Trigger & Data Acquisition (at the LHC)

Paris Sphicas
CERN/PH and Univ. of Athens
CERN Summer Student Lectures
August 2006

- **Introduction**
  - **LHC: The machine and the physics**
  - **Trigger/DAQ architectures and tradeoffs**
- **Level-1 Trigger**
  - **Architectures, elements, performance**
- **DAQ**
  - **Readout, Event-Building, Control & monitor**
- **High-Level trigger**
  - **Farms, algorithms**

Introduction:
Mission Make-it-Possible

Collisions at the LHC: summary

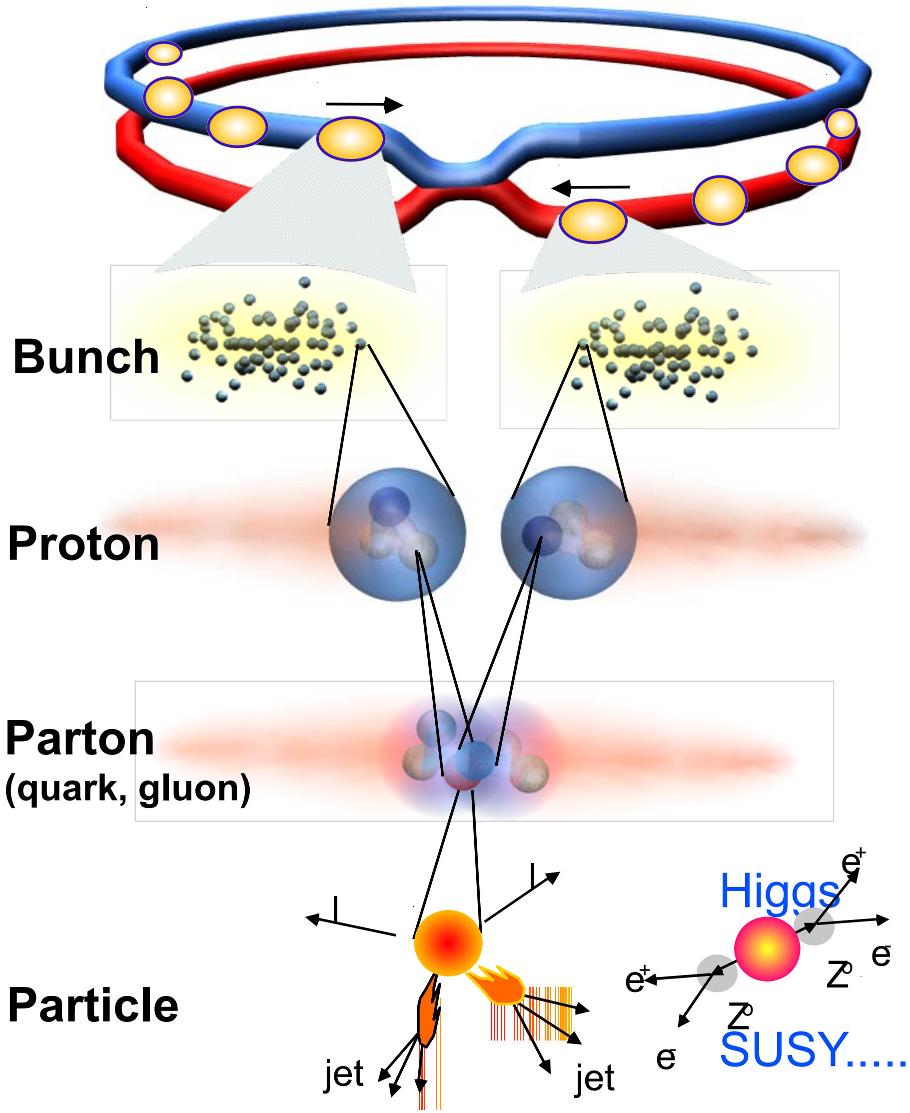

Proton - Proton	2804 bunch/beam
Protons/bunch	10^{11}
Beam energy	7 TeV (7×10^{12} eV)
Luminosity	$10^{34} cm^{-2} s^{-1}$
Crossing rate	40 MHz
Collision rate ≈	10^7-10^9

New physics rate ≈ .00001 Hz

Event selection:
1 in 10,000,000,000,000

Trigger and Data Acquisition System

- **Mandate:**

  "Look at (almost) all bunch crossings, select the most interesting ones, collect all detector information for them and store it for off-line analysis"

    - P.S. For a reasonable number of CHF

- **The photographer analogy:**
  - Trigger: the photographer/camera push-button combination
  - DAQ: burning the film, rolling out the picture, storing film
  - Quality of shot: number of pictures/second, number of pixels
    - And of course the photographer
  - Cost of shot: the camera (one-time); film (recurring); the shot itself (cannot take another picture for a short time after we push on the camera button)

- **Trigger/DAQ: the HEP experiment photographer. All physics analysis runs off of the film (s)he produces**

LHC: physics goals and machine parameters

Higgs boson production at LHC

- **Primary physics goal: explore the physics of Electroweak symmetry breaking.**
  - In the SM: the Higgs
  - Energy of the collider: dictated by machine radius and magnets
  - Luminosity: determine from requirements
- **Higgs mass: unknown; could be up to ~1TeV/c².**
  - Wish: ~20-30 events/year at highest masses
- **Luminosity needed: $10^{34} cm^{-2}s^{-1}$**
  - At 10^{11} protons/bunch, 27 km (i.e. 90 μs), need ~3000 bunches

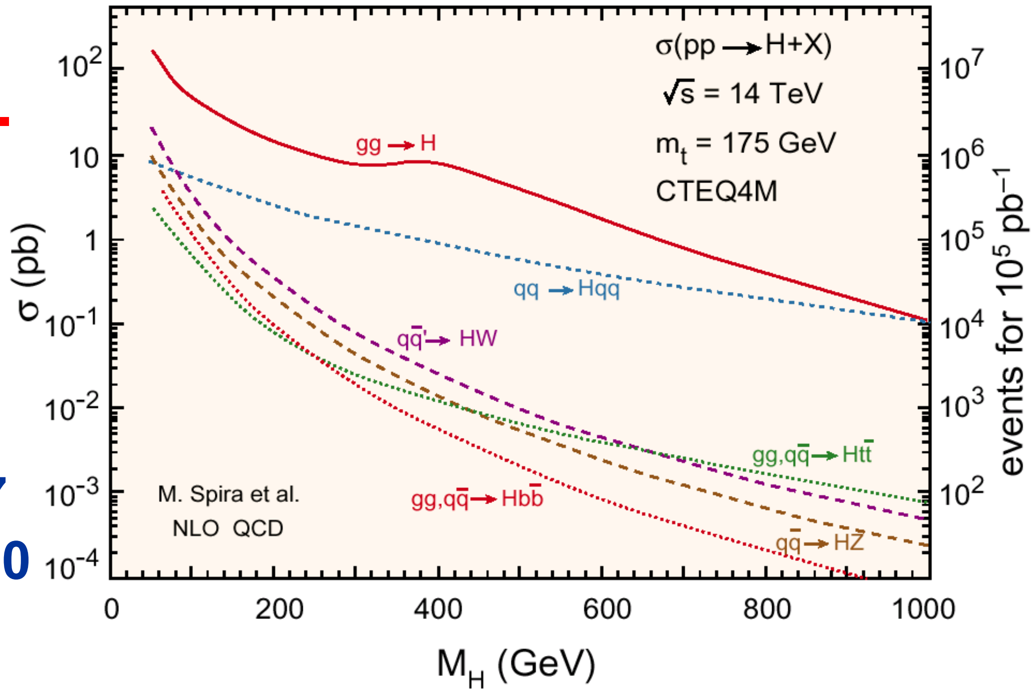

Beam crossings: LEP, Tevatron & LHC

- **LHC will have ~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**

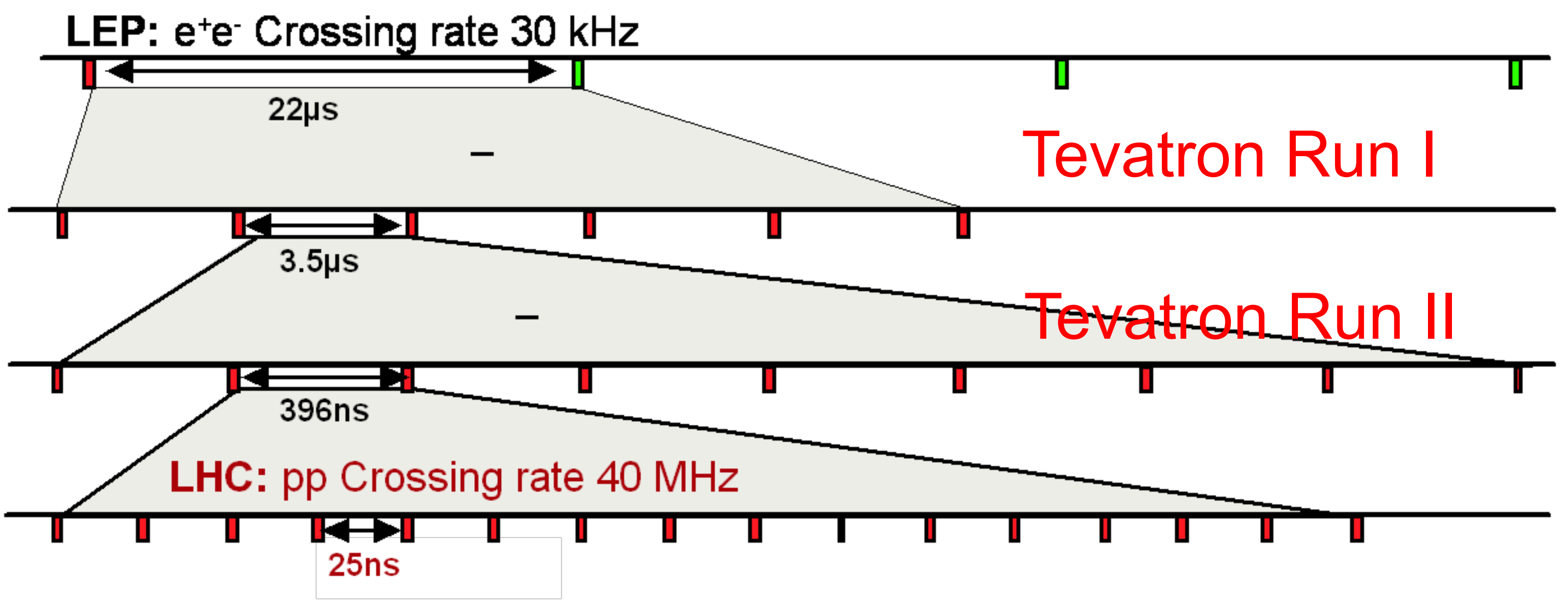

pp cross section and min. bias

- **# of interactions/crossing:**
  - **Interactions/s:**
    - Lum = 10^{34} cm^{-2}s^{-1}=10^7mb^{-1}Hz
    - σ(pp) = 70 mb
    - Interaction Rate, R = 7x10^8 Hz
  - **Events/beam crossing:**
    - Δt = 25 ns = 2.5x10^{-8} s
    - Interactions/crossing=17.5
  - **Not all p bunches are full**
    - 2835 out of 3564 only
    - Interactions/"active" crossing = 17.5 x 3564/2835 = 23

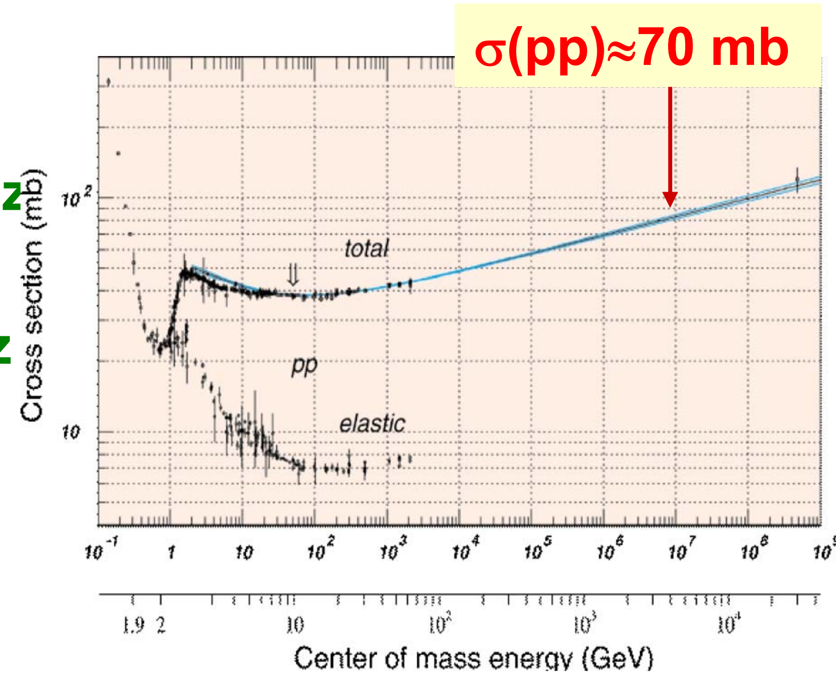

Operating conditions (summary):
1) A "good" event containing a Higgs decay +
2) ≈ 20 extra "bad" (minimum bias) interactions

pp collisions at 14 TeV at 10^{34} cm^{-2}s^{-1}

- **20 min bias events overlap**
- **H→ZZ**

Z →μμ

H→ 4 muons: the cleanest (“golden”) signature

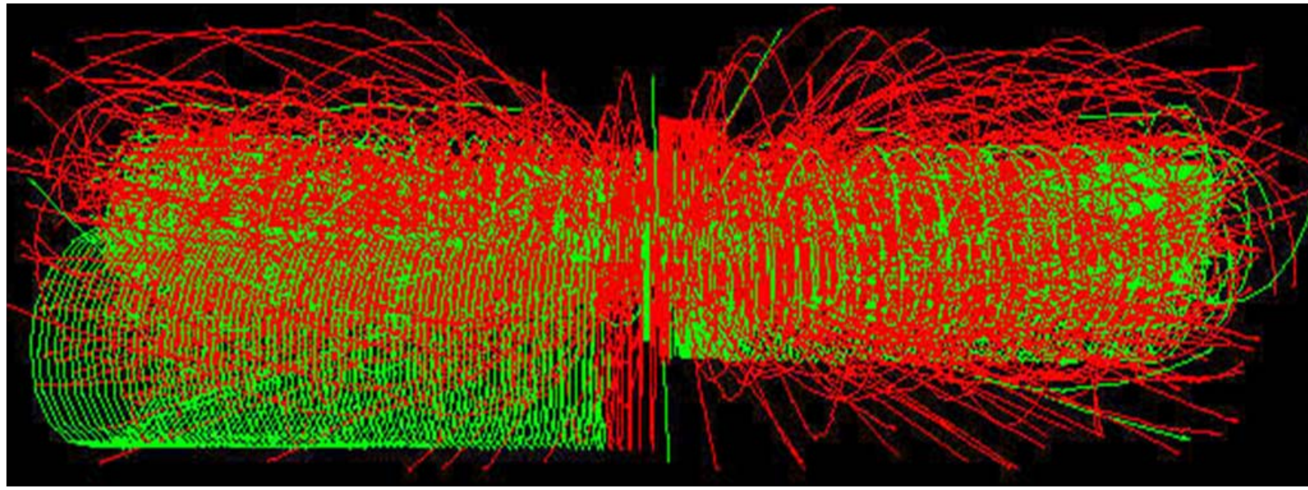

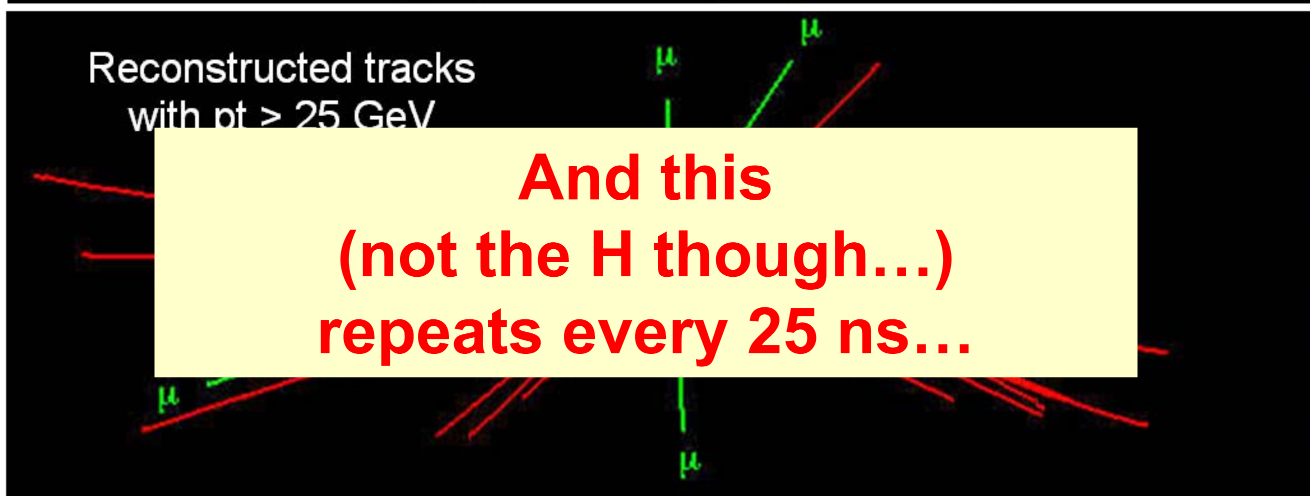

And this
(not the H though…)
repeats every 25 ns…

Impact on detector design

- **LHC detectors must have fast response**
  - Avoid integrating over many bunch crossings ("pile-up")
  - Typical response time : 20-50 ns

    $\rightarrow$ integrate over 1-2 bunch crossings $\rightarrow$ pile-up of 25-50 min-bias events $\rightarrow$ very challenging readout electronics
- **LHC detectors must be highly granular**
  - Minimize probability that pile-up particles be in the same detector element as interesting object (e.g. γ from $H \rightarrow \gamma\gamma$ decays)

    $\rightarrow$ large number of electronic channels
- **LHC detectors must be radiation resistant:**
  - high flux of particles from pp collisions $\rightarrow$ high radiation environment e.g. in forward calorimeters:
    - up to 10^{17} n/cm^2 in 10 years of LHC operation
    - up to 10^7 Gy (1 Gy = unit of absorbed energy = 1 Joule/Kg)

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”**

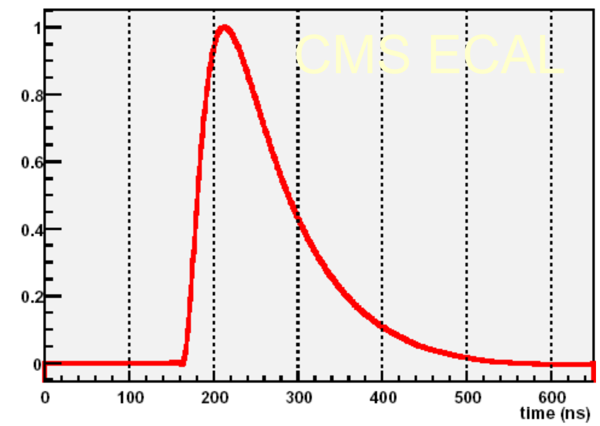

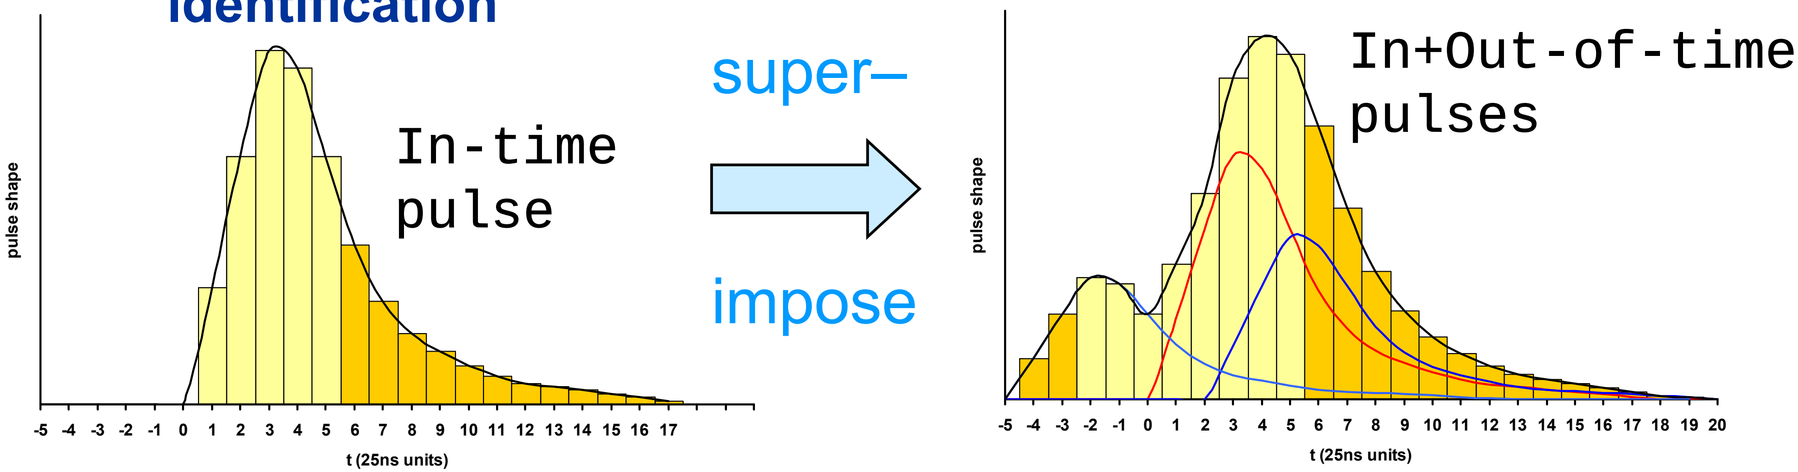

Time of Flight

c=30cm/ns; in 25ns, s=7.5m

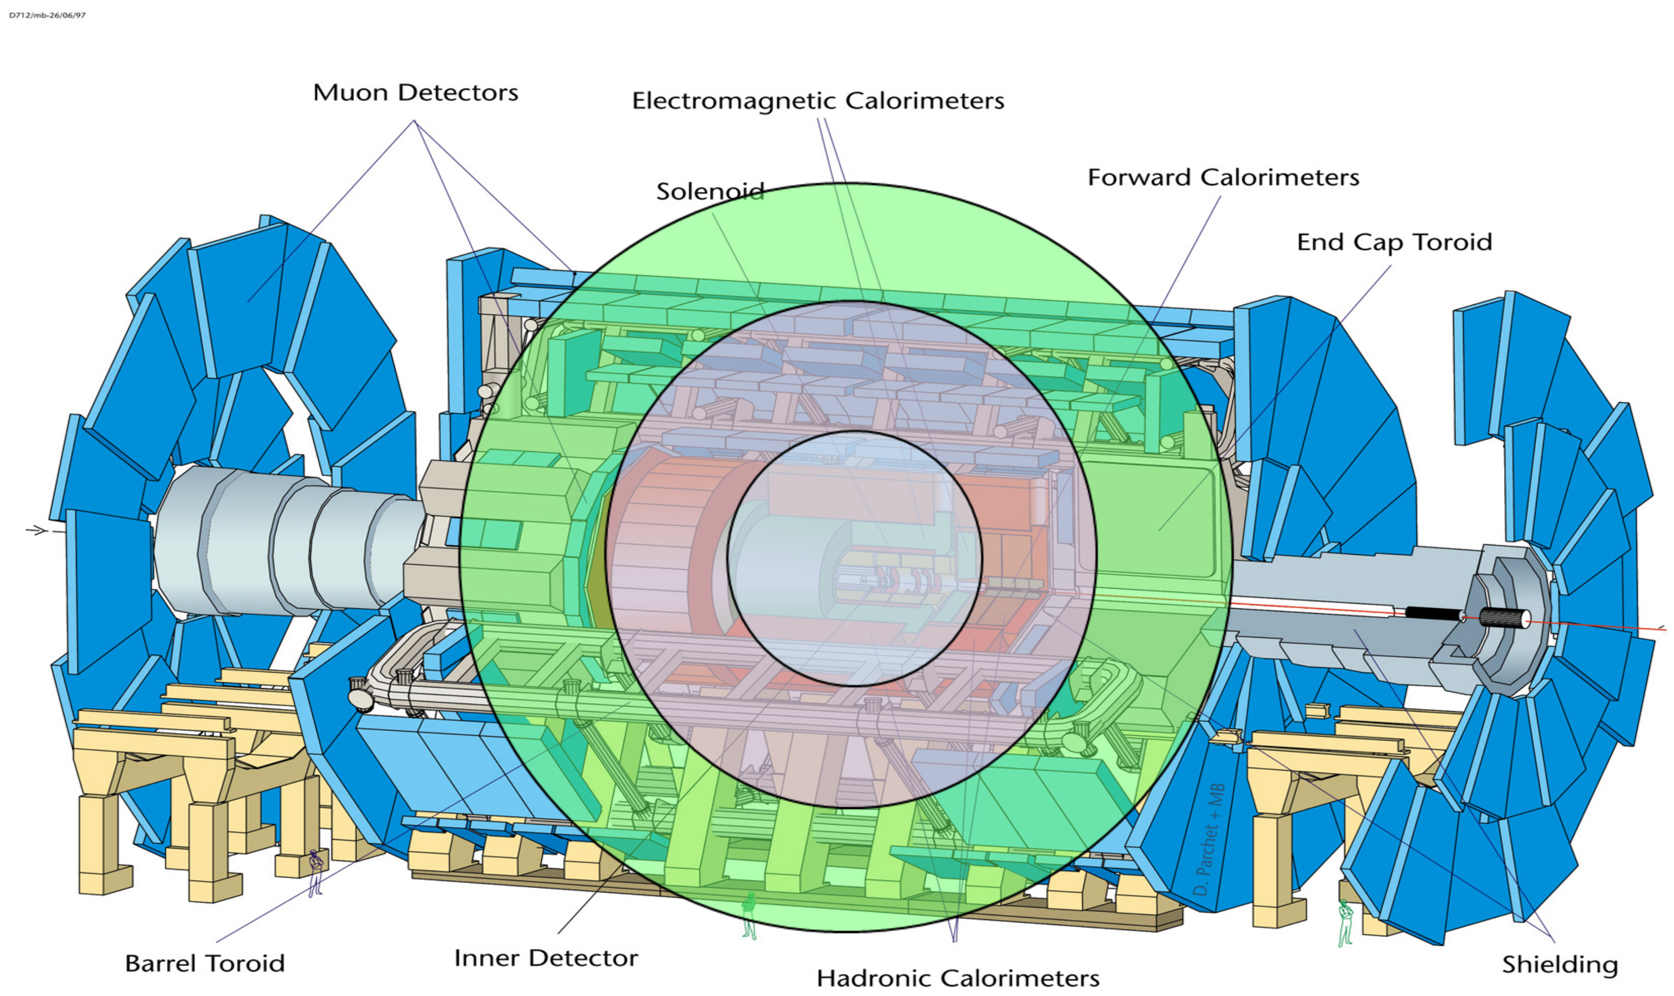

Selectivity: the physics

- **Cross sections of physics processes vary over many orders of magnitude**
  - Inelastic: 10^9 Hz
  - $W \rightarrow \ell\, \nu$: 10^2 Hz
  - t t production: 10 Hz
  - Higgs (100 GeV/c^2): 0.1 Hz
  - Higgs (600 GeV/c^2): 10^{-2} Hz
- **QCD background**
  - Jet E_T ~250 GeV: rate = 1 kHz
  - Jet fluctuations → electron bkg
  - Decays of K, π, b → muon bkg
- **Selection needed: 1:10^{10-11}**
  - Before branching fractions...

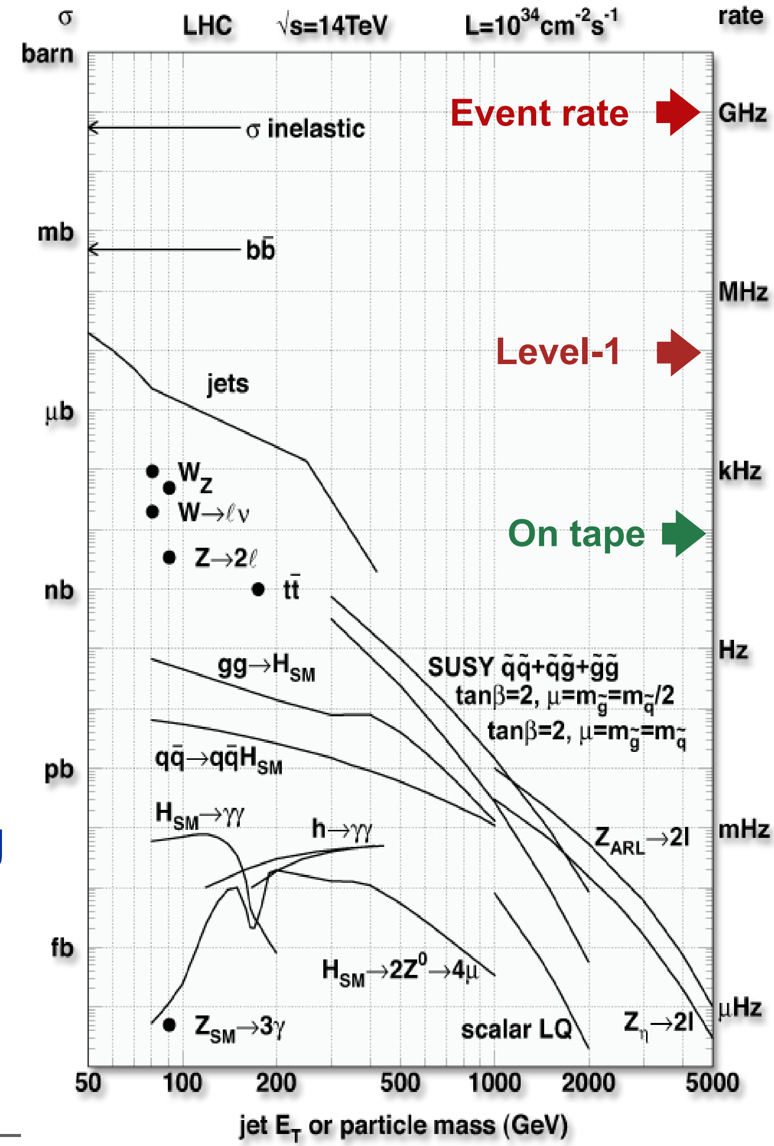

Physics selection at the LHC

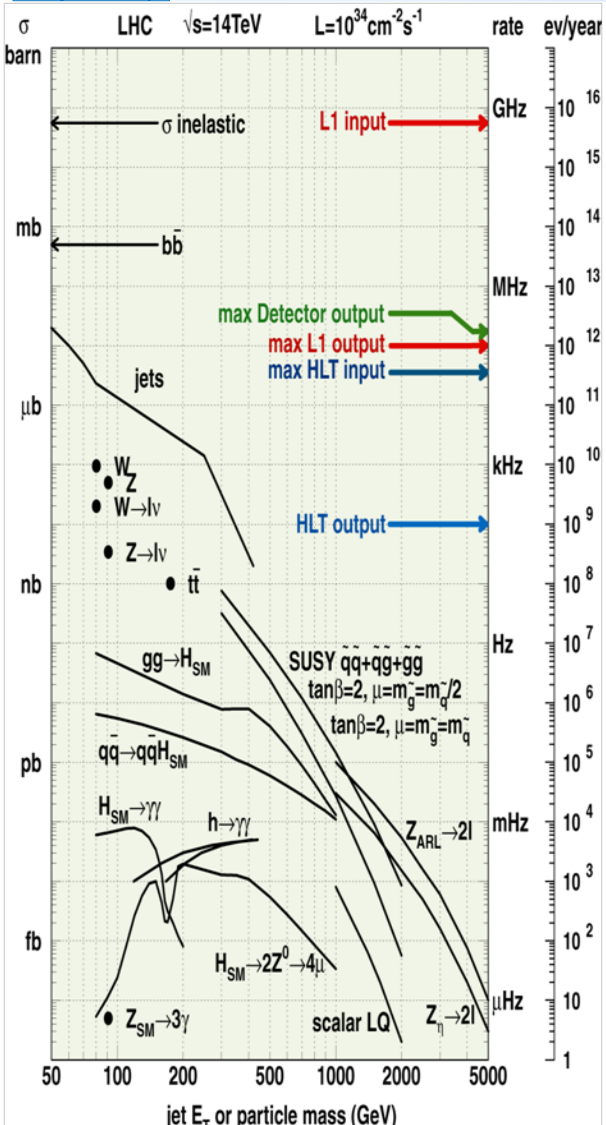

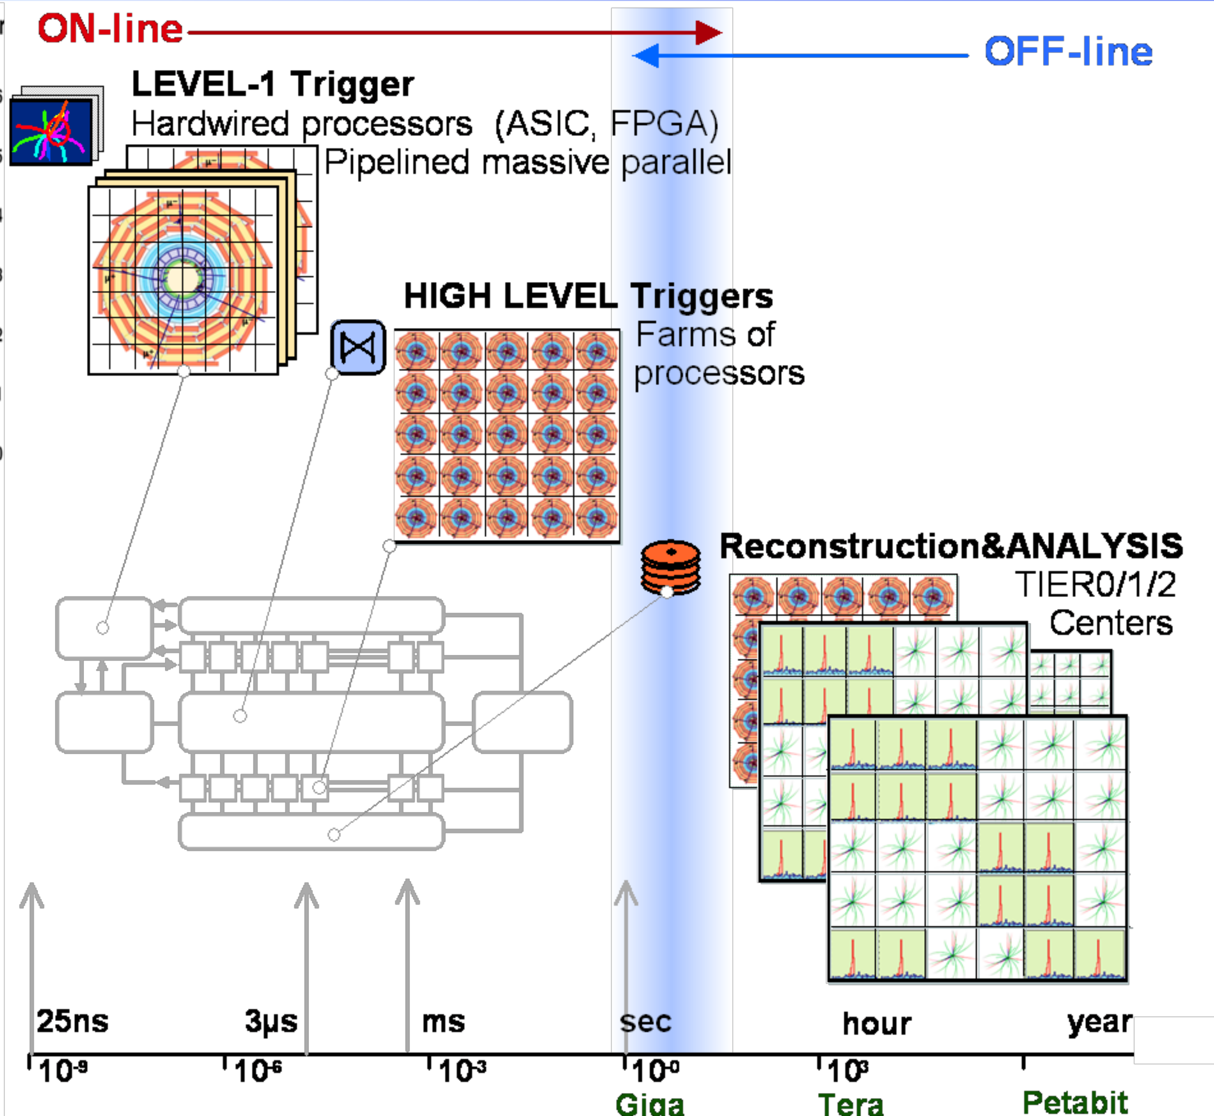

Trigger/DAQ requirements/challenges

- N (channels) ~ O(10^7); ≈20 interactions every 25 ns
  - need huge number of connections
  - need information super-highway
- Calorimeter information should correspond to tracker info
  - need to synchronize detector elements to (better than) 25 ns
- In some cases: detector signal/time of flight > 25 ns
  - integrate more than one bunch crossing's worth of information
  - need to identify bunch crossing...
- Can store data at ≈ 10^2 Hz
  - need to reject most interactions
- It's On-Line (cannot go back and recover events)
  - need to monitor selection

Trigger/DAQ: architectures

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 :

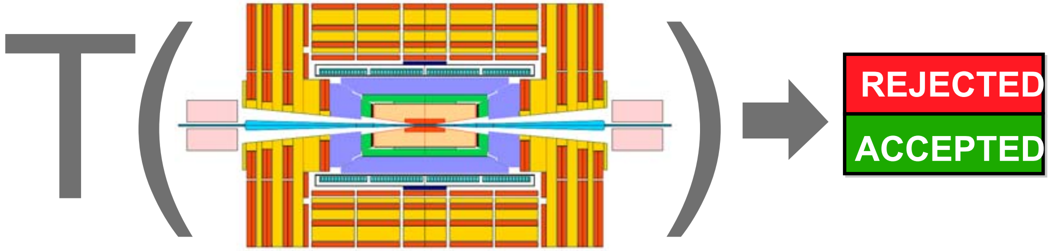

Event data & Apparatus
Physics channels & Parameters

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

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

TRIGGER LEVELS

(possibly with zero dead time)

Online Selection Flow in pp

- **Level-1 trigger: reduce 40 MHz to 10^5 Hz**
  - **This step is always there**
  - **Upstream: still need to get to 10^2 Hz; in 1 or 2 extra steps**

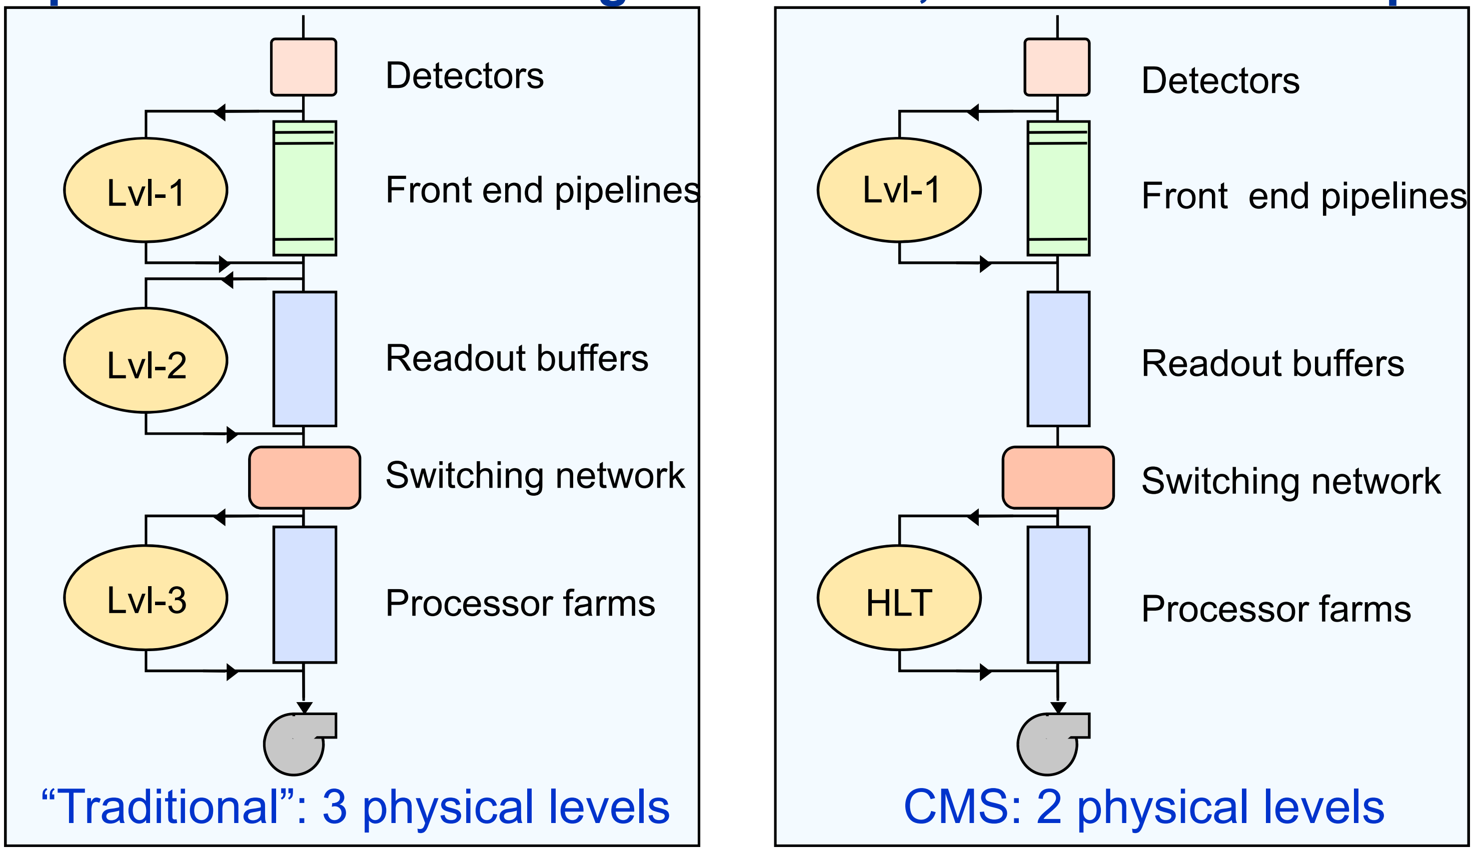

Three physical entities

- **Additional processing in LV-2: reduce network bandwidth requirements**

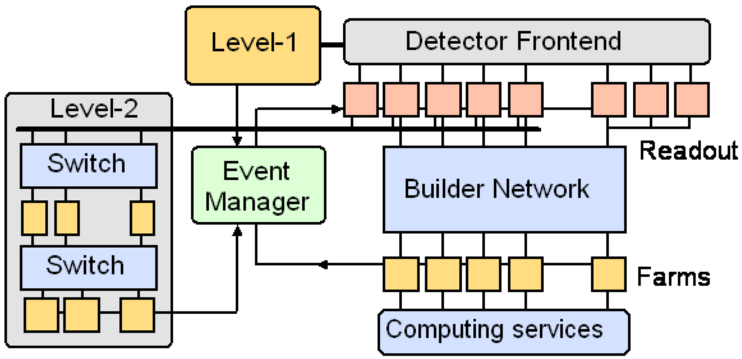

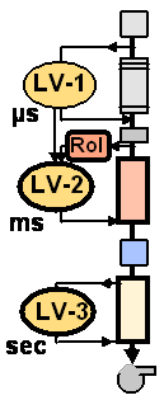

40 MHz

10^5 Hz

10^3 Hz

10 Gb/s

10^2 Hz

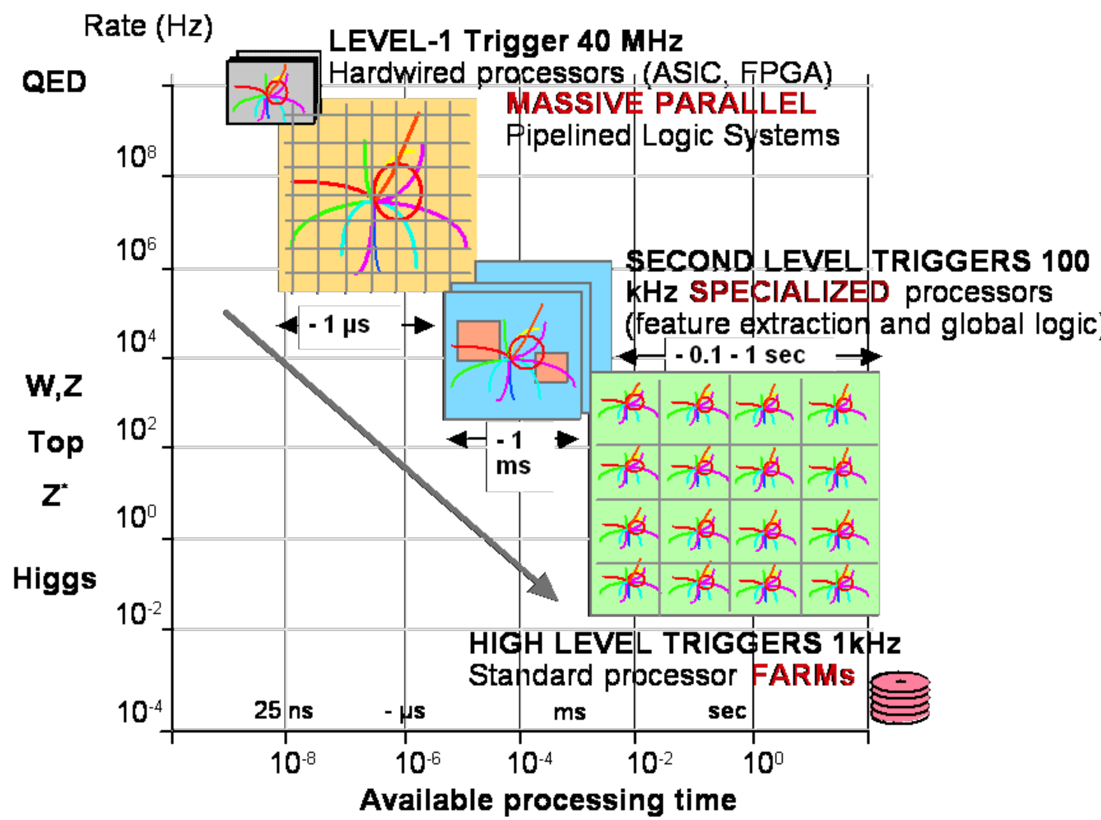

Two physical entities

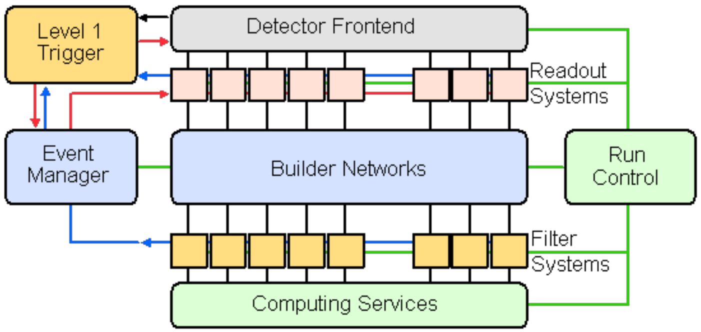

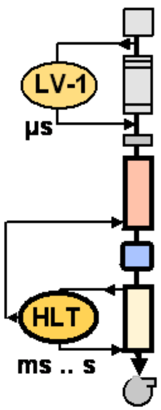

40 MHz

10^5 Hz

1000 Gb/s

10^2 Hz

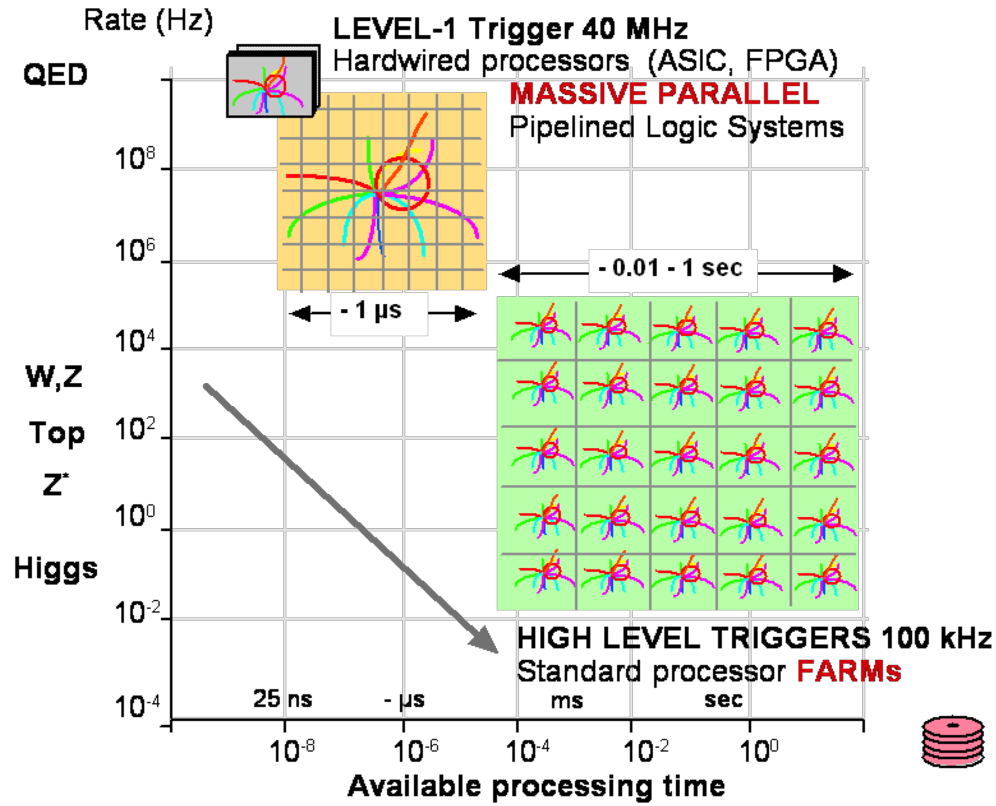

- Reduce number of building blocks
- Rely on commercial components (especially processing and communications)

Comparison of 2 vs 3 physical levels

- **Three Physical Levels**
  - **Investment in:**
    - **Control Logic**
    - **Specialized processors**

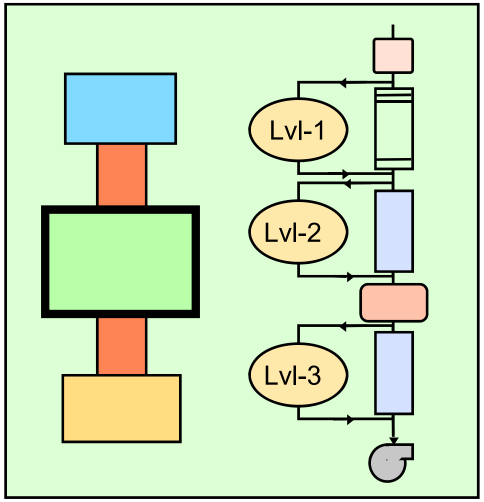

Model

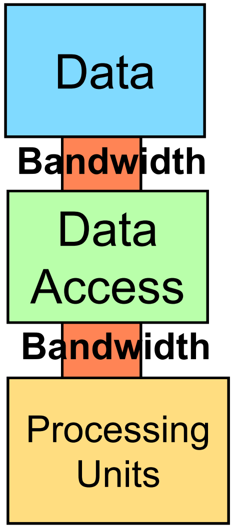

- **Two Physical Levels**
  - **Investment in:**
    - **Bandwidth**
    - **Commercial Processors**

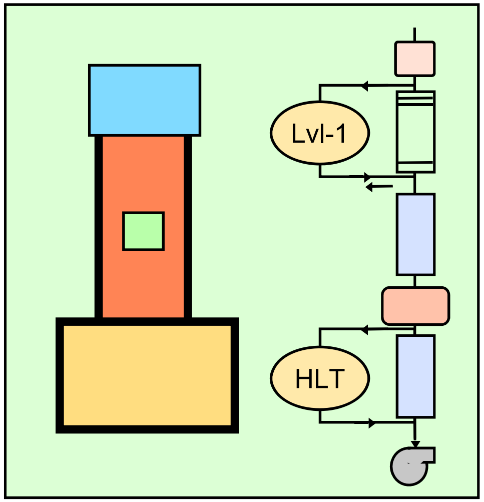

Trigger/DAQ parameters: summary

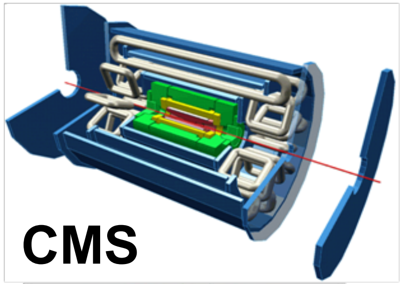

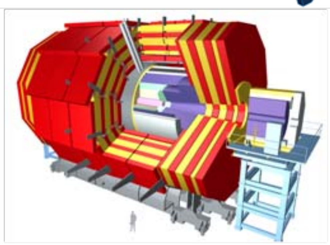

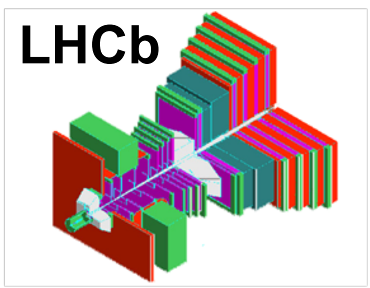

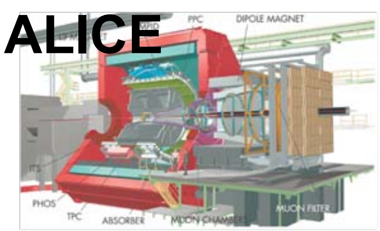

	No.Levels Trigger	**Level-1** Rate (Hz)	**Event** Size (Byte)	**Readout** Bandw.(GB/s)	**Filter Out** MB/s (Event/s)
ATLAS	**3**	**10^5** LV-2 **10^3**	**10^6**	**10**	**100** (10^2)
CMS	**2**	**10^5**	**10^6**	**100**	**100** (10^2)
LHCb	**3**	LV-0 **10^6** LV-1 **$4\ 10^4$**	**$2x10^5$**	**4**	**40** ($2x10^2$)
ALICE	**4**	Pp-Pp **500** p-p **10^3**	**$5x10^7$** **$2x10^6$**	**5**	**1250** (10^2) **200** (10^2)

Trigger/DAQ systems: present & future

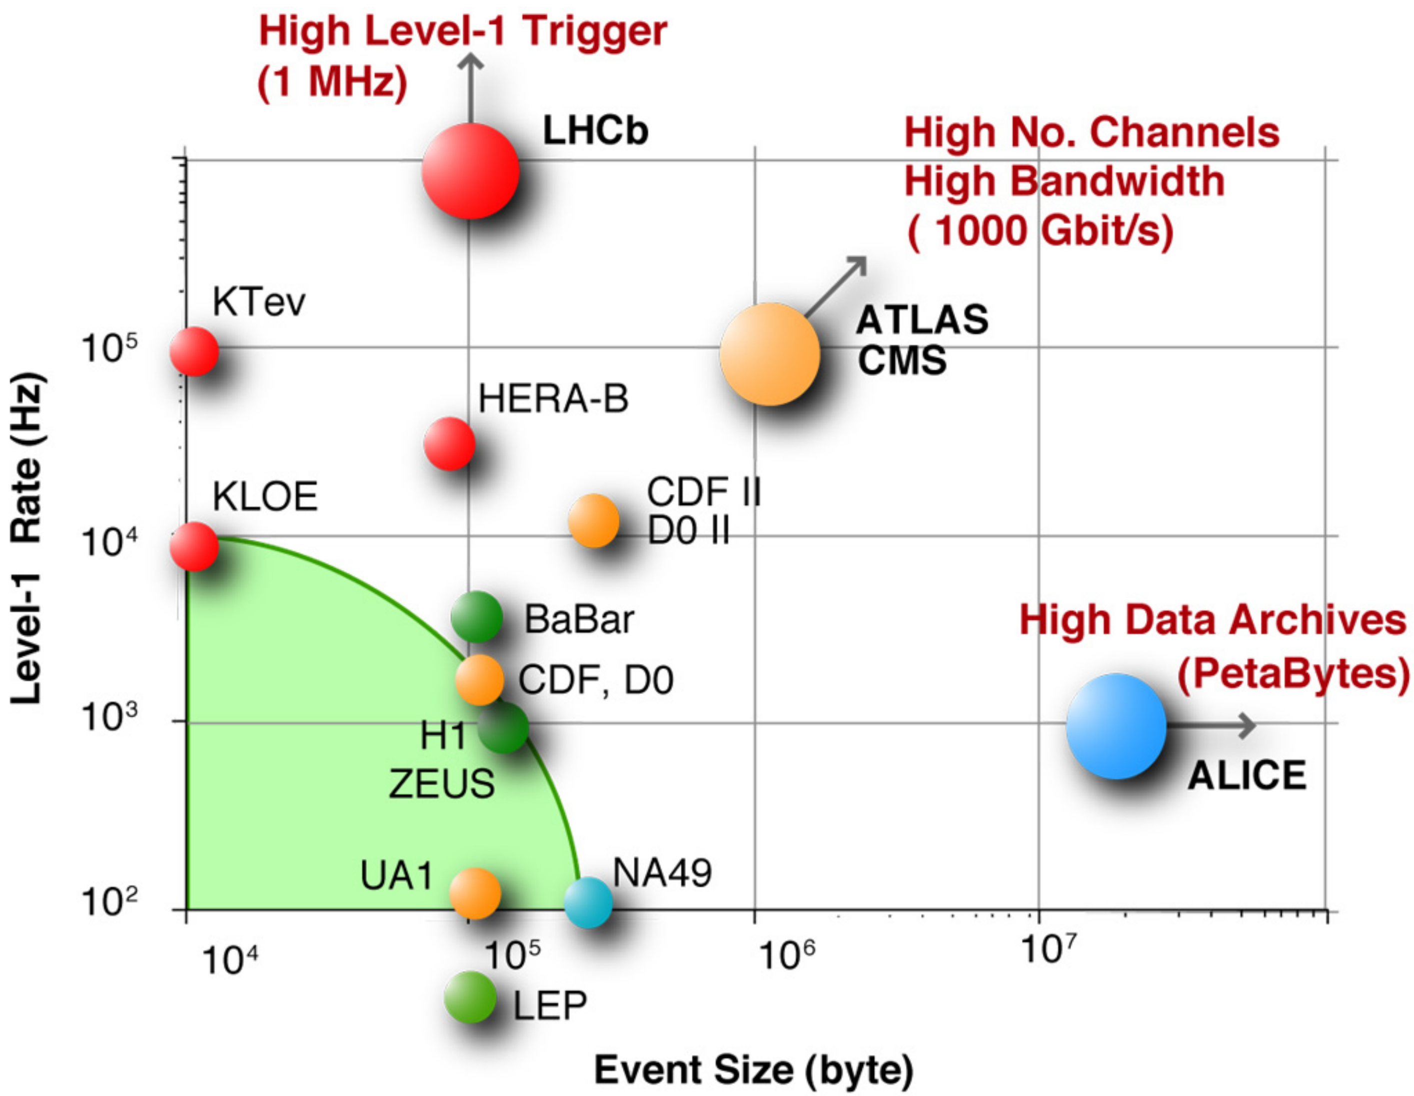

Trigger/DAQ systems: grand view

ATLAS

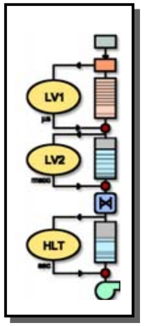

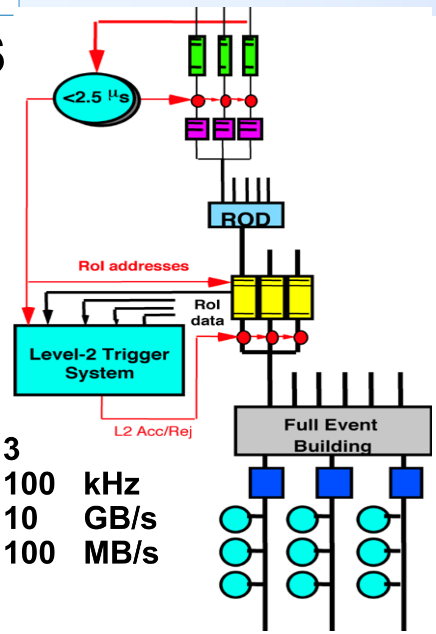

Levels	3	
LV-1 rate	100	kHz
Readout	10	GB/s
Storage	100	MB/s

ALICE

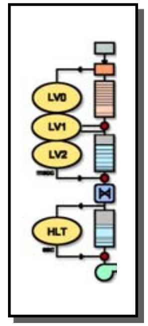

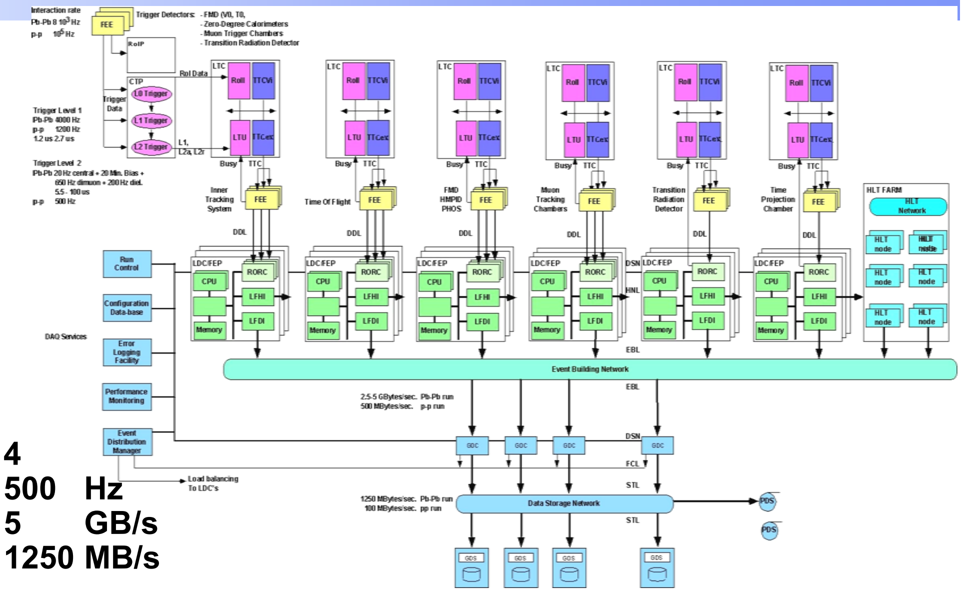

Levels	4	
LV-1 rate	500	Hz
Readout	5	GB/s
Storage	1250	MB/s

CMS

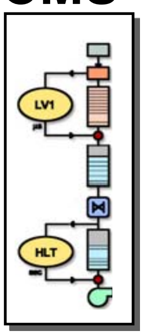

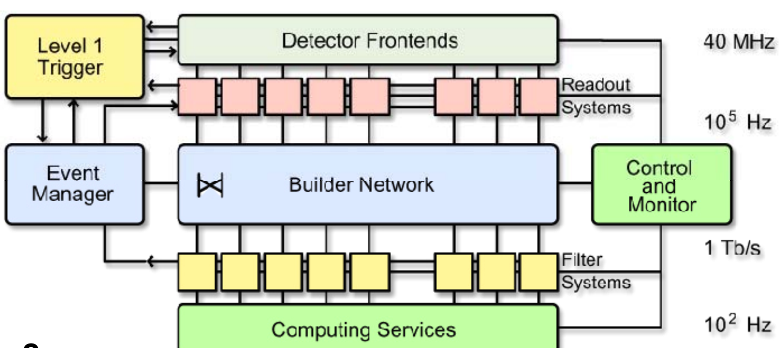

Levels	2	
LV-1 rate	100	kHz
Readout	100	GB/s
Storage	100	MB/s

LHCb

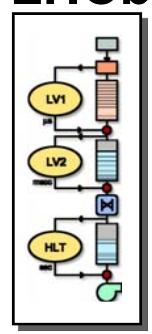

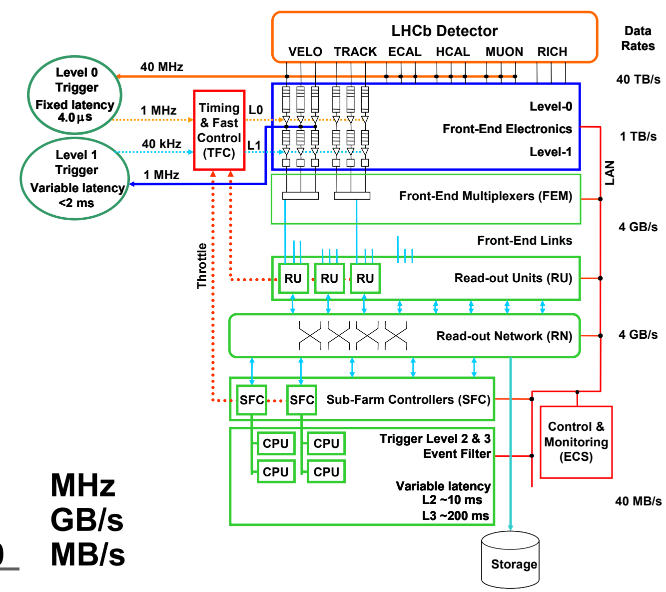

Levels	3	
LV-1 rate	1	MHz
Readout	4	GB/s
Storage	40	MB/s

Level-1 Trigger

Physics selection at the LHC

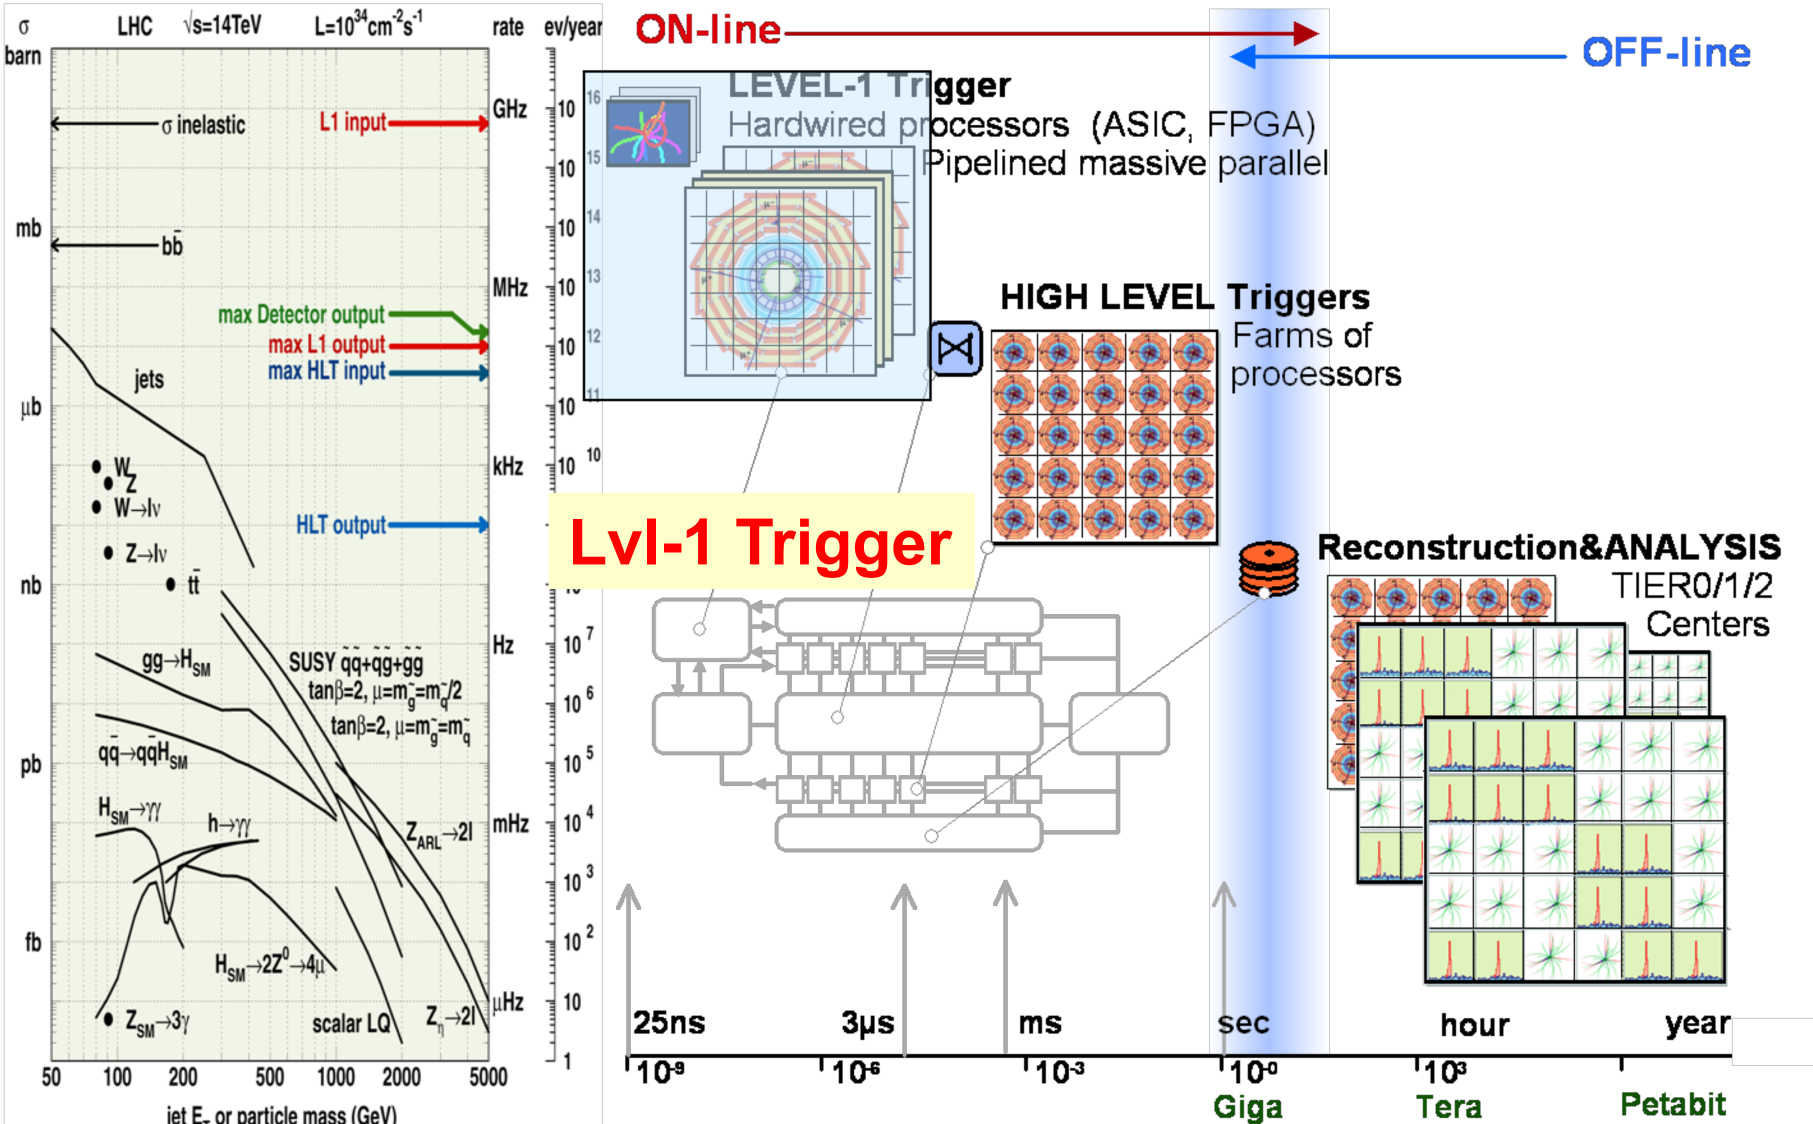

Level-1 trigger algorithms

- **Physics facts:**
  - pp collisions produce mainly hadrons with P_T~1 GeV
  - Interesting physics (old and new) has particles (leptons and hadrons) with large transverse momenta:
    - $W \rightarrow e\nu$: $M(W)=80\ GeV/c^2$; $P_T(e) \sim 30\text{-}40$ GeV
    - H(120 GeV)$\rightarrow\gamma\gamma$: $P_T(\gamma) \sim 50\text{-}60$ GeV
- **Basic requirements:**
  - Impose high thresholds on particles
    - Implies distinguishing particle types; possible for electrons, muons and "jets"; beyond that, need complex algorithms
  - Typical thresholds:
    - Single muon with P_T>20 GeV (rate ~ 10 kHz)
      - ➔ Dimuons with P_T>6 (rate ~ 1 kHz)
    - Single e/γ with P_T>30 GeV (rate ~ 10-20 kHz)
      - ➔ Dielectrons with P_T>20 GeV (rate ~ 5 kHz)
    - Single jet with P_T>300 GeV (rate ~ 0.2-0.4 kHz)

Particle signatures in the detector(s)

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

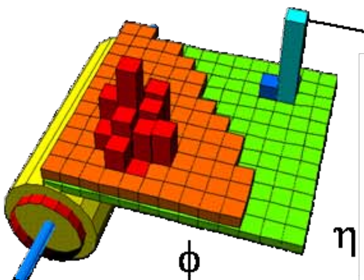

CALORIMETERs
Cluster finding and energy deposition evaluation

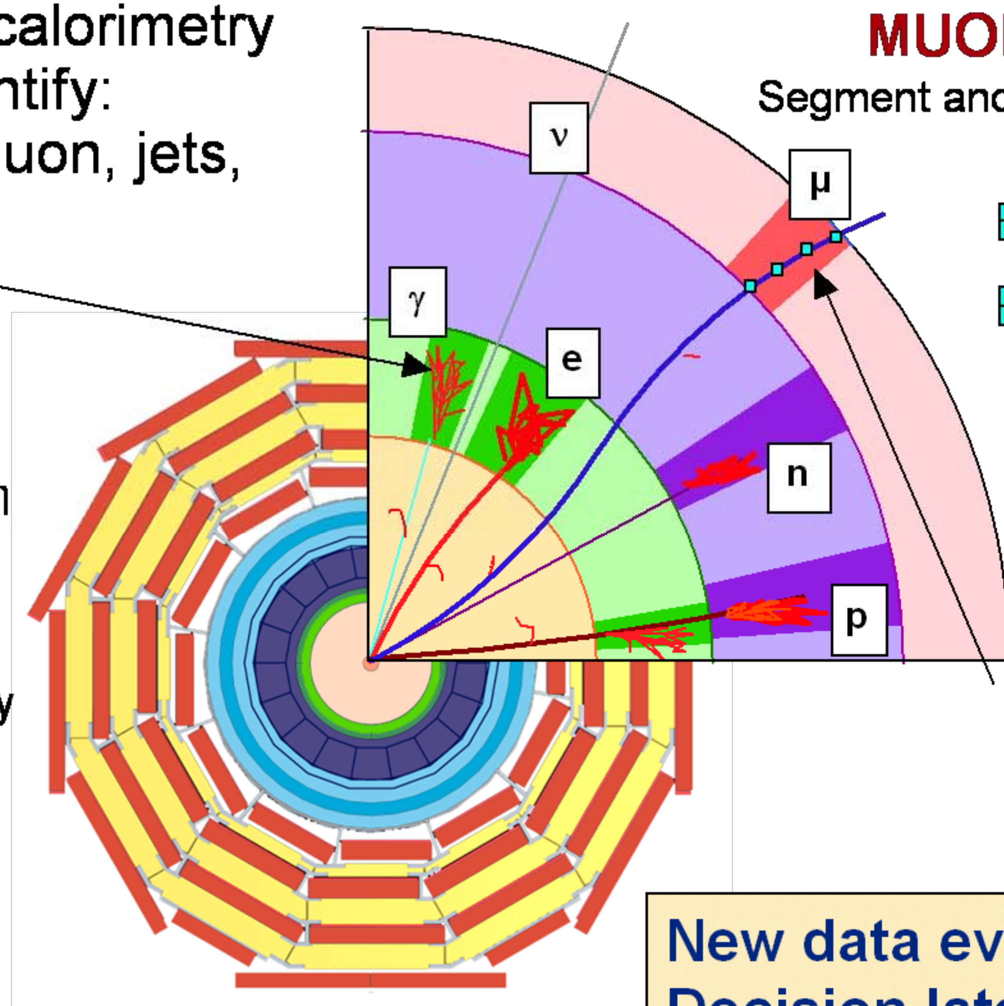

MUON System
Segment and track finding

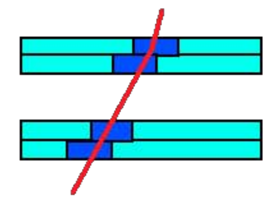

New data every 25 ns
Decision latency ~ μs

At Level-1: only calo and muon info

- **Pattern recognition much faster/easier**

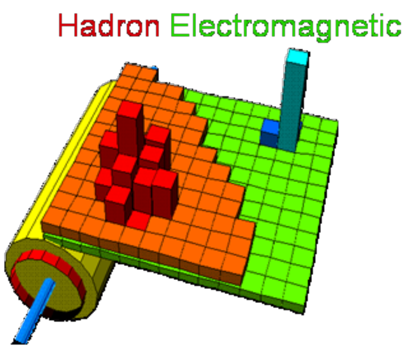

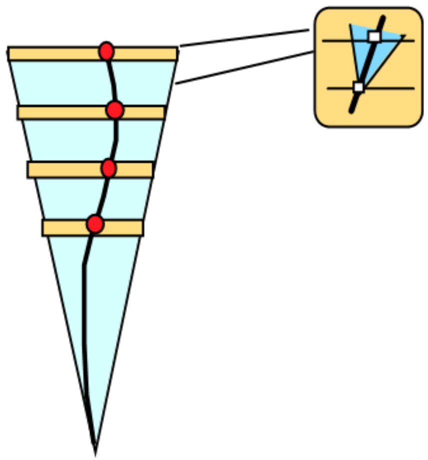

- **Simple algorithms**
- **Small amounts of data**
- **Local decisions**

- **Compare to tracker info**

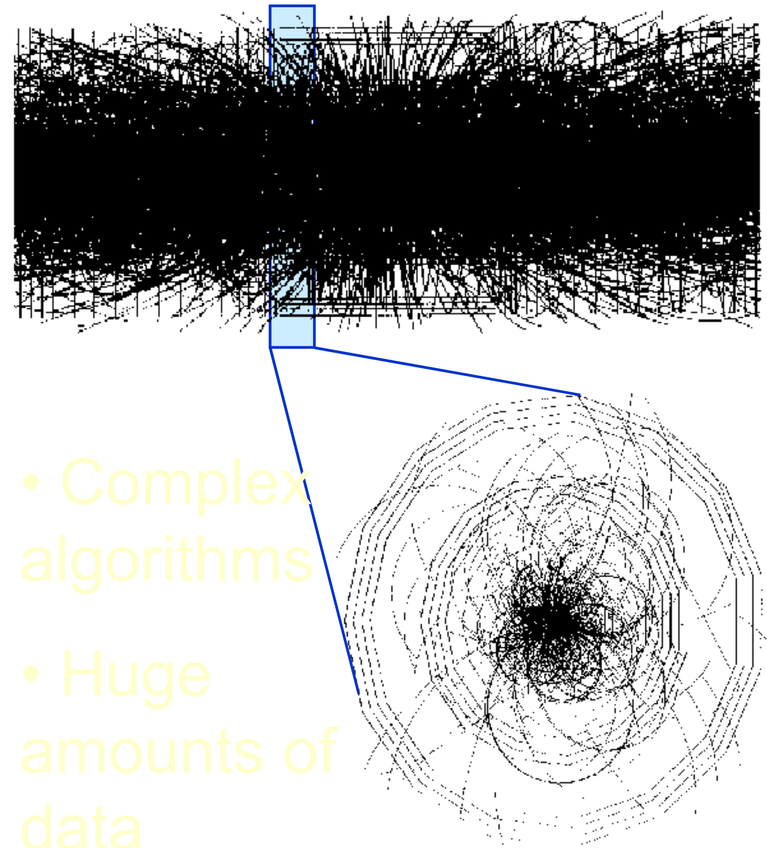

- Complex algorithms
- Huge amounts of data
- Need to link sub-detectors

Level-1 Trigger: decision loop

- **Synchronous 40 MHz digital system**
  - **Typical: 160 MHz internal pipeline**
  - **Latencies:**
    - **Readout + processing: < 1μs**
    - **Signal collection & distribution: ≈ 2μs**
- **At Lvl-1: process only calo+μ info**

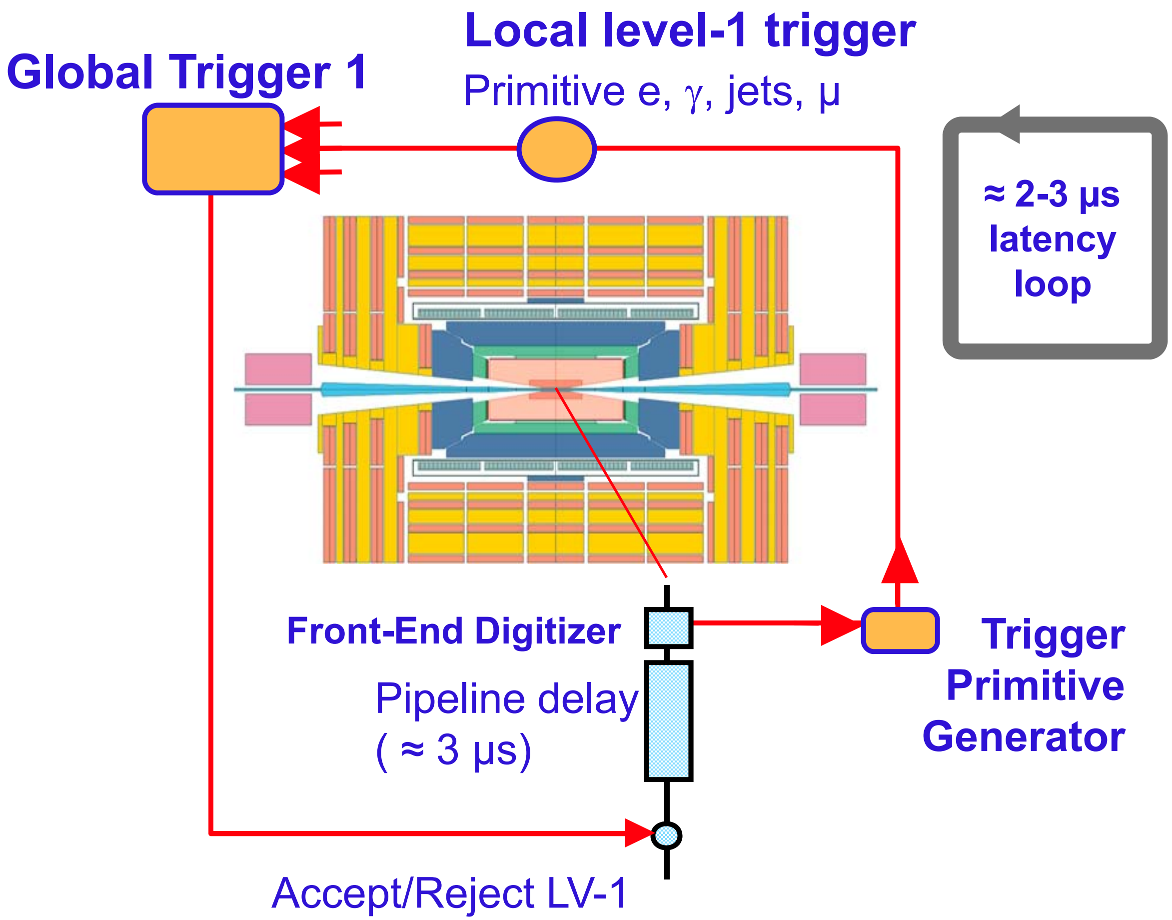

Signaling and pipelining

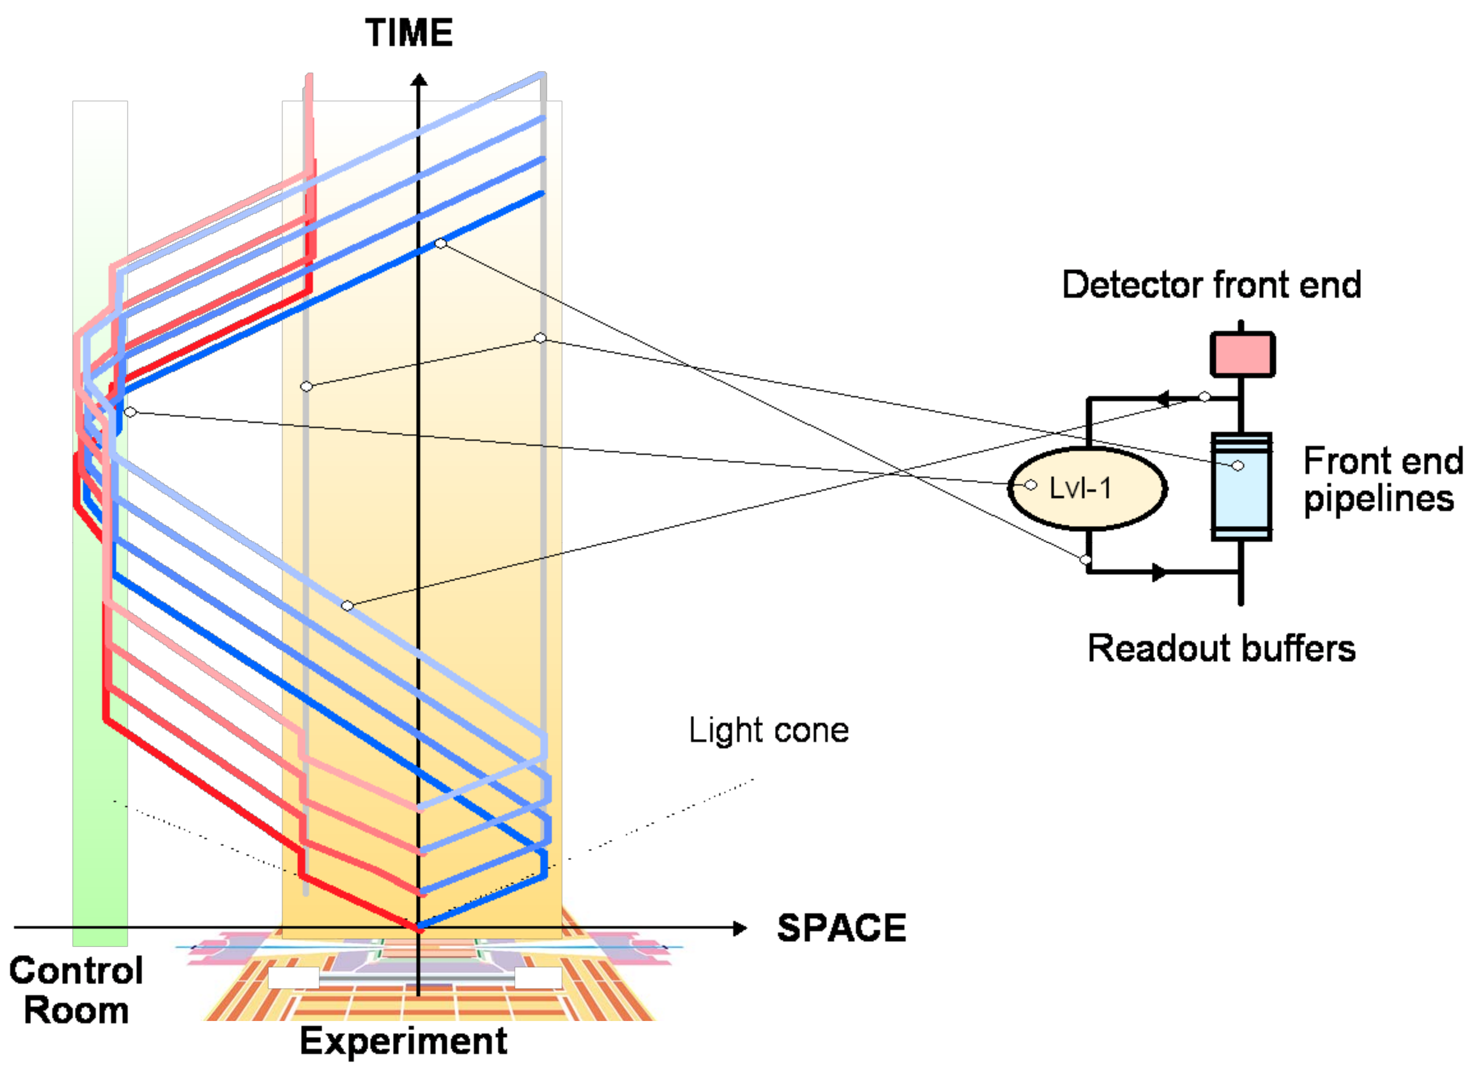

Detector Readout: front-end types

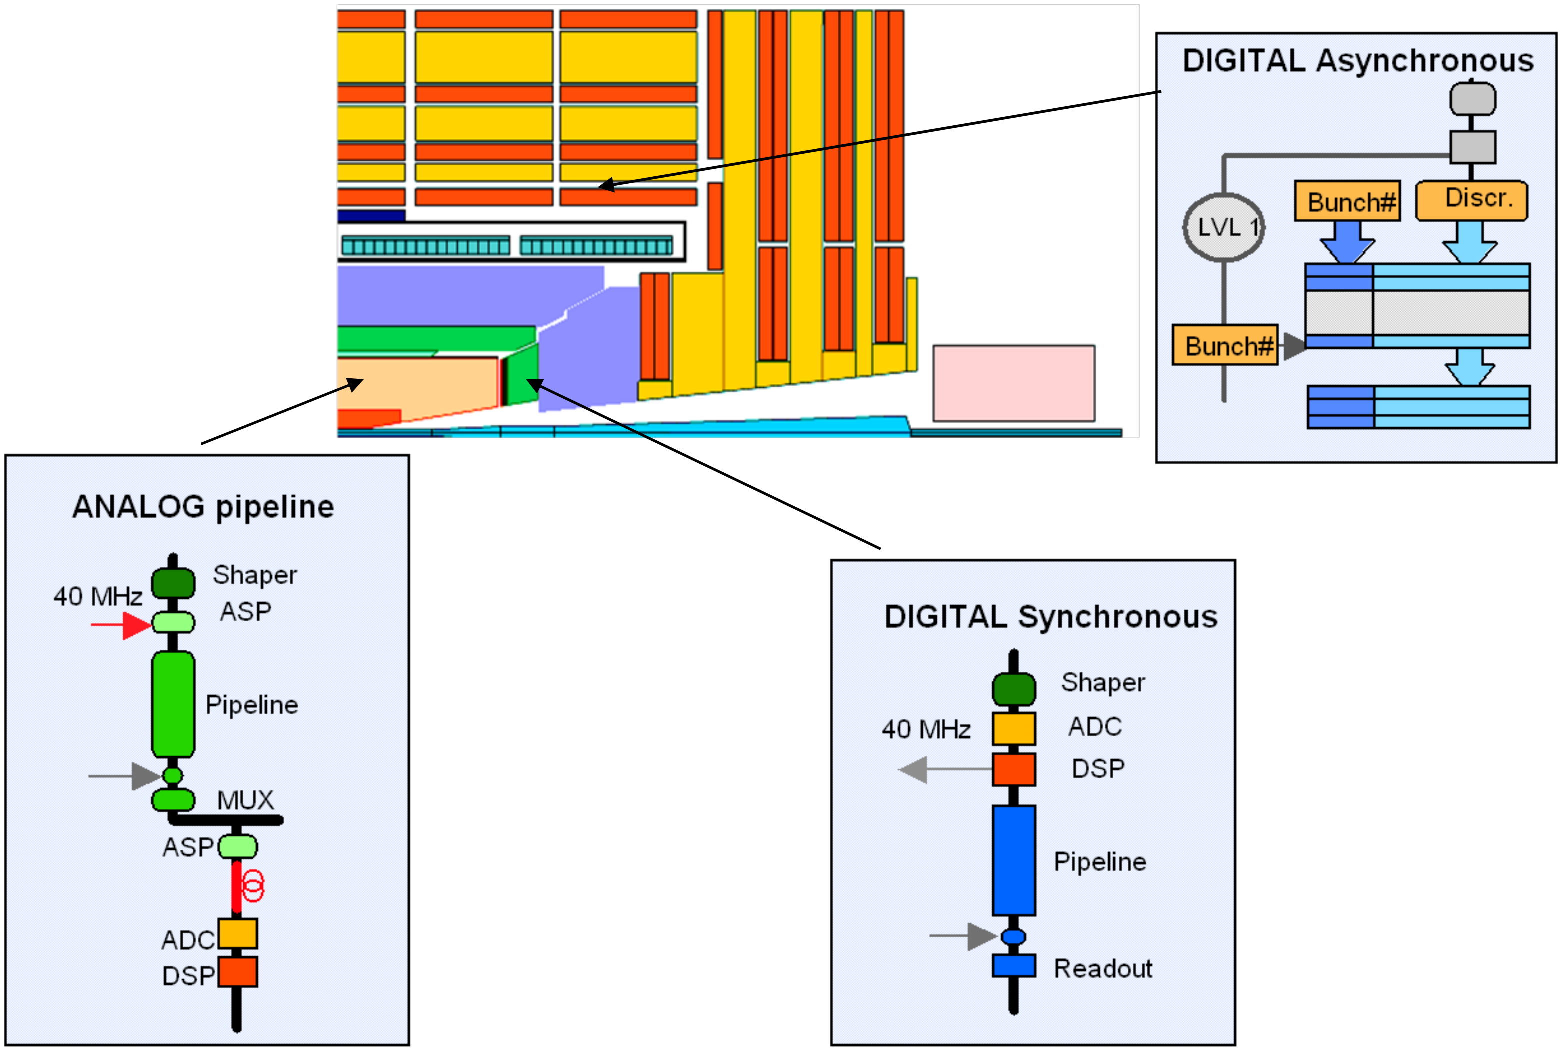

Clock distribution & synchronization

- **Trigger, Timing & Control (TTC); from RD12**

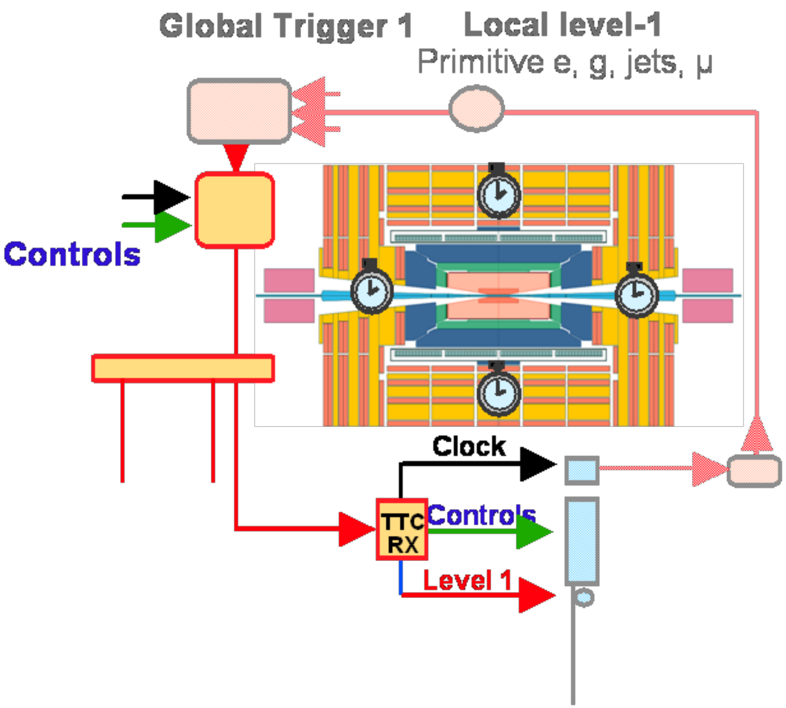

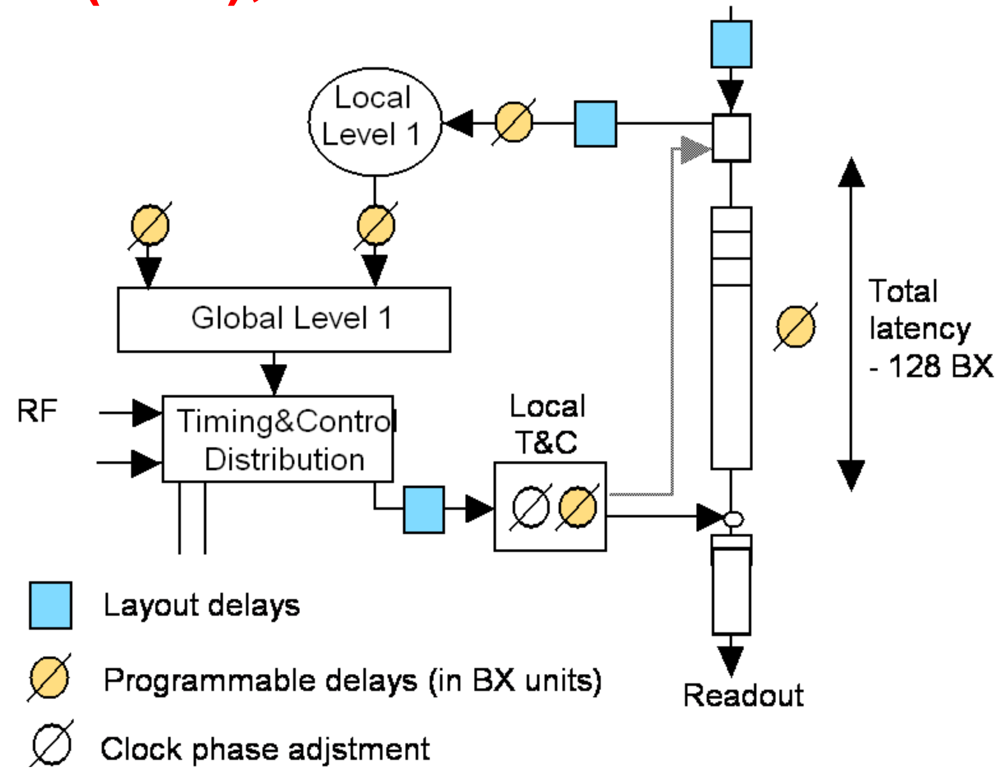

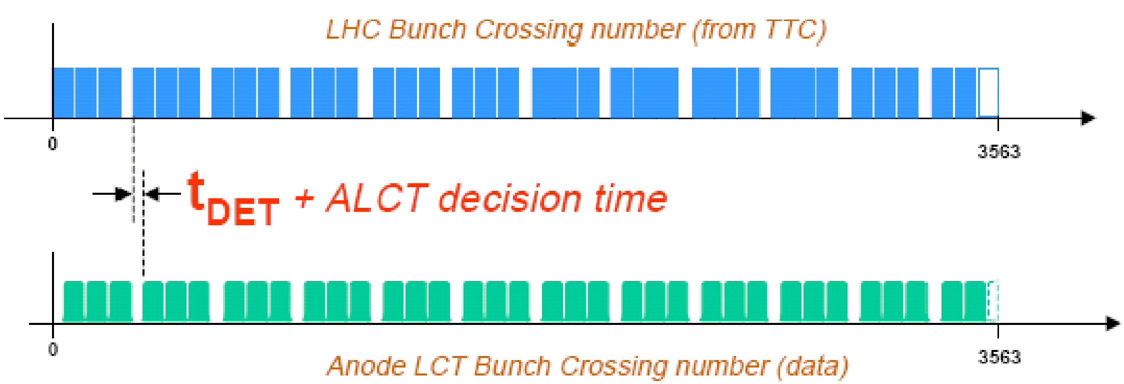

Lvl-1 trigger architecture: ATLAS

CMS ~ similar

CAVERN

Radiation tolerance, cooling, grounding, magnetic field, no access

~7000 calorimeter trigger towers (analogue sum on detectors)

O(1M) RPC/TGC channels

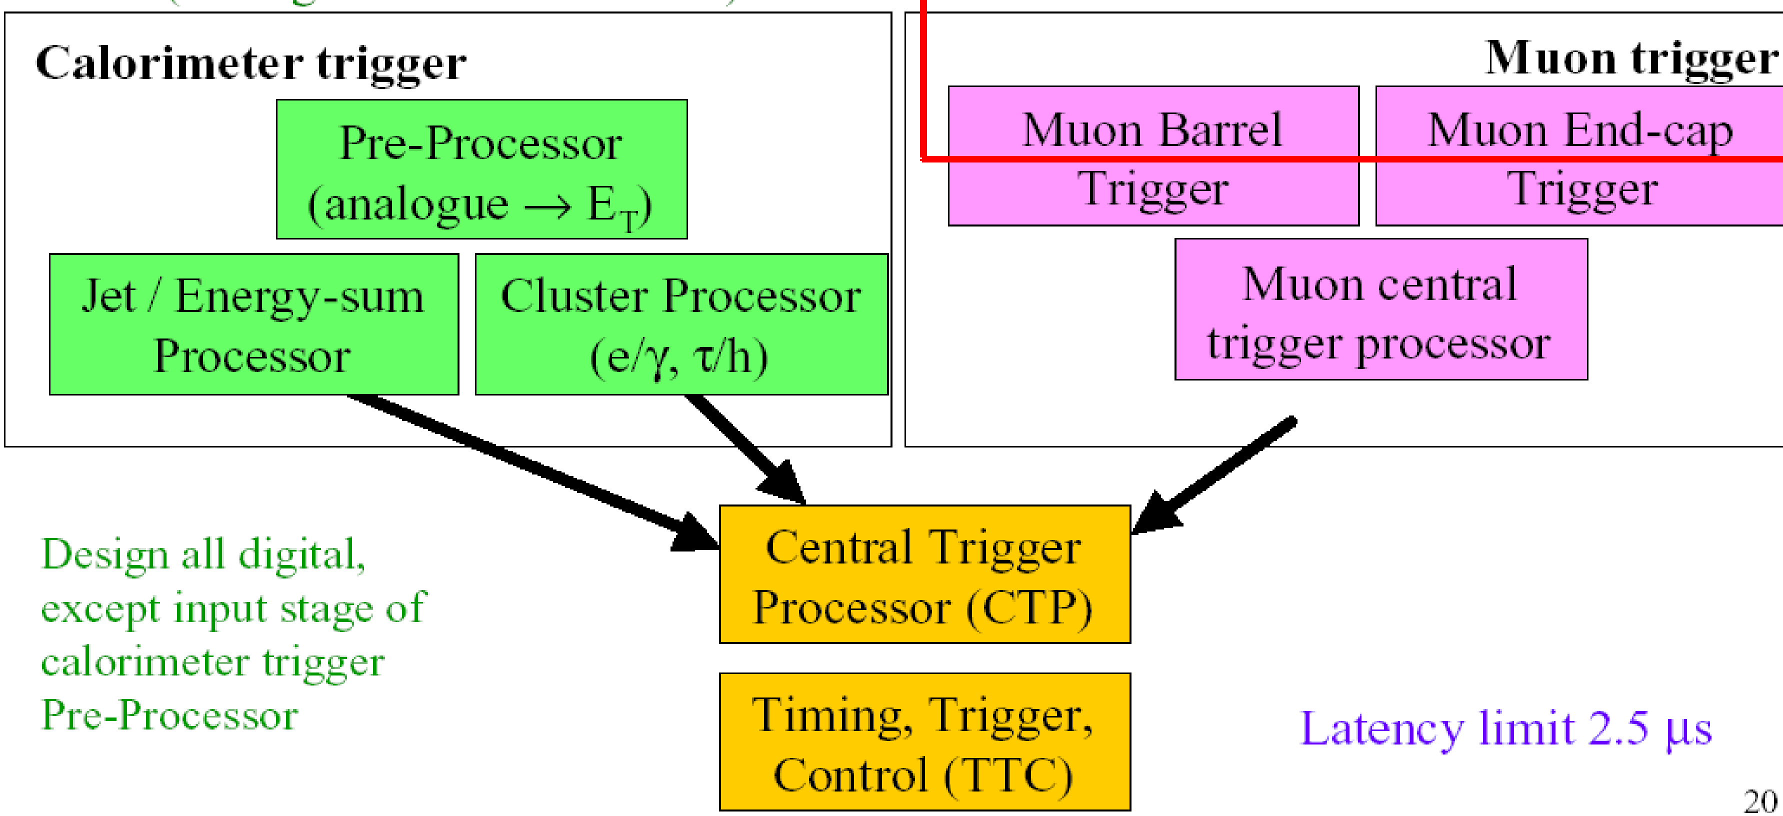

Design all digital, except input stage of calorimeter trigger Pre-Processor

Latency limit 2.5 μs

20

Lvl-1 trigger data flow: ATLAS

- **On-detector:**
  - analog sums to form trigger towers
- **Off-detector:**
  - Receive data, digitize, identify bunch crossing, compute E_T
  - Send data to Cluster Processor and Jet Energy Processor crates
- **Local processor crates:**
  - Form sums/comparisons as per algorithm, decide on objects found
- **Global Trigger: decision**

Level-1 Calorimeter Trigger Architecture

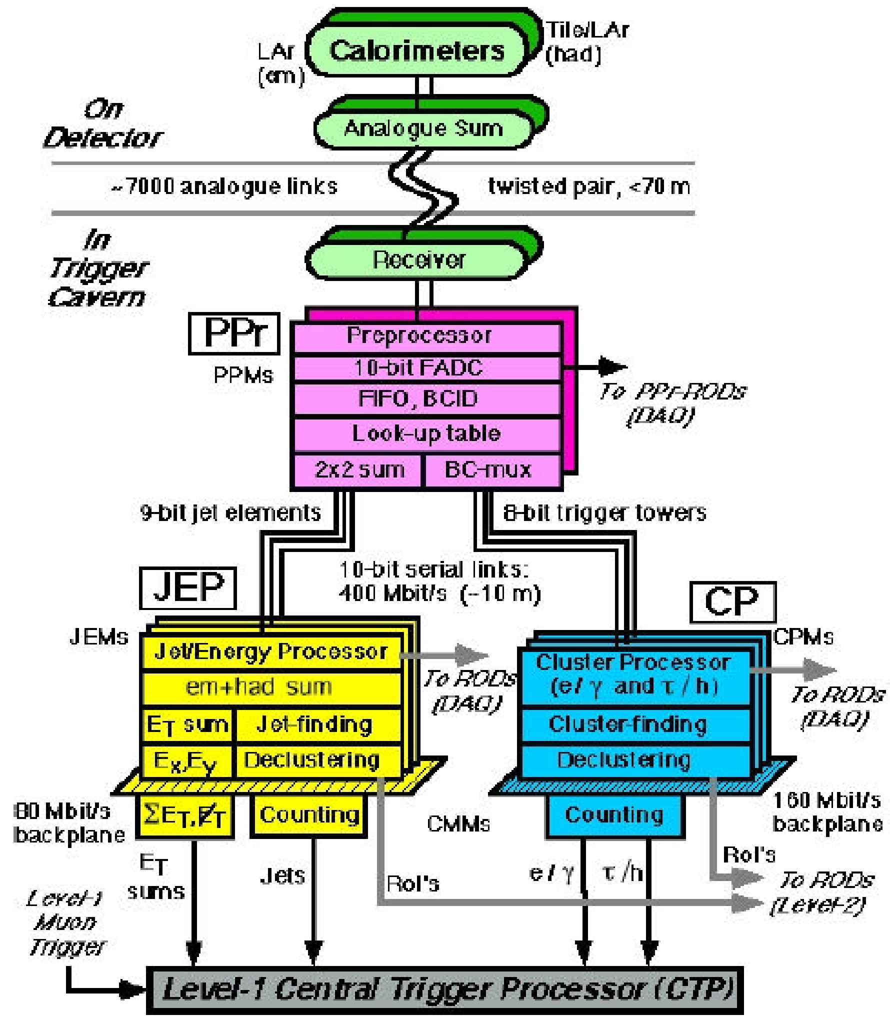

Lvl-1 Calo Trigger: e/γ algorithm (CMS)

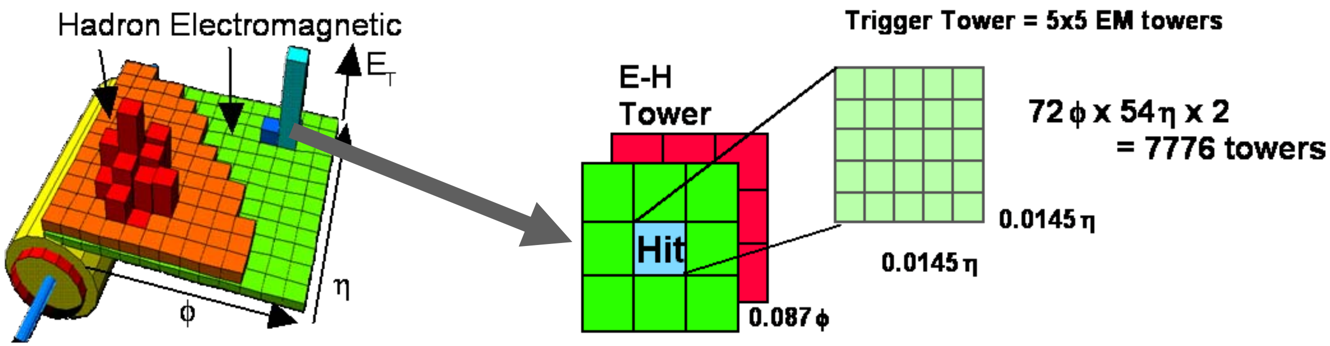

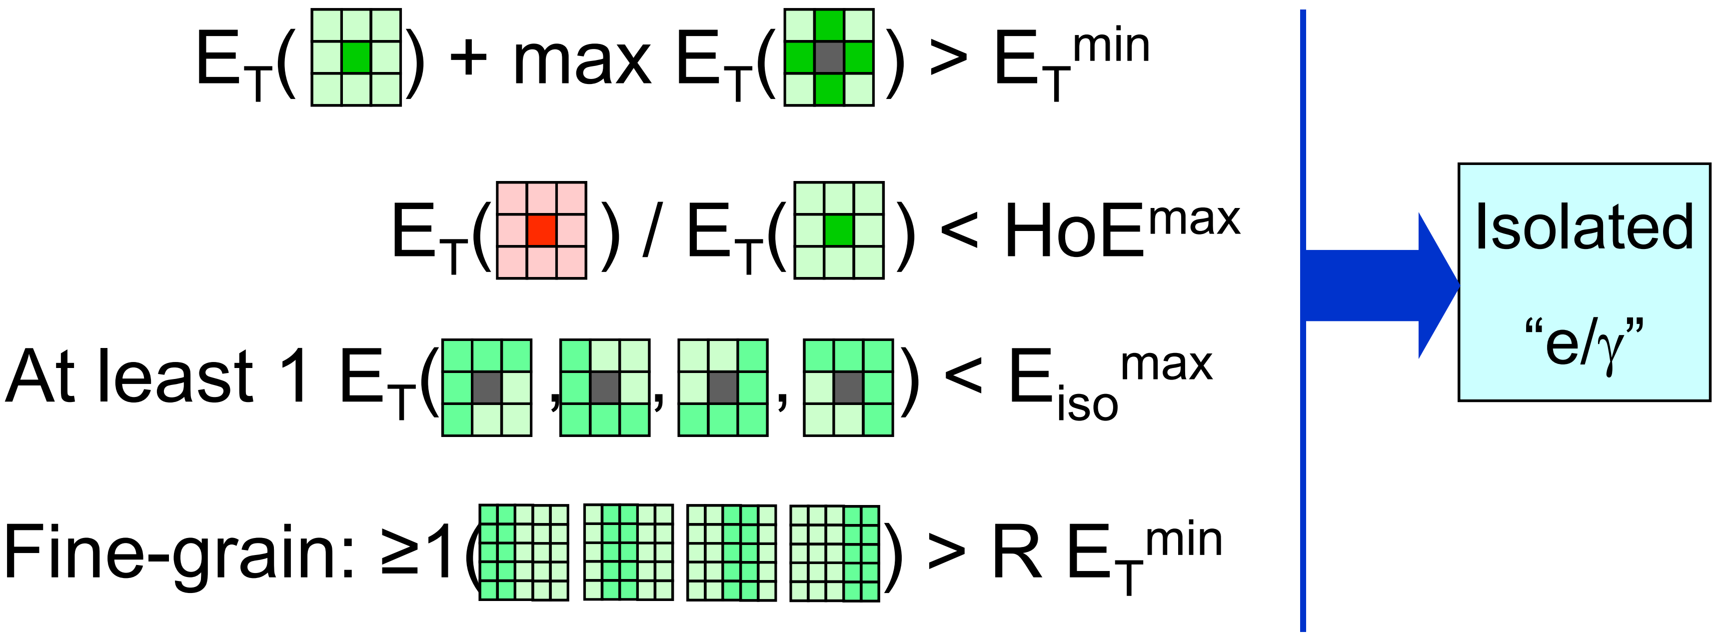

Lvl-1 Calo e/γ trigger: performance

- **Efficiencies and Trigger Rates**

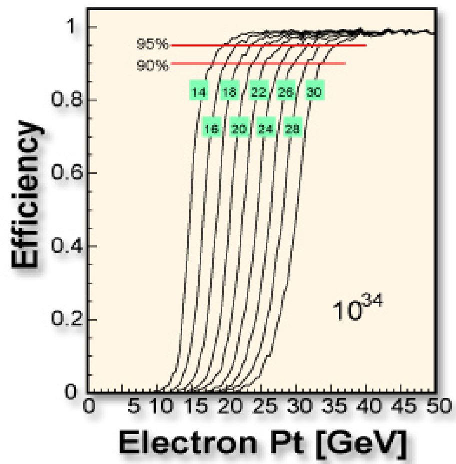

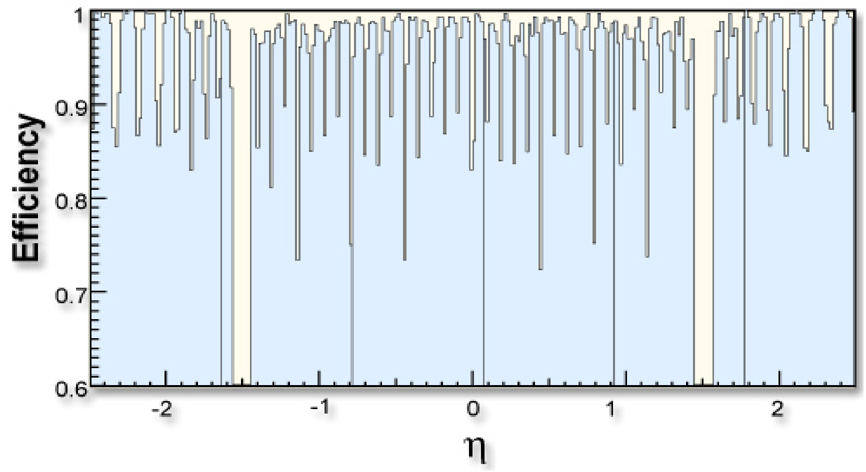

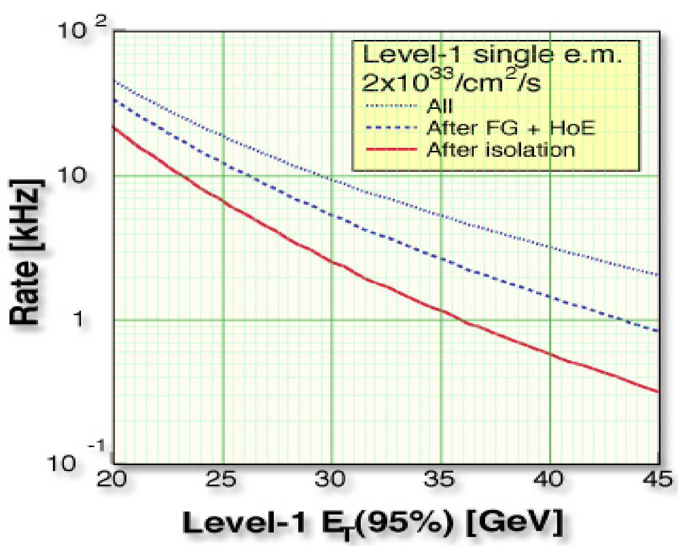

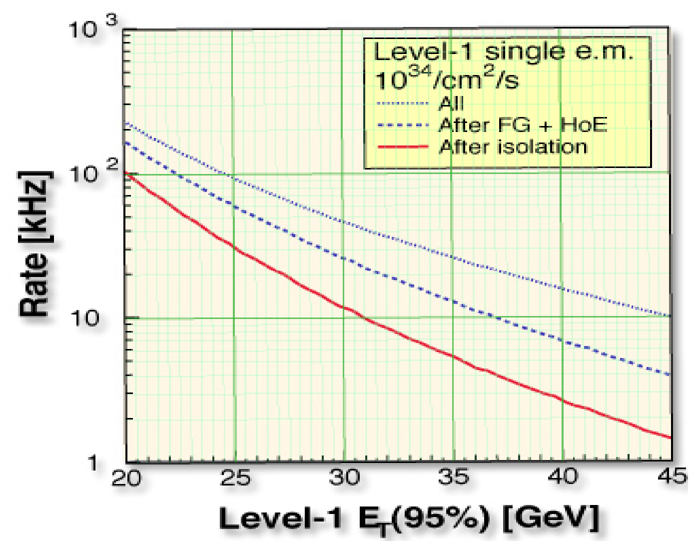

Lvl-1 jet and τ triggers

- **Issues are jet energy resolution and tau identification**
  - **Single, double, triple and quad thresholds possible**
  - **Possible also to cut on jet multiplicities**
  - **Also ETmiss, SET and SET(jets) triggers**

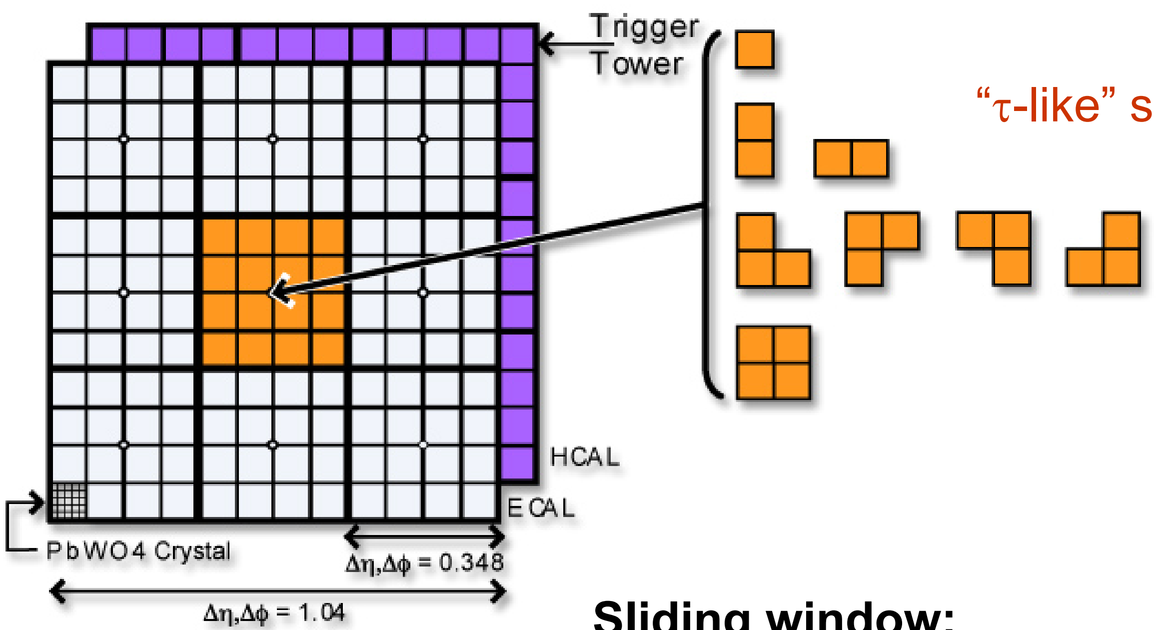

"τ-like" shapes identified for τ trigger

Sliding window:
- **granularity is 4x4 towers = trigger region**
- **jet E_T summed in 3x3 regions $\Delta\eta,\Delta\phi = 1.04$**

Lvl-1 muon trigger

- **The goal: measure momentum online**
  - Steeply falling spectrum; resolution costs!
- **The issue: speed**
  - ATLAS: dedicated muon chambers (RPC and TGC)
  - CMS: RPC added to DT and CSC (which provide standalone trigger)

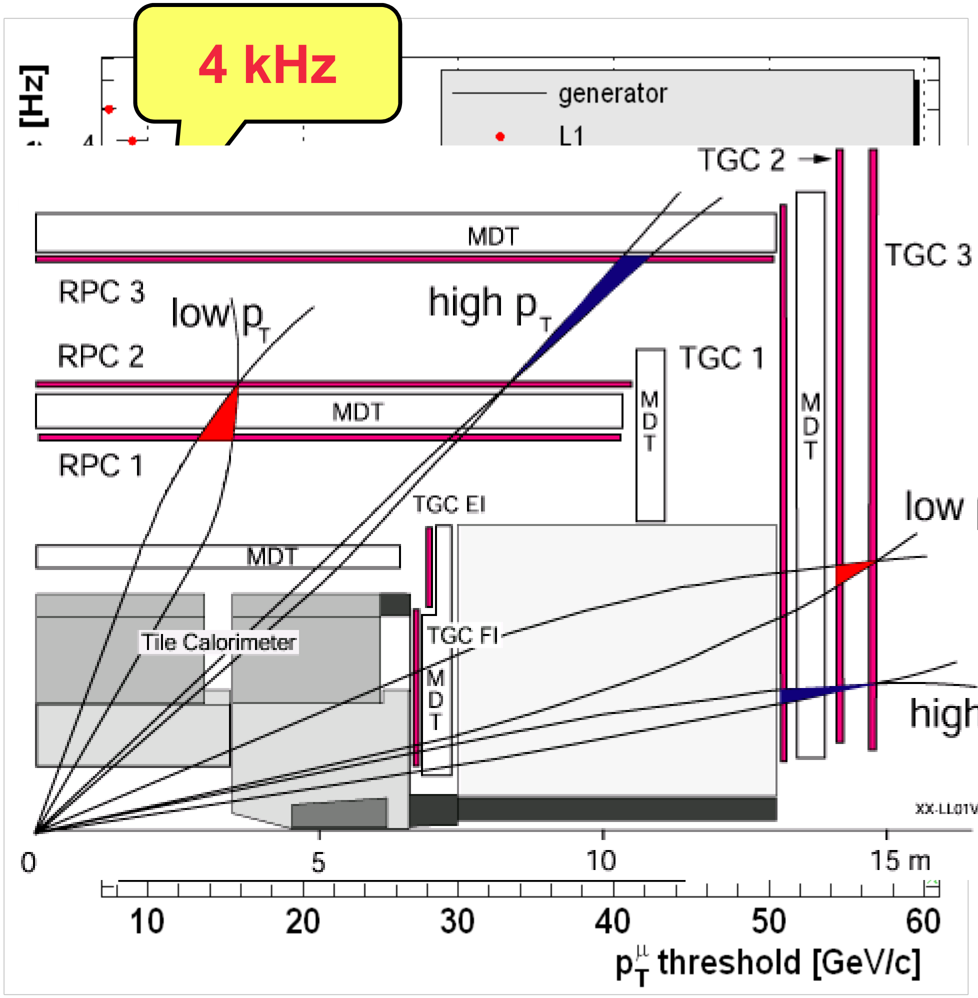

Lvl-1 muon trigger (CMS)

Drift Tubes

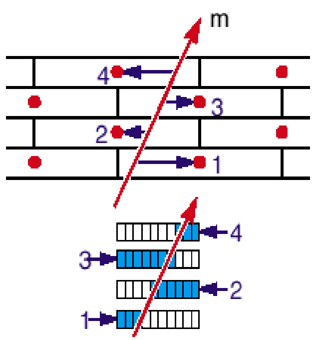

Meantimers recognize tracks and form vector / quartet.

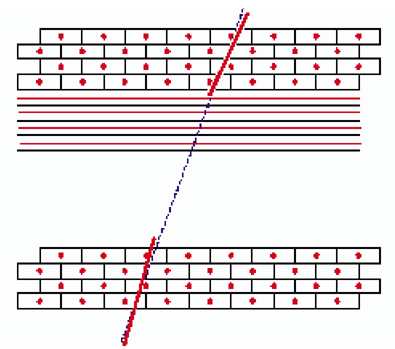

Correlator combines them into one vector / station.

CSC

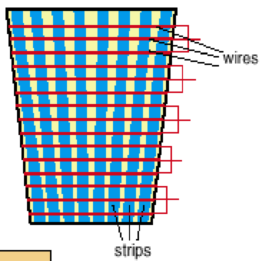

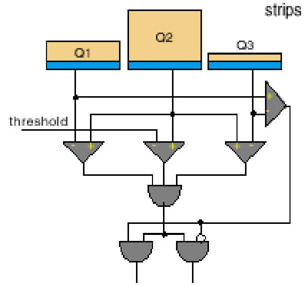

Comparators give 1/2-strip resol.

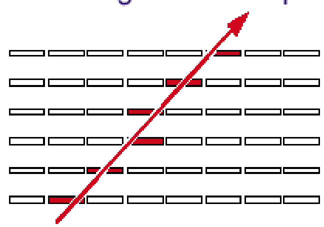

Hit strips of 6 layers form a vector.

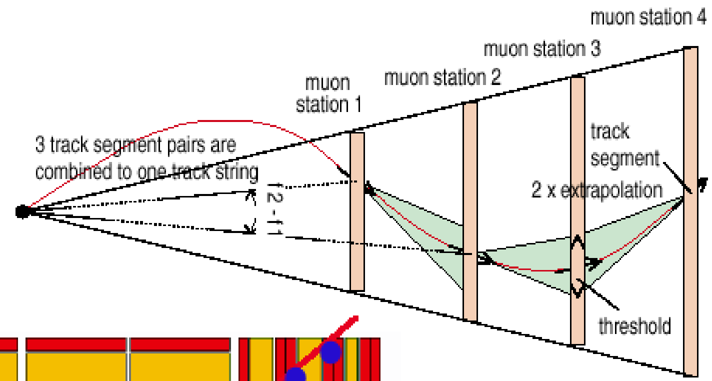

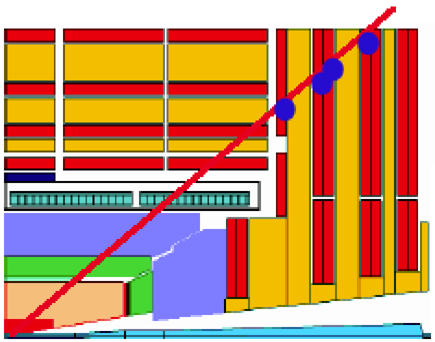

- Extrapolation: using look-up tables
- Track Assembler: link track segment-pairs to tracks, cancel fakes
- Assignment: P_T (5 bits), charge, η (6 bits) , φ(8 bits), quality (3 bits)

Hardware implementation:
ASICs for Trigger Primitive Generators
FPGAs for Track Finder processors

Lvl-1 muon trigger (CMS)

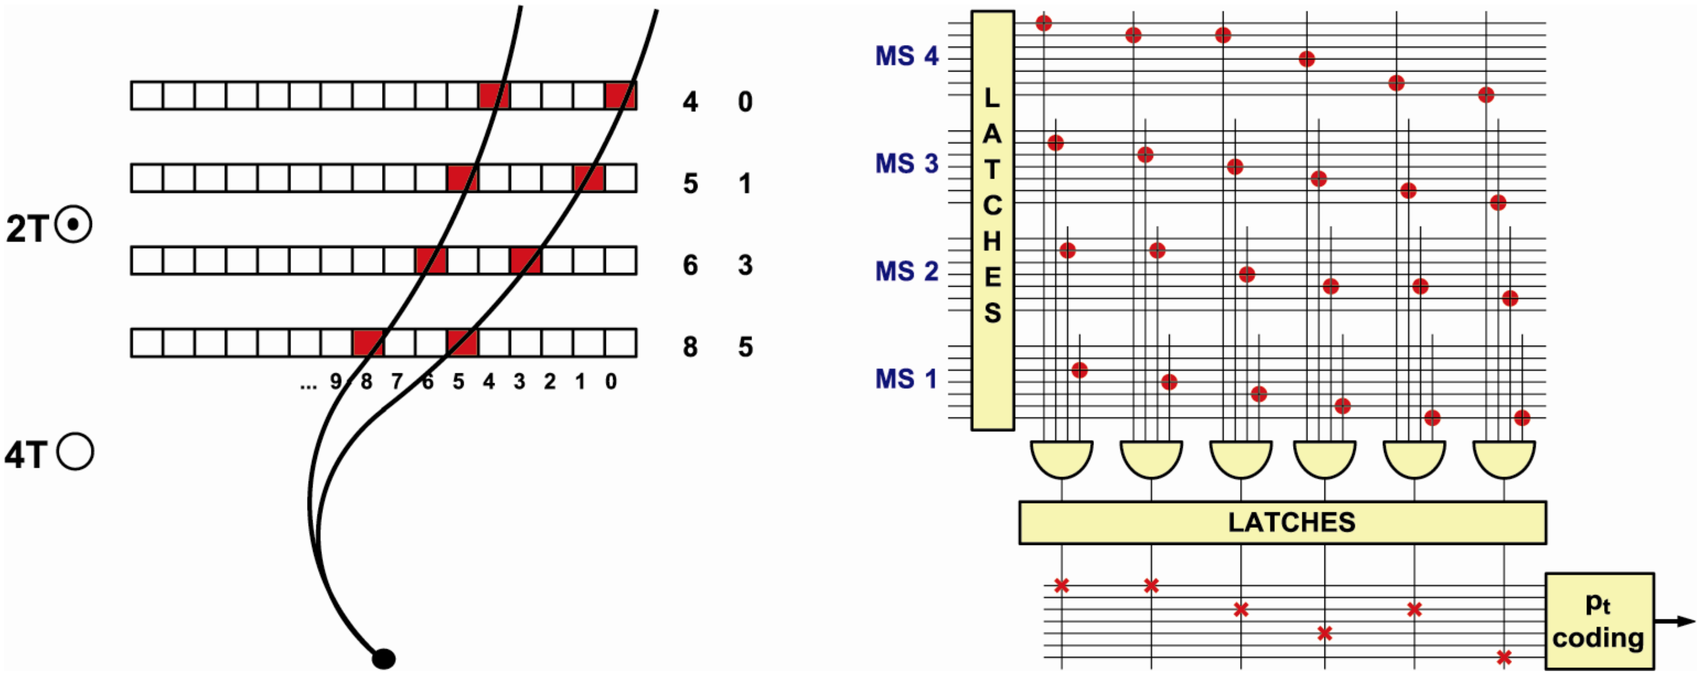

Pattern of strips hit:
Compared to predefined patterns corresponding to various p_T

Implemented in FPGAs

Global muon trigger (CMS)

- Combine results from RPC, CSC and DT triggers
- Match muon candidates from different trigger systems; use complementarity of detectors
- improve efficiency and rate
- assign muon isolation
- deliver the 4 best (highest P_T, highest-quality) muons to Global Trigger
- Pt resolution:
  - 18% barrel
  - 35% endcaps
- Efficiency: ~ 97%

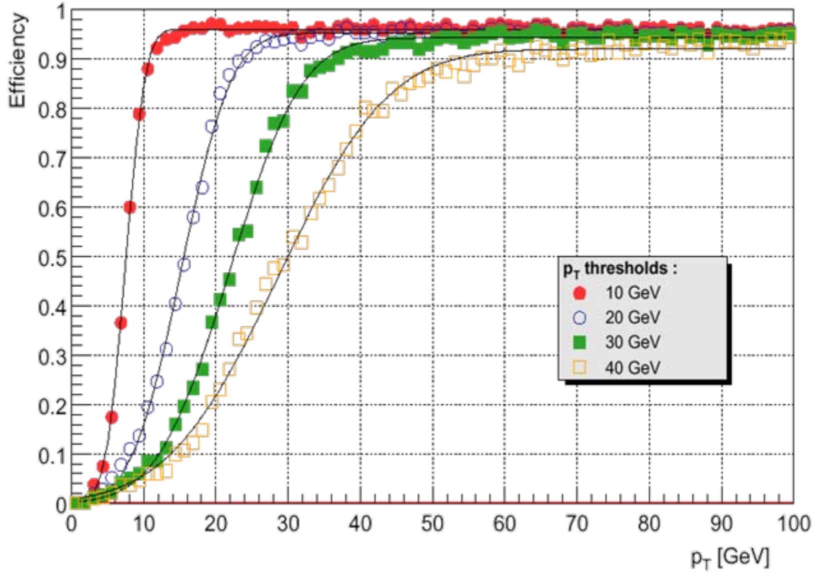

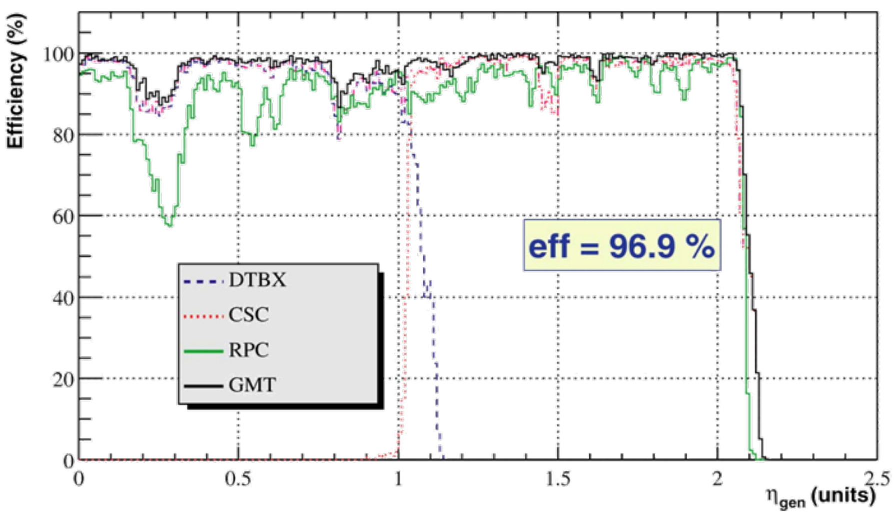

Technologies in Level-1 systems

- **ASICs (Application-Specific Integrated Circuits) used in some cases**
  - Highest-performance option, better radiation tolerance and lower power consumption (a plus for on-detector electronics)
- **FPGAs (Field-Programmable Gate Arrays) used throughout all systems**
  - Impressive evolution with time. Large gate counts and operating at 40 MHz (and beyond)
  - Biggest advantage: flexibility
    - Can modify algorithms (and their parameters) in situ
- **Communication technologies**
  - High-speed serial links (copper or fiber)
    - LVDS up to 10 m and 400 Mb/s; HP G-link, Vitesse for longer distances and Gb/s transmission
  - Backplanes
    - Very large number of connections, multiplexing data

      ➔ operating at ~160 Mb/s

Lvl-1 Calo Trigger: prototypes

Trigger Crate
(160 MHz backplane)

Back

Front

Receiver Card

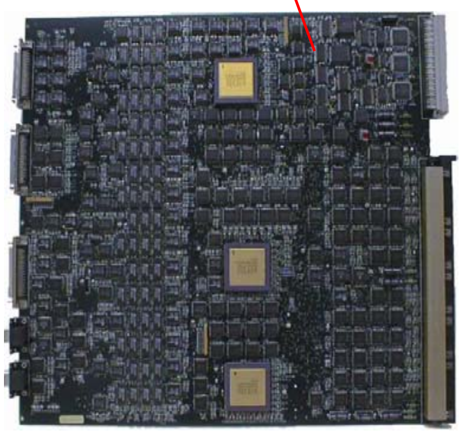

Links

Electron
(isolation)
Card

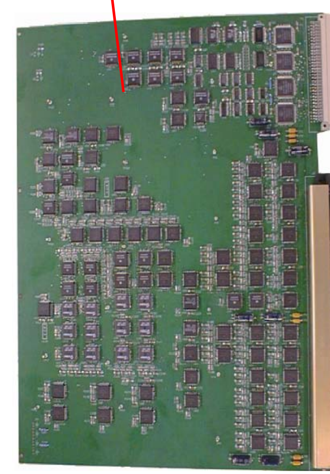

Bunch-crossing identification

- **Need to extract quantities of the bunch-crossing in question (and identify the xing)**
- **FIR (finite impulse response filter)**
  - **Feed LUT to get E_T**
  - **Feeds peak-finder to identify bunch-xing**
  - **Special handling of very large pulses (most interesting physics…)**
- **Can be done in an ASIC (e.g. ATLAS)**

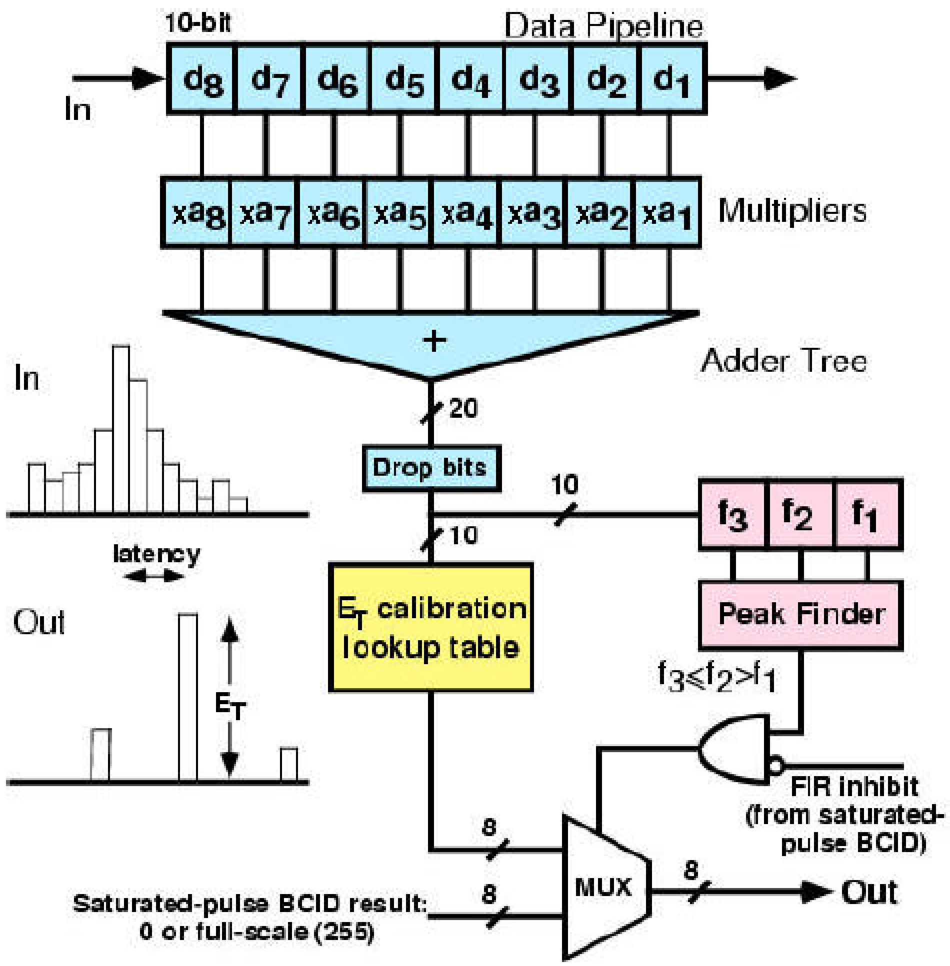

Global Trigger

- **A very large OR-AND network that allows for the specification of complex conditions:**
  - 1 electron with P_T>20 GeV OR 2 electrons with P_T>14 GeV OR 1 electron with P_T>16 and one jet with P_T>40 GeV…
  - The top-level logic requirements (e.g. 2 electrons) constitute the "trigger-table" of the experiment
    - Allocating this rate is a complex process that involves the optimization of physics efficiencies vs backgrounds, rates and machine conditions
      - → More on this in the HLT part

Lvl-1 trigger: summary

- Some challenges of unprecedented scale
  - Interaction rate and selectivity
  - Number of channels and synchronization
  - Pile-up and bunch-crossing identification
  - Deciding on the fate of an event given ~3 μs
    - Of which most is spent in transportation
- Trigger levels: the set of successive approximations (at the ultimate save-or-kill decision)
  - Number of physical levels varies with architecture/experiment
- Level-1 is always there, reduces 40 MHz to 40-100 kHz
  - Level-0 may be used to (a) reduce initial rate to ~ 1MHz allow for slightly more complex processing (e.g. simple tracking)

DAQ system

Physics selection at the LHC

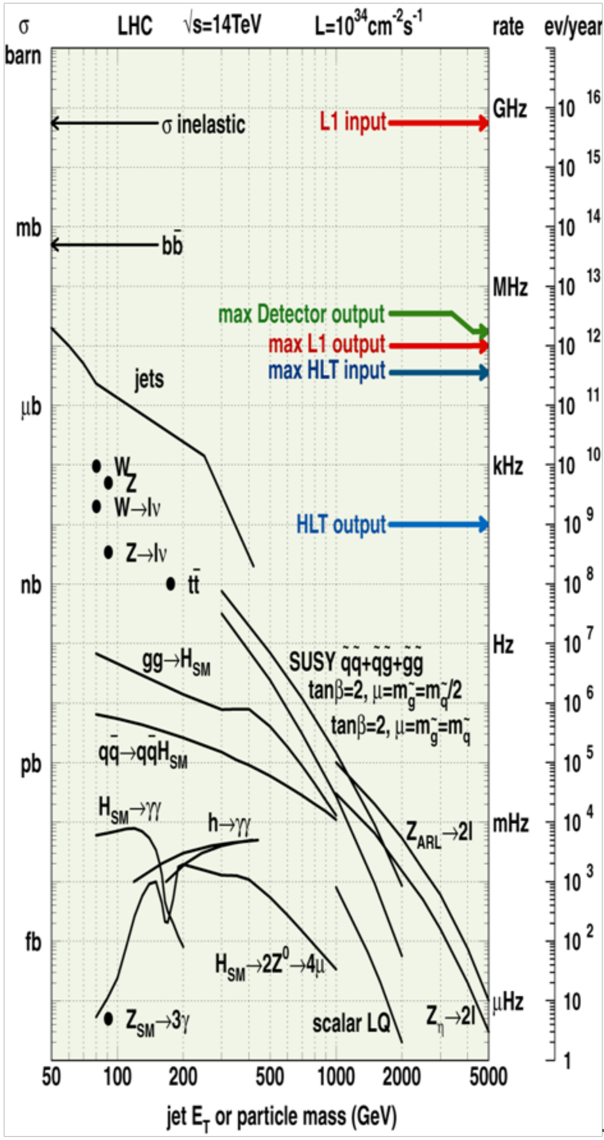

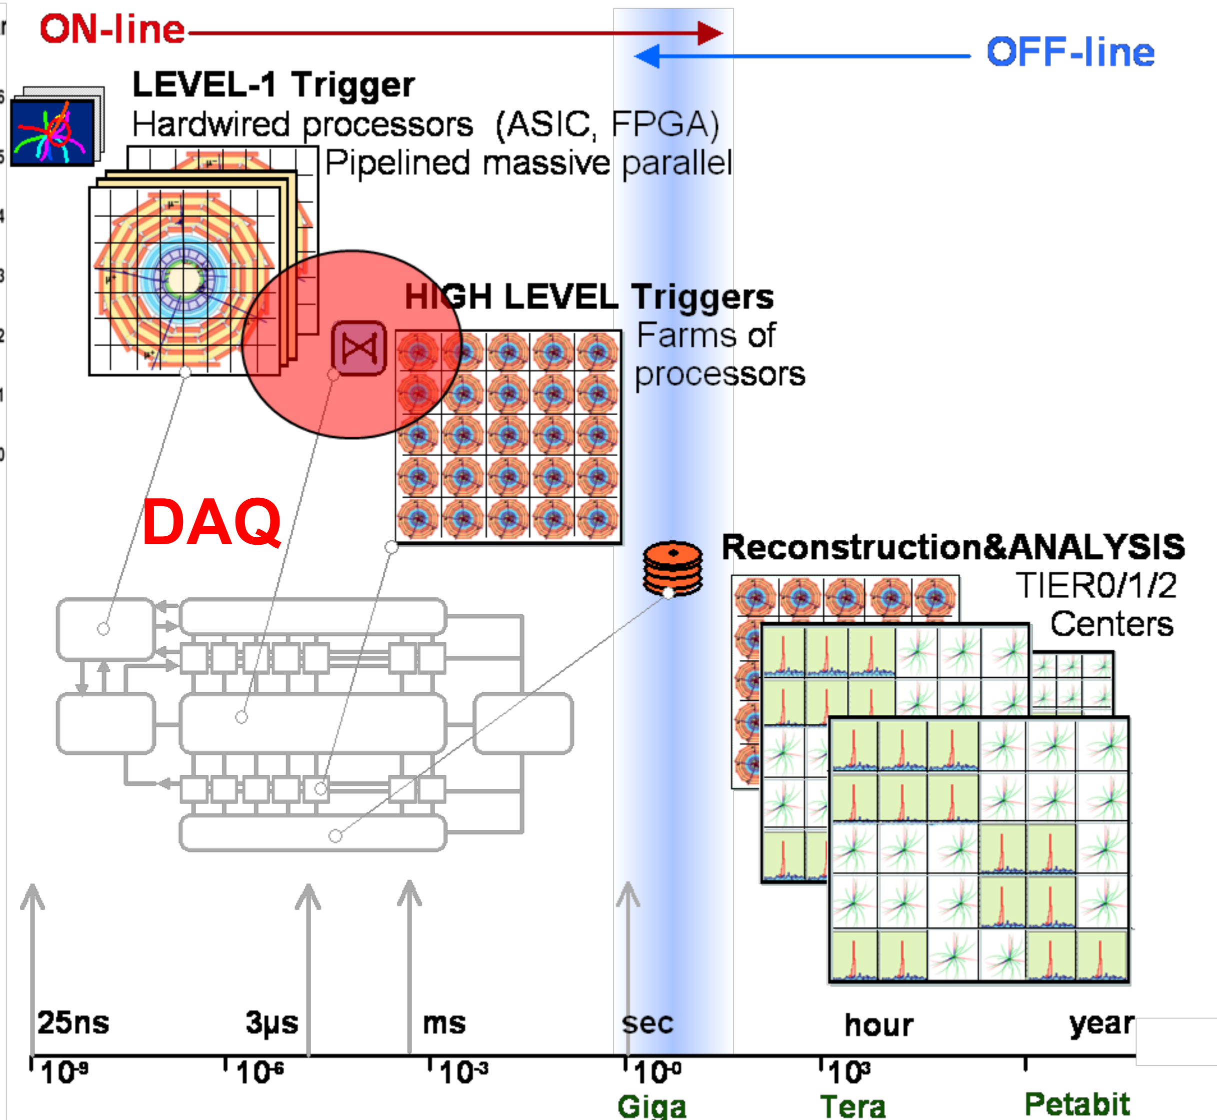

Technology evolution

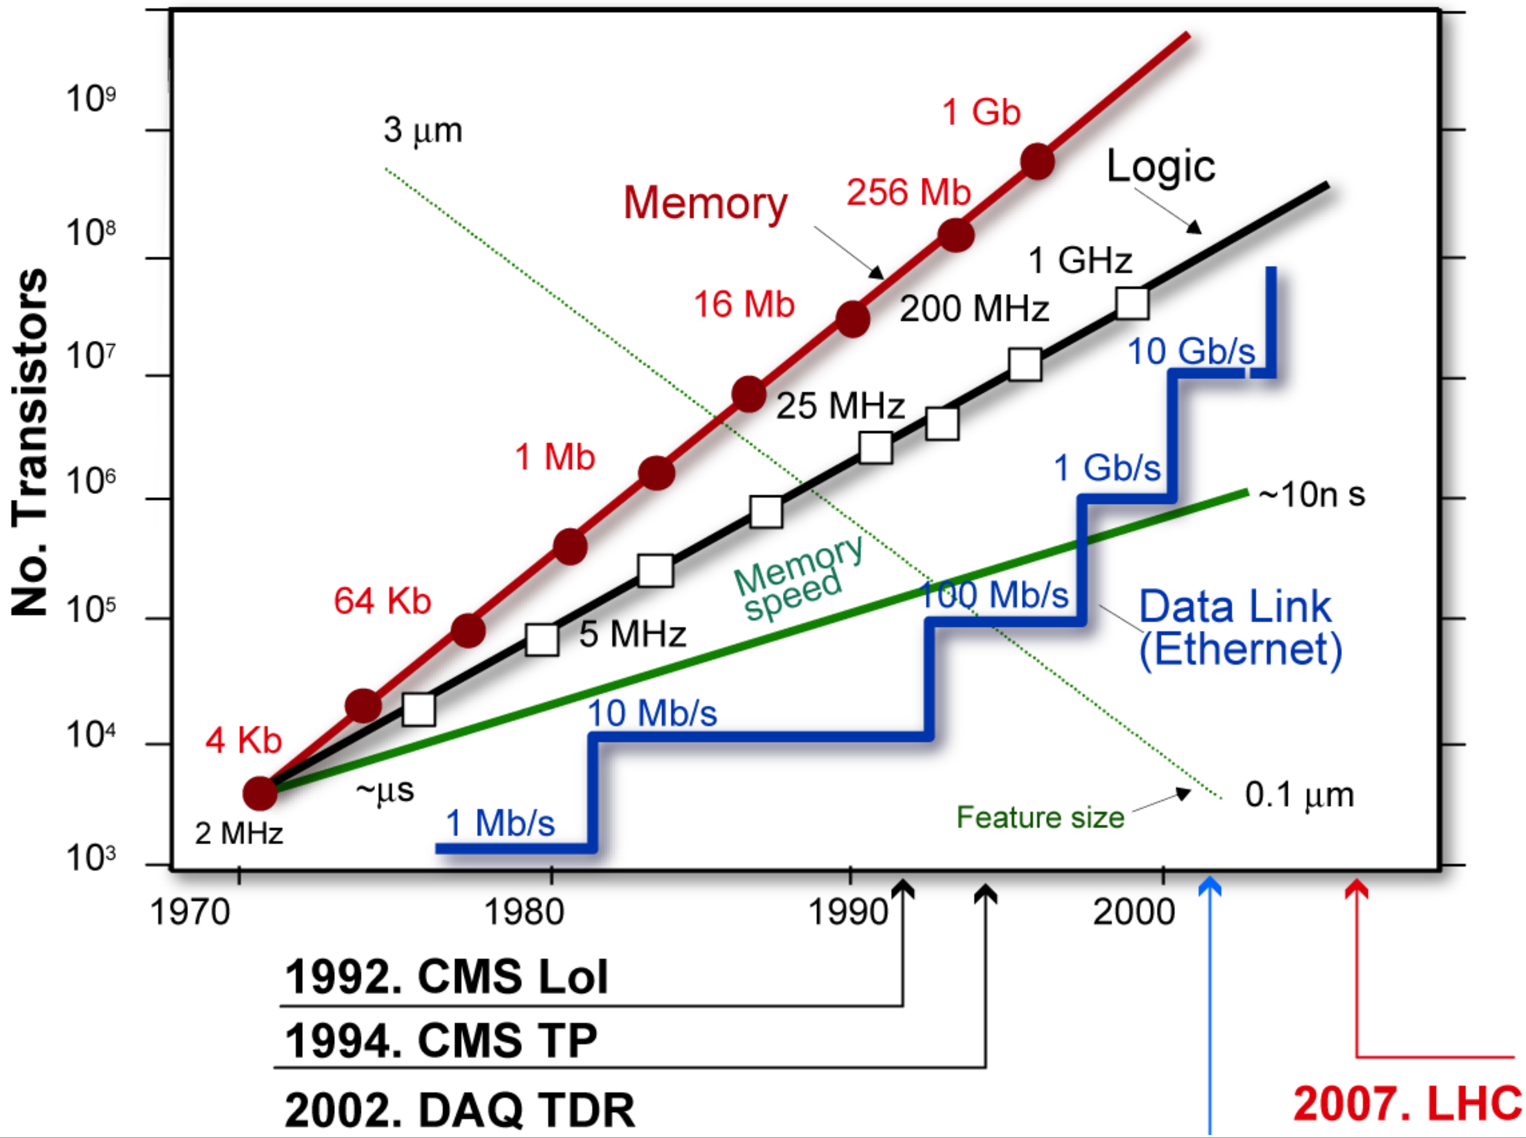

Online Selection Flow in pp

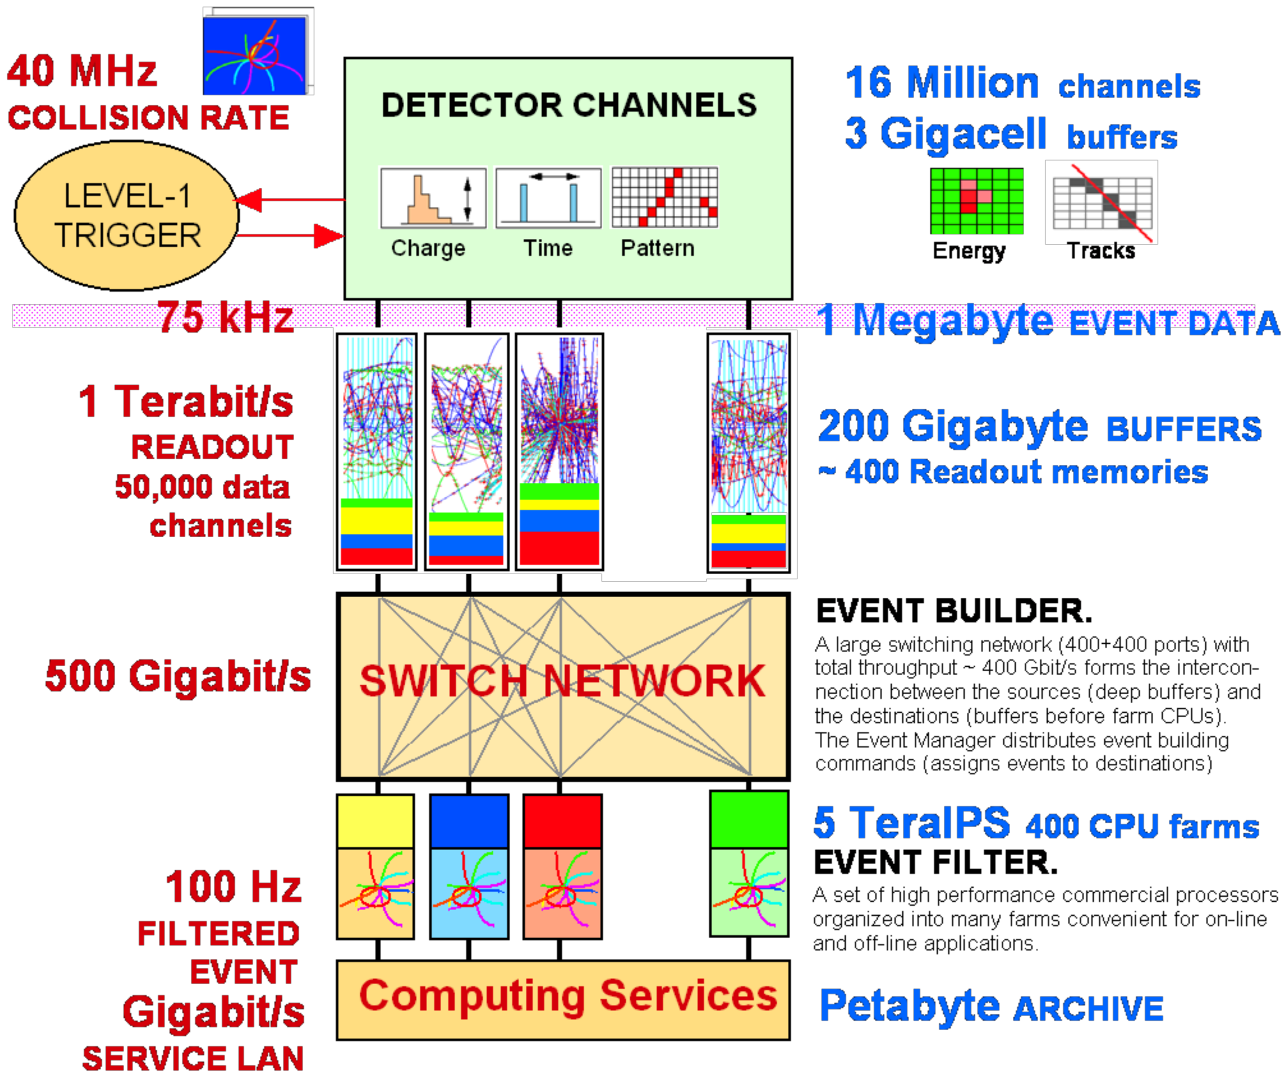

Trigger/DAQ systems: grand view

ATLAS

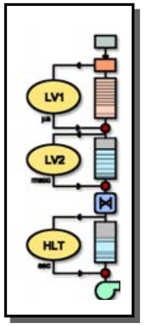

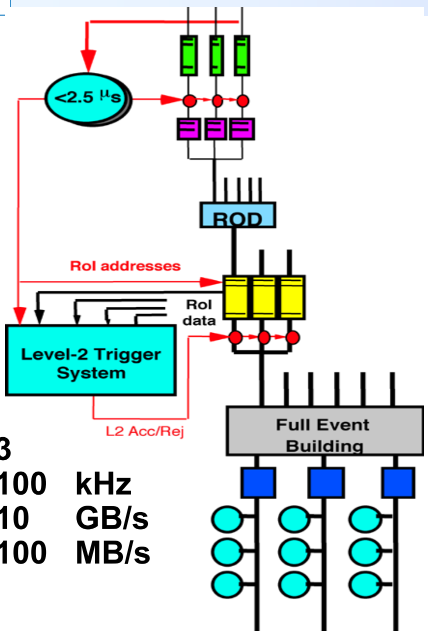

Levels	3	
LV-1 rate	100	kHz
Readout	10	GB/s
Storage	100	MB/s

ALICE

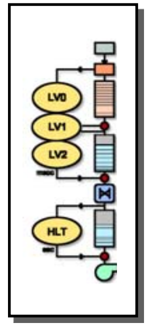

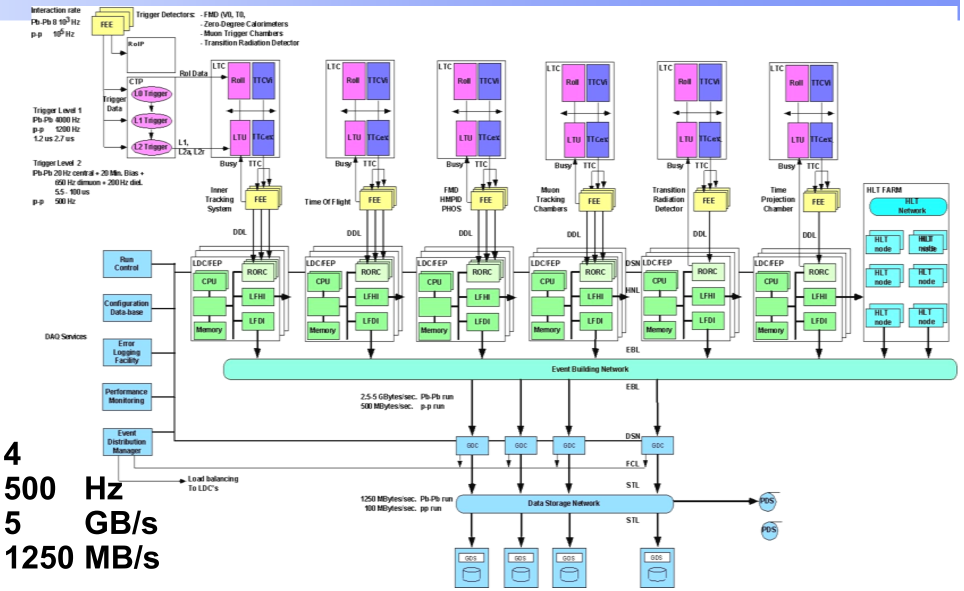

Levels	4	
LV-1 rate	500	Hz
Readout	5	GB/s
Storage	1250	MB/s

CMS

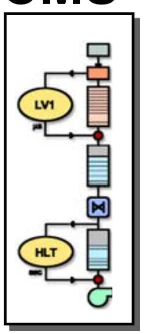

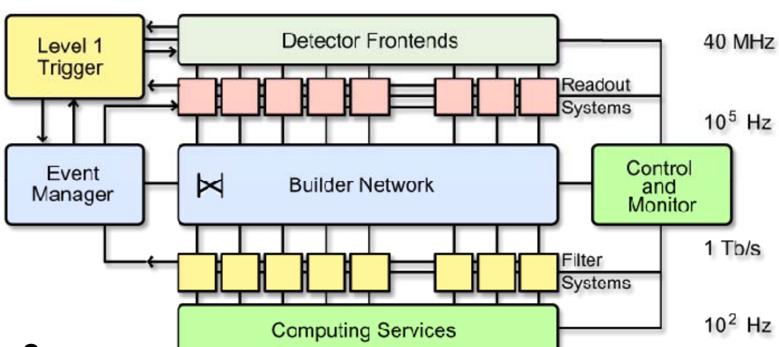

Levels	2	
LV-1 rate	100	kHz
Readout	100	GB/s
Storage	100	MB/s

LHCb

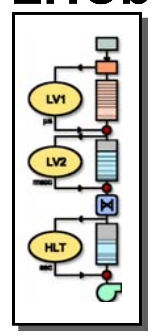

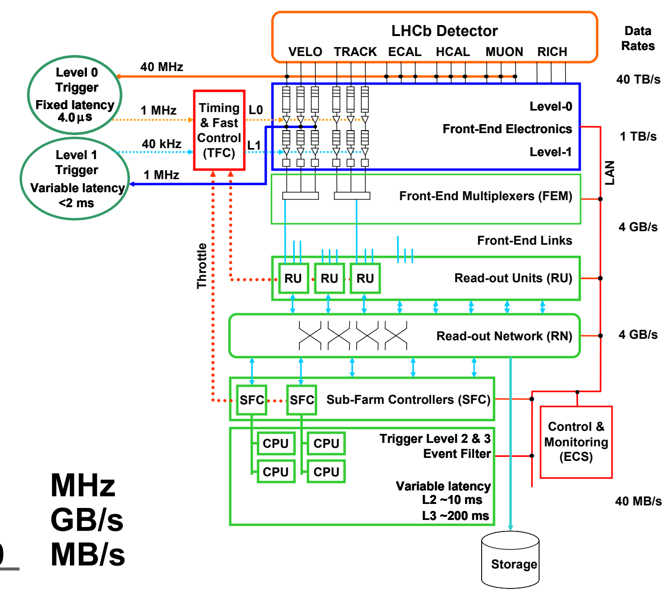

Levels	3	
LV-1 rate	1	MHz
Readout	4	GB/s
Storage	40	MB/s

Trigger/DAQ: basic blocks

- **Current Trigger/DAQ elements**

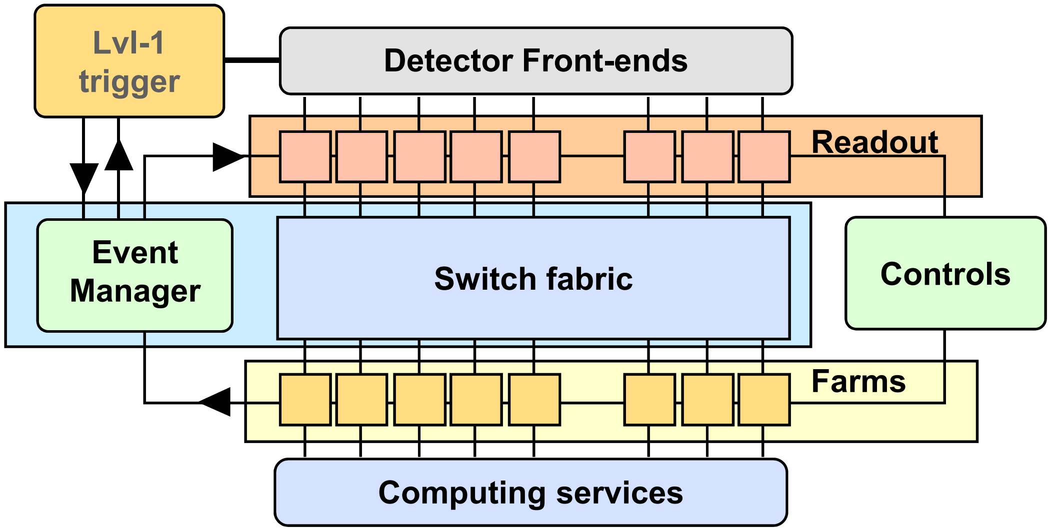

Detector Front-ends, feed Lvl-1 trigger processor

Readout Units: buffer events accepted by Lvl-1 trigger

Switching network: interconnectivity with HLT processors

Processor Farm

+ control and monitor

Readout types

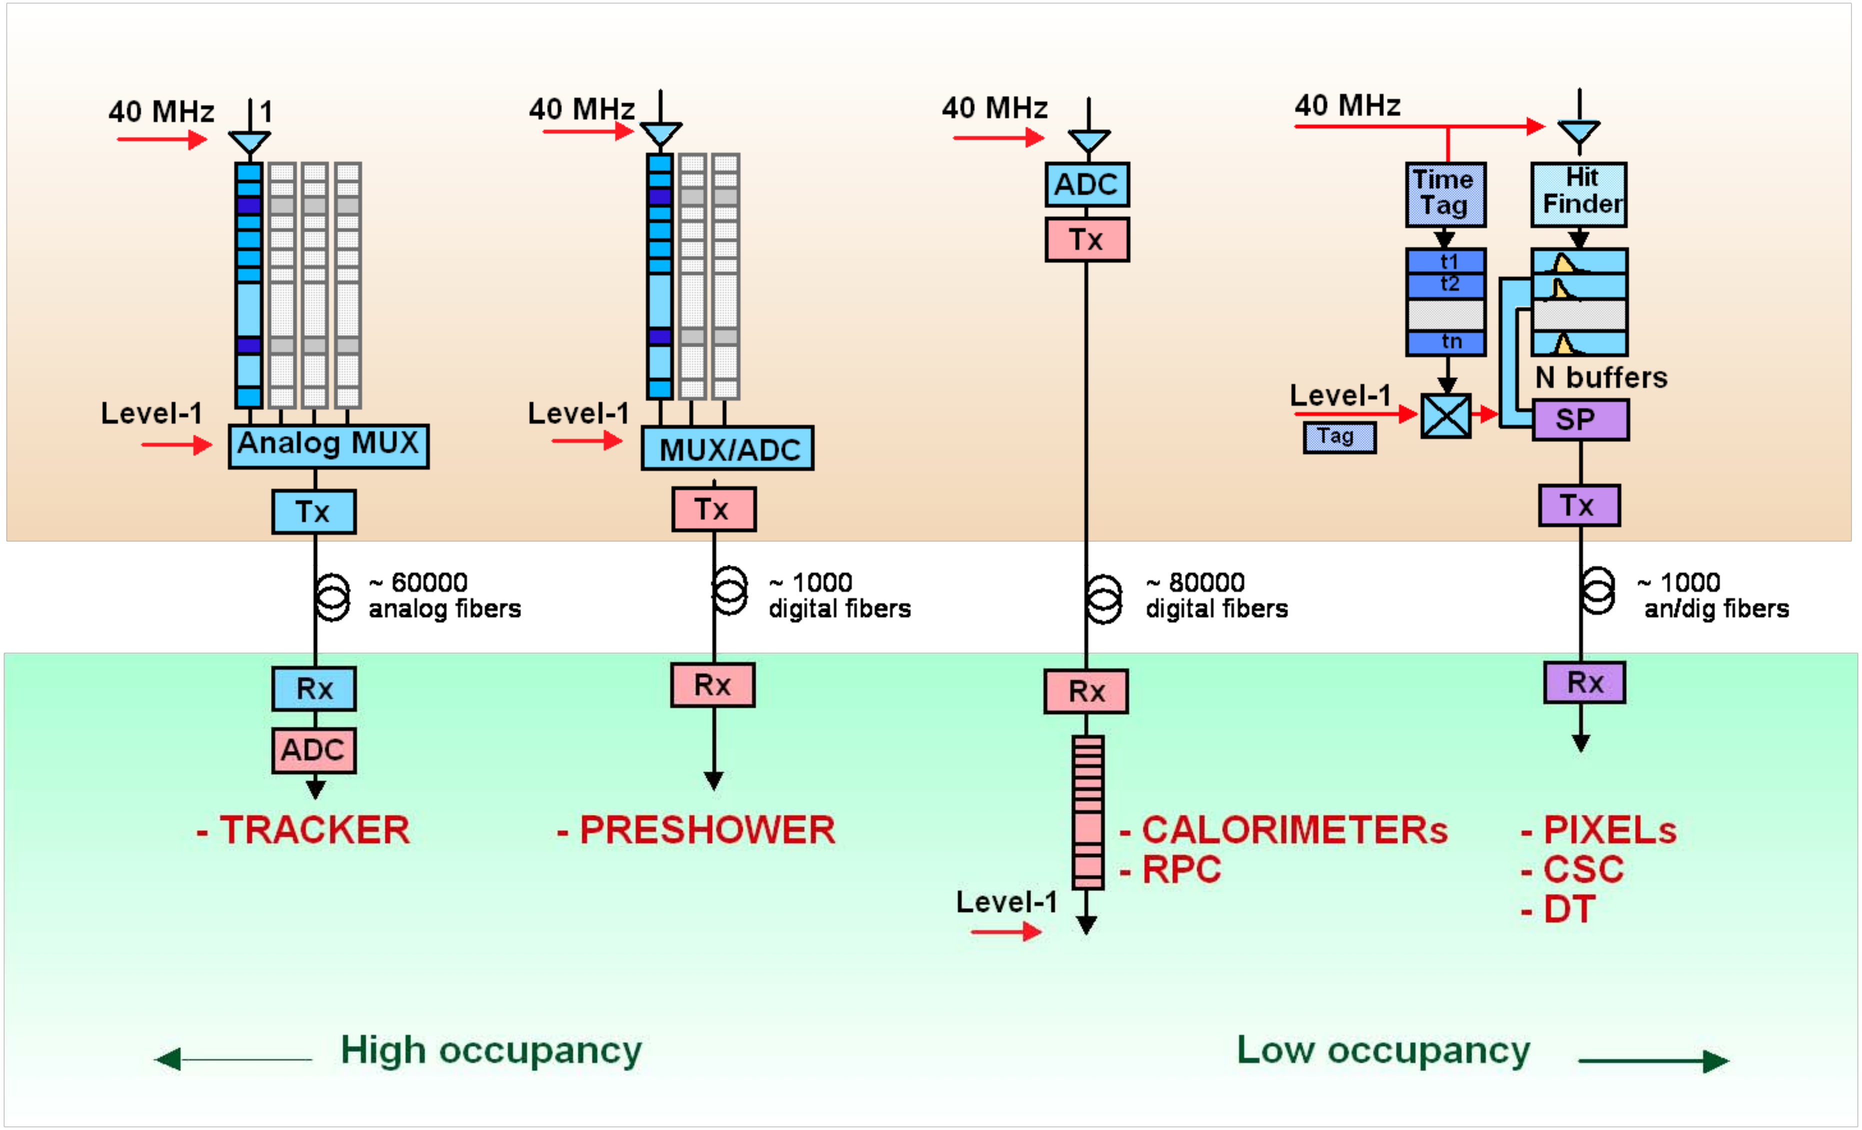

Need standard interface to front-ends

- **Large number of independent modules**

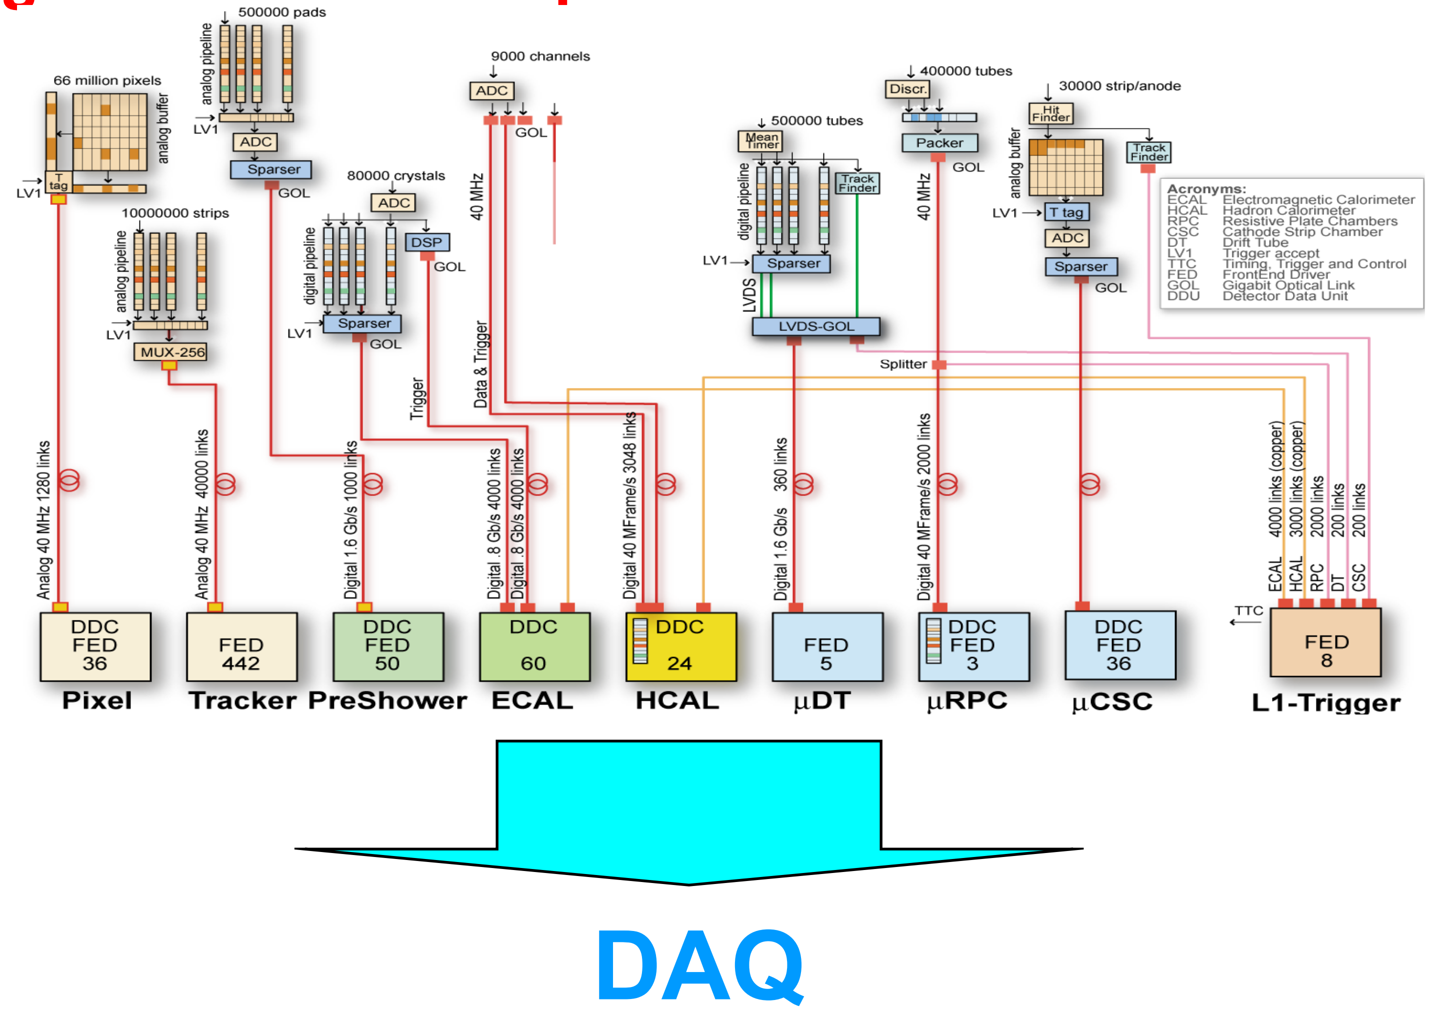

Event Building

- **Form full-event-data buffers from fragments in the readout. Must interconnect data sources/destinations.**

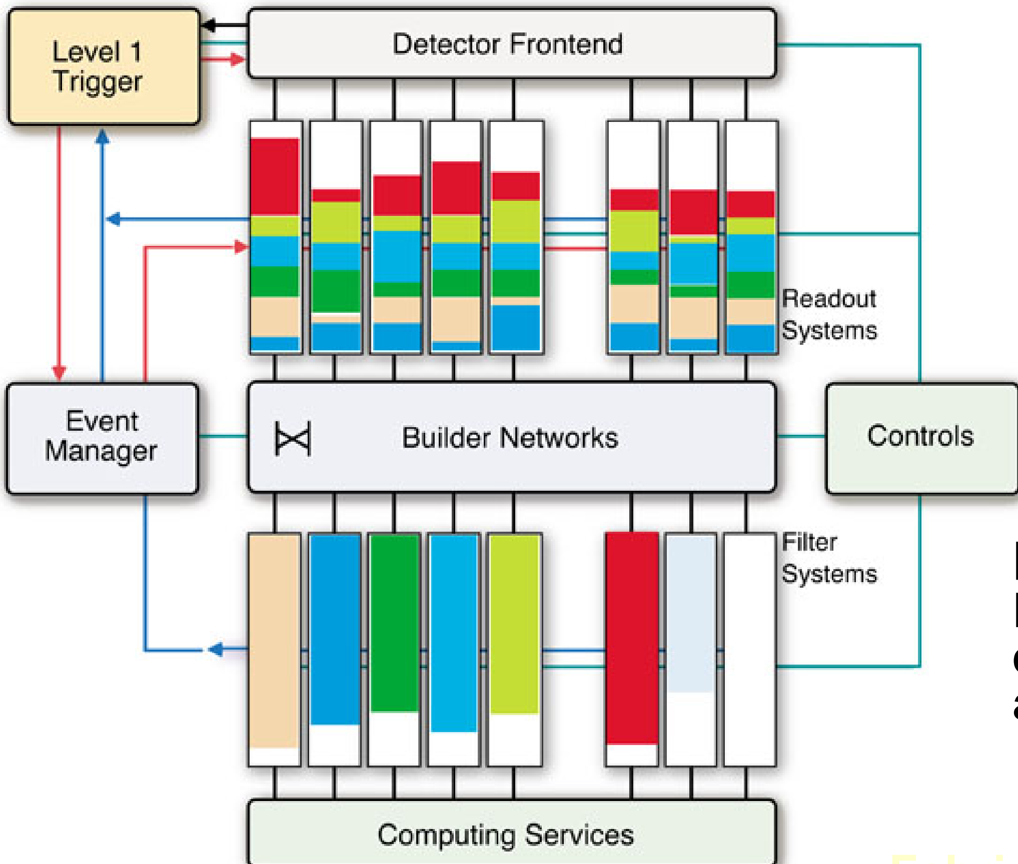

Event fragments :
Event data fragments are stored in separated physical memory systems

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

Event Building via a Switch

- **Three major issues:**
  - **Link utilization**
  - **The bottleneck on the outputs**
  - **The large number of ports needed**
- **Space-division: crossbar**
  - **Simultaneous transfers between any arbitrary set of inputs and outputs**
    - **Can be both self-routing and arbiter-based (determine connectivity between S's and D's for each cycle); the faster the fabric, the smaller the arbitration complexity**
    - **Does not solve Output Contention issue**
    - **Need *Traffic Shaping***

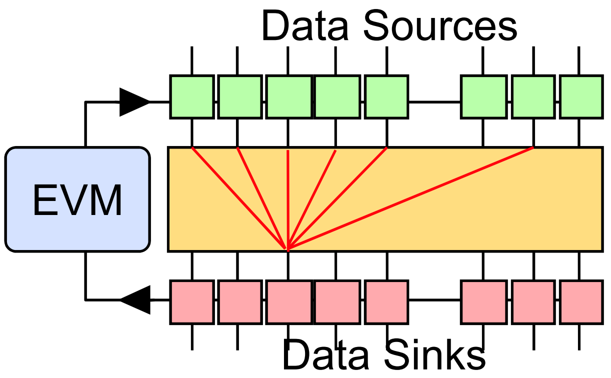

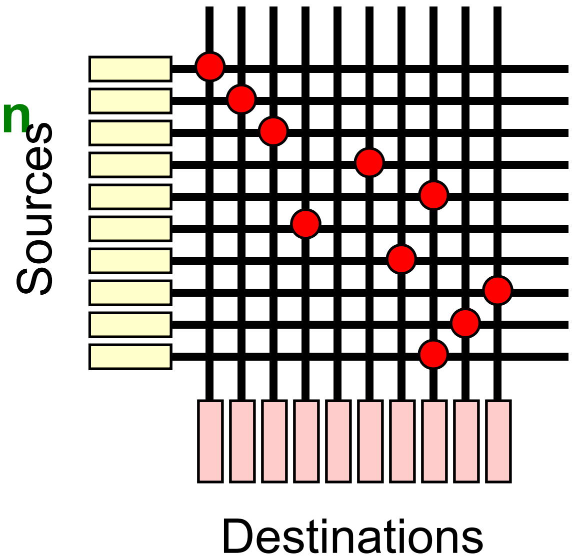

EVB traffic shaping: barrel shifter

- **Barrel-shifter: principle**

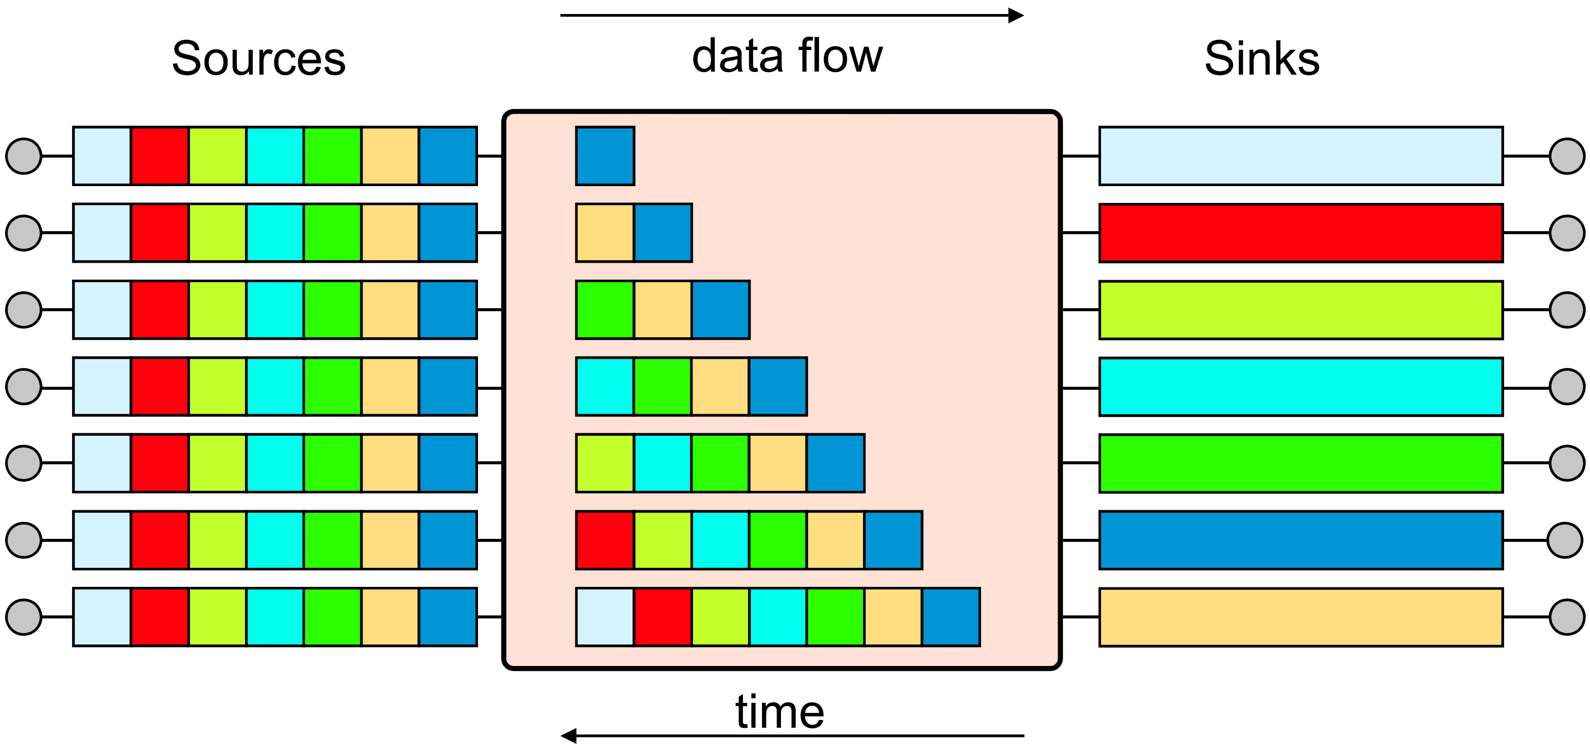

Barrel-shifting with variable-size events

- **Demonstrator**
  - Fixed-block-size with barrel-shifter
  - Basic idea taken from ATM (and time-division-muxing)
  - As seen in composite-switch analysis, this should work for large N as well
  - Currently testing on 64x64… (originally: used simulation for N≈500; now ~obsolete)

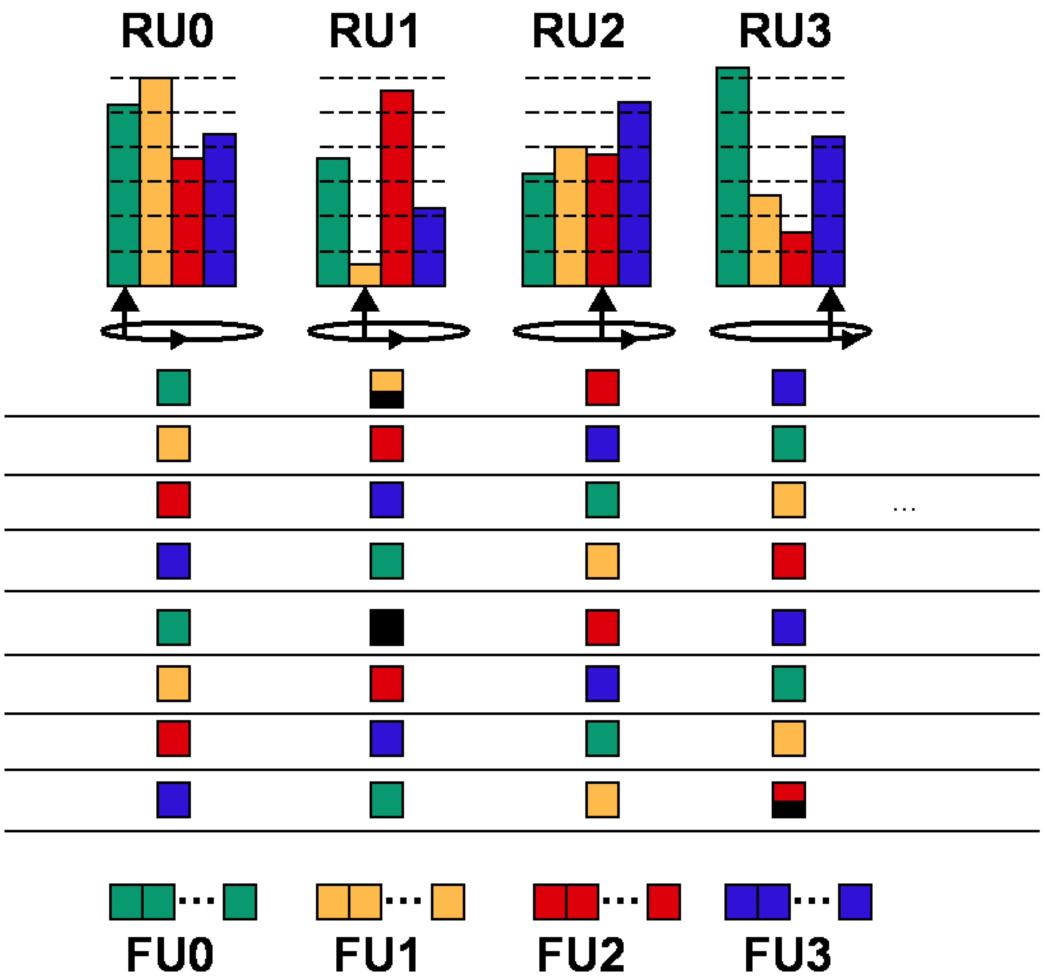

Detector readout & 3D-EVB

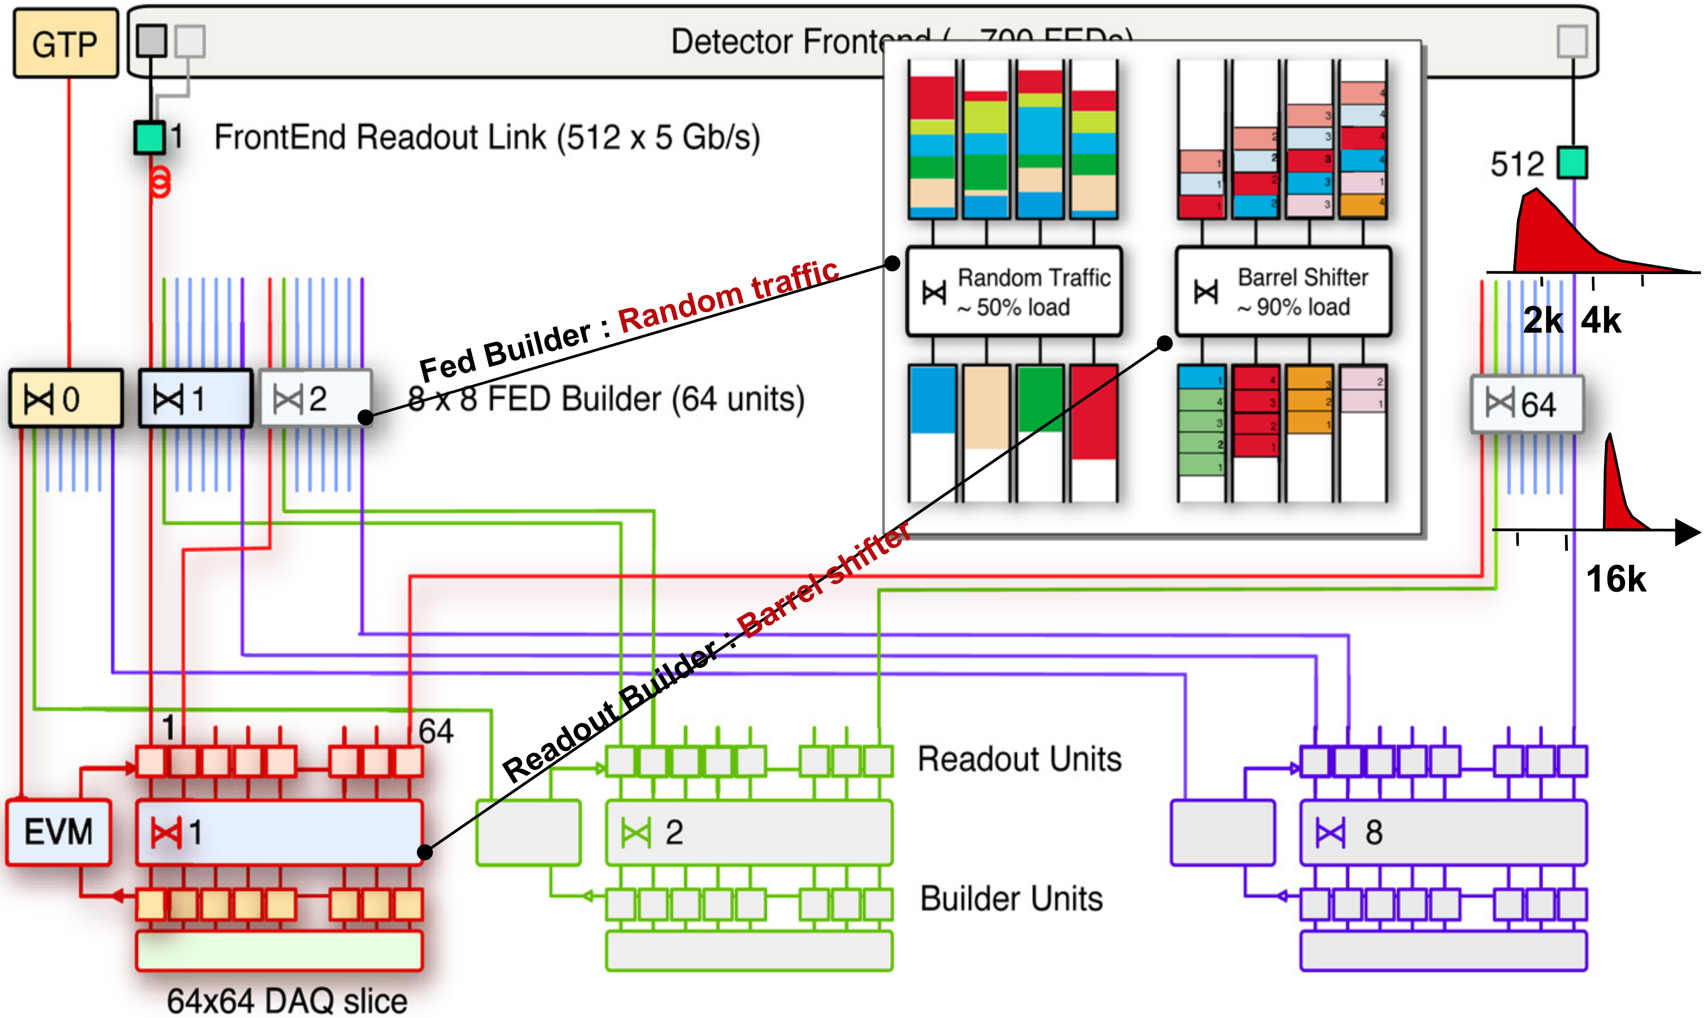

Control & Monitor

- **Challenges:**
  - Large N (on everything)
  - Disparity in time scales (μs–s; from readout to filtering)
  - Need to use standards for
    - Communication (Corba? Dead! "now": SOAP!)
    - User Interface (is it the Web? Yes…)
  - Physics monitoring complicated by factor 500 (number of sub-farms);
    - Need merging of information; identification of technical, one-time problems vs detector problems
- **Current work:**
  - Create toolkits from commercial software (SOAP, XML, HTTP etc); integrate into packages, build "Run Control" on top of it;
- **Detector Control System: DCS. All of this for the ~10^7 channels… SCADA (commercial, standard) solutions**

High-Level Trigger

Physics selection at the LHC

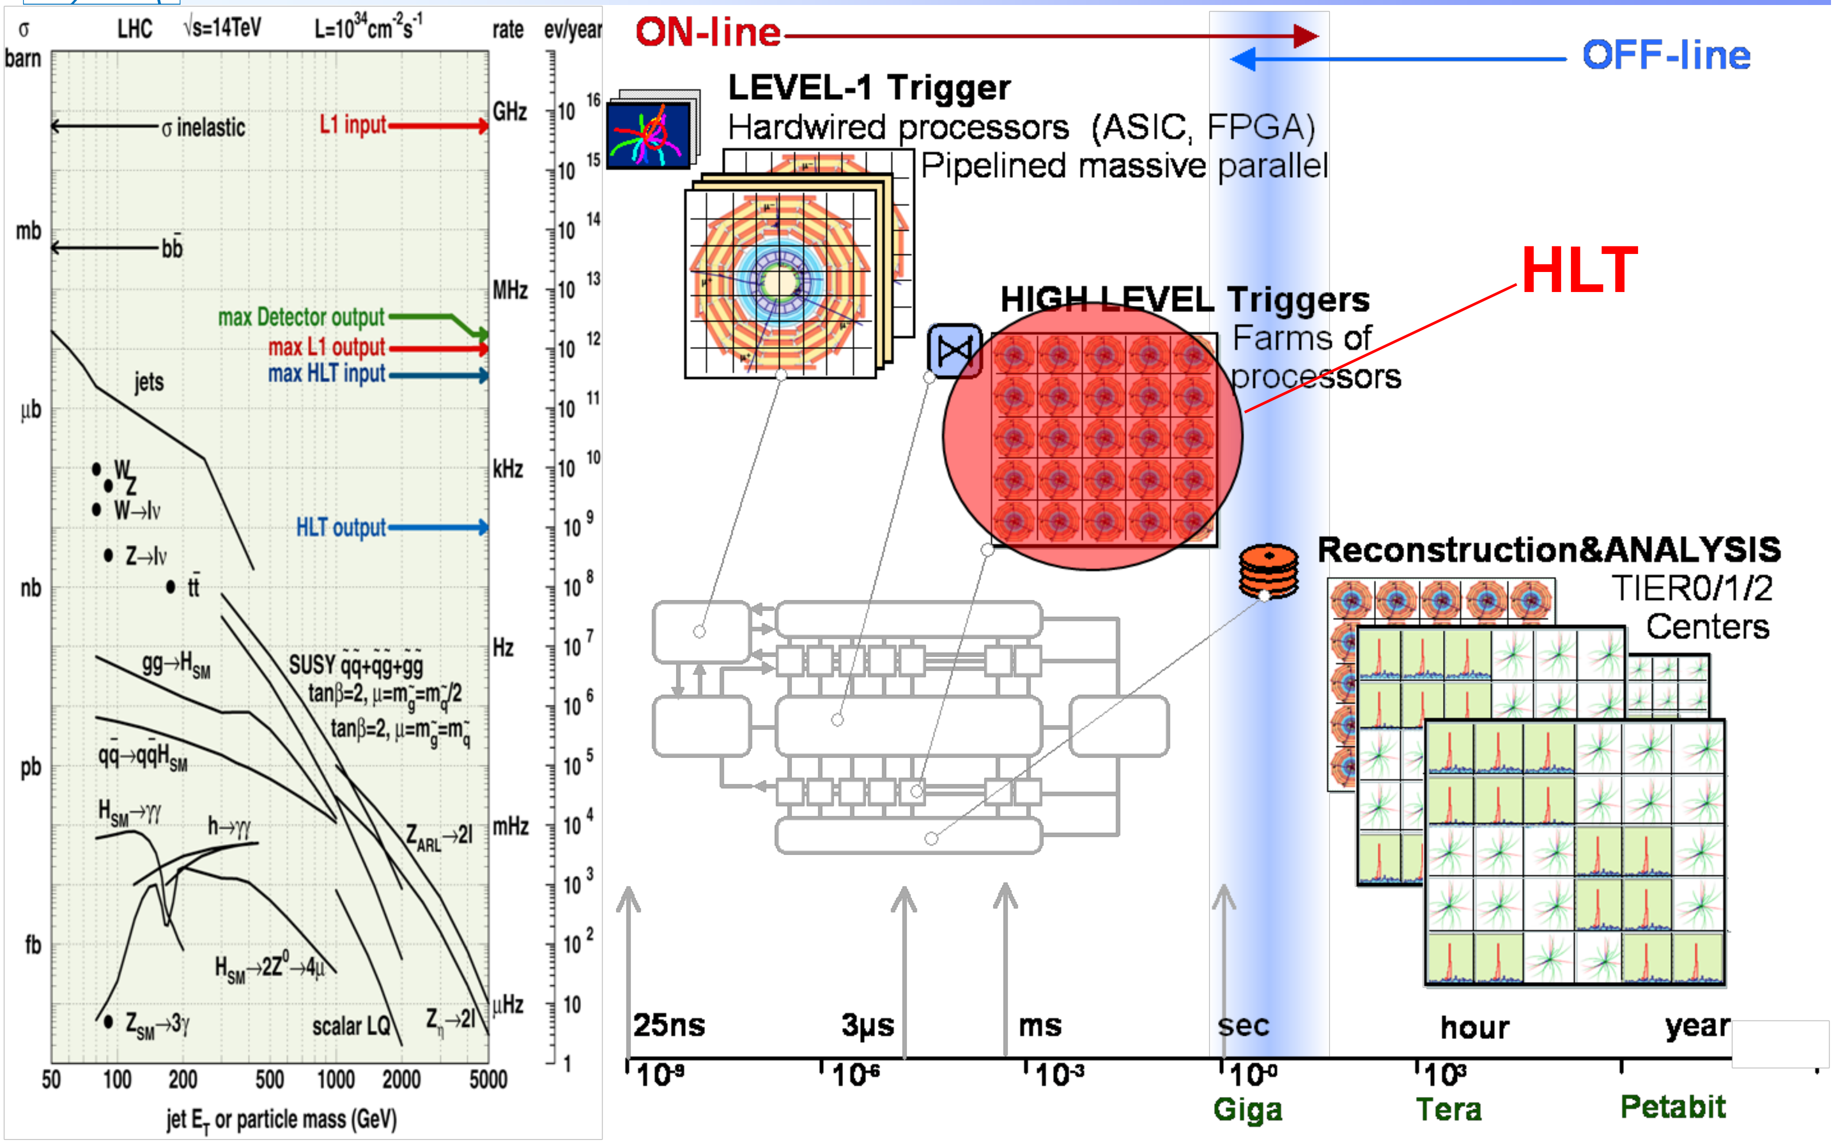

Branches

1. **Throughput of ~32 Gb/s is enough (ALICE)**
   - ALICE needs 2.5 GB/s of "final EVB"
   - Then proceed no further; software, control and monitor, and all issues of very large events (storage very important)
2. **Need more bandwidth, but not much more (e.g. LHCb; event size ~100 kB @ 40 kHz = 4 GB/s = 32 Gb/s)**
   - Implement additional capacity
3. **Need much more than this; CMS+ATLAS need 100 GB/s = 800Gb/s**
   - Two solutions:
     - Decrease rate by using a Level-2 farm (ATLAS)
       - ➔ Thus, two farms: a Level-2 and Level-3 farm
     - Build a system that can do 800 Gb/s (CMS)
       - ➔ Thus, a single farm

100 GB/s case: Level-2/Level-3 vs HLT

- **Level-2 (ATLAS):**
  - Region of Interest (ROI) data are ~1% of total
  - Smaller switching network is needed (not in # of ports but in throughput)
  - But adds:
    - Level-2 farm
    - "ROB" units (have to "build" the ROIs)
    - Lots of control and synchronization
  - Problem of large network → problem of Level-2
- **Combined HLT (CMS):**
  - Needs very high throughput
  - Needs large switching network
  - But it is also:
    - Simpler (in data flow and in operations)
    - More flexible (the entire event is available to the HLT – not just a piece of it)
  - Problem of selection → problem of technology

ATLAS: from demonstrator to full EVB

- **With Regions of Interest:**
  - If the Level-2 delivers a factor 100 rejection, then input to Level-3 is 1-2 kHz.
  - At an event size of 1-2 MB, this needs 1-4 GB/s
    - An ALICE-like case in terms of throughput
    - Dividing this into ~100 receivers implies 10-40 MB/s sustained – certainly doable
  - Elements needed: ROIBuilder, L2PU (processing unit),

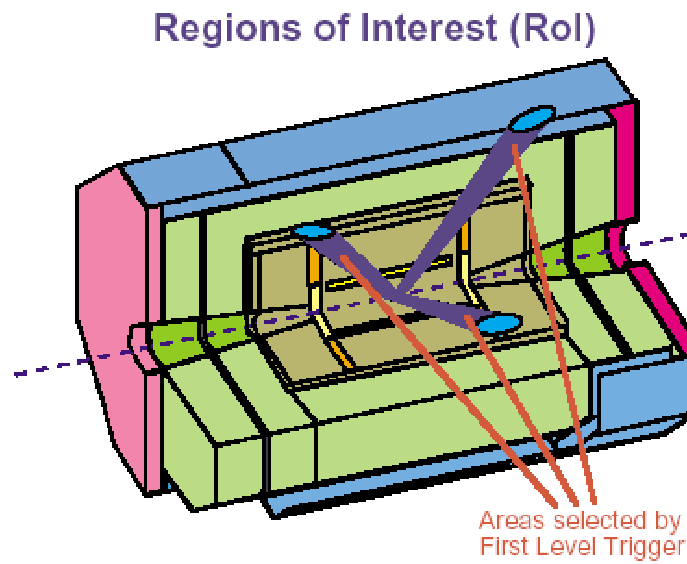

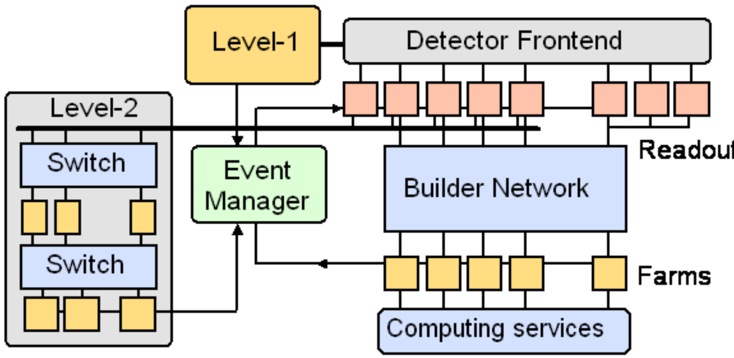

3D-EVB: DAQ staging and scaling

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

Data to surface:

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

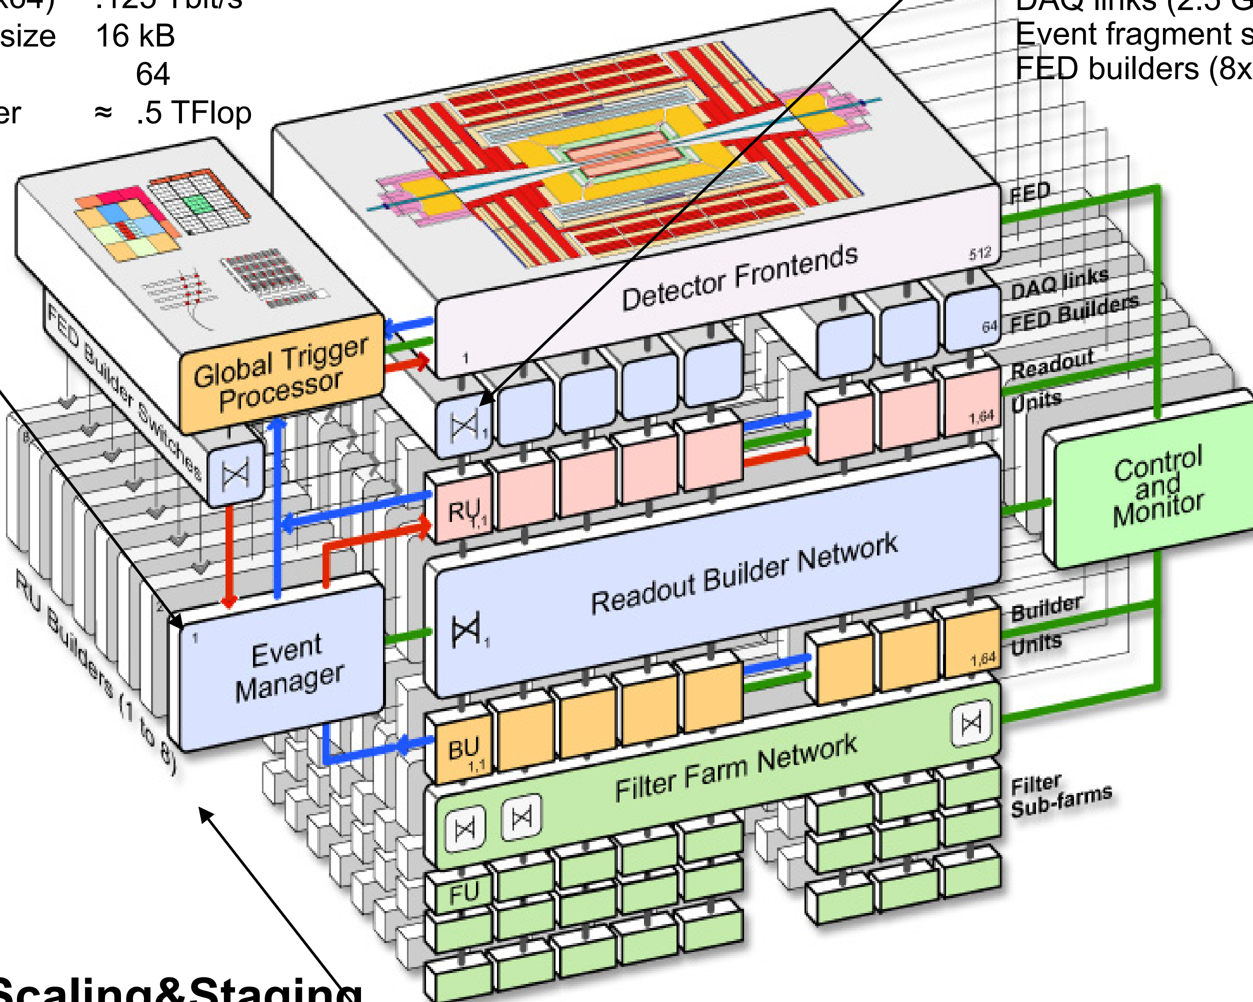

DAQ Scaling&Staging

Processor Farm: the 90's super-computer; the 2000's large computer

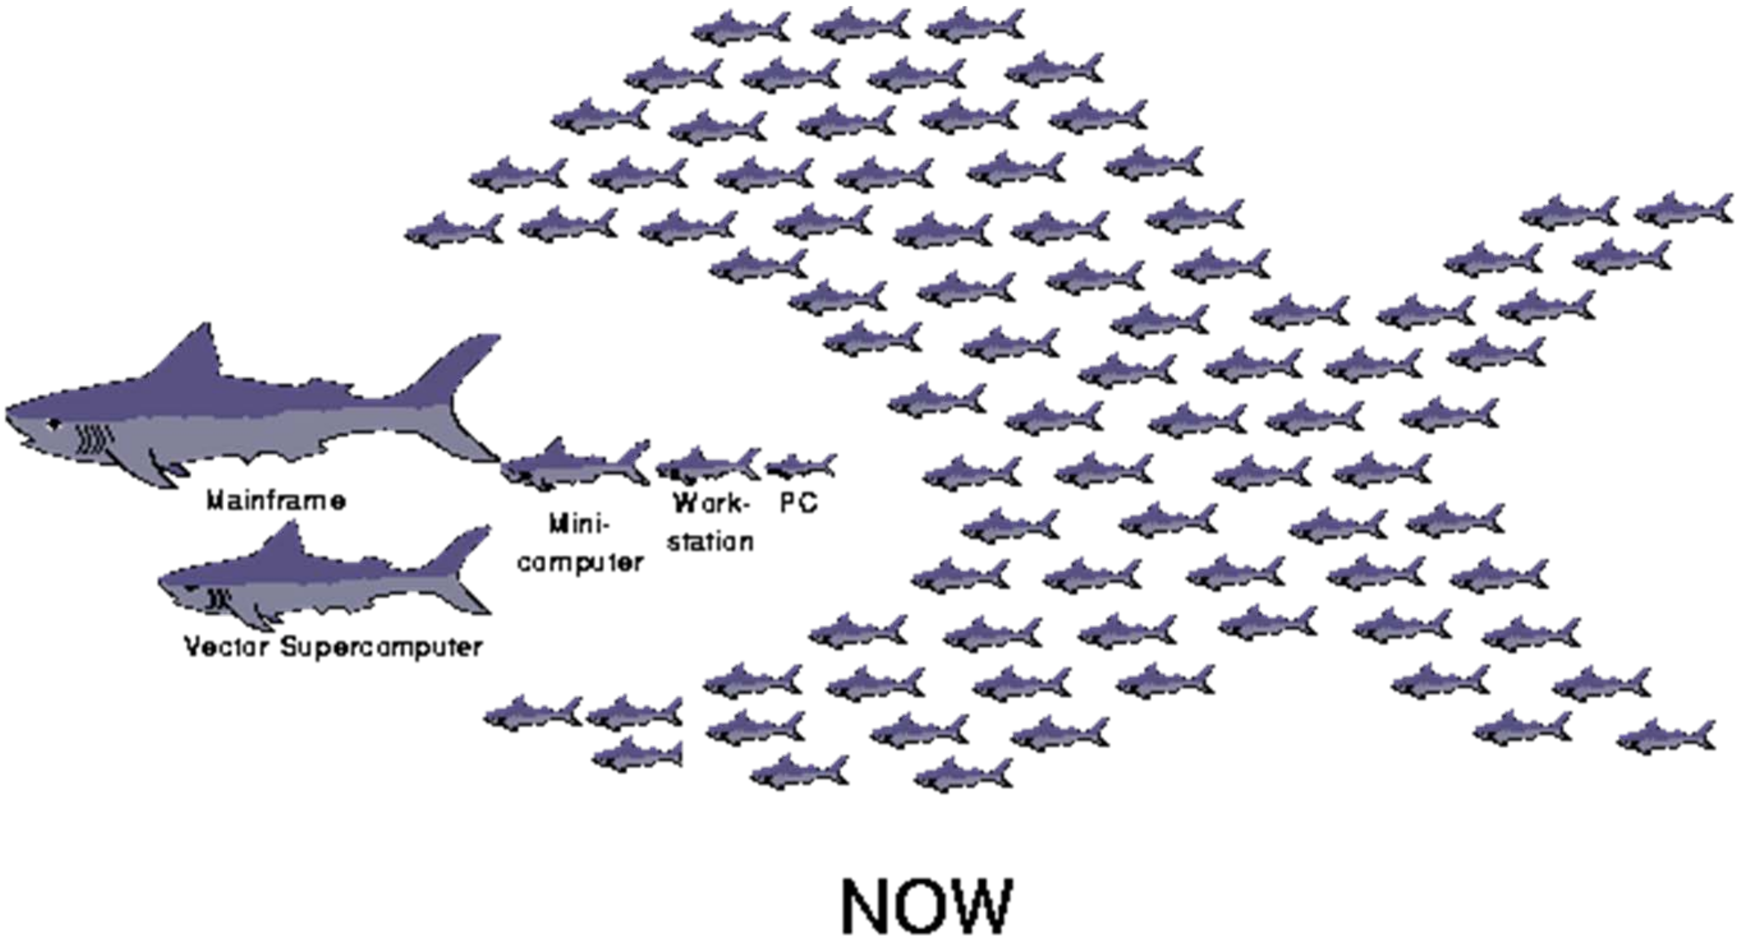

Found at the NOW project (http://now.cs.berkeley.edu)

Processor Engines

- **Final stage of the filtering process: almost an offline-quality reconstruction & selection**
  - **Need real programmable processors; and lots of them**
  - **(Almost) all experiments in HEP: using/will use a processor farm**

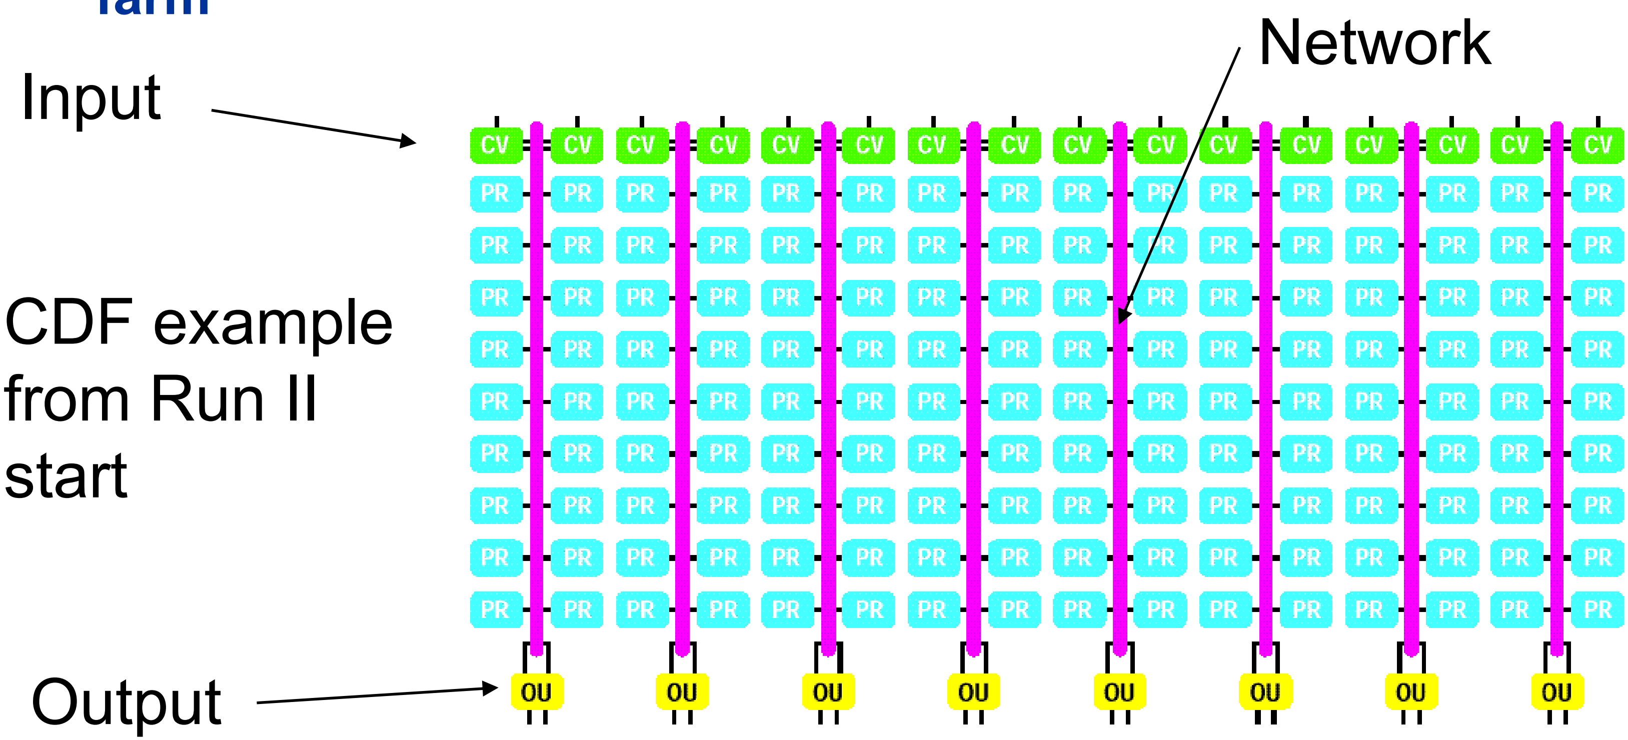

CDF example from Run II start

Processor Engine (II)

- **PC+Linux: the new supercomputer for scientific applications**

obswww.unige.ch/~pfennige/gravitor/gravitor_e.html

www.cs.sandia.gov/cplant/

Event Filter (a processor farm)

- **Explosion of number of farms installed**
  - Very cost-effective
    - Linux is free but also very stable, production-quality
    - Interconnect: Ethernet, Myrinet (if more demanding I/O); both technologies inexpensive and performant
  - Large number of message-passing packages, various API's on the market
    - Use of a standard (VIA?) could be the last remaining tool to be used on this front
  - Despite recent growth, it's a mature process: basic elements (PC, Linux, Network) are all mature technologies. Problem solved. What's left: Control & Monitor.
    - Lots of prototypes and ideas. Need real-life experience.
      - ➔ Problem is human interaction

HLT requirements and operation

- **Strategy/design guidelines**
  - Use offline software as much as possible
    - Ease of maintenance, but also understanding of the detector
- **Boundary conditions:**
  - Code runs in a single processor, which analyzes one event at a time
  - HLT (or Level-3) has access to full event data (full granularity and resolution)
  - Only limitations:
    - CPU time
    - Output selection rate (~10^2 Hz)
    - Precision of calibration constants
- **Main requirements:**
  - Satisfy physics program (see later): high efficiency
  - Selection must be inclusive (to discover the unpredicted as well)
  - Must not require precise knowledge of calibration/run conditions
  - Efficiency must be measurable from data alone
  - All algorithms/processors must be monitored closely

HLT (regional) reconstruction (I)

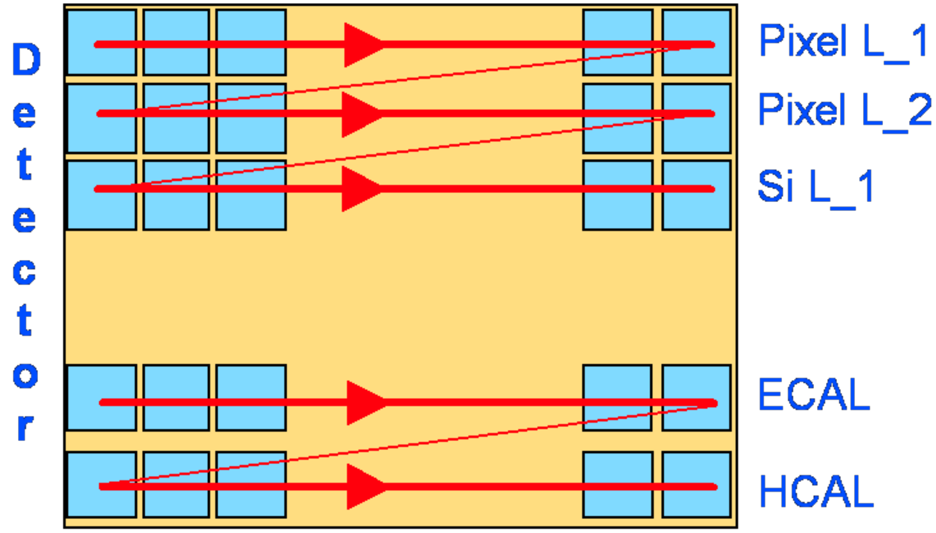

Global
- process (e.g. DIGI to RHITs) each detector fully
- then link detectors
- then make physics objects

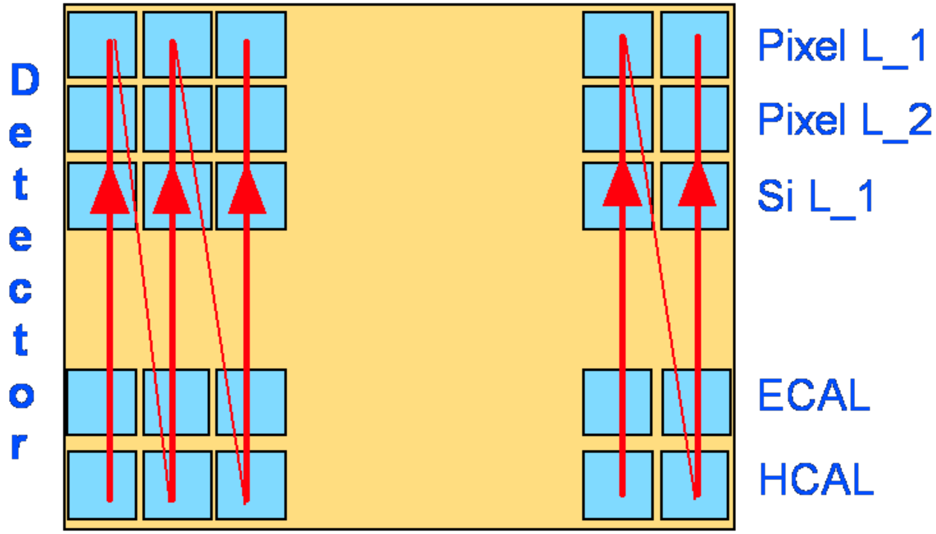

Regional
- process (e.g. DIGI to RHITs) each detector on a "need" basis
- link detectors as one goes along
- physics objects: same

HLT (regional) reconstruction (II)

- **For this to work:**
  - Need to know where to start reconstruction (seed)
- **For this to be useful:**
  - Slices must be narrow
  - Slices must be few
- **Seeds from Lvl-1:**
  - e/γ triggers: ECAL
  - μ triggers: μ sys
  - Jet triggers: E/H-CAL

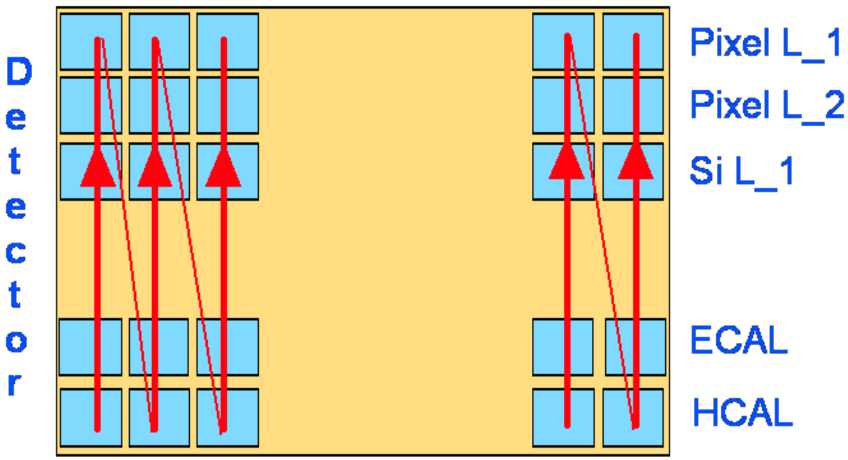

- **Seeds ≈ absent:**
  - Other side of lepton
  - Global tracking
  - Global objects (Sum E_T, Missing E_T)

Example: electron selection (I)

- **“Level-2” electron:**
  - 1-tower margin around 4x4 area found by Lvl-1 trigger
  - Apply “clustering”
  - Accept clusters if H/EM < 0.05
  - Select highest E_T cluster

- **Brem recovery:**
  - Seed cluster with $E_T > E_T^{min}$
  - Road in ϕ around seed
  - Collect all clusters in road
  - → “supercluster”

and add all energy in road:

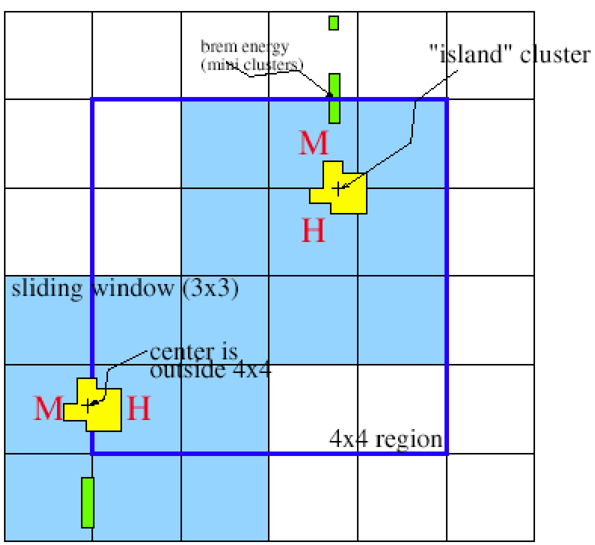

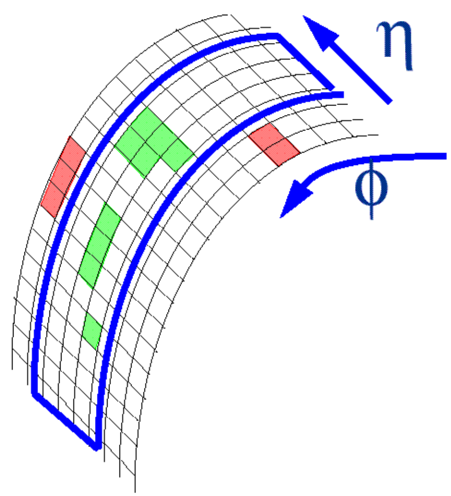

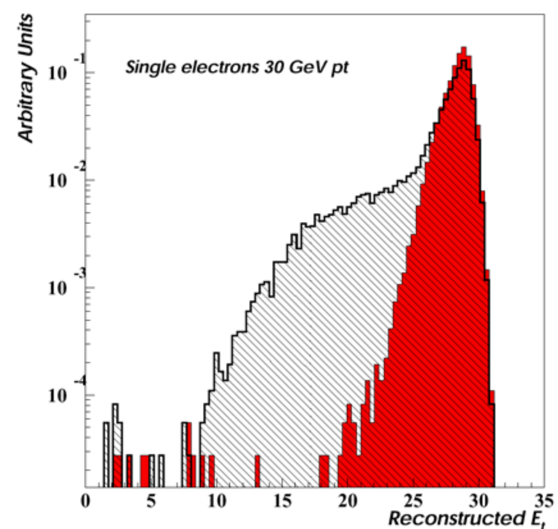

Example: electron selection (II)

- **"Level-2.5" selection: add pixel information**
  - **Very fast, high rejection (e.g. factor 14), high efficiency (ε=95%)**
    - **Pre-bremsstrahlung**
    - **If # of potential hits is 3, then demanding ≥ 2 hits quite efficient**

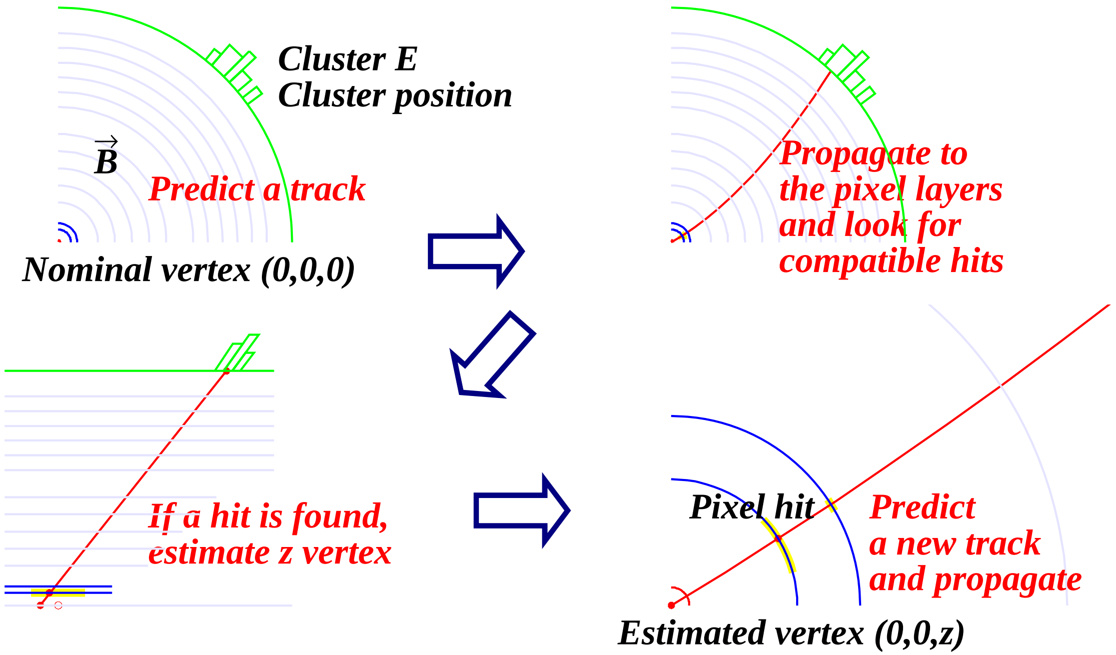

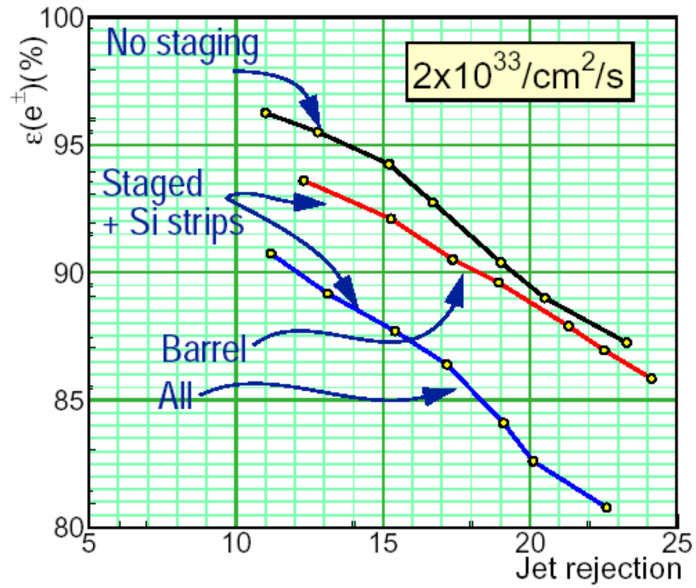

No staging: 3 cylinders + 2 disks
Staged: 2 cylinders + 1 disk

Example: electron selection (III)

- **"Level-3" selection**
  - Full tracking, loose track-finding (to maintain high efficiency):
  - Cut on E/p everywhere, plus
    - Matching in η (barrel)
    - H/E (endcap)
  - Optional handle (used for photons): isolation

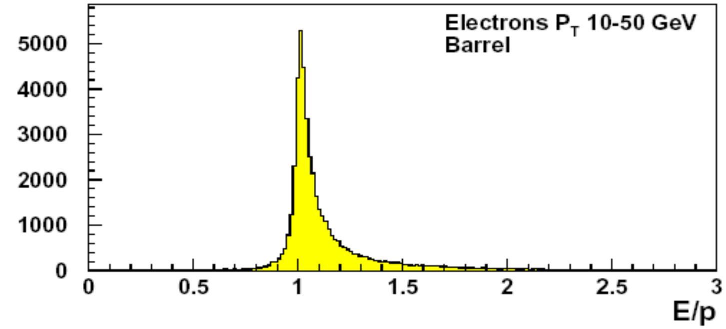

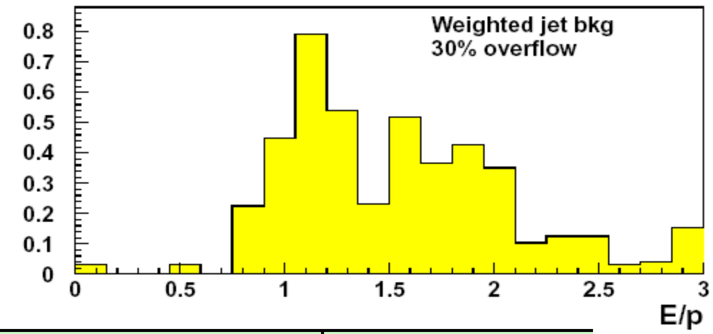

	Signal	Background	Total
Single e	$W \to e\nu$: 10 Hz	$\pi^{\pm}/\pi^0$ overlap: 5 Hz π^0 conversions: 10 Hz b/c $\to$ e: 8 Hz	33 Hz
Double e	$Z \to ee$: 1 Hz	~0	1 Hz
Single γ	2 Hz	3 Hz	5 Hz
Double γ	~0	5 Hz	5 Hz
			44 Hz

Online Physics Selection: summary

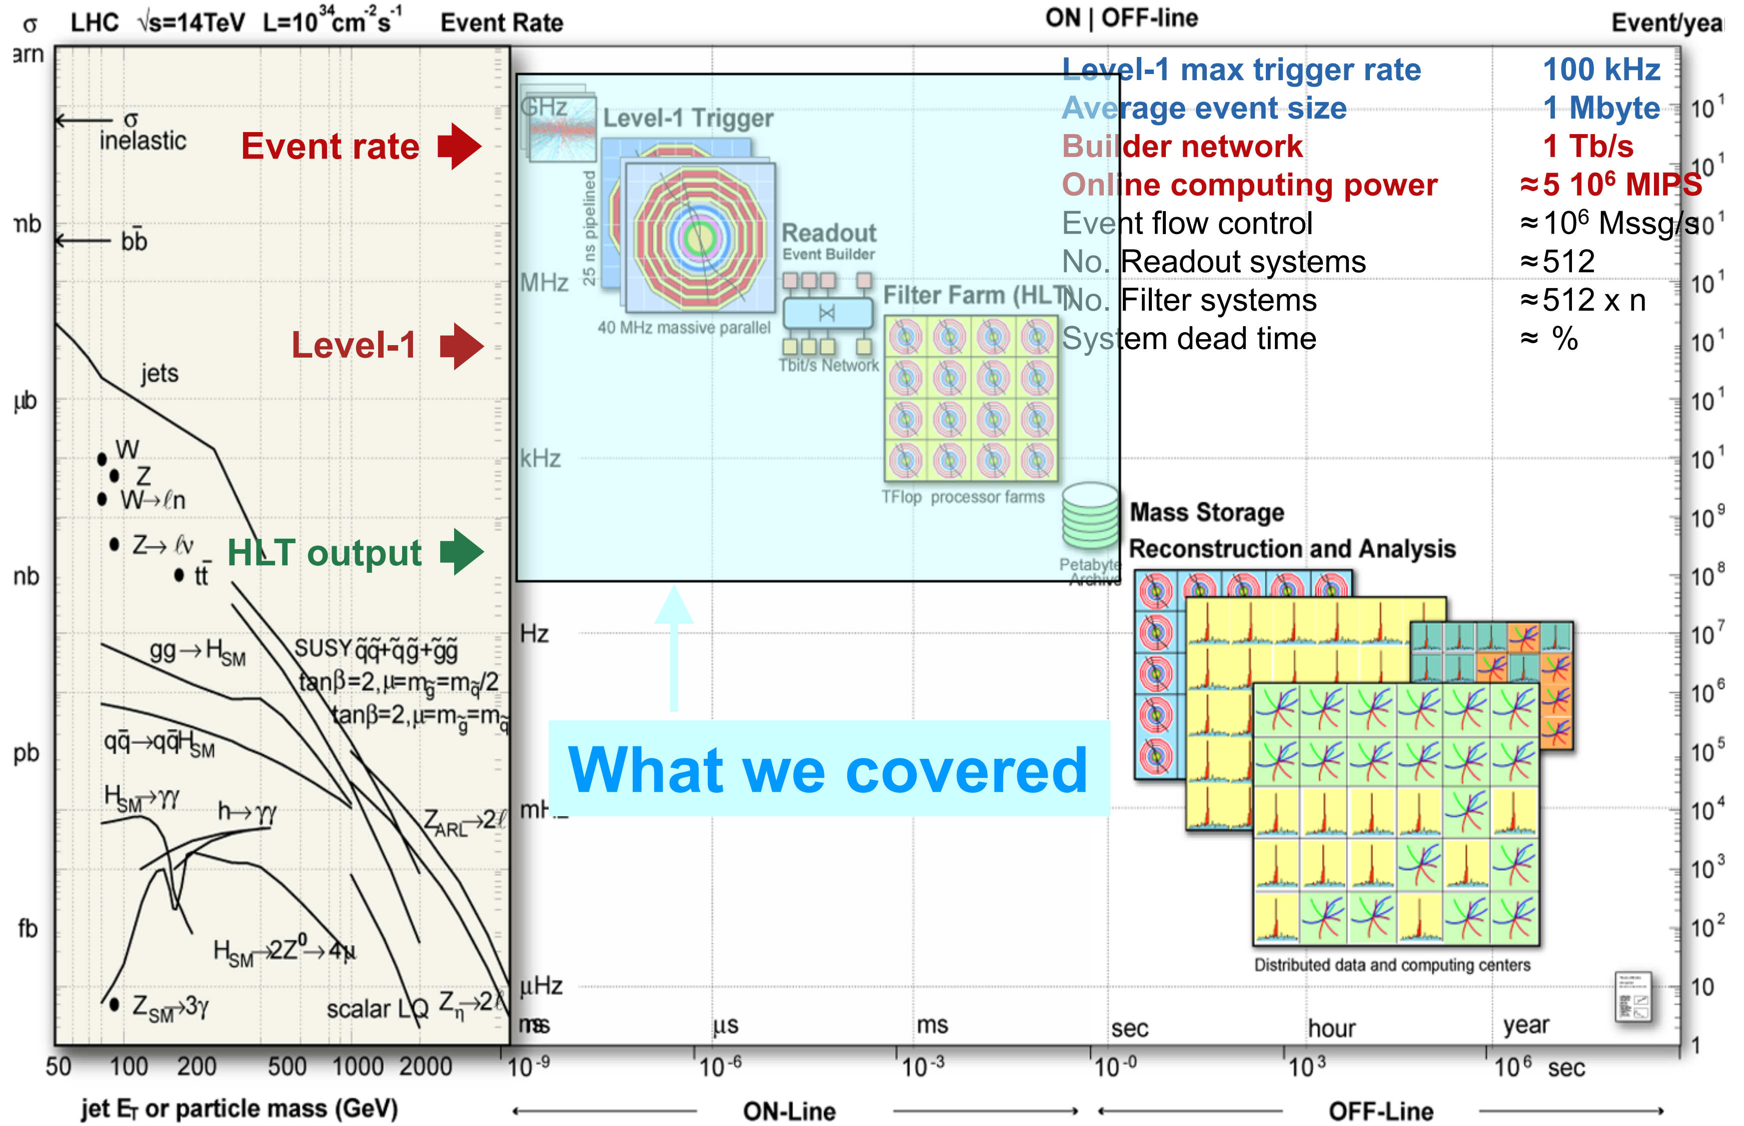

After the Trigger and the DAQ/HLT

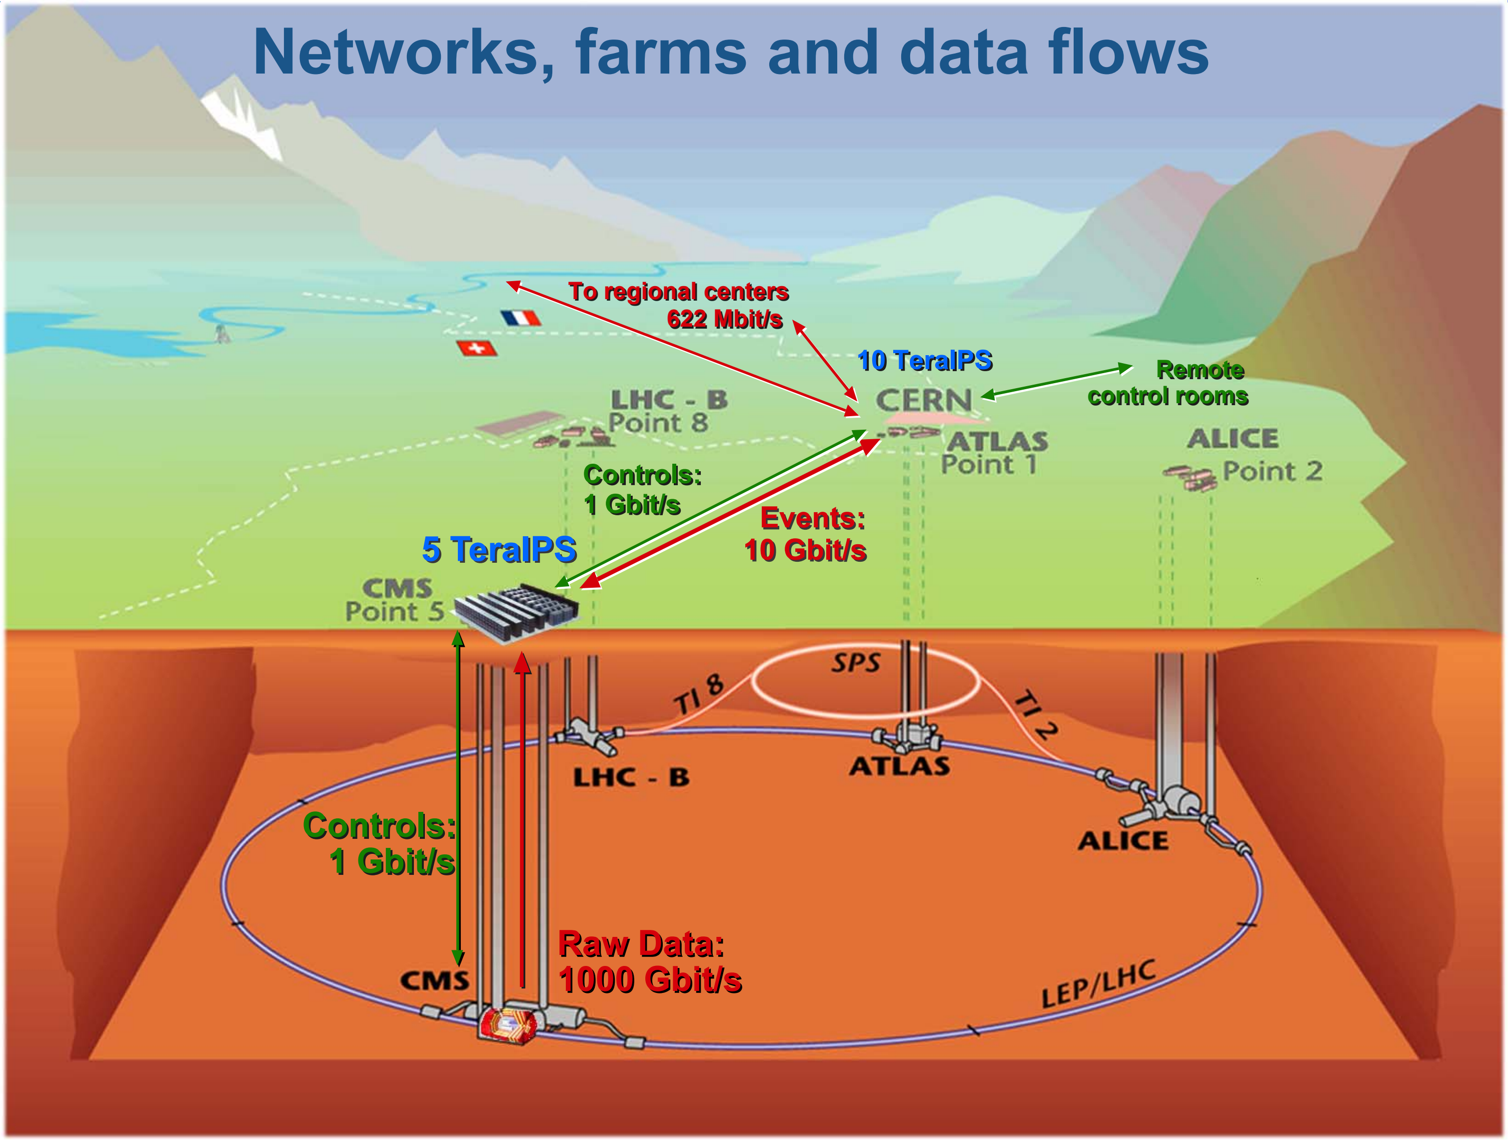

(Grand) Summary

- The Level-1 trigger takes the LHC experiments from the 25 ns timescale to the 10-25 μs timescale
  - Custom hardware, huge fanin/out problem, fast algorithms on coarse-grained, low-resolution data
- Depending on the experiment, the next filter is carried out in one or two (or three) steps
  - Commercial hardware, large networks, Gb/s links.
  - If Level-2 present: low throughput needed (but need Level-2)
  - If no Level-2: three-dimensional composite system
- High-Level trigger: to run software/algorithms that are as close to the offline world as possible
  - Solution is straightforward: large processor farm of PCs
  - Monitoring this is a different issue
- All of this must be understood, for it's done online.

A parting thought

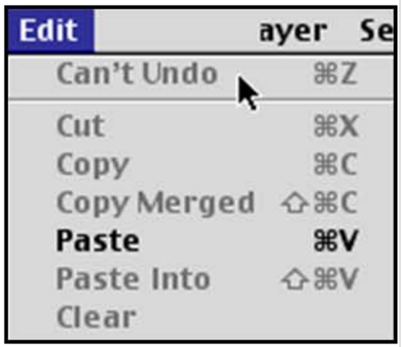

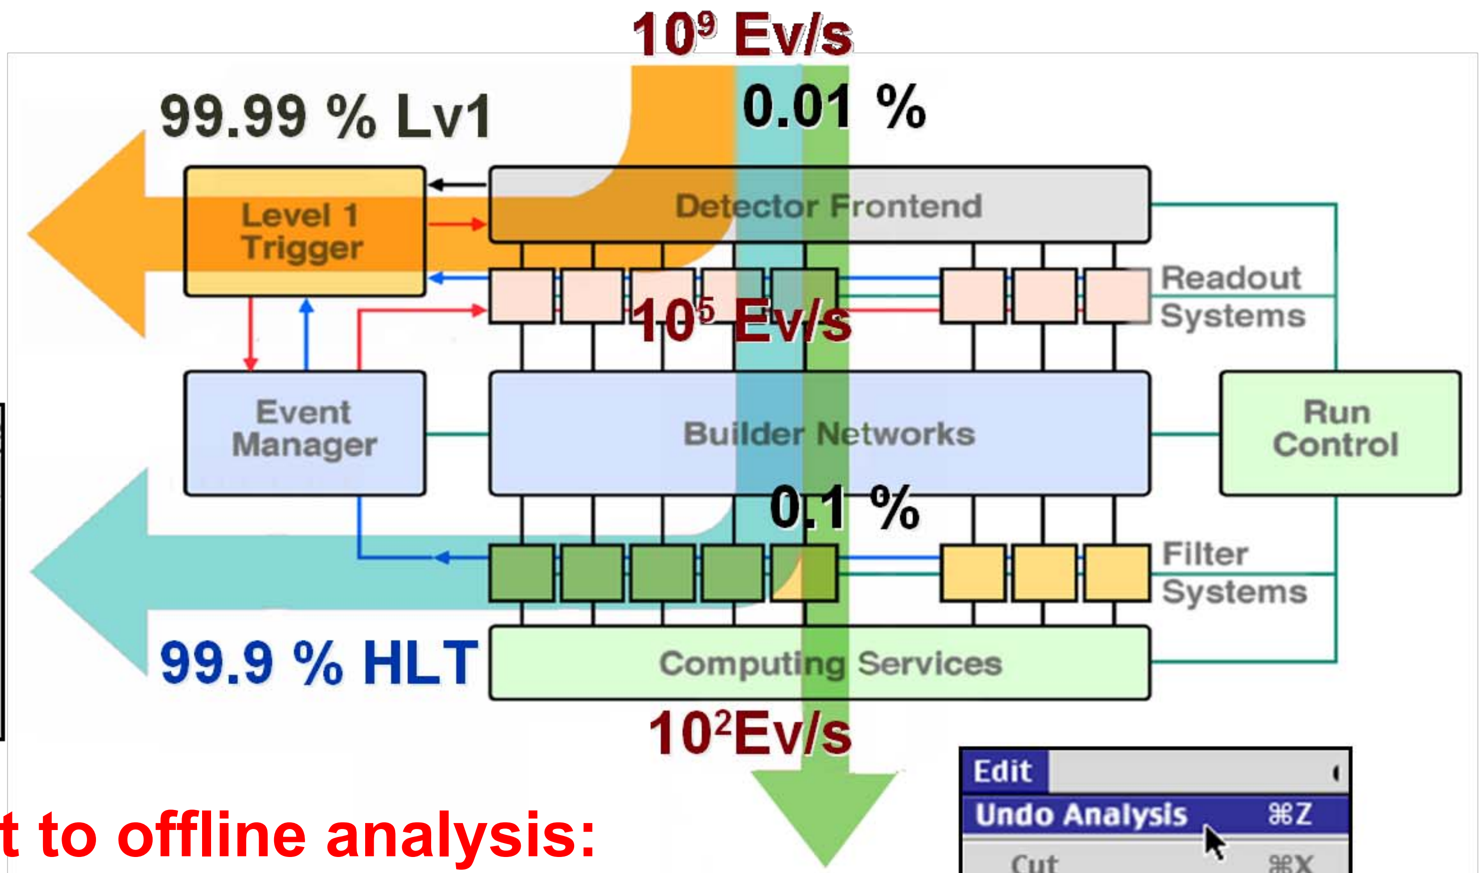

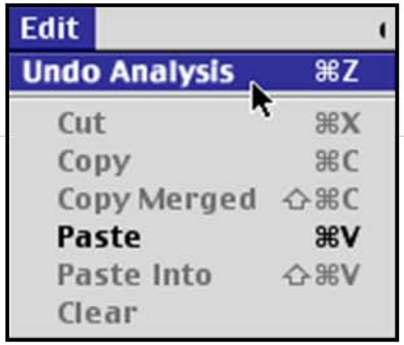

With respect to offline analysis:

Same hardware (Filter Subfarms)

Same software (

But different situations