Future Trigger & DAQ

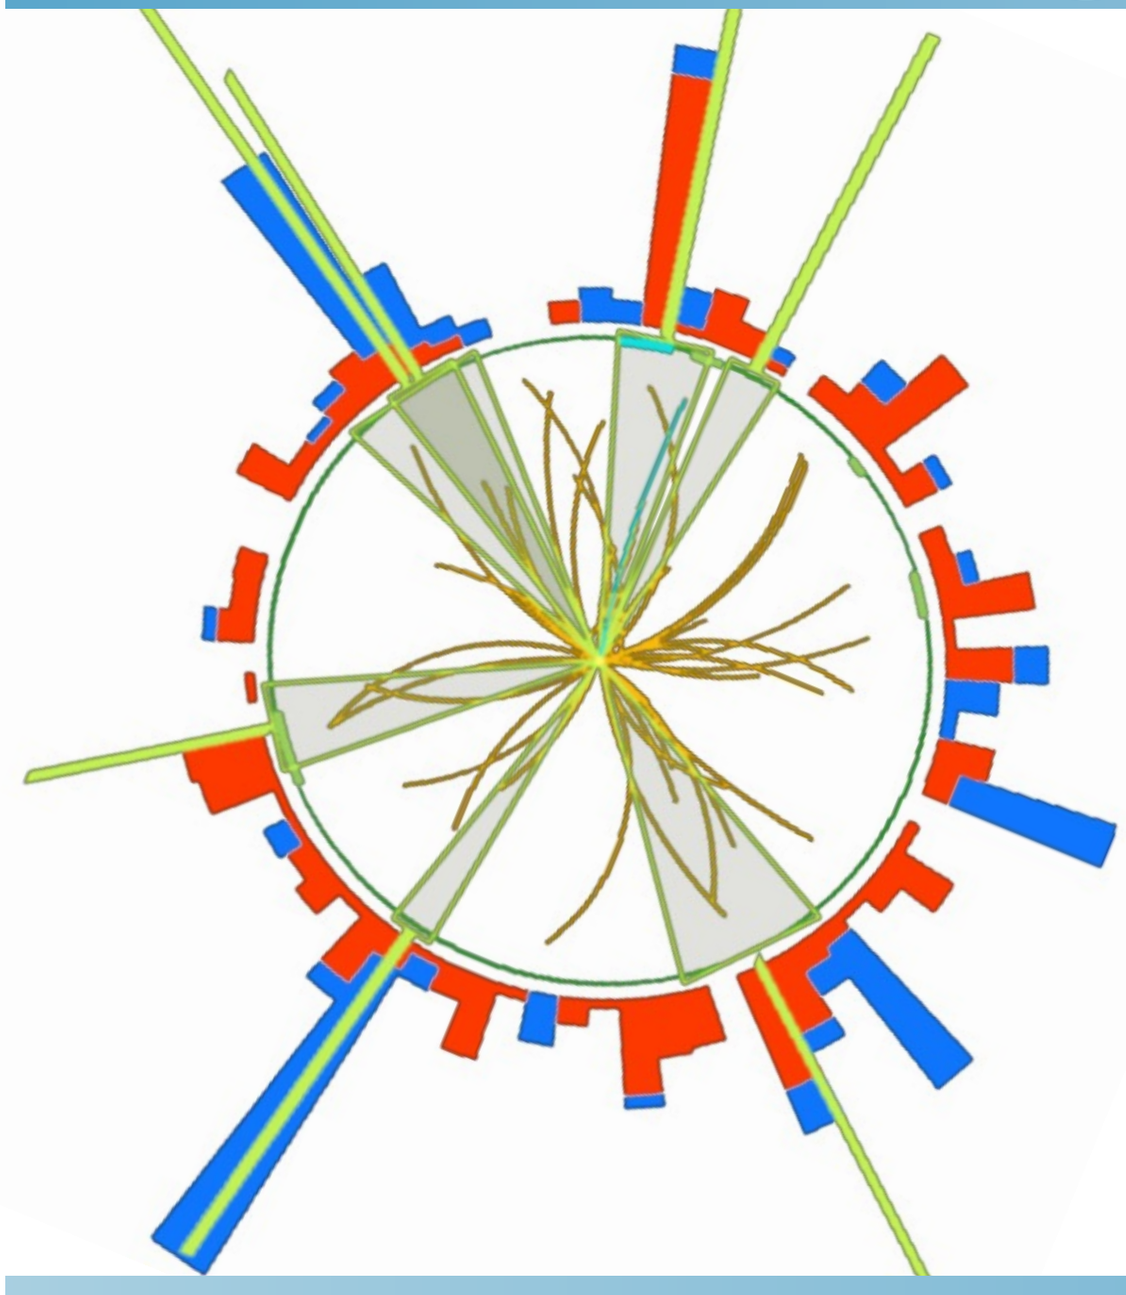

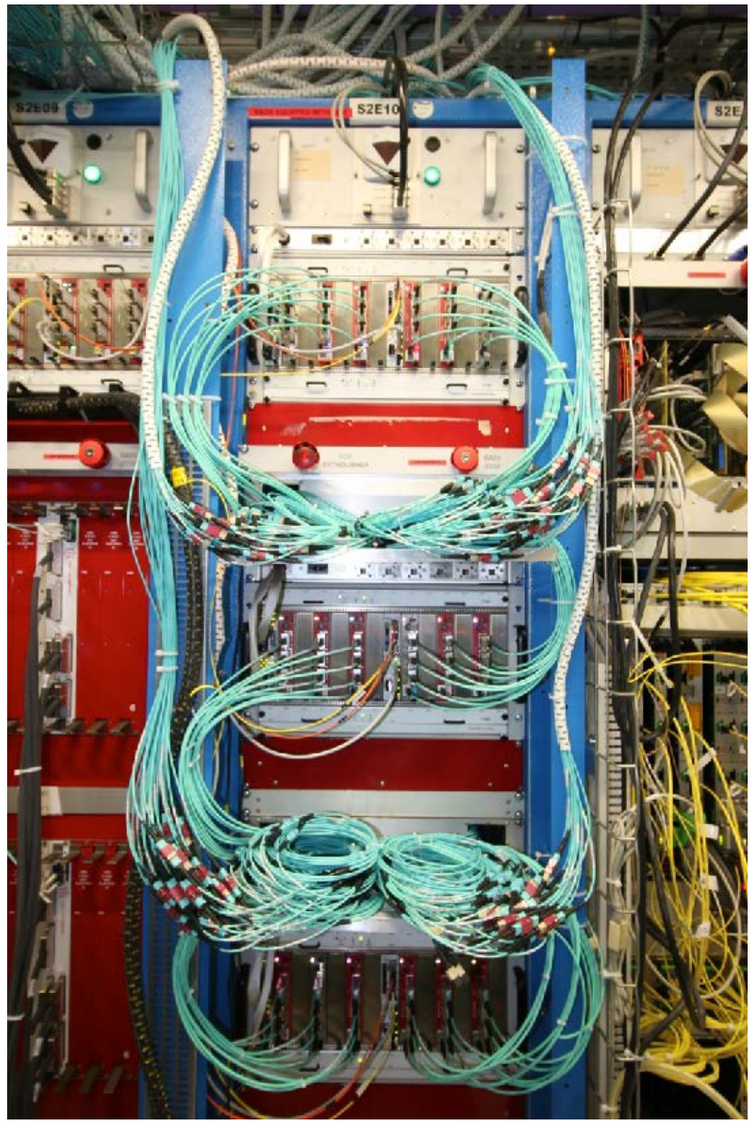

CMS Phase-1 Layer-1 Calo Trigger (U. Wisconsin)

Dave Newbold

U. Bristol / Rutherford Laboratory

Thanks to: P. Bloch, D. Cussans, M. Kocian, A. Marchioro, M. Raymond

Plan

Q: 'Can we build an FCC-hh GPD with triggerless readout?'

Plan

Q: 'Can we build an FCC-hh GPD with triggerless readout?'

- A: 'No'
  - Recording the full data stream is inconceivable even in 203x
  - But 'trigger' may mean something different to today
- Plan of talk
  - The task at hand
  - GPD trigger evolution
  - Implications for future systems
  - Technology issues
  - Next steps

- Some pre-apologies
  - This talk concerns FCC-hh primarily; it is the hard problem for TDAQ
  - More trigger more than DAQ
  - I extrapolate mainly from CMS; ATLAS performance & strategy is similar

Trigger Requirements

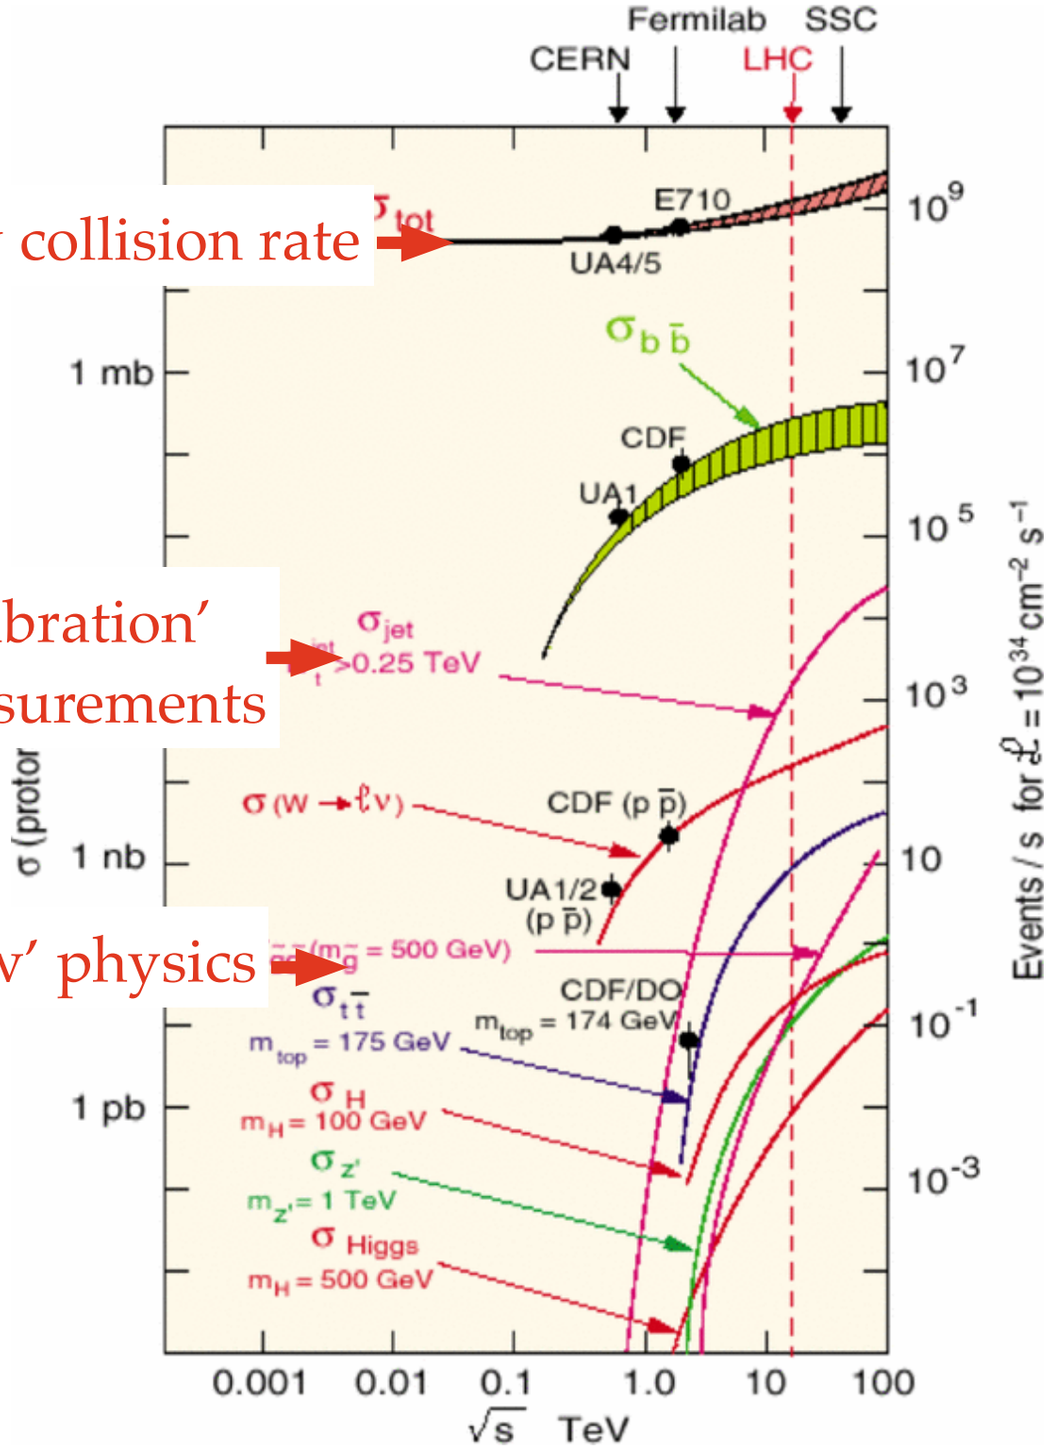

- A GPD trigger must
  - Reduce recorded event rate by factor of ~200,000
    - i.e. reduce detector readout bandwidth and storage requirements
  - Conform to stringent technical requirements
  - Use limited subset of detector data
- Name of the trigger game:
  - Pushing edge of technology
  - Continuous evolution / upgrade
  - Guaranteed reliability
  - Quantifiable, measured efficiency
- Physics reach is ultimately dictated by the trigger

LHC Trigger Strategy: Recap

- We are working at the electroweak scale and above
  - Need good sensitivity for decays of objects of mass O(100GeV)
- Basic trigger objects
  - Leptons – with or without isolation criteria
    - Isolation suffers badly as pileup increases
  - Photons – hard to distinguish from electrons without tracking
  - Jets of sufficiently high E_t; hadronic tau, typically with poor efficiency in Run 1
  - Global event variables: total E_t, missing E_t, H_t, etc
- Backgrounds
  - Real leptons from heavy flavour production, typically soft
  - Mis-measured jets; photons in jets; fake electrons from conversions
  - Punch-through in muon system
- Making the decision
  - Require combinations of trigger objects above a set of E_t thresholds
  - More clever combinations are possible with sufficient trigger resolution
    - Event topology cuts, invariant mass, etc

CMS L1 Trigger Evolution: Run I

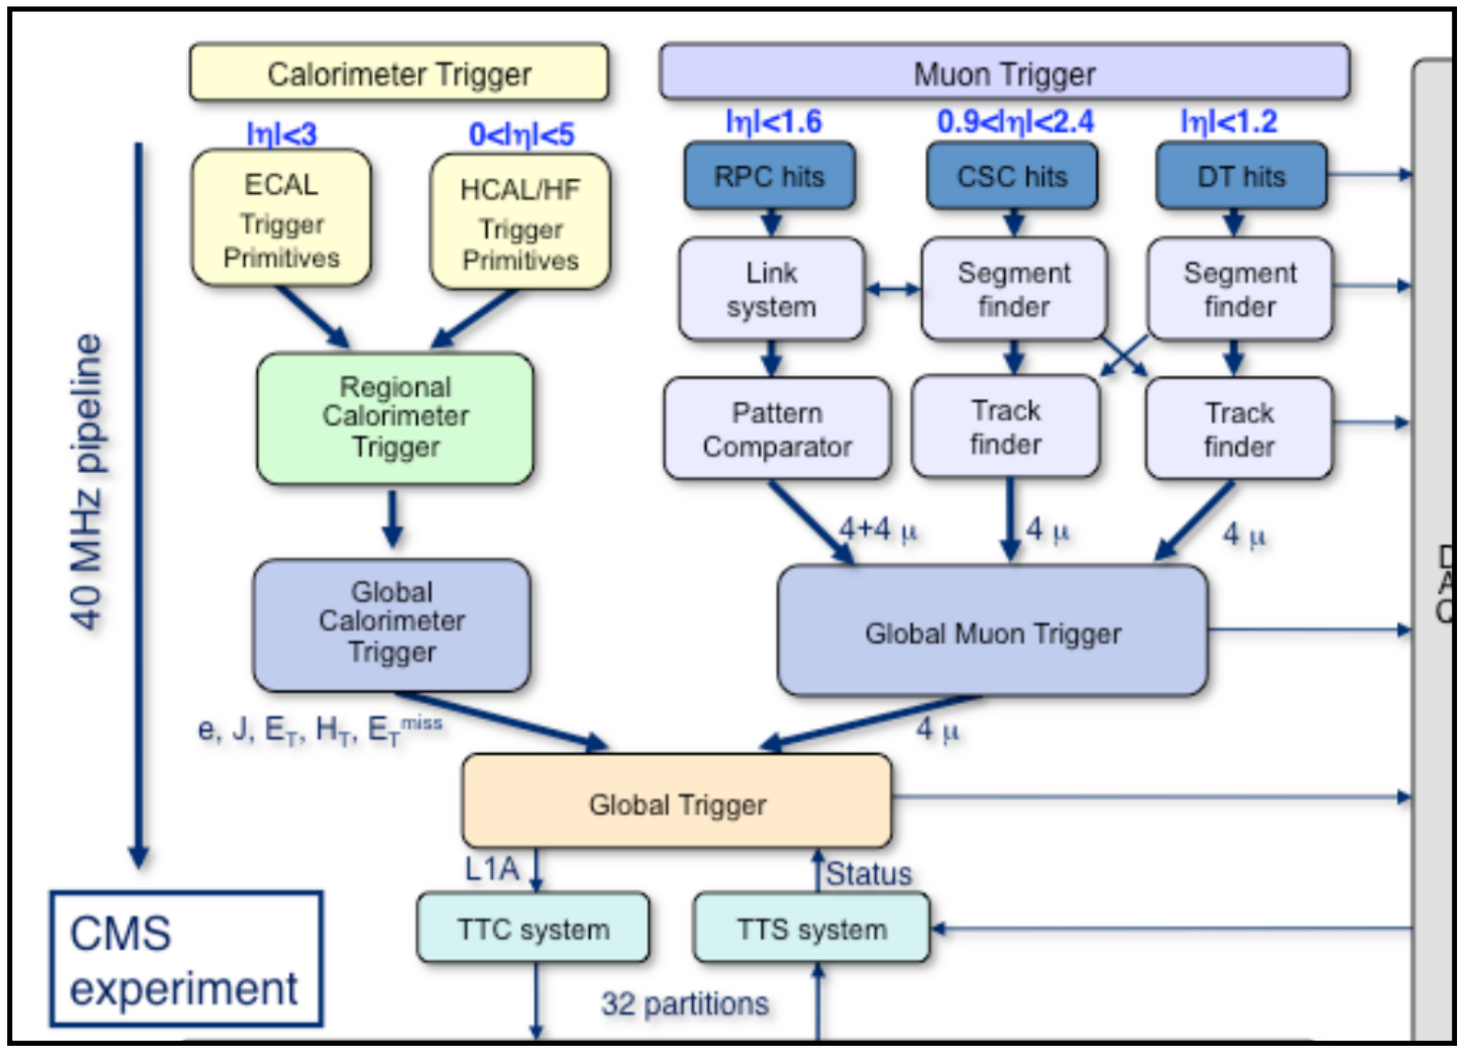

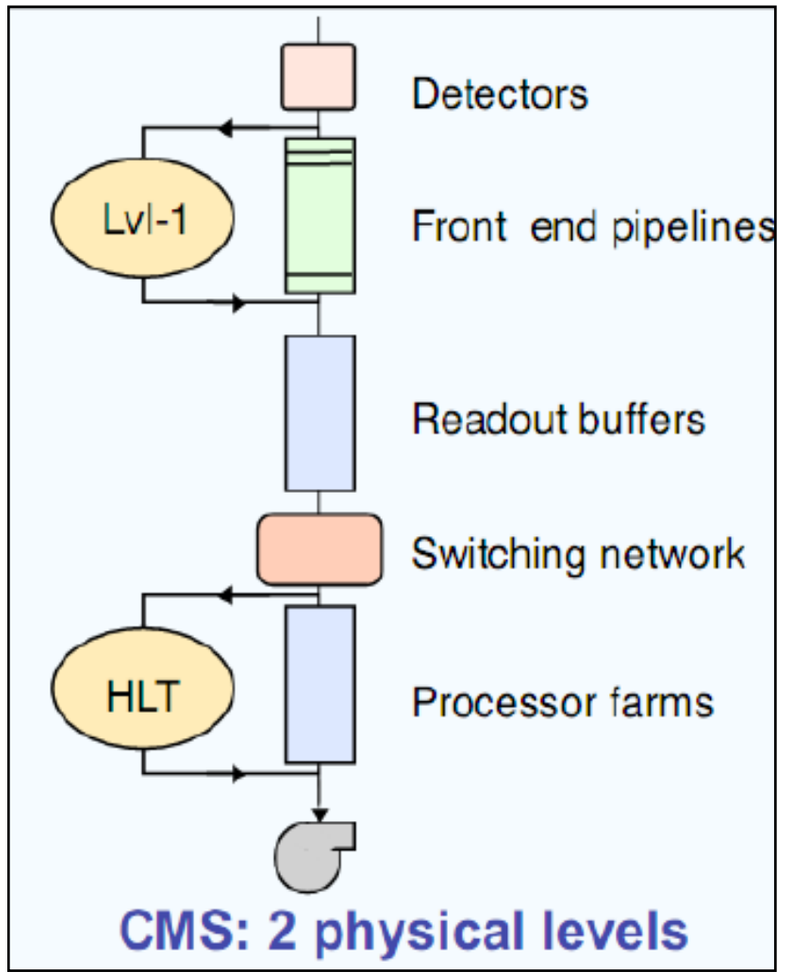

- Performance of L1 (hardware) trigger dictated by:
  - Processing capacity of data reduction logic
  - Affordable link bandwidth from detector
- L1 decision based on number and type of locally-found 'trigger objects'

CMS L1 Trigger Evolution: Run II

~10Tb/s 'event builder network' for calo trigger

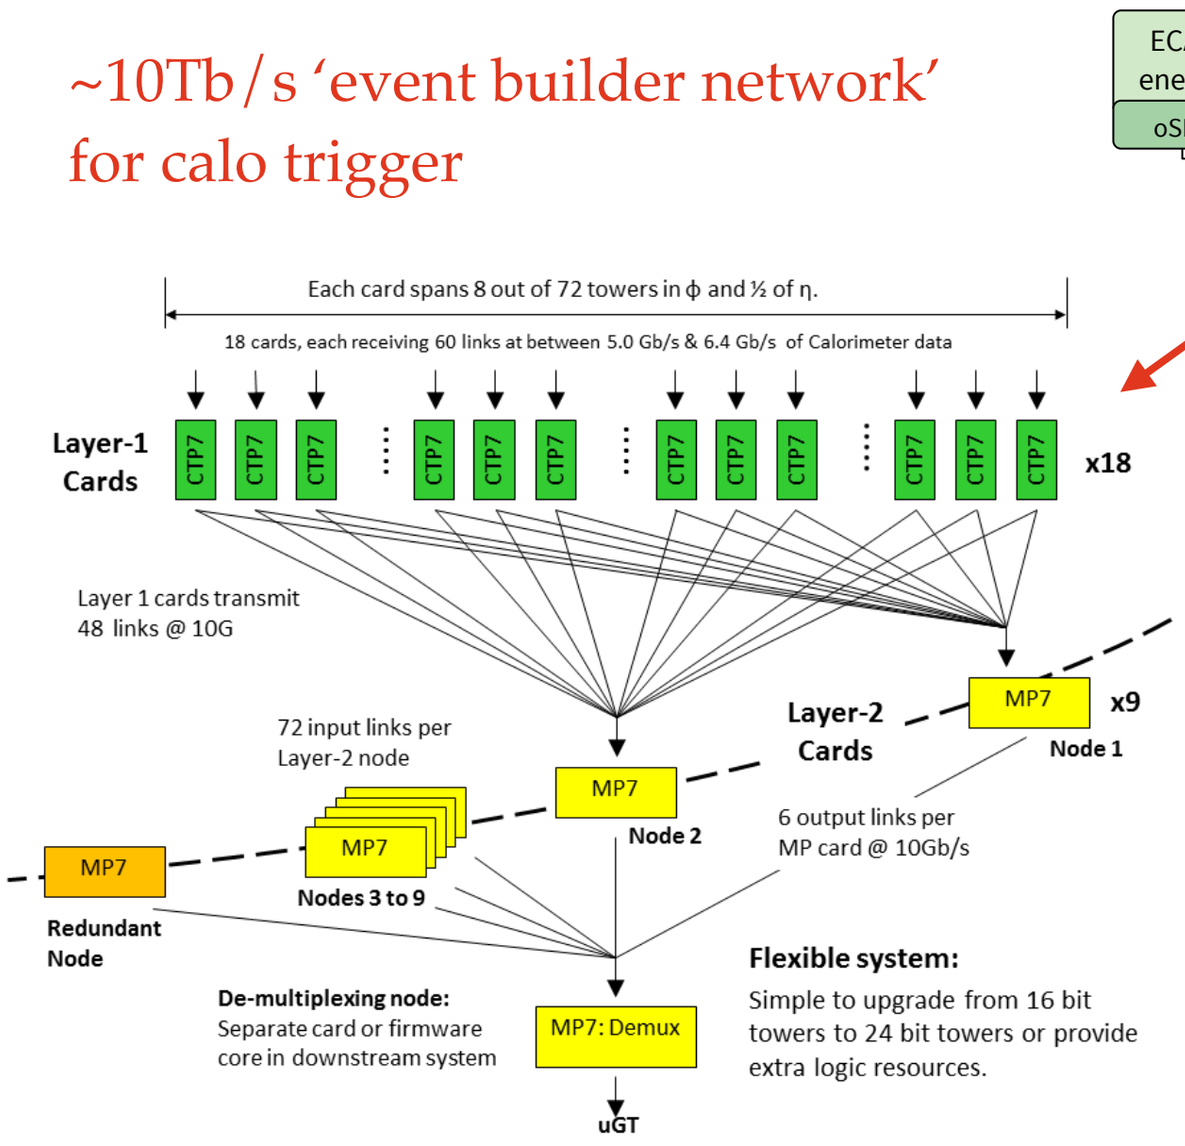

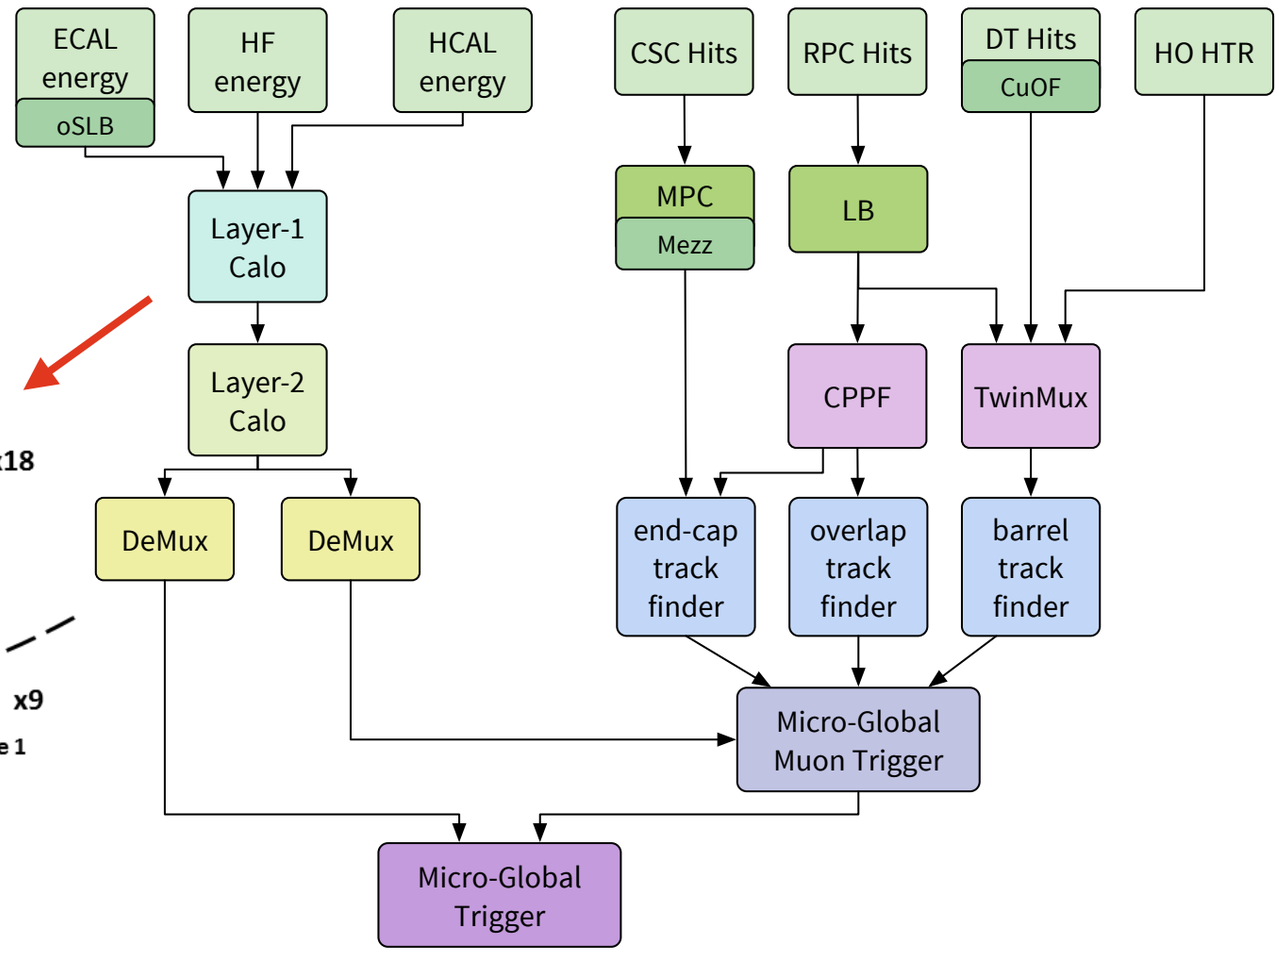

- Running in CMS right now; three years development
- Key feature: 'time multiplexing' of *entire* events into single processor
  - Allows more efficient trigger object ID, some redundancy

CMS L1 Trigger Roadmap

	Run 1 (2009)	Run 2 (2015)	Run 3 (2020)	HL-LHC (2025)
ECAL / HCAL granularity	Regions / Regions	Towers / Towers	Towers / Towers	Crystals / Towers
Detector information	Calo + muon	Enhanced calo / unganged muons	+ additional muon coverage	+ inner tracking
L1 Trigger rate	100kHz	100kHz	?	1MHz
GT algorithms	Cut and count + topological	+ Invariant mass	?	Particle flow, track isolation

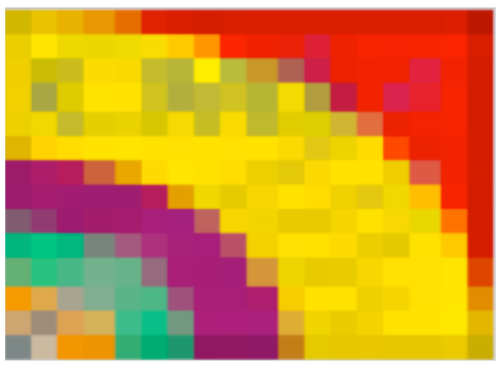

Region Granularity

Trigger Tower Granularity

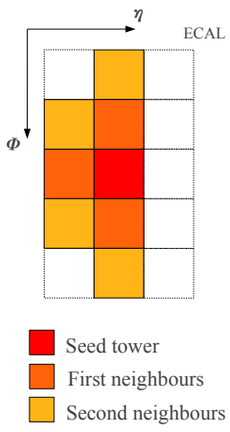

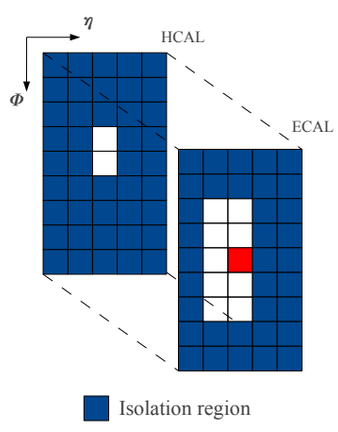

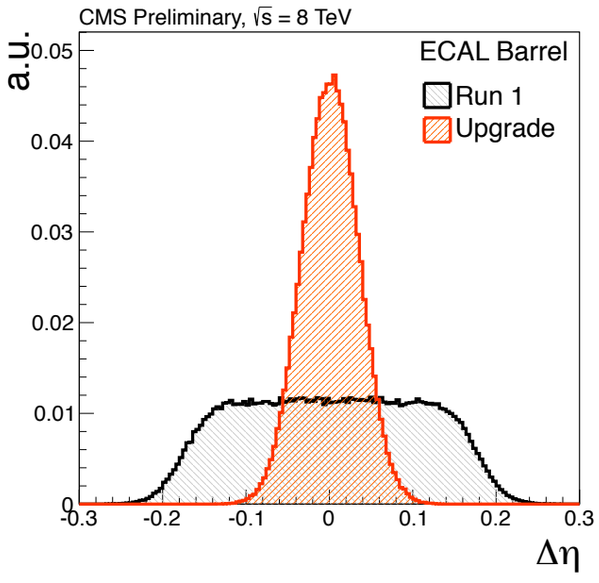

CMS L1 Trigger Evolution: HL-LHC

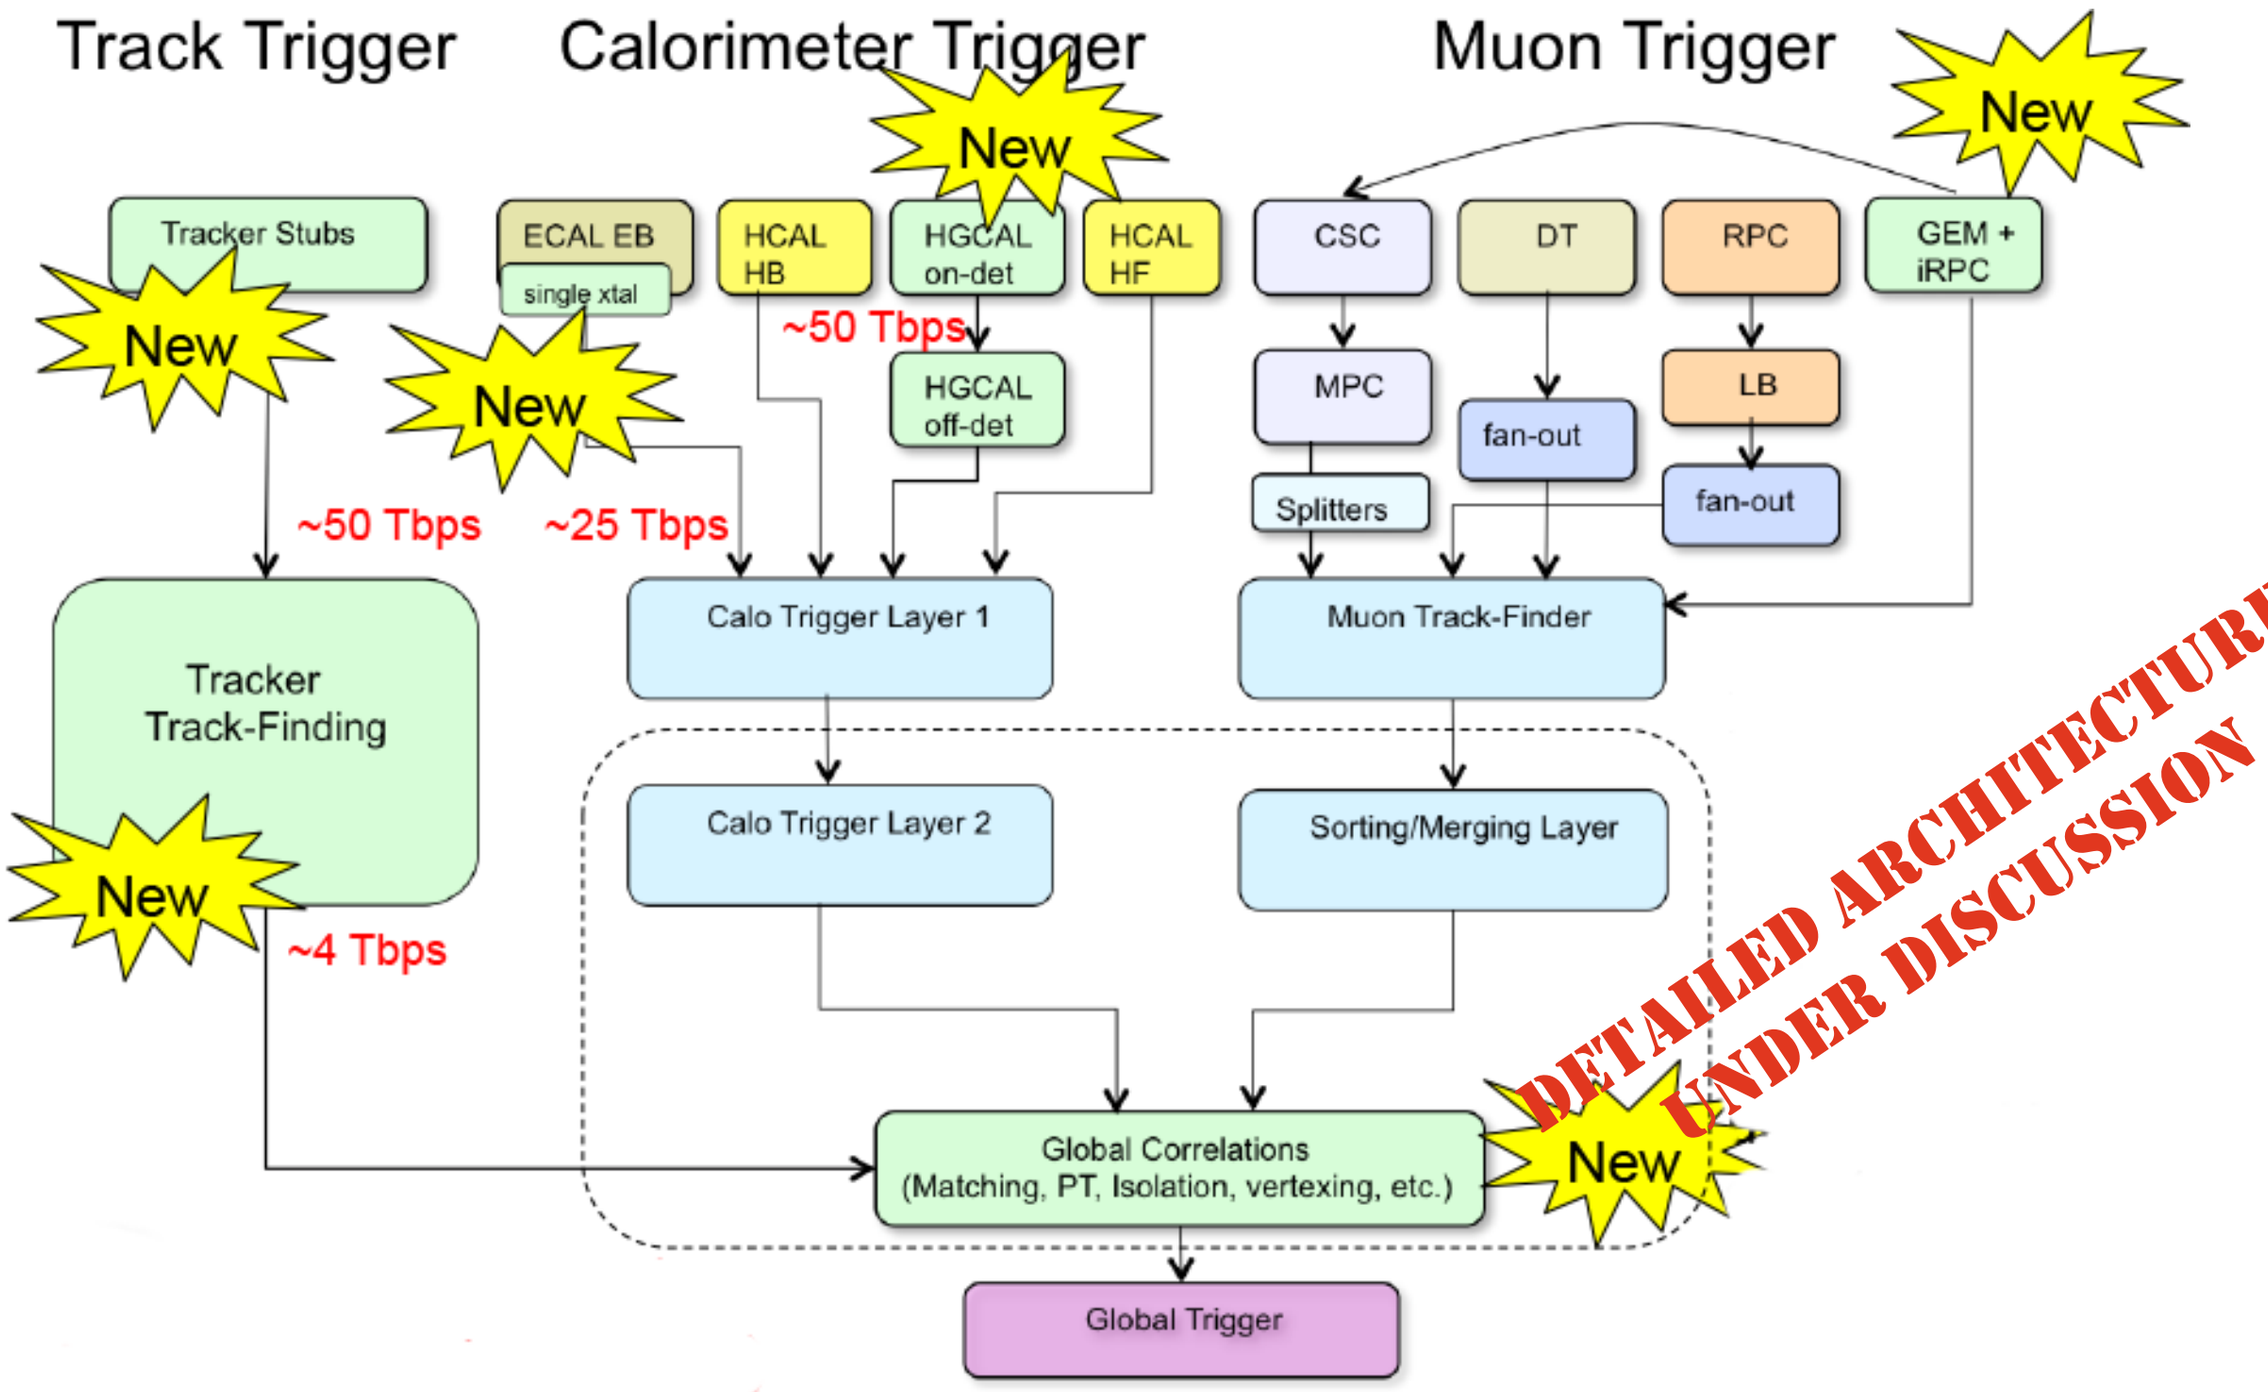

- Message: hardware L1 is already becoming more HLT-like

L1 Tracking Concept

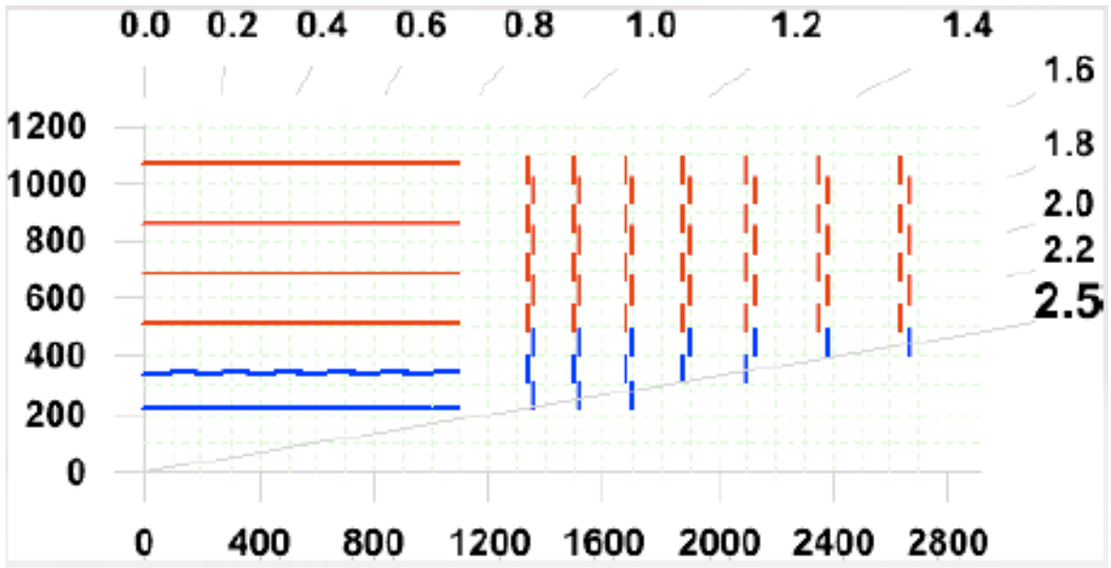

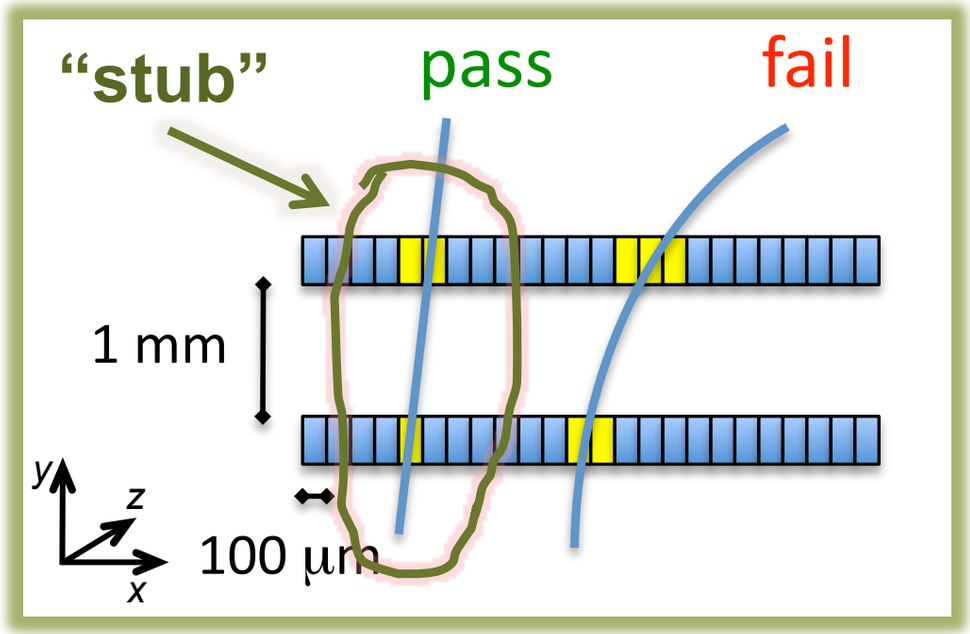

- Upgrade tracker design driven by triggering requirements
  - Full detector readout at 40MHz technologically impossible today
    - Power and mass considerations
  - Geometry of tracker design to provide in situ data reduction
  - ‘Stub’ primitives read out at full bandwidth to off-detector track finder
- Relevance for FCC-hh
  - Detector design must take into account trigger requirements from the start
  - This applies even to ‘software trigger’ and reconstruction

Where does the Bandwidth go?

Rate division between broad trigger classes for Phase-2 (total 180kHz)

Full menu is complex (>128 entries)

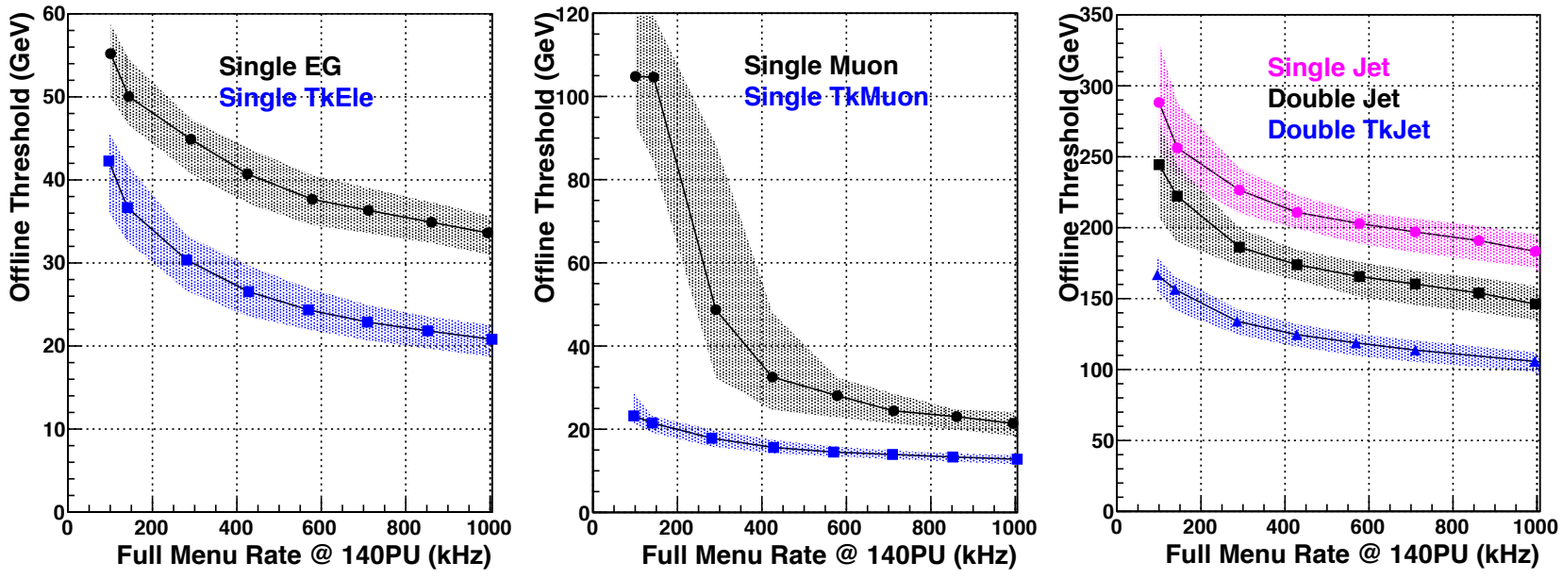

- Rate of triggers does *not* indicate their 'importance for physics'
  - Rather, represents the difficulty of using different objects

High Level Trigger

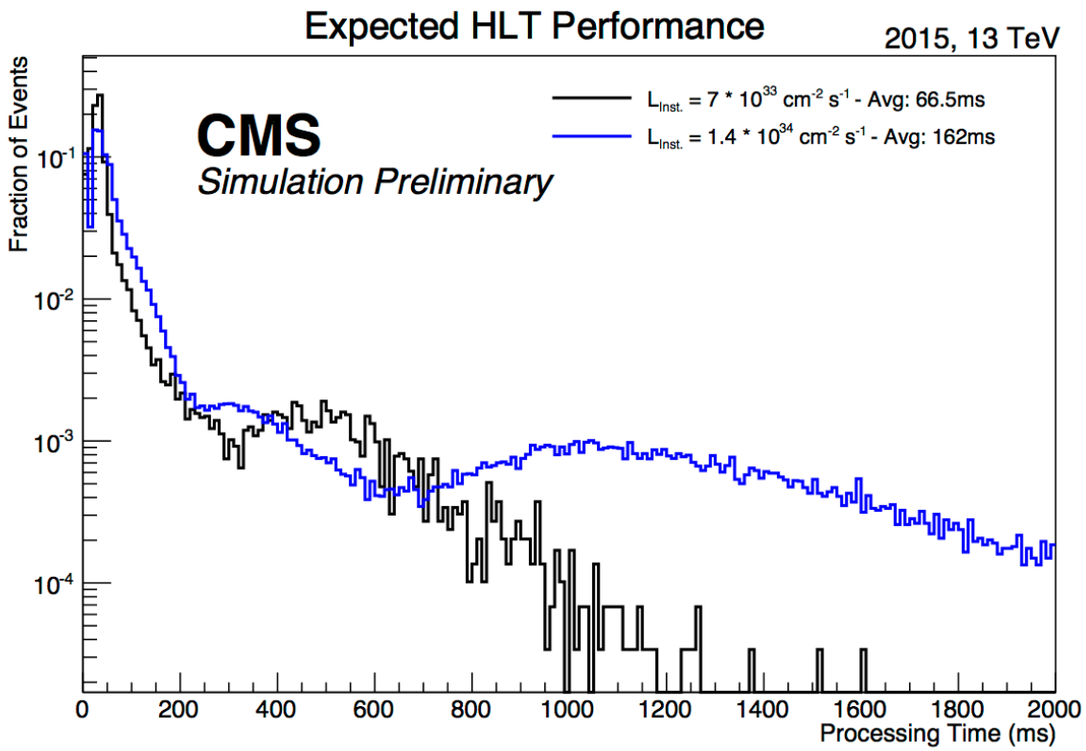

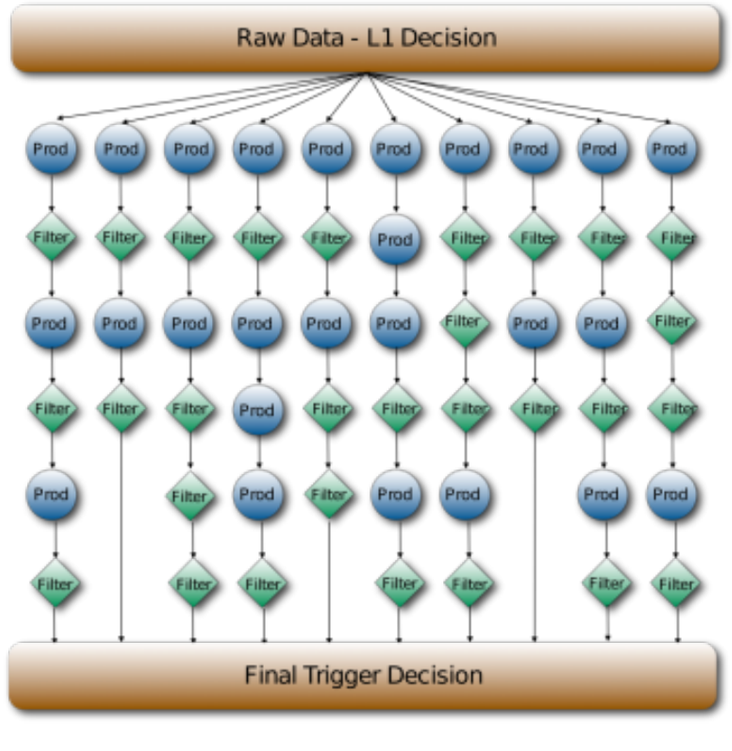

- Multiple higher trigger levels are implemented in software
  - More 'funnel' than 'filter' – implemented on CPUs (~2MW power by 2025)
  - ATLAS has an additional L2 stage using full granularity in RoI
- Note that few events are 'fully reconstructed'
  - Throw everything away as fast as possible; 'no trigger' is not an option
  - Unpacking and basic manipulation of raw data is a large overhead
- Future HLT upgrades will involve co-processing of various types

Triggering at FCC-hh

- Higher CoM energy
  - Resolution for high pt leptons
    → high-precision primitives
  - Boosted final states → calo reco hard
  - High occupancy in forward region
    → better granularity
- Higher luminosity (25ns scenario)
  - Vertex count / merging → vertex tagging
  - Calo-based isolation ineffective → track-based isolation
  - Poor calo-based MET / MHT resolution → track-calo correlation
- Higher luminosity (5ns scenario)
  - Bunch crossing ID becomes challenging
    - Multiple showers in the calorimetry simultaneously; enormous pile-up in tracking; TOF effects
  - Precision timing important for trigger → more data
- All of these effects go in the same direction
  - More resolution; more granularity; more data; more processing

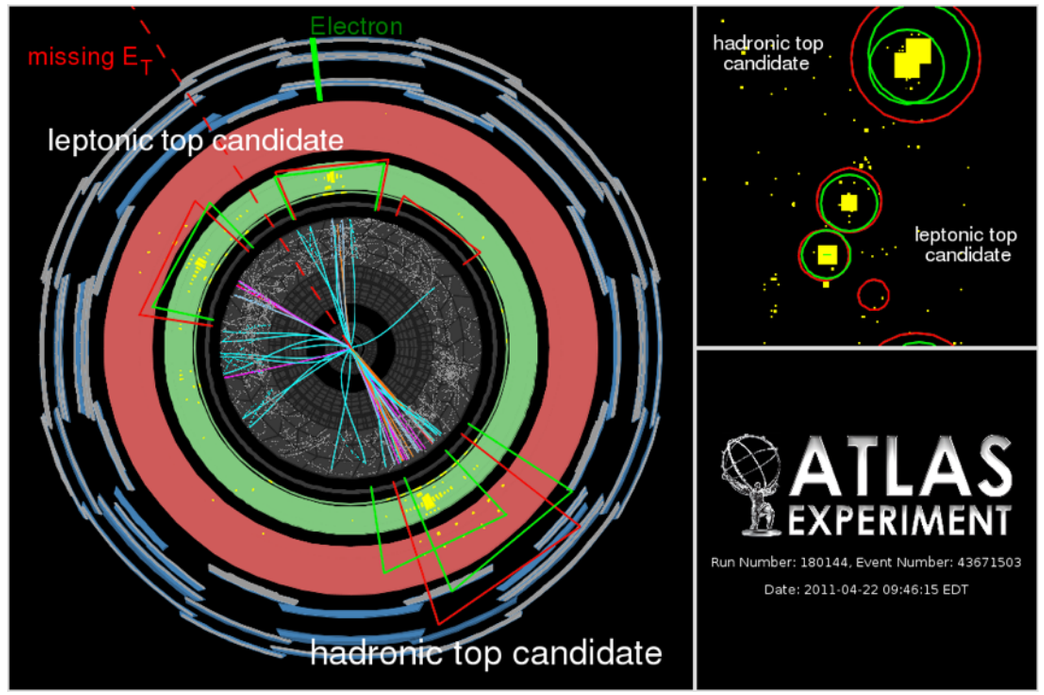

Implications

- For any trigger / filter system:
  - Trigger to be based on all sub-detectors (incl. inner tracking)
  - Detector modularity / alignment should allow for fast basic reconstruction
  - All the data for a crossing needs to be promptly in one place
    - This has significant implications for detectors which are 'slow' wrt BX frequency
- Is a 'traditional' pipelined hardware trigger feasible?
  - Today's LHC Run-II systems are (intentionally) pushing the envelope
    - Technology marches on; FPGAs are densest, lowest power, processing technology available
    - However, the costs are becoming significant, and modularity is not always suitable
  - Extrapolation from CMS / ATLAS does seem possible, but hard and costly
  - Key problem: data flow management into processors
    - 'Getting all the data into one place' requires significant on-detector processing
- All considerations motivate the use of a software-based trigger
  - Flexibility for physics; increase capacity with luminosity; commodity hardware
  - Note that 'software' does not imply 'general purpose Intel CPU'
  - Is this feasible?

‘Triggerless’ Approach – LHCb

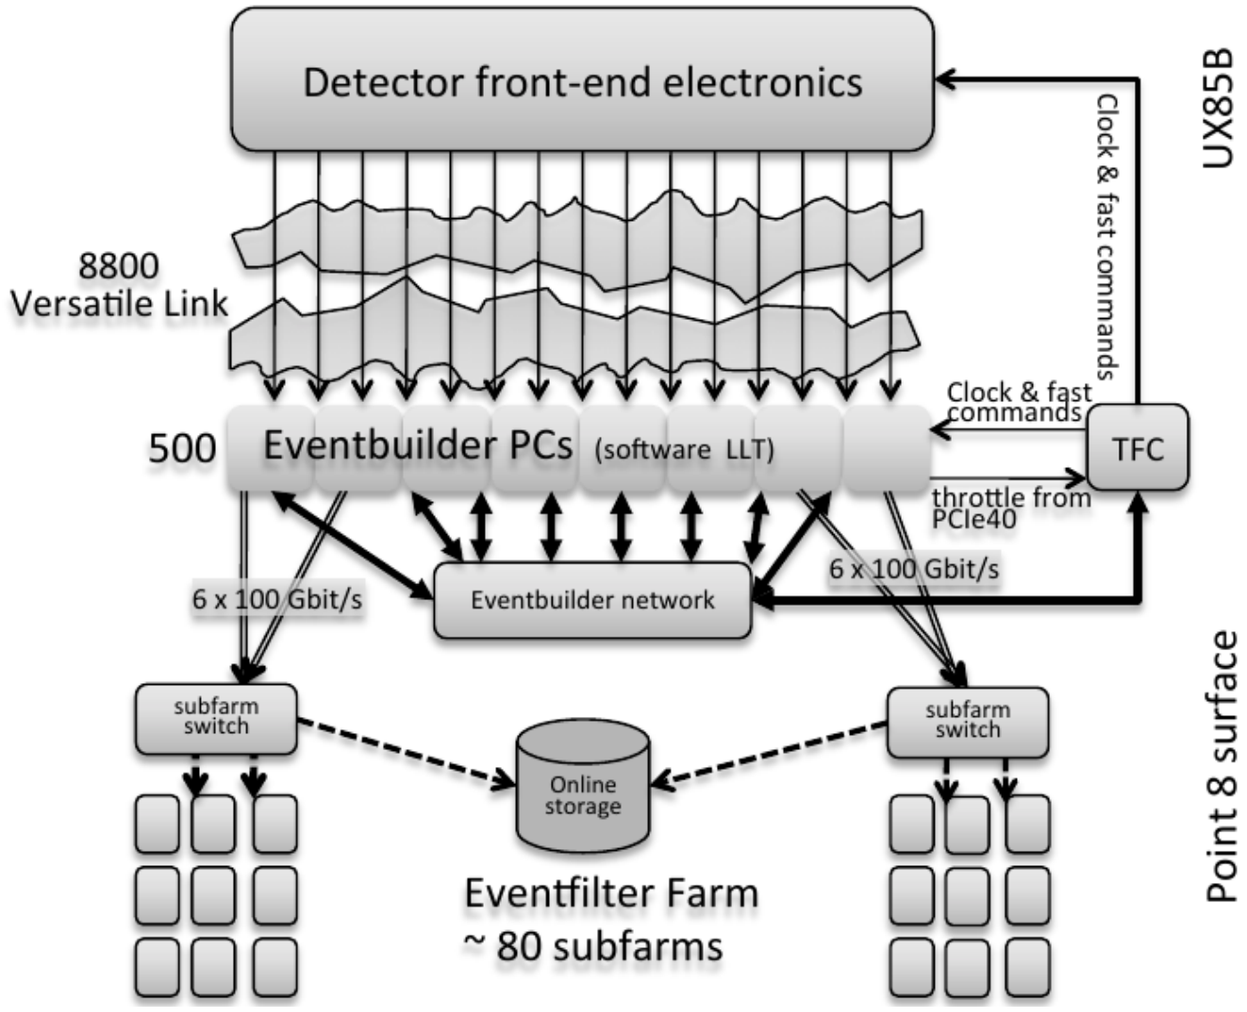

- LHCb are already going down this path for LHC Run III!
  - Full software trigger, all data read from detector at 40MHz rate
  - Some ‘hardware assist’ for data unpacking seems to be also proposed
  - 30Tb/s COTS (ish) network (10x CMS / ATLAS) – feasible in 2020

Data Rates

- An estimation of the detector data rates is now required
- Assumptions
  - Tracking detector extrapolated from CMS Phase-2 binary readout design
  - Calorimeter extrapolated from CMS Phase-2 HGC (ECAL+HCAL)
  - Tracking data rate goes as acceptance × sample_rate × pitch × occupancy
  - Calo data rate goes as acceptance × sample_rate × occupancy
  - No increase in resolution, granularity, timing precision from CMS
    - This is almost certainly therefore a lower limit, with most ZS 'tricks' already played
- First rough estimates (within factor of two)
  - Tracking and calo *each* have raw data rates of ~2000TB/s
  - Using 10Gb/s modularity links, this is 4M optical fibres
    - Also implies an event-building network of 40Pb/s capacity
- For comparison:
  - 'Entire internet WAN' today is ~500Tb/s; largest Google data centre is ~1Pb/s
- A very scary number, even for 2035, but perhaps not impossible

The Price of Bandwidth

WARNING: Britain faces internet RATIONING as UK grid struggles to cope with web demand

BRITONS could see the unwelcome return of RATIONING as the UK power grid and communications network struggles to cope with the country's growing demand for internet access, an expert has warned.

By AARON BROWN
PUBLISHED: 07:42, Wed, May 6, 2015 | UPDATED: 16:51, Fri, May 8, 2015

SHARE TWEET 425 22

The days of unlimited internet access may be numbered

- Internet communications now use 2% of UK electricity supply
  - The servers use less power than the infrastructure…

Technology Issues

- How much power do rad-hard data links take?
  - 'Best in class' *today* seems to be lpGBT at ~500mW for 5Gb/s (plus laser)
    - If no progress on this, indicates a power budget of 2MW for links alone – infeasible
  - There are no commercial applications for these links, so no COTS
  - Technology will improve, but there are some fundamental limits
    - New ideas for power saving are coming forward, but may not be applicable for us
- What are the limits?
  - Electrical signalling places a fundamental limit of ~10mW per link
    - But Shannon's limit also mandates a move to PAM / FEC → more tx and rx power
    - ~10mW for 5Gb/s in lab with 'fancy technology' (high mass, expensive, not rad hard)
  - Reducing to this level would require substantial investment in R&D
    - Not clear when / if we will have access to the required technology nodes
  - The real limit is likely to be cost
    - Also bearing in mind that COTS rx ports are ~$100, and not decreasing
    - This implies aggregation onto fast (100Gb/s+) fibres from lower speed local links
- Cost & power budget of on-detector electronics is *the* problem

Possible Approaches

- ‘Conventional trigger’
  - Extreme processor performance
  - On-detector primitives logic
  - On-detector front end buffers

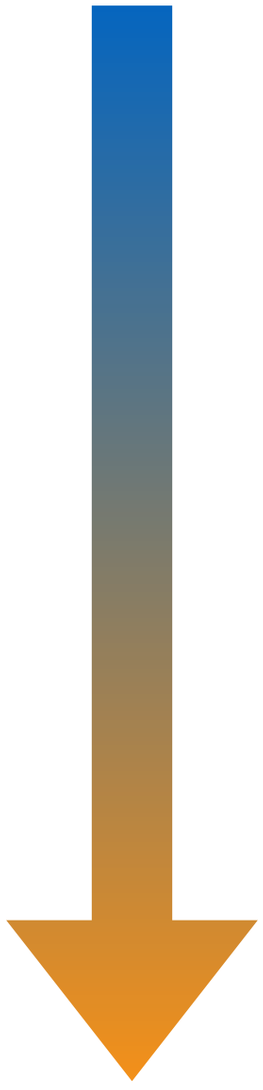

- ‘Triggerless’
  - Massive bandwidth
  - Little on-detector logic
  - Small front end buffers

Possible Approaches

- ‘Conventional trigger’
  - Extreme processor performance
  - On-detector primitives logic
  - On-detector front end buffers
  - Emphasis on on-detector processing

- ‘Triggerless’
  - Massive bandwidth
  - Little on-detector logic
  - Small front end buffers
  - Emphasis on data transmission

- ‘Sequential readout’
  - Stage out event to multi-level trigger
  - Successive levels of details with time
  - All data through event-builder network
  - Trigger implemented in software
  - Implement large ‘bulk memory’ in low radiation zone of detector
  - Emphasis on on-detector buffer

Why Worry About This Now?

- Which approach should be pursued?
  - Depends on cost / performance evolution of links, front end digital logic, bulk memory
  - My bet: bulk ECC memory cheaper than links or rad-hard front end logic
- Who cares? This is all in 2035…
  - As we've seen, trigger /readout strongly affects detector design
  - Basic R&D on new front-end technologies needs to start now
  - Factorisation of 'local' and 'global' on-detector electronics could help

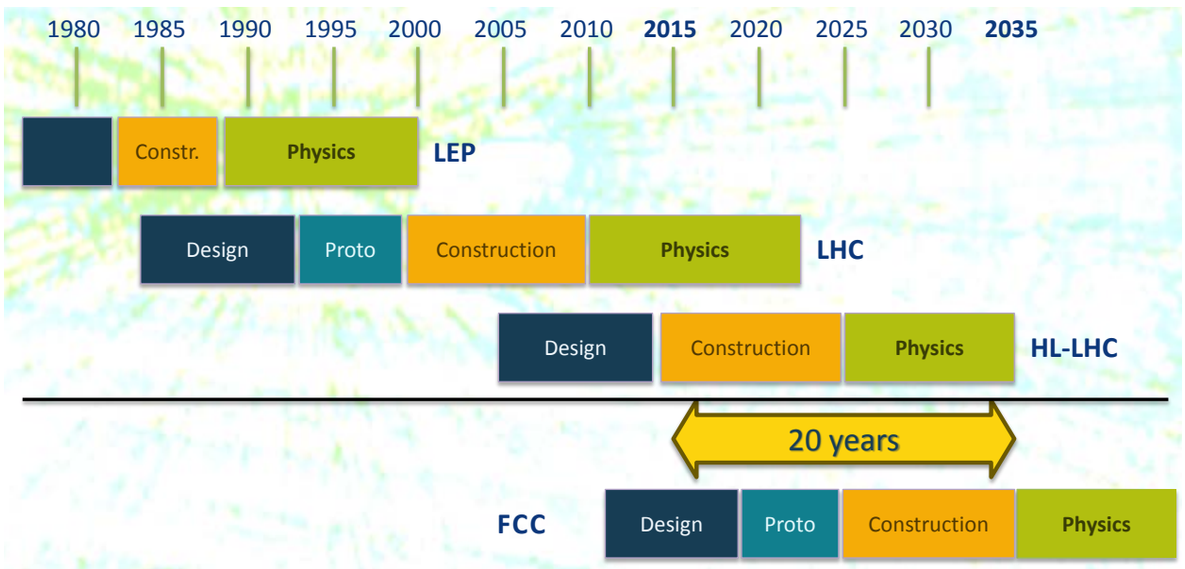

Conclusions

- Too early to propose a trigger / readout architecture
  - But examination of the constraints needs to continue
  - Don't depend too strongly on 'magical' technology improvement
    - We might not be able to afford the development or purchase cost; the market may go in a different direction
- Trigger performance studies
  - Reconstruction studies need to take into account trigger needs from start
  - Can things be sped up with 'helpful' detector modularity / configuration?
  - 'Feasibility check' of traditional trigger approach needs to be made
- R&D is needed
  - Front-end electronics (links, processing, memory)
  - Technology evolution for bulk packet switching (e.g. data centre ethernet, etc)
  - Accelerated computing architectures (e.g. OpenPower/CAPI)
- Detector design
  - Do not ignore the problem of trigger and readout for your detector! It will bite you later.
  - Design for the worst case; you will never rebuild the FE electronics
- Watch LHC Phase-2 upgrades carefully
  - FCC-hh developments are likely to be a direct 'follow-on'