CMS HL-LHC Trigger Upgrade

Wesley H. Smith

U. Wisconsin - Madison

Instrumentation Frontier Meeting 2016:
NEW TECHNOLOGIES FOR DISCOVERY II

Caltech, October 9, 2016

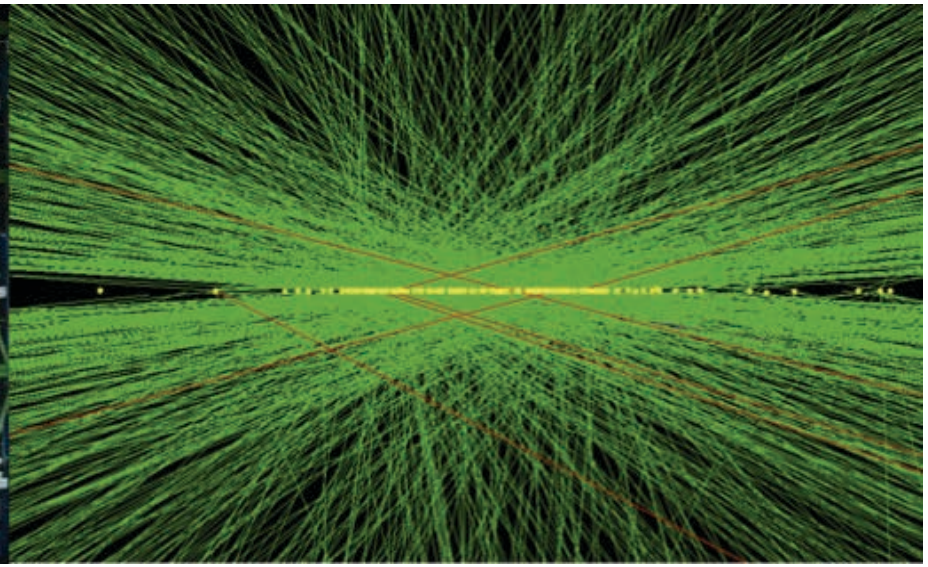

CMS HL-LHC Trigger Scenario

L1 Accept rate up to 750 kHz

- Plan for L1A rate up to 500 kHz @ 140 PU and 750 kHz @ 200 PU
- Driven by keeping acceptance consistent with Run 1 and Phase 1 Upgrade
- Includes acceptance of consecutive 25 ns beam crossings
- Limited by pixel readout, impacts on DAQ readout, EVB, & HLT CPU

Tracking Trigger

- Driven by keeping acceptance consistent with Run 1 and Phase 1 Upgrade
- Leptons: P_T cut & isolation, Jets: Vertex
- Improves E_T sums and enables Track MET

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 $|\eta| > 1.5$ (e.g. GEMs)

Latency of 12.5 μsec

- Driven by Tracking Trigger and logic to incorporate it.
- Limit from complications for outer tracker readout

HLT Output Rate up to 7.5 kHz

- Plan for 5 kHz @ 140 PU and 7.5 kHz @ 200 PU
- Limit from Computing (e.g. cost), no limit from DAQ
- Same reduction of L1 rate (~100) as present

CMS HL-LHC L1 Trig. Components

Calorimeter Trigger

- Process individual readout granularity cells for optimal matching to track trigger
- Data processed by input Layer 1 and then final Layer 2 providing the output. Similar to Phase 1 upgrade calorimeter trigger, essentially scaled to higher number of channels involved.

Endcap Muon Trigger

- Covering $|\eta|$ from 1.6 to 2.5: rebuilt to incorporate additional chambers in endcap and to provide input to the tracking correlator.

Overlap & Barrel Muon Triggers

- Modifications of existing muon triggers covering the barrel and overlap regions to provide input to tracking correlator.

Track Trigger Correlator

- L1 Track Finding is contained within the Tracker, with L1 Trigger performing correlation of produced track with muon and calorimeter trigger information.
- Logic is based on adaptation of Particle Flow ideas to L1 Trigger.
- Input trigger data is processed by an input Layer 1 and then final Layer 2 providing output to Global trigger

Global Trigger

- Process more information than Phase 1 upgrade from many more objects with additional Tracking Trigger load. Design scales by ratio of data volume from Phase 1 upgrade.

HL-LHC L1 Trigger Architecture

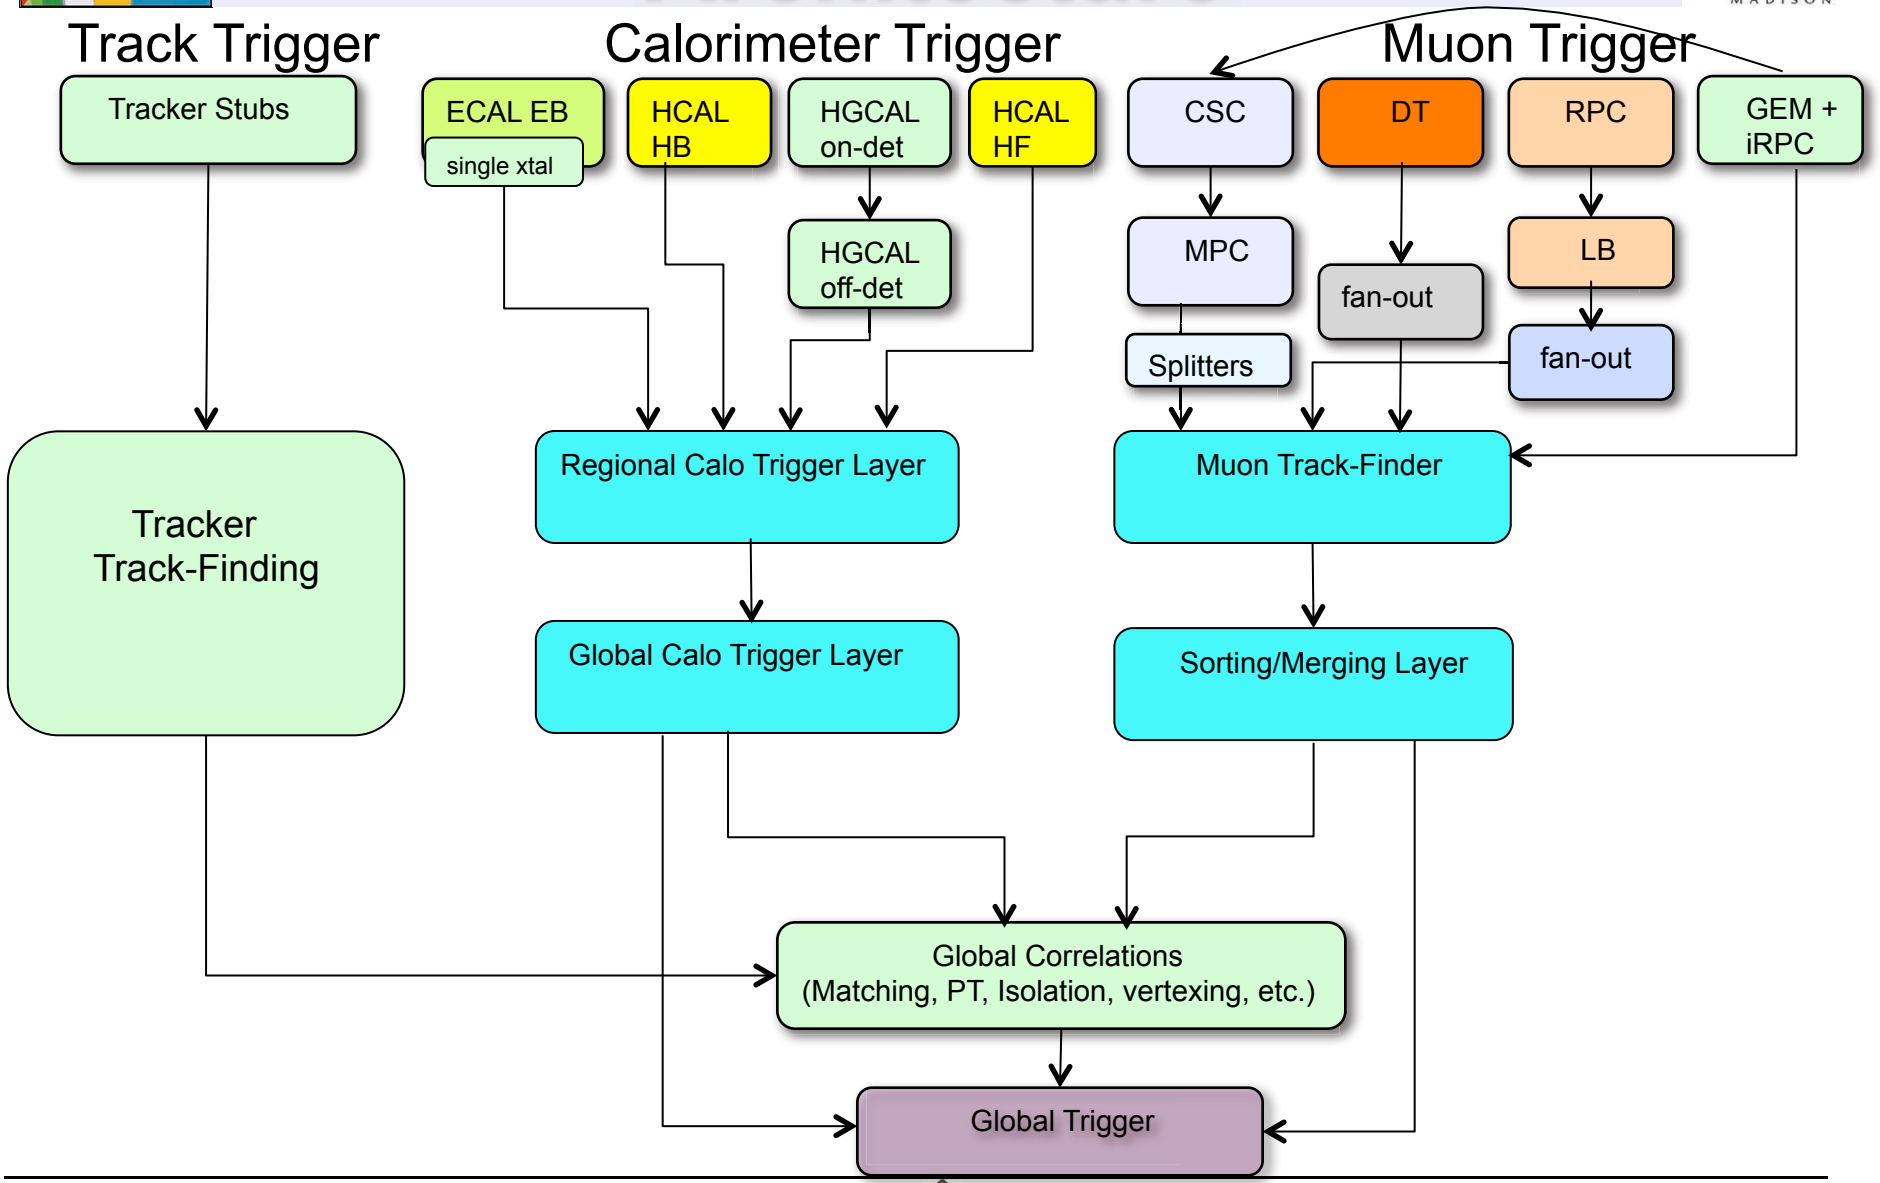

Model for HL-LHC L1 Cal. Trigger

cluster ECAL using fine granularity information for e/γ candidates for track matching/veto + track isolation, and use wider H clusters behind for veto, etc.

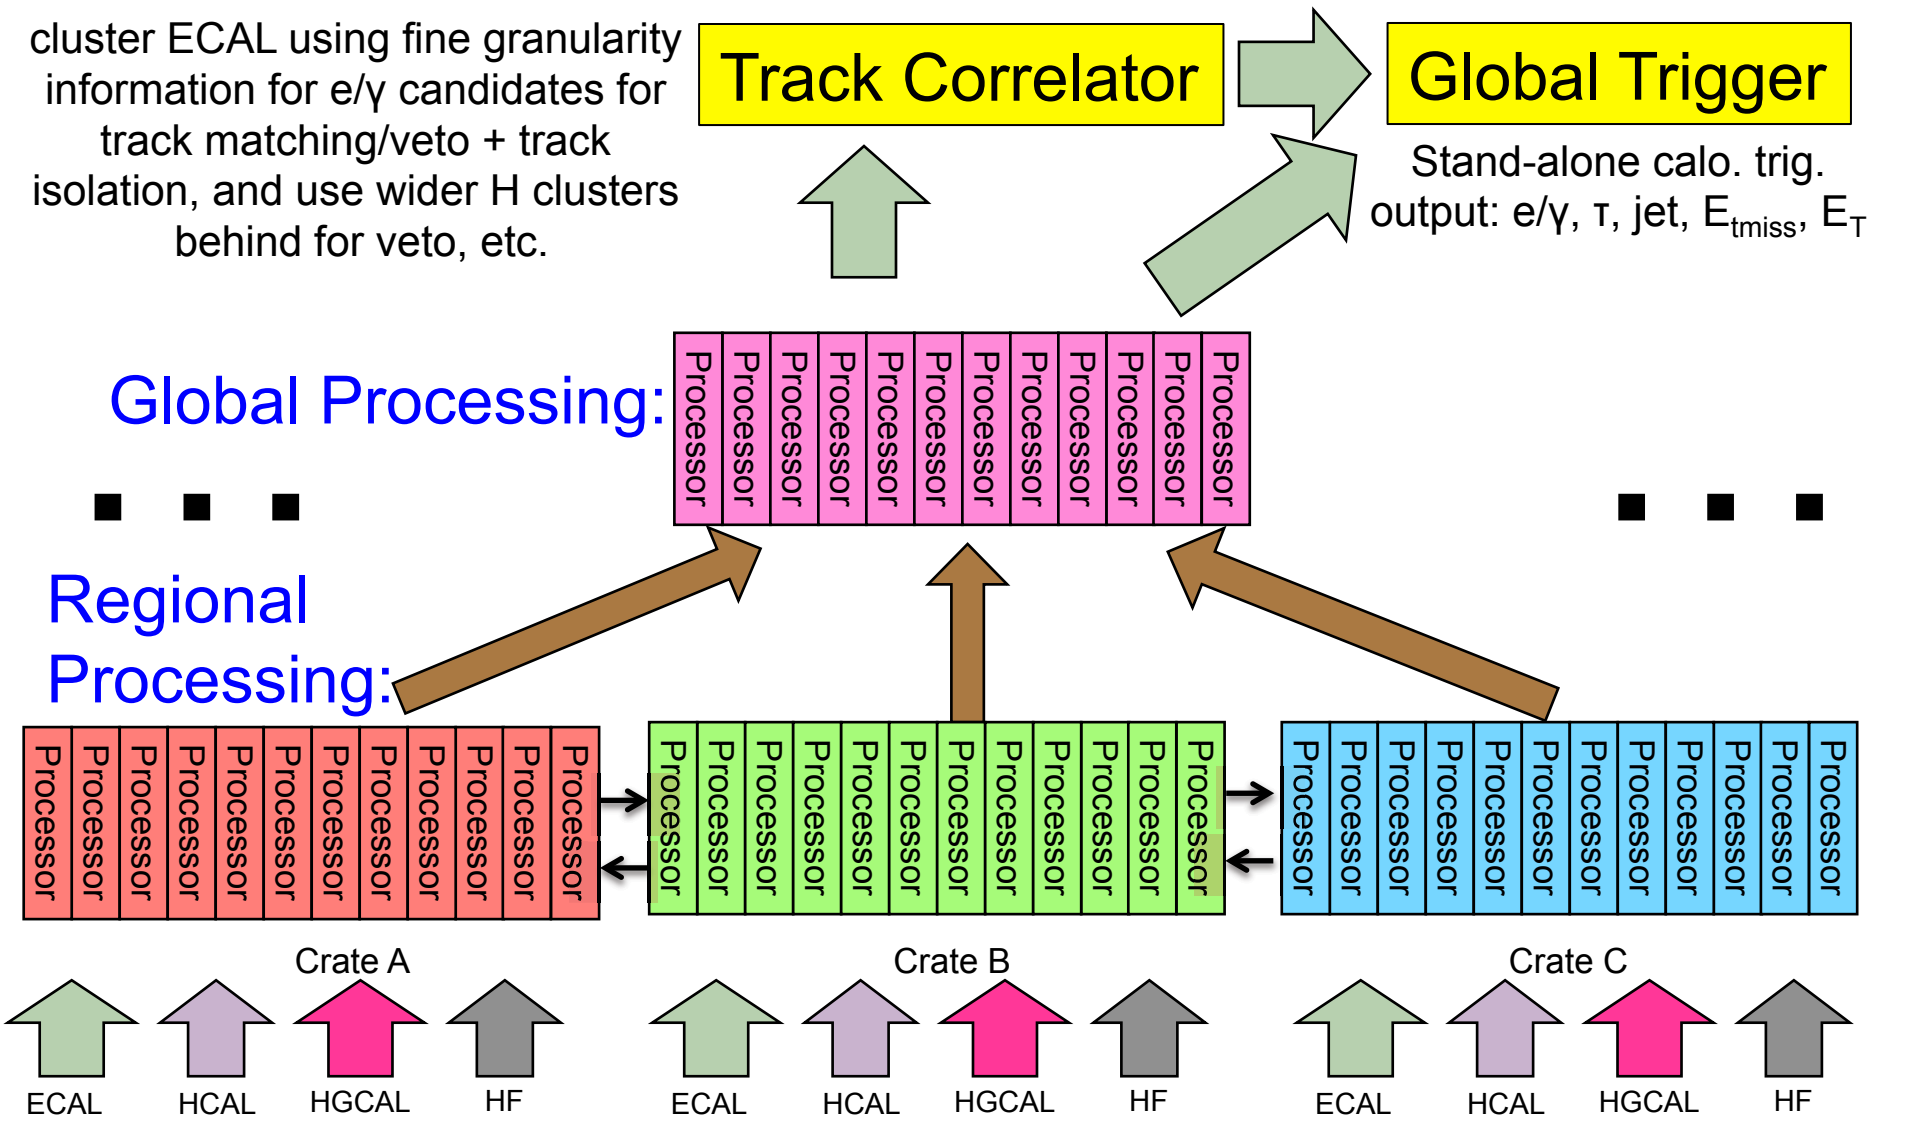

L1 Calo Trigger HL-LHC

- **ECAL Barrel: process individual crystal energies instead of present 5x5 crystal towers**
  - **$\eta \times \phi$ = .0875×.0875 → .0175×.0175**
- **HCAL Barrel: keeps .0875×.0875**
- **HGCAL: produce energy information on ~ same scale (.0175×.0175) as Barrel ECAL.**
- **Improvement in stand-alone electron trigger efficiency + rate example from Barrel →**
- **Also provides higher resolution for matching to tracks: ΔR < 0.006**
  - **See later slide**

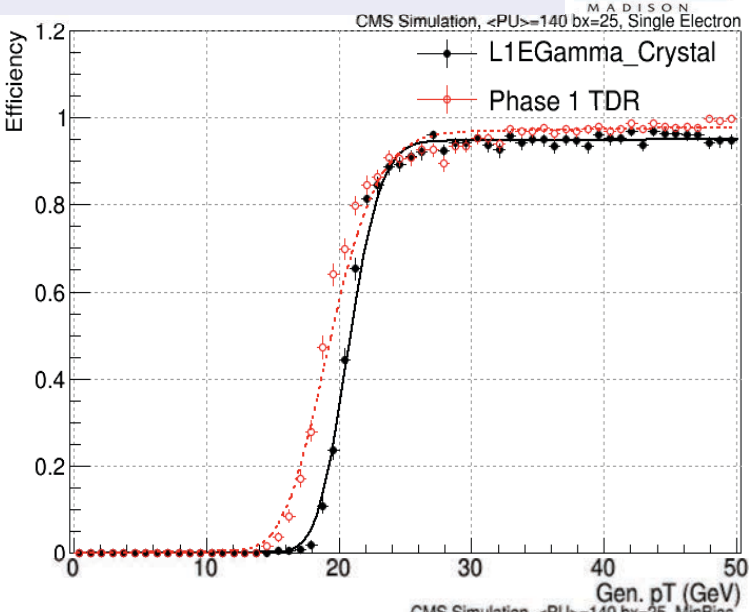

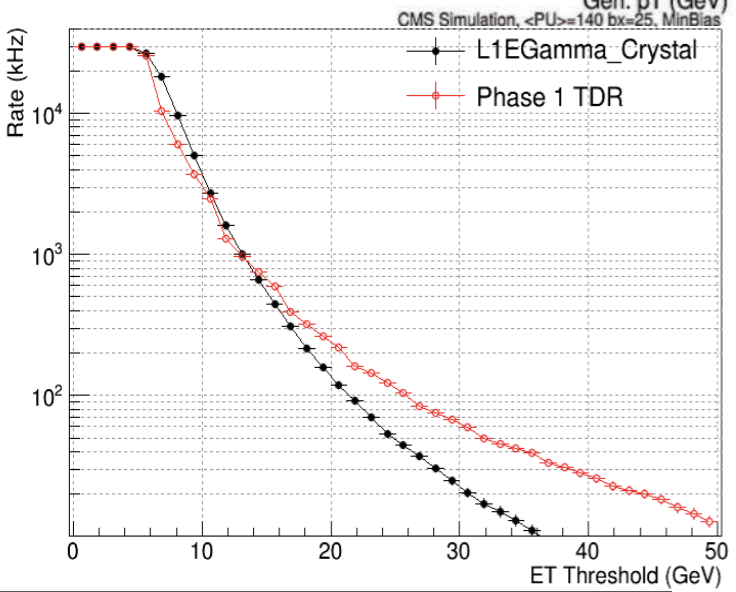

HL-LHC Muon Trigger

- New detectors in forward region to add redundancy and enhance trigger
  - GEM in ME0, GE1/1, GE2/1
  - Improved RPC in outer layers
- Additional inputs and improved angular information (bend angle) for track-finding for better efficiency and rate rejection in stand-alone muon trigger.

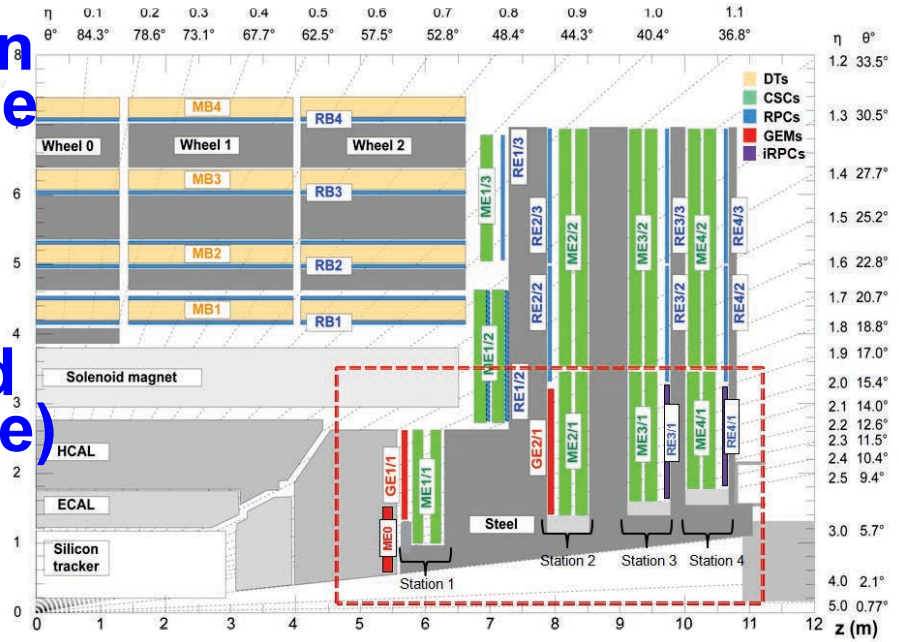

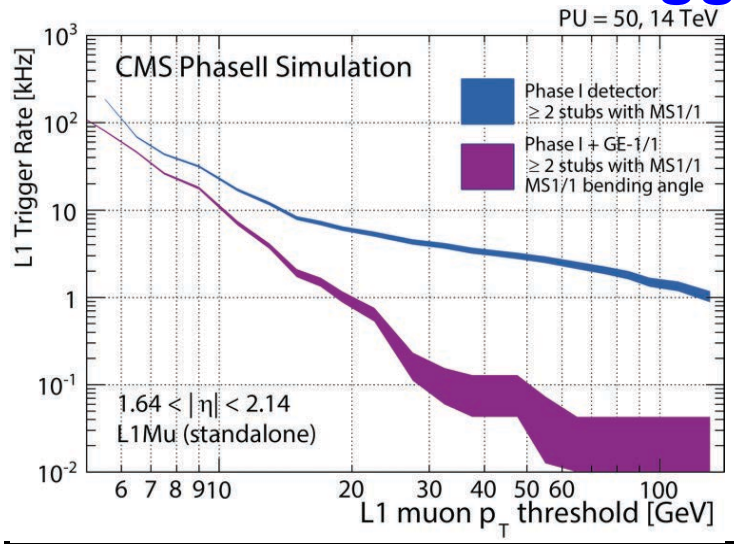

Rate only in GE1/1 region

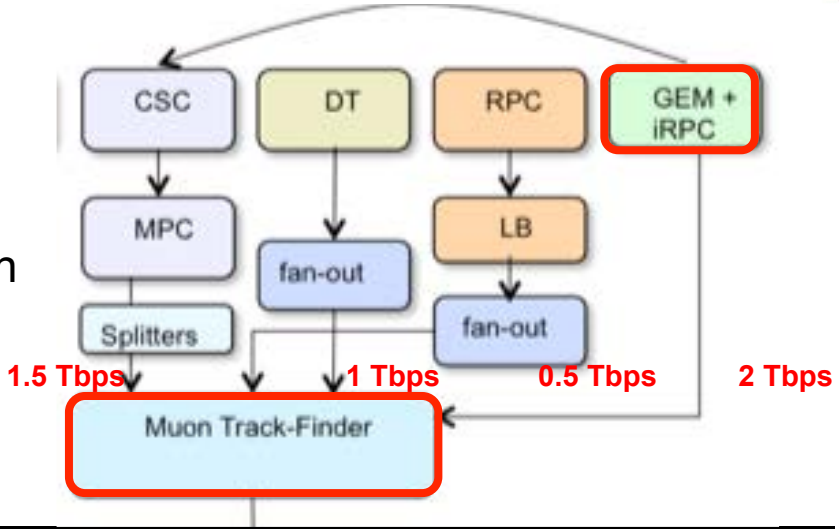

HL-LHC L1 Tracking Trigger

Require:

- Highest possible efficiency over all η for isolated high P_T tracks
- Good efficiency for tracks in jets for vertex identification
- P_T > 2-3 GeV (small difference within this range)
  - Expect ~ 115 charged tracks with P_T > 2 GeV at PU = 140
  - Design for 350 tracks per bunch crossing
- Z vertex position resolution ~ 1 mm

Use:

- Charged Lepton ID
- Improve P_T resolution of muons
- Determine isolation of leptons and photons
- Determine vertex of charged leptons and jet objects
- Determine primary vertex and MET from L1 Tracks from this vertex

HL-LHC L1 Track Trigger Use

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.

L1 Track Trigger Design

Self-seeded Level-1 Track Trigger

- Relies on local p_T reconstruction
- Reconstructs tracks with p_T > 2 GeV
- Identifies z-vertex location with ~ 1 mm precision
  - Similar to average vertex separation at PU ~ 140.

p_T modules provide p_T discrimination through hit correlations between closely spaced sensors

- Correlate pairs of clusters consistent w/track > 2 GeV
- In minimum bias events, ~ 95% of tracks have p_T < 2GeV

Example: CMS Track Trigger design finds stubs with $p_T \gtrsim 2$:

Thanks to CMS 3.8 T magnetic field!

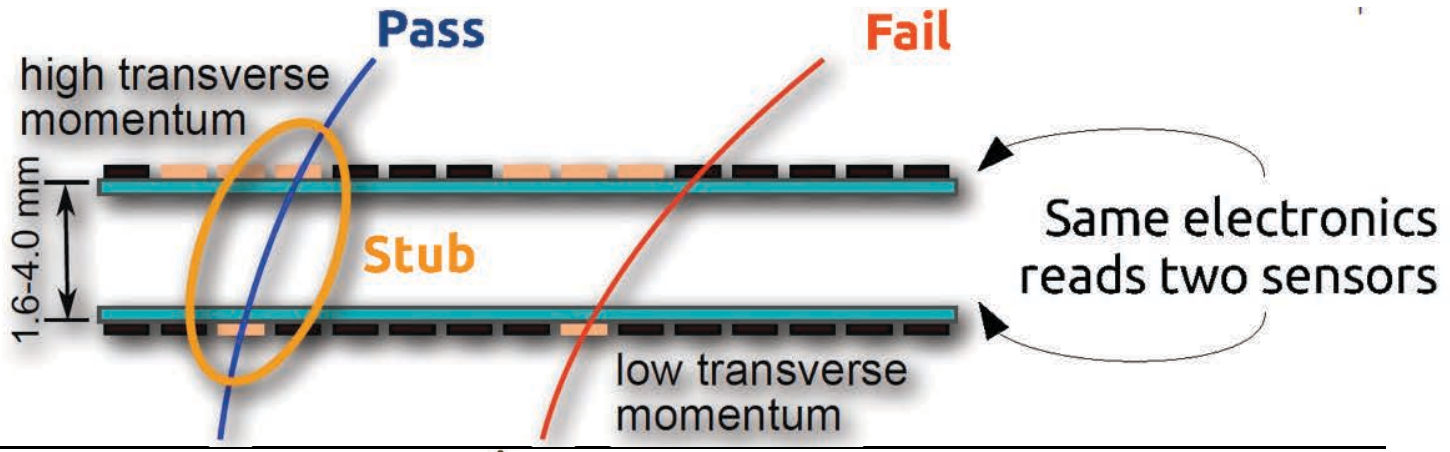

Track Stub Finding

Push" design: all stubs forwarded to Track-finder

Stub efficiency vs. p_T for various layers and disks:

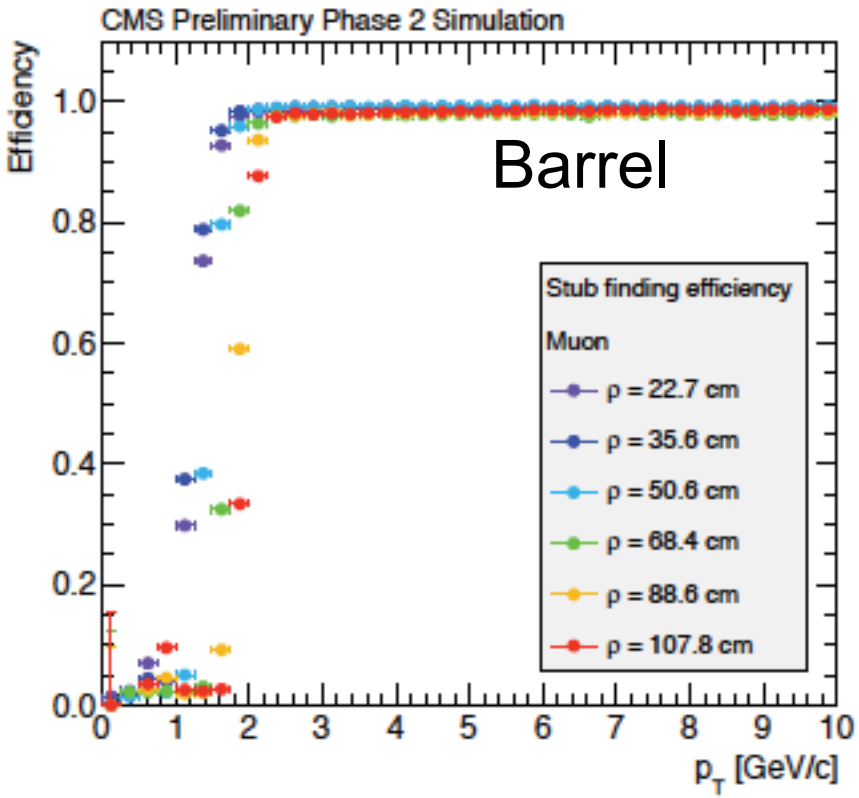

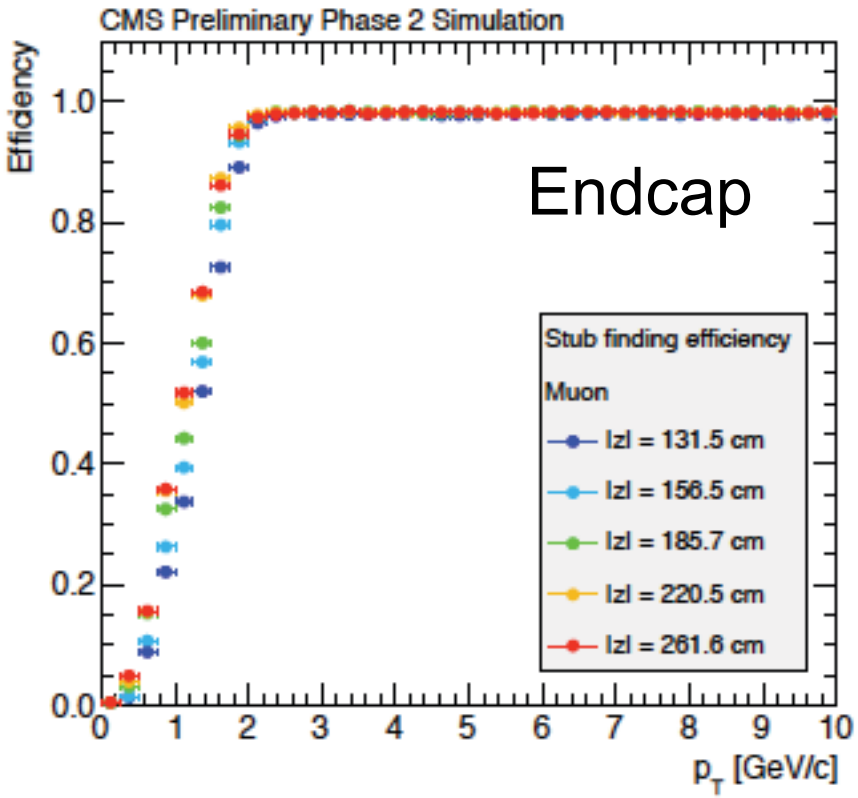

1 BX = 25ns

Track-Finding

"Push" design, not region of interest

- Input: Expect ~10K stubs/BX @ PU~140, of which 5-10% belong to tracks with P_T>2 GeV
- Output: ~100 bits/track → several Tbps to track correlators
- Latency: ~5 μs allocated

Multiple approaches under consideration to find tracks, examples:

- Pattern-based: Track patterns stored in Associative Memory chips
  - Target implementation in custom ASICs (Associative Memories, e.g. VIPRAM) to store large number of patterns (~100M overall)
- "Tracklet" approach: Track building from stubs with pair-wise layer extrapolations
  - Target implementation in FPGAs
- More details in Tracking Parallel Session

Generally implemented with time-multiplexing of the input data (round-robin of event data to processors)

All followed by a (generally FPGA-based) track-fitting stage to extract track parameters

Track Trigger Performance

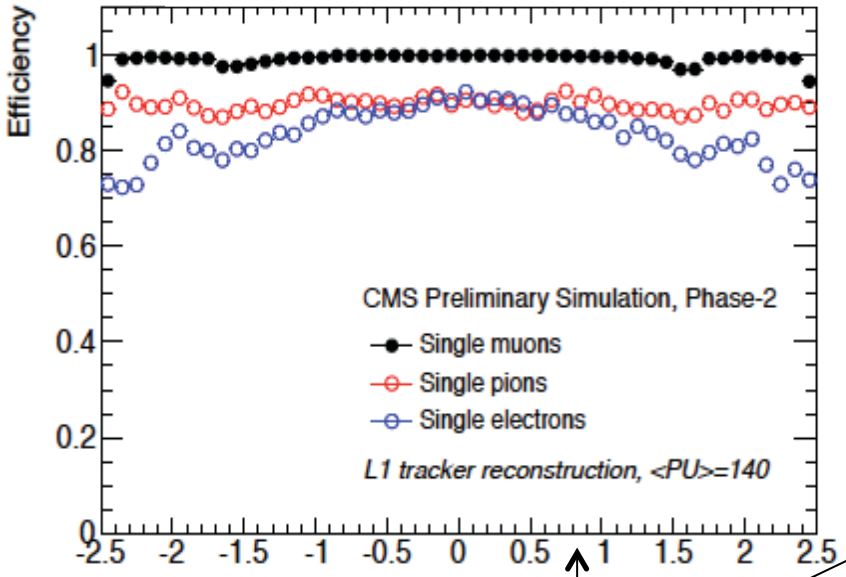

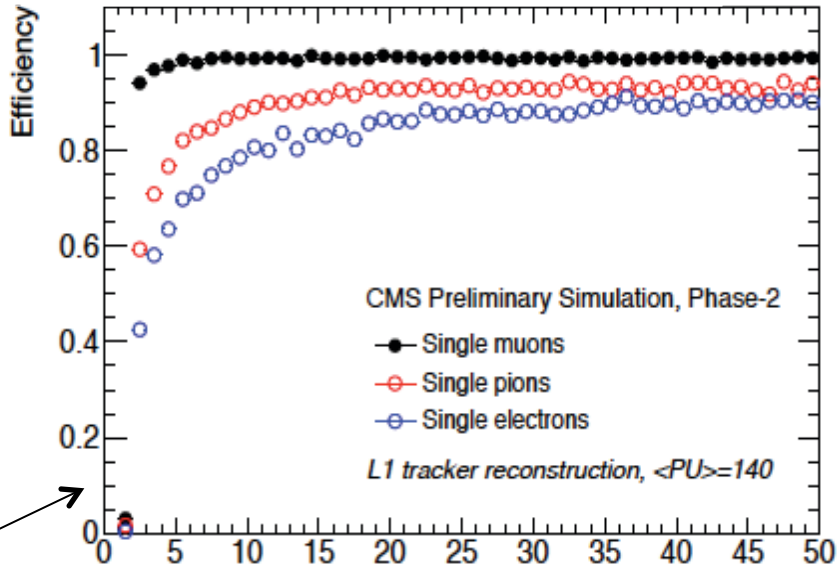

Efficiency curves vs. η & p_T for PU=140

Z_0 resolution:

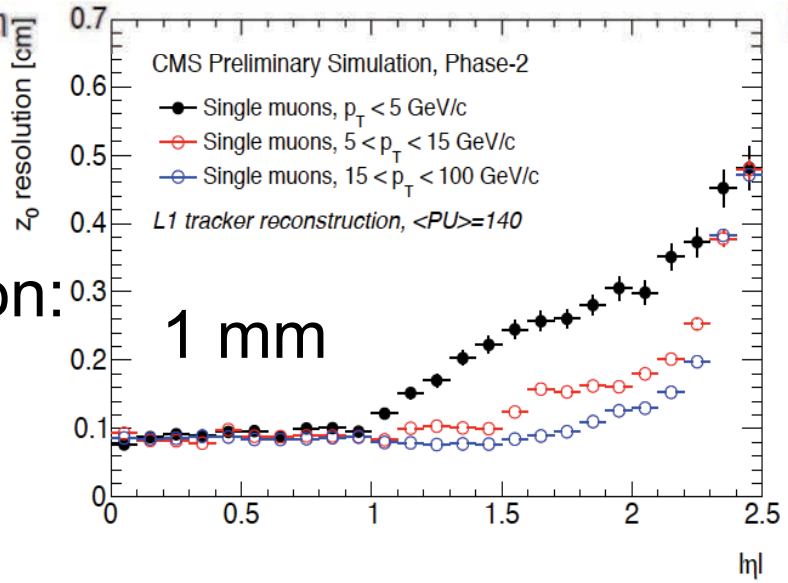

Correlator Conceptual Design

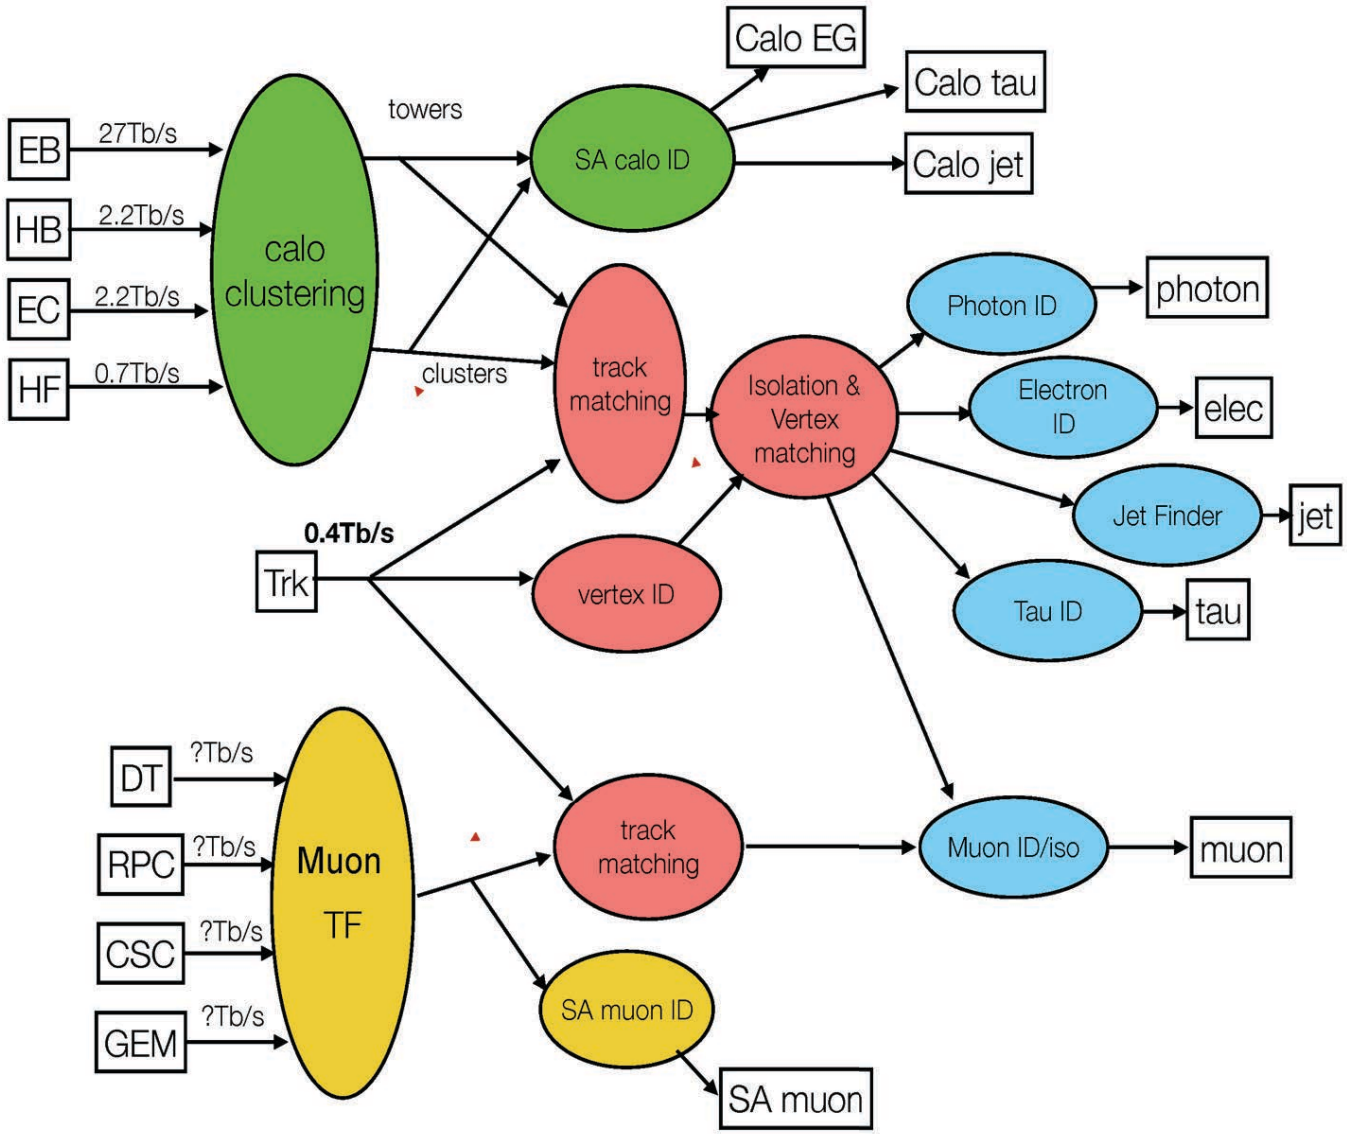

Model for L1 Correlator Trigger

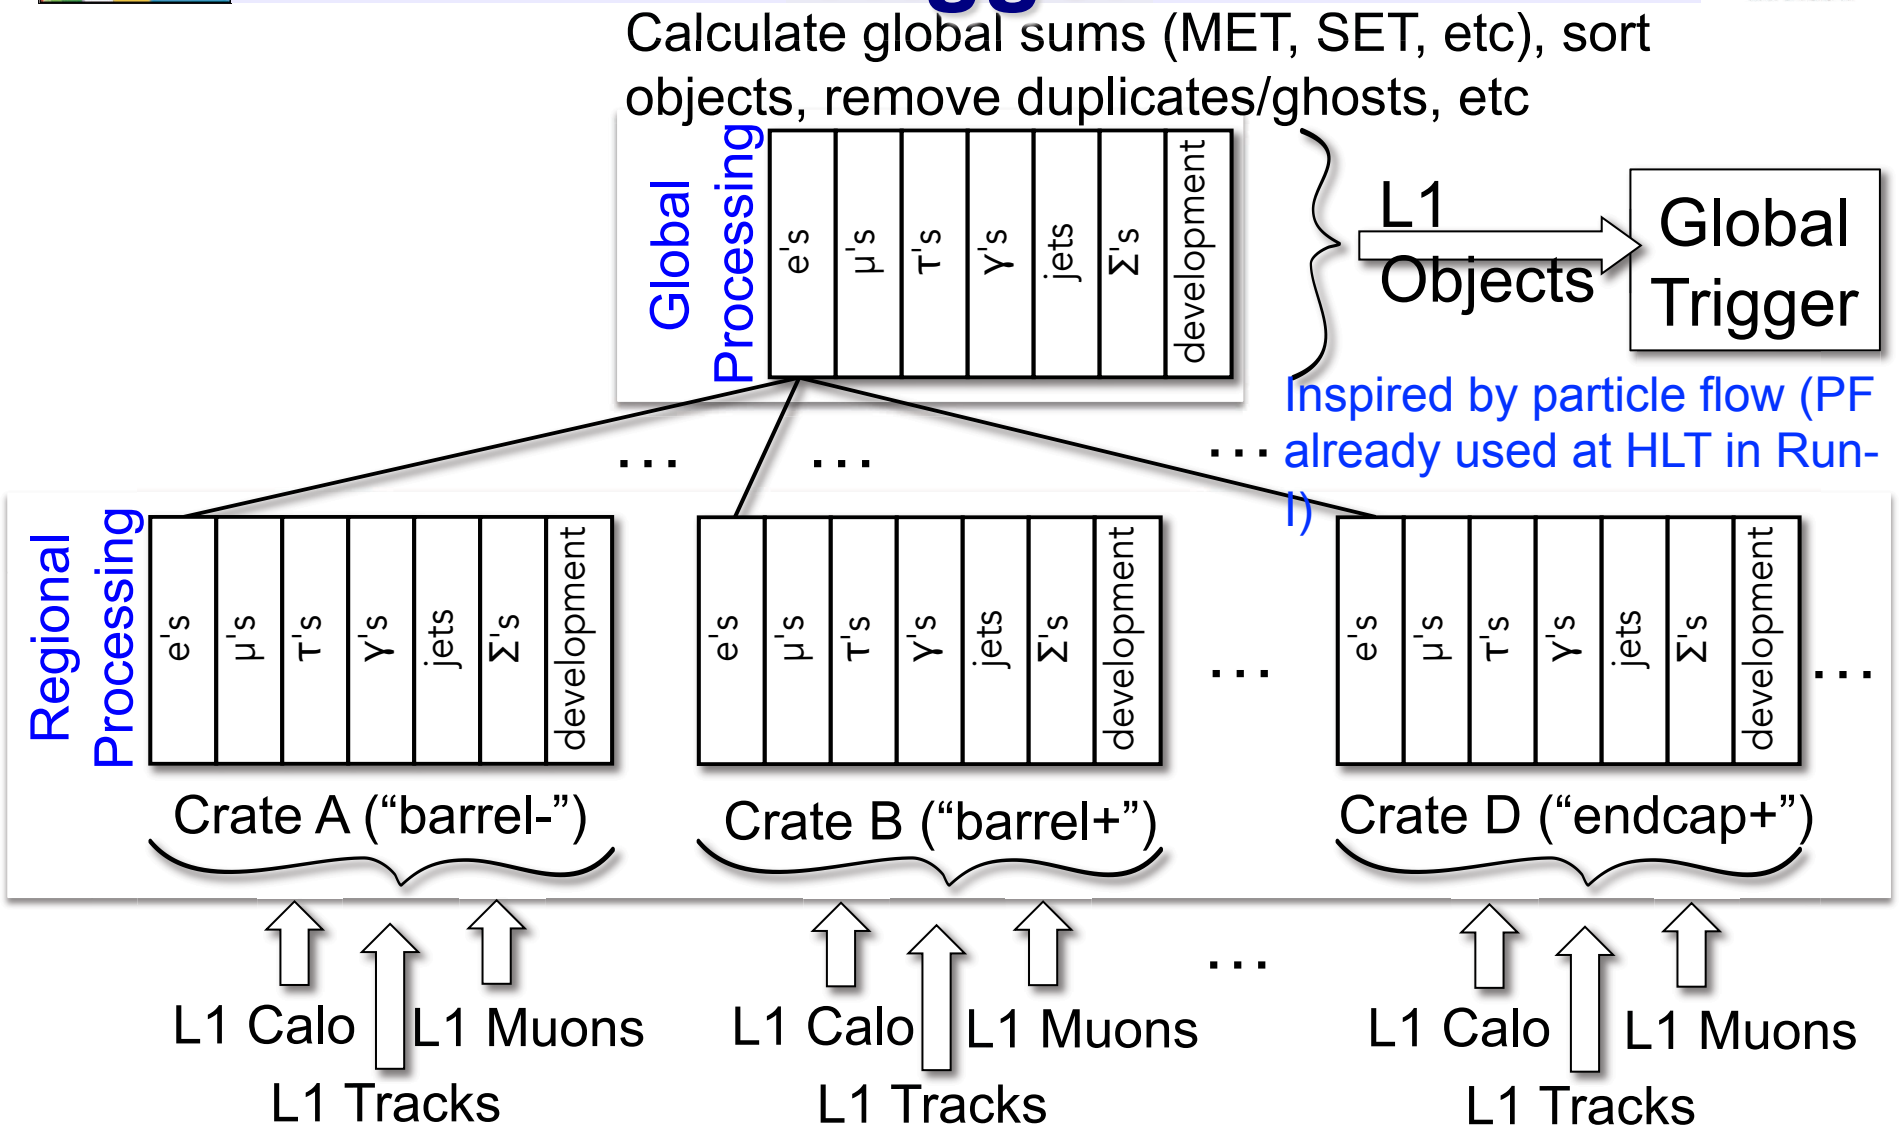

CMS HL-LHC L1 Trigger Latency
(preliminary estimate)

Estimate L1 Track Trigger result available 5 μs after interaction occurred.

- L1 Track trigger is a "push" system, not region of interest

Regional processing (Correlation Logic) involves (+2.5 μs ⇒ 7.5 μs) :

- Preprocessing + distribution of L1 Cal, Mu, Track Trigger Primitives – 0.5 μs
- Using tracks to find primary vertex – 0.5 μs
- Associating tracks with the primary vertex – 0.5 μs
- Associating tracks with calorimeter objects
- Associating tracks with and fitting tracks with muon tracks
- Calculating track-correlated L1-objects + their characteristics – 1.0 μs
- Using tracks to calculate isolation of cal. and muon objects – 0.5 μs.

In parallel with 0.5 μs; 2.0 μs

Global processing involves +(1 μs ⇒ 8.5 μs):

- Global sums, kinematic calculations, correlations between trigger objects, trigger decision logic (incl. trigger rules)

Propagation back to detector front ends (+1 μs ⇒ 9.5 μs)

- 38 bx in present system (0.95 μs)

Safety Factor of 30% ⇒ 12.5 μs

- NB: Original Trig. TDR Latency: 127 bx, Run 1: 157 bx (+ 0.8 μs), req'd: 168 bx
  - Run 2 Latency now up to 164 bx.

Set 12.5 μs as minimum design latency

- Once 6 μs exceeded, 12.5 μs is the next point where consequences

Track Trig. Correlator: muons

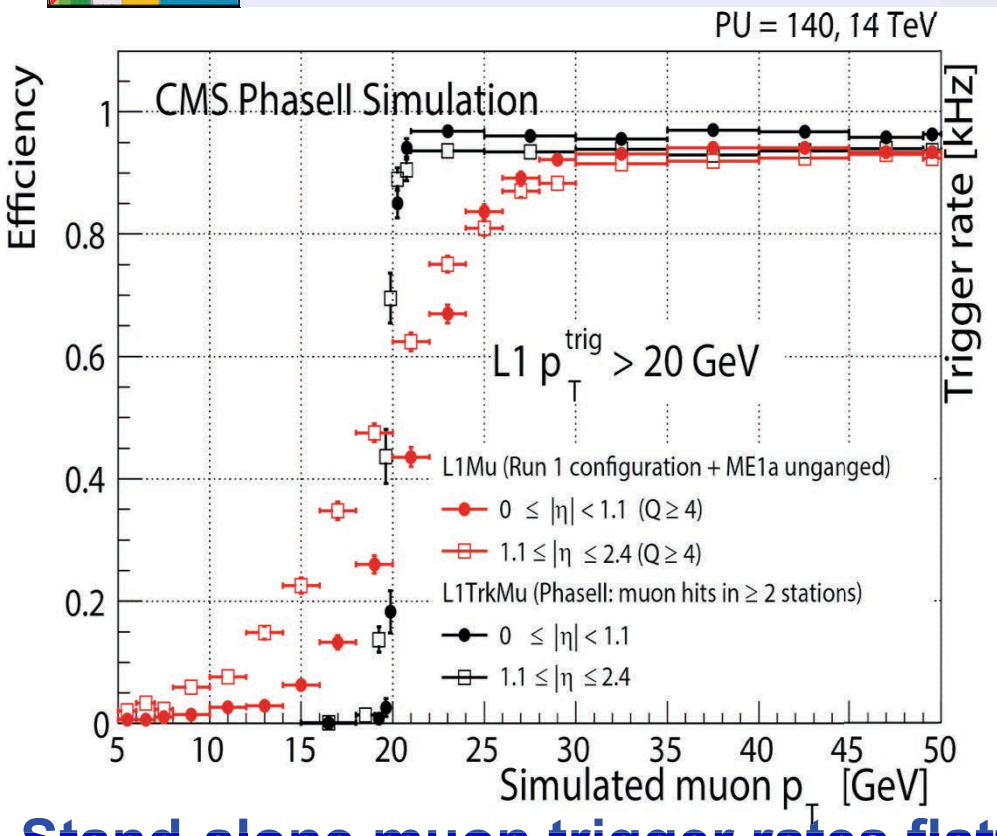

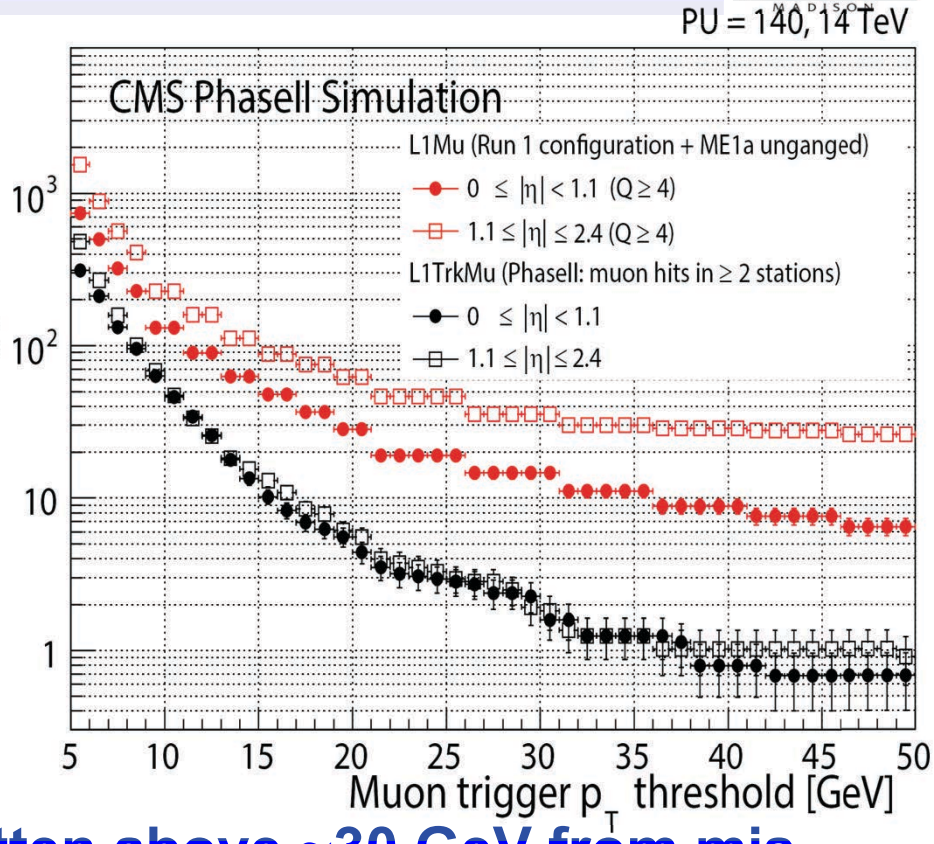

Stand-alone muon trigger rates flatten above ~30 GeV from mis-assignments of low p_T muons

Match L1 tracks + L1 muons (inside-out vs outside-in)

Sharpens turn-on curve, improves efficiency

For a threshold of ~ 20 GeV, rate of single muon trigger can be reduced by a factor of O(10)

Track Trig. Correlator: electrons

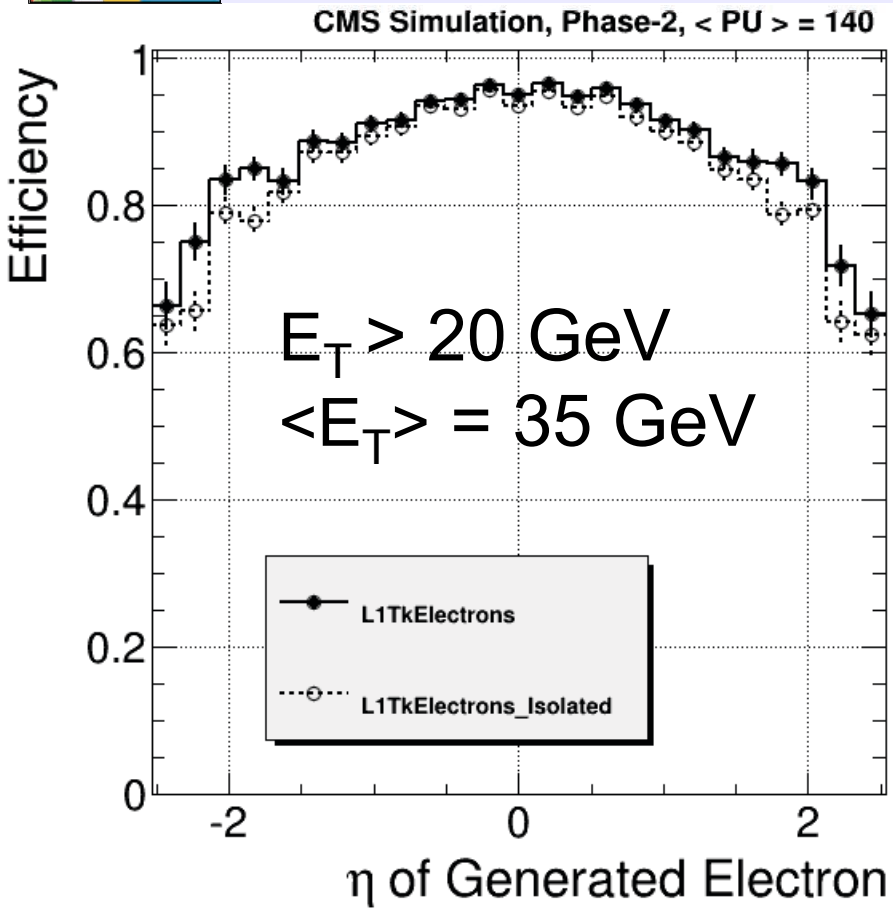

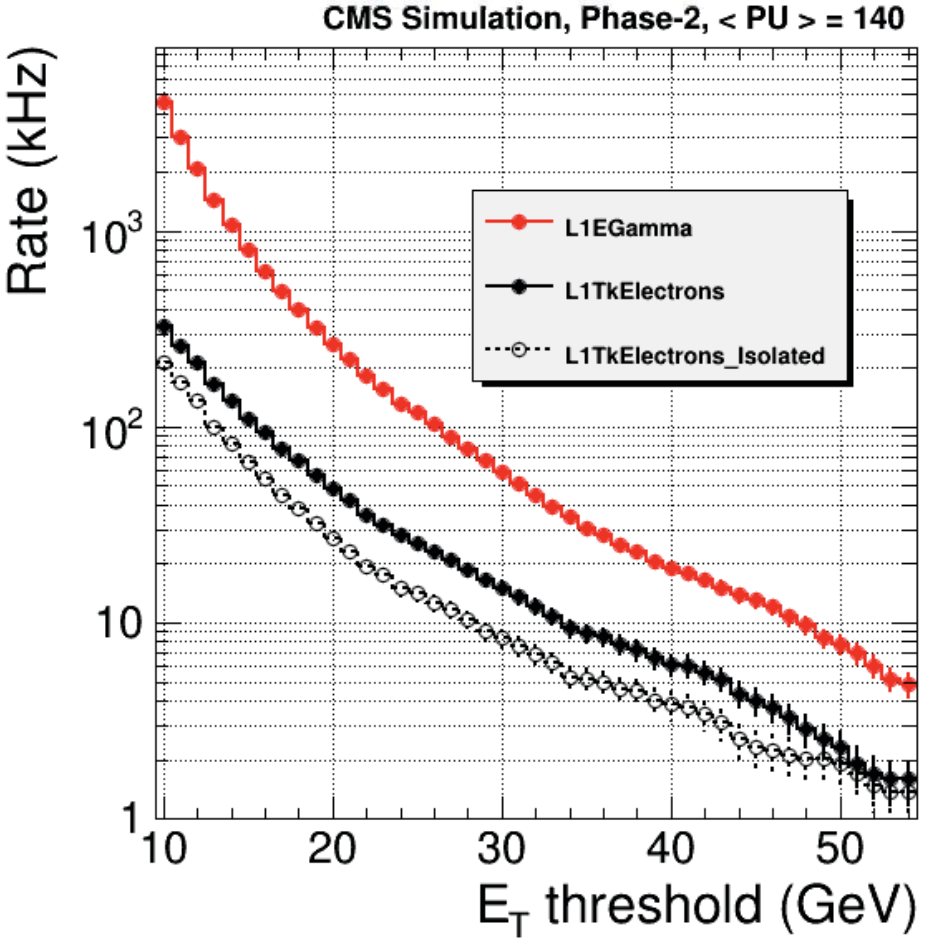

Efficiency exceeds 90% in central region

Match with Track Trigger plus track-based isolation retains good efficiency with rate reduction of O(10)

Track Trig. Correlator: photons

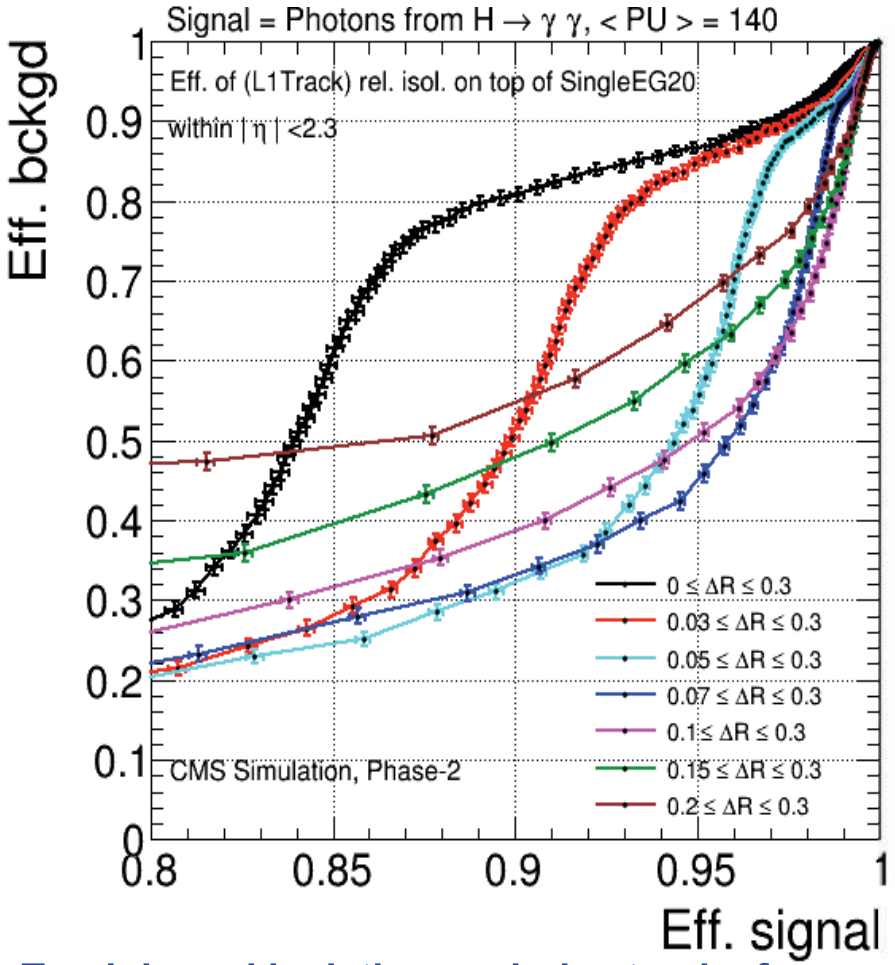

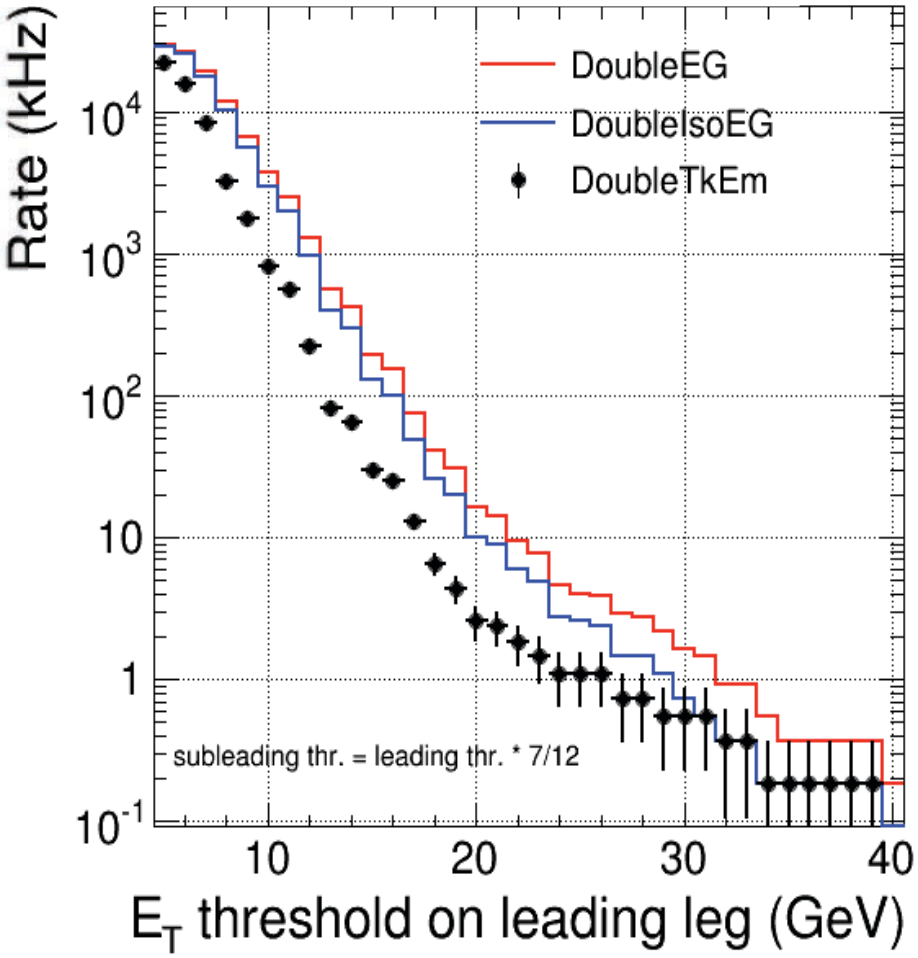

Track-based isolation excludes tracks from conversions

Reduce rate of Single Photon trigger by 2 (3) for an efficiency on H → γγ of 95% (90%)

For diphoton thresholds of ~ 18,10 GeV on leading, subleading legs, the rate can be reduced by a factor of > 6

Track Trig. Correlator: taus

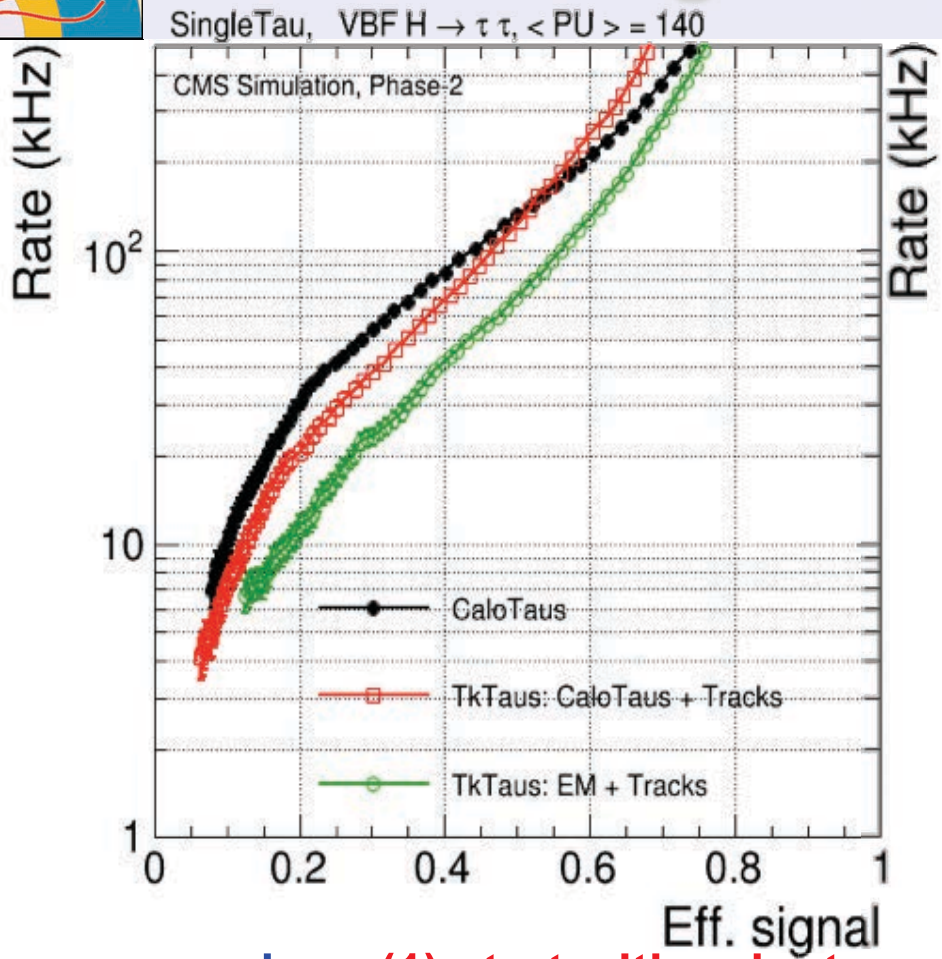

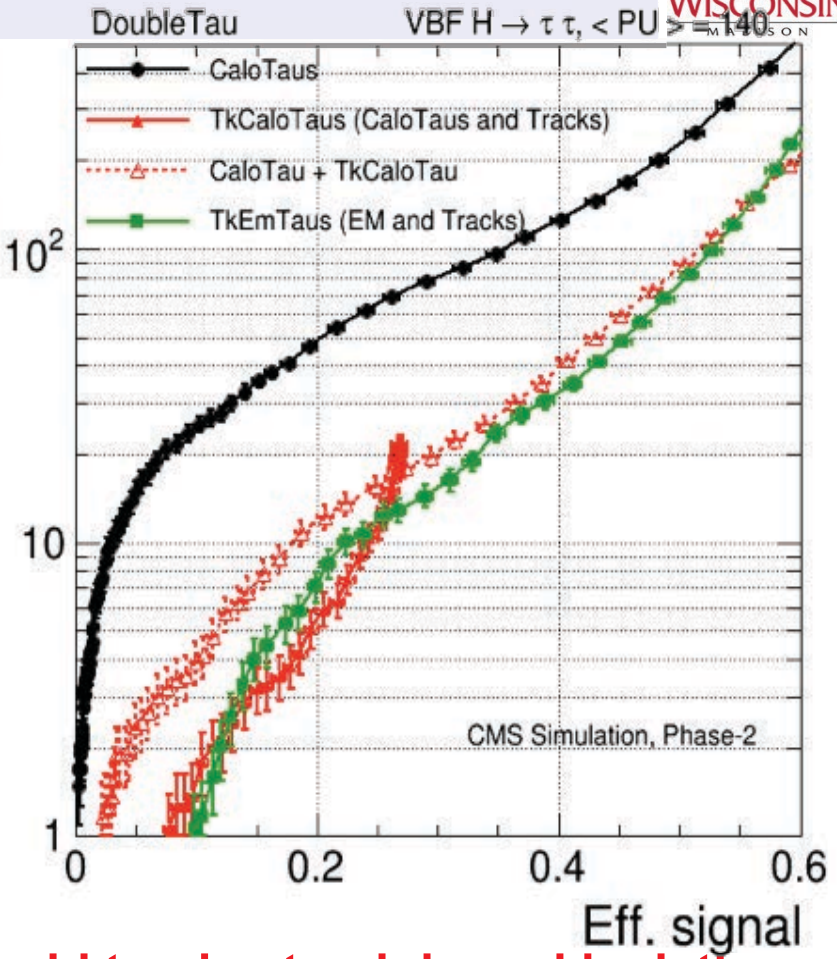

Two approaches: (1) start with calo. taus, add tracks, track-based isolation or (2) start w/tracks, add crystal L1 EM objects.

With either approach can maintain ~ 50 kHz trigger rate with ~ 50% efficiency for VBF H →ττ

Track Trig. Correlator: MET, MHT

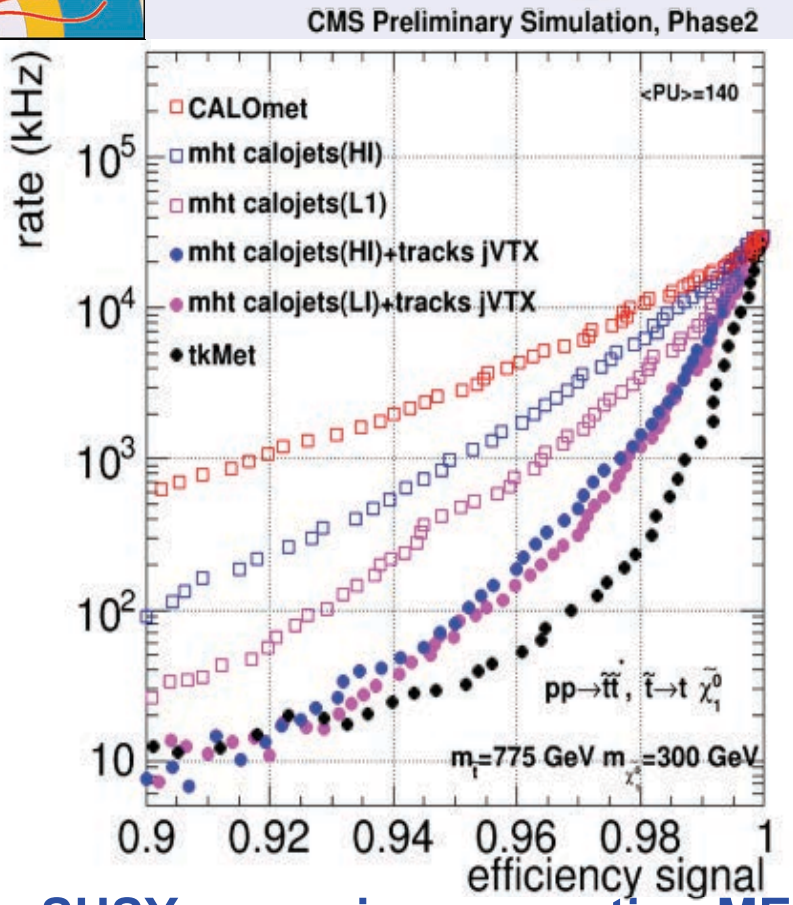

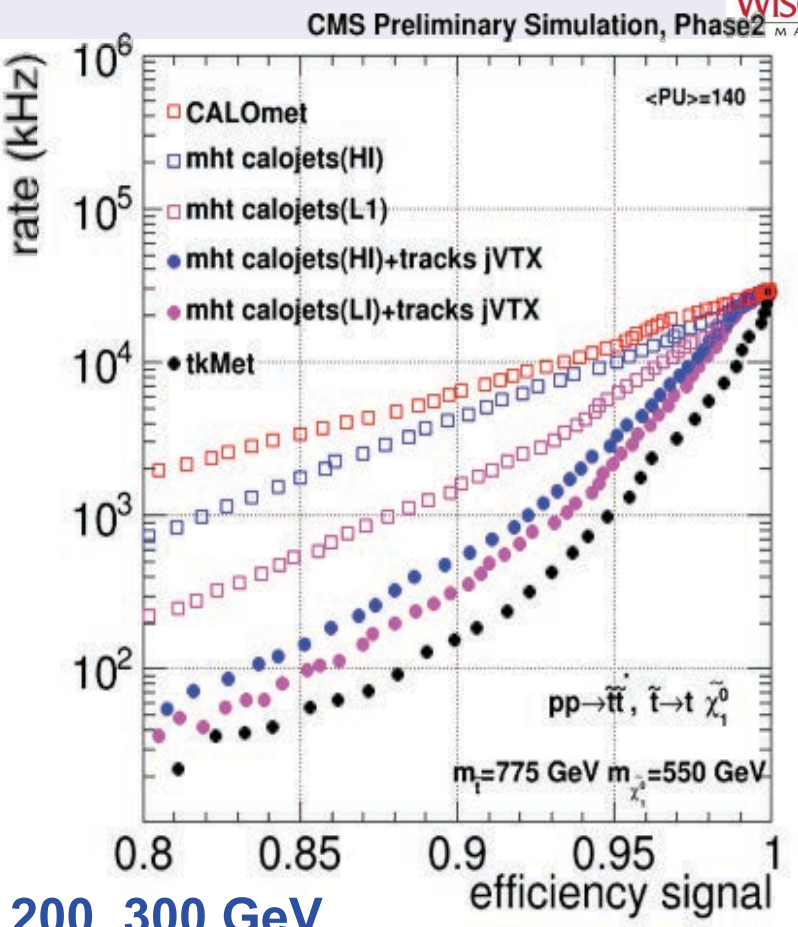

Two SUSY scenarios generating MET of 200, 300 GeV

- **Stop pair decays to top and neutralino**

Signal defined to be events with gen. MET > 100 GeV

MHT plus vertex constraint 95% efficient with few 10's kHz rate

Track-based MET reduces rate over calo MET by 2 orders of magnitude while maintaining 80-85% efficiency in "Low MET" scenario

L1 Menu Studies

Goal: maintain overall physics performance + acceptance of L1 Trigger Upgrade TDR (e.g. thresholds near the end of run 1)

- **Study a sample menu using ~70% of bandwidth**
  - Not included: any prescaled trigger (minbias, triggers with lower thresholds), triggers involving forward calorimeter, MET/MHT, three lepton triggers, other acceptance, diagnostic and calibration triggers.
- **Compensate by increasing found total rate by 30%**

Benchmark conditions: 140 PU

- **Use Minbias Sample with PU = 140**
- **Assume bunch spacing at 25 ns: L ~ 5.6 E34**

Use Tracking Trigger

- **Assume "perfect" performance**

High PU Conditions: 200 PU as a check

- **Use unofficial production of 200 PU events**
- **Assume bunch spacing at 25 ns: L ~ 8 E34**

Uncertainties: Need Safety factor of at least 1.5

- **Simulation uncertainty**
- **Readout underachievement**
- **Realistic Tracking Trigger performance**

L1 Trigger Menu at 140 PU with L1 Tracking Trigger

Menu w/o Tracking Trigger produces a total rate of 1.5 MHz* at 140 PU

- Tracking trigger provides factor of 5.5 reduction to 260 kHz →
- w/1.5 safety factor: 390 kHz
- Use 500 kHz as benchmark

Menu w/Tracking Trigger produces total rate of 500 kHz at 200 PU

- w/1.5 safety factor: 750 kHz
- w/o track trig., approach 4 MHz!
  - (6 MHz w/safety!)

Warning: this menu is just a sample table only for evaluating bandwidth

$L = 5.6 \times 10^{34}$ cm^{-2}s^{-1} $\langle PU \rangle = 140$		Level-1 Trigger with Level-1 Tracks		No Track Trigger
Trigger Algorithm (Menu with L1 Track Trigger)		Rate [kHz]	Offline Threshold(s) [GeV]	Rate [kHz]
Single Mu (tk)		14	18	139
Double Mu (tk)		1.1	14 10	177
ele (iso tk) + Mu (tk)		0.7	19 10.5	160
Single Ele (tk)		16	31	78
Single iso Ele (tk)		13	27	89
Single γ (tk-iso)	(tk-veto)	31	31	70
ele (iso tk) + e/γ		11	22 16	88
Double γ (tk-iso)	(tk-veto)	17	22 16	53
Single Tau (tk)		13	88	34
Tau (tk) + Tau		32	56 56	55
ele (iso tk) + Tau		7.4	19 50	42
Tau (tk) + Mu (tk)		5.4	45 14	52
Single Jet		42	173	185
Double Jet (tk)		26	2@136	144
Quad Jet (tk)		12	4@72	175
Single ele (tk) + Jet		15	23 66	60
Single Mu (tk) + Jet		8.8	16 66	64
Single ele (tk) + H_T^{miss} (tk)		10	23 95	73
Single Mu (tk) + H_T^{miss} (tk)		2.7	16 95	
H_T (tk)		13	350	
Rate for above Triggers		180		1000
Est. Total Level-1 Menu Rate		**260**		1500 *

L1 Thresholds vs. Bandwidth: Single Lepton Triggers

Allocate initial BW for each algorithm in full menu according to BW fractions used in sample menus with and without L1 Tracking

- **Determine trigger thresholds necessary to fit assigned rate**
- **Repeat for a range of total bandwidths**

Plots show rates for full menus with and without L1 Tracking @ 140 PU

- **No safety factors applied**

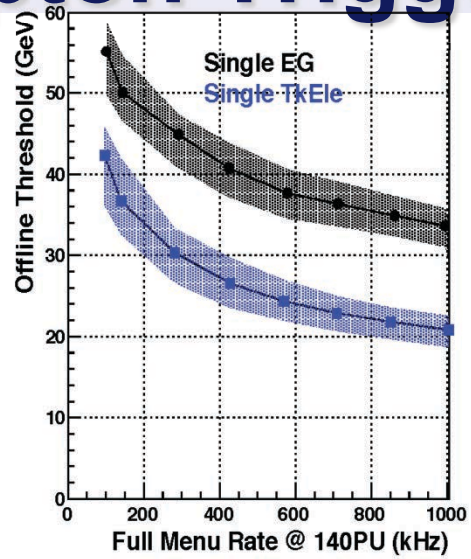

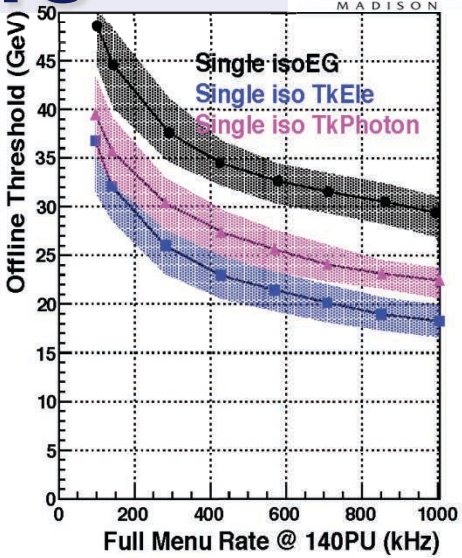

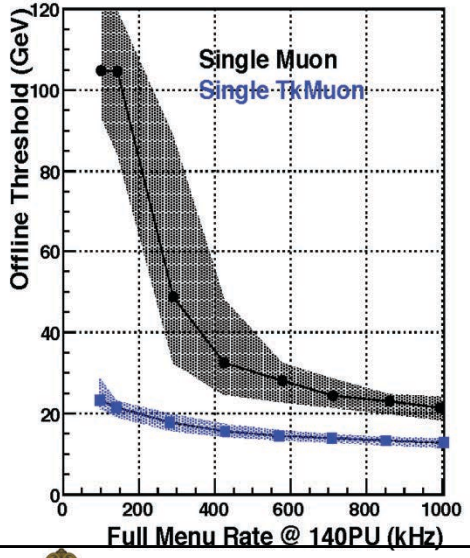

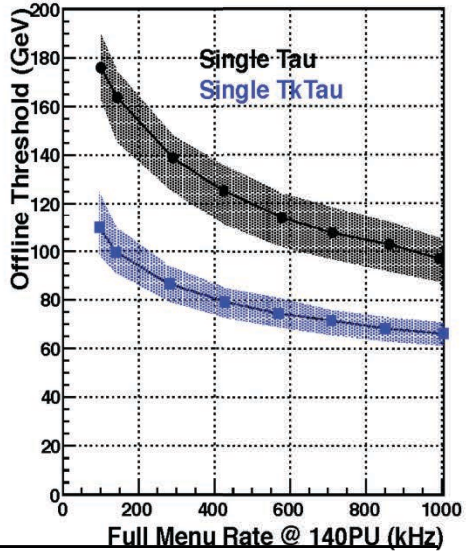

CMS HL-LHC HLT & DAQ

	Run-I	Phase-I upgr.	Phase-II upgr.	
Energy	7-8 TeV	13 TeV	13 TeV	
Peak Pile Up (Av./crossing)	35	50	140	200
Level-1 accept rate (maximum)	100 kHz	100 kHz	500 kHz	750 kHz
Event size (design value)	1 MB	1.5 MB	4.5 MB	5.0 MB
HLT accept rate	1 kHz	1 kHz	5 kHz	7.5 kHz
HLT computing power	0.21 MHS06	0.42 MHS06	5.0 MHS06	11 MHS06
Storage throughput (design value)	2 GB/s	3 GB/s	27 GB/s	42 GB/s

HLT: x24/51 at 500/750 kHz+140/200 PU increased processing power v. Run 1

- **x 5/7.5 increase in L1 input rate (500 kHz @ PU=140/750 kHz @ PU=200)**
- **x 4/5.7 for increased complexity due to higher PU, x 1.5 for higher energy**
  - Estimated from observed CPU dependence on PU and MC studies
  - x 0.8 mitigation due to prompt access to tracking info. from L1 Track Trigger at HLT
- **Affordable in 2024 assuming extrapolation of ~1 CHF/HS06 in 2024**

DAQ: For 5 MB events at 750 kHz, total required throughput is 30 Tbps.

- **Similar event builder and storage as present 800 links x 100 Gbps with 30% eff. will provide 30 Tbps event building throughput**
- **InfiniBand switch today provides 32 Tbps unidirectional bandwidth. Costs of this infrastructure are likely below those of existing system.**

HL-LHC Trigger R&D Goals

Challenges:

- **Large increase in trigger input data**
  - e.g. present EB 5x5 trigger towers vs. full xtal granularity – x25 increase, also new HGCal data volume w/ similar granularity
- **Large increase in processing complexity**
  - Tracking information
  - Fine grain calorimeter information
  - Fitting Muon and Tracking data together
  - More complex objects, conditions and algorithms

Phases:

- **Establish algorithms, techniques and feasibility**
- **Decide on a hardware framework**
- **Build prototypes to test functionality**
- **Build individual parts of the system**
- **Construct demonstrators**
- **Connect to detector prototypes to validate designs**

Schedule and Milestones

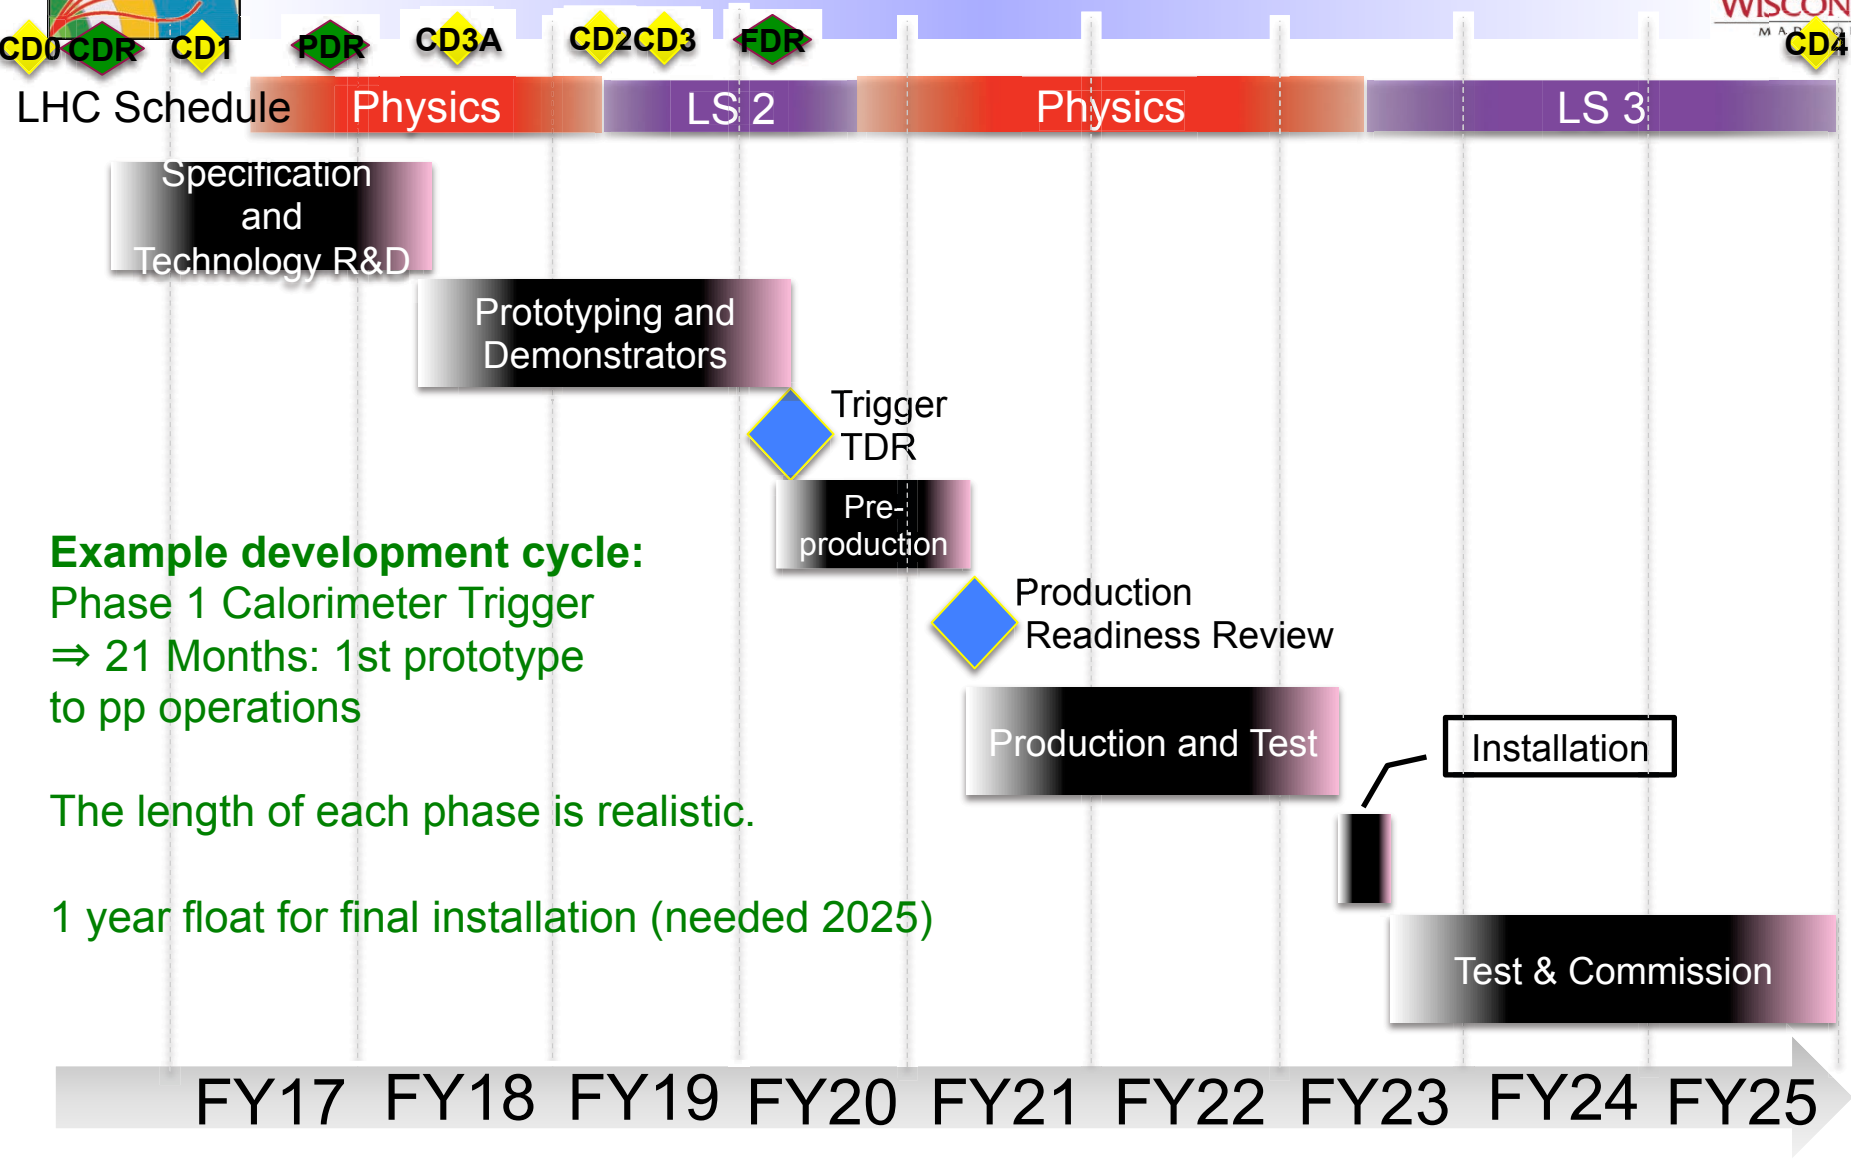

Example development cycle:
Phase 1 Calorimeter Trigger
⇒ 21 Months: 1st prototype
to pp operations

The length of each phase is realistic.

1 year float for final installation (needed 2025)

Example Starting Point: Phase 1 Cal. Trig. Card

In use for Tracklet Trigger, GEM, EB, Correl, Calo, Muon Trigger HL-LHC R&D:

CTP7:

- Xilinx Virtex-7 690T for data processing
- ZYNQ `045 System -on-Chip (SoC) Device (embedded Linux control platform)
- 67 10Gbps optical Input links
- 48 10Gbps optical Output links

JTAG/USB Console Interface Mezzanine

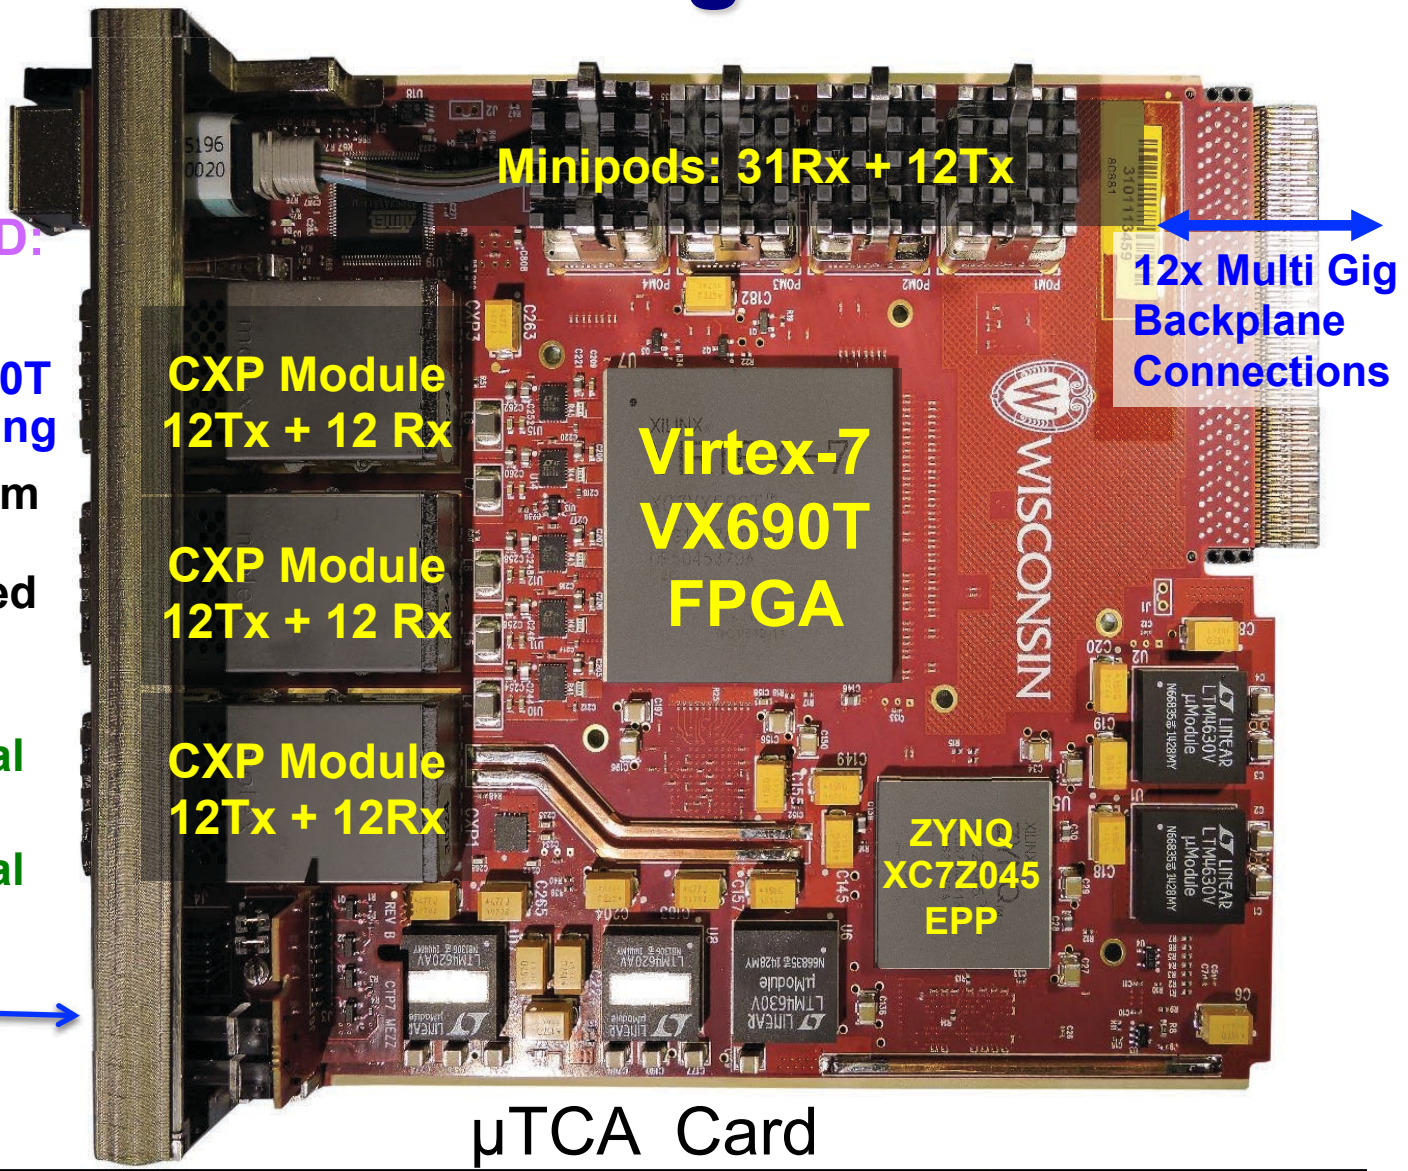

μTCA Card

Example: Hardware Framework R&D: Advanced Processor (APx) Concept

APx Top level diagram: Ultrascale+ processing ZYNQ SoC control, expandable memory, high speed links

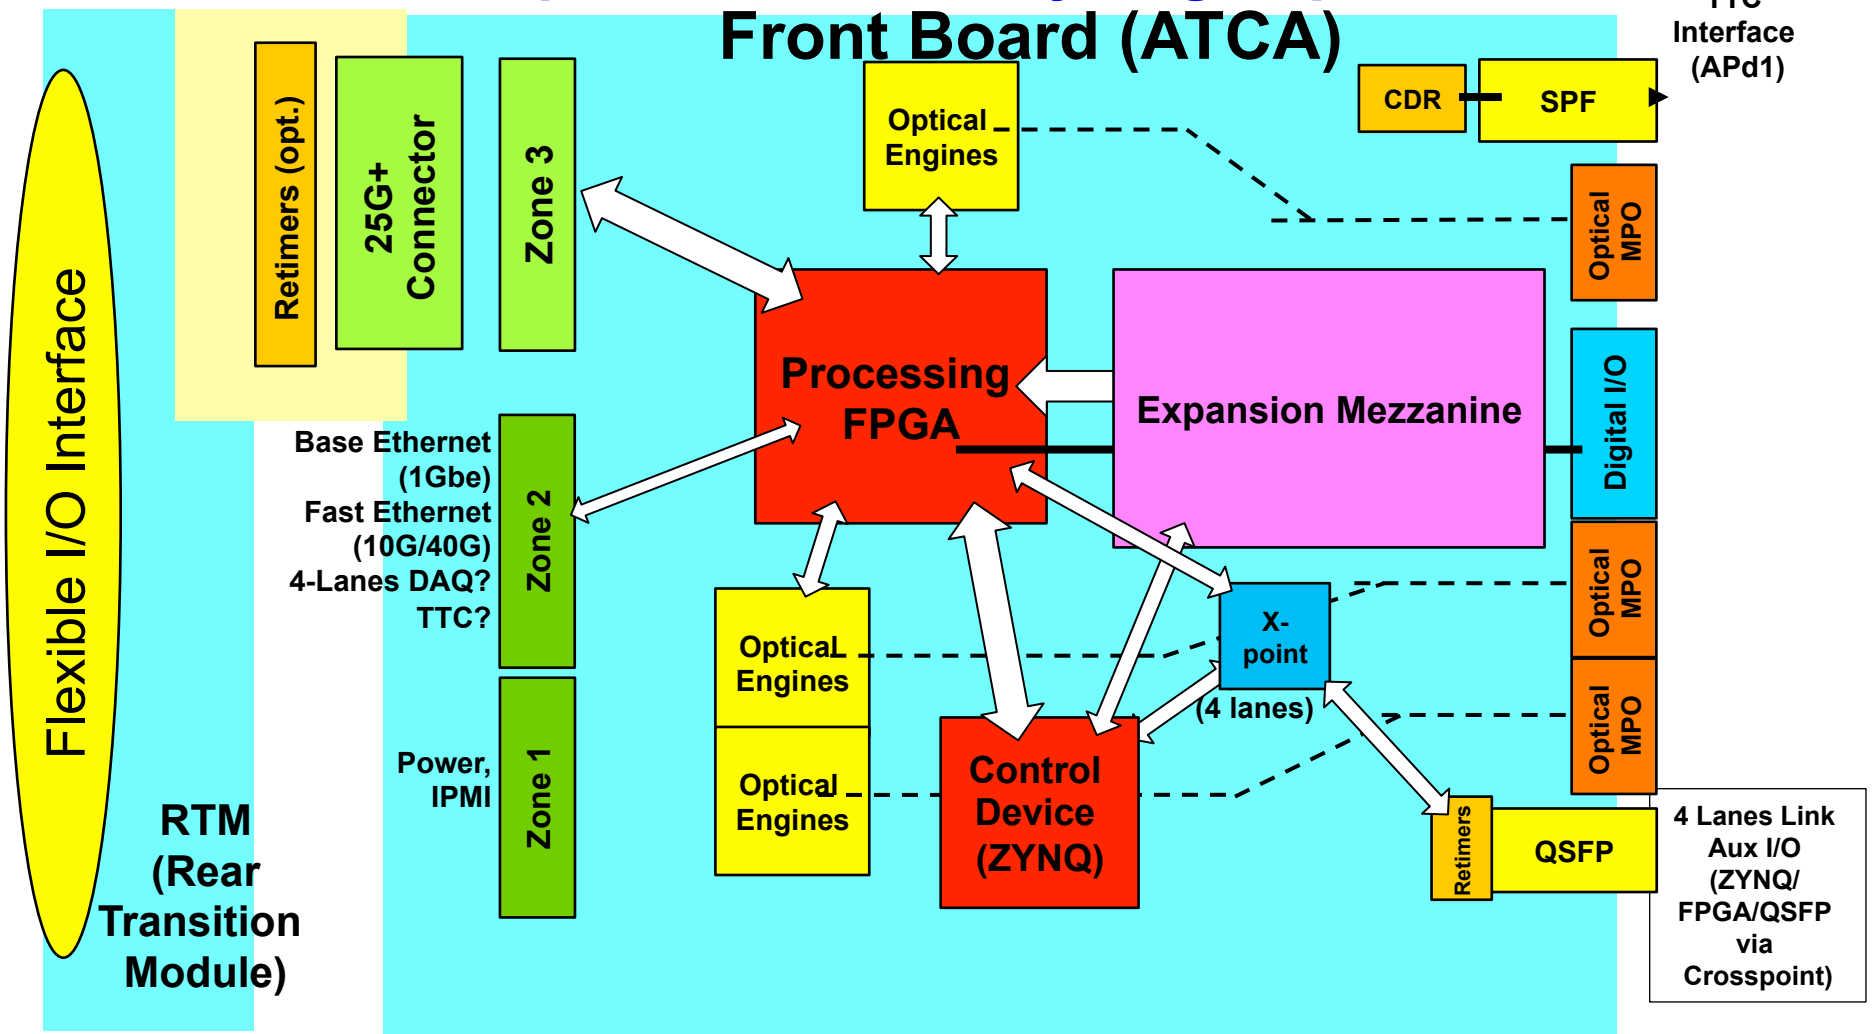

CMS HL-LHC Trigger Summary

Architecture:

- Level-1 Accept: 500 kHz @ 140 PU, 750 kHz @ 200 PU
- Level-1 Latency: 12.5 μsec
- HLT Output rate: 5 kHz @ 140 PU, 7.5 kHz @ 200 PU

Features:

- Level-1 Track Trigger
- Upgraded Calorimeter, Muon, Global Triggers

Performance:

- Track Trigger reduces L1 Menu rate by factor of 5.5
- Track Trigger combined with L1A rate maintains overall physics performance + acceptance of L1 Trigger Phase 1 Upgrade (e.g. thresholds near what are running now)

Schedule:

- R&D Program underway now…important to validate design, techniques & technologies

Backup

HL-LHC ATLAS & CMS Triggered vs. Triggerless Architectures

1 MHz (Triggered) - planned:

- **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) – not planned:

- **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 …**

Tools for Triggers: FPGAs

Example: Xilinx Virtex 7 (28 nm), Ultrascale (20 nm), Ultrascale + (16 nm)

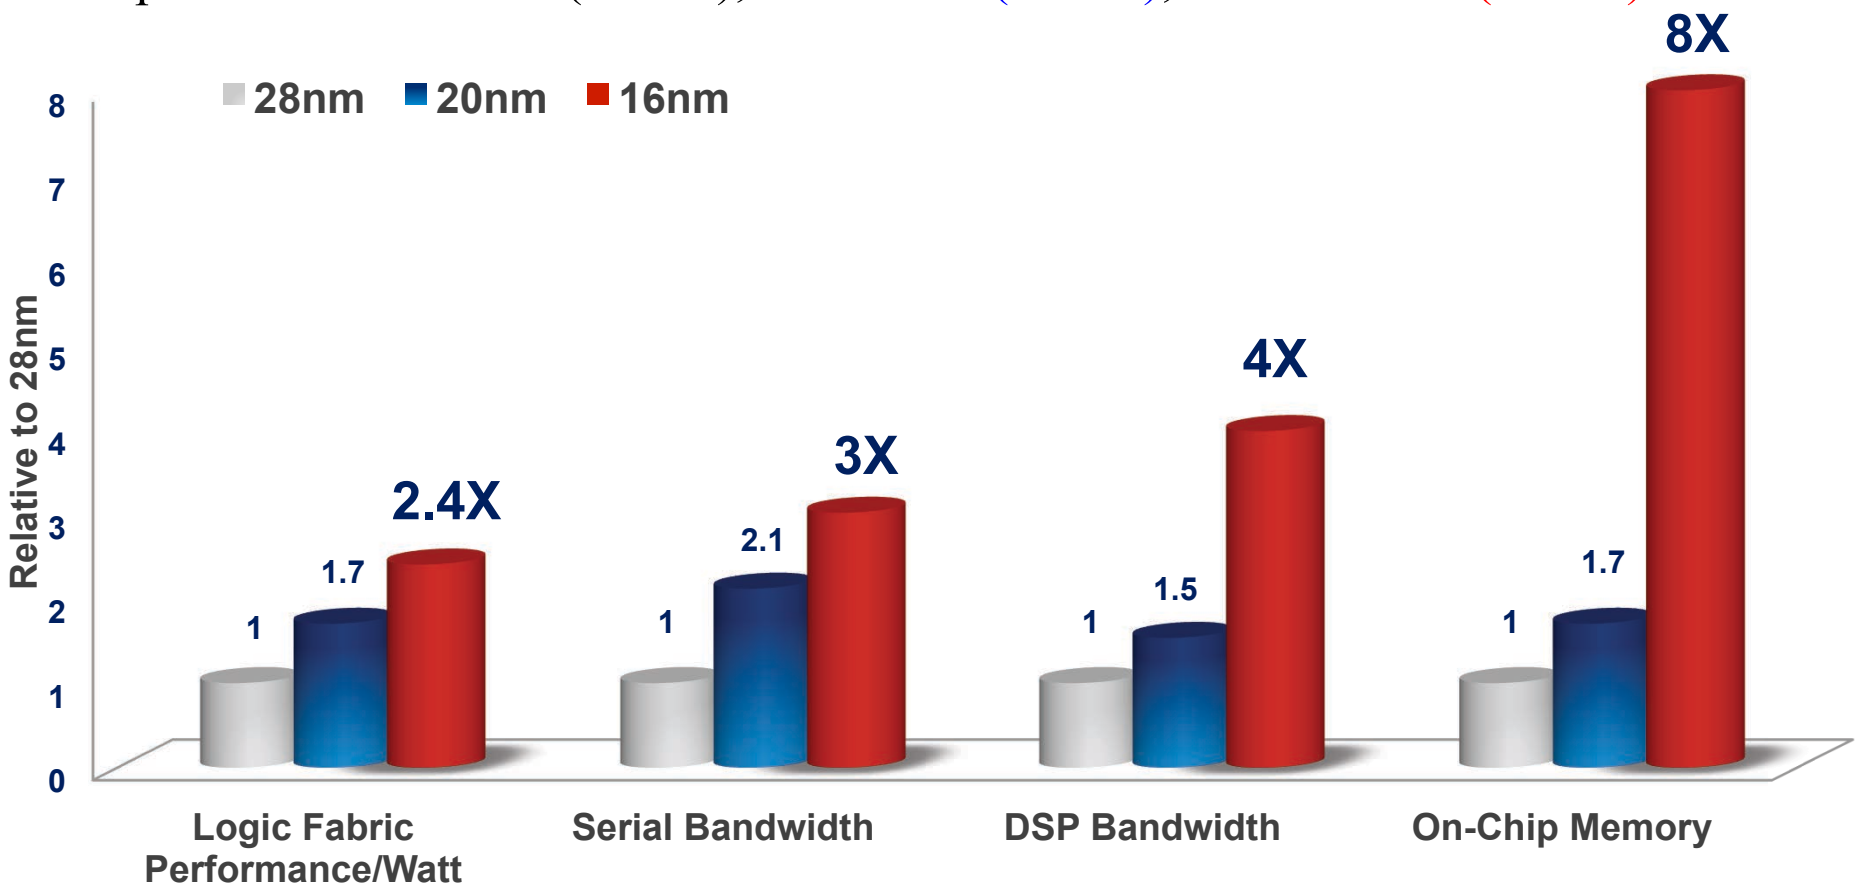

Enhanced Fabric with FinFET performance	Up to 128 transceivers at up to 32.75 Gb/s	~12,000 DSP slices running at ~900 MHz	UltraRAM for SRAM device replacement

FPGA Examples: Xilinx devices

	KINTEX-7	KINTEX UltraSCALE	VIRTEX-7	VIRTEX UltraSCALE
Logic Cells (LC)	478	1,161	1,995	4,407
Block RAM (BRAM) (Mbits)	34	76	68	132
DSP-48	1,920	5,520	3,600	2,880
Peak DSP Performance (GMACs)	2,845	8,180	5,335	4,268
Transceiver Count	32	64	96	104
Peak Transceiver Line Rate (Gb/s)	12.5	16.3	28.05	30.5
Peak Transceiver Bandwidth (Gb/s)	800	2,086	2,784	5,886
PCI Express Blocks	1	6	4	6
Memory Interface Performance (Mb/s)	1,866	2,400	1,866	2,400
I/O Pins	500	832	1,200	1,456

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

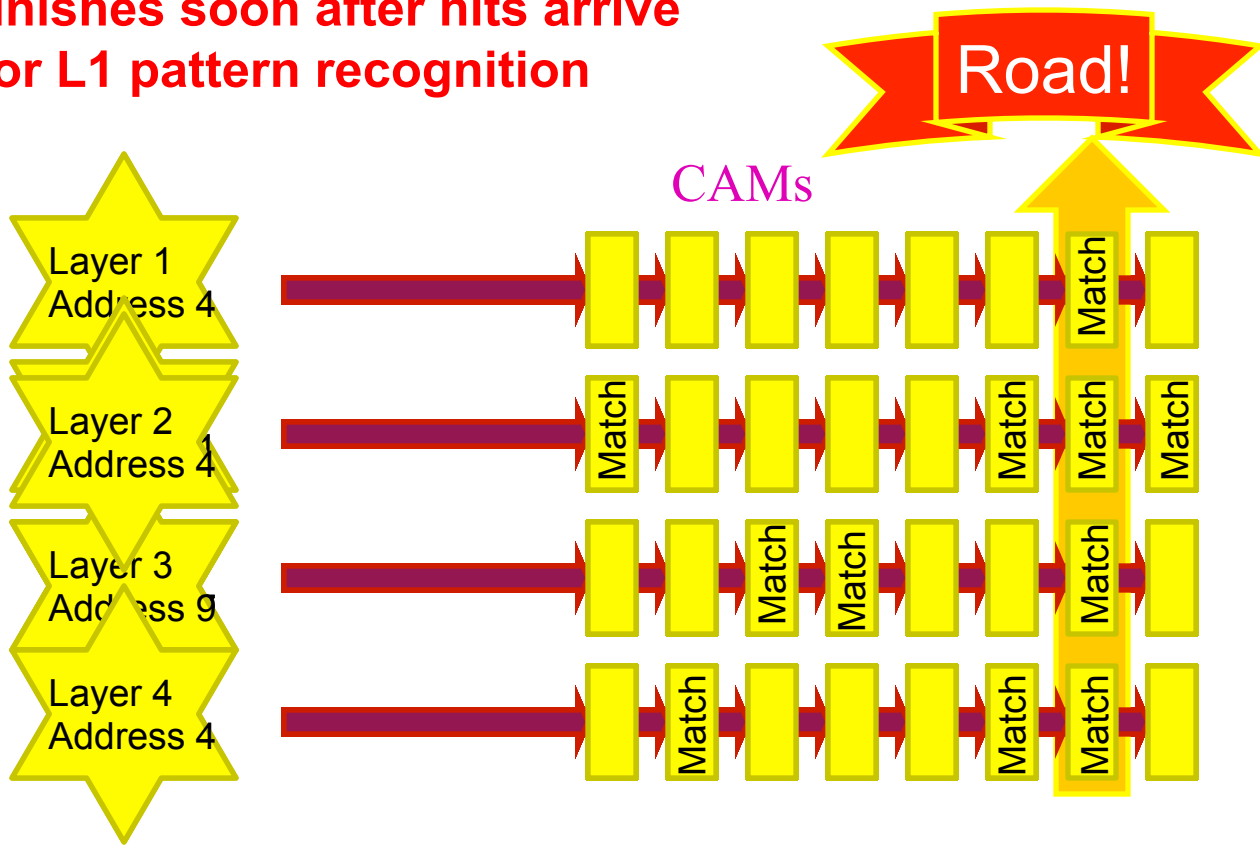

Tools for Trigger/DAQ: ATCA

- Advanced Telecommunications Computing Architecture
- Example: Pulsar Card (FNAL, UIC, Northwestern) for CMS Track Trigger
  - Use FPGAs for low latency
  - FPGAs are directly connected to the full mesh fabric channels
  - No network switch
  - Low overhead serial protocols
  - High bandwidth I/O via serial links on Rear Transition Module and mezzanines

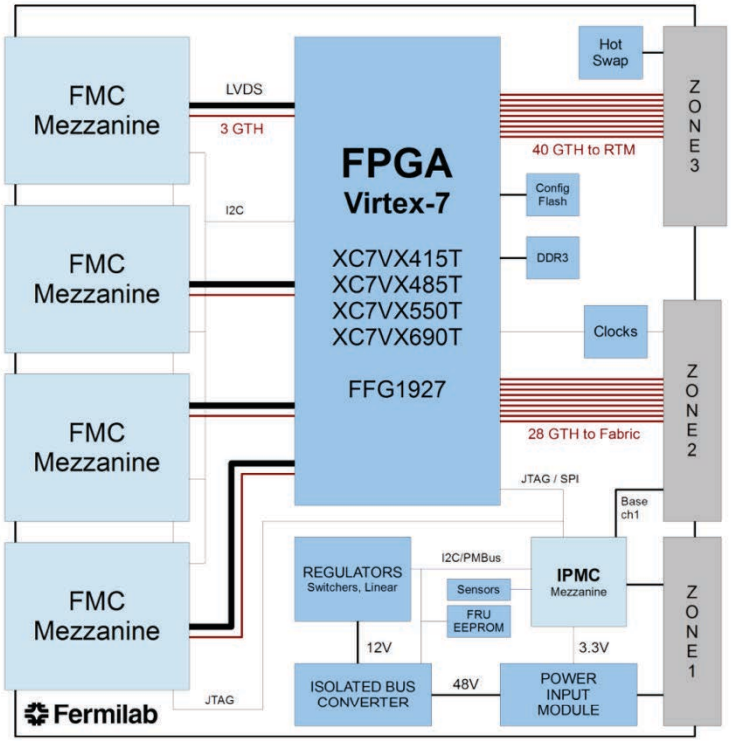

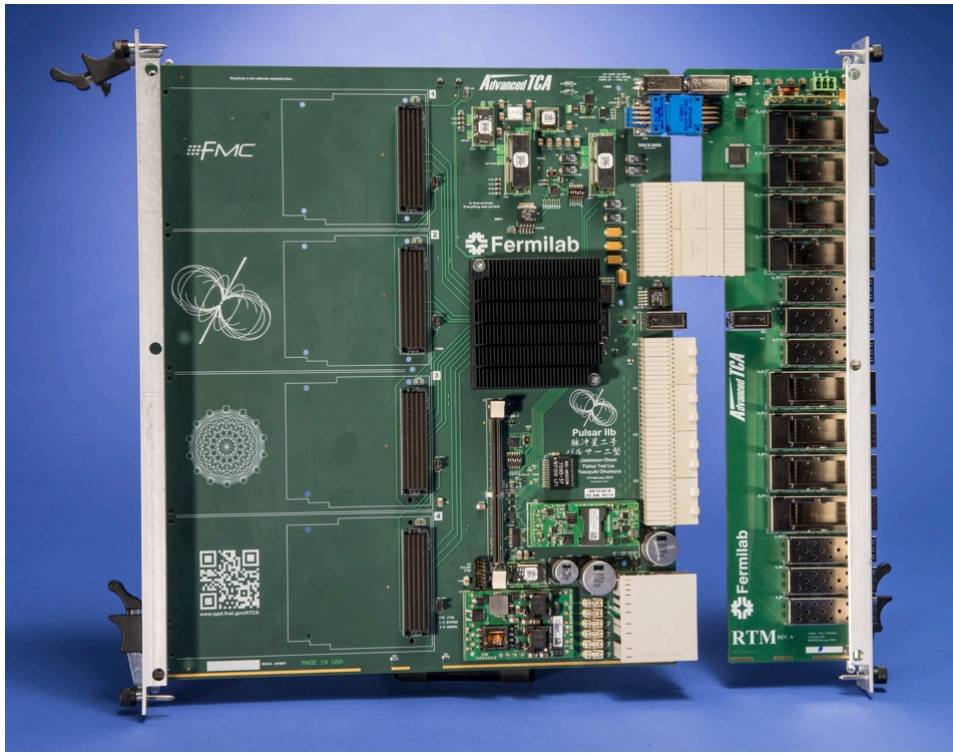

ATCA Backplane example: Pulsar IIb

- Full shelf tests with all lanes running at 10 Gbps
  - BER = 2×10^{-16}
- Evaluating latest high performance 40G+ full mesh backplanes from ASIS-PRO, COMTEL, and Pentair/Schroff

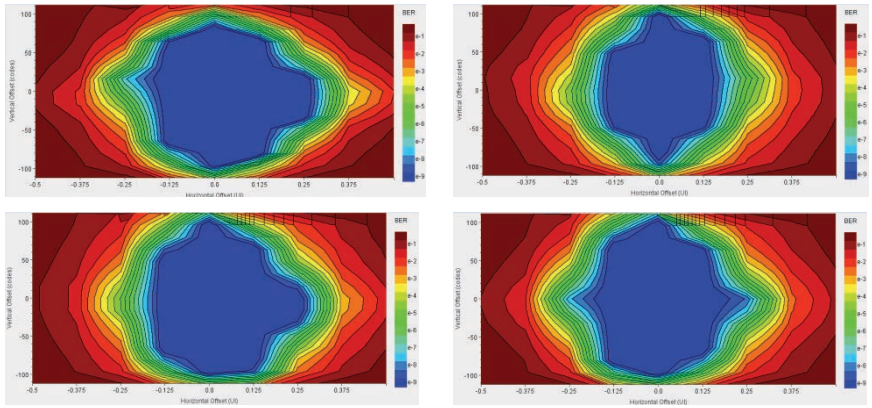

HL-LHC Trigger Technologies (under test with CTP7)

Embedded Linux

- Functional Linux system (network, file system, shell)
- Low latency access point tightly integrated with workhorse FPGA
- Basic card level infrastructure with very little new code—Ethernet, I2C, USART, GPIO drivers, ssh, file system, etc.—all standard
- Paid for itself in time saved in the first project cycle
  - First CTP7 proto. power-up to integrated operation in CMS pp runs in 21 months

AXI Architecture

- Industry standard on-chip interconnection scheme for FPGAs
- Straightforward to implement AXI interfaces for registers and memory
- AXI infrastructure bridged into the Virtex-7 ("Chip2Chip" core), a single integrated address space for both devices
- 95% of CTP7 generic infrastructure from ZYNQ hard cores and Library IP catalog, no custom HDL needed—it's in the tools
- Improved status and access fo advances applications
  - Real-time link eye-diagrams for all channels while taking data available online!

XVC – Embedded Linux Xilinx Virtual Cable (e.g. JTAG)

- Debug Card at P5 via TCP/IP just as if on the bench in the lab

HL-LHC R&D Technology Examples

Link monitoring/tuning using embedded Linux (ZYNQ):

Monitor individual link performance by running continuous eye-scans while taking data.

- Used to predict BER and proactively detect weak links before data transmission errors appear – manage large #'s of links

Automatically and quickly tune parameters to optimize link performance while link is operating

- Allows use of higher bandwidth links and moving boards while minimizing performance differences

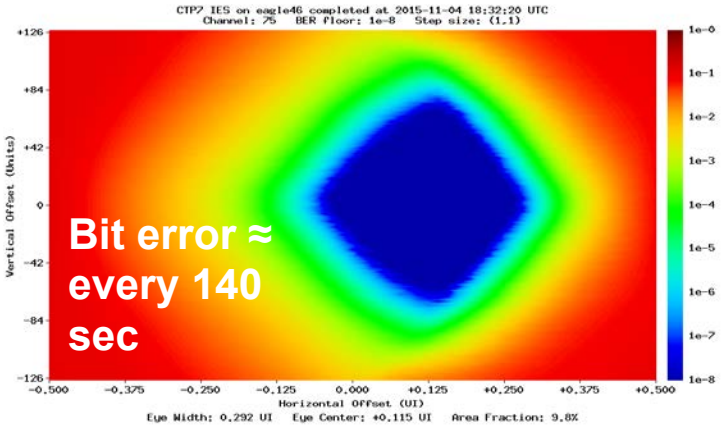

CTP7 using DFE* program on ZYNQ ➔

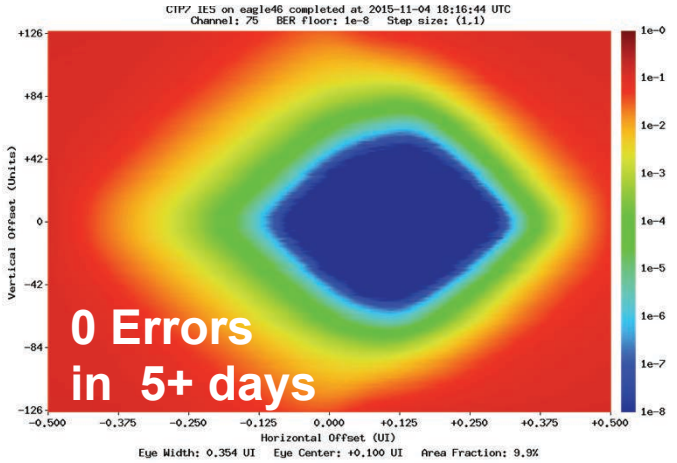

*Xilinx DFE = Decision Feedback Equalization plus Cont. Time Linear Equalization