Upgrade of LHC Online Systems

Wesley H. Smith

U. Wisconsin – Madison

CHEP12, New York

May 21, 2012

Outline:

- LHC Experiments Trigger & DAQ
- LHC Luminosity Upgrade & Challenges
- Upgrade Tools: FPGAs, ATCA, Transceivers, GPU, PCIe
- Upgrades for LHC Experiments' Trigger & DAQ

LHC now

	Design	2010	2011	2012
Beam Energy (TeV)	7	3.5	3.5	4
Bunches/ Beam	2835	368	1380	1380
Proton/Bunch (10^{11})	1.15	1.3	1.5	1.5
Peak Lumi. (10^{32} cm^{-2} s^{-1})	100	2	30	60
Integrated Lumi. (fb^{-1})	100/yr	0.036	6	15*
Pile-Up	23	~1	20	30

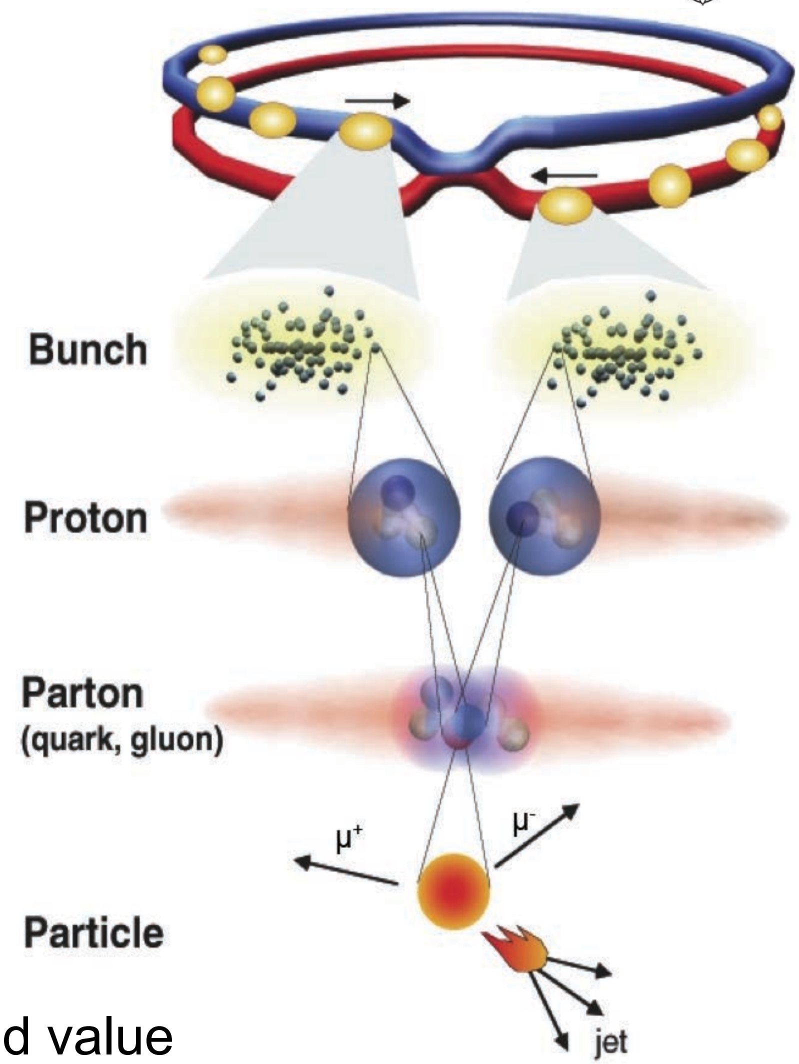

Pile-Up – the number of proton interactions occurring during each bunch crossing

*expected value

ATLAS Three Level Trigger Architecture

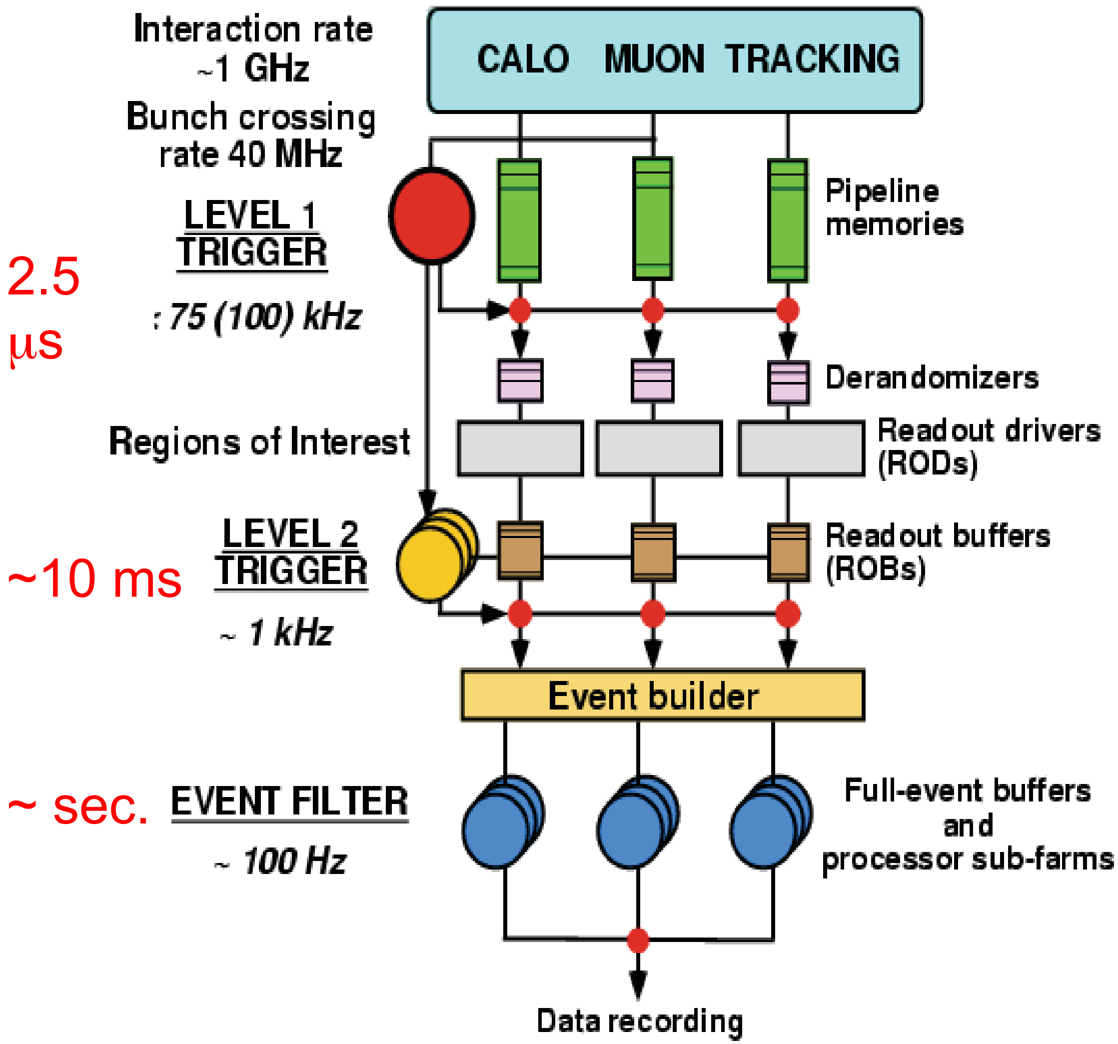

- **LVL1 decision** made with **calorimeter** data with coarse granularity and **muon trigger chamber** data.
  - Buffering on detector
- **LVL2 uses Region of Interest data** (ca. 2%) with full granularity and combines information from all detectors; performs fast rejection.
  - Buffering in Readout Buffers
- **EventFilter** refines the selection, can perform **event reconstruction** at full granularity using latest alignment and calibration data.
  - Buffering in EB & EF

CMS 2-Level Trigger & DAQ

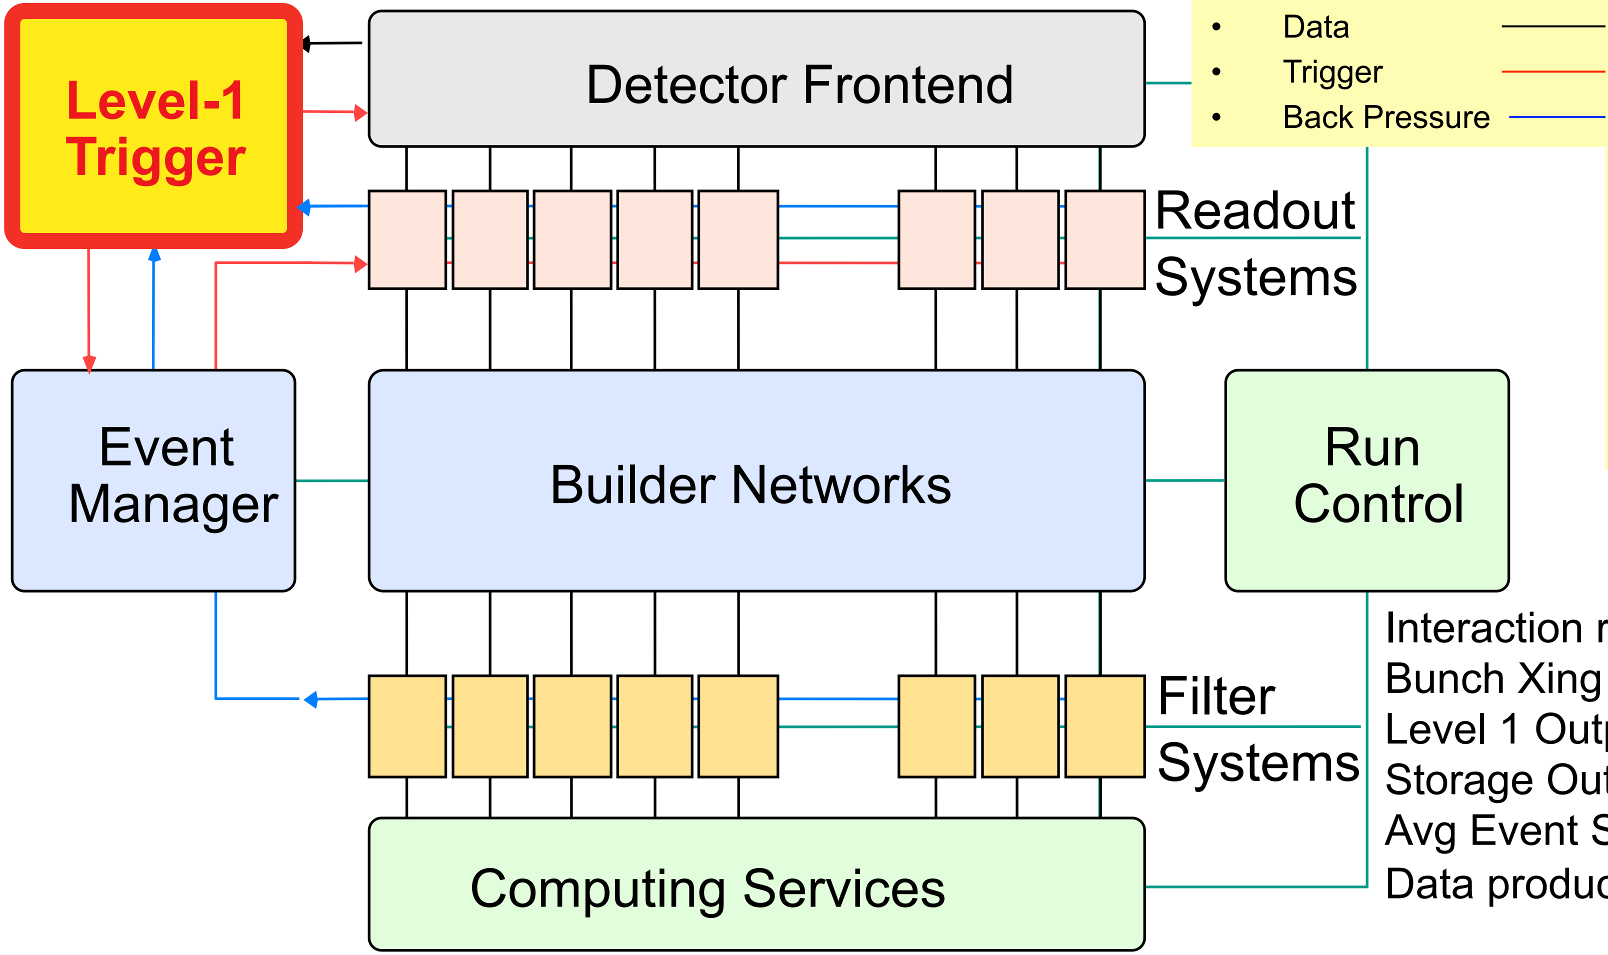

The Lv1-Accept decision causes the readout system to respond with backpressure information.

Interaction rate: 1 GHz
Bunch Xing rate: 40 MHz
Level 1 Output: 100 kHz
Storage Output: ~ 400 Hz
Avg Event Size: ~ 0.5 MB
Data production ~ 1 TB/day

Lv1 decision is distributed to front-ends & readout via TTC system (red).

Readout buffers designed to accommodate Poisson fluctuations from 100 KHz Lv1 trigger rate.

Present ATLAS & CMS Level 1 Trigger Data

Use prompt data (calorimetry and muons) to identify:
High p_t electron, muon, jets, missing E_T

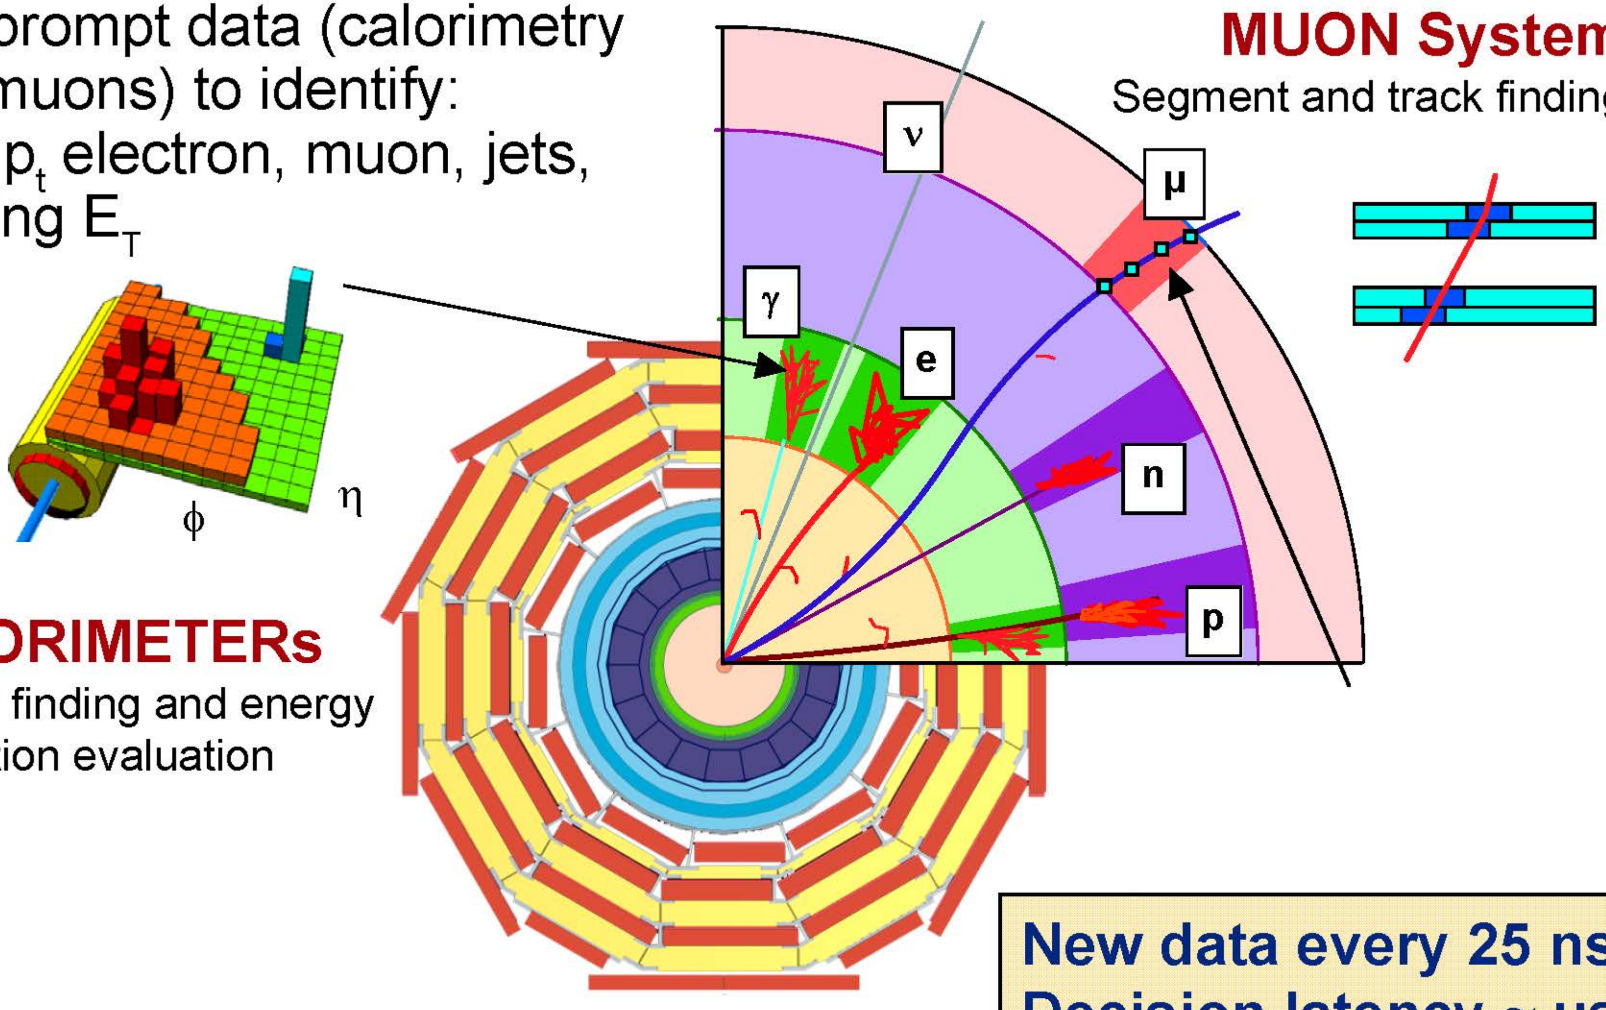

CALORIMETERs

Cluster finding and energy deposition evaluation

New data every 25 ns
Decision latency ~ μs

Present ATLAS & CMS L1: Only Calorimeter & Muon

High Occupancy in high granularity tracking detectors

- **Pattern recognition much faster/easier**

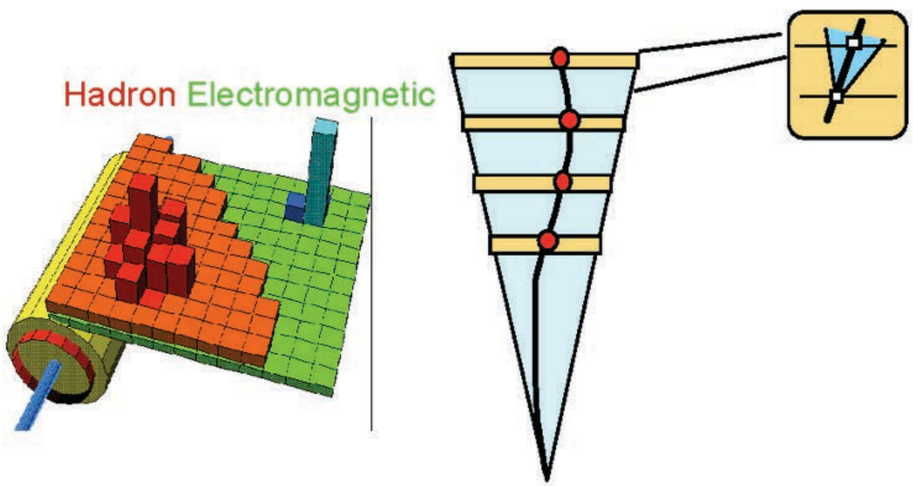

Simple Algorithms

Small amounts of data

- **Compare to tracker info**

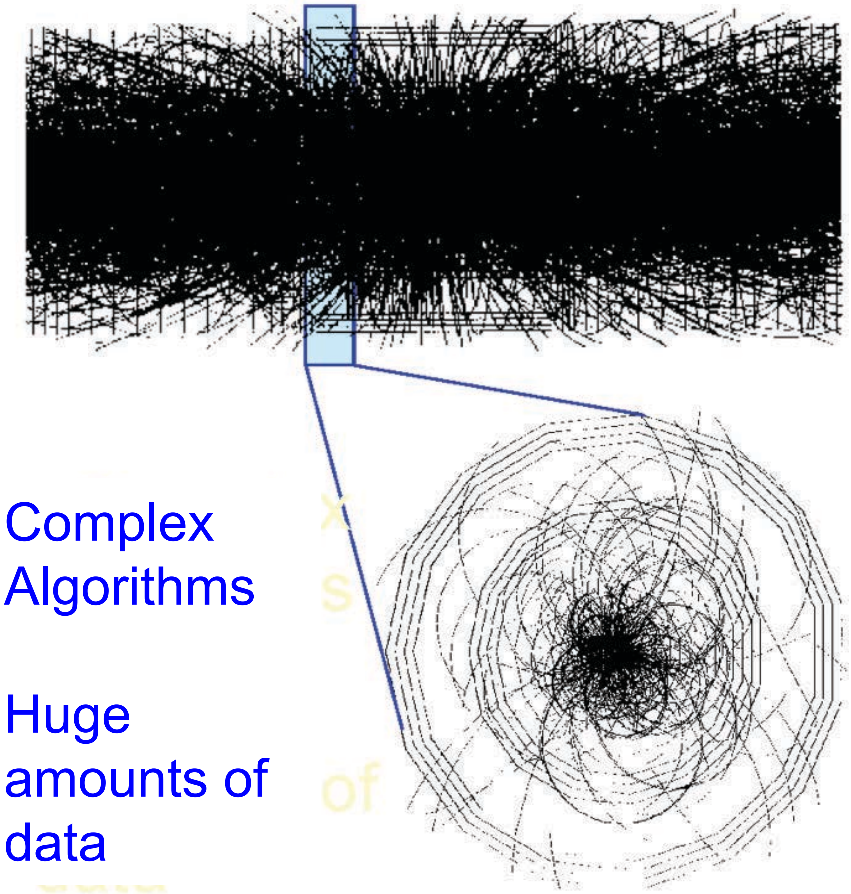

Complex Algorithms

Huge amounts of data

Present ATLAS L1 Trigger

- Process reduced granularity data from calorimeter and muon detectors
- Trigger decision based on object multiplicities
- Generate L1A and send via TTC distribution to detector front-ends to initiate readout
- Maximum round-trip latency 2.5 us
  - Data stores in on-detector pipelines
- Identify regions-of-interest (RoI) to seed L2 trigger
- Custom built electronics
- Synchronous, pipelined processing system operating at the bunch crossing rate

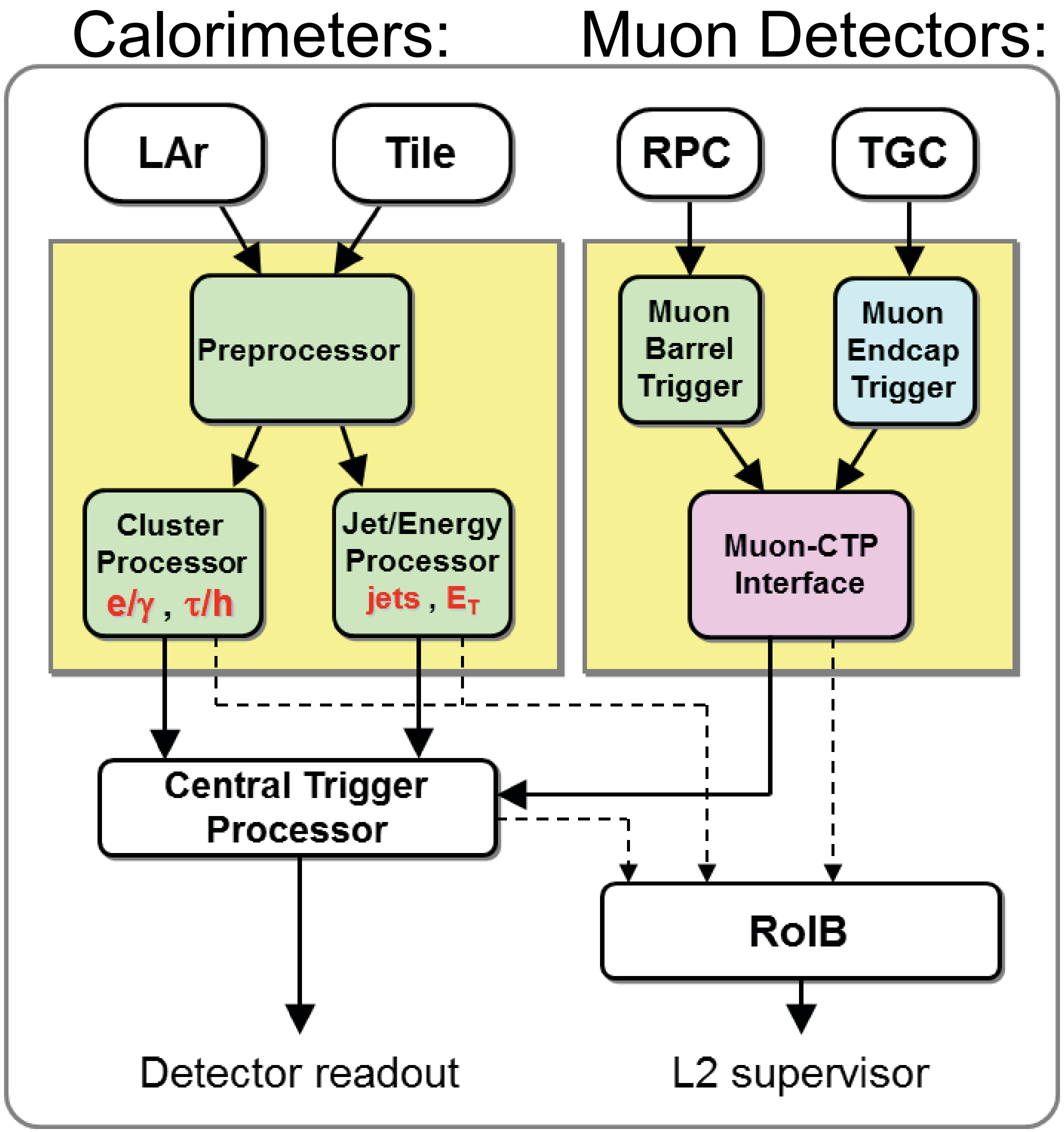

Present CMS L1 Trigger System

Lv1 trigger is based on calorimeter & muon detectors.

At L1 trigger on:

- **4 highest E_t $e^{\pm}/\gamma$**
- **4 highest E_t central jets**
- **4 highest E_t forward jets**
- **4 highest E_t tau-jets**
- **4 highest P_t muons**

For each of these objects rapidity, η, and φ are also transmitted to Global Trigger for topological cuts & so Higher Level Triggers can seed on them.

Also trigger on inclusive triggers:

- **E_t , MET, H_t**

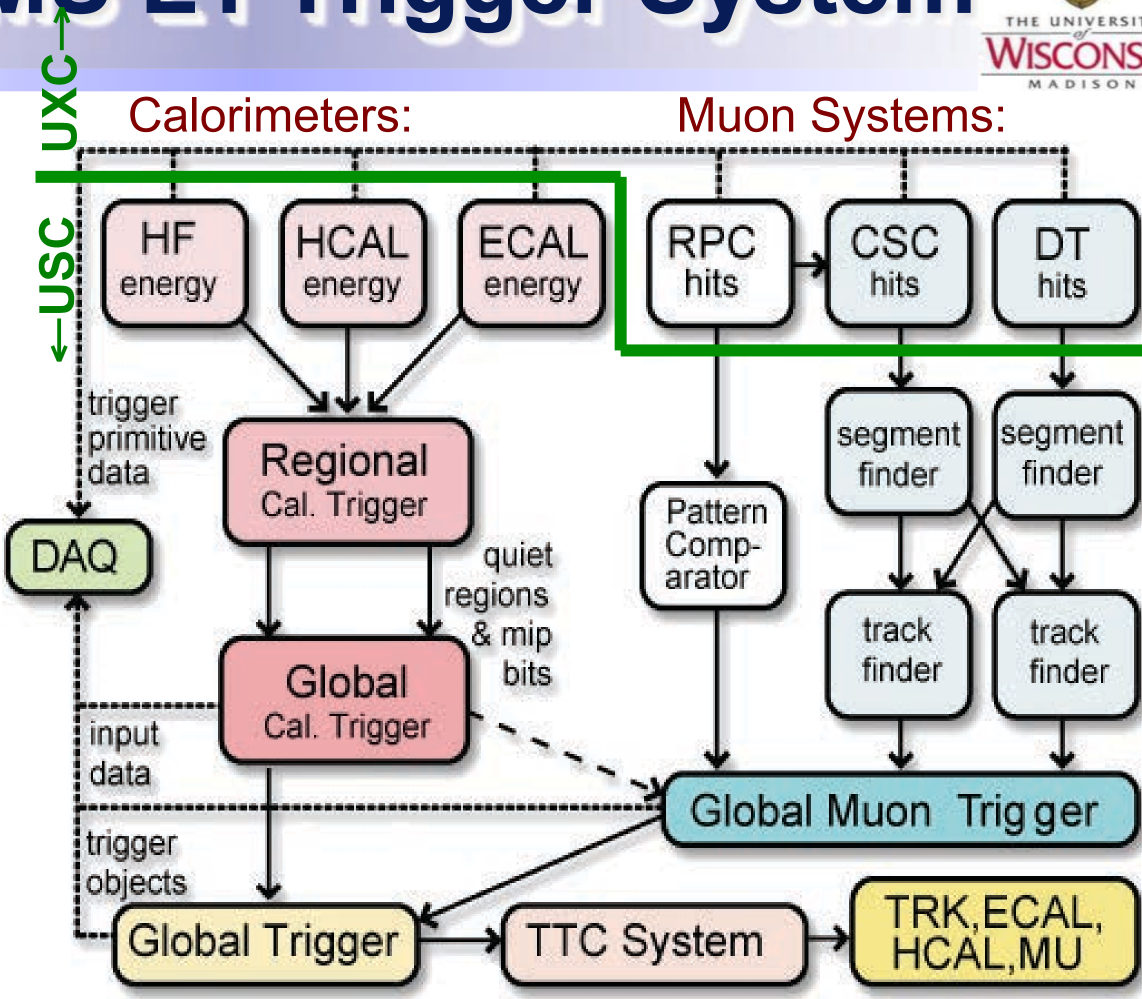

Generate L1A and send via TTC distribution to detector front-ends to initiate readout

Maximum round-trip latency 4 μs

Data stored in on-detector pipelines

Present LHCb Trigger & DAQ

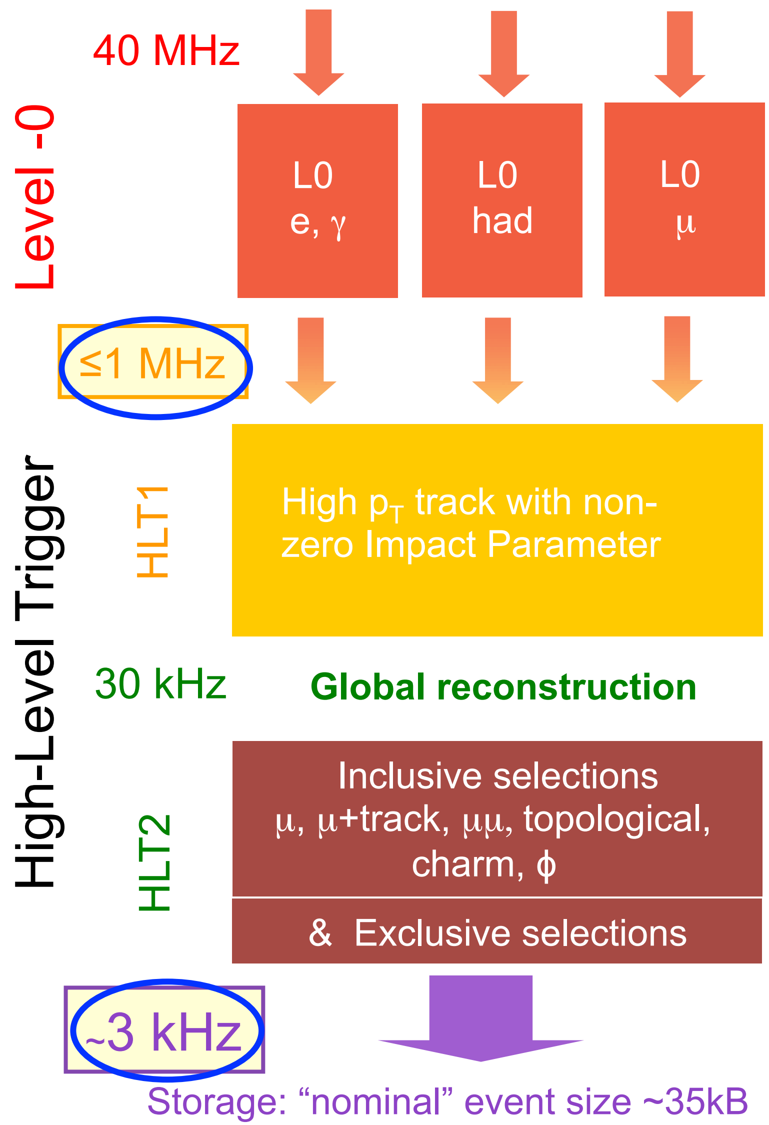

Level 0: Hardware

Both Software Levels run on commercial PCs

Level-1:

- Input: 4.8 kB @ 1.1 MHz
- uses reduced data set: only part of the sub-detectors (mostly Vertex-detector and some tracking) with limited-precision data
- reduces event rate by selecting events with displaced secondary vertices

High Level Trigger (HLT)

- Input: 38 kB @ 30 kHz
- uses all detector information
- Output 3 kHz for permanent storage

Present Alice Trigger & DAQ

3 decision levels: L0: 1.2 μs, L1: 6.5 μs, L2: 88 μs

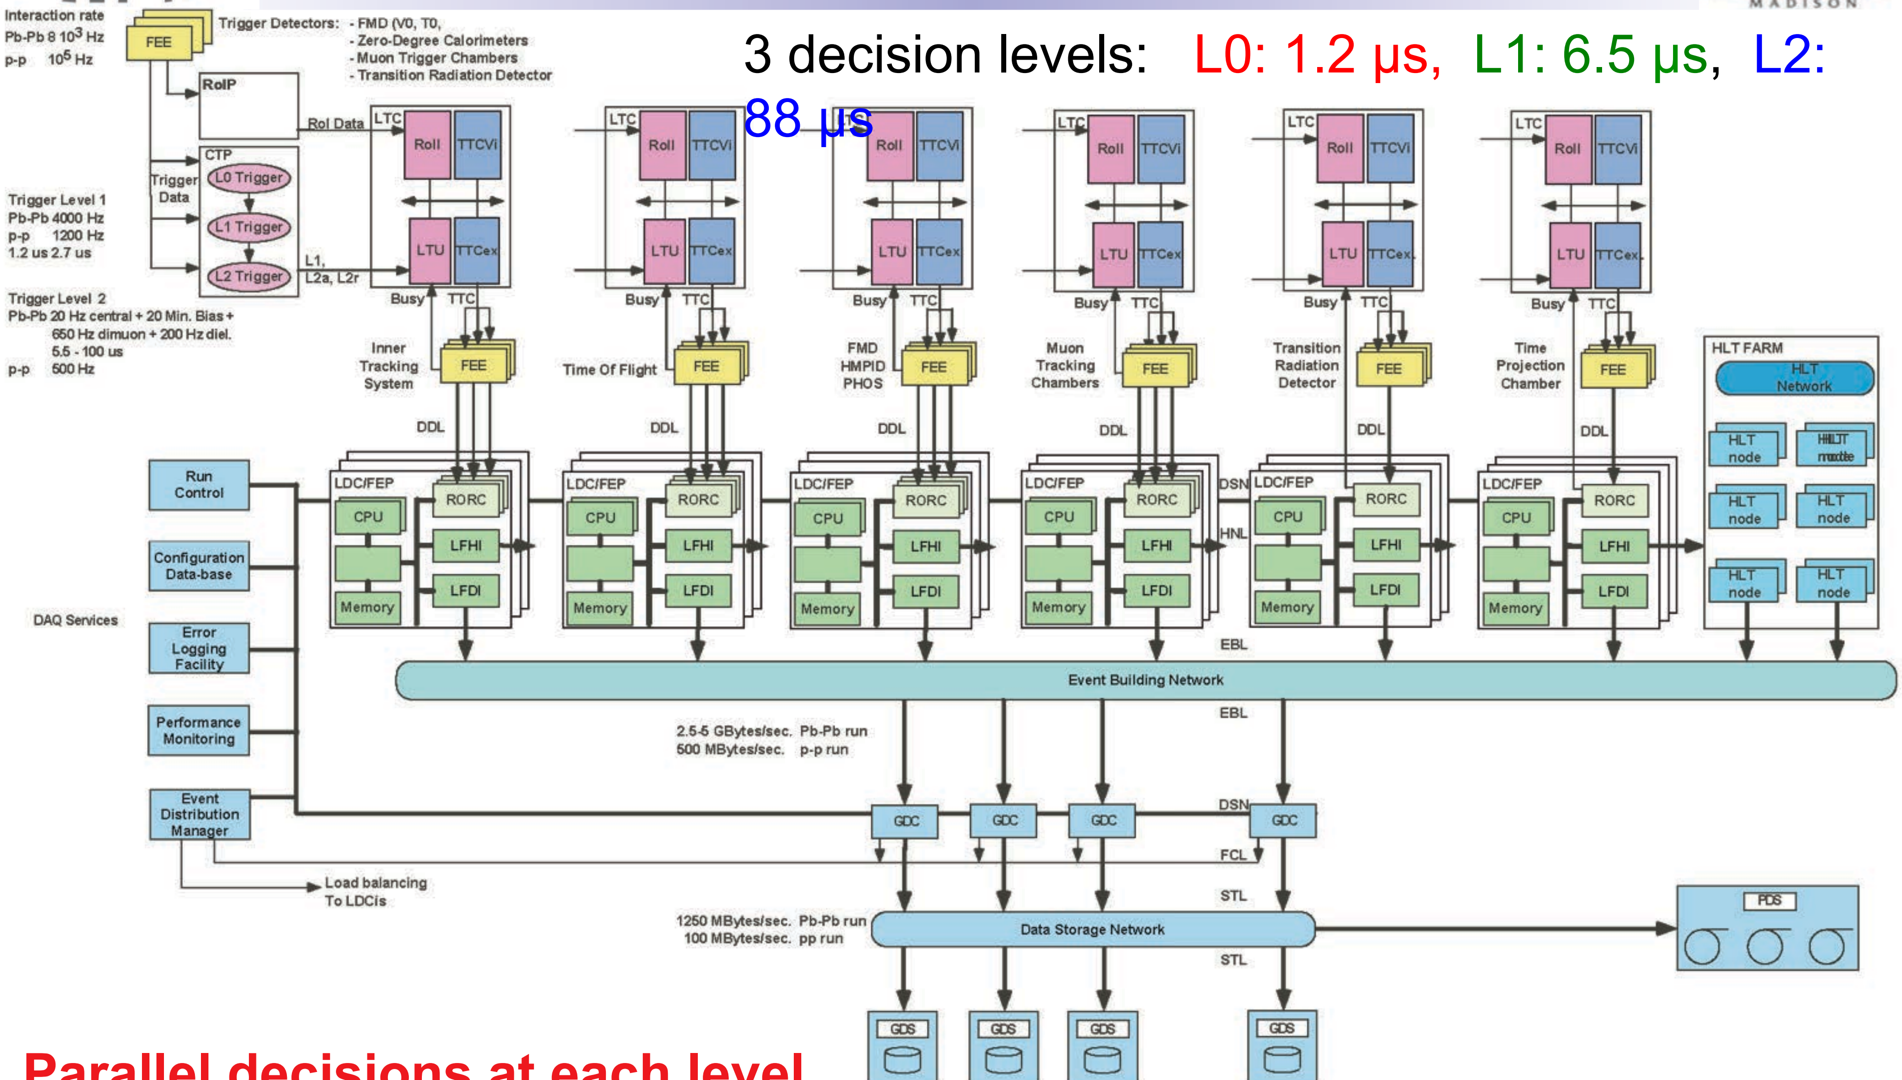

Parallel decisions at each level

- different groups of detectors (clusters) are reading out different events at same time

ATLAS & CMS High Luminosity Motivation

Establish nature of Higgs boson and of EWSB:

- fundamental or composite?
- how many doublets? singlets? charged H's?

Need to measure, as accurately as possible:

- Higgs couplings to fermions, gauge bosons & self-couplings
- Rare decay modes, possible Flavor Changing Neutral Current
- WW scattering at high E
- Gauge boson self-couplings

Example:

- For light Higgs: $H \rightarrow Z\gamma$ @ 3.5/11 σ with 600/6000 fb^{-1} or $H \rightarrow \mu^{+}\mu^{-}$ < 3.5σ for 600 fb^{-1} and ~ 7σ for 6000 fb^{-1}

Requirements for LHC phases of the upgrades: ~2010-2030

Phase 1: (until 2021)

- Goal of extended running in second half of decade to collect ~100s/fb
- 80% of this luminosity in last three years of this decade
- About half the luminosity would be delivered at luminosities above the original LHC design luminosity
- Trigger & DAQ systems should be able to operate with a peak luminosity of up to 2 x 10^{34}

Phase 2: High Lumi LHC (2023+)

- Continued operation of the LHC beyond a few 100/fb will require substantial modification of detector elements
- Goal is to achieve 3000/fb in phase 2
- Need to be able to integrate ~300/fb-yr
- Will require new tracking detectors for ATLAS & CMS
- Trigger & DAQ systems should be able to operate with a peak luminosity of up to 5 x 10^{34}

Detector Luminosity Effects

$H \rightarrow ZZ \rightarrow \mu\mu ee$, M_H= 300 GeV for different luminosities in CMS

10^{32} cm^{-2}s^{-1}

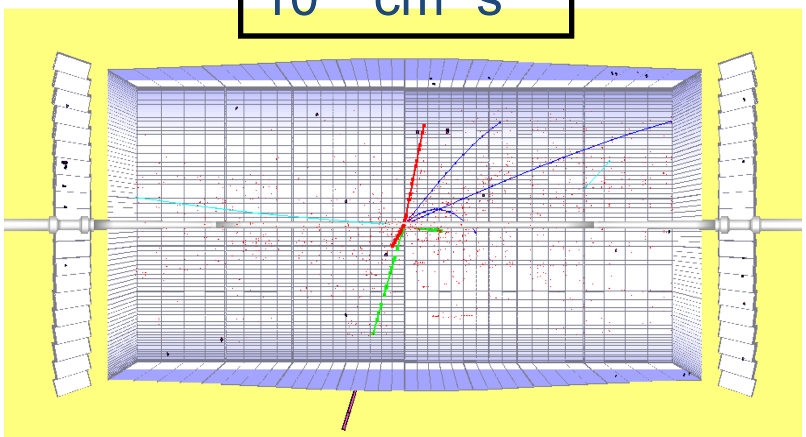

10^{33} cm^{-2}s^{-1}

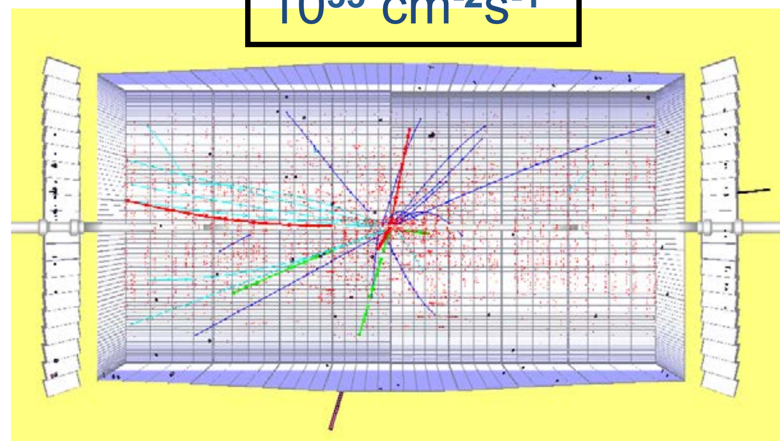

10^{34} cm^{-2}s^{-1}

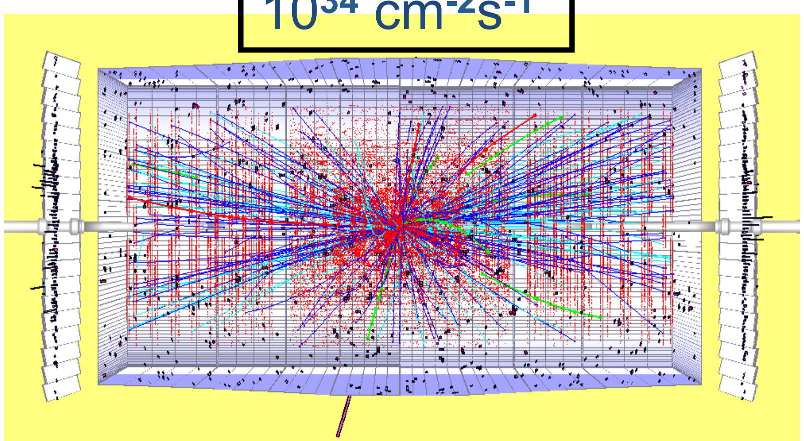

10^{35} cm^{-2}s^{-1}

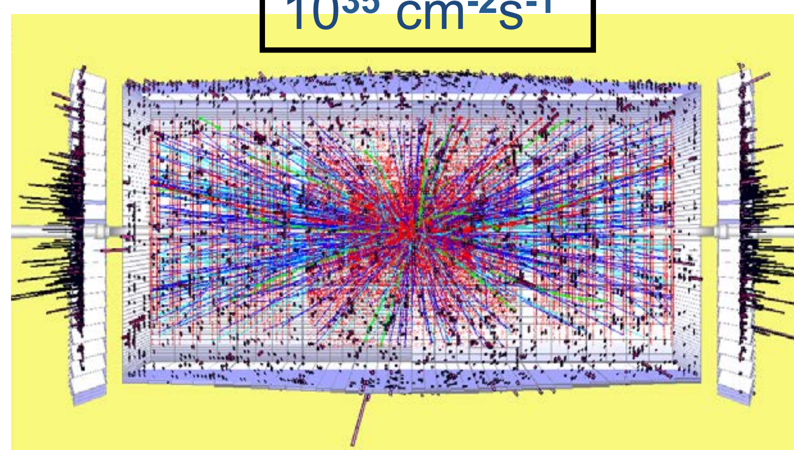

Tools for Upgrades: FPGAs

Logic Cells

- 28 nm: > 2X gains over 40 nm→

On-Chip High Speed Serial Links:

- Connect to new compact high density optical connectors (SNAP-12…)

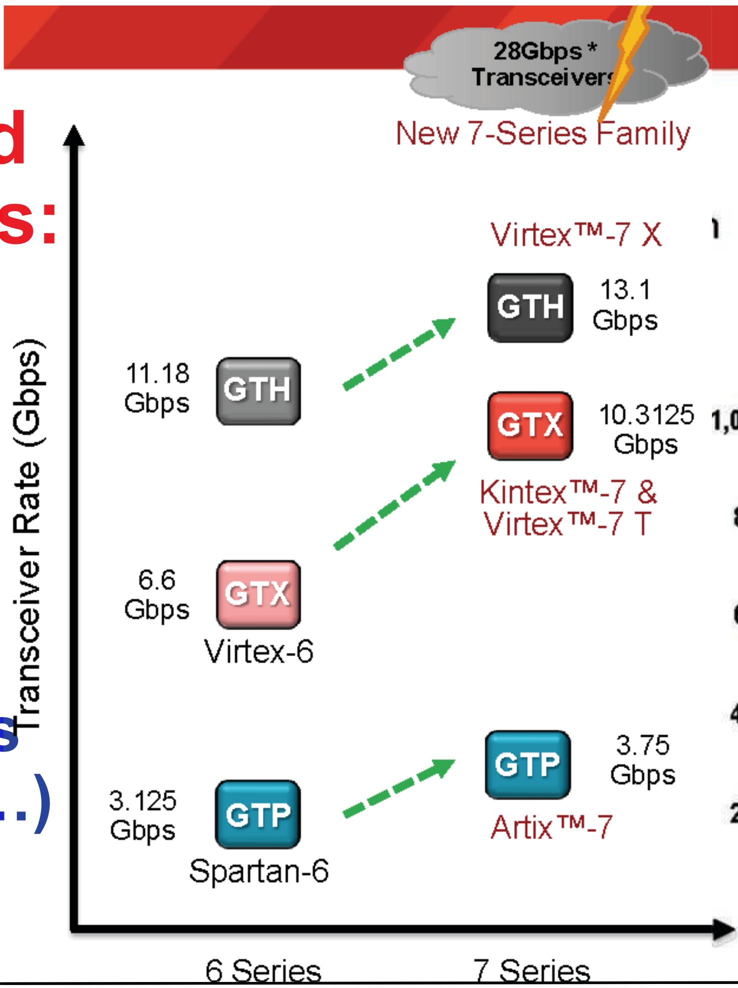

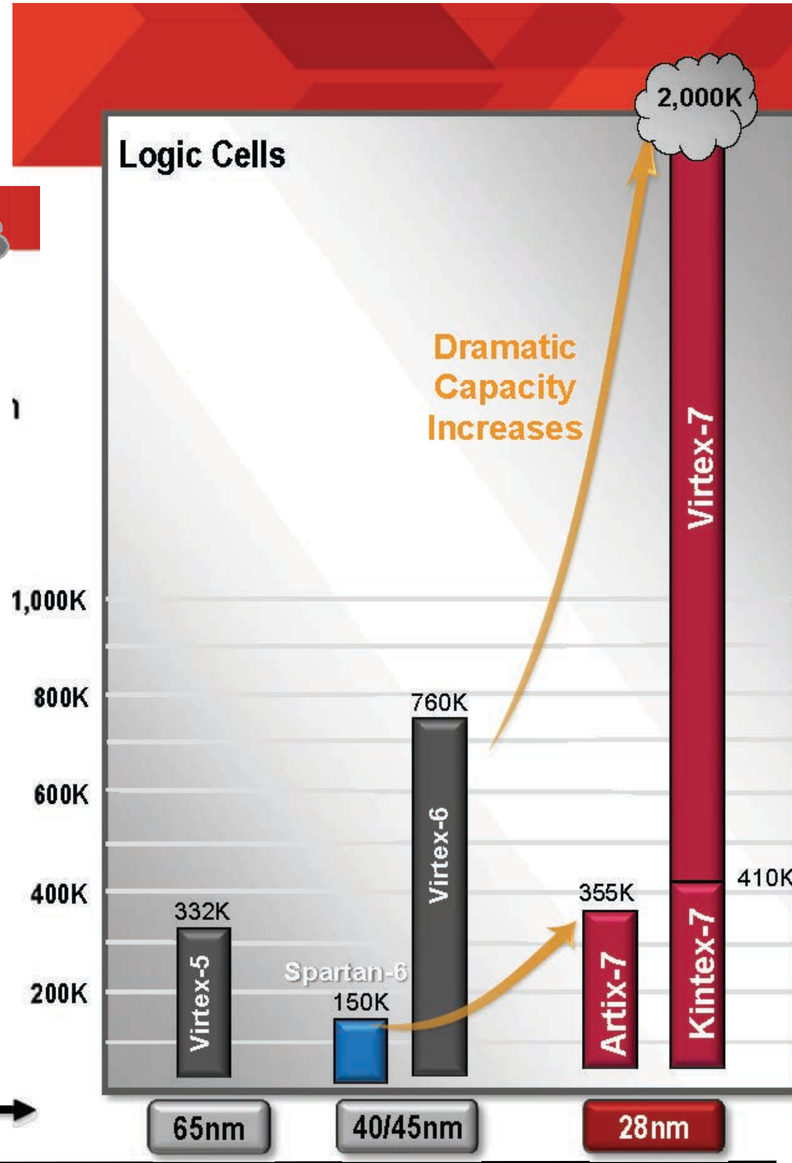

Tools for upgrades: ATCA

- **Advanced Telecommunications Computing Architecture ATCA**
- **Example: ATLAS Upgrade Calorimeter Trigger Topological Processor Card**
  - **12-chan. ribbon fiber optic modules**
  - **Backpl. opt. ribbon fiber connector**

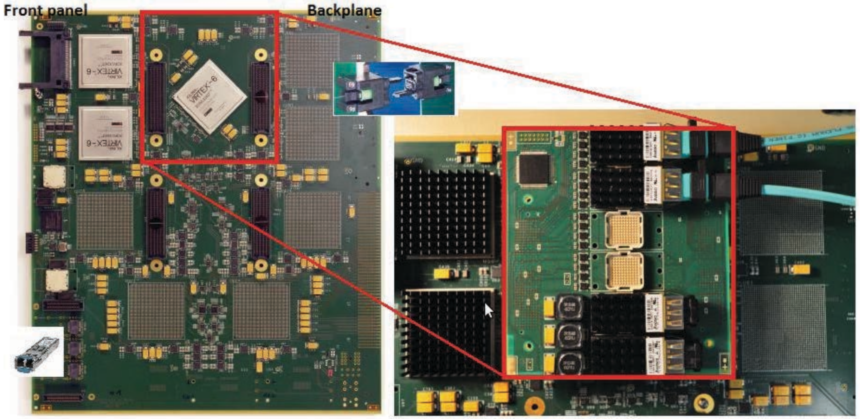

- **Example: µTCA derived from AMC std. used by CMS HCAL, Trig.**
  - **Advanced Mezzanine Card**
  - **Up to 12 AMC slots**
    - *Processing modules*
  - **6 standard 10Gb/s point-to -point links slot to hub slots (more cavailable)**
  - **Redundant power, controls,clocks**
  - **Each AMC can have in principle (20) 10 Gb/sec ports**
  - **Backplane customization is routine & inexpensive**

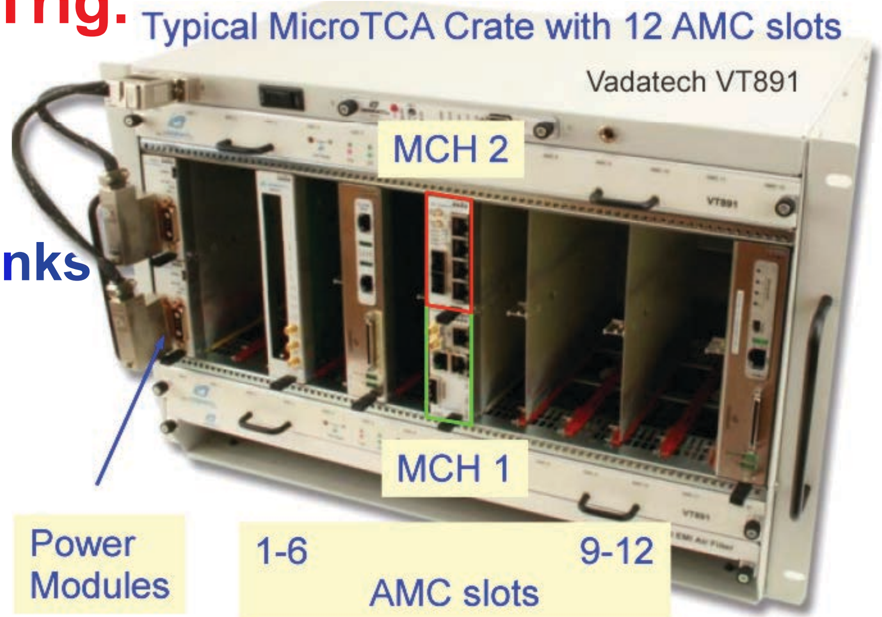

Tools cont'd: CPU, GPU, PCIe

CPU Gains for High Level Triggers: Moore's Law

GPU Enhancement of HLT →

- GPU performance tracks Moore's Law, since GPU architecture is scalable:
  - Large Increase in memory bandwidth x10 in Gbytes/s
  - Power efficient x3 with latest GPU card
  - Well suited to tracking, fitting algorithms

Peak Double Precision FP

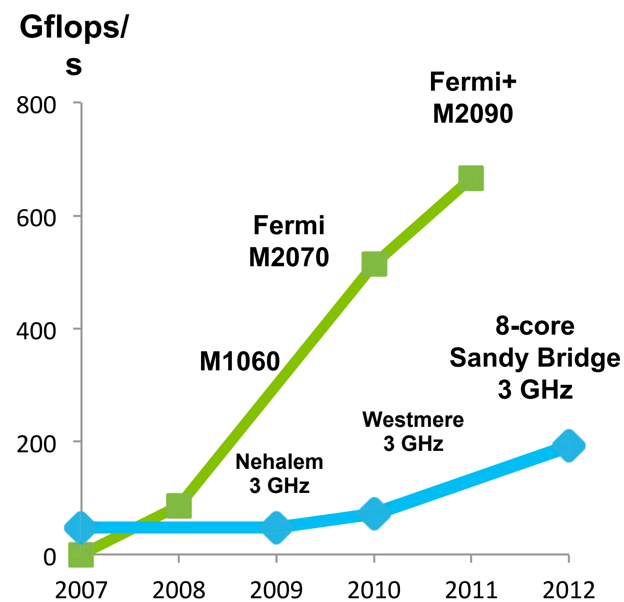

Enhancement of detector to DAQ readout:

- **PCI Express Gen3 Cards now available**
- **Up to 56Gb/s InfiniBand or 40 Gigabit Ethernet per port**

ATLAS Upgrade Trigger Strategy

Near Future (2014)

- New calorimeter trigger & central trigger processor modules provide topological triggers, more triggers

Phase 1:

- "New Small Wheel" provides inner track segments to reduce endcap muon trigger rate
- Muon trigger upgrade to provide topological triggers
- Calorimeter trigger digital "preprocessor" & feature extractors allow use of finer granularity information
- Latency & L1 Trigger Rate stay same through phase 1

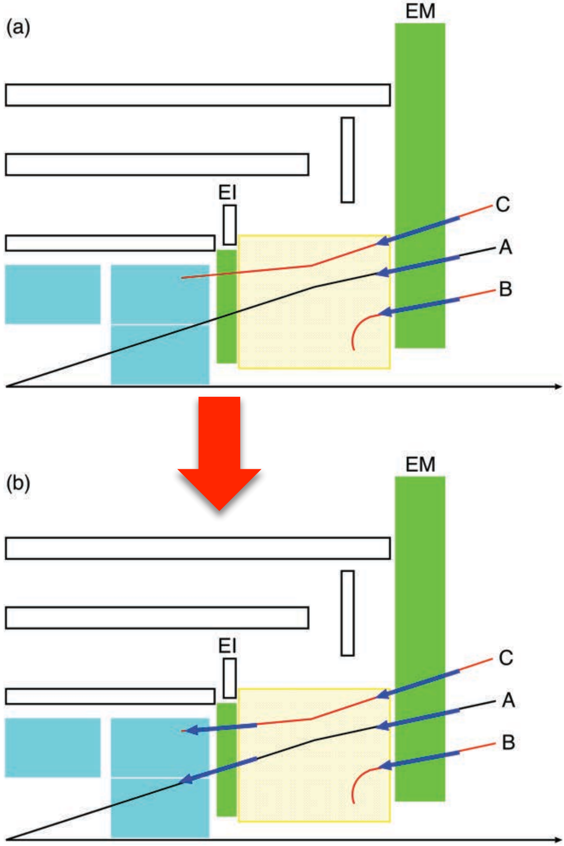

Phase 2 Options:

- Divide L1 trigger into L0, L1 of latency 5, 20 μsec, rate < 500, 200 kHz
- L0 uses Calo & Muon Triggers, generates track trigger seeds
- L1 uses track trigger & more muon detectors & more fine-grained calorimeter trigger information.

CMS Upgrade Trigger Strategy

Constraints

- **Output rate at 100 kHz**
- **Input rate increases x2/x10 (Phase 1/Phase 2) over LHC design (10^{34})**
  - Same x2 if crossing freq/2, e.g. 25 ns spacing → 50 ns at 10^{34}
- **Number of interactions in a crossing (Pileup) goes up by x4/x20**
- **Thresholds remain ~ same as physics interest does**

Example: strategy for Phase 1 Calorimeter Trigger

- **Present L1 algorithms inadequate above 10^{34}**
  - Pileup degrades object isolation
- **More sophisticated clustering & isolation deal w/more busy events**
  - Process with full granularity of calorimeter trigger information
- **Should suffice for x2 reduction in rate as shown with initial L1 Trigger studies & CMS HLT studies with L2 algorithms**

Potential new handles at L1 needed for x10 (Phase 2: 2023+)

- **Tracking to eliminate fakes, use track isolation.**
- **Vertexing to ensure that multiple trigger objects come from same interaction**
- **Requires finer position resolution for calorimeter trigger objects for matching (provided by use of full granularity cal. trig. info.)**

CMS Phase 1 Upgrade Calorimeter Trigger

- **Particle Cluster Finder**
  - Applies tower thresholds to Calorimeter
  - Creates overlapped 2x2 clusters
- **Cluster Overlap Filter**
  - Removes overlap between clusters
  - Identifies local maxima
  - Prunes low energy clusters
- **Cluster Isolation and Particle ID**
  - Applied to local maxima
  - Calculates isolation deposits around 2x2,2x3 clusters
  - Identifies particles
- **Jet reconstruction**
  - Applied on filtered clusters
  - Groups clusters to jets
- **Particle Sorter**
  - Sorts particles & outputs the most energetic ones
- **MET,HT,MHT Calculation**
  - Calculates Et Sums, Missing Et from clusters

Rate reductions x4 w/improved efficiency

Implemented in 4 μTCA Crates

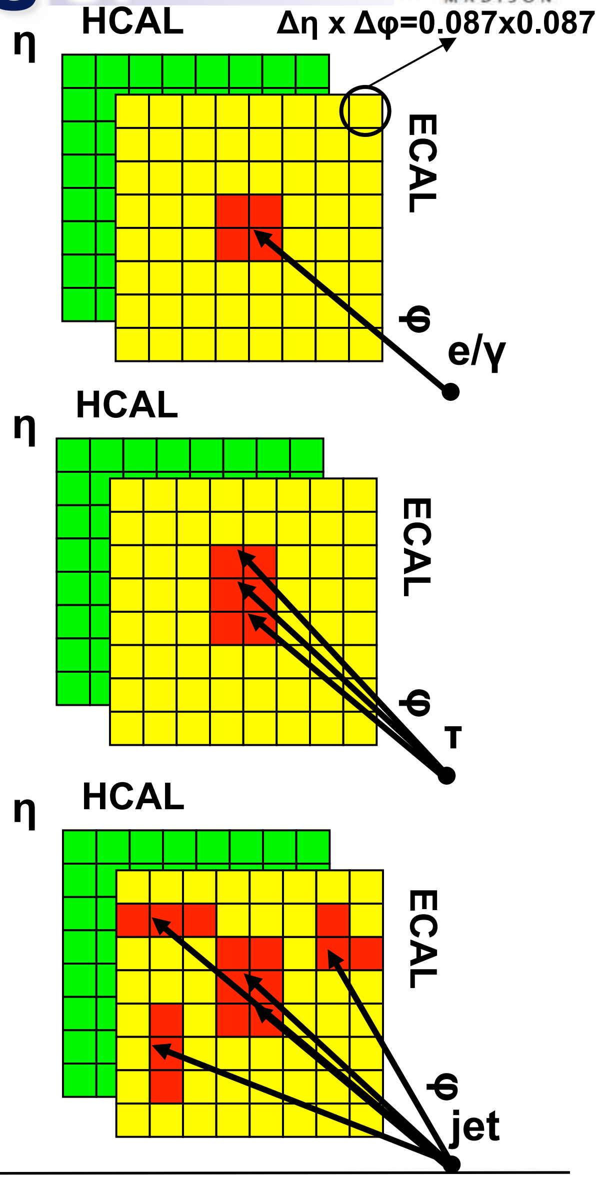

CMS Phase 2: Tracker input to L1 Trigger

Use of Tracker input to Level-1 trigger

- μ, e and jet rates would exceed 100 kHz at high luminosity
  - Even considering "phase-1" trigger upgrades
- Increasing thresholds would affect physics performance
  - Performance of algorithms degrades with increasing pile-up
    - Muons: increased background rates from accidental coincidences
    - Electrons/photons: reduced QCD rejection at fixed efficiency from isolation
- Add tracking information at Level-1
  - Move part of HLT reconstruction into Level-1!

Full-scope objectives:

- Reconstruct "all" tracks above 2 - 2.5 GeV
- Identify the origin along the beam axis with ~ 1 mm precision

CMS Track Trigger Architectures: Phase 2

"Push" path:

- L1 tracking trigger data combined with calorimeter & muon trigger data regionally with finer granularity than presently employed.
- After regional correlation stage, physics objects made from tracking, calorimeter & muon regional trigger data transmitted to Global Trigger.

"Pull" path:

- L1 calorimeter & muon triggers produce a "Level-0" or L0 "pre-trigger" after latency of present L1 trigger, with request for tracking info at ~1 MHz. Request only goes to regions of tracker where candidate was found. Reduces data transmitted from tracker to L1 trigger logic by 40 (40 MHz to 1 MHz) times probability of a tracker region to be found with candidates, which could be less than 10%.
- Tracker sends out info. for these regions only & this data is combined in L1 correlation logic, resulting in L1A combining track, muon & cal. info..
- Only on-detector tracking trigger logic in specific region would see L0

"Afterburner" path:

- L1 Track trigger info, along with rest of information provided to L1 is used at very first stage of HLT processing. Provides track information to HLT algorithms very quickly without having to unpack & process large volume of tracker information through CPU-intensive algorithms. Helps limit the need for significant additional processor power in HLT computer farm.

CMS Track Trigger: General concept

Silicon modules provide at same time "Level-1 data" (@ 40 MHz) & "readout data" (@ 100 kHz, upon Level-1 trigger)

- The whole tracker sends out data at each BX: "push path"

Level-1 data require local rejection of low-p_T tracks

- To reduce the data volume, and simplify track finding @ Level-1
  - Threshold of ~ 1-2 GeV ⇒ data reduction of > one order of magnitude

Design modules with p_T discrimination ("p_T modules")

- Correlate signals in two closely-spaced sensors
  - Exploit CMS strong magnetic field

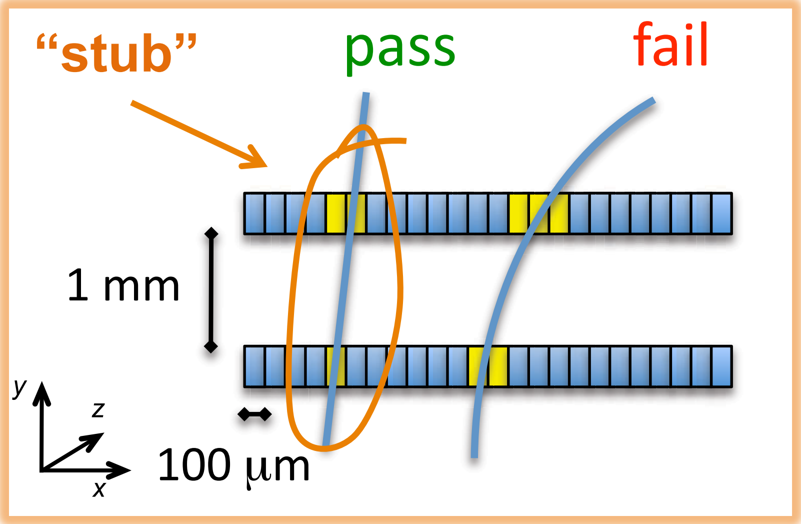

Level-1 "stubs" processed in back-end

- Form Level-1 tracks, p_T above 2-2.5 GeV
  - Improve different trigger channels

CMS Track Trigger p_T modules: working principle

Sensitivity to p_T from measurement of $\Delta(R\varphi)$ over a given ΔR

For a given p_T, $\Delta(R\varphi)$ increases with R

- **Same geometrical cut, corresponds to harder p_T cuts at large radii**
- **At low radii, rejection power limited by pitch**
- **Optimize selection window and/or sensors spacing**
  - To obtain consistent p_T selection through tracking volume

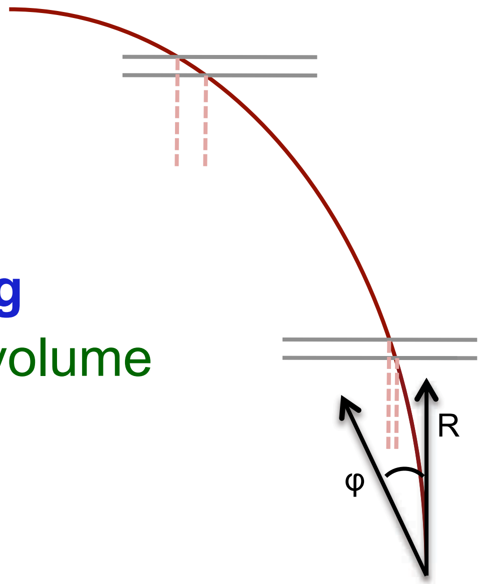

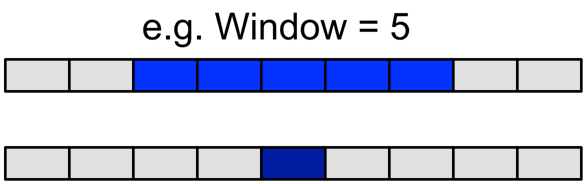

In the barrel, ΔR is given directly by the sensors spacing

In the end-cap, it depends on the location of the detector

- **End-cap configuration typically requires wider spacing**

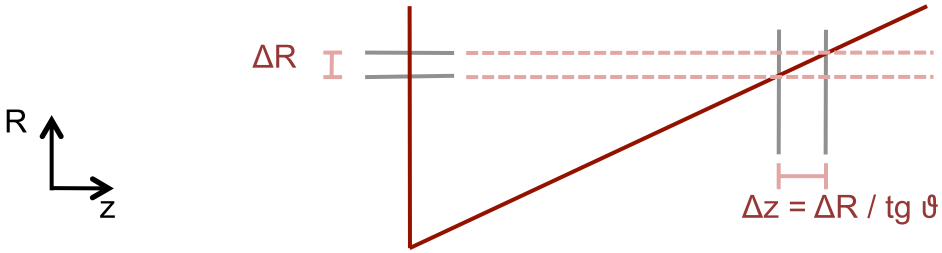

ATLAS Trigger Upgrades

Various projects being pursued:

- **Track trigger**
  - Fast Track Finder (FTK), hardware track finder for ATLAS (at L1.5)
    - **Phase 1**
  - ROI based track trigger at L1 – Phase 2
  - Self seeded track trigger at L1 – Phase 2

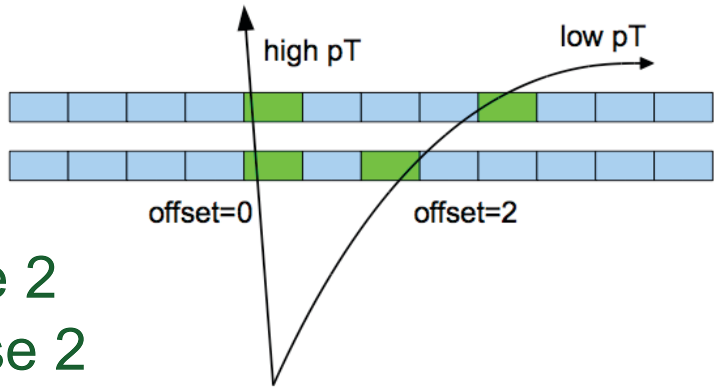

- **Combining trigger objects at L1 & topological "analysis"**
  - Phase 1 & 2
- **Full granularity readout of calorimeter**
  - requires new electronics – Phase 2
- **Changes in muon systems (small wheels), studies of an MDT based trigger & changes in electronics – Phase 1**
- **Upgrades of HLT farms**

Some of the changes are linked to possibilities that open when electronics changes are made (increased granularity, improved resolution & increased latency)

ATLAS FastTracKer (FTK)

For Phase 1:

Dedicated hardware processor completes GLOBAL track reconstruction by beginning of level-2 processing.

- **Allows very rapid rejection of most background, which dominates the level-1 trigger rate.**
- **Frees up level-2 farm to carry out needed sophisticated event selection algorithms.**

Addresses two time-consuming stages in tracking

- **Pattern recognition – find track candidates with enough Si hits**
  - 10^9 prestored patterns simultaneously see each silicon hit leaving the detector at full speed.
- **Track fitting – precise helix parameter & χ^2 determination**
  - Equations linear in local hit coordinates give near offline resolution

ATLAS FTK Approach

Use hardware to perform the global tracking in two steps
pattern recognition and track fit

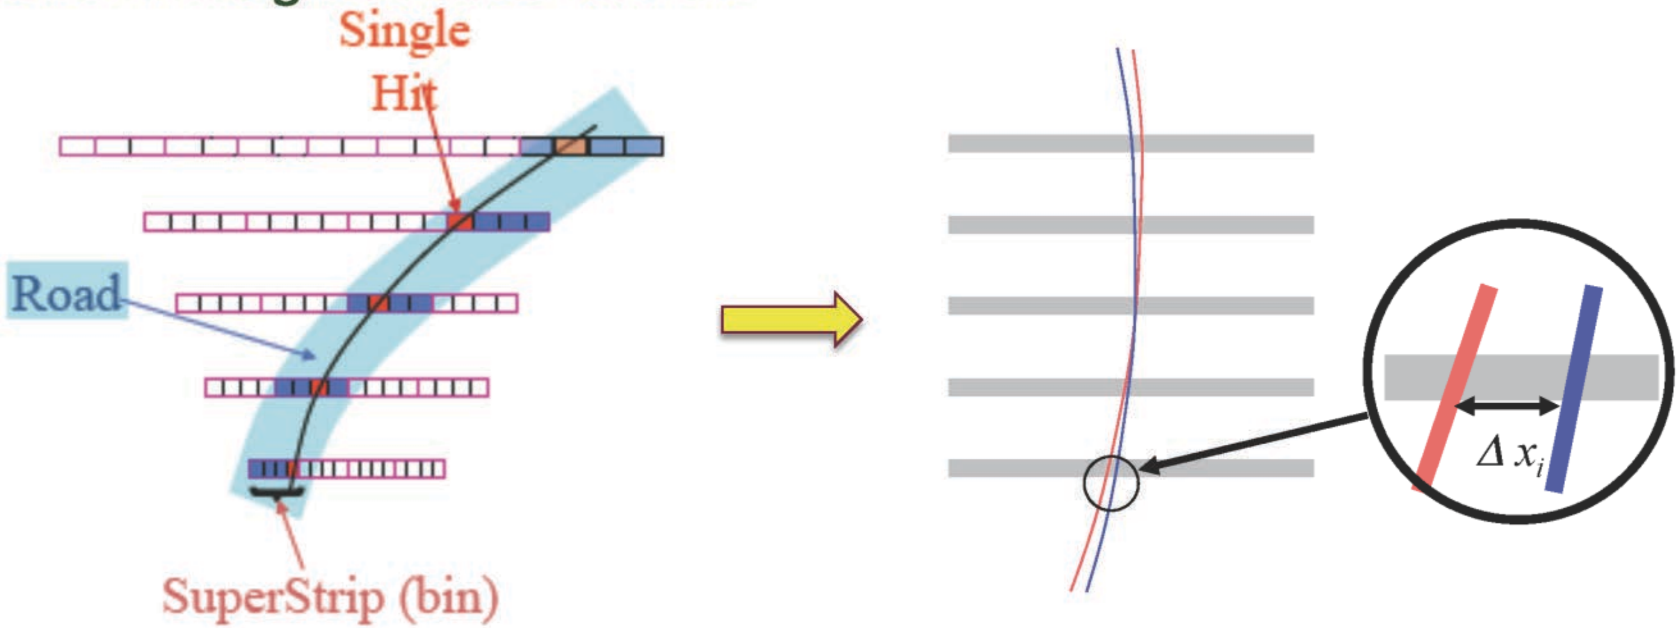

Pattern recognition in coarse resolution (superstrip➔road)

Track fit in full resolution (hits in a road)

$F(x_1, x_2, x_3, \ldots) \sim a_0 + a_1\Delta x_1 + a_2\Delta x_2 + a_3\Delta x_3 + \ldots = 0$

Design: FTK completes global tracking in 25 μsec at 3×10^{34}.

Current level-2 takes 25 msec per jet or lepton at 3×10^{34}.

ATLAS L1 Track Trigger Design Options for Phase 2

Region Of Interest based Track Trigger at L1

- **uses ROIs from L1Calo & L1Muon to seed track finding**
- **has a large impact on the Trigger architecture**
  - requires significantly lengthened L1 pipelines and fast access to L1Calo and L1Muon ROI information
  - could also consider seeding this with an early ("Level-0”) trigger, or sending a late ("Level-1.5") track trigger
- **smaller impact on Silicon readout electronics**

Self-Seeded Track Trigger at L1

- **independent of other trigger information**
- **has a large impact on Silicon readout electronics**
  - requires fast access to Silicon detector data at 40 MHz
- **smaller impact on the Trigger architecture**

ATLAS Self-Seeded L1 Track Trigger with Doublet Layers

- Moderate pT dependent discrimination of hits using coincidences in closely spaced double layers
- High pT discrimination using coincidences between several doublet layers
- Has to operate at full BCO frequency (40 MHz)

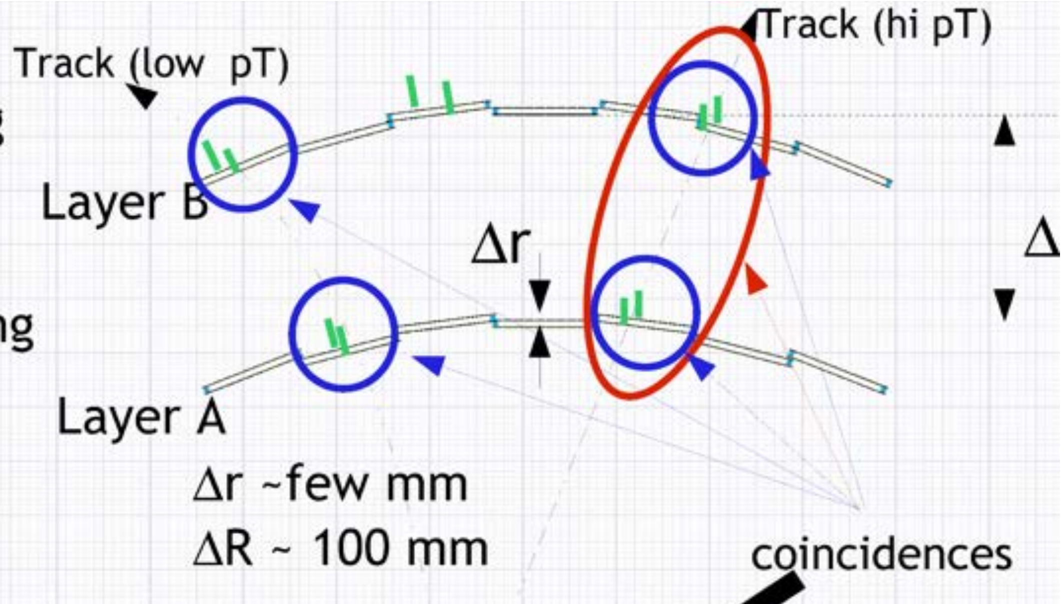

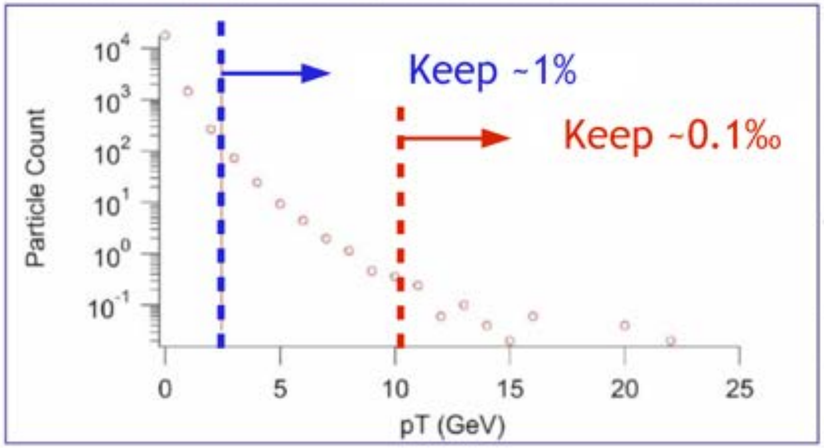

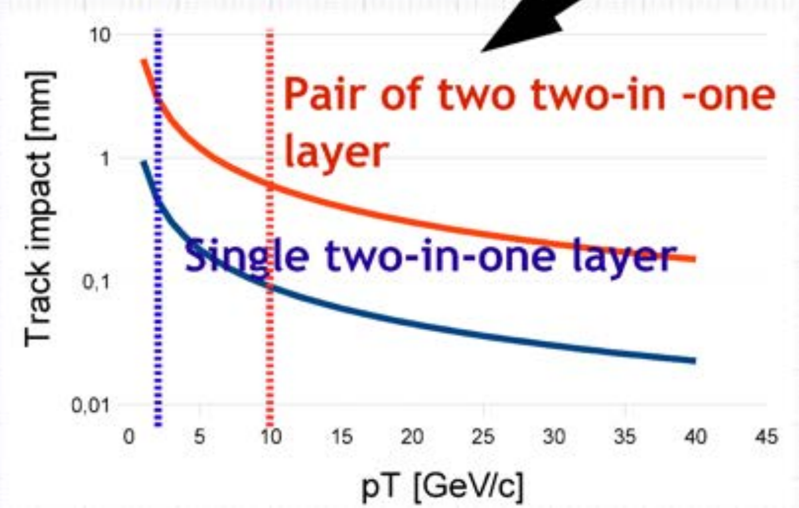

LHCb Upgrade Trigger & DAQ

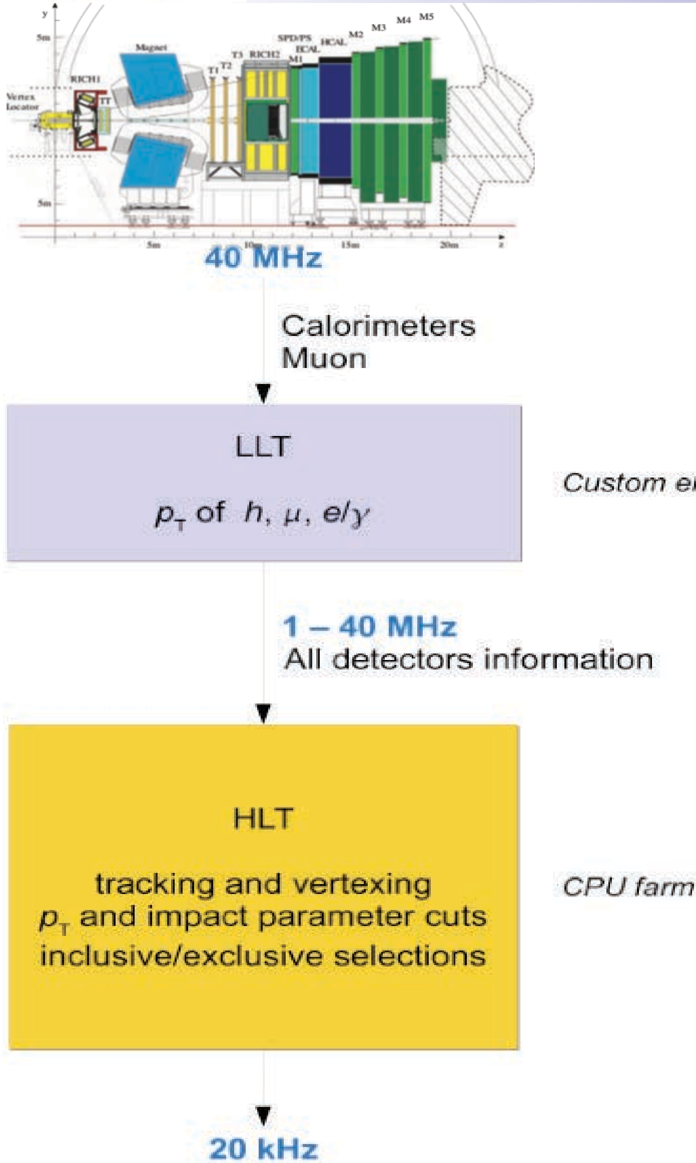

Execute whole trigger on CPU farm

➔Provide ~40 MHz detector readout

- **Cannot satisfy present 1 MHz requirement w/o deeply cutting into efficiency for hadronic final states**
  - worst state is $\phi\phi$, but all hadronic modes are affected
  - Can ameliorate this by reading out detector & then finding vertices

Upgrade Trigger & DAQ

- **flexible software trigger with up to 40 MHz input rate and 20 kHz output rate**
- **run at ~ 5-10 times nominal LHCb luminosity → L ~ 1-2 · 10^{33} cm^{-2} s^{-1}**
- **big gain in signal efficiency (up to x7 for hadron modes)**
- **upgrade electronics & DAQ architecture**
- **collect ≥ 5/fb per year and ~ 50/fb in 10 years**

ALICE Upgrade

Run at high rates, 50 kHz Pb-Pb (*i.e.* L = 6×10^{27} $cm^{-1}s^{-1}$), with minimum bias (pipeline) readout (max readout with present ALICE set-up ~500Hz)

- Factor 100 increase in recorded luminosity
- Improve vertexing and tracking at low p_t

Pb-Pb run complemented by p-Pb & pp running

Entails building High-rate upgrade for readout of TPC, TRD, TOF, CALs, Muons, DAQ/HLT

Two HLT scenarios for the upgrade:

- Partial event reconstruction (clustering and tracking): Factor of ~20 → Rate to tape: 20 kHz
  - clusters (associated with tracks) information recorded on tape
- Full event reconstruction: additional reduction factor ~3 → Rate to tape > 50 kHz
  - track parameters recorded on tape

ALICE Upgrade Readout & Online Systems Architecture

Readout scheme

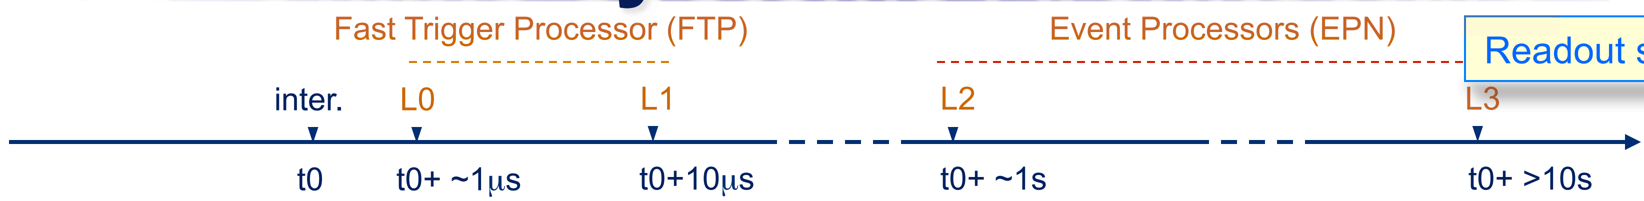

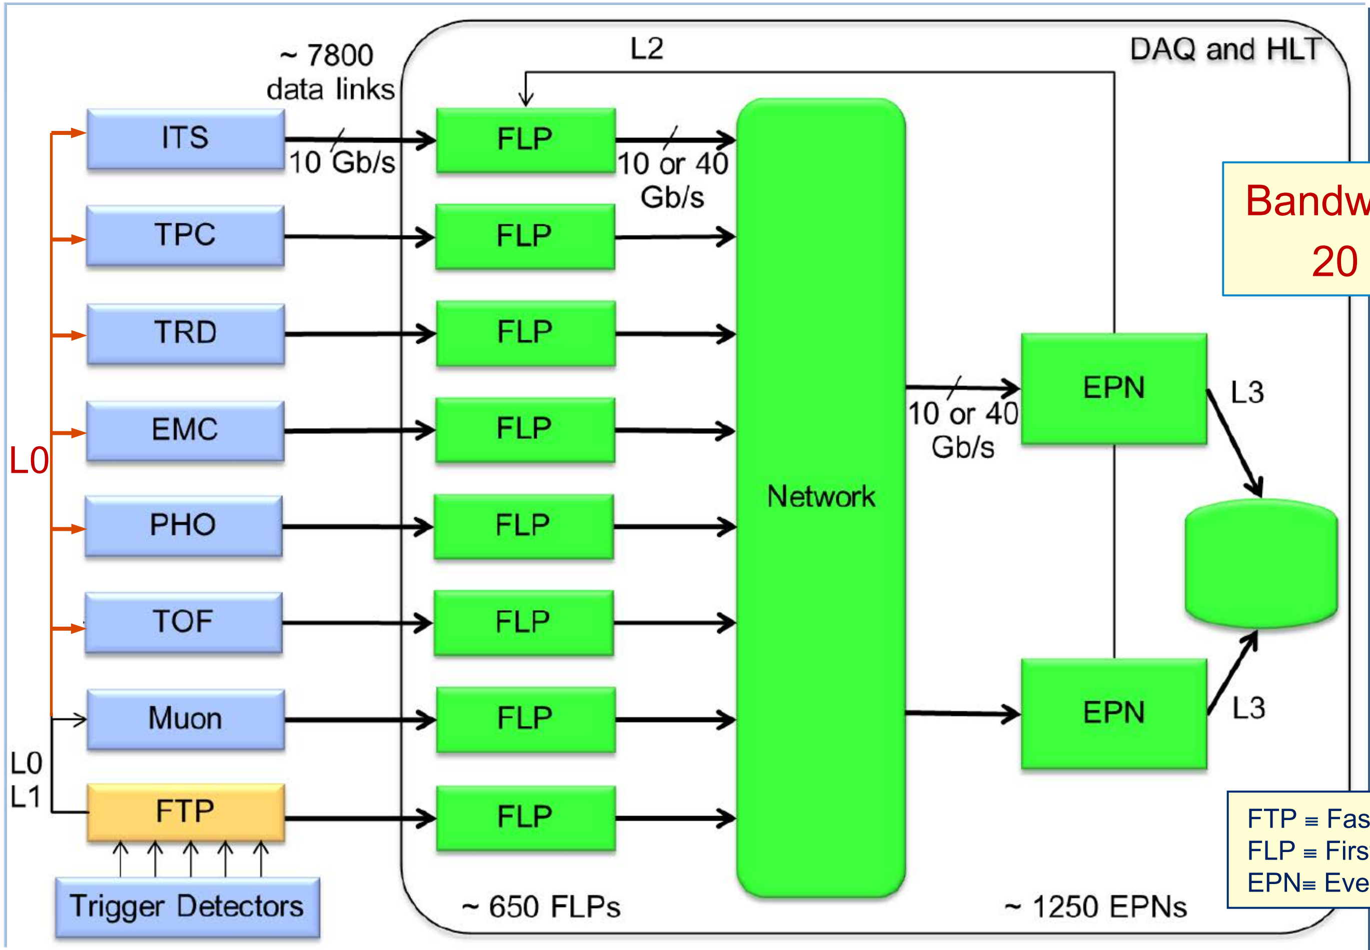

Bandwidth to tape
20 Gbyte/s

FTP ≡ Fast Trigger Processor
FLP ≡ First Level Processor
EPN≡ Event processor Node

LHC Online Systems Summary

- Very significant challenges to operate trigger & DAQ systems for high rate experiments.
- Very substantial assets to bring to bear on these challenges from commercial world: ATCA, FPGAs, high speed links (transceivers), optical connectors …
- Exploiting these assets enables physics input to drive much more precise selection of events and processing of a much higher volume of data.
  - e.g. a level-1 tracking trigger for ATLAS & CMS
- There is considerable technical difficulty involved in successfully exploiting these advances in technology and implementing them in running experiments in a controlled and adiabatic manner.

Backup

LHC Luminosity until 2021: Phase 1

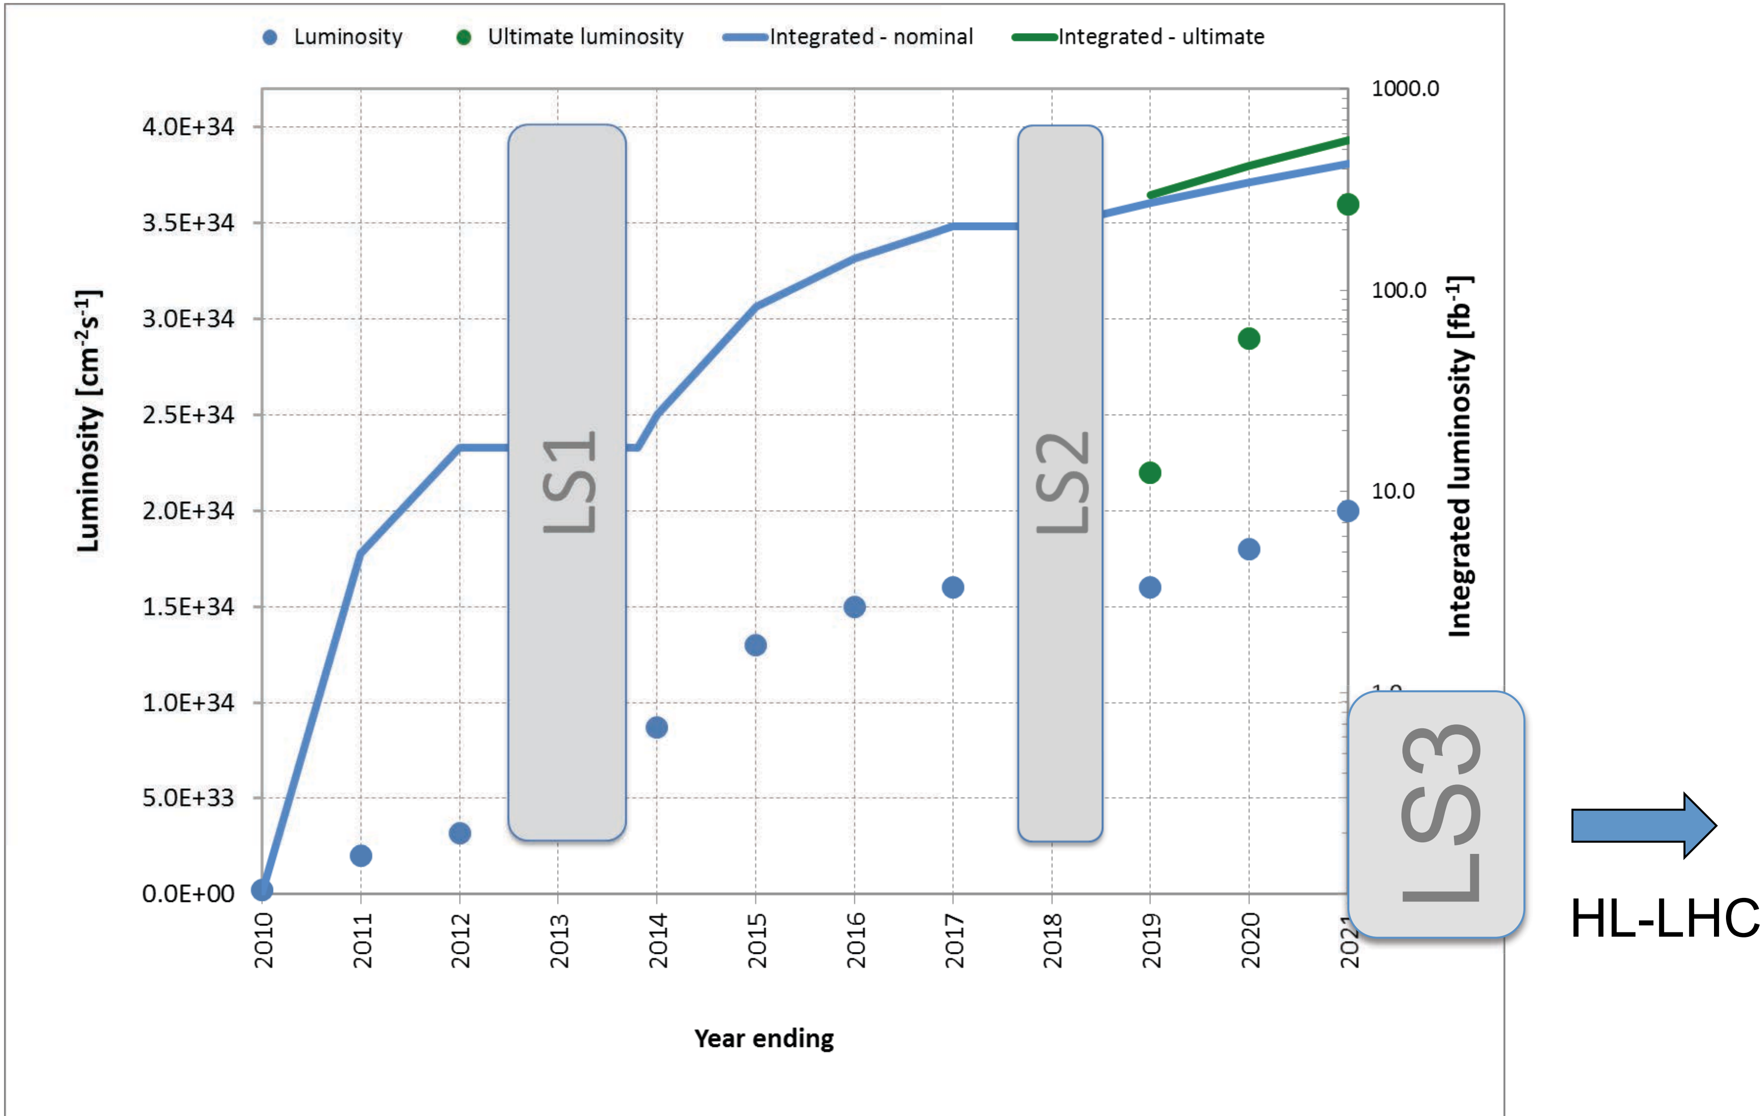

HL-LHC: 2023 onward: Phase 2

3000 fb^{-1} in a 10-12 year run

Peak Luminosity: 7×10^{34}

Pile-up = 95 events @ 25 ns xing freq.

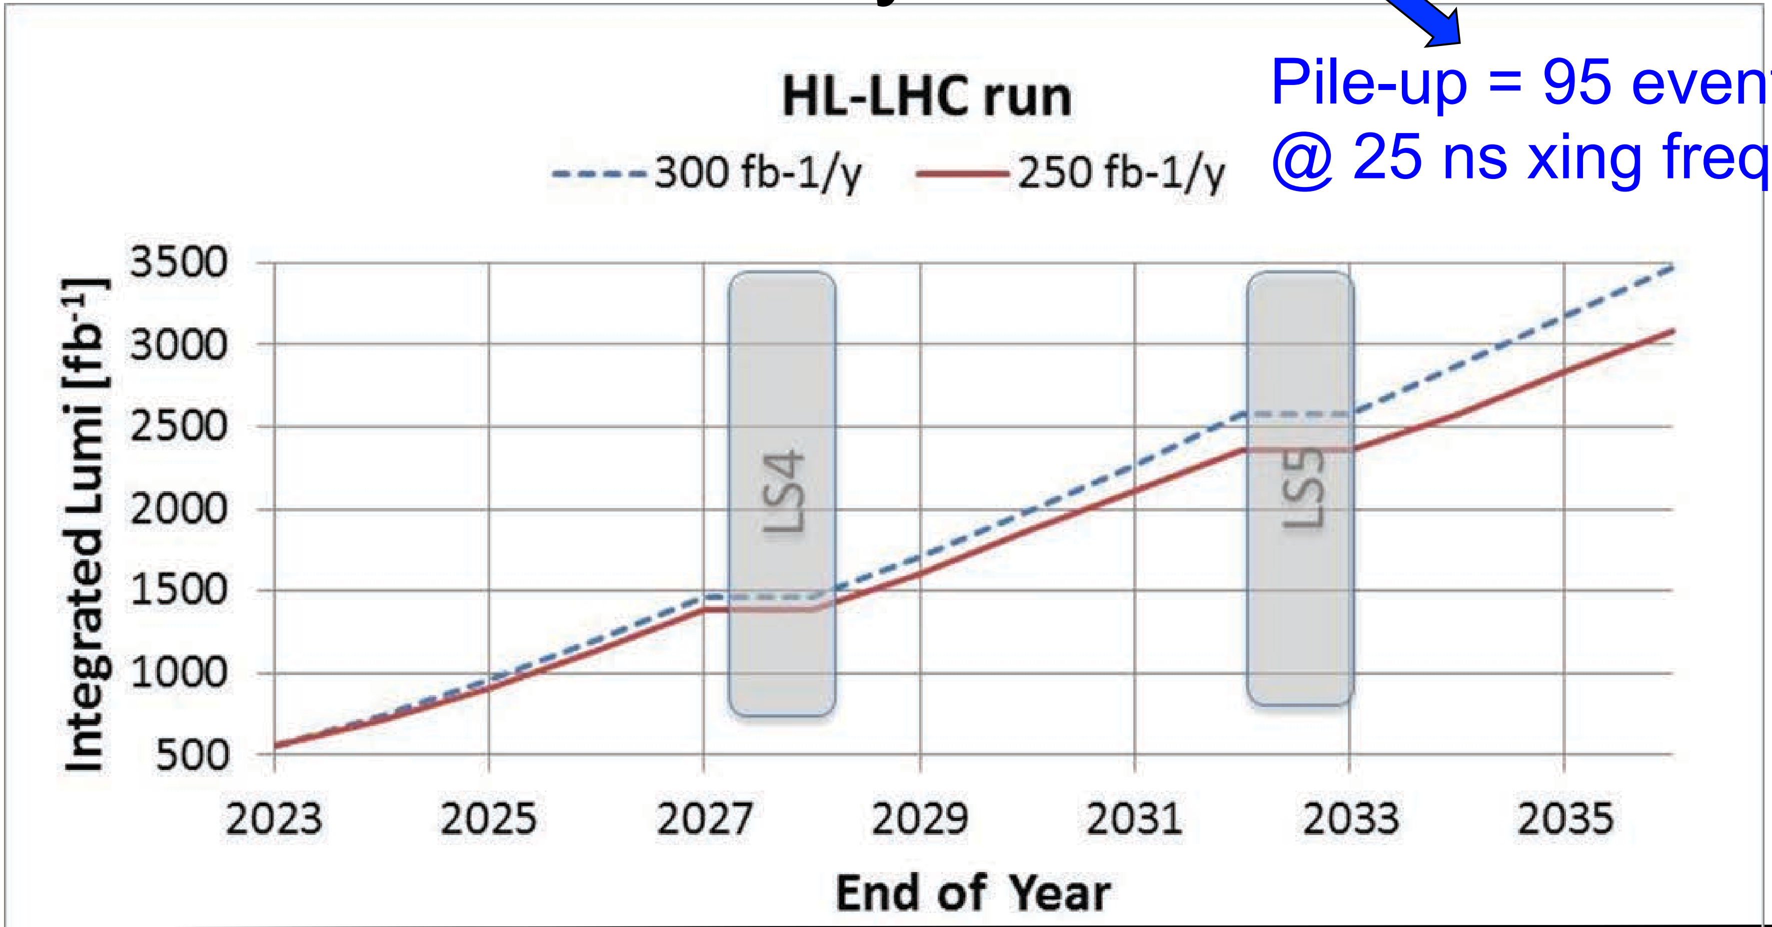

The Track Trigger Problem

- Need to gather information from 10^8 pixels in $200m^2$ of silicon at 40 MHz
- Power & bandwidth to send all data off-detector is prohibitive
  - Local filtering necessary
  - Smart pixels needed to locally correlate hit P_t information
- One option is to use 3D electronics to provide ability to locally correlate hits between two closely spaced layers

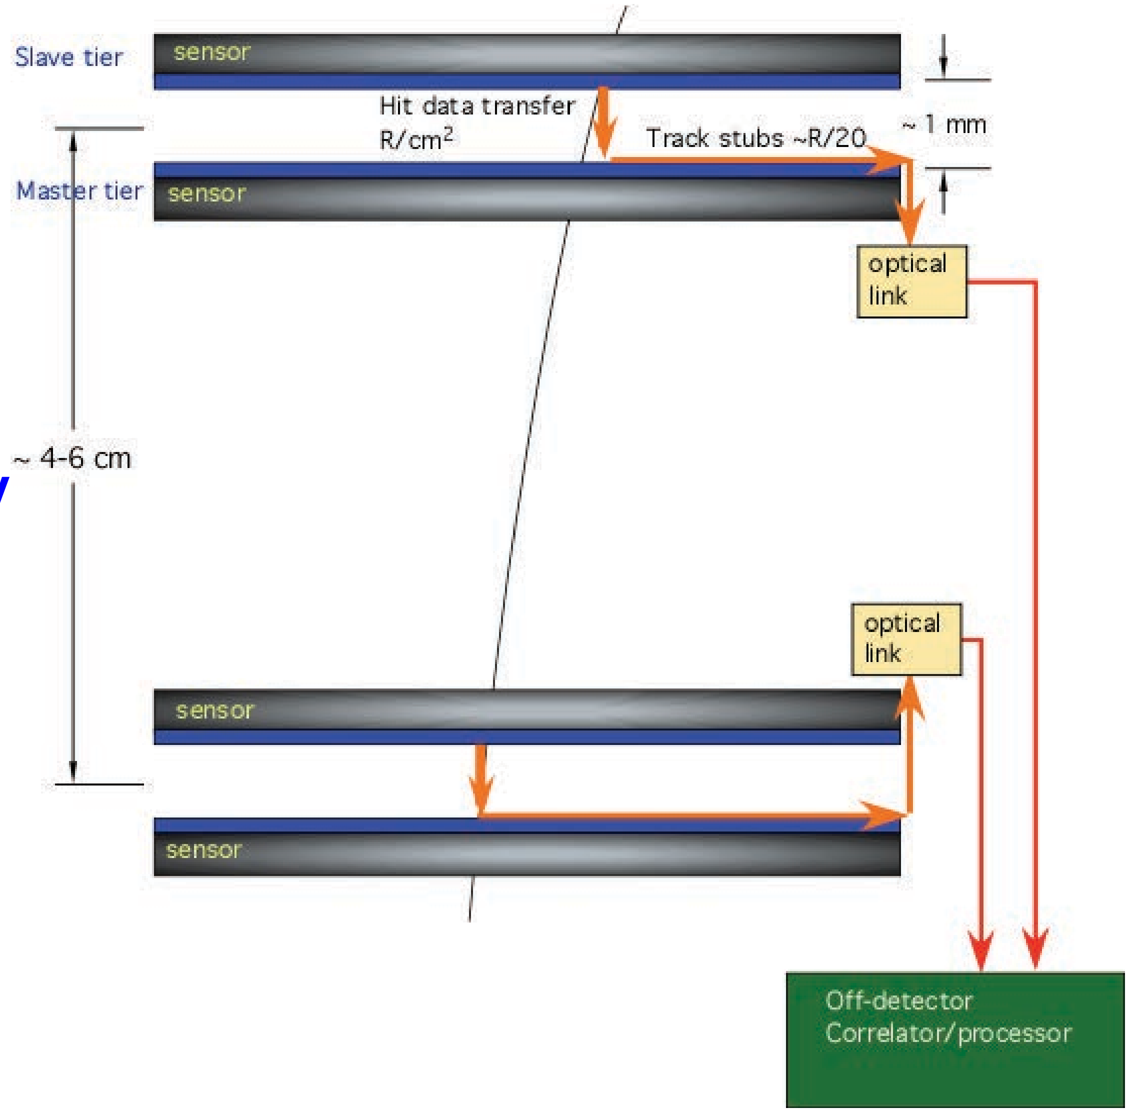

Tracking for electron trigger

Present CMS electron HLT

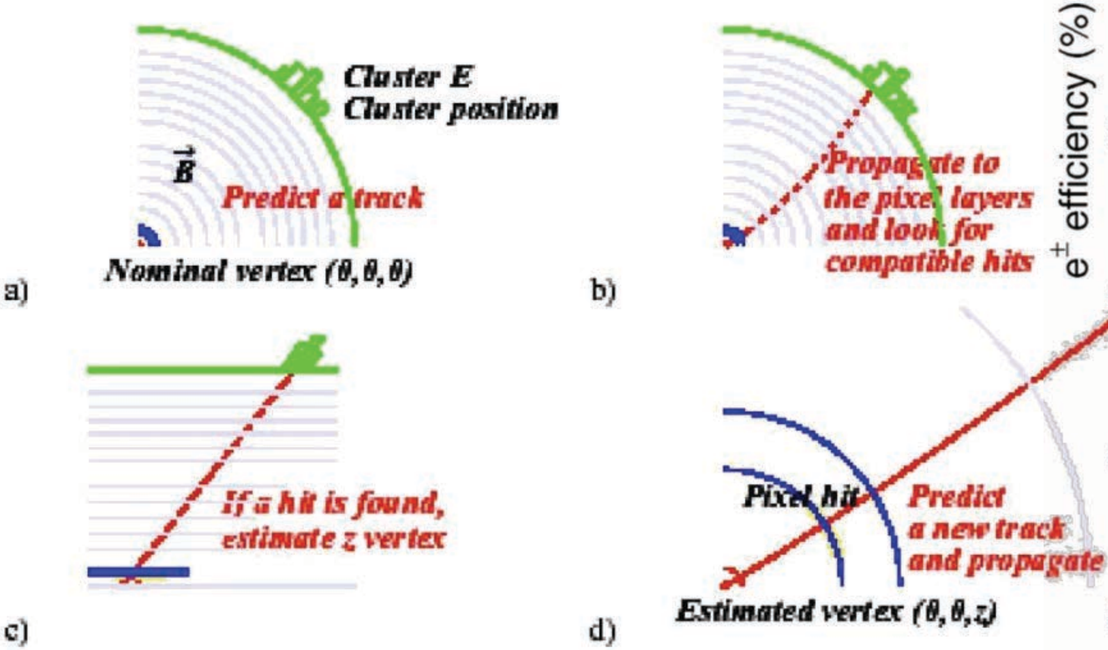

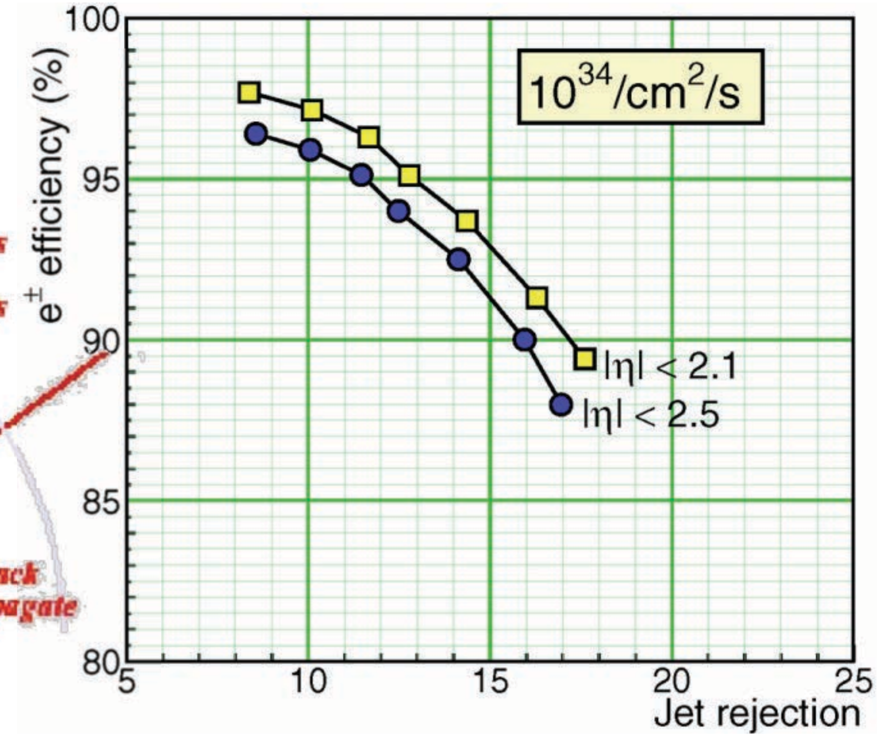

Factor of 10 rate reduction

γ: only tracker handle: isolation

- **Need knowledge of vertex location to avoid loss of efficiency**

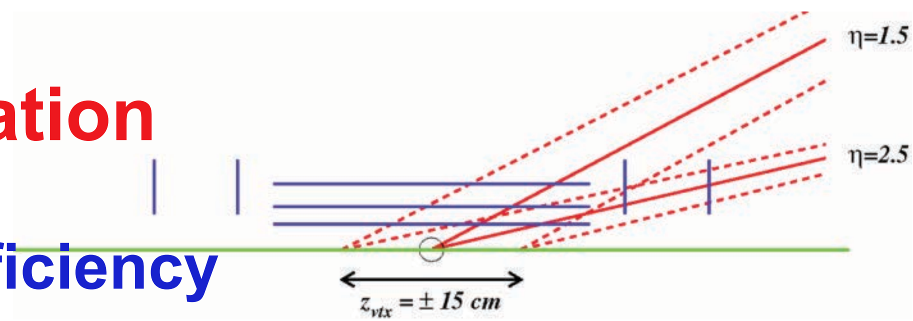

Tracking for τ-jet isolation

τ-lepton trigger: isolation from pixel tracks outside signal cone & inside isolation cone

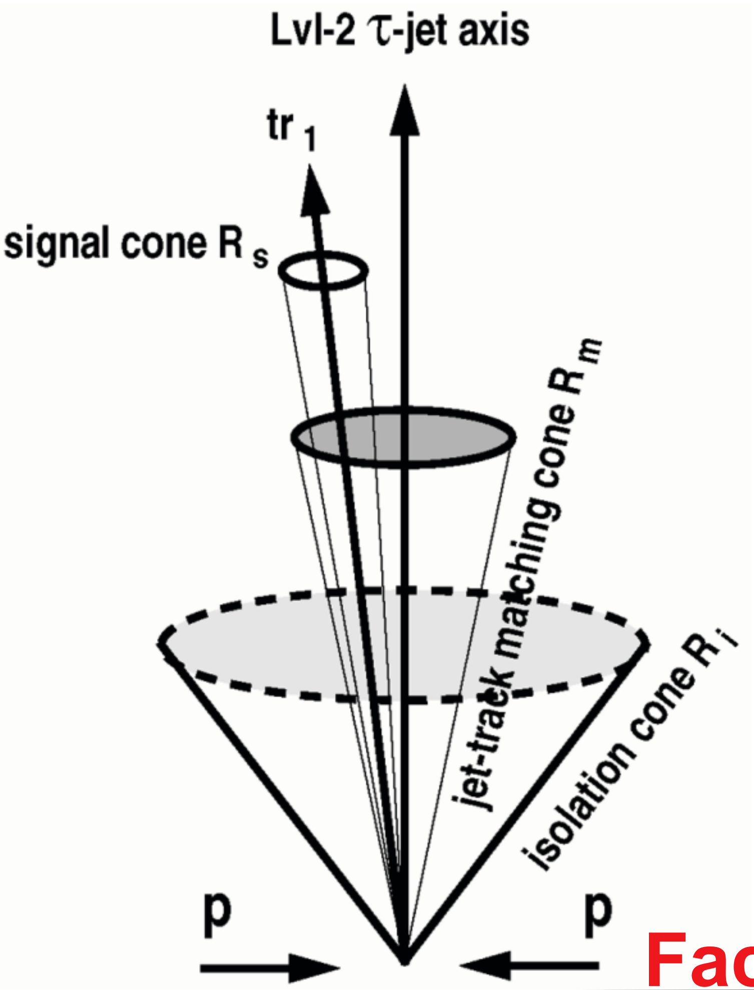

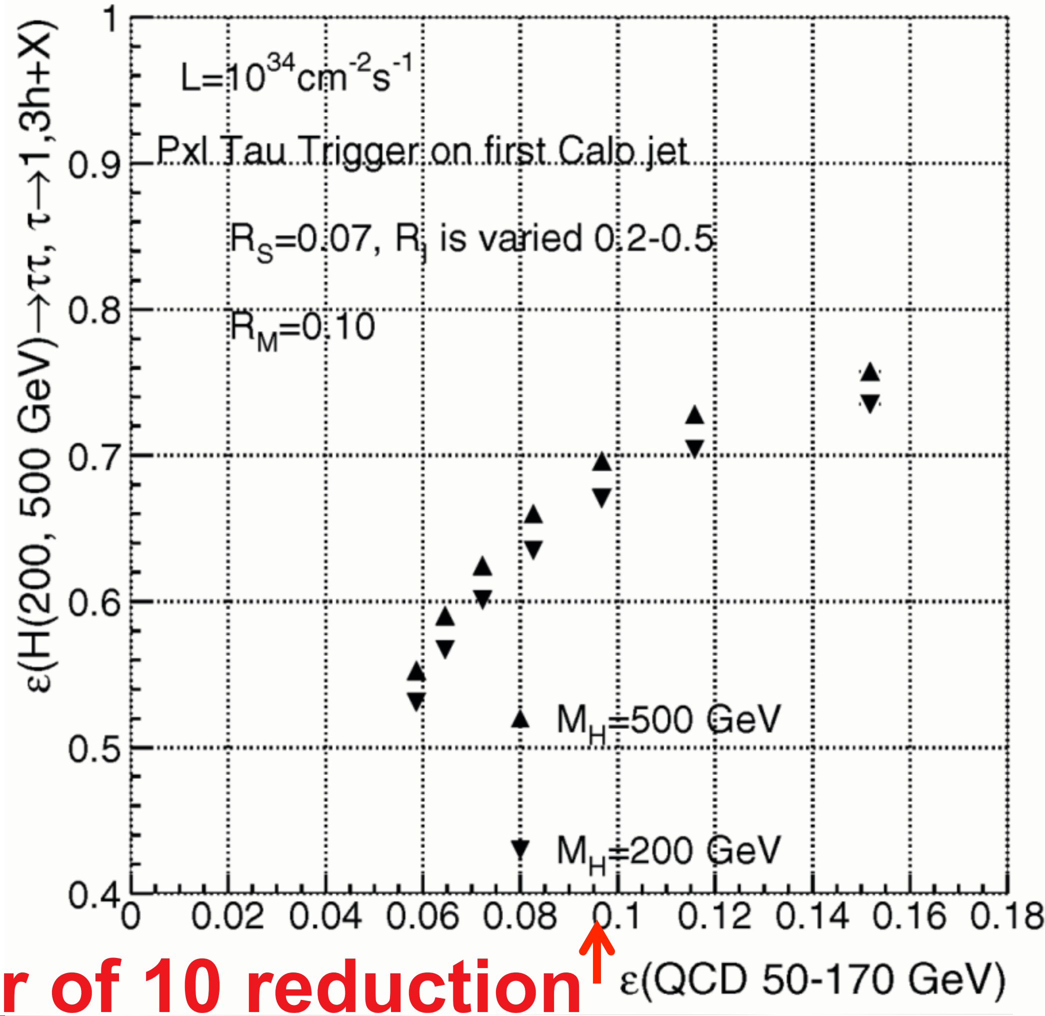

Factor of 10 reduction

CMS L1 Track Trigger for Muons

Combine with L1 μ trigger as is now done at HLT:

- **Attach tracker hits to improve P_T assignment precision from 15% standalone muon measurement to 1.5% with the tracker**
  - Improves sign determination & provides vertex constraints
- **Find pixel tracks within cone around muon track and compute sum P_T as an isolation criterion**
  - Less sensitive to pile-up than calorimetric information *if* primary vertex of hard-scattering can be determined (~100 vertices total at SLHC!)

To do this requires η-ϕ information on muons finer than the current 0.05-2.5°

- **No problem, since both are already available at 0.0125 and 0.015°**