CTEQ

Data Selection & Collection: Trigger & DAQ

CTEQ Summer School

Wesley H. Smith

U. Wisconsin - Madison

July 19, 2011

Outline:

Introduction to LHC Trigger & DAQ

Challenges & Architecture

Examples: ATLAS & CMS Trigger & DAQ

The Future: LHC Upgrade Trigger & DAQ (if time)

Triggering

- Task: inspect detector information and provide a first decision on whether to keep the event or throw it out

The trigger is a function of :

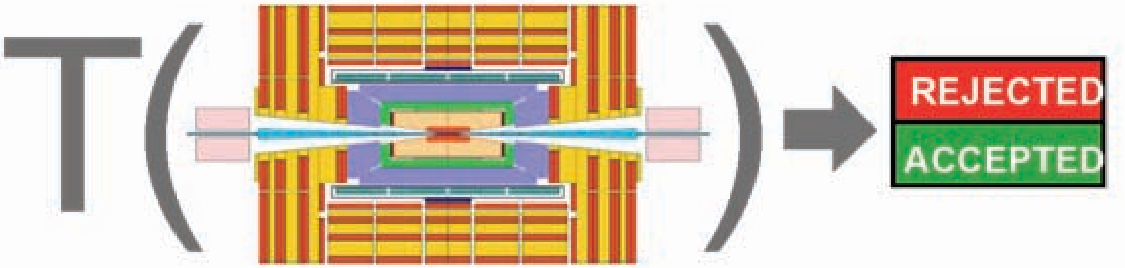

Event data & Apparatus
Physics channels & Parameters

- Detector data not (all) promptly available
- Selection function highly complex

⇒T(...) is evaluated by successive approximations, the

TRIGGER LEVELS

(possibly with zero dead time)

LHC Collisions

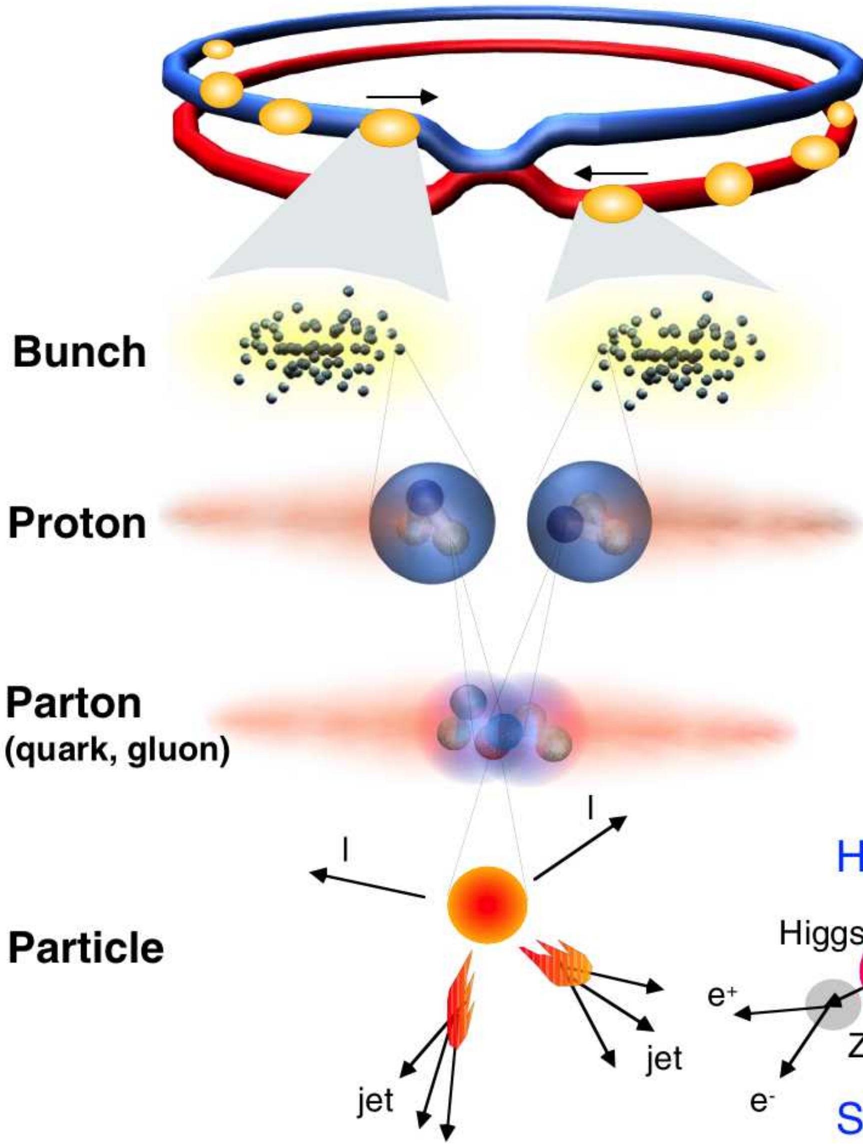

Proton-Proton	2835 bunch/beam
Protons/bunch	10^{11}
Beam energy	7 TeV (7×10^{12} eV)
Luminosity	10^{34} cm^{-2} s^{-1}
Crossing rate	40 MHz

with every bunch crossing
23 Minimum Bias events
with ~1725 particles produced

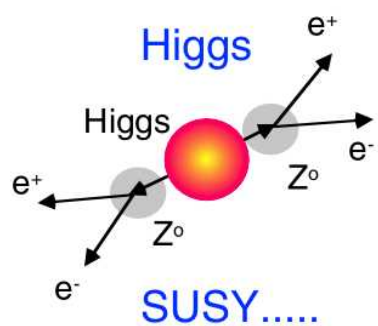

Selection of 1 in
10,000,000,000,000

Beam Xings: LEP. TeV, LHC

LHC has ~3600 bunches

- And same length as LEP (27 km)
- Distance between bunches: 27km/3600=7.5m
- Distance between bunches in time: 7.5m/c=25ns

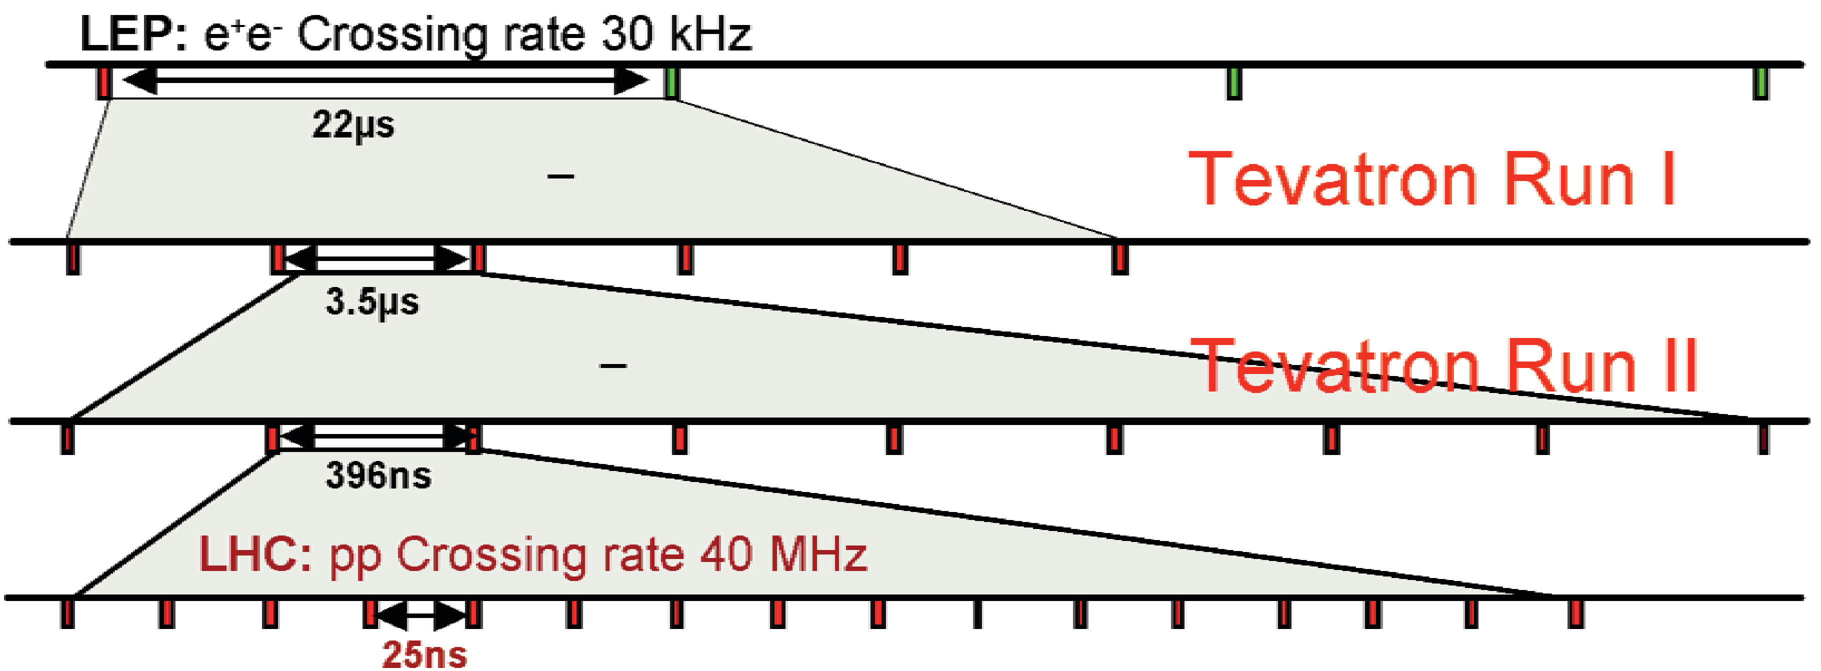

CTEQ

LHC Physics & Event Rates

THE UNIVERSITY of WISCONSIN MADISON

At design *L* = 10^{34}cm^{-2}s^{-1}

- **23 pp events/25 ns xing**
  - ~ 1 GHz input rate
  - "Good" events contain ~ 20 bkg. events
- **1 kHz W events**
- **10 Hz top events**
- **< 10^4 detectable Higgs decays/year**

Can store ~ 300 Hz events

Select in stages

- **Level-1 Triggers**
  - 1 GHz to 100 kHz
- **High Level Triggers**
  - 100 kHz to 300 Hz

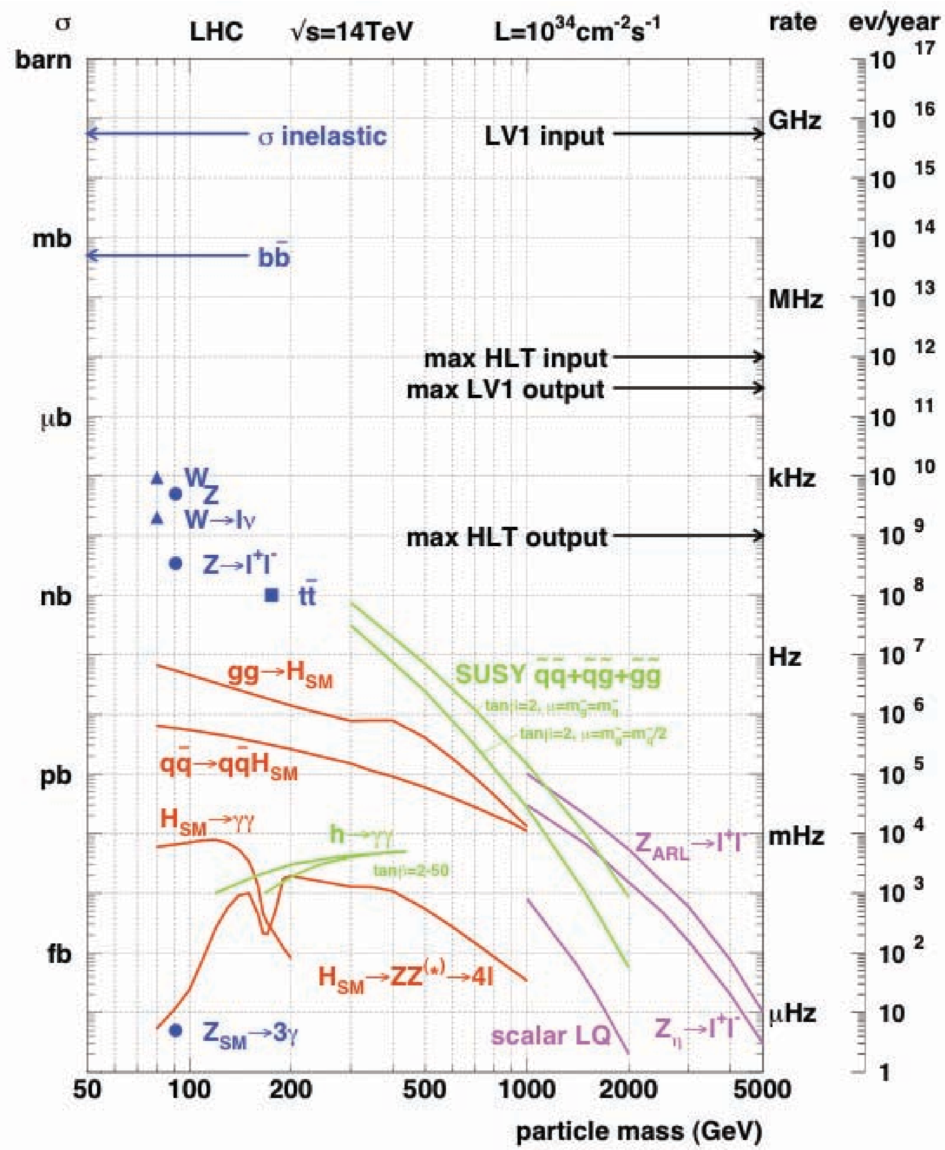

Collisions (p-p) at LHC

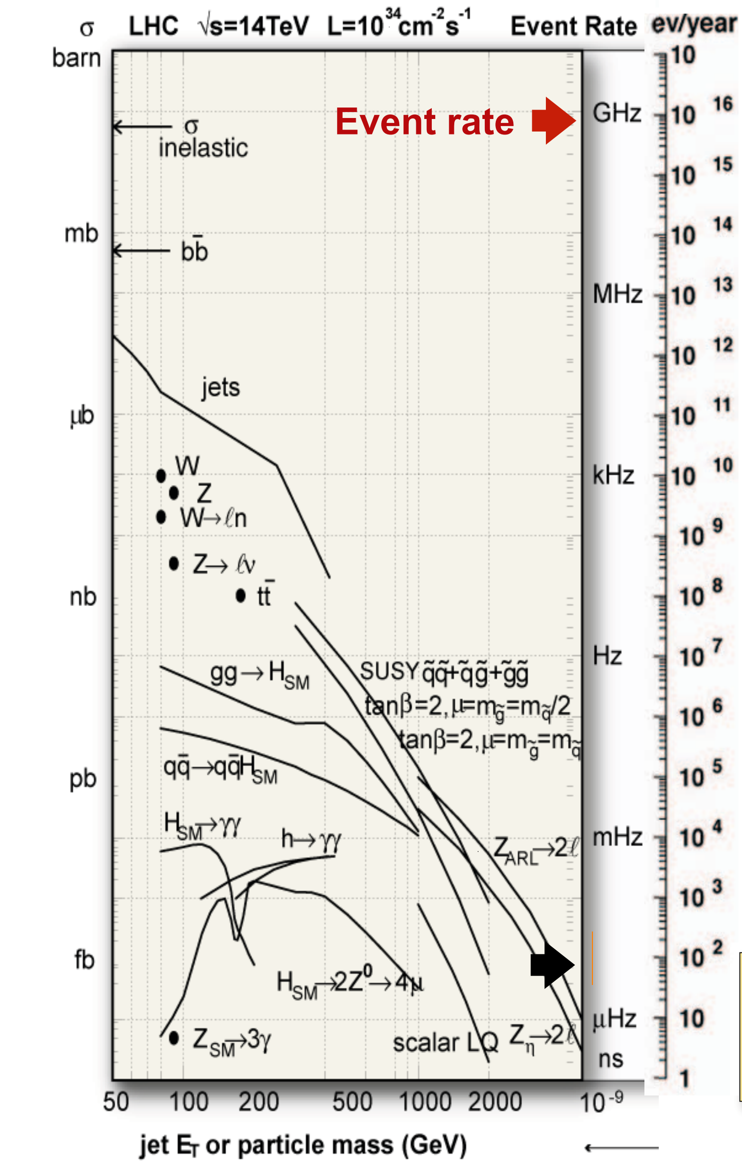

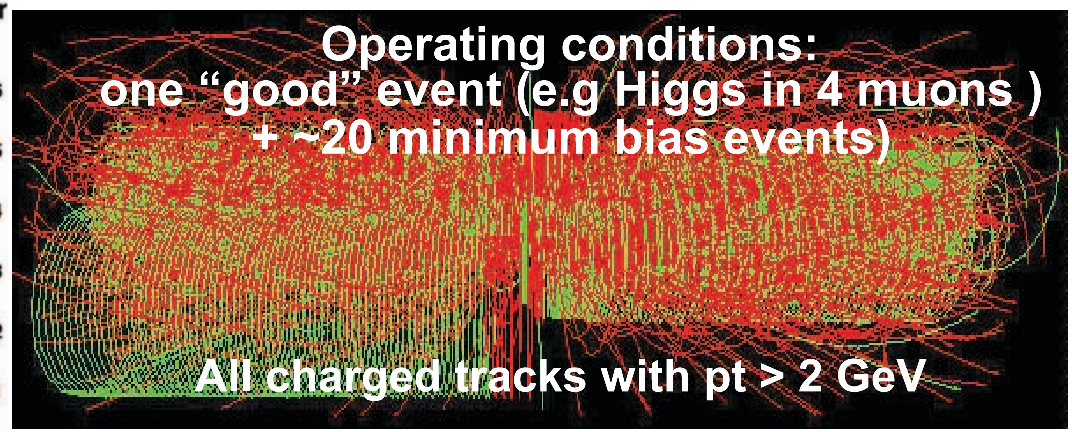

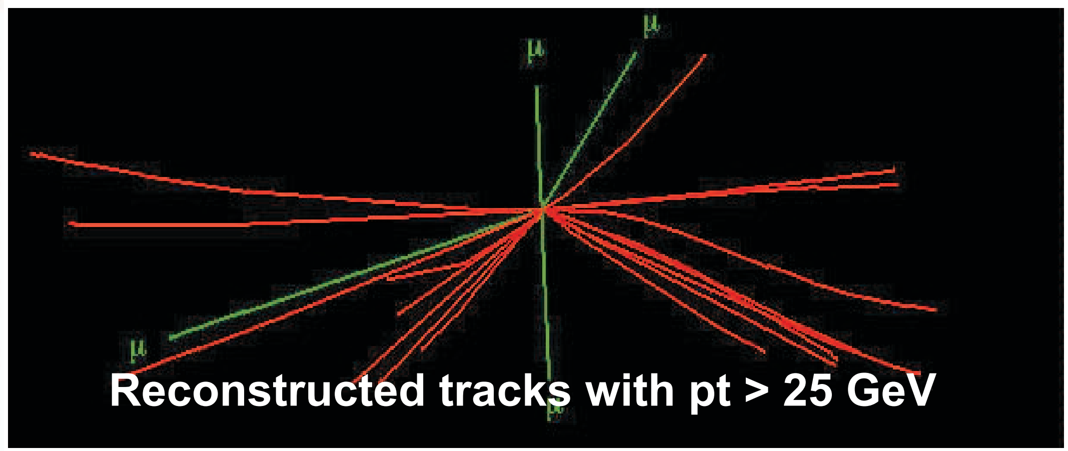

Event size:	**~1 MByte**
Processing Power:	**~X TFlop**

Processing LHC Data

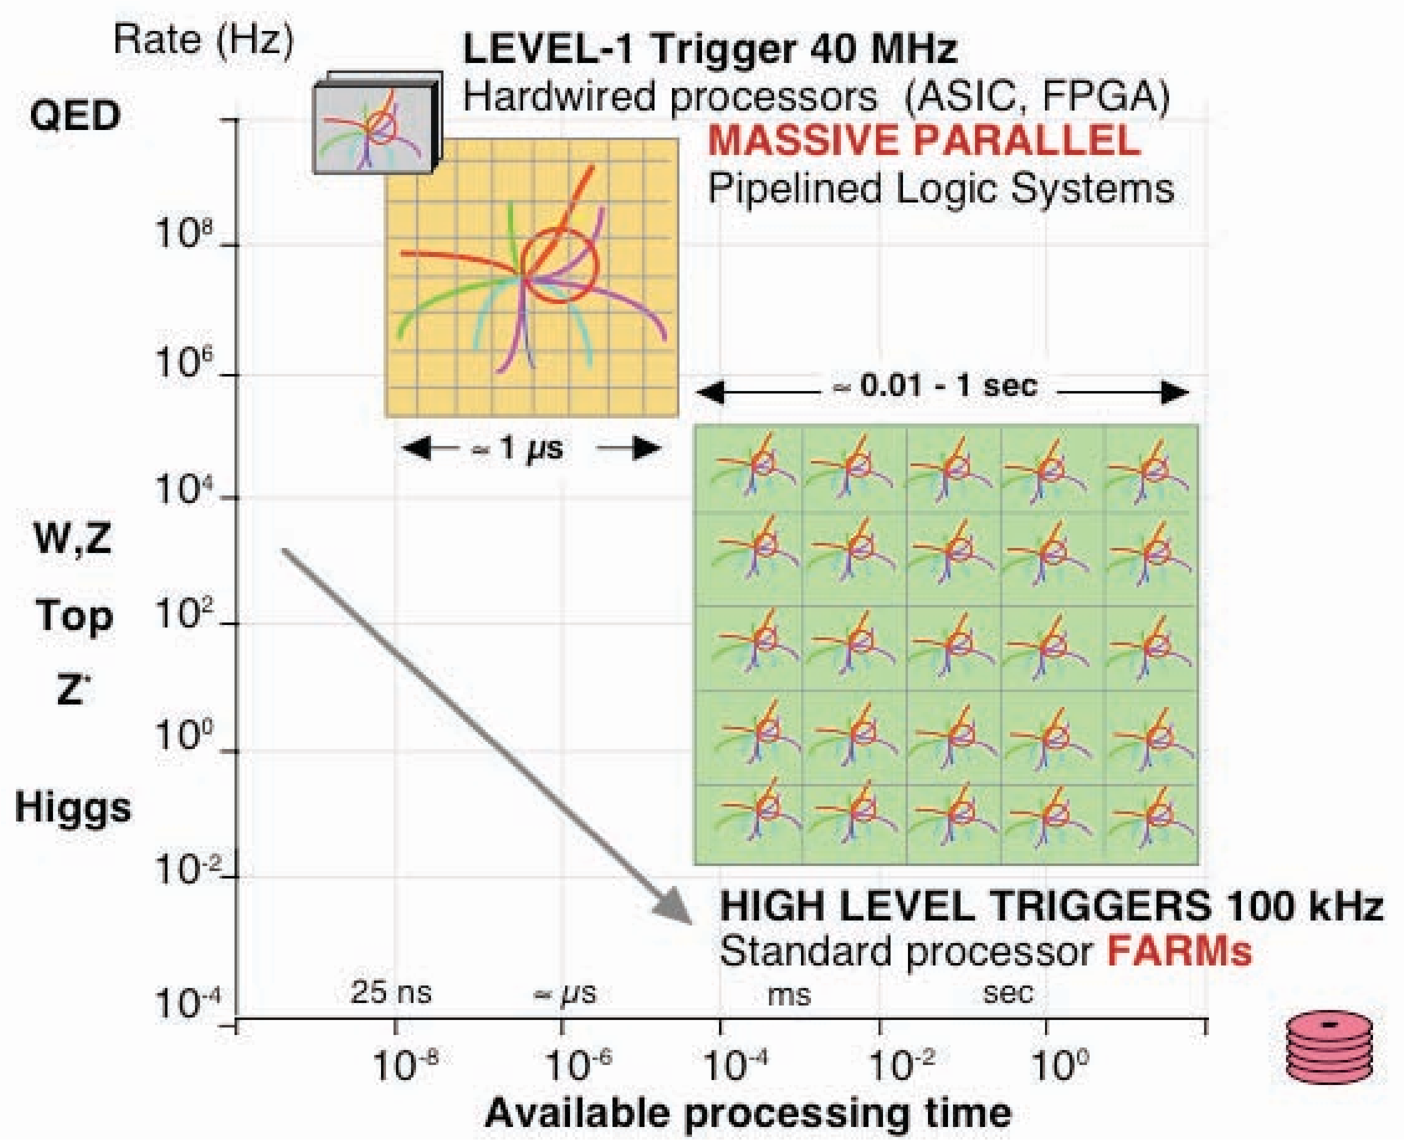

C T E Q

LHC Trigger & DAQ Challenges

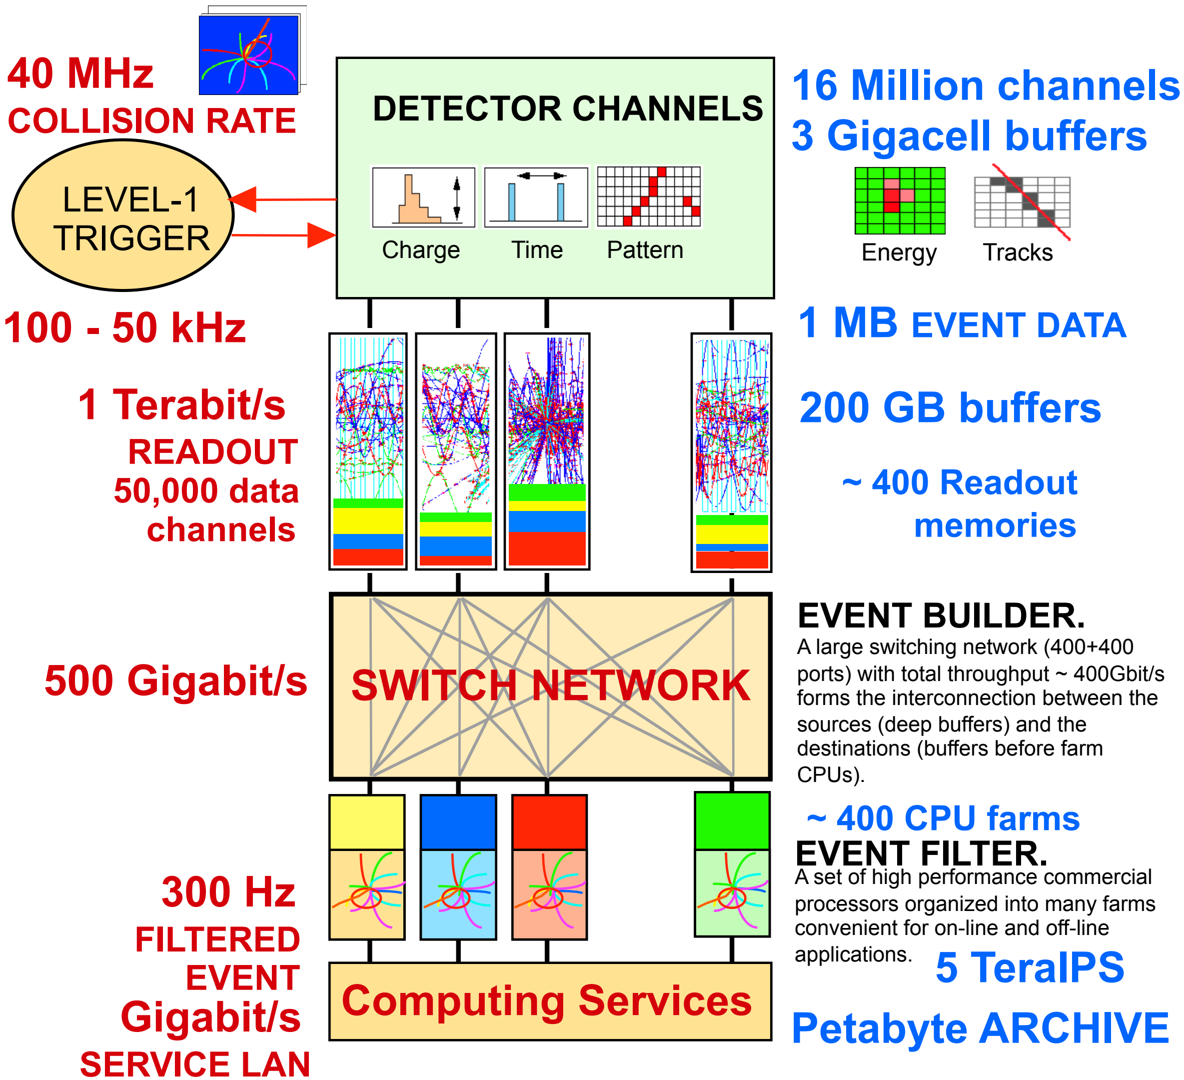

Challenges:

1 GHz of Input Interactions

Beam-crossing every 25 ns with ~ 23 interactions produces over 1 MB of data

Archival Storage at about 300 Hz of 1 MB events

Challenges: Pile-up

- "In-time" pile-up: particles from the same crossing but from a different pp interaction
- Long detector response/pulse shapes:
  - "Out-of-time" pile-up: left-over signals from interactions in previous crossings
  - Need "bunch-crossing identification"

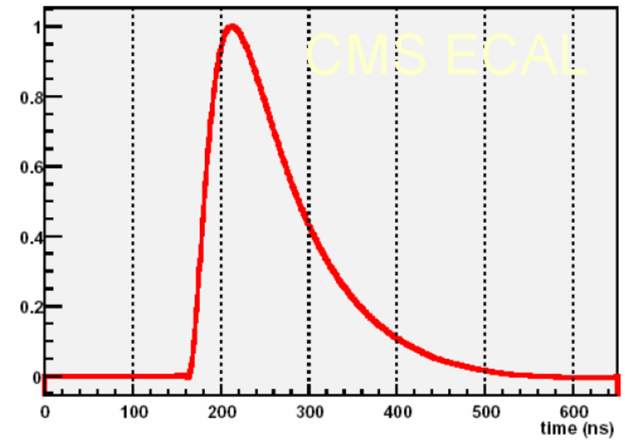

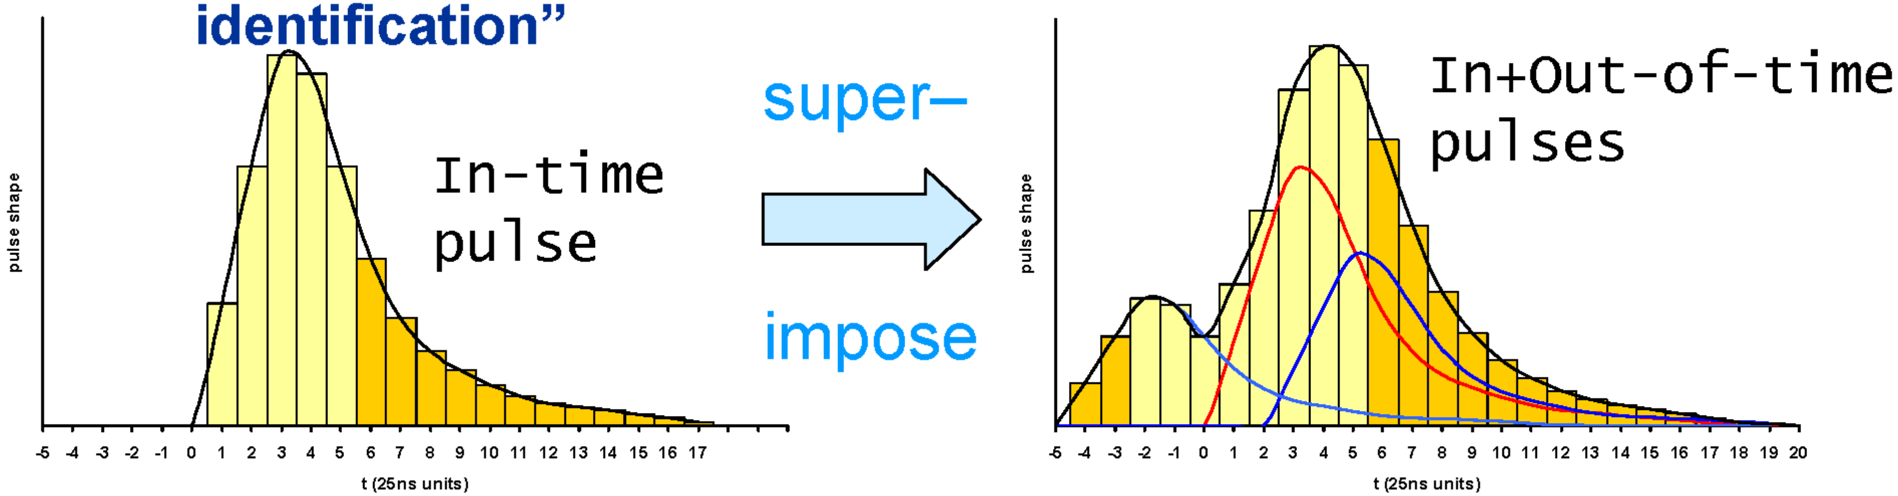

CTEQ

Challenges: Time of Flight

c = 30 cm/ns → in 25 ns, s = 7.5 m

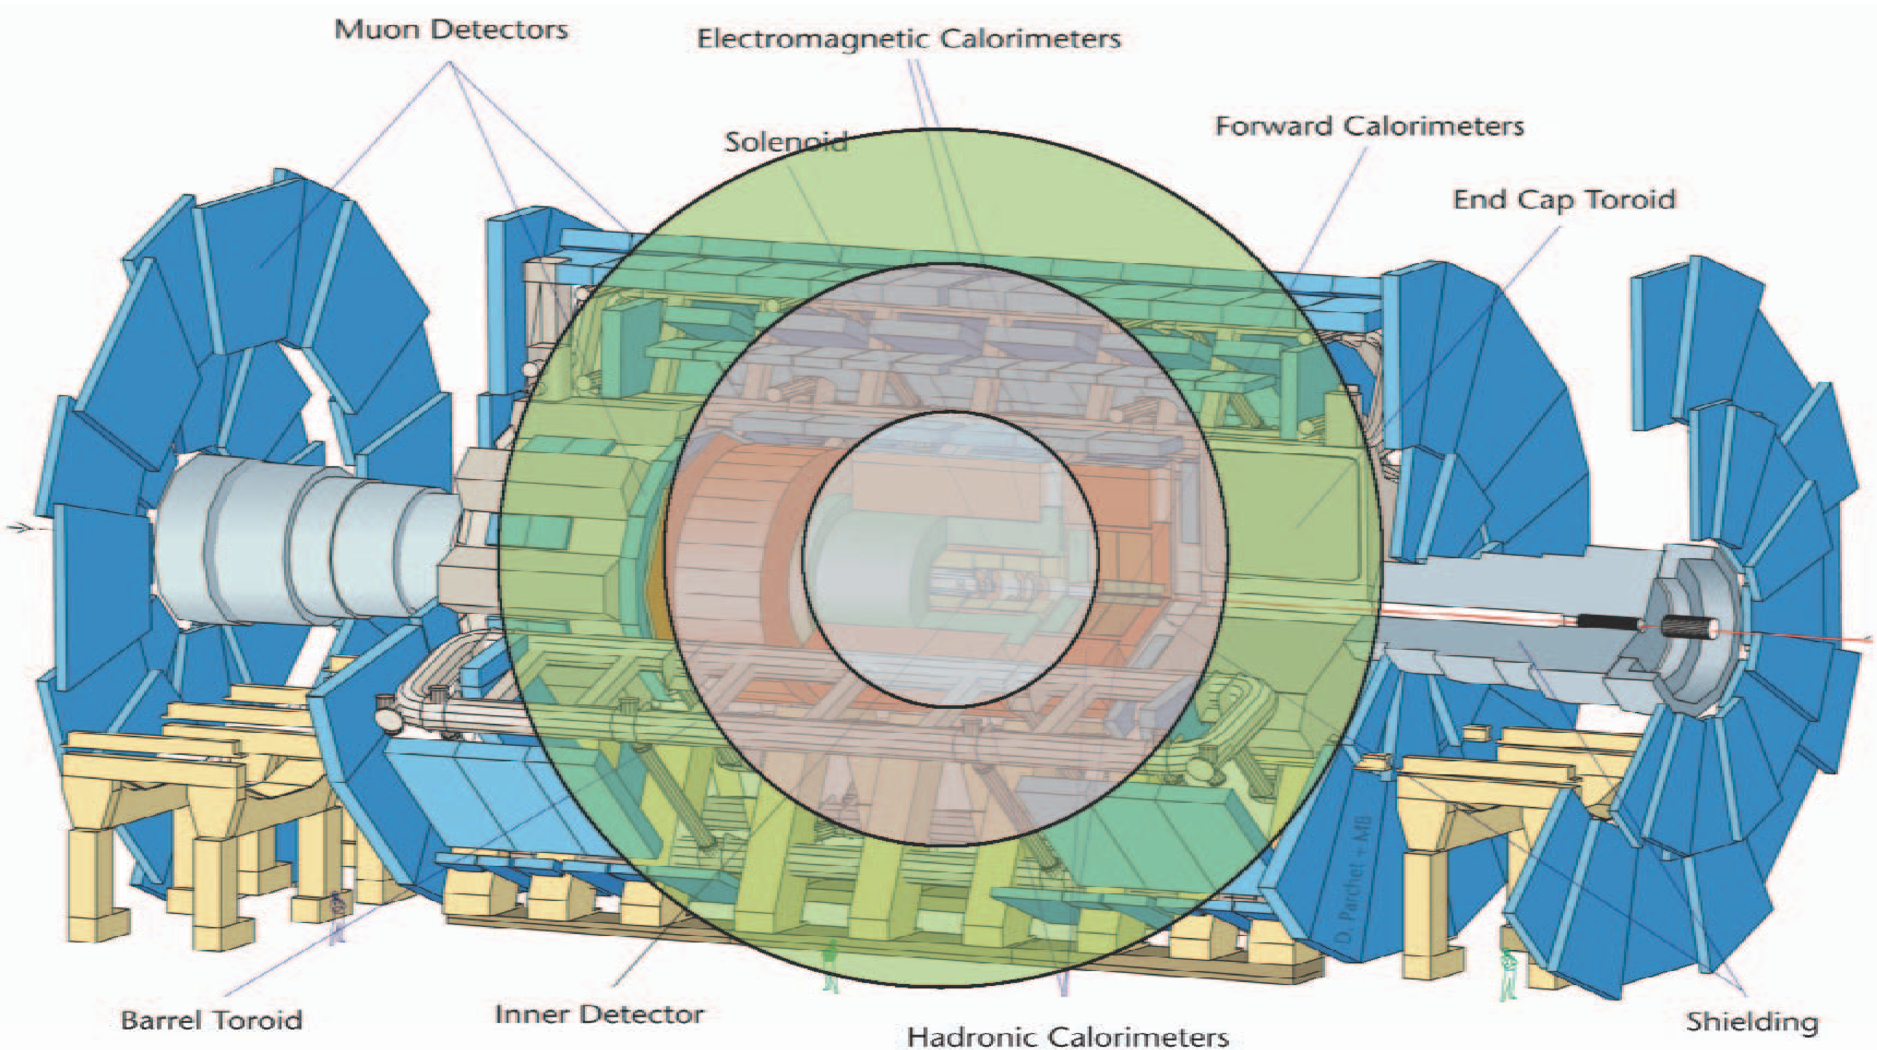

LHC Trigger Levels

10^{-7} s

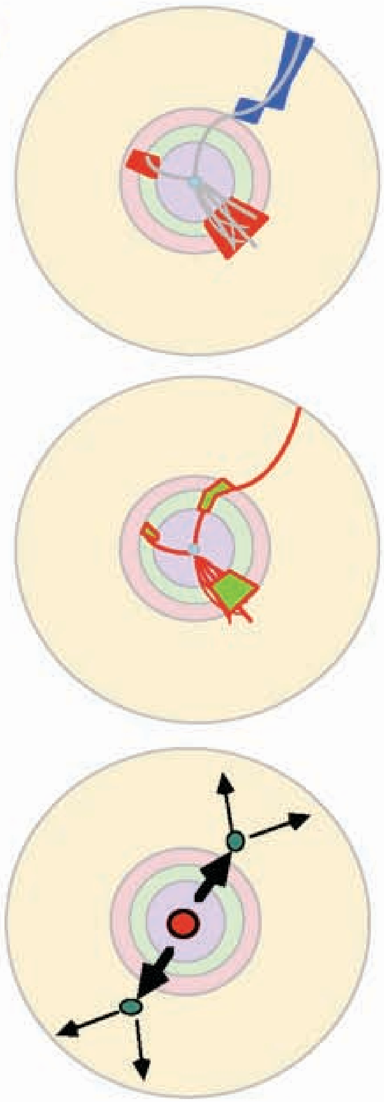

Collision rate 10^9 Hz

Channel data sampling at 40 MHz

Level-1 selected events 10^5 Hz

Particle identification (High p_T e, μ, jets, missing E_T)

- Local pattern recognition
- Energy evaluation on prompt macro-granular information

10^{-6} s

Level-2 selected events 10^3 Hz

Clean particle signature (Z, W, ..)

- Finer granularity precise measurement
- Kinematics. effective mass cuts and event topology
- Track reconstruction and detector matching

10^{-3} s

Level-3 events to tape 100- 300 Hz

Physics process identification

- Event reconstruction and analysis

10^{-0} s

Level 1 Trigger Operation

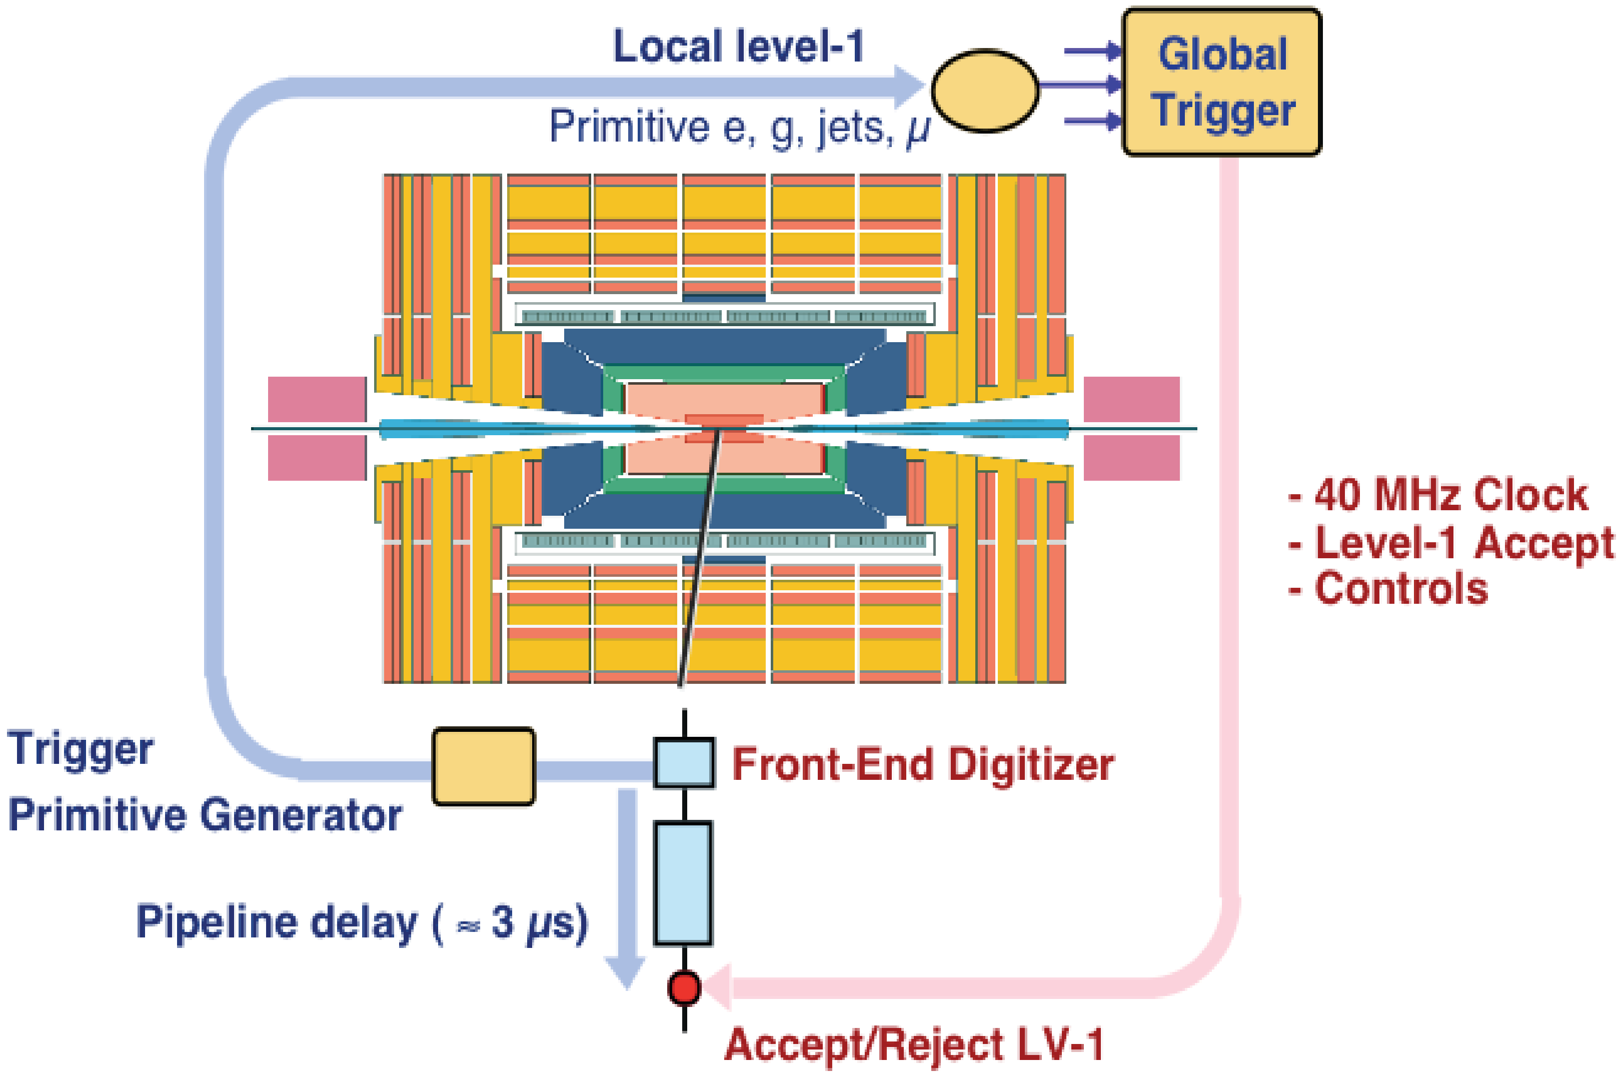

CTEQ

Level 1 Trigger Organization

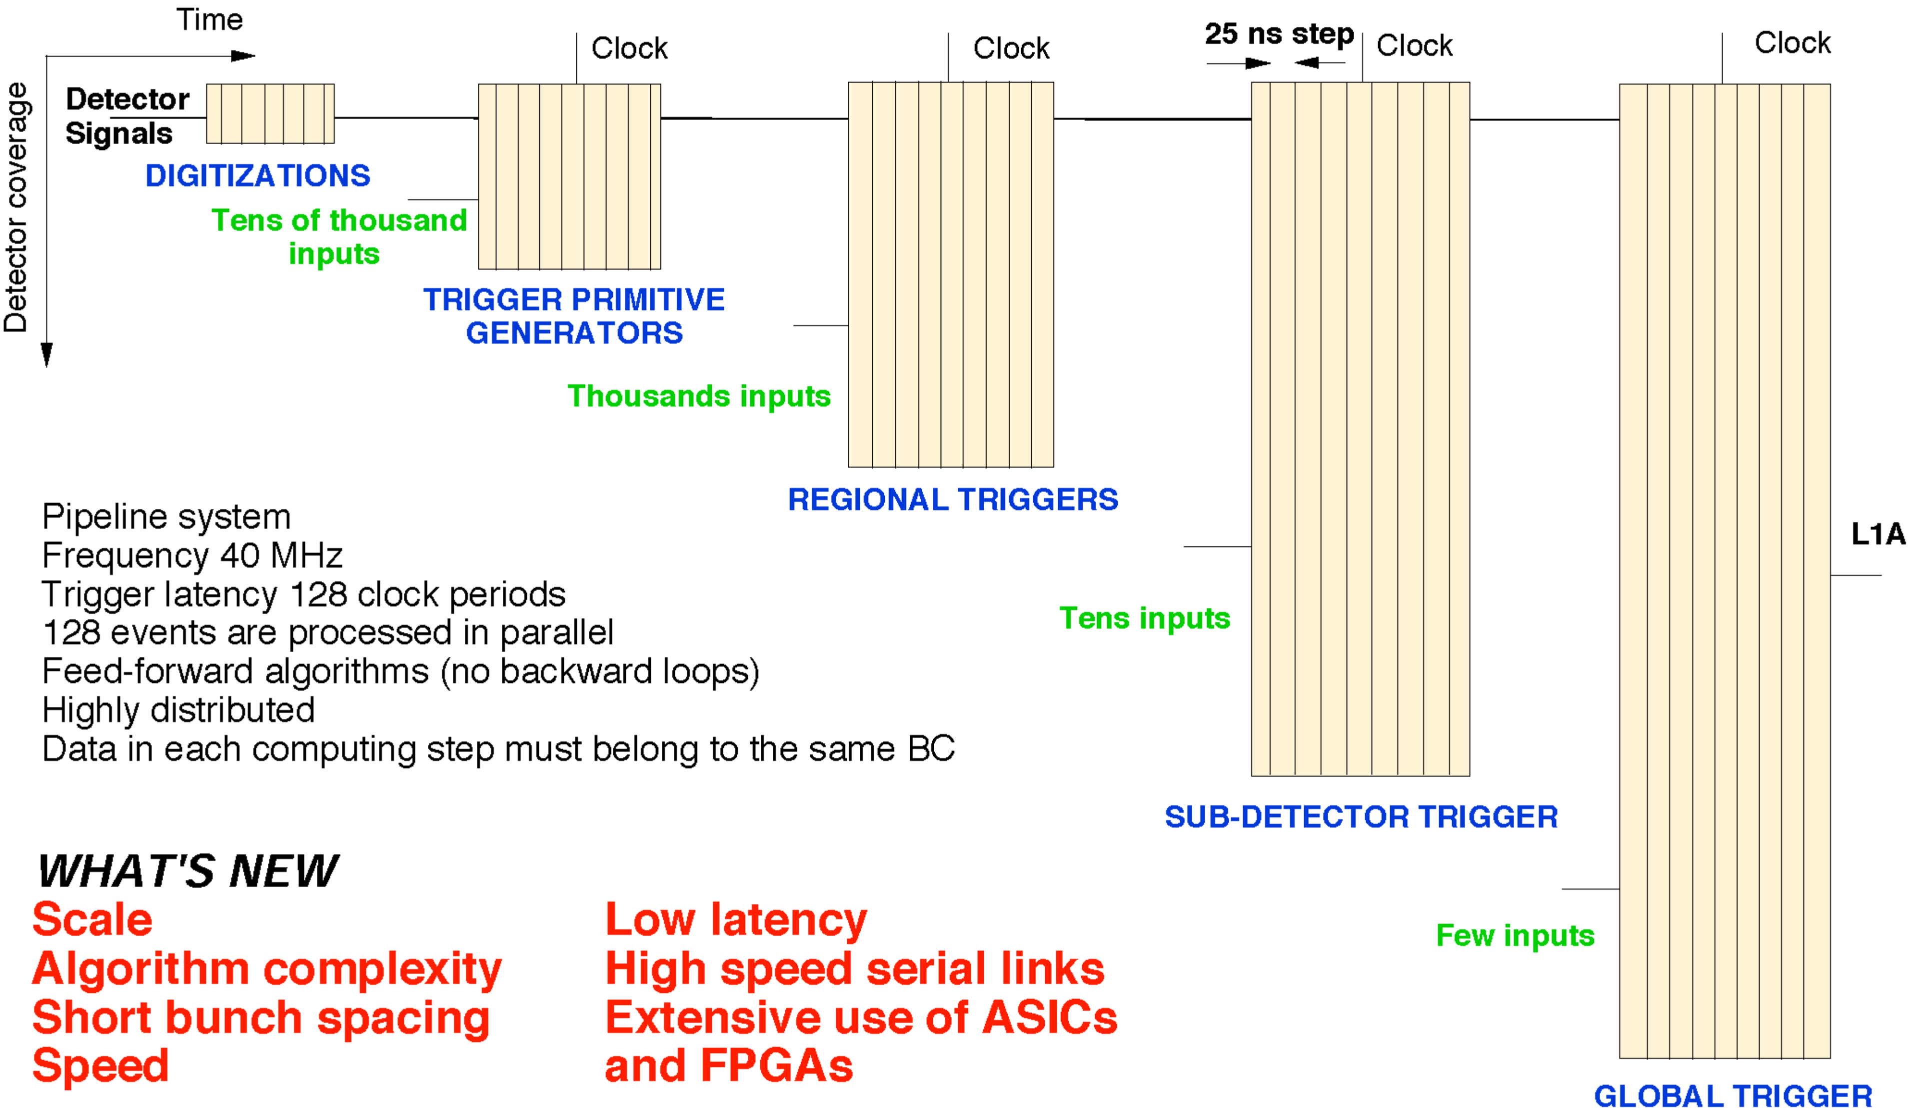

Pipeline system
Frequency 40 MHz
Trigger latency 128 clock periods
128 events are processed in parallel
Feed-forward algorithms (no backward loops)
Highly distributed
Data in each computing step must belong to the same BC

WHAT'S NEW

Scale
Algorithm complexity
Short bunch spacing
Speed

Low latency
High speed serial links
Extensive use of ASICs and FPGAs

Trigger Timing & Control

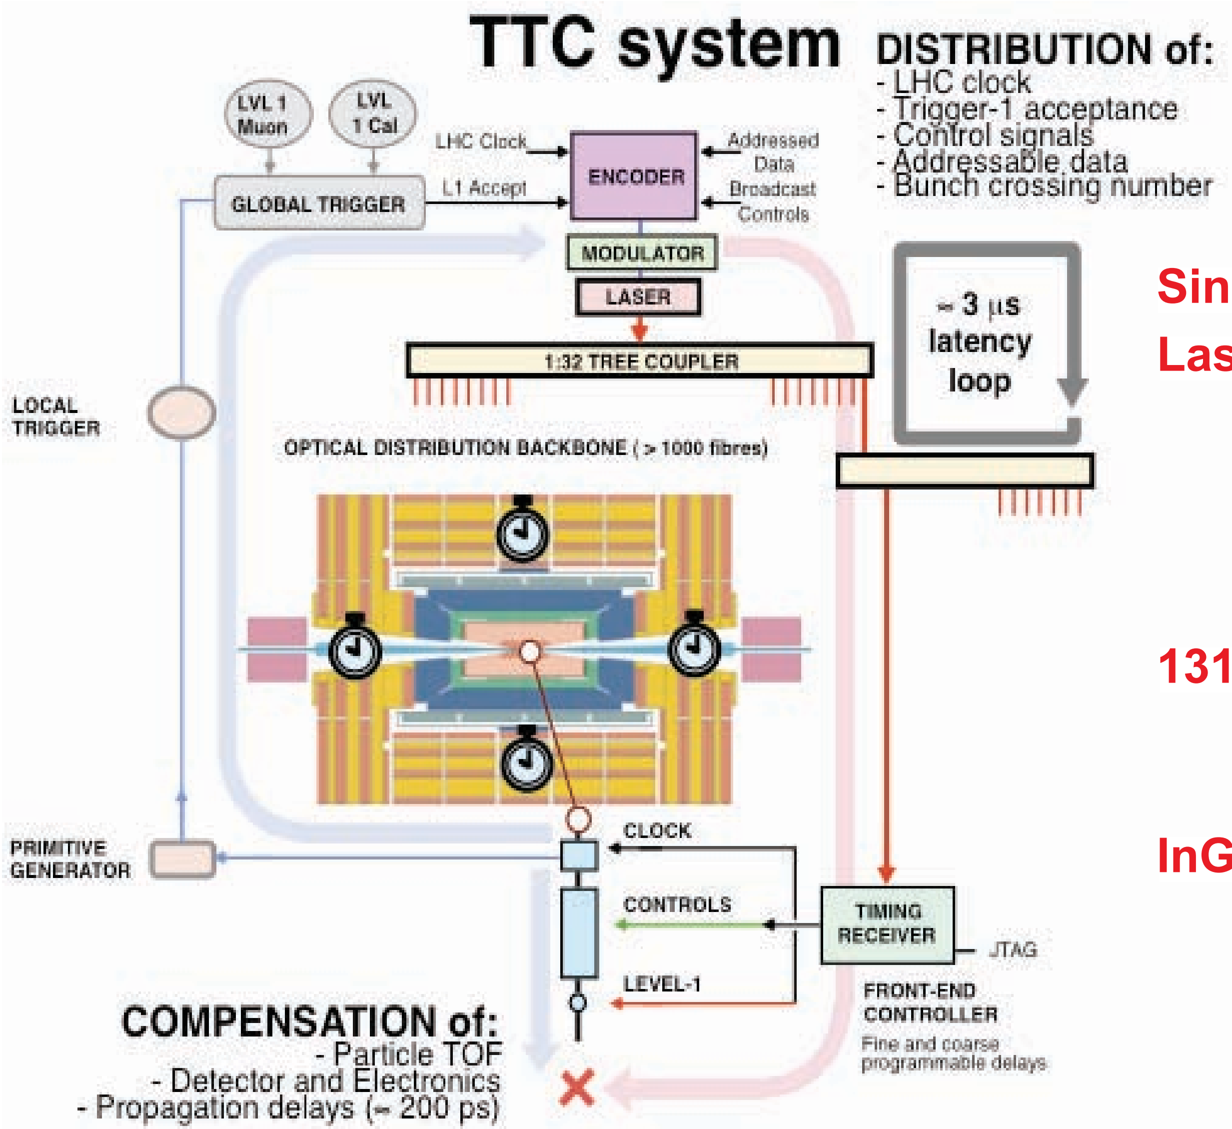

Optical System:

Single High-Power Laser per zone

- **Reliability, transmitter upgrades**
- **Passive optical coupler fanout**

1310 nm Operation

- **Negligible chromatic dispersion**

InGaAs photodiodes

- **Radiation resistance, low bias**

Detector Timing Adjustments

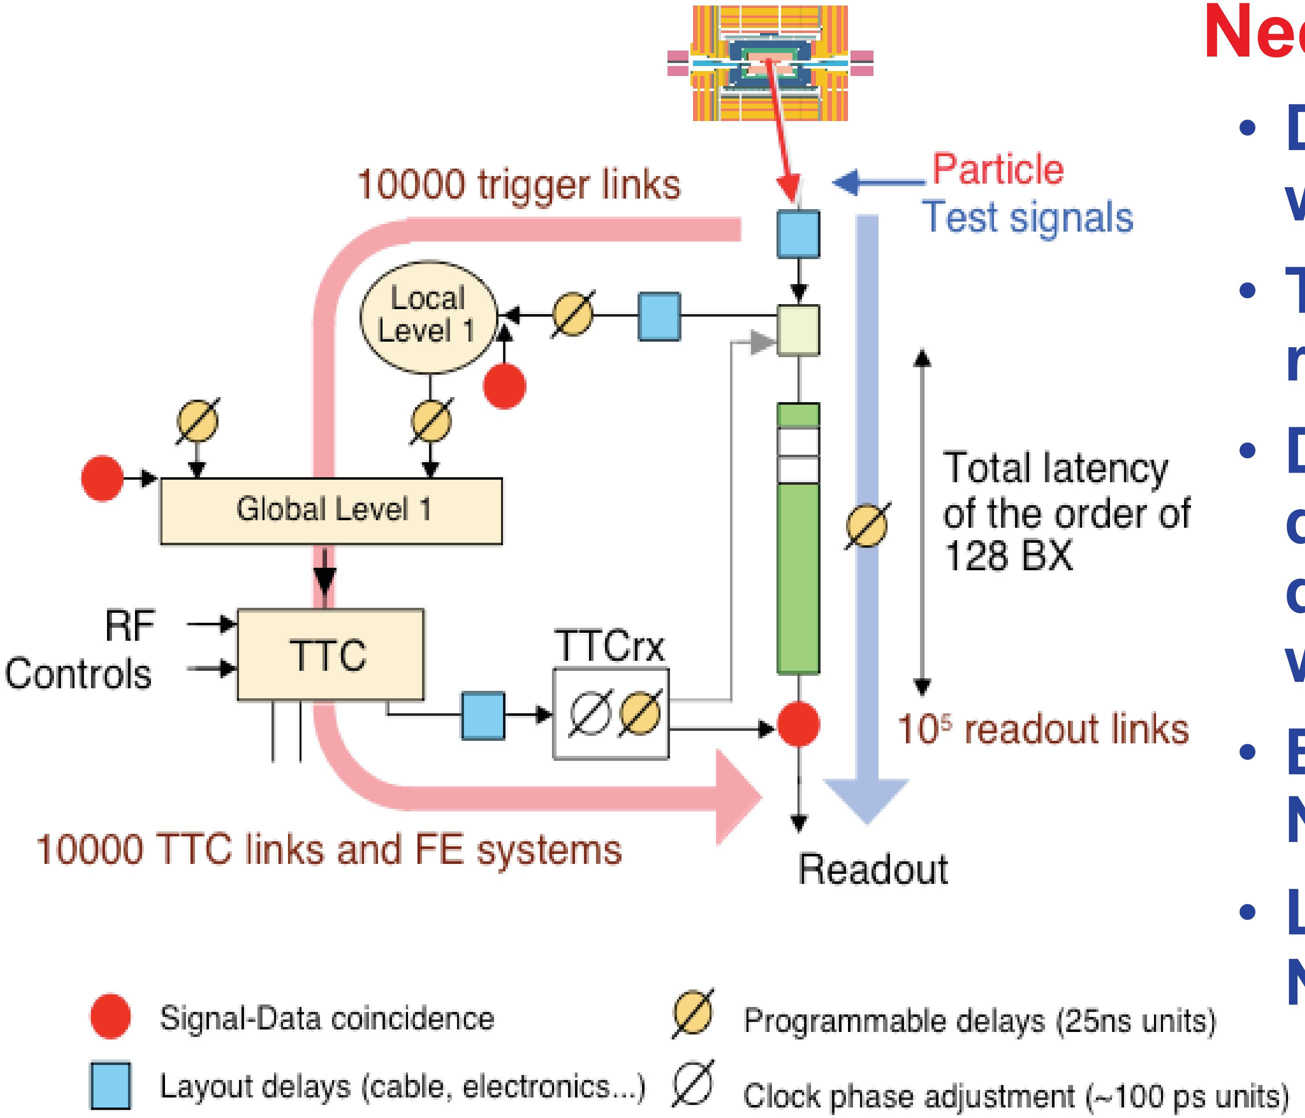

Need to Align:

- Detector pulse w/collision at IP
- Trigger data w/ readout data
- Different detector trigger data w/each other
- Bunch Crossing Number
- Level 1 Accept Number

Synchronization Techniques

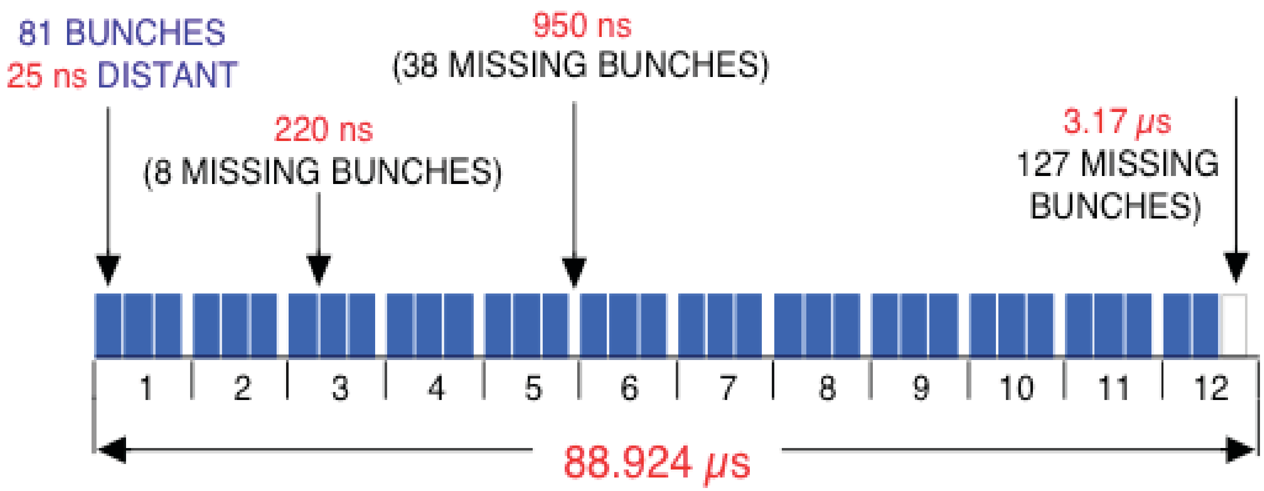

2835 out of 3564 p bunches are full, use this pattern:

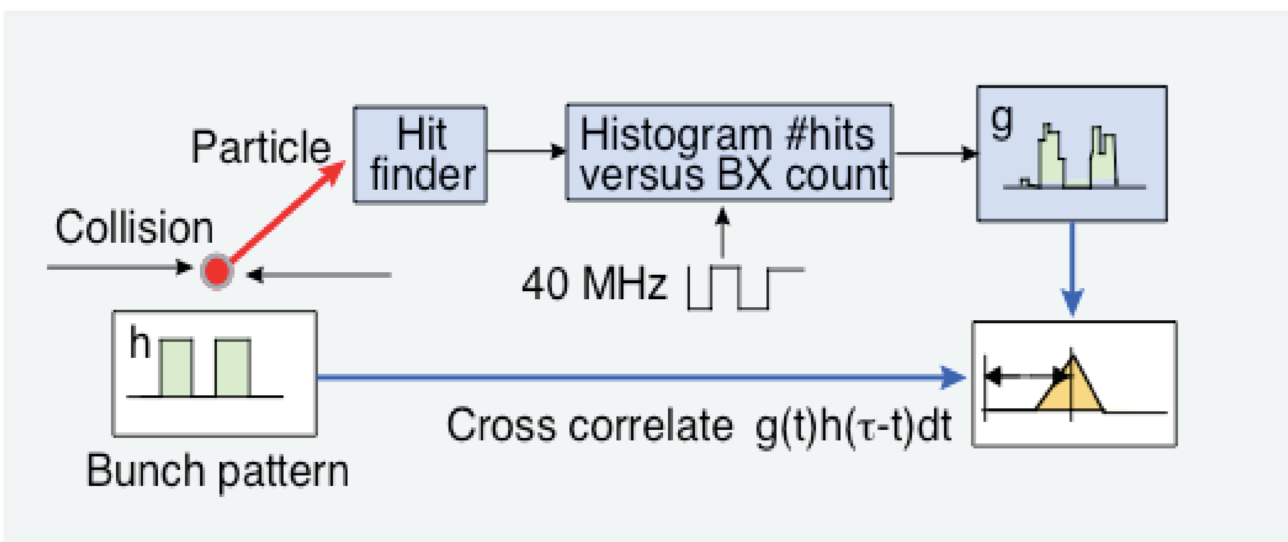

ATLAS Detector Design

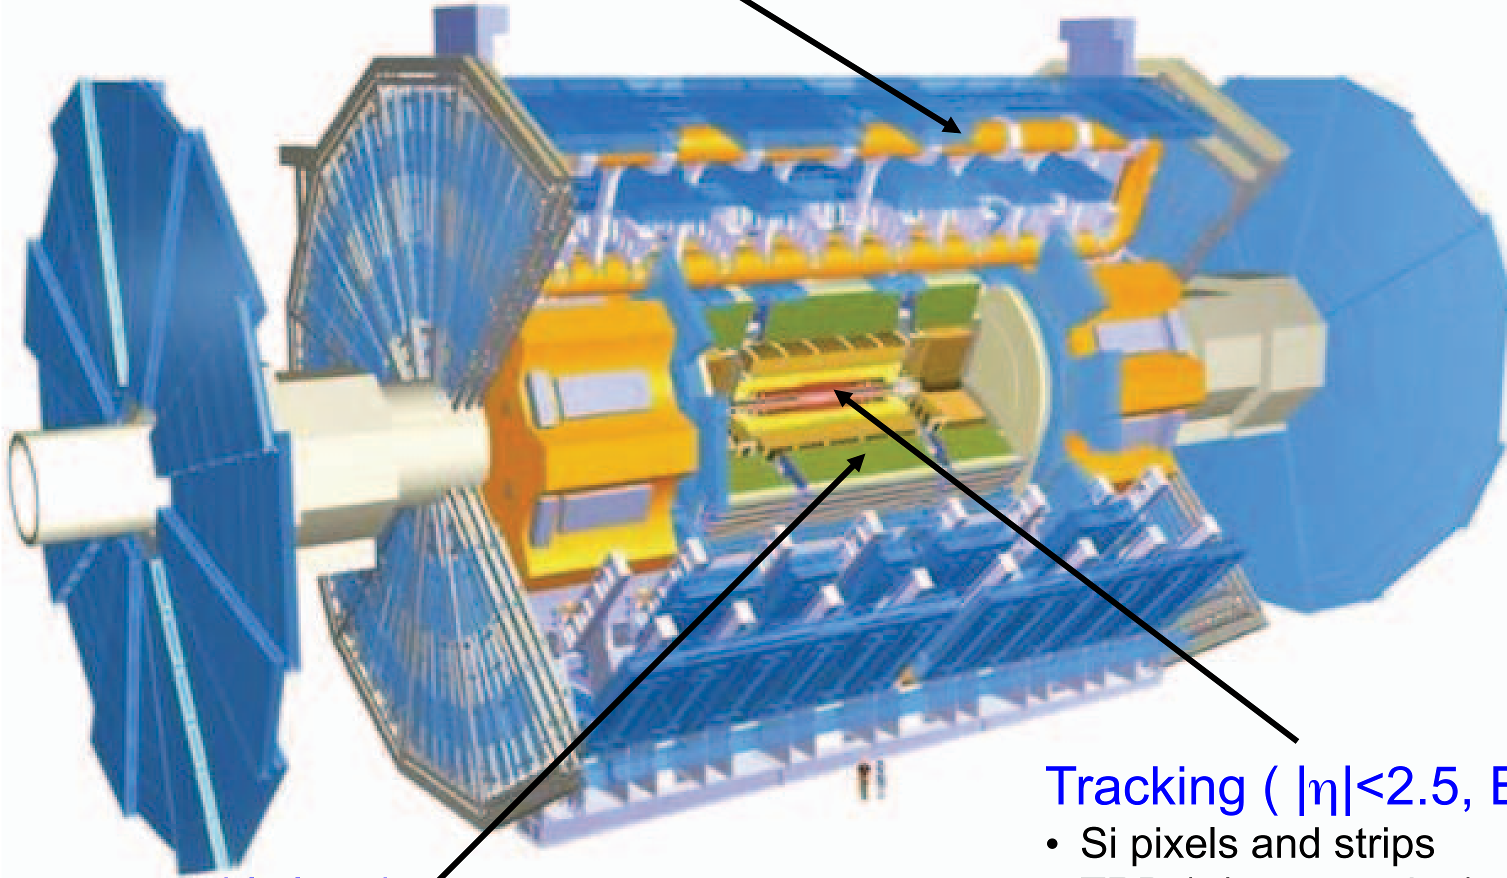

Muon Spectrometer ($|\eta|<2.7$)
- air-core toroids with muon chambers

Tracking ($|\eta|<2.5$, B=2T)
- Si pixels and strips
- TRD (e/π separation)

Calorimetry ($|\eta|<5$)
- EM : Pb-LAr
- HAD : Fe/scintillator (central), Cu/W-Lar (fwd)

CMS Detector Design

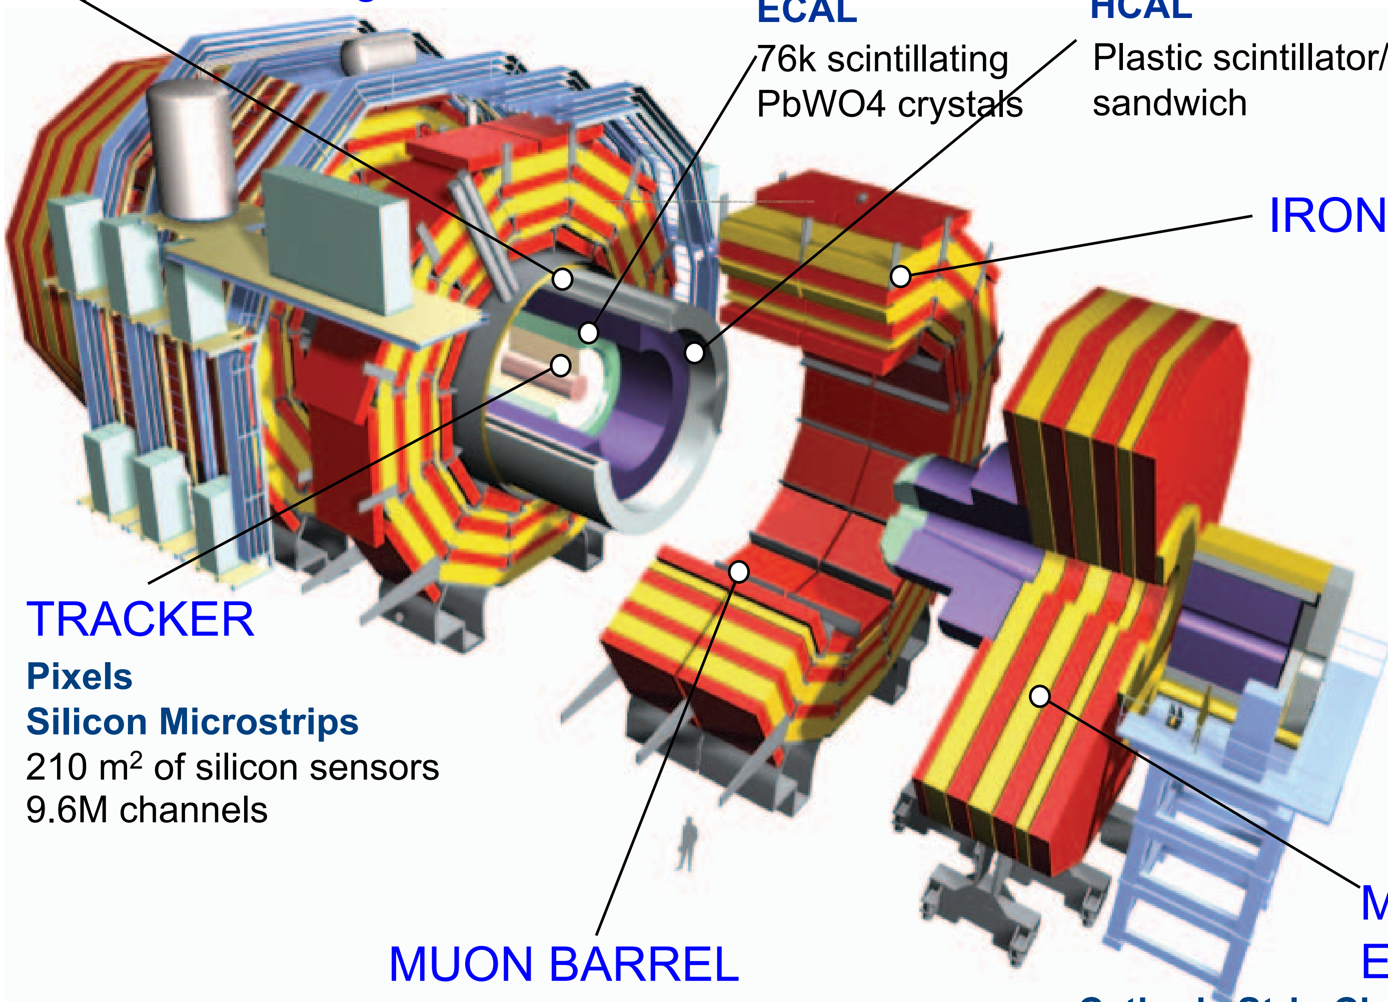

CTEQ

ATLAS & CMS Trigger & Readout Structure

≈ 30 Collisions/25ns
(10^9 event/sec)

10^7 channels
(10^{16} bit/sec)

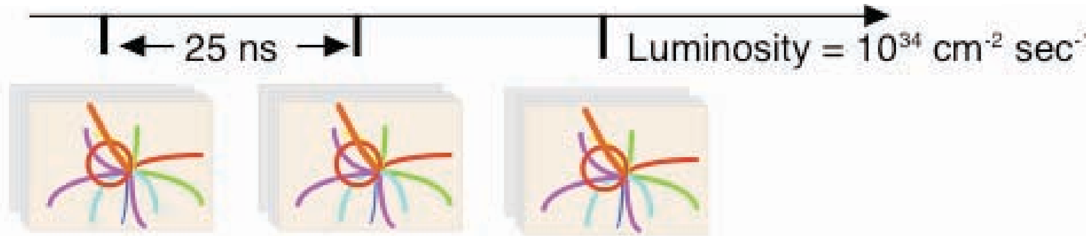

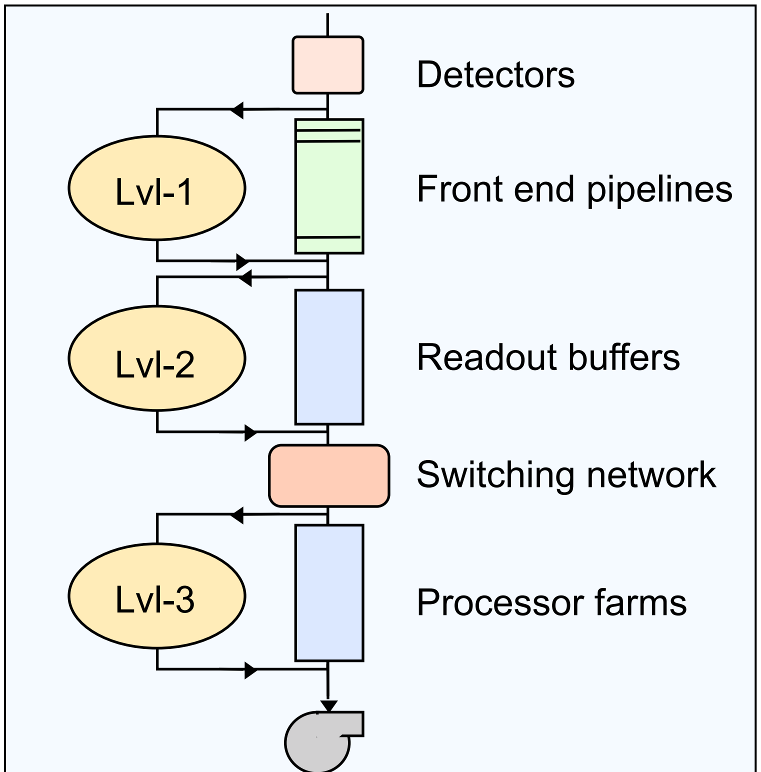

ATLAS: 3 physical levels

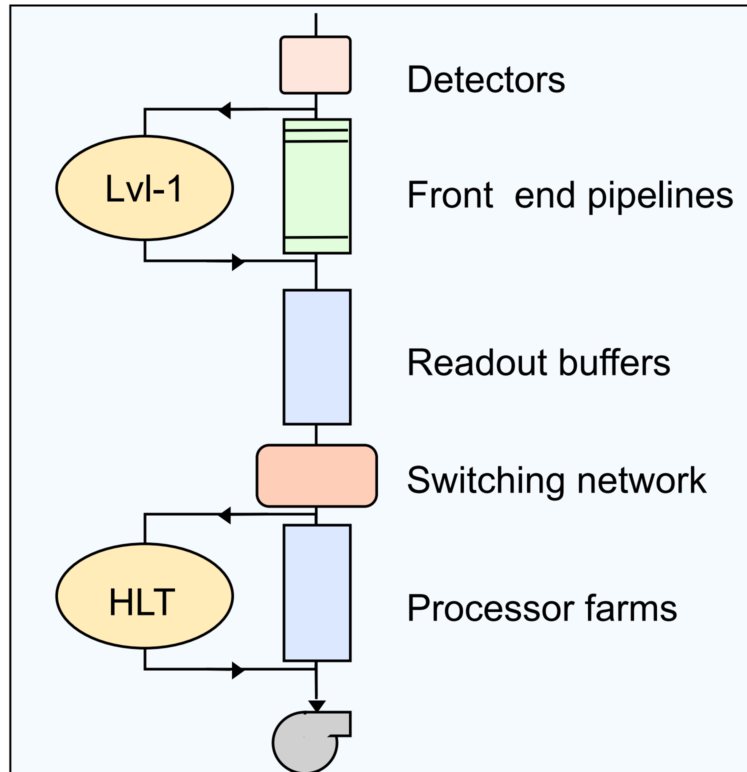

CMS: 2 physical levels

ATLAS & CMS Trigger Data

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

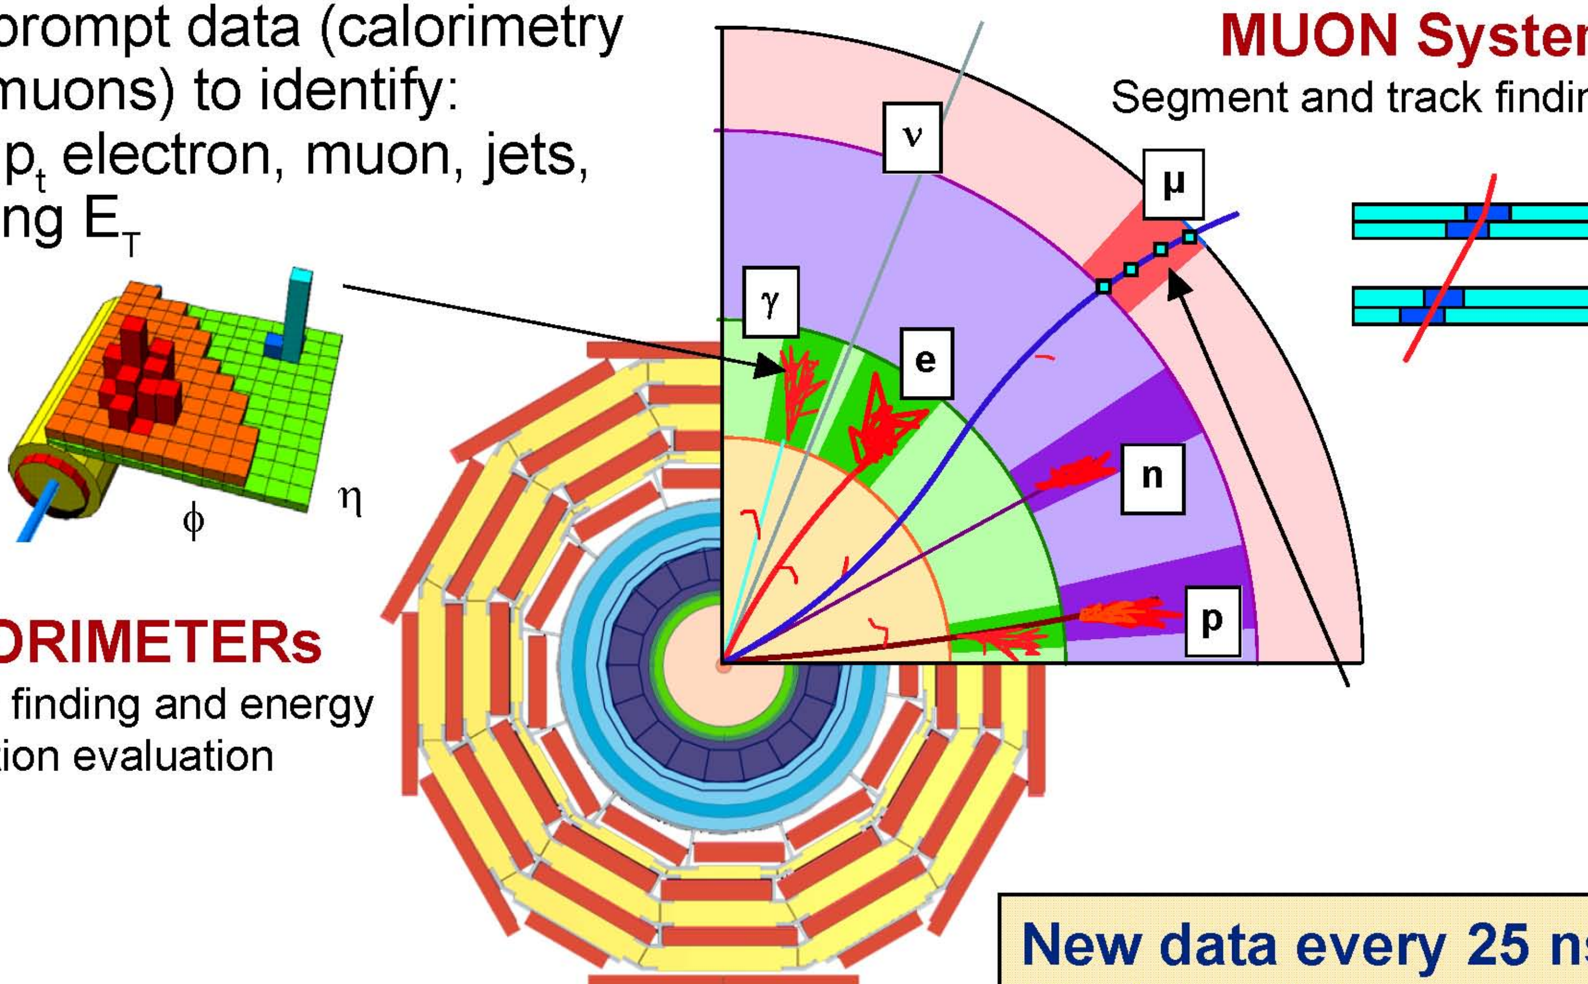

MUON System
Segment and track finding

CALORIMETERs
Cluster finding and energy deposition evaluation

New data every 25 ns
Decision latency ~ μs

ATLAS & CMS Level 1: Only Calorimeter & Muon

High Occupancy in high granularity tracking detectors

- **Pattern recognition much faster/easier**

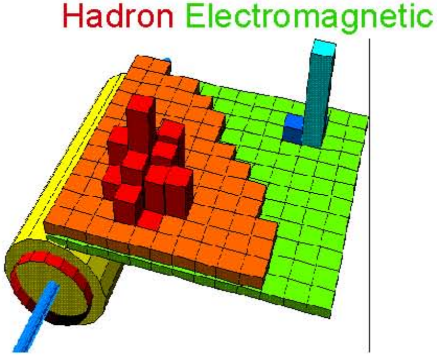

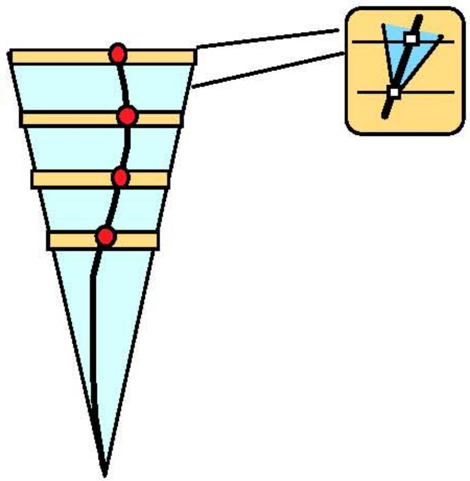

Simple Algorithms

Small amounts of data

- **Compare to tracker info**

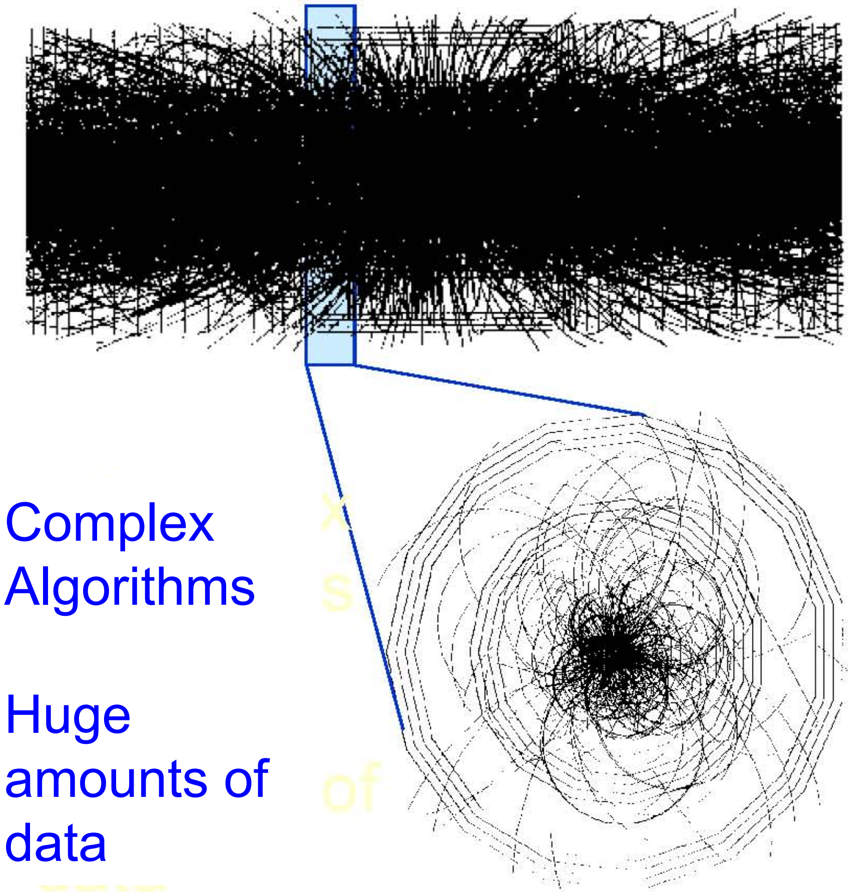

Complex Algorithms

Huge amounts of data

ATLAS Trigger & DAQ Architecture

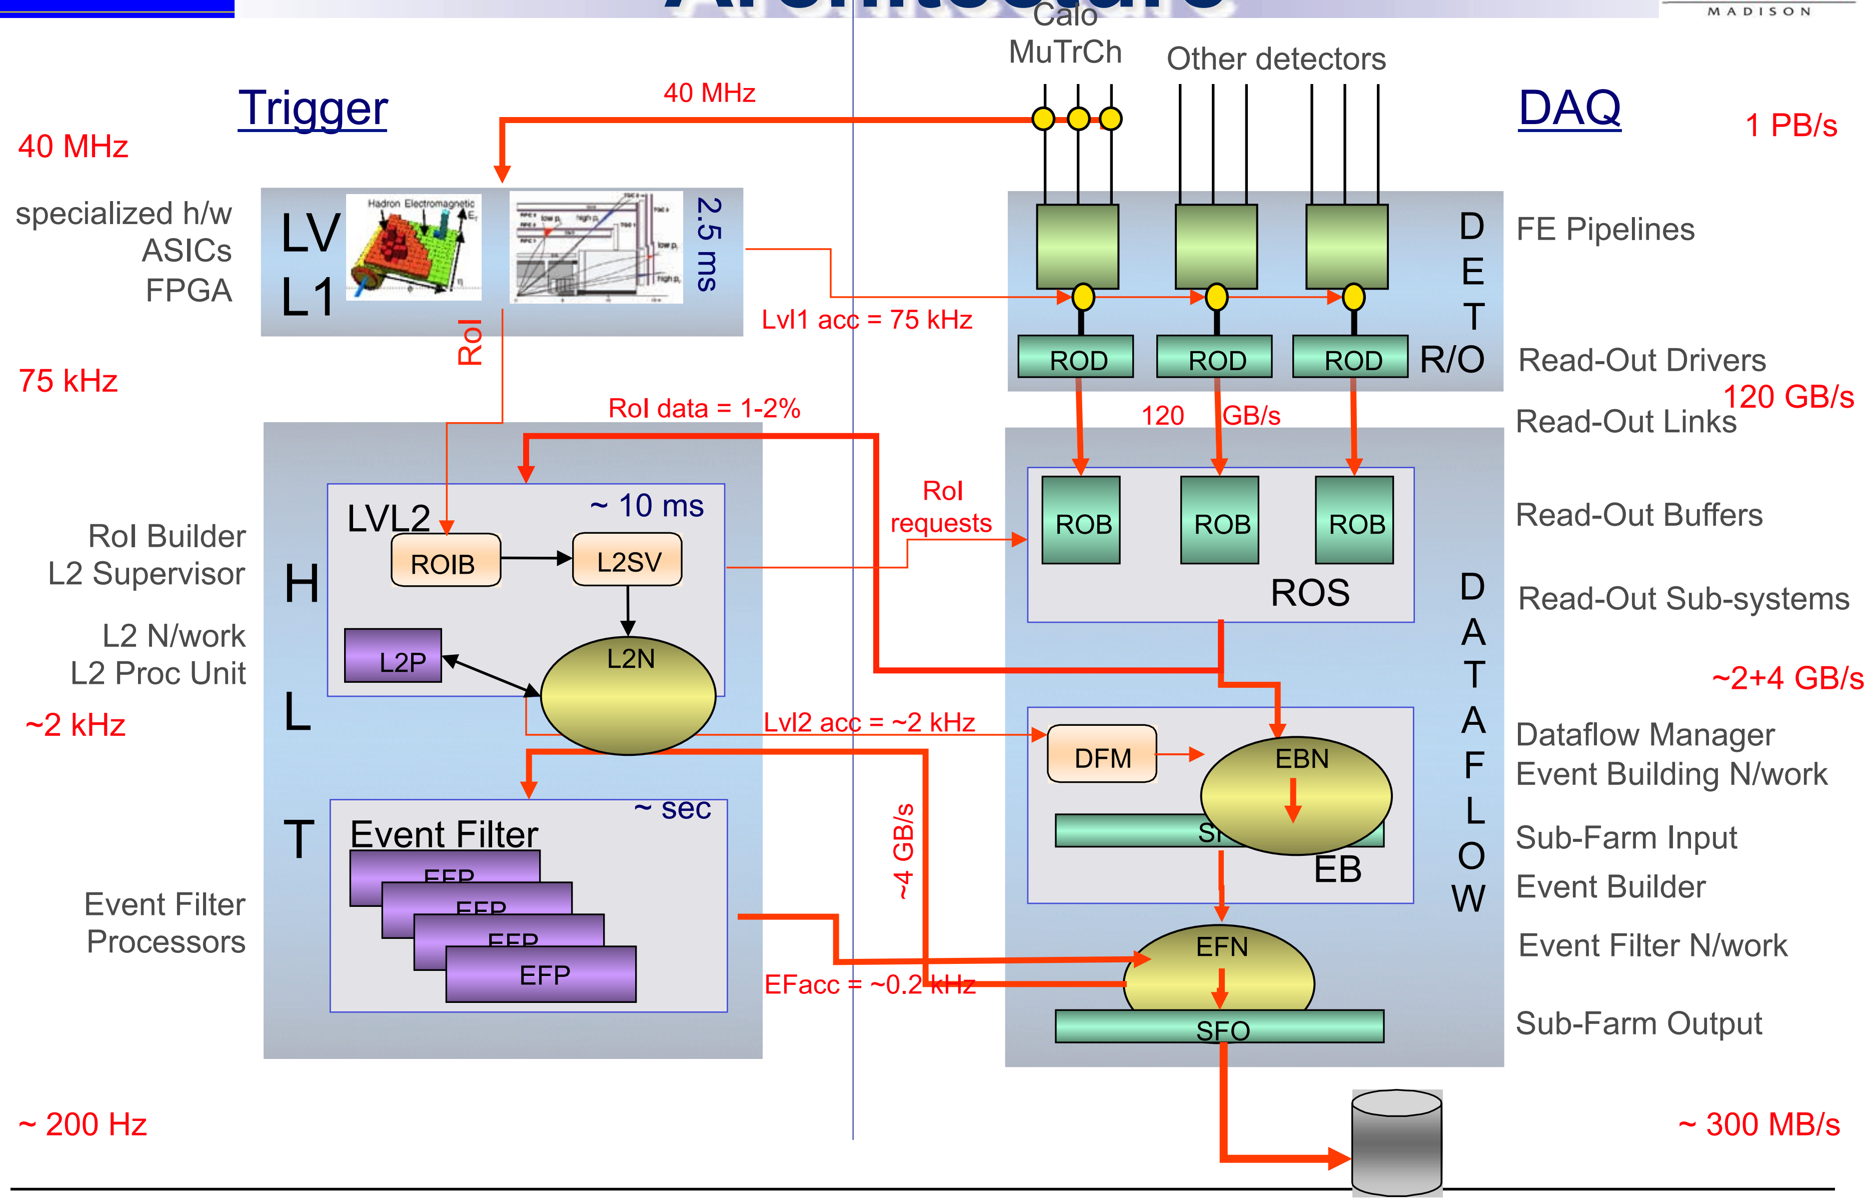

ATLAS Three Level Trigger Architecture

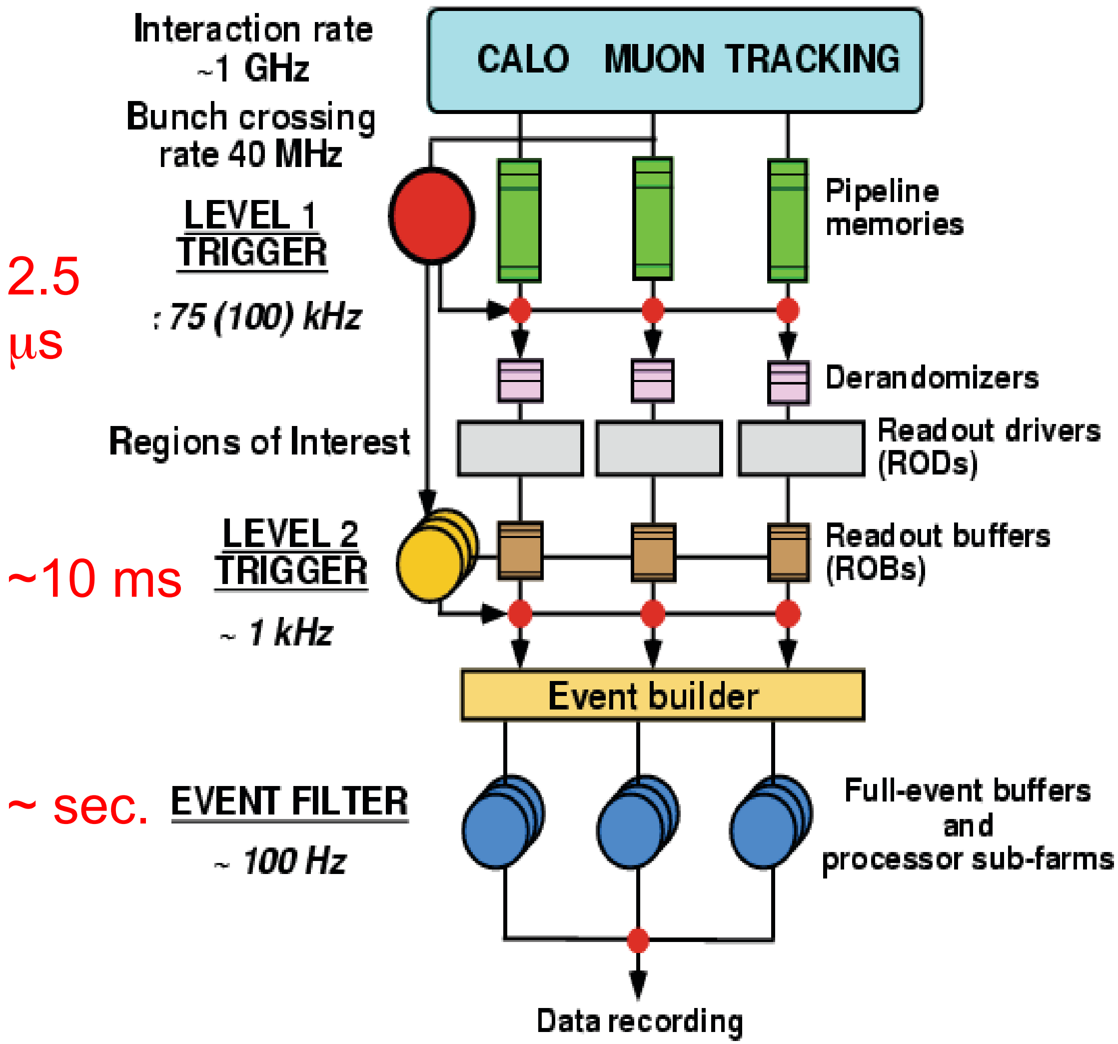

- **LVL1 decision** made with **calorimeter** data with coarse granularity and **muon trigger chambers** data.
  - Buffering on detector
- **LVL2 uses Region of Interest data** (ca. 2%) with full granularity and combines information from all detectors; performs fast rejection.
  - Buffering in ROBs
-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

ATLAS Level-1 Trigger - Muons & Calorimetry

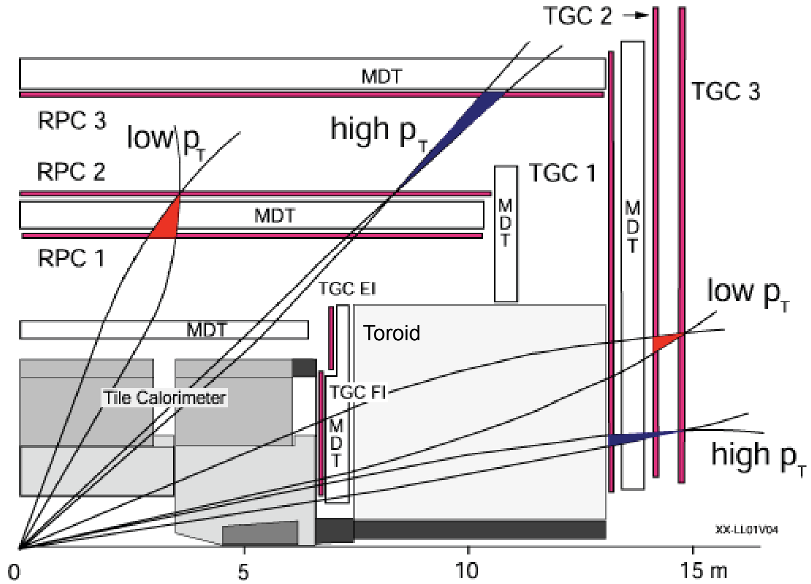

Muon Trigger looking for coincidences in muon trigger chambers
2 out of 3 (low-p_T; >6 GeV) and
3 out of 3 (high-p_T; > 20 GeV)

Trigger efficiency 99% (low-p_T) and 98% (high-p_T)

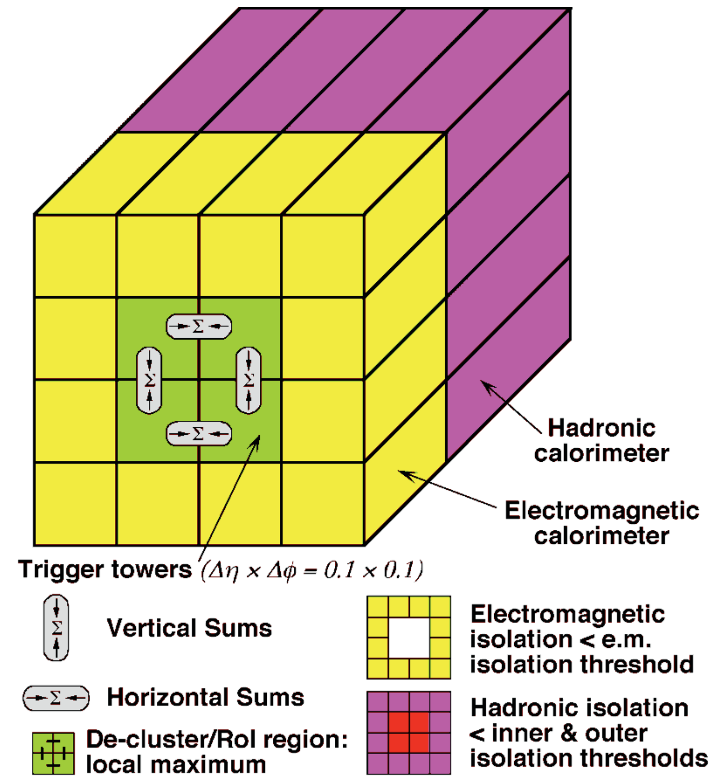

Calorimetry Trigger looking for e/γ/τ + jets

- Various combinations of cluster sums and isolation criteria
- $\Sigma E_T^{em,had}$, E_T^{miss}

ATLAS LVL1 Trigger

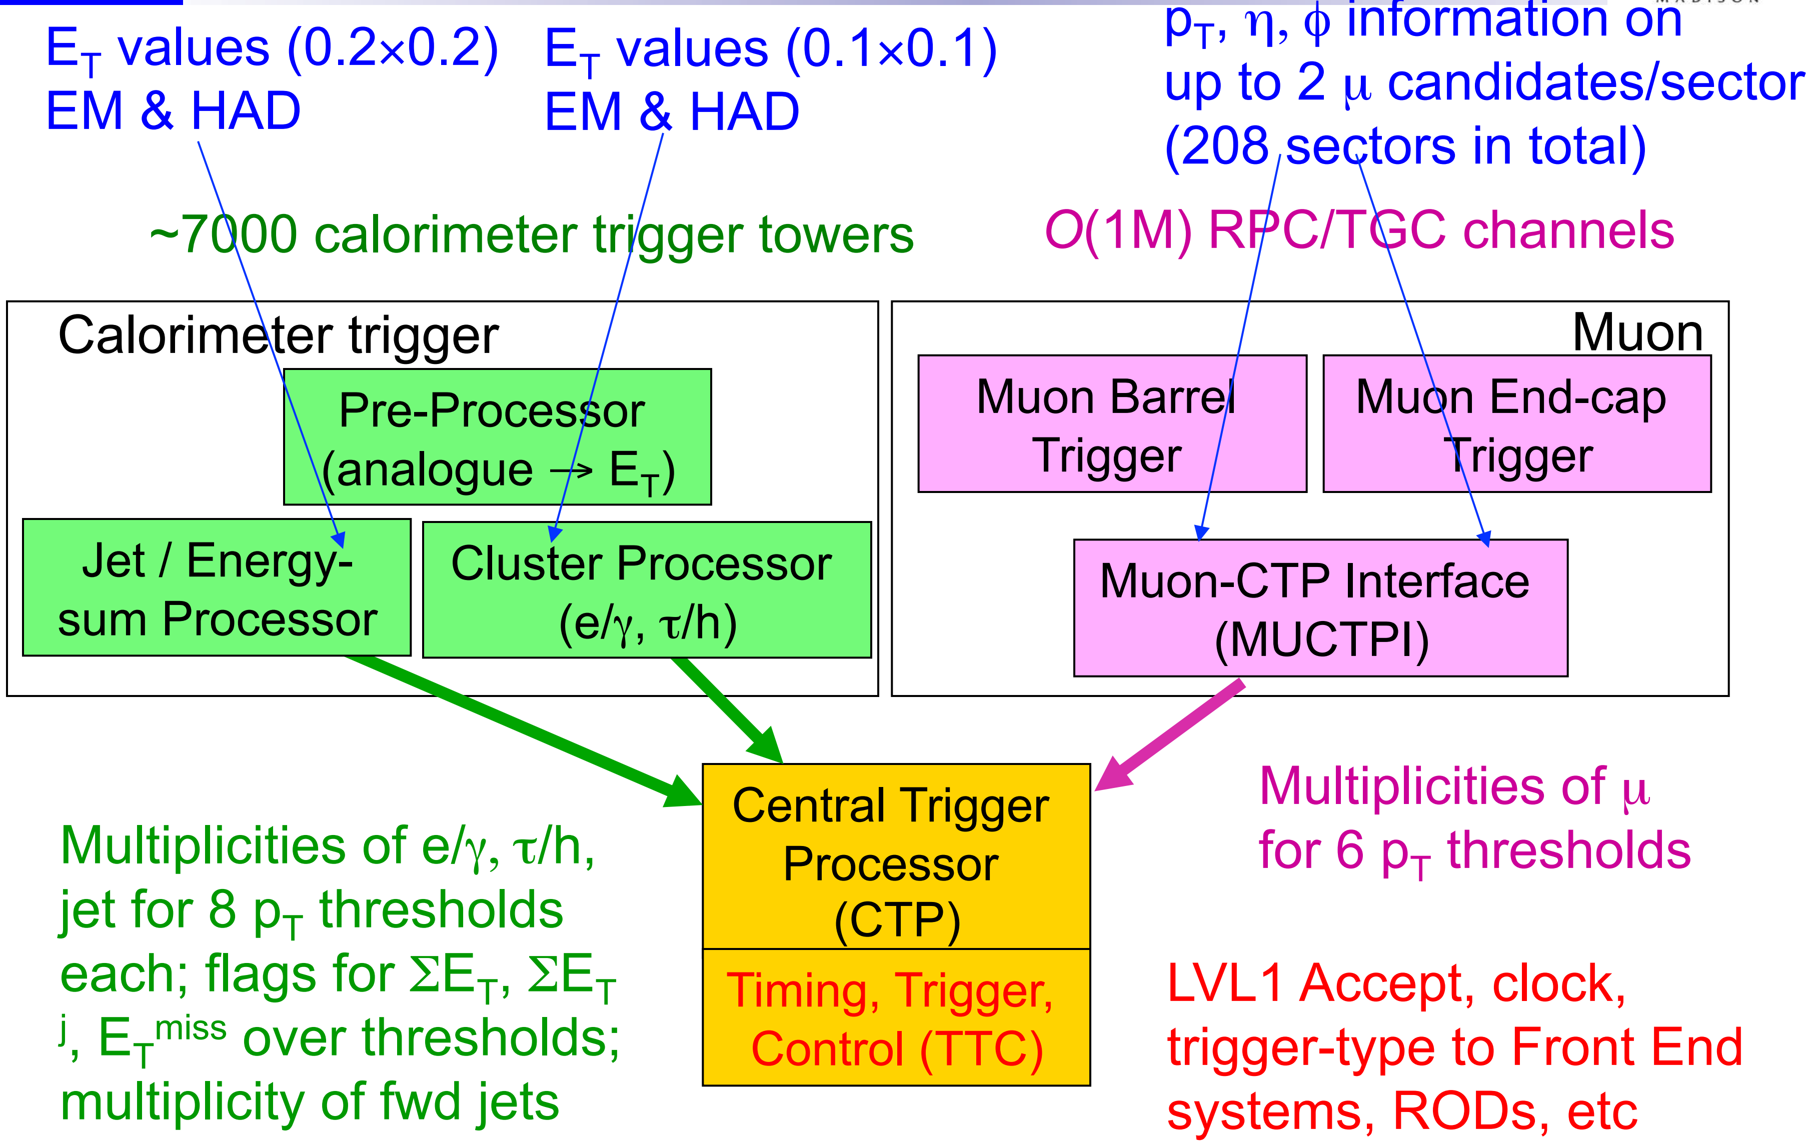

RoI Mechanism

LVL1 triggers on high p_T objects

- **Caloriemeter cells and muon chambers to find e/γ/τ-jet-μ candidates above thresholds**

LVL2 uses Regions of Interest as identified by Level-1

- Local data reconstruction, analysis, and sub-detector matching of RoI data

The total amount of RoI data is minimal

- ~2% of the Level-1 throughput but it has to be extracted from the rest at 75 kHz

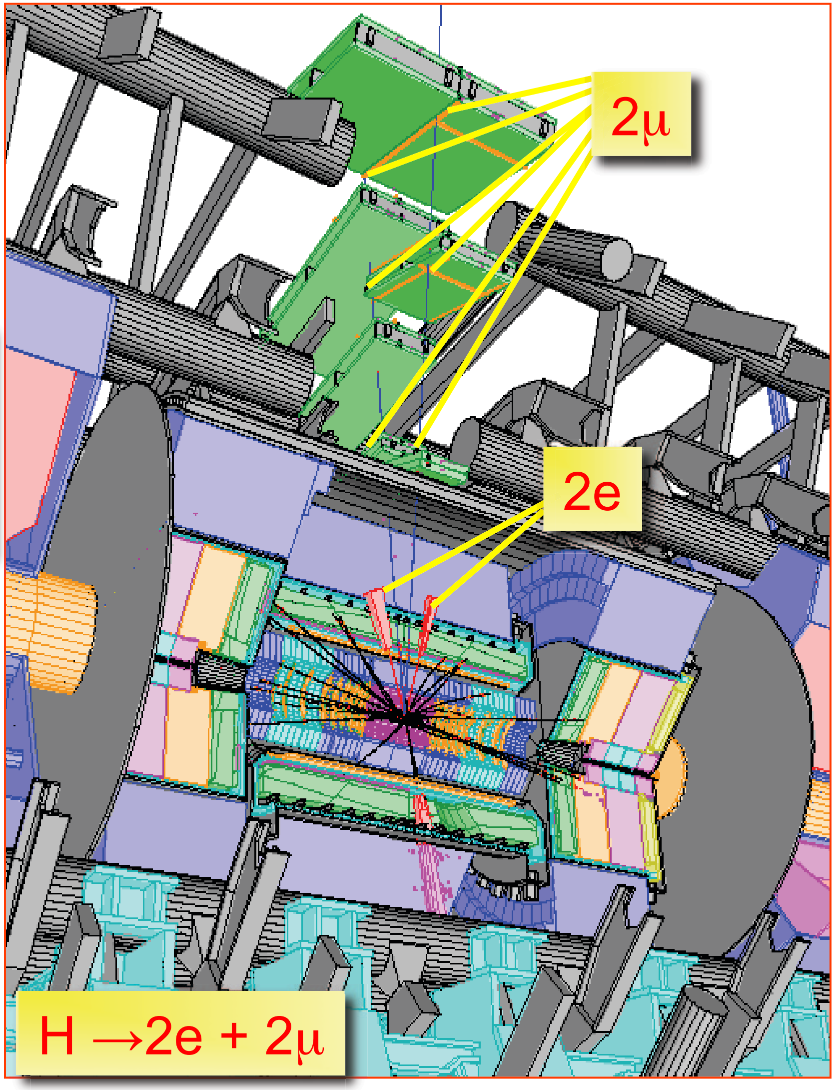

C T E Q

CMS Trigger Levels

THE UNIVERSITY of WISCONSIN MADISON

40 MHz

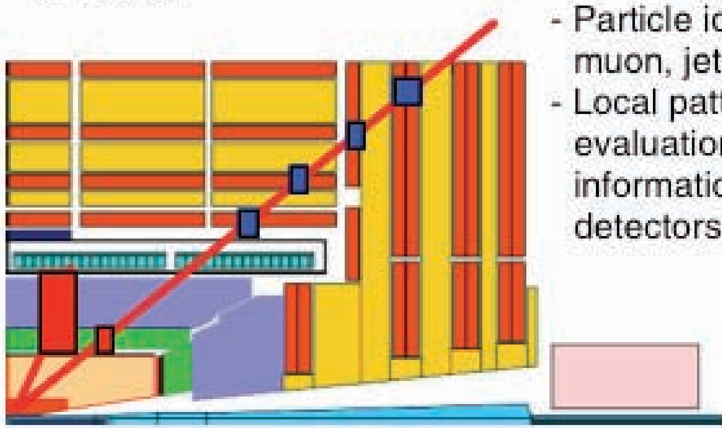

Level-1. Specialized processors

- Particle identification: high p_T electron, muon, jets, missing E_T
- Local pattern recognition and energy evaluation on prompt macro-granular information from calorimeter and muon detectors

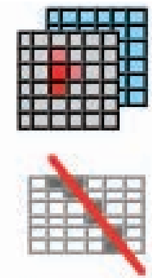

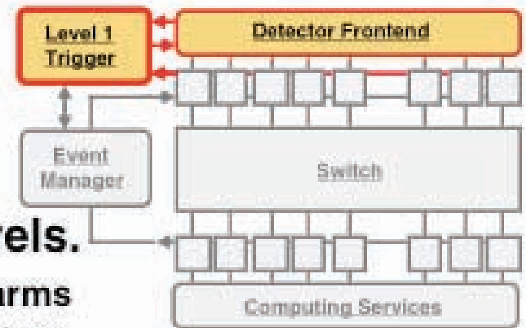

Up to 100 kHz

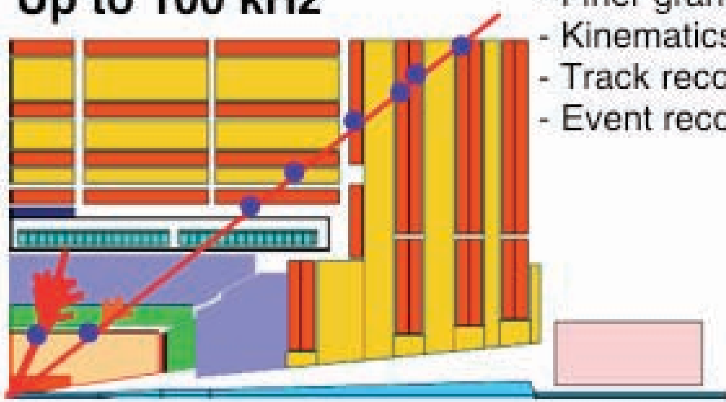

High trigger levels.

Network and CPU farms

- Clean particle signature
- Finer granularity precise measurement
- Kinematics. effective mass cuts & event topology
- Track reconstruction and detector matching
- Event reconstruction and analysis

≈ 100 Hz

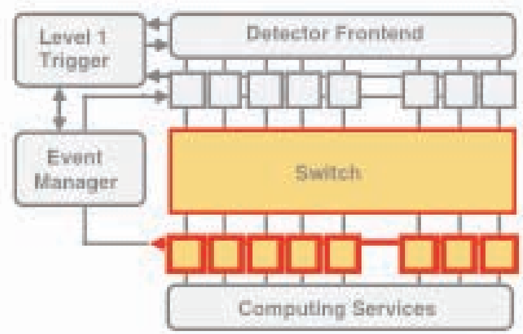

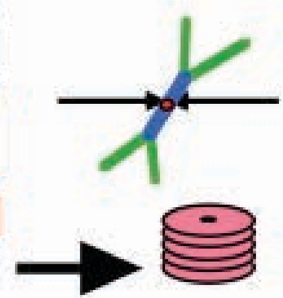

CMS Level-1 Trigger & DAQ

Overall Trigger & DAQ Architecture: 2 Levels:

Level-1 Trigger:

- 25 ns input
- 3.2 μs latency

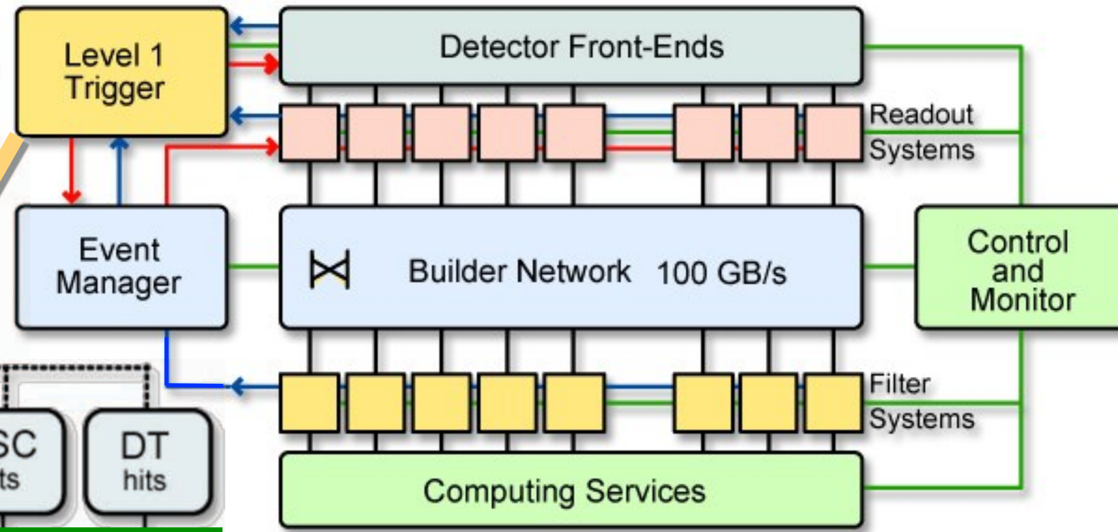

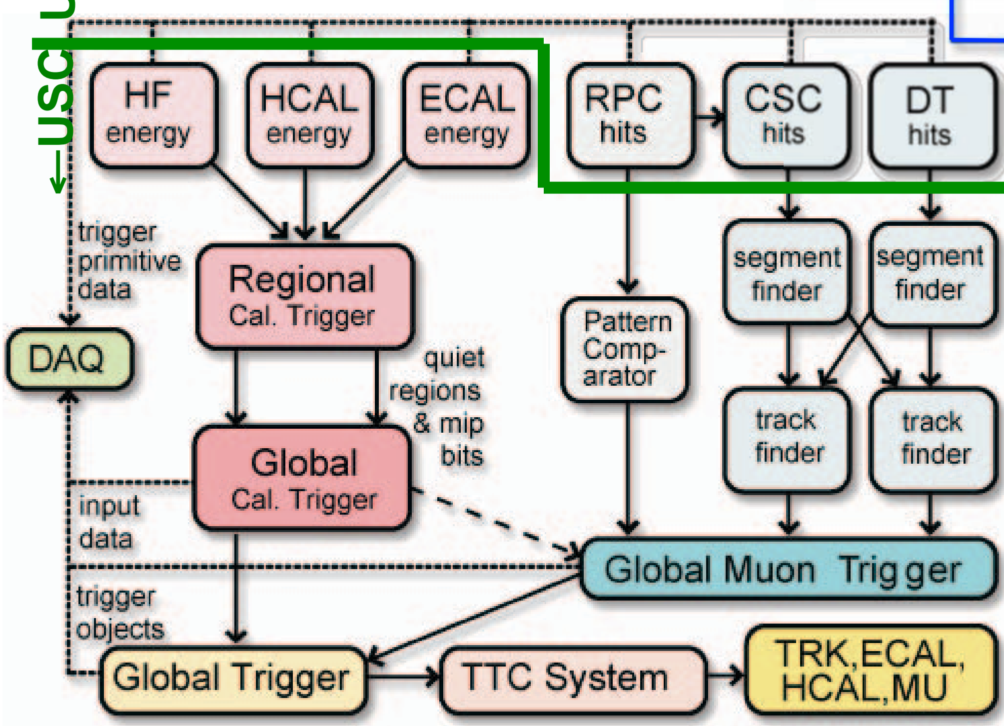

Interaction rate: 1 GHz

Bunch Crossing rate: 40 MHz

Level 1 Output: 100 kHz (50 initial)

Output to Storage: 100 Hz

Average Event Size: 1 MB

Data production 1 TB/day

Calorimeter Trigger Processing

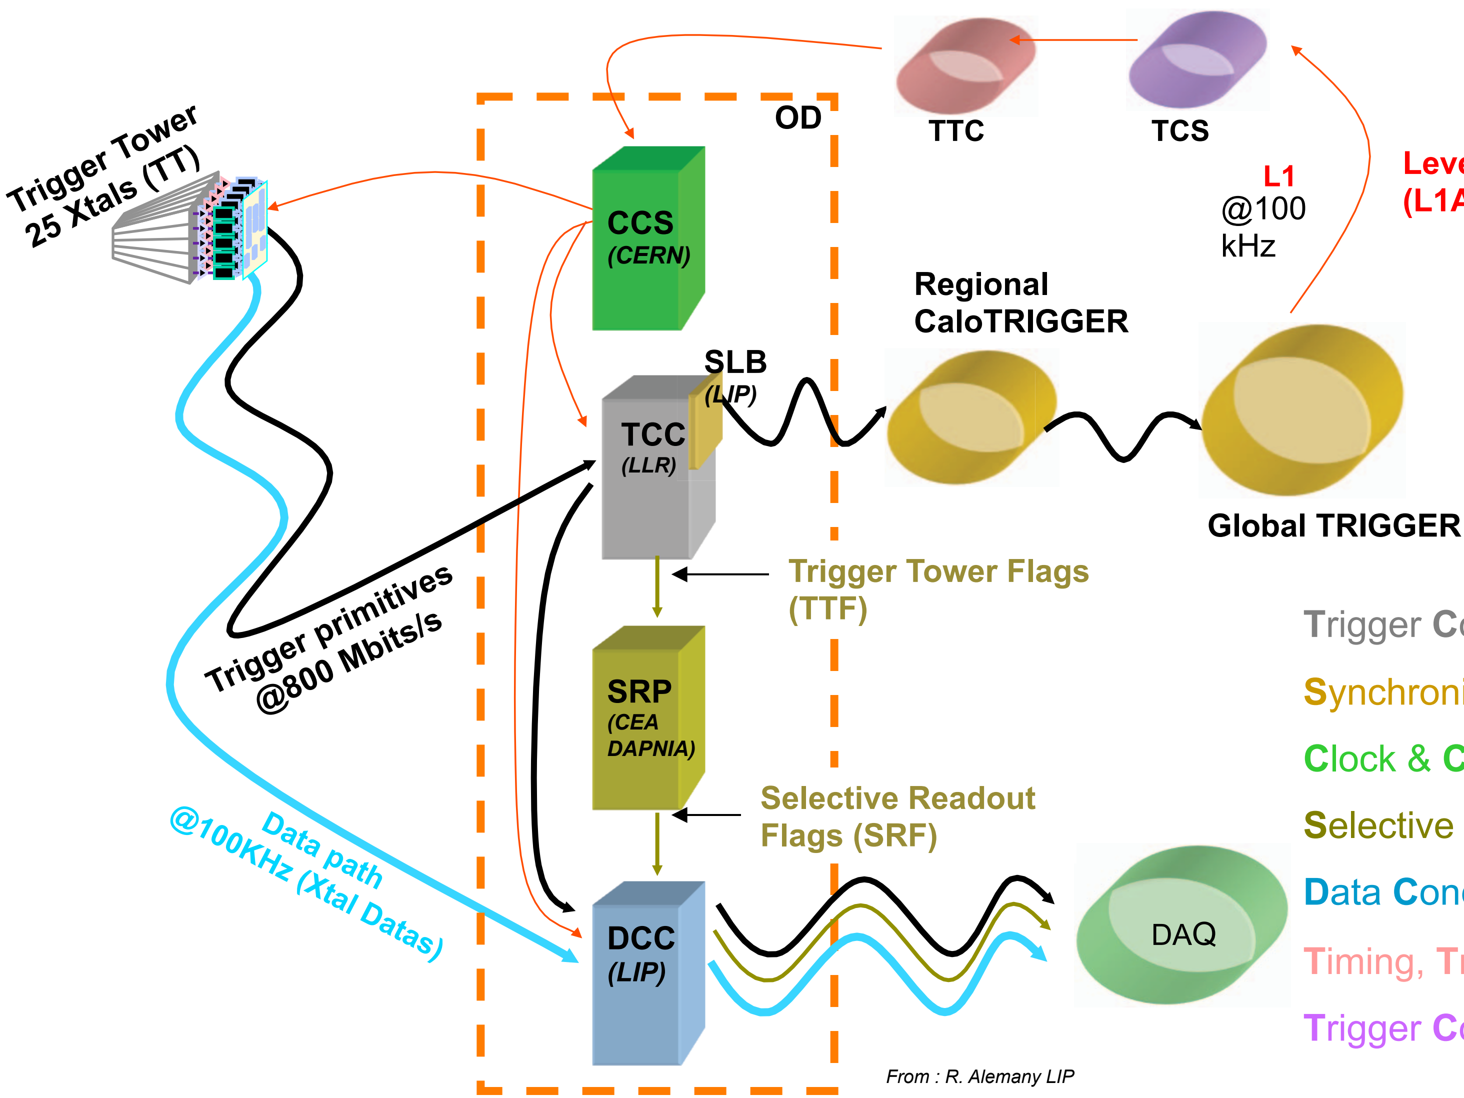

Trigger Concentrator Card

Synchronisation & Link Board

Clock & Control System

Selective Readout Processor

Data Concentrator Card

Timing, Trigger & Control

Trigger Control System

ECAL Trigger Primitives

In the trigger path, **digital filtering** followed by a **peak finder** is applied to energy sums **(L1 Filter)**

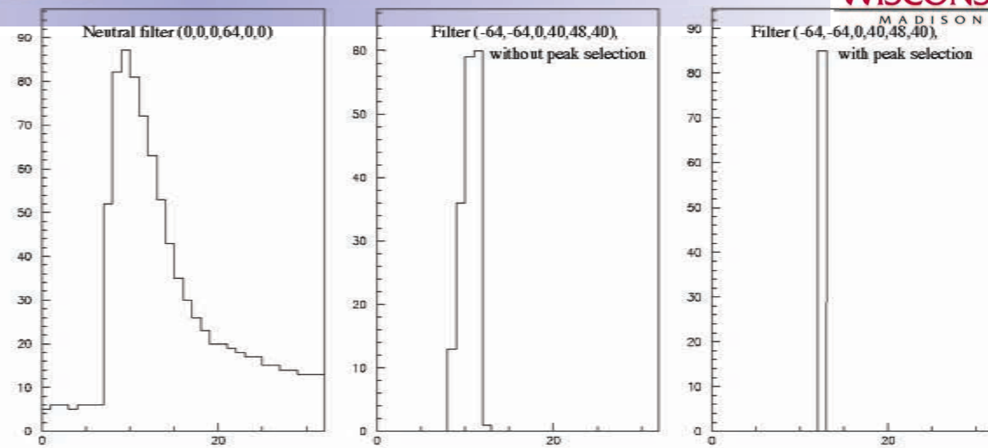

Efficiency for energy sums above 1 GeV should be close to 100% (depends on electronics noise)

Pile-up effect: for a signal of 5 GeV the efficiency is close to 100% for pile-up energies up to 2 GeV (CMS)

Test beam results (45 MeV per xtal):

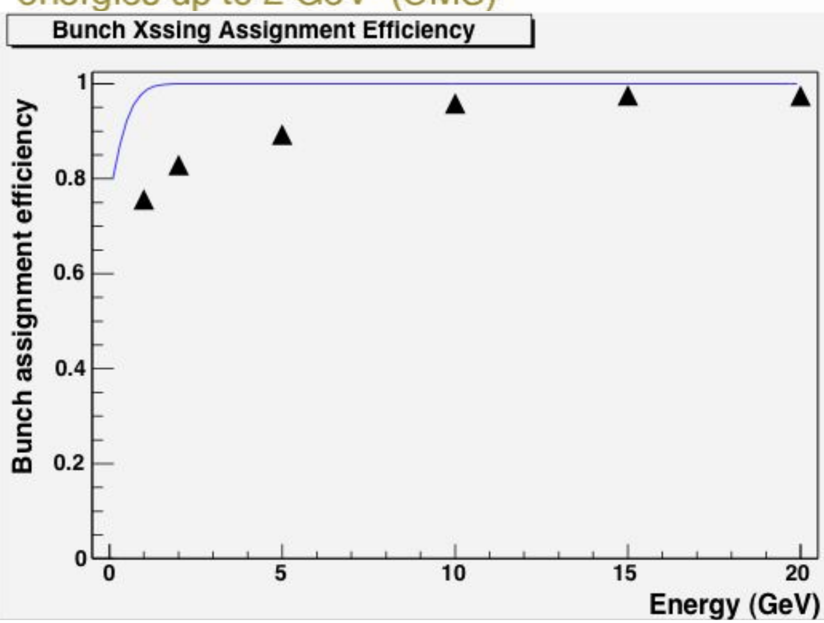

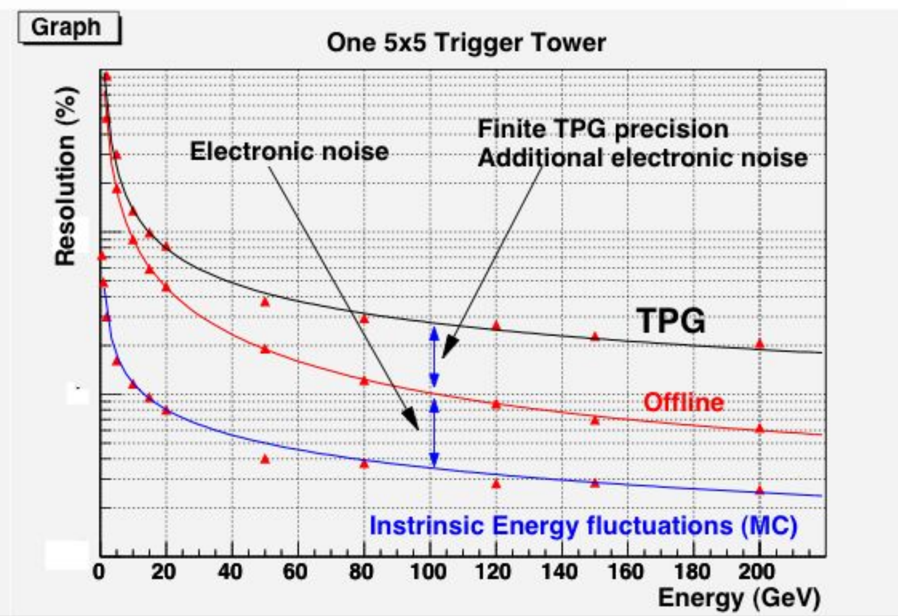

C T E Q

CMS Electron/Photon Algorithm

Trigger Primitive Generator

Fine grain — Flag Max of (, , ,) & Sum ET

Regional Calorimeter Trigger

E_T cut — + Max ( ) > Threshold

AND

Longitudinal cut (H/E) — / < 0.05

AND

Isolation, Hadronic & EM — < 2 GeV

AND

One of (, , ,) < 1 GeV

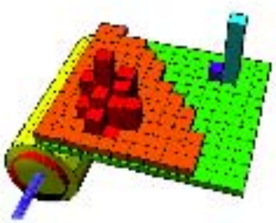

ELECTRON or PHOTON

CMS τ / Jet Algorithm

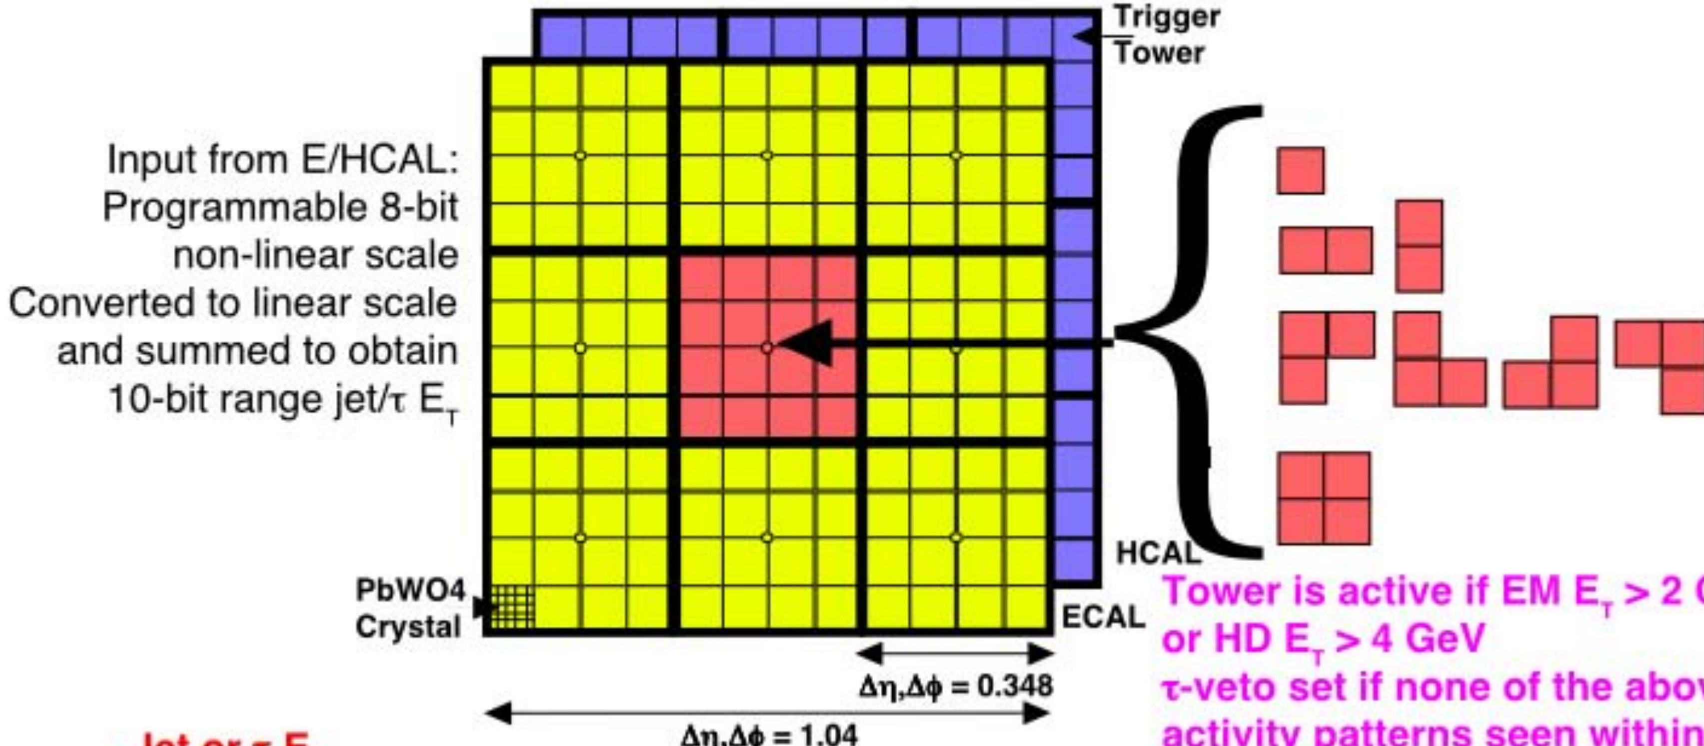

Tower is active if EM E_T > 2 GeV or HD E_T > 4 GeV

τ-veto set if none of the above activity patterns seen within 4x4

Jet or τ E_T

- 12x12 trigger tower E_T sums in 4x4 region steps with central region > others
- Larger trigger towers in HF but ~ same jet region size, 1.5 η x 1.0 φ

τ algorithm (isolated narrow energy deposits), within -2.5 < η < 2.5

- Redefine jet as τ jet if none of the nine 4x4 region τ-veto bits are on

Output

- Top 4 τ-jets and top 4 jets in central rapidity, and top 4 jets in forward rapidity

CMS Muon Chambers

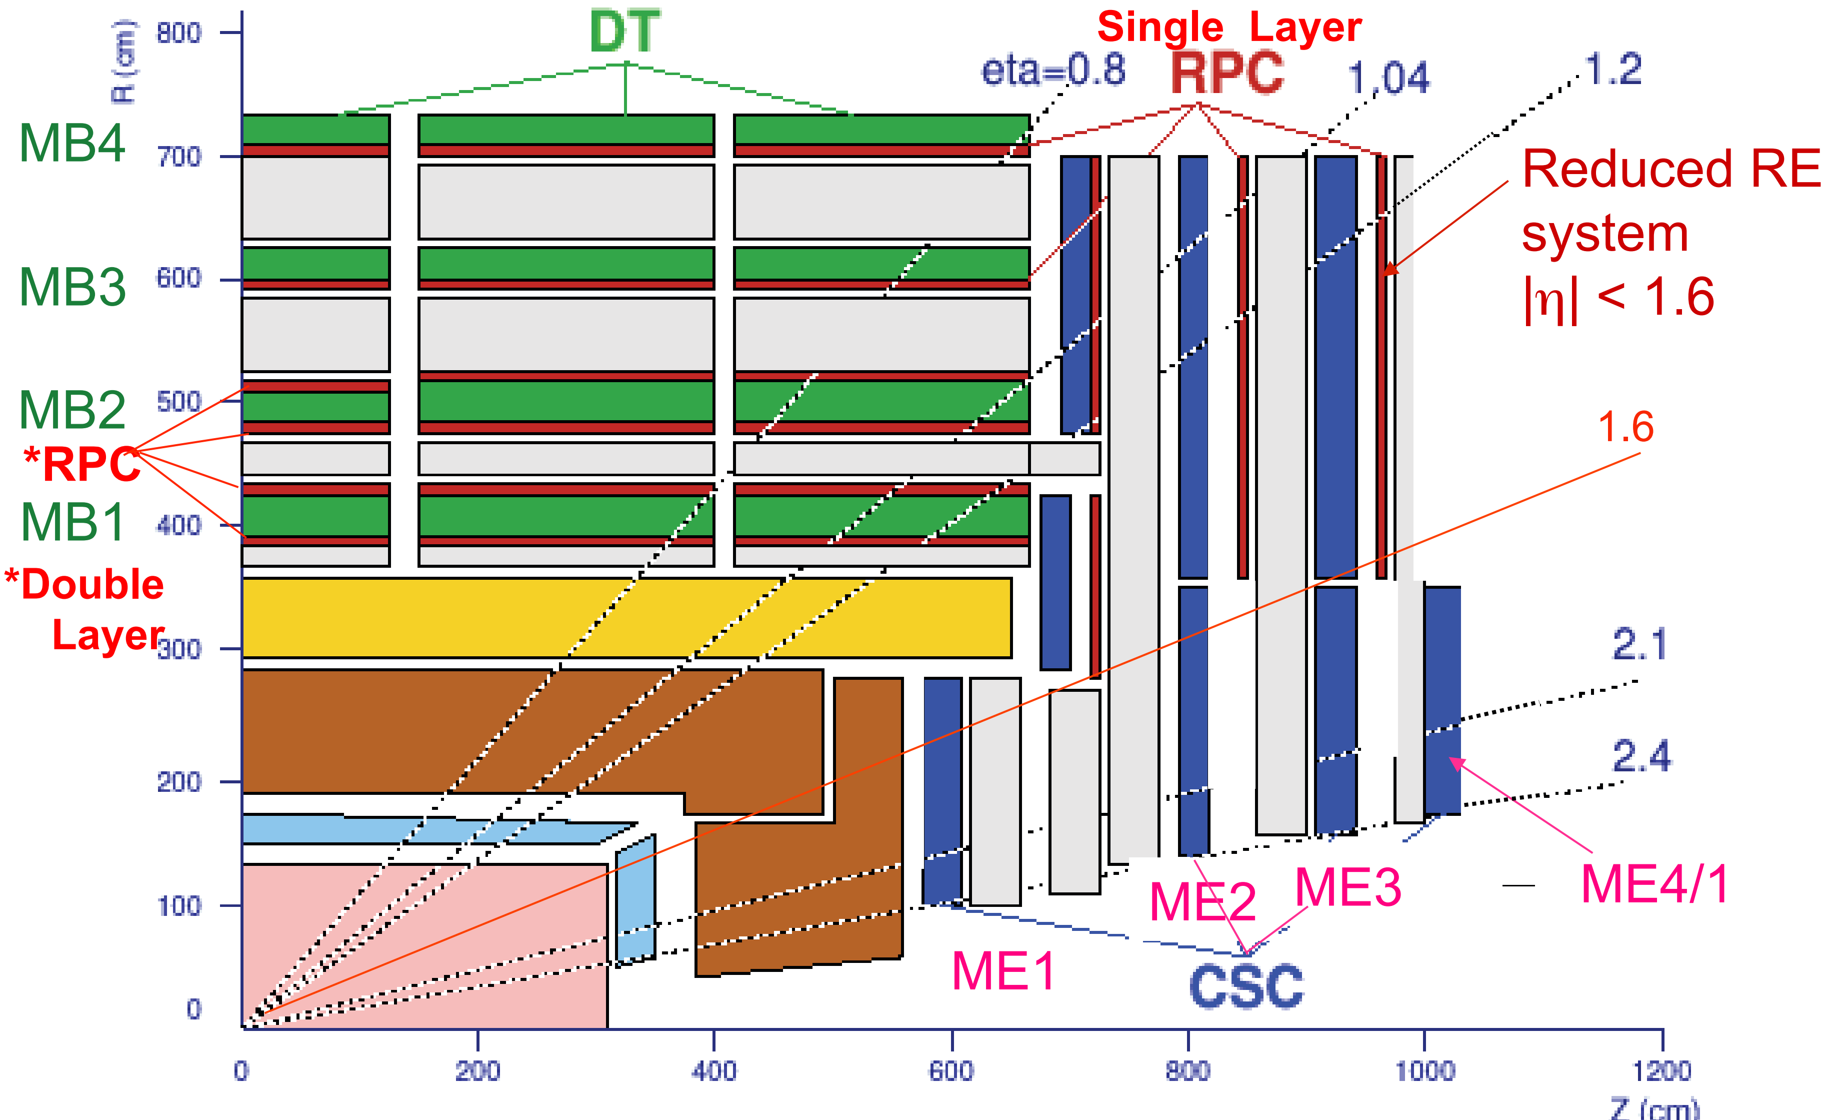

Muon Trigger Overview

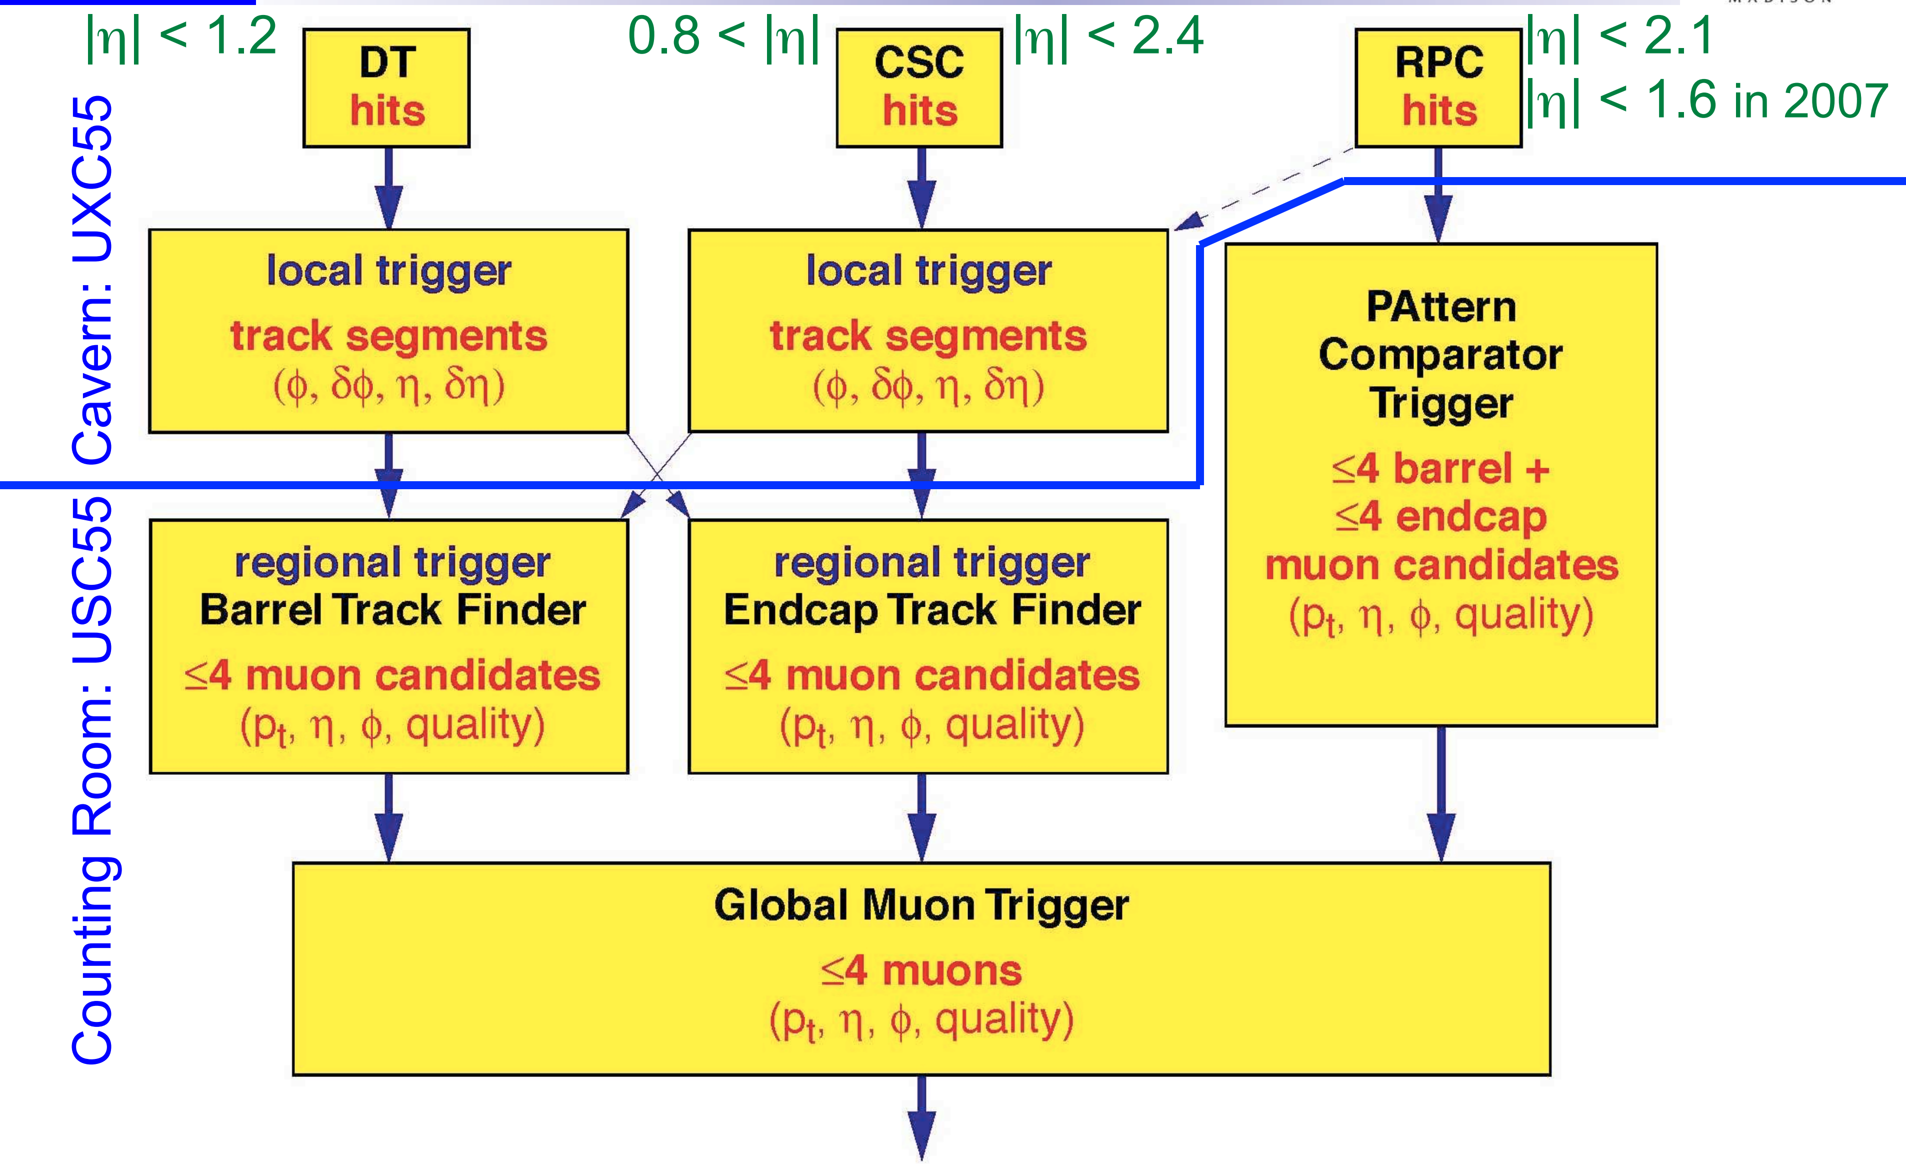

CMS Muon Trigger Primitives

RPC pattern recognition
- Pattern catalog
- Fast logic

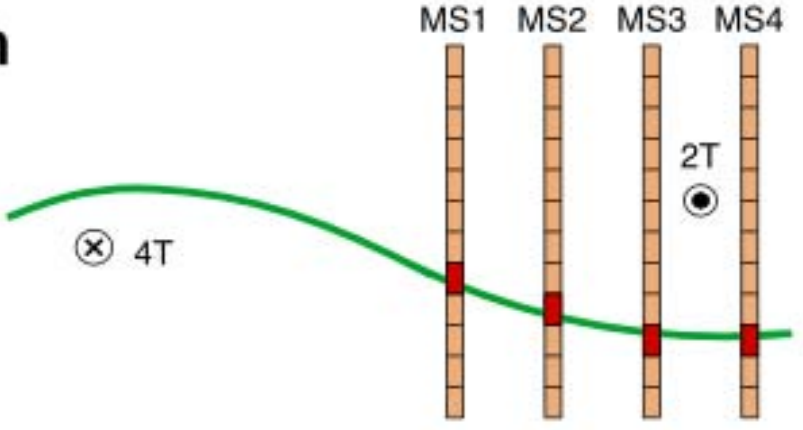

Memory to store patterns

Fast logic for matching

FPGAs are ideal

DT and CSC track finding:
- Finds hit/segments
- Combines vectors
- Formats a track
- Assigns p_t value

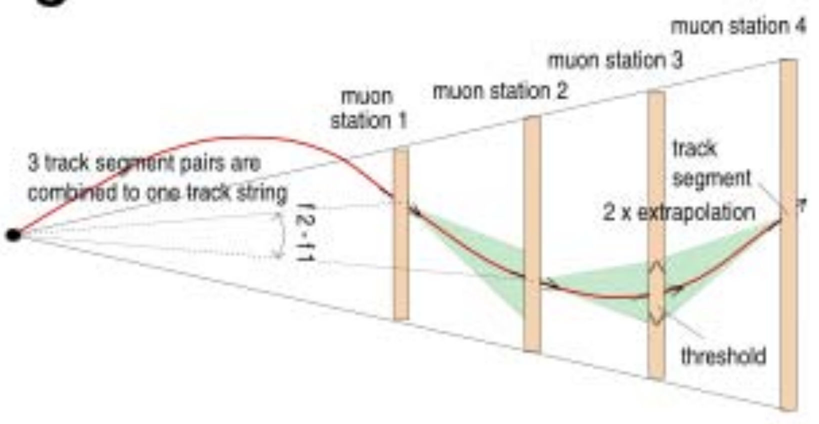

CMS Muon Trigger Track Finders

Drift Tubes (DT)

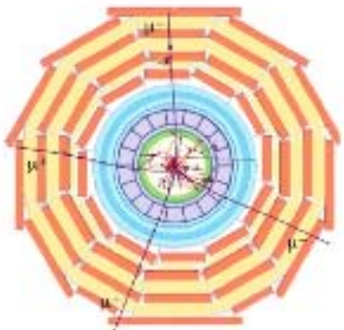

Drift Tubes

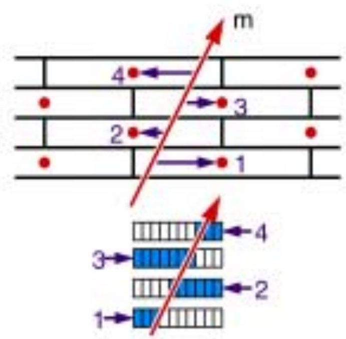

Meantimers recognize tracks and form vector / quartet.

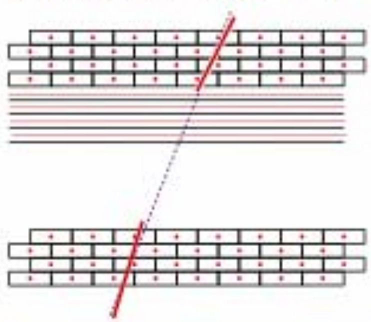

Correlator combines them into one vector / station.

Cathod Strip Chambers (CSC)

CSC

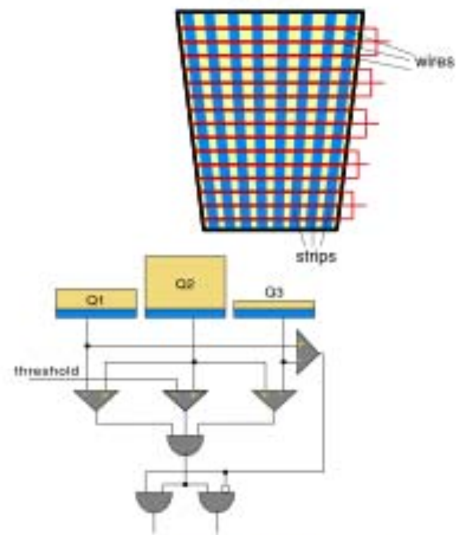

Comparators give 1/2-strip resol.

Hit strips of 6 layers form a vector

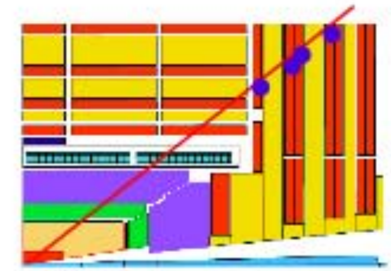

Sort based on P_T, Quality - keep loc.

Combine at next level - match

Sort again - Isolate?

Top 4 highest P_T and quality muons with location coord.

Match with RPC

Improve efficiency and quality

CMS Global Trigger

Input:

- Jets: 4 Central, 4 Forward, 4 Tau-tagged, & Multiplicities
- Electrons: 4 Isolated, 4 Non-isolated
- 4 Muons (from 8 RPC, 4 DT & 4 CSC w/P_t & quality)
    - All above include location in η and ϕ
- Missing E_T & Total E_T

Output

- L1 Accept from combinations & proximity of above

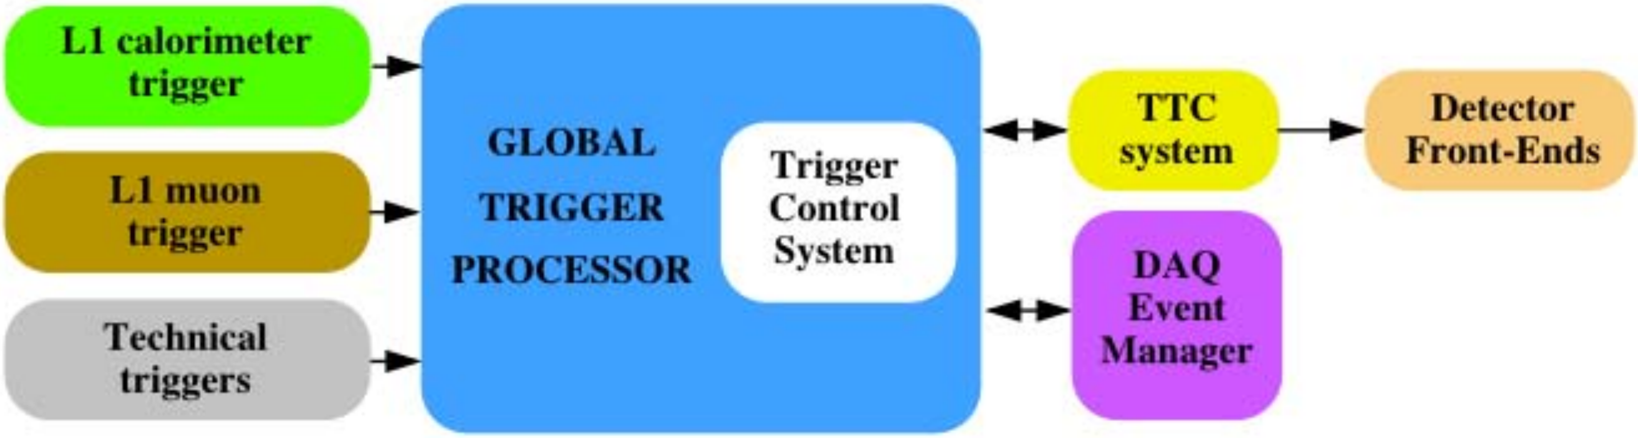

Global L1 Trigger Algorithms

Particle Conditions

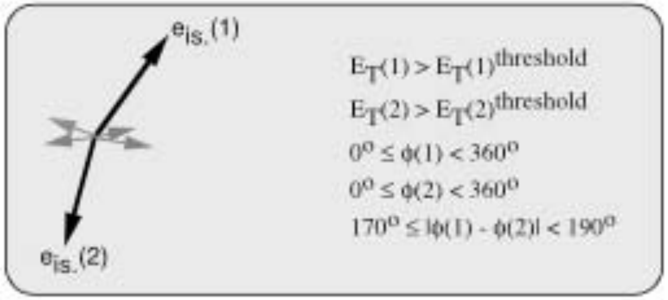

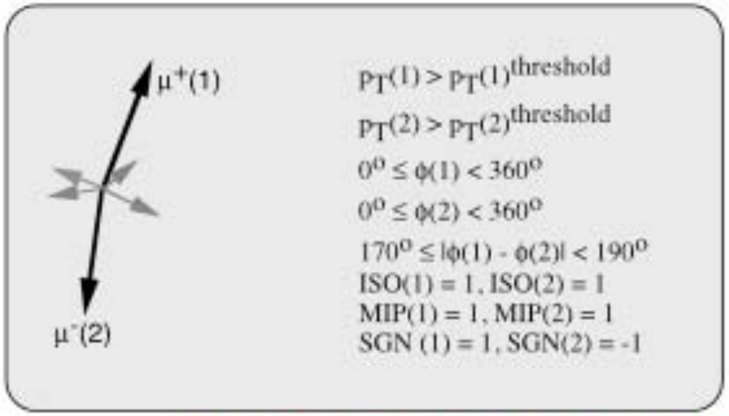

Logical Combinations

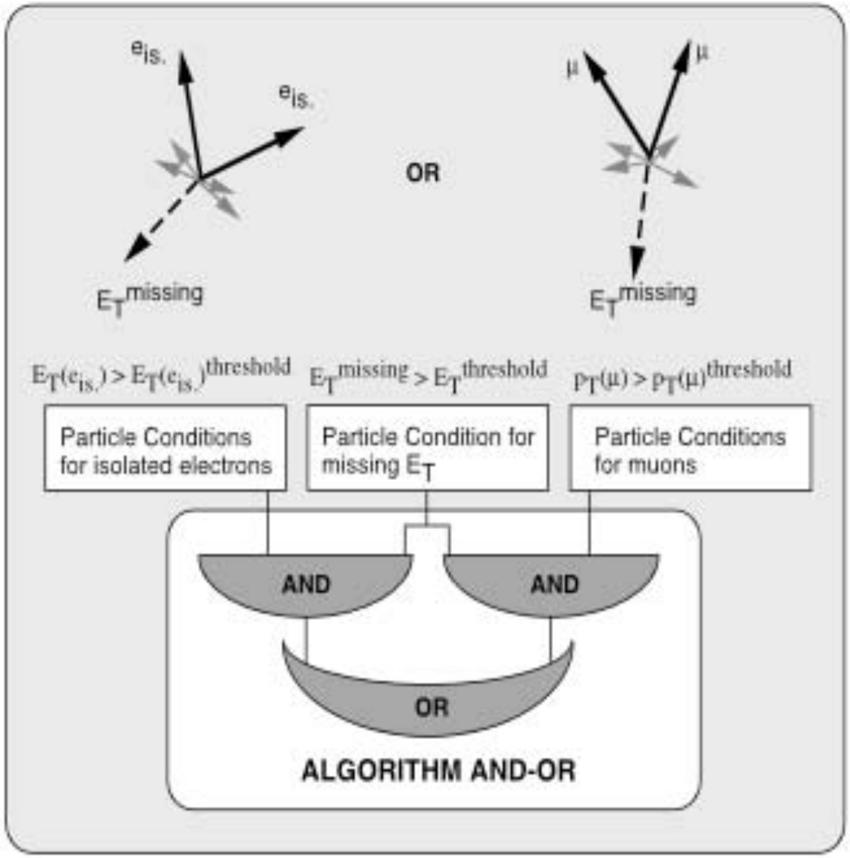

Flexible algorithms implemented in FPGAs
100s of possible algorithms can be reprogrammed

High Level Trigger Strategy

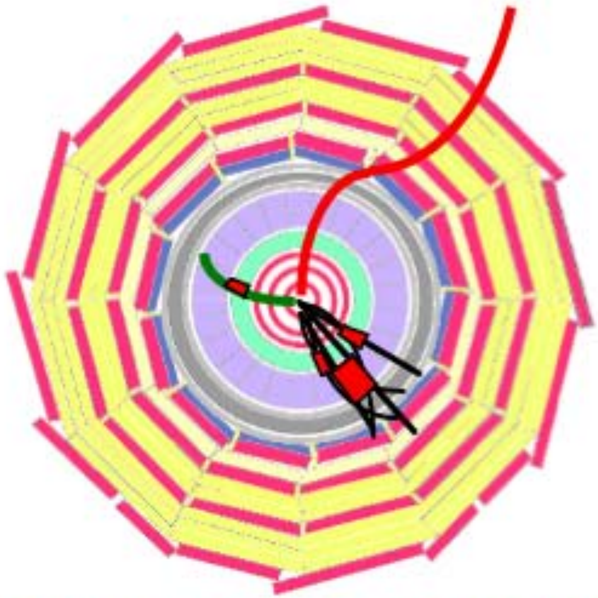

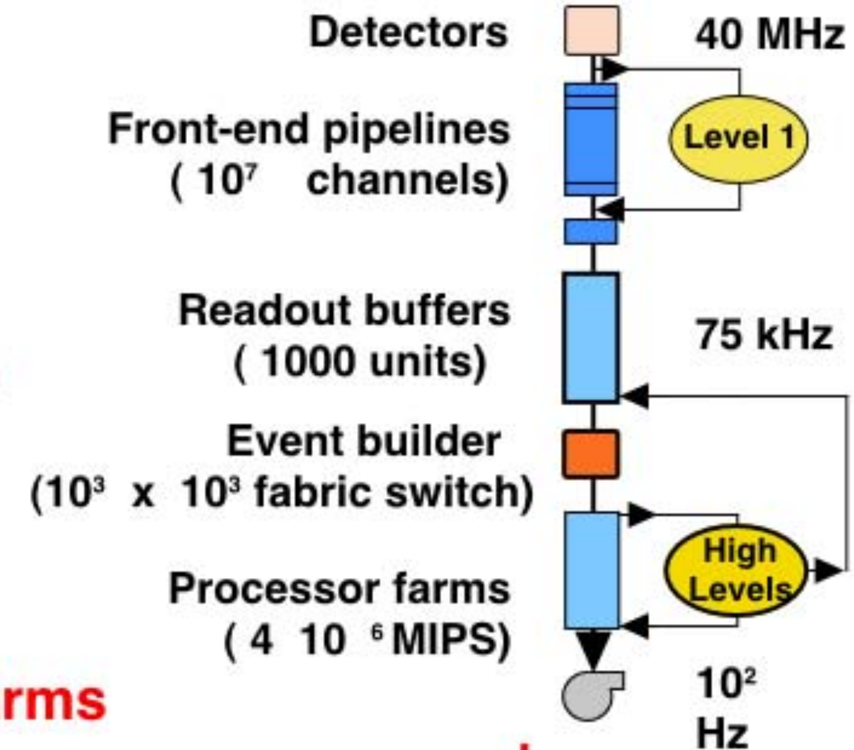

High level triggers. CPU farms

- Finer granularity precise measurement
- Clean particle signature (π^0-γ, isolation, ...)
- Kinematics. Effective mass cuts and topology
- Track reco and matching, b,τ-jet tagging
- Full event reconstruction and analysis

Successive improvements : background event filtering, physics selection

CMS DAQ & HLT

DAQ Scaling & Staging

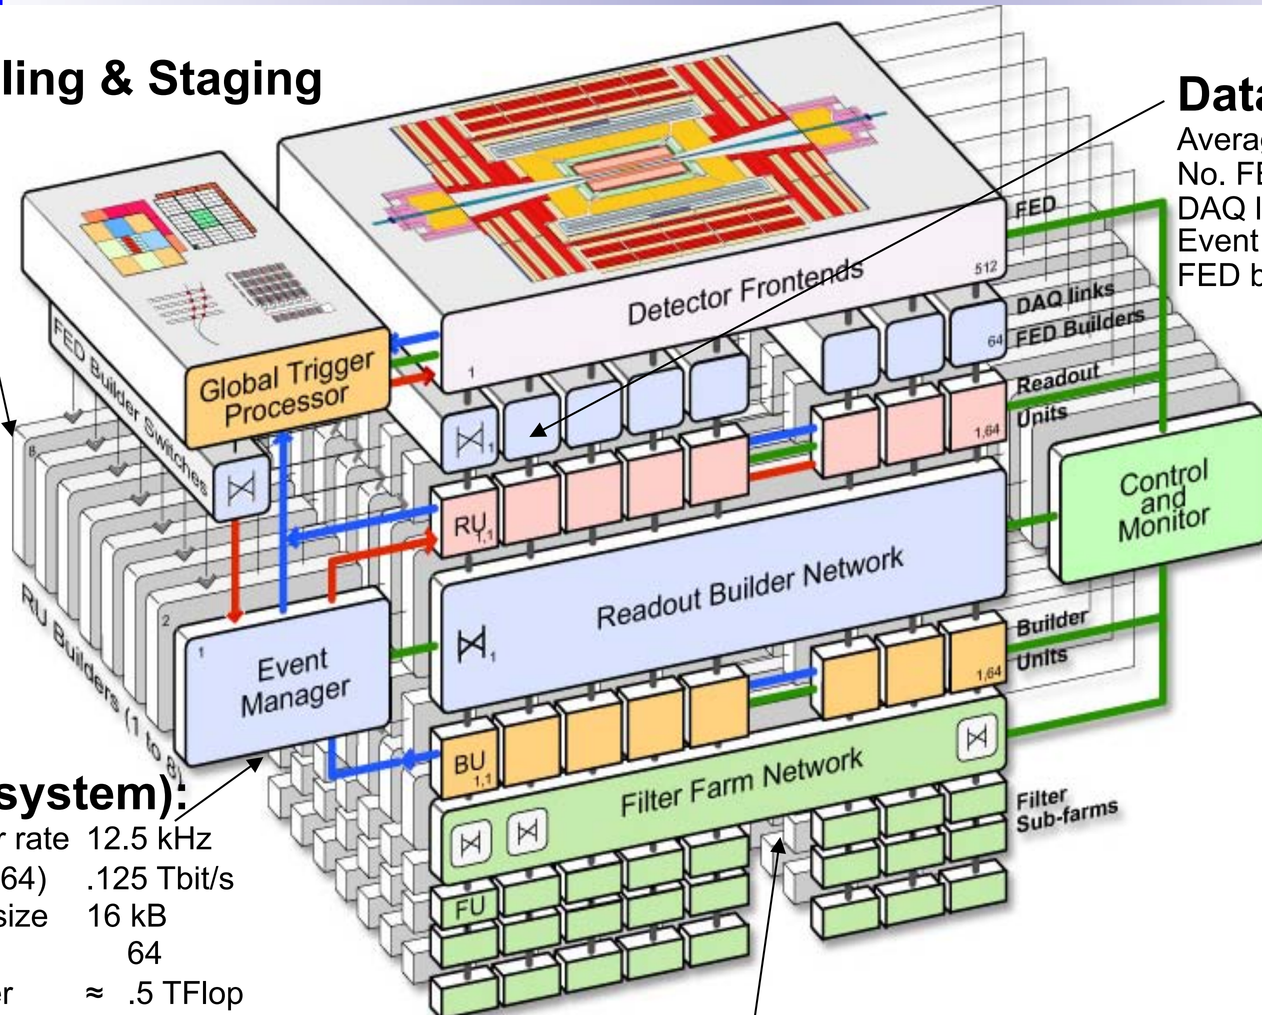

Data to surface:

Average event size		1 Mbyte
No. FED s-link64 ports	>	512
DAQ links (2.5 Gb/s)		512+51
Event fragment size		2 kB
FED builders (8x8)	≈	64+64

DAQ unit (1/8th full system):

Lv-1 max. trigger rate		12.5 kHz
RU Builder (64x64)		.125 Tbit/s
Event fragment size		16 kB
RU/BU systems		64
Event filter power	≈	.5 TFlop

HLT: All processing beyond Level-1 performed in the Filter Farm

Partial event reconstruction "on demand" using full detector resolution

Start with L1 Trigger Objects

Electrons, Photons, τ-jets, Jets, Missing E_T, Muons

- HLT refines L1 objects (no volunteers)

Goal

- Keep L1T thresholds for electro-weak symmetry breaking physics
- However, reduce the dominant QCD background
  - From 100 kHz down to 100 Hz nominally

QCD background reduction

- Fake reduction: e±, γ, τ
- Improved resolution and isolation: μ
- Exploit event topology: Jets
- Association with other objects: Missing E_T
- Sophisticated algorithms necessary
  - Full reconstruction of the objects
  - Due to time constraints we avoid full reconstruction of the event - L1 seeded reconstruction of the objects only
  - Full reconstruction only for the HLT passed events

Electron selection: Level-2

"Level-2" electron:

- Search for match to Level-1 trigger
  - Use 1-tower margin around 4x4-tower trigger region
- Bremsstrahlung recovery "super-clustering"
- Select highest E_T cluster

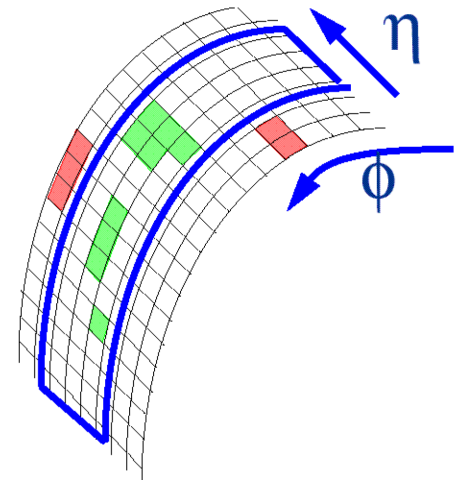

Bremsstrahlung recovery:

- Road along ϕ — in narrow η-window around seed
- Collect all sub-clusters in road → "super-cluster"

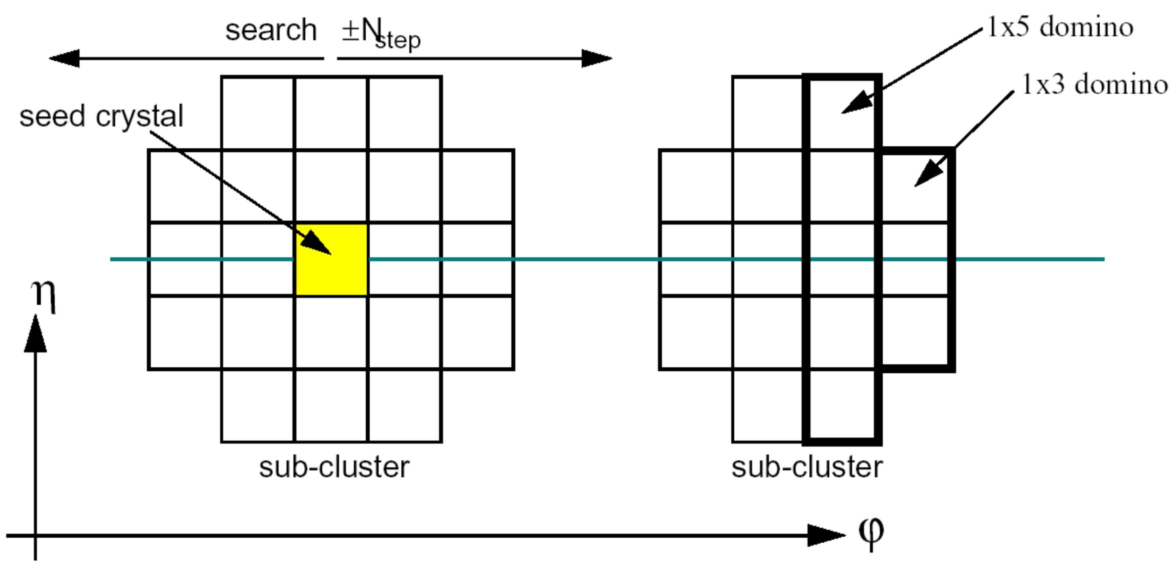

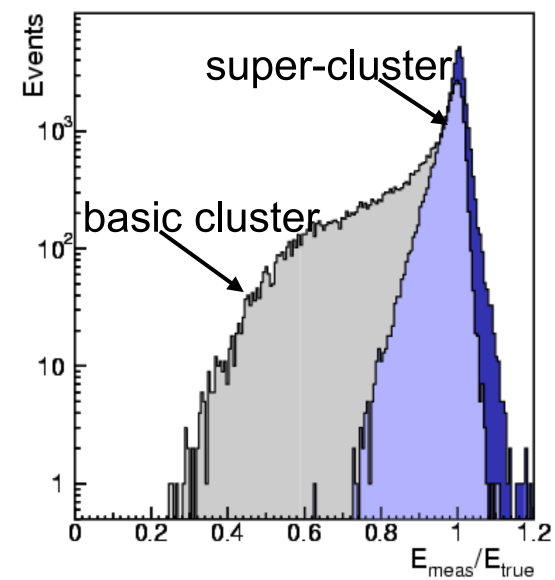

CMS tracking for electron trigger

Present CMS electron HLT

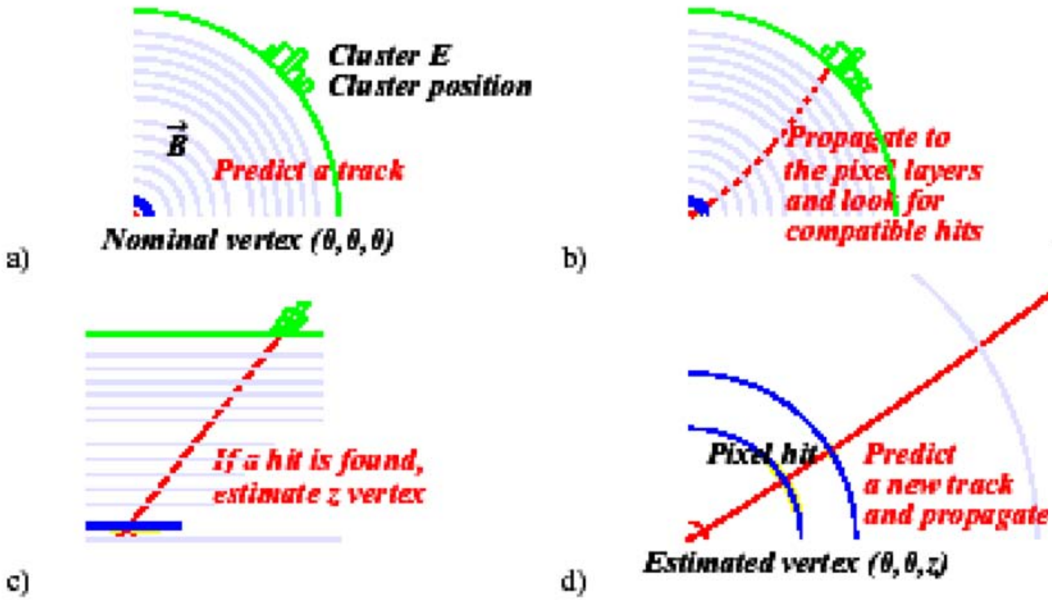

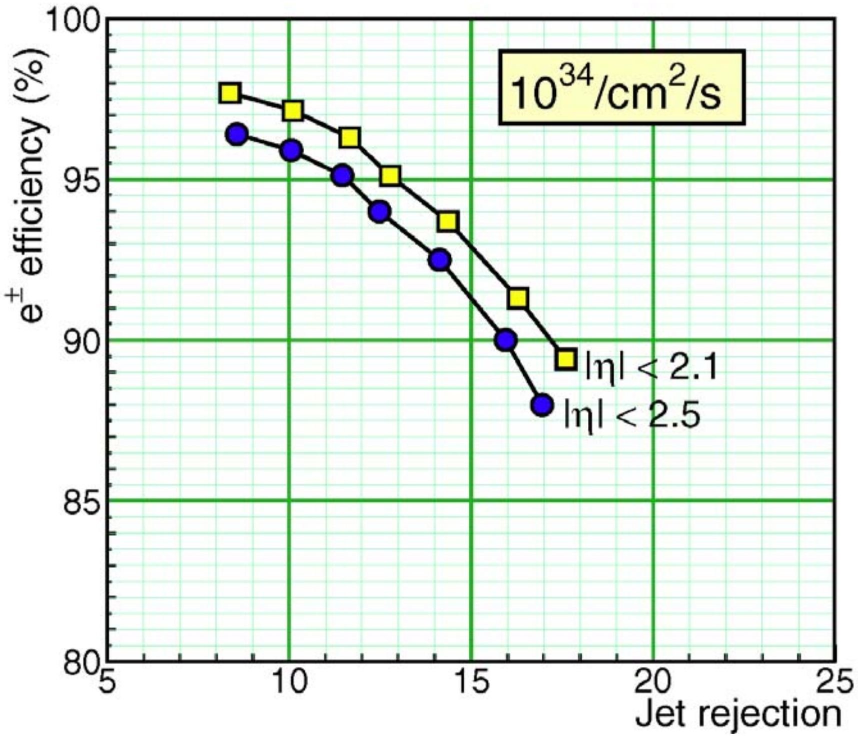

Factor of 10 rate reduction

γ: only tracker handle: isolation

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

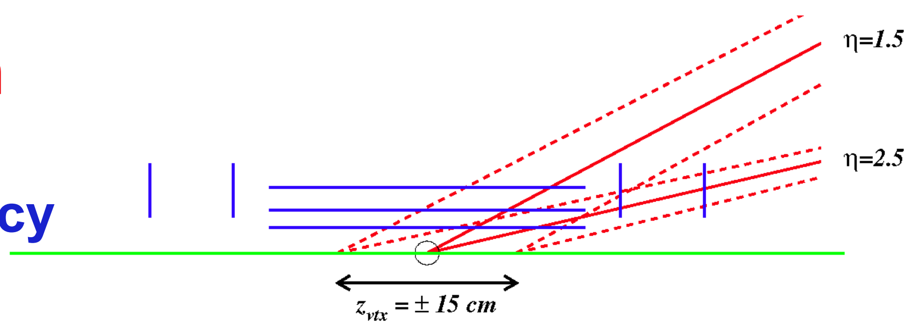

τ-jet tagging at HLT

τ-jet ($E_t^{\tau\text{-jet}} > 60$ GeV) identification (mainly) in the tracker:

- Hard track, $p_t^{max} > 40$ GeV, within $\Delta R < 0.1$ around calorimeter jet axis
- Isolation: no tracks, $p_t > 1$GeV, within $0.03 < \Delta R < 0.4$ around the hard track
- For 3-prong selection 2 more tracks in the signal cone $\Delta r < 0.03$

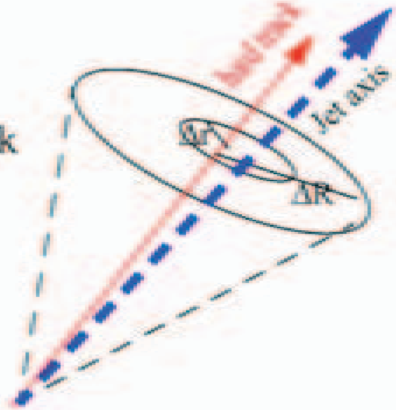

QCD jet rejection from isolation and hard track cuts

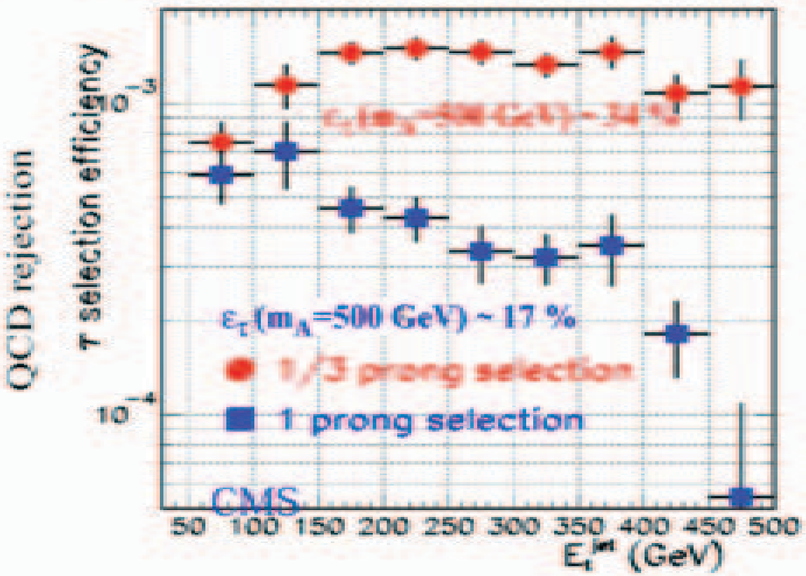

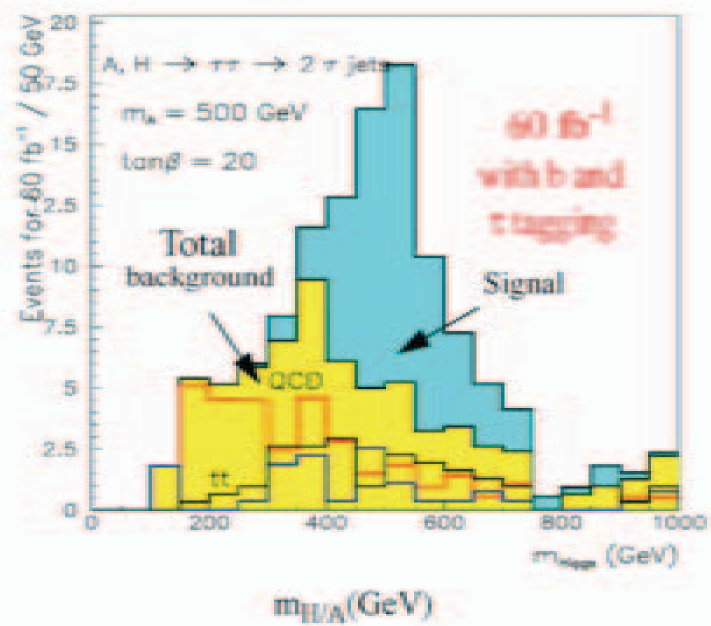

Further reduction by ~ 5 expected for 3-prong QCD jets from τ vertex reconstruction (CMS full simulation)

B and τ tagging

Soft b-jets with a wide η-range:

Efficiency to tag one b-jet ~ 35% for ~1% mistagging rate (CMS)

τ - tagging with impact parameter measurement

combining the ip measurements of the hard tracks in the two τ's (τ -> hadron, τ -> lepton) into one variable: $\sqrt{\sigma_{ip}(\tau_1)^2 + \sigma_{ip}(\tau_2)^2}$

CMS full simulation for

H -> ττ -> 2 τ–jets and QCD events

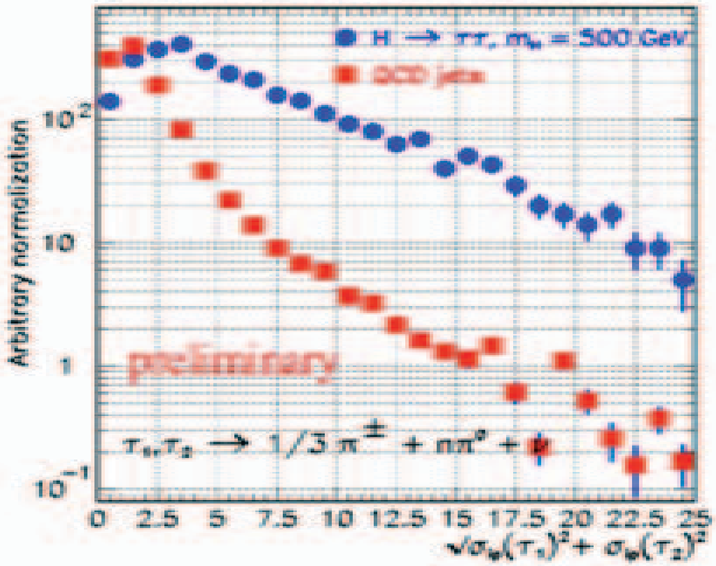

Expect rejection of 5 - 10 against QCD background and backgrounds with W -> lν, Z -> ll

Signal superimposed on the total background for m_A = 200 GeV, tanβ =20

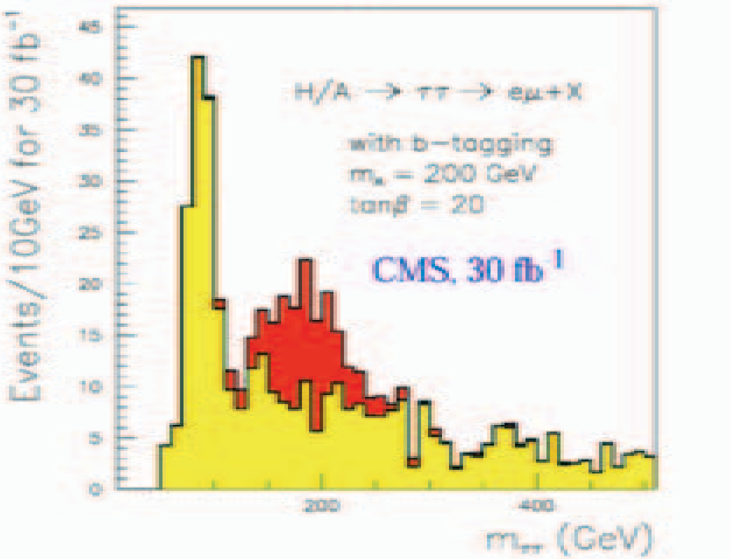

CMS HLT Time Distribution

Prescale set used: 2E32 Hz/cm²
Sample: MinBias L1-skim 5E32 Hz/cm² with 10 Pile-up

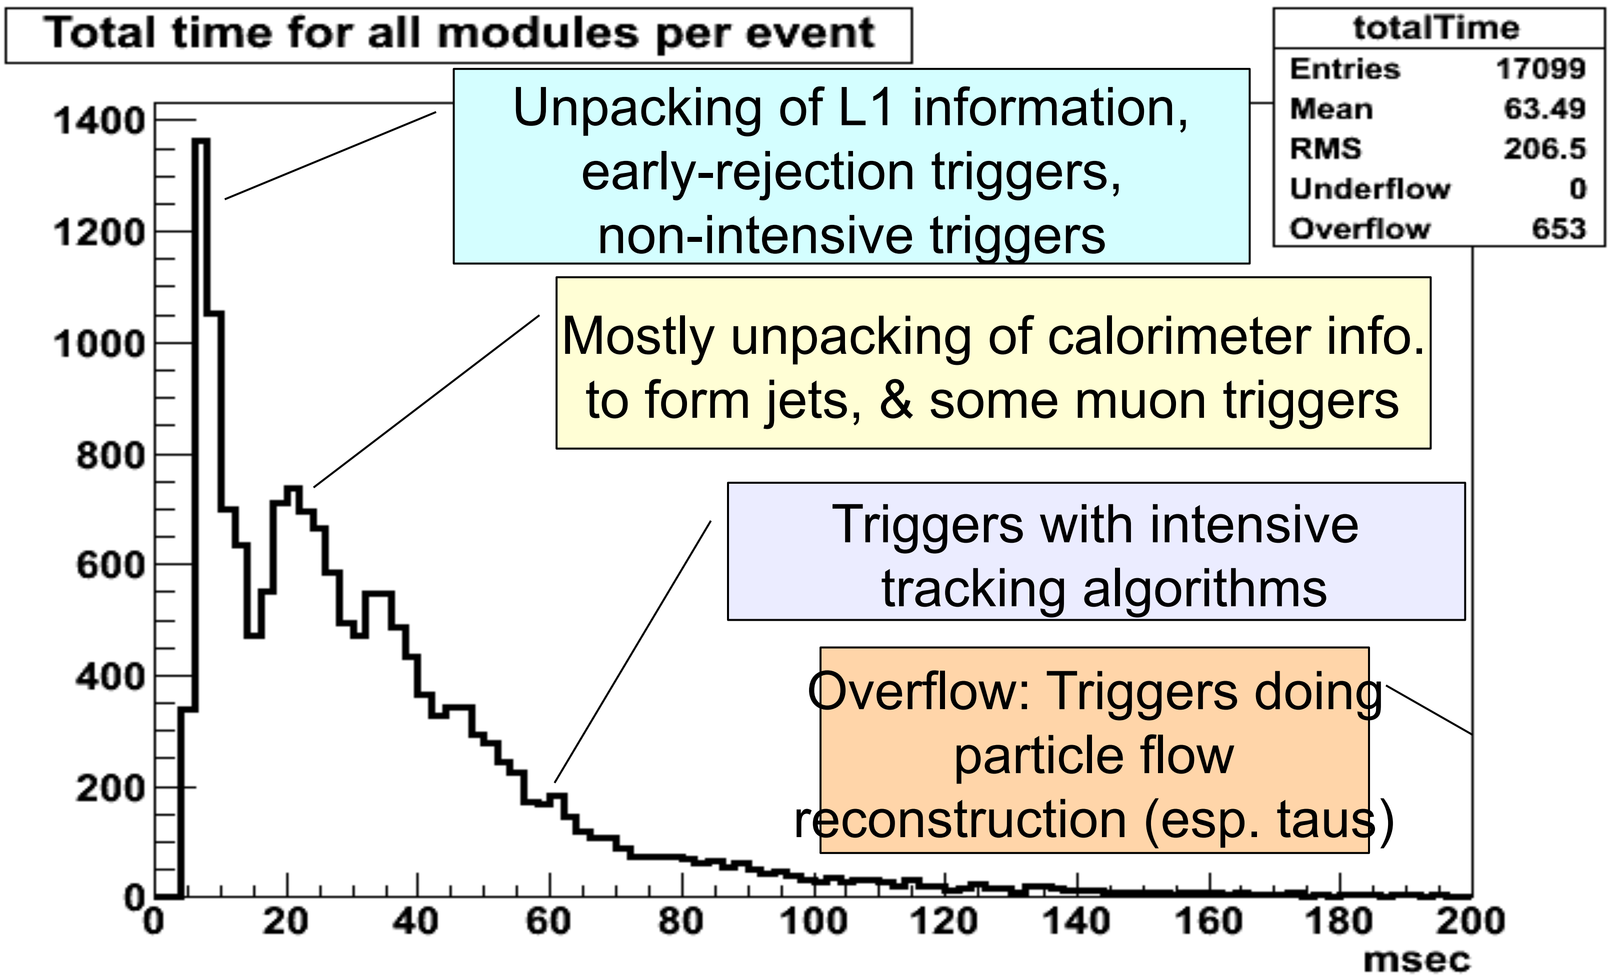

CMS DAQ

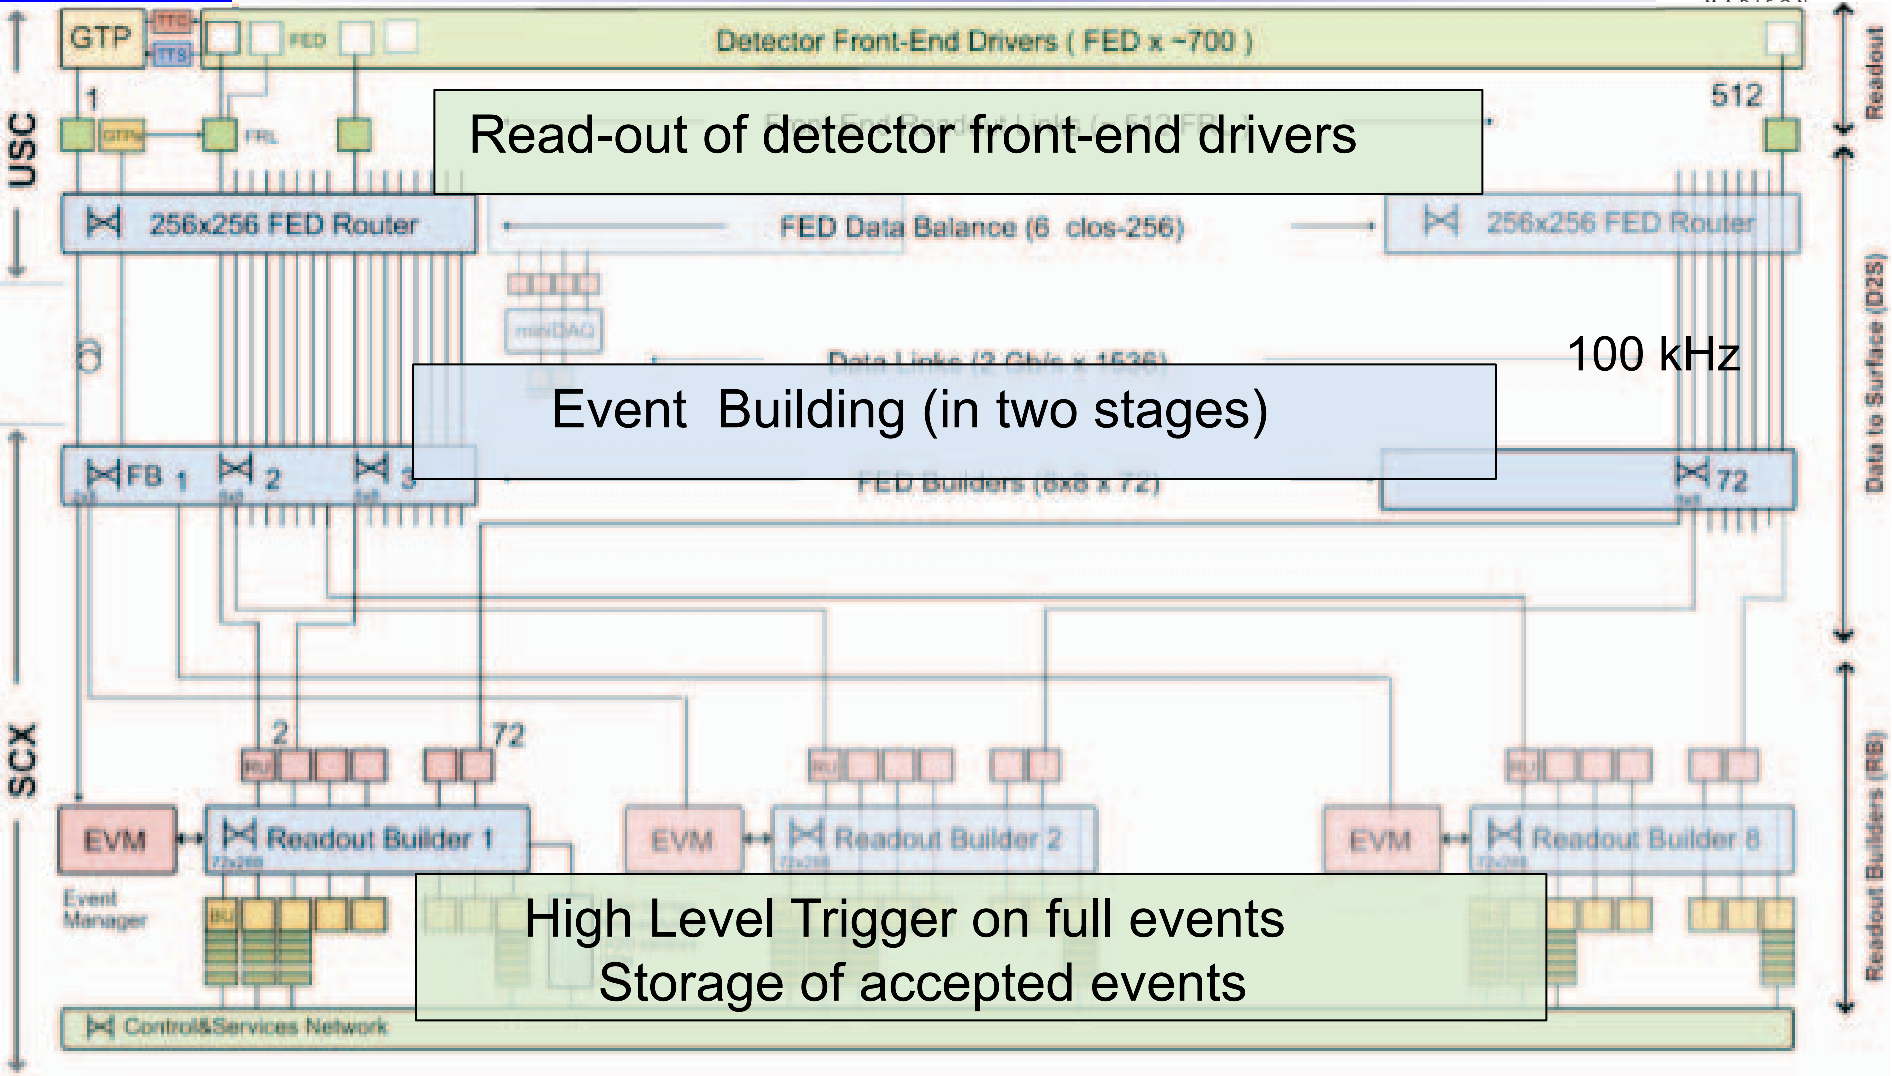

12.5 kHz 12.5 kHz … 12.5 kHz

2-Stage Event Builder

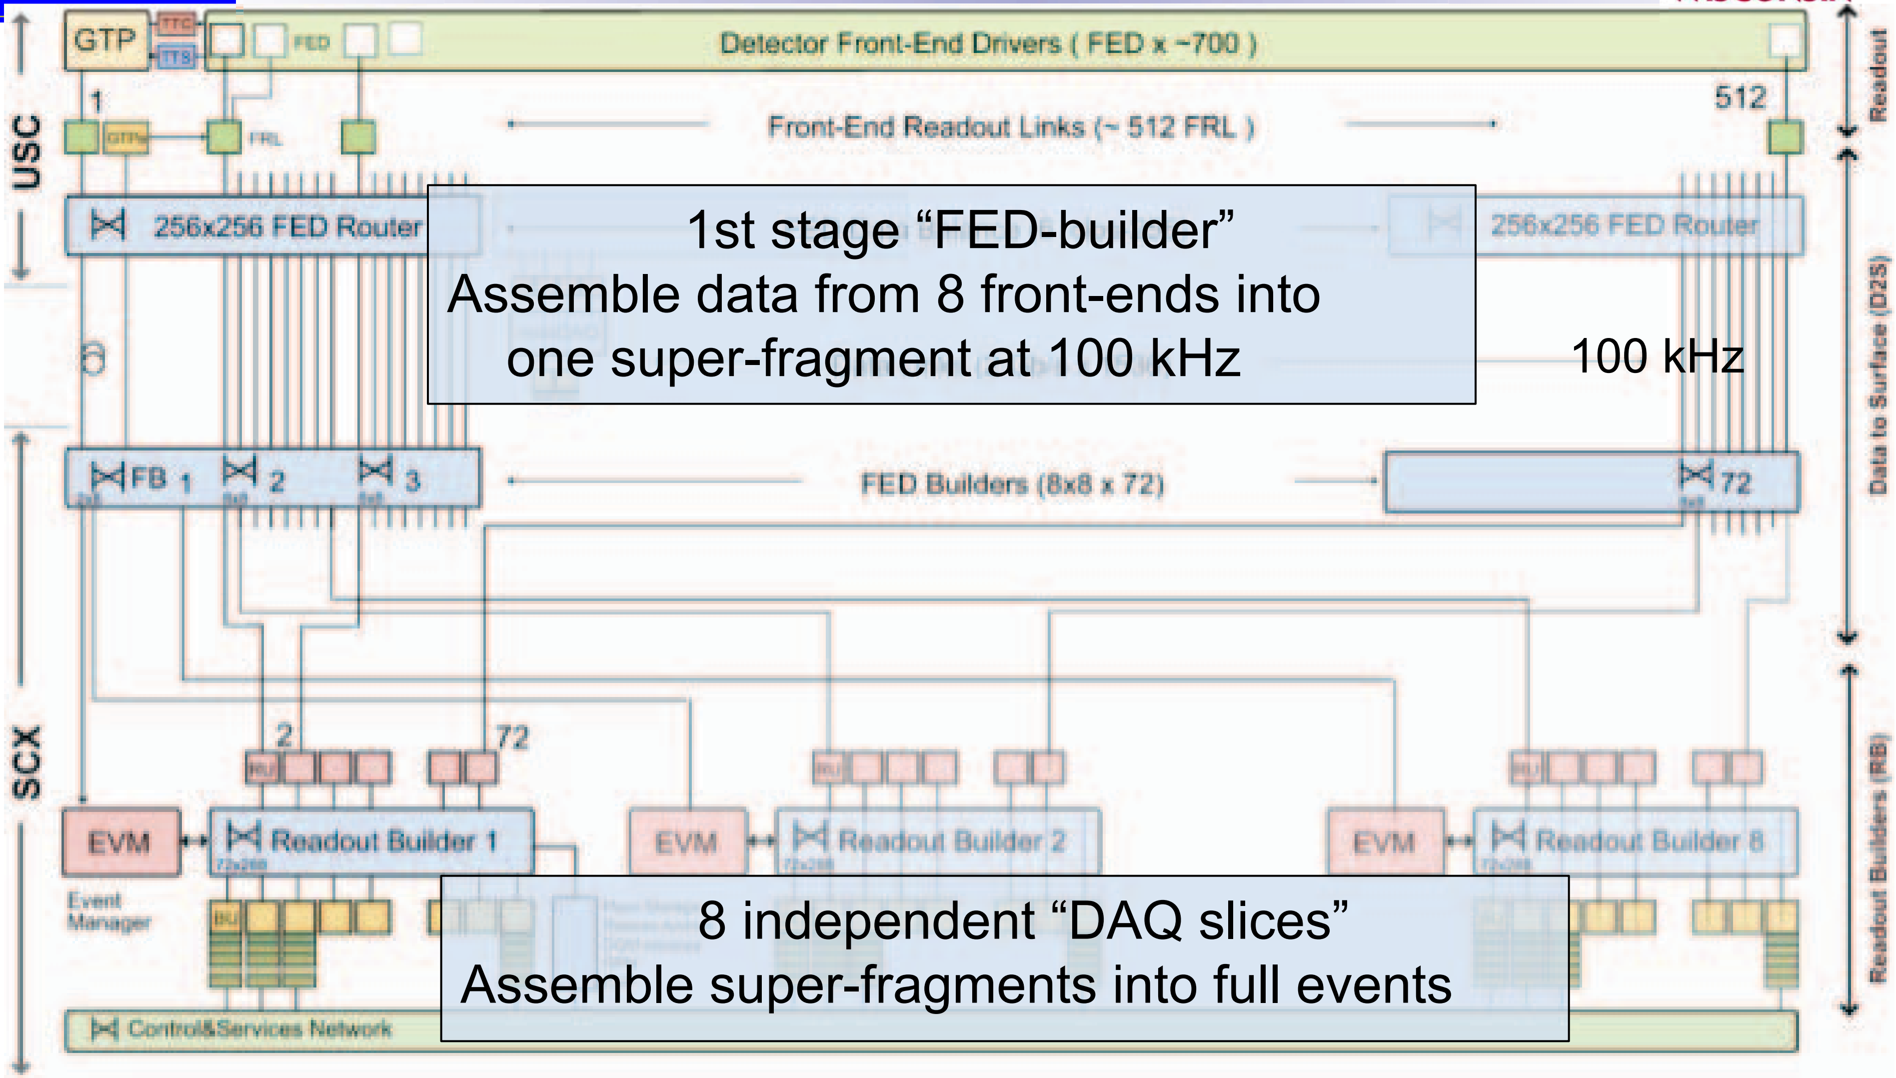

1st stage "FED-builder"
Assemble data from 8 front-ends into one super-fragment at 100 kHz

100 kHz

8 independent "DAQ slices"
Assemble super-fragments into full events

12.5 kHz 12.5 kHz … 12.5 kHz

Building the event

Event builder :
Physical system interconnecting data sources with data destinations. It has to move each event data fragments into a same destination

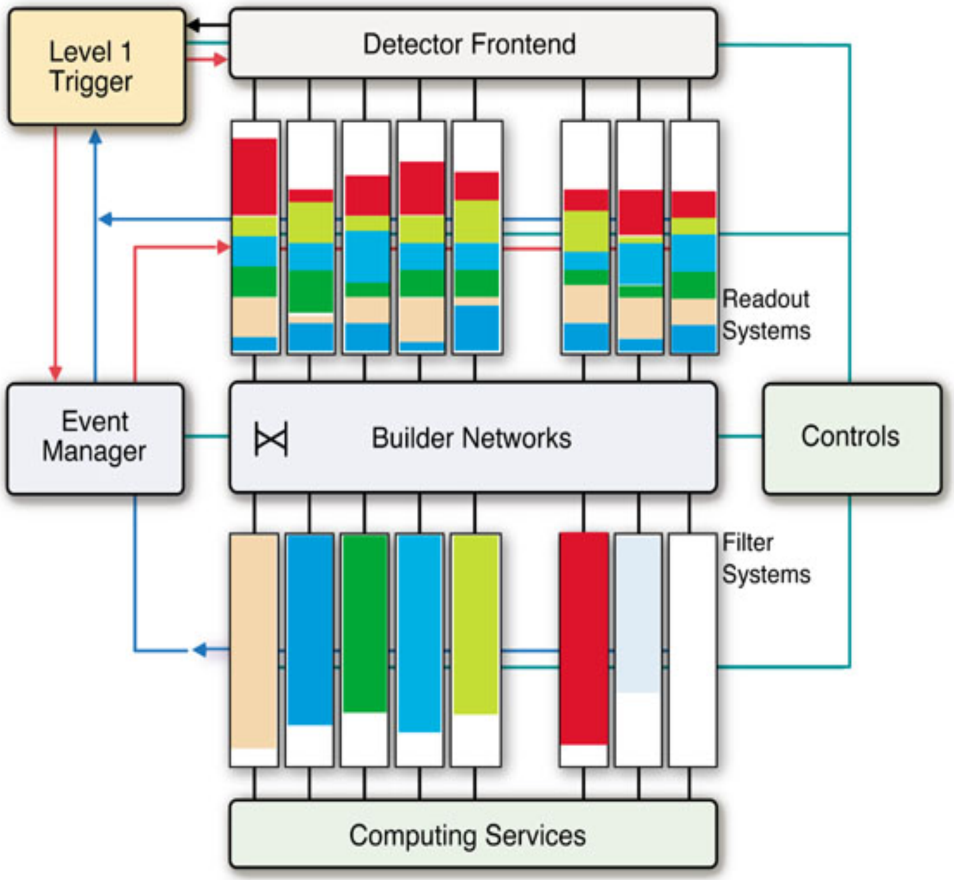

Event fragments :
Event data fragments are stored in separated physical memory systems

Full events :
Full event data are stored into one physical memory system associated to a processing unit

Hardware:
Fabric of switches for builder networks
PC motherboards for data Source/Destination nodes

Myrinet Barrel-Shifter

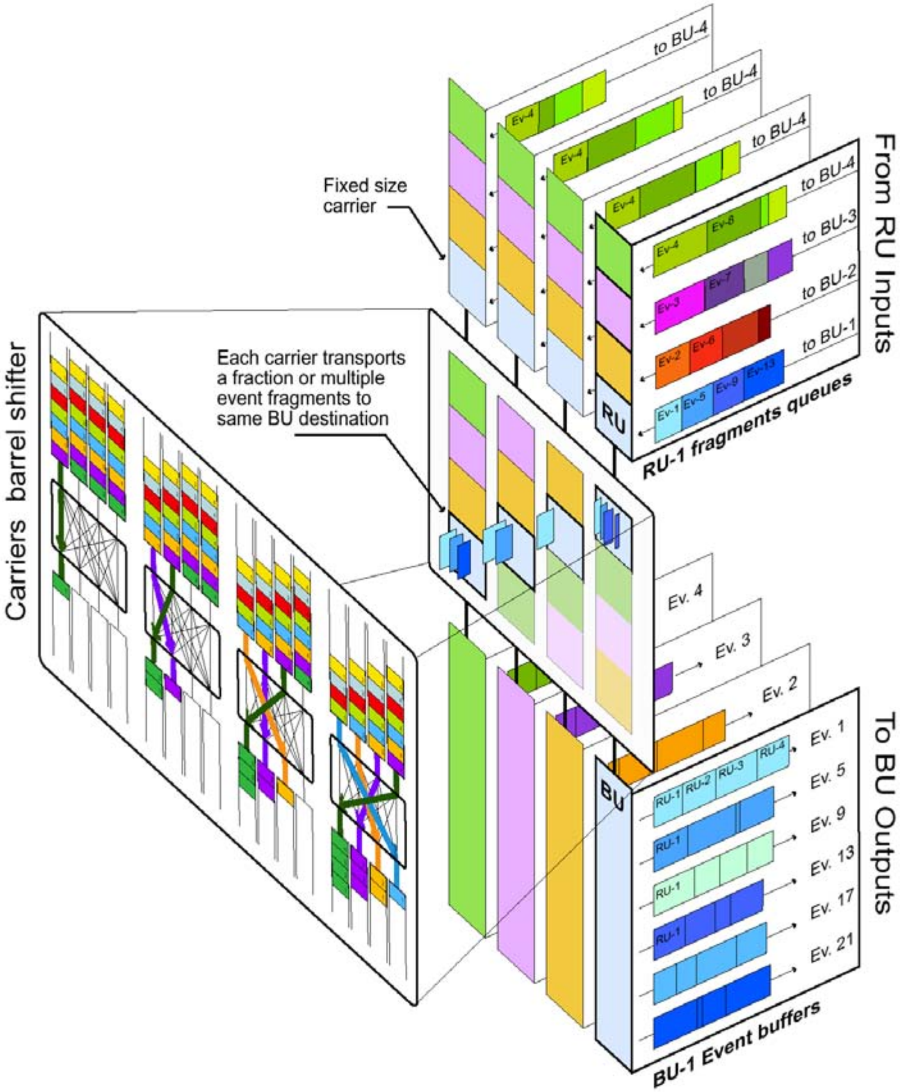

BS **implemented in firmware**

- Each source has message queue per destination
- Sources divide messages into fixed size packets (carriers) and cycle through all destinations
- Messages can span more than one packet and a packet can contain data of more than one message
- No external synchronization (relies on Myrinet back pressure by HW flow control)

zero-copy, **OS-bypass**

principle works for multi-stage switches

EVB – HLT installation

- EVB – input “RU” PC nodes
  - 640 times dual 2-core E5130 (2007)
  - Each node has 3 links to GbE switch
- Switches
  - 8 times F10 E1200 routers
  - In total ~4000 ports
- EVB – output + HLT node (“BU-FU”)
  - 720 times dual 4-core E5430, 16 GB (2008)
  - 288 times dual 6-core X5650, 24 GB (2011)
  - Each node has 2 links to GbE switch

HLT Total: 1008 nodes, 9216 cores, 18 TB memory
@100 kHz: ~90 ms/event
Can be easily expanded by adding PC nodes and recabling EVB network

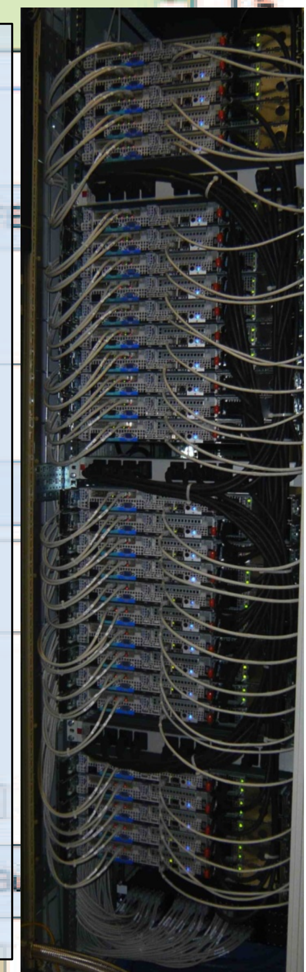

C T E Q

Requirements for LHC phases of the upgrades: ~2010-2020

Phase 1:

- Goal of extended running in second half of the decade to collect ~100s/fb
- 80% of this luminosity in the 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

- Continued operation of the LHC beyond a few 100/fb will require substantial modification of detector elements
- The 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}

Expected Pile-up at High Lumi LHC in ATLAS at 10^{35}

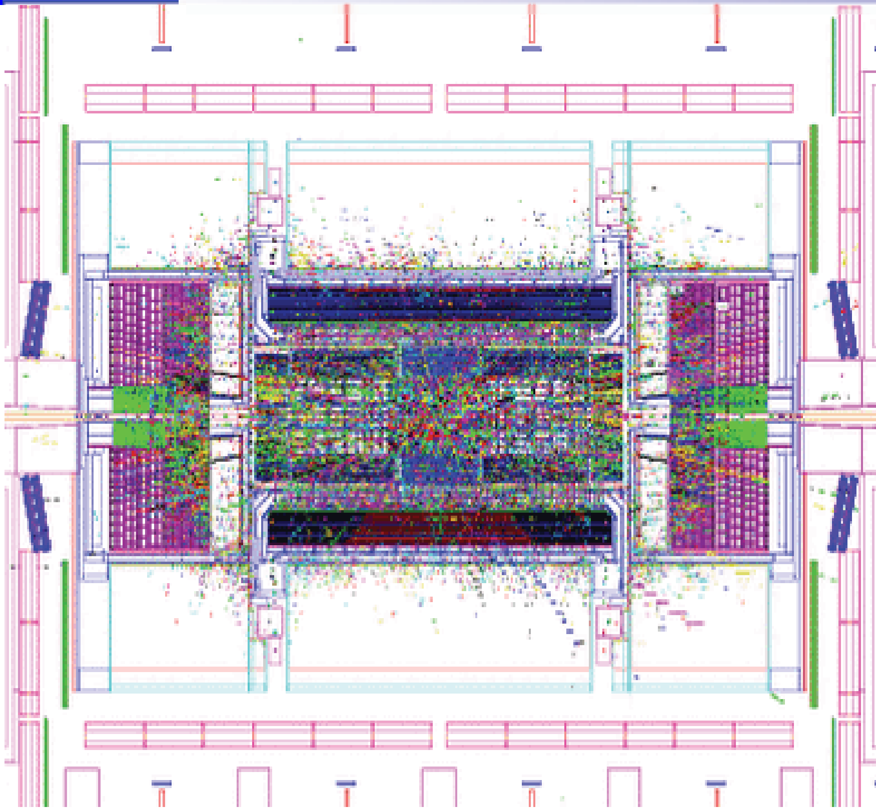

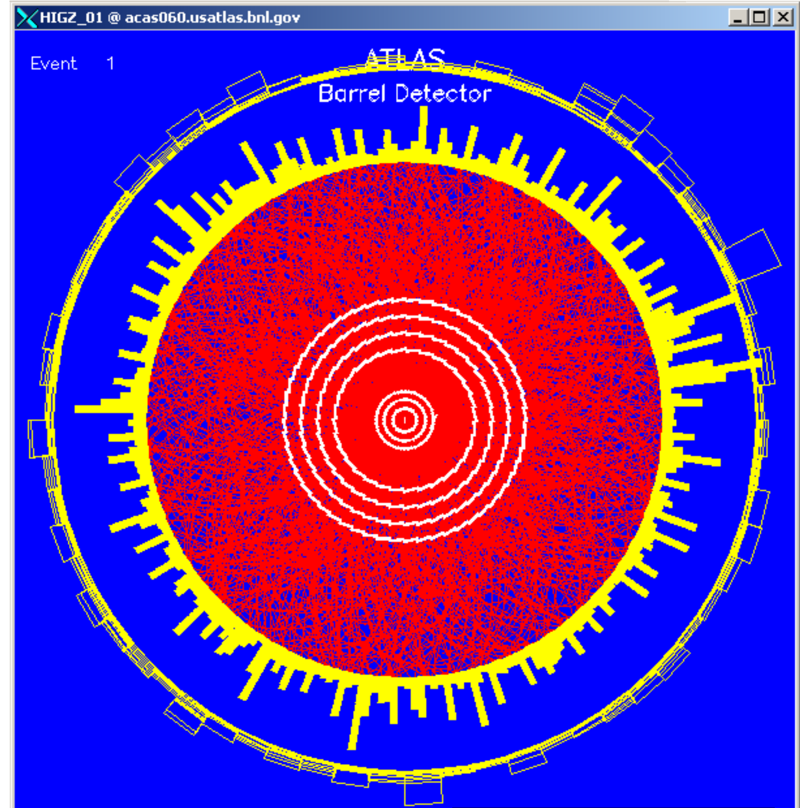

$N_{ch}(|y| \leq 0.5)$

- 230 min.bias collisions per 25 ns. crossing
- ~ 10000 particles in $|\eta| \leq 3.2$
- mostly low p_T tracks
- requires upgrades to detectors

Detector Luminosity Effects

H→ZZ → μμee, M_H= 300 GeV for different luminosities in CMS

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

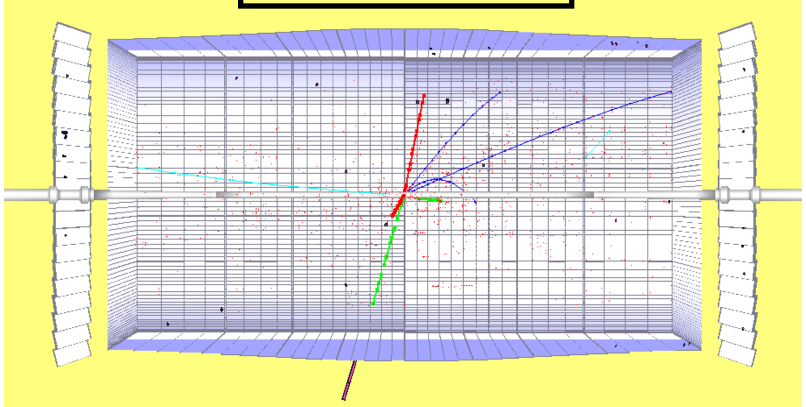

10^{33} cm^{-2}s^{-1}

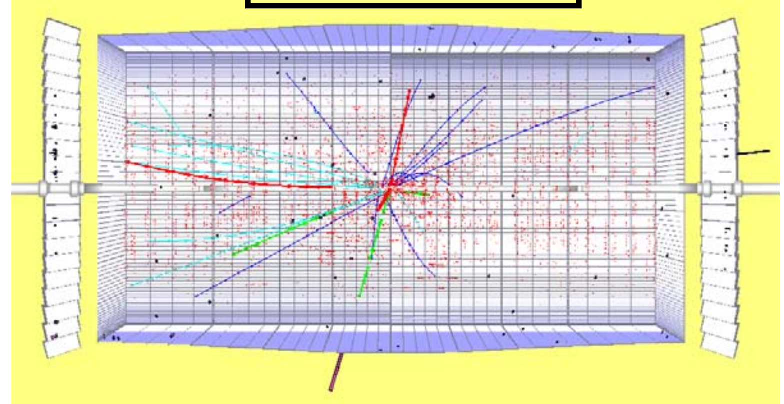

10^{34} cm^{-2}s^{-1}

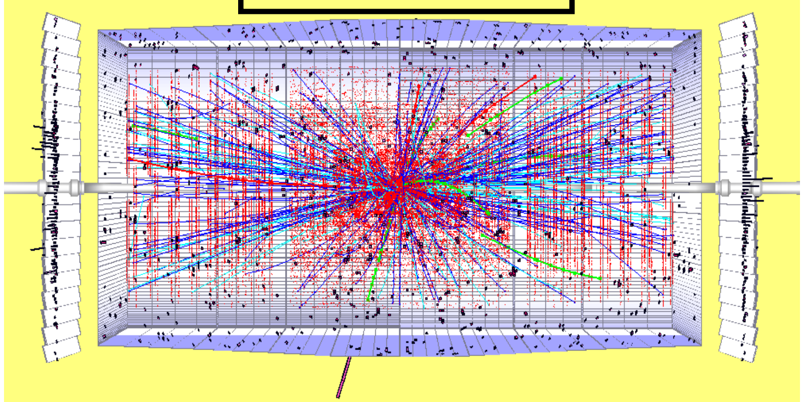

10^{35} cm^{-2}s^{-1}

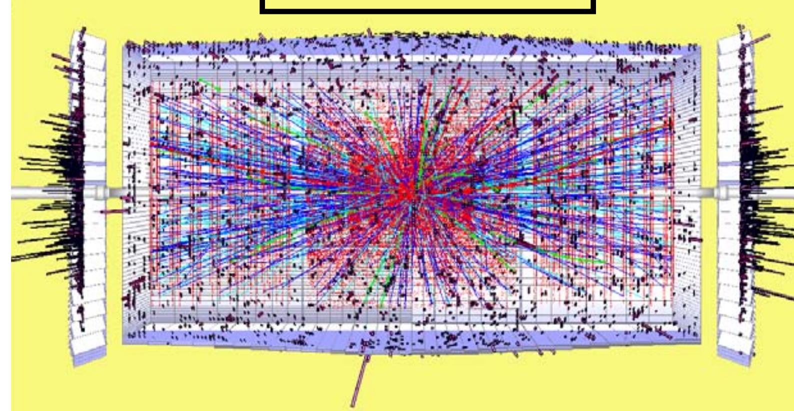

CMS Upgrade Trigger Strategy

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 (operating 2016+):

- Present L1 algorithms inadequate above 10^{34} or 10^{34} w/ 50 ns spacing
  - 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: 2020+)

- 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.)

Tracking needed for L1 trigger

Single electron trigger rate

Isolation criteria are insufficient to reduce rate at L = 10^{35} *cm*$^{-2}$*s*$^{-1}$

(Or $5x10^{34}$ at 50 ns)

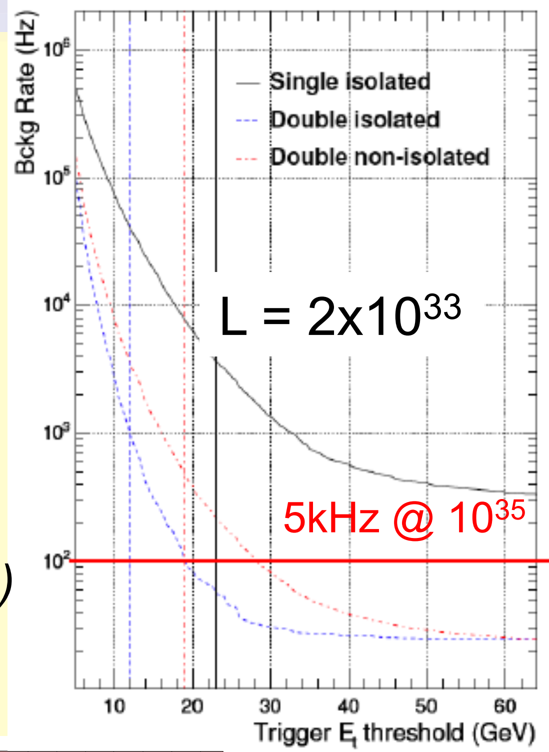

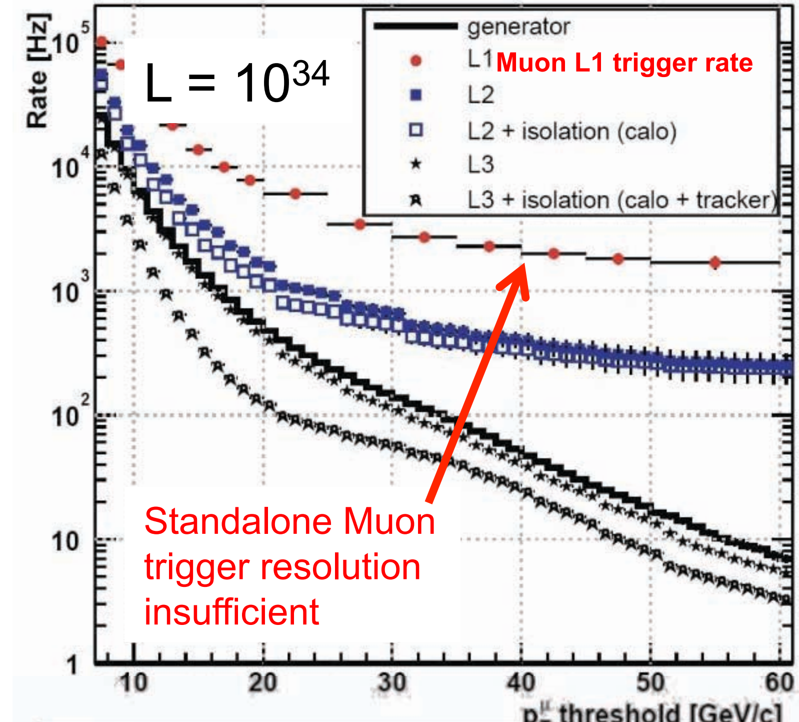

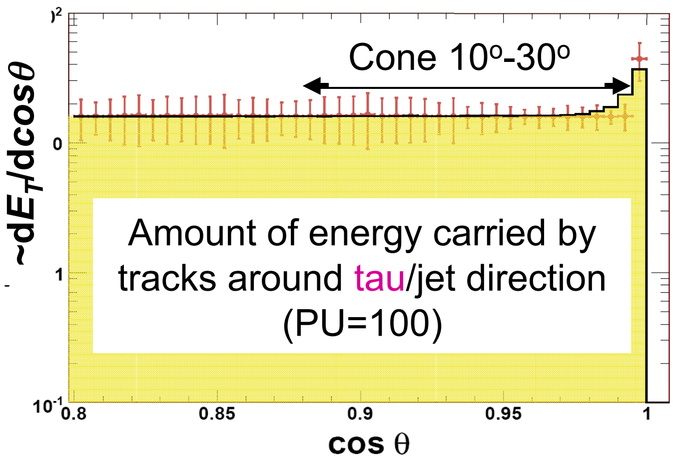

We need to get another x200 (x20) reduction for single (double) tau rate!

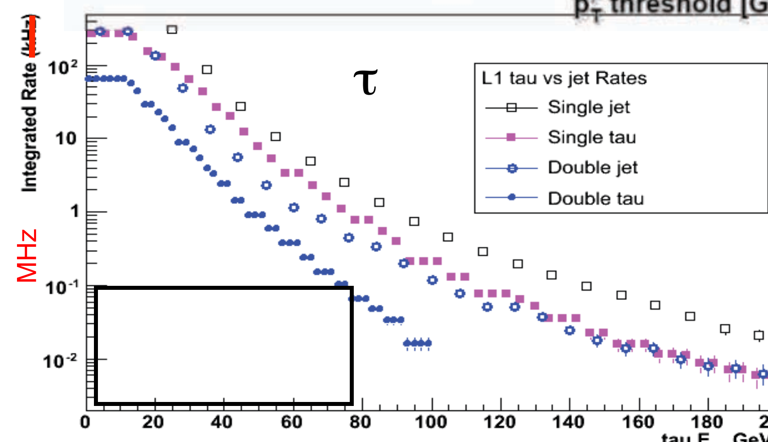

The Track Trigger Problem

- Need to gather information from 10^8 pixels in $200 m^2$ 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
- Studying the use of 3D electronics to provide ability to locally correlate hits between two closely spaced layers

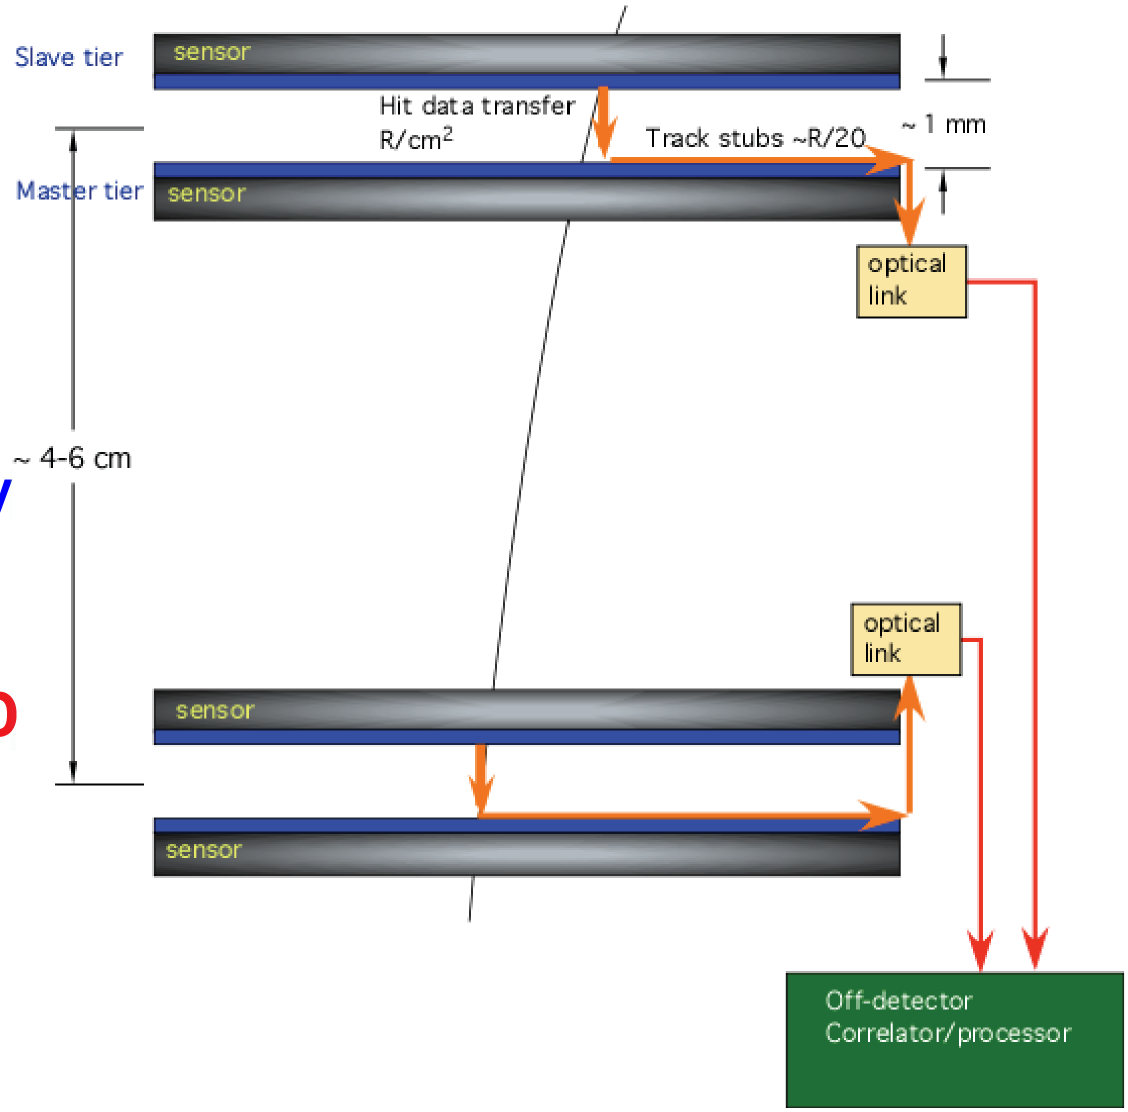

3D Interconnection

Key to design is ability of a single IC to connect to both top & bottom sensor

Enabled by “vertical interconnected” (3D) technology

A single chip on bottom tier can connect to both top and bottom sensors – locally correlate information

Analog information from top sensor is passed to ROIC (readout IC) through interposer

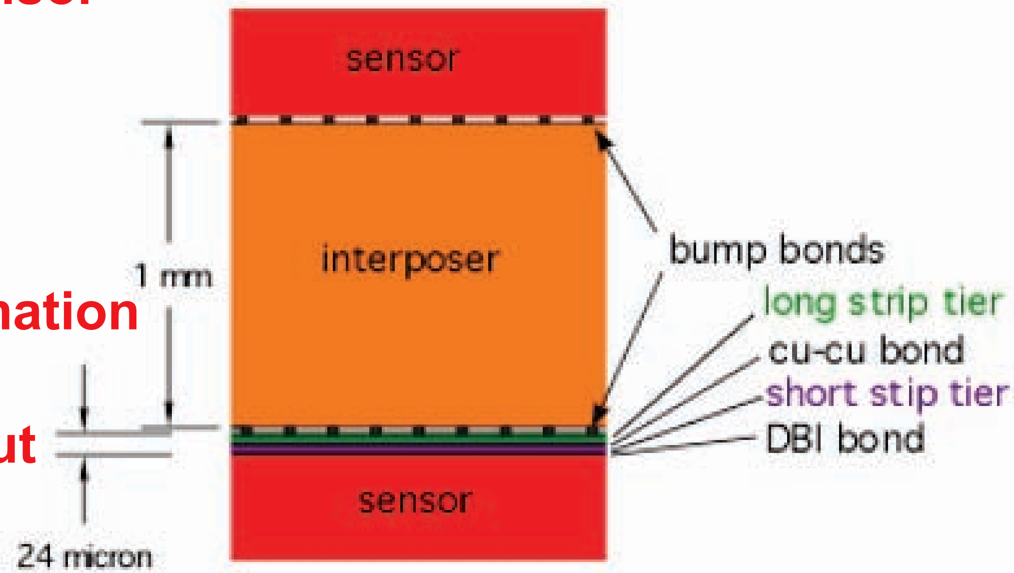

One layer of chips

No “horizontal” data transfer necessary – lower noise and power

Fine Z information is not necessary on top sensor – long (~1 cm vs ~1-2 mm) strips can be used to minimize via density in interposer

Track Trigger Architecture

CTEQ

THE UNIVERSITY of WISCONSIN MADISON

Readout designed to send all hits with P_t>~2 GeV to trigger processor

High throughput – micropipeline architecture

Readout mixes trigger and event data

Tracker organized into phi segments

- Limited FPGA interconnections
- Robust against loss of single layer hits
- Boundaries depend on p_t cuts & tracker geometry

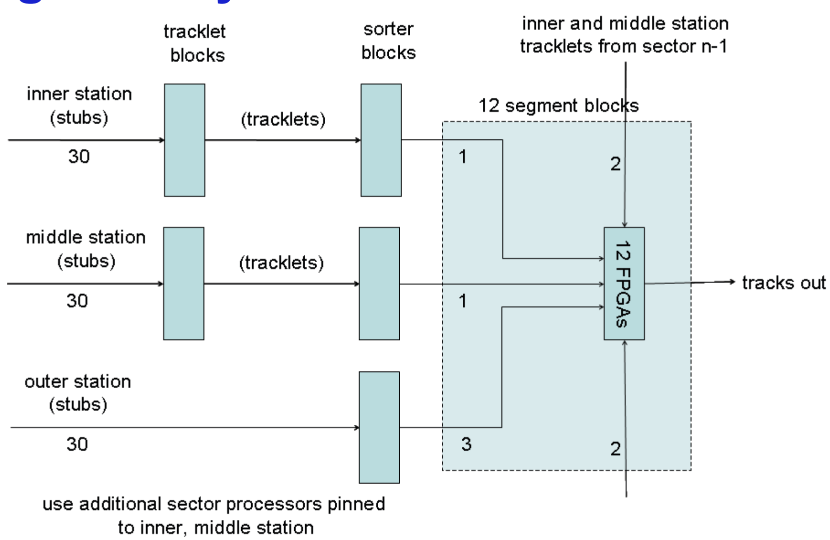

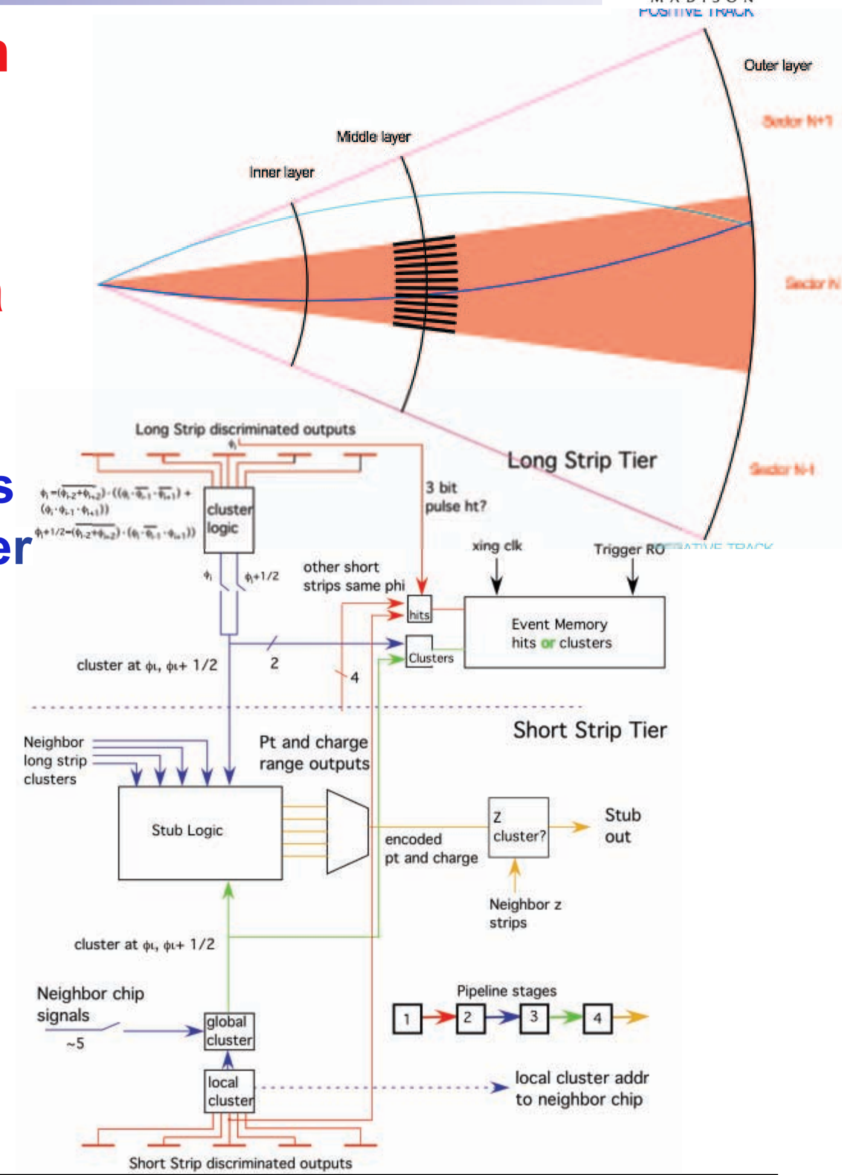

ATLAS Trigger Upgrades

Various projects being pursued:

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

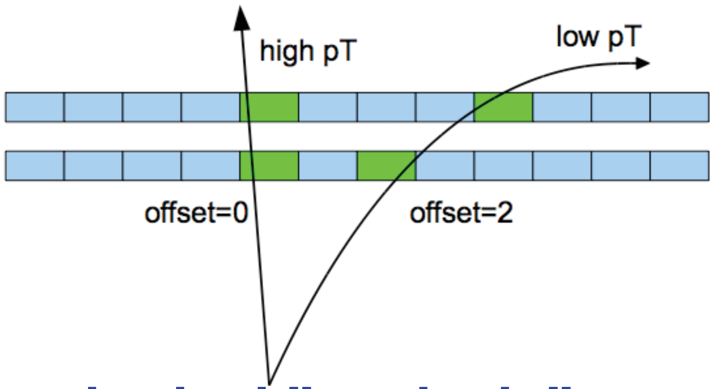

- **Combining trigger objects at L1 & topological "analysis"**
- **Full granularity readout of calorimeter**
  - requires new electronics
- **Changes in muon systems (small wheels), studies of an MDT based trigger & changes in electronics**
- **Upgrades of HLT farms**

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

Upgrade CMS DAQ

Phase 2 Network bandwidth at least 5-10 times LHC

- Assuming L1 trigger rate same as LHC
- Increased Occupancy
- Decreased channel granularity (esp. tracker)

CMS DAQ Component upgrades

- Readout Links: replace existing SLINK (400 MB/s) with 10 Gbit/s
- Present Front End Detector Builder & Readout Unit Builder replaced with updated network technology & mult-gigabit link network switch
- Higher Level Trigger CPU Filter Farm estimates:
  - 2010 Farm = 720 Dual Quad Core E5430 16 GB (2.66 GHz)
  - 2011 Farm = add 288 Dual 6-Core X5650 24 GB (2.66 GHz)
    - 1008 nodes, 9216 cores, 18 TB memory @100 kHz: ~90 ms/event
  - 2012 Farm = 3$\times$ present farm
  - 2016 Farm = 3$\times$ 2012 farm
    - Requires upgrades to network (40 Gbps links now affordable)

C T E Q

Extrapolating PC performance

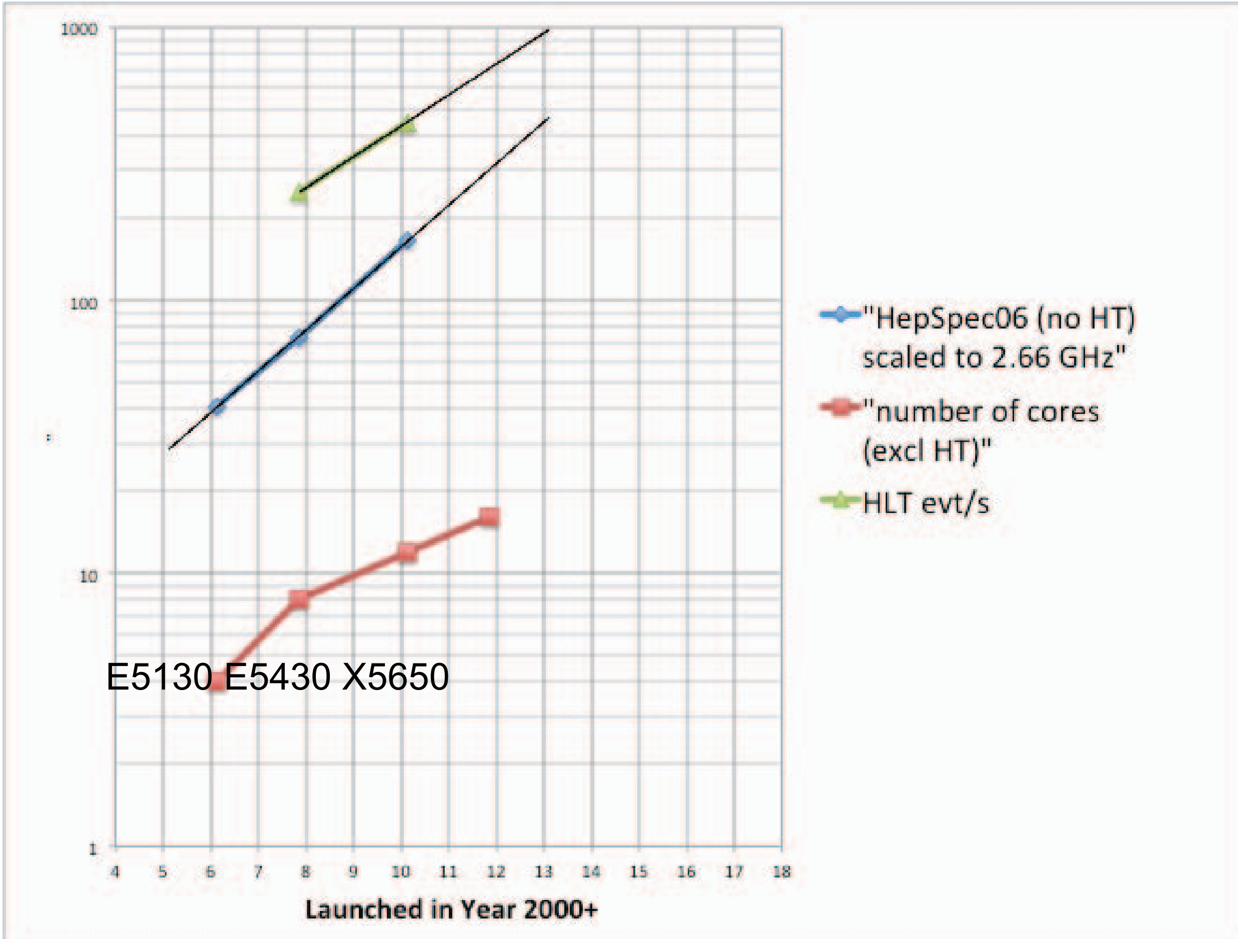

Extrapolate performance dual-processor PCs

In 2014 could have same HLT performance with 100 – 200 nodes

Likely to have 10 GbE onboard

HEP tools for high rate experiments: μTCA

- Advanced Telecommunications Computing Architecture ATCA
- μTCA Derived from AMC std.
  - Advanced Mezzanine Card
  - Up to 12 AMC slots
    - *Processing modules*
  - 1 or 2 MCH slots
    - *Controller Modules*
- 6 standard 10Gb/s point-to -point links from each 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

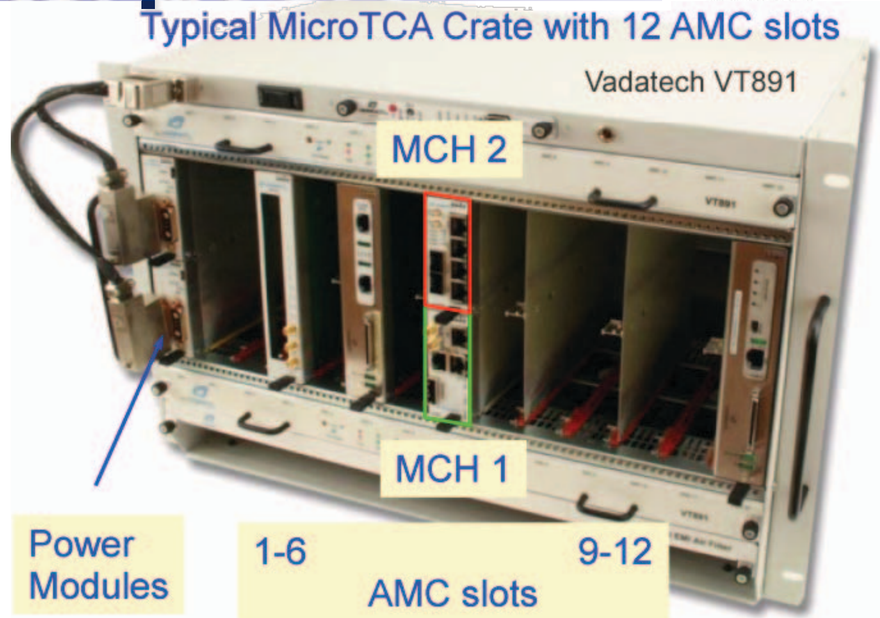

Single Module (shown): 75 x 180 mm
Double Module: 150 x 180mm

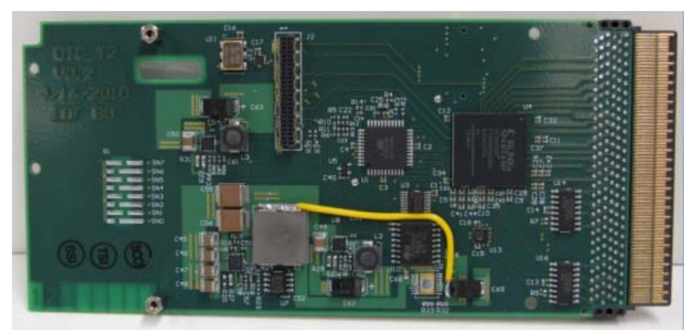

FPGAs: Logic Cells

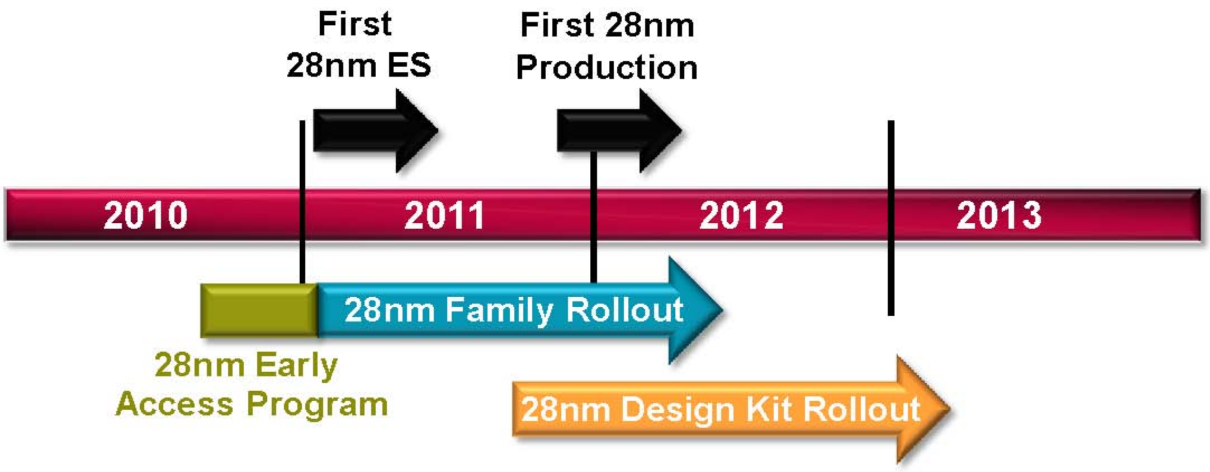

Next generation 28 nm:

> 2x capacity gains over 40nm devices

Family	Capacity Range
Artix-7	20K – 355K LCs
Kintex-7	30K – 410K LCs
Virtex-7	285K – 2,000K LCs

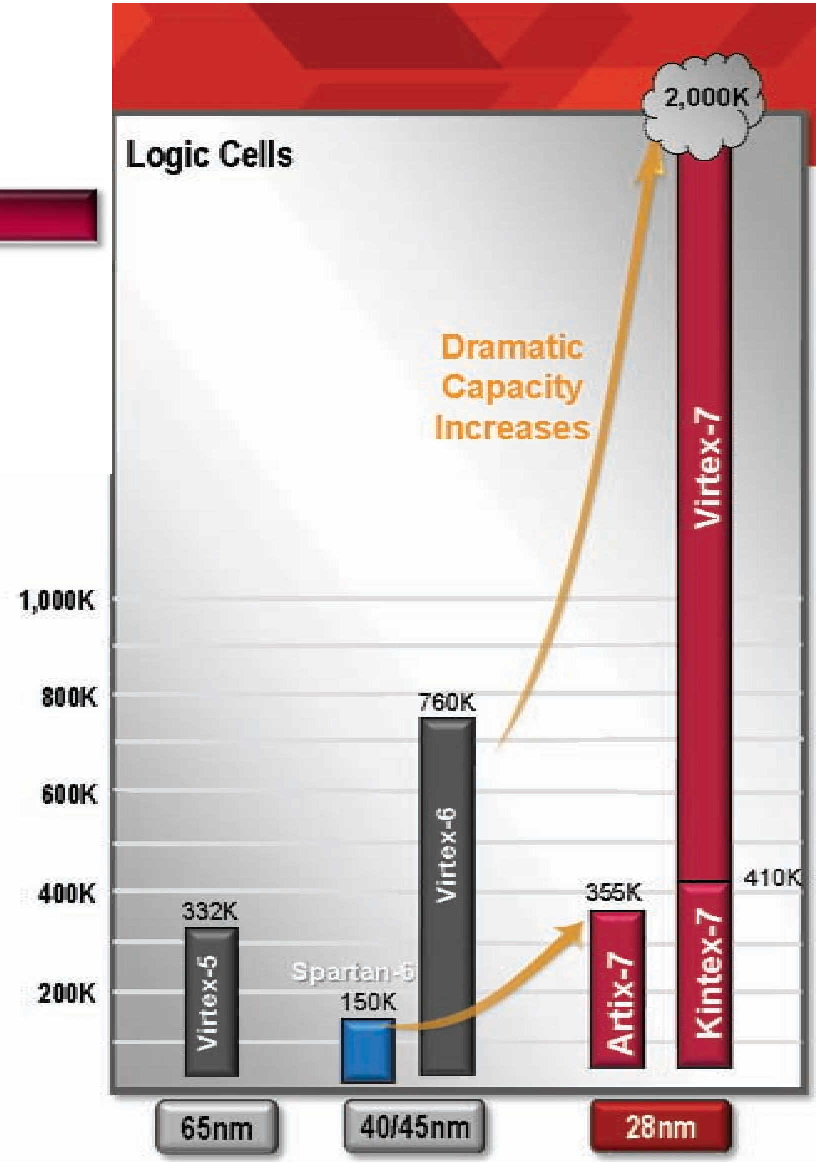

XILINX

FPGAs: Transceivers

XILINX

- **Challenge:**
  - Increase device BW
  - No increase in total device power
  - XCVR gains from scaling: negligible
- **Solution:**
  - Careful circuit design throughout XCVR
  - Increased Gbps / XCVR
  - More XCVR / Device
  - Low power mode for short channels
  - Lanes share a PLL vs PLL per lane
- **Result:**
  - 60% Increased max device BW
  - Device XCVR power unchanged

	GTP	GTX	GTH	GT28
Max Rate (Gbps)	3.75	10.3125	13.1	28
Relative Power (Per GT)	.35x	.7x	1x	-
Max GTs per Device	4	56	72	-

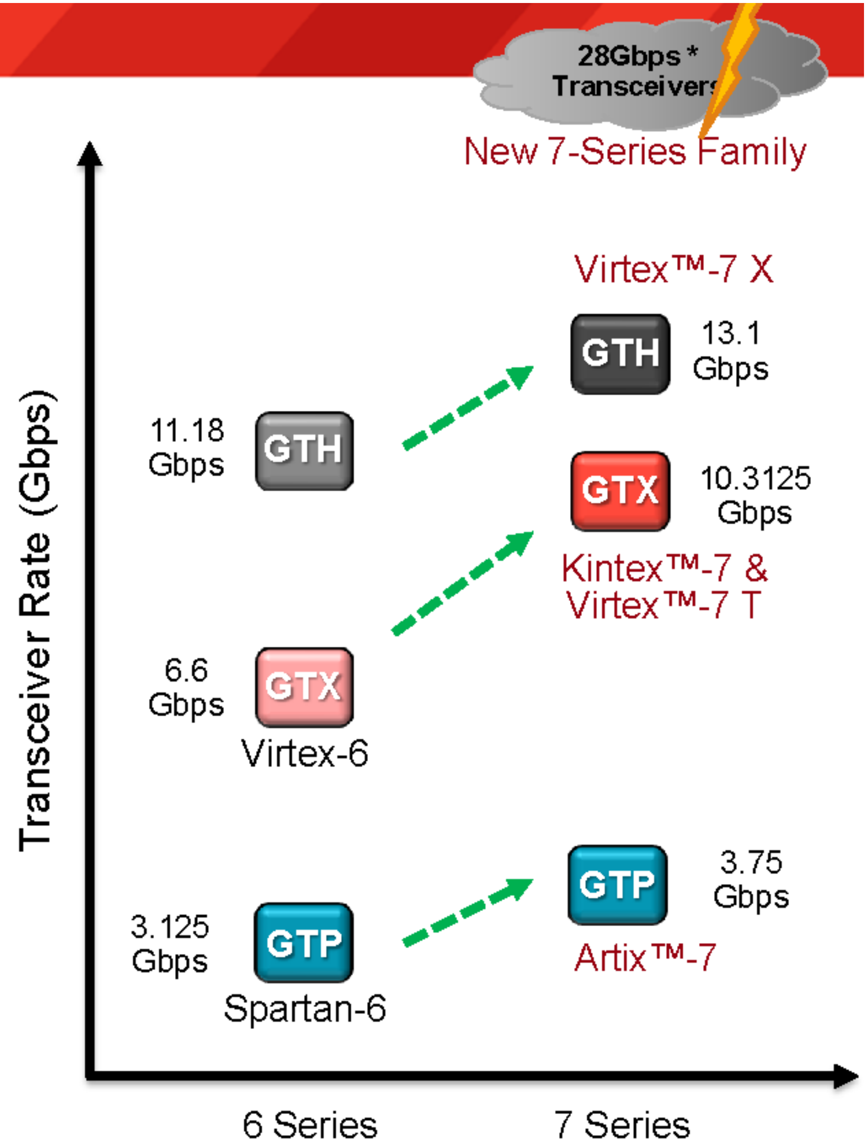

Trigger & DAQ Summary:
LHC Case

Level 1 Trigger

- Select 100 kHz interactions from 1 GHz (10 GHz at SLHC)
- Processing is synchronous & pipelined
- Decision latency is 3 μs (x~2 at SLHC)
- Algorithms run on local, coarse data
  - Cal & Muon at LHC (& tracking at SLHC)
  - Use of ASICs & FPGAs (mostly FPGAs at SLHC)

Higher Level Triggers

- Depending on experiment, done in one or two steps
- If two steps, first is hardware region of interest
- Then run software/algorithms as close to offline as possible on dedicated farm of PCs