The LHCb upgrade

U. Marconi, INFN Bologna

IFD2014

Trento, March 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;
  - being sensitive to new particles above the TeV scale not accessible to direct searches.
- 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 (MFV).

The LHCb Detector

Experiment optimized for beauty and charm physics:
– efficient trigger for hadronic and leptonic modes
– precision tracking and vertexing (mass, proper time)
– excellent particle identification

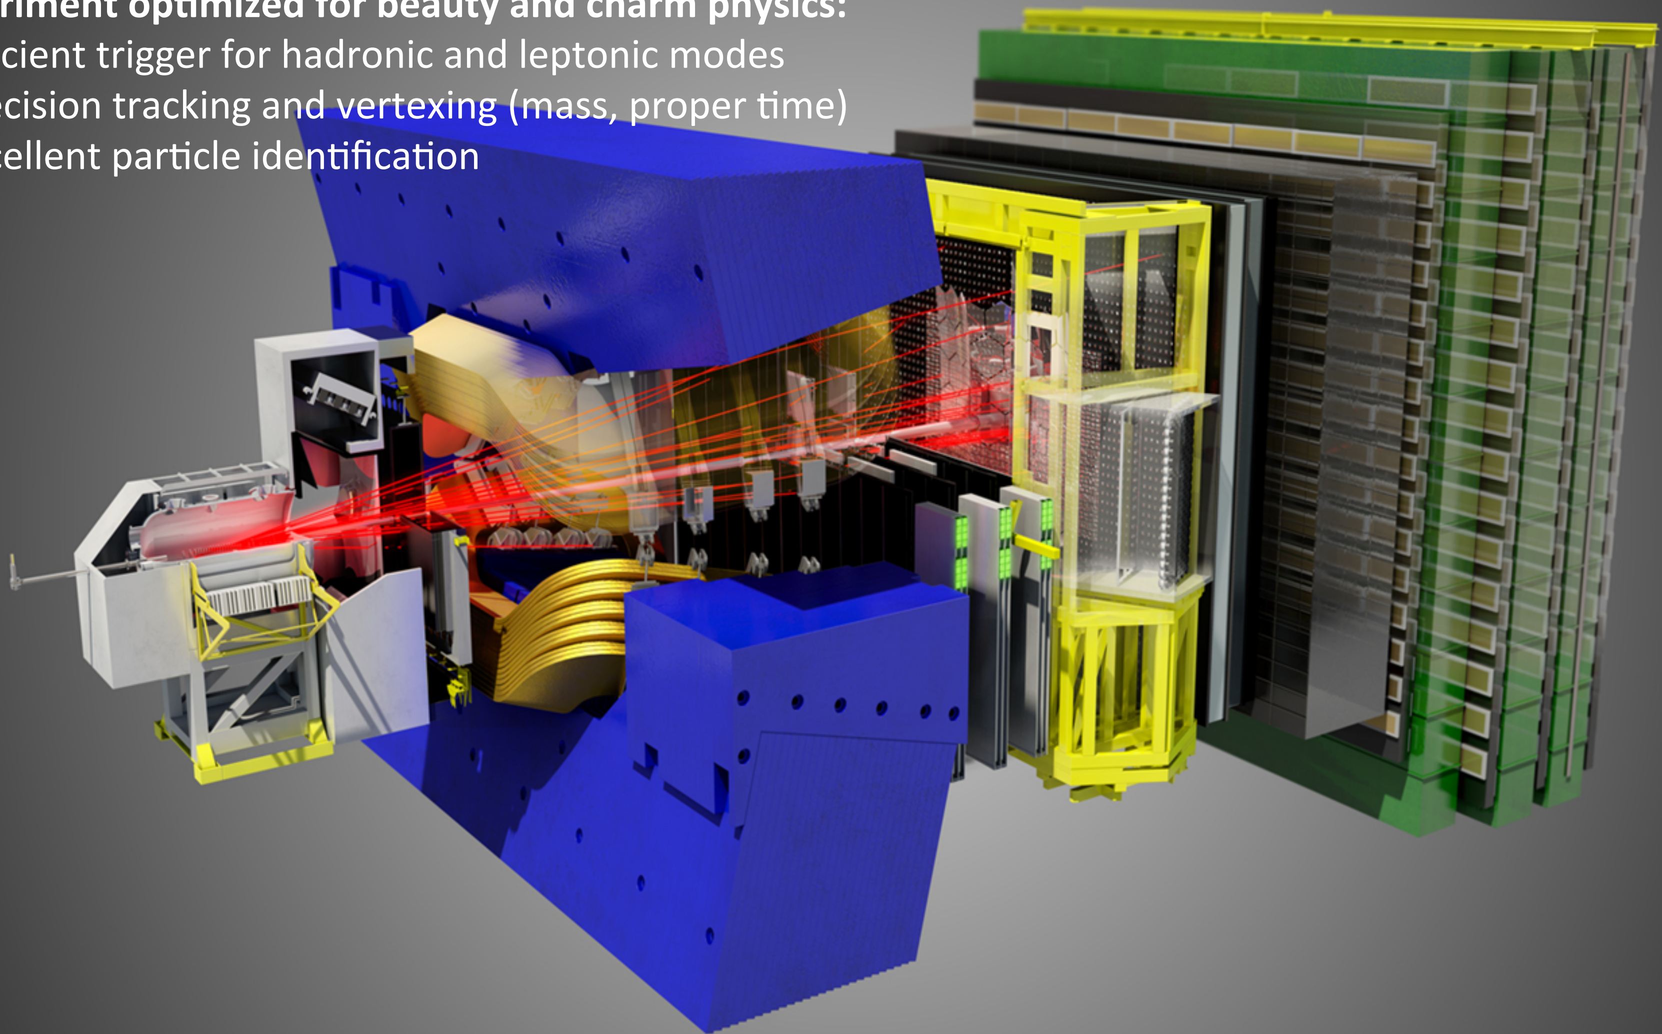

The LHCb detector

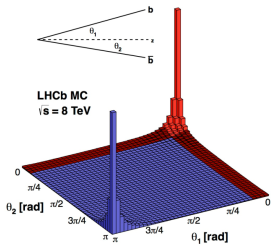

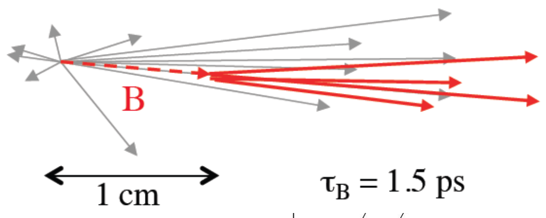

- Decay time resolution **Δ**t: 30-50 fs
- **Δ**p/p = 0.4-0.6%
- **Δ**m = 10-20 MeV/c^2
- Muon ID: **ε**(μ/μ) = 95%, mis-ID **ε(π/**μ) ~ 1%
- RICH ID **π**/K: **ε**(K/K) = 95%, mis-ID **ε(K/π)** ~ 5%
- CALO ID: : **ε**(e/e)=90%, mis-ID ~ 5%

B field polarity can be reversed

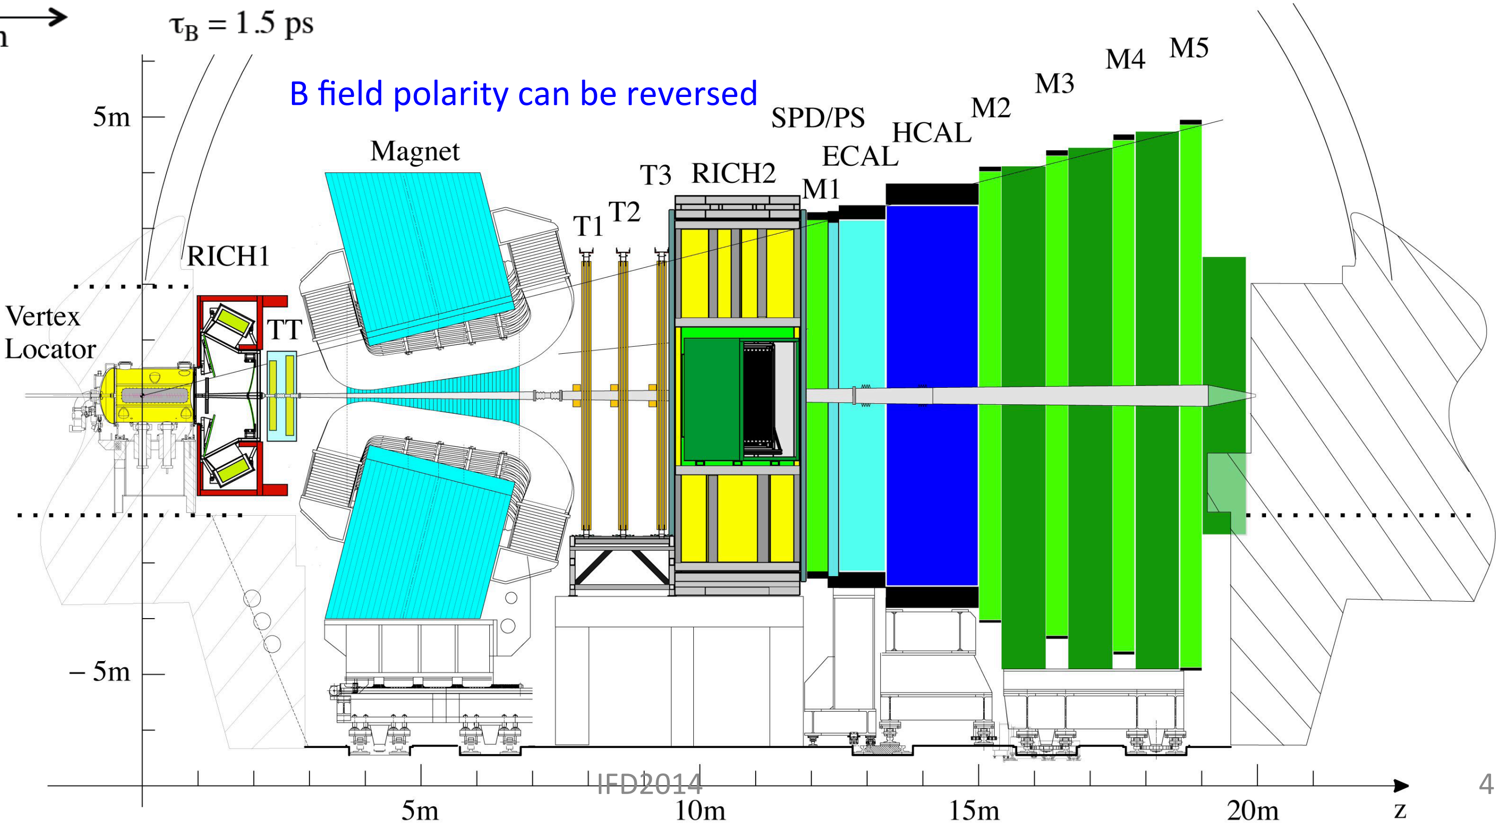

"The LHCb Detector at CERN"; The LHCb Collaboration, JINST 3, S08005 (2008)

LHCb events

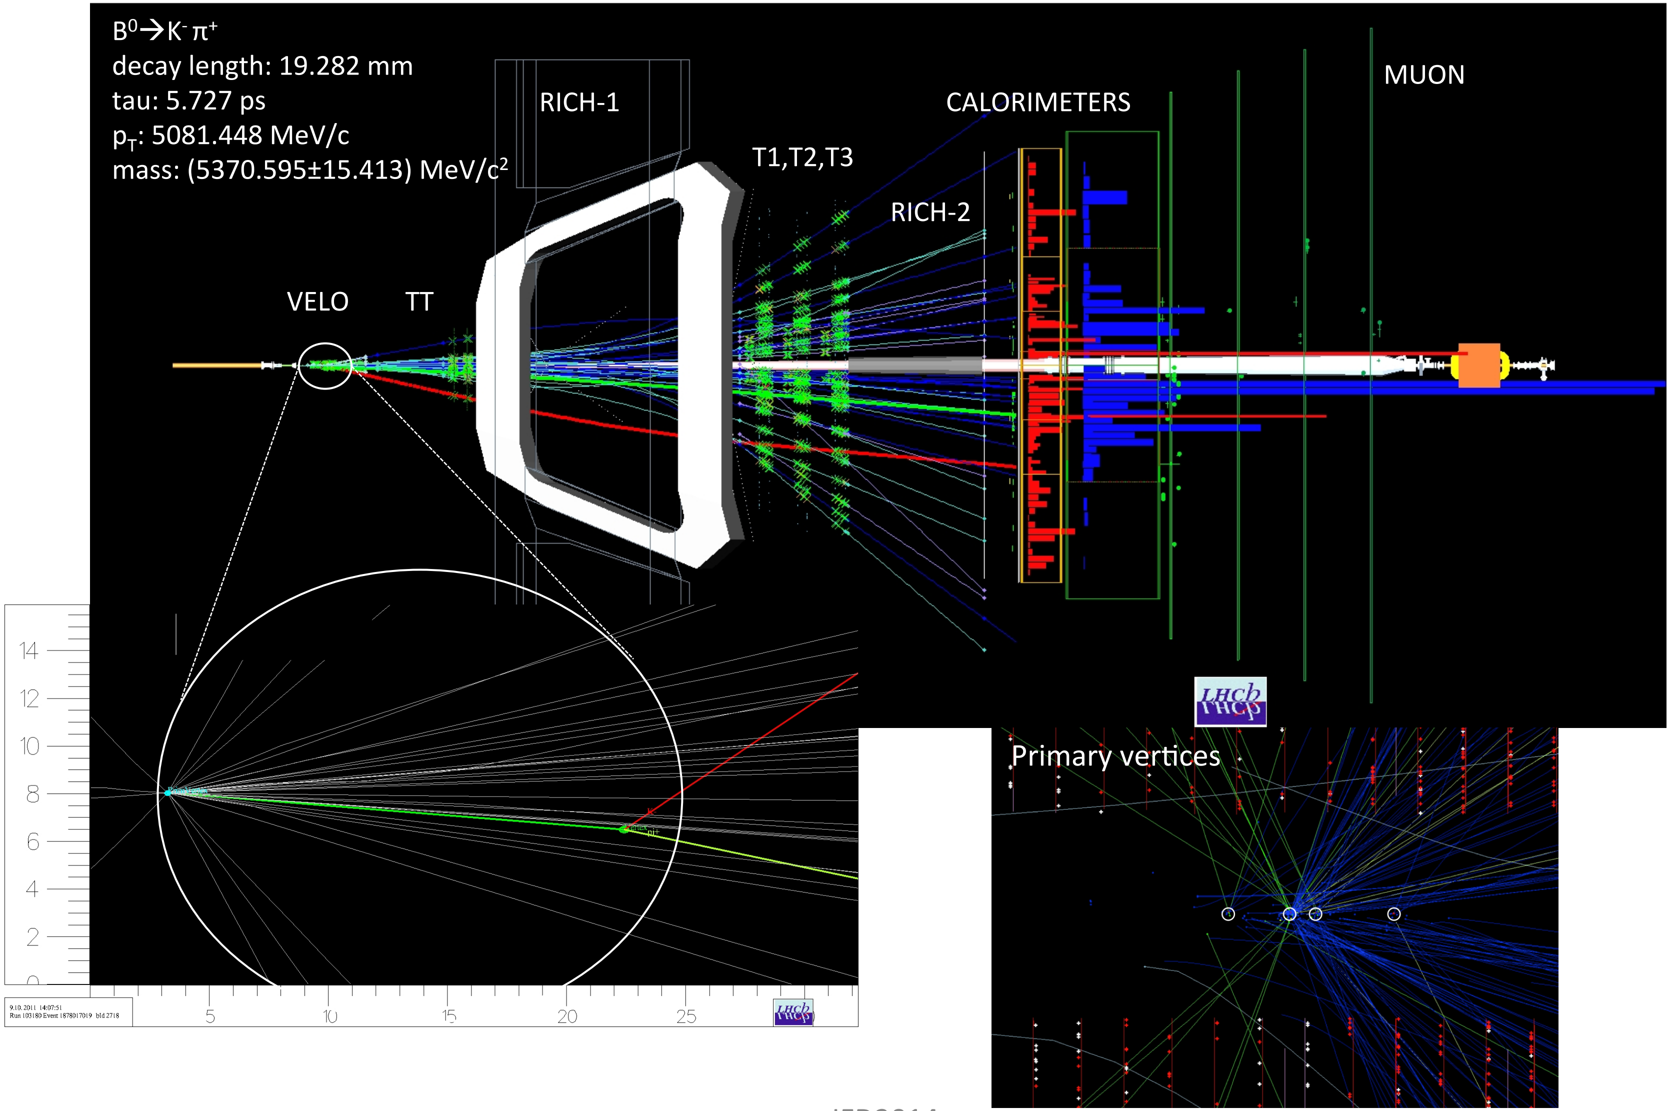

$B^0 \rightarrow K^- \pi^+$
decay length: 19.282 mm
tau: 5.727 ps
p_T: 5081.448 MeV/c
mass: (5370.595±15.413) MeV/c^2
RICH-1
T1,T2,T3
CALORIMETERS
MUON
RICH-2
VELO
TT
14
12
10
8
6
4
2
0
5
10
15
20
25
Primary vertices

Instantaneous Luminosity

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

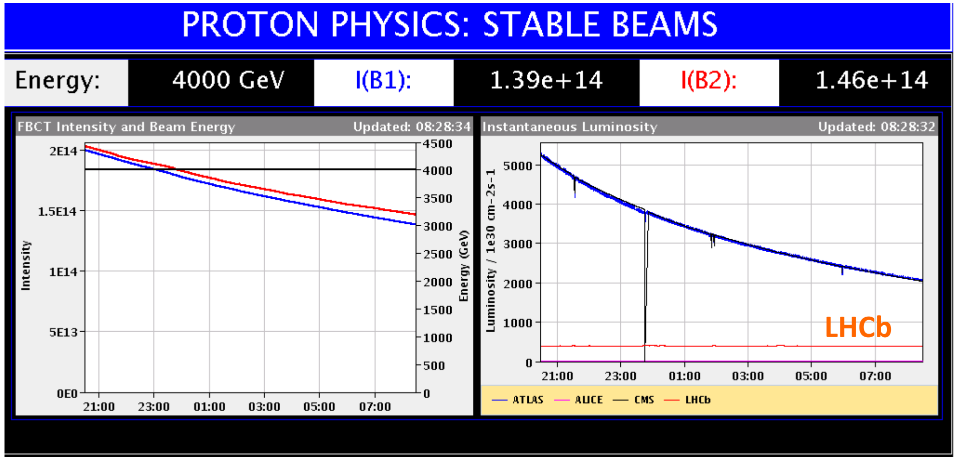

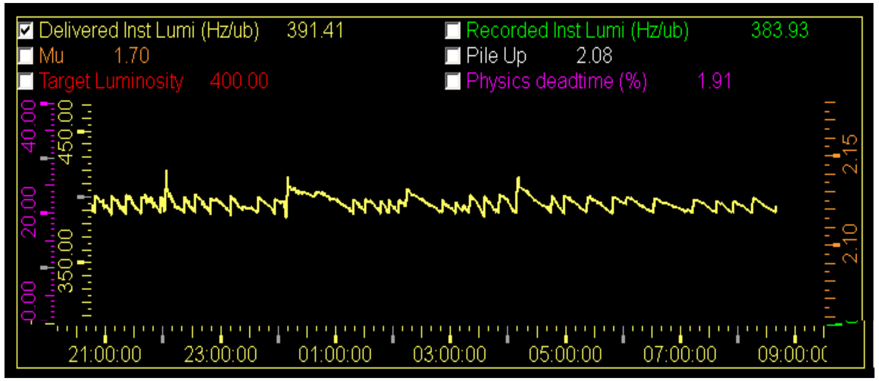

- LHCb was designed to operate with a single collision per bunch crossing, running at a instantaneous luminosity of **2. × 10^{32} cm^{-2} s^{-1}** (assuming about **2700** circulating bunches).
  - At the time of design there were worries about possible ambiguities in assigning the B decay vertex to 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 **μ > 2.5** **(μ: average n. of visible interactions)**

The LHCb Trigger

- The **Level-0 trigger** based on the signals from ECAL, HCAL and MUON detectors read at 40 MHz, operates on custom electronics, with a **maximum output rate limited to 1.1 MHz**.
  - Fully pipelined, constant latency of about 4 µs.
  - Bandwidth to the HLT ~ 4 Tb/s, GOL serdes, optical links.
  - High p_T muon (1.4 GeV) or di-muon.
  - High p_T local cluster in HCAL (3.5 GeV) or ECAL (2.5 GeV)
  - **25%** of the events are **deferred**: temporarily stored on disk and processed with the HLT farm during the inter-fills.
- **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**.

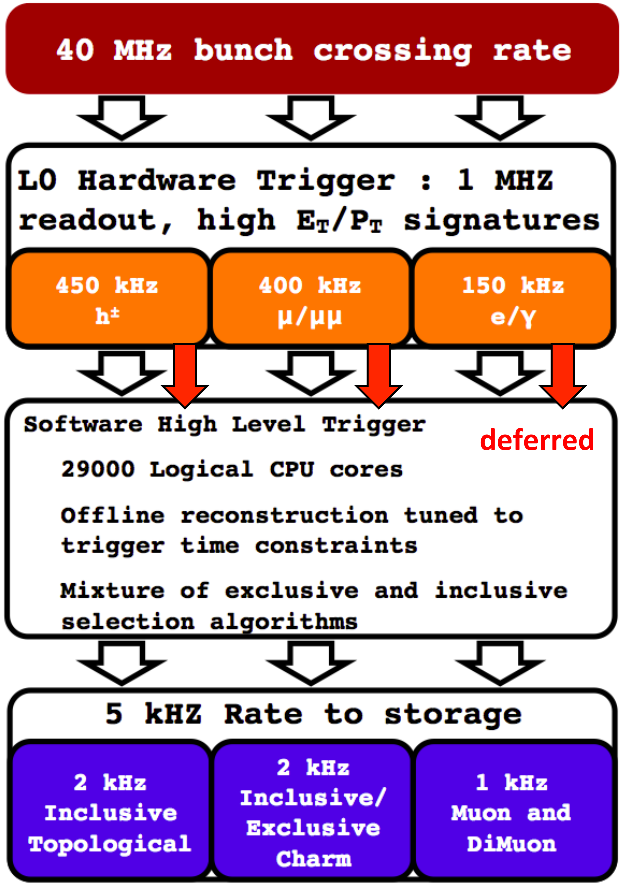

The 1MHz L0 rate limitation

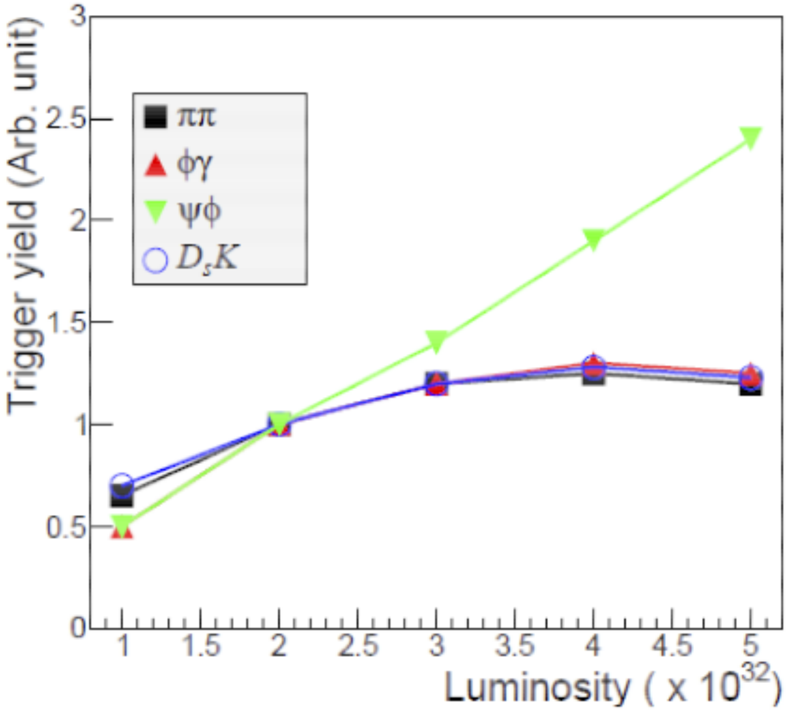

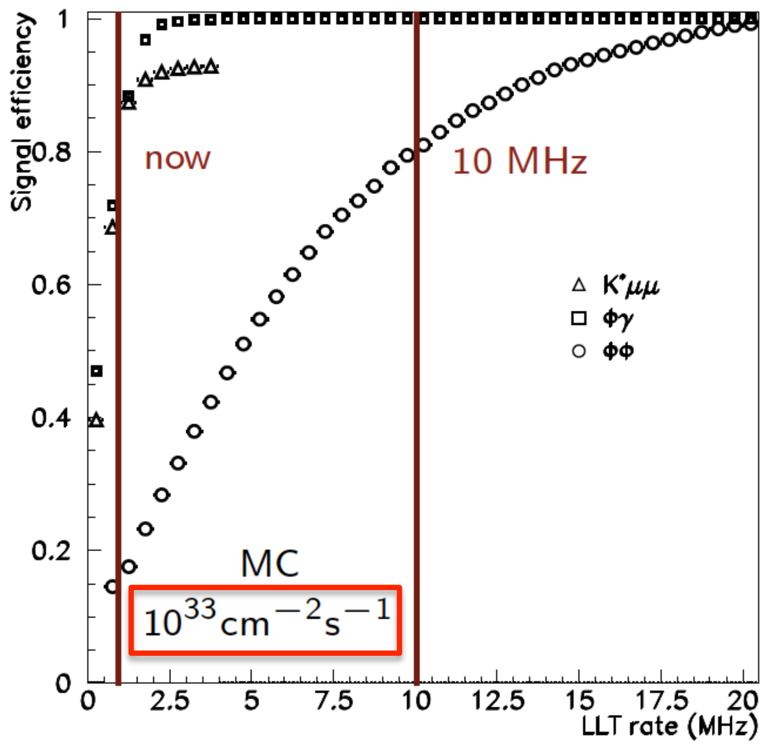

"Framework TDR for the LHCb Upgrade: Technical Design Report", The LHCb Collaboration,CERN/LHCC-2012-007, LHCb-TDR-12 - 2012

- 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.

The LHCb upgrade

- Readout the whole detector at 40 MHz.
- **Trigger-less data acquisition system**.
  - Use a Low Level Trigger as a throttle mechanism, while progressively increasing the power of the event filter farm to run the HLT up to 40 MHz.
- 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.

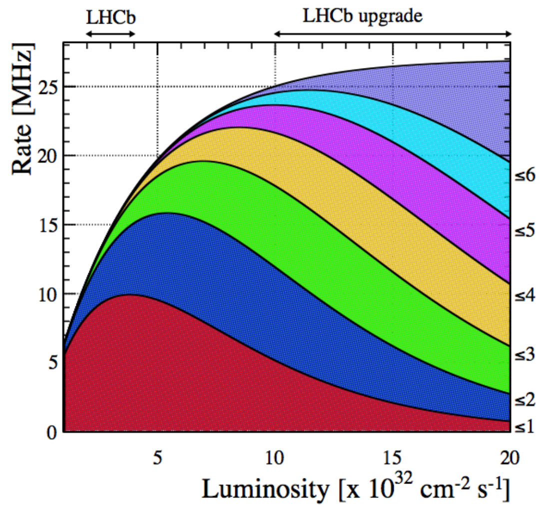

Mean visible interactions per crossing

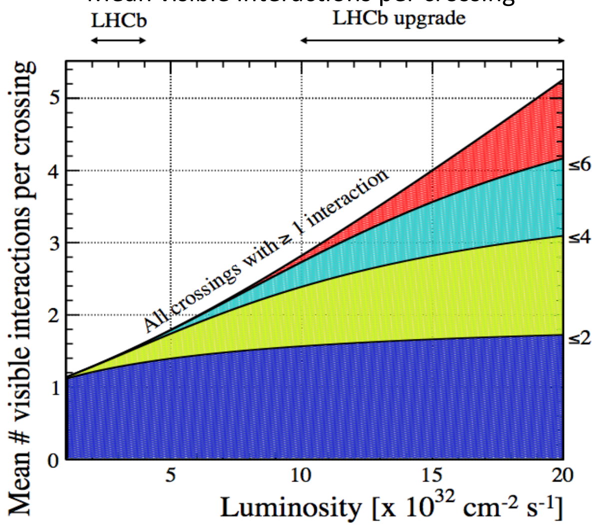

LHCb upgrade: consequences

- The detector front-end electronics has to be entirely rebuilt, because of the output rate requirements.
  - No more buffering in the front-end electronics boards.
  - Data need to be zero suppressed and formatted before transmission to optimize the number of required links
  - A lot more optical links to get the required bandwidth, needed to transfer data from the front-end to the read-out boars at 40 MHz.
- 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}

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
- Online and Trigger TDR in preparation.

LHCb DAQ today

Push-protocol with centralized flow-control

Readout boards: 313 x TELL1

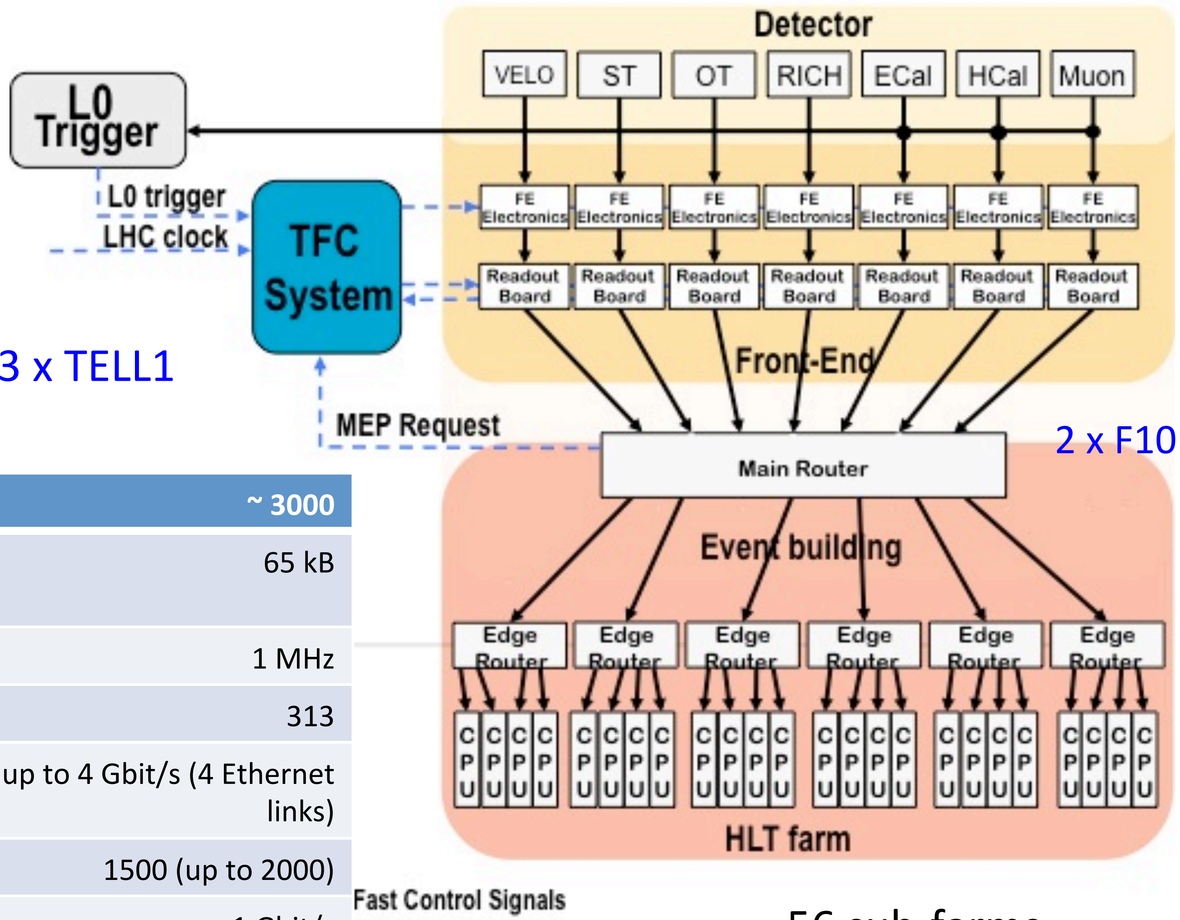

2 x F10 E1200i

56 sub-farms

# 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

LHCb upgrade: DAQ

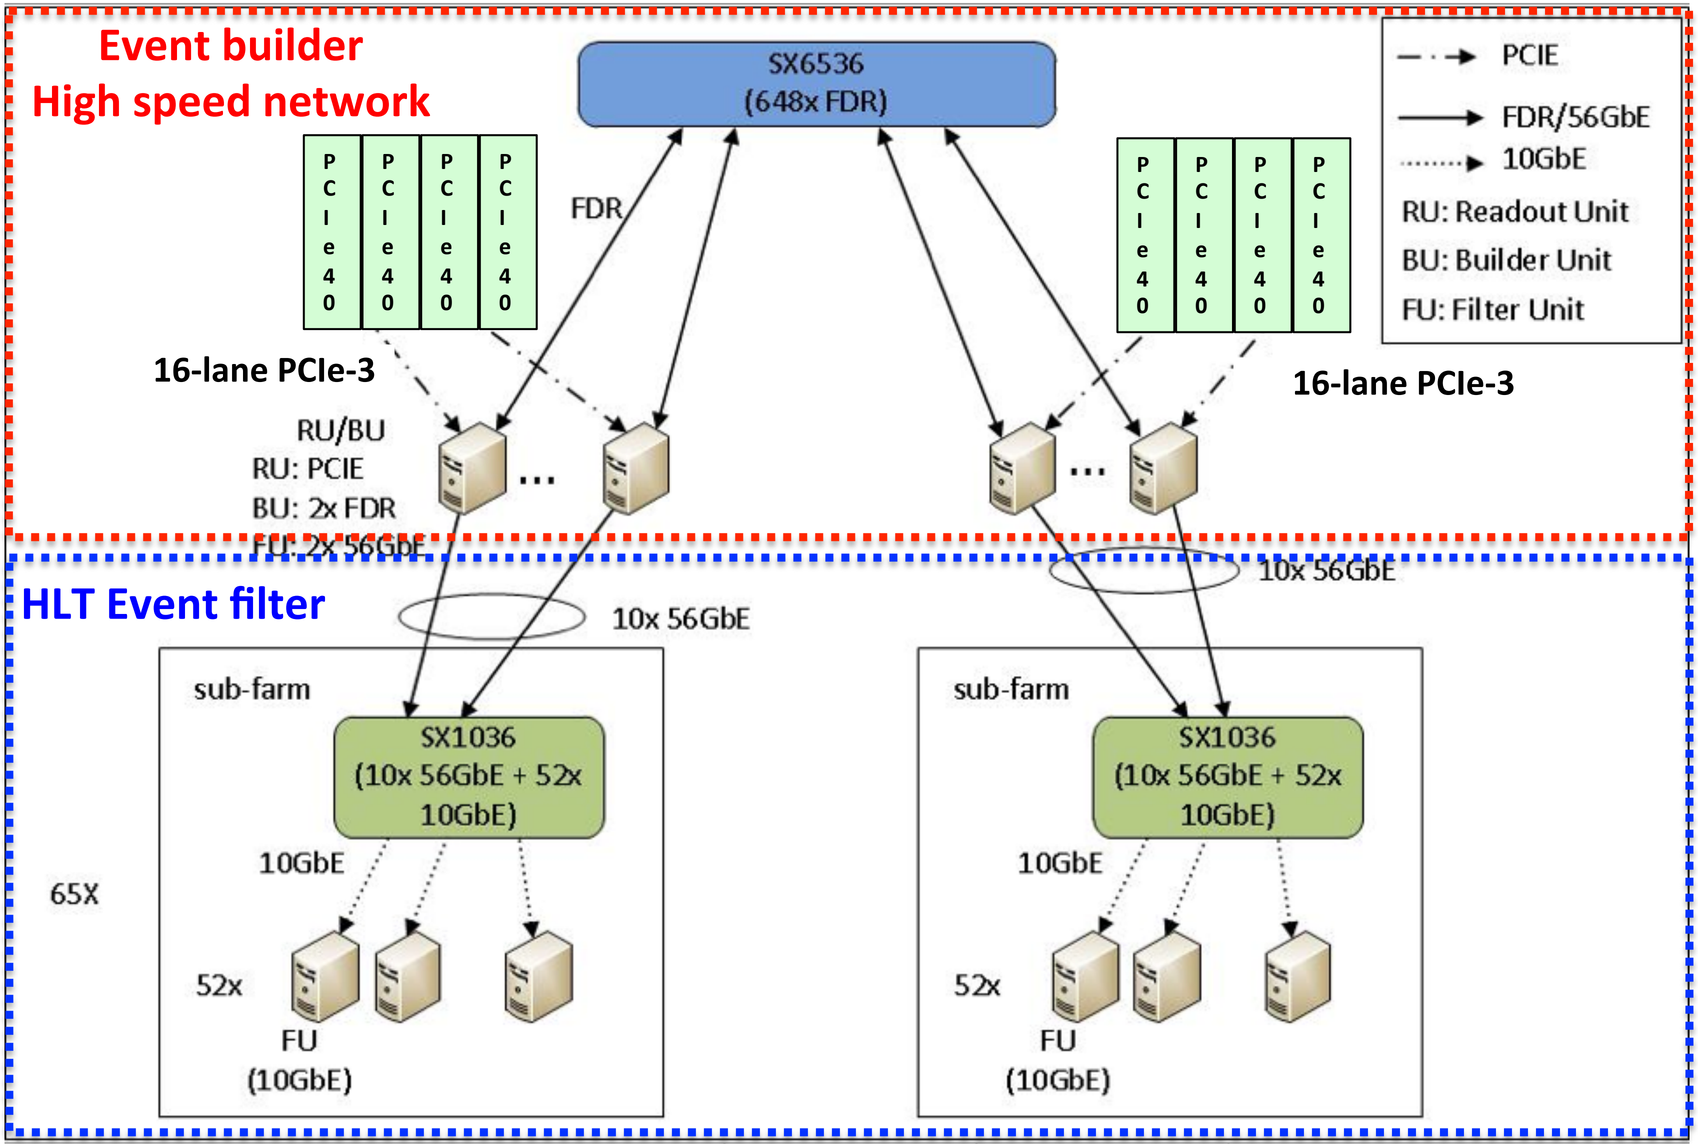

U. Marconi et al. The PCIe Gen3 readout for the LHCb upgrade, CHEP2013

PCIe Gen3 based readout

- A main FPGA manages the input streams and transmits data to the event-builder server using PCIe Gen3.
- PCIe Gen3 throughput: 16-lane × 8 Gb/s/lane = 128 Gb/s
- The readout version of the board uses two de-serializers.

24 optical links from the FEE

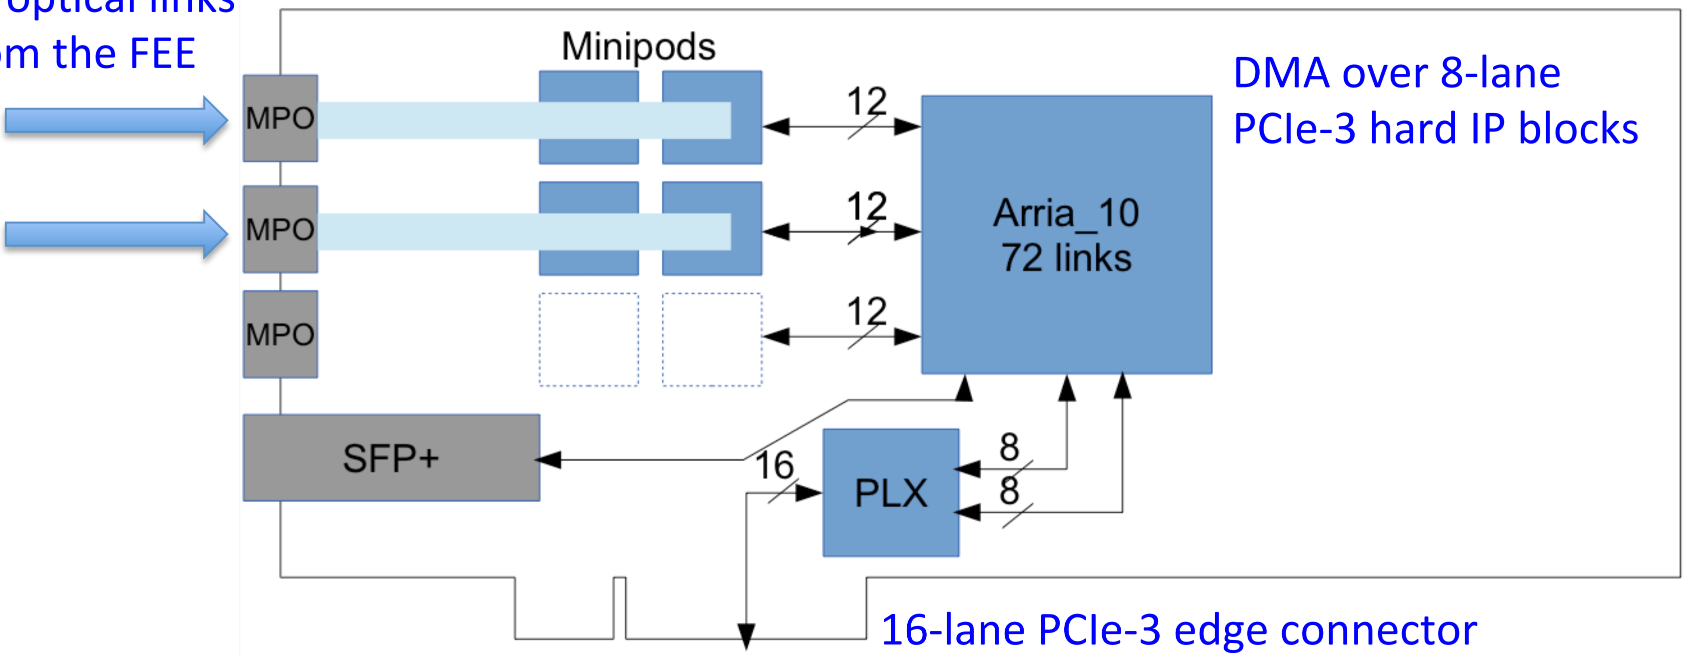

DMA over 8-lane PCIe-3 hard IP blocks

16-lane PCIe-3 edge connector to the motherboard of the host PC

Event builder node

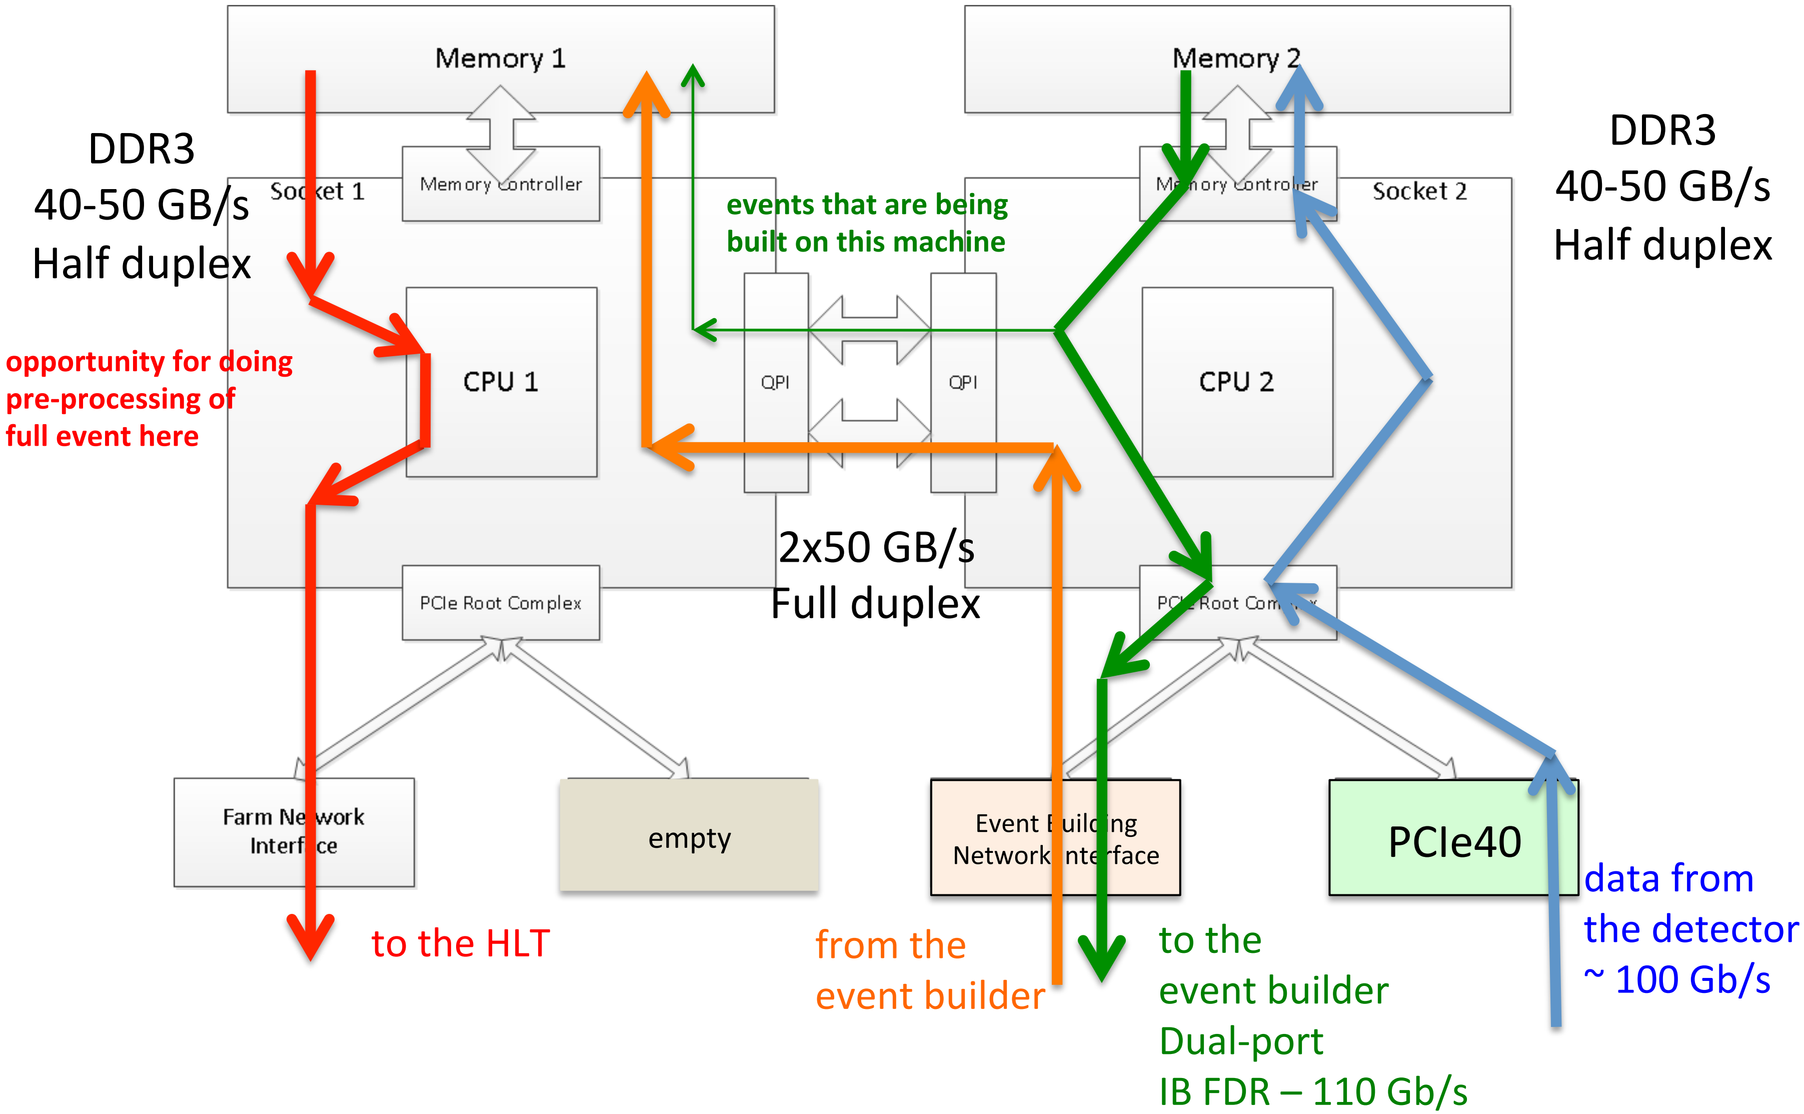

LHCb upgrade: HLT farm

- 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: **$4.\times 10^5$**
- Number of core/CPU available in 2018: ~ **200**
  - Intel tick-tock plan: 7nm technology available by 2018-19, the n. of core accordingly scales as $12. \times (32\ \mathrm{nm}/\ 7\ \mathrm{nm})^2 = 250$
- Number of computing nodes required: ~ **1000**

TPU: HLT assisted tracking

- "A specialized processor for track reconstruction at the LHC crossing rate". Presented by G. Punzi, INSTR 2014
- 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."

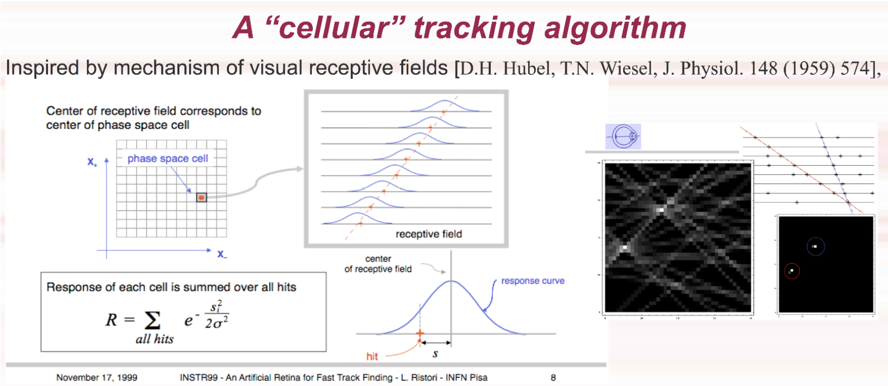

Implementation

- 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.

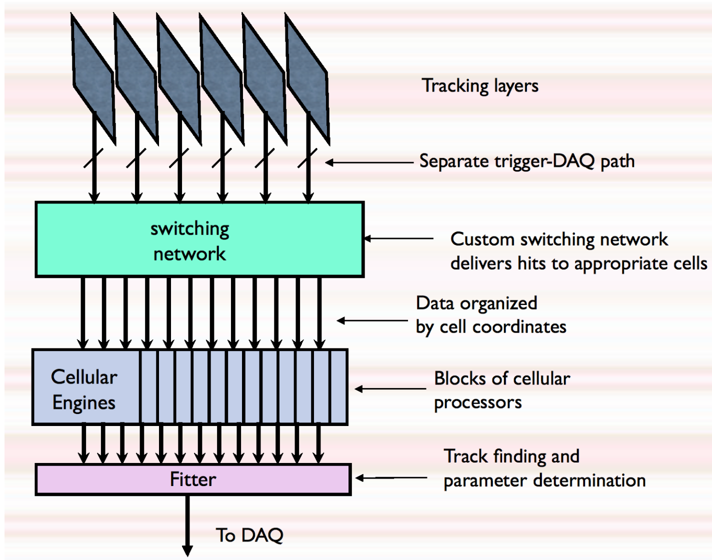

Upgrade of the tracking system

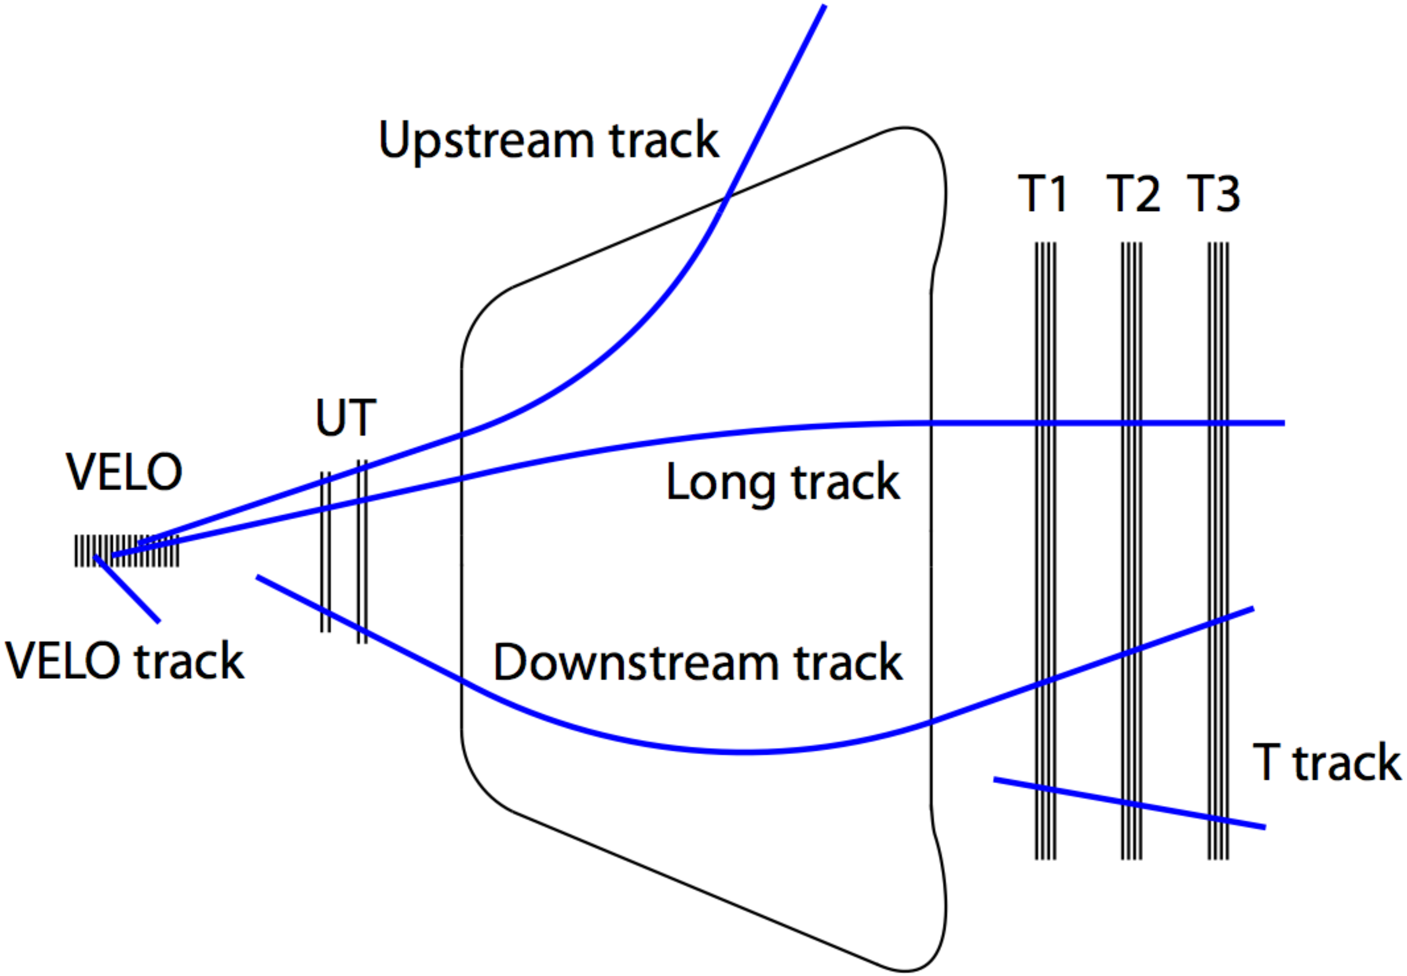

The current VELO

- 21 stations of silicon micro-strip sensors, located nearby the interaction region.
  - R and ϕ geometry modules
  - 44 mm external radius
- Left and right halves of the modules can be moved inward/outward the beam lines: modules open and close at each fill. Distance from the beam line 5.5 mm
  - Position centred around the current beam position. Modules do not move during a fill.
  - Positions reproducible within 5 μm: measured as the average distance between right-side PV and left-side PV.
- Primary (beam) and secondary vacuum are separated by a thin Al box ("RF foil").

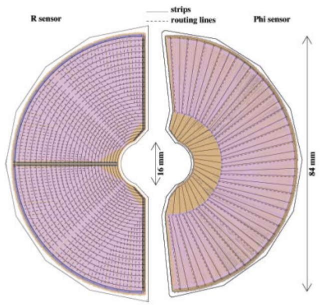

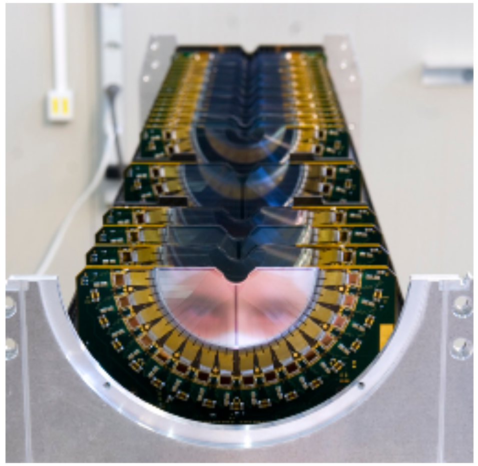

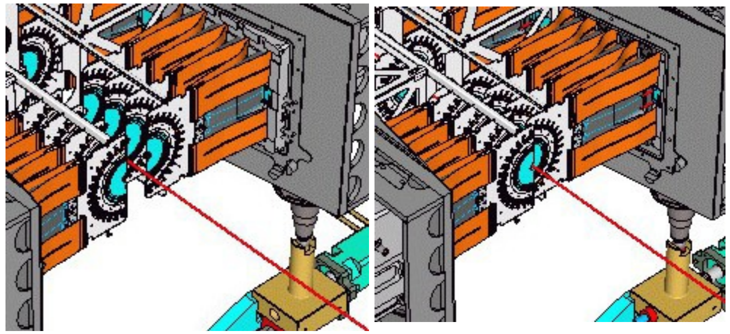

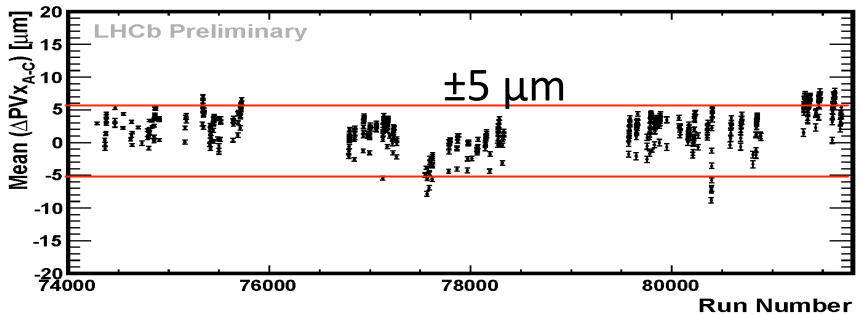

The VELO upgrade

http://cds.cern.ch/record/1624070/files/LHCB-TDR-013.pdf

- Silicon **pixels** with a surface area of 55x55 μm^2
- **Lower material budget**:

Sensor thickness: 300 μm → 200 μm

RF Aluminium foil thickness: 300 → ≤ 250 μm,

- **Enlarge acceptance**

Inner foil radius: 5.5 mm → 3.5 mm

first sensitive pixel edge at a radius of 5.1 mm, which is about the limit for the pixel chip as determined by the data rate from detailed simulations

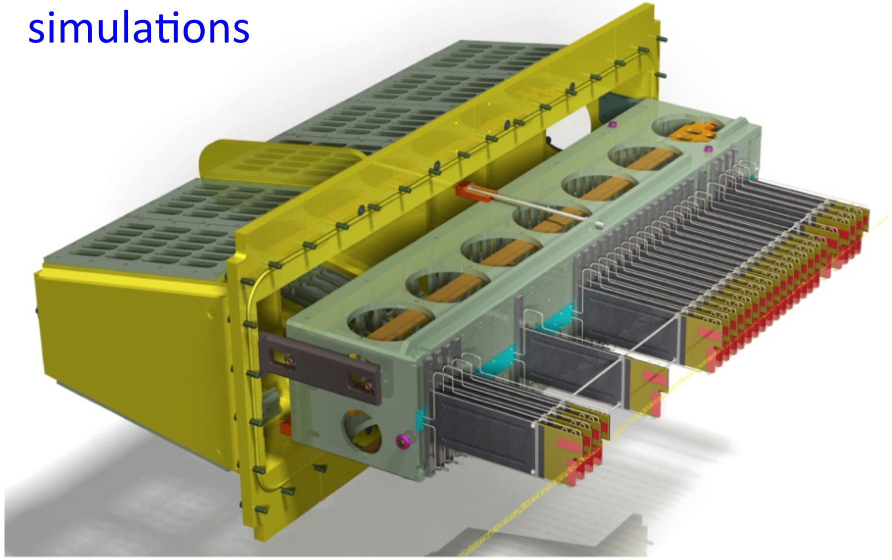

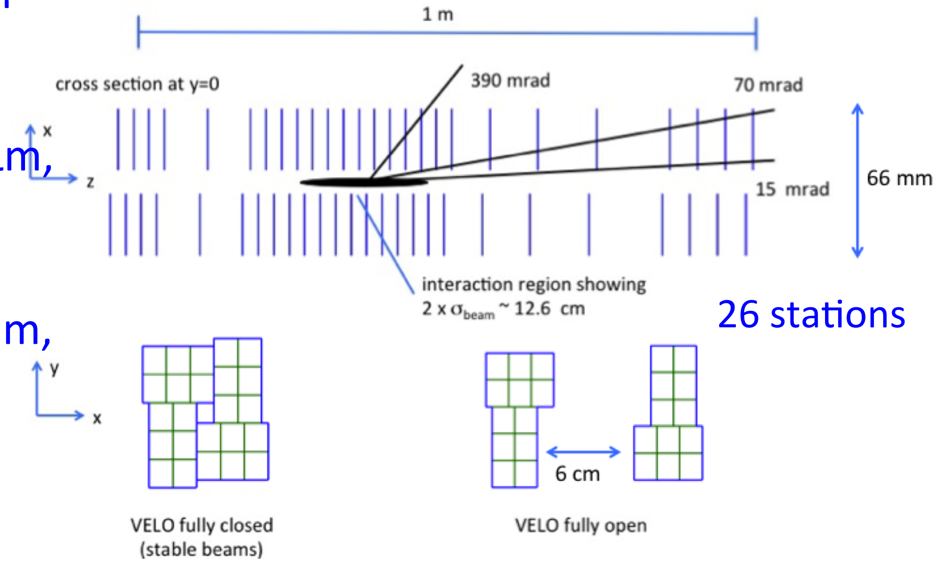

26 stations

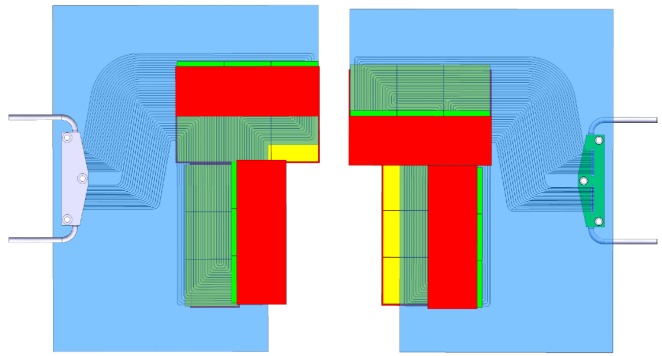

Figure 39: Front and rear side of a module containing two 3 × 1 tiles on either side. Sensors are drawn in red and the microchannel cooling substrate in blue. The bottom sides of the ASICs are drawn in yellow, while the area where the bond pads are located is made green. The overall horizontal and vertical dimensions of the substrate are 80 and 104 mm, respectively.

The VELO upgrade

PV resolution in x as function of the number of reconstructed tracks in the vertex

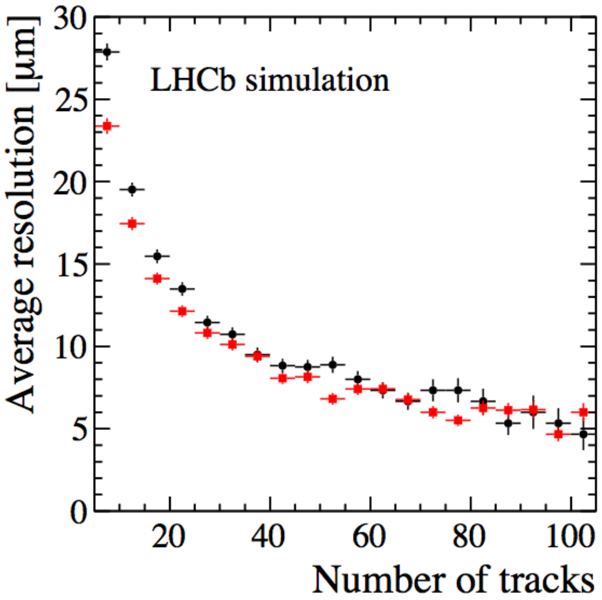

IP resolution in x
(resolutions in x and y are similar)

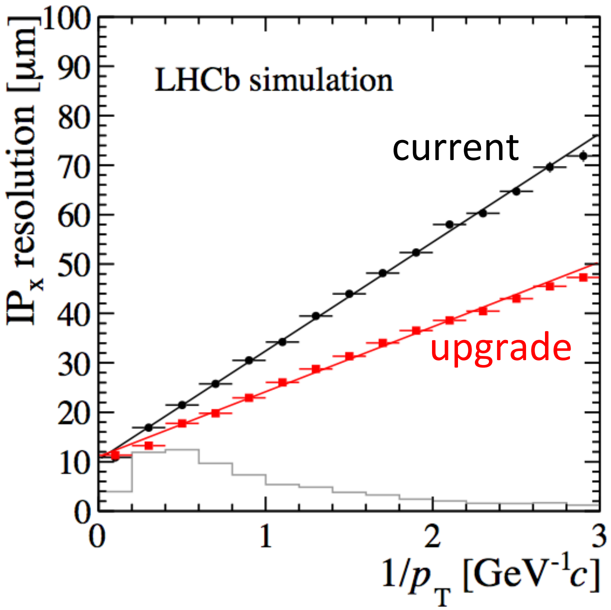

Tracking efficiency

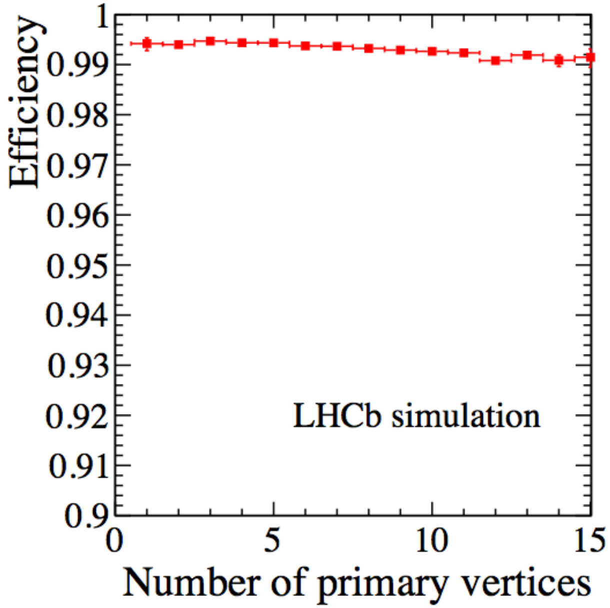

ν: pp interactions per crossing
μ: visible interactions

	Existing VELO [%] $\nu = 2$	Existing VELO [%] $\nu = 7.6$	Upgraded VELO [%] $\nu = 7.6$
Ghost rate	6.2	25.0	2.5
Clone rate	0.7	0.9	1.0
Reconstruction efficiency			
VELO, $p > 5\,\mathrm{GeV}/c$	95.0	92.7	98.9
long	97.9	93.7	99.4
long, $p > 5\,\mathrm{GeV}/c$	98.6	95.7	99.6
b-hadron daughters	99.0	95.4	99.6
b-hadron daughters, $p > 5\,\mathrm{GeV}/c$	99.1	96.6	99.8

strip pixel

Ghost rate

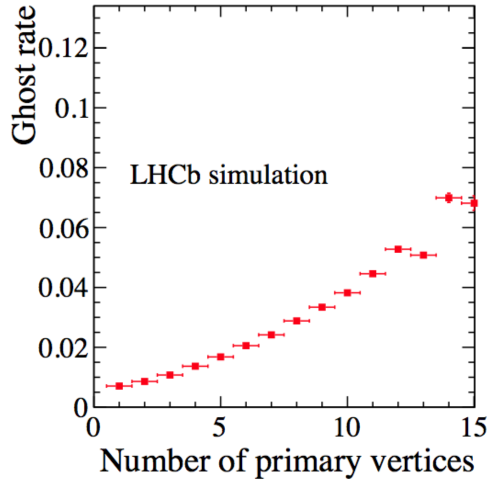

Total output data rate: 3 Tb/s

Upstream Tracker

Silicon Strip detector

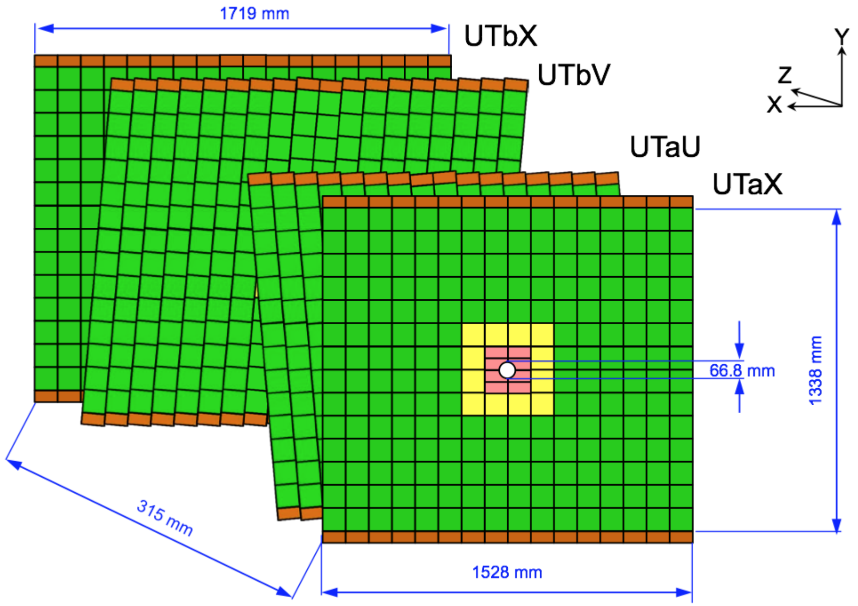

Property	Sensors B,(C,D)	Sensors A
Technology	n^{+}-in-p	p^{+}-in-n
Thickness	250 µm	250 µm
Physical dimensions	98 mm X 98 (49) mm	98 mm X 98 mm
Length of read-out strip	98 (49) mm	98 mm
Number of read-out strips	1024	512
Read-Out strip pitch	95 µm	190 µm
Sensor number (needed)	48 (16,16)	888

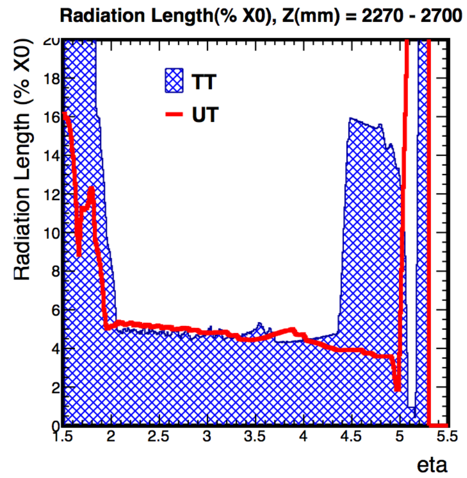

- The first and the last planes have vertical strips, whereas the middle two are at ±5◦.
- The current TT read-out strip geometries (length of about 10 cm and a strip pitch of 183 µm) will lead to unacceptably high occupancies under the foreseen running conditions.

Upstream Tracker

Downstream Tracks

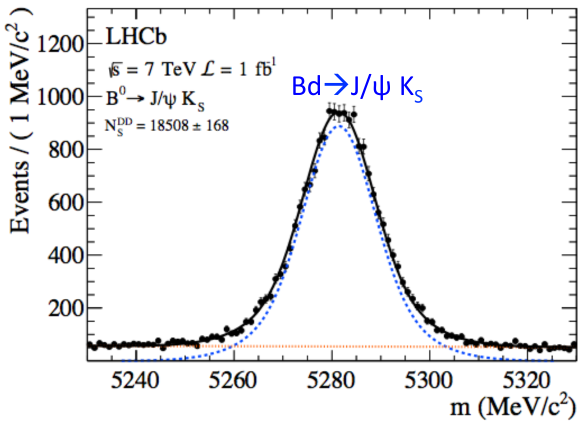

Long Tracks

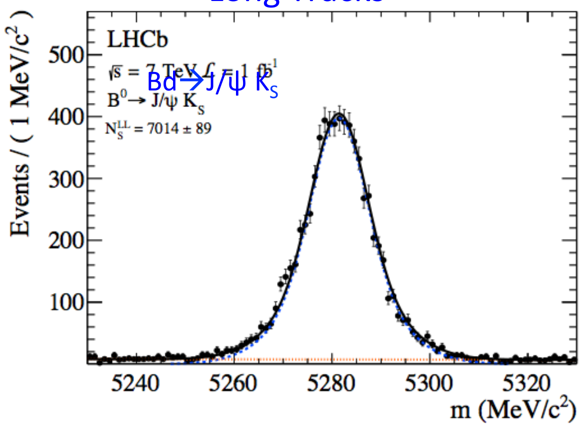

Trigger: p_T

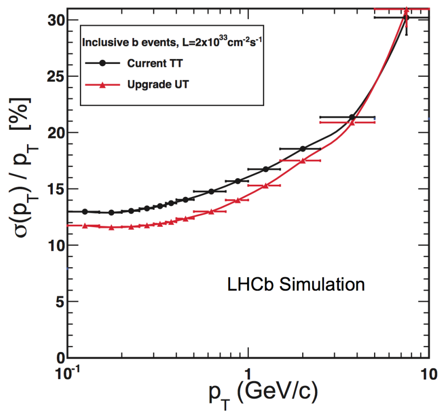

M. Ullan et al., Embedded pitch adapters for the ATLAS Tracker Upgrade, Nucl. Instrum. Meth. A732 (2013) 178.

- Most of the B→J/ψ K_S events, 73%, are reconstructed from decays **downstream** using the UT and the tracker. Momentum resolution of long tracks is significantly improved, by ~25%, if the tracks have UT hits.
- **For triggering** it is possible to use the stray magnetic field of 0.02 T between the VELO and the UT to measure the track momentum $\sigma(p_T)/p_T \approx 15\%$, good enough to measure the sign of the charge.

Scintillator Fibre Tracker

- The SCSF-78MJ, 250 μm diameter scintillating fibre, by Kuraray, produces a sufficient light yield and has a long enough attenuation length for this detector.
- Radiation damage to the fibre material results in degraded transmission, but it is expected to result in less than 40% loss of signal from the worst regions near the beam-pipe after 50 fb^{-1}
- The fast signal response and recovery time of the detector makes the SiPM particularly suitable for the fast interaction rate.

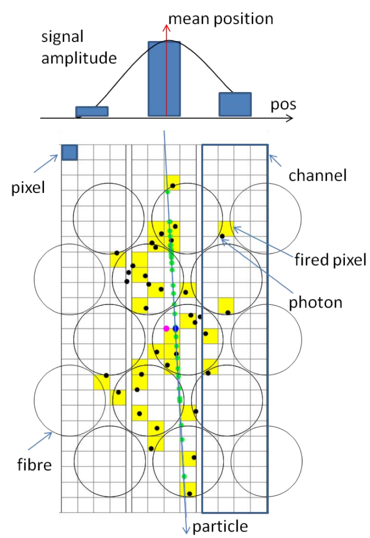

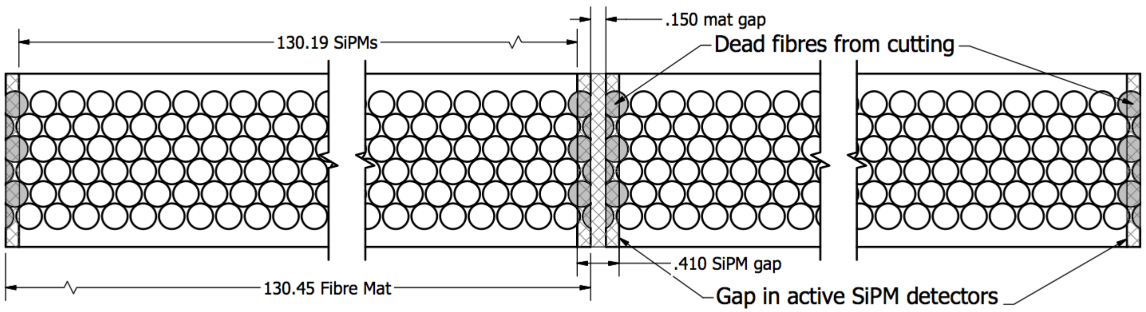

Tracking performance

Bs → ϕϕ events

	Current LHCb [%]	Upgrade LHCb [%]	
	$\nu = 2$	$\nu = 3.8$	$\nu = 7.6$
Ghost rate	13.1	14.7	25.5
Reconstruction efficiency			
long	90.9	86.9	84.5
long, $p > 5\,\mathrm{GeV}/c$	95.4	92.9	91.5
b-hadron daughters	93.9	91.9	90.6
b-hadron daughters, $p > 5\,\mathrm{GeV}/c$	96.1	95.1	94.2

Clear degradation of the performance as function of the number of primary vertices is observed

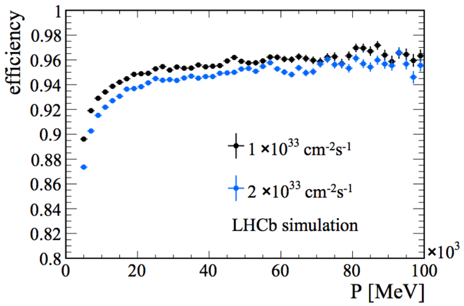

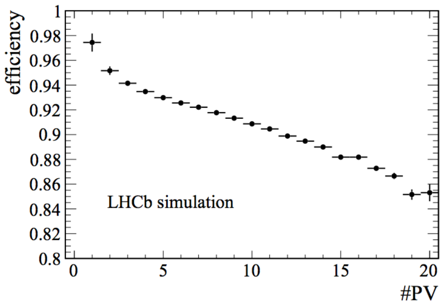

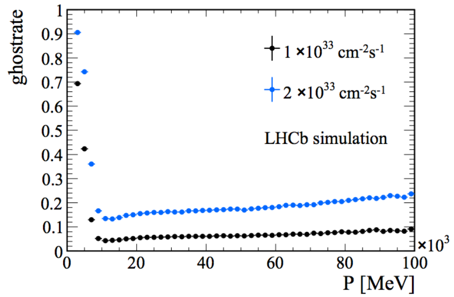

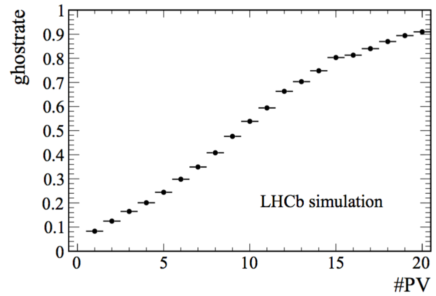

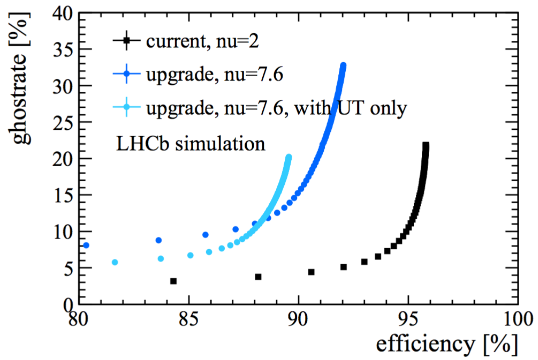

Different points varying the χ^2 cuts after the Kalman filter

To achieve similar ghost rates for the upgrade experiment at ν = 7.6 as in the current experiment at ν = 2, a drop in efficiency of 5% would be required

PID performance

- Photon detectors must be replaced, since the current hybrid photon detectors (HPDs) have their 1 MHz read-out electronics encapsulated within the tube.
- The overall structure of both RICH detectors will remain unchanged.
- The optical layout of RICH 1 has to be modified to reduce the otherwise prohibitively large hit occupancy in the central region of the detector.
  - Increasing the the focal length of the spherical mirrors halving the occupancy.
  - The gain is also achieved by removing the Aerogel radiator material (~3.5% of X_0). K (π) threshold in C_4F_{10} is 9.3 (2.6) GeV/c .
- It is proposed to replace the HPDs with commercial multi-anode photomultipliers (MaPMT) with external readout electronics (1920 in RICH 1 and 2560 in RICH 2).

The customized ASIC readout chip, "CLARO": comprised of an analogue pulse shaper amplifier and a binary discriminator.

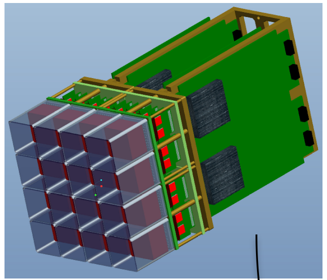

R11265 MaPMT from Hamamatsu.

PDM size of 116 × 116 mm^2

PDM matrix is 6 columns × 12 rows.

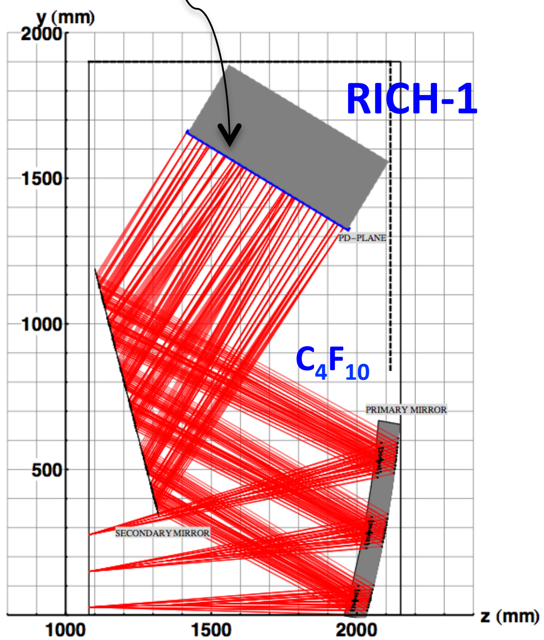

PID performance

Running the full simulation and reconstruction chain in $B_s \to \phi\phi$

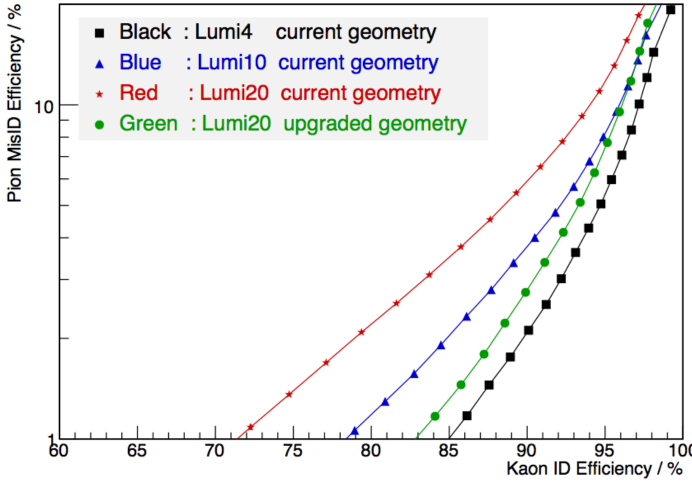

The PID performance of the current geometrical layout and the upgraded superimposed

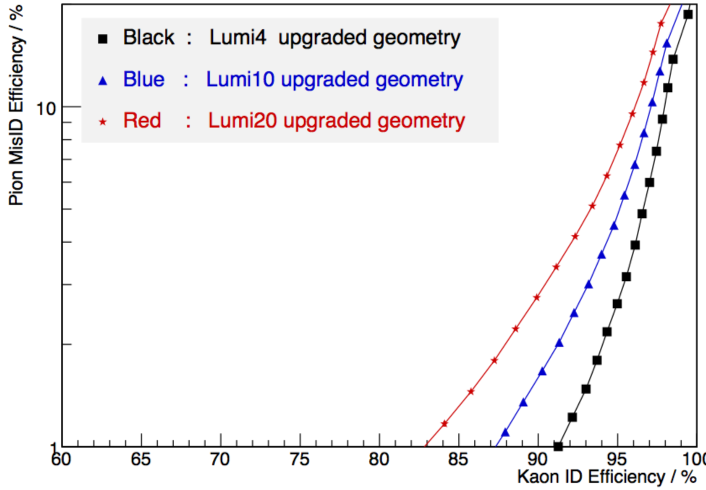

The PID performance of the upgraded geometrical layout at various luminosities.

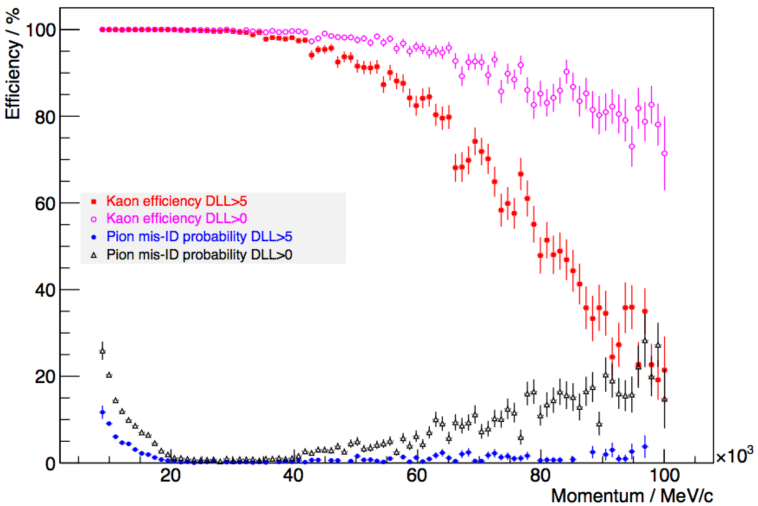

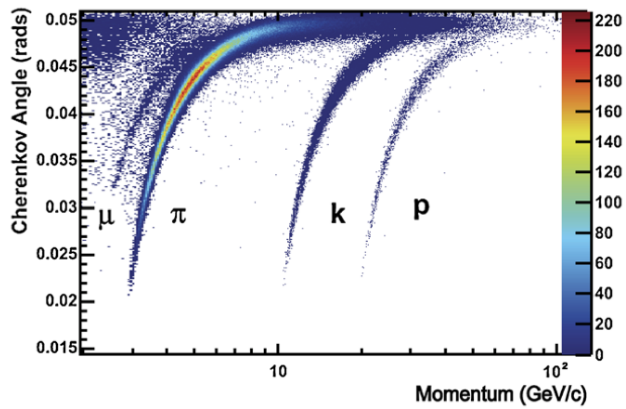

Fig. 14 Reconstructed Cherenkov angle as a function of track momentum in the C_4F_{10} radiator

Performance of the LHCb RICH detector at the LHC
The European Physical Journal CMay 2013, 73:2431.

MUON system

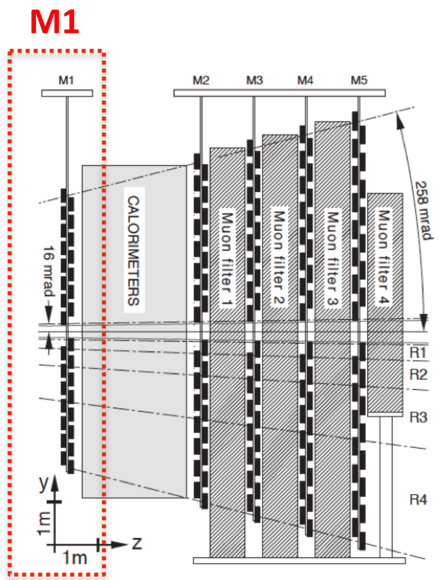

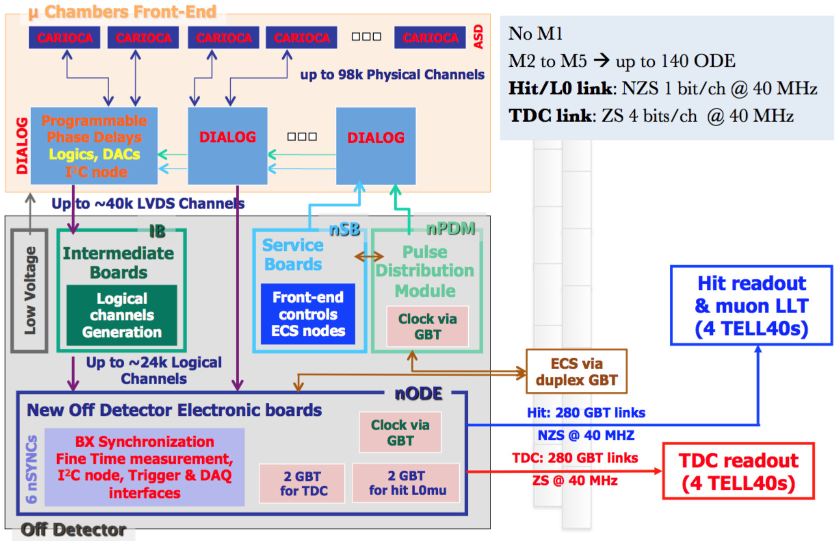

The baseline detectors currently foreseen are anode-pad triple-GEM detectors for the R1 regions and cathode-pad MWPCs for the external regions.

The MWPCs chambers are composed of four gaps, which provide the necessary redundancy in the system.

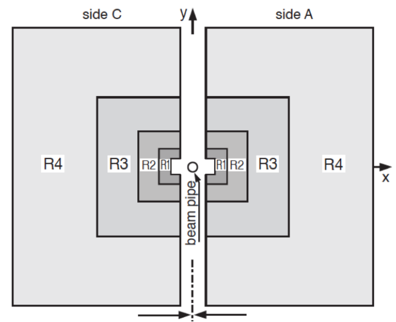

- M1 will be removed in the upgraded experiment.
- The particle flux in the innermost regions of station M2 is expected to be very high, and so additional shielding will be installed around the beam-pipe behind the HCAL to reduce the occupancy in these regions.
- Design of new off-detector readout electronics compliant with full 40 MHz readout and the new GBT–based communication protocol: nODE, new radiation tolerant custom ASIC.
- R&D of new high granularity detectors for regions R1-R2 of stations M2 and M3 and of their front-end electronics.

MUON ID Performance

Efficiency

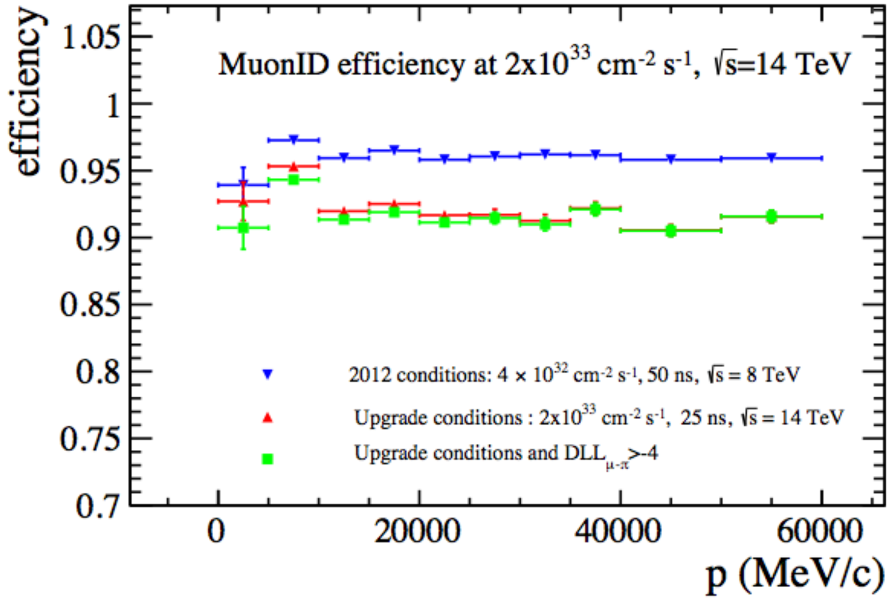

Pion MisID

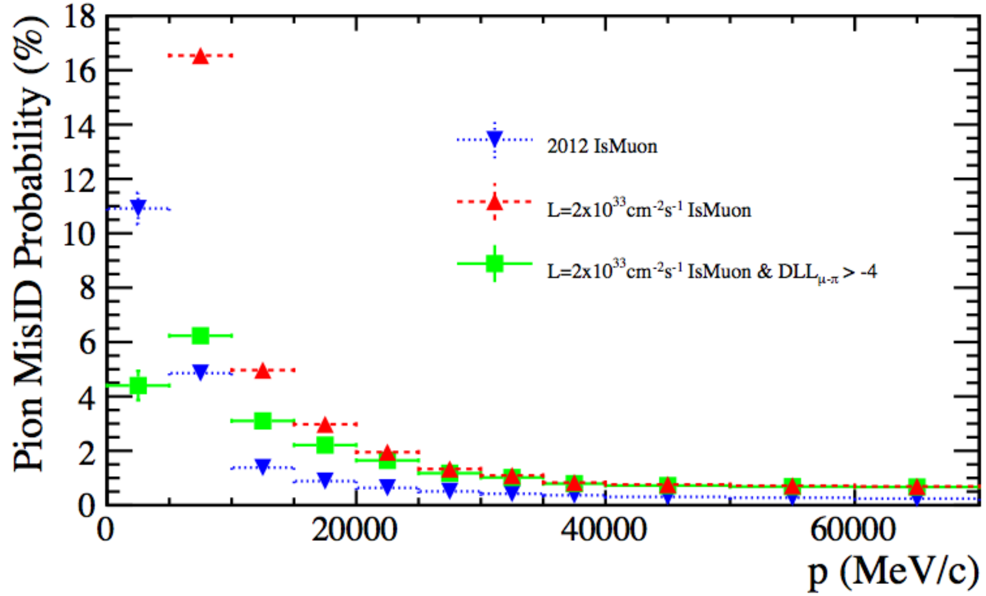

- The efficiency is obtained from a simulated sample of $B_s \rightarrow \mu^+\mu^-$ events and is evaluated for single muon.
- $DLL_{\mu\pi}$ variable based on the distance of matching hits from the extrapolated track in the muon stations combined with the information coming from RICH and calorimeter detectors

Calorimeters

- The scintillating pad detector (SPD) and the pre-shower (PS) of the current detector will be removed.
- The "shashlik" ECAL calorimeter modules, are expected to suffer from radiation damage as the integrated luminosity received by the LHCb detector increases.
- The effect of the removal of the SPD/PS on photon selection for candidates with $p_T \geq$ 200 MeV/c: an absolute reduction of 10 to 15% in the efficiency is observed at a fixed background retention
- For the higher electron momentum sample p > 10 GeV/c the absence of the SPD/PS has very little effect.

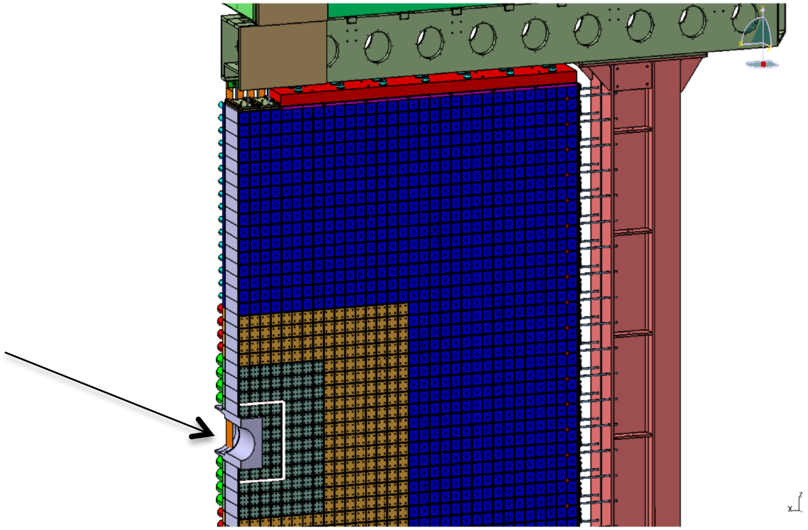

Photon selection

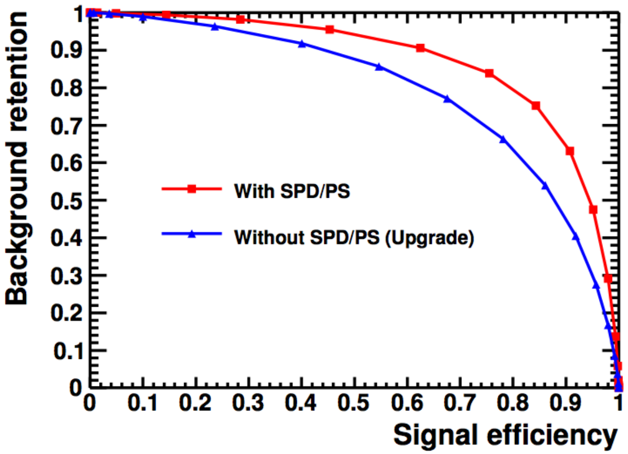

Table 3.7: Typical electron performance in terms of mis-identification rate (in %) at $\varepsilon = 80$ and 90% for $\nu = 2.0$, 3.8 and 7.6 samples and the two studied geometries.

Momentum (GeV/c)	SPD/PS $\nu = 2.0$	SPD/PS $\nu = 3.8$	no SPD/PS $\nu = 3.8$	SPD/PS $\nu = 7.6$	no SPD/PS $\nu = 7.6$
selection efficiency $\varepsilon = 80\%$					
$0 < p < 10$	0.62	0.57	4.6	3.2	9.0
$p > 10$	0.16	0.12	0.16	0.29	0.32
selection efficiency $\varepsilon = 90\%$					
$0 < p < 10$	2.1	2.5	11	12	18
$p > 10$	1.1	0.73	0.72	1.3	1.4

INFN group items

- MUON detector:
  - Cagliari, Ferrara, Firenze, Frascati (LNF), Roma1, Roma2
- Online, readout electronics, software trigger:
  - Bologna, Padova
- Computing:
  - Bologna, Ferrara, Padova
- RICH:
  - Ferrara, Genova, Milano-Bicocca, Padova
- Tracking:
  - Milano
- Trigger GPU, TPU:
  - Padova, Pisa

The LHCb upgrade: physics

EPJ C (2013) 73:2373

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_{sl}^s	6.4×10^{-3} [44]	0.6×10^{-3}	0.2×10^{-3}	0.03×10^{-3}
Gluonic penguins	$2\beta_s^{\text{eff}}(B_s^0 \to \phi\phi)$	–	0.17	0.03	0.02
	$2\beta_s^{\text{eff}}(B_s^0 \to K^{*0}\overline{K}^{*0})$	–	0.13	0.02	< 0.02
	$2\beta^{\text{eff}}(B^0 \to \phi K_S^0)$	0.17 [44]	0.30	0.05	0.02
Right-handed currents	$2\beta_s^{\text{eff}}(B_s^0 \to \phi\gamma)$	–	0.09	0.02	<0.01
	$\tau^{\text{eff}}(B_s^0 \to \phi\gamma)/\tau_{B_s^0}$	–	5 %	1 %	0.2 %
Electroweak penguins	$S_3(B^0 \to K^{*0}\mu^+\mu^-;\, 1 < q^2 < 6\,\text{GeV}^2/c^4)$	0.08 [68]	0.025	0.008	0.02
	$s_0 A_{\text{FB}}(B^0 \to K^{*0}\mu^+\mu^-)$	25 % [68]	6 %	2 %	7 %
	$A_{\text{I}}(K\mu^+\mu^-;\, 1 < q^2 < 6\,\text{GeV}^2/c^4)$	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}	–

Conclusions

- The concept of the LHCb experiment, we believe, is definitely proved: a dedicated experiment for heavy flavour physics, exploiting a forward spectrometer at a hadron collider.
  - Many world leading results and many more to come with the 3.2 fb^{-1} full data set collected.
    - Standard Model still survives.
    - We are now on probing regions where new physics effects might appear.
- LHCb plans the upgrade, to be installed in 2018: the upgrade is an essential next step forward for flavour physics.
- TDRs are all ready, but one, concerning "Trigger and DAQ", on the way to come: is planned for the next June.

Features of current LHCb DAQ

- Push-protocol with centralized flow-control (throttle)
- Connection-less, unreliable protocol from FPGA source directly through DAQ network (2 hops) to PCs → losses in network must be very small → buffering important
- Very lean cost-effective system. Total DAQ cost < 2 MCHF (2009) excluding readout-boards and farm)

# 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

- Very compact system all installed in UX85A on 3 floors
- Max cable distance 36 m –> uses cheap UTP Ethernet everywhere, practically no optical links in the DAQ

The DMA PCIe-3 effective bandwidth

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

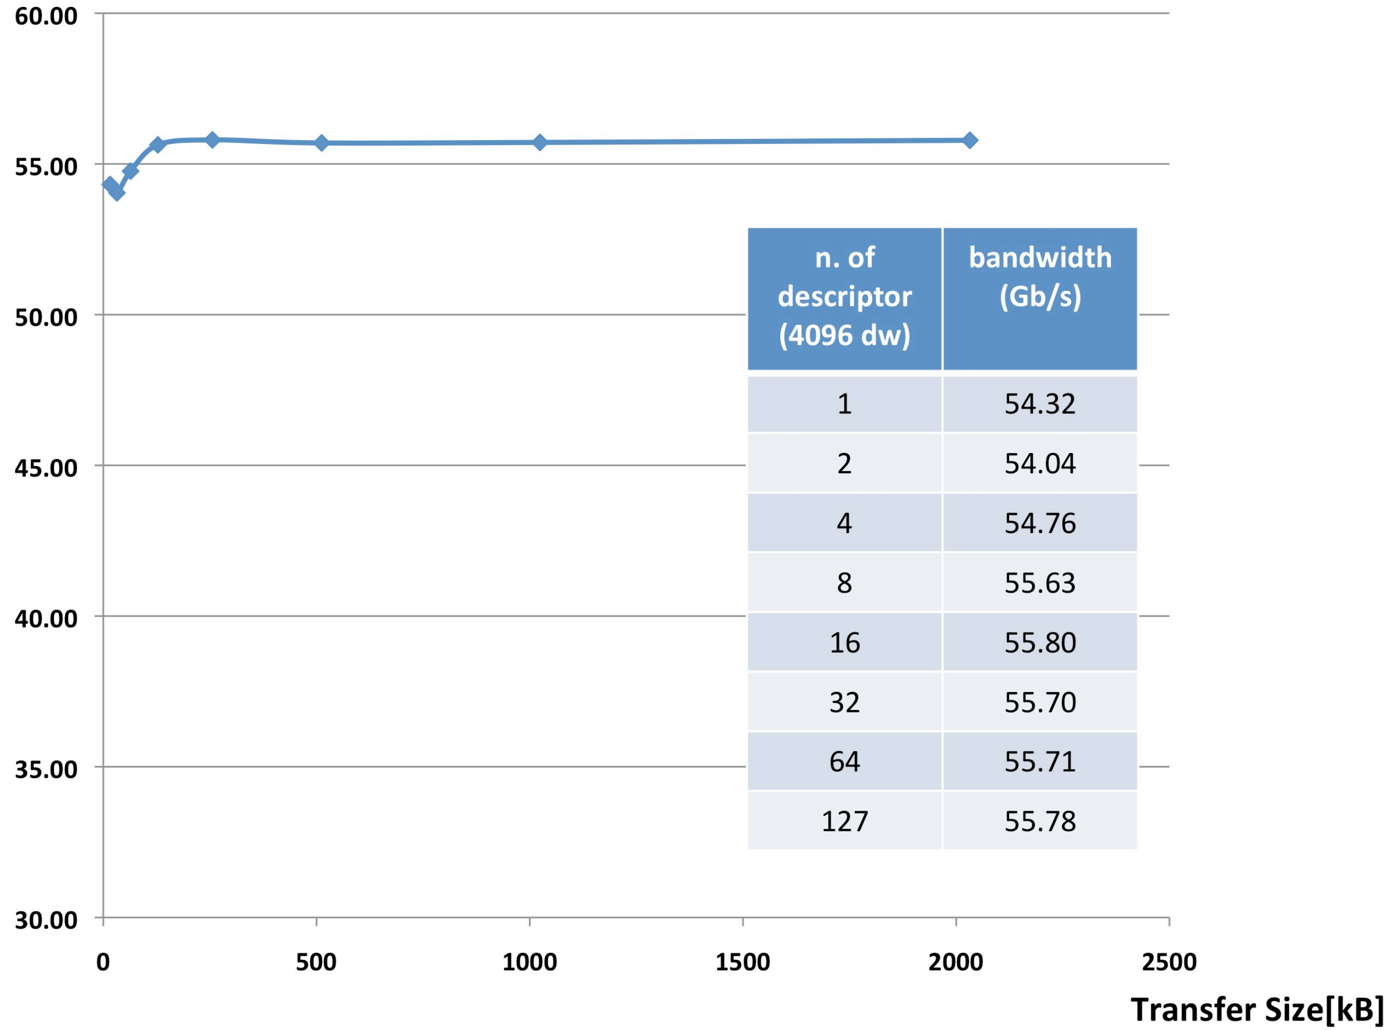

n. of descriptor (4096 dw)	bandwidth (Gb/s)
1	54.32
2	54.04
4	54.76
8	55.63
16	55.80
32	55.70
64	55.71
127	55.78