The CMS High Level Trigger System

Vuko Brigljević
Institut Ruđer Bošković

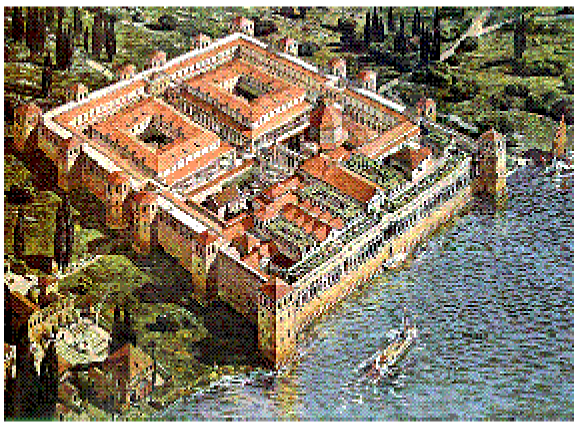

LHC Days in Split
5-9 October 2004

HLT: why should even theorists care?

Plitvice Lakes National Park

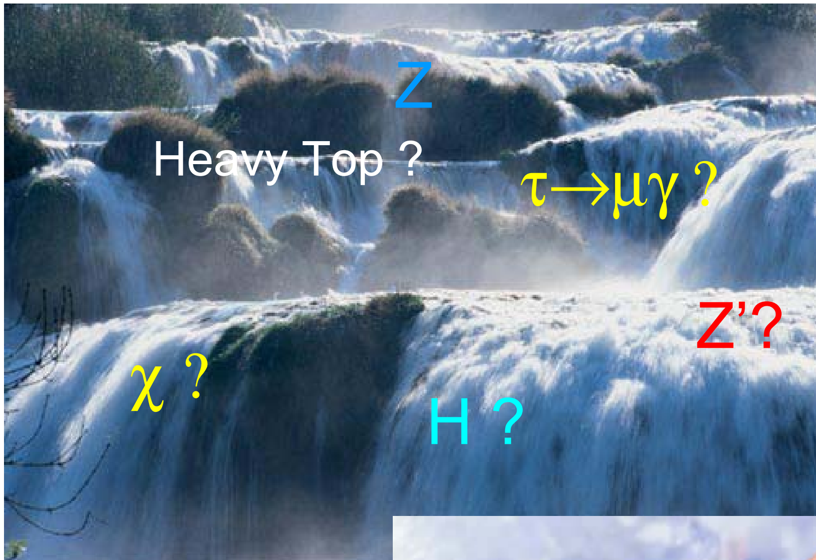

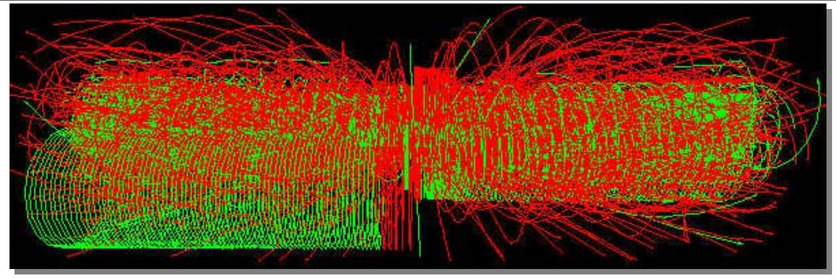

A lot of physics will pour out of pp collisions at the LHC!

may be even your preferred new physics signal; yes, but…

… will it be in the tiny fraction that we will keep?

Physics selectivity at LHC

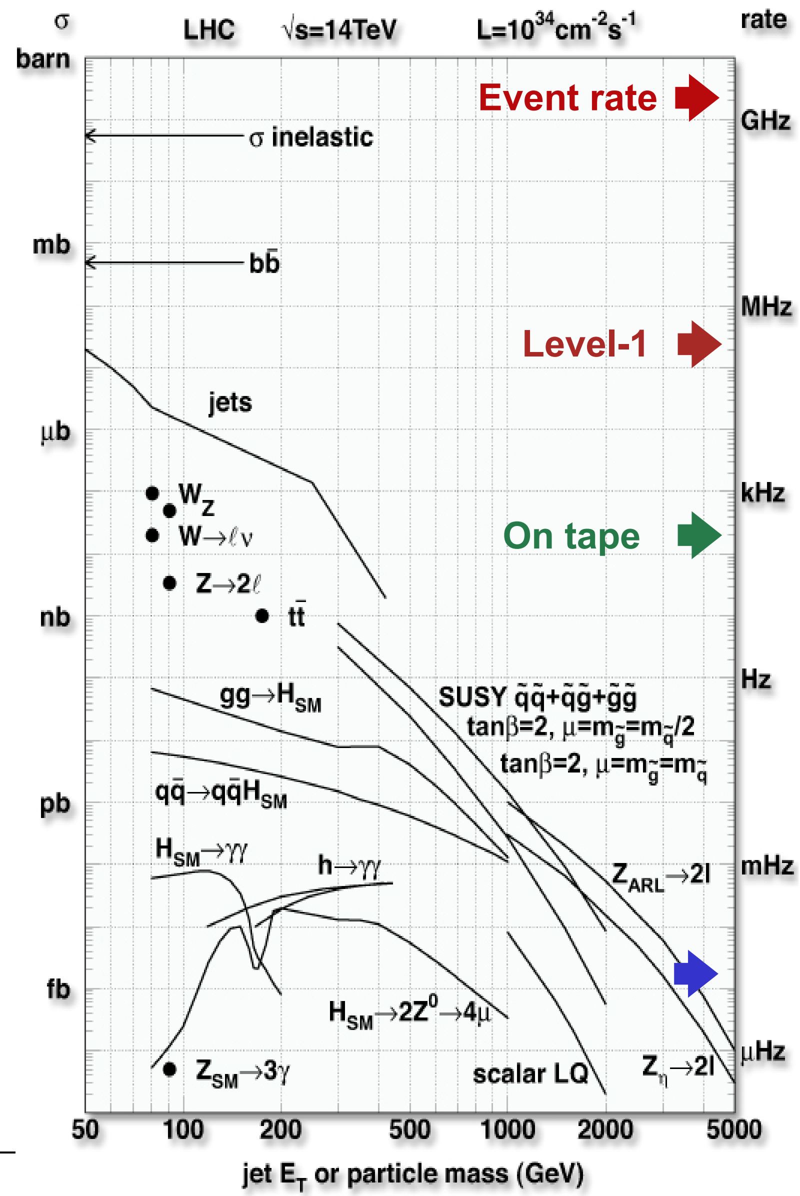

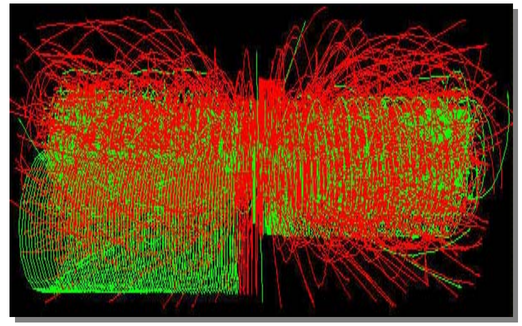

Operating conditions:
Higgs in 4 muons
+ ~20 minimum bias

All charged tracks with pt > 2 GeV

Event Rates:	**~10^9 Hz**
Event size:	**~1 MByte**

Level-1 Output	**100 kHz**
Mass storage	**10^2 Hz**
Event Selection:	**~$1/10^{13}$**

Reconstructed tracks with pt > 25 GeV

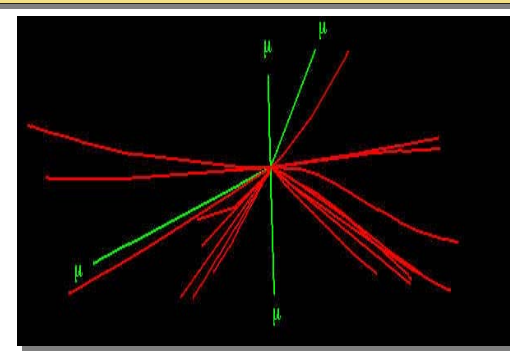

HLT in CMS: the grand picture

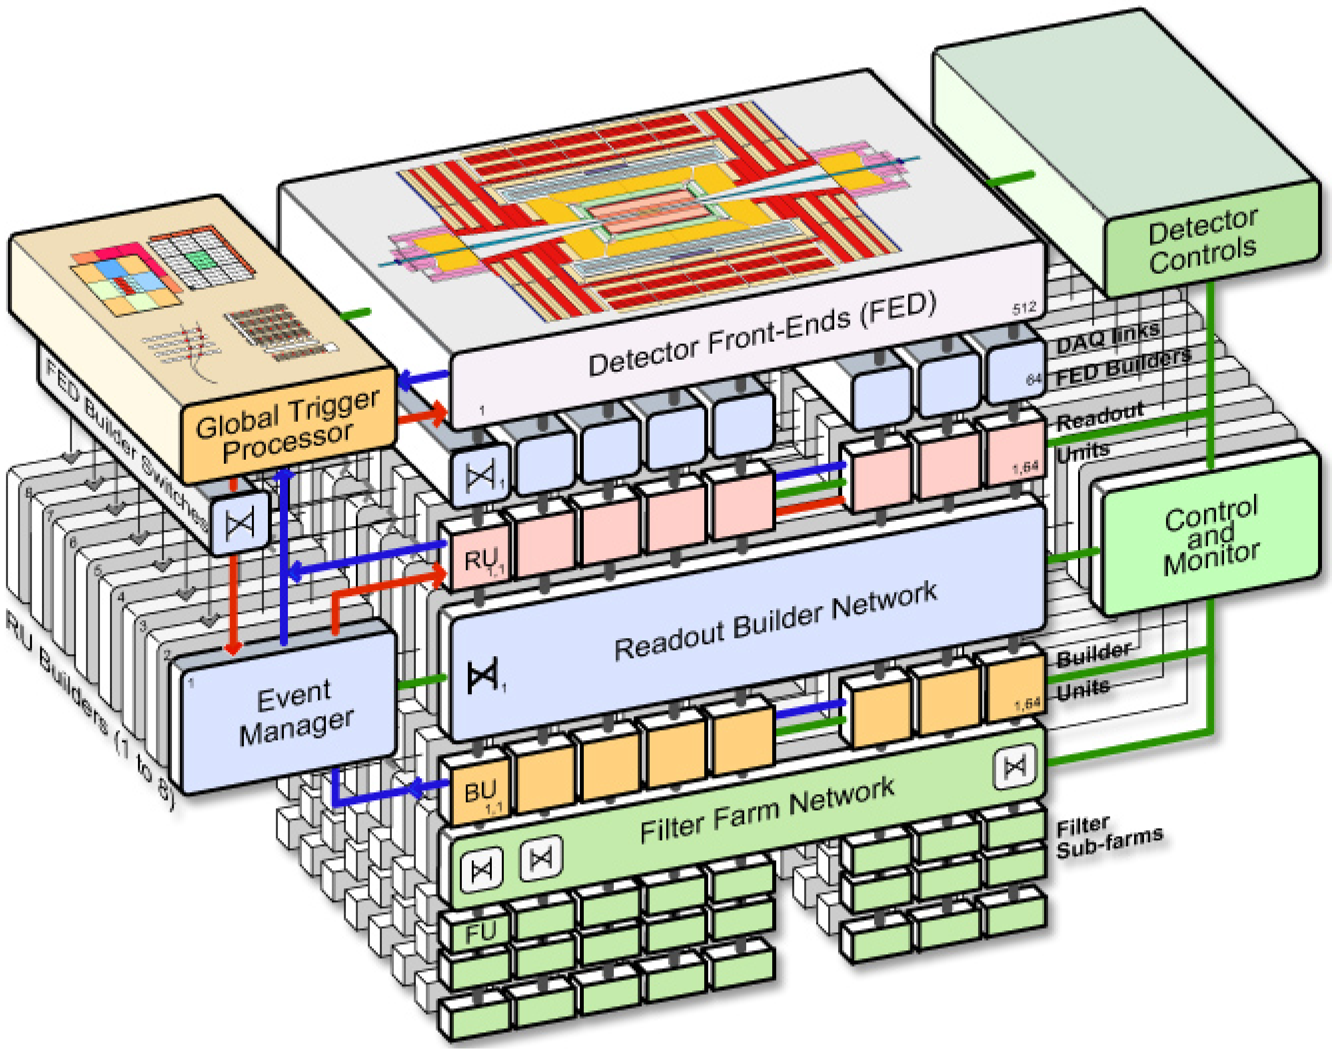

The HLT in the CMS DAQ

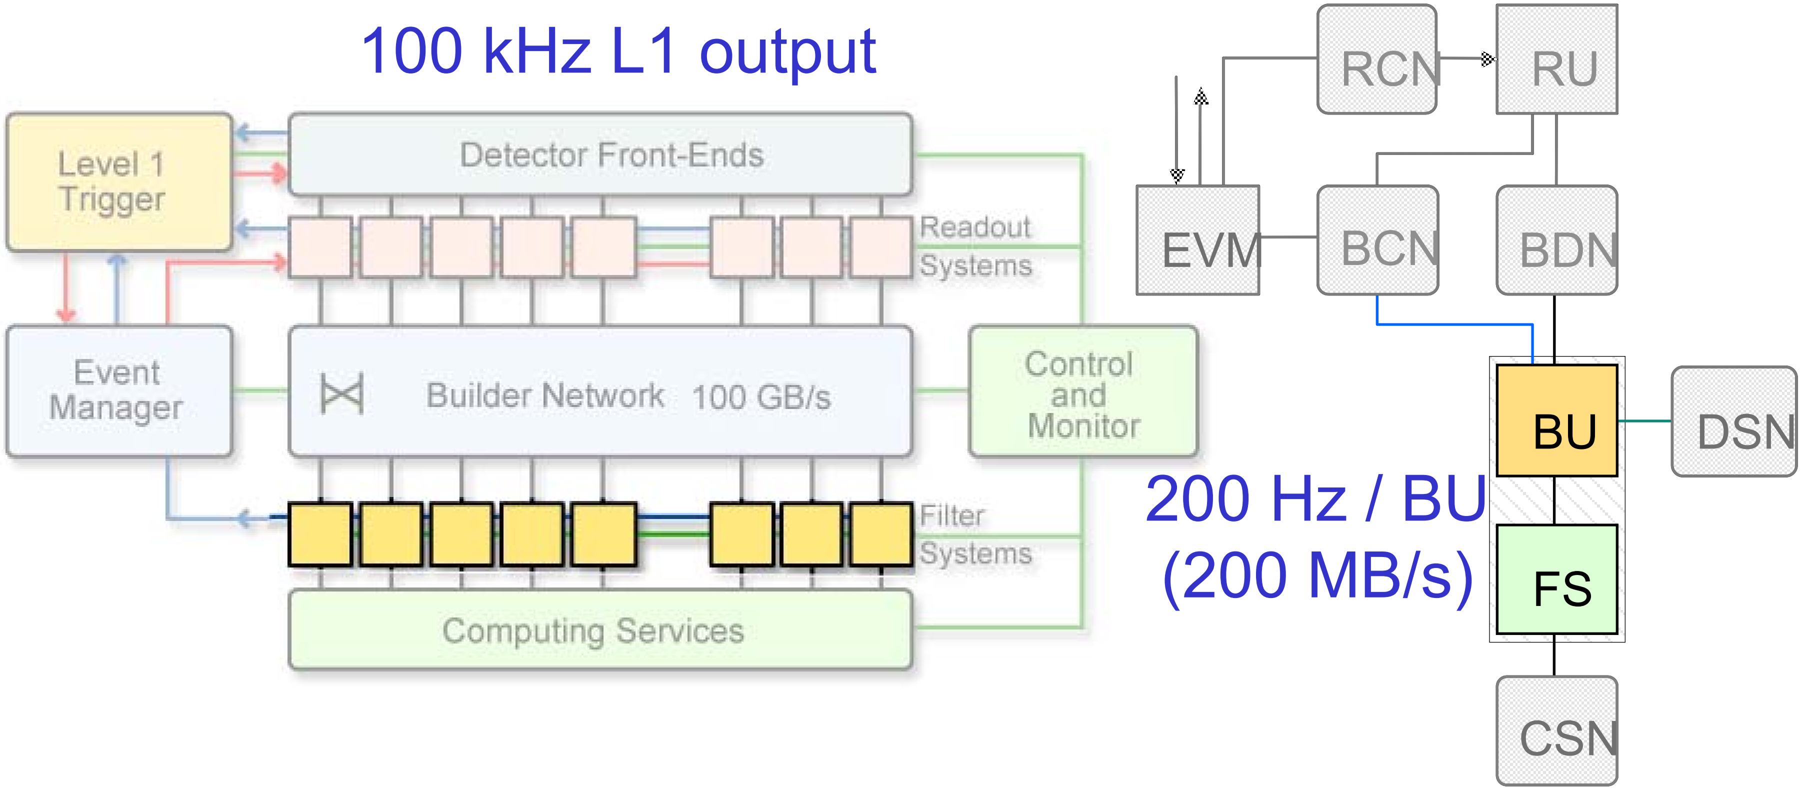

- Builder Unit (BU) connects to switch and distributes fully built events to a collection of Filter Units (FU)
- The FU's run the HLT algorithms and ask for data on a need basis

HLT requirements and operation

- 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
  - Only limitations:
    - CPU time: guarantee deadtimeless operation at nominal L1 output rate
    - Output selection rate ($\sim 10^2$ Hz)
- Main requirements:
  - Satisfy physics program: high efficiency
  - Selection must be inclusive (to discover the unpredicted as well)
  - Allow complete freedom of HLT algorithms
  - Must not require precise knowledge of calibration/run conditions
  - All algorithms/processors must be monitored closely

CPU for the HLT: Filter FARM

- Final stage of the filtering process: almost an offline-quality reconstruction & selection
  - Need real programmable processors; and lots of them
- PC+Linux: the new supercomputer for scientific applications
- CMS full DAQ system:
  - ~ 2’000 dual CPU PCs
  - = 4’000 Filter Units
  - = ~ 40 ms / event

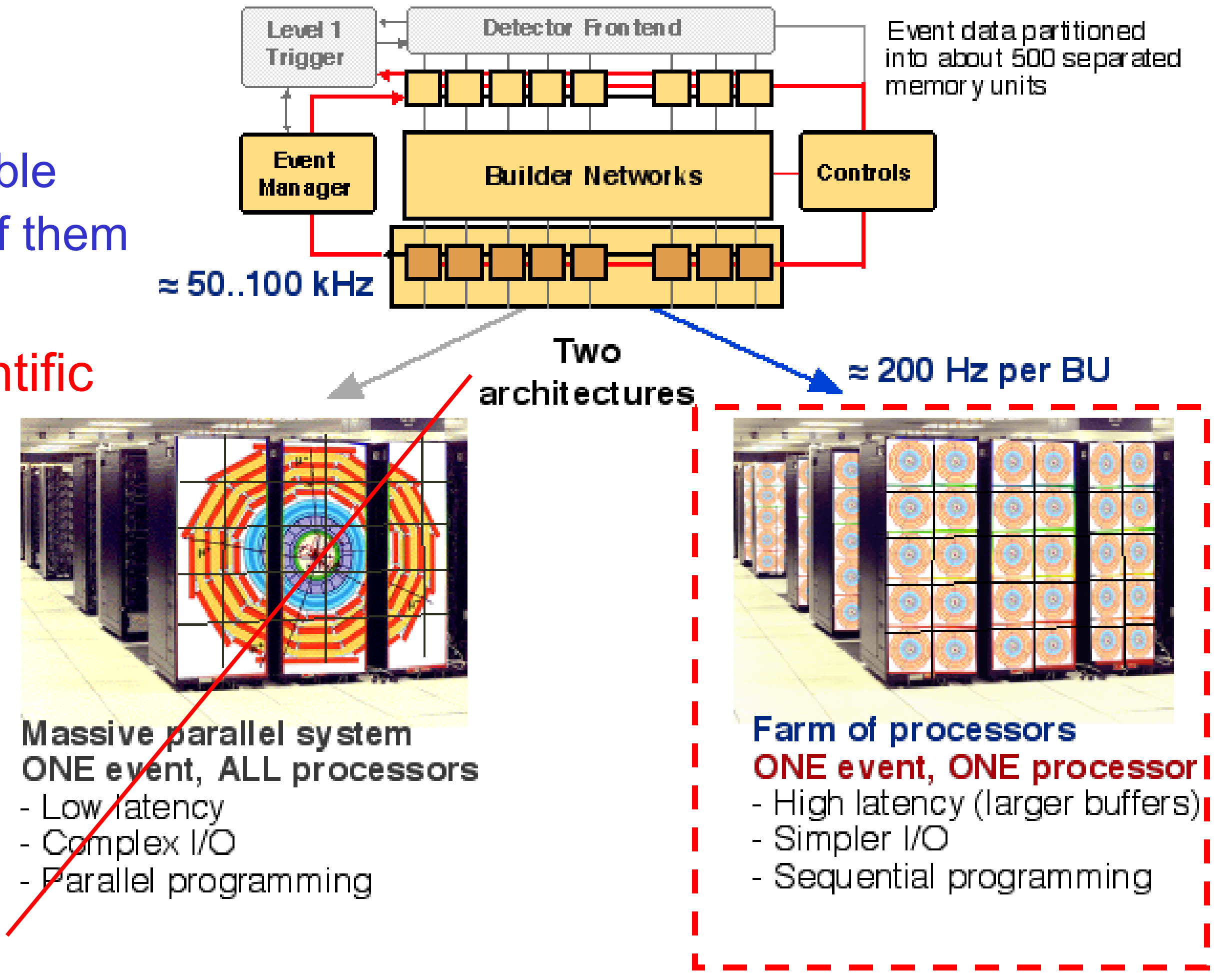

Managing complexity: Divide et impera

Filter Farm divided in subfarms controlled by a Subfarm Manager headnode

- Facilitates installation staging
- Isolates problems
- Allows DAQ subpartitions
- Test of different SW version

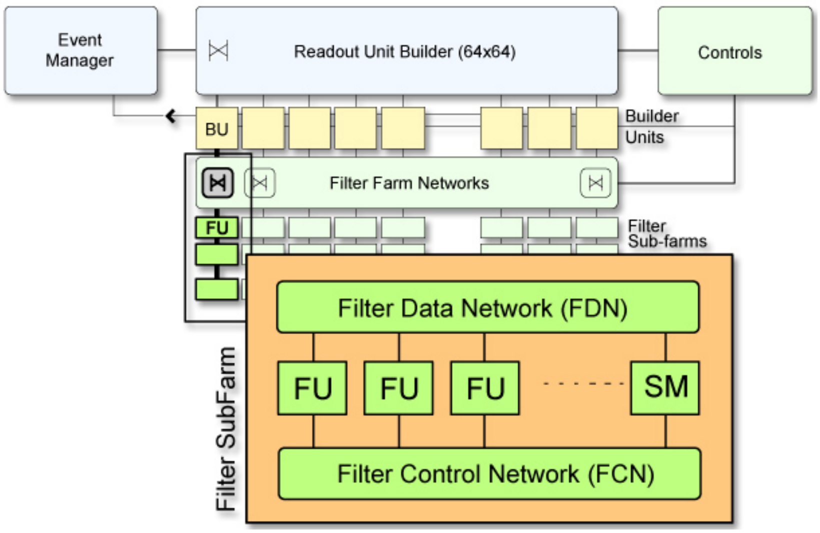

Communication protocols:

- Data (BU-FU): low level TCP
- Control & Monitoring: http, SOAP, XML

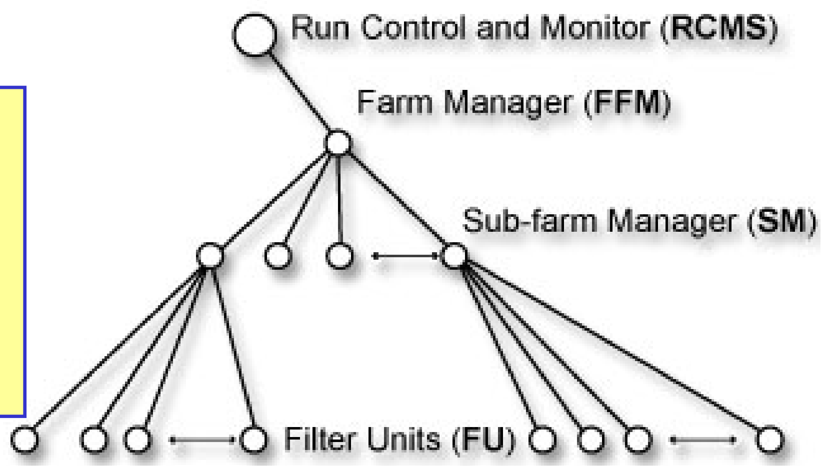

HLT Algorithms

- Strategy/design guidelines
  - Use offline software as much as possible (only specific I/O)
    - Ease of maintenance, but also understanding of the detector
    - Make use of large developer community
    - But tight quality requirements
- Flexibility & freedom to change Trigger table
- Reconstruct ALL and ONLY what is needed to decide quickly:
  - Unpack only needed raw data (also reduces BU output)
  - Regional reconstruction
  - Intelligent steering of algorithm sequence: use L1 input

  All of this is made possible thanks to the

  "Reconstruction on demand"

  Design built in the CMS Reconstruction software

HLT (regional) reconstruction (I)

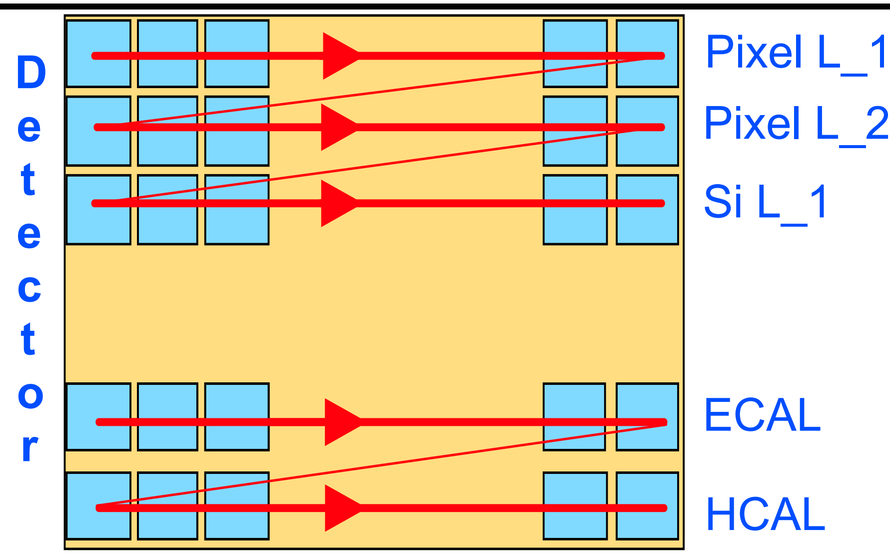

Global

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

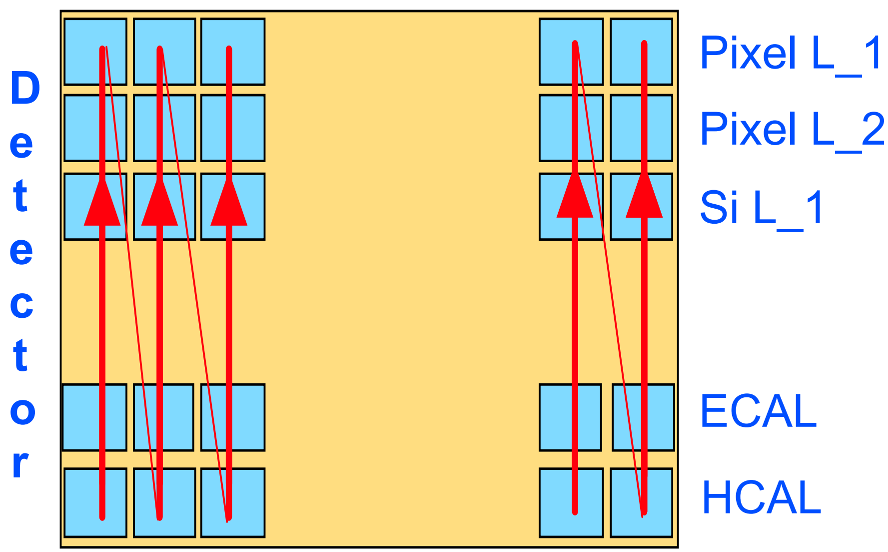

Regional

- 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

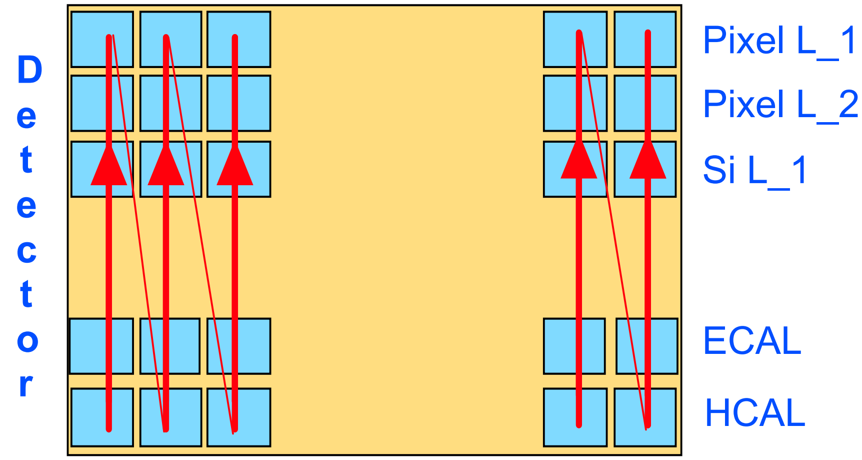

- Seeds from Lvl-1:
  - e/γ triggers: ECAL
  - μ triggers: μ sys
  - Jet triggers: E/H-CAL

- 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:

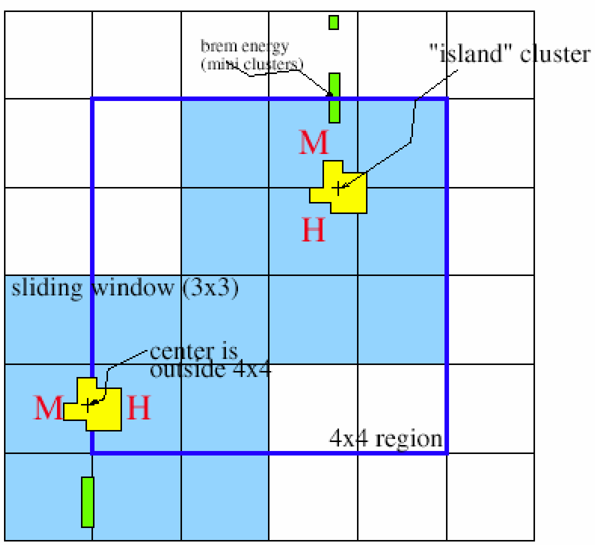

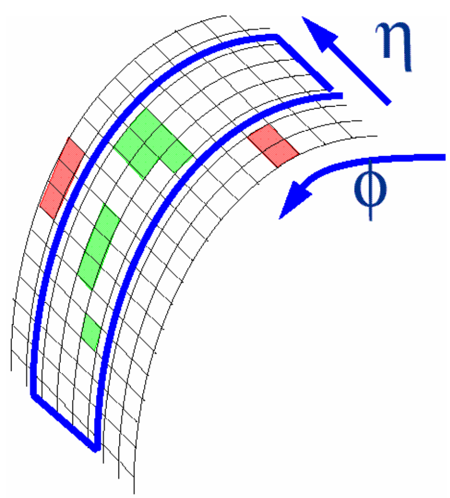

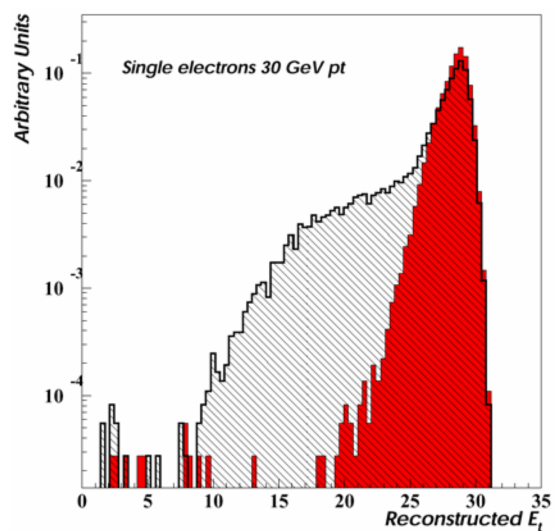

Example: electron selection (II)

- "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

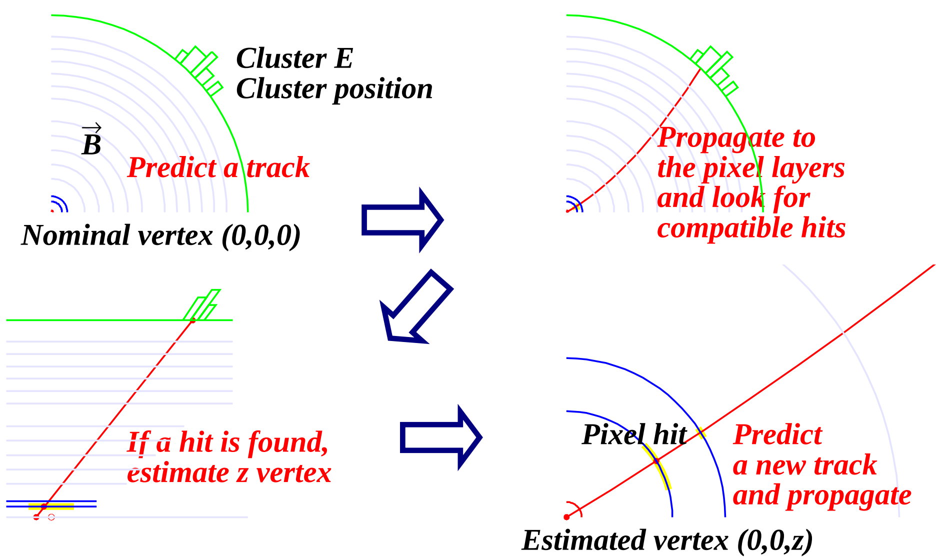

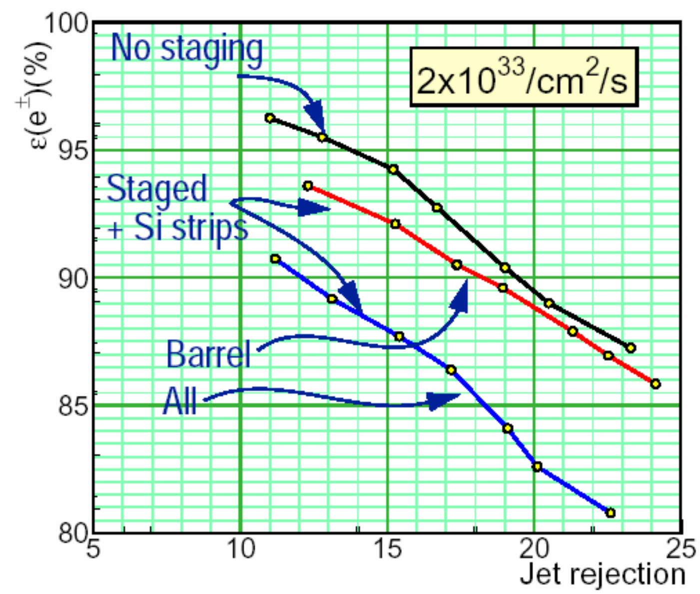

No staging: 3 cylinders + 2 disks
Staged: 2 cylinders + 1 disk

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

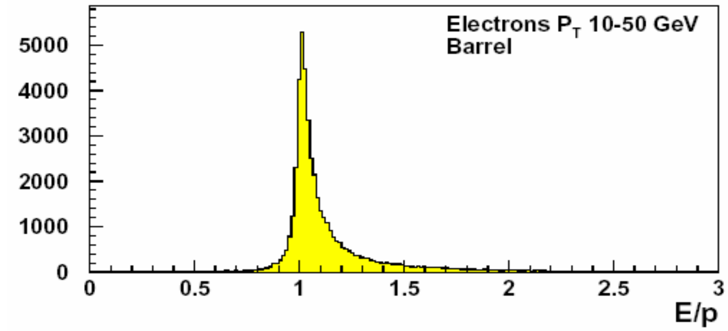

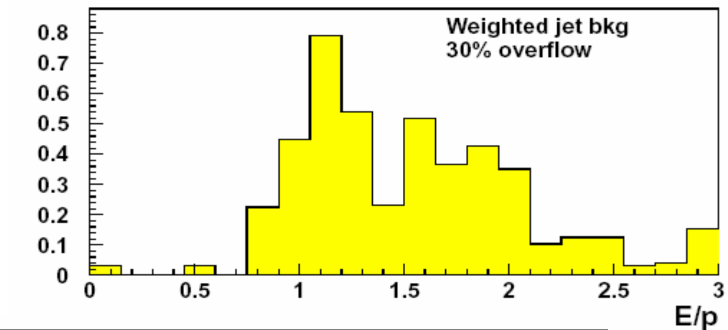

	Signal	Background	Total
Single e	$W \to e\nu$: 10 Hz	$\pi^{\pm}/\pi^{0}$ overlap: 5 Hz π^{0} conversions: 10 Hz b/c $\to$ e: 8 Hz	33 Hz
Double e	$Z \to ee$: 1 Hz	~0	1 Hz
Single γ	2 Hz	3 Hz	5 Hz
Double γ	~0	5 Hz	5 Hz
			44 Hz

HLT Steering

$$\overbrace{E_t(e) > 23} \vee \overbrace{E_t(\mu) > 18} \vee \overbrace{(E_t(e) > 14 \wedge E_t(\tau) > 48)}$$

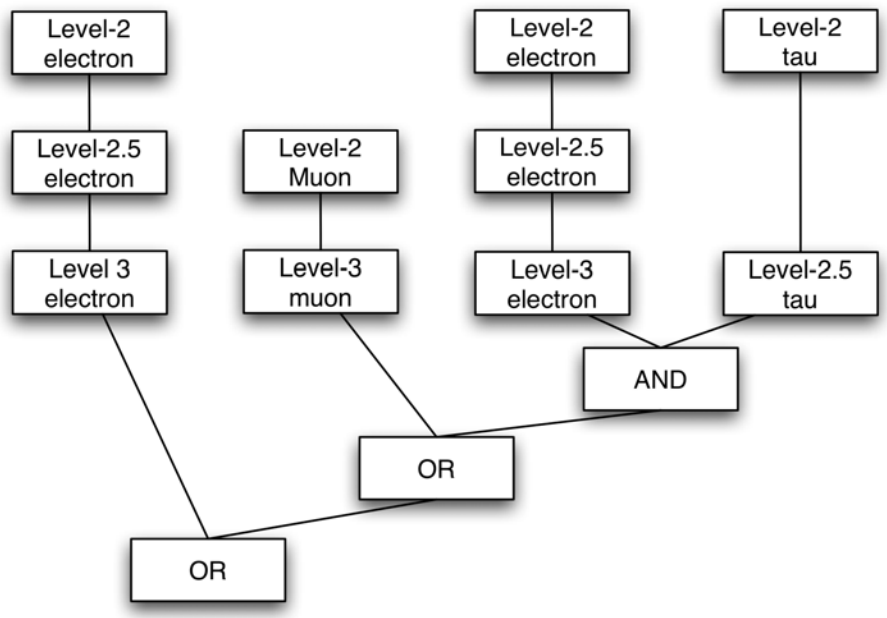

HLT table can be dynamically loaded / modified during running (XML Document)

- HLT Trigger table is equivalent to a logical decision tree
- Evaluation sequence optimized to minimize computation time
- Allow Veto mode: HL subtriggers computed only if corresponding L1 accept on
- Mean rejection time dominates the computation time

Physics Plan and Trigger Table
(as of DAQ TDR)

Trigger table determination (I)

- Startup configuration: don't need 100 kHz on day 1
  - Machine conditions non-optimal
  - Funds for completion of DAQ will be present later
  - Exploit technological developments – buy ALAP
- Startup setup:
  - Physics startup assumptions: $2 \times 10^{33} cm^{-2} s^{-1}$, and a DAQ with 4 RU builders, i.e. 50 kHz throughput
- Starting point: 50kHz/3 $\rightarrow$16kHz to allocate
  - Factor 3 is safety: accounts for all processes that have not been simulated, uncertainties in generator/simulation and beam conditions
    - This factor varies across experiments
  - Initial step: equal allocation across (1&2e/γ), (1&2μ), (1&2τ) and jets/cross channels (e&τ, μ*jet, etc)
  - Get thresholds, efficiencies; look at physics cost; iterate

Trigger table determination (II)

- Deciding Lvl-1 cuts: 1e/γ vs 2e /γ, 1μ vs 2μ, 1τ vs 2τ
  - Create iso-rate plot (contours of "equal cost")
  - For each contour (in relevant range, e.g. 2kHz, 3kHz, 4kHz) get efficiency of physics channel in 1-obj vs 2-obj requirement

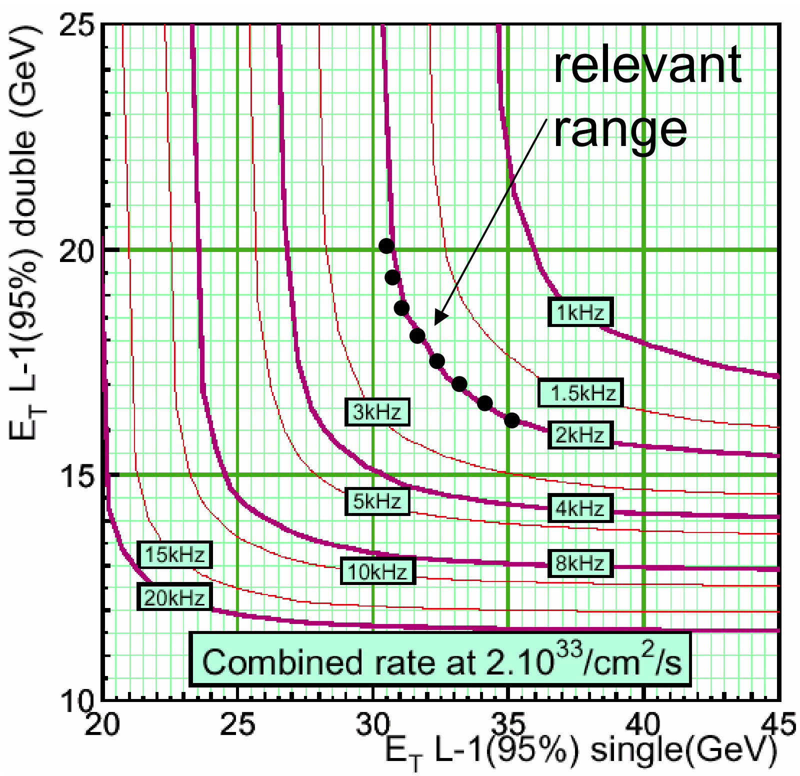

(and of course: operate at point of rapid slope change)

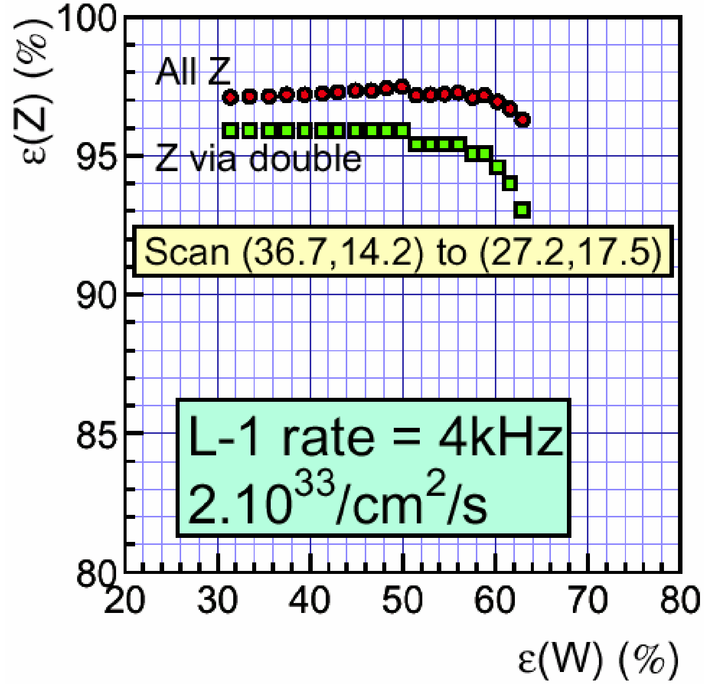

Level-1 trigger table (low lumi)

- Total Rate: 50 kHz. Factor 3 safety, allocate 16 kHz

Trigger	Threshold (ε=90-95%) (GeV)	Indiv. Rate (kHz)	Cumul rate (kHz)
1e/γ, 2e/γ	29, 17	4.3	4.3
1μ, 2μ	14, 3	3.6	7.9
1τ, 2τ	86, 59	3.2	10.9
1-jet	177	1.0	11.4
3-jets, 4-jets	86, 70	2.0	12.5
Jet * Miss-E_T	88 * 46	2.3	14.3
e * jet	21 * 45	0.8	15.1
Min-bias		0.9	16.0

HLT table (low luminosity)

- Total Rate: 105 Hz

Trigger	Threshold (ε=90-95%) (GeV)	Indiv. Rate (Hz)	Cumul rate (Hz)
1e, 2e	29, 17	34	34
1γ, 2γ	80, (40*25)	9	43
1μ, 2μ	19, 7	29	72
1τ, 2τ	86, 59	4	76
Jet * Miss-E_T	180 * 123	5	81
1-jet, 3-jet, 4-jet	657, 247, 113	9	89
e * jet	19 * 52	1	90
Inclusive b-jets	237	5	95
Calibration/other		10	105

HLT table

- Issues to "fight"
  - Purity of streams is not the same (e.g. electrons vs muons)
  - Overlap (kinematically) is necessary; but also: redundancy
    - Question most asked in large analysis meetings, when a problem is under investigation in W->eν: do we see this in the muons?
  - But, above all, comparison of unlike things:
    - How much more bandwith should go to lower-P_T muons than to electrons?
    - How should one share the bandwidth between jet*missE_T and di-electrons?
  - Only guidance in the end of the day is efficiency to all the known channels
    - While keeping the selection INCLUSIVE
    - For this is online. Events rejected are lost forever.

HLT performance

- With previous selection cuts

Channel	Efficiency (for fiducial objects)
H(115 GeV)$\rightarrow\gamma\gamma$	77%
H(160 GeV)$\rightarrow$WW* $\rightarrow 2\mu$	92%
H$\rightarrow$ZZ$\rightarrow 4\mu$	92%
A/H(200 GeV)$\rightarrow 2\tau$	45%
SUSY (~0.5 TeV sparticles)	~60%
With R_P-violation	~20%
W$\rightarrow e\nu$	67% (fid: 60%)
W$\rightarrow\mu\nu$	69% (fid: 50%)
Top$\rightarrow\mu$ X	72%

HLT: CPU usage

- All numbers for a 1 GHz, Intel Pentium-III CPU

Trigger	CPU (ms)	Rate (kHz)	Total (s)
1e/γ, 2e/γ	160	4.3	688
1μ, 2μ	710	3.6	2556
1τ, 2τ	130	3.0	390
Jets, Jet * Miss-E_T	50	3.4	170
e * jet	165	0.8	132
B-jets	300	0.5	150

- Total: 4092 s for 15.1 kHz $\rightarrow$ 271 ms/event
  - Therefore, a 100 kHz system requires 1.2×10^6 SI95
- Expect improvements, additions. Time completely dominated by muons (GEANE extrapolation) – will improve
- This is "current best estimate", with ~50% uncertainty.

CPU Usage

- Today: need ~300 ms on a 1GHz Pentium-III CPU
  - For 50 kHz, need 15,000 CPUs
  - Moore's Law: 2x2x2 times less time (fewer CPUs) in 2007
    - Central estimate: 40 ms in 2007, i.e. 2,000 CPUs
    - Thus, basic estimate of 1,000 dual-CPU boxes in TDR
    - (Note: not an excess of CPU, e.g. no raw-data handling)
  - Start-up system of 50kHz (Level-1) and 105 Hz (HLT) can satisfy basic "discovery menu"
    - Some Standard Model physics left out; intend to do it, at lower luminosity and pre-scales as luminosity drops through fill
      - Examples: inclusion of B physics (can be done with high efficiency and low CPU cost; limitation is Level-1 bandwidth); details in TDR. Also low-mass dijet resonances.
- Single-farm design works.

FAQ

- What happens if we turn on and we only need 42kHz (i.e. safety factor is <3)?
  - We lower thresholds, add triggers,etc to use full bandwidth available
- What happens if we turn on and we need 70 kHz?
  - The Level-1 trigger is programmable, it can, e.g. mask hot regions, etc etc. Requirement is to stay within 50 kHz.
    - Must look carefully at beam-gas etc
- Can we add triggers?
  - All tables: just indications of type of combinations and requirements we can have on "day-1". (Actually at a lumi of $2x10^{33}cm^{-2}s^{-1}$)
    - Much will depend on the Tevatron, on when we turn on, on actual beam conditions, on actual event size, on actual DAQ system…

Summary

- CMS HLT implemented on a farm of PCs
  - Farm design scales with CPU needs
  - Running offline quality selection code
  - As for DAQ, we have a working design, the specific implementation will follow needs & technology
- HLT framework allows flexible and efficient algorithm implementation
- DAQ TDR shows alpha version HLT trigger table
  - Certainly not the final thing, will be moving target anyway
  - Will follow input from HERA, Tevatron, theory,…

My question to the offline community:

Why not more than 100 Hz?

A parting thought

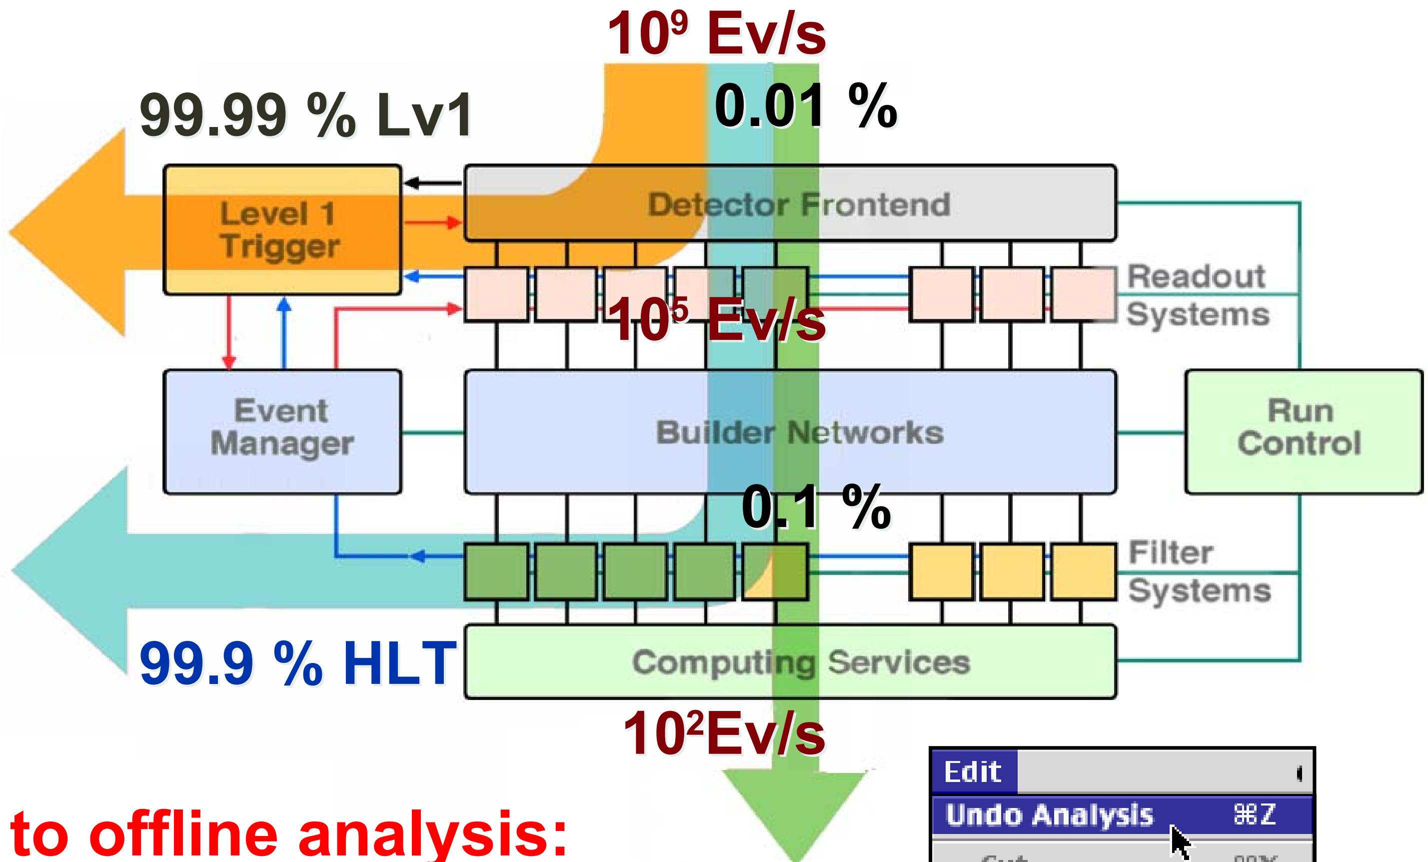

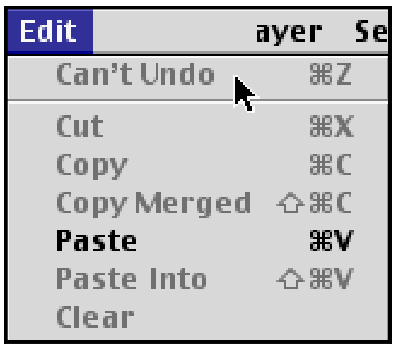

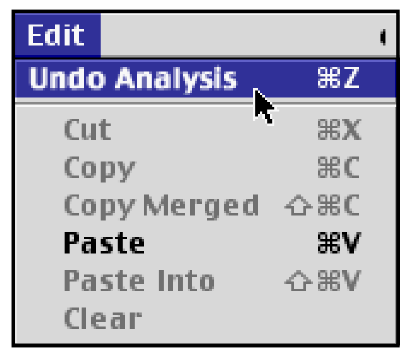

With respect to offline analysis:

Same hardware (Filter Subfarms)

Same software

But different situations

So make sure it ends up in there!