Trigger-less DAQ
How to transform a detector in a PC peripheral

Umberto Marconi
INFN Bologna
Catania, Novembre 2014

The LHCb experiment

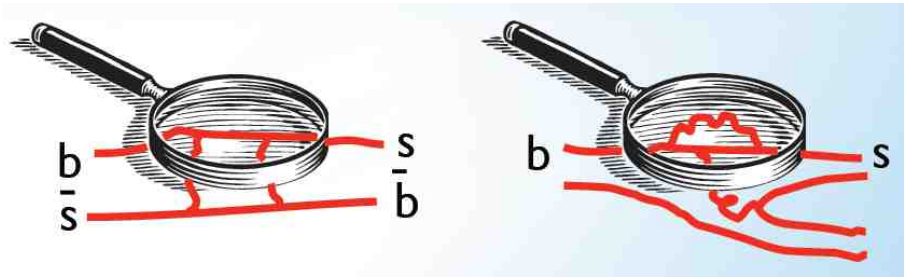

- LHCb is a **high-precision experiment** devoted to the search for New Physics beyond the Standard Model:
  - By studying CP violation and rare decays in the b and c-quark sectors.
  - Searching for deviations from the SM due to virtual contributions of new heavy particles in loop diagrams.
- Past and running experiments have shown that:
  - Flavour changing processes are consistent with the CKM mechanism.
  - **Large sources of flavour symmetry breaking are excluded at the TeV scale**.
  - The flavour structure of the NP, if it exists, would be very peculiar at the TeV scale.

INFN
Istituto Nazionale
di Fisica Nucleare

The LHCb Detector

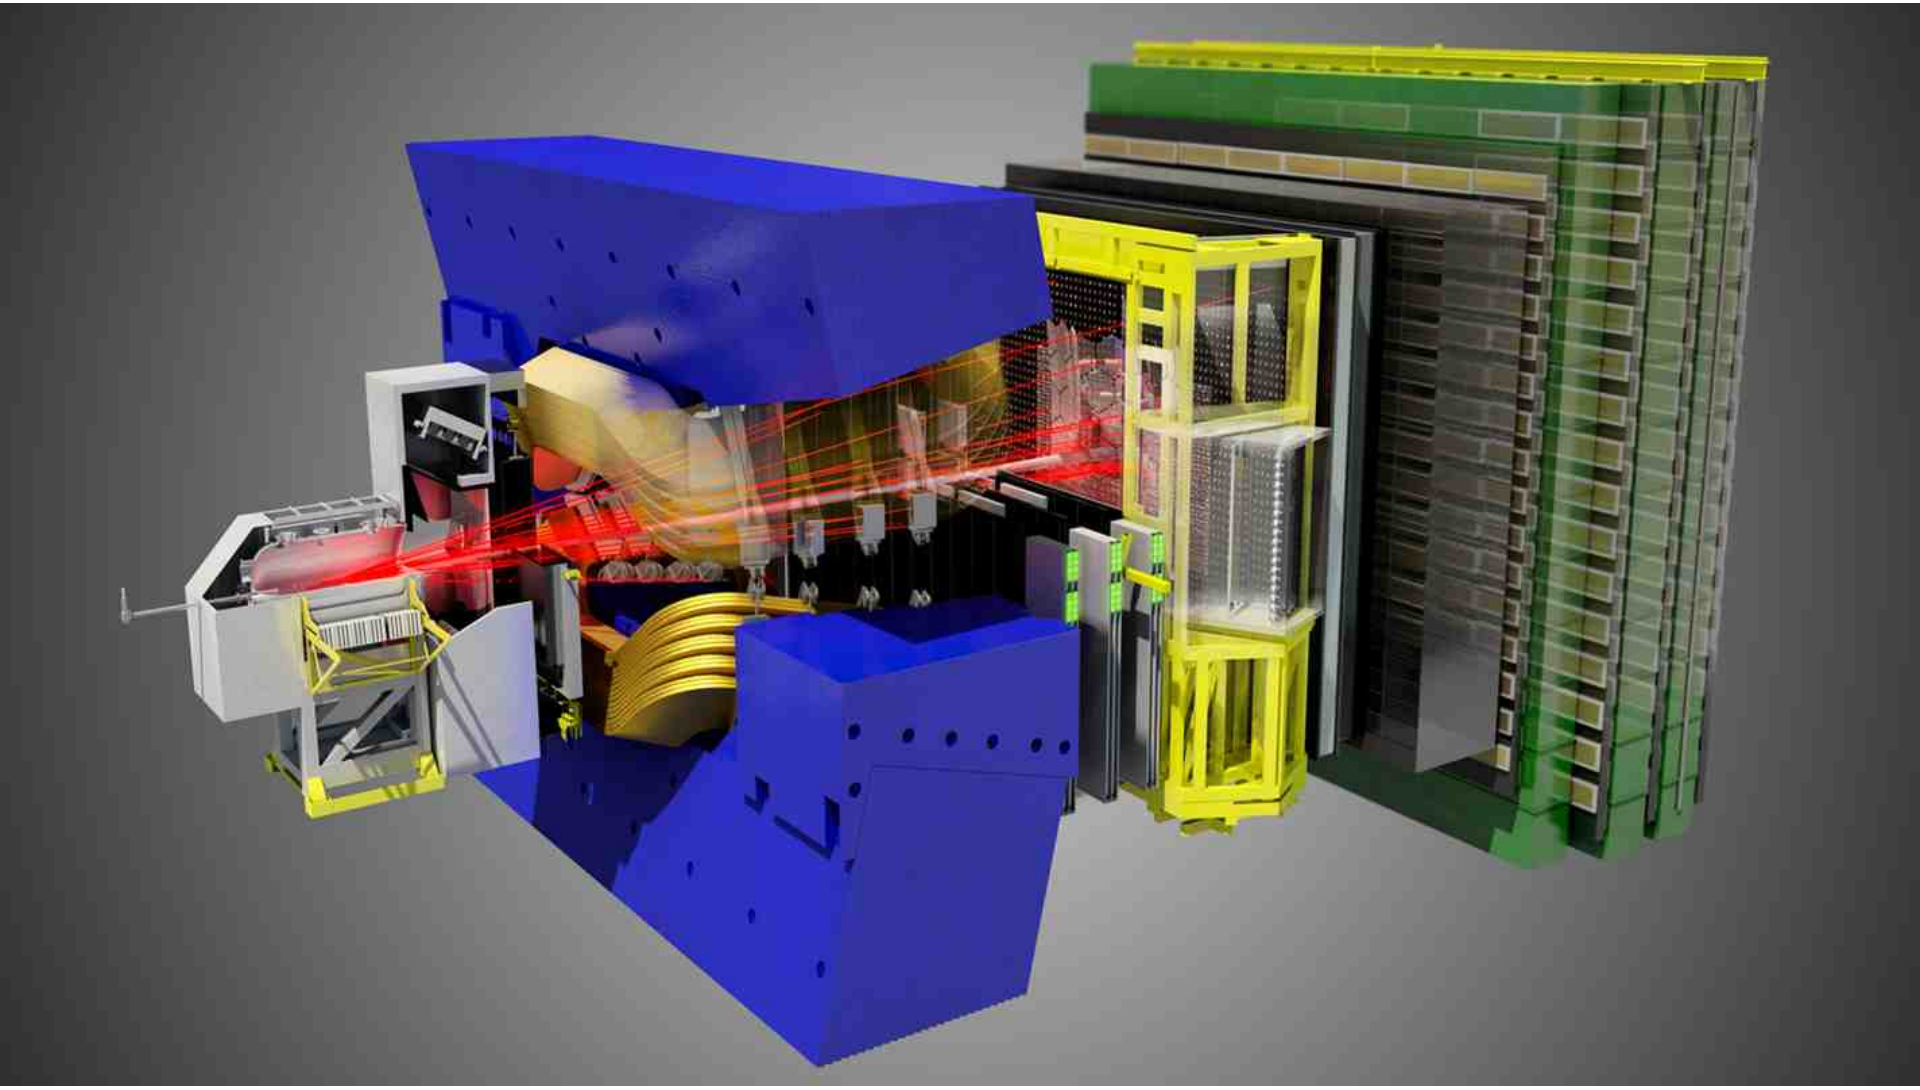

The LHCb detector

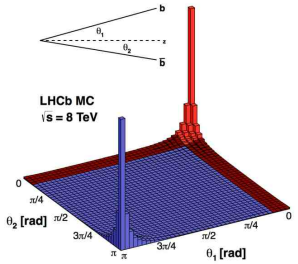

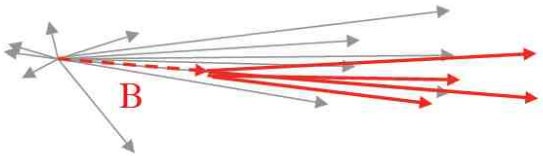

- Decay time resolution **Δ**t: 30-50 fs
- **Δ**p/p = 0.4-0.6%
- **Δ**m = 10-20 MeV/c^2
- Muon ID: **ε**(μ/μ) = 95%, mis-ID **ε**(**π**/μ) ~ 1%
- RICH ID **π**/K: **ε**(K/K) = 95%, mis-ID **ε**(**K**/**π**) ~ 5%
- CALO ID: **ε**(e/e)=90%, mis-ID ~ 5%

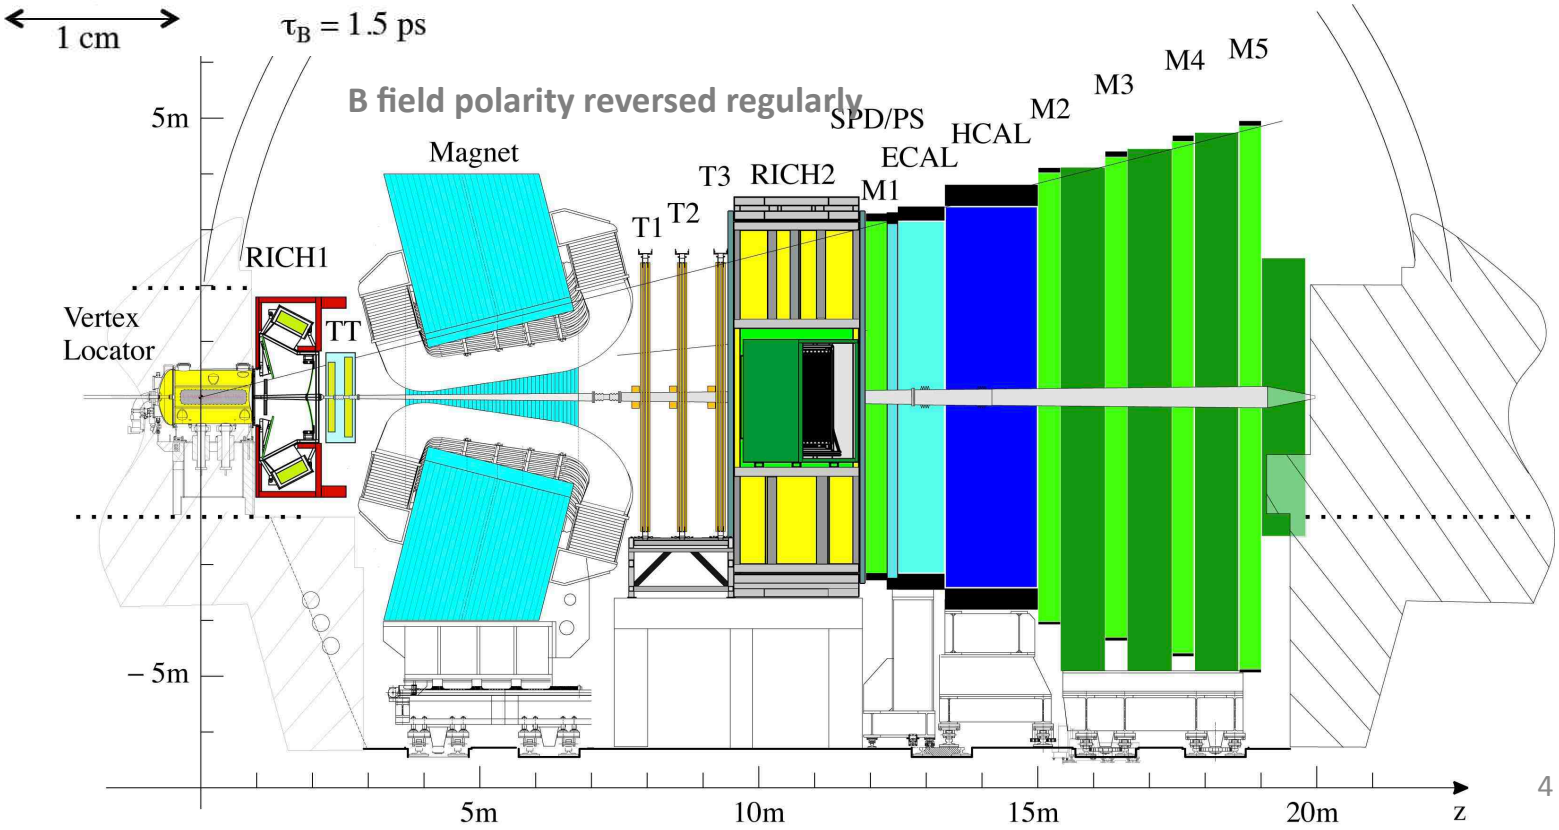

"The LHCb Detector at CERN"; The LHCb Collaboration, JINST 3, S08005 (2008)

LHCb events

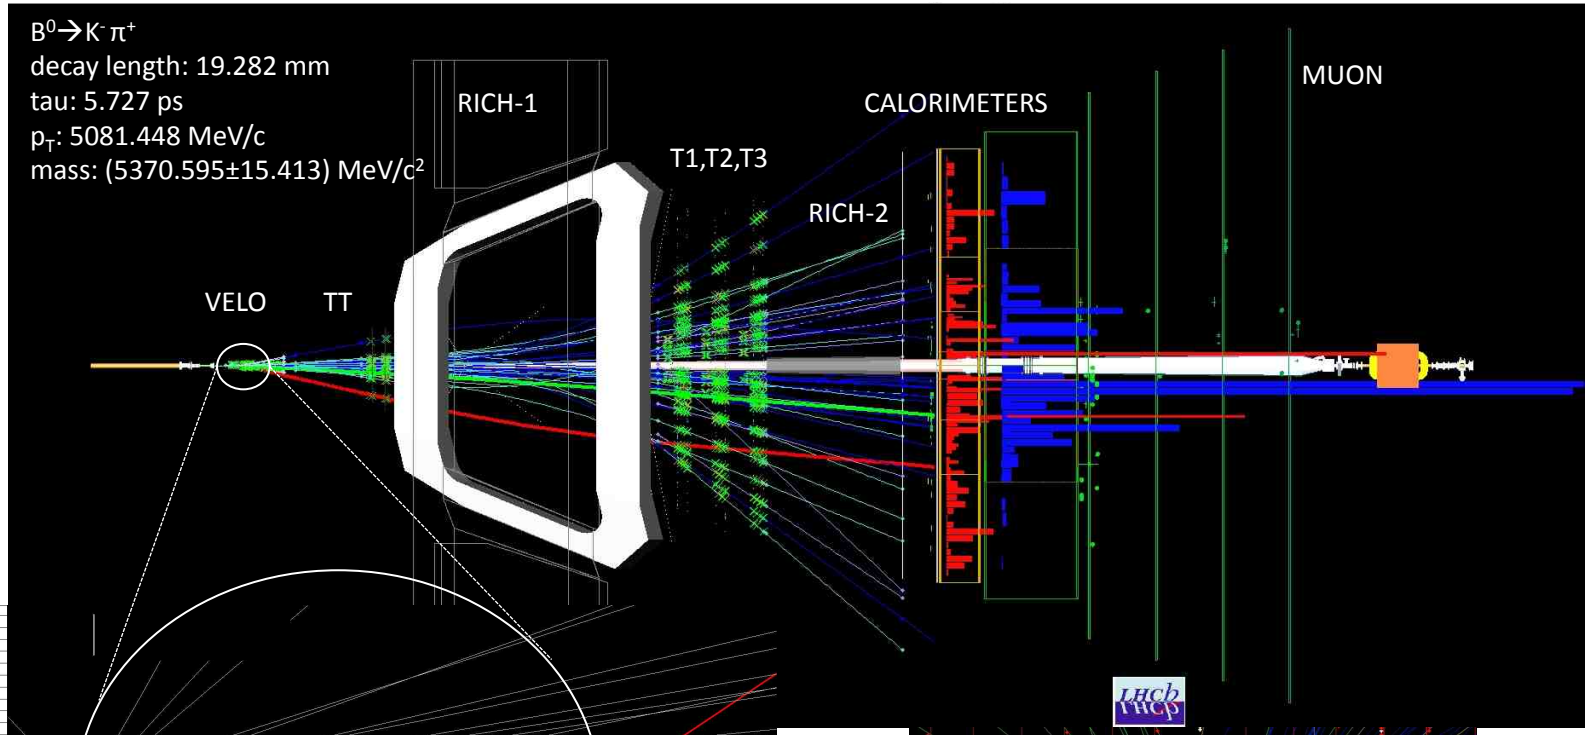
$B^0 \rightarrow K^- \pi^+$
decay length: 19.282 mm
tau: 5.727 ps
p_T: 5081.448 MeV/c
mass: (5370.595±15.413) MeV/c^2
RICH-1
T1,T2,T3
RICH-2
CALORIMETERS
MUON
VELO
TT

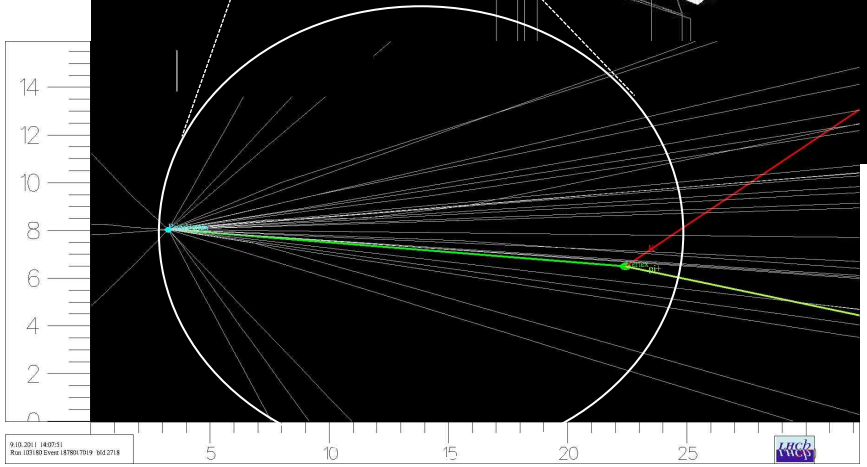
14
12
10
8
6
4
2
0
5
10
15
20
25

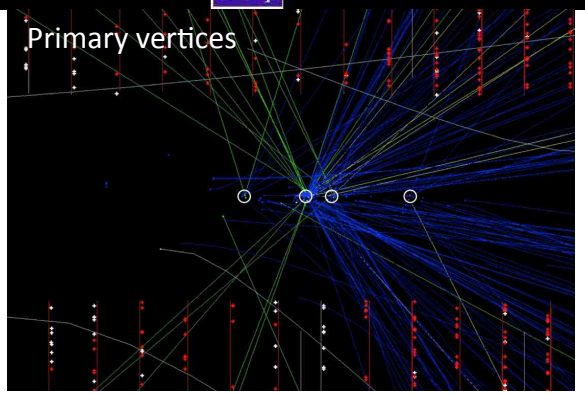
Primary vertices

Instantaneous Luminosity

- Instantaneous luminosity leveling at **4×10^{32} cm^{-2} s^{-1}**, ±3% around the target value

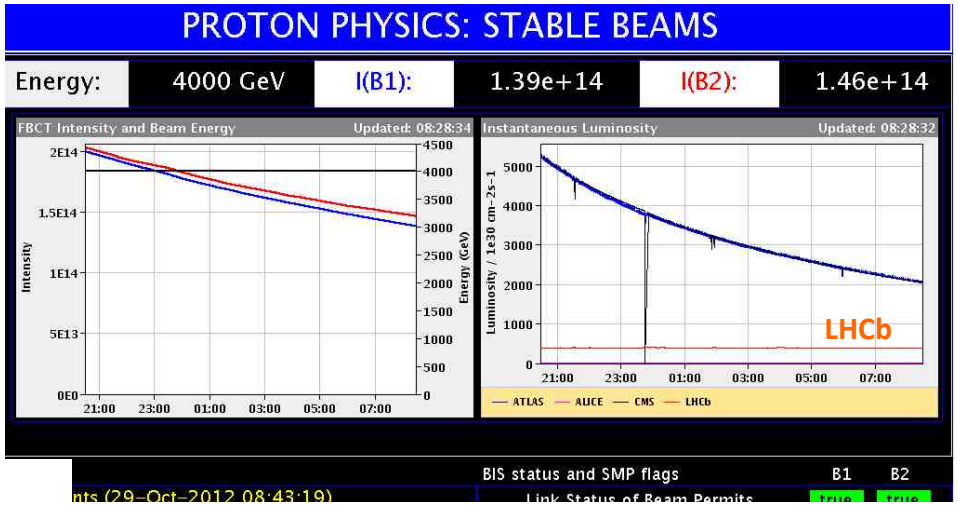

- LHCb was designed to operate with a single collision per bunch crossing, running at a instantaneous luminosity of **2×10^{32} cm^{-2} s^{-1}** (assuming about **2700** circulating bunches).
  - At the time of design there were worries about possible ambiguities in assigning the B decay vertex to the proper primary vertex among many.
- Soon LHCb realized that running at higher multiplicities would have been possible. In 2012 we run at **4×10^{32} cm^{-2} s^{-1}** with only **1262** colliding bunches.
  - 50 ns separation between bunches while the nominal 25 ns (will available by 2015).
  - **4 times more collisions per crossing than planned in the design**.
  - The average number of visible collisions per crossing in 2012 raised up to **$\mu > 2.5$**
    (μ: average n. of visible interactions)

Heavy flavour signatures

Beauty hadrons

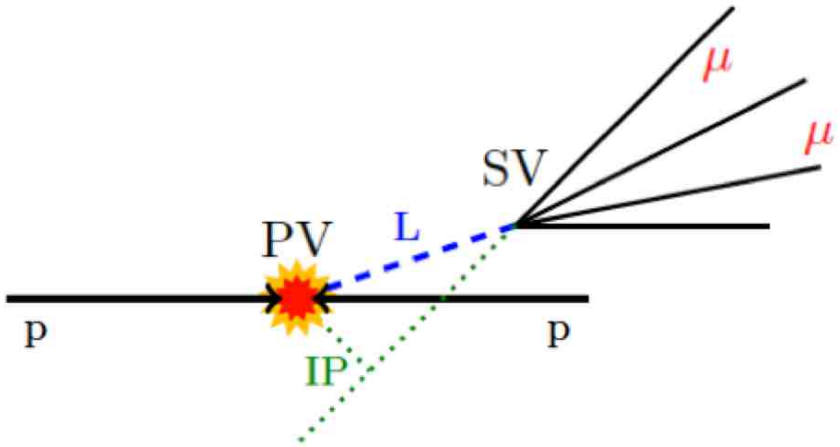

- B^+ mass 5.28 GeV, daughter p_T $\mathcal{O}$(1 GeV)
- lifetime ≈ 1.6 ps ⇒ flight distance ≈ 1 cm
- common signature: detached $\mu\mu$

Charmed hadrons

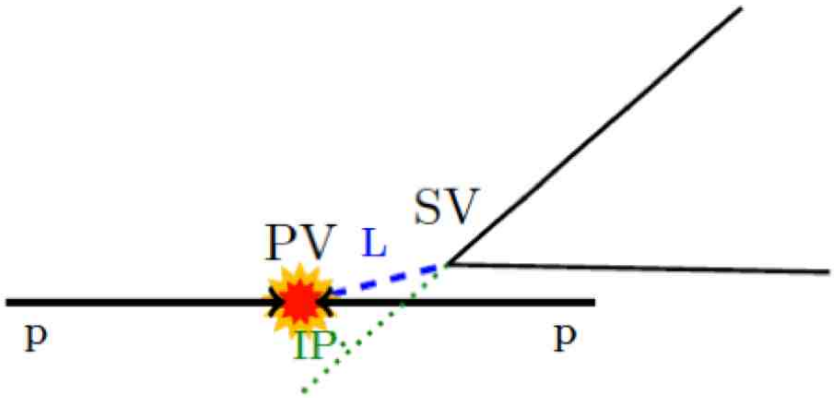

- D^0 mass 1.86 GeV, sizeable daughter p_T
- lifetime ≈ 0.4 ps ⇒ flight distance ≈ 4 mm
- can be produced in B decays

The present LHCb Trigger

Muon Based L0 Triggers

Reconstructs muon track segments

- $\Delta p/p \approx 20\%$

Two L0 muon triggers:

- single muon, $p_T > 1.76\,\text{GeV}$
- dimuon, $p_{T1} \times p_{T2} > (1.6\,\text{GeV})^2$
- total rate $\approx 400\,\text{kHz}$

Typically over 90% efficiency (wrt offline selected $B \to J/\psi K$)

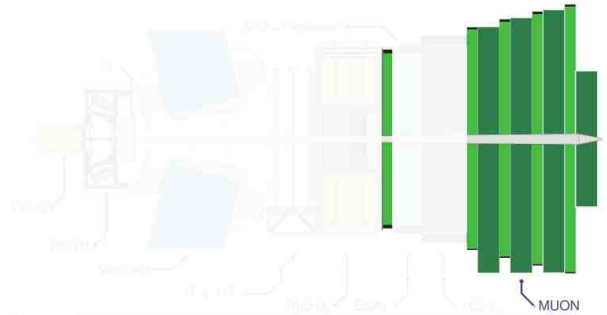

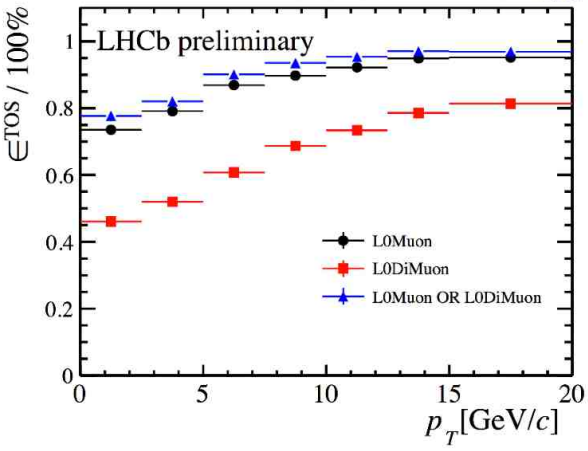

The present LHCb Trigger (2)

Calorimeter Based L0 Triggers

Select high E_T hadrons, electrons and photons

- Preshower and SPD discriminate between electrons and photons

L0 hadron trigger:

- threshold 3.6 GeV
- rate ≈ 490 kHz

L0 electron, photon:

- threshold 3 GeV
- rate ≈ 150 kHz
- ≈ 80% efficient for $B \to X\gamma$ (wrt offline selection)

Total L0 rate ≈ 1 MHz

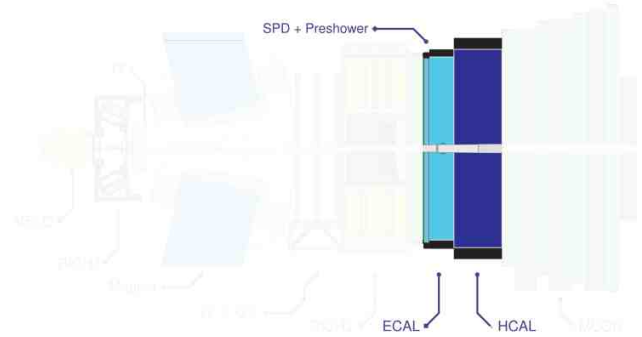

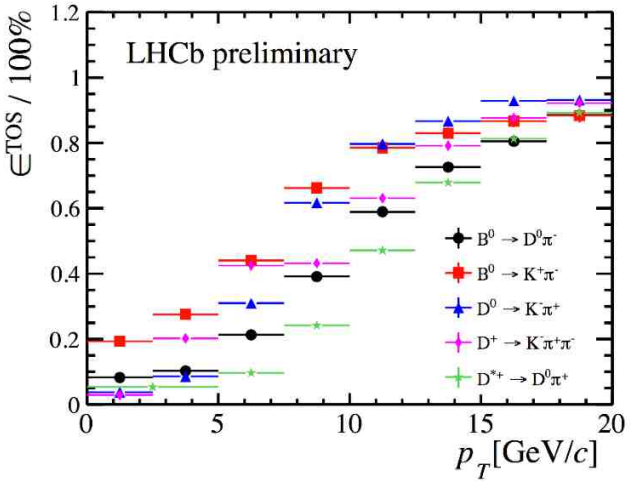

The present LHCb Trigger

- The **Level-0 trigger** operates at 40 MHz, with a **maximum output rate limited to 1.1 MHz**.
  - Custom electronics, fully pipelined, constant latency of about **4 µs**.
  - **Bandwidth to the HLT** ~ 1.5 Tb/s, through GOL serializers ~1 Gb/s and optical links.
  - **25%** of the events are **deferred**: temporarily stored c disk and processed with the HLT farm during the inter-fills.
- The **HLT** is a software trigger.
  - Reconstruct VELO tracks and primary vertices
  - Select events with at least one track matching p, p_T , impact parameter and track quality cuts.
  - At around **50 kHz** performs inclusive or exclusive selections of the events.
  - Full track reconstruction, without particle-identification.
  - Total accept rate to disk for offline analysis is **5 kHz**.

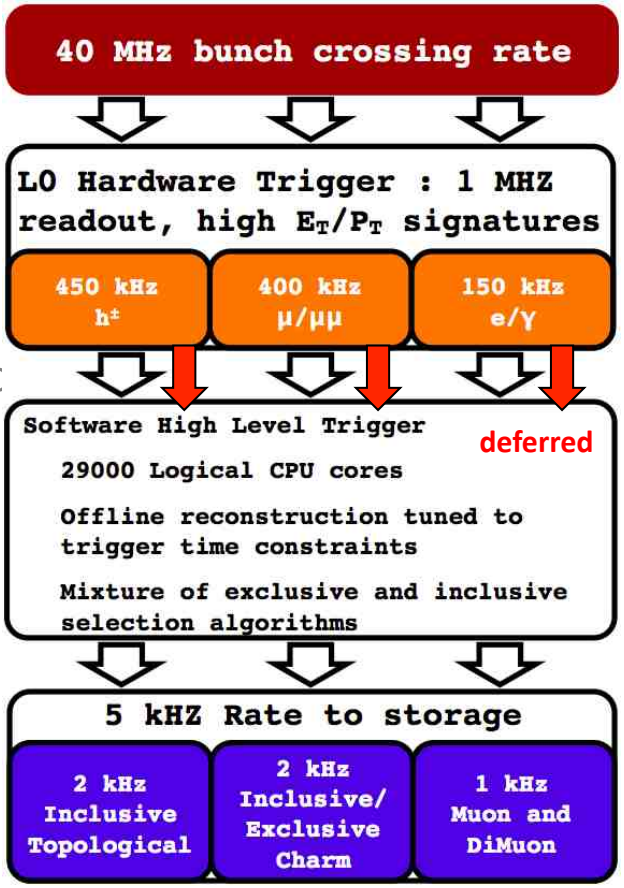

LHCb DAQ today

Switched packed network, push protocol,
with centralized flow-control

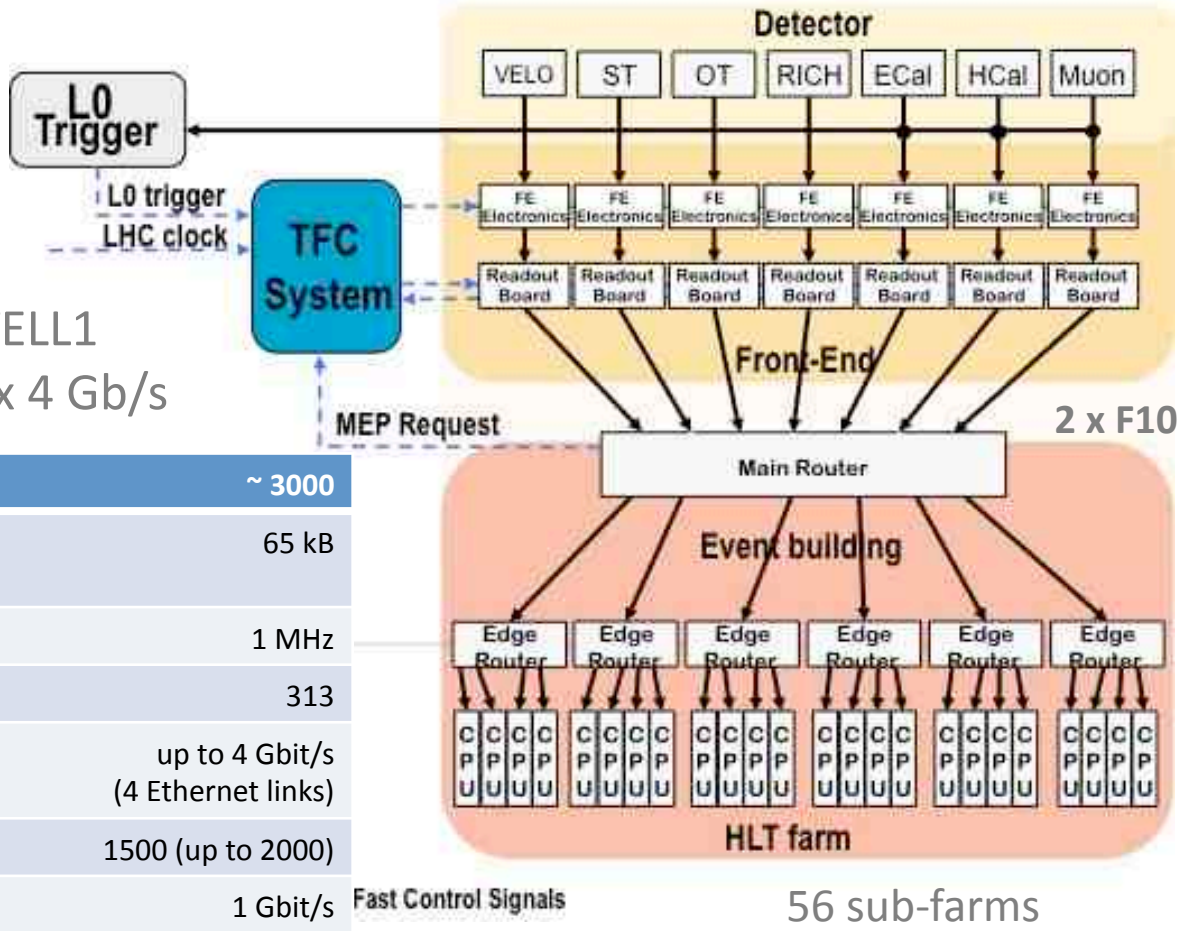

Readout boards: TELL1
Throughput: 313 x 4 Gb/s

# links (UTP Cat 6)	~ 3000
Event-size (total – zero-suppressed)	65 kB
Read-out rate	1 MHz
# read-out boards	313
output bw / read-out board	up to 4 Gbit/s (4 Ethernet links)
# farm-nodes	1500 (up to 2000)
max. input bw / farm-node	1 Gbit/s
# core-routers	2
# edge routers	56

JINST 8 (2013) P04022

Deferred trigger

- 2 TB on 1000 HLT servers

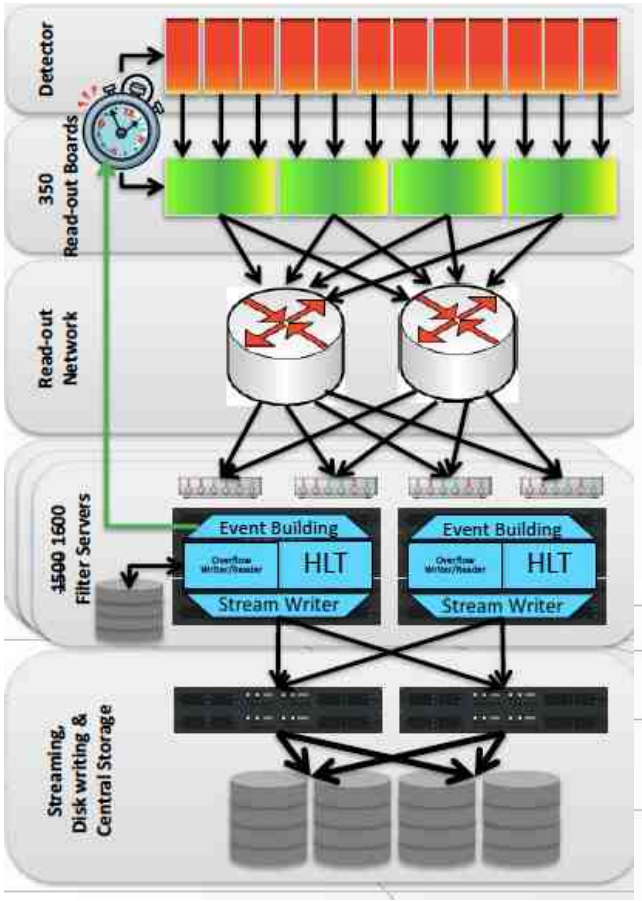

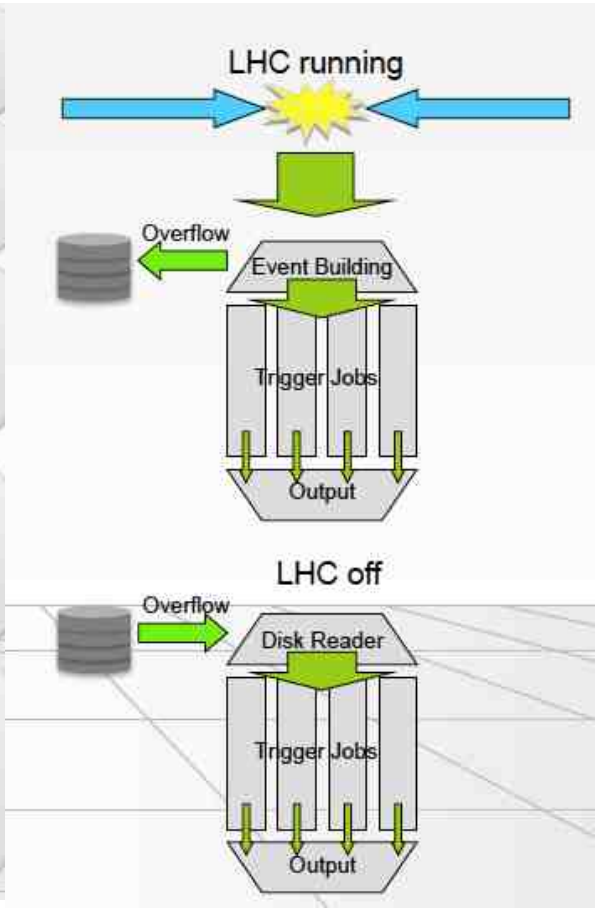

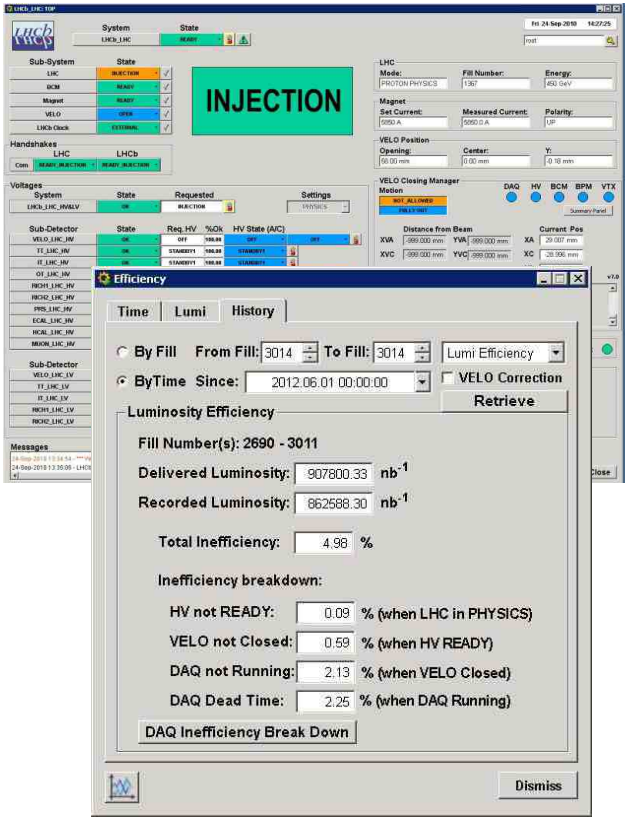

Push based event building

- Run at > 90% link load.
- Multi Event Packets: aggregate O(10) event fragments and just drop them onto the wire.
- **IP Datagrams**: no resend, no explicit flow control.
- Rely on network hardware with large buffers to handle traffic flow pattern.
- Link aggregation load balancing is not 100% fair nor standardized.
- 1 Gb/s ≠ 1 Gb/s: there is a small ± which can cause trouble at very high link loads.
- Commercial network stuff most of the time does not work off the shelf.

The core router

The readout-board TELL1

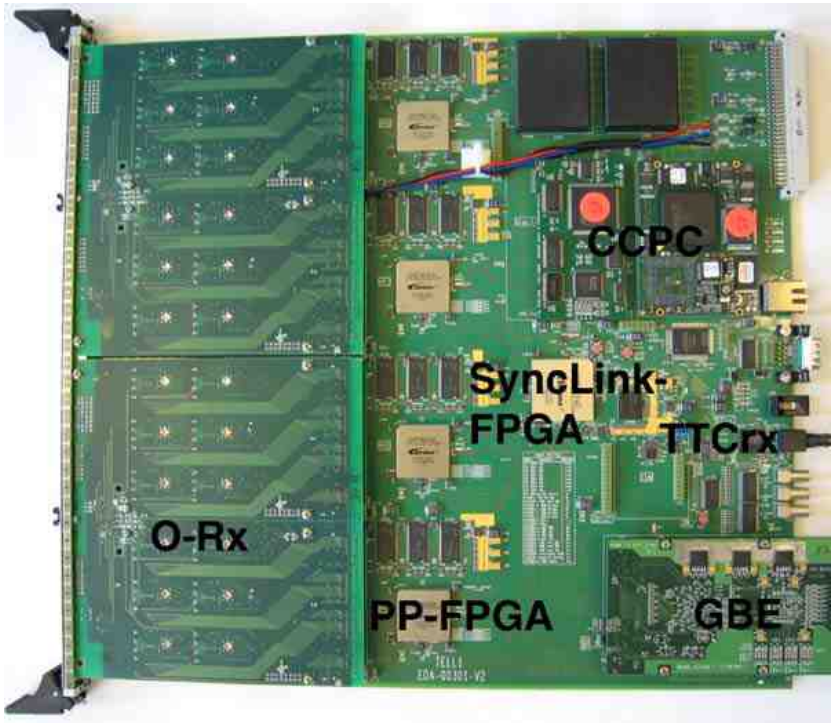

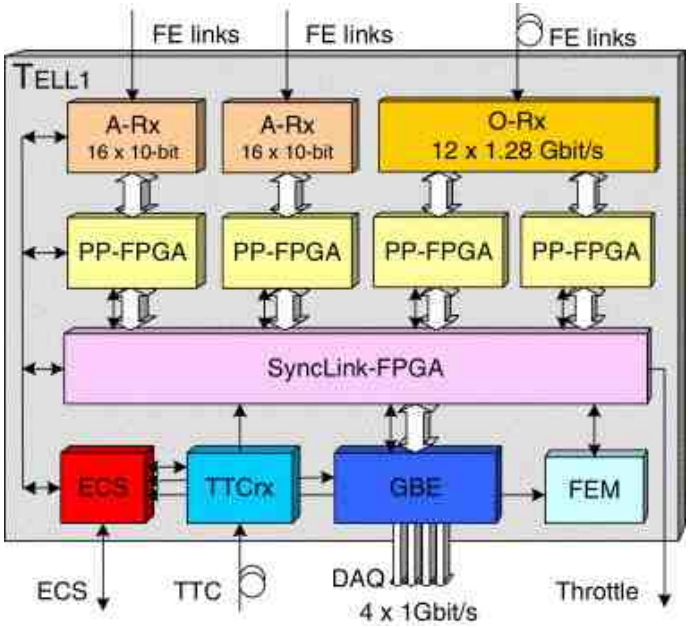

TTC in
Throttle out

4 x 1 GbE (UTP)

Central FPGA
pushing data via UDP

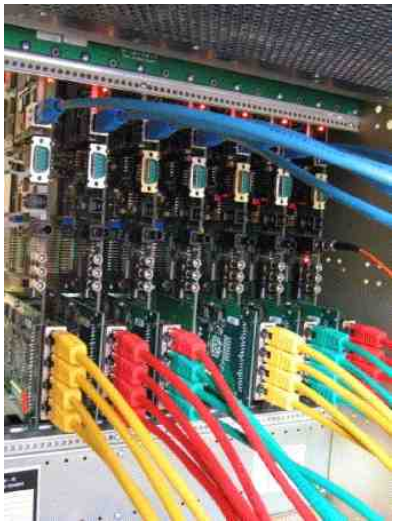

The 1MHz L0 rate limitation

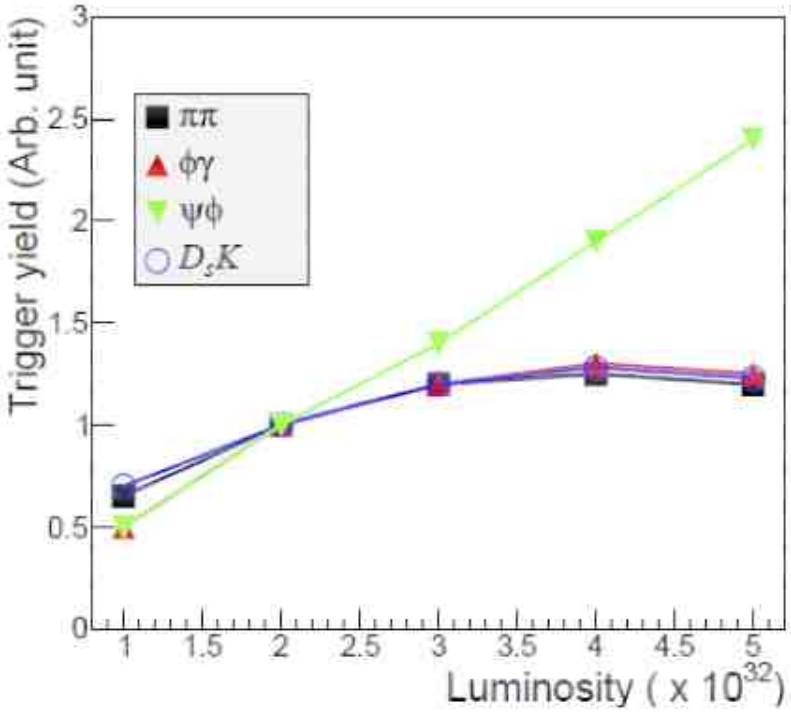

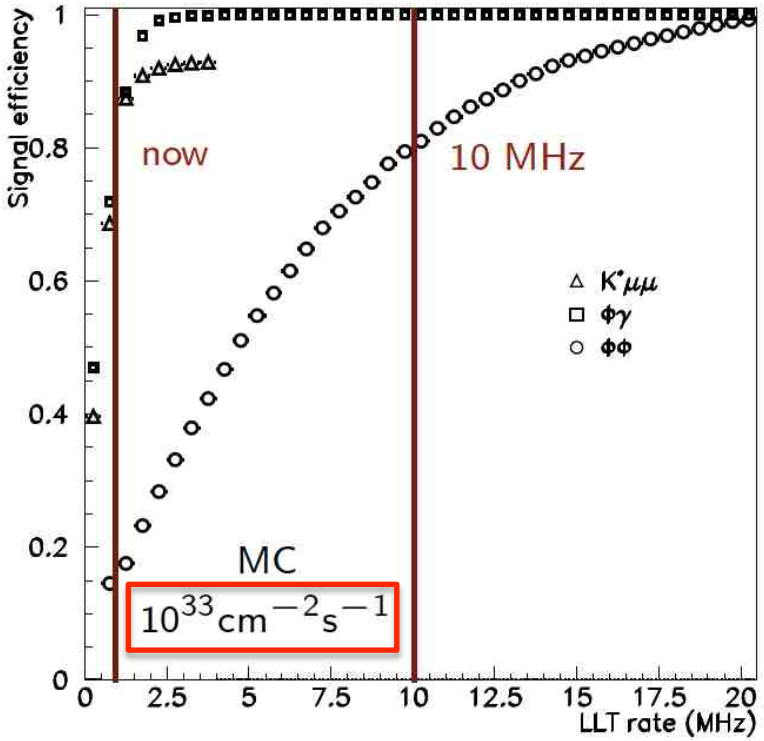

- Due to the available bandwidth and the limited discrimination power of the hadronic L0 trigger, LHCb experiences the saturation of the trigger yield on the hadronic channels around 4 $\times 10^{32}$ $cm^{-2}s^{-1}$
- Increasing the first level trigger rate considerably increases the efficiency on the hadronic channels.

"Framework TDR for the LHCb Upgrade: Technical Design Report"; The LHCb Collaboration,CERN/LHCC-2012-007, LHCb-TDR-12 - 2012

The LHCb upgrade

- Readout the whole detector at 40 MHz.
- Use a Low Level Trigger as a **throttle mechanism.**
- **Trigger-less data acquisition system** running at the highest possible frequency.
- We have foreseen to reach **20×10^{32} $cm^{-2}s^{-1}$** and therefore to prepare the sub-detectors on this purpose.
  - pp interaction rate 27 MHz
  - At 20. × 10^{32} cm^{-2} s^{-1}pile up μ ≅5.2
  - Increase the yield in the decays with muons by a factor five and the yield of the hadronic channels by a factor ten.
- Collect **50 fb^{-1}** of data over ten years.
  - **8 fb^{-1}** is the integrated luminosity target, to reach by 2018 with the present detector; **3.2 fb^{-1}** collected so far.
- The upgrade shall take place during the Long Shutdown 2 (LS2) in 2018/19.

Running Conditions

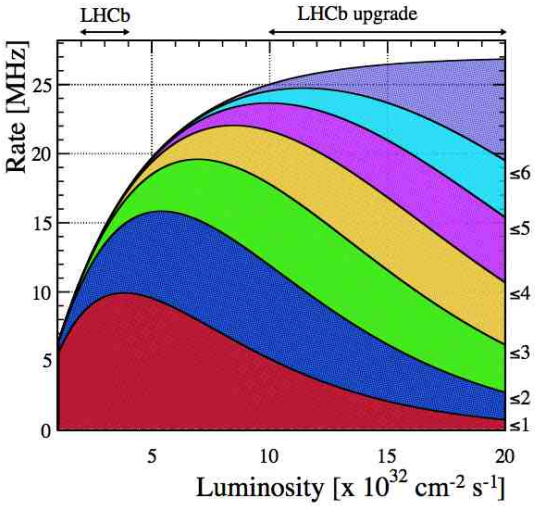

27 MHz

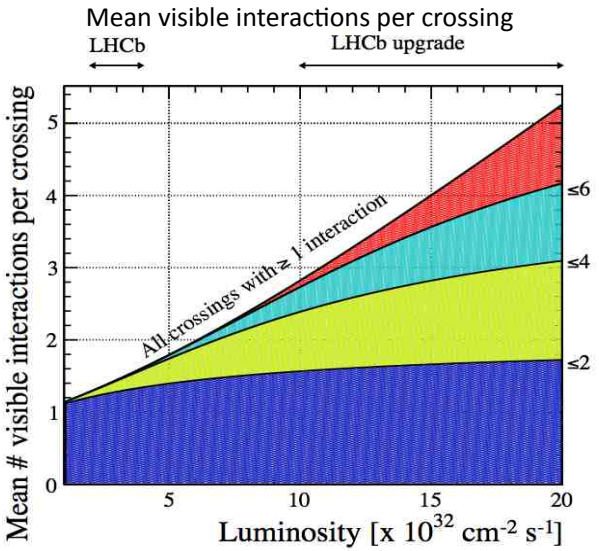

μ =5.2

LHCb upgrade: consequences

- The detector front-end electronics has to be entirely rebuilt, because of the current readout speed is limited to 1 MHz.
  - Synchronous readout, no trigger.
  - No more buffering in the front-end electronics boards.
  - Zero suppression and data formatting before transmission to optimize the number of required links.
    - **Average event size 100 kB**
  - Three times the optical links as currently to get the required bandwidth, needed to transfer data from the front-end to the read-out boards at 40 MHz.
    - **GBT links simplex (DAQ) 9000, GBT duplex (ECS/TFC) 2400**
- New HLT farm and network to be built by exploiting new LAN technologies and powerful many-core processors.
- Rebuild the current sub-detectors equipped with embedded front-end chips.
  - Silicon strip detectors: VELO, TT, IT
  - RICH photo-detectors: front-end chip inside the HPD.
- Consolidate sub-detectors to let them stand the foreseen luminosity of $20. \times 10^{32}$ cm^{-2} s^{-1}

DAQ upgrade: First idea

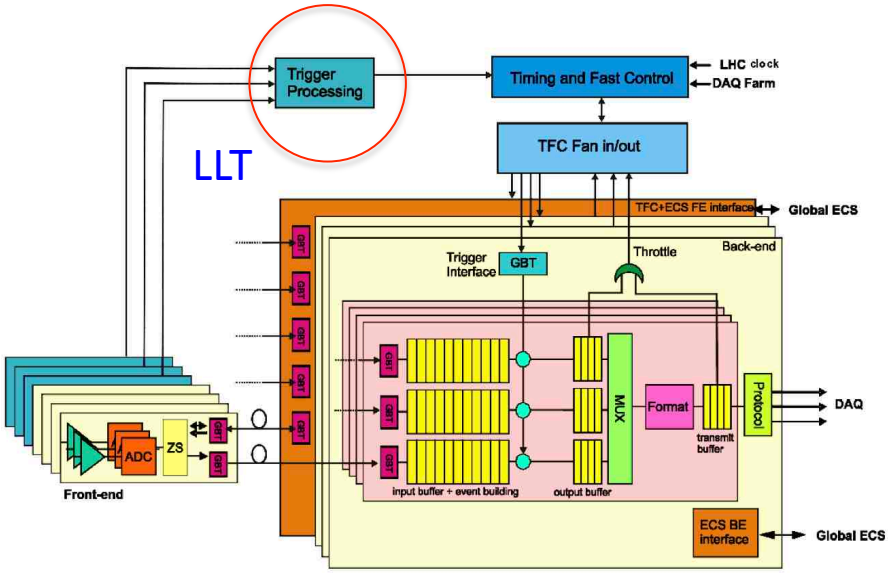

- **Hardware LLT** to provide a reduction in the rate of input events to be processed by the Event Filter Farm (EFF),

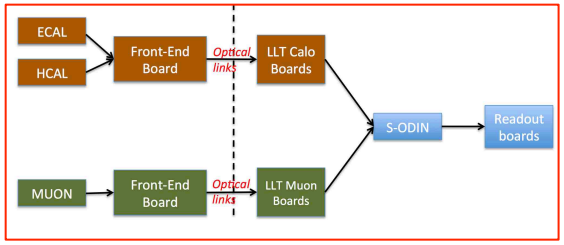

LLT performances at $L=2\times10^{33}$ $cm^{-2}s^{-1}$

Hadron Trigger

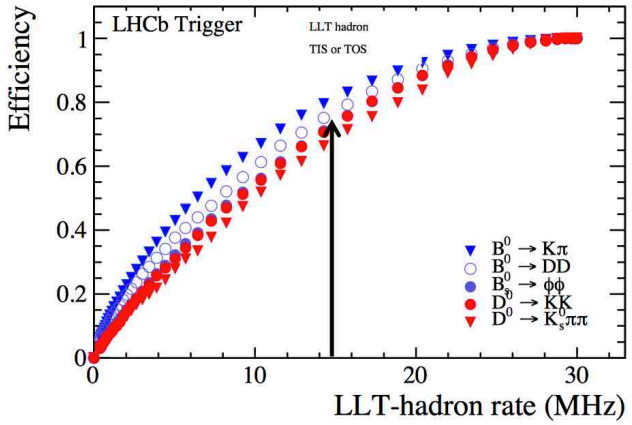

Muon Trigger

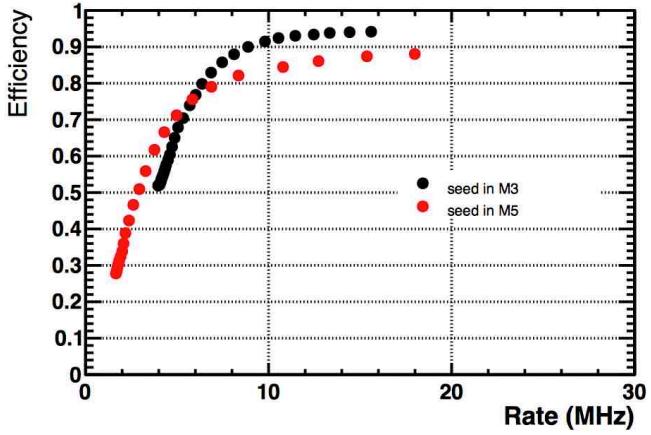

At 15 MHz, **efficiency** between 65% and 80% for hadronic channels, 85%-95% for muons

DAQ upgrade: First idea (2)

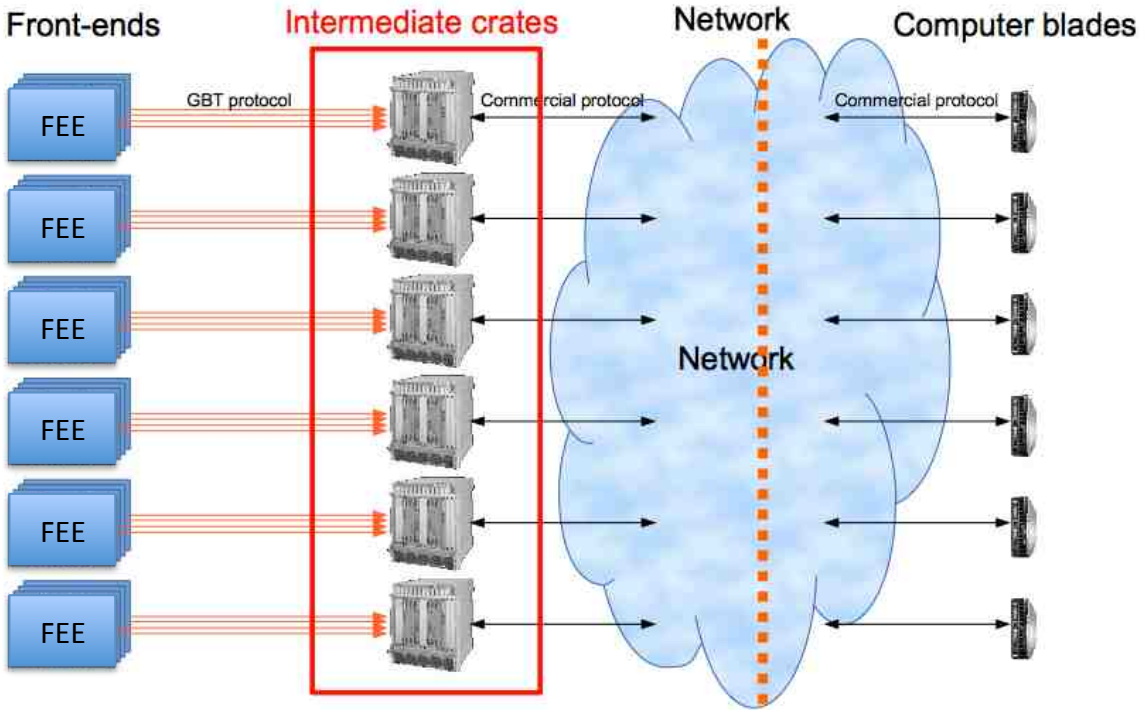

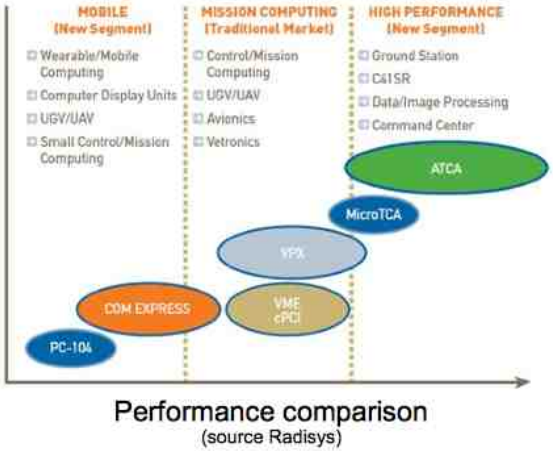

Performance comparison
(source Radisys)

Standard	Power consumption per slot
VPX 3U	75W
VPX 6U	150 W
µTCA	50 to 80 W
ATCA 10G	200W
ATCA 40G	400W

- Intermediate layer of electronics boards arranged in crates to decouple FEE and PC farm: for buffering and data format conversion.
- The optimal solution with this approach: ATCA, µTCA crates, ATCA carrier board hosting AMC standard mezzanine boards.
- AMC boards equipped with FPGAs to de-serialize the input streams and transmit event-fragments to the farm, using a standard network protocol, **using 10 GbEthernet**.

DAQ upgrade: First idea (3)

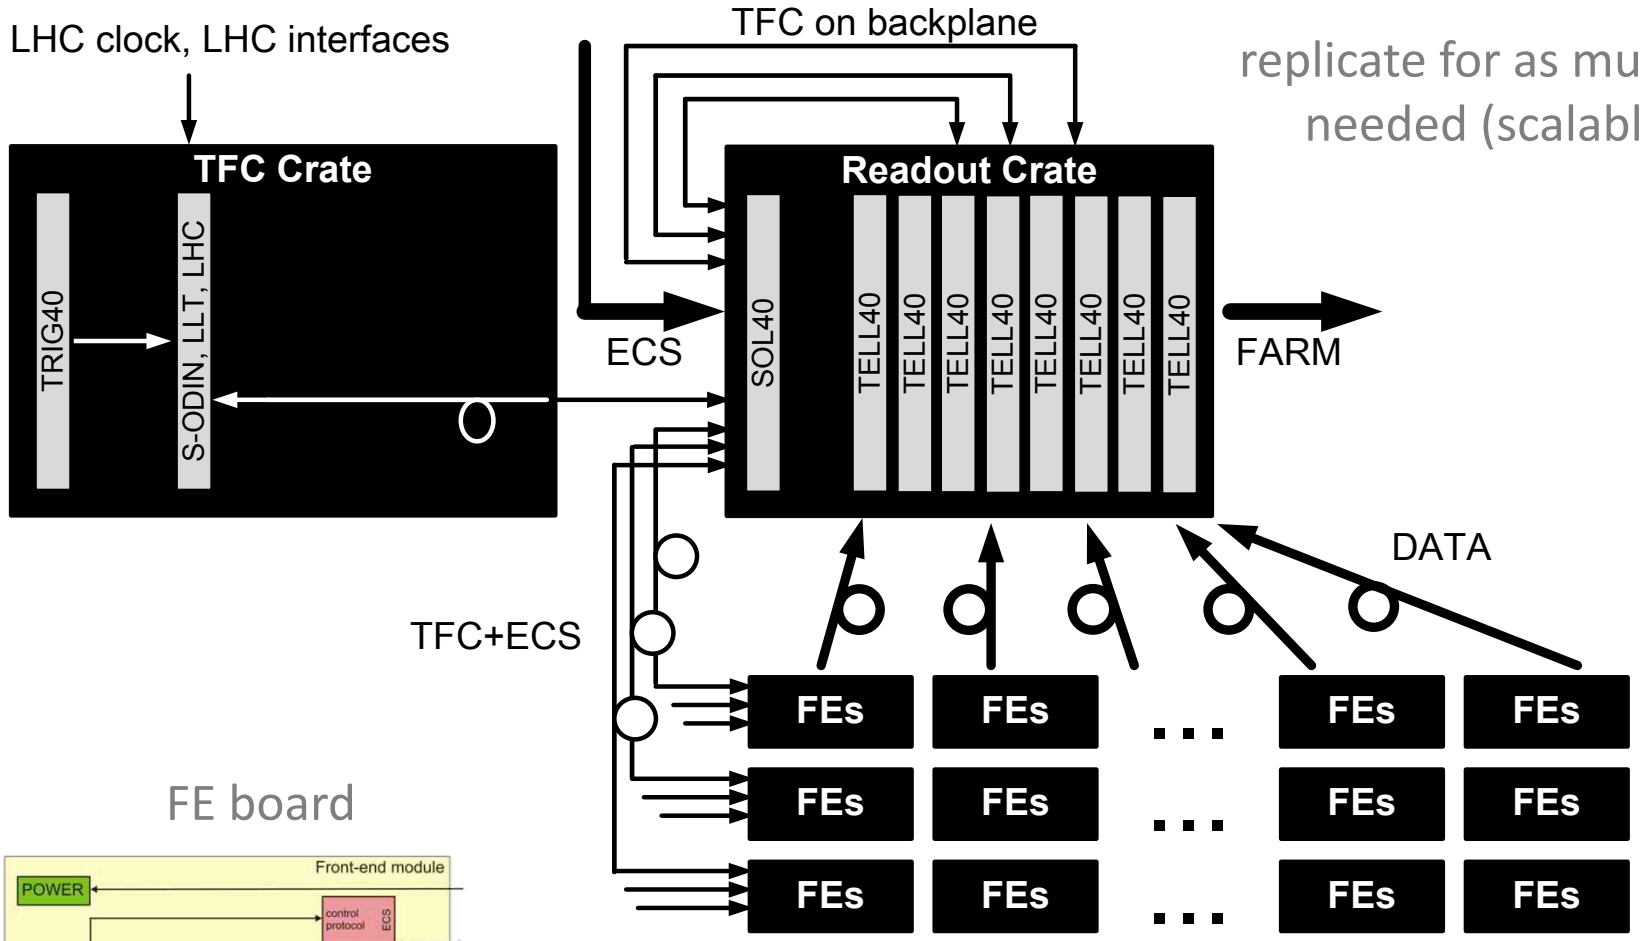

replicate for as much as needed (scalable)

FE board

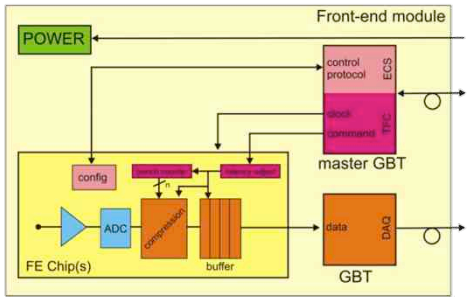

Same ATCA board used in different flavors: same generic hardware, different firmware

TELL40

Classical approach: ~0.5 Tb/s data aggregator board with 10 GbE to FARM

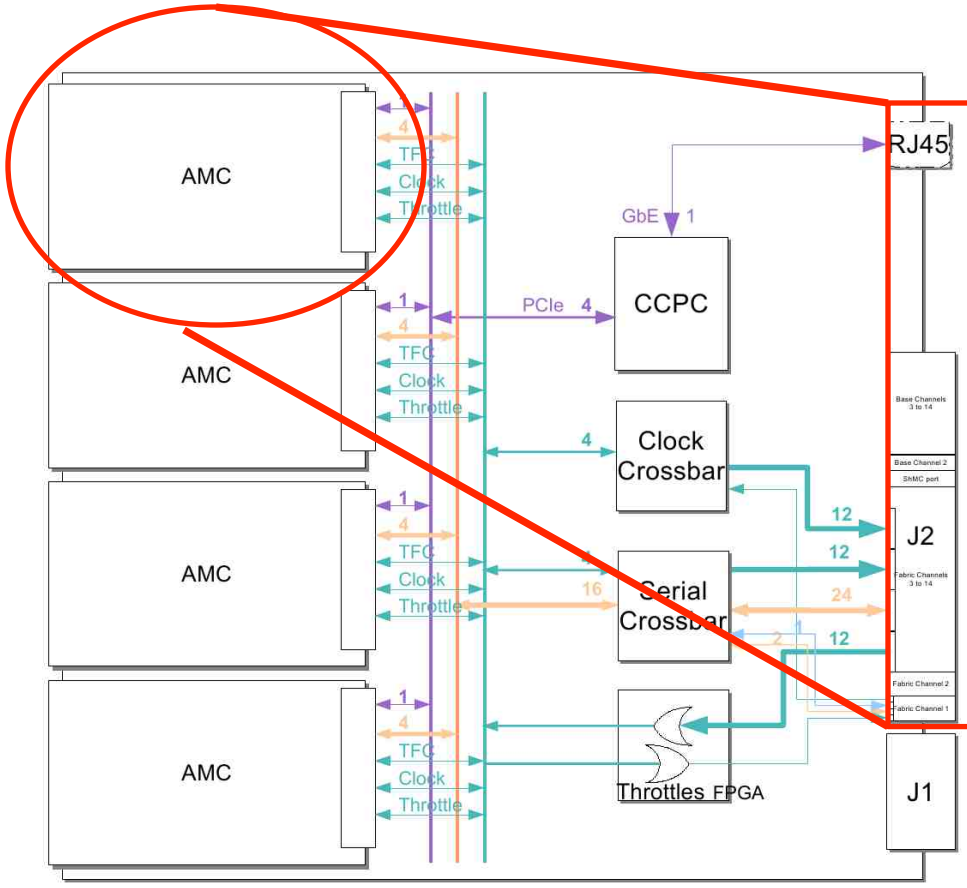

24 inputs
@ 4.8 Gb
GBT format

12 outputs/
inputs
@ 10Gb
Ethernet

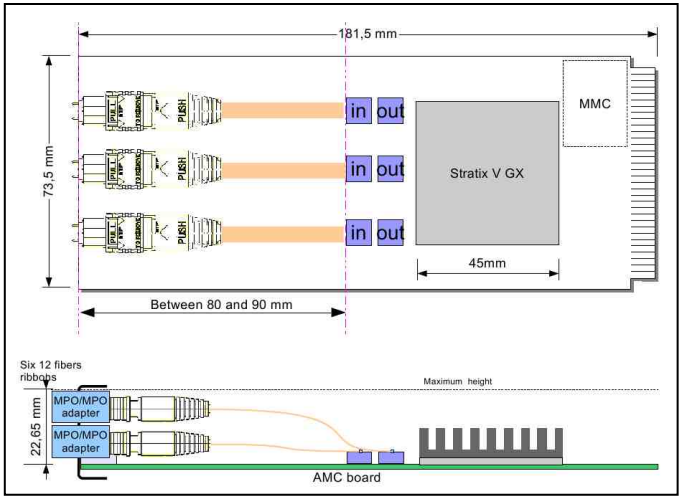

AMC40

AMC40 Input: 24 inputs @ 4.8 Gb GBT format, **Output**: 12 outputs/inputs @ 10GbE
TELL40 96 inputs @ 4.8 Gb → processing in FPGA → 48 x 10 Gb Ethernet ports

TELL40

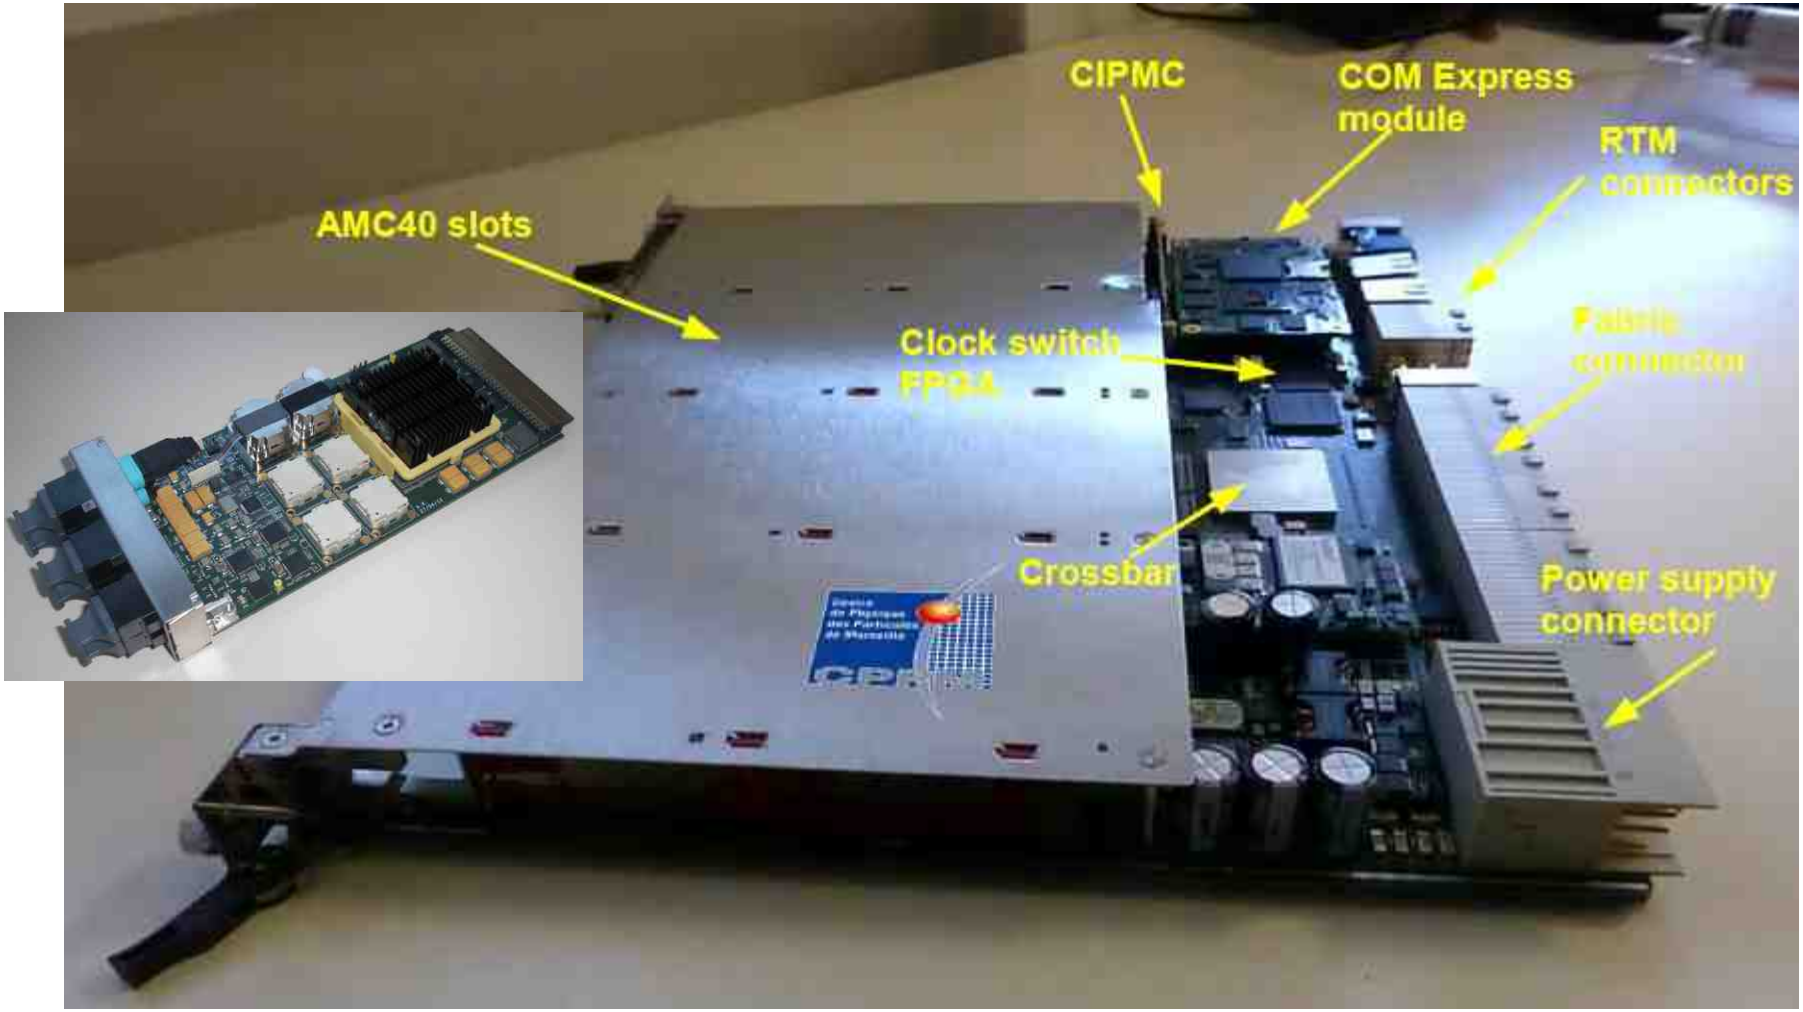
CIPMC
COM Express module
RTM connectors
AMC40 slots
Clock switch
Crossbar
Power supply connector

TPU: HLT assisted tracking

- "A specialized processor for track reconstruction at the LHC crossing rate". https://indico.inp.nsk.su/contributionDisplay.pycontribId=129&sessionId=6&confId=0
- "We have shown with a realistic detector arrangement that it is possible to reconstruct tracks and measure their parameters very well with a "brain inspired" cell-matrix method."

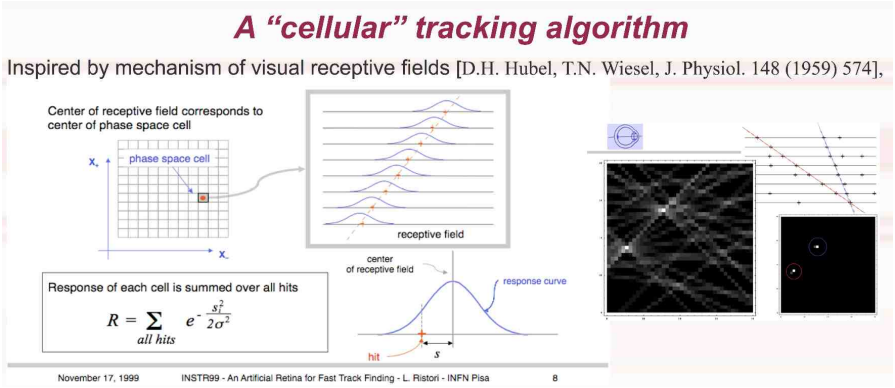

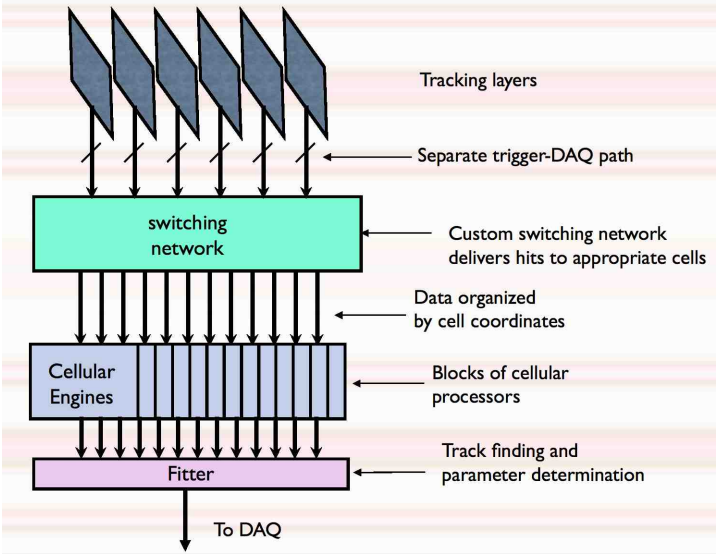

Implementation

- Use modern, large FPGA devices.
  - Large I/O capabilities: now O(Tb/s) with optical links !
  - Large internal bandwidth – a must !
  - Fully flexible, easy to program and simulate
  - Steep Moore's slope, and easy to upgrade
  - Highly reliable, easy to maintain and update
  - Industry's method of choice for complex project with a small number of pieces (CT scanners, high-end radars...)
- We used Altera's Stratix V
  - Same device used elsewhere in LHCb readout system.

DAQ present view

- Use PCIe Generation 3 as communication protocol to inject data from the FEE directly into the event-builder PC …

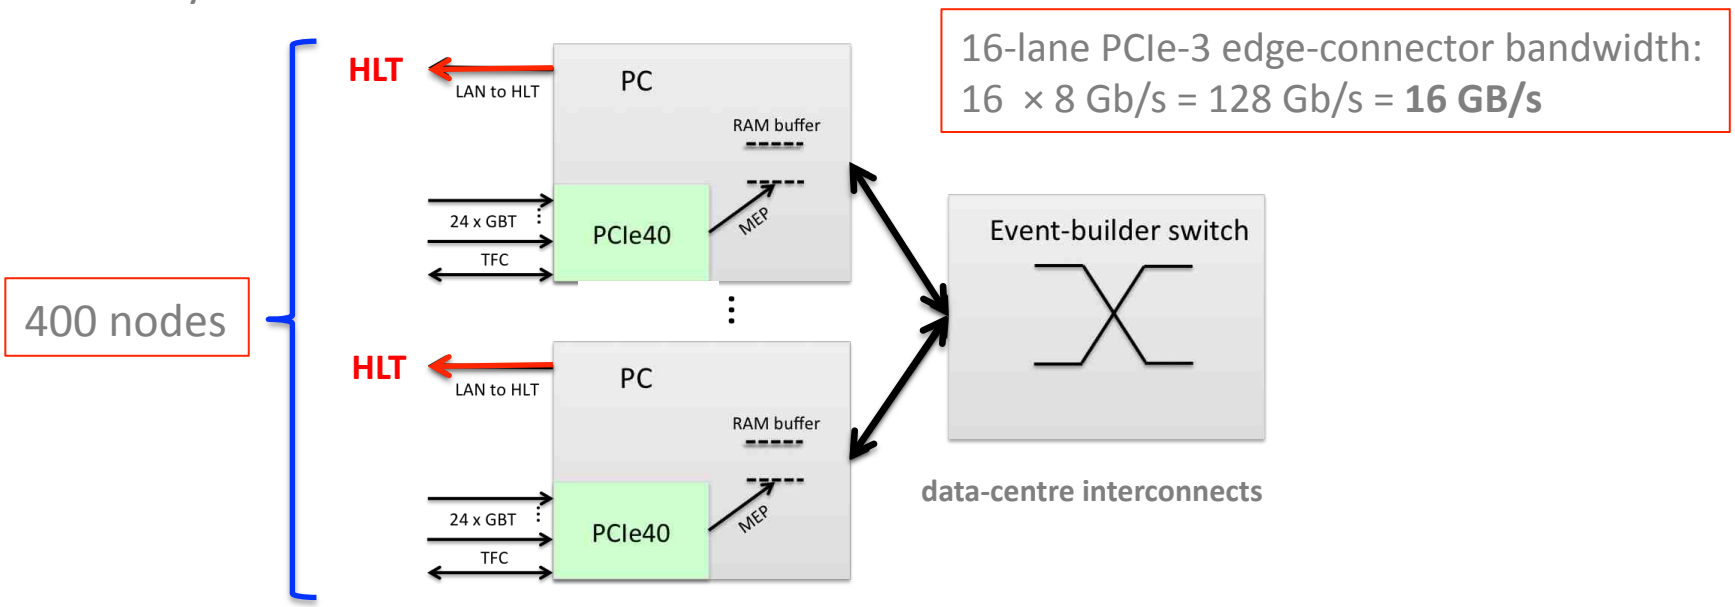

- A **much cheaper event-builder network** because **data-centre interconnects** can be used on the PC, which are not realistically implementable on an FPGA (large software stack, lack of soft IP cores,…)
- Moreover PC provides: huge memory for buffering, OS and libraries.

  Up to date network adapter cards and drivers available as pluggable modules.

DAQ network

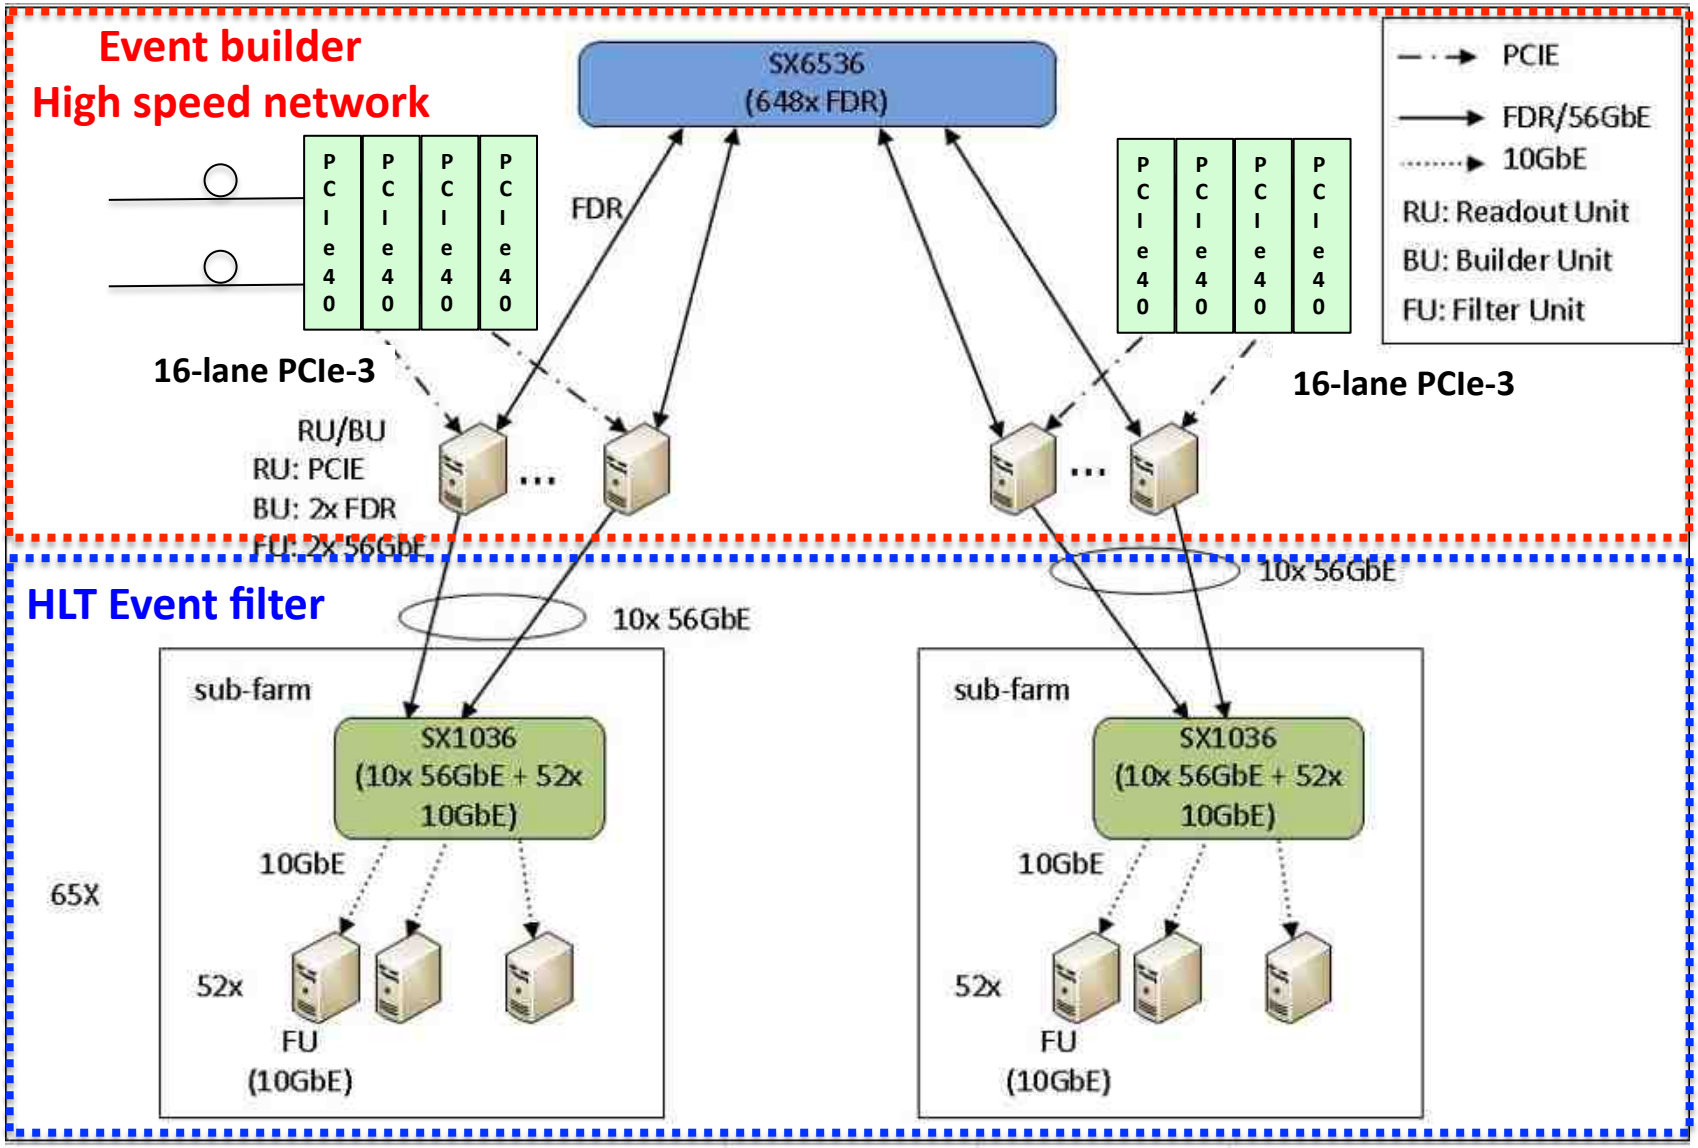
Event builder
High speed network
SX6536
(648x FDR)
PCIe40
FDR
16-lane PCIe-3
16-lane PCIe-3
RU/BU
RU: PCIE
BU: 2x FDR
FU: 2x 56GbE
10x 56GbE
10x 56GbE
HLT Event filter
sub-farm
sub-farm
SX1036
(10x 56GbE + 52x 10GbE)
SX1036
(10x 56GbE + 52x 10GbE)
10GbE
10GbE
65X
52x
52x
FU
(10GbE)
FU
(10GbE)
PCIE
FDR/56GbE
10GbE
RU: Readout Unit
BU: Builder Unit
FU: Filter Unit
E
B
F
E
F
F

EB Network Topology

- **Fat-Tree** (Constant Bisectional Bandwidth).
  - The fat tree topology maintains identical bandwidth at each level of the network.
  - The Fat Tree topologies do not scale linearly with the cluster size. Cabling and switching become increasingly difficult and expensive as cluster size grows, with very large core switches required for larger clusters
- In the simple **binary tree**, the number of links and thus the aggregate bandwidth is reduced by half at each stage of the network.

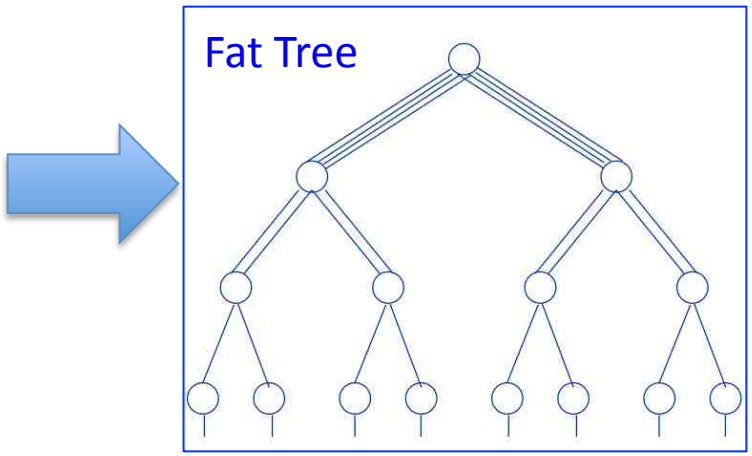

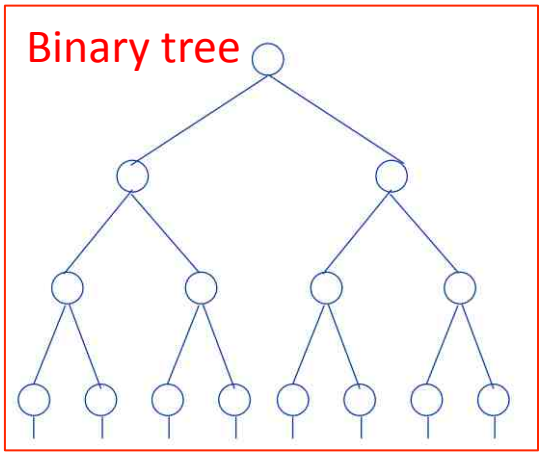

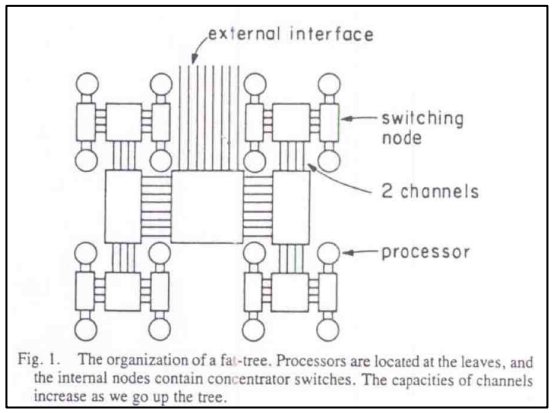

Fig. 1. The organization of a fat-tree. Processors are located at the leaves, and the internal nodes contain concentrator switches. The capacities of channels increase as we go up the tree.

Fat-Trees: Universal Networks for Hardware-Efficient Supercomputing, C.E. Leiserson. IEEE Transactions on Computers, Vol 34, October 1985, pp 892-901

Note that the links are bidirectional so the notion of upstream and downstream describes the direction of the interconnect topology towards or away from the shortest path to an end node, rather than the actual data flow.

Long distance fibres

Locate the **Event Builder Farm** and the **Event Filter Farm** for the HLT **outside** the cavern.

GBT links simplex (DAQ)	**9000**
GBT duplex (ECS/TFC)	**2400**

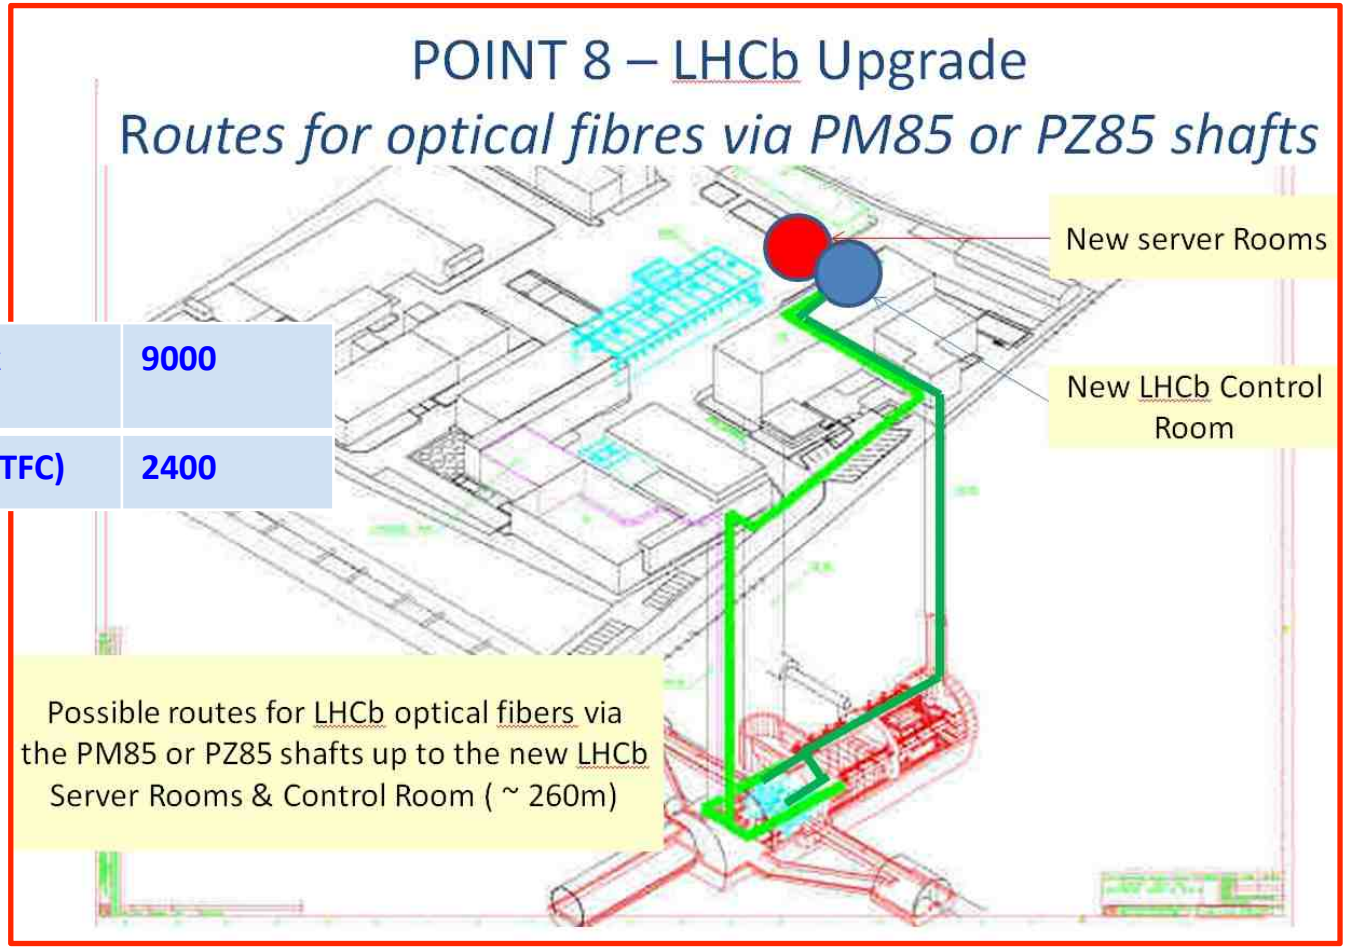

The distance to cover with 850 nm OM multimode optical cables, from underground to the surface, is **300 m**.

PCIe Gen3 based readout

- A main FPGA manages the input streams and transmits data to the event-builder PC by using DMA over PCIe Gen3.
- The readout version of the board uses two de-serializers.
- The same board can be used to clock and control distribution.

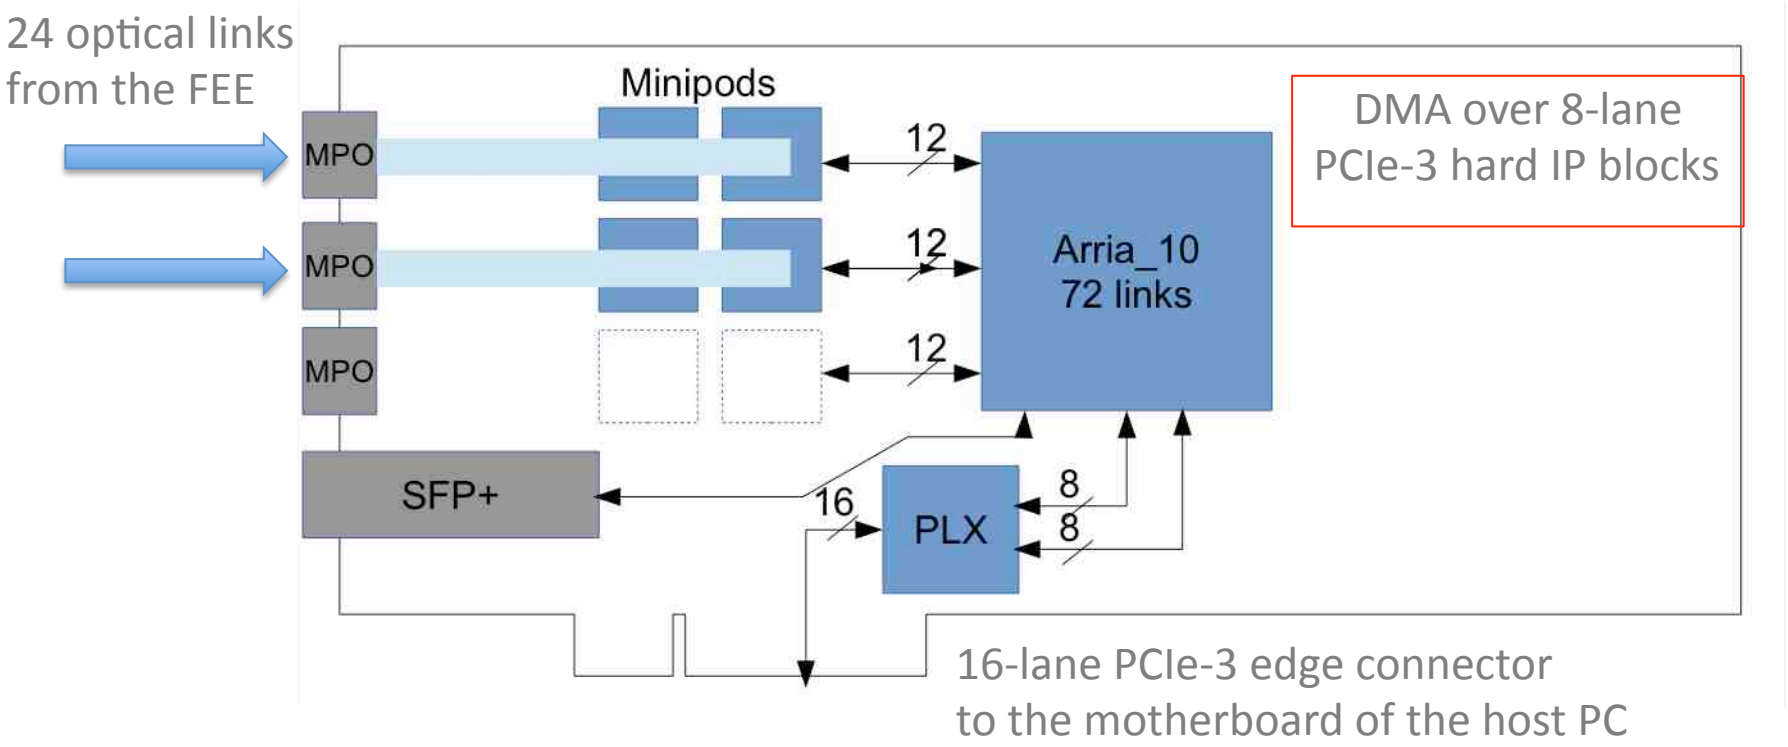

The PCIe-3 test setup

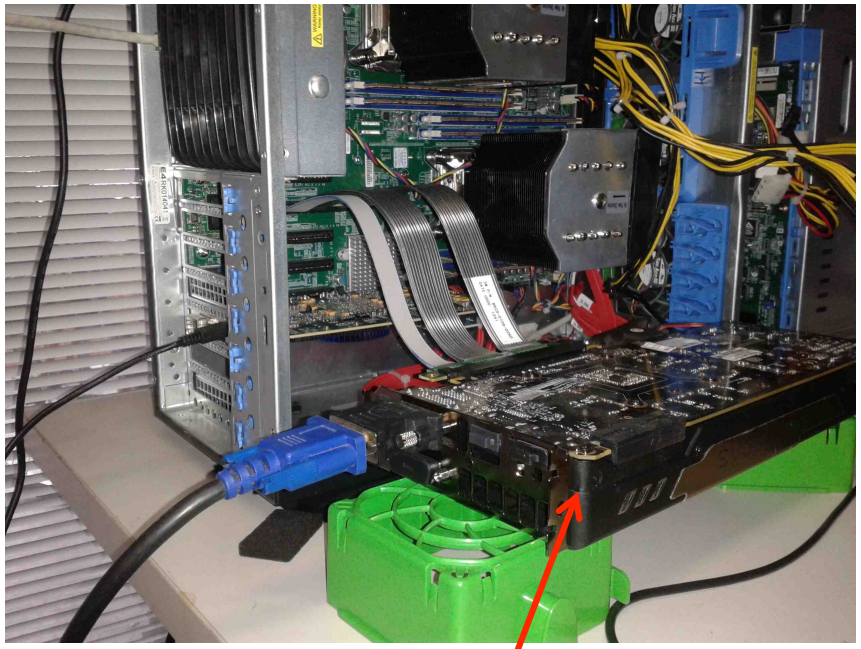

GPU used to test 16-lane PCIe-3 data transfer between the device and the host memory

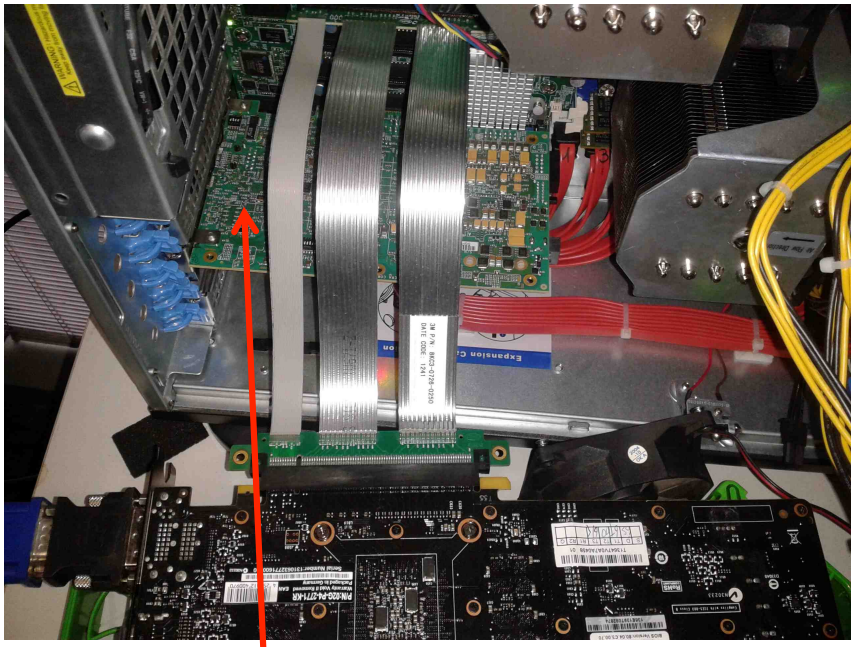

The FPGA provides 8-lane PCIe-3 hard IP blocks and DMA engines.

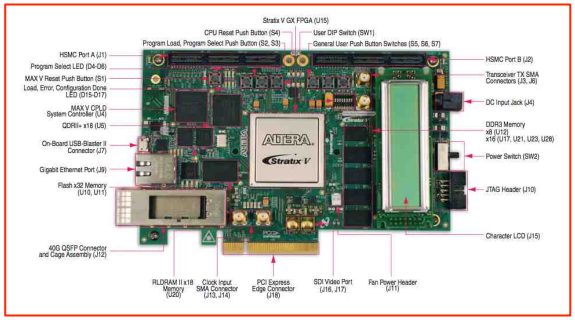

- ALTERA development board, Stratix V GX FPGA, model 5SGXEA7K2F40C2N

DMA PCIe-3 effective bandwidth

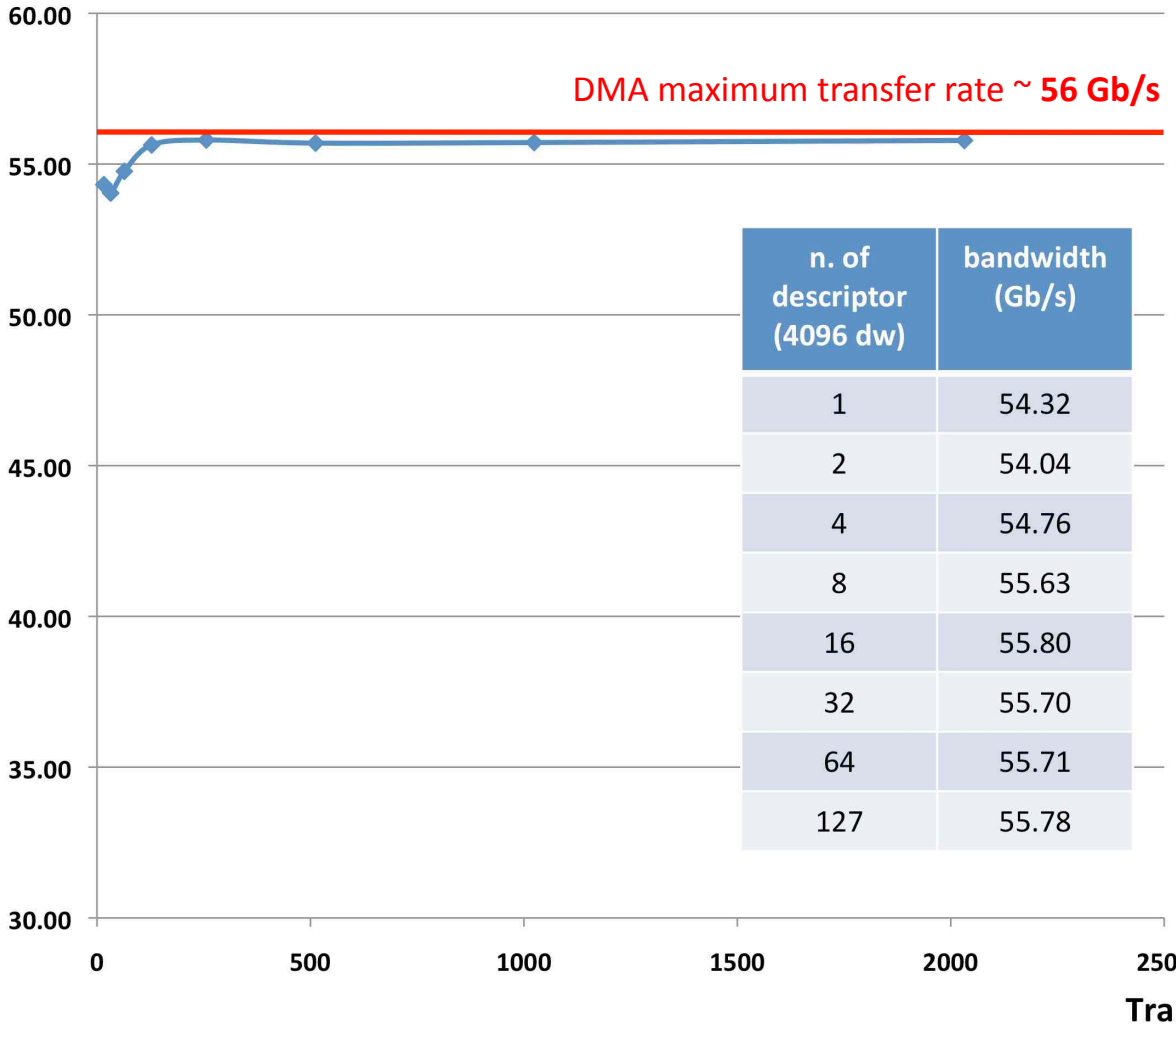

n. of descriptor (4096 dw)	bandwidth (Gb/s)
1	54.32
2	54.04
4	54.76
8	55.63
16	55.80
32	55.70
64	55.71
127	55.78

DMA over 8-lane PCIe-3 hard IP blocks ALTERA Stratix V

Test of PLX bridge

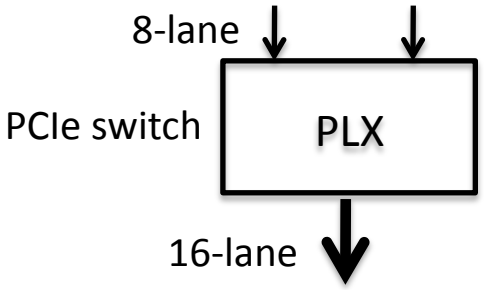

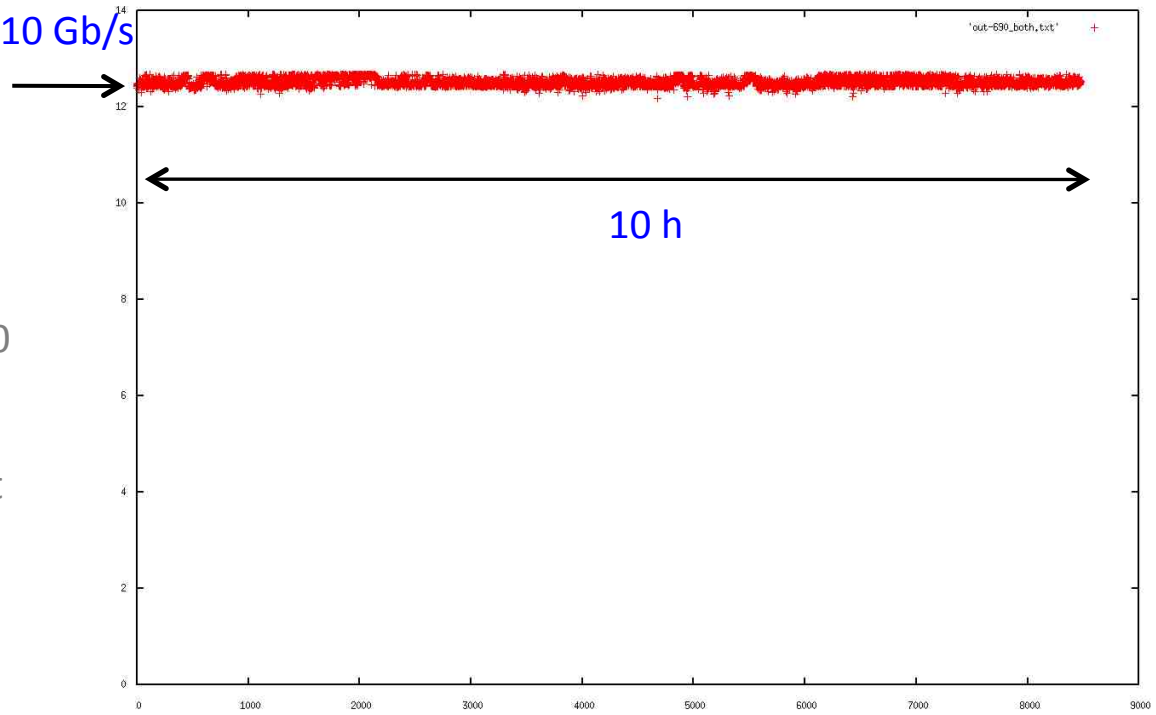

- Long-term test using GTX690 card / PLX 8747 bridge.
- Zero impact of using a bridge and two independent PCIe targets pushing data into a PC. Consistently around 110 Gb/s over long-term.
- No load balancing issues between the two competing links observed.
- Details at: https://lbonupgrade.cern.ch/wiki/index.php/I/O_performance_of_PC_servers#Upgrade_to_GTX690

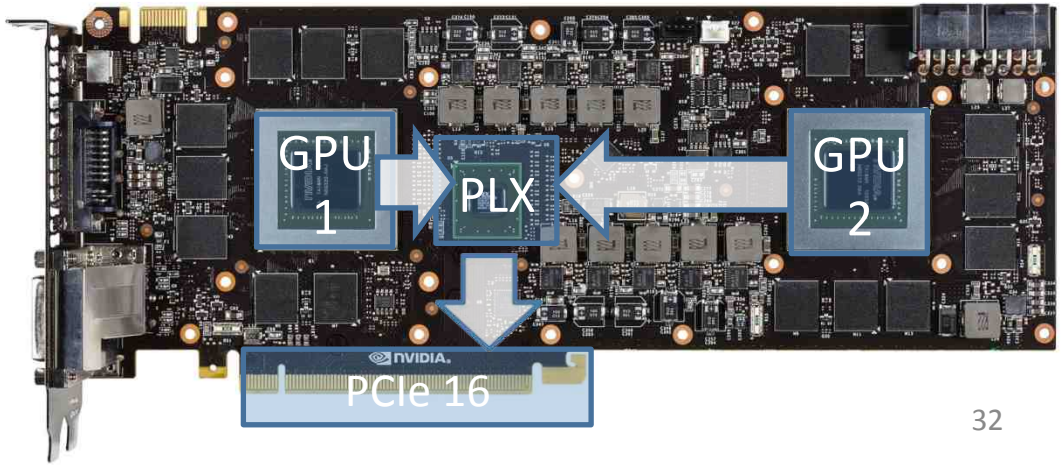

Event builder fluxes: 400 Gb/s

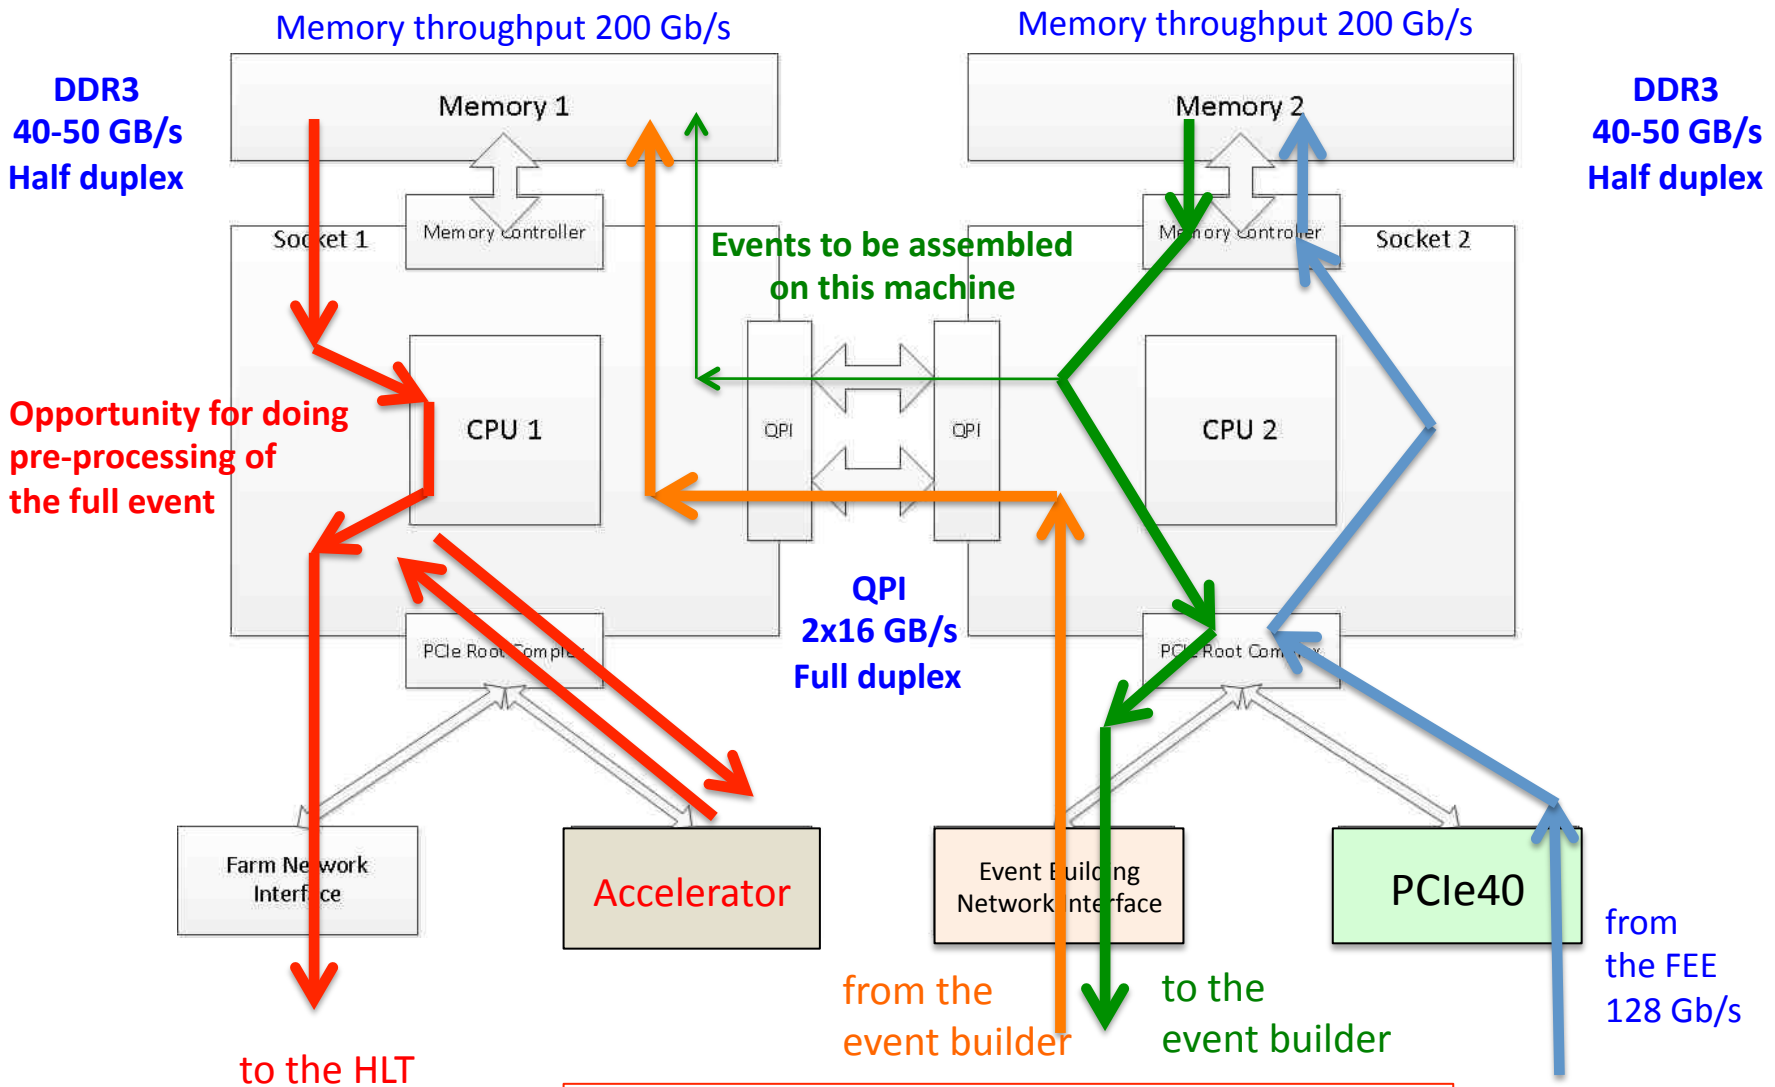

PC event-builder tests

Setup

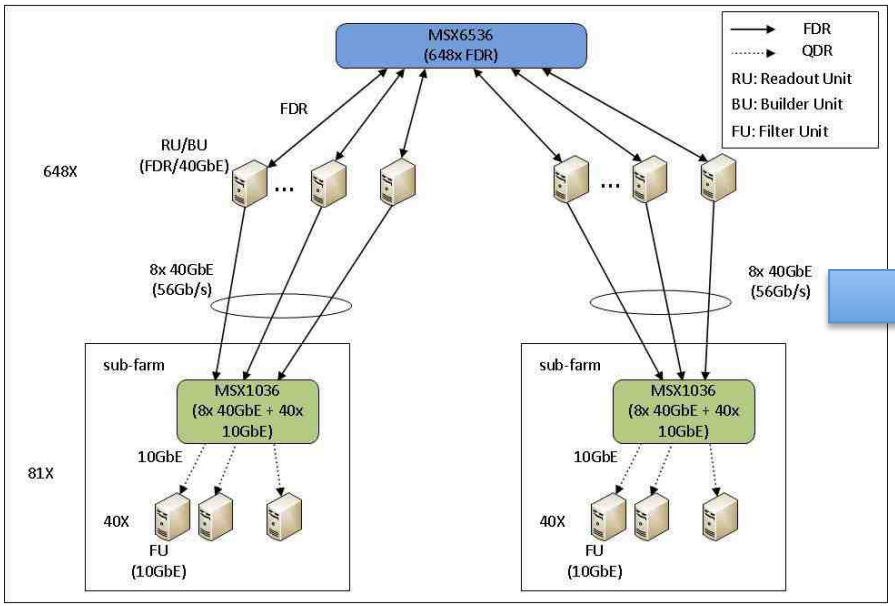

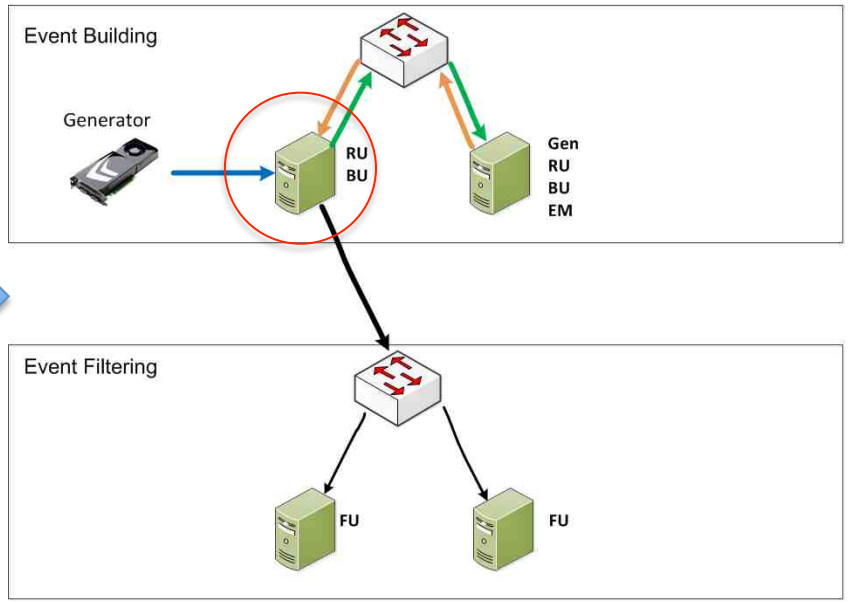

Protocol

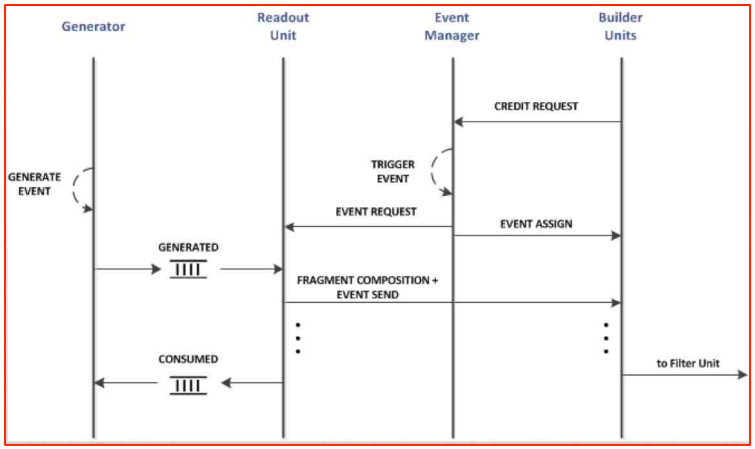

- Test with a Ivy Bridge Intel dual CPU, 12 cores, PC.
- InfiniBand FDR connections
- OpenMPI-based application

EVB performance

At about 400 Gb/s more than 80% of the CPU resources are free

CPU consumption

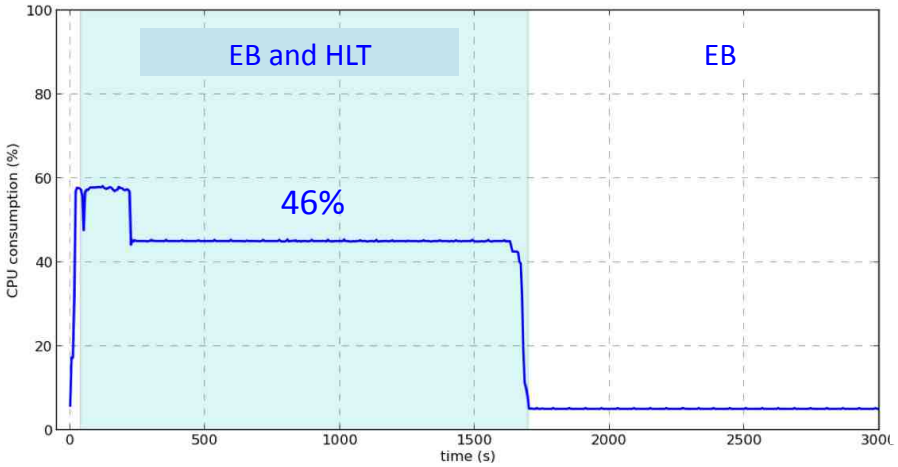

Memory I/O bandwidth

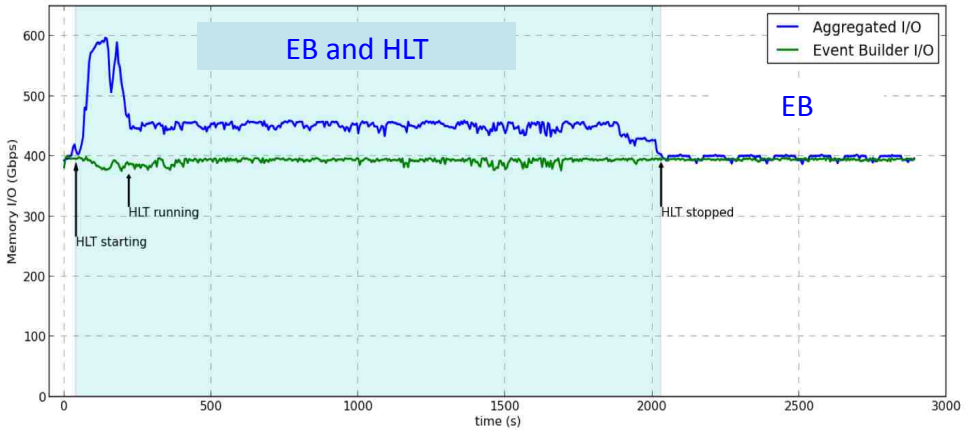

Memory consumption

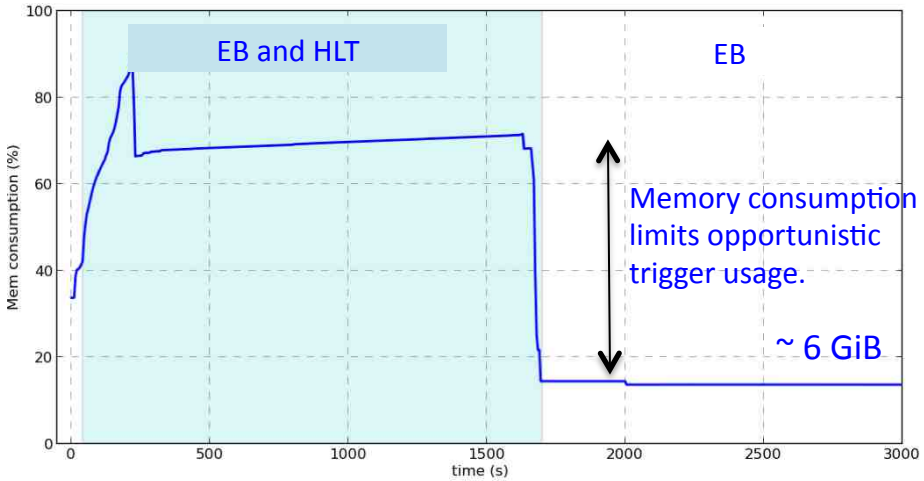

- PC sustains the **event building at 100 Gb/s today**.
- The Event Builder performs stably at **400 Gb/s**
- Aggregated CPU utilization of EB application and trigger 46%
- We currently observe **50% free resources for opportunistic triggering on EB nodes**: event builder execution requires about 6 logical core. Additional 18 instances of the HLT software running simultaneously.

The CPUs used in the test are Intel E5-2670 v2 with a C610 chipset. The servers are equipped with 1866 MHz DDR3 memory in optimal configuration. Hyper-threading has been enabled.

Software LLT

- The LLT algorithms can be executed in the event builder PC after the event building.
- Preliminary studies show that the LLT runs in less than 1 ms, if the CALO clusters are built in the FEE.
- Assuming 400 servers, 20 LLT processes running per PC, and a factor 8 for the CPU power from the Moore Law, the time budget available turns out to be safely greater then 1ms:

$$\frac{1}{40\,MHz} \times 400 \times 20 \times 8 \approx 3.2\,ms$$

$$\text{processing time budget} = \frac{1}{\text{event rate}} \times \text{nodes} \times \text{cores per node} \times \text{task per node}$$

DAQs Upgrade

Future LHC DAQs in numbers

	Event-size [kB]	Rate [kHz]	Bandwidth [Gb/s]	Year [CE]
ALICE	20000	50	8000	2019
ATLAS	4000	200	6400	2022
CMS	2000	200	3200	2022
LHCb	100	40000	32000	2019

- Some overlapping trends across experiments, at least conceptually
  - custom-made Readout Boards with fast optical links and big&powerful FPGAs
    - ✓ ideally with fast interface to PCs (PCIe Gen3 or future…)
    - ✓ ideally with some co-processing (Xeon, GPUs…)
  - commercial network technologies following market trends in terms of BW & costs
    - ✓ distributed vs data-center-like network.
    - ✓ network technologies: Ethernet vs InfiniBand.

LHCb upgrade: HLT farm

- Trigger-less system at **40 MHz**:
  A selective, efficient and adaptable software trigger.
- Average event size: **100 kB**
- Expected data flux: **4 TB/s**
- Total HLT trigger process latency: ~**15 ms**
  - Tracking time budget (VELO + Tracking + PV searches): 50%
  - Tracking finds **99%** of offline tracks with p_T >500 MeV/c
- Number of running **trigger process** required: **4×10^5**
- Number of core/CPU available in 2018: ~ **200**
  - Intel tick-tock plan: 7nm technology available by 2018-19, the number of core accordingly scales as $12 \times (32\ nm/\ 7\ nm)^2 = 250$, equivalent 2010 cores.
- Number of computing nodes required: ~ **1000**

The most powerful CPU ever built: IBM Power8 CPU, http://www.extremetech.com/computing/181102-ibm-power8-openpower-x86-server-monopoly

DAQ numbers

GBT links simplex (DAQ)	9000
GBT duplex (ECS/TFC)	2400
PCIe40 (DAQ)	400
SOL40 (ECS/TFC)	66
Estimated total mean event-size	100 kB
100G links for event-building network @ 40 MHz	400
Dual-socket servers for High Level Trigger	2000 – 4000

Conclusions

- The concept of the LHCb experiment has been definitely proved: exploiting a forward spectrometer at a hadron collider to perform a dedicated experiment for heavy flavour physics.
  - Many world leading results and many more to come with the 3.2 fb^{-1} full data set collected.
- LHCb plans the upgrade, to be installed in 2018: the upgrade is an essential next step forward for flavour physics.
- The DAQ system we envisage for the upgrade will allow us to feed the HLT trigger and exploit its capabilities at 40 MHz.
- TDRs are all ready.

Spare material

Prospects

	LHC era		HL-LHC era		
	2010-2012	2015-2017	2019-2021	2024-2026	2028-2030+
ATLAS & CMS	25 fb^{-1}	100 fb^{-1}	300 fb^{-1}	→	3000 fb^{-1}
LHCb	3 fb^{-1}	8 fb^{-1}	23 fb^{-1}	46 fb^{-1}	100 fb^{-1}
Belle II	-	0.5 ab^{-1}	25 ab^{-1}	50 ab^{-1}	-

Expected precision on BR(Bd→μ+μ-)/BR(Bs→μ+μ-)

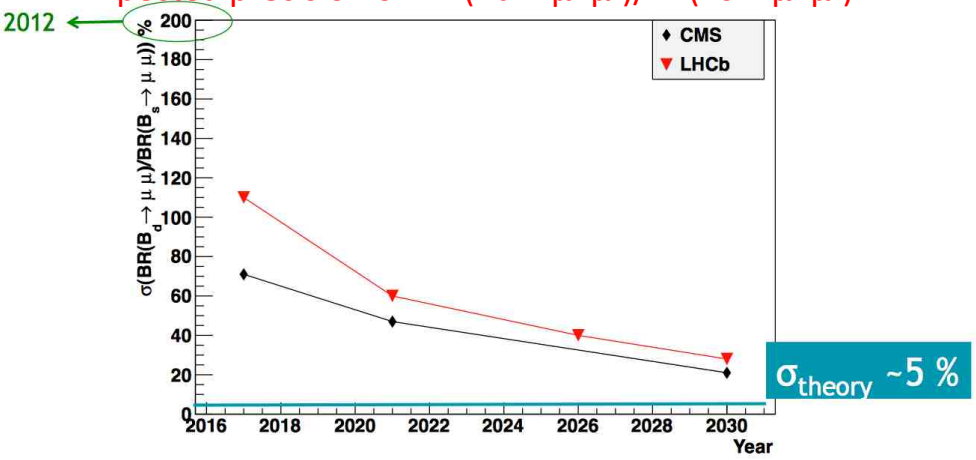

Expected precision on ϕ_s (rad)

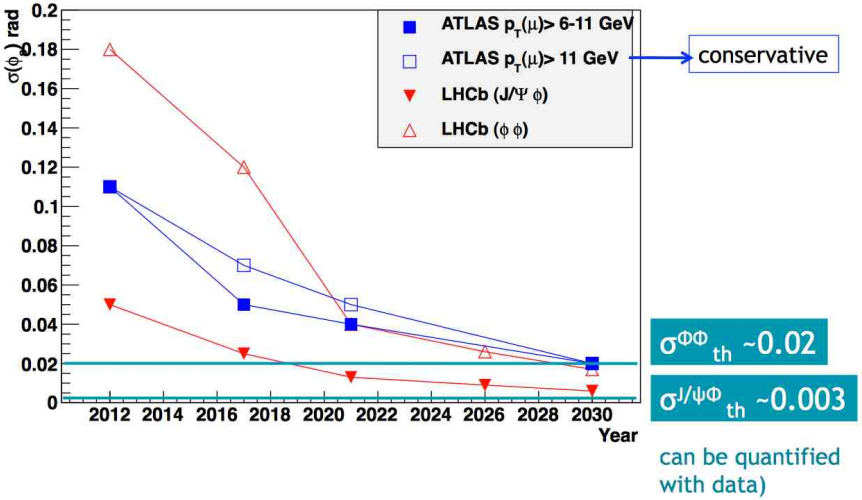

Prospects (II)

The "ultimate" γ-from tree-decays precision will be reached through many individual measurements, with very different sensitivities (due to different *b→u to b→c* amplitudes ratio)

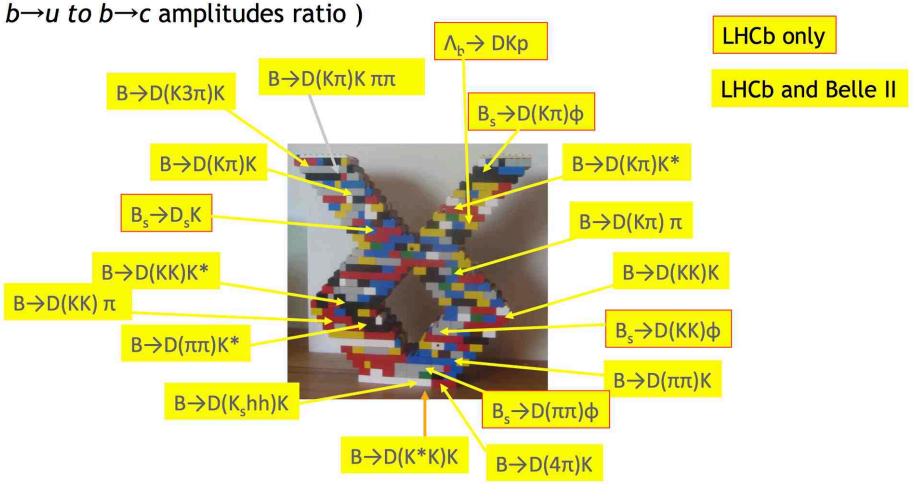

Expected precision on γ from tree decays

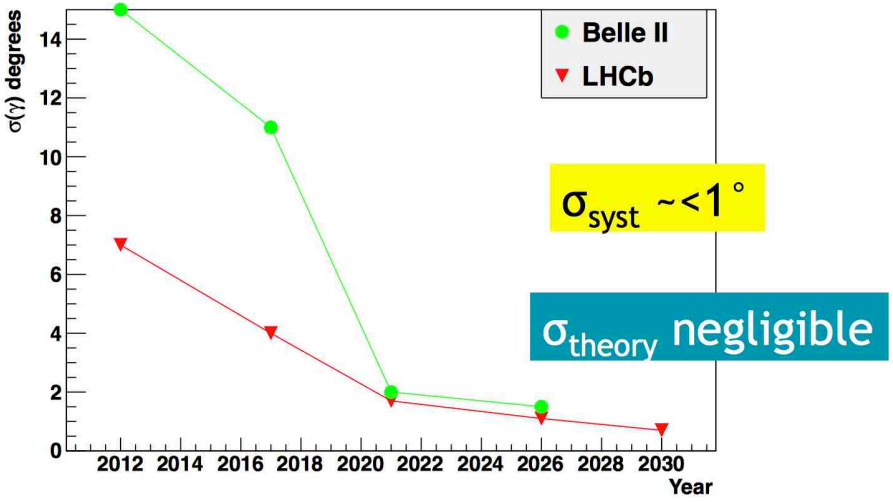

Evolution of LHC luminosity

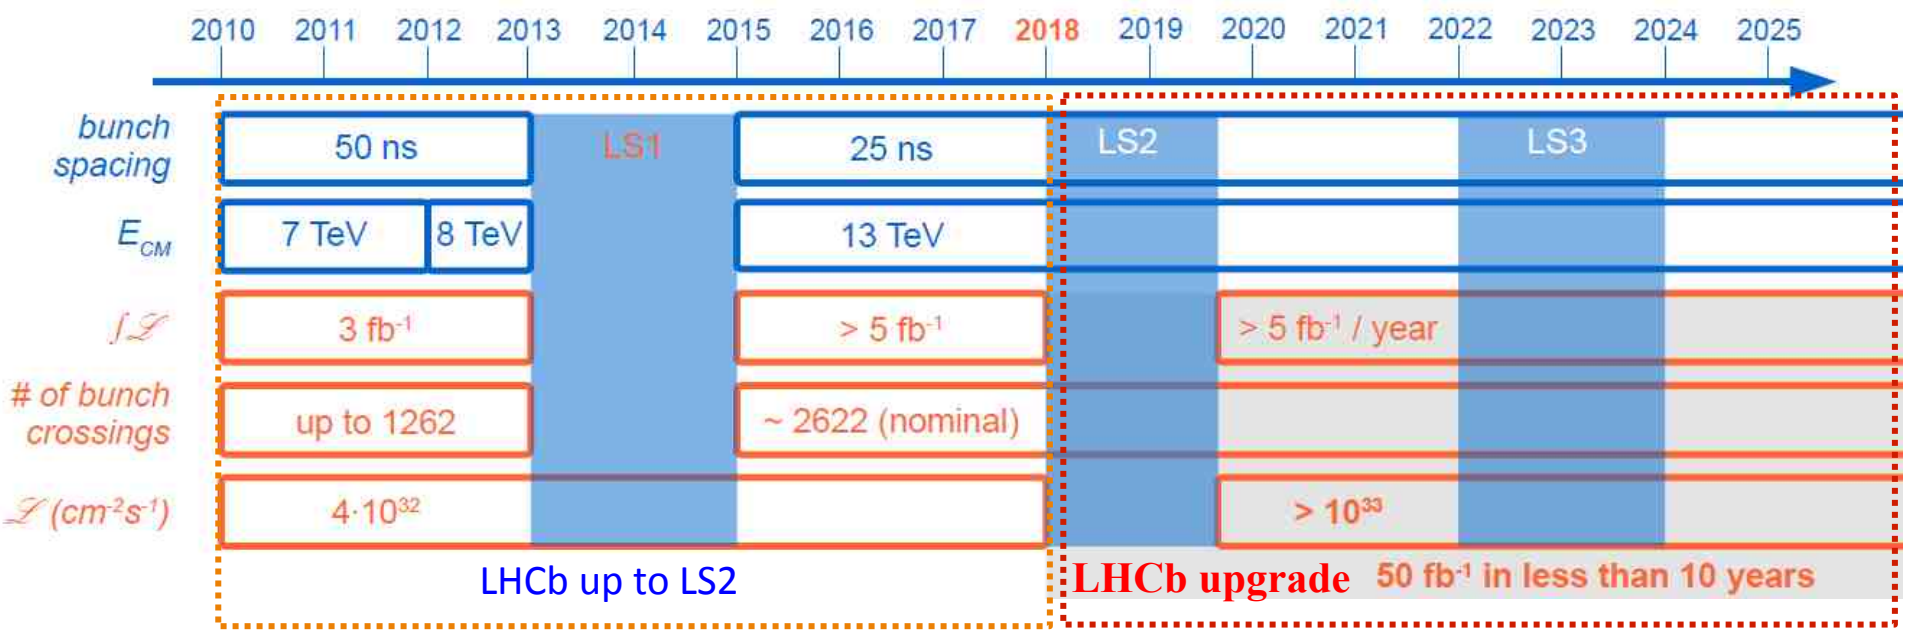

LHCb up to 2018 → ~ 8-10 fb^{-1} :

- find or rule-out large sources of flavour symmetry breaking at the TeV scale

LHCb upgrade → ≥ 50 fb^{-1}:

- increase precision on quark flavour physics observables
- aim at experimental sensitivities comparable to theoretical uncertainties
- reinforce LHCb as a general purpose forward detector

Long distance fibres (II)

- 4.8 Gbit/s signal produced on the detector by Versatile Link transmitters.
- VTTx to MiniPod for data acquisition.
- MiniPod to VTRx for control, configuration.

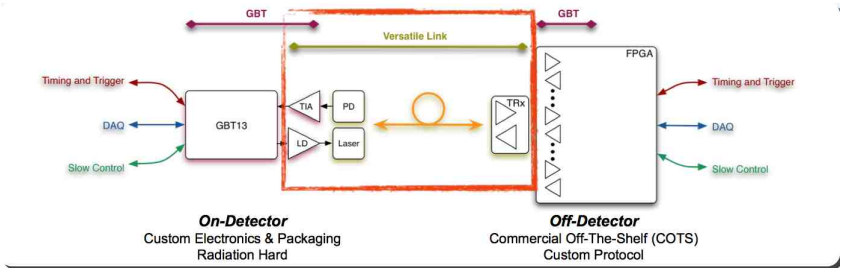

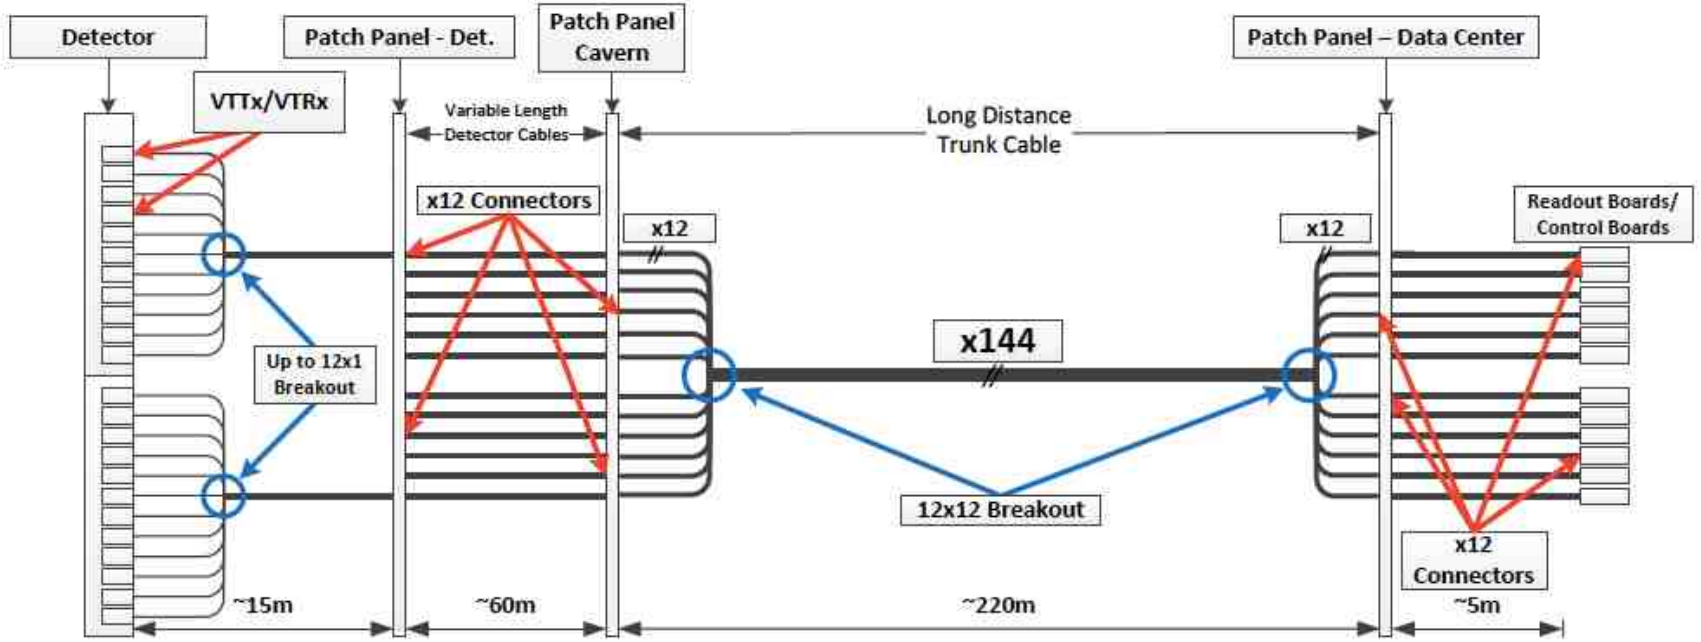

- 144 fibres per cable. A total of 120 such cables.
- 3 patch panels (breakpoints) foreseen: expected attenuation ~ 3dB

MiniPod

Optical fibres studies

850 nm High Performance EMB (MHz/km)
OM3: 2000

OM4: 4700

850-nm Ethernet Distance
1-GbE
OM3: 1000 m

OM4: 1000 m

10-GbE
OM3: 300 m

OM4: 550 m

40-GbE
OM3: 100 m

OM4: 150 m

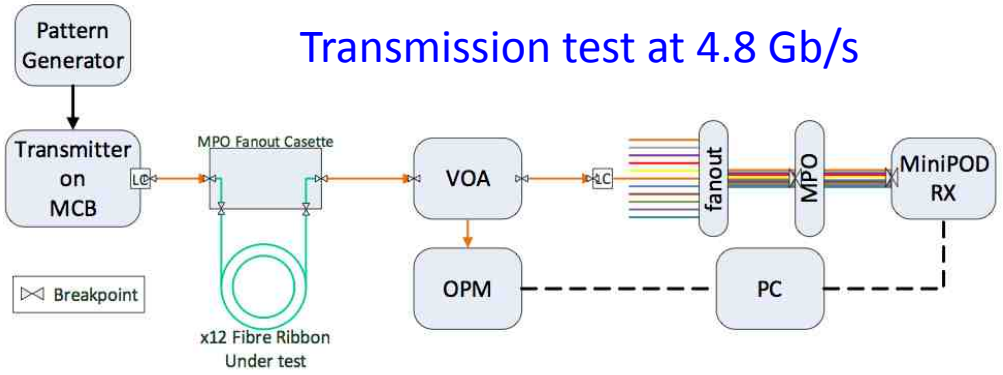

Transmission test at 4.8 Gb/s

Figure 3: Schematic of Setup B, used for measuring BER vs. OMA for the DAQ direction.

Measurement of BER vs. receive OMA(*) on different OM3 and OM4 fibres.
(*) optical modulation amplitude

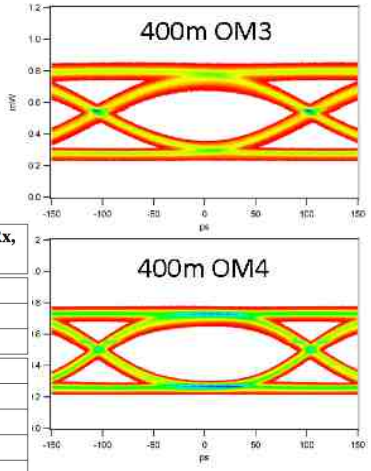

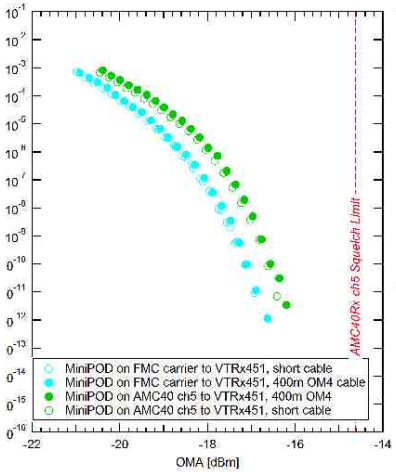

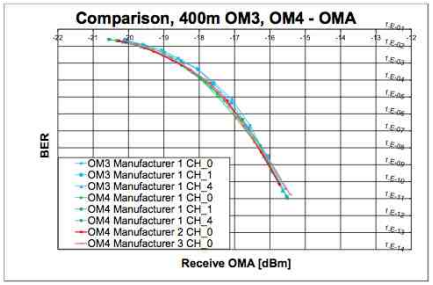

Description	Unit	VTTx to MP, spec.	VTTx to MP, meas.	MP to VRx, spec.	MP to VRx, meas.
Transmitter OMA	dBm	-5.2	NM	-3.2	NM
Receiver sensitivity	dBm	-11.1	-14.2	-13.1	NM
Power budget	dB	5.9	9.0	9.9	9.9
Fibre loss (2.3 dB/km)	dB	0.9	NM	0.9	NM
Connectors (0.5 dB/pair)	dB	1.5	NM	1.5	NM
Disp. (400 m, 4.8 Gbit/s)	dB	2.4	0.5	2.4	0.5
TX Radiation penalty	dB	0	NM	-	-
RX Radiation penalty	dB	-	-	2.5	NM
Fibre Radiation penalty	dB	0.1	NM	0.1	NM
Margin	dB	1.0	6.0	2.5	4.4

The target value of 3 dB on OM3 using the Versatile Link and MP is reachable.

Fast Control

Clock distribution and phase/latency control

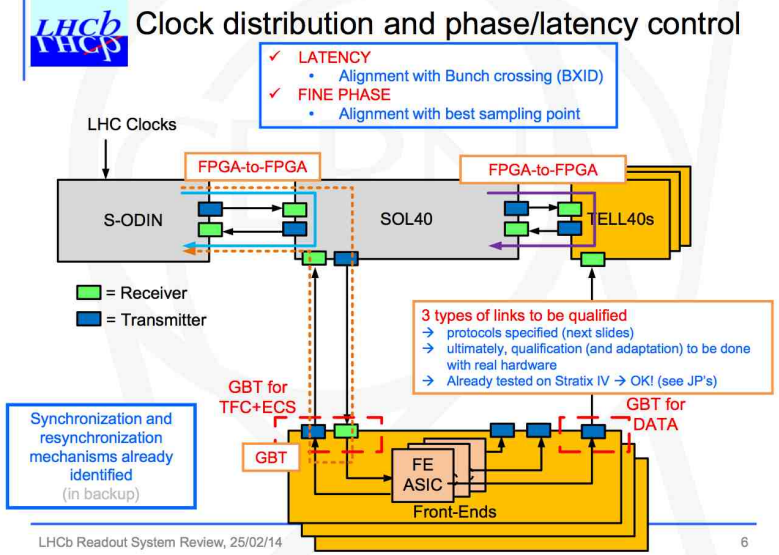

TFC in the upgrade – PCIe case

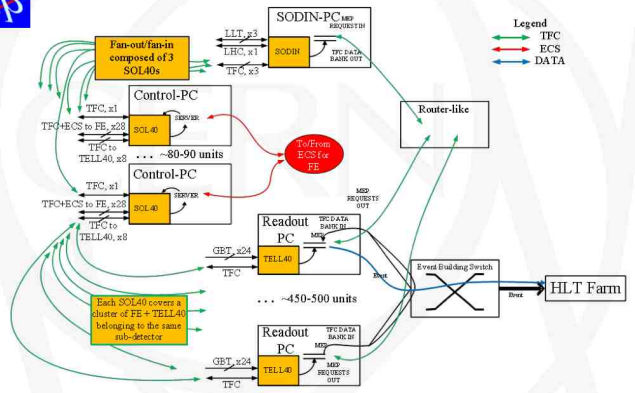

Common electronics board for upgraded readout system: PCIe-Gen3 card fitted on a host PC
- S-ODIN & SOL40 → PCIe card
- TELL40 → PCIe card
- LHC Interfaces → specific PCIe card (or might simply be Ethernet - WhiteRabbit)
- ECS goes via PCIe, TFC to TELL40 goes via fibers

General considerations and conclusions

From TFC point of view:

✓ the choice of PCIe does not pose any problems in terms of hardware, functional and firmware implementation (of TFC). ATCA was already demonstrated to be adequate.
- Additional number of cards in PCIe and additional number of fibers.
- No crates. Distributed system + scalability at its best.
- No need to test backplane if going for PCIe. One solution fits all.
- No need of a motherboard if going for PCIe. Commercial PCs.
- Output bus is flexible and bigger bandwidth for ECS to FE in PCIe.
- PCIe seems more flexible in terms of future development and optimizations.
  - ✓ Fully 40 MHz streamlined system (data center-like).

My personal point of view:

✓ ATCA is a more compact solution, perfect for a triggered system (with LLT)

✓ PCIe is a more distributed and streamlined solution, perfect for a trigger-less system
- *Technology at the service of development (and efficiency)!*

Numbers: ATCA vs PCIe

Most of the numbers from TDRs ...
(in parenthesis is the # of additional SOL40 in PCIe vs ATCA)

	# of Master GBTs	# of TELL40s for data (24 input)	# SOL40s (24 input)	# SOL40s (36 input)
VELO	156	104 (16 input)	11 (+4)	8 (+3)
UT	168	55	10 (+3)	7 (+1)
RICH	1030	77	47 (+4)	31 (+2)
Sci-Fi	678	113	33 (+4)	22 (+3)
CALO	72	53	6 (+3)	4 (+1)
MUON	140	24	7 (+1)	5 (+1)
TOTAL	2244	426	114 (+19)	77 (+11)

Only additional complication is the ACTIVE fan-in/fan-out of TFC and throttle information to TELL40s and FE (> 3 boards needed of 36 input/output)
→ While maintaining partitioning and logical separations

DAQ event-builder

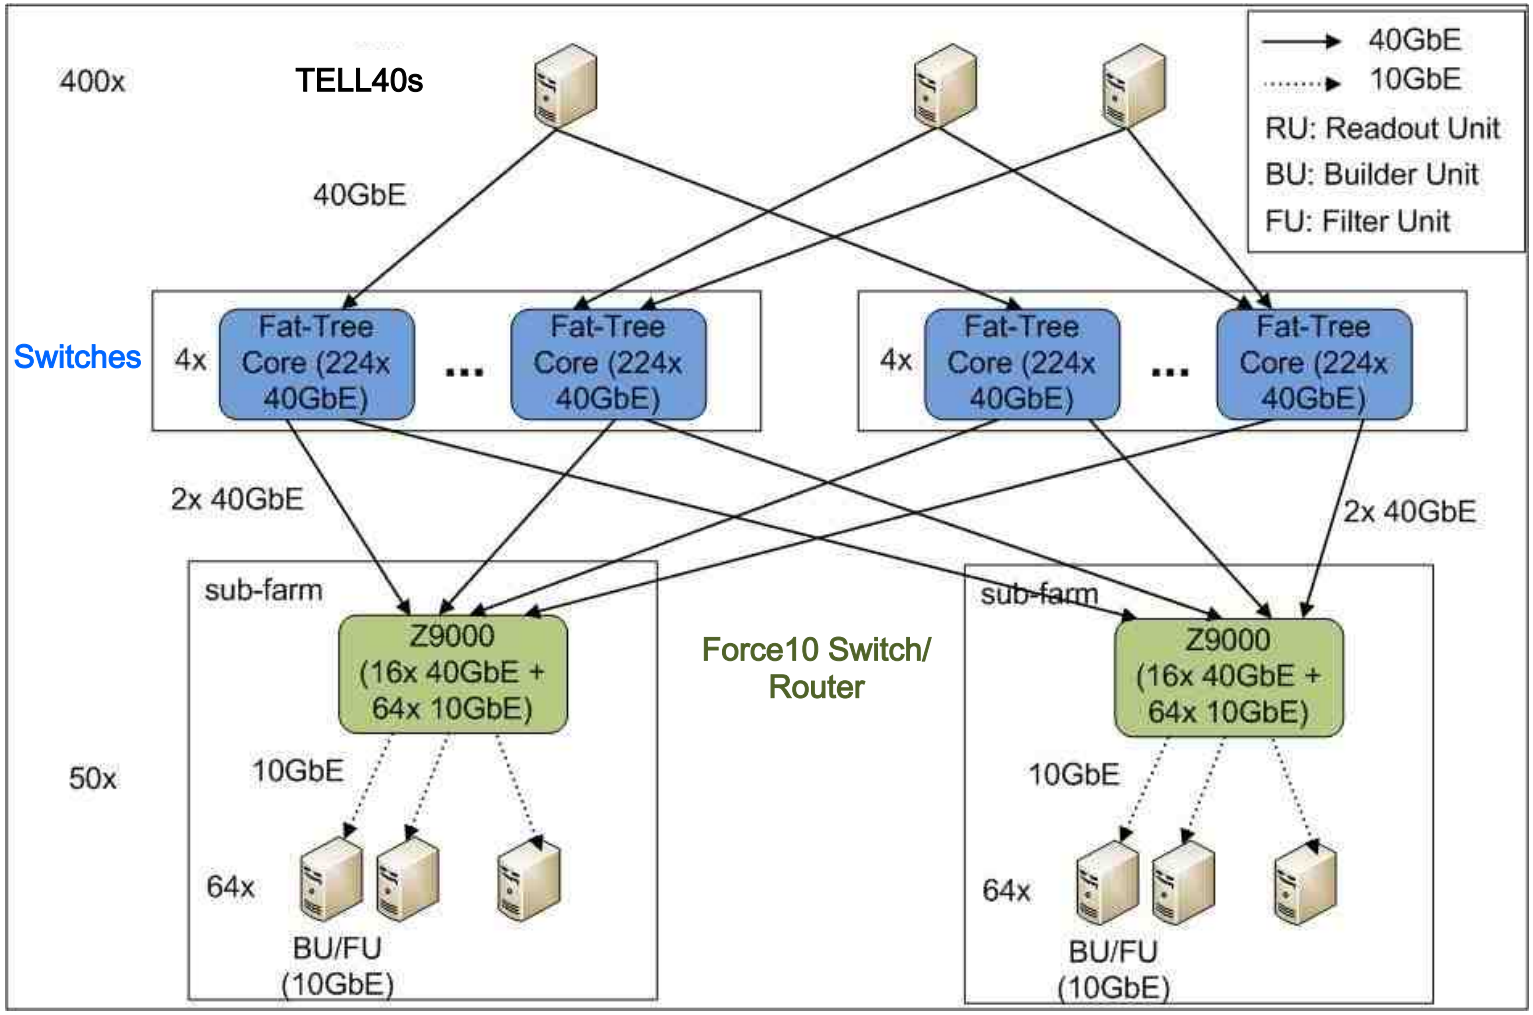

The LHCb upgrade: physics

LHCb upgrade: statistical sensitivity to key observables

Type	Observable	Current precision	LHCb 2018	**Upgrade** (50 fb^{-1})	Theory uncertainty
B_s^0 mixing	$2\beta_s(B_s^0 \to J/\psi\phi)$	0.10 [139]	0.025	0.008	~0.003
	$2\beta_s(B_s^0 \to J/\psi f_0(980))$	0.17 [219]	0.045	0.014	~0.01
	$a_{\rm sl}^s$	6.4×10^{-3} [44]	0.6×10^{-3}	0.2×10^{-3}	0.03×10^{-3}
Gluonic penguins	$2\beta_s^{\rm eff}(B_s^0 \to \phi\phi)$	–	0.17	0.03	0.02
	$2\beta_s^{\rm eff}(B_s^0 \to K^{*0}\overline{K}^{*0})$	–	0.13	0.02	< 0.02
	$2\beta^{\rm eff}(B^0 \to \phi K_S^0)$	0.17 [44]	0.30	0.05	0.02
Right-handed currents	$2\beta_s^{\rm eff}(B_s^0 \to \phi\gamma)$	–	0.09	0.02	<0.01
	$\tau^{\rm eff}(B_s^0 \to \phi\gamma)/\tau_{B_s^0}$	–	5 %	1 %	0.2 %
Electroweak penguins	$S_3(B^0 \to K^{*0}\mu^+\mu^-; 1 < q^2 < 6\ {\rm GeV}^2/c^4)$	0.08 [68]	0.025	0.008	0.02
	$s_0 A_{\rm FB}(B^0 \to K^{*0}\mu^+\mu^-)$	25 % [68]	6 %	2 %	7 %
	$A_{\rm I}(K\mu^+\mu^-; 1 < q^2 < 6\ {\rm GeV}^2/c^4)$	0.25 [77]	0.08	0.025	~0.02
	$\mathcal{B}(B^+ \to \pi^+\mu^+\mu^-)/\mathcal{B}(B^+ \to K^+\mu^+\mu^-)$	25 % [86]	8 %	2.5 %	~10 %
Higgs penguins	$\mathcal{B}(B_s^0 \to \mu^+\mu^-)$	1.5×10^{-9} [13]	0.5×10^{-9}	0.15×10^{-9}	0.3×10^{-9}
	$\mathcal{B}(B^0 \to \mu^+\mu^-)/\mathcal{B}(B_s^0 \to \mu^+\mu^-)$	–	~100 %	~35 %	~5 %
Unitarity triangle angles	$\gamma(B \to D^{(*)}K^{(*)})$	~10–12° [252, 266]	4°	0.9°	negligible
	$\gamma(B_s^0 \to D_s K)$	–	11°	2.0°	negligible
	$\beta(B^0 \to J/\psi K_S^0)$	0.8° [44]	0.6°	0.2°	negligible
Charm *CP* violation	A_Γ	2.3×10^{-3} [44]	0.40×10^{-3}	0.07×10^{-3}	–
	ΔA_{CP}	2.1×10^{-3} [18]	0.65×10^{-3}	0.12×10^{-3}	–

EPJ C (2013) 73:2373

LHCb upgrade: TDRs

- **Letter of Intent for the LHCb Upgrade.**
  CERN-LHCC-2011-001 ; LHCC-I-018. - 2011.
- **Framework TDR for the LHCb Upgrade : Technical Design Report**
  CERN-LHCC-2012-007 ; LHCb-TDR-12. - 2012.
- **LHCb VELO Upgrade Technical Design Report**
  CERN-LHCC-2013-021 ; LHCB-TDR-013. - 2013.
- **LHCb PID Upgrade Technical Design Report**
  CERN-LHCC-2013-022 ; LHCB-TDR-014. - 2013.
- **LHCb Tracker Upgrade Technical Design Report**
  CERN-LHCC-2014-001; LHCB-TDR-015. – 2014
- **LHCb Trigger and Online Upgrade Technical Design Report**
  CERN-LHCC-2014-016 ; LHCB-TDR-016. - 2014.