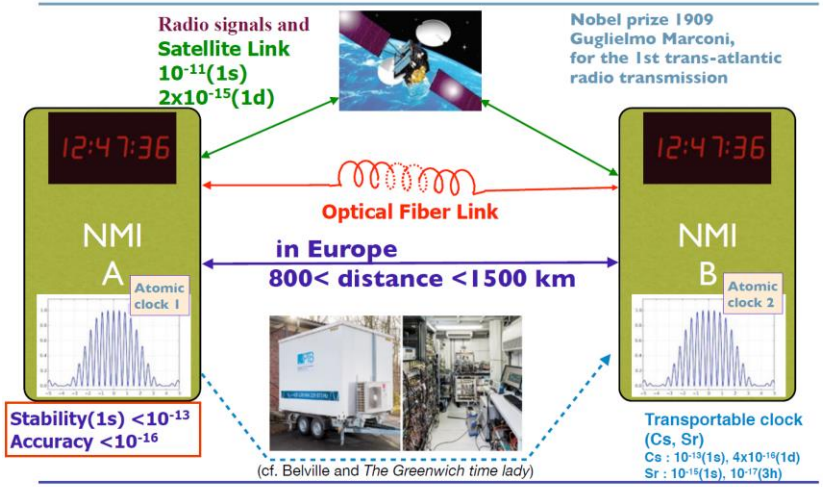

Synchronisation des canaux (transmission du temps)

PTP

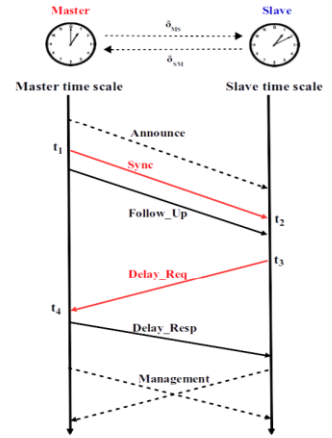

FIGURE 2.1: PTP two way messsage exchange mechanism [36].

PTP accounts for instrumental asymmetries.

Round trip time (RTT) = (t2−t1) + (t4−t3)

Clock offset = t2 − t1 + δMS

In case of asymmtery (δMS ≠ δSM):

error = (δMS −δSM) /2

PTP

Results from test set-up using M1000 units

Test set-up: A master M1000 unit with a 1PPS and NTP reference from UTC(NPL) transmitting PTP to a slave M1000 unit over two 50km fibre spools with long range SFPs.

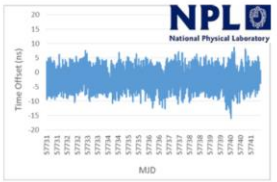

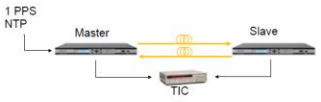

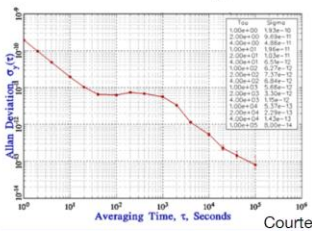

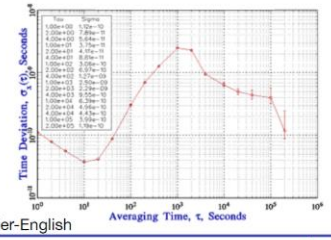

Courtesy E. Laier-English

WR-PTP

Synchronous Ethernet (SyncE)

Layer-1 syntonization
A common frequency reference for the entire network
All nodes of the network are locked to the frequency of the System timing master

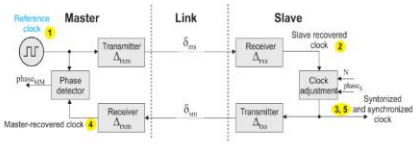

Digital Dual Mixer Time Difference (DDMTD)

Precise phase measurement
A phase compensated clock signal for the slave

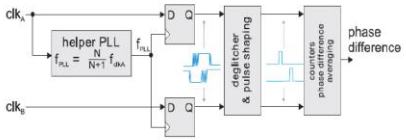

Asymmetry compensation

Sources of propagation asymmetry in a White Rabbit link:
- Chromatic dispersion
- Unequal fiber lengths

'Static' correction of propagation asymmetry possible with WR.

G. Daniluk,(CERN). Nuclear Instruments and Methods in Physics Research **725**, 187–190 (2013).

Synchronisation des canaux (transmission du temps)

Optical time transfer

Amplitude or Phase modulation of the optical carrier

O.Lopez, et al. Applied Physics B 110, 3–6 (2013).

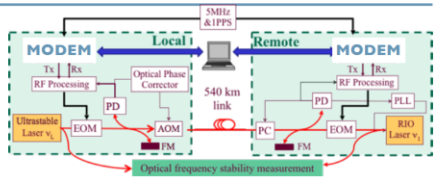

Optical demodulation

F. Frank, et al. IEEE T-UFFC 1–1 (2018)

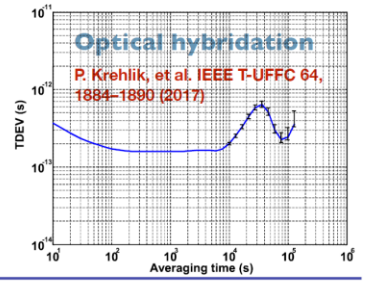

Techniques not yet as mature as frequency transfer

Cascaded links

- Multi-segments approach :
  - Link is divided into a few segments, depending on noise and losses
  - → shorter delay
  - → larger bandwidth and better noise rejection
- Repeater stations are needed :
  - Repeater station N : send back signal to station N-1, amplify and filter, correct the noise of next link N

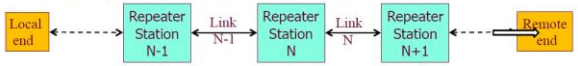

Repeater laser station commercially available
F.Guillou-Camargo et al, Appl. Opt., AO 57, 7203–7210 (2018).

25 years of range improvement

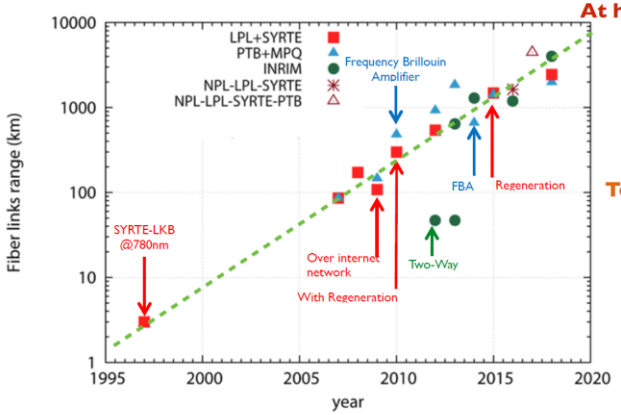

At horizon 2020 : 8000 km

GÉANT
connect • communicate • collaborate

Towards a large research infrastructure ?

RENATER, CESNET, PSNC, GARR JISC/JANET, DFN, SURFNET, NORDUNET…

Future prospects

Towards Research Infrastructure

- Work with Network for Education and Research Industry to make the technology available
- Ways to access the network
- Compatibility with TelCo

Project CLONETS involved 16 partners from 7 European countries. Partners represent 4 main areas:
- National Measurement Institutes: OBS PARIS (FR), NPL (UK), PTB (DE), INRIM (IT)
- National Research and Education Network: RENATER (FR), CESNET (CZ), PSNC (PL), GARR* (IT),
- Academic Laboratories: AGH (PL), UP13 (FR), UCL (UK), ISI (CZ), CNRS* (FR)
- Industrial: MUQUANS (FR), MENLO (DE), PIKTIME (PL), SEVEN SOL (SP), OPTOKON (CZ), TOP-IX* (IT)

* Third-party member

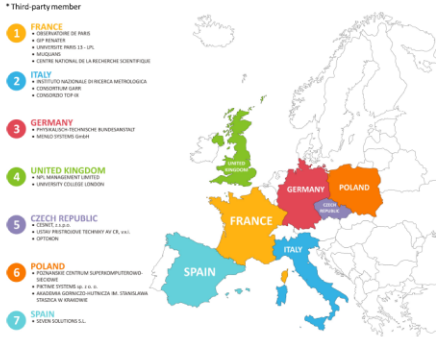

Les Technologies émergentes pour les DAQ

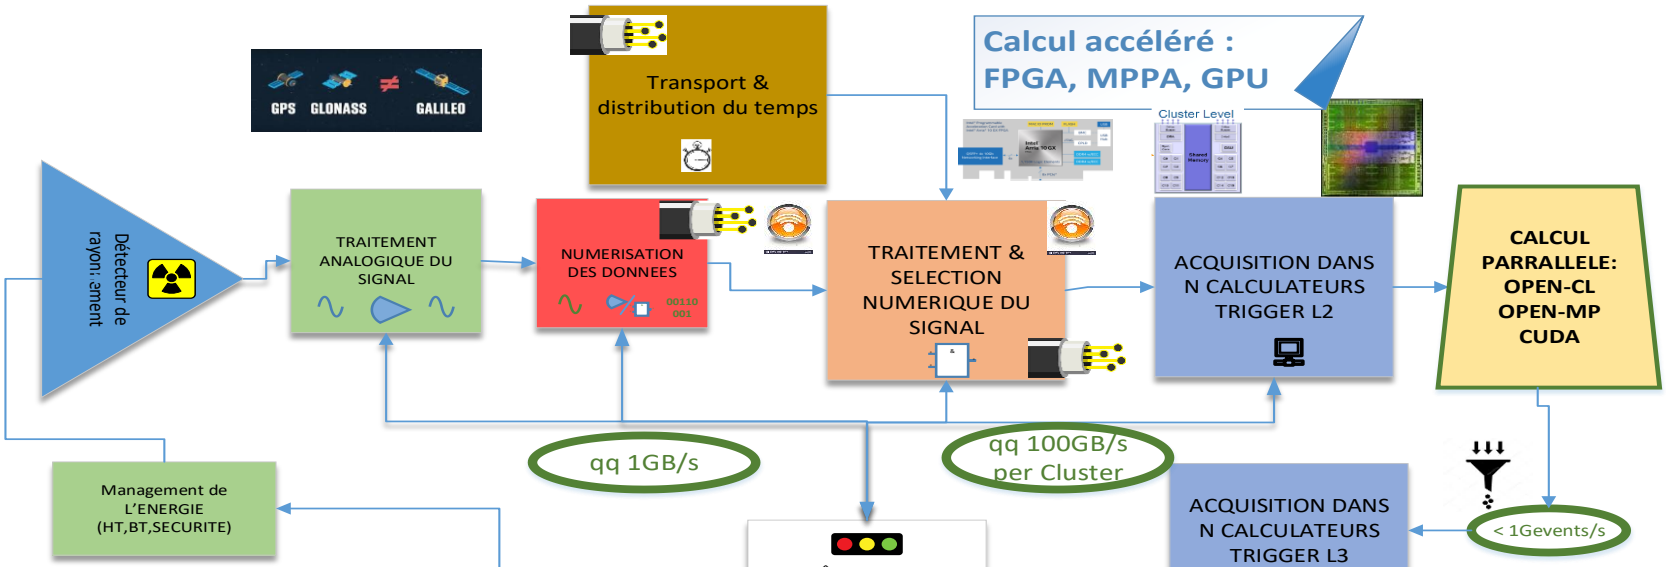

Les points de rupture :

- Transmission Optique et RF des FE
- Distribution d'horloges précises (qq ps)
- Calcul accéléré et massivement parallèle
- Algorithmes IA (*embarqués*)

Importance de la hiérarchie mémoire

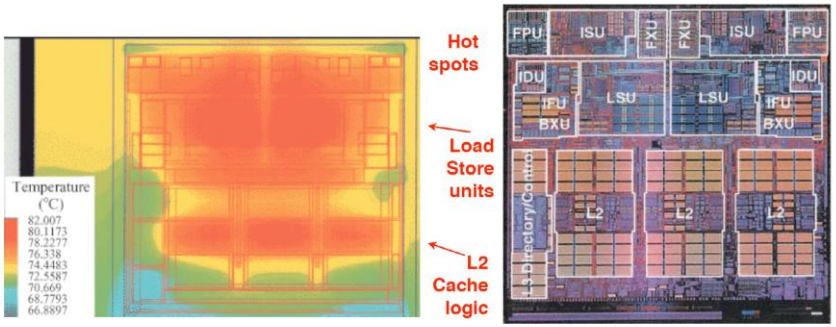

Memory Wall

Memory access latency is a "wall"

– Fundamental barrier to further improvement in singlethread computation performance

- Attempts to mitigate latency

– Caching (memory hierarchy), out-of-order execution, prefetching

– … but can only go so far

- Migration towards multi-processing

– Provide other threads of execution while waiting for memory

– Increase memory bandwidth to compensate for long latency

DDR memory : latency/bandwidth

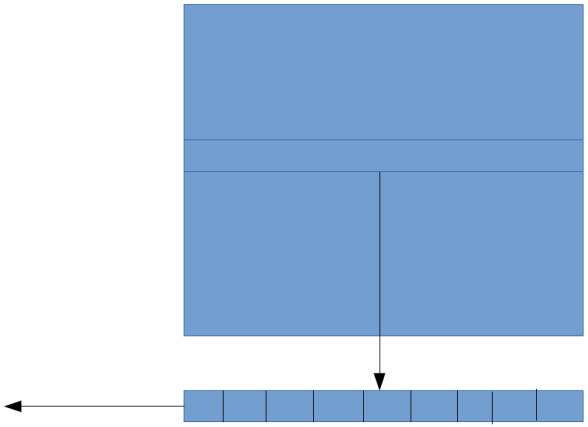

Classic Multicore Memory Hierarchy

Challenge: managing interference between cores

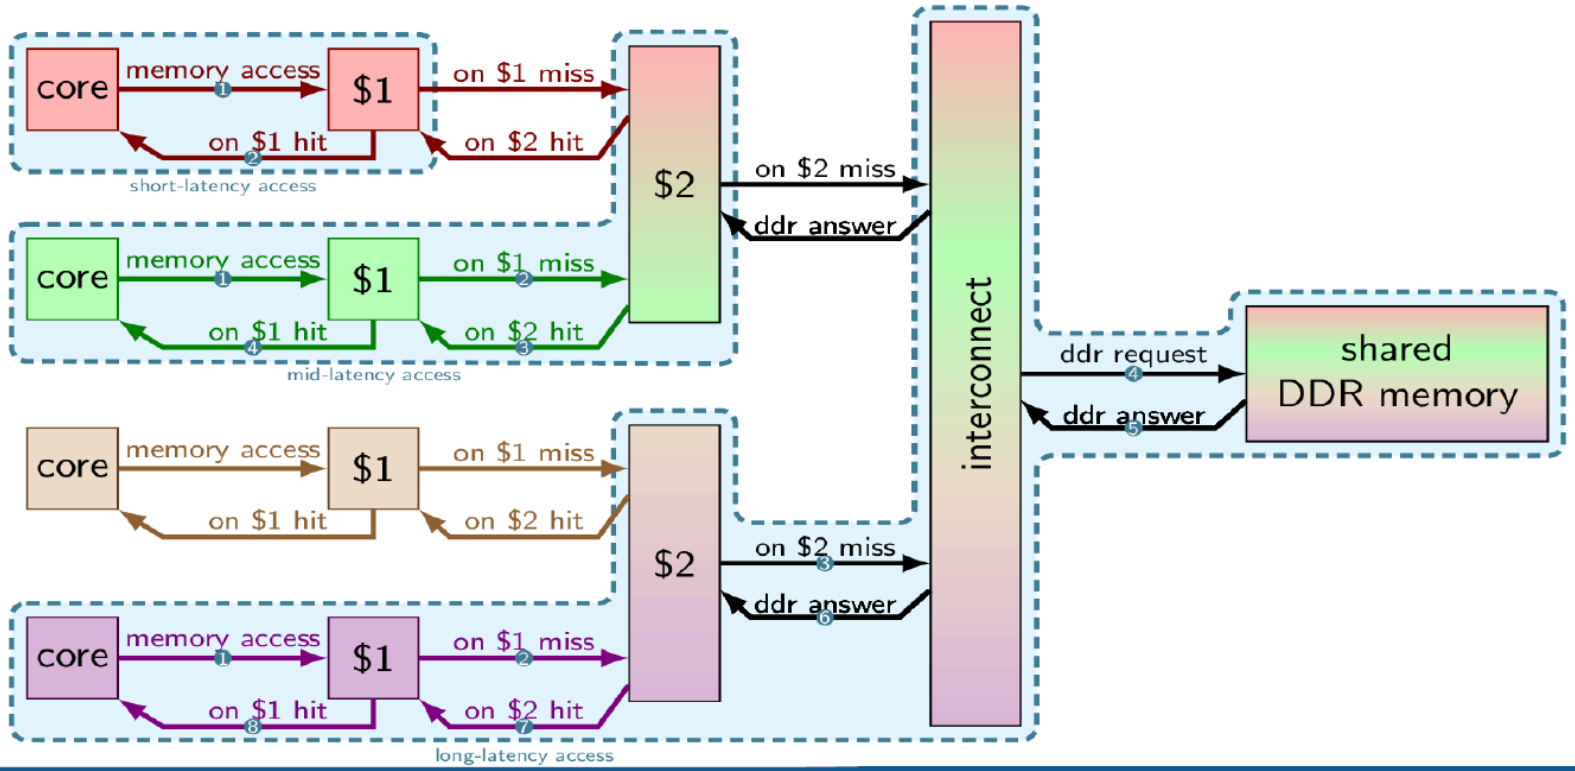

Embedded Multicore Memory Hierarchy

Challenge: programmability of DMA and private memories

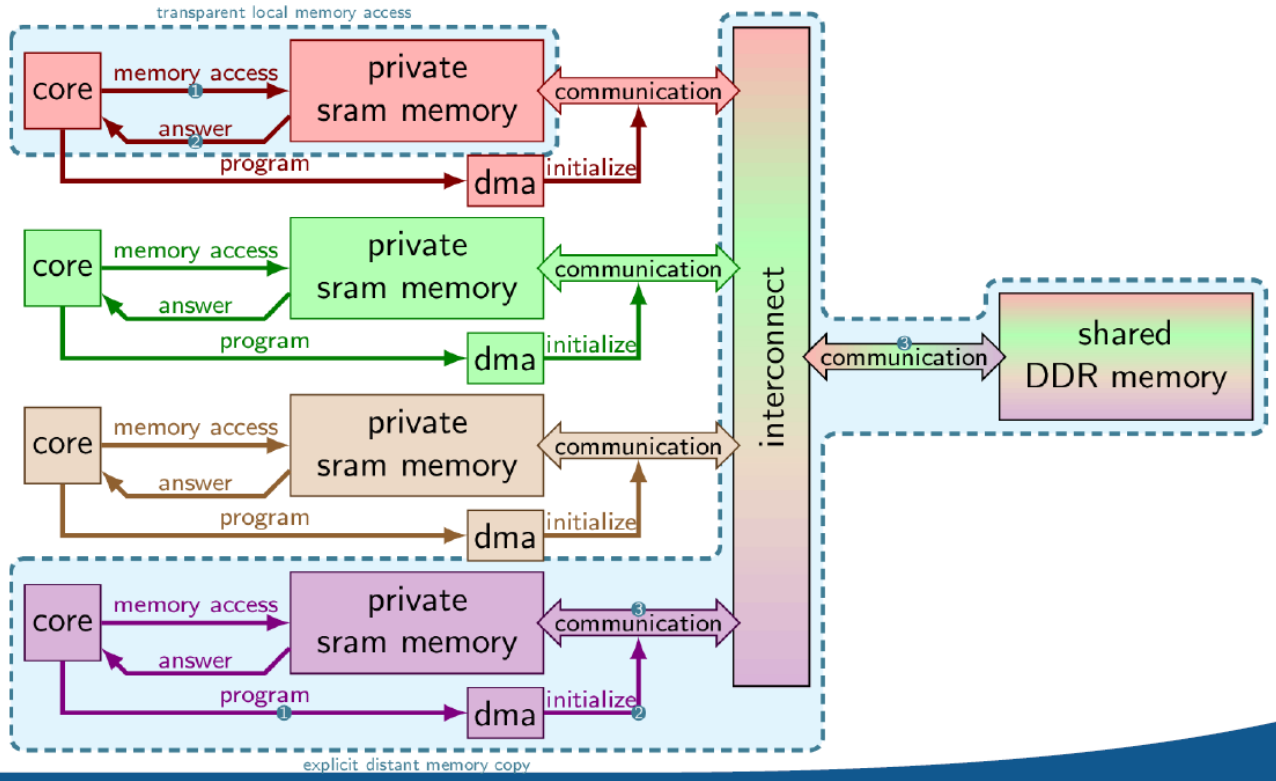

Calcul accéléré : aspect Matériel

Graphical Processing Unit

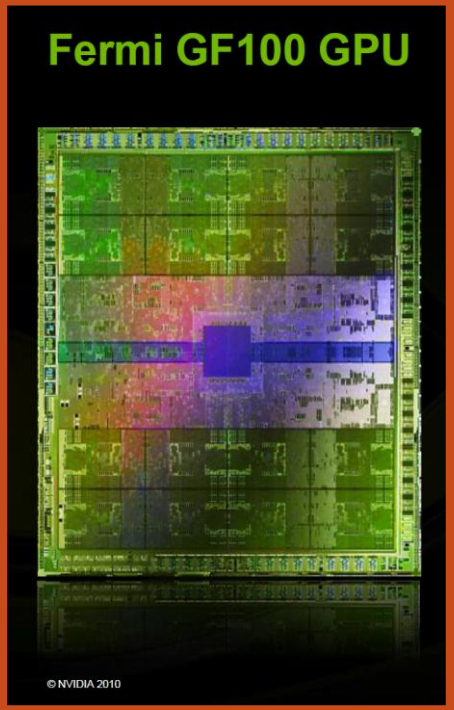

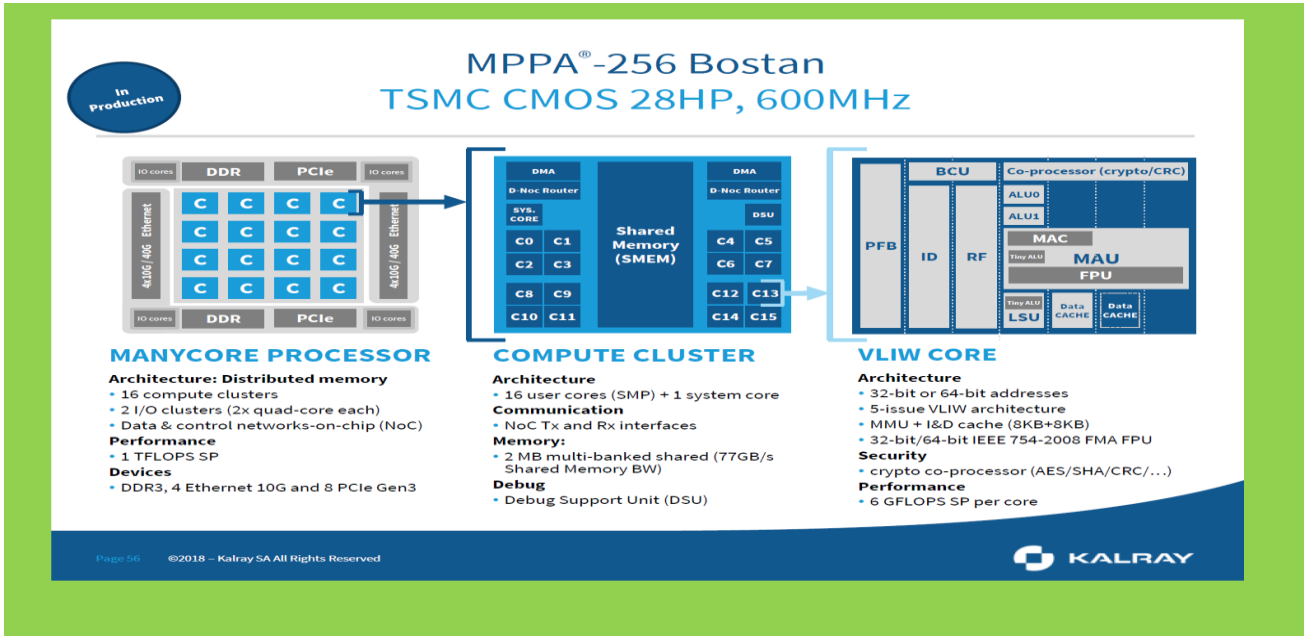

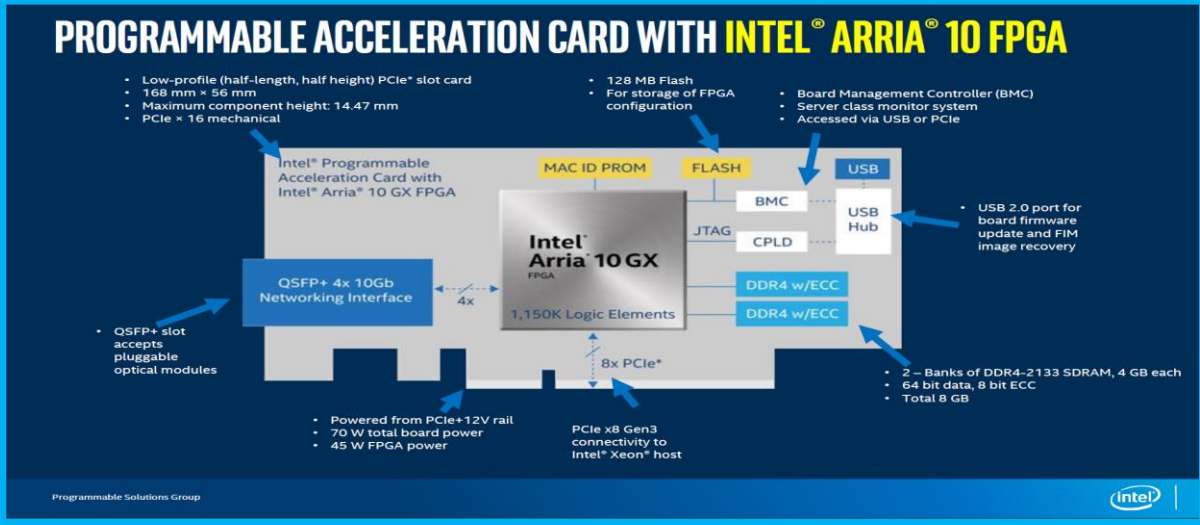

GPUs: Throughput Oriented Design

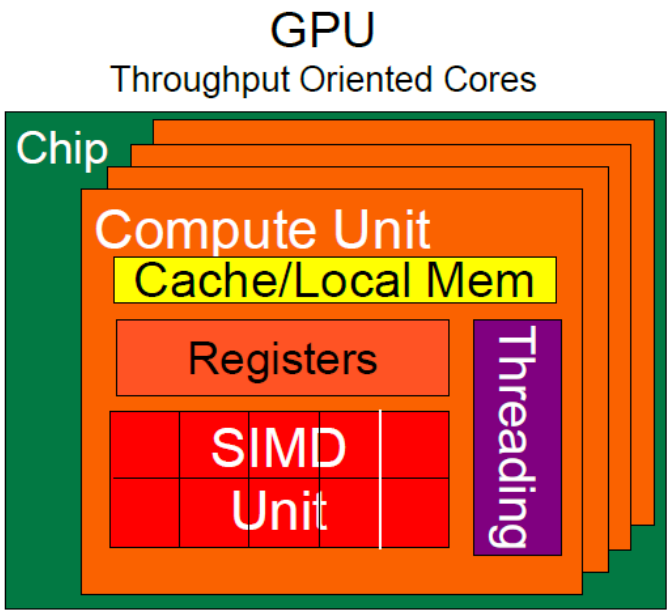

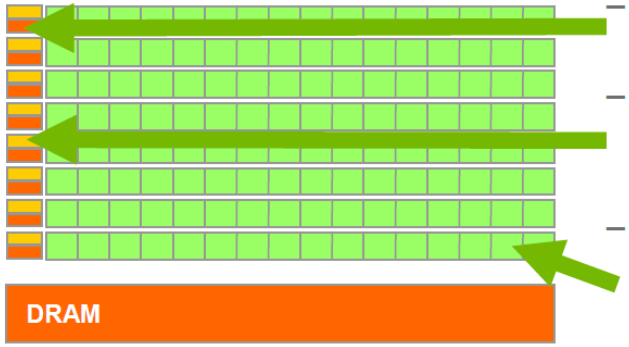

- Small caches
  - To boost memory throughput
- Simple control
  - No branch prediction
  - No data forwarding
- Energy efficient ALUs
  - Many, long latency but heavily pipelined for high throughput
- Require massive number of threads to tolerate latencies
  - Threading logic
  - Thread state

GPU permet le parallélisme et donc booste les performances des calculs matriciels (Image, FFT, Reseau de Neurones)

Calcul accéléré : Architecture GPU

Fermi SM

- **Objective – DX11 support**
  - Polymorph engine
- **Objective – Optimize for GPU computing**
  - New ISA
  - Revamp issue / control flow
  - New CUDA core architecture
- **32 cores per SM (512 cores total)**
- **64KB configurable L1$ / shared memory**

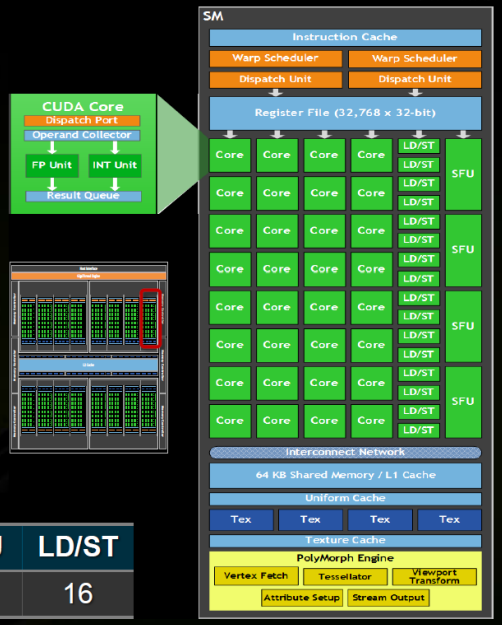

	FP32	FP64	INT	SFU	LD/ST
Ops / clk	32	16	32	4	16

logiBRICKS™ Designed by XYLON

logiREF-ZGPU-ZED GPU Reference Design User's Manual

Xylon

April 8th, 2015 | Version: v2.02.a

Figure 1: The ZedBoard Development Kit Running Xylon's 3D Graphics Demo

(Video clip: http://www.logicbricks.com/logicBRICKS-IP-Library/Video-Galleries/logicBRICKS-Demos-ZedBoard-Clip.aspx)

Calcul accéléré : Architecture Multi-Purpose Processor Array

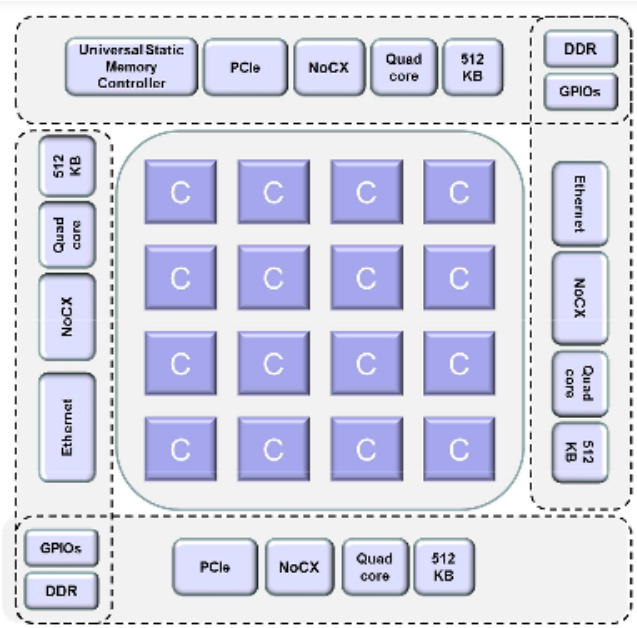

Figure 1 – MPPA®-256 block diagram

- **Image and Audio processing:** HD encoding, broadcasting, video surveillance, augmented reality
- **Signal processing:** Radar, telecom, medical
- **Intensive Computing:** Oil & Gas, finance, video live streaming, numerical simulation, Bio Sciences
- **Control Command:** Aeronautics, industrial automation
- **Telecom:** Routers, cryptography, software defined radio, base station

Core architecture

The MPPA® core is a 32-bit Very Long Instruction Word (VLIW) processor made of:

- One Branch/Control Unit
- Two Arithmetic Logic Units
- One Load/Store Unit including simplified ALU
- One Multiply-Accumulate (MAC) / FPU including a simplified ALU
- Standard IEEE 754-2008 FPU with advanced Fused Multiply-Add (FMA) and dot product operators
- One Memory Management Unit (MMU)

Compute Cluster

Each compute cluster is composed of:

- 16 identical cores with private FPU and MMU
- Dynamic Voltage and Frequency Scaling (DVFS) and Dynamic Power Switch off (DPS) support
- 1 system core with private FPU and MMU
- An instruction and data L1-cache per core
- 1 smart Direct Memory Access (DMA)
- A shared memory
- 1 Debug Support Unit

The cores are connected to a multibank memory enabling low latency access or bank private access depending on the configuration.

Calcul accéléré : Architecture Multi-Purpose Processor Array

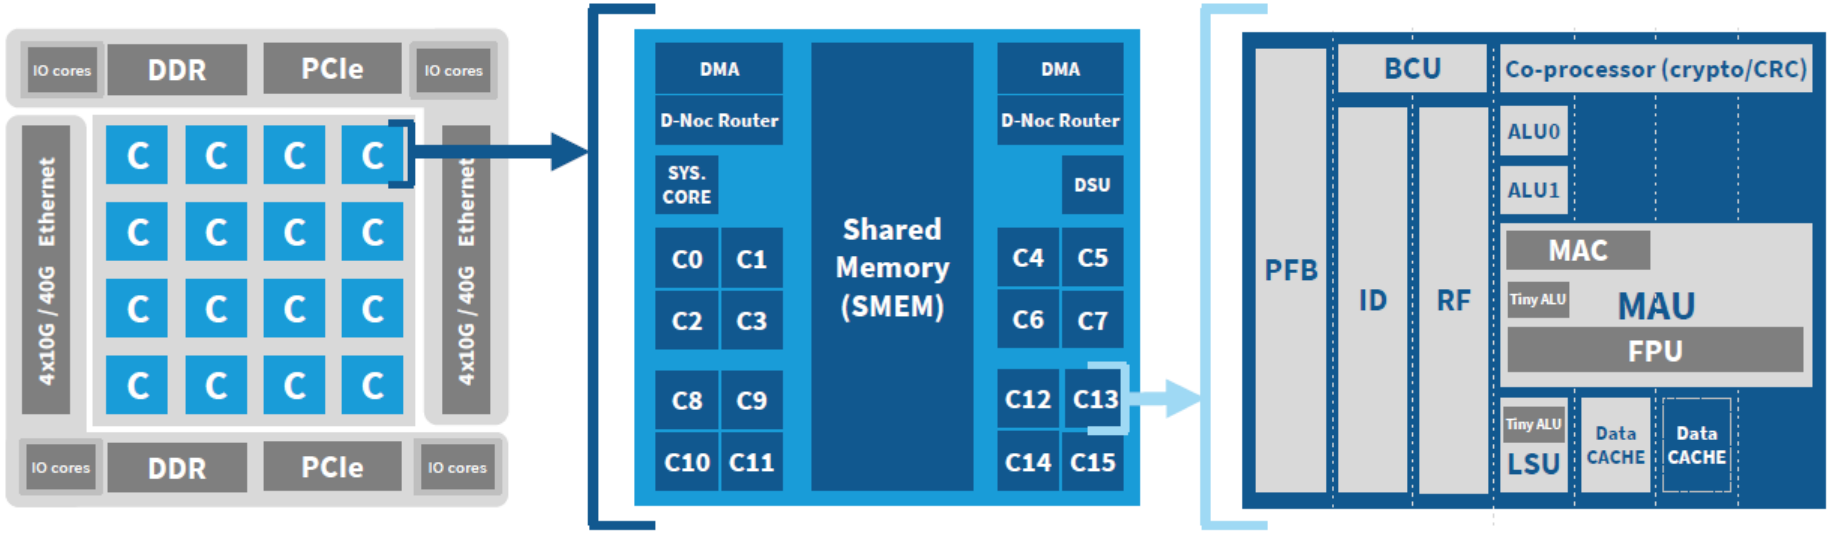

MANYCORE PROCESSOR

Architecture: Distributed memory

- 16 compute clusters
- 2 I/O clusters (2x quad-core each)
- Data & control networks-on-chip (NoC)

Performance

- 1 TFLOPS SP

Devices

- DDR3, 4 Ethernet 10G and 8 PCIe Gen3

COMPUTE CLUSTER

Architecture

- 16 user cores (SMP) + 1 system core

Communication

- NoC Tx and Rx interfaces

Memory:

- 2 MB multi-banked shared (77GB/s Shared Memory BW)

Debug

- Debug Support Unit (DSU)

VLIW CORE

Architecture

- 32-bit or 64-bit addresses
- 5-issue VLIW architecture
- MMU + I&D cache (8KB+8KB)
- 32-bit/64-bit IEEE 754-2008 FMA FPU

Security

- crypto co-processor (AES/SHA/CRC/…)

Performance

- 6 GFLOPS SP per core

KALRAY

Calcul accéléré : Architecture CPU + FPGA

https://www.microsoft.com/en-us/research/wp-content/uploads/2016/10/Cloud-Scale-Acceleration-Architecture.pdf

Idée : ***Accélérer les transmissions de données et les calculs sur les nœuds des data center de Microsoft***

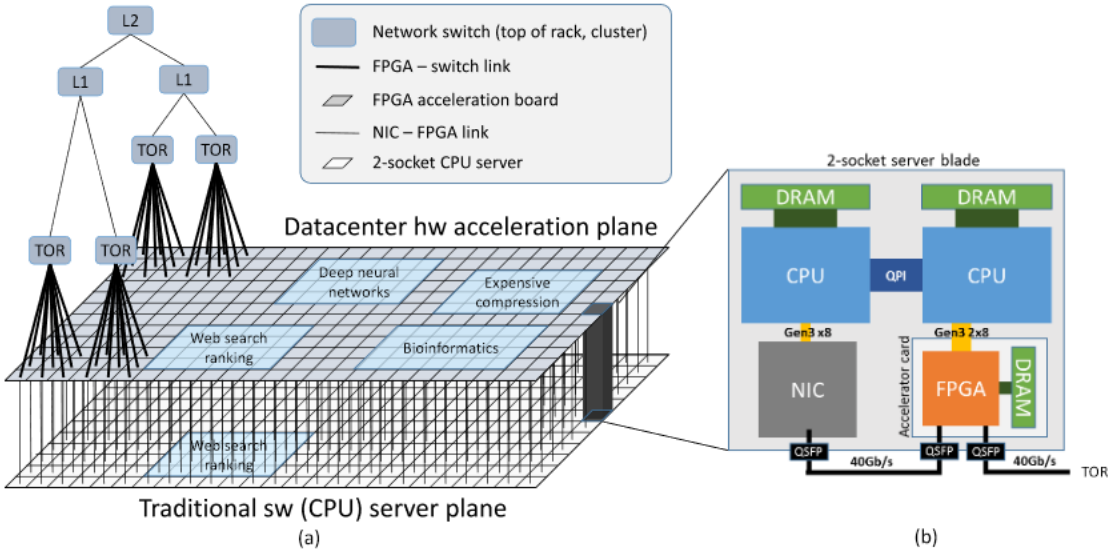

Fig. 1. (a) Decoupled Programmable Hardware Plane, (b) Server + FPGA schematic.

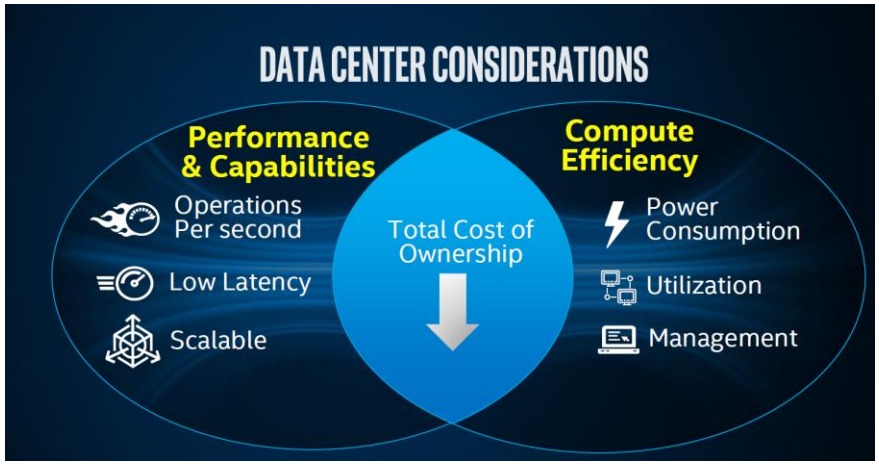

Calcul accéléré : Architecture CPU + FPGA

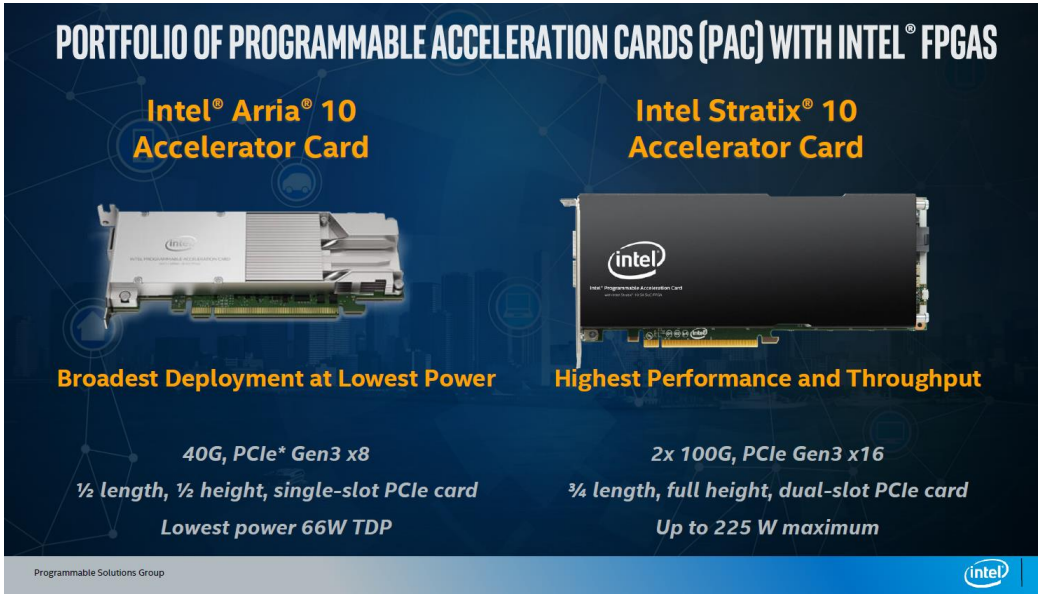

Calcul accéléré : Architecture CPU + FPGA

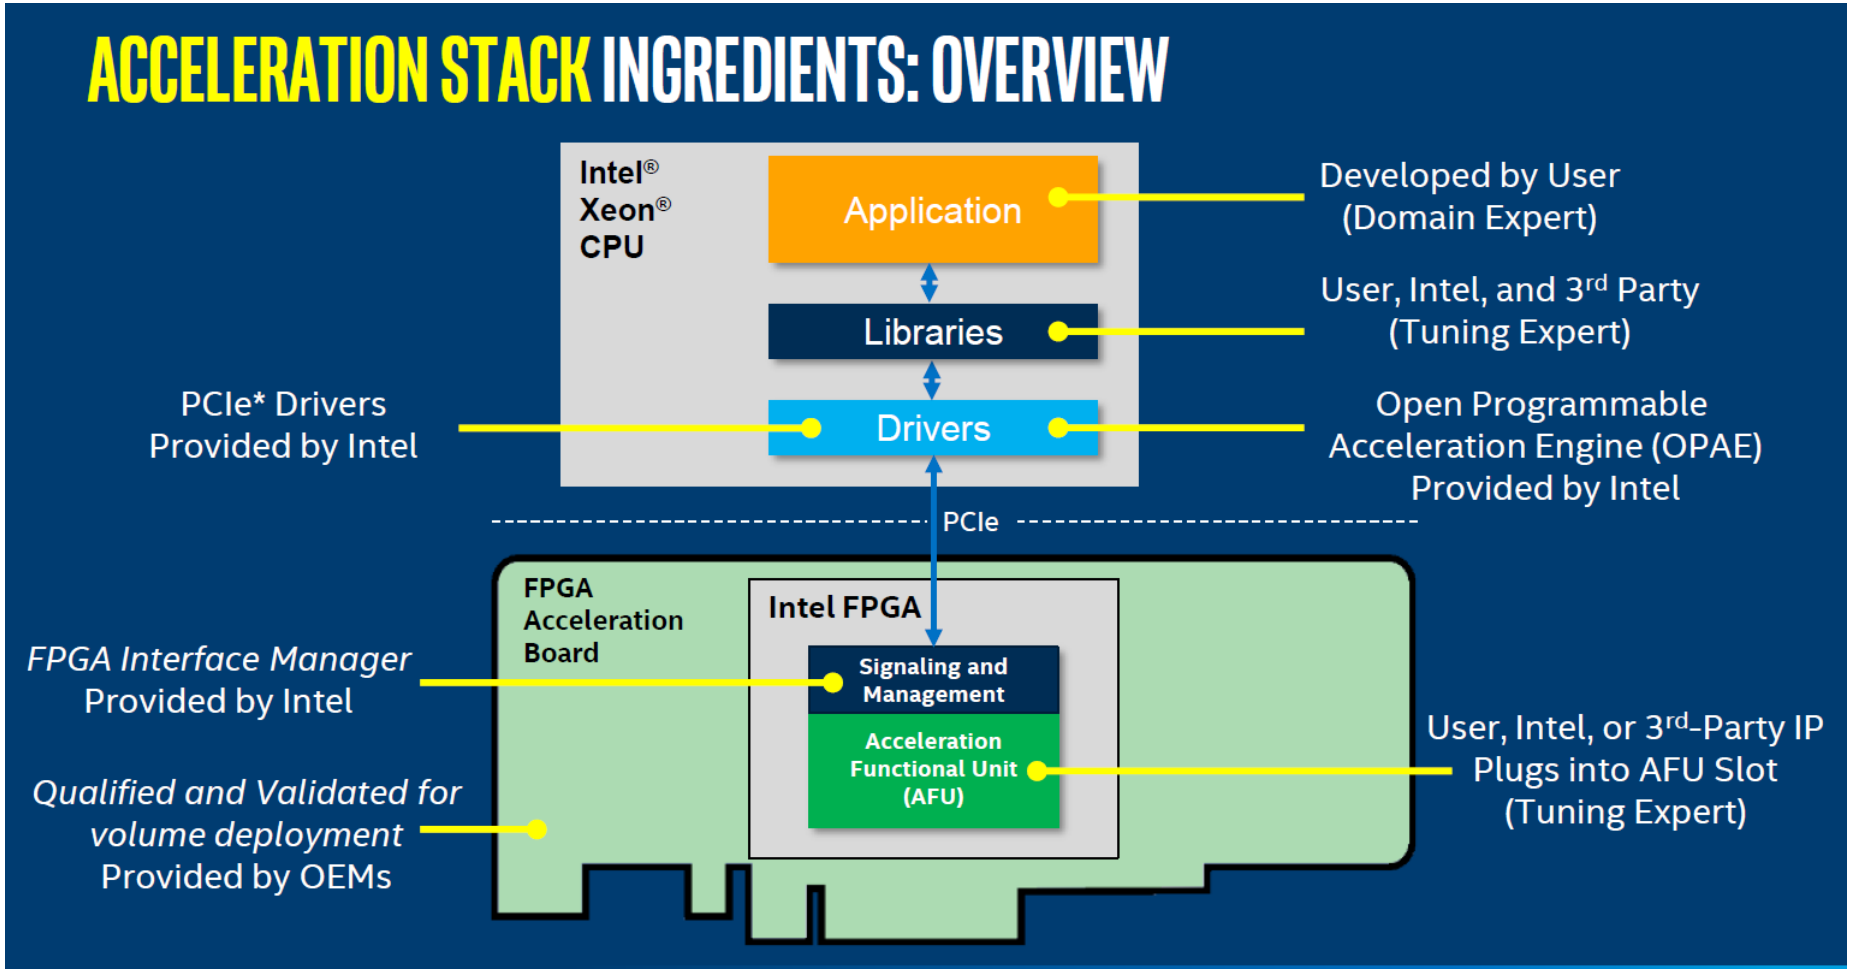

Calcul accéléré : aspect logiciel

Langage avec #PRAGMA et bibliotheques dédiées

CUDA
OPENMP
OPENCL

Couches C, C++

Calcul GPU :
Langage CUDA
Langage OpenCL

Calcul MPPA :
Langage OpenMP 3
Langage OpenCL 1.2
Dedicated Libraries (BLAS, DL, FFT,CV

Calcul CPU + FPGA :
Langage HLS
Langage OpenCL

Implementation :
Mathlab, Mathematica, Labview, CUDA Fortran, CUDA C & C++, PyCUDA…

Tiled Matrix Multiplication Kernel

```
__global__ void MatrixMulKernel(float* M, float* N, float* P, Int Width)
{
  __shared__ float ds_M[TILE_WIDTH][TILE_WIDTH];
  __shared__ float ds_N[TILE_WIDTH][TILE_WIDTH];

  int bx = blockIdx.x;  int by = blockIdx.y;
  int tx = threadIdx.x; int ty = threadIdx.y;

  int Row = by * blockDim.y + ty;
  int Col = bx * blockDim.x + tx;
  float Pvalue = 0;

// Loop over the M and N tiles required to compute the P element
for (int p = 0; p < n/TILE_WIDTH; ++p) {
   // Collaborative loading of M and N tiles into shared memory
   ds_M[ty][tx] = M[Row*Width + p*TILE_WIDTH+tx];
   ds_N[ty][tx] = N[(t*TILE_WIDTH+ty)*Width + Col];
   __syncthreads();

   for (int i = 0; i < TILE_WIDTH; ++i)Pvalue += ds_M[ty][i] * ds_N[i][tx];
   __synchthreads();
}
P[Row*Width+Col] = Pvalue;
}
```

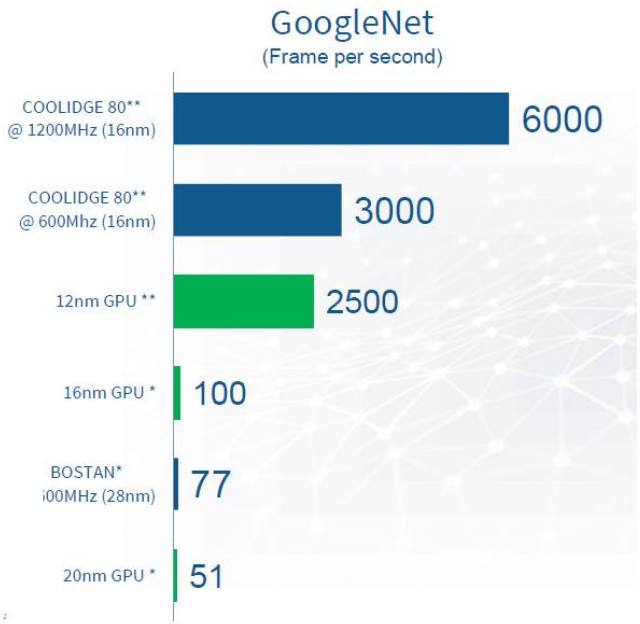

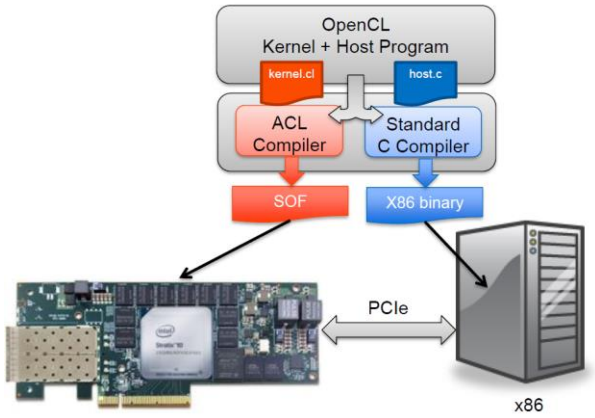

Traditional FPGA Design Flow

Design specification

```
module dut(…);
  always @(…)
  …
endmodule;
```

Design entry/RTL coding
- Behavioral or structural description of design
- Possibly with the help of high-level tools

RTL functional simulation
- Use 3rd party simulators
- Verify logic model & data flow

Synthesis (Mapping)
- Translate design into device-specific primitives
- Design optimized
- Quartus® synthesis or other 3rd party tools

Place & route (Fitting)
- Assign primitives to locations
- Route the resources

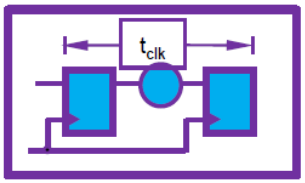

Static timing analysis
- Verify performance specs can be met

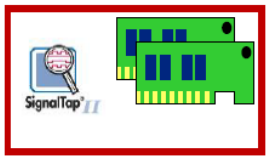

Board simulation & test
- Simulate board design
- Program & test device on board

Calcul accéléré : aspect logiciel pour les FPGA

FPGA High Level Design with OpenCL™

Goal: Design FPGA custom hardware with C-based software language

```
__kernel void _foo (__global float *x) {
  int i …
}
```

Intel® FPGA SDK for OpenCL™

- Benefits
  - Makes FPGA acceleration available to software engineers
  - Debug and optimize in a software-like environment
  - Significant productivity gains compared to hardware-centric flow
  - Easier to perform design exploration
  - Abstracts away FPGA design flow and FPGA hardware

OpenCL™ Platform Model

- One Host with one or more OpenCL™ Devices
  - Each Device is composed of one or more compute units
- Memory divided into Host Memory and various types of Device Memory

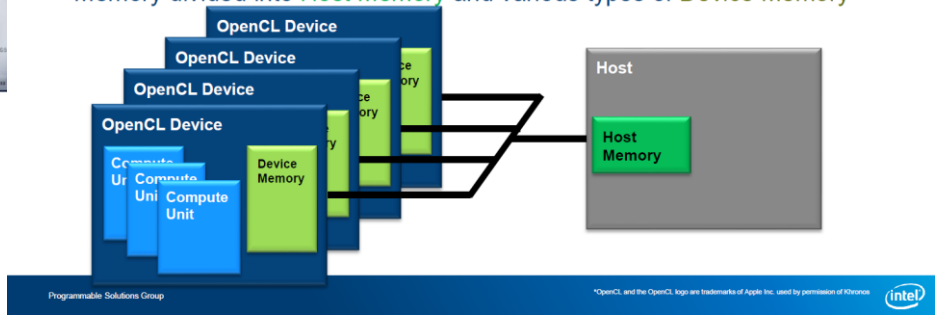

Altera's OpenCL Flow

Intel's OpenCL SDK for FPGA takes a system level view

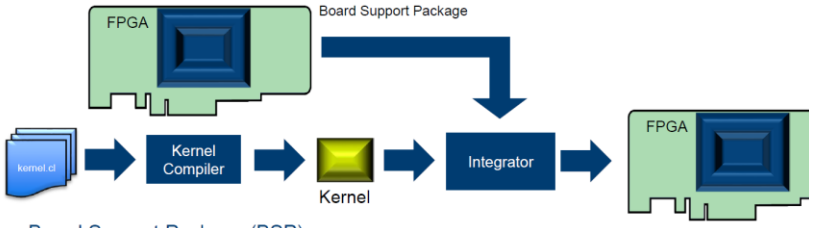

- Board Support Package (BSP)
  - "Chassis" to hold the newly created kernel
- Kernel Compiler
  - Optimized pipelines from C
- System Integrator
  - Merge all together and generate partial reconfiguration files for FPGA

Intel® FPGA OpenCL™ Device

Each device is made of many independent compute units

- Each compute unit is custom built from kernel code

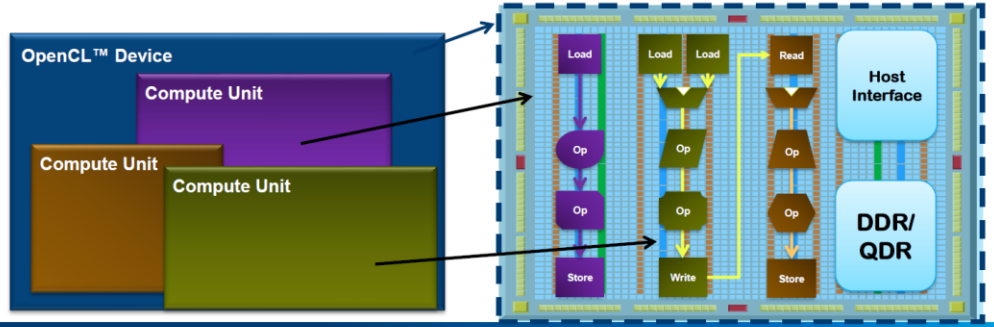

Les Technologies émergentes pour les DAQ

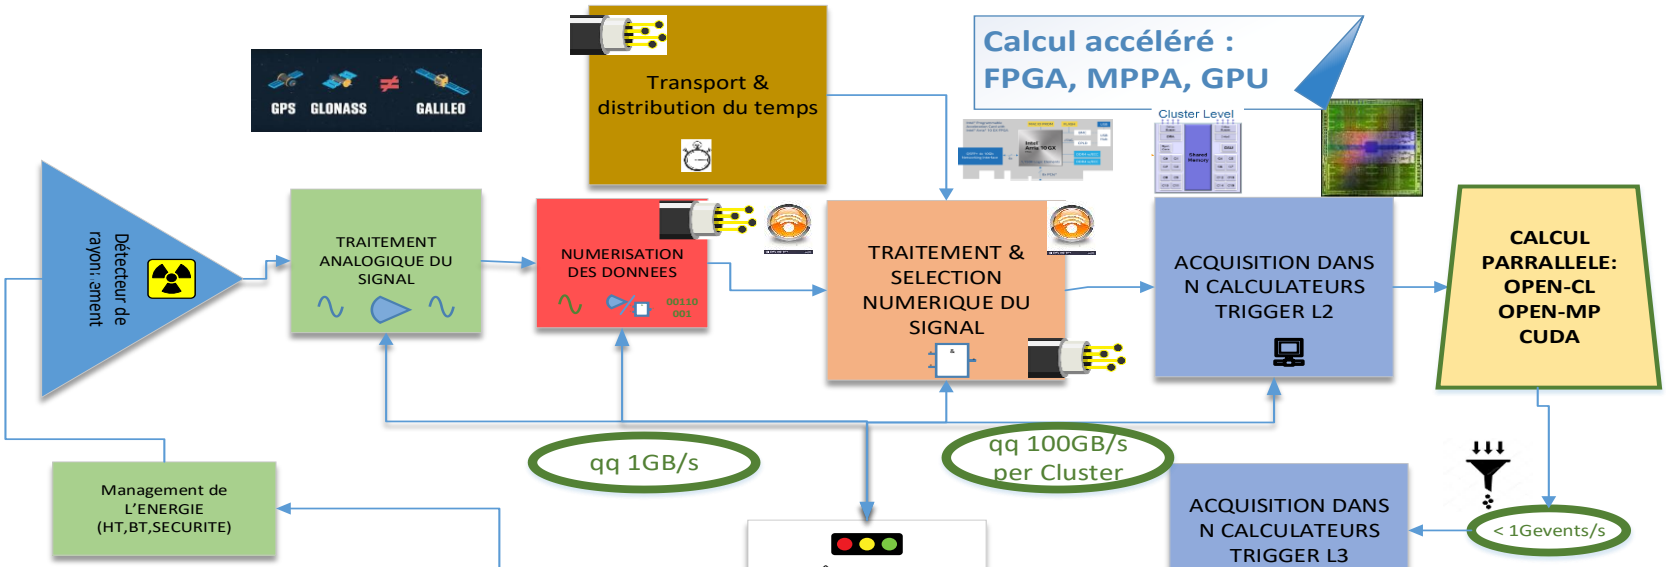

Les points de rupture :

- Transmission Optique et RF des FE
- Distribution d'horloges précises (qq ps)
- Calcul accéléré et massivement parallèle
- Algorithmes IA (*embarqués*)

Algorithmes IA: Deep Learning pour la physique

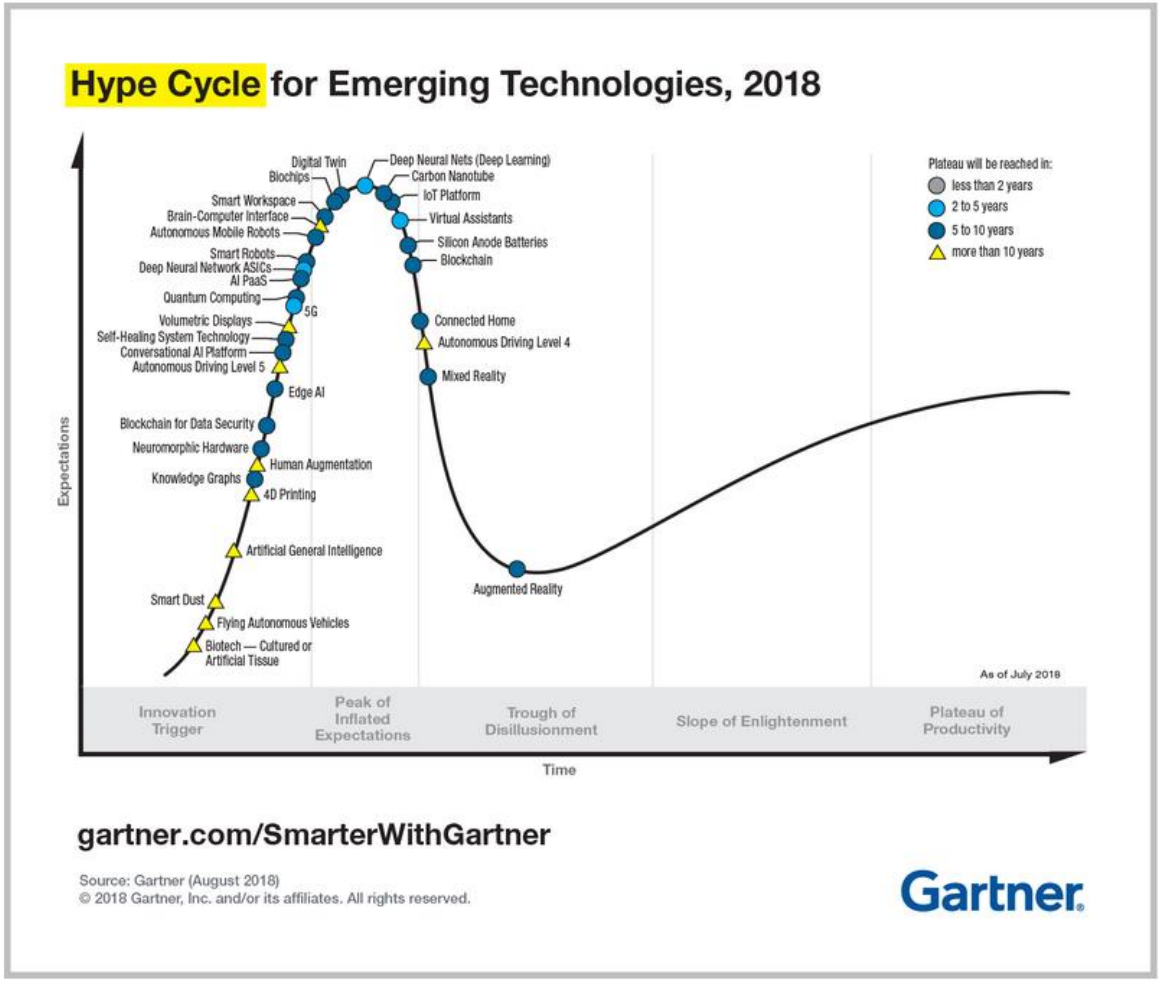

Tendances :

1. Democratisation de l'AI
2. Ecosystemes numeriques
3. Bio-evolution (biochips) humain augmenté
4. Experiences immersives (capteurs, Plateformes IT
5. Nouvelles Infrastructures (Ordinateur quantique…)

Machine Learning : Ensemble d'algorithmes pour résoudre des problèmes:

- Identification d'objet
- Classification
- Detection…

Deep Learning (IA) : Technologies liées aux réseaux de Neurones.
Phase d'apprentissage
Fonctionne bien dans des cas particuliers

Algorithmes IA: Deep Learning pour la physique

One Neuron

- Elementary computation

$$\text{activation} = w^T.x = \sum_j w_j x_j + w_0$$

$$\text{output} = ReLU(a(x))$$

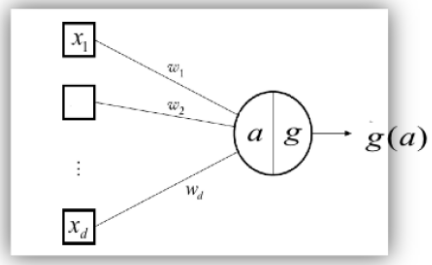

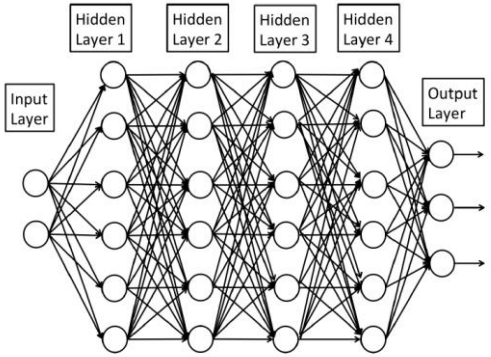

Autoencoders

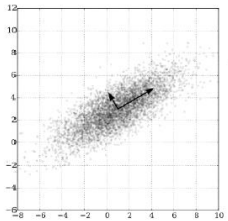

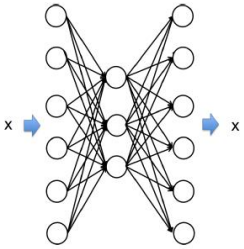

LeNet [LeCun and al., 1997]

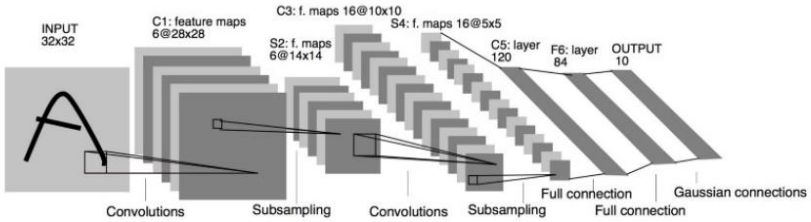

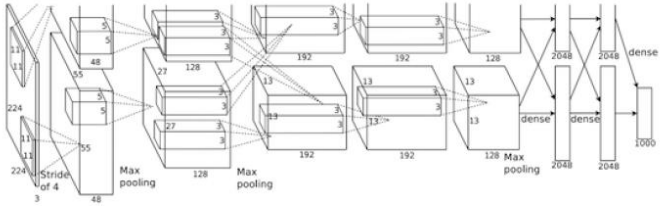

AlexNet [Krizhevsky and al., 2012]

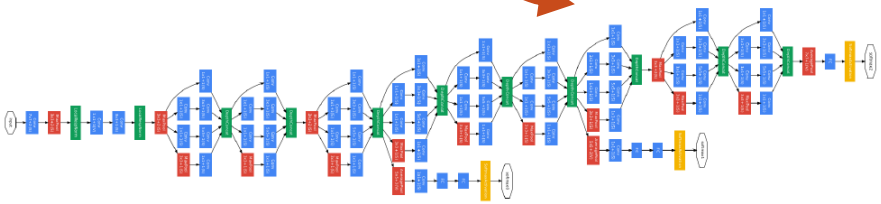

Algorithmes IA: Deep Learning pour la physique

Why now ?

Huge training resources for huge models

- Huge volumes of training data
- Huge computational ressources (clusters of GPUs)

Advances in understanding optimizing NNs

- Regularization (Dropout...)
- Making gradient flow (ResNets, LSTM, ...)

Faster diffusion than ever

- Softwares
- Results
  - Publications (arxiv publication model) + codes
  - Architectures, weights (3 python lines for loading a state of the art computer vision model!)

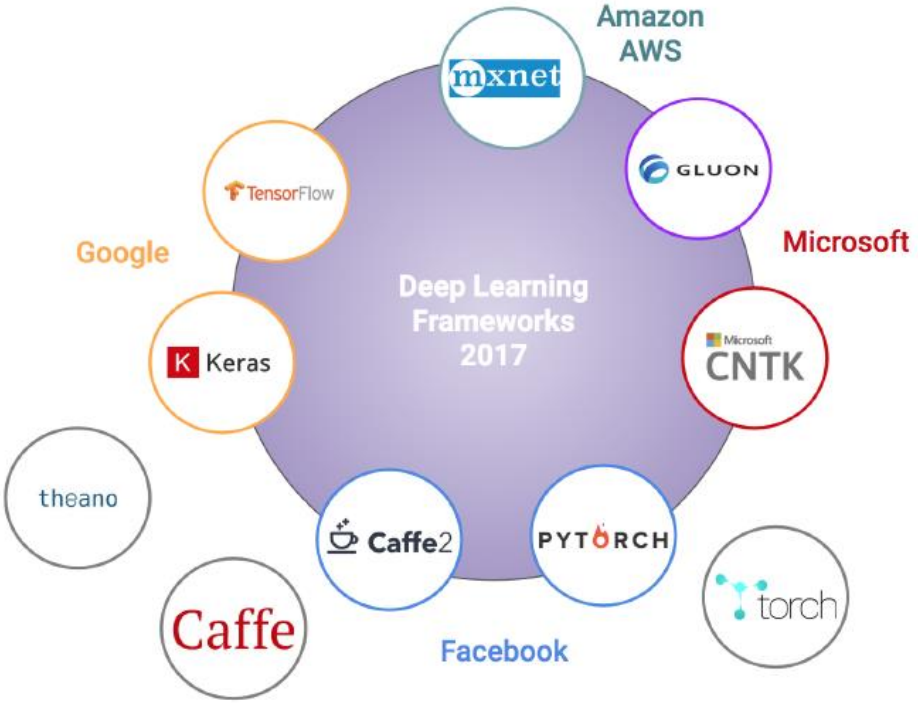

Algorithmes IA: Deep Learning pour la physique

Pivot adversarial learning [Louppe et al., 2017]

- Use adversarial learning to align simulated and systematic data

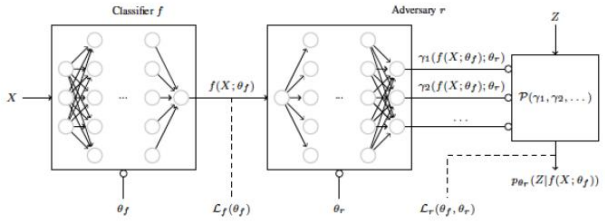

Figure 1: Architecture for the adversarial training of a binary classifier f against a nuisance parameters Z. The adversary r models the distribution $p(z|f(X;\theta_f)=s)$ of the nuisance parameters as observed only through the output $f(X;\theta_f)$ of the classifier. By maximizing the antagonistic objective $\mathcal{L}_r(\theta_f,\theta_r)$, the classifier f forces $p(z|f(X;\theta_f)=s)$ towards the prior $p(z)$, which happens when $f(X;\theta_f)$ is independent of the nuisance parameter Z and therefore pivotal.

Deep and HEP

Ziyu Guo's thesis (with Y. Coadou)

- Deep learning in the search for ttH with the ATLAS experiment at the Large Hadron Collider
- 1. Replace a reconstruction BDT + classification BDT with a end to end learned joint model

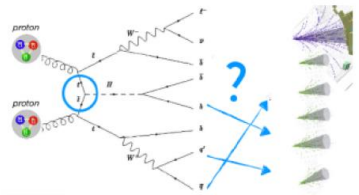

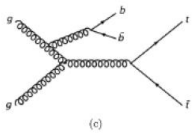

Deep and HEP

QCD-aware Recursive Neural Networks for Jet Physic [Louppe et al., 2018]

- Learn to aggregate features for jets using a tree structure inspired from data knowledge

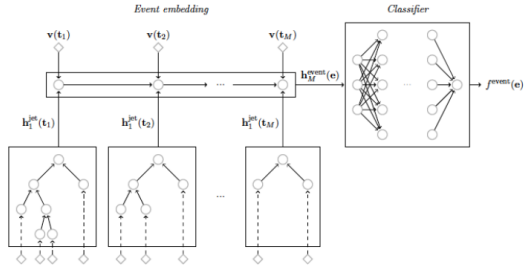

Deep and HEP

Ziyu Guo's thesis (with Y. Coadou)

- Deep learning in the search for ttH with the ATLAS experiment at the Large Hadron Collider
- 2. Rely on the physical process to design the NN structure
- → Better results than DNN. Comparing now with BDTs

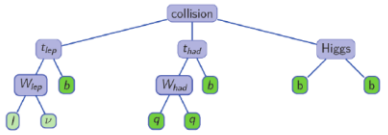

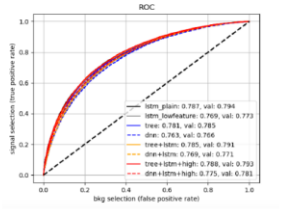

➔ **Les nouvelles collaborations ont lancé un groupe de travail sur le sujet**

Fast inference of Deep Learning Applications with FPGAs

Maurizio Pierini

CERN

hls 4 ml

erc European Research Council

Deep Learning and LHC Big Data

- *Possible solution to the HL-LHC problem: modern Machine Learning to be faster and better in what we do today, freeing resources for new ideas*
- *This ML deployment need to happen in between collisions and data analysis (trigger, reconstruction, …), where freeing resources will make a difference*

13

Porting Deep Learning to Trigger/DAQ system

erc European Research Council

HLS4ML

- *HLS4ML aims to be this automatic tool*
  - *reads as input models trained on standard DeepLearning libraries*
  - *comes with implementation of common ingredients (layers, activation functions, etc)*
  - *Uses HLS softwares to provide a firmware implementation of a given network*
  - *Could also be used to create co-processing kernels for HLT environments*

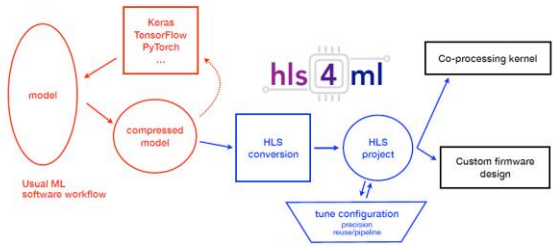

39

Why do we need to process data in real time?

What are the relationships between physics, dataflow, and optimal processing architectures?

How do we calibrate and monitor real-time processing systems?

Data @ some current/future experiments

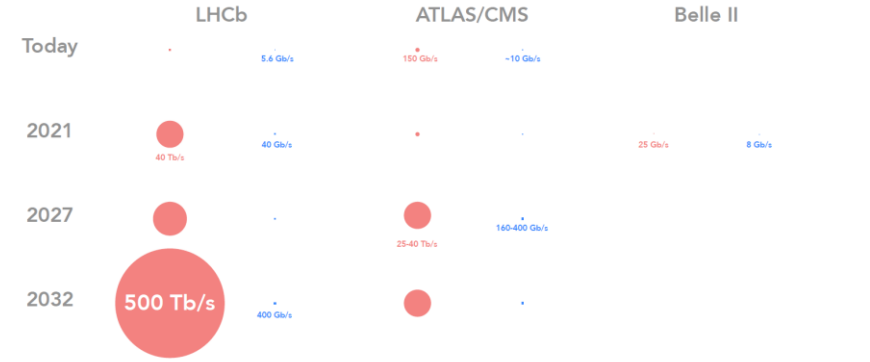

Data rate read out from detectors by the DAQ system for further online or offline processing

Data written to permanent storage for long-term analysis

8

Putting that DAQ into context

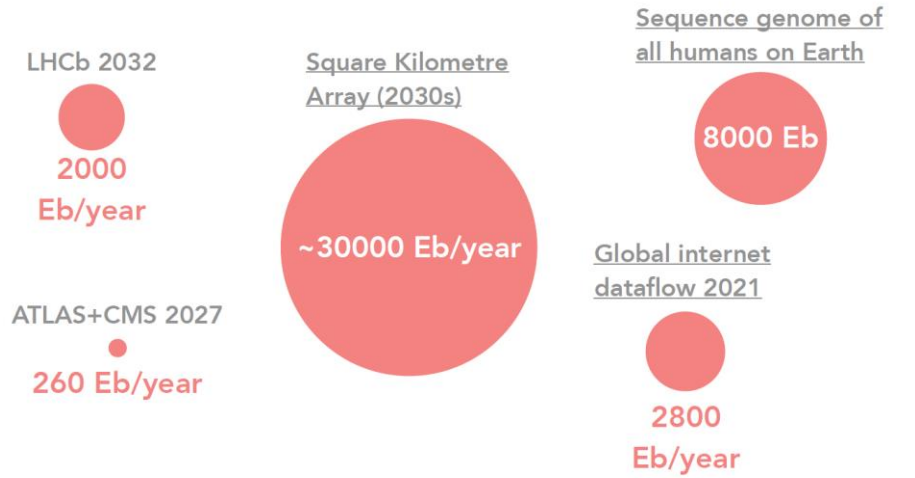

Enjeux : Réduction des données en temps réel pour éviter le stockage sur disque (très couteux) :

➔Embarqué les algorithmes dans les nœuds de decisions

➔Optimiser les traitements (classifications….DL)

➔Utiliser un mixte GPU, MPPA, FPGA

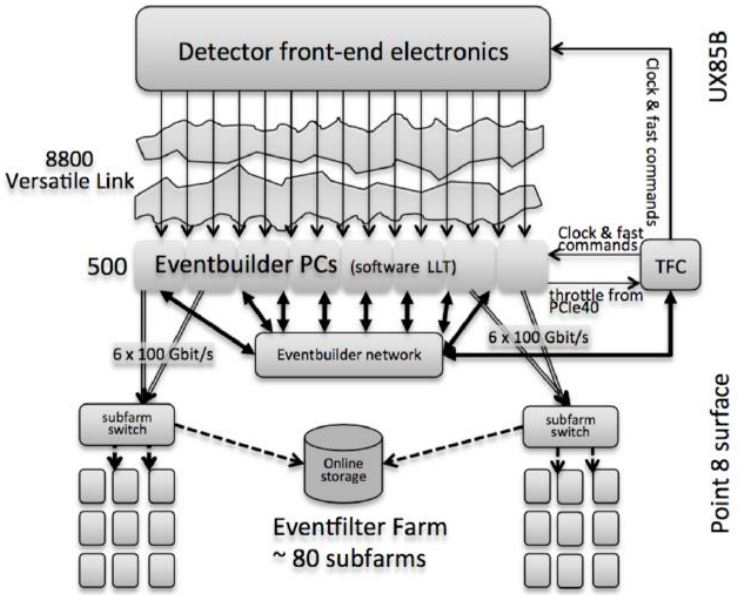

Keep in mind that to reduce this data volume you can either select a subset of events, compress the events, or do a mixture of these two things.

(Annual) cost of storing data to disk

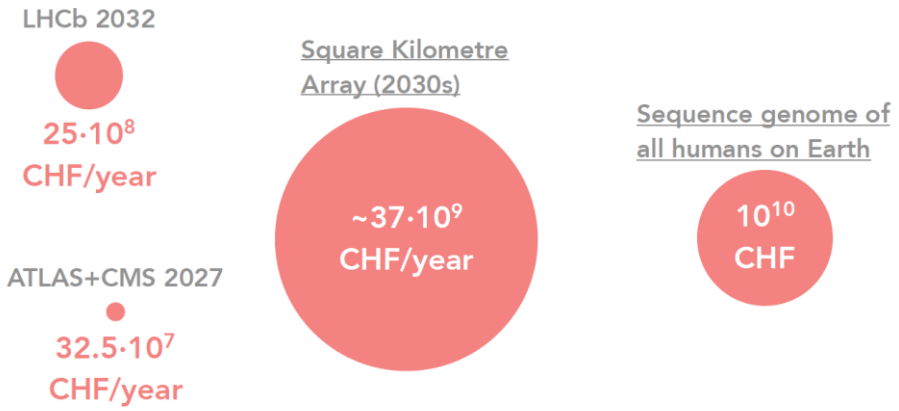

Haven't even discussed the cabling/networking: game over

FPGA Compute Acceleration in HEP

Christian Färber
CERN Openlab Fellow
LHCb Online group

On behalf of the LHCb Online group and the HTC Collaboration

Technologies émergentes pour systèmes DAQ
IN2P3 School, Marseille
16.11.2018

Summary

- Results are very encouraging to use FPGA acceleration in the HEP field
- Comparing the energy consumption with CPUs show better performance for FPGAs (getting a greener CERN computing ?)
- Programming model with OpenCL very attractive and convenient for HEP field, HLS also available
- Also other experiments want to test the usage of the Intel® Xeon®+FPGA with Arria10
- High bandwidth interconnect coupled with Arria® 10 FPGA suggests excellent performance per Joule for HEP algorithms! Don't forget Stratix® 10 … !

Algorithmes IA: Nouveaux Ordinateurs !!!

End of Moore's Law: What Next ?

Evolution or Revolution ?

- Improved Von Neumann (Adiabatic logic, computing in memory)
- Improved switching device (new transistor, other technology)
- New state variable (beyond hardware)

→ New information processing paradigm

What are the needs ?

- High computing power

- New fonctionnalities. More 'intelligence'. Processing of natural data (image, sound…..)

• Ultra-low power dissipation
• High energy efficiency

Algorithmes IA: Nouveaux Ordinateurs !!!

Impressive Recent Progress

IBM Raises the Bar with a 50-Qubit Quantum Computer

Google quantum computer test shows breakthrough is within reac

- Qbits

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

$$\alpha, \beta \in \mathbb{C} \quad |\alpha|^2 + |\beta|^2 = 1$$

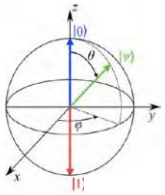

But…

- Several orders of magnitude slower than expected
- No quantum algorithm library
- Quantum state coherence time limited (100 □S only)
- Operation temperature often close to 0°K (-273 °C) !

- Quantum computers often resemble fundamental physics experiments !

Neuromorphic architectures

- The most complex information processing system known to date.
- 2% of the human weight but 20 % of the consumed energy
- In human, ~ 10^{11} neurons and ~ 10^{4} synapses/neurons

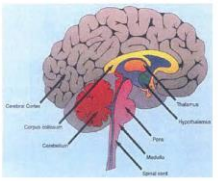

1 mm^3 cortex

$5x10^4$ to 10^5 neurons
$3 x10^8$ to $7x10^8$ synapses
4 km of axon
0,5 km of dendrites

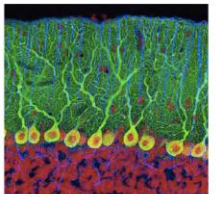

IBM TrueNorth architecture

Loihi Intel

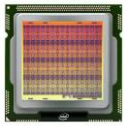

130 000 neurons and 130 M synapses

IBM TrueNorth architecture

1M spiking neurons and 256 M plastic synapses

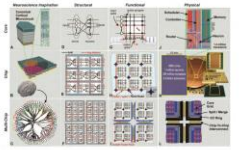

1 M neurons and 256 M synapses

Choix de méthode conceptuelle

Conception par la biologie

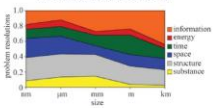

Biological effects arranged according to size/hierarchy.

Minimisation de l'énergie, environ 2% quelque soit l'échelle de réalisation tout en maintenant un % de temps respectable.

Conception par l'humain

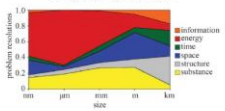

Engineering TRIZ solutions arranged according to size/hierarchy.

Minimisation du temps, quelques % dépendant de l'échelle de réalisation, la dépense en énergie est très importante.

Une méthode bio-inspiré donnera de meilleurs résultats.

Biomimetics: its practice and theory (2006)
Julian F. V. Vincent*, Olga A. Bogatyreva, Nikolaj R. Bogatyrev, Adrian Bowyer and Anja-Karina Pahl

Anotherbrain 06/12/2018

Attracteur dynamique: application

Changement de fonctionnalité par le choix de l'information en entrée

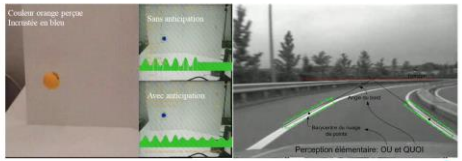

Anotherbrain 06/12/2018

Perspectives à l'IN2P3

C'est une révolution en marche qui va faire évoluer nos métiers

- **Expériences au CERN :**
  - Réduction des données sur CMS, ATLAS, LHCb
  - Test RN dans les FPGA HLS4ML
  - Test le calcul accéléré

- **Produit en Physique Nucléaire:**
  - LPC Caen : Test MPPA avec une machine KALRAY
    → filtrage numérique par voies
    → Evolution du Produit DAQ FASTER
  - IPN Lyon: DL pour AGATA (protons, gamma..)

- **R2D2 au CENBG**
  - Classifier (déconvolution) embarqué sur système SAM (GET)
  - Logiciel DAVIS – évolution vers SoC-

Remerciements

« Si j'ai vu plus loin, c'est en montant sur les épaules de géants. »
Isaac Newton – 1675 –

- **Fred Goulding (1925-2013)** was a pioneer in the application of modern electronics to nuclear physics research He was internationally recognized for the development and application of semiconductor radiation detectors, and was among the earliest to implement computer control and data acquisition methods for science. He served as group leader in the Nuclear Chemistry Division (now Nuclear Science), and a department head in the Engineering Division.

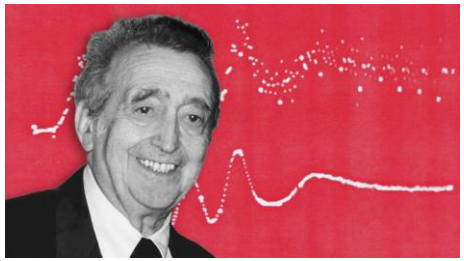

- **Emilio Gatti (Torino, 18 marzo 1922 - Milano, 9 luglio 2016)** began his activity in 1948 in Nuclear Electronics at CISE, Milan Italy. Professor at Politecnico di Milano since 1951 (full professor since 1957 and emeritus in 1998). Life-long associated to the Italian Institute for Nuclear Physics, INFN (Istituto Nazionale di Fisica Nucleare), that he contributed to found in 1951. Since 1973 Senior Visiting Scientist at Brookhaven National Laboratory, for 30 years collaborating with Veljko Radeka and the late Pavel Rehak. Among his most brilliant achievement which played a revolutionary role:

- **the charge preamplifier** to perform minimum noise measurements of the charge deposited in radiation detectors (at the time ionization chambers)
- the conception jointly with the friend and colleague Pavel Rehak of the Sideward Depletion together with the invention of the Silicon Drift Detector
- the **sliding scale technique** to improve the differential non-linearity in ADCs.
- countless relevant studies on **optimum filtering** that resulted also in a lengthy monograph jointly written with the late friend and colleague Franco Manfredi.

Remerciements

« Si vous n'êtes pas capable d'expliquer quelque chose à un enfant de 6 ans, c'est que vous ne le comprenez pas vous-même»

Albert Einstein

- **Pier Francesco Manfredi (1935-2015) :**

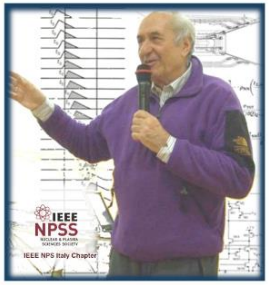

Franco Manfredi gave an essential contribution to the understanding of radiation effects on the noise performance of front-end devices and to the development of techniques enabling readout electronics to operate with a large signal-to-noise ratio at high radiation levels

Pavel Rehak 1945 - 2009

has had a deep interest in physics, but most of his work was motivated by his belief that detector developments are among the main forces responsible for the progress in physics and in other natural sciences. He was truly a renaissance physicist, in that he could delve deeply into various areas of physics.

Veljko Radeka 1930 -

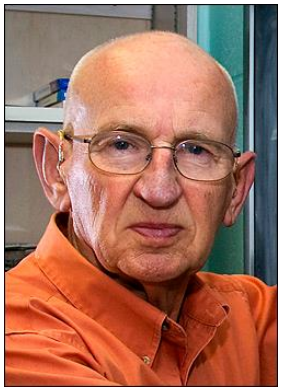

LOW-NOISE TECHNIQUESIN DETECTORS
Brookhaven National Laboratory, Upton, New York 11973

Remerciements

Pour ces présentations, j'ai utilisé les documents de plusieurs écoles IN2P3 (Microélectronique, Electronique analogique et Data Acquisition :

- **Laurent Leterrier – LPC Caen -**
- **Christophe de La Taille – Omega -**
- **Philippe Vallerand – LAL -**
- **Laurent Royer – LPC Clermont-Ferrand**
- **Shébli Anvar – CEA/IRFU -**
- **Olivier Gevin – CEA/IRFU -**
- **Pascal Baron – CEA/IRFU -**
- **Selma Conforti – OMEGA -**
- **Cedric Cerna – CENBG -**
- **Federic Perrot – CENBG -**

Et l'ensemble des **intervenants** de l'ANF ***DAQ émergents*** du mois de Novembre 2018 et ***Informatique avancée*** du CNRS du mois de Decembre 2018

Les Technologies émergentes pour les DAQ

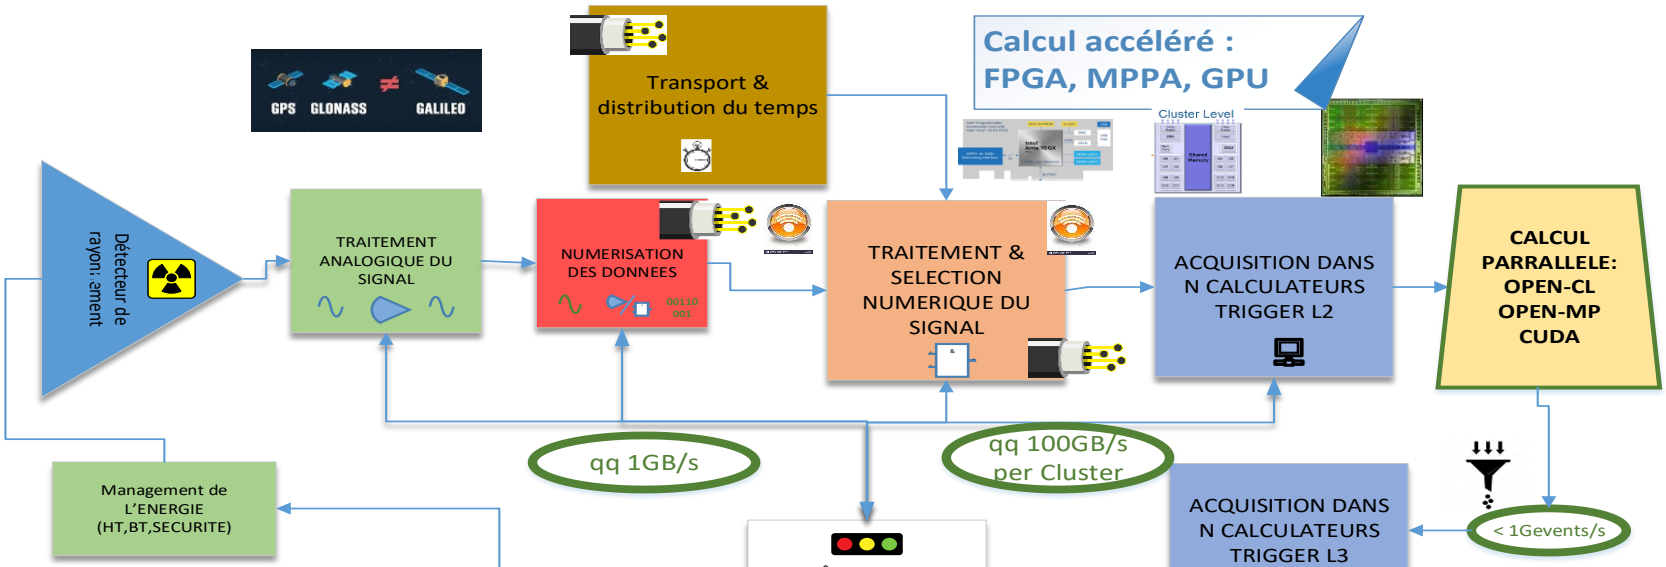

Les points de rupture :

- Transmission Optique et RF des FE
- Distribution d'horloges précises (qq ps)
- Calcul accéléré et massivement parallèle
- Algorithmes IA (*embarqués*)