Real-time Processing for HEP

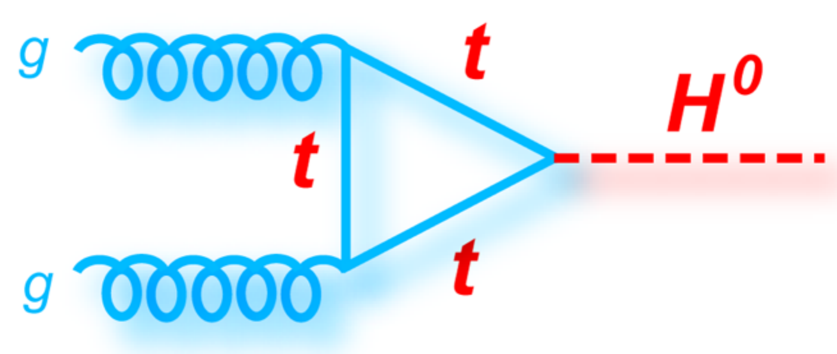

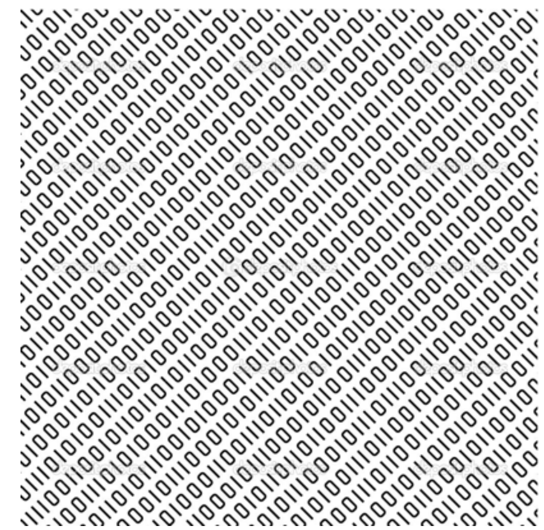

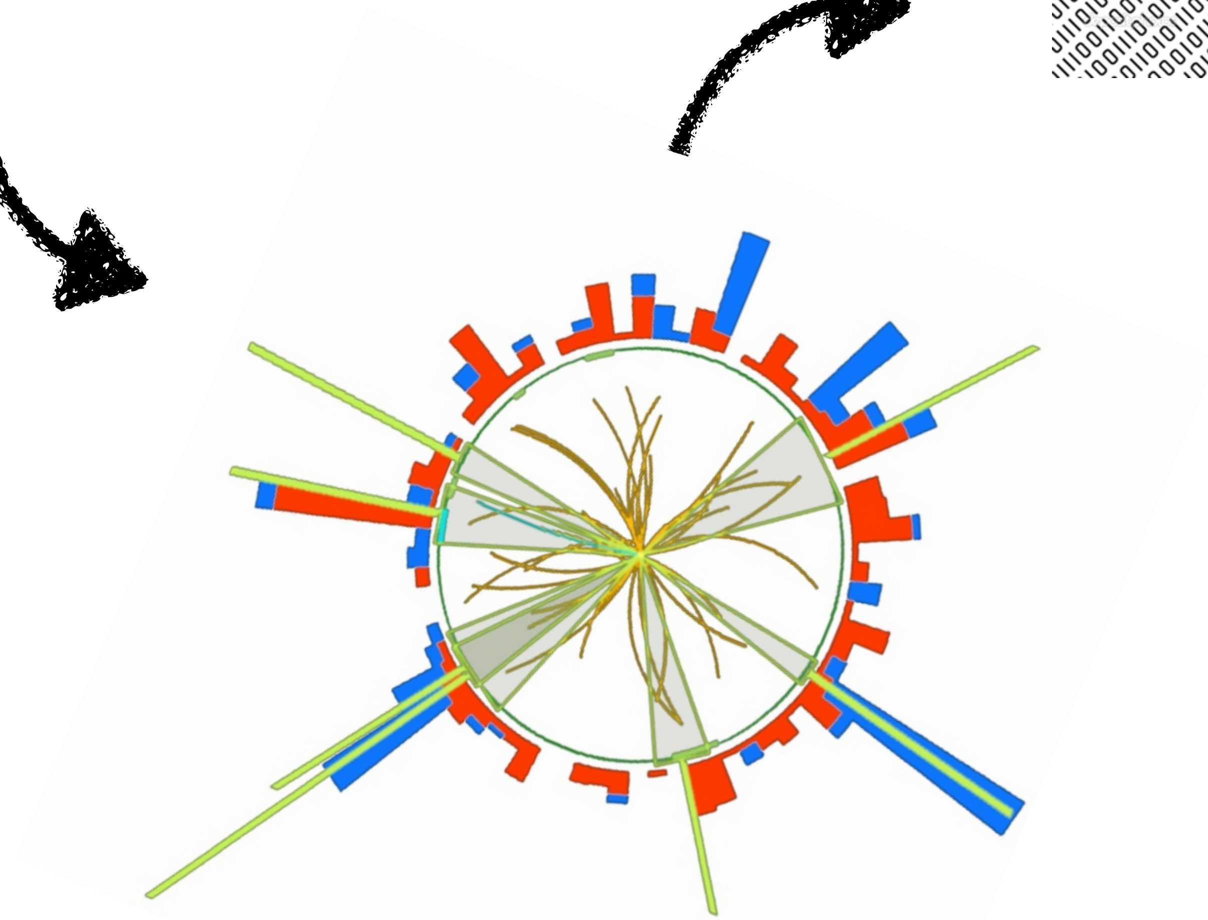

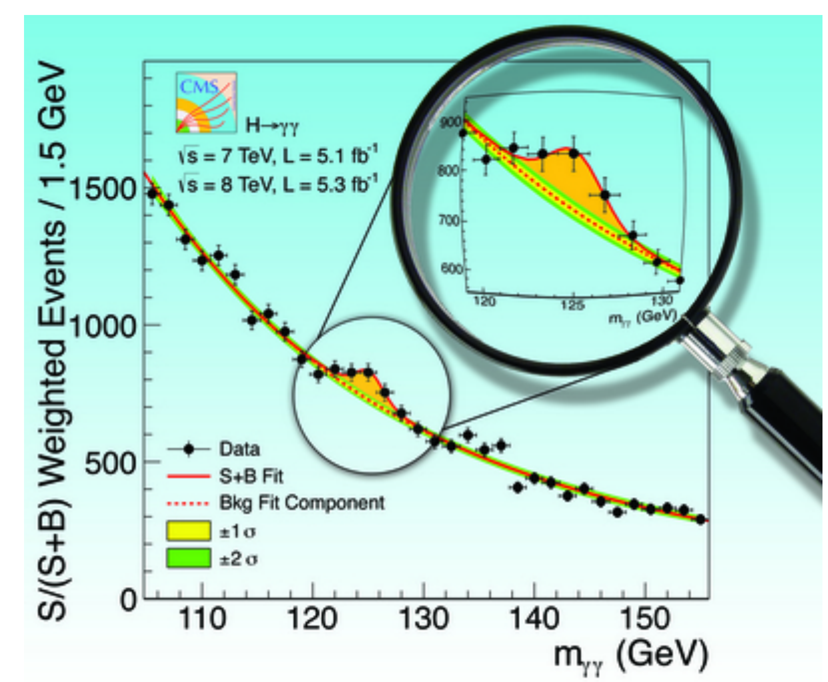

Dave Newbold, Bristol / RAL

Science Motivation

(100k) Haystacks

(interaction probability)

Cross-section

Needle

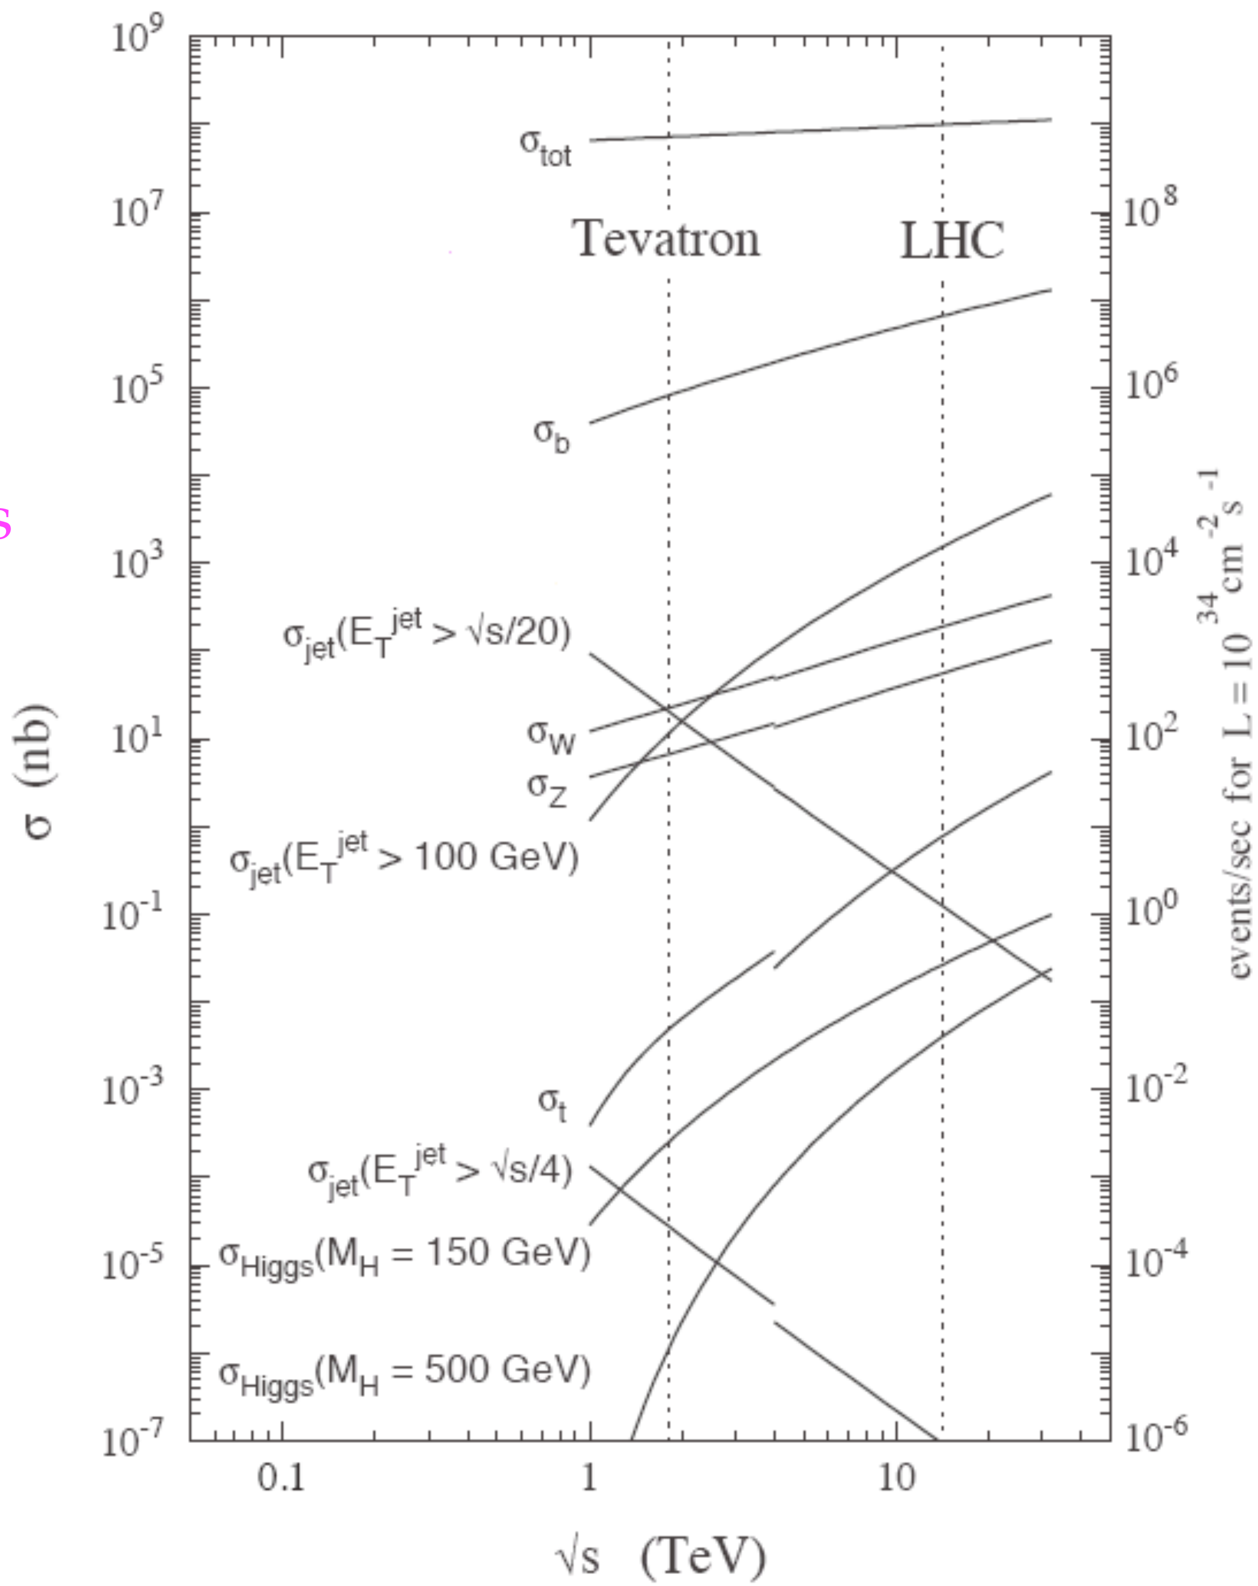

1GHz

1MHz

1kHz

1Hz

1/hr

1/wk

Event rate

2008

2009

2010

2011-2

‘Typical’ LHC Bunch Crossing

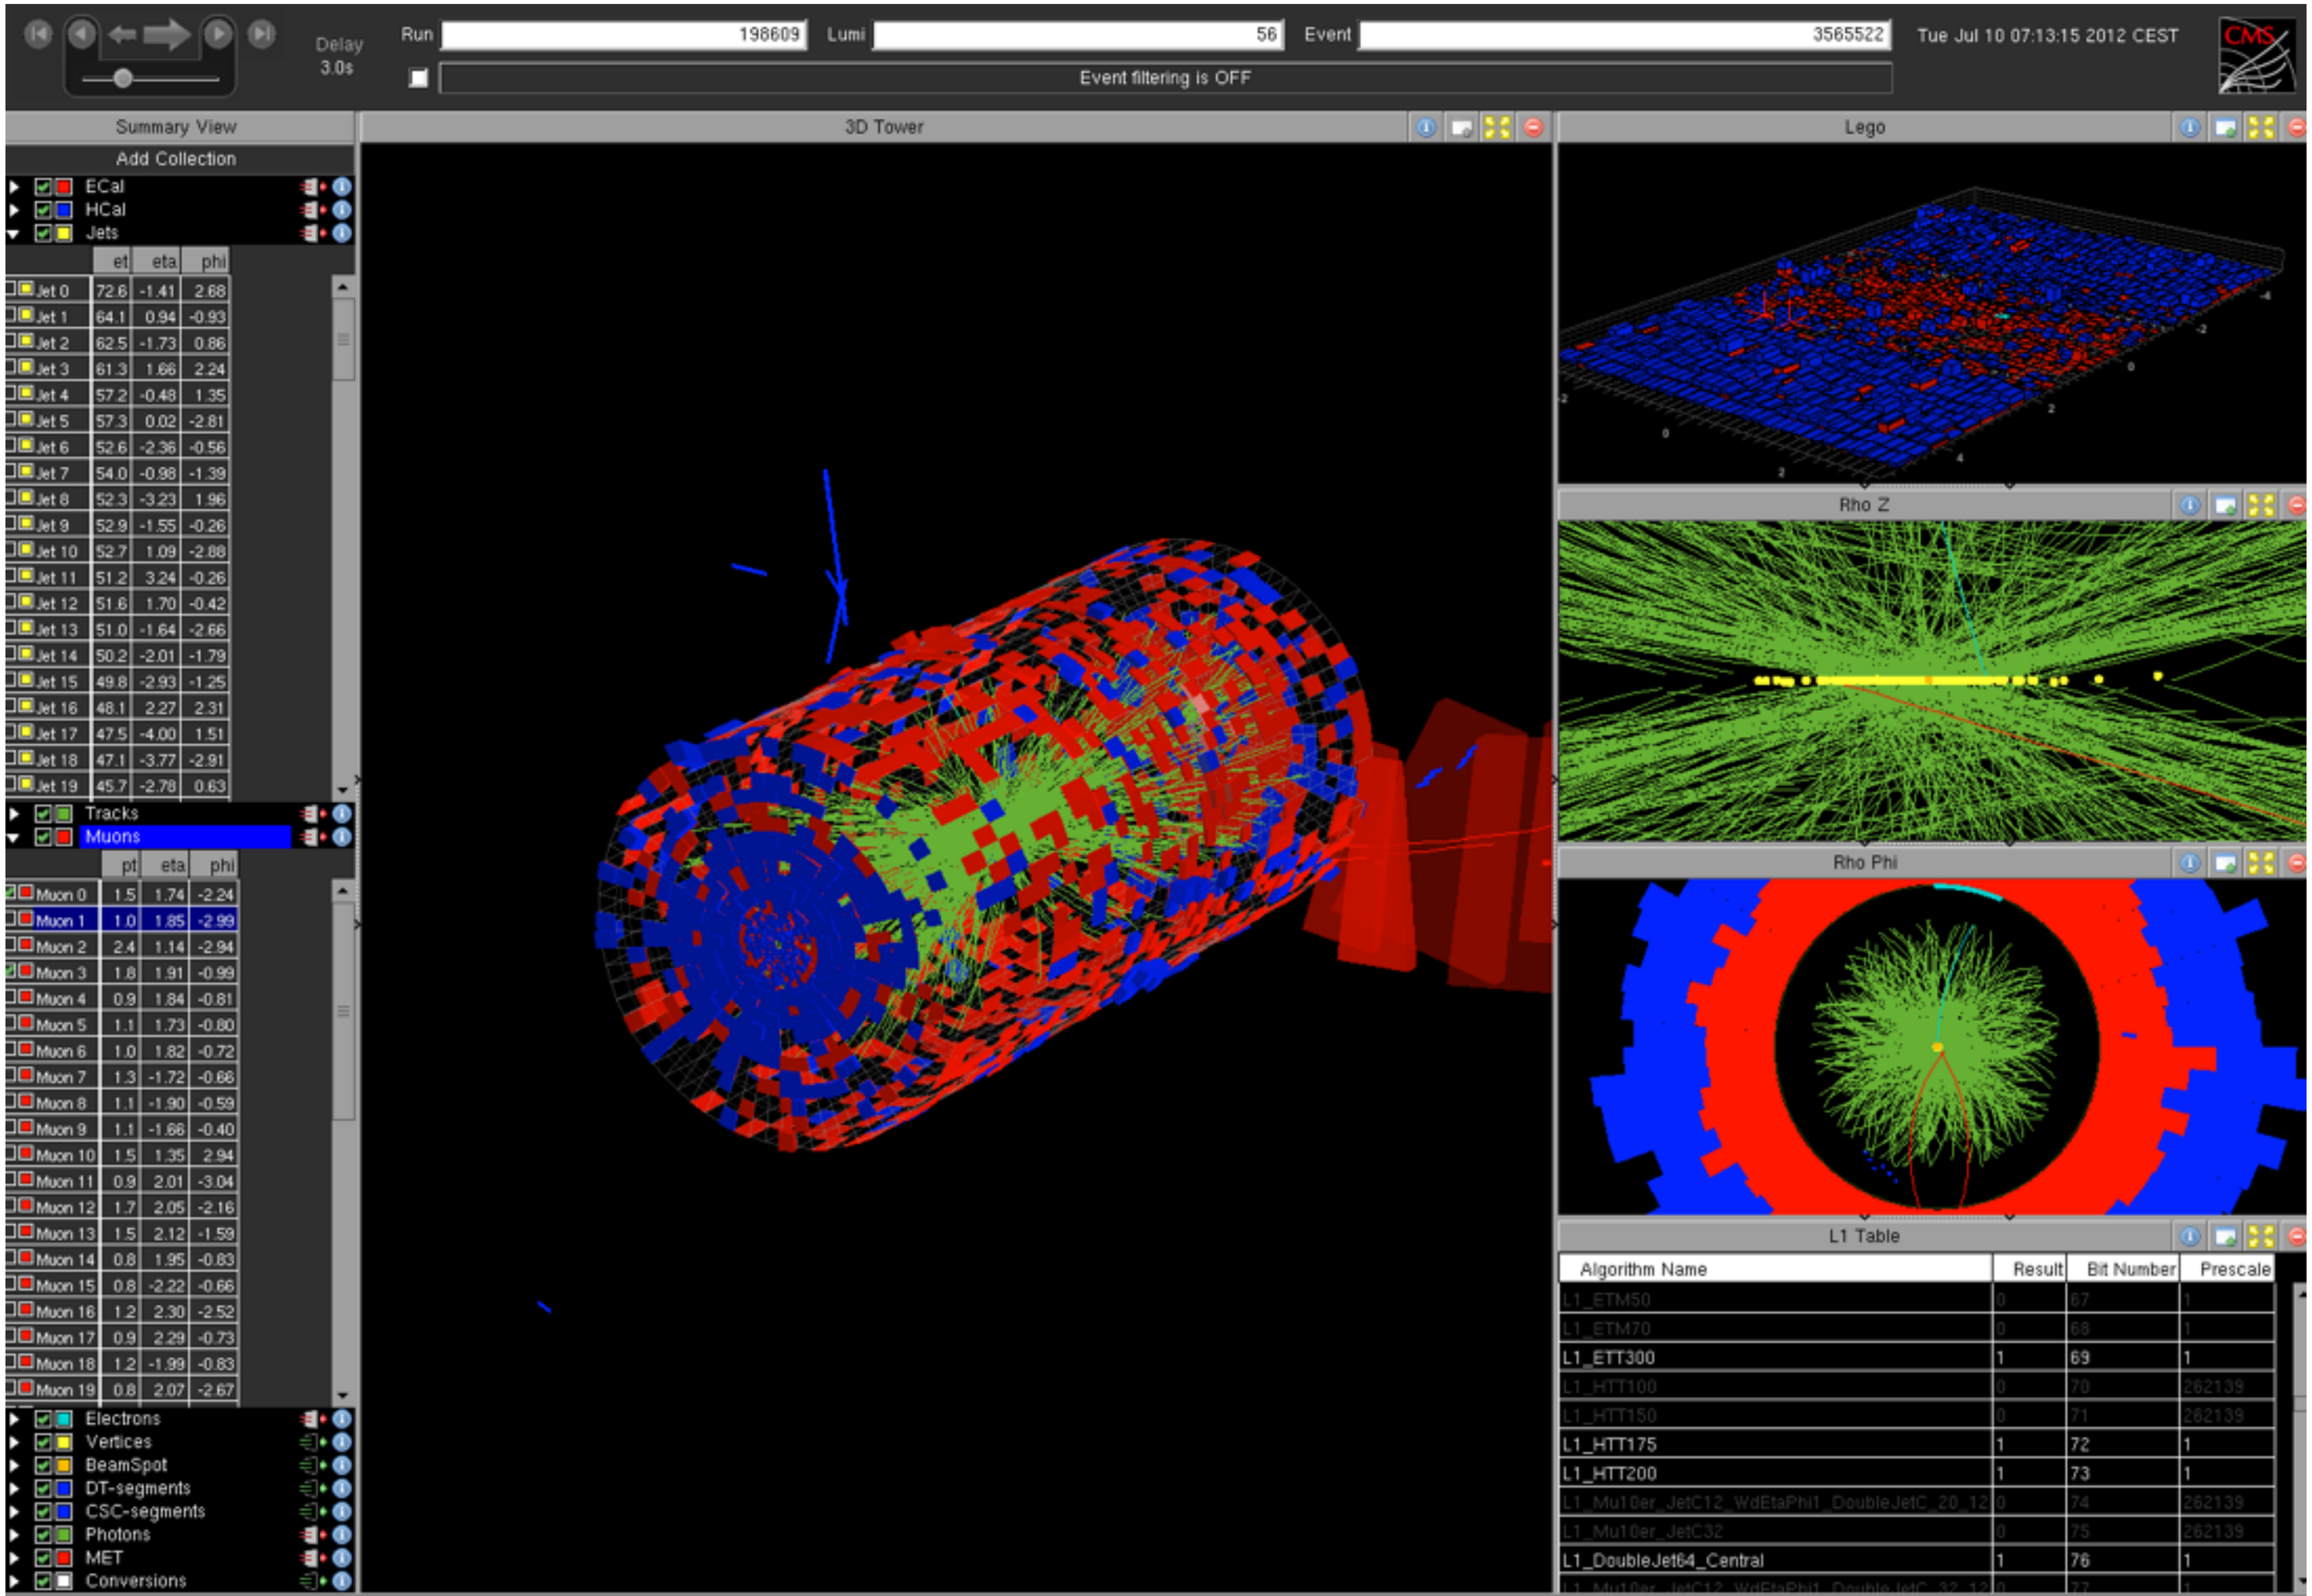

Typical Detector Concept

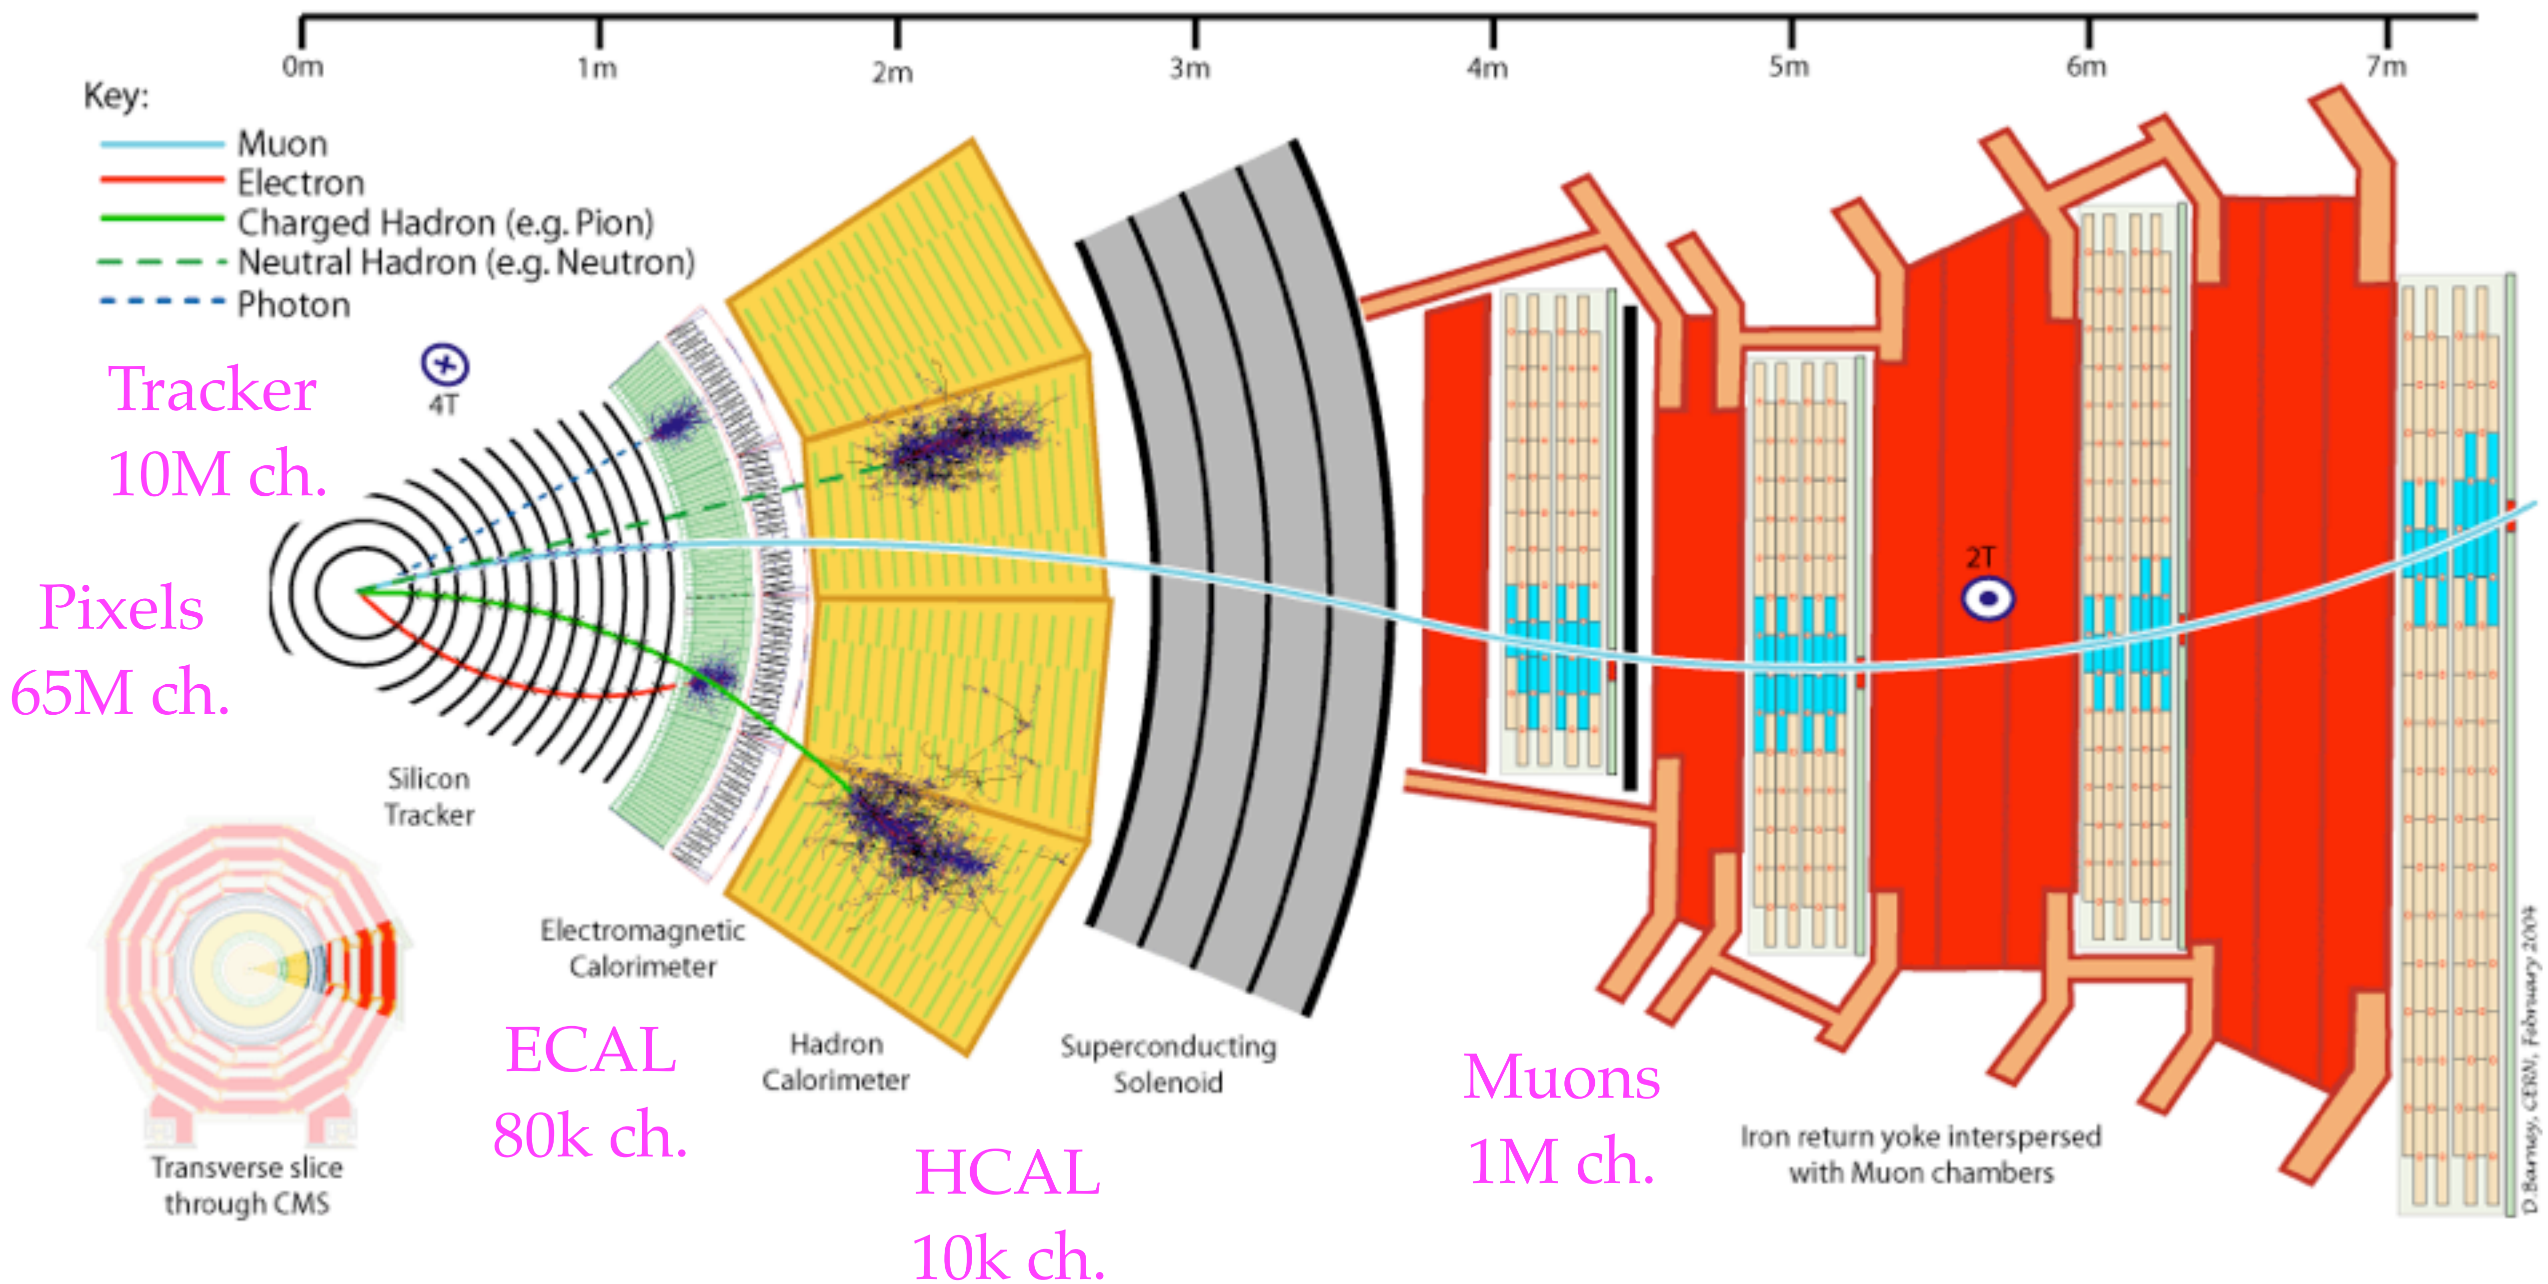

- Modern (LHC) detectors are 20-year, ~$0.5G projects
  - Operate for ~20 year lifetimes with progressive upgrades

Science & Technology Facilities Council Rutherford Appleton Laboratory
University of BRISTOL

Detector Design

- LHC detector mission
  - Find and measure incredibly rare events... (1/hr)
  - Against almost indistinguishable background of common events (1kHz)
  - In an environment of incredibly high-rate background (1GHz)
- Detector characteristics
  - Fast response time
    - Unique crossing-ID required -> 25ns time resolution
  - Large area and hermeticity; lowest possible material for inner detectors
  - High granularity
    - Efficient pattern recognition -> For low occupancy -> 10k's to M's of channels
  - Good resolution, low noise, high dynamic range
    - Energy resolution in calorimetry; (interpolated) position information in tracking
- The environment
  - Highly constrained in terms of space, cooling, access, services
  - Electromagnetically noisy & high radiation dose in places

Science & Technology Facilities Council Rutherford Appleton Laboratory
University of BRISTOL

Data Flow & Triggering

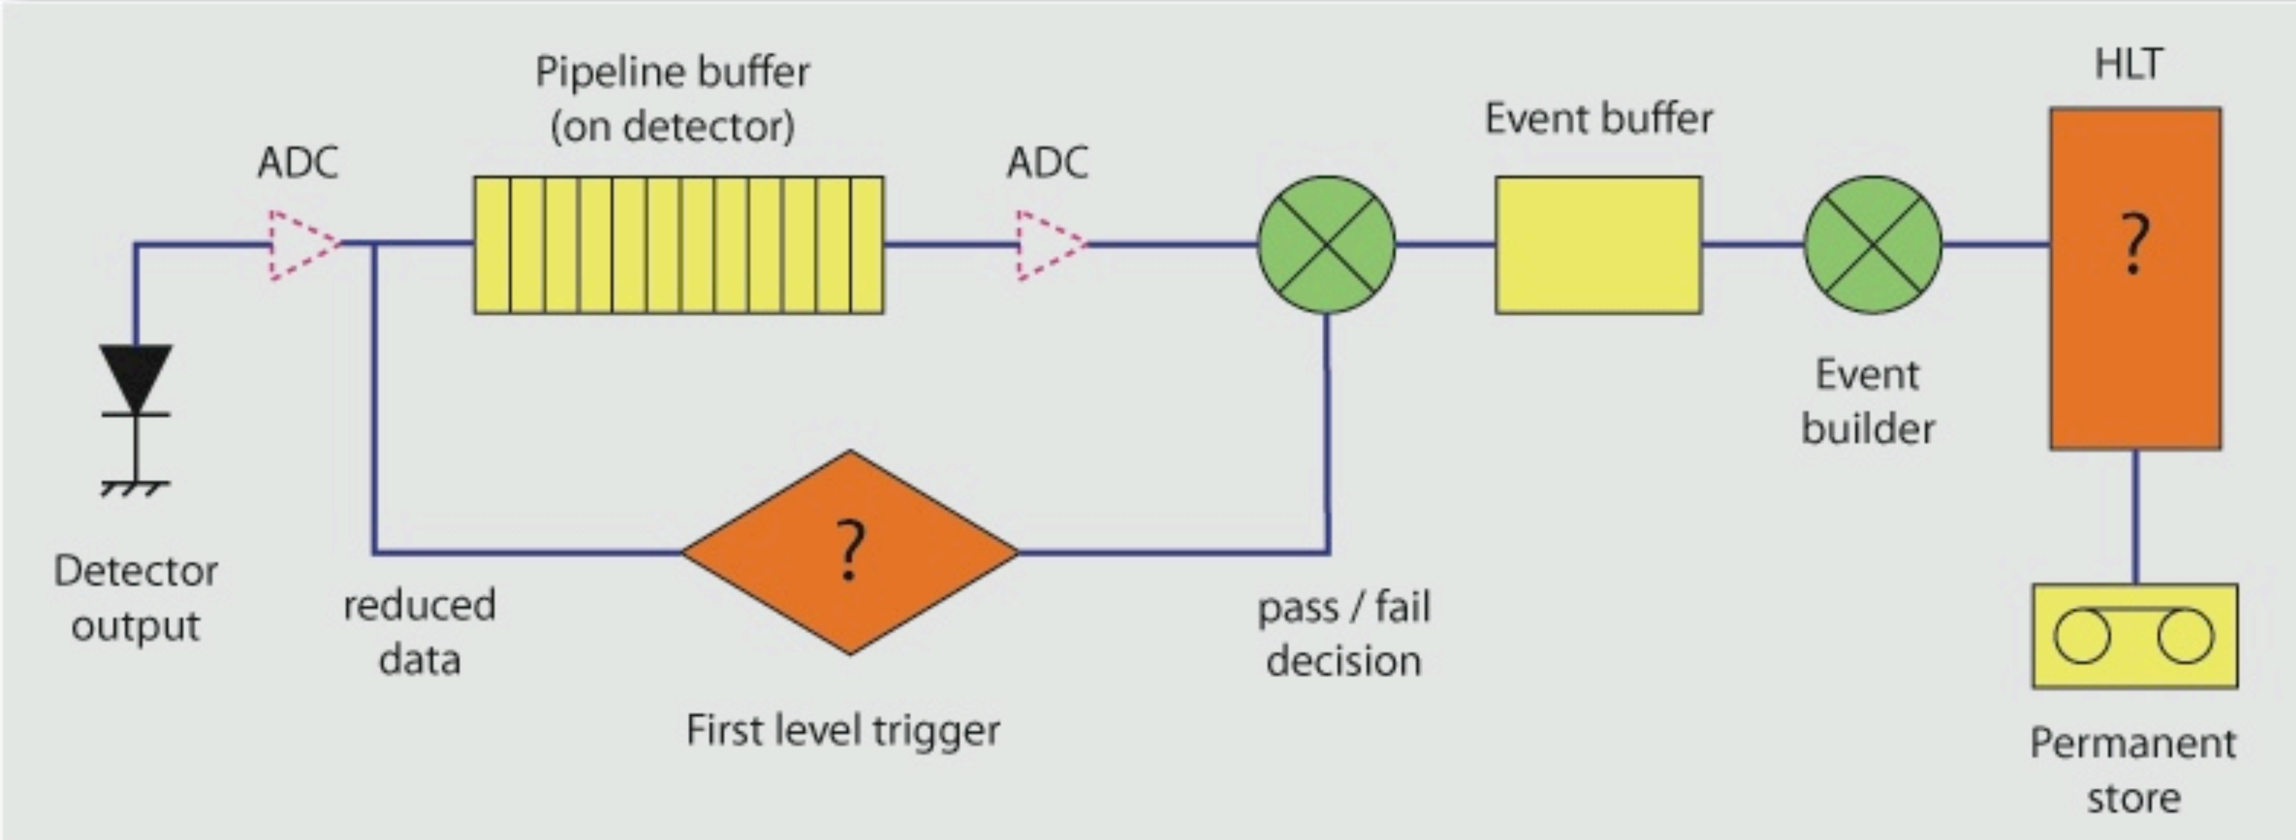

- Impossible to capture all detector output (10's of TB/s)
- Online event selection ('triggering') is required
  - Conceptually part of on-detector system, though usually 'close to' detector
  - Important metrics are: accept rate, efficiency, dead-time
- Events not selected within ~3us are permanently lost

Front-End Electronics

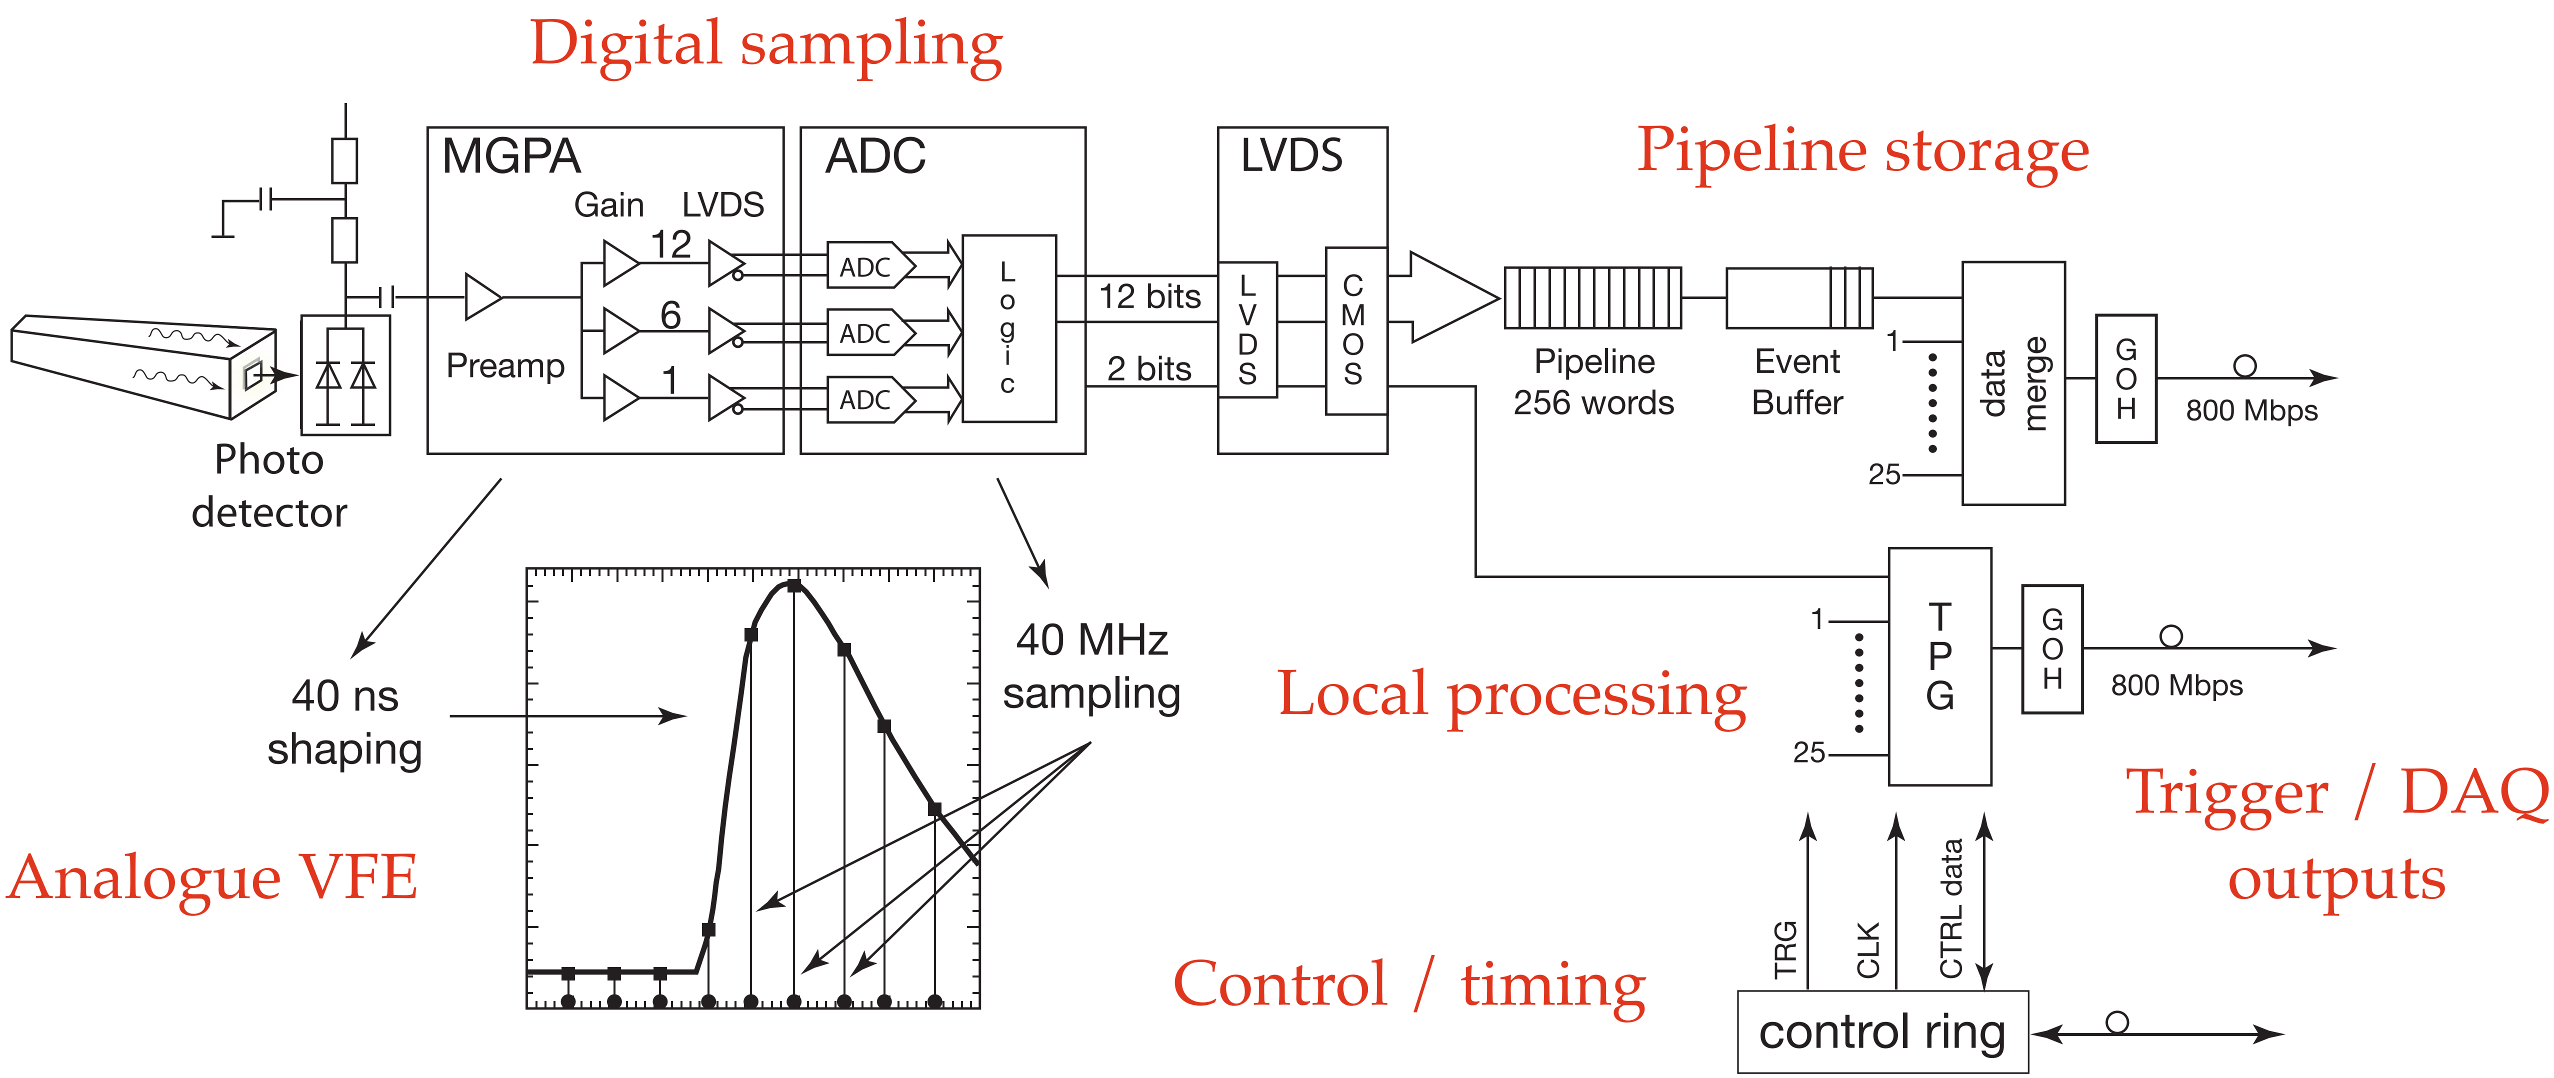

- e.g. CMS ECAL front end electronics (UK development)
  - ~80000 channels, 40Ms/s, 12b resolution, 16 b dynamic range
- Based on two custom rad-hard CMS ASICs, 0.25u technology
  - Along with carefully qualified commercial optoelectronics, sensors

The Reality

- Fully integrated into CMS detector + inaccessible thereafter

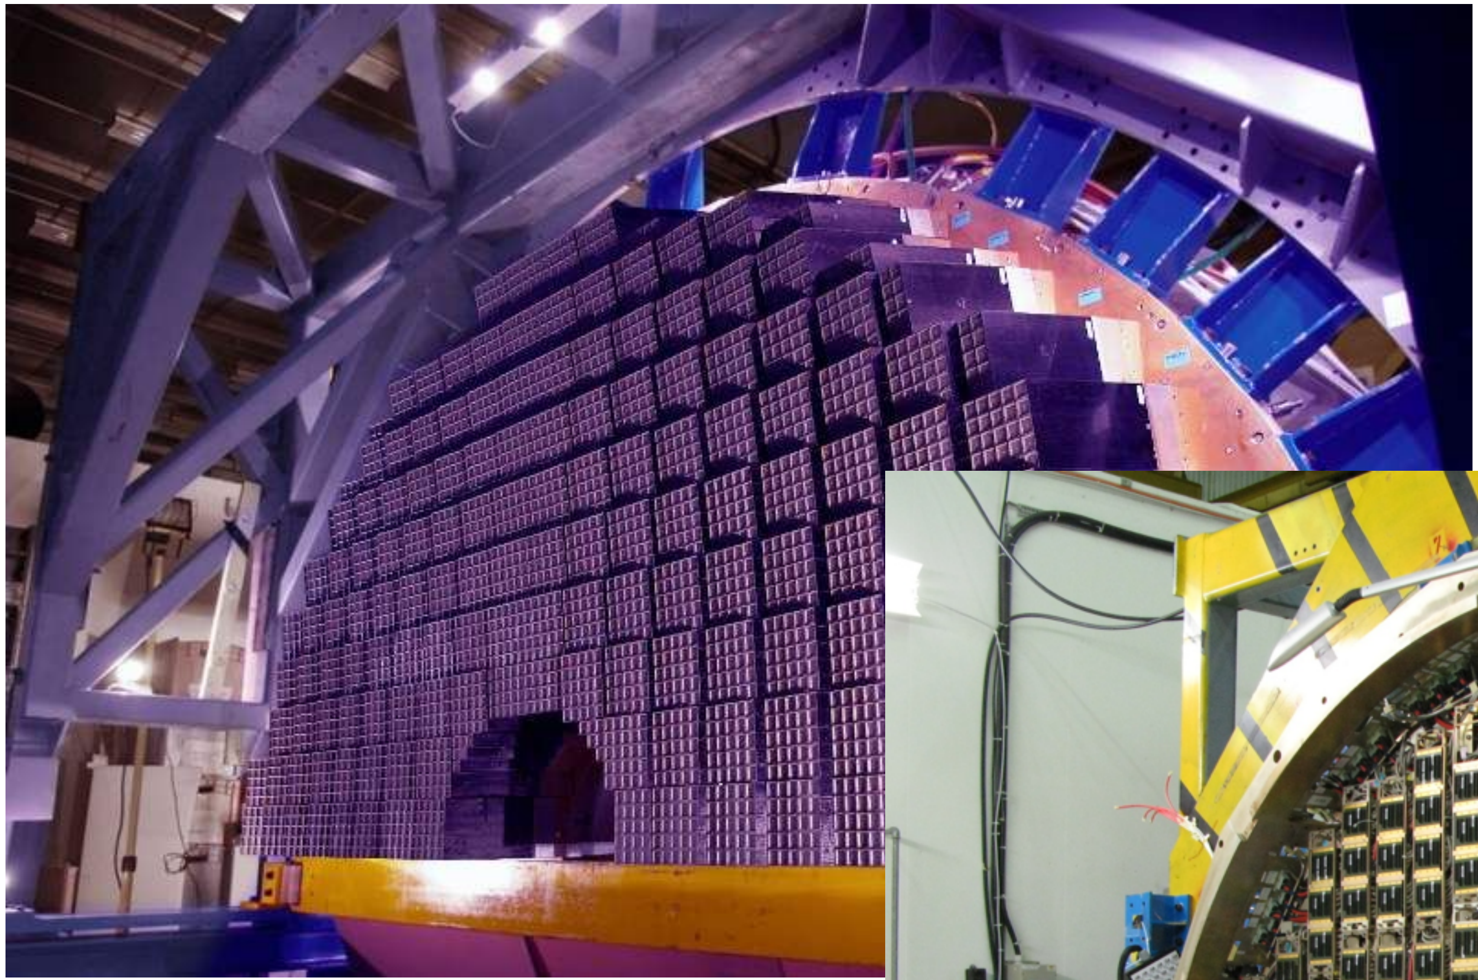

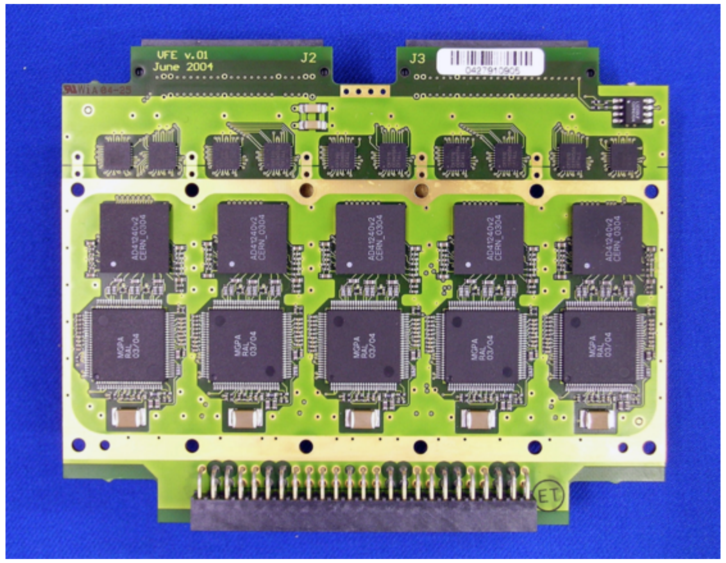

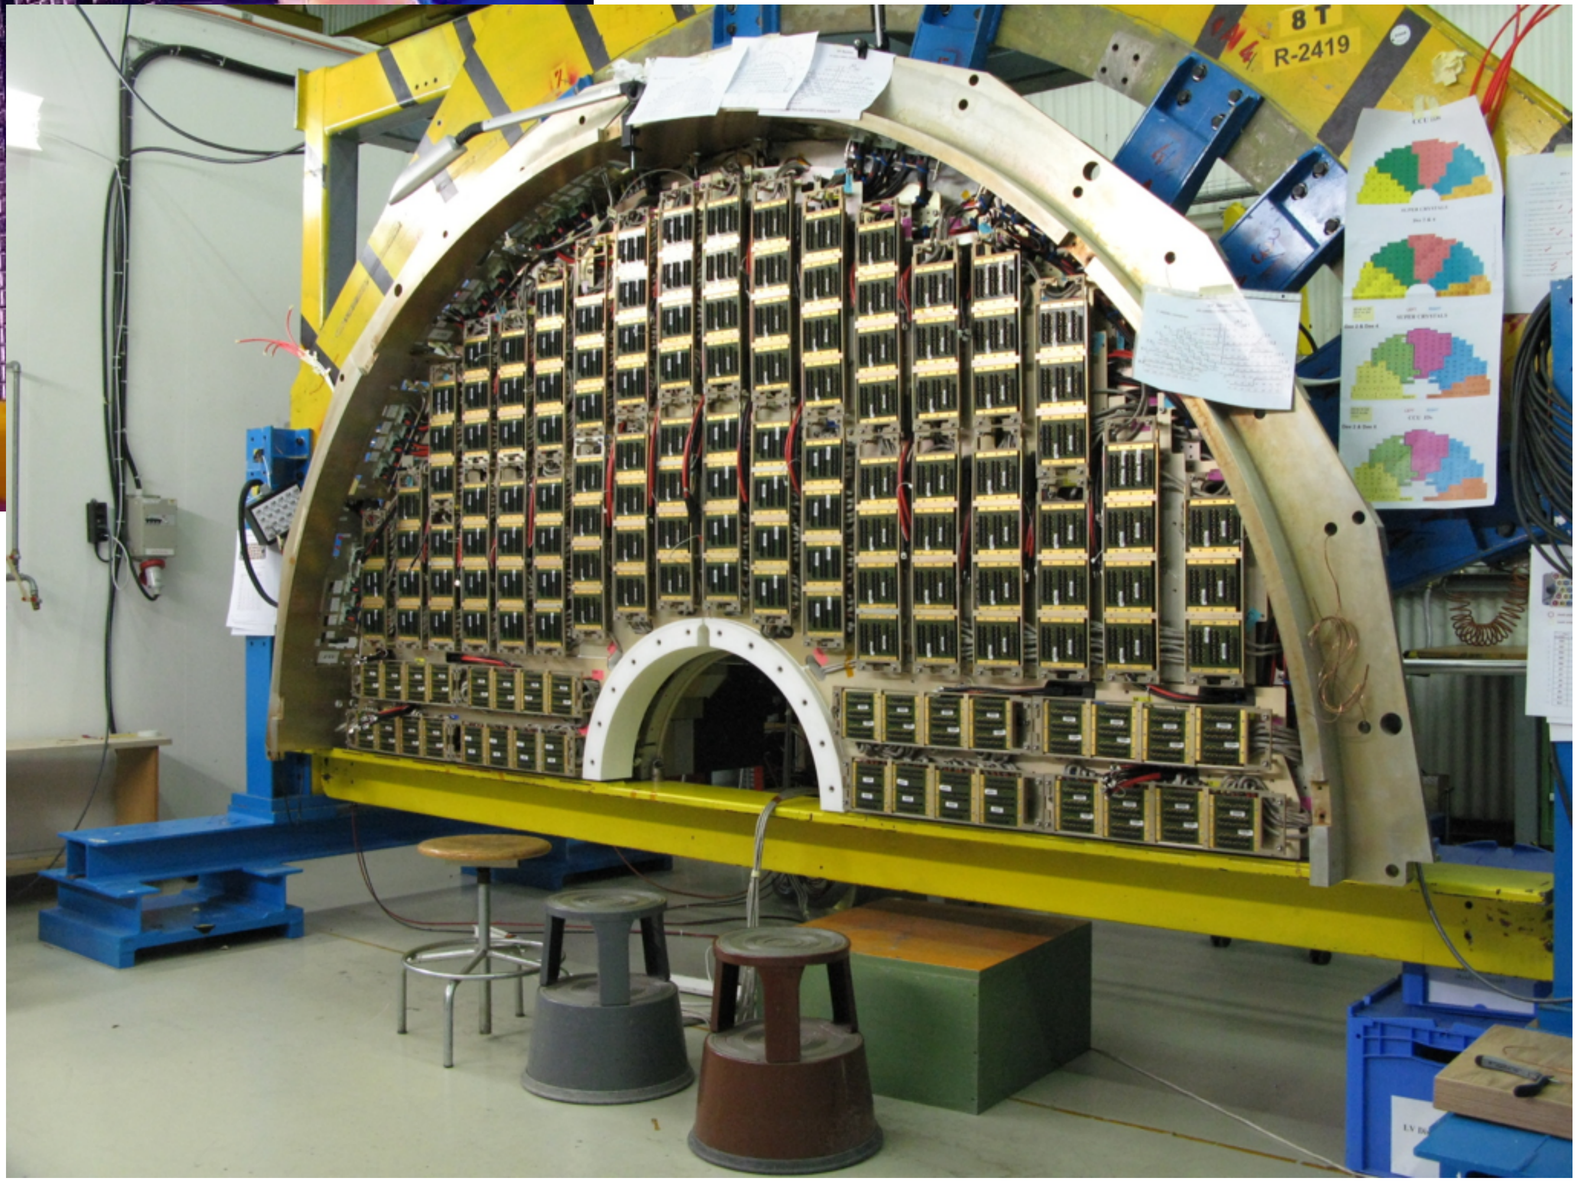

Science & Technology Facilities Council Rutherford Appleton Laboratory
University of BRISTOL

Front-End Electronics

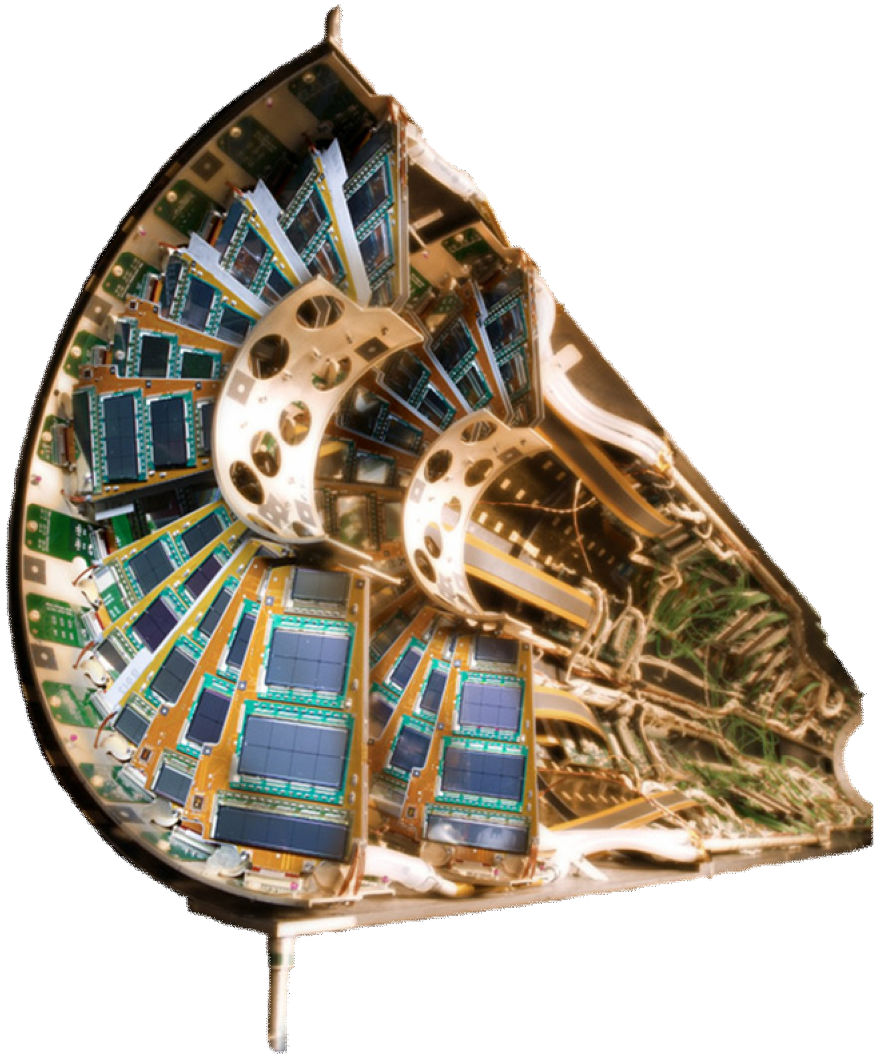

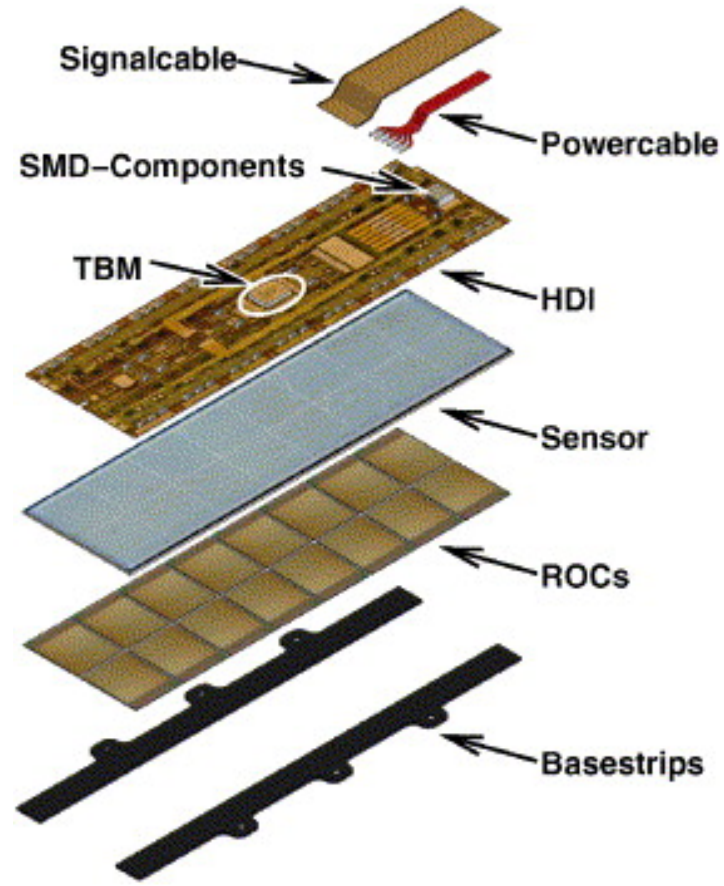

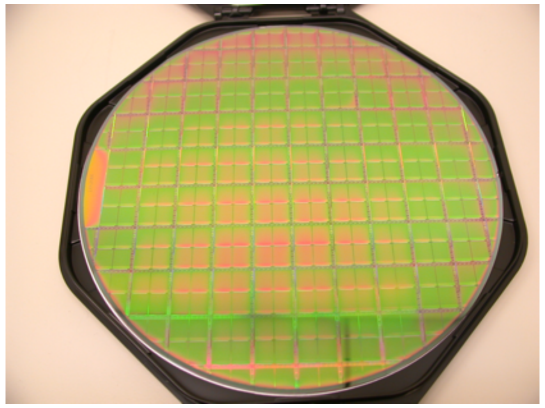

- Other end of the spectrum: pixel ROC
  - Readout is essentially integrated with sensor element
  - Emphasis on huge data reduction & multiplexing on detector
  - Readout bandwidth is the overriding concern (power / space limited)

Front-End Technical Challenges

Performance

- Data reduction / storage capacity
- On-detector local pattern recognition
- Timing accuracy

Power consumption

- Strictly limited in inner detectors
- Detectors are environmentally sensitive

Well-characterised technology

- Acceptable ASIC yields
- Known radiation tolerance

Trade-offs & choices

Cost

- Dictates affordable granularity

Flexibility

- Programmable local processing for the unexpected
- Adaptable for changing backgrounds

Robustness

- Operate for ~decade without intervention

Trigger Functionality

- Mission statement
  - Decide (in hard real time) whether to keep or drop each set of samples
  - Maximise acceptance for interesting physics; minimise rate of background
  - Provide means of monitoring and checking performance
- In practical terms
  - Hardware processor filtering the event stream based on a ‘quick look’
  - Can make use of a very limited subset of the recorded data
- This is a tricky business
  - Physics performance of experiment is dictated by performance of trigger!
    - Especially at hadron collider experiments – 99% of event sample selection is done in real time
  - No room for error, as discarded events are gone for ever
    - Accelerator time costs ~$M per day
  - We usually do not know exactly what to expect in advance
    - We are searching for the unknown, often against unknown background
  - Technology is usually at the limit of what can be done

Trigger Algorithms I

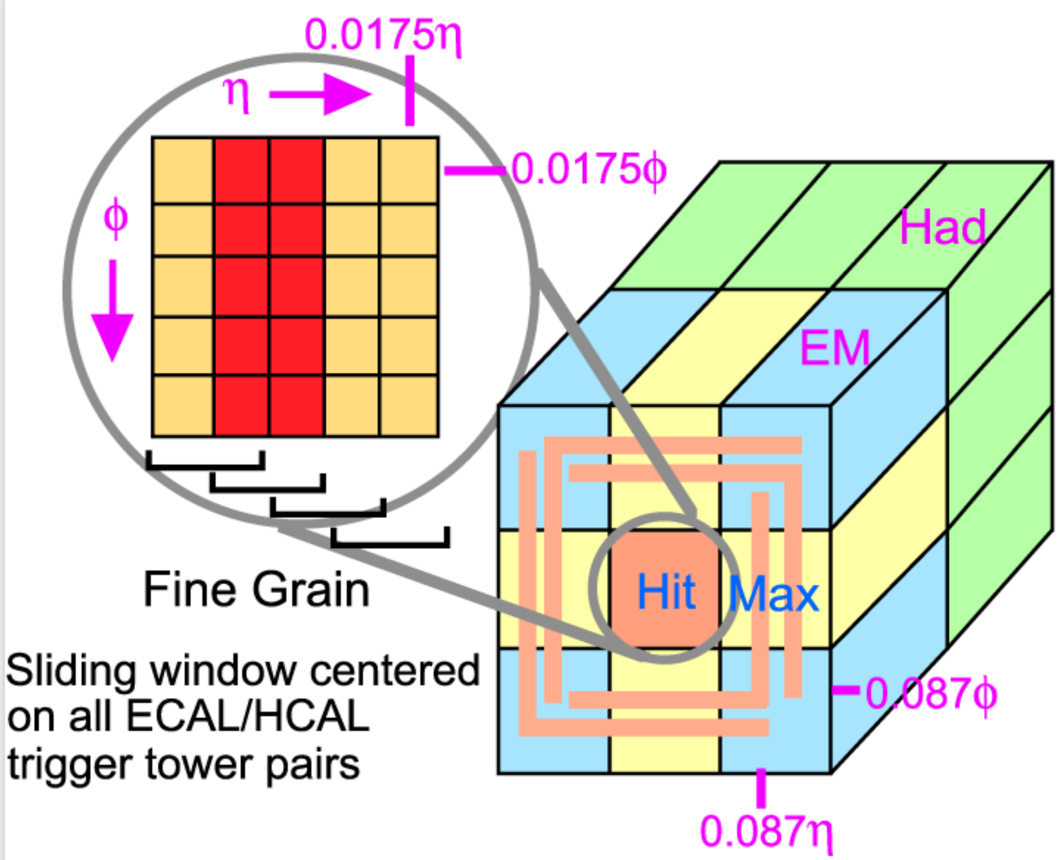

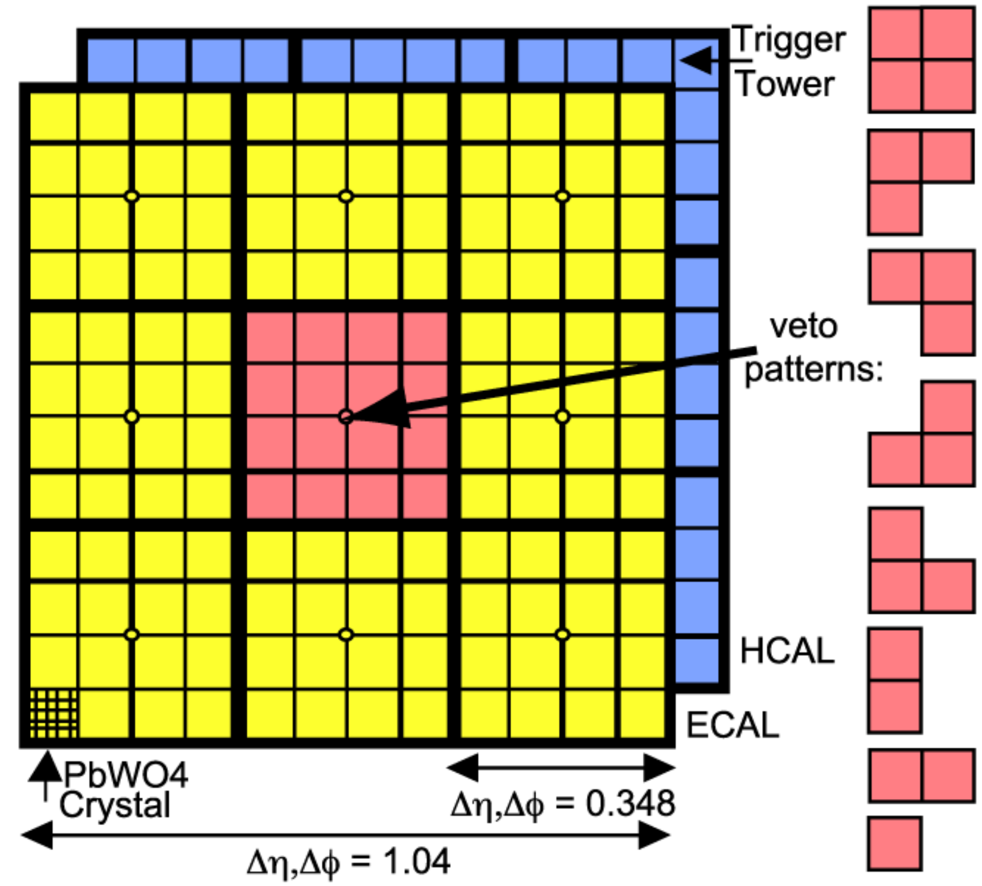

e/γ (hit tower + max neighbour):
- 2-tower Et; hit tower passes H/E cut
- Hit tower: 2x5 strip with >90% Et in 5x5 (FG)

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

- e.g. CMS calorimeter trigger – electron / photon ID

Science & Technology Facilities Council Rutherford Appleton Laboratory University of BRISTOL

Trigger Algorithms II

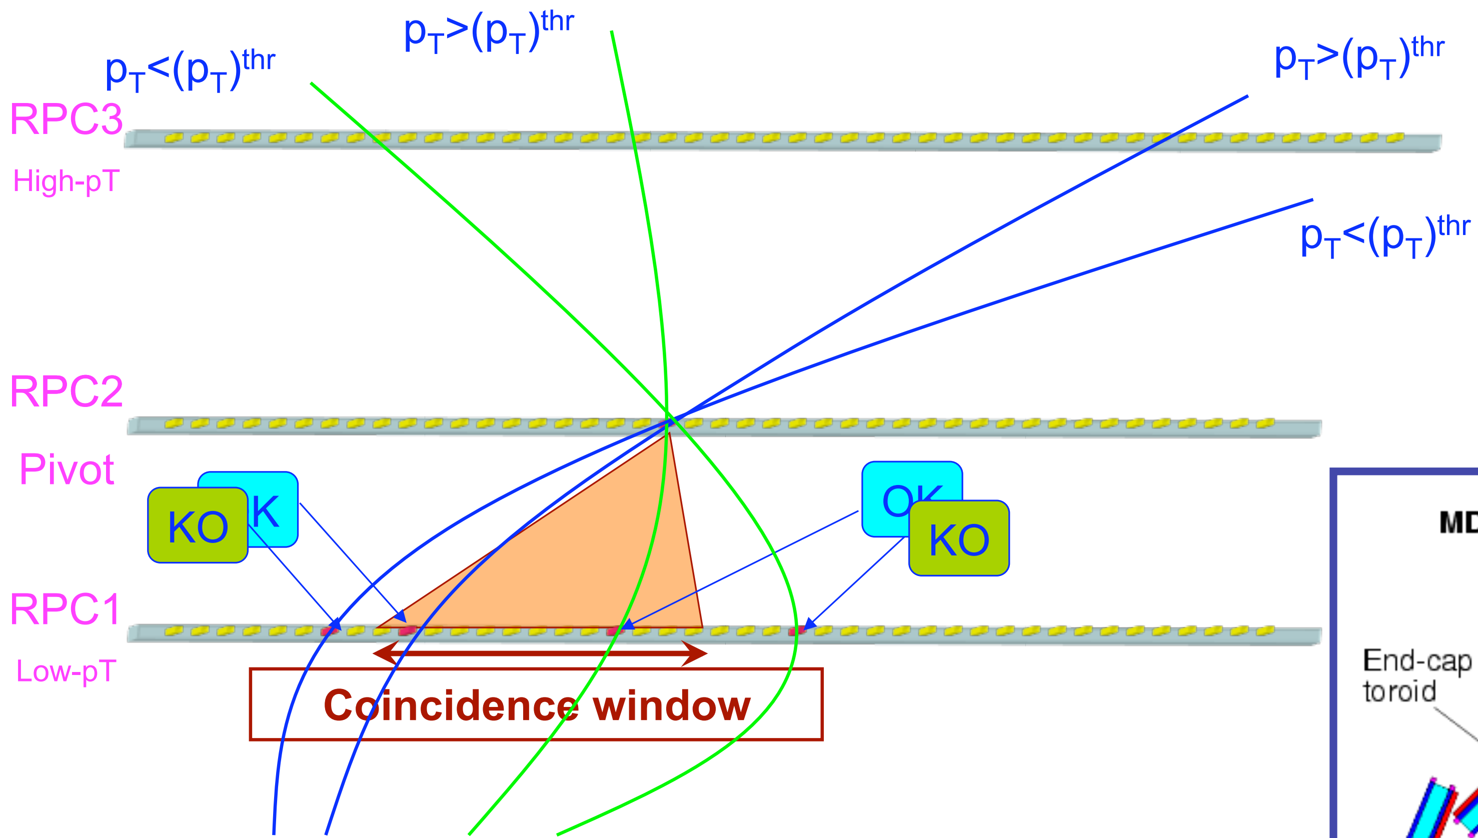

- e.g. ATLAS barrel muon

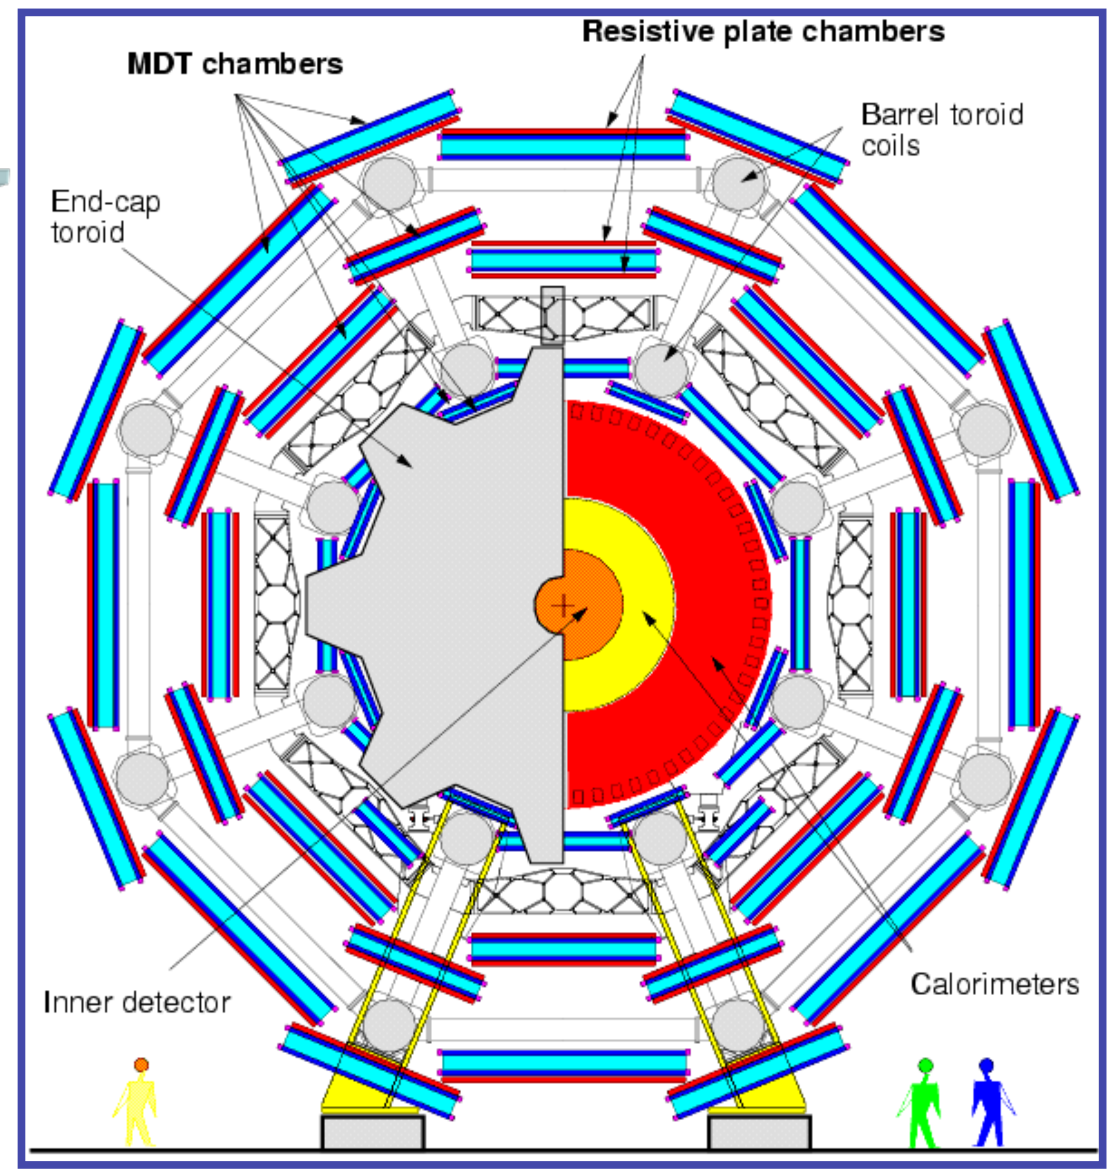

- Not as simple as it looks!
  - Hit correlation in 4D is necessary
  - Muon detector spacing is large compared to time-of-flight
  - Detectors with very good time resolution required for bunch-crossing assignment

Trigger Algorithms III

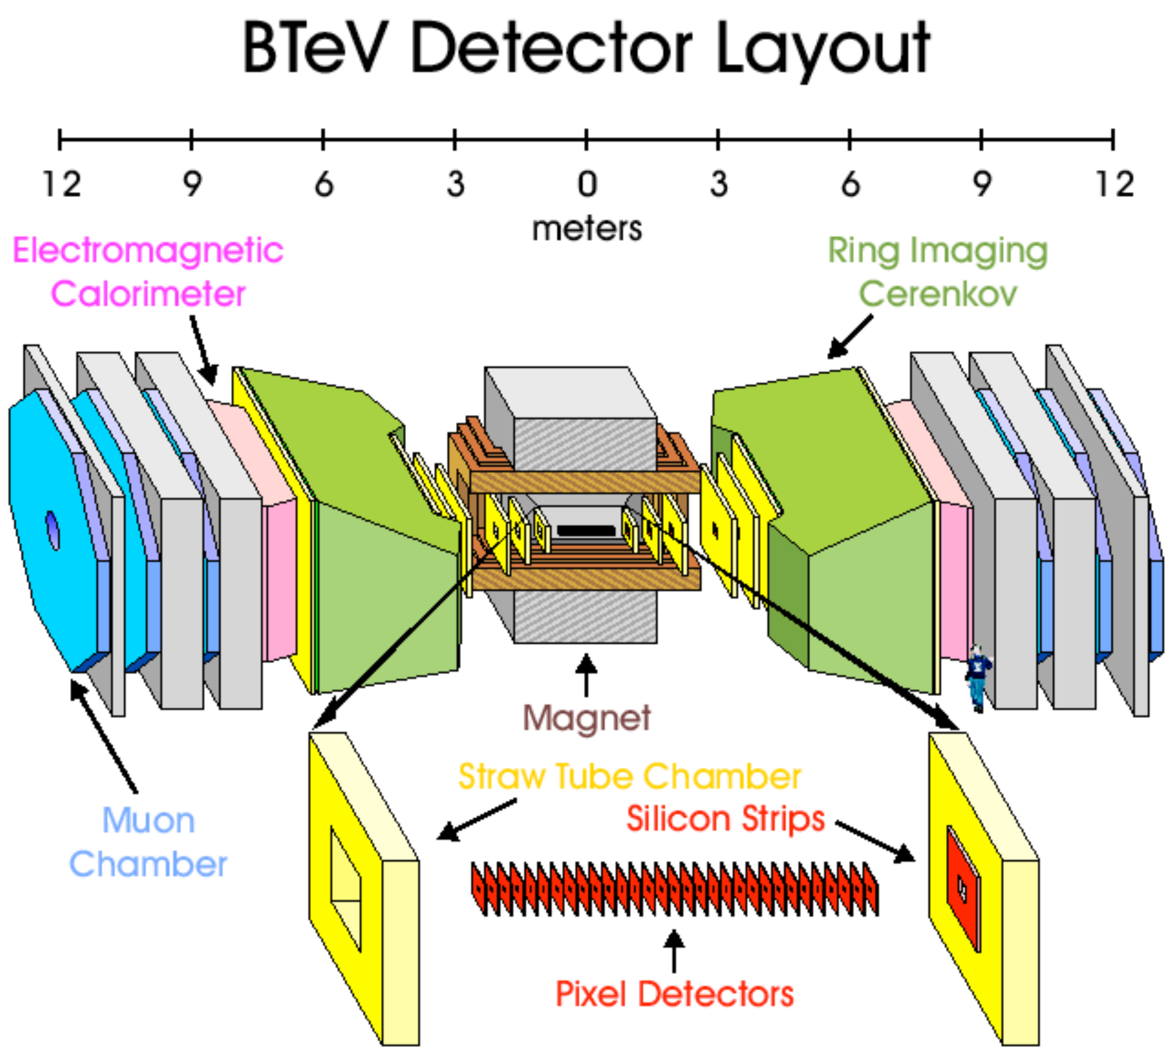

- BTeV pixel trigger
  (with apologies to LHCb)

Blue segments are ‘entering’ detector

Green segments are ‘leaving’ detector

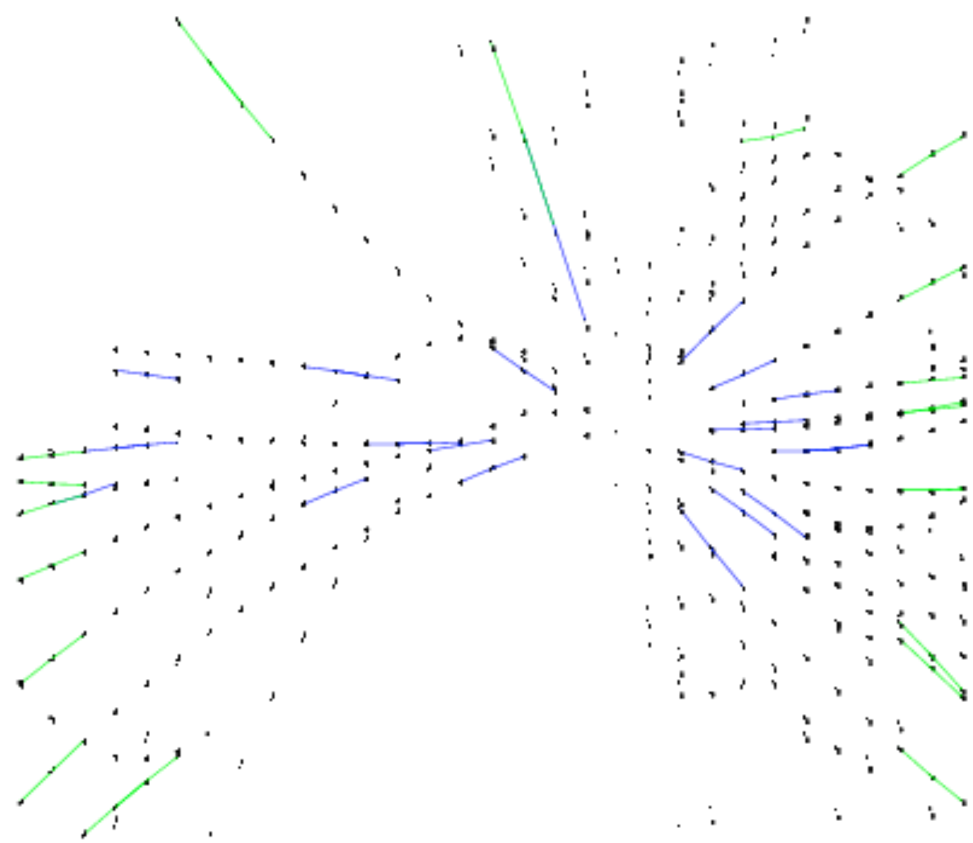

- Based upon triplet-finding approach - rather neat
- Finds number of displaced vertices
- Rejects pile-up and high-multiplicity events

Technology Challenges

- The key problem
  - Data reduction, so that a single yes / no decision can be made
  - Building a trigger is essentially a problem in reliable data transfer
  - State-of-the-art systems have 10-100Tb/s input rates
- In numbers
  - 10-100Tb/s input rates (for next generation of trigger systems)
  - 1us processing time; <0.5% deadtime; $10M cost envelope
  - Note that this rules out all current general-purpose processors
- Technology choices
  - Analogue vs digital? Centralised vs distributed?
  - Data transmission: electrical vs optical, serial vs parallel?
  - Processing elements: custom ASIC vs FPGA vs DSP?
    - Construction of 1Mloc firmware systems is a significant challenge
  - Timing: mesochronous vs plesiosynchronous vs asynchronous?
    - Distribution of <1ns-accurate, <100ps jitter accelerator-locked clock is non-trivial

The Big Picture

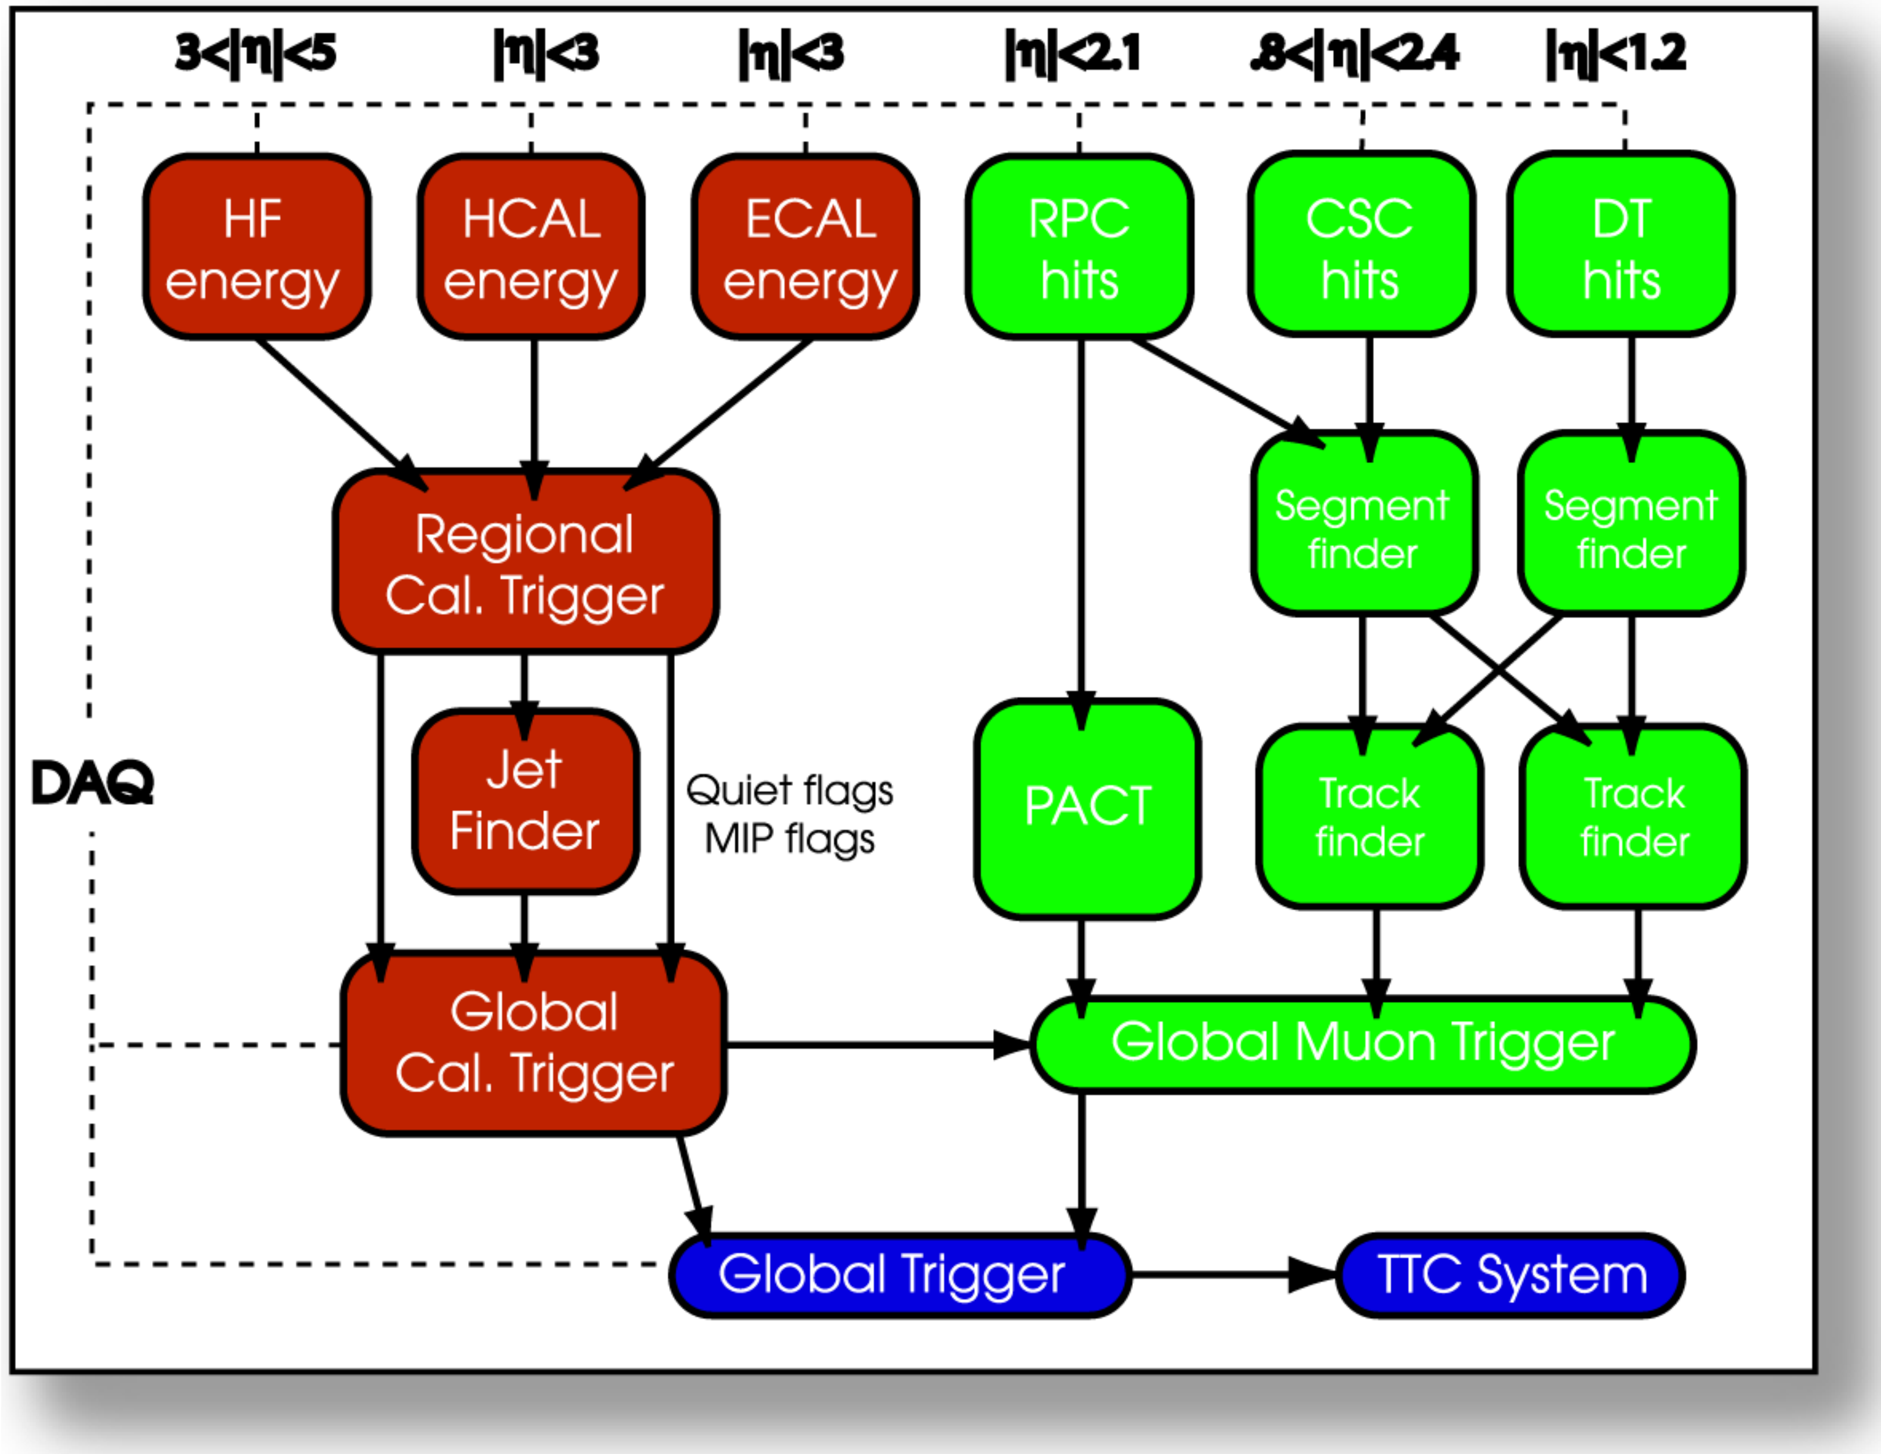

Example Trigger Menu

```
   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
```


Science & Technology Facilities Council Rutherford Appleton Laboratory
University of BRISTOL

Sometimes it Even Works

- ~60 seconds after first LHC collisions, 2009

Future Developments

- HEP front-end / trigger is the state of the art
  - The most complex electronic systems yet assembled for science
  - This may not be true for much longer (see astro talks at this school!)
- What next for LHC?
  - Progressive upgrades, culminating in up to 10x average collision rate
  - Triggering with 10x background will be *the* problem
  - New concepts being developed (in INFIERI project and elsewhere)
    - Centre around tracking for L1 trigger – order-of-magnitude data rate increase
- Other facilities: ILC
  - (up to) 1TeV electron-positron collider on 2030 timescale
  - Benign environment removes need for fast L1 trigger a la LHC
  - Front-end and readout is a much more complex problem however
  - Focus will move to cost-optimisation of very complex readout architecture

SLHC Environment & Hardware

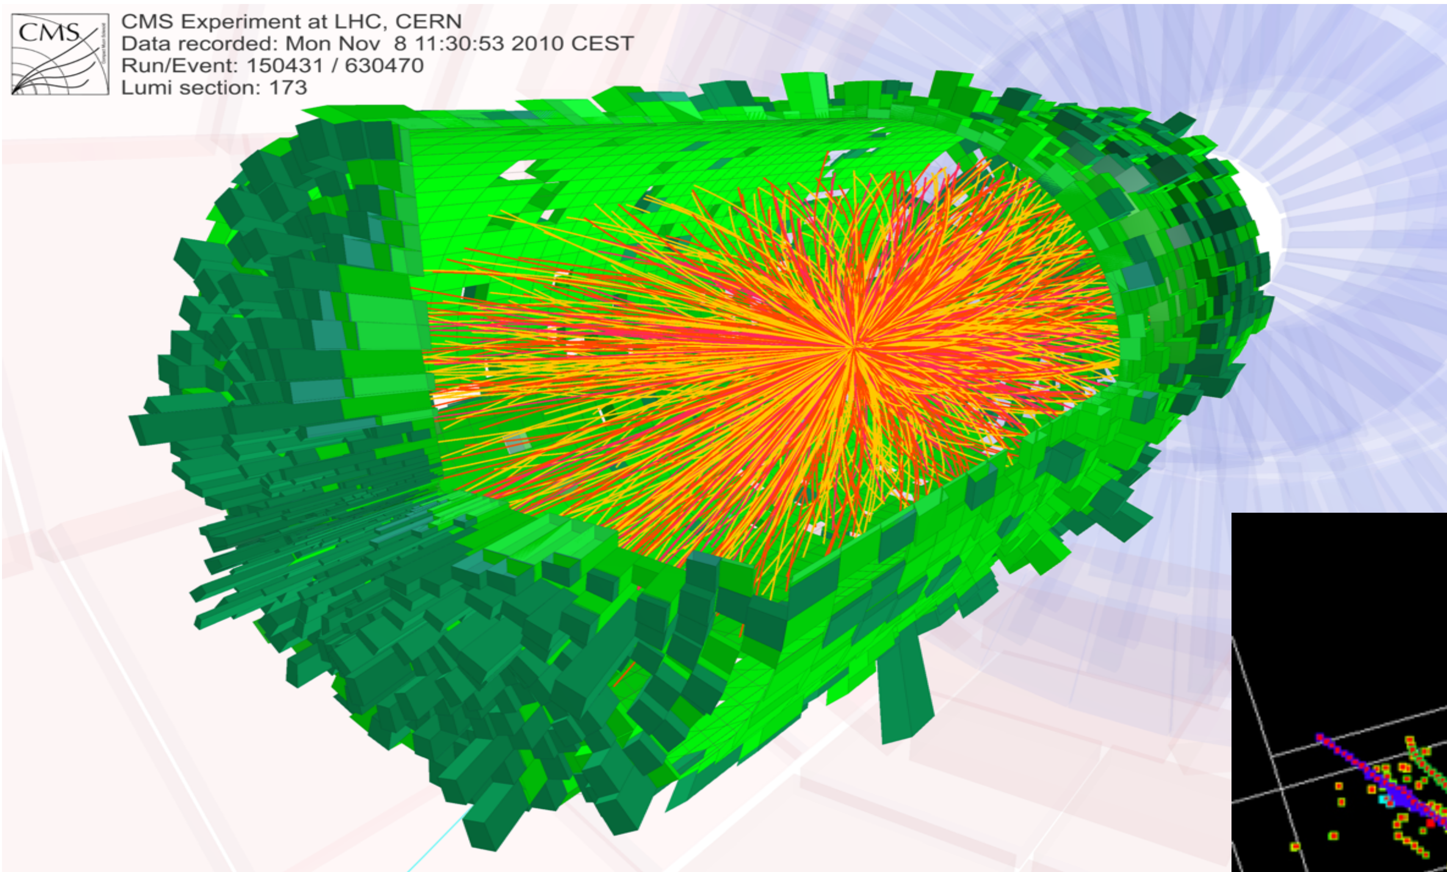

CMS heavy-ion collision
Track density similar to SLHC

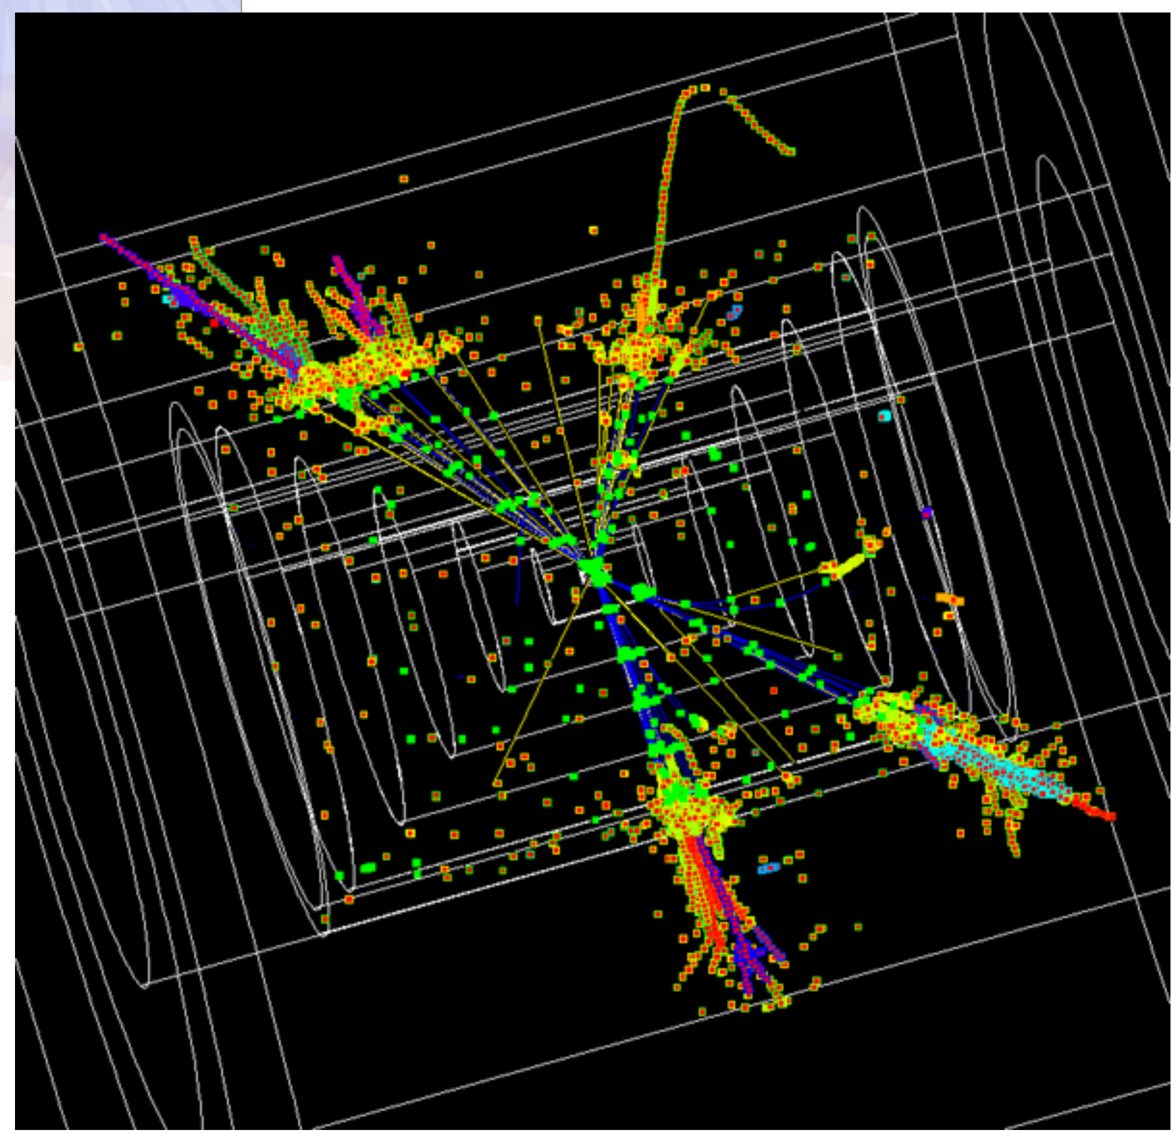

ILC tracking calorimeter
10^{12} channels!!!

('Terrorpixel'?)

State of the Art - ?

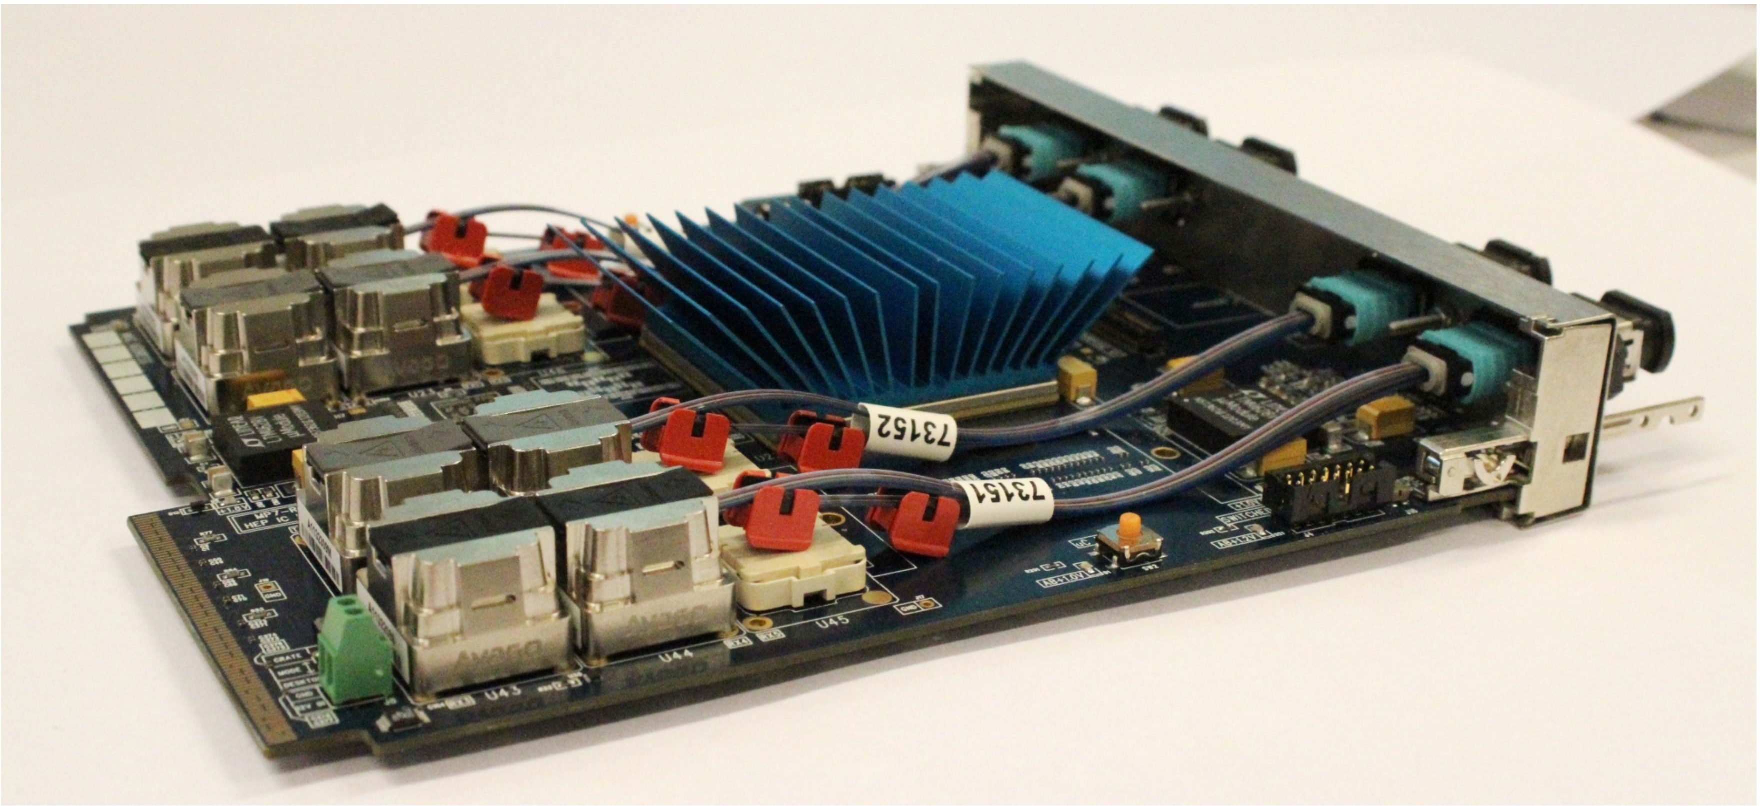

- MP7 card: building block for L1 and pixel systems
  - Large Virtex-7 series FPGA (6B transistors); 144Mb fast RAM
  - 1.4Tb/s of low-latency IO on optical links; 50Gb/s backplane IO
  - Integrated into industry-standard uTCA software / hardware environment
- Will future L1 / FE look more like a commercial switch fabric?
  - This is what happened for the last generation of event builders

Conclusions

- Real-time processing is a vital component of modern HEP
  - All modern experiments use high performance digital readout systems
  - Real-time event filtering is a key aspect of physics analysis
- Technical challenges
  - Design and optimisation of such systems is a tough 10-year task
  - Many difficult technical constraints, unique to HEP environment
  - Failure to meet specifications results in degradation of science output
- Still significant work to come
  - Need to repeat success of current concepts at upgraded LHC
  - The ILC will bring a new set and different of challenges
- A great place to work as a student / postdoc
  - Difficult technical problems needing novel solutions
  - Requires continuous interplay of technical and physics insight
  - Ideal forum to learn fast and make a contribution!

Backup

Science & Technology Facilities Council Rutherford Appleton Laboratory
University of BRISTOL

Stacked Tracking Concept

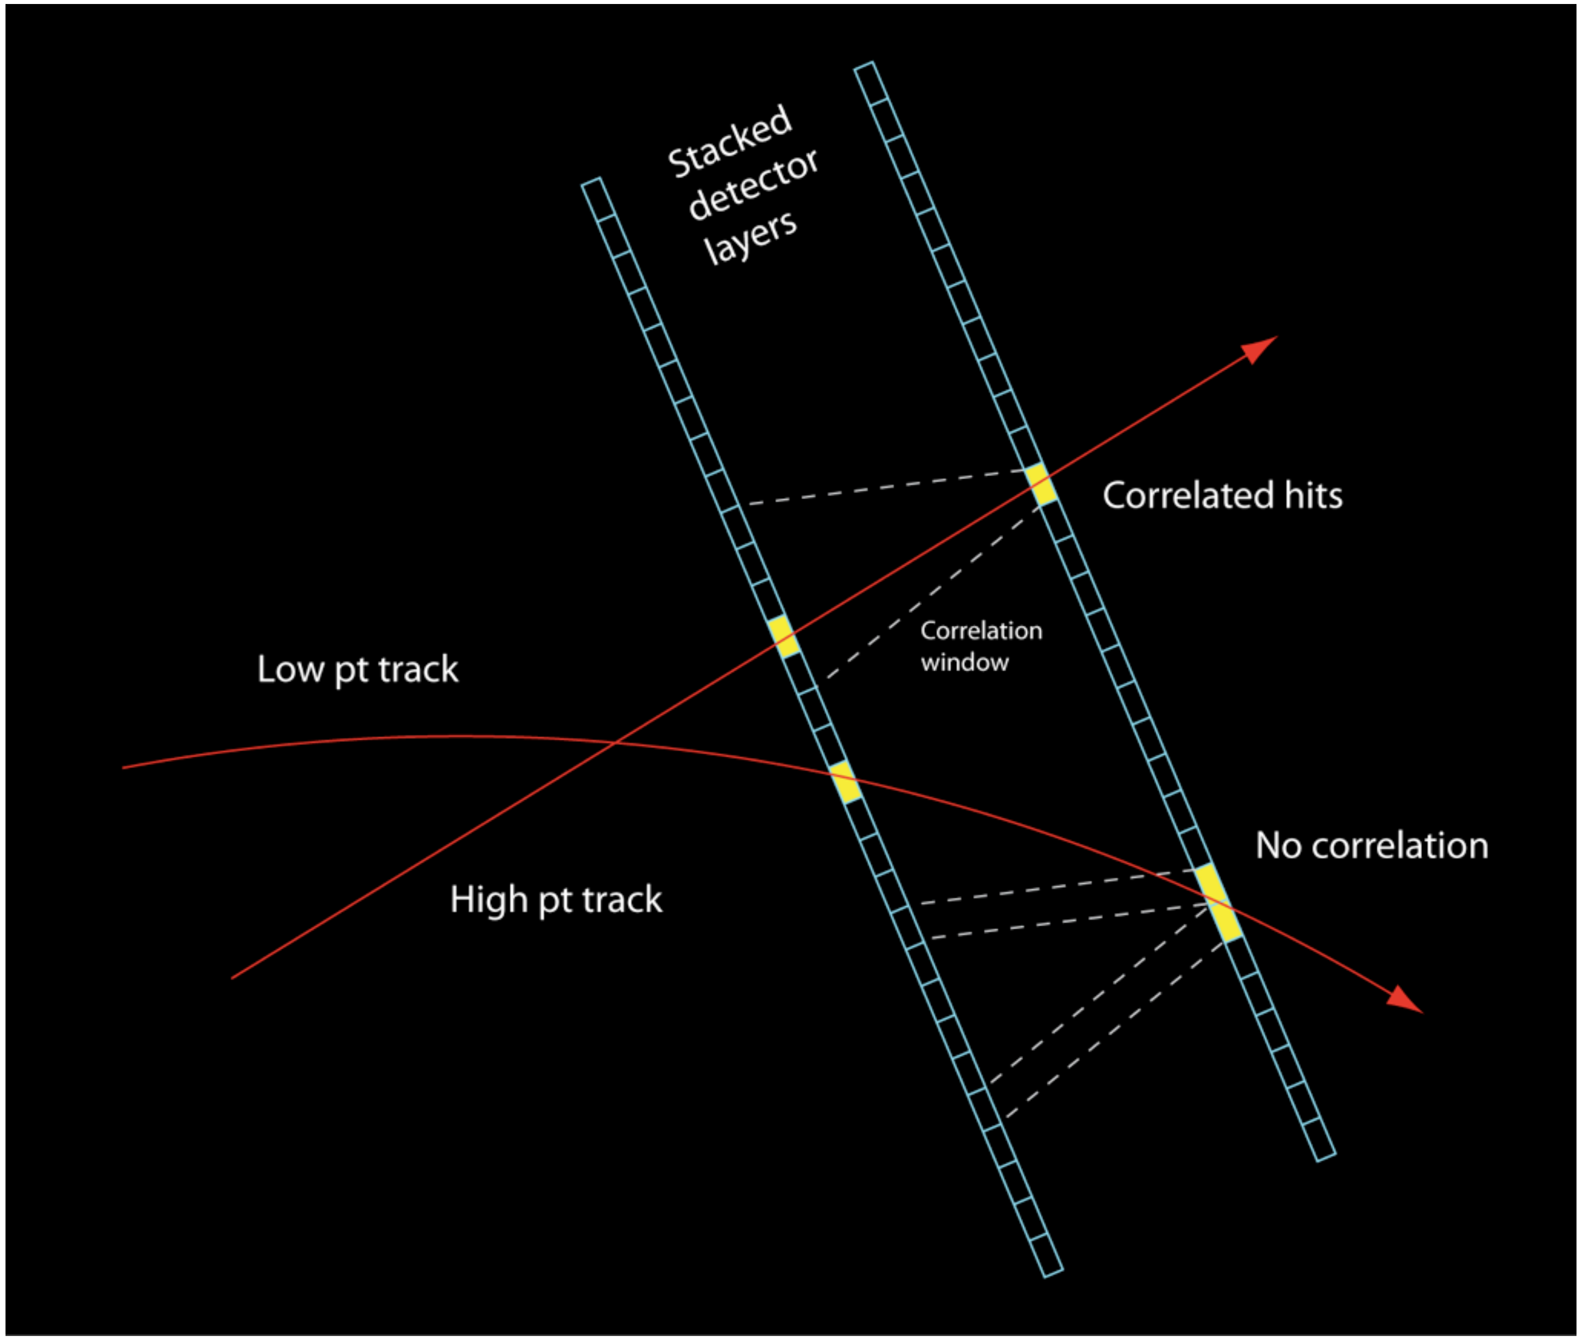

Science & Technology Facilities Council Rutherford Appleton Laboratory
University of BRISTOL

Firmware / Software Stack

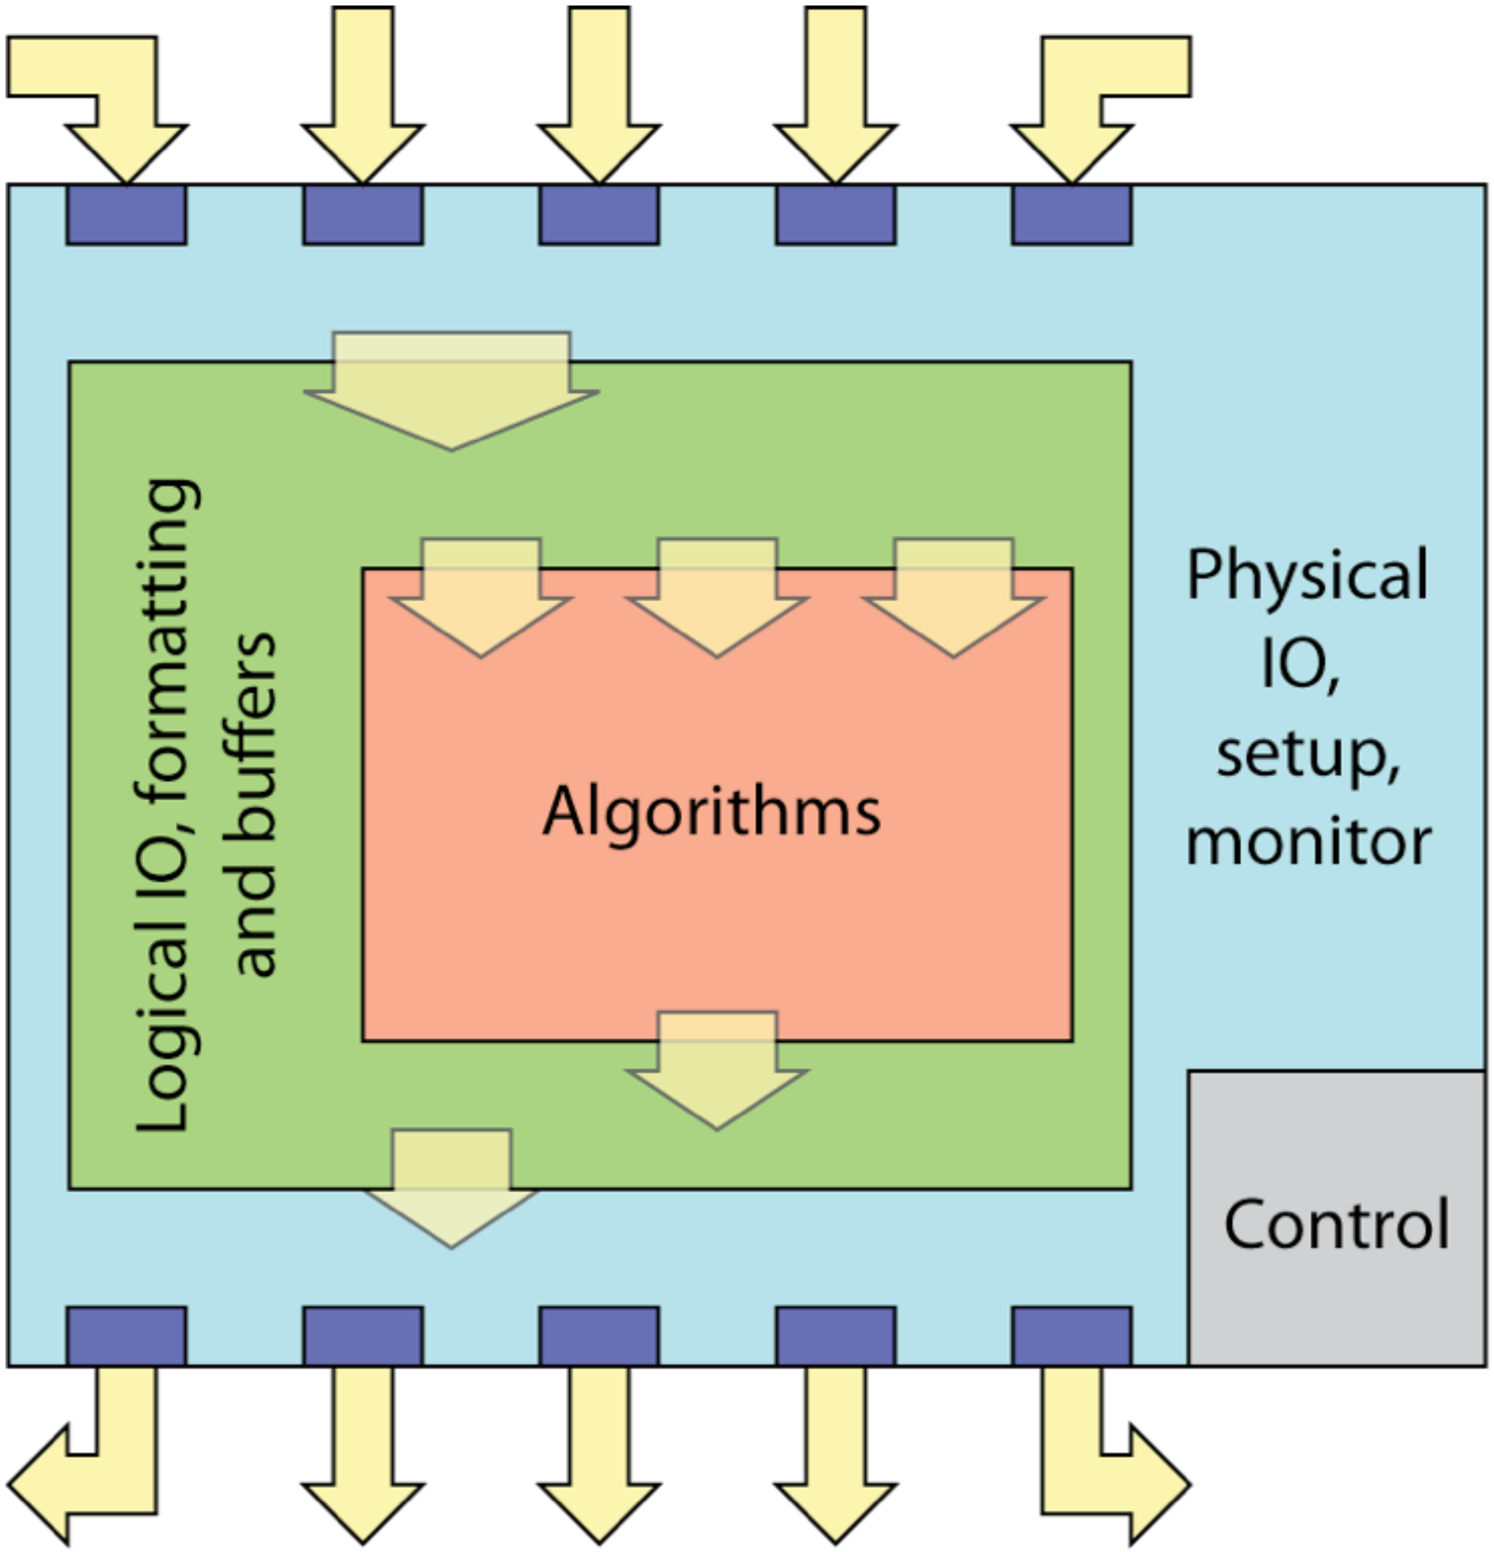

Trigger emulator
Open Development

System setup and test
Common across trigger

Low-level control
Hardware-specific development

uTCA infrastructure
CMS common standard

The DAQ View

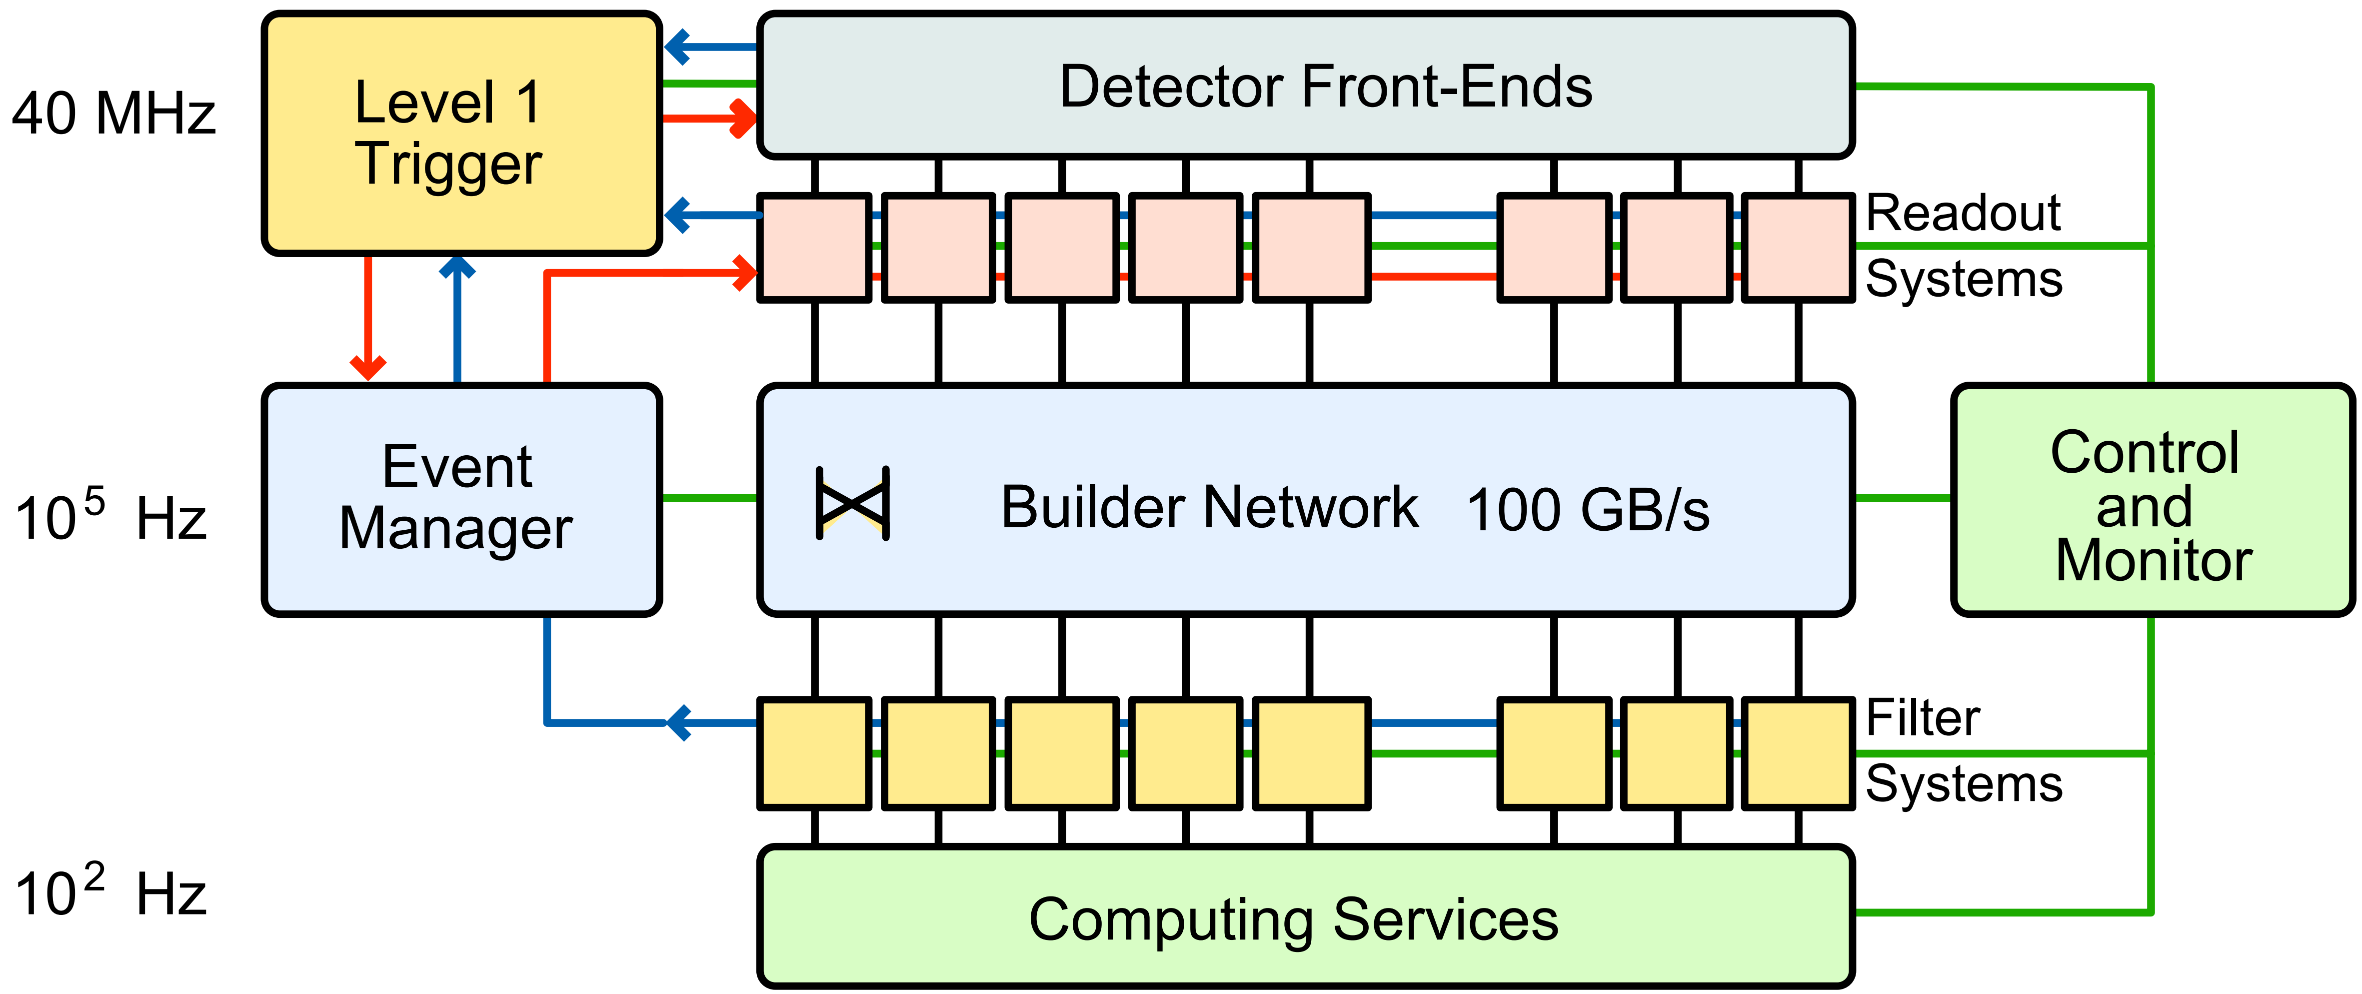
40 MHz
10^5 Hz
10^2 Hz
Level 1 Trigger
Detector Front-Ends
Readout Systems
Event Manager
Builder Network 100 GB/s
Control and Monitor
Filter Systems
Computing Services