


Concepts and Technologies for Data Acquisition in Particle Physics Experiments


Outline:

- Introduction
- Some historical remarks
- Challenges of modern experiments
- Trigger, data transport and data filtering
- FPGAs and High Speed Data Transport
- Case study: the Belle II Experiment
- Outlook and some final words

Wolfgang Kühn, Universität Giessen




ON THE AUTOMATIC REGISTRATION OF α -PARTICLES, β -PARTICLES AND γ -RAY AND X-RAY PULSES



A device is described which makes it possible for detected particle pulses to make records automatically on chronograph paper, which records can be conveniently and accurately studied at leisure.

Phys. Rev. 13, 272 (1919)



Alois F. Kovarik
Sheffield Scientific School
Yale University
New Haven, Conn.
January 25, 1919



LHCb

ATLAS

CERN Meyrin

PS 628 m

CERN Prévessin

SPS 7 km

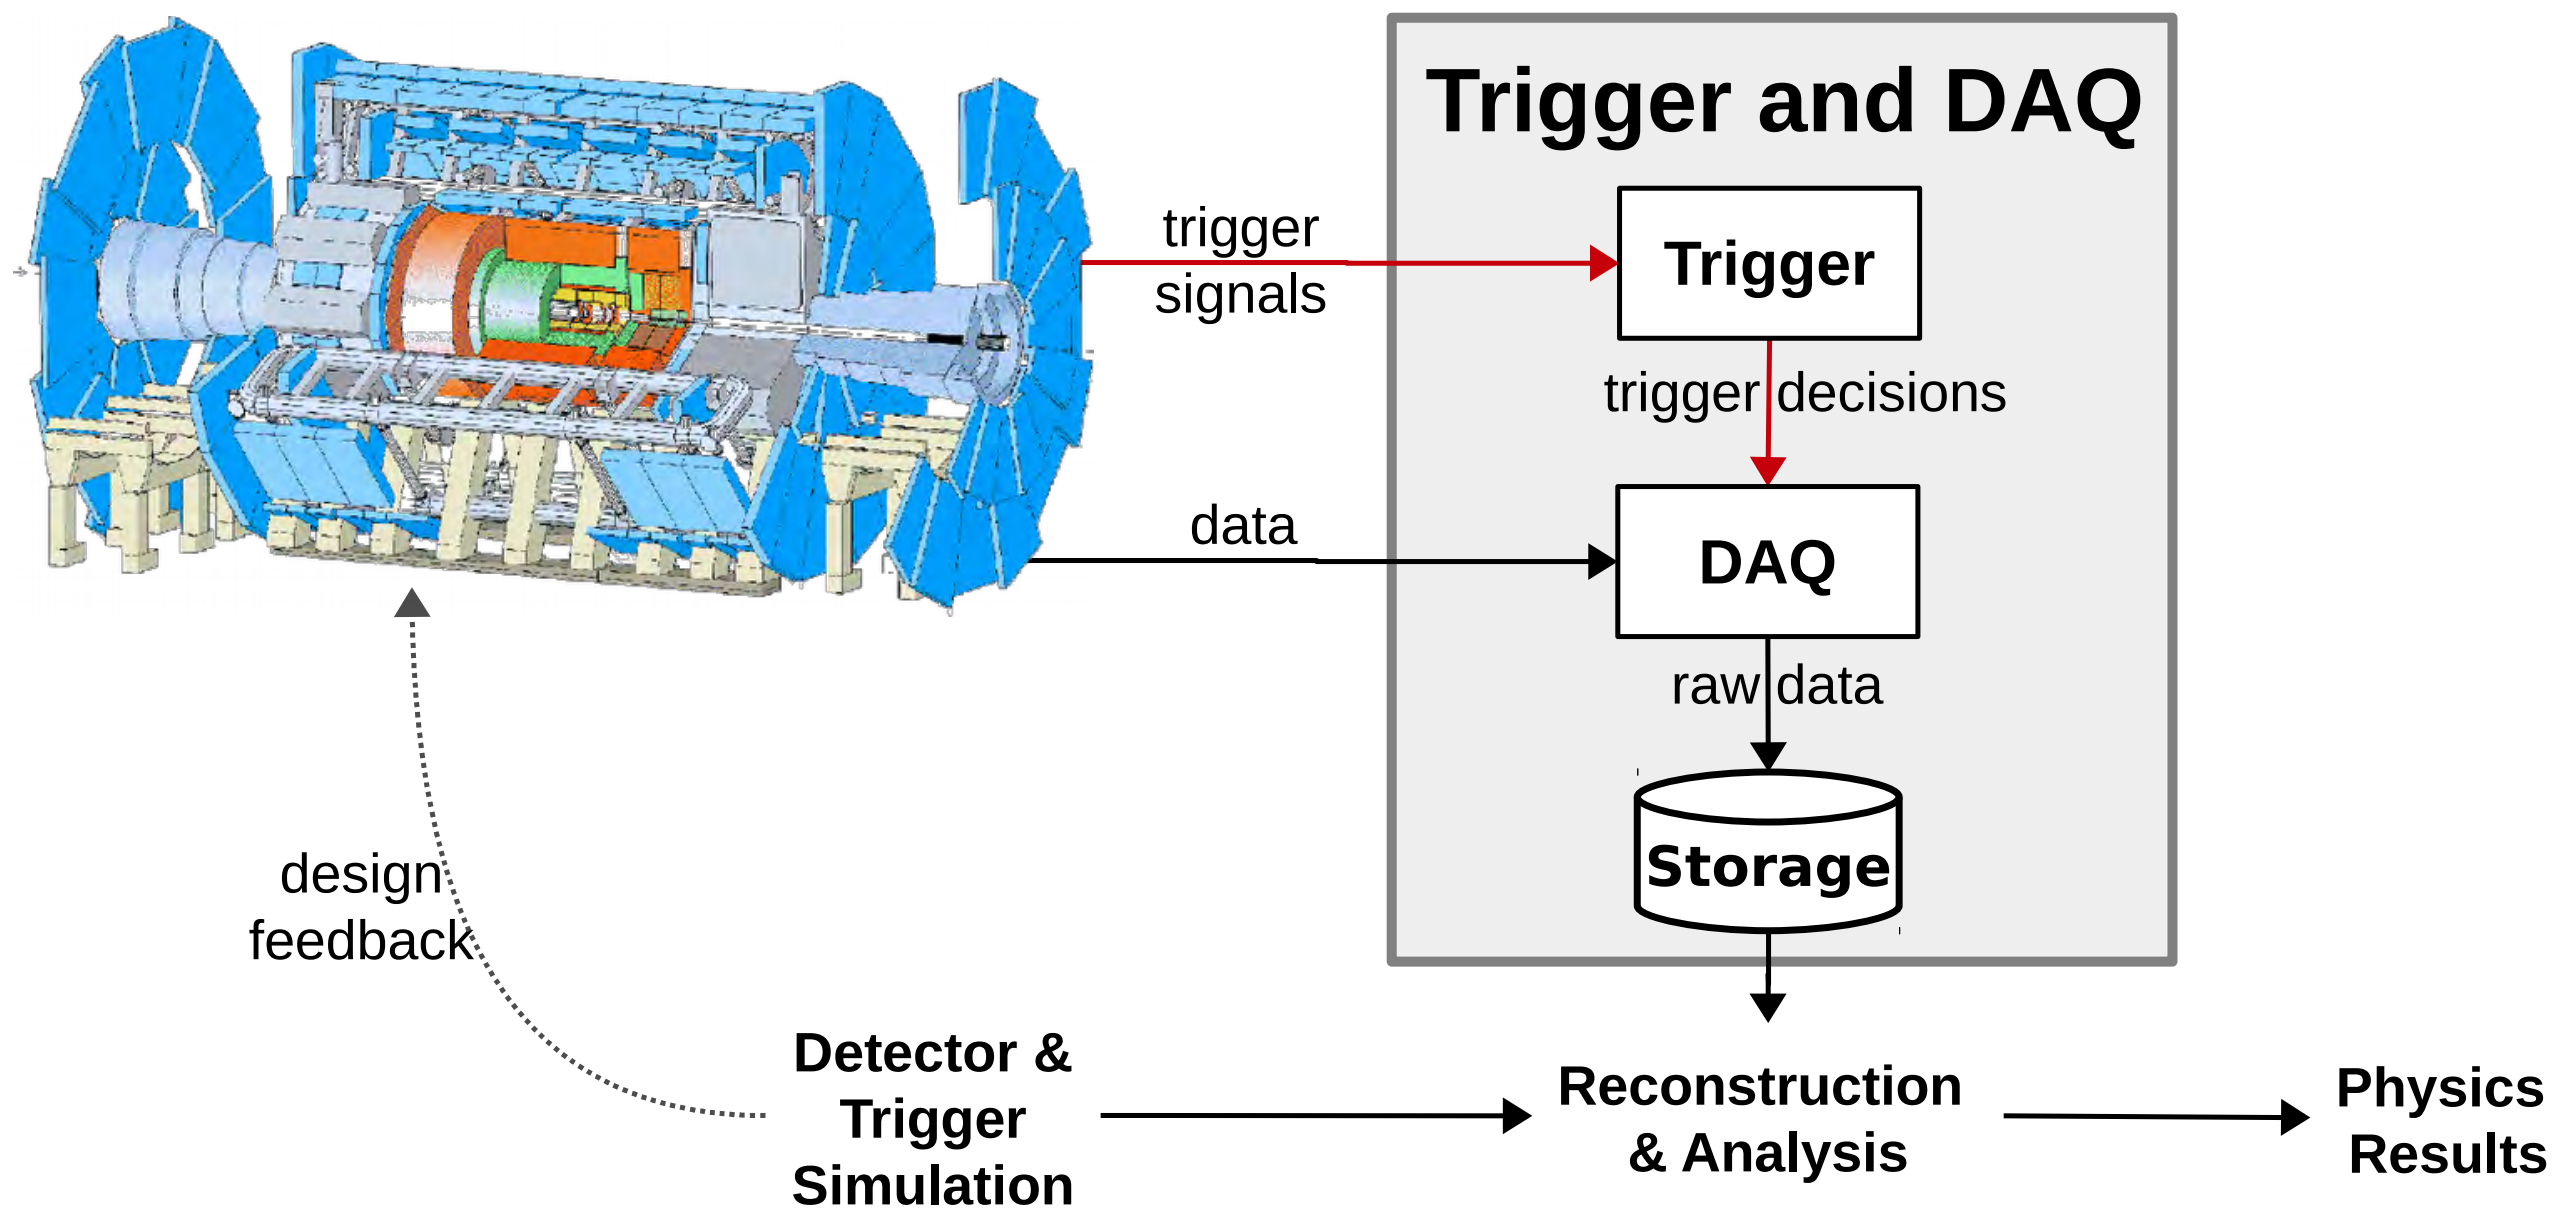
ALICE

CMS

LHC 27 km

SUISSE
FRANCE

Trigger and Data Acquisition




A perfect detector would be able to


- detect **charged particles**
 - charged leptons, charged hadrons, ...
- detect **neutral particles**
 - Photons, neutral hadrons, neutrinos
- perform **particle identification**
- **precisely measure the energy and/or the momentum** of each particle
 - allow to construct **4-vectors** for all particles produced in an interaction
 - do so even at very **high interaction rates** (> 20 MHz ?)



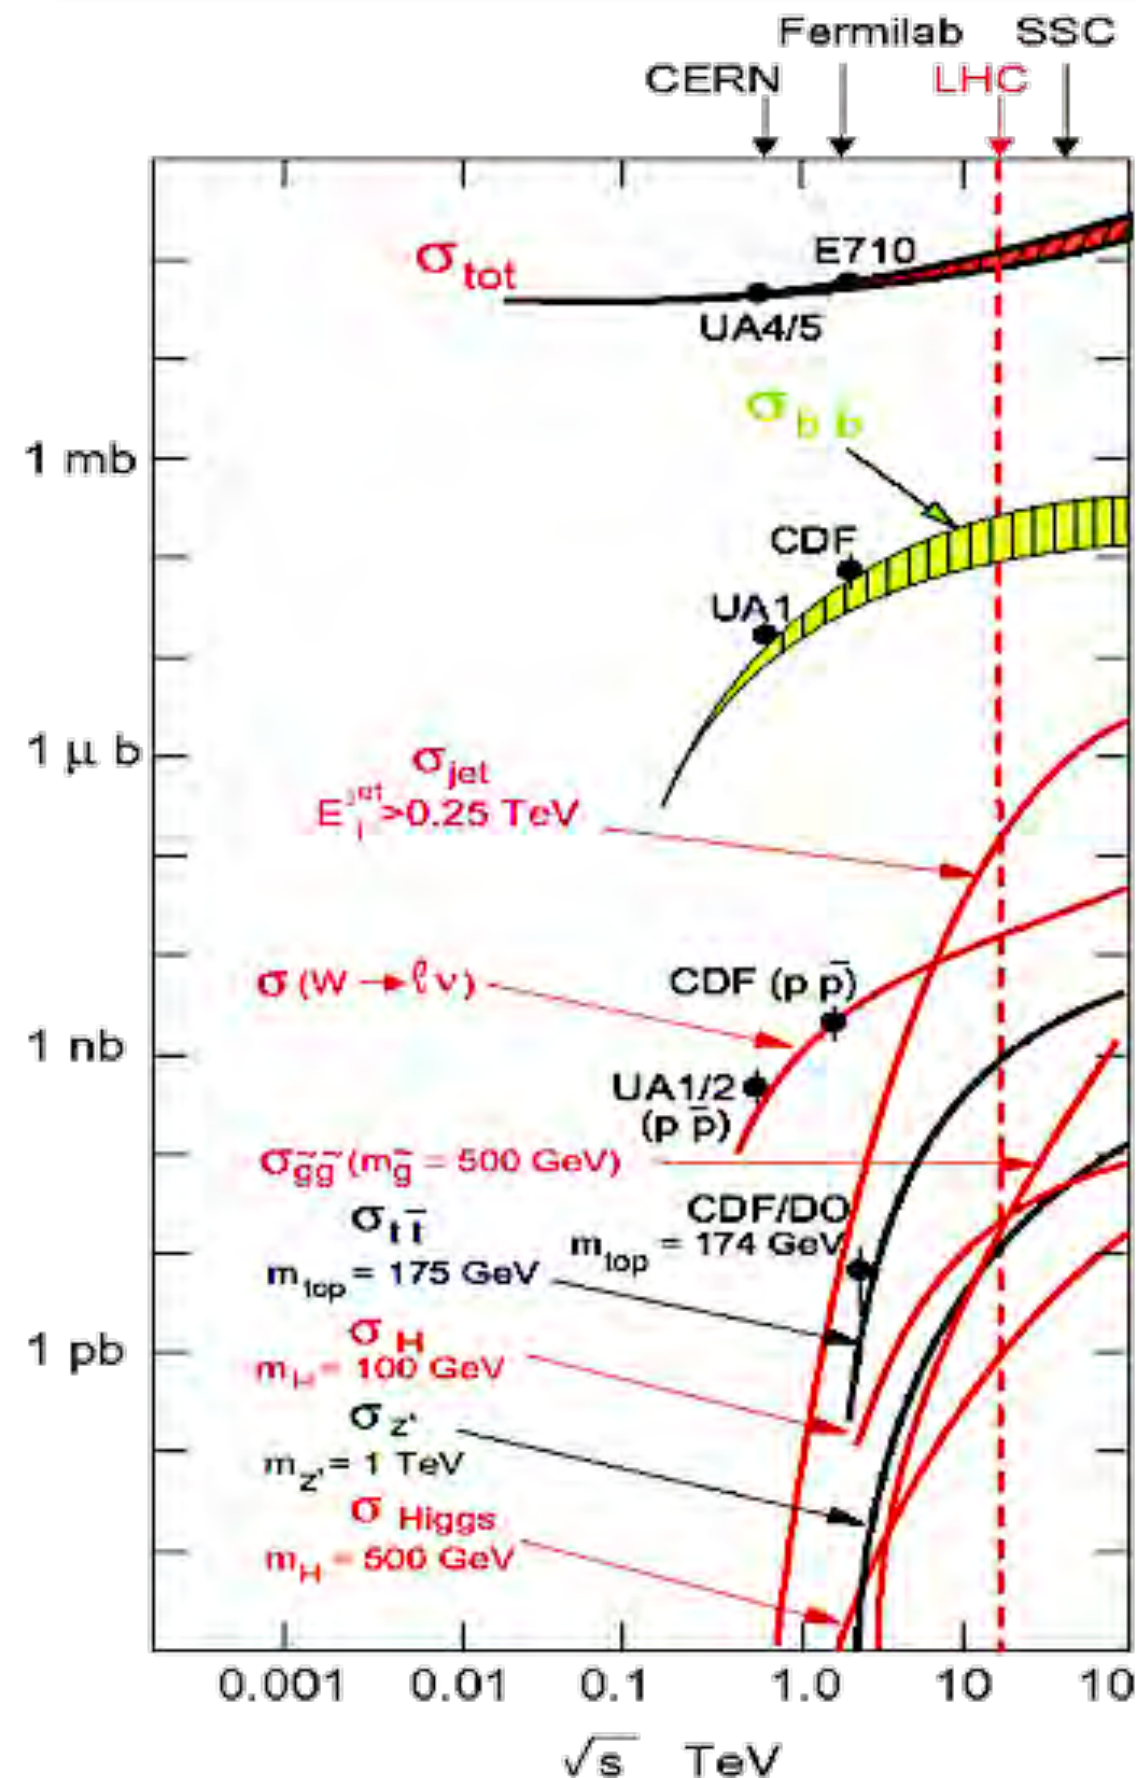
Typical Particle Physics Detector: Onion Shell Principle



Example: CMS (CERN)



Physics with Hadron Colliders



Process	Cross section (nb)	Production rates (Hz)
	@ 14 TeV	@ $\mathcal{L} = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$
inelastic	10^8	10^9
$b\bar{b}$	5×10^5	5×10^6
$W \rightarrow \ell \nu$	15	150
$Z \rightarrow \ell \ell$	2	20
$t\bar{t}$	1	10
H	0.05	0.5

Problem:

- Interesting physics channels have production rates as low as 1 event in 10^9 interactions

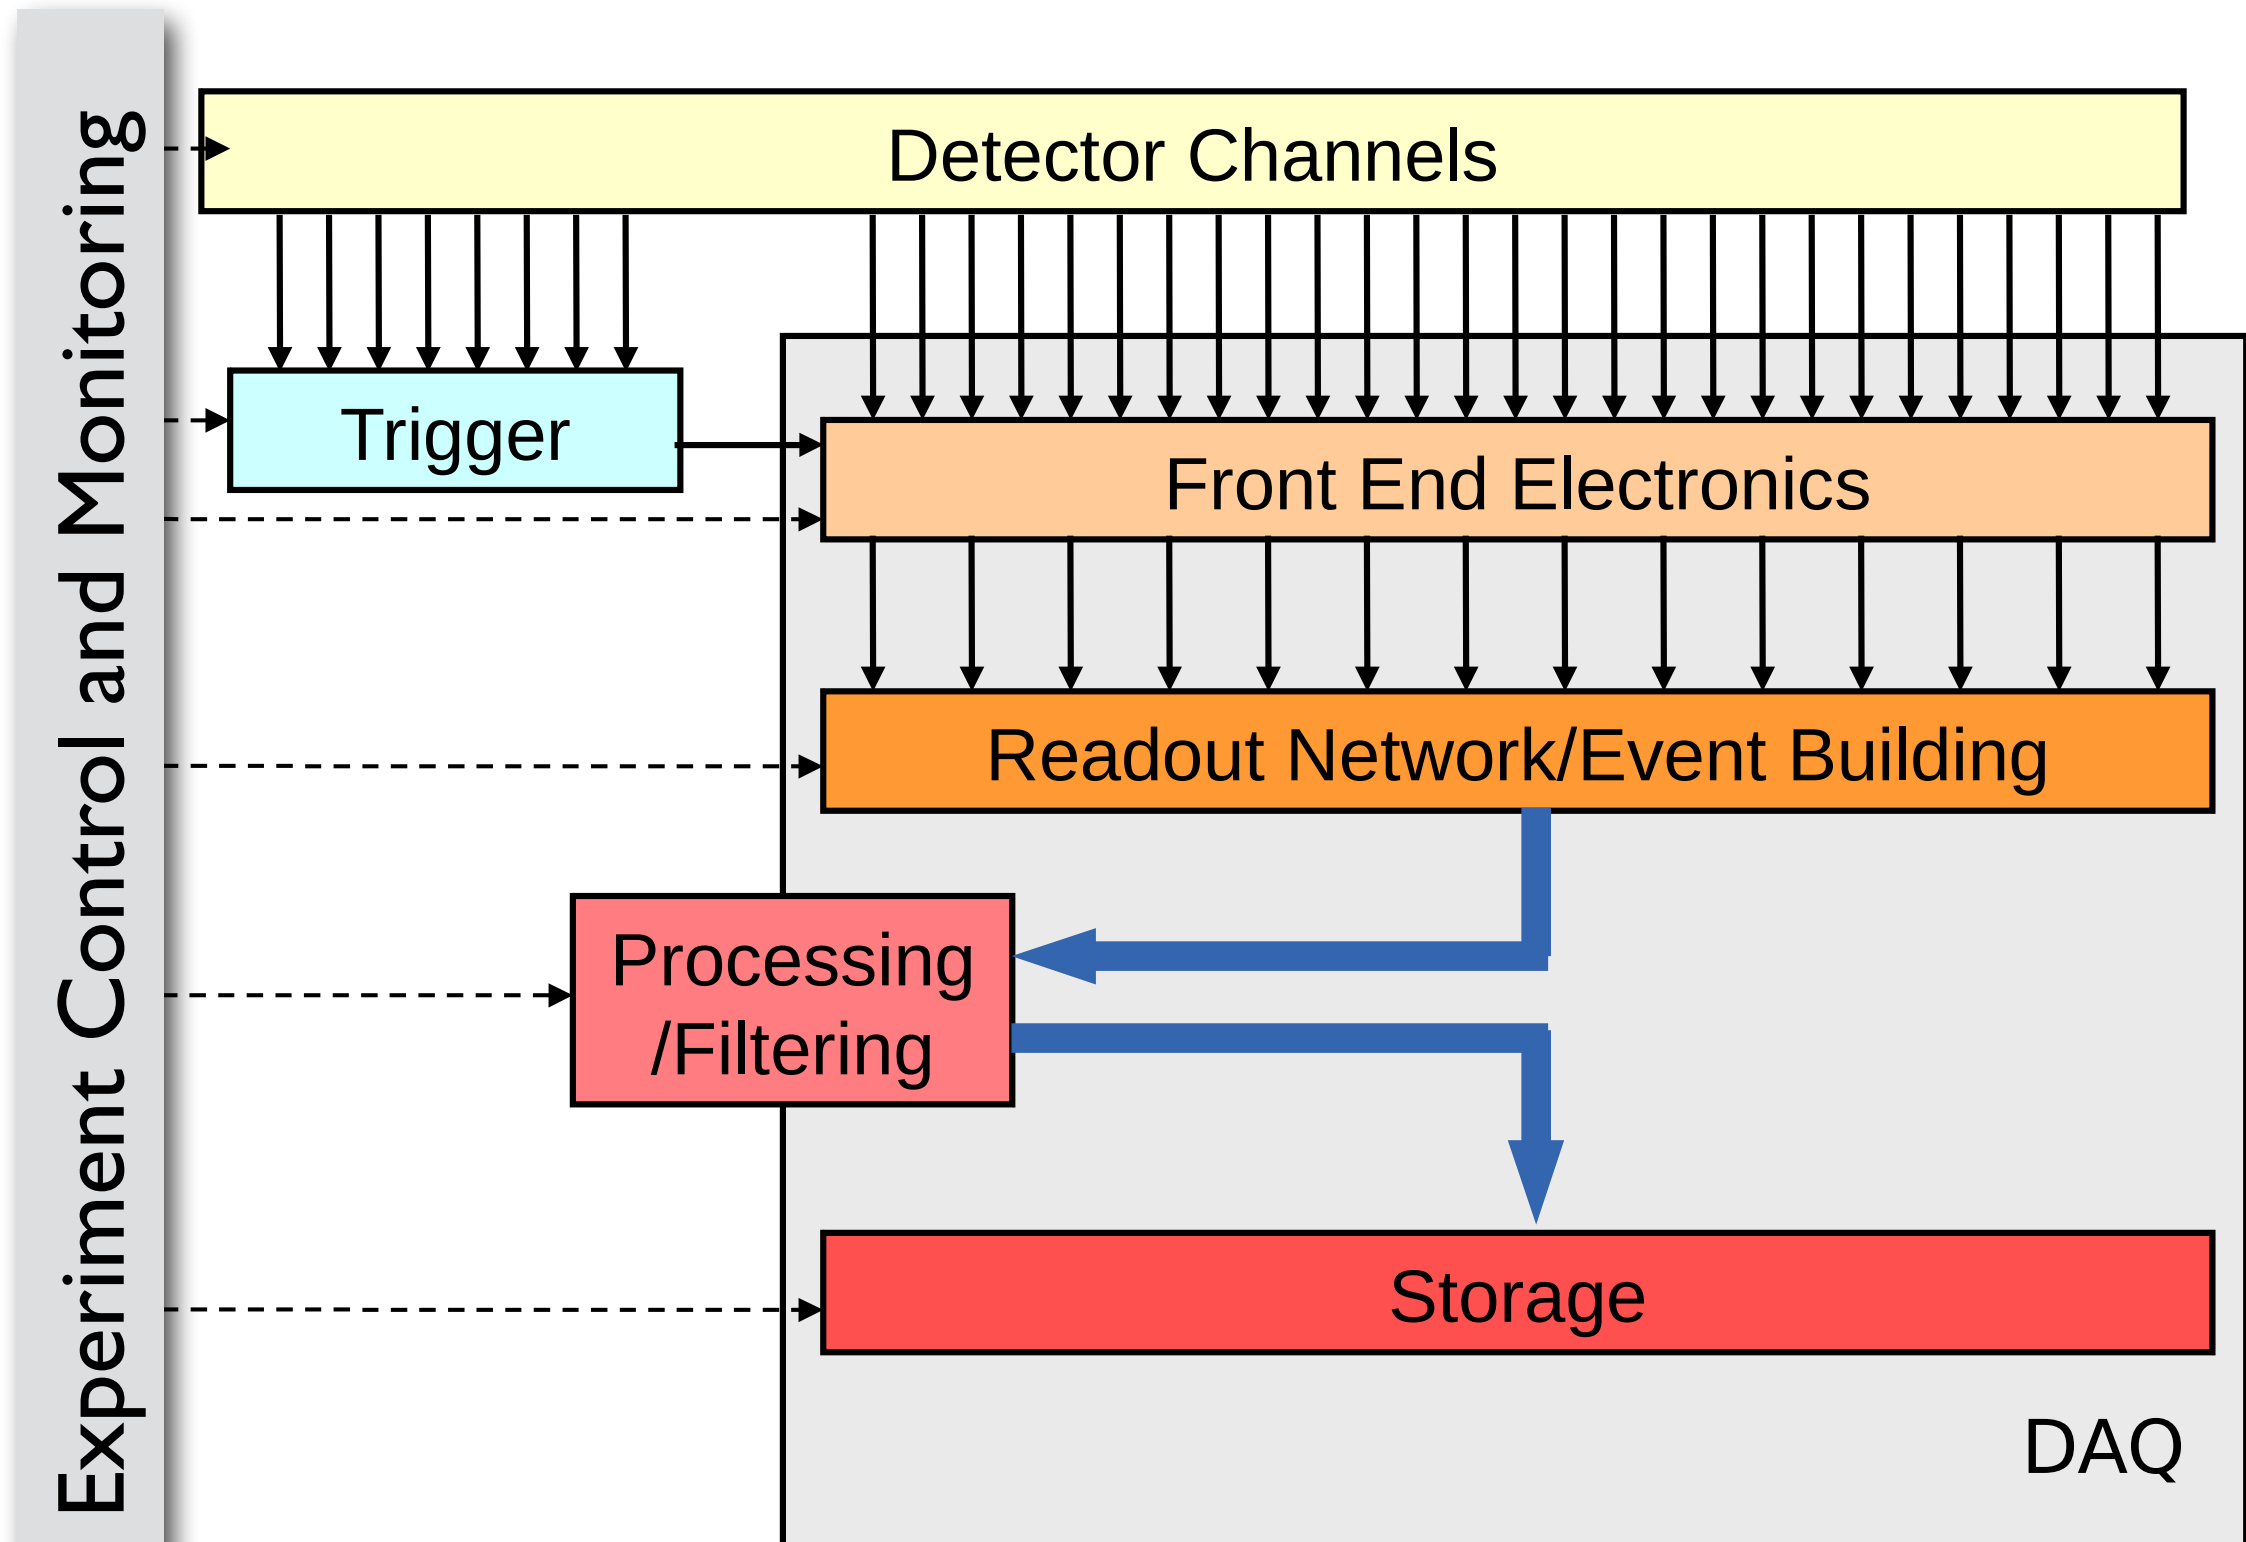
Options:

- Store everything and do the selection in the offline analysis
- Find selective triggers allowing to record only interesting physics
- Event filtering in order to effectively suppress events which are not interesting

Options for triggering / event filtering

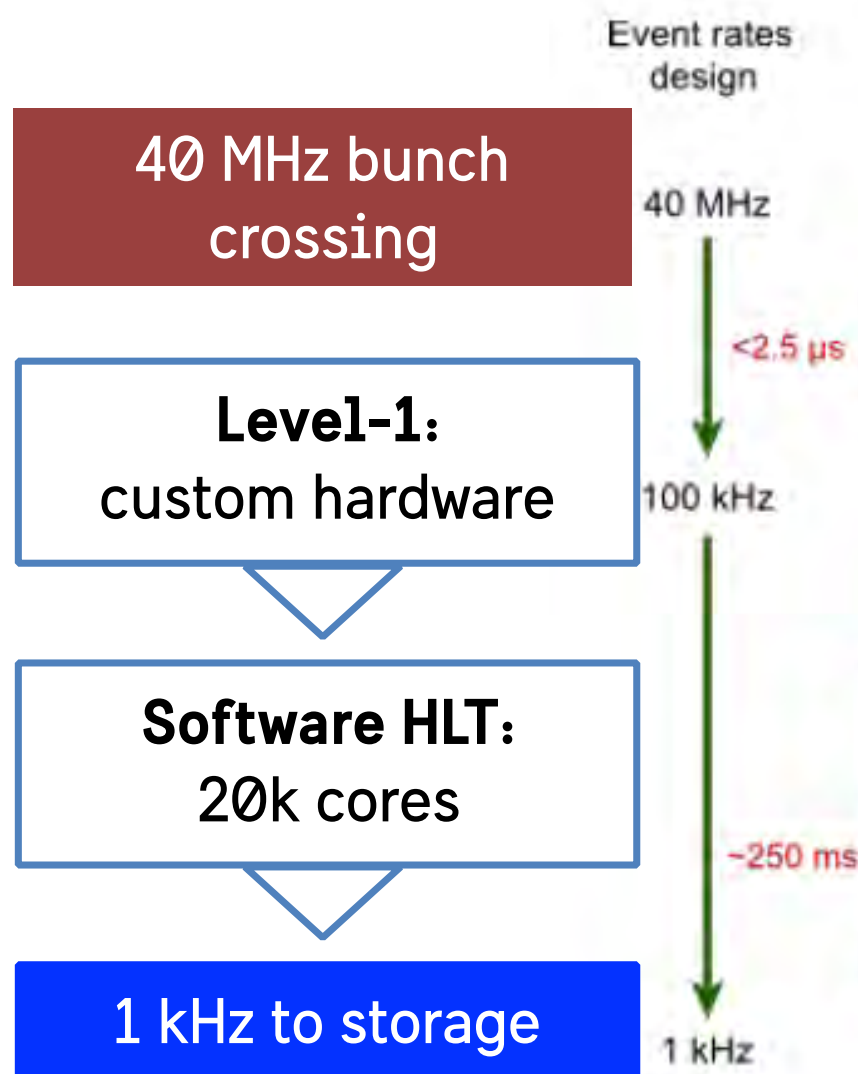
- **Store everything** and do the selection in the offline analysis ?
 - Limitation: storage capacity and processing resources for offline analysis
 - LHC: 40 MHz bunch crossing,
 - ATLAS/CMS : 60 TB/s raw data rate !!!
 - **Impossible to store !**
 - LHCb: running at lower luminosity
 - possible (but expensive !)
- **Find simple, selective triggers allowing to record only interesting physics**
 - Suitable only for experiments with a very specific focus
 - **There is no simple Higgs trigger !**
- **Event filtering** in order to effectively suppress events which are not interesting
 - The most general approach
 - Can be done by a combination of veto-triggers and more sophisticated online-selection
 - Typically multi-level system, combination of hardware/firmware/software
 - First-level trigger, second level trigger (hardware)
 - High level event filtering software

Typical Triggered System




Trigger systems in ATLAS/CMS/LHCb

ATLAS

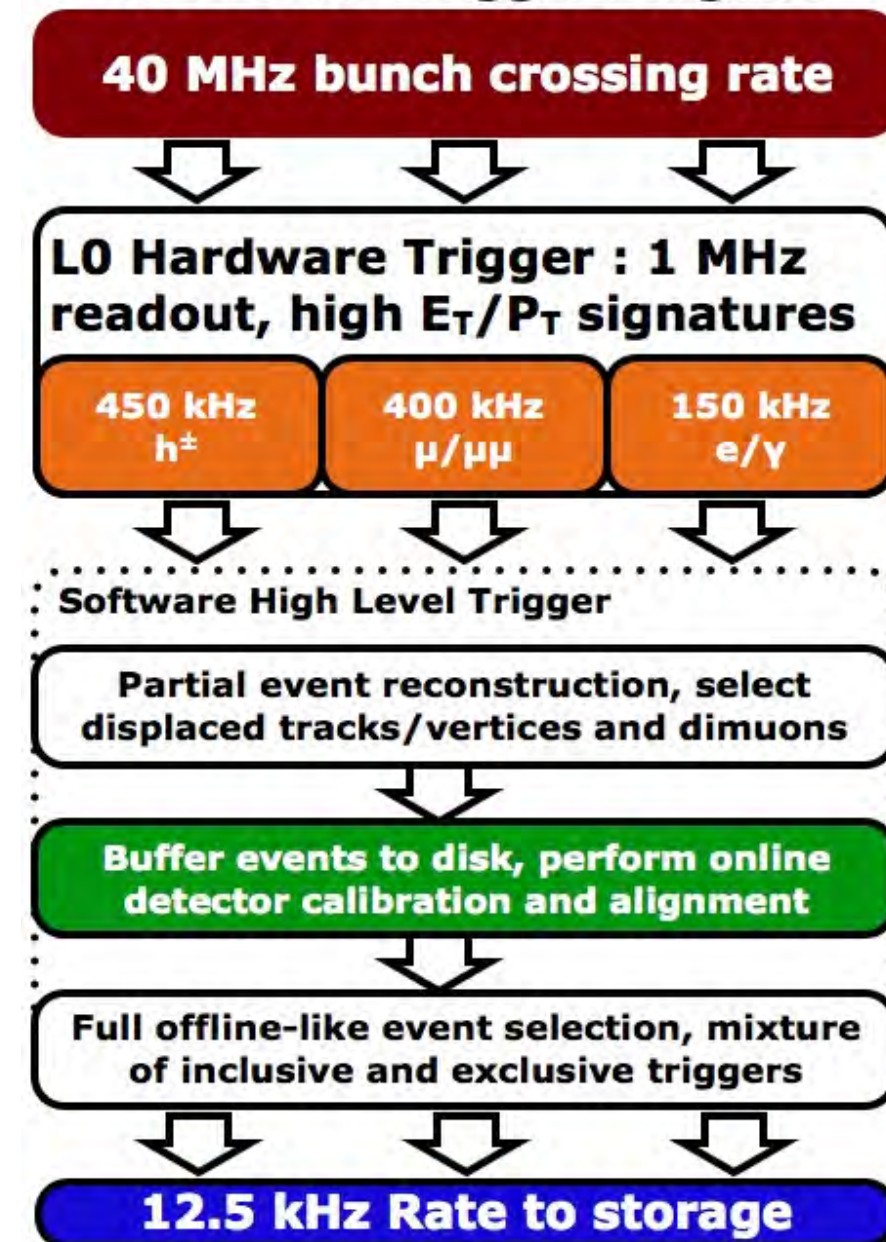


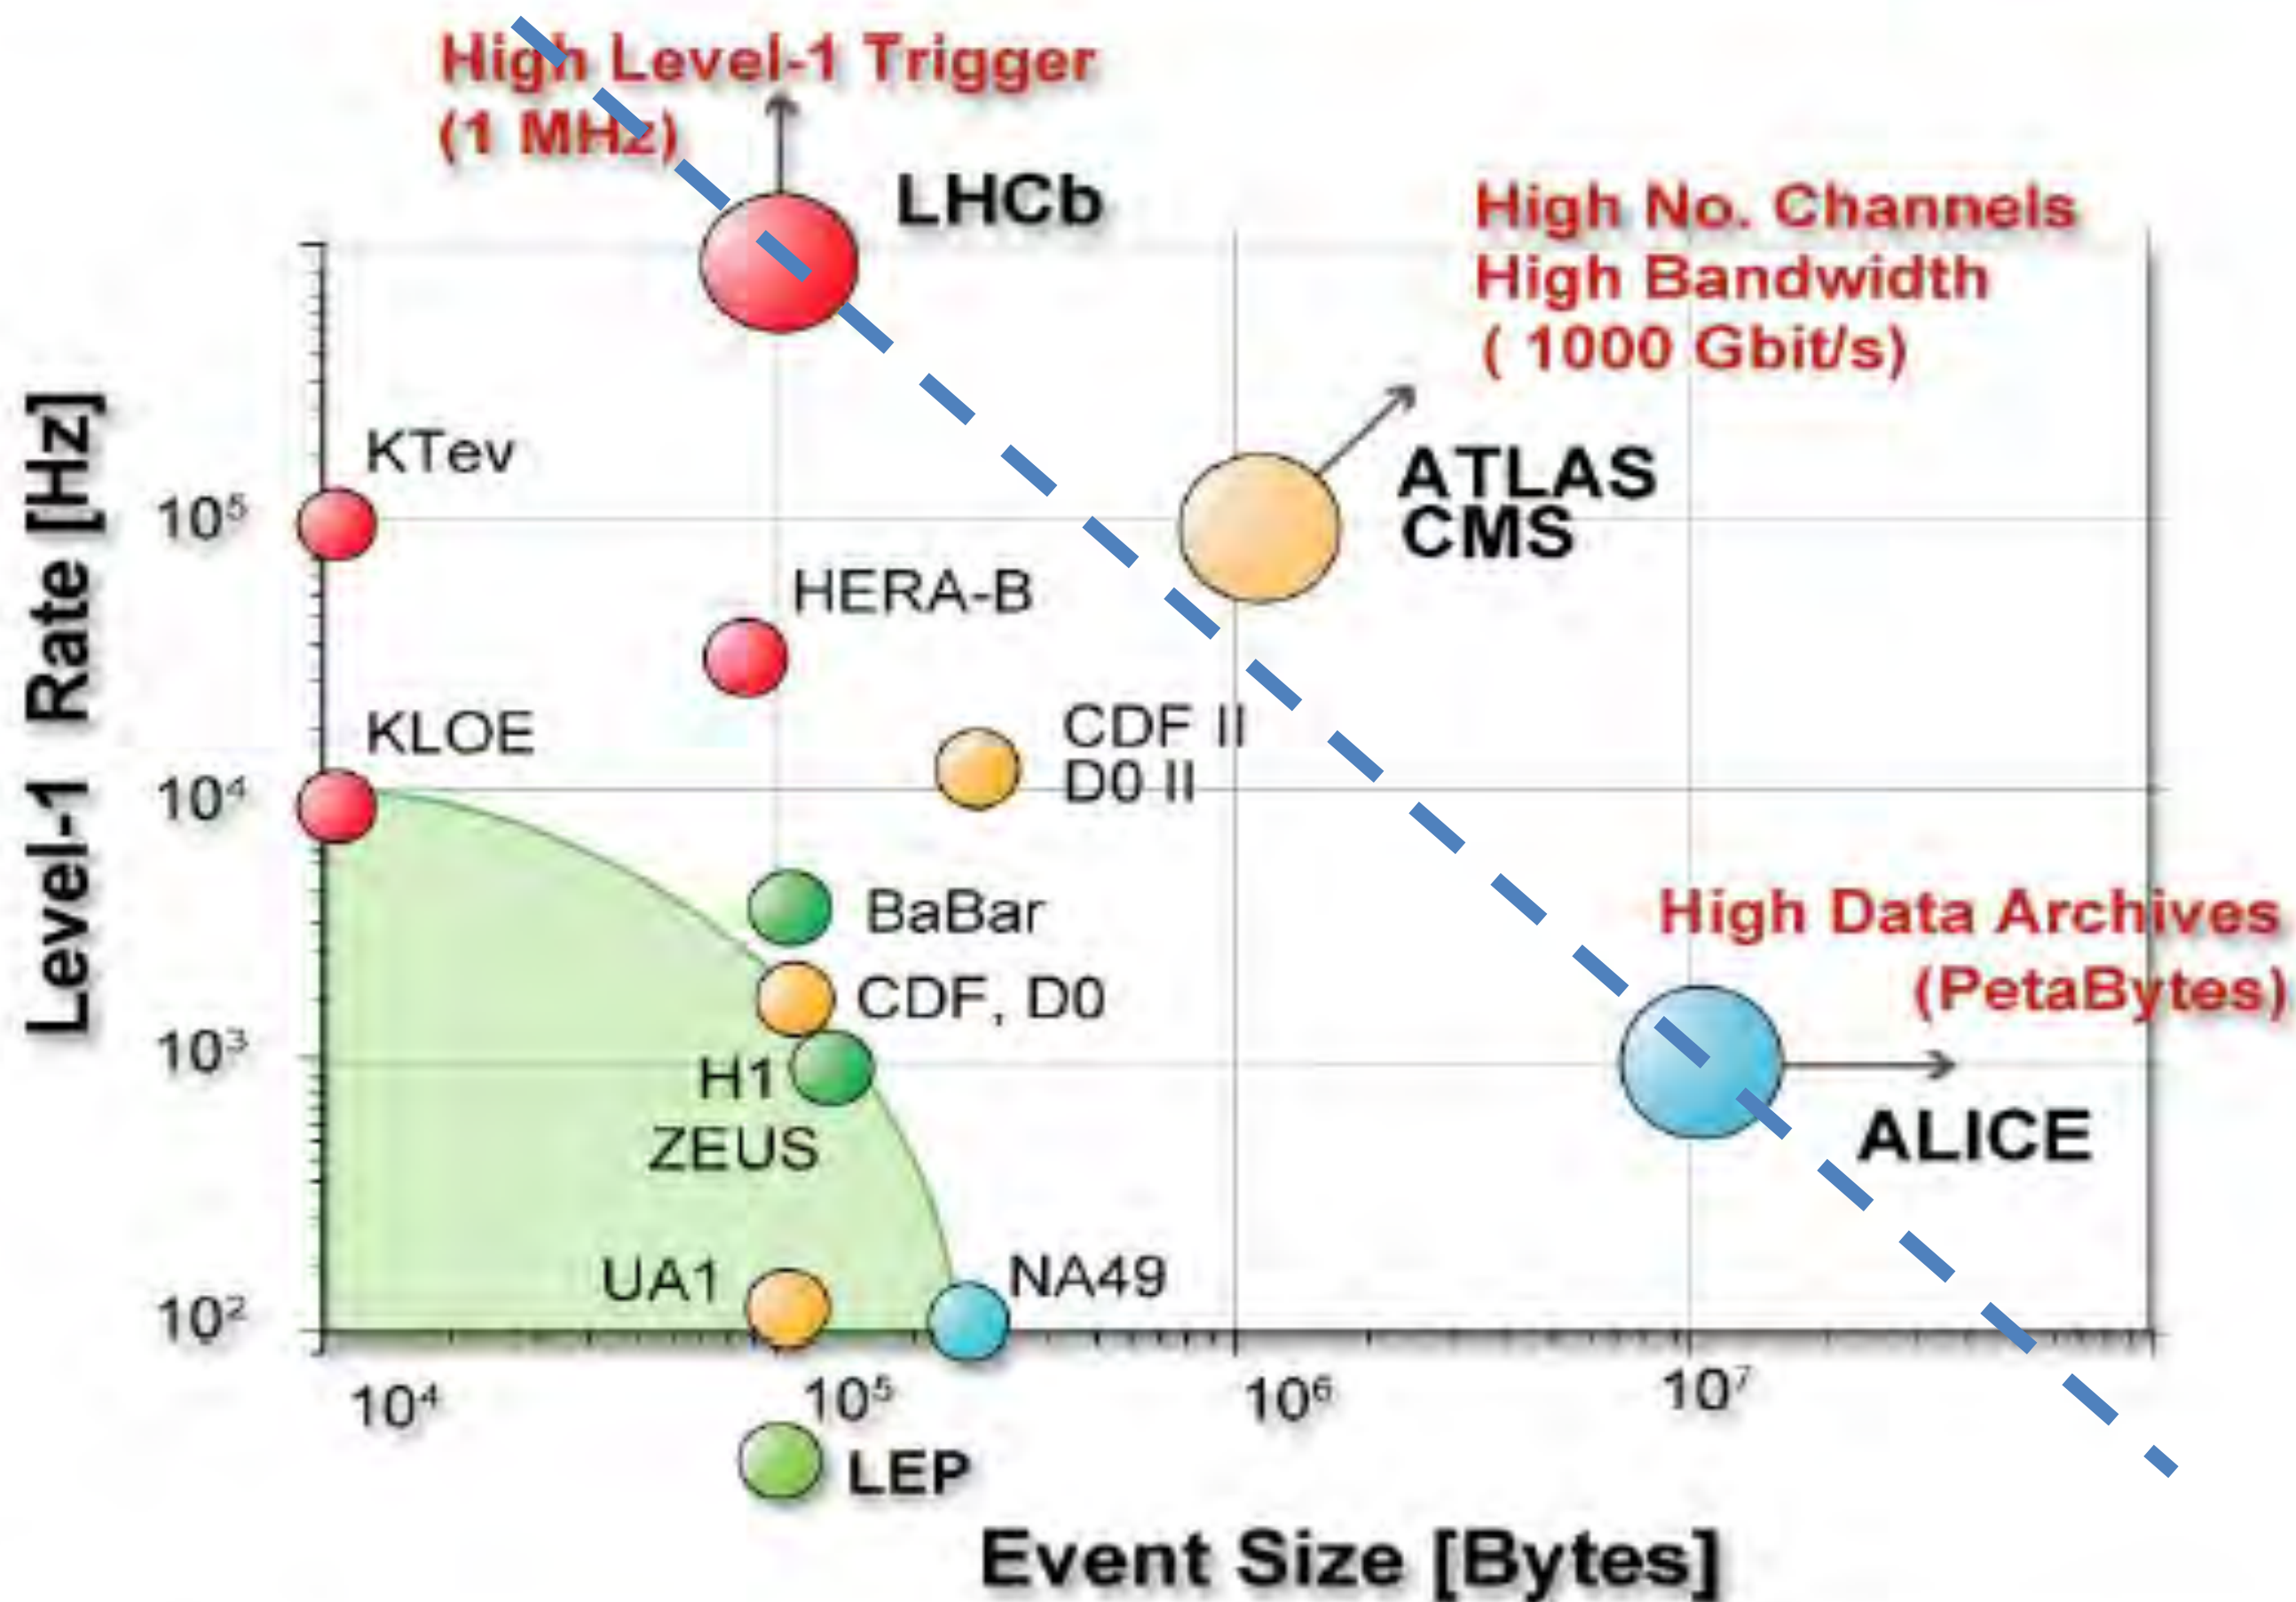
CMS



LHCb

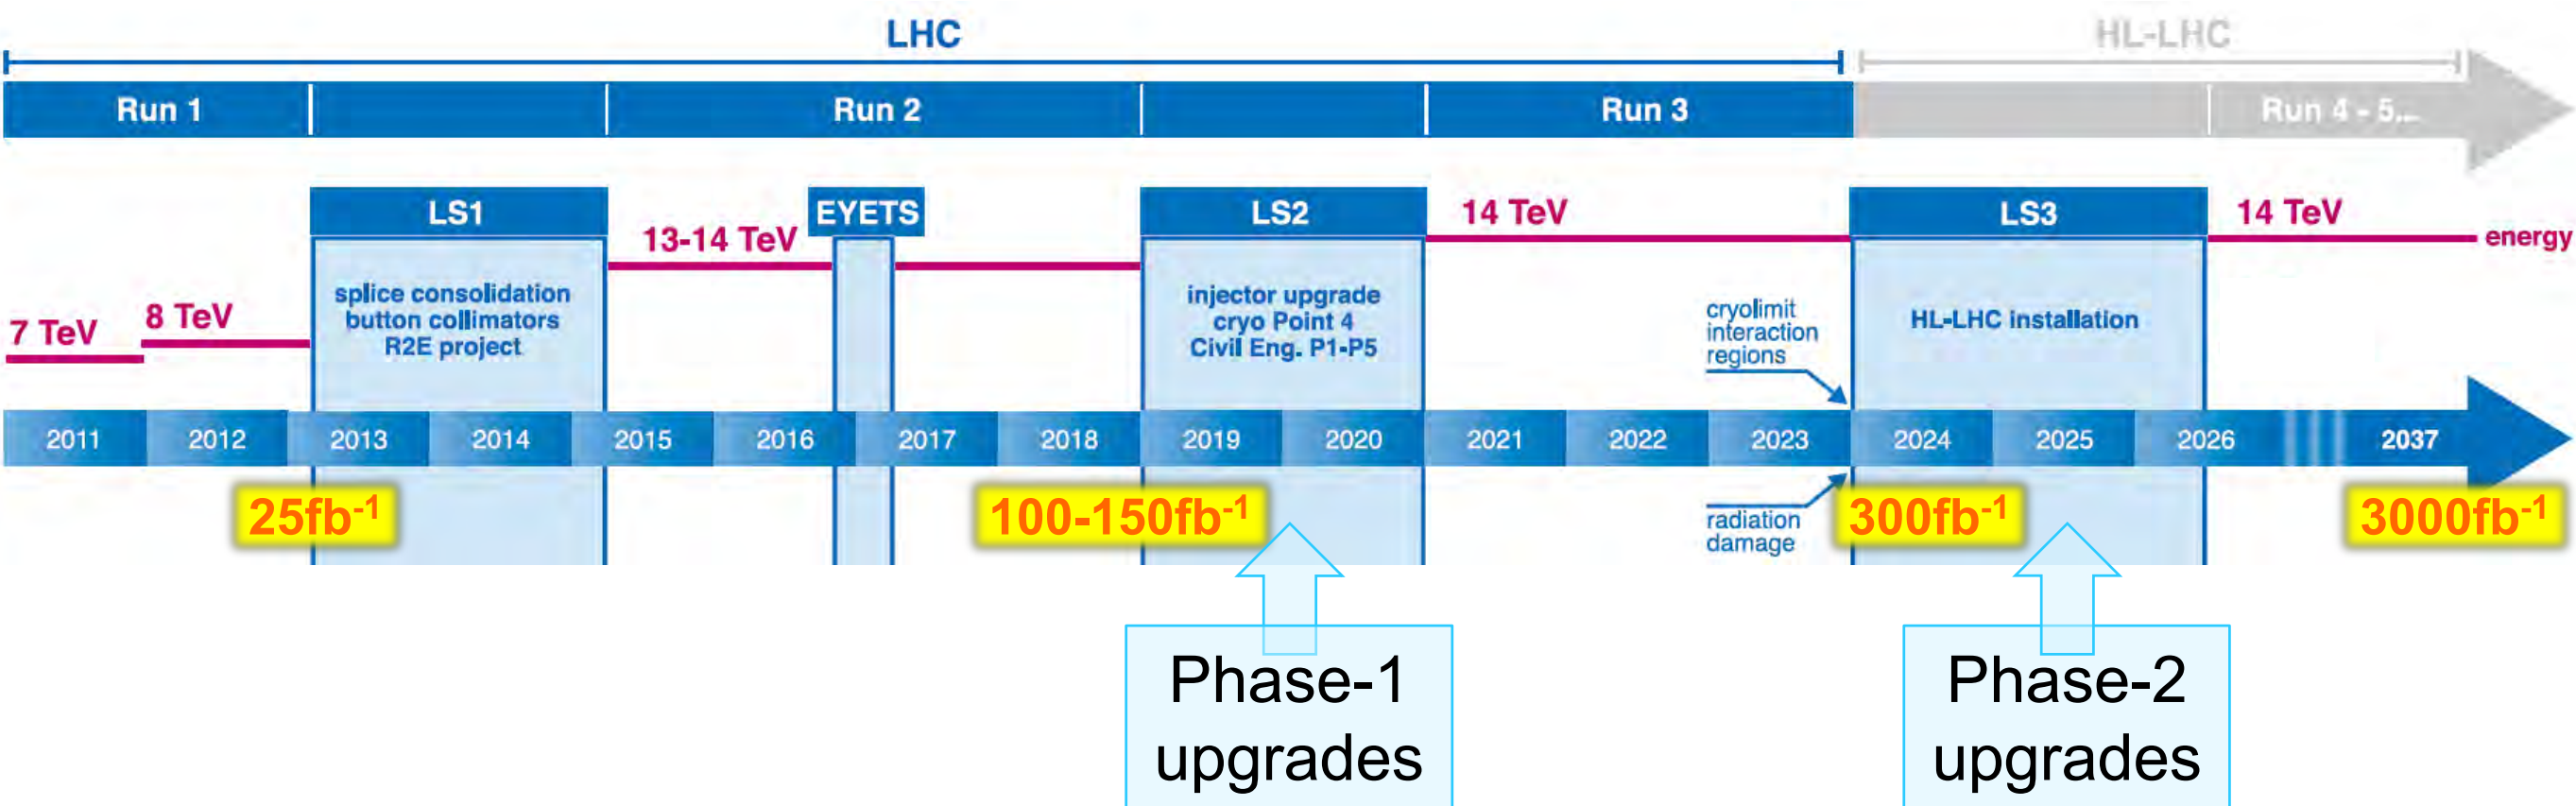
LHCb 2015 Trigger Diagram







Timelines towards HL-LHC



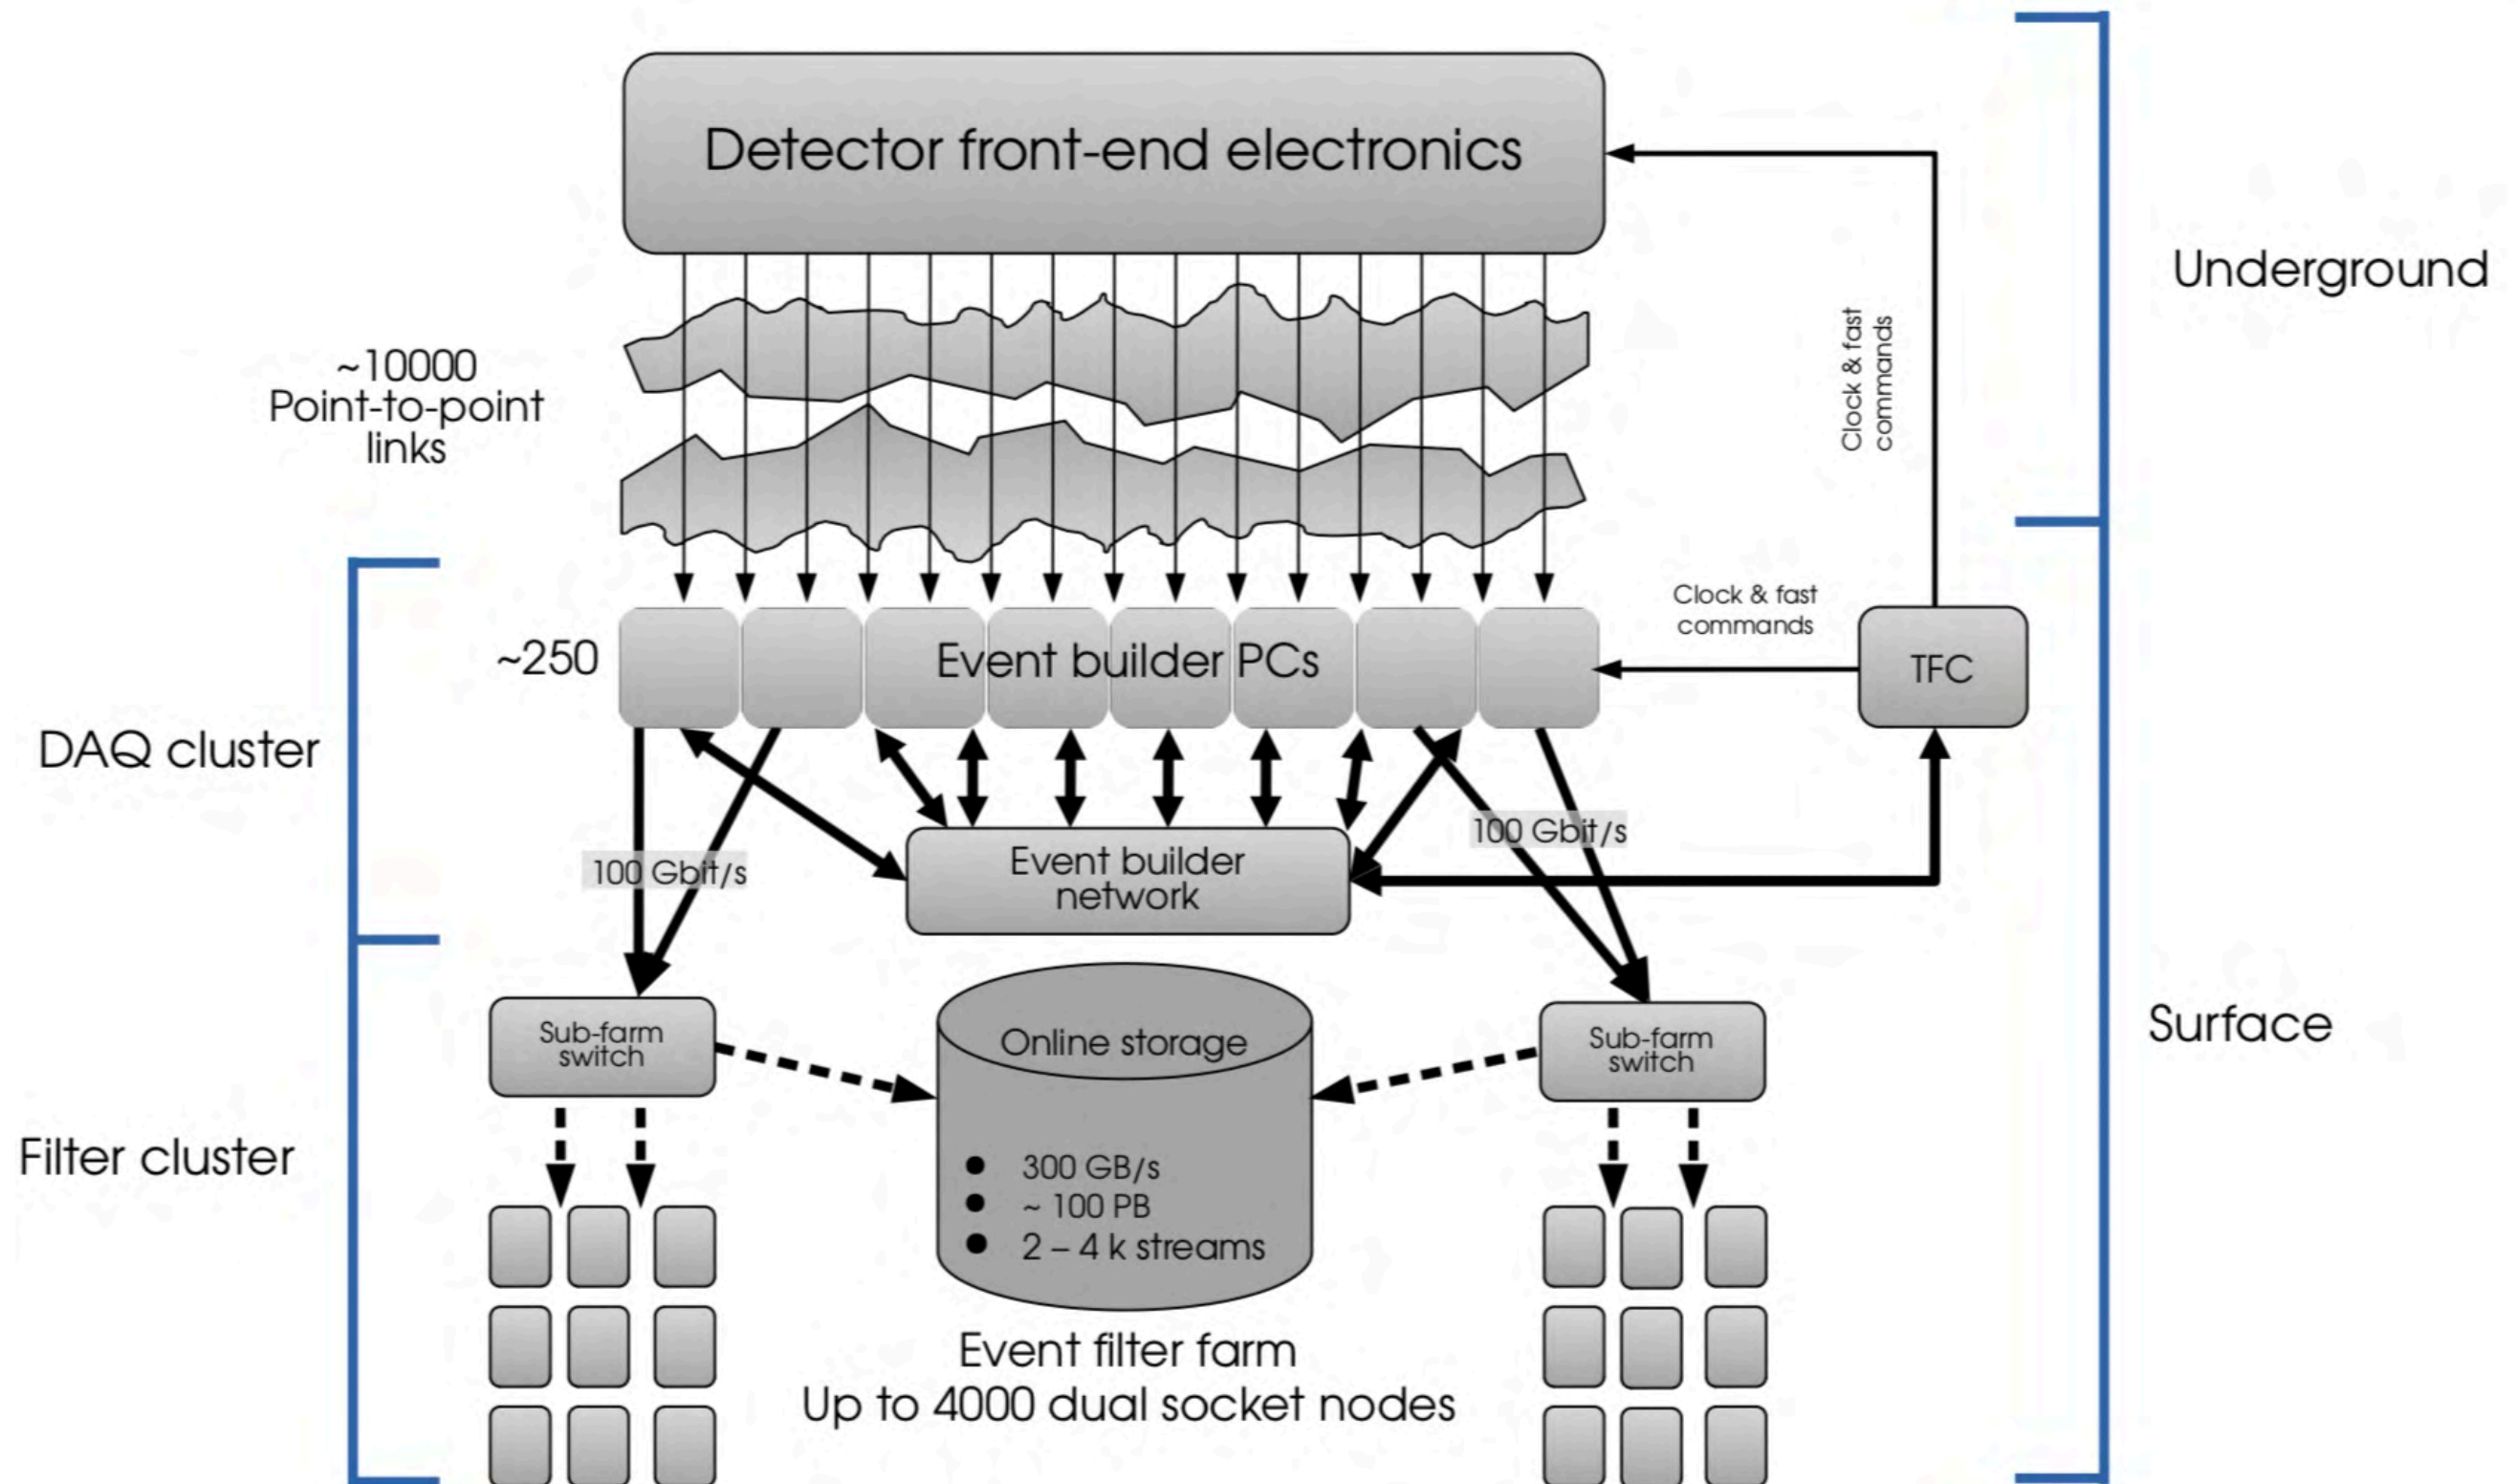
Run 3

- Peak lumi up to $\sim 3 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$
 - Peak $\langle \text{pile-up} \rangle \sim 80$
 - $\sim 100 \text{fb}^{-1}/\text{year} \rightarrow \sim 300 \text{fb}^{-1}$ by LS3
- ==> Upgrades mainly to cope with higher L1 trigger rates

HL-LHC

- Peak lumi up to $\sim 7.5 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$
 - Peak $\langle \text{pile-up} \rangle \sim 200$
 - $\sim 300 \text{fb}^{-1}/\text{year} \rightarrow \sim 3000 \text{fb}^{-1}$ in 10 years
- ==> Upgrades to cope with occupancies, radiation levels, trigger rates etc...

LHCb Run 3



Triggered vs. Streaming DAQ Systems


Triggered Systems

- Hardware/firmware based first level trigger
 - Example: at least two high pT muons
Digitisation happens after trigger signal is received
 - Analog pipelines in Frontend electronics to wait for trigger latency
 - Latency cannot be very large (think: microseconds)
 - Algorithms cannot be very sophisticated
- Hardware/Firmware/software based second level (or high level) trigger
 - Digitised raw data from Frontend electronics need to be stored in buffer memory
 - Latency and complexity of algorithms depends on available buffer space

Streaming Systems

- Data from all sub-systems is permanently digitised
- Since there is no trigger, typically sampling ADCs are used
 - Zero-suppression and feature extraction on the front end
- High-precision time distribution system required to correlate the data fragments from the sub-systems (event building)
- Data concentrators and event building network
- Event filtering by FPGA based systems or on PC/GPU server farms
 - Latency of algorithms determined by buffer space in the various processing stages

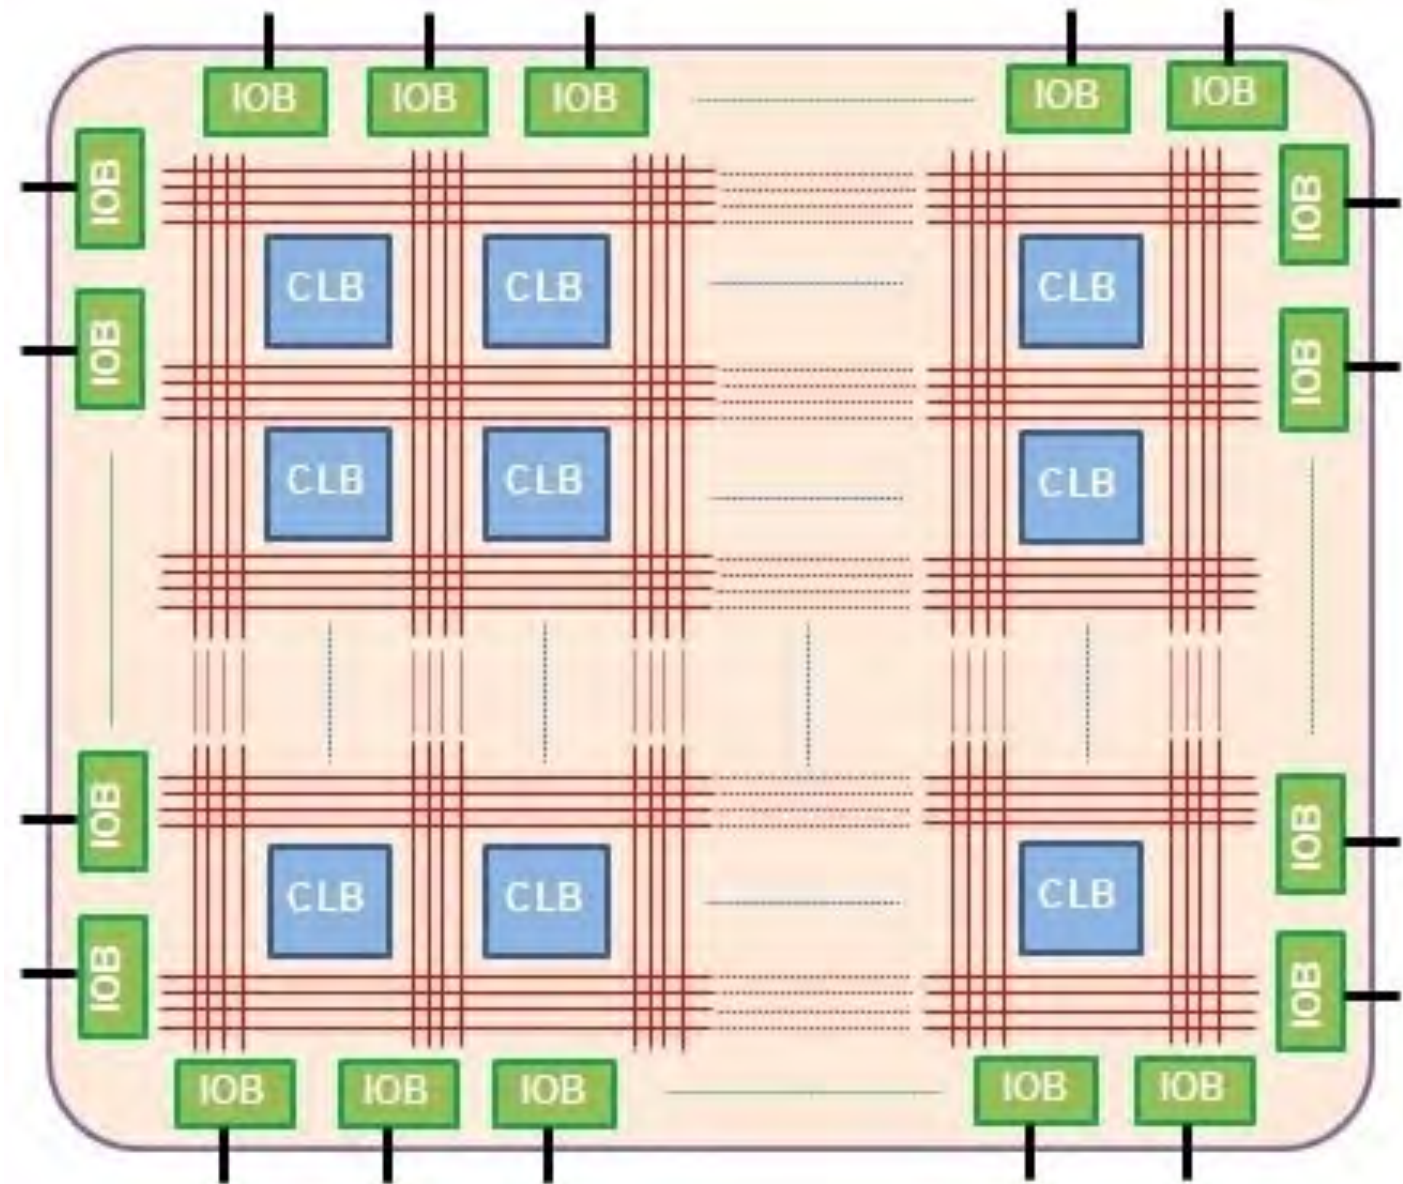
Example 1: CMS Global Muon Trigger




- Input: ~1000 bits @ 40 and 80 MHz
- Output: ~50 bits @ 80MHz
- Processing time: 250 ns
- Pipelined logic one new result every 25 ns
- 10 Xilinx Virtex-II FPGAs
- up to 500 user I/Os per chip
- Up to 25000 LUTs per chip used
- Up to 96 x 18kbit RAM used
- In use in the CMS trigger 2008-2015
- The CMS Global Muon trigger received 16 muon candidates from the three muon systems of CMS
- It merged different measurements for the same muon and found the best 4 over-all muon candidates

FPGA: Field-Programmable Gate Array

- Array of configurable logic blocks (CLB) with configurable interconnects
- Configurable input/output blocks (IOB)
- Configuration is defined using a hardware description language (VHDL or Verilog)
- A development tool interprets the code and creates a bit stream which is loaded into SRAM-cells in the FPGA, creating the designed configuration




Modern FPGAs



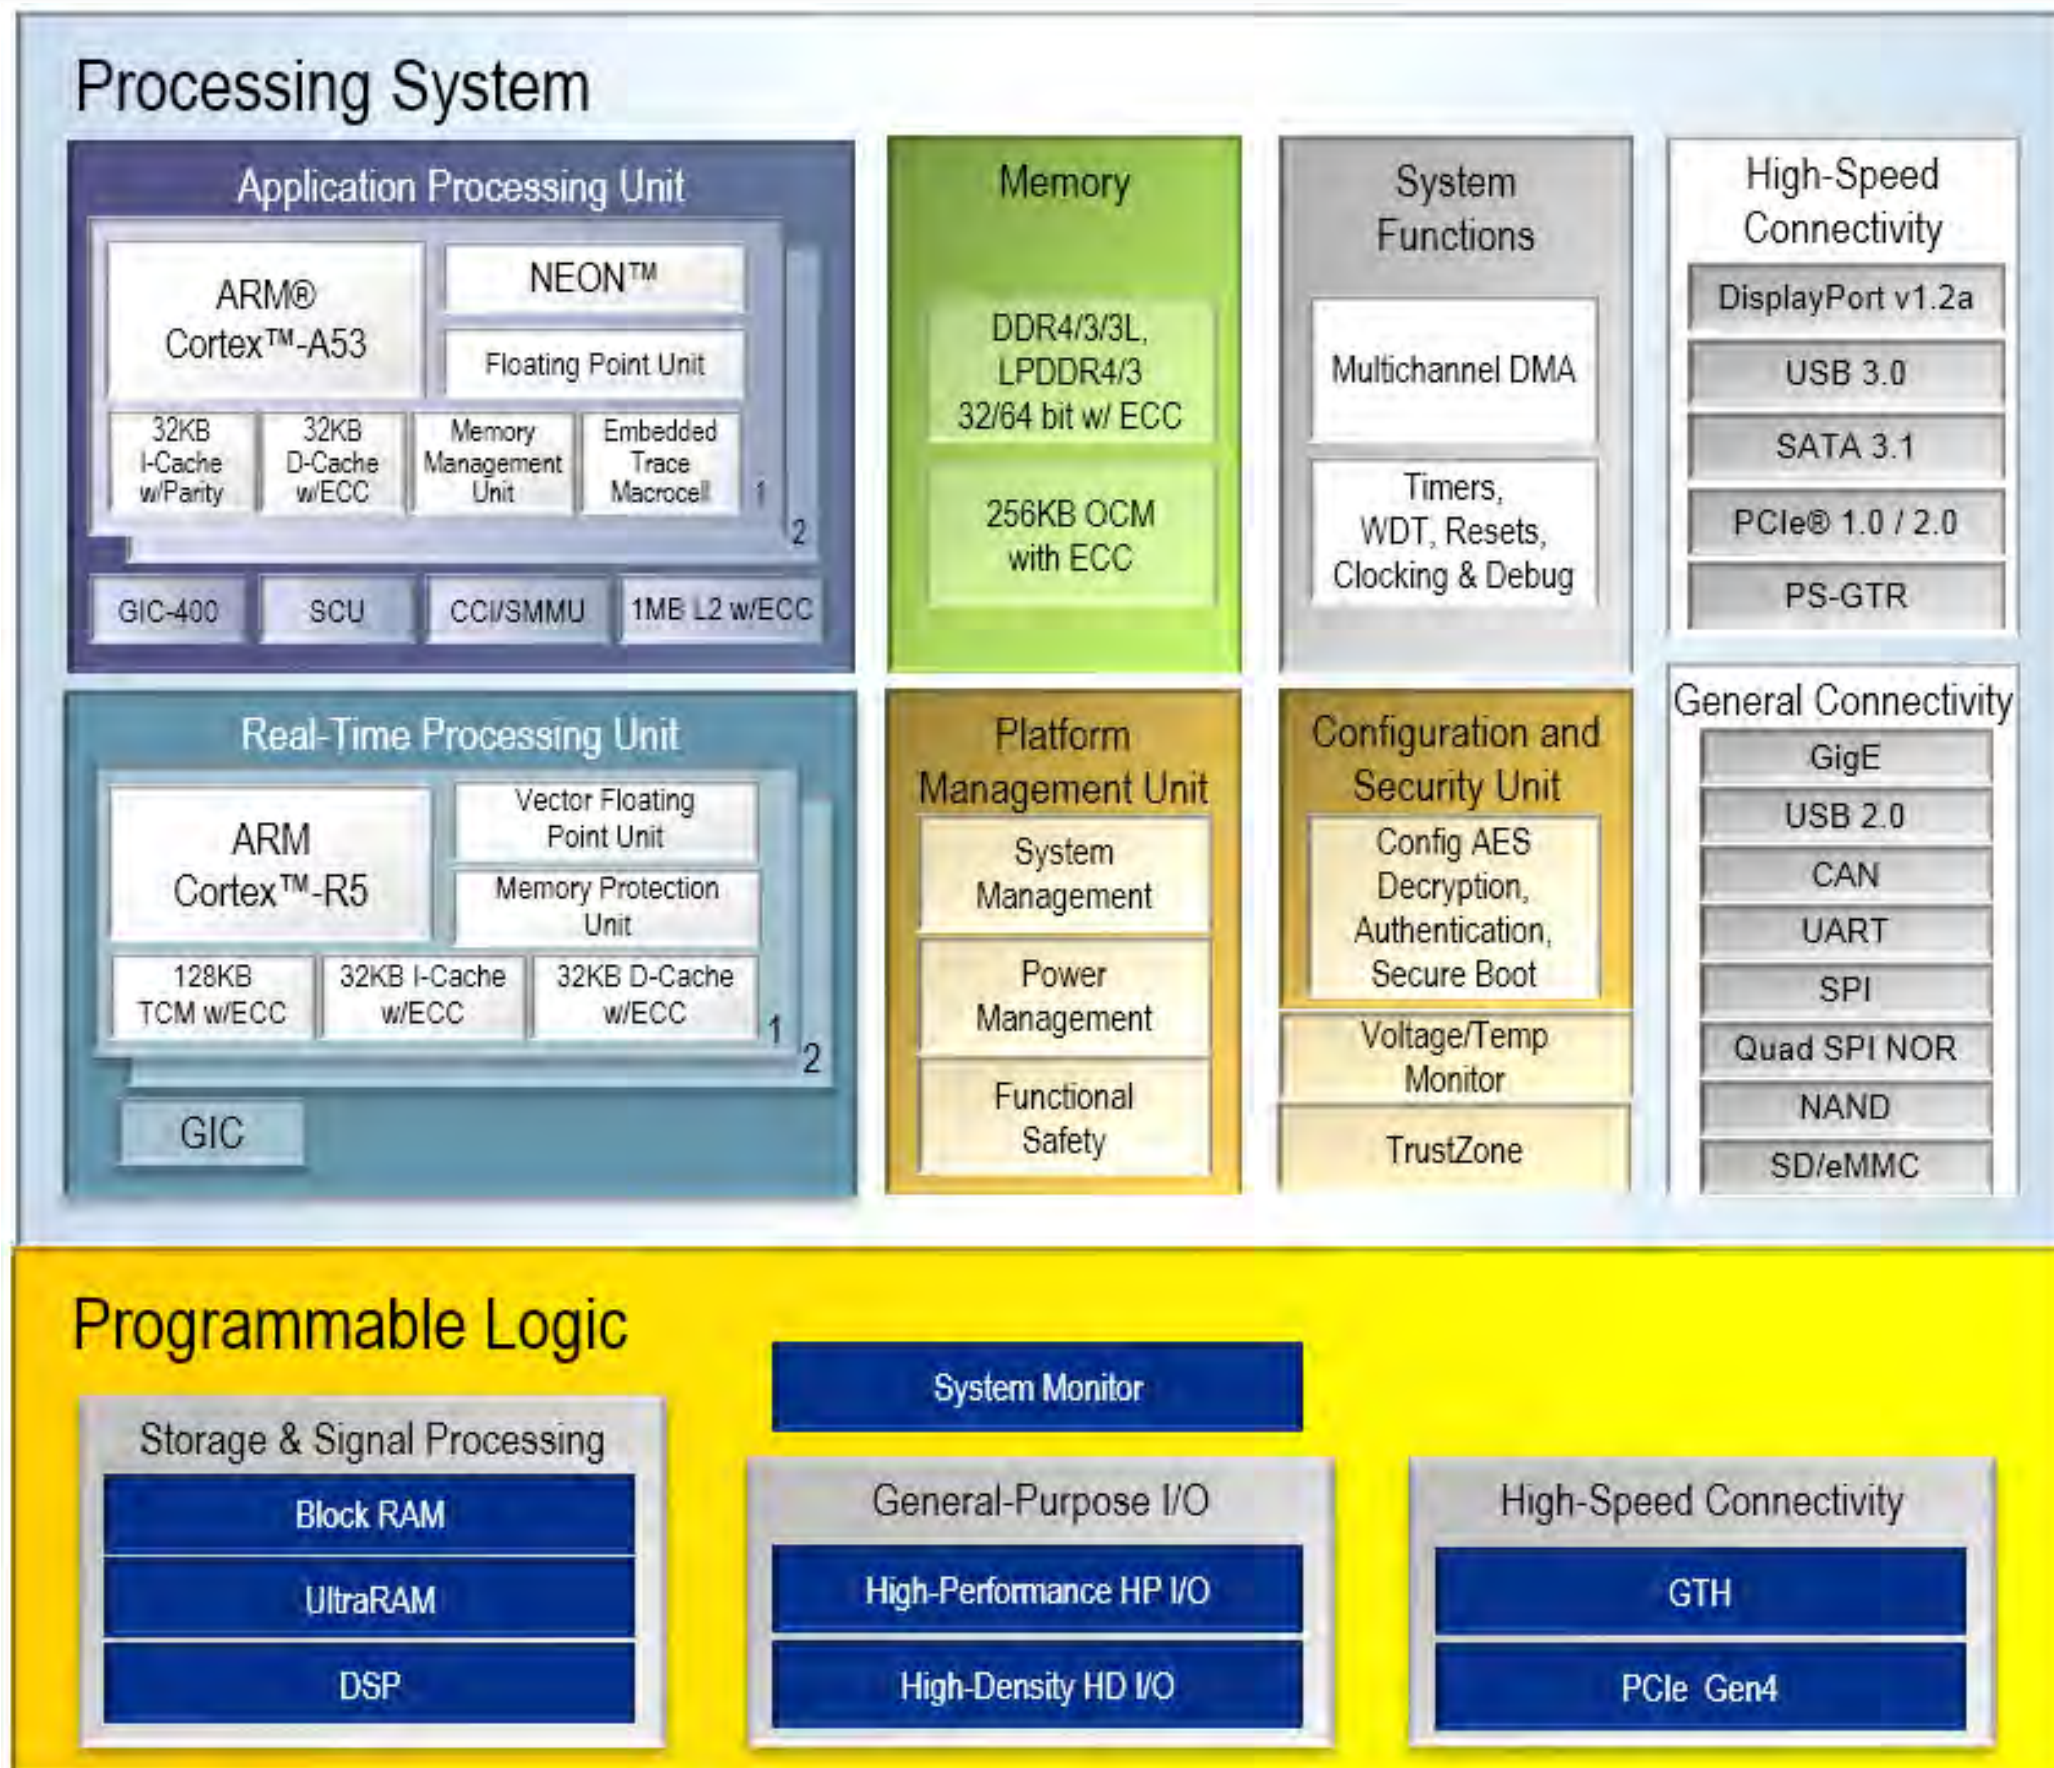
FPGAs vs. ASICs

- Application-specific circuits and FPGAs can perform similar functions
 - ASICs are not configurable. If you do not like their features, you have to design and manufacture a new one
 - FPGA designs can be changed very quickly and without cost for new hardware



FPGA vs ASIC COST ANALYSIS

XILINX Zynq Ultrascale+ Architecture




Case Study


The Belle II Experiment at KEK

Belle II Detector

TDR: arXiv:1011.0352




Belle II Vertex Detector (4 Layers SiStrip + 2 Layers PXD)



The DEPFET Ladder



SwitcherB

Row control




AMS/IBM HVCMOS 180 nm
Size $3.6 \times 1.5 \text{ mm}^2$
Gate and Clear signal
Fast HV ramp for Clear

PXD




DCDB (Drain Current Digitizer)
Analog frontend




UMC 180 nm
Size $5.0 \times 3.2 \text{ mm}^2$
TIA and ADC
Pedestal compensation

DHP (Data Handling Processor)
First data compression



TSMC 65 nm
Size $4.0 \times 3.2 \text{ mm}^2$
Stores raw data and pedestals
Common mode and pedestal correction
Data reduction (zero suppression)
Timing signal generation

Central Drift Chamber (CDC)



Silicon Vertex Dector


- **Silicon Vertex Detector (SVD)**

4 layers of DSSD

$r = 3.8 \text{ cm}, 8.0 \text{ cm}, 11.5 \text{ cm}, 14 \text{ cm}$

$L = 60 \text{ cm}$

$\sim 1 \text{ m}^2$



- **Pixel Detector (PXD)**


2 layers of DEPFET pixels

$r = 1.4 \text{ cm}, 2.2 \text{ cm}$





$L = 12 \text{ cm}$

$\sim 0.027 \text{ m}^2$


Charged Particle Track




Belle II SVD Module



DEPFET Pixel Sensor




- 768x250 DEPFET Pixels
- $50 \times 75 \mu\text{m}^2$ pixel pitch
- $75 \mu\text{m}$ thickness




Requires dedicated ASIC development

SwitcherB
Row control




AMS/IBM HVCMOS 180 nm
Size $3.6 \times 1.5 \text{ mm}^2$
Gate and Clear signal
Fast HV ramp for Clear
Rad. Hard proved (36 Mrad)


DCDB (Drain Current Digitizer)
Analog frontend



UMC 180 nm
Size $5.0 \times 3.2 \text{ mm}^2$
TIA and ADC
Pedestal compensation
Rad. Hard proved (20 Mrad)




DHP (Data Handling Processor)
First data compression




TSMC 65 nm
Size $4.0 \times 3.2 \text{ mm}^2$
Stores raw data and pedestals
Common mode and pedestal correction
Data reduction (zero suppression)
Timing signal generation
Rad. Hard proved (100 Mrad)

PXD



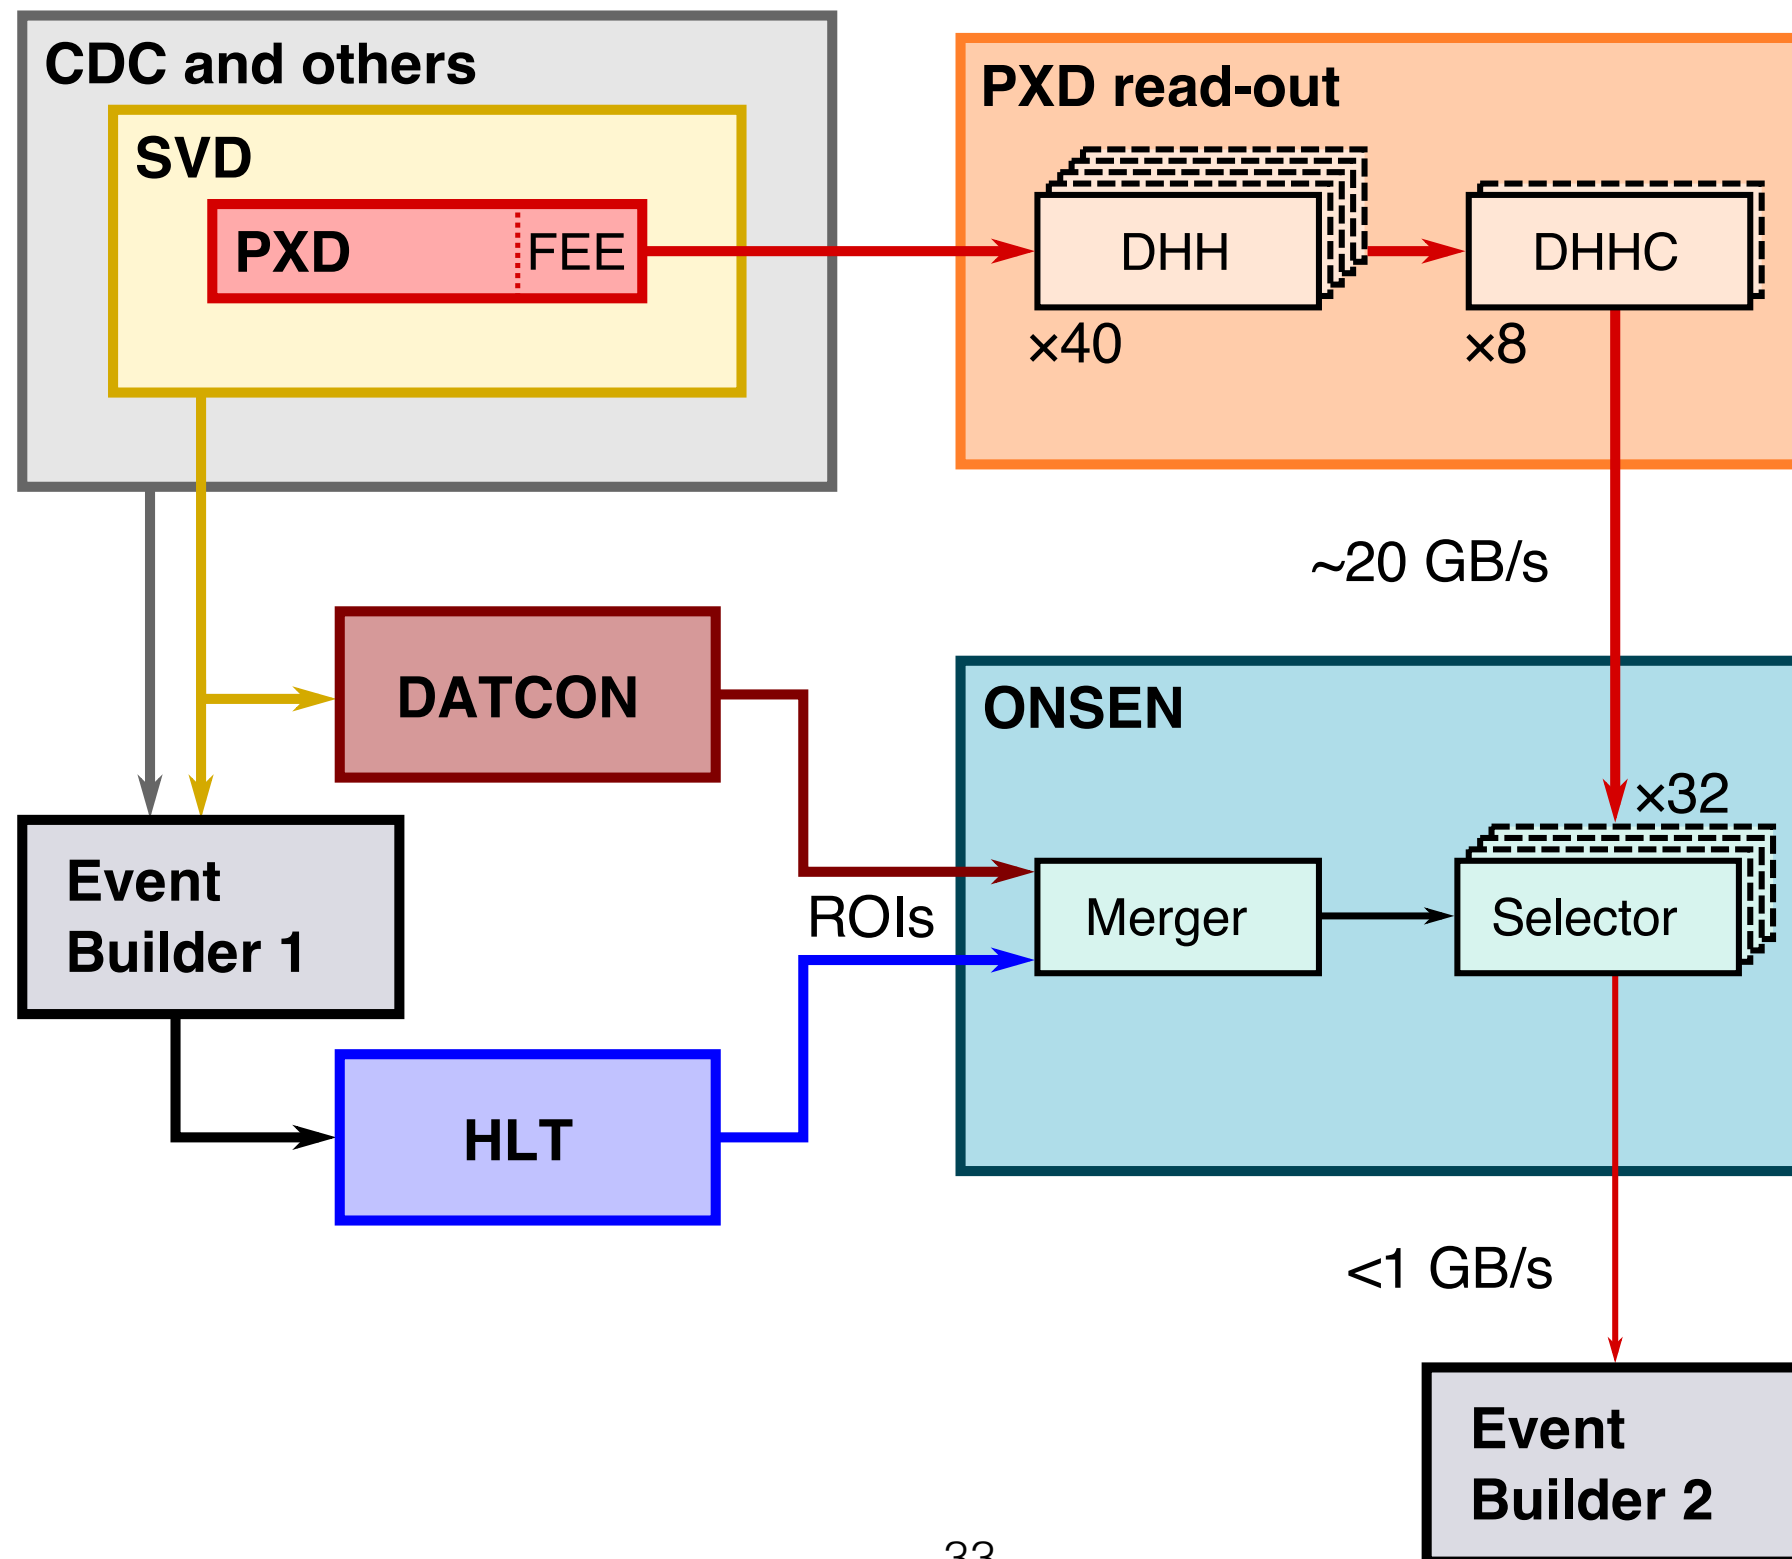
PXD - Detector - small, fragile and expensive (This is a full-size model)






- Amount of data created by PXD is larger than the data generated by all other subdetectors
- Only reduced PXD data is written to tape
- Use tracks in SVD (and CDC) to find PXD regions of interest


Belle II Data Acquisition System



FPGA based Embedded Systems (Trigger Lab, IHEP)




(a)

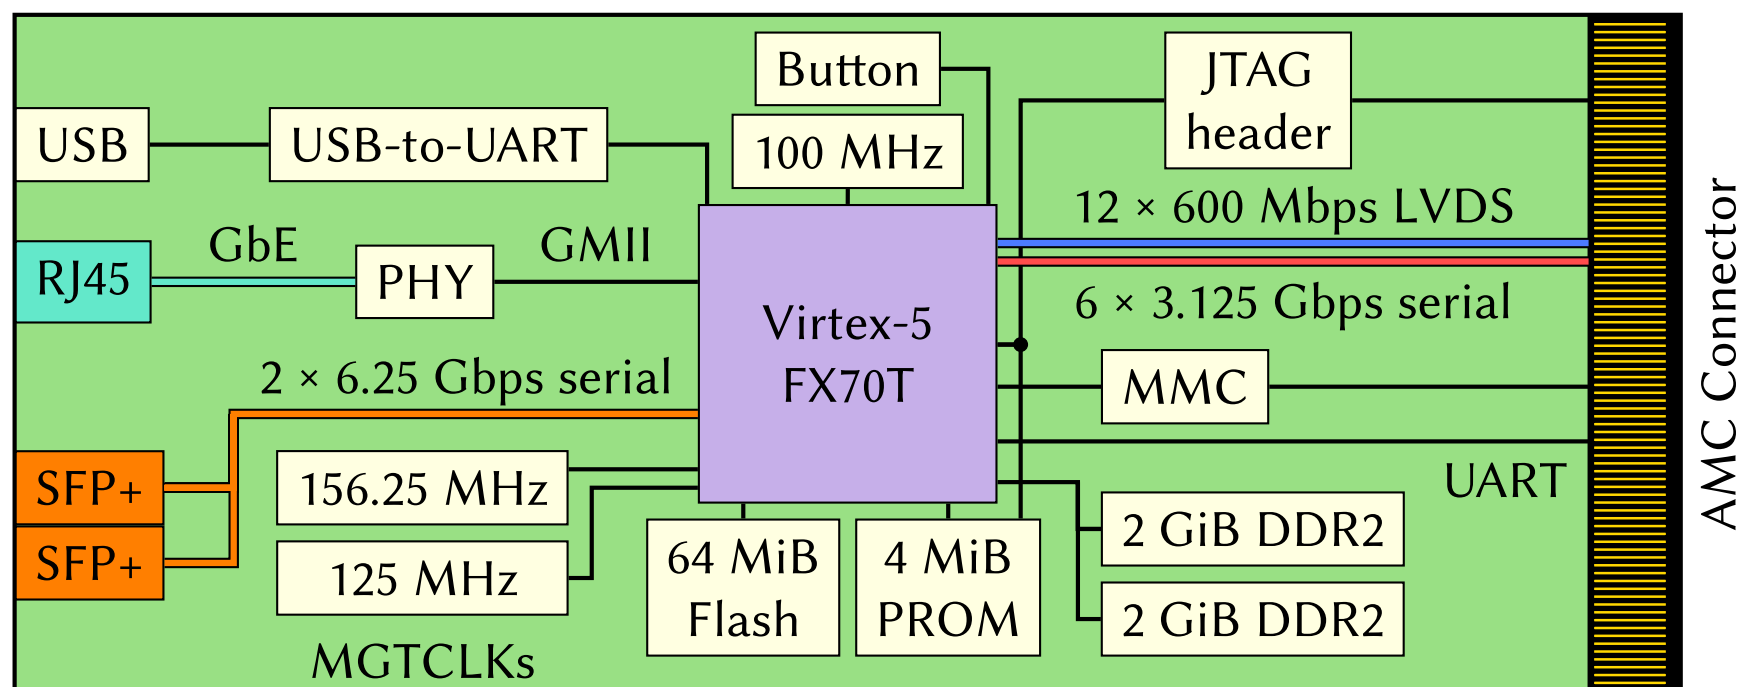


(b)³⁴

xFP Card (Xilinx Virtex 5)




(a)

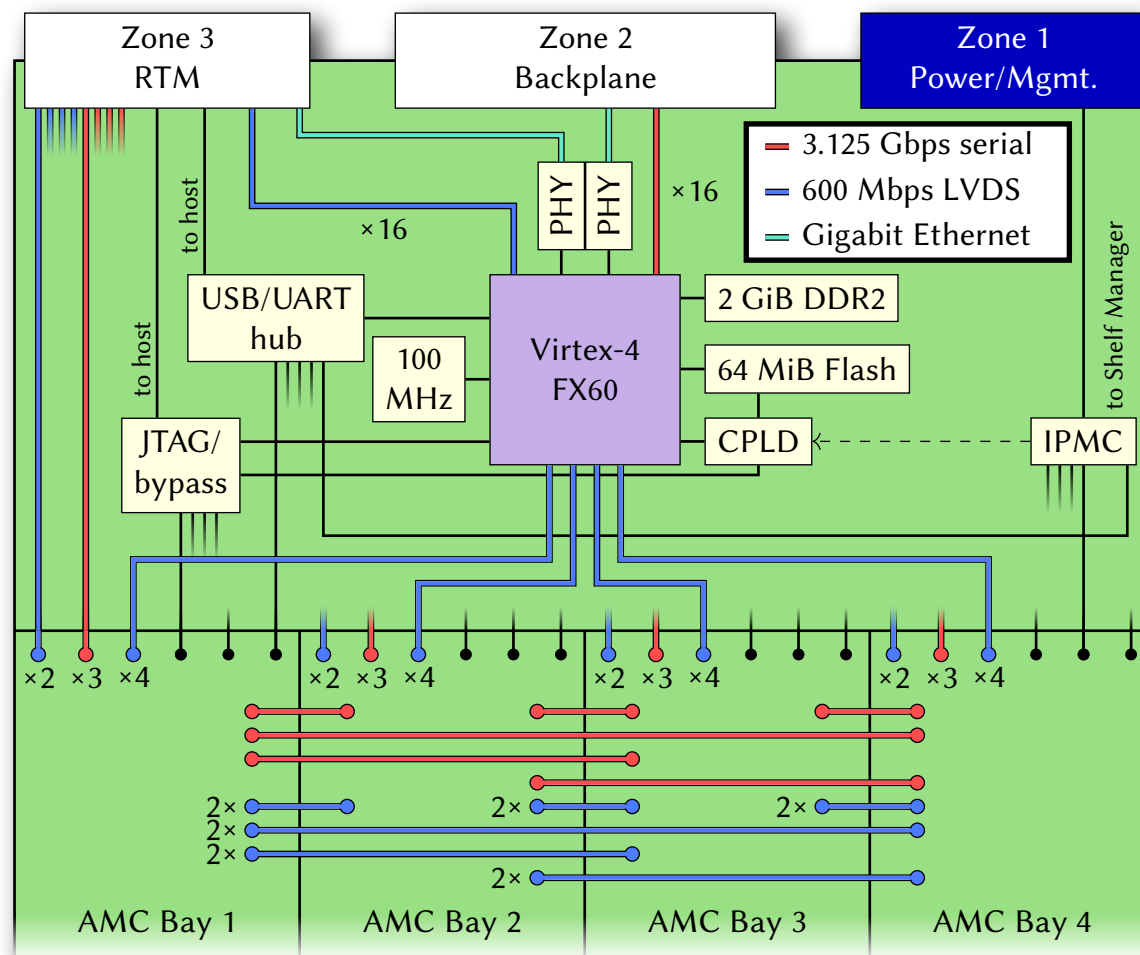


(b) 35

ATCA Carrier Board and ATCA Shelf



(a)




Outlook and Conclusion

Technological Challenges:


Example: Data Links

- Radiation hard links for data transport from the frontend electronics
- Power Consumption is a problem
 - 1 W per 10 Gb/s
 - For 10 Pb/s: 1 MW !!



Microprocessor Trends

Power wall and memory wall



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Laborte, O. Shacham, K. Olukotun, L. Hammond, and C. Baller
New plot and data collected for 2010-2015 by K. Rupp


Data Transport: High Speed Serial Links

- Xilinx Virtex Ultrascale+
 - Up to 128 serial links
 - Up to 8.38 Tb/s
 - What to do with so much data ?
 - Store in attached DDR memory for further processing ?

Table 3: Range of Support within UltraScale and UltraScale+ Families

	Type	Max Performance	Max Transceivers ⁽¹⁾	Peak Bandwidth ⁽²⁾
Virtex UltraScale+	GTY	32.75Gb/s	128	8,384Gb/s
Kintex UltraScale+	GTH/GTY	16.3 / 32.75Gb/s	44/32	3,530Gb/s
Zynq UltraScale+	PS-GTR/GTH/GTY	6.0 / 16.3 / 32.75Gb/s	4/44/28	3,316Gb/s
Virtex UltraScale	GTH/GTY	16.3 / 30.5Gb/s	60/60	5,616Gb/s
Kintex UltraScale	GTH	16.3Gb/s	64	2,086Gb/s

Memory Bandwidth Bottleneck



WP485_01_051217

Figure 1: Relative Memory Bandwidth Requirements

DRAM Memory - FPGA Hybrids

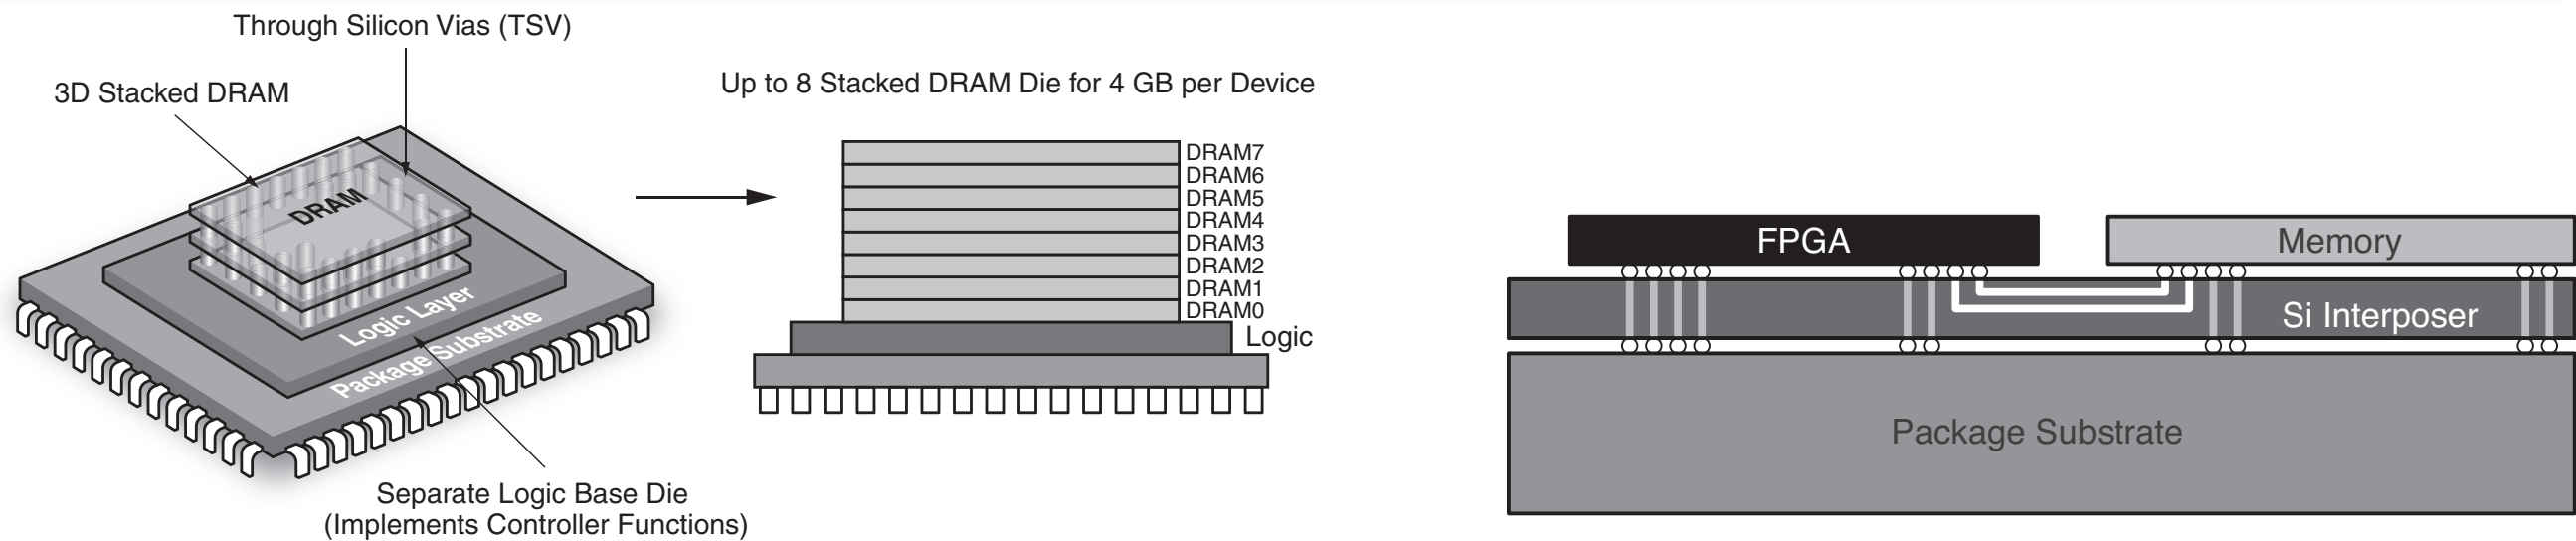


Figure 3: Illustration of the Structure of HMC

WP456_03_100114

WP485_02_051217

Table 1: Comparison of Key Features for Different Memory Solutions

	DDR4 DIMM	RLDRAM-3	HMC	HBM
Description	Standard commodity memory used in servers and PCs	Low latency DRAM for packet buffering applications	Hybrid memory cube serial DRAM	High bandwidth memory DRAM integrated into the FPGA package
Bandwidth	21.3GB/s	12.8GB/s	160GB/s	460GB/s
Typical Depth	16GB	2GB	4GB	16GB
Price / GB	\$	\$\$	\$\$\$	\$\$
PCB Req	High	High	Med	None
pJ / Bit	~27	~40	~30	~7
Latency	Medium	Low	High	Med

Some final thoughts ...

- Emerging technologies will make DAQ for next generation hadron collider experiments feasible
- Data storage and data processing resources might be the final bottleneck
 - Possible solutions, however there is a price to pay
 - Decide not to store (all) the detector (raw) data but instead perform high level of feature extraction online and store only the extracted data
 - Track parameters instead of hits in tracking detectors
 - Cluster energy instead of energies of individual modules
 - Ultimately: store only 4-vectors (not very realistic)
 - Decide to run more selective filtering algorithms to cut down raw data rate
 - Dangerous, when looking for New Physics