ECFA - TDOC

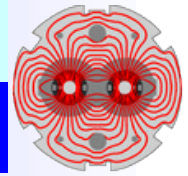

Trigger and Data Acquisition at the HL-LHC

Wesley H. Smith

U. Wisconsin – Madison

On behalf of the Trigger/DAQ/Offline/Computing Preparatory Group

ECFA High Luminosity LHC Experiments Workshop

October 3, 2013

TDOC Membership

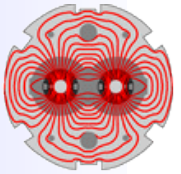

ALICE: Pierre Vande Vyvre, Thorsten Kollegger, Predrag Buncic

ATLAS: David Rousseau, Benedetto Gorini, Nikos Konstantinidis

CMS: Wesley Smith, Christoph Schwick, Ian Fisk, Peter Elmer

LHCb: Renaud Legac, Niko Neufeld

Journey to HL-LHC

2012-2013 run:

- Lumi = 7 x 10^{33}, PU = 30, E = 7 TeV, 50 nsec bunch spacing
- 2012 ATLAS, CMS operating:
  - L1 Accept ≤ 100 kHz,
  - Latency ≤ 2.5 (AT), 4 μsec (CM)
  - HLT Accept ≤ 1 kHz

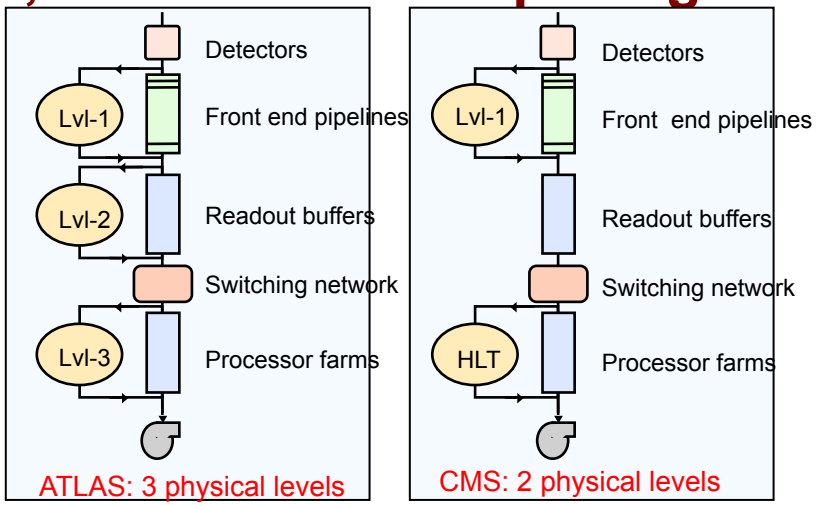

Where ATLAS & CMS will be:

- Lumi = 5 x 10^{34}
- <PU> = 140, Peak PU = 192 (increase × 6)
- E = 14 TeV (increase × 2)
- 25 nsec bunch spacing (reduce × 2)
- Integrated Luminosity > 250 fb^{-1} per year

Need to establish scenario for L1 Accept, Latency, HLT Accept & new trigger "features" (e.g. tracking trigger)

LHC Experiment Scenarios

An introductory Summary

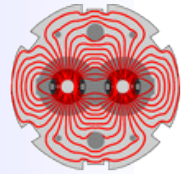

ALICE (post-LS2):Triggerless

- Readout 50 kHz Pb-Pb (*i.e.* L = $6x10^{27}$ $cm^{-1}s^{-1}$), with minimum bias (pipeline) readout (max readout at present ~500 Hz)

ATLAS (post-LS3):Triggered

- Divide L1 Trigger into L0, L1 of latency 6, 20 μsec, rate ≥ 500, ≥ 200 kHz, HLT output rate of 5 - 10 kHz
  - Calorimeter readout at 40 MHz w/backend waveform processing (140 Tbps)
- L0 uses Cal. & μ Triggers, which generate track trigger seeds
- L1 uses Track Trigger & more muon detectors & more fine-grained calorimeter trigger information.

CMS (post LS3):Triggered

- Considering L1 Trigger latency, rate: 10 – 20 μsec, 0.5 – 1 MHz
- L1 uses Track Trigger, finer granularity μ & calo. Triggers
- HLT output rate of 10 kHz

LHCb (post LS2):Triggerless

- Execute whole trigger on CPU farm ⇒ ~40 MHz readout

ATLAS & CMS Triggered vs. Triggerless Architectures

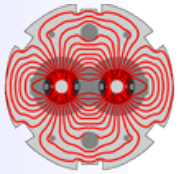

1 MHz (Triggered):

- **Network:**
  - 1 MHz with ~5 MB: aggregate ~40 Tbps
  - Links: Event Builder-cDAQ: ~ 500 links of 100 Gbps
  - Switch: almost possible today, for 2022 no problem
- **HLT computing:**
  - General purpose computing: 10(rate)x3(PU)x1.5(energy)x200kHS6 (CMS)
    - **Factor ~50 wrt today maybe for ~same costs**
  - Specialized computing (GPU or else): Possible

40 MHz (Triggerless):

- **Network:**
  - 40 MHz with ~5 MB: aggregate ~2000 Tbps
  - Event Builder Links: ~2,500 links of 400 Gbps
  - Switch: has to grow by factor ~25 in 10 years, difficult
- **Front End Electronics**
  - Readout Cables: Copper Tracker! – Show Stopper
- **HLT computing:**
  - General purpose computing: 400(rate) x3(PU)x1.5(energy)x200kHS6 (CMS)
    - **Factor ~2000 wrt today, but too pessimistic since events easier to reject w/o L1**
    - **This factor looks impossible with realistic budget**
  - Specialized computing (GPU or …)
    - **Could possibly provide this …**

Trigger Challenges at HL-LHC: ATLAS & CMS

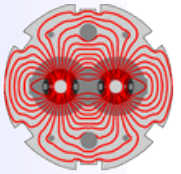

Goals:

- **Study with high precision properties of Higgs with focus on self-couplings and precision measurements of couplings**
  - Keep trigger acceptance for Higgs at least as high as in 2012.
- **Keep same sensitivity for SUSY and Exotic searches as in 2012.**

Challenges:

- **Higher Interaction Rates**
  - For physics of interest and backgrounds!
  - ~ 6k primary tracks per bunch crossing within $|\eta|<2.5$ plus conversions and nuclear interactions ~ one order of magnitude larger than 2012
- **Occupancy causes degraded performance of algorithms**
  - Electrons: reduced rejection at fixed efficiency from isolation
  - Muons: increased background rates from accidental coincidences
- **Implies raising E_T thresholds on electrons, photons, muons, jets and use of less efficient multi-object triggers, unless we have new information ⇒Tracker at L1**
  - Compensate for larger interaction rate & degradation in algorithm performance

ATLAS & CMS L1 Tracking Trigger

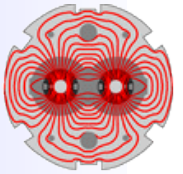

Reduces Leptonic Trigger Rate

- Validate calorimeter or muon trigger object, e.g. discriminating electrons from hadronic ($\pi^0 \rightarrow \gamma\gamma$) backgrounds in jets
- Addition of precise tracks to improve precision on p_T measurement, sharpening thresholds in muon trigger
- Degree of isolation of e, γ, μ or τ candidate
- Requires calorimeter trigger trigger at the finest granularity to reduce electron trigger rate

Other Triggers

- Primary z-vertex location within 30 cm luminous region derived from projecting tracks found in trigger layers,
- Provide discrimination against pileup events in multiple object triggers, e.g. in lepton plus jet triggers.

HL-LHC Track Trigger Architectures:

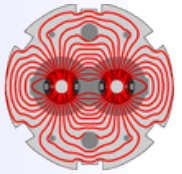

"Push" path (CMS Tracker Approach):

- 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 (ATLAS Tracker Approach):

- L1 calorimeter & muon triggers produce a "Level-0" or L0 "pre-trigger" after latency of present L1 trigger, with request for tracking info at ~0.5 MHz. Request only goes to regions of tracker where candidate was found. Reduces data transmitted from tracker to L1 trigger logic by < 80 (40 MHz to > 0.5 MHz) times probability of a tracker region to be found with candidates, which could be less than 10%, (e.g. 50 kHz, < speed of ATLAS FTK)
- Tracker sends out info. for these regions only & this data is combined in L1 correlation logic, resulting in L1A combining track, muon & cal. info..

"Afterburner" path (both ATLAS & CMS):

- 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.

ATLAS "Double buffer" readout

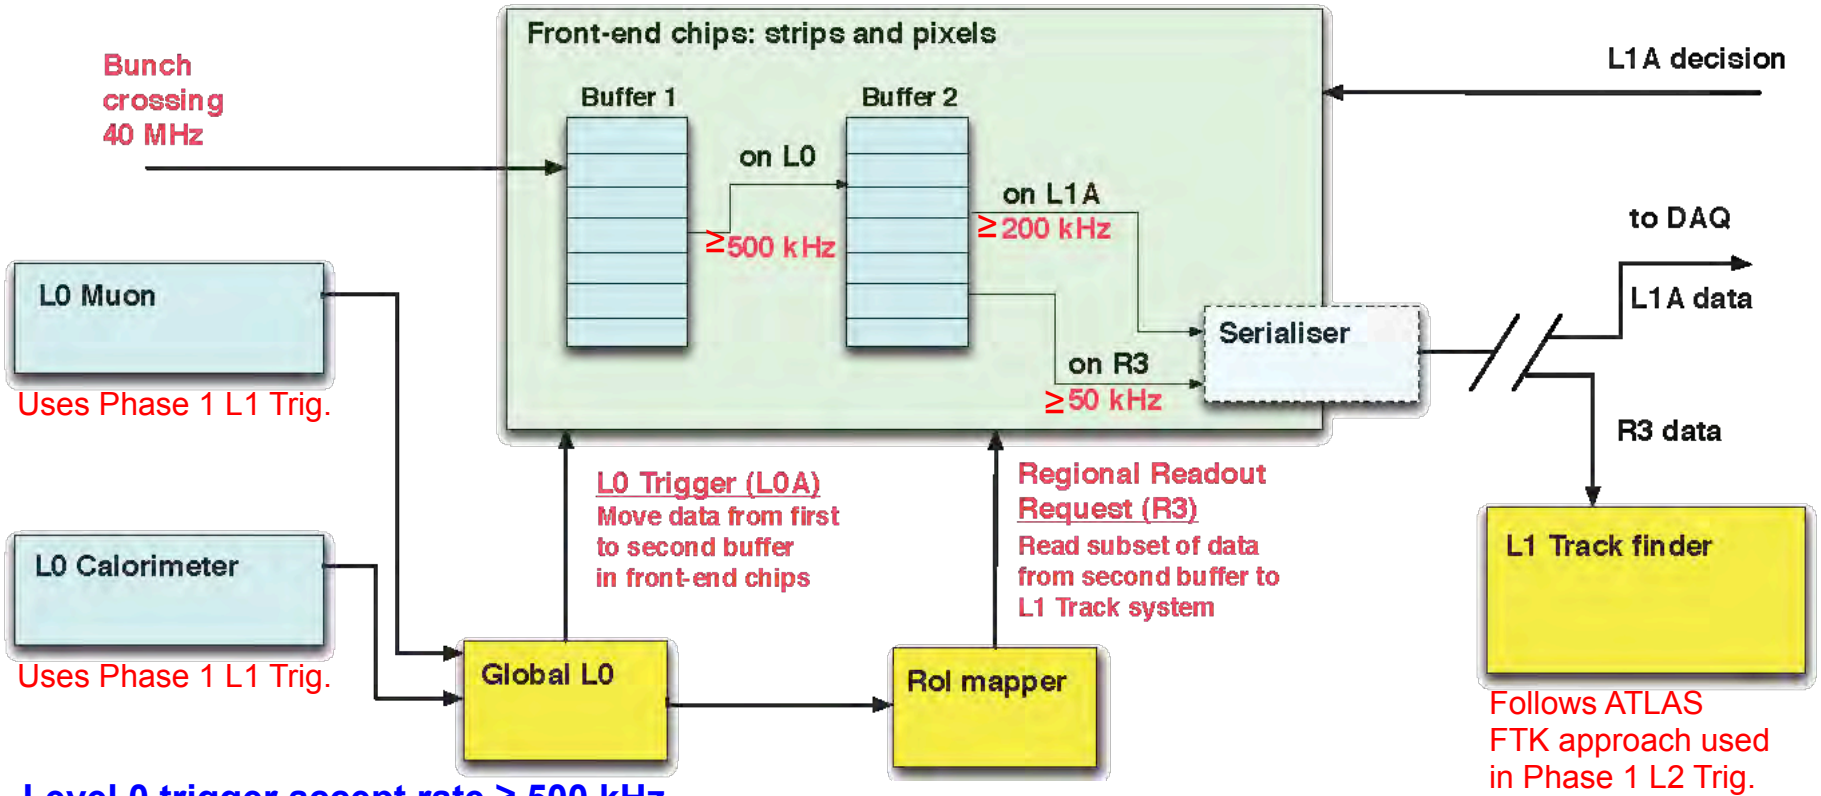

- **Level 0 trigger accept rate ≥ 500 kHz**
  - **On an L0 accept, copy data from primary to secondary buffer**
  - **Identify "Regions" in detector (1-10% of the detector on each L0 accept) like L1 RoI**
  - **Generate "Regional Readout Request" (R3) - modules in "Region" read out subset of their data**
- **On an L1 accept (≥ 200 kHz), all modules read out event from Secondary buffer**
- **Since only ~10% of the detector (the "Regions") will be read out on the Level 0 accept, R3 request rate for any specific part of the detector will be ≥ 50 kHz**

ATLAS Gains from Track Trigger

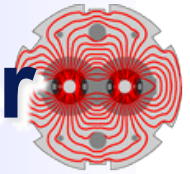

- **Matching tracks to Level 1 objects (electrons, taus and muons) can significantly reduce rate**
  - **Remove mis-reconstructed or fake objects**
  - **Ensure objects come from the vertex**
- **Potential benefits have been studied for electrons, muons, taus and jets, single and multiple/combined object triggers using both smeared offline tracks, and smeared truth particles**
  - **Even modest resolution tracking information (p_T, η, φ) can provide sufficient rejection**
  - **Factors of between 3 and 5 for electrons, taus, muons (p_T > 20) with only small efficiency losses (~5%) wrt. Phase 1 Trigger system.**
    - Taus on next slide

ATLAS Gain from Track Trigger

Rate vs. tau finding efficiency curves for taus from the decay of a 120 GeV Higgs boson for the inclusive tau trigger at $7x10^{34}$ cm^{-2} s^{-1} for different track multiplicity and minimum track p_T requirements.

The bands show the rate vs. efficiency parameterized for different L1 cluster E_T thresholds, shown as the small numbers next to the corresponding points on each band.

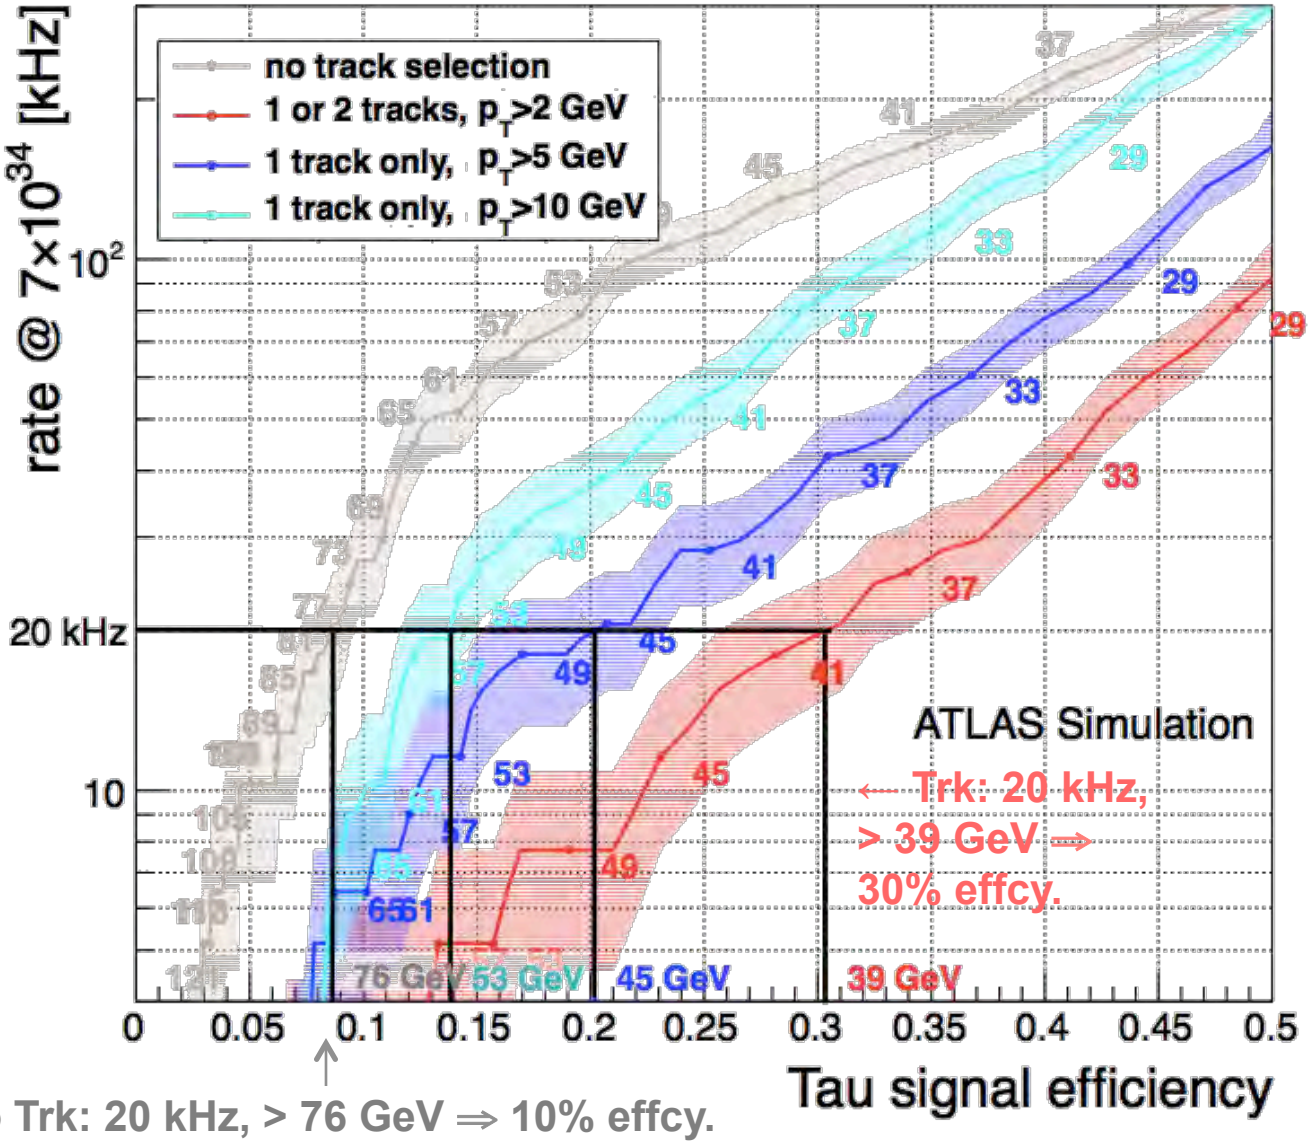

CMS Phase 2 Trigger Scenario

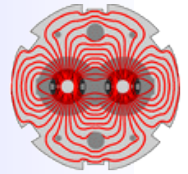

Replace ECAL Barrel and Endcap Front End electronics

- Allows L1 latency & accept rate increases (below)
- Includes providing individual crystal level (not 5x5 sums) trigger information
  - Resolution based on $\Delta\eta \times \Delta\phi = 0.087 \times 0.087 \rightarrow 0.017 \times 0.017$
  - Improved spike rejection in EB
- Assume: EE electronics replaced with EE replacement

Latency of 10 μsec

- Limit from Endcap Muon Cathode Strip Chamber Front End Electronics

L1 Accept rate of up to 1 MHz (expect to operate at 500 kHz)

- Provides more acceptance and lower thresholds
- Limit provided by DAQ readout, EVB, & HLT CPU
- Requires: Drift Tube Readout Electronics replacement (planned)

Tracking Trigger

- Leptons (P_T cut & isolation)
- Vertex for jets

New L1 Trigger (Calorimeter, Muon, Global) to incorporate Track Trigger

- Finer calorimeter cluster trigger, muon & calorimeter seeds for track match
- Also incorporate additional muon chambers for $1.6 < |\eta| < 2.4$.

HLT Output Rate of 10 kHz

- Limit from Downstream Computing

CMS Tracking Trigger

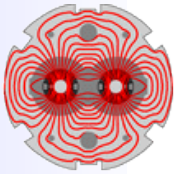

Outer Tracker "Baseline"

- Lighter Tracker, with better overall Tracking and Calorimetry performance compared to the present systems
- Level-1 Tracking Trigger including all tracks with p_T > 2 GeV, well measured & with ~ 1mm primary vertex resolution
- Pursuing a "Push" Architecture based on
  - **Module filtering of hits from tracks with p_T above ~ 2 GeV**
  - **Low power (low mass) 5 GHz optical links**
  - **Lower latency, less hits produced up front**

Inner Pixel Option

- **Usable for B-tags, Taus, c, electron-ID, added vertex info.**
- **Exploring a Region of Interest "Pull" architecture**
- **As a possible complement to the L1 "Push" Tracking Trigger and/or HLT pre-processor**

CMS Gains from Track Trigger

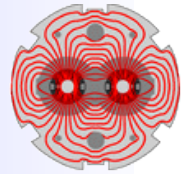

Preliminary simulation studies demonstrate addition of L1 tracking trigger provides significant gains in rate reduction with good efficiency for physics objects. Note these results are "work in progress".

Trigger, Threshold	Algorithm	Rate reduction	Full eff. at the plateau	Comments
Single Muon, 20 GeV	Improved Pt, via track matching	~ 13 ($\|\eta\| < 1$)	~ 90 %	Tracker isolation may help further.
Single Electron, 20 GeV	Match with cluster	> 6 (current granularity) >10 (crystal granularity) ($\|\eta\| < 1$)	90 %	Tracker isolation can bring an additional factor of up to 2.
Single Tau, 40 GeV	CaloTau – track matching + tracker isolation	O(5)	O(50 %) (for 3-prong decays)	
Single Photon, 20 GeV	Tracker isolation	40 %	90 %	Probably hard to do much better.
Multi-jets, HT	Require that jets come from the same vertex			Performances depend a lot on the trigger & threshold.

CMS Gains for μ, e Triggers

(over Phase 1 Trigger, |η| < 1)

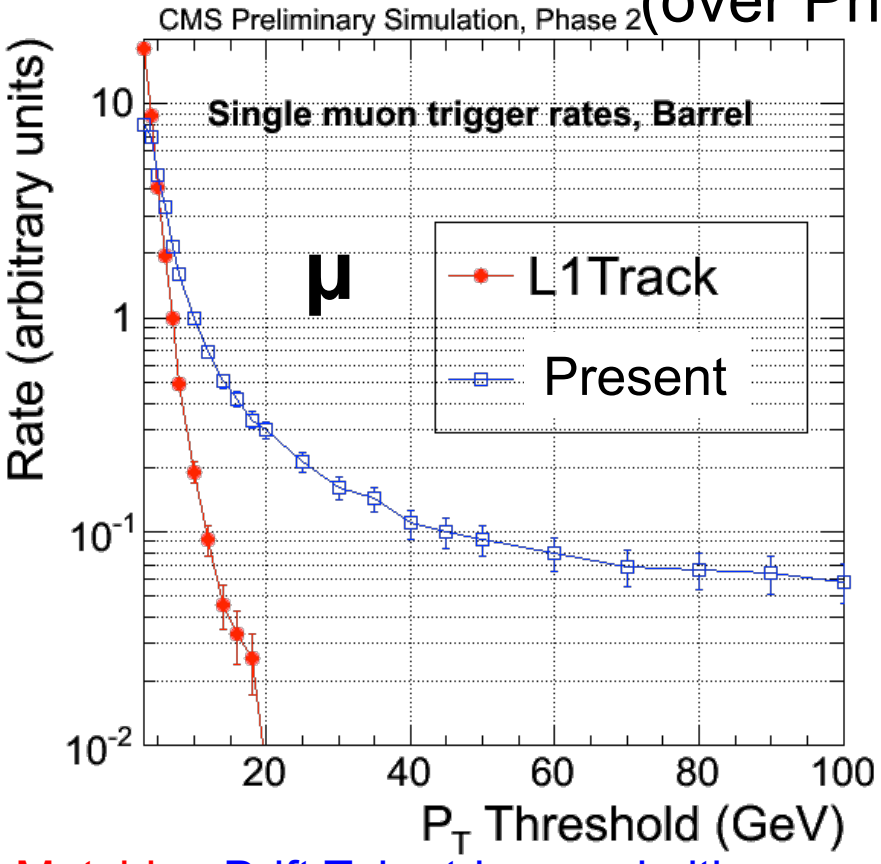

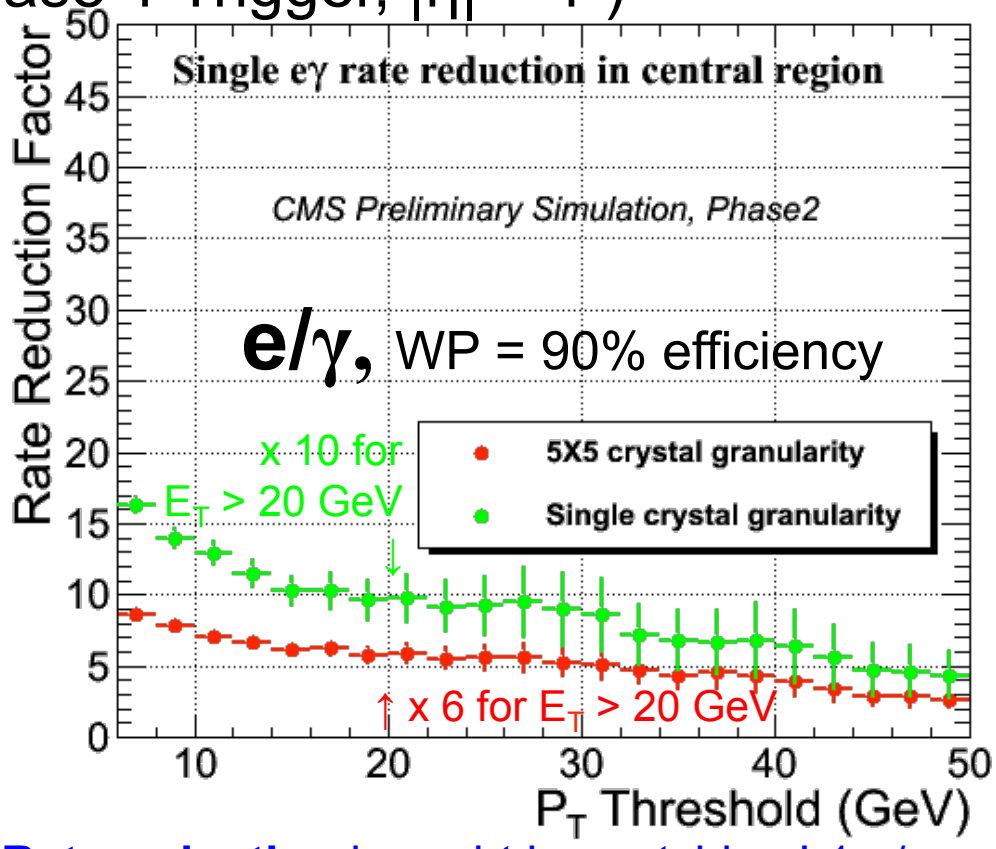

Matching Drift Tube trigger primitives with L1Tracks: **large rate reduction: > 10 at threshold > ~ 14 GeV.** Normalized to present trigger at 10 GeV. Removes flattening at high P_t

Rate reduction brought by matching L1 e/γ to L1Track stubs for | η | < 1.

Red: with current (5x5 xtal) L1Cal granularity.

Green : using single crystal-level position resolution improves matching

CMS HL-LHC L1 Trigger Upgrade

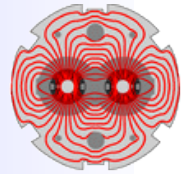

Integration of Track & Pixel Trigger information into L1 Trigger requires upgrade of rest of L1 Trigger

- **Calorimeter trigger should use full information to provide smallest resolution for combination with a tracking trigger**
  - Resolution based on calorimeter readout towers
  - Also improves calorimeter trigger pattern recognition (b
  - Increases input data but can mitigate by compressed input scale, EM pre-clustering, taking advantage of newer technology higher speed links (presently 13 Gbps, guess at least ×2 for 2023)
- **Muon Triggers will need to calculate results on a finer scale for combination with a tracking trigger**
- **Muon triggers may integrate track trigger information into muon track-finders**
- **Global Trigger will be processing coincidences on a finer resolution**

CMS HL-LHC L1 Trig. Latency, Rate

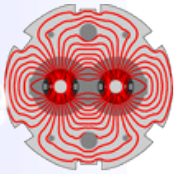

Extended Latency: Simplifies tracking trigger

- **Timing is tight for tracking trigger**
  - Including processing & use of track trigger information
- **Makes design of tracking trigger easier**
  - Relaxed constraints: reduces power, transmission bandwidth…

Extended Latency: Provides option of pixel tracking trigger

- **Pixel trigger requires "pull" architecture**
- **Required for b-tags in L1 Trigger**
  - Along with 0.5-1 MHz L1 bandwidth

Higher Rate: Reduces Thresholds for physics signals

- **Can set thresholds comparable to present ones when coupled with tracking triggers**

Higher Rate: Needed for Hadronic Triggers

- **Track Trigger helps leptonic triggers**
- **Less of an impact on hadronic triggers**
  - Vertex for jets

Higher Rate: Needed for b-tags

- **Pixel trigger may not reduce rate sufficiently**

CMS HL-LHC HLT Output Rate

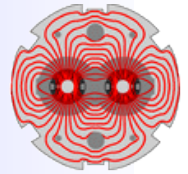

Processing 0.5-1 MHz Input

- **DAQ hardware & HLT processing compatible with Moore's Law scaling until 2023 & estimated x4 longer reconstruction time, event size for PU ≥ 140 (must cope with peak luminosity)**
  - Prediction of HLT CPU time/event = 600 ms at PU=125 (200 now)
    - **Issue of complexity of events passing L1 w/tracking trigger**
- **Use of L1 Track Trigger information as input allows immediate, fast use of tracking information.**
- **Possibility to share resources with Tier-0 (Cloud computing)**
  - Goes both ways
- **If need more CPU, we can bring more online rapidly (if can afford)**

5-10 kHz Output Rate

- **1 MHz L1 Accept Rate → 10 kHz HLT output rate keeps same reduction of L1 rate (x100) as present HLT design (100 kHz → 1 kHz)**
- **Output to Computing**
  - Compatible with Moore's Law scaling (with SW work) until 2023 & estimated X3 longer reconstruction time, event size (avg'd over year)

CMS DAQ after LS3

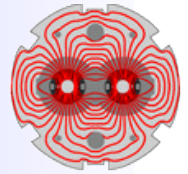

Level 1 rate	1 MHz
Event Size	4 MB
HLT Accept rate (recording rate)	10 kHz
HLT computing power	10 M Hep-Spec-06
EVB architecture	Full event building at L1 rate

Remarks

- **1 MHz L1 rate**
  - allows for flexible physics trigger
  - Feasible for front end electronics
- **Event Size 4MB**
  - Estimated from linear pile-up extrapolation to PU=140
    - **Need simulation work to back up this assumption**
- **HLT accept rate:**
  - Requires factor 100 suppression in HLT as today
- **Computing power: next slides**

CMS Estimation of required HLT CPU power

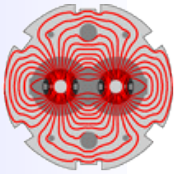

Observation so far

- **Required HLT power scales linearly with pile-up**
  - This has been observed for PU in the range of 10-40
    - **Conservatively assume this continues – needs verification**

Assuming

- **Linear scaling with average PU up to 140**
- **A factor 1.5 due to energy increase to 8 TeV**
  - Also conservative – takes into account complexity of events selected by L1 Trigger scaling with energy
  - Operation after LS1 with 7.5 TeV per beam will quickly allow refining this estimate
- **10 times higher L1 rate**

A total factor of 50 increase of HLT power would be needed wrt. today's farm.

- **This results in 10 M HEP-SPEC-06**

LHCb Upgrade Trigger & DAQ

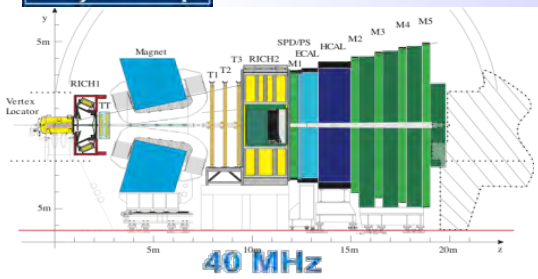

40 MHz

Calorimeters
Muon

LLT

p_T of h, μ, e/γ

Custom electronics

1 – 40 MHz
All detectors information

HLT

tracking and vertexing
p_T and impact parameter cuts
inclusive/exclusive selections

CPU farm

20 kHz

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

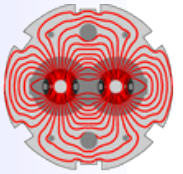

Run at high rates, 50 kHz Pb-Pb (*i.e.* L = 6x10^{27} cm^{-1}s^{-1}), with minimum bias (pipeline) readout (max readout with present ALICE set-up ~ 500 Hz)

- **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

HLT scenario:

- **Partial event reconstruction (clustering and tracking): Factor of ~15 in data reduction → Rate to tape: 50 kHz**
  - clusters (associated with tracks) information recorded on tape

Tools for Triggers: 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…)

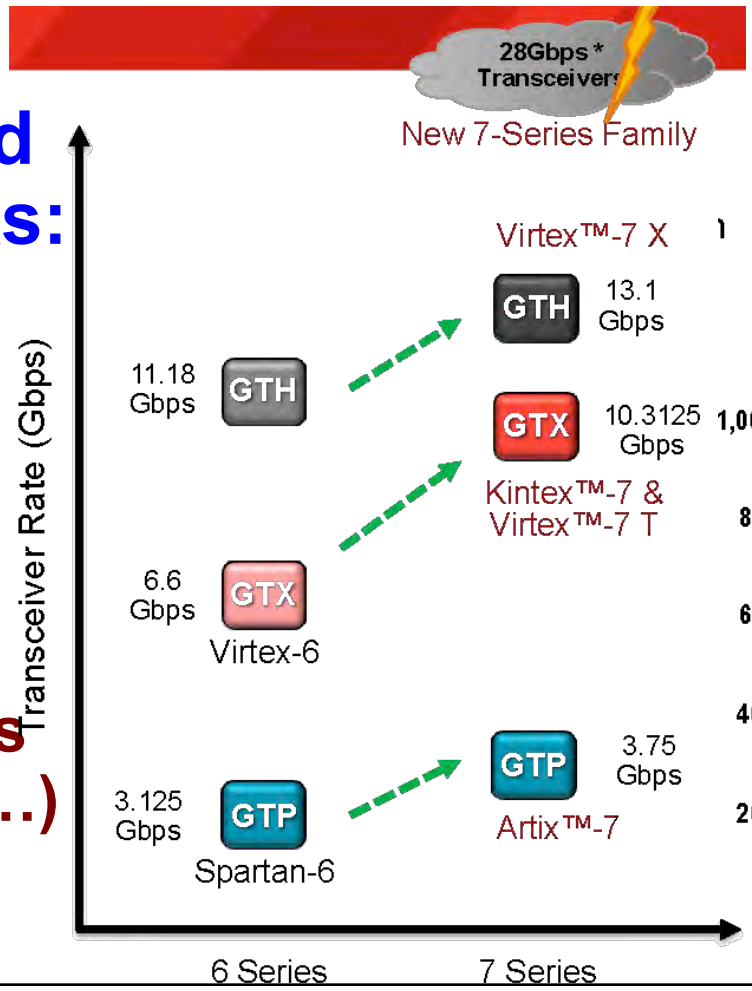

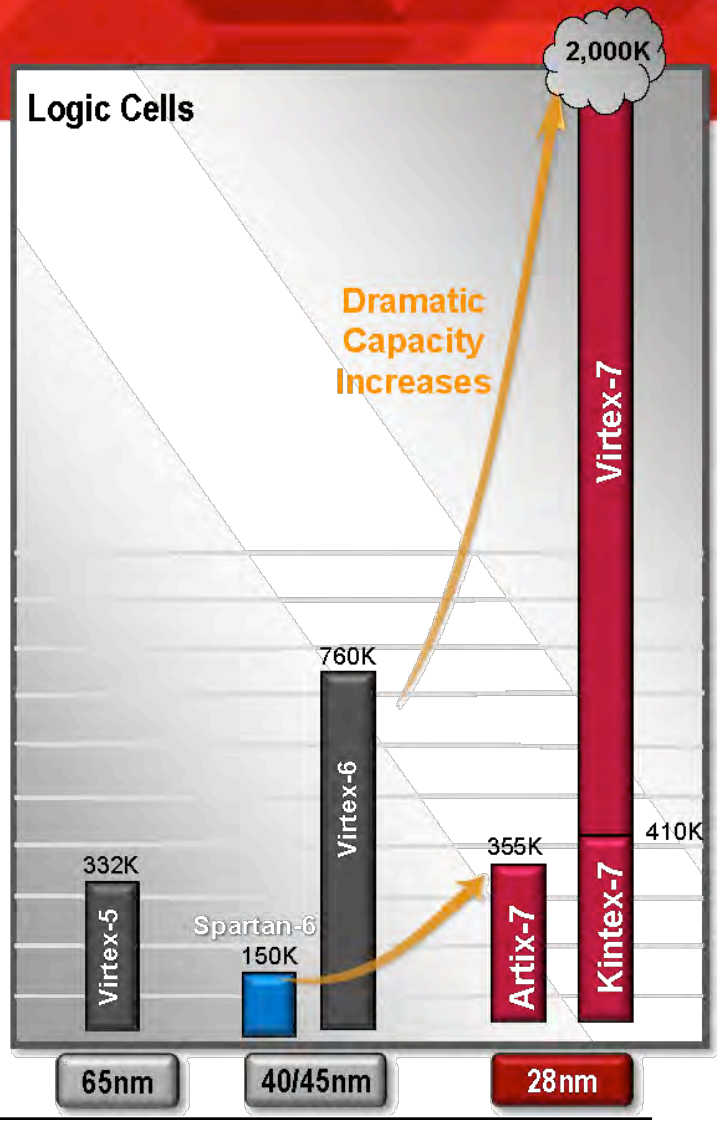

FPGA Challenges/Opportunities

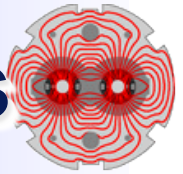

Latest generation FPGAs create complex placement issues that are difficult for Xilinx tool algorithms to resolve

- Build times getting in excess of 24 hours
- Need to perform smart explorer builds to achieve timing closure
- Can use batch systems (e.g. Condor) to perform multiple builds in parallel

Designs must be heavily floor-planned

- Similar to ASIC layout process
- Needs detailed knowledge of routing structure & alternative Xilinx tool flows

Embedded Processors

- Move many tasks from FPGA design to SW design
- Shortens design cycle
- Remove FPGA design integration burden for commodity interface cores
  - Utilize proven and FREE embedded system IP
- Interfaces more flexible under software control
  - Conform to industry standards at core
  - Add software application above core to specialize
- Ensures FPGA design focus on custom logic
  - Custom high speed communication interfaces and custom physics algorithms
- Example: Xilinx ZYNQ
  - Runs PetaLinux (also runs on Microblaze on Xilinx V6 & V7)
  - Write Communications and Control functions in Linux

Tool for Tracking Triggers: Associative Memories

Pattern Recognition Associative Memory (PRAM)

- **Based on CAM cells to match and majority logic to associate hits in different detector layers to a set of pre-determined hit patterns**
  - Example of FTK planned for ATLAS Level 2 Trigger in Phase 1
- **highly flexible/configurable, much less demand on detector design**
- **Pattern recognition finishes soon after hits arrive**
- **Potential candidate for L1 pattern recognition**
- **However: Latency**
- **Challenges:**
  - Increase pattern density by 2 orders of magnitude
  - Increase speed x 3
  - Same Power
  - Use 3D architecture: Vertically Integrated Pattern Recognition AM - VIPRAM

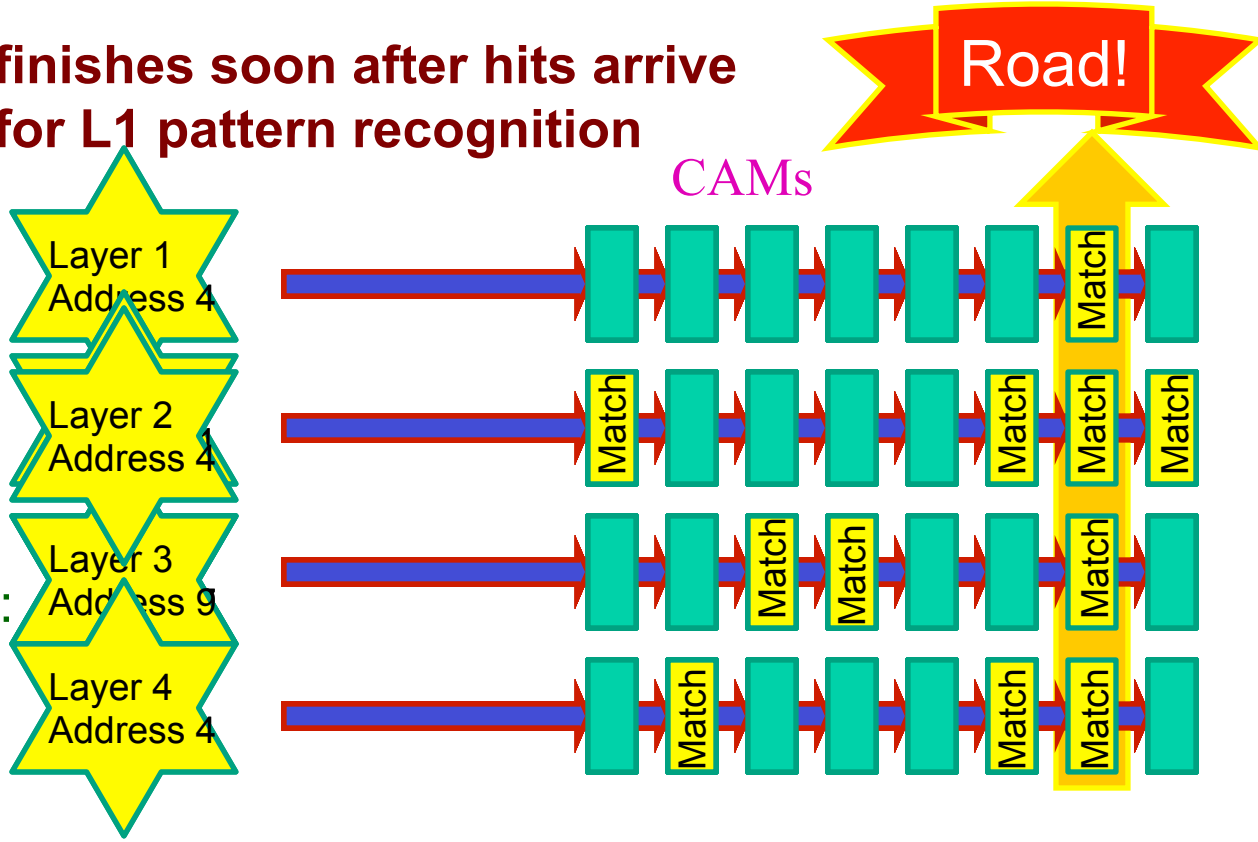

Tools for Trigger/DAQ: xTCA

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

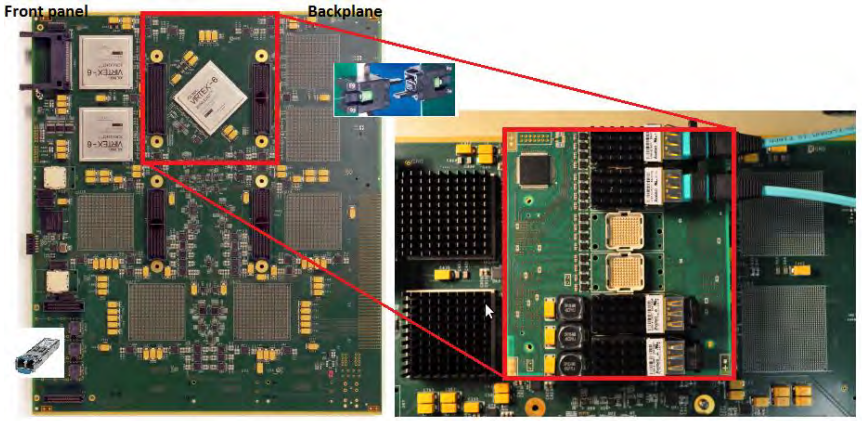

- Example: µTCA derived from AMC std. used by CMS HCAL, Trig.
  - Advanced Mezzanine Card
  - Up to 12 AMC slots
    - *Processing modules*
  - 6 standard 10Gb/s point-to -point links slot to hub slots (more available)
  - Redundant power, controls,clocks
  - Each AMC can have in principle (20) 10 Gb/sec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

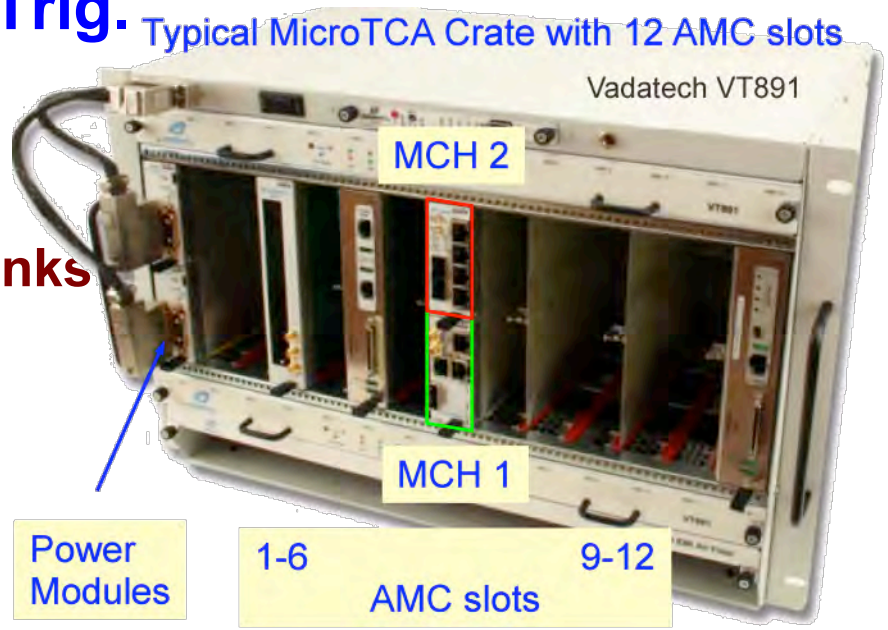

ATCA Example: RCE System

Developed at SLAC

- **Integrated hardware + software entity where generic core firmware & software infrastructure are common & provided.**
- **ATCA infrastructure used**
- **Xilinx ZYNQ series with ARM processors that can run either RTEMS or LINUX.**
- **Has three principal components:**
  - **Programmable FPGA *Fabric***
  - **Programmable *Cluster-Element* (CE).**
  - ***Plugins***
- **Currently being used in:**
  - **ATLAS CSC (proposed: Small Wheel), Darkside, Heavy Photon Search, LBNE, LSST, LCLS, nEXo…**

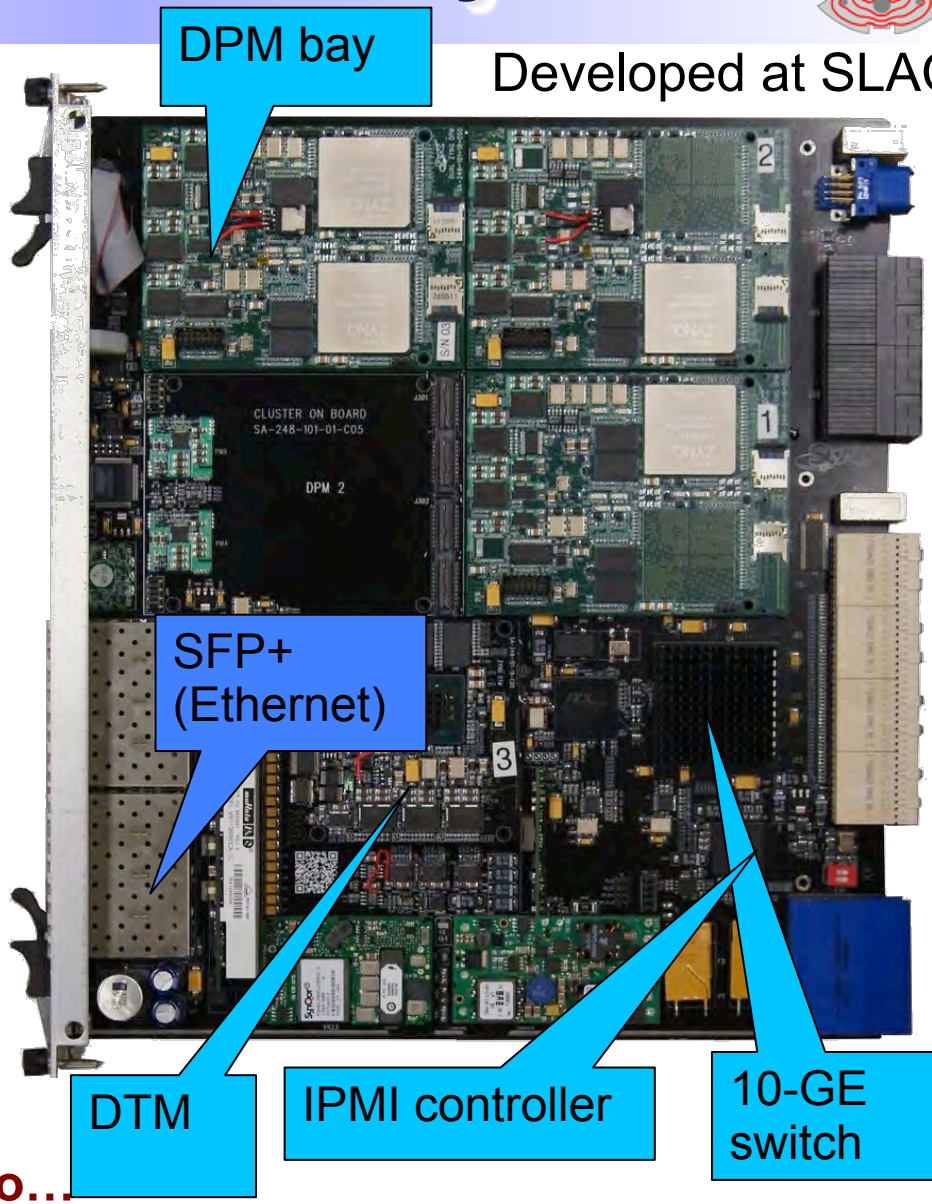

Tools cont'd: CPU, GPU, PCIe

CPU Gains for High Level Triggers: Moore's Law

- **e.g. Xeon Phi Co-processor**
  - 1.2 TeraFlop/s double precision peak performance today

GPU Enhancement of HLT:

- GPU uses a highly scalable architecture that closely tracks Moore's Law
- High performance memory system with ≥ 5x bandwidth vs. CPU
- Better performance / Watt vs. CPU
- Hardware and software support for moving data directly from network interface to GPU memory

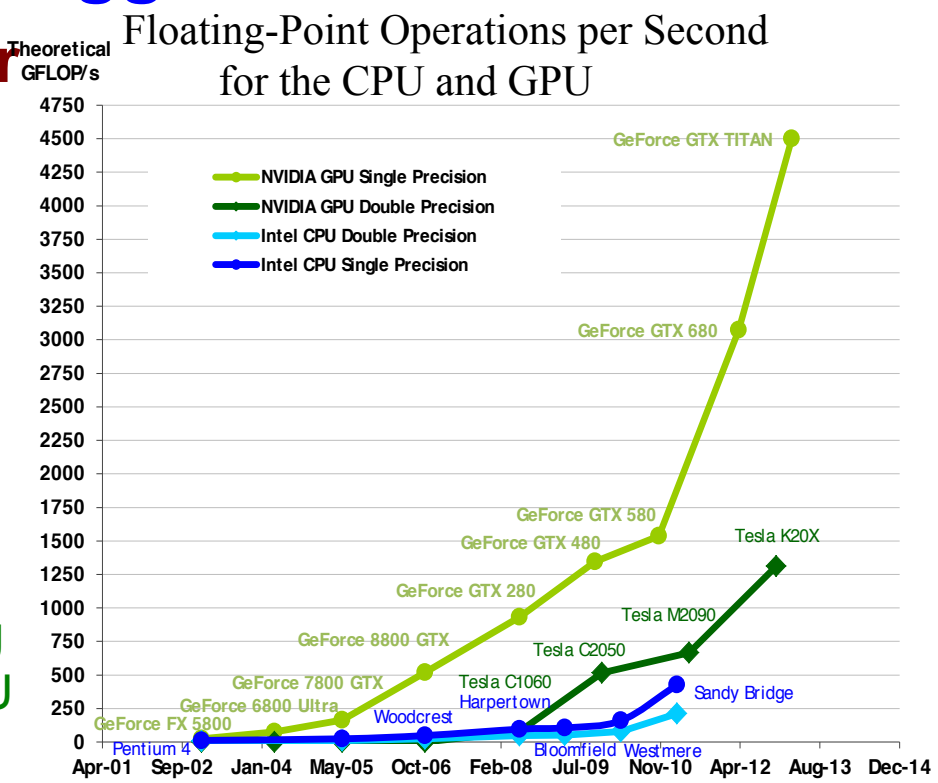

Enhancement of detector to DAQ readout:

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

R&D Topics: Trigger

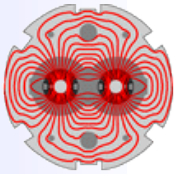

Increase of rate from Level-0 to HLT to read out

- Absolute rate & balance between levels
- L1 complexity vs. HLT input rates

L1 Trigger Latency

- How much is needed & consequences on electronics

L1 Track Triggers

- Associative Memories
- Study techniques: sharpen p_T threshold, e- & μ- ID, Isolation, primary vertex for jets, multi-object triggers, possibility of pixel b-tag.
- Interplay with tracker design

Improvements to L1 Calo. & Muon Triggers

- Processing of much finer-grain, higher bandwidth information

Impact of higher bandwidth links & denser optical interconnects

New packaging & interconnect technologies

- ATCA, μTCA

Use of FPGAs in L1 Trigger

Trigger Timing & Control Networks

- e.g. use of Passive Optical Networks (PON)

Impact of detector timing improvements (~100 ps)

- e.g. crystal calorimeters (CMS: PbWO3 has ~ 150 ps, LYSO < 100 ps)

R&D Topics: HLT & DAQ

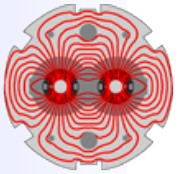

New packaging & interconnect technologies
- ATCA, µTCA

Event building architectures

Future of Server PC architecture

Network Switches

Clock & Control Networks

HLT on the Cloud
- e.g. share resources between HLT & Tier-0

HLT Specialized Track Processing
- e.g. GPU
- depends on resources available: cpu but also link speed

Simulation of HLT
- More sophisticated algorithms, increased occupancy

Use of New Processors in HLT
- ARM, Nvidia Tesla (GPU), Xeon Phi…
  - Just a list of what we can use in the next 2 years!
  - Eventually: heterogeneous mixtures of cores: general & specialized?
- Applies also to computing & software topics

Merging of HLT & offline software development

Conclusions

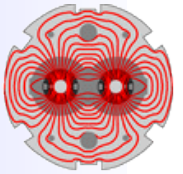

ATLAS & CMS L1 Trigger Scenario:

- 10 - 20 μs latency & L1/0 Accept rates of 0.2/0.5 – 1 MHz.
- L1 Track Trigger

ATLAS & CMS Phase 2 DAQ

- HLT design to accept < 1 MHz of 5 MB events w/PU ≥ 140
- Output of 10 kHz.

ALICE & LHCb Trigger & DAQ:

- Moving to “triggerless” architecture

R&D Program

- FPGAs, Links, Telcom Tech., Associative Mem., GPU, New Processors, Architectures (heterogeneous)
- More powerful tools require more investment to exploit!

Backup

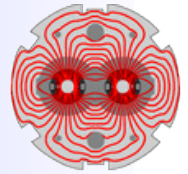

Additional Slides

2012 ATLAS Trigger & DAQ Architecture

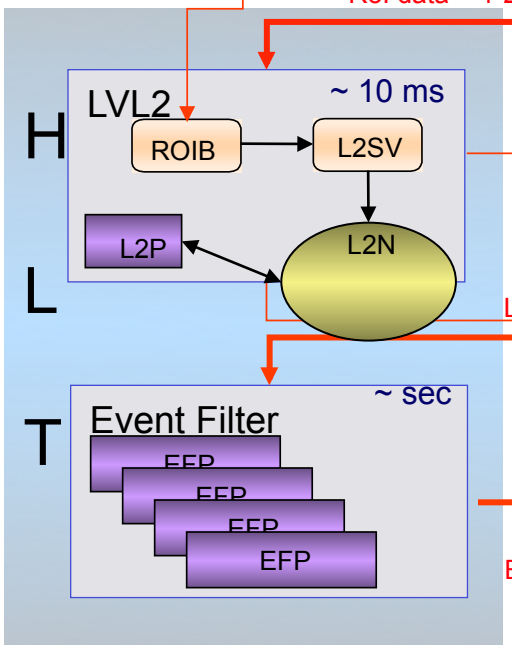

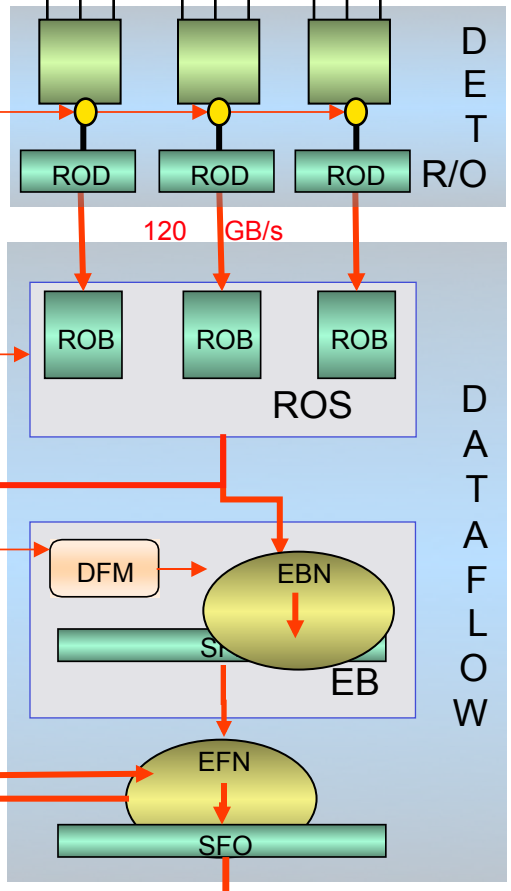

Trigger

DAQ

40 MHz

40 MHz

specialized h/w
ASICs
FPGA

RoI

75 kHz

Lvl1 acc = 75 kHz

RoI data = 1-2%

120 GB/s

120 GB/s

Read-Out Links

RoI Builder
L2 Supervisor

RoI requests

Read-Out Buffers

Read-Out Sub-systems

L2 N/work
L2 Proc Unit

~2+4 GB/s

~2 kHz

Lvl2 acc = ~2 kHz

Dataflow Manager
Event Building N/work

~4 GB/s

Sub-Farm Input

Event Builder

Event Filter
Processors

Event Filter N/work

EFacc = ~0.2 kHz

Sub-Farm Output

~ 200 Hz

~ 300 MB/s

1 PB/s

FE Pipelines

Read-Out Drivers

ATLAS Split Level 0/1 Architecture

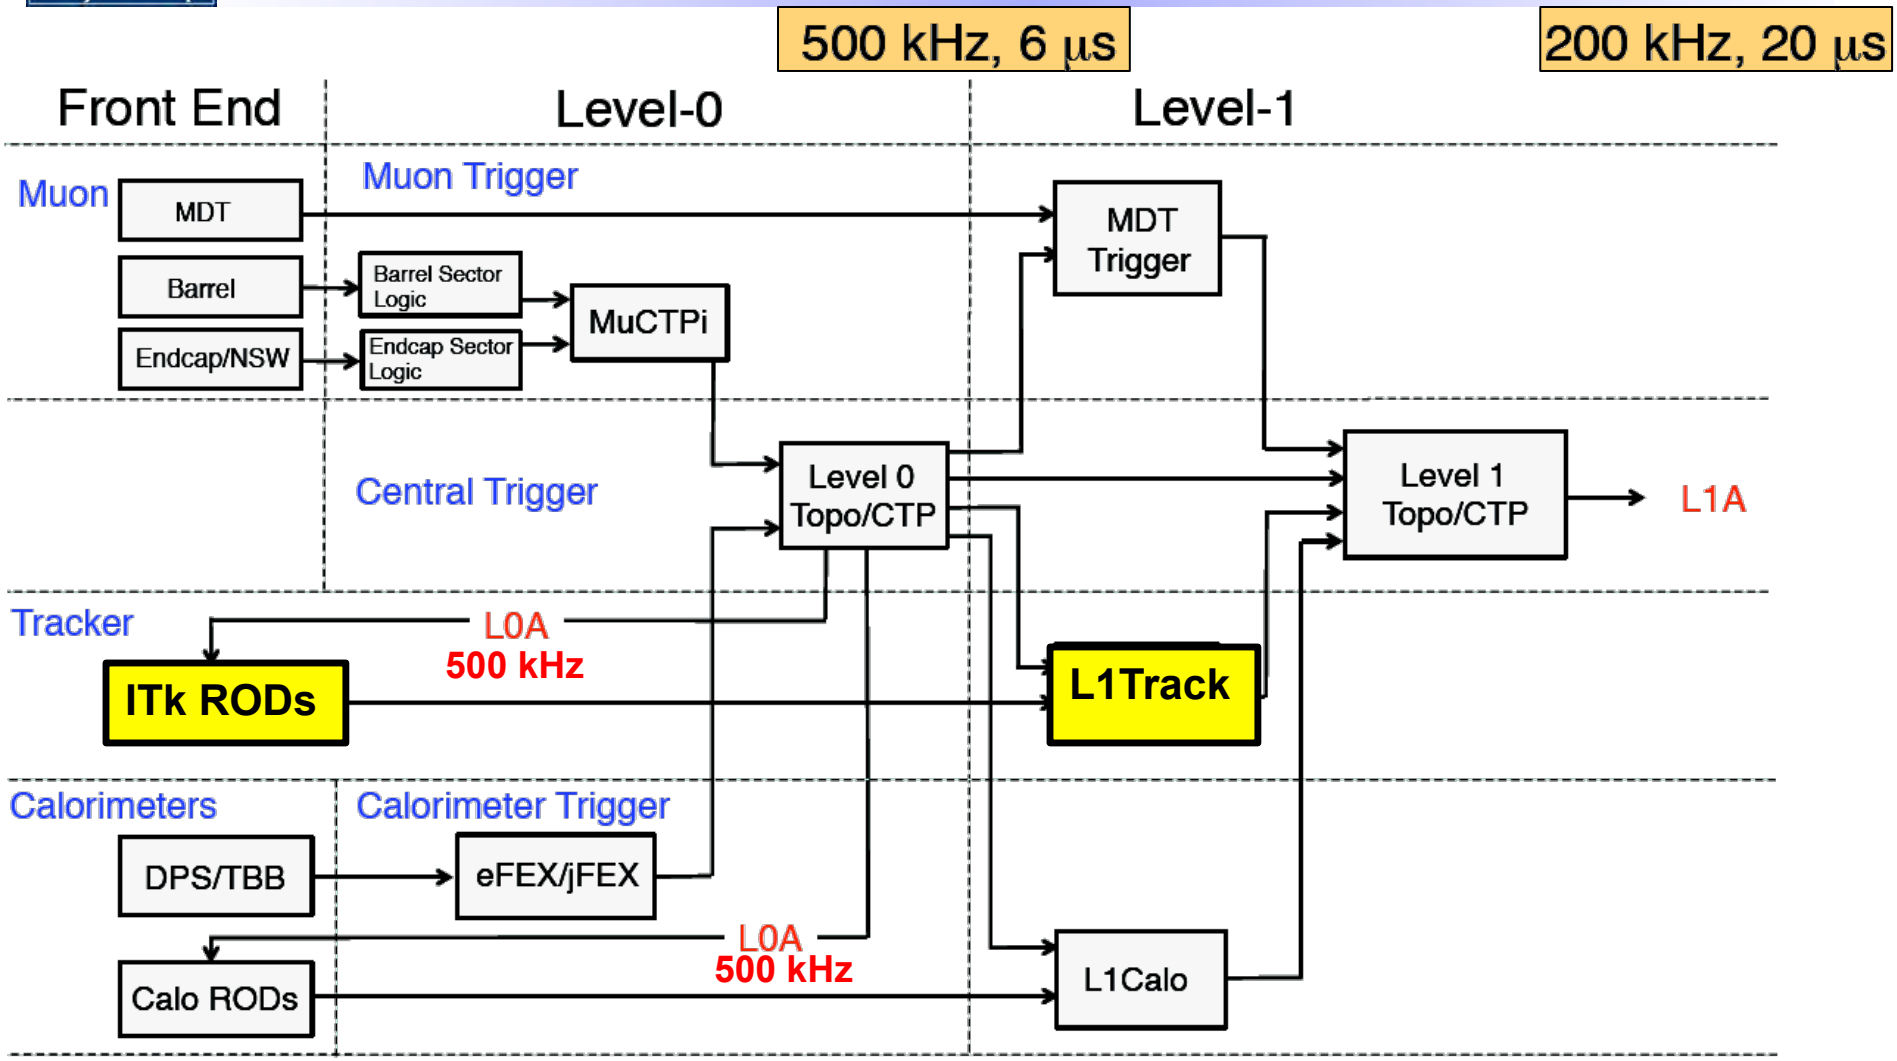

L0 within 6 μs; L1 must be distributed to detector front-ends within 20 μs

- **Leaves only ~ 6 μs to readout tracker and ~ 7 μs for L1 track finding**

CMS 2012 L-1 Trigger & DAQ

Overall Trigger & DAQ Architecture: 2 Levels:

Level-1 Trigger:

- 25 ns input
- 3.2 μs latency

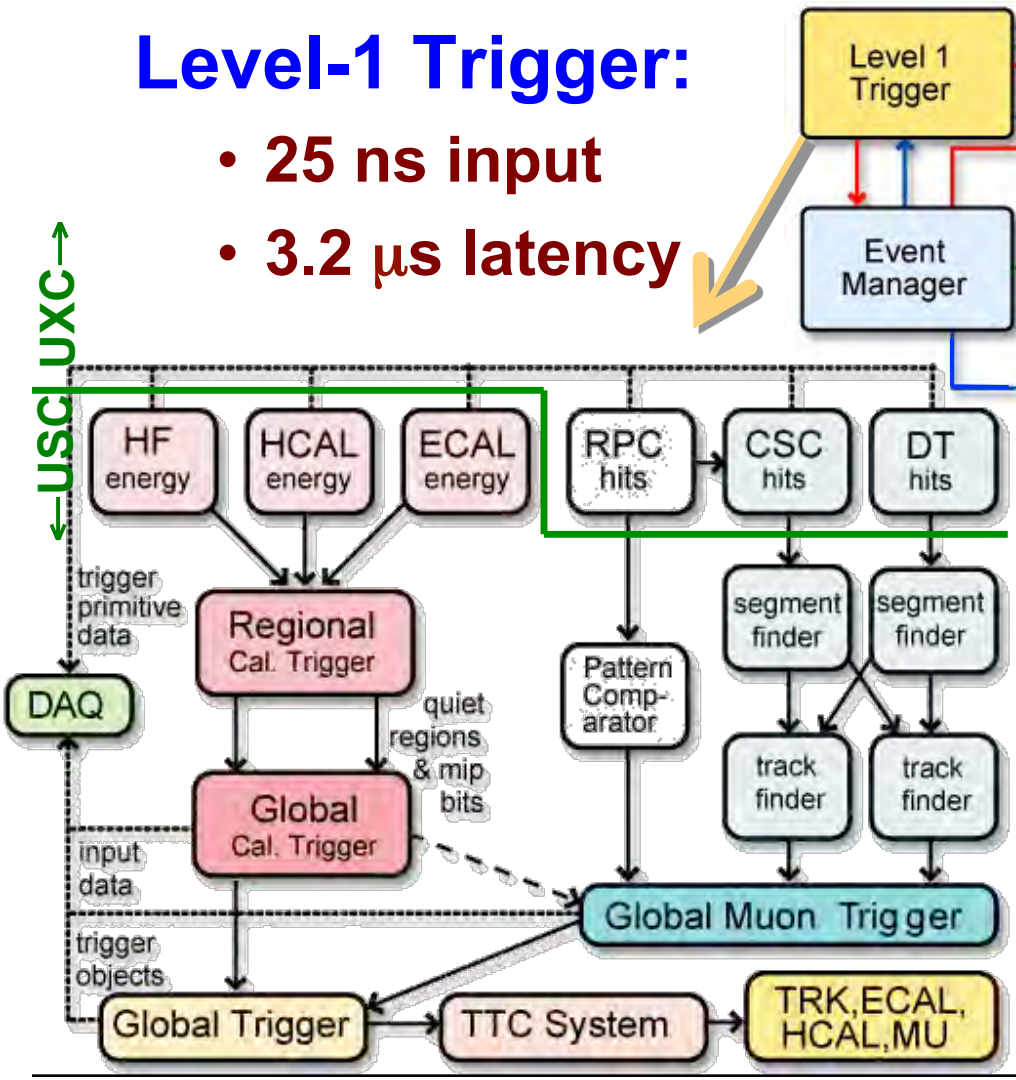

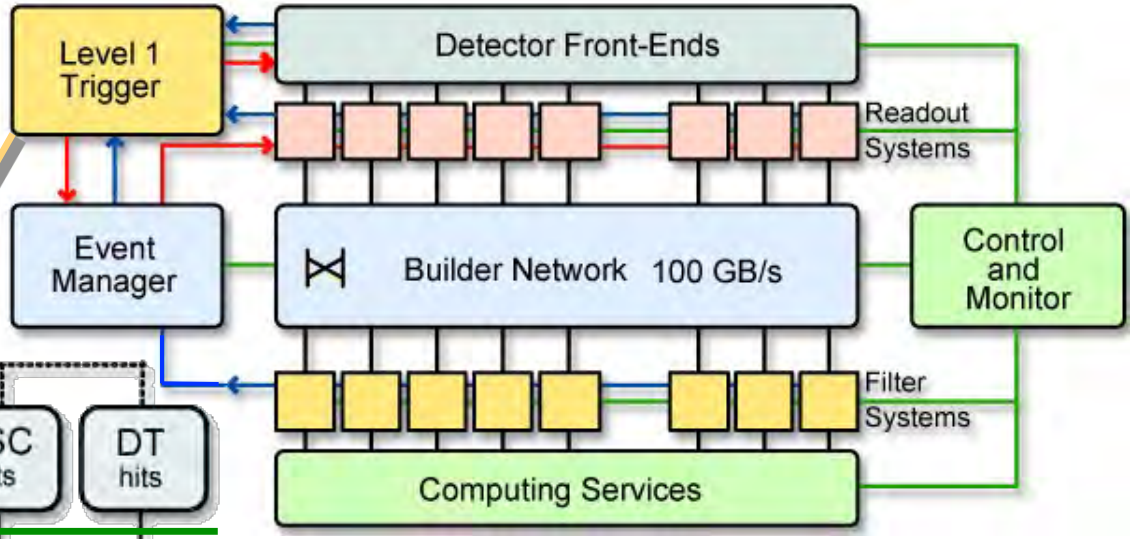

Interaction rate: 1 GHz

Bunch Crossing rate: 40 MHz

Level 1 Output: 100 kHz

Output to Storage: 400 Hz

Average Event Size: 1 MB

Data production 1 TB/day