Introduction to First-Level Triggering

- Today's trigger menu:
  - Trigger architecture
  - Requirements on first-level triggers
  - Trigger approaches at LHC
  - Algorithms
  - Implementation
- We will focus mainly on LHC triggers
- Please stop me to ask questions...
  - Or ask for jargon to be explained
- Who am I?
  - Reader in Bristol experimental group
  - ~15 years on fixed-target spectroscopy and CMS trigger systems

What is a Trigger?

- Simplest definition:
  - A system that decides in real time whether to retain or discard the measurements corresponding to each observed interaction
- Practical definition:
  - Hardware/software processor filtering the event stream based upon a ‘quick look’ at the data
  - It must keep ~all the interactions of interest for later analysis
  - It must accept interactions at a rate low enough for storage and reconstruction
- This is a risky business!
  - Little room for error, as events discarded can never be recovered
  - We usually do not know what to expect in advance
    - “If we knew what we were doing, it wouldn’t be called research” - AE
  - Often a (the?) determining factor in physics reach of an experiment
    - Especially true of ‘energy frontier’ hadron collider experiments

Modern DAQ Architecture

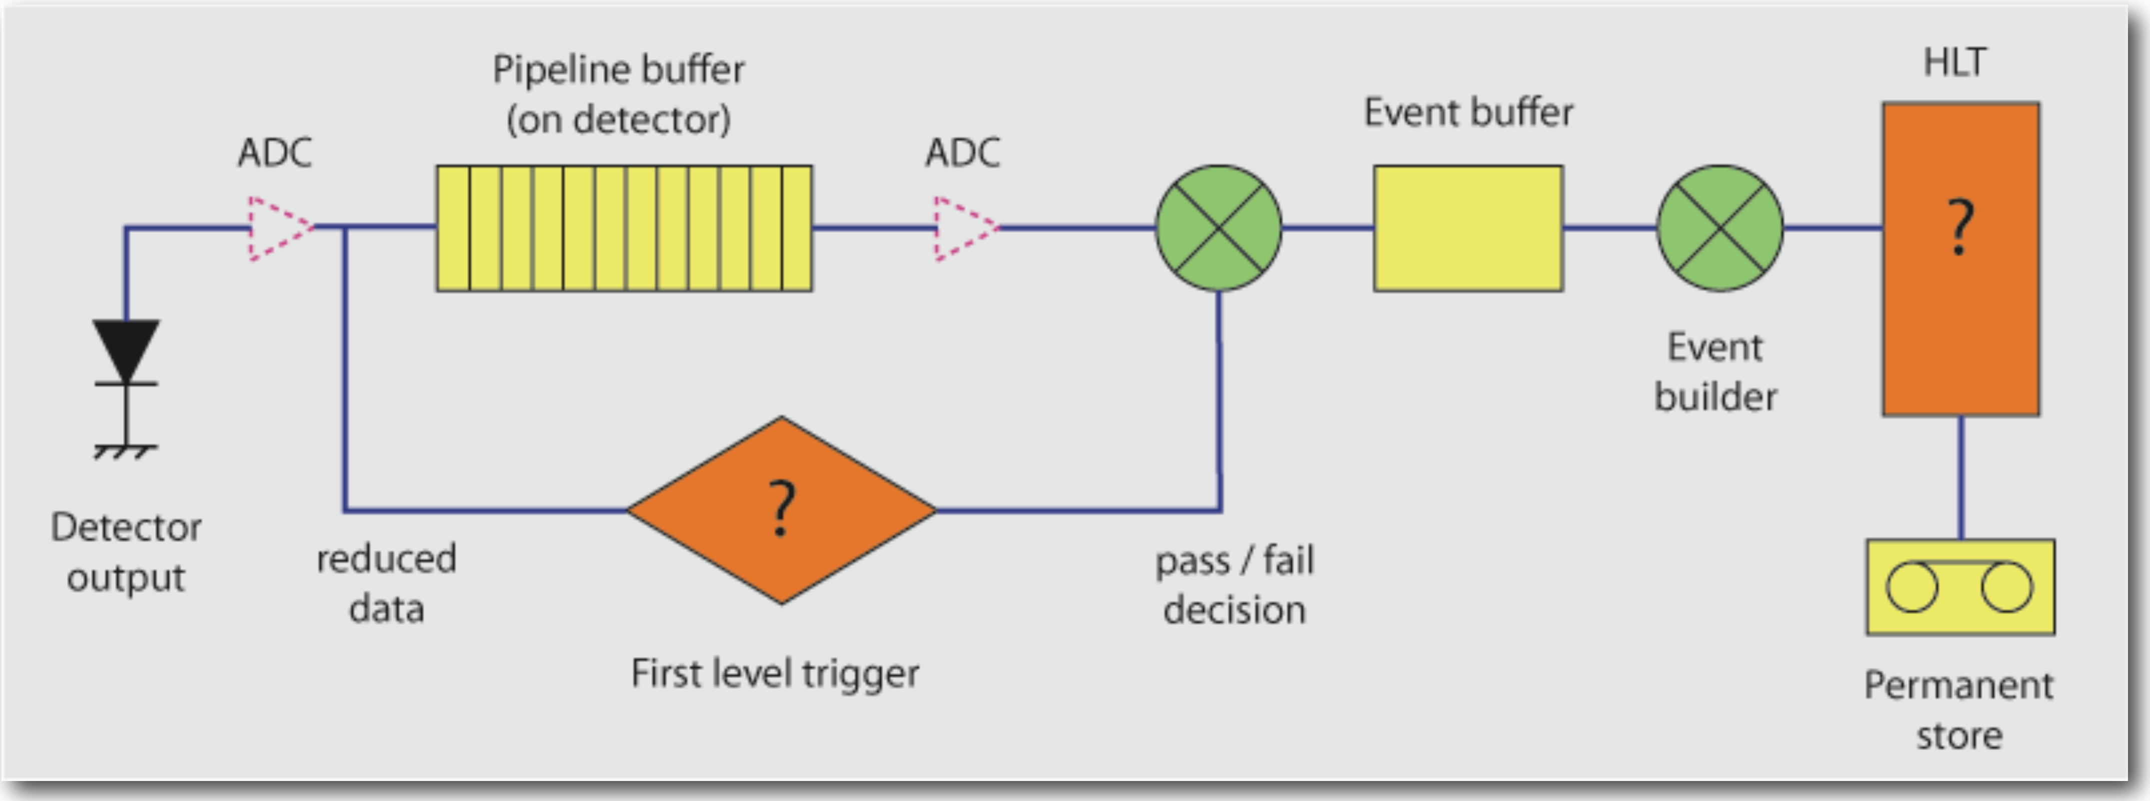

- 20+ years ago
  - The key problem was often *readout time*
  - Slow detectors need to be explicitly cleared if no trigger decision is made
  - *Dead time* [time unavailable for recording signals] a major issue
- The most recent experiments
  - Key problem is *data volume* - local storage gives very low deadtime
  - e.g. ~100M ATLAS channels at ~12 bits each; how much data per year?

Dave.Newbold@cern.ch

LHC Cross-Sections

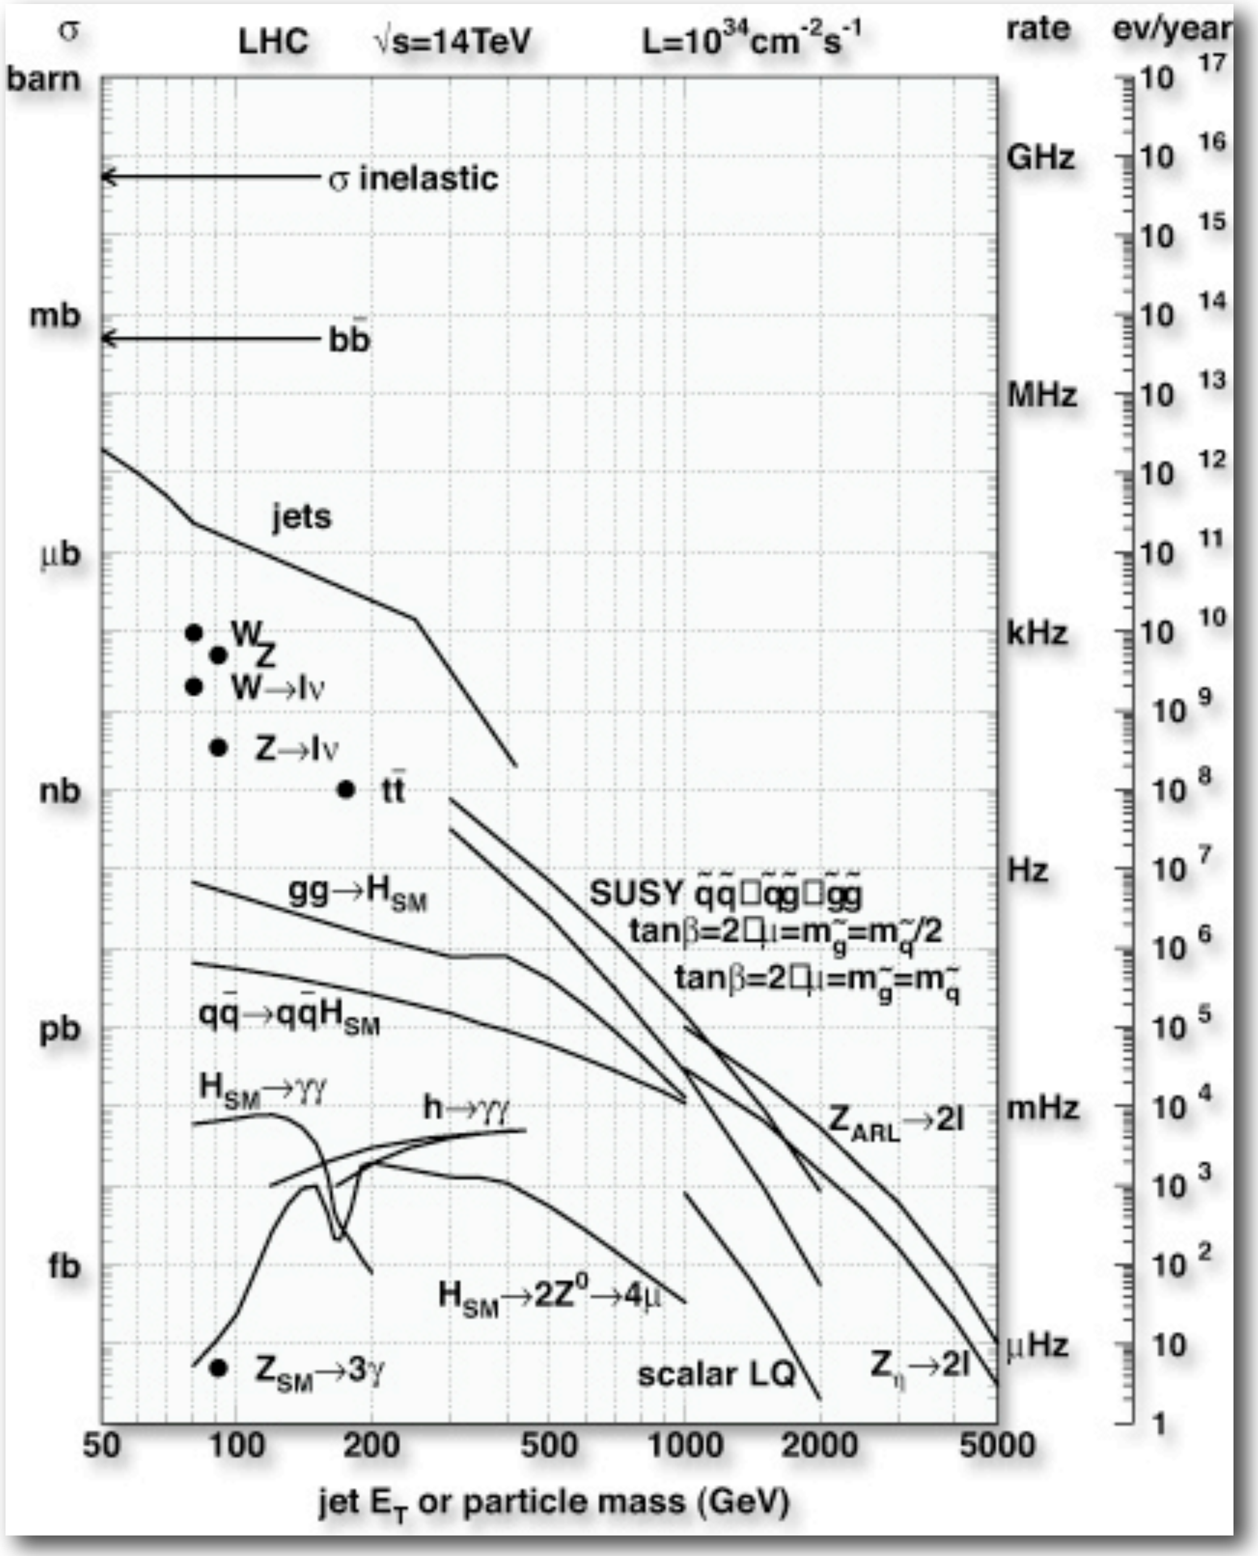

Reqts: Rate Reduction & Selectivity

- Typical requirement (LHC / Tevatron)
  - Maximum event accept rate of O(50kHz)* => rate reduction of O(1000)
    - *For pedants: strictly 50kBq
  - At high-lumi machines, the trigger of course works on *crossings*
    - $\langle n_{inelastic} \rangle$ = ~20 at 10^{34} $cm^{-2}s^{-1}$ lumi; crossing rate is 40MHz (LHC) => 1GHz of events
- Basic strategy for 'energy frontier' experiments
  - We are looking for heavy states with short lifetimes [not always!]
  - Try to identify their decay products in the detector
    - Separate from background by imposing a transverse momentum threshold
    - Recall: in hadron collisions, p_z tells you very little due to asymmetry of collision
  - At hadron machines, avoid overwhelming QCD background where possible
    - i.e. our 'trigger objects' are leptons, photons (often in pairs), and global event variables
    - Very large rate of light (and heavy!) quark production -> QCD-based signatures are buried
- Important: trigger rates dominated by background
  - The key trade-off is between acceptance and trigger rate
  - Whilst respecting all the constraints of implementation...

Turn-on Curve

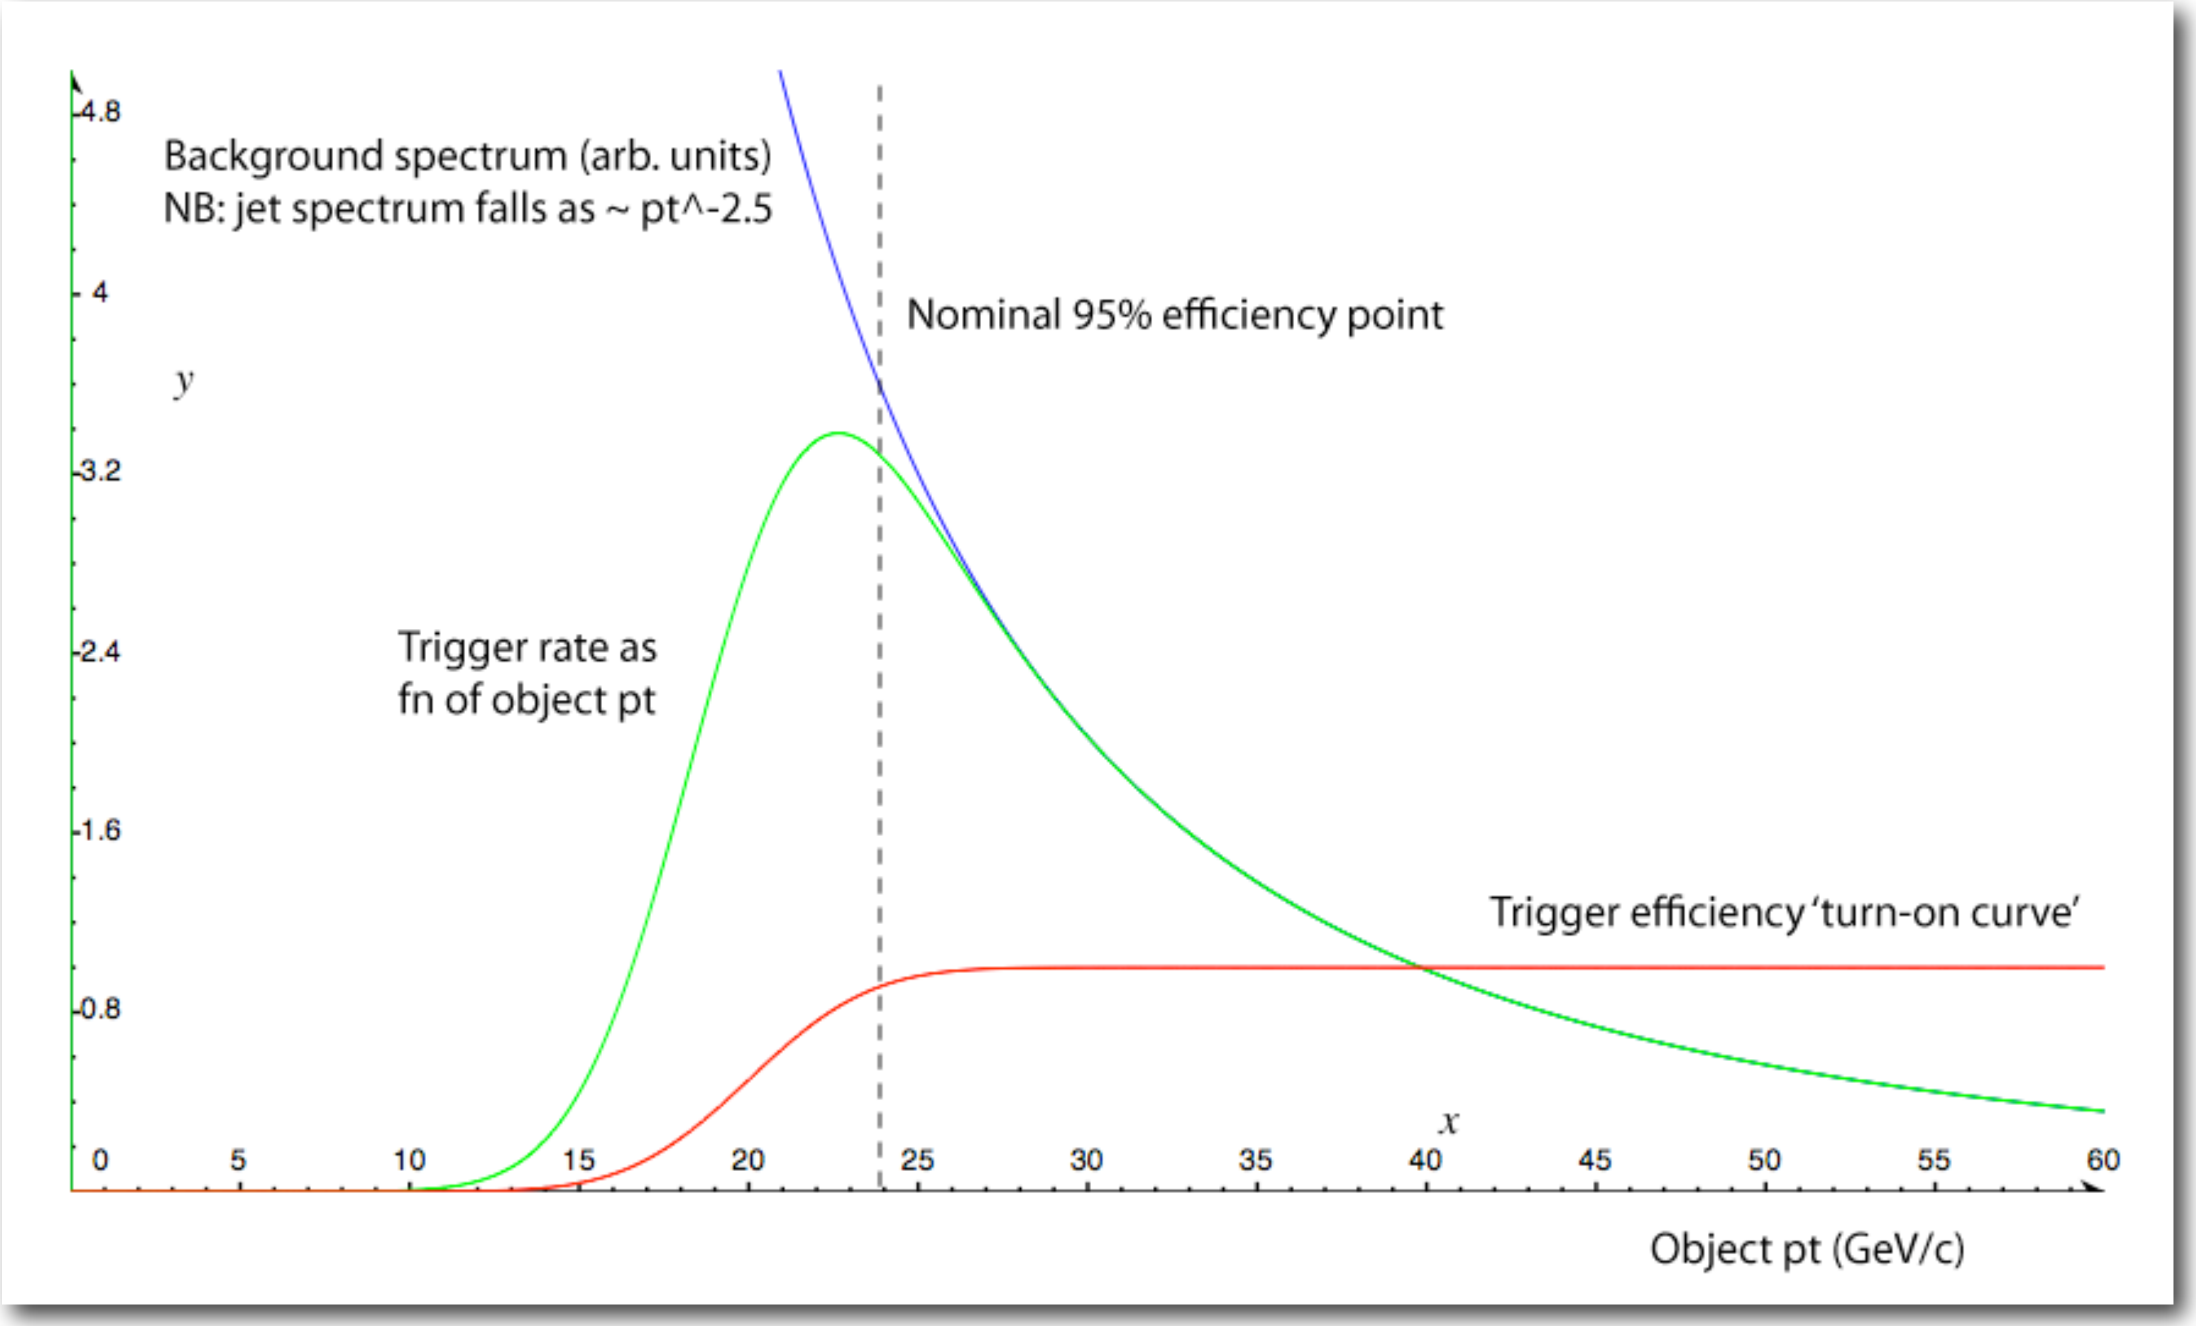

- Key issue - steeply falling p_t spectrum of b/g
  - Rate can be dominated by 'turn-on curve' - sharp turn-on is essential

Reqts: Acceptance & Bias

- Trigger acceptance directly affects the measurements
  - Ideal: complete acceptance for all events of interest
  - In practice, aim for: trigger thresholds lower than any conceivable analysis cut
  - At LHC, these goals are not always achievable at reasonable cost
    - e.g. rate of b-physics events is limited by available trigger rate / storage
- Plan trigger strategy according to physics goals
  - Electroweak physics (Higgs, TGCs etc) suggest lepton / W / Z
    - CMS and ATLAS aim for ~full efficiency for inclusive W & Z leptonic decays
  - SUSY suggests high-pt jets, missing transverse energy (not easy)
  - b-physics requires early (~Level 1) displaced vertex trigger
  - Direct BSM searches not usually limited by trigger (but there are subtleties...)
- Trigger bias
  - Thresholds close to analysis cuts will bias distributions - must account for this
  - Topological / impact factor cuts can have more subtle bias
    - e.g. for overlapping trigger objects

Example: Higgs Boson @ LHC

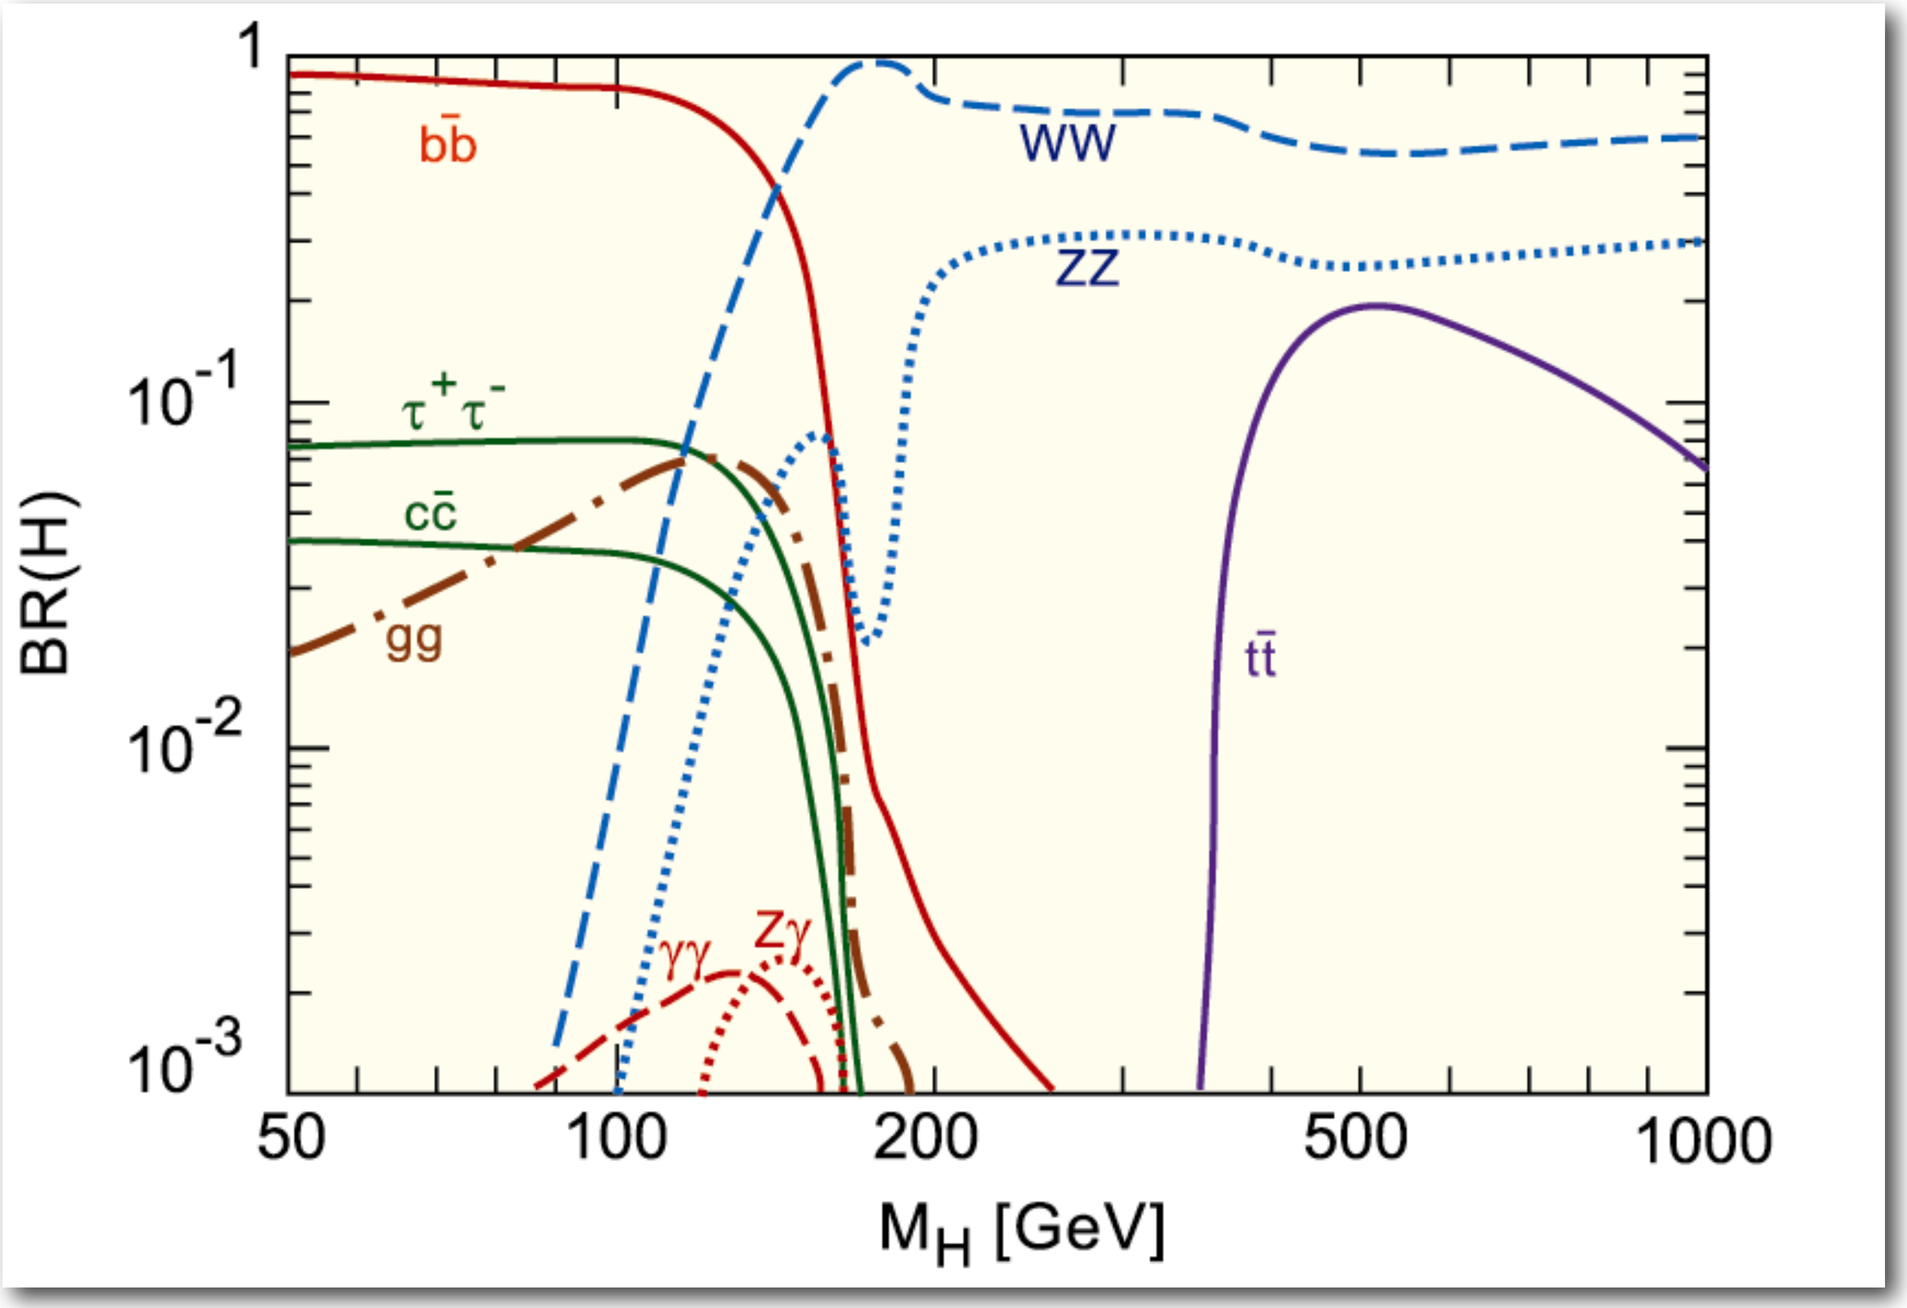

- Trigger strategy must support all possible analysis paths
  - H -> γγ ; H -> WW; H -> ZZ*; H -> WW [H -> bb impossible with baseline trigger]

Reqts: Latency & Robustness

- What stops us building the perfect trigger? Latency.
  - Latency = time taken to reach a decision
  - During this time, all data must be held (usually on-detector) in memories
- Timescales
  - Detector memories have typically ~128BX depth = 3.2μs
    - But: distance from detector to trigger (and back - for decision) is ~200m
  - => Available time for trigger algorithms is more like 1.5μs
    - This is beyond any general-purpose computer - e.g. this is a few 1000 cycles of a modern CPU
    - For comparison, HLT has ms - s timescale to do its job; can use DSP or commodity CPU
  - In practice, requires *pipelined digital logic*
- Robustness is essential
  - Trigger is a 'mission critical' system; no data can be taken without it
  - It must function, and function predictably, under all experimental conditions
  - Mechanisms must be devised for monitoring and self-test of trigger functions

Reqts: Control & Understanding

- At a practical level
  - Trigger is a complex system - there is a lot to go wrong
  - Technical monitoring and trigger path verification are essential
    - Many systems have the ability to 'play through' analogue or digital data to check performance
    - The effects of dead components, dead/noisy detector channels, etc, must be recorded
- Trigger performance measurement
  - How to quantify the efficiency / bias of the trigger?
    - Remember, the trigger has an irrevocable effect on *all* statistical analysis / counting studies
  - Monte Carlo simulation is one approach
    - This is not sufficient for analysis purposes; performance varies with lumi, time, detector performance
  - Possible to monitor the trigger performance from 'minimum bias' data
    - Minimum-bias triggers [triggers with no requirements at all] provide a neutral dataset
    - Can assess trigger result on this data offline, estimate trigger performance
  - Pre-scaled triggers [take 1-in-n triggers at a lower threshold] are also necessary
    - Minbias may not contain a sufficiently large control sample of triggerable events
    - Pre-scaled data also forms part of the trigger menu [see later] for calibration and physics

Data Reduction

- How does the trigger receive input data?
  - Typically 'parasitic' on the main detector readout system
  - Exception is when dedicated trigger detectors are used (e.g. RPCs for muons)
  - Do not need and cannot handle all detector data in trigger system
    - Only a subset of detectors used - calorimeter, muon system, [sometimes] inner tracking
  - NB: trigger needs data *promptly*; significant bandwidth required for this data
    - In many cases, more than the detector readout itself
- Data volume reduction
  - Zero-suppression not typically used (no-hit cells are important for trigger!)
  - Elements are grouped (e.g. addition of calorimeter cells) => granularity reduction
  - Detector signals truncated / compressed / delinearised => resolution reduction
  - Some trigger functions are performed on detector (e.g. hit correlation)
    - Often perform filtering to extract timing information 'close to' the detector
  - Timing information (explicit or implicit) must be preserved!
- Input data to trigger known as 'trigger primitives'

Objects & Algorithms

- What is the trigger looking for?
  - Evidence of decay products in the detector; each has a characteristic signature
  - Trigger may count objects-above threshold, or sort objects by pt and pass on
  - Typically a limit on the number of objects in given detector region
    - This is due to the limitations of trigger electronics – fixed-size internal busses, etc
- Trigger algorithms
  - Operate on trigger primitives information from subdetector(s) to find objects
  - Generally, several algorithms operate in parallel to find different objects
    - e.g. calorimeter information used to find electrons + jets in parallel
  - Algorithms must cover whole detector in an unbiased way
    - So watch out for edges where different systems overlap, e.g. in 'sliding window' algos
  - Output is a count or list of trigger objects, possibly with additional information
    - Object pt, position, charge, 'quality', etc
- Some algorithms are 'global' over the whole detector
  - Examples: Missing Et, Total Et [is this useful?], Ht, global object counts

Calorimeter Trigger Algorithms

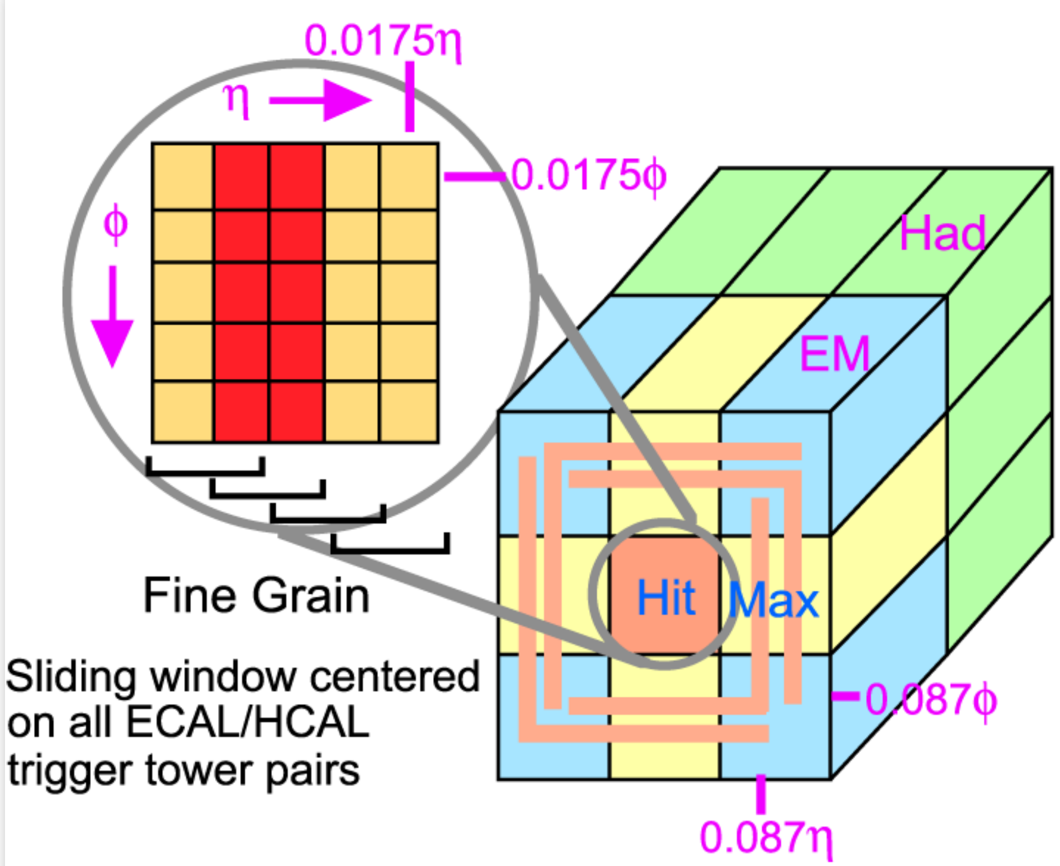

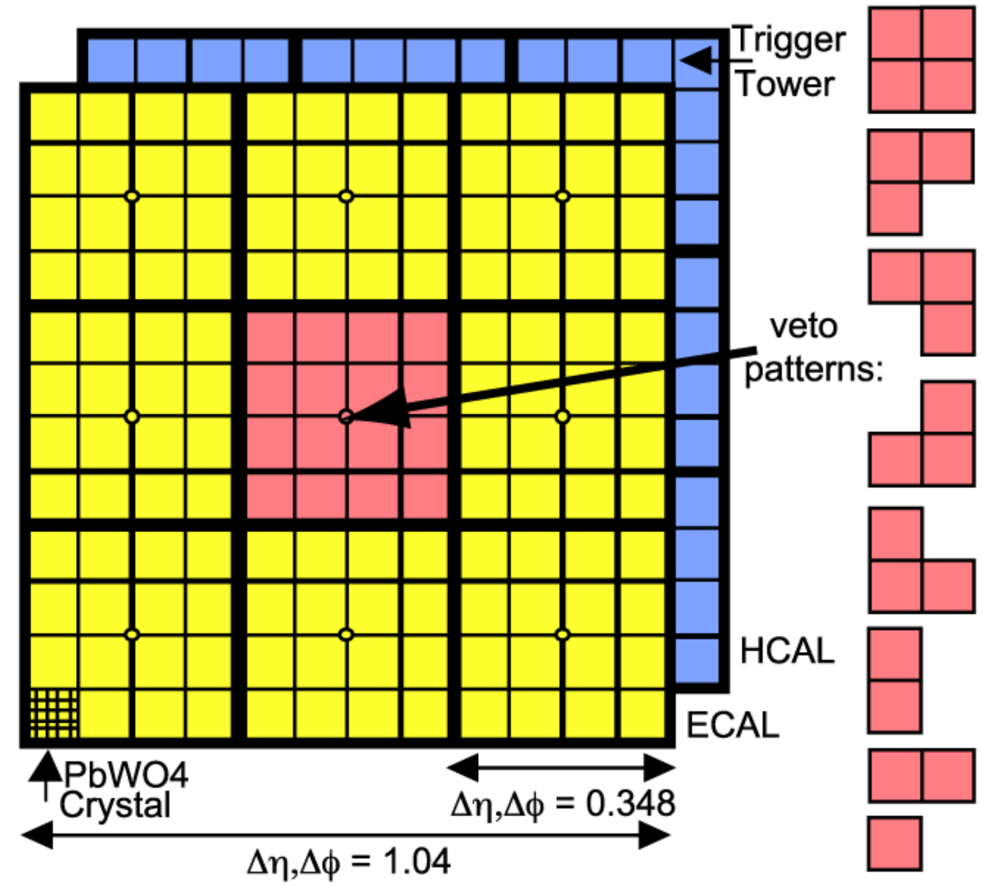

e/γ (hit tower + max neighbour):

- 2-tower Et; hit tower passes H/E cut
- Hit tower: 2x5 strip with >90% Et in 5x5 (FG)

Isolated e/γ added criteria:

- All 9 towers pass FG and H/E
- One 'corner' group of EM towers < Thr

Jet or τ:

- ΣEt of 12x12 trig tower sliding window
- Central 4x4 Et > each neighbour

τ (isolated narrow deposit) added criteria:

- all 9 regions have 'τ pattern' deposit

Total / missing Et uses 4x4 granularity

Total "Ht" uses found jets only

- CMS calorimeter trigger

Dave.Newbold@cern.ch

Muon Trigger Algorithms

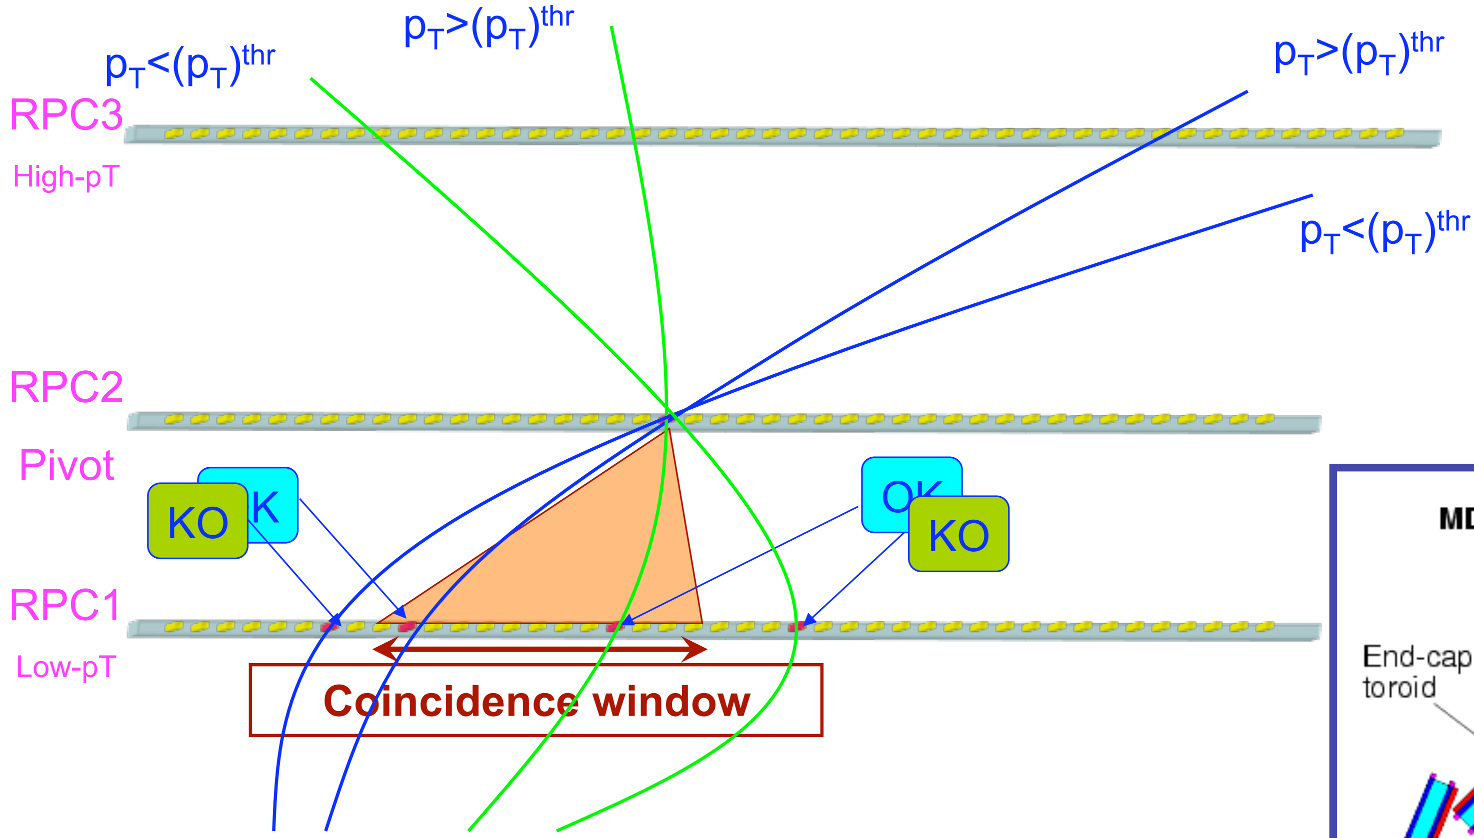

- ATLAS barrel muon trigger

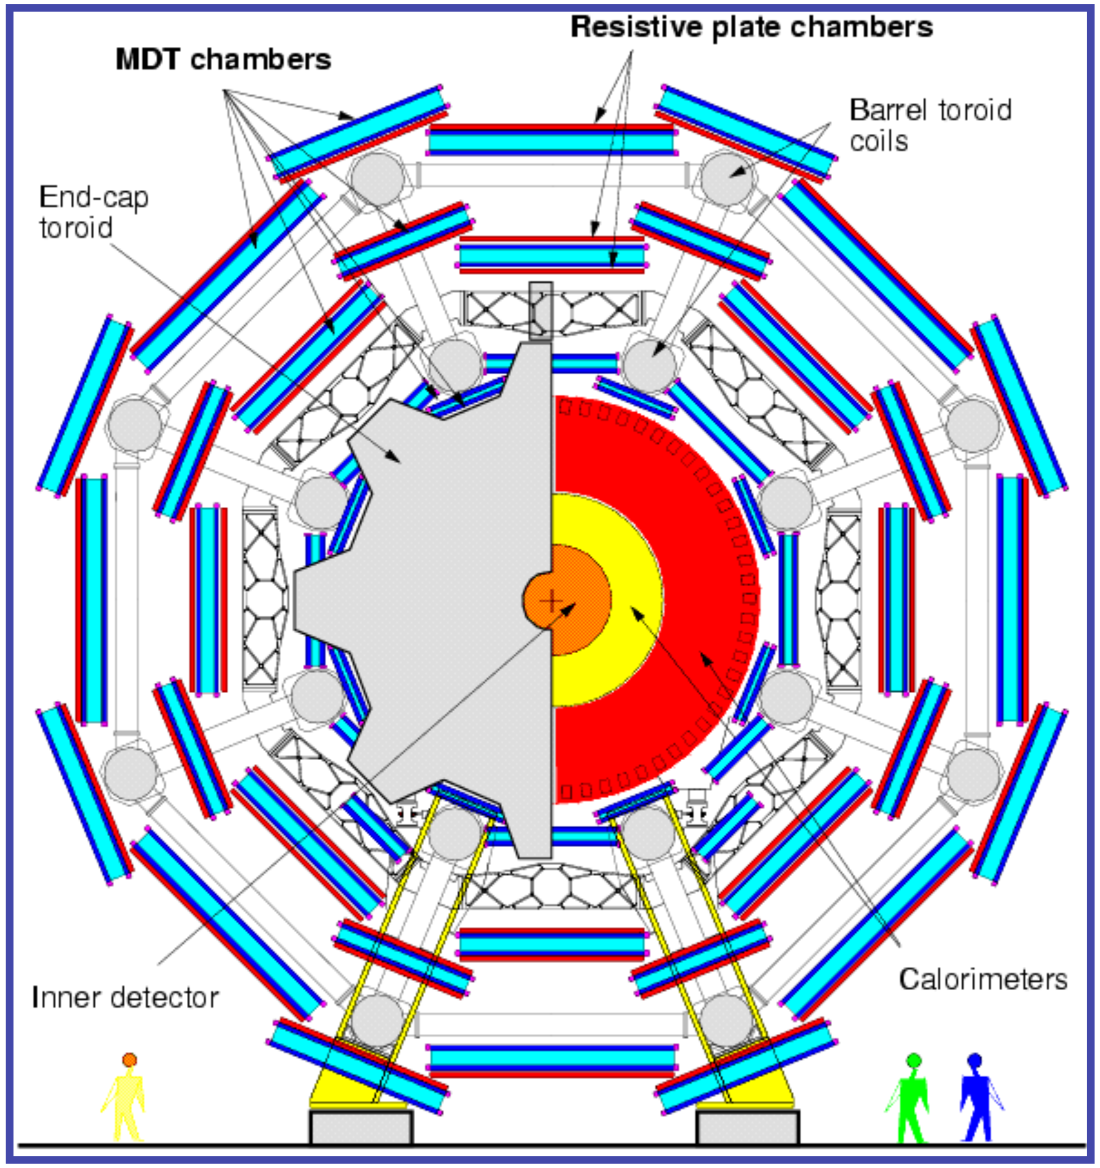

- Not as simple as it looks!
  - Hit correlation in 4D is necessary
  - Muon detector spacing is large compared to time-of-flight
  - Detectors with very good time resolution required for bunch-crossing assignment

Tracking Trigger Algorithms

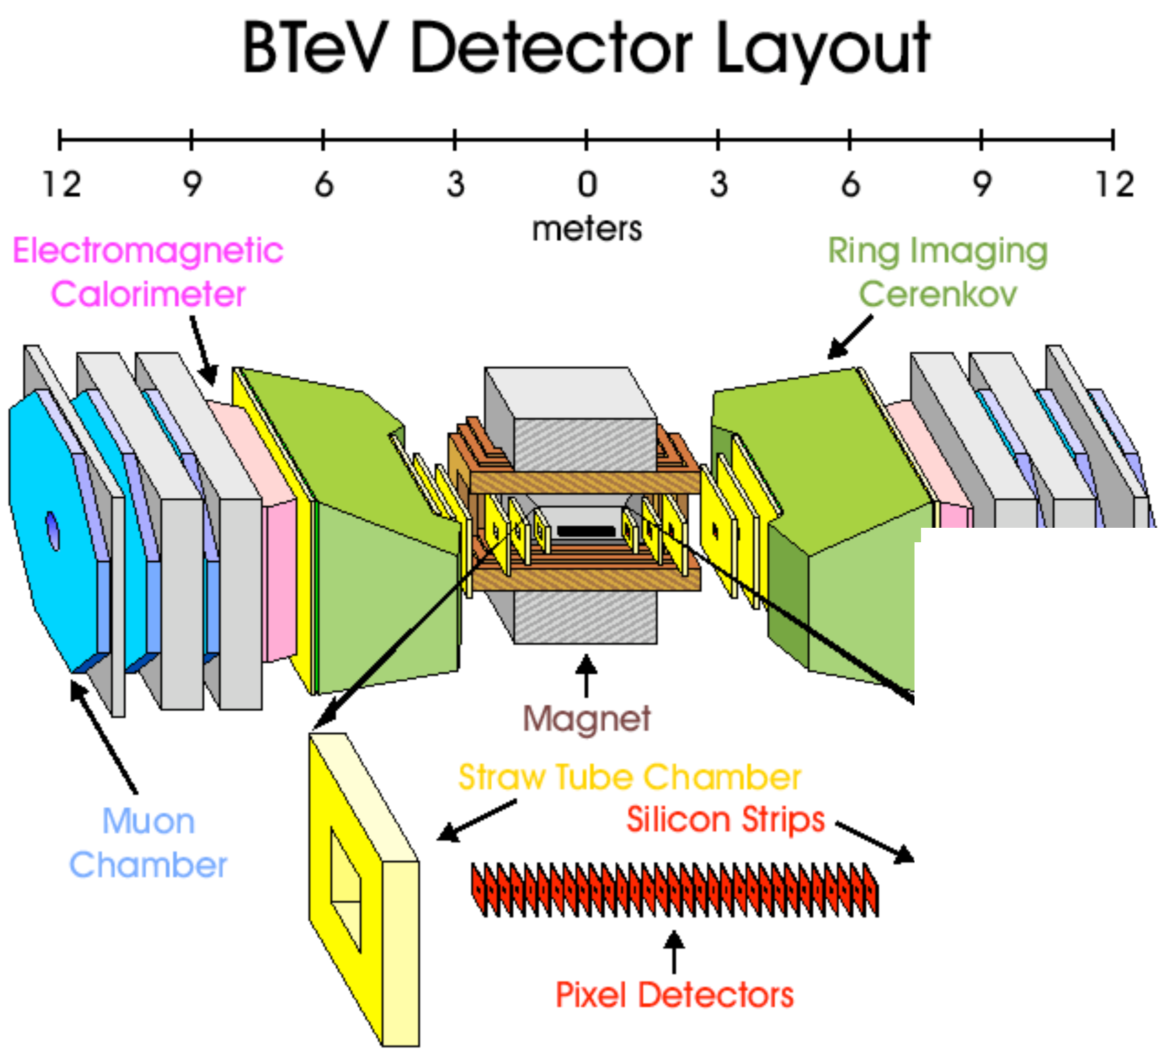

- BTeV pixel trigger

Blue segments are 'entering' detector

Green segments are 'leaving' detector

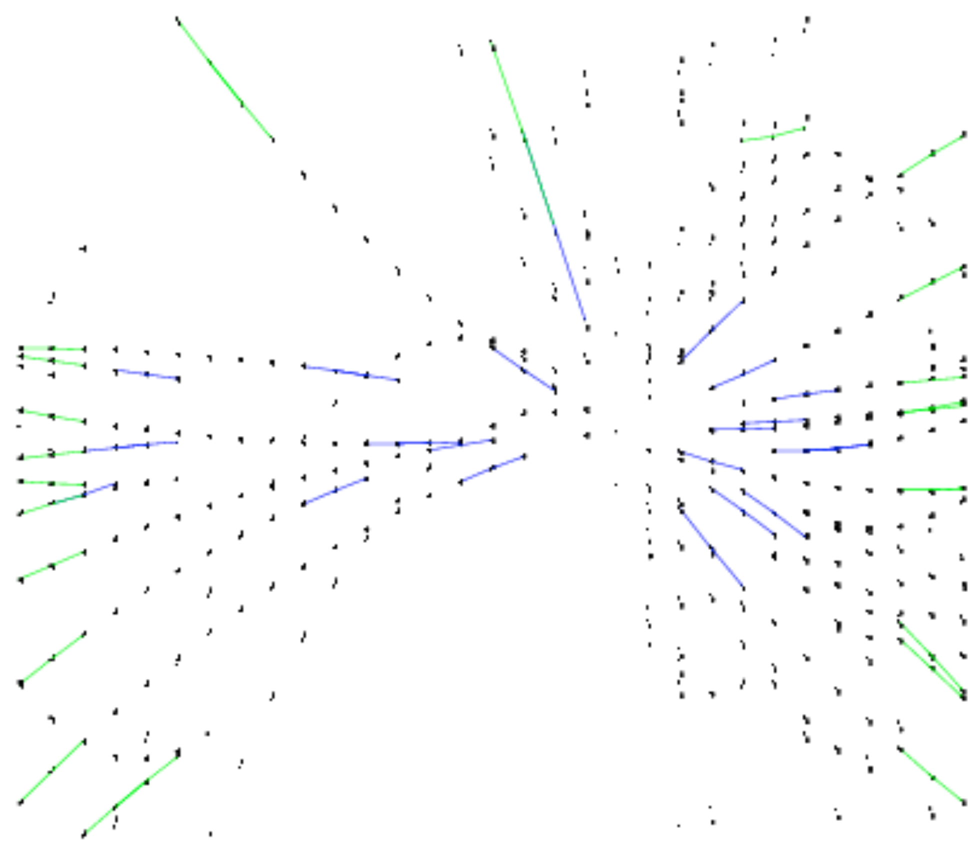

- Based upon triplet-finding approach - rather neat
- Finds number of displaced vertices
- Rejects pile-up and high-multiplicity events

Decision-Making: Global Triggers

- The decision-making process
  - Global trigger objects or object counts from trigger subsystems
    - Typically, calorimetry, muons and tracking are in different subsystems
  - Applies a set of criteria to make a final yes / no 'Level-1 Accept' decision
  - These criteria are flexible and programmable
    - The criteria will evolve as the experiment goals mature; they may even change during a run
  - Global trigger typically also generates technical / calibration triggers
    - These include detector monitoring (calibration pulses, etc), 'empty bunch' monitoring, system tests
- Decision logic
  - Typically arranged in an 'evaluate -> and -> or' tree
  - Object energies and counts are evaluated against a large set of possible criteria
    - e.g. 'two electrons above 40GeV; jet and muon in opposite hemispheres'
  - Criteria are merged to form trigger paths (the jargon varies)
    - e.g. two leptons, two jets and missing energy
  - The paths are or'ed such that a trigger is issued if any are active
    - The set of all possible paths is sometimes known as the 'trigger menu'

Overall First-Level Architecture: CMS

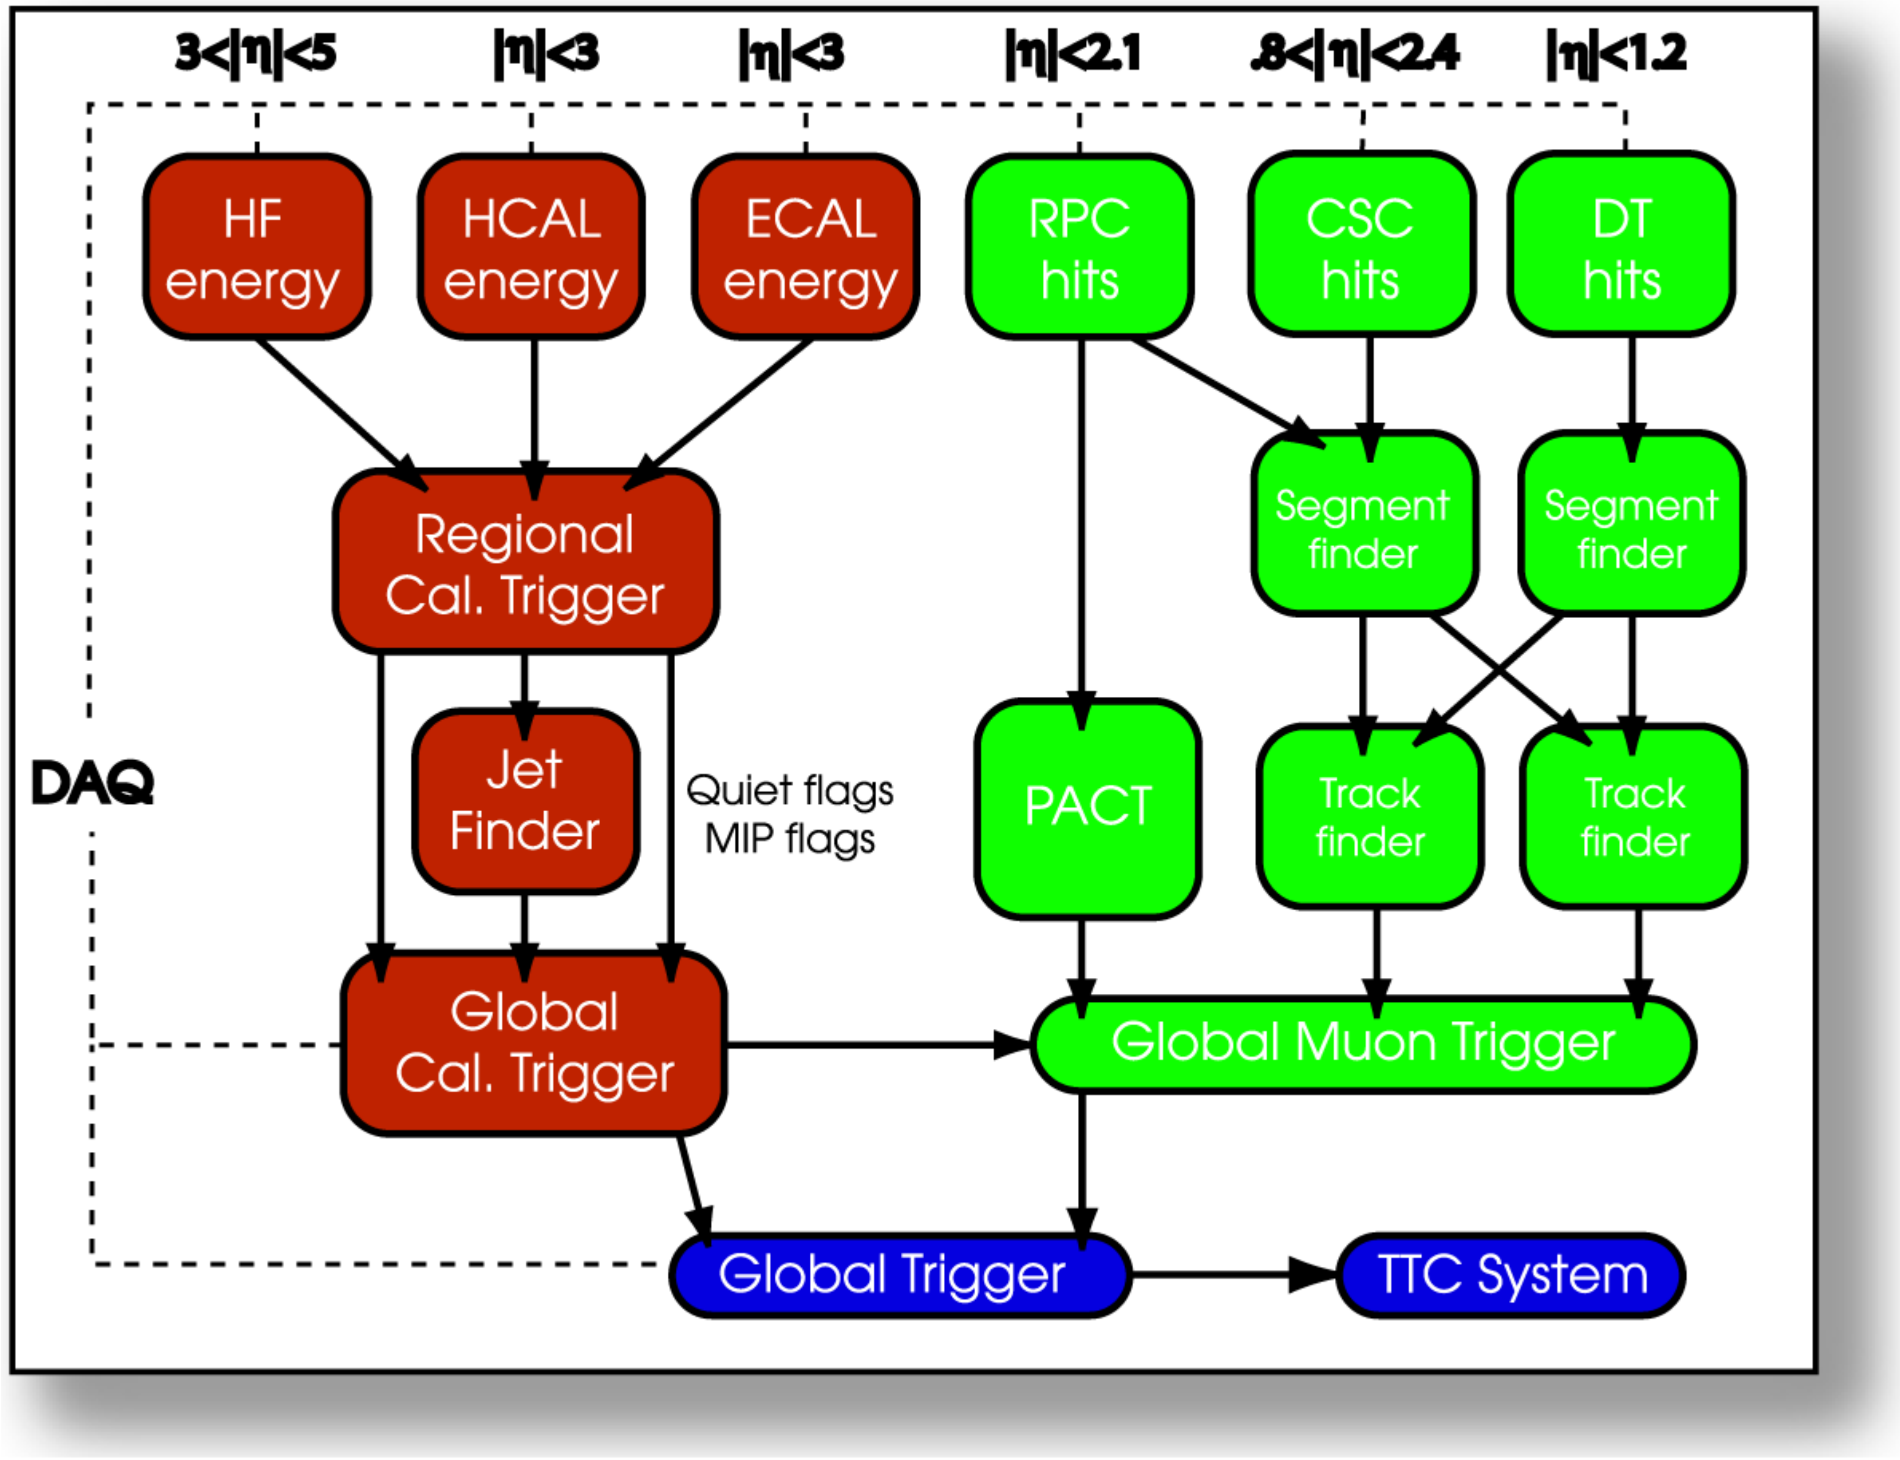

Example Trigger Menu

```
   L1_SingleMu3 (4000) : Indiv.: 3.2 +/- 2.5
L1_SingleMu5 (2000) : Indiv.: 3.2 +/- 2.5
L1_SingleMu10 (1) : Indiv.: 496.7 +/- 17.1
L1_DoubleMu3 (1) : Indiv.: 316.1 +/- 20.3
L1_TripleMu3 (1) : Indiv.: 7.0 +/- 2.5
L1_Mu3_Jet15 (20) : Indiv.: 200.0 +/- 17.1
L1_Mu5_Jet20 (1) : Indiv.: 1282.5 +/- 36.0
L1_Mu3_IsoEG5 (1) : Indiv.: 922.0 +/- 35.6
L1_Mu5_IsoEG10 (1) : Indiv.: 57.4 +/- 7.0
L1_Mu3_EG12 (1) : Indiv.: 82.9 +/- 9.2
L1_SingleIsoEG8 (1000) : Indiv.: 19.2 +/- 6.5
L1_SingleIsoEG10 (100) : Indiv.: 82.8 +/- 13.5
L1_SingleIsoEG12 (1) : Indiv.: 4003.4 +/- 93.0
L1_SingleIsoEG15 (1) : Indiv.: 1757.9 +/- 61.3
L1_SingleIsoEG20 (1) : Indiv.: 574.8 +/- 34.8
L1_SingleIsoEG25 (1) : Indiv.: 232.1 +/- 22.0
L1_SingleEG5 (10000) : Indiv.: 13.3 +/- 5.5
L1_SingleEG8 (1000) : Indiv.: 21.9 +/- 7.0
L1_SingleEG10 (100) : Indiv.: 99.8 +/- 14.8
L1_SingleEG12 (100) : Indiv.: 53.4 +/- 10.7
L1_SingleEG15 (1) : Indiv.: 2471.9 +/- 72.3
L1_SingleEG20 (1) : Indiv.: 925.5 +/- 43.7
L1_SingleEG25 (1) : Indiv.: 456.7 +/- 30.7
L1_SingleJet15 (100000) : Indiv.: 10.3 +/- 4.9
L1_SingleJet30 (10000) : Indiv.: 18.7 +/- 6.5
L1_SingleJet70 (100) : Indiv.: 34.2 +/- 8.5
L1_SingleJet100 (1) : Indiv.: 588.3 +/- 34.7
L1_SingleJet150 (1) : Indiv.: 66.4 +/- 11.0
L1_SingleJet200 (1) : Indiv.: 19.5 +/- 6.0
L1_SingleTauJet40 (1000) : Indiv.: 0.0 +/- 0.0
L1_SingleTauJet80 (1) : Indiv.: 723.1 +/- 38.4
L1_SingleTauJet100 (1) : Indiv.: 214.5 +/- 20.8
```

```
 L1_HTT100 (10000) : Indiv.: 16.3 +/- 6.0
 L1_HTT200 (1000) : Indiv.: 22.3 +/- 7.0
 L1_HTT250 (100) : Indiv.: 60.6 +/- 11.3
 L1_HTT300 (1) : Indiv.: 1739.1 +/- 59.8
 L1_HTT400 (1) : Indiv.: 158.5 +/- 17.4
 ETM45 (1) : Indiv.: 527.6 +/- 33.8
 ETM45_Jet30 (1) : Indiv.: 511.6 +/- 33.3
 ETM50 (1) : Indiv.: 190.0 +/- 20.0
 L1_DoubleIsoEG8 (1) : Indiv.: 740.4 +/- 39.2
 L1_DoubleEG10 (1) : Indiv.: 0.0 +/- 0.0
 L1_DoubleJet70 (1) : Indiv.: 733.9 +/- 38.8
 L1_DoubleJet100 (1) : Indiv.: 150.3 +/- 17.4
 L1_DoubleTauJet40 (1) : Indiv.: 2970.4 +/- 78.9
 L1_IsoEG10_Jet15 (20) : Indiv.: 345.4 +/- 27.4
 L1_IsoEG10_Jet30 (1) : Indiv.: 3990.7 +/- 92.2
 L1_IsoEG10_Jet70 (1) : Indiv.: 472.8 +/- 31.0
 L1_IsoEG10_TauJet20 (1) : Indiv.: 3697.9 +/- 88.7
 L1_IsoEG10_TauJet30 (1) : Indiv.: 2389.5 +/- 70.9
 L1_TauJet30_ETM30 (1) : Indiv.: 3570.6 +/- 88.3
 L1_TauJet30_ETM40 (1) : Indiv.: 587.7 +/- 35.4
 L1_HTT100_ETM30 (1) : Indiv.: 0.0 +/- 0.0
 L1_TripleJet50 (1) : Indiv.: 349.7 +/- 26.1
 QuadJet40 (1) : Indiv.: 192.9 +/- 19.3
 QuadJet50 (1) : Indiv.: 43.7 +/- 8.9
 L1_ExclusiveDoubleIsoEG6 (1) : Indiv.: 467.1 +/- 32.3
 L1_ExclusiveDoubleJet60 (1) : Indiv.: 158.5 +/- 18.6
 L1_ExclusiveJet25_Gap_Jet25 (1) : Indiv.: 776.4 +/-
42.7 seqPure:
 L1_IsoEG10_Jet20_ForJet10 (1) : Indiv.: 2130.9 +/-
67.6
 L1_MinBias_HTT10 (1) : Indiv.: 0.4 +/- 0.1
 L1_ZeroBias (1) : Indiv.: 0.6 +/- 0.1
```

- Example CMS L1 trigger menu for 10^{32} luminosity, 17kHz L1A rate
  - Entries are trigger path (corresponds to global trigger logic), prescale, and MC predicted rate in Hz

Implementation: Processing

- Trigger systems typically require complex custom hardware
  - No ‘off the shelf’ system means current needs
  - Need to use ‘off the shelf’ component whenever possible
    - Many commercial processing and communications technologies used in imaginative ways
- Analogue or digital?
  - Analogue processing is fast, low power, performance good enough for trigger
  - Digital electronics is easier to design and test, less risky
  - Most systems use some combination of both
    - Most often analogue front-end, digital algorithms and pipeline storage
- Processing devices
  - Custom designed ASICs have many advantages - typically used on-detector
    - Low cost in bulk, rad hard, high density, analogue functions, latest technologies (if you have the $$$)
    - Some ASICs used in trigger logic for LHC
  - FPGAs are now the dominant technology in triggering
    - (Re)programmable for any logic function (>1M gates), low risk, ‘easy’ to design for and test
    - Remember: flexibility is key for a first-level hardware trigger: FPGAs can help provide this

Implementation: Data Communication

- The key problem for trigger implementation
  - Processing technology has now reached very high densities and speeds
  - Communications technology lags behind
    - It has not been driven so hard by the consumer electronics market
  - Data density is the main issue in system design
- Data transmision: electrical or optical?
  - Large complex, distributed systems: interconnections are always an issue
    - Clock distribution, noise, ground loops, power consumption, cable plant bulk
  - Optical communication has many nice features, including noise immunity
    - But is still low-density compared to copper: no demand for parallel optical communication
  - Copper serial links are the current state of the art
    - Can move 10Gb/s via a 6mm diameter copper parallel-pair cable (infiniband 4x standard)
- Timing and control
  - LHC triggers were entirely 'synchronous systems', at every pipeline step
  - This may not have solved as many problems as it caused - yet to see!

The Future?

- The state of the art
  - LHC systems are the culmination of 40 years progress in triggering
  - Possibly the most complex custom electronics systems in use in science today
  - Their performance is (predicted to be) good, for reasonable cost
- Linear Collider
  - The ILC works entirely differently to LHC
    - Readout time is not an issue due to very low duty cycle and low occupancies
    - High-pt background is not an issue due to clean initial state
  - A traditional first-level trigger is probably not needed
- SLHC – and future energy-frontier machines
  - SLHC is a 2x, then 10x, upgrade to LHC luminosity by ~2015
  - Backgrounds are up to 20x worse than at LHC - 20000 charged tracks per BX
  - Trigging will once again be *the* problem in this environment
  - Imaginative thinking already under way to solve the problems
    - e.g. development of track-based first-level triggering for CMS. Possible? We will see.

Summary

- Trigger functions:
  - Filter very large amounts of detector data to an acceptable rate
  - Keep everything that matters, throw away most of the rest
  - LHC triggering is based around identification of high-pt leptons, photons, jets
    - b-physics experiments also depend upon displaced vertex triggers
- First level triggers:
  - Complex custom hardware systems, mostly digital logic based
  - Carry out parallel algorithms on reduced detector data
    - Identify trigger objects corresponding to idealised leptons, photons, jets, etc
  - A yes/no decision is made based upon a set of trigger criteria
- Triggers must be:
  - Highly selective, efficient, robust, well-understood, controllable
- Triggers are essential for physics, and challenging to build
  - But also fun!