Triggering at LHC

The Challenge &
the CMS Strategy

Christos Leonidopoulos

CERN-PH

Particle Physics and Photon Science
Joint Instrumentation Seminar
7 May 2010
DESY, Hamburg University, XFEL

News from CERN

After 20 years of R&D, Detector Building, Commissioning and Preparation

- LHC is up & running
- The Experiments:
  - ➢ Are in physics commissioning mode
  - ➢ Trigger and detector performance according to design
- First Collision Results

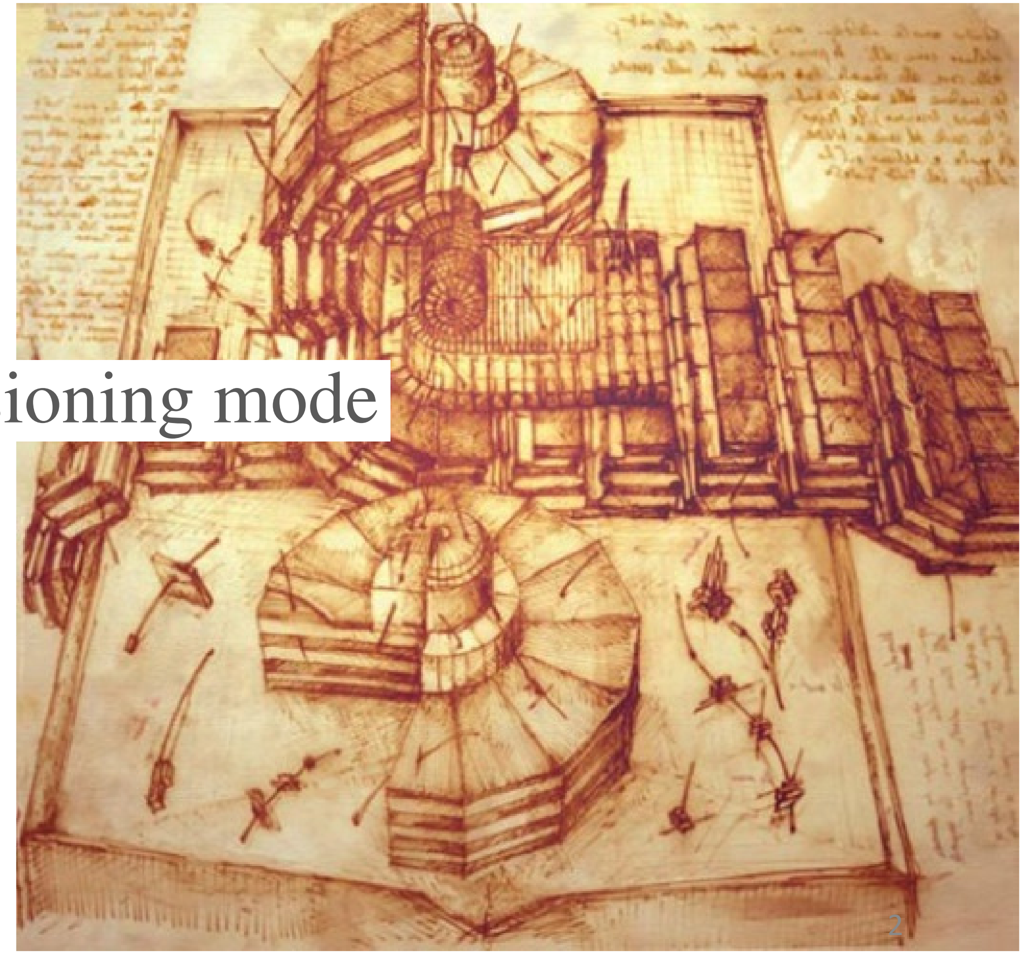

How did it all begin?

And what is the big deal about the trigger?

What are we trying to do?

- Find the most interesting physics signals at LHC
- Store them for off-line processing

What do we expect to see?

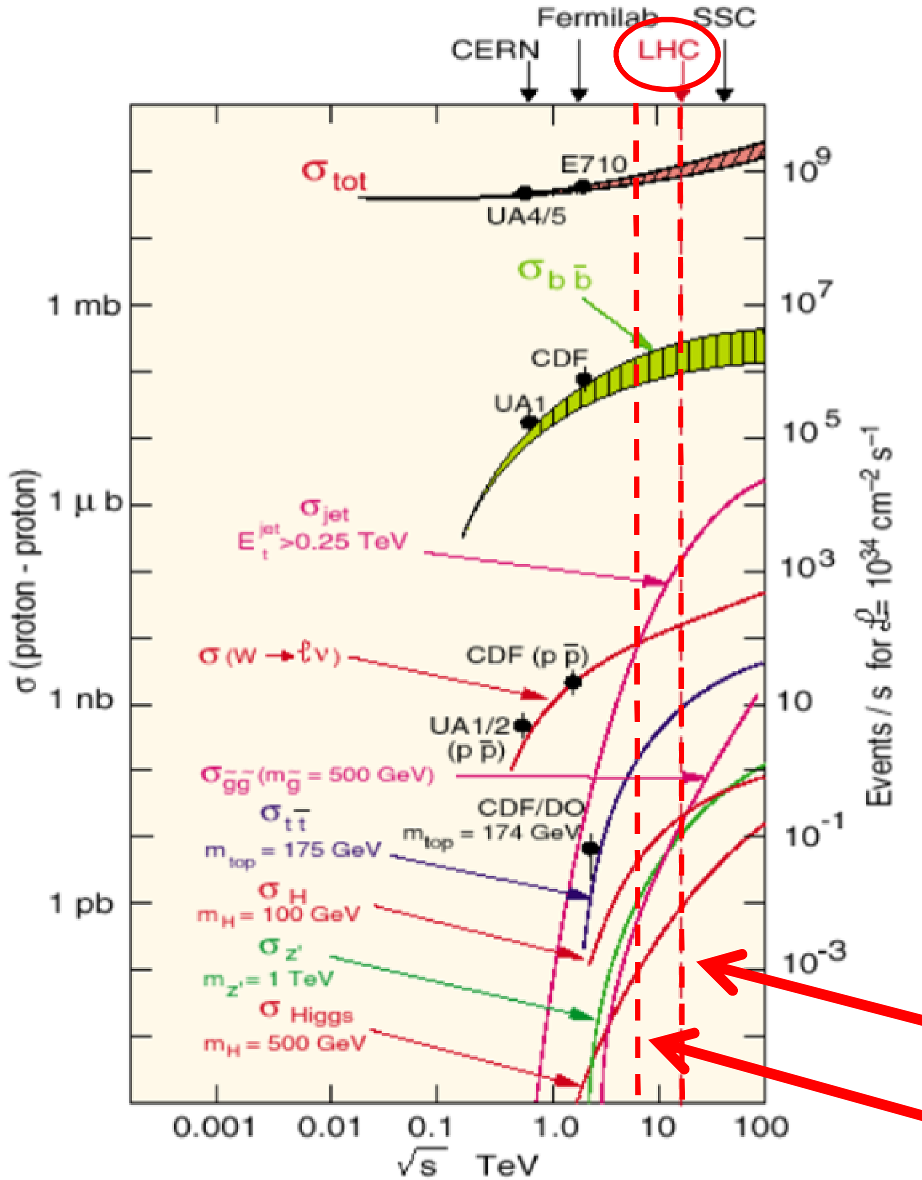

$\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$

Process	σ (nb)	Production rates (Hz)
Inelastic	10^8	10^9
$b\bar{b}$	5×10^5	5×10^6
$W \to \ell\nu$	15	100
$Z \to \ell\ell$	2	20
$t\bar{t}$	1	10
H(100 GeV)	0.05	0.1
Z'(1 TeV)	0.05	0.1
$\tilde{g}\tilde{g}$ (1 TeV)	0.05	0.1
H(500 GeV)	10^{-3}	10^{-2}

You are here

Well, actually at 7 TeV…

What is the problem?

1) We don’t keep all these events → Selection
2) Old Physics happens more often than New Physics
3) New Physics buried under a ton of Old Physics

We don’t keep all these events

- How many do we keep? About 150-300 Hz

- Why only so few? Not enough resources!
  - 300 Hz at 1-2 MB/event → Up to 36 GB per minute
  - Up to 6´000´000 GB of storage needed per year
  - Plus: about 30 secs to reconstruct every event off-line

- “Interesting” physics occurs at ~10, 1 or < 1 Hz
  - We are only interested in a (tiny) fraction of all events
  - We *don’t* really want to keep all these events

Old Physics: more likely than New Physics

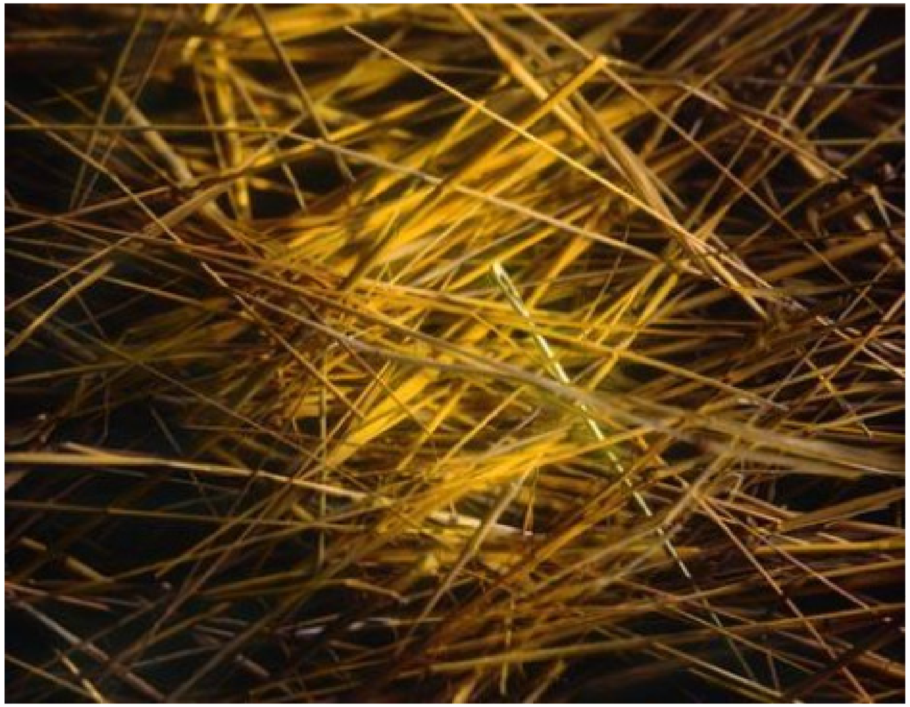

$$\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$$

Process	σ (nb)	Production rates (Hz)
Inelastic	10^8	10^9
$b\bar{b}$	5×10^5	5×10^6
$W \to \ell\nu$	15	100
$Z \to \ell\ell$	2	20
$t\bar{t}$	1	10
$H(100\ \mathrm{GeV})$	0.05	0.1
$Z'(1\ \mathrm{TeV})$	0.05	0.1
$\tilde{g}\tilde{g}\ (1\ \mathrm{TeV})$	0.05	0.1
$H(500\ \mathrm{GeV})$	10^{-3}	10^{-2}

It is challenging (to say the least)
to find these rare exciting events

LHC reference numbers

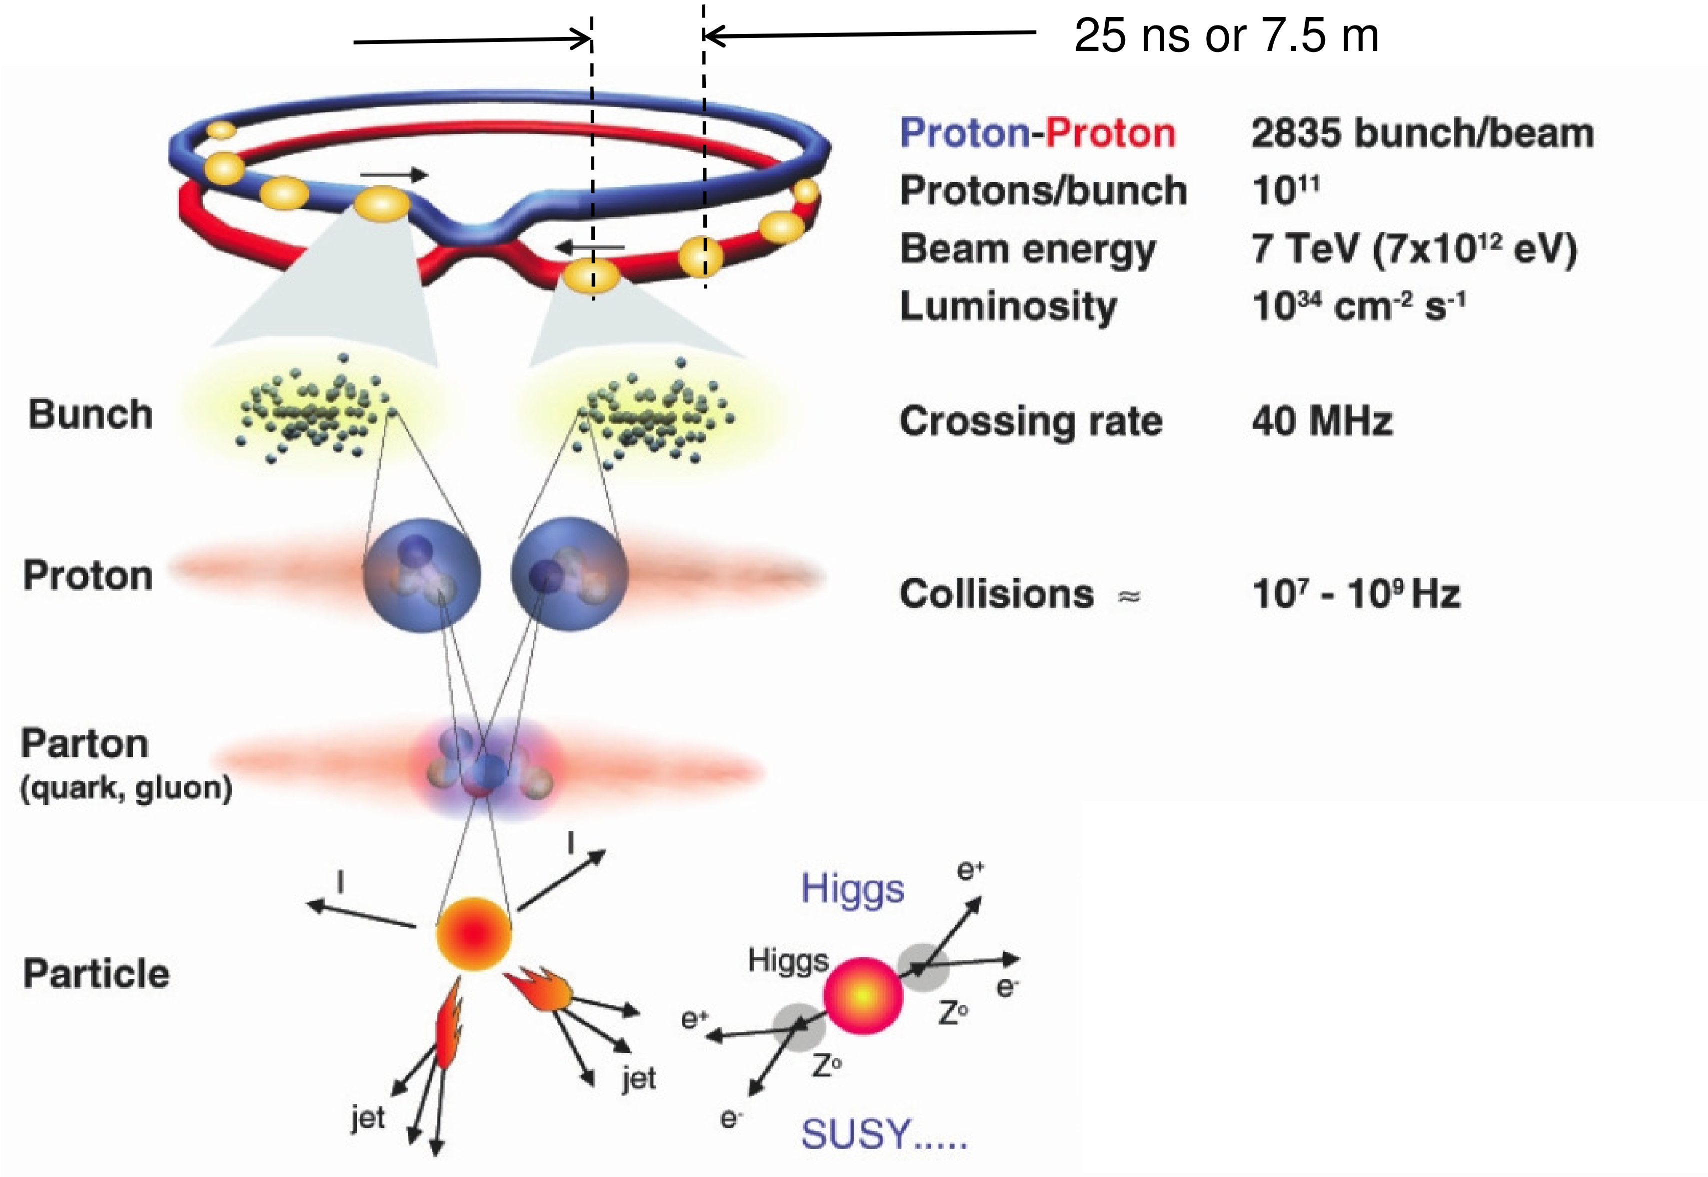

New Physics buried under Old Physics

- Interaction rate:

$$R = \mathcal{L} \times \sigma_{\text{tot}} = 10^{34}\ \text{cm}^{-2}\text{s}^{-1} \times 80\ \overset{(*)}{\text{mb}} \sim 0.8\ \text{GHz}$$

(*) Total inelastic cross section at 14 TeV (±20%)

- Distance between bunch crossings:

$$\Delta t = 25\ \text{ns (or 7.5 m)}$$

- Non-empty bunch crossings:

$$2835 \text{ out of } 3564 \text{ (or } \epsilon = 79.5\% \text{)}$$

- Average # of interactions per (non-empty) crossing:

$$\bar{n} = R \times \Delta t / \epsilon \sim 25$$

New Physics buried under Old Physics

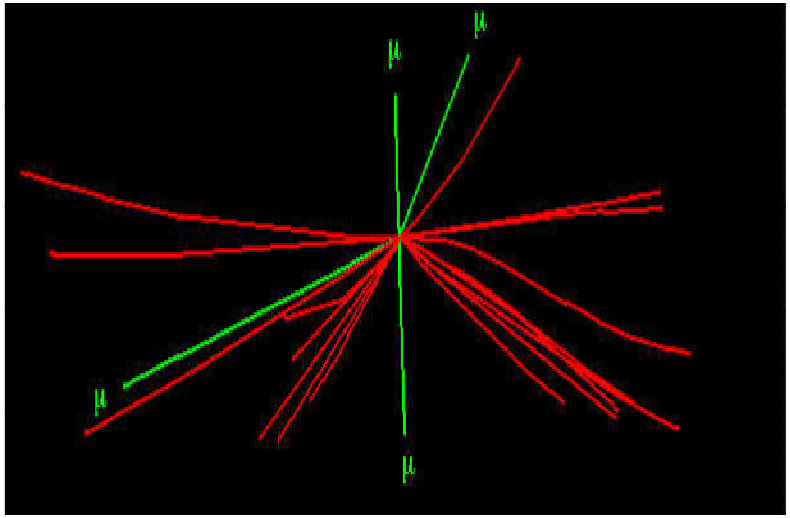

For every exciting interaction…

$$H \to ZZ \to 4\mu$$

Reconstructed tracks with $p_T > 25$ GeV

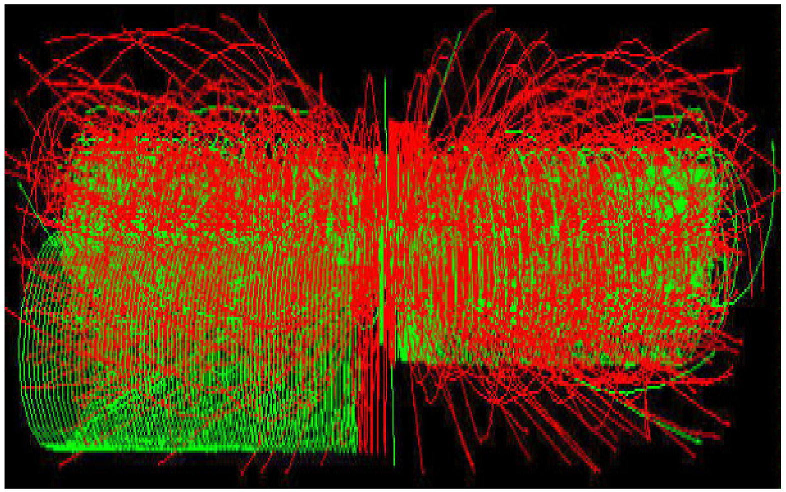

…expect 25 non-exciting *overlaid* interactions (at ~1000 tracks per event)

Reconstructed tracks with $p_T > 2$ GeV

Pileup: serious problem at LHC at high luminosities

The 25 ns challenge

Credit: Daniel Froidevaux

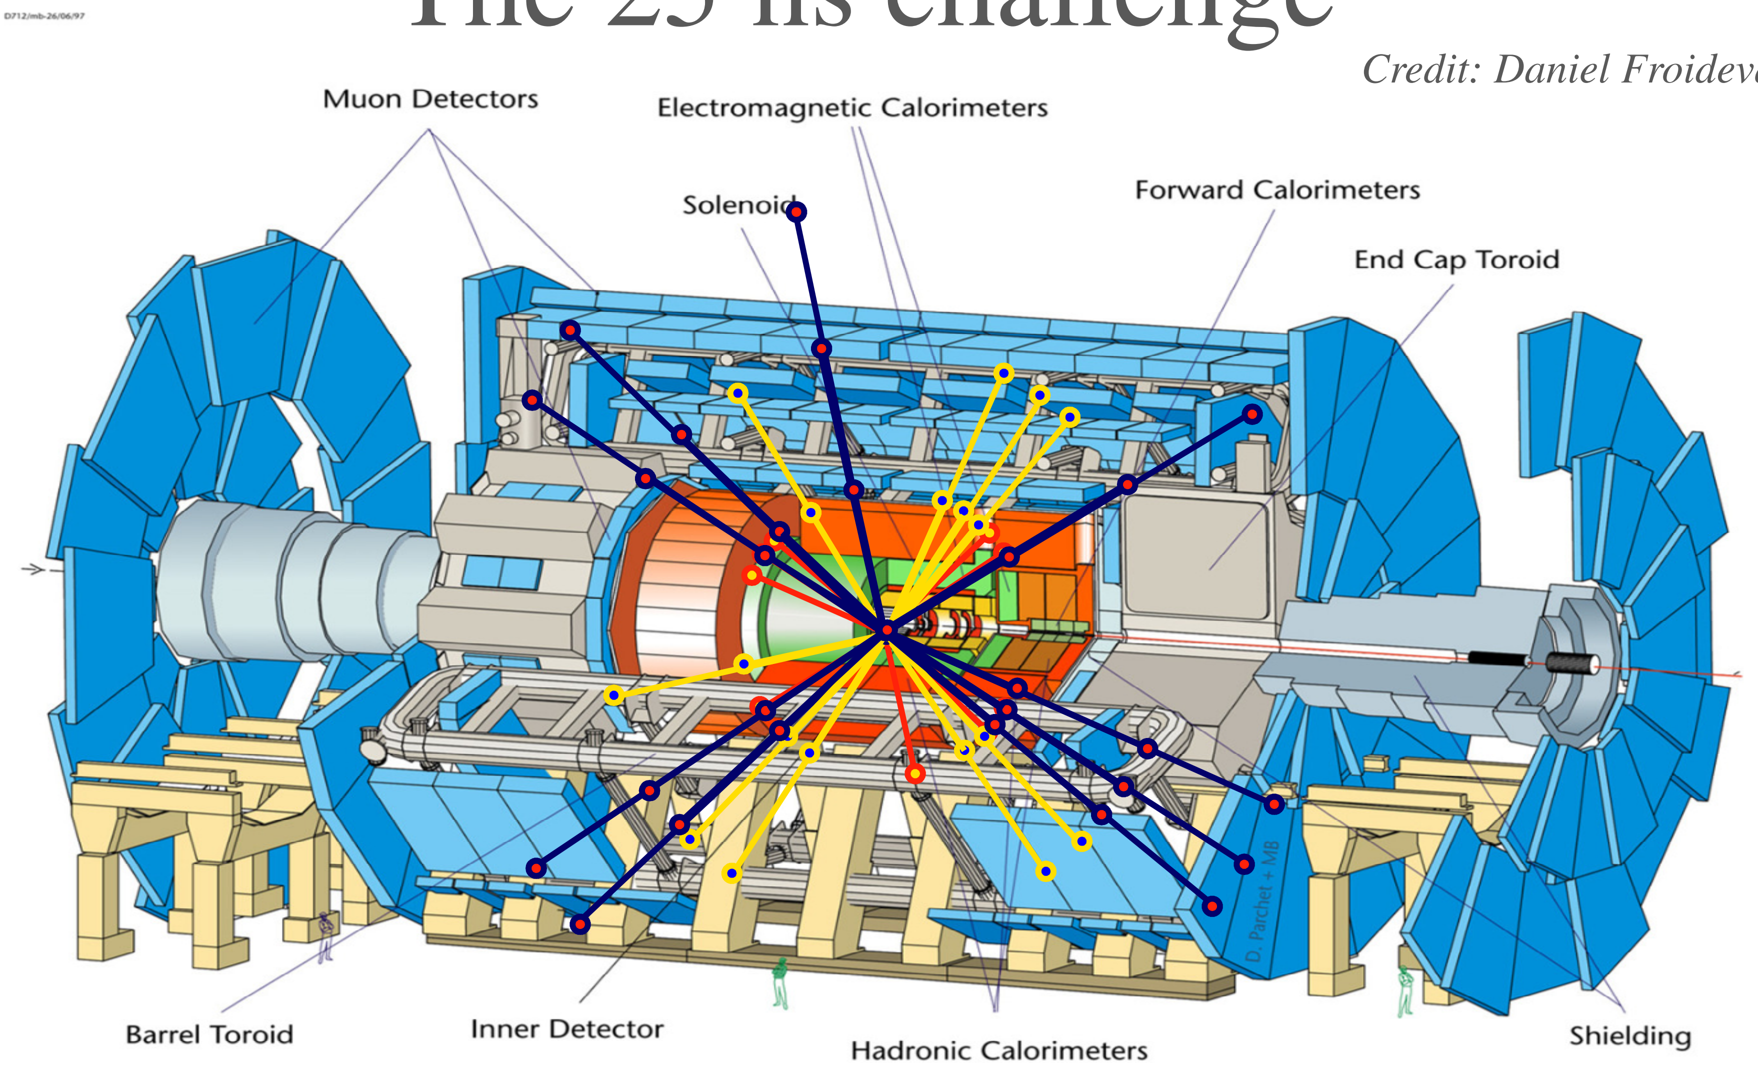

Interactions every 25 ns …

- In 25 ns particles travel 7.5 m

Cable length ~100 meters …

- In 25 ns signals travel 5 m

What are we trying to do? (v.2)

- Select the most interesting physics signals at LHC
  - 150-300 Hz out of ~ 1 GHz of “noise” (selection: 10^{-7})
- In real time
- Store them for off-line processing

Background is a Disease

Meet the Cure

ATLAS and CMS triggers

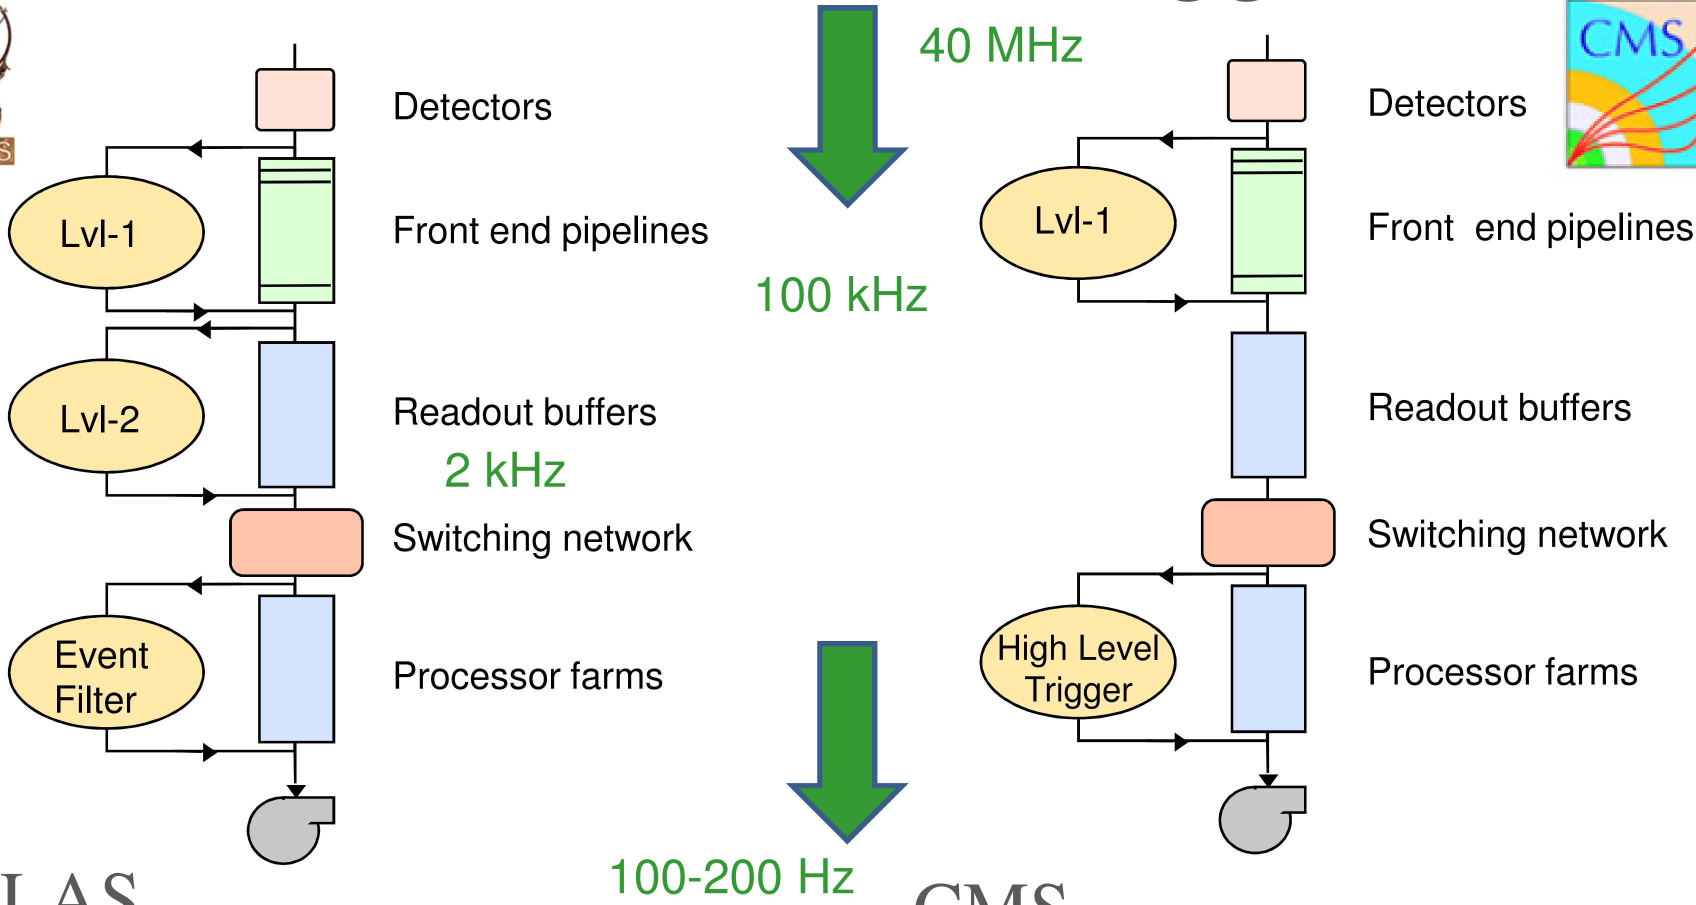

<u>ATLAS</u>

- 3 levels (traditional design)
- L1: hardware, firmware
- L2 + EvF: high-level software

<u>CMS</u>

- L2, L3: merged into HLT
- L1: hardware, firmware
- HLT: high-level software

Particle-id at Level-1

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

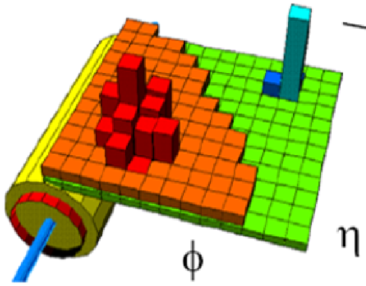

CALORIMETERs

Cluster finding and energy deposition evaluation

MUON System

Segment and track finding

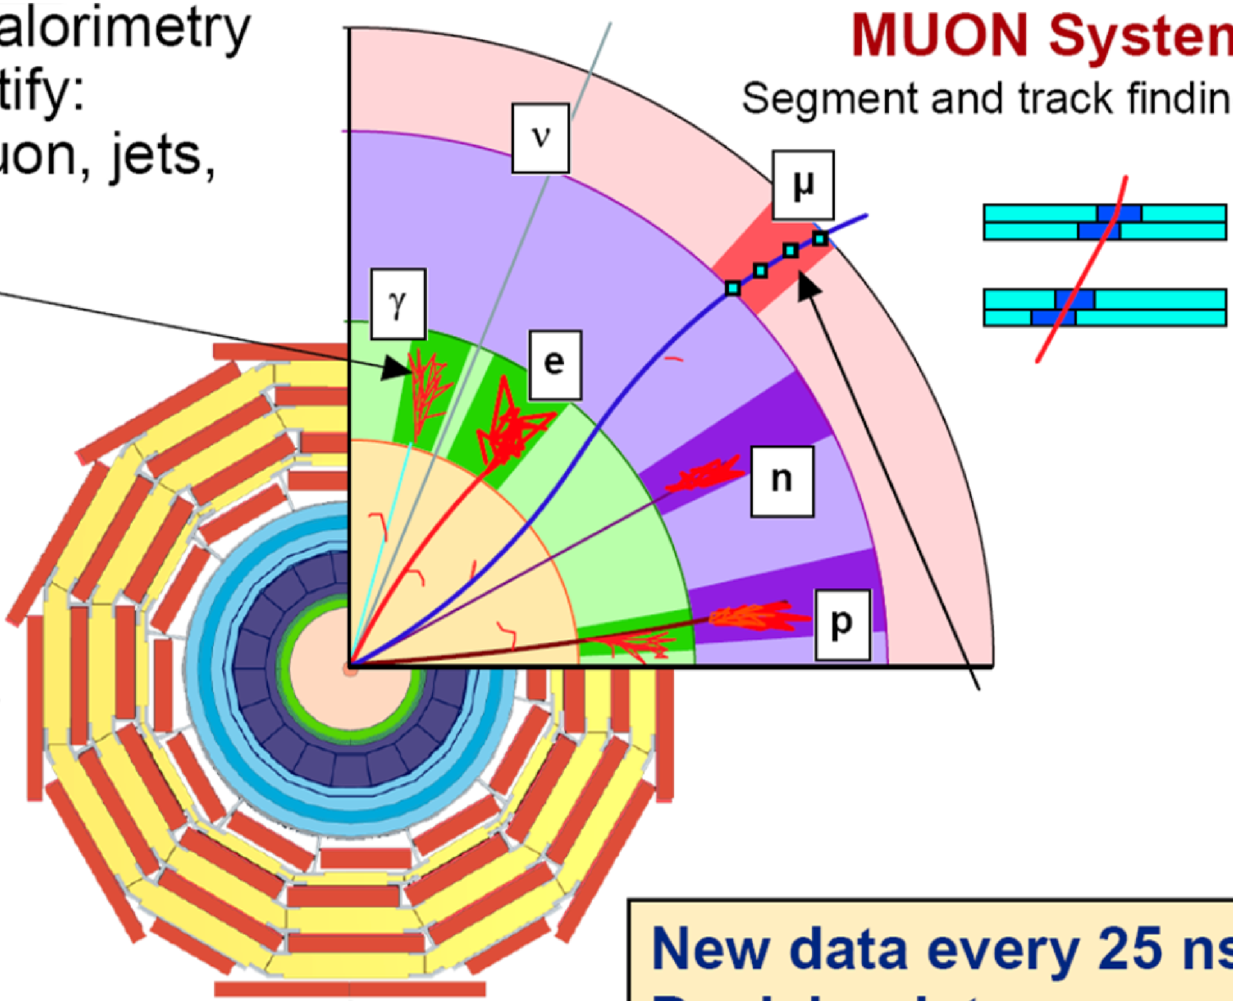

New data every 25 ns
Decision latency ~ μs

Why not use tracker info at Level-1?

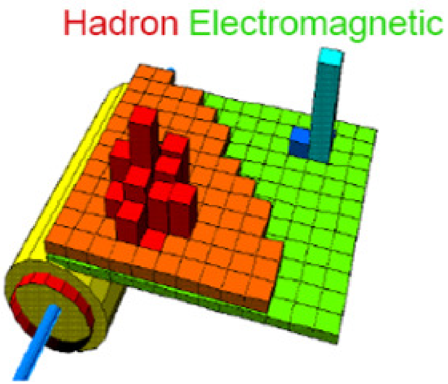

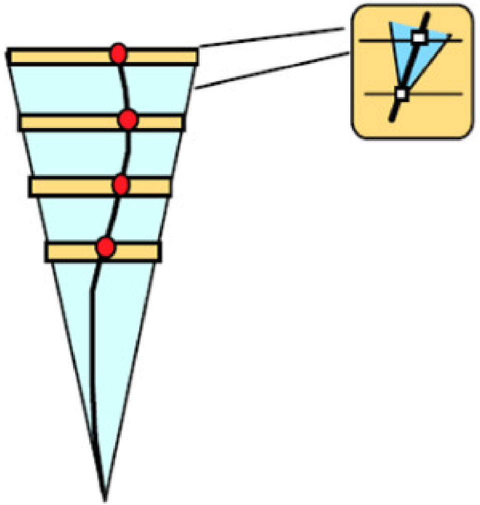

Plans of including tracker info at L1 for SLHC

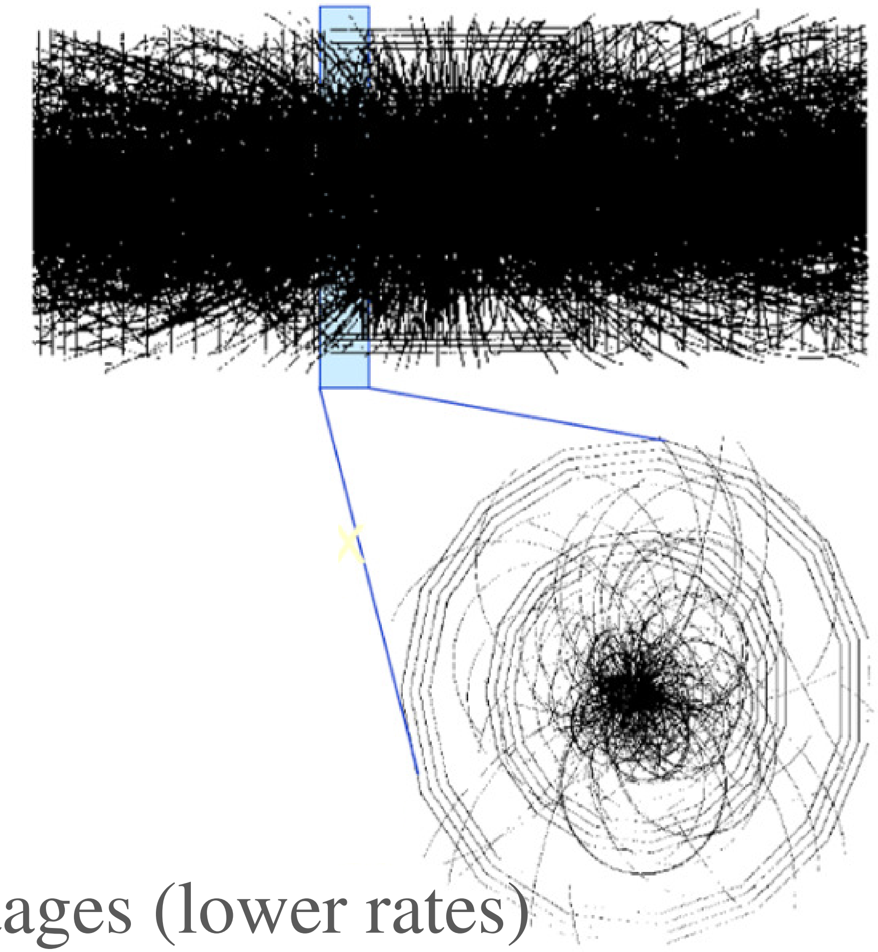

Calorimeter, muon detectors:

- Thousands of channels
- Patter recognition fast

Tracking, vertexing detectors:

- Millions of channels
- Patter recognition slow
- Reserved for later triggering stages (lower rates)

ATLAS High Level Trigger

- L2 and L3 (Event Filter) form High Level Trigger (HLT)
- L2 (~500 CPUs) accesses ~10% of event info (full granularity) seeded by L1 objects
- Event Filter (~2000 CPUs) accesses full event using "off-line quality" algorithms
- Custom L2-steering system
- L1: 2.5 μs, L2: 40 ms, L3: 4s

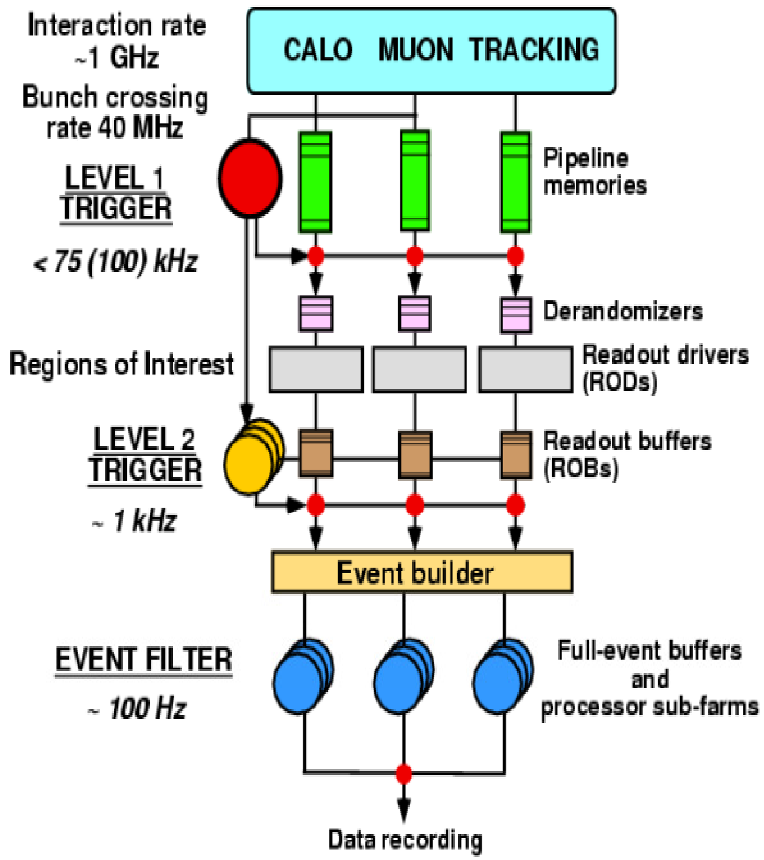

CMS High Level Trigger

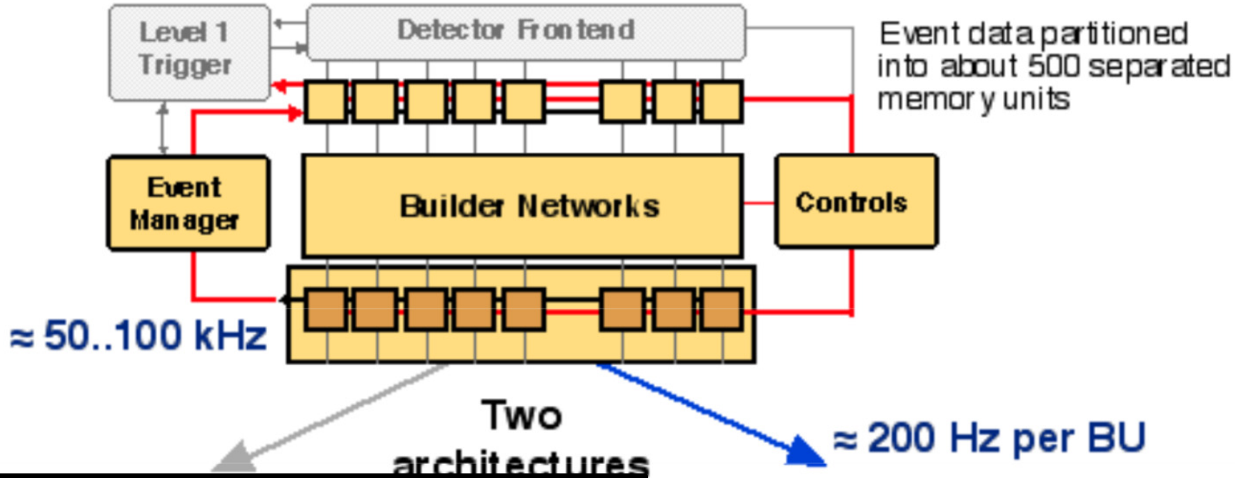

- L2 and L3 merged into High Level Trigger (HLT)
- HLT (~5000 CPUs) accesses full event info (full granularity) seeded by L1 objects using “off-line quality” algorithms
- L1: 3.2 μs, HLT: 40 ms

Farm of processors

ONE event, ONE processor

- High latency (larger buffers)
- Simpler I/O
- Sequential programming

ATLAS *vs.* CMS Triggers

- More traditional, safer design
- Concrete steps & requirements for each of Level-2, Level-3 steps of selection
- Accesses fraction of event at L2 (small throughput)
- But: Custom controls and separate farms for L2, L3

- More flexibility
  - Full event info (and offline reconstruction) as early as L2
  - HLT: continuous software environment in single farm
- But:
  - Large data throughput (and switching network) needed
  - Risky design decision (at the time)

ATLAS *vs*. CMS Triggers

Overall:

- Very similar performances
- Trigger bandwidth determined by detectors and physics programs, not trigger design
- Systems still differ (two farms vs. single farm at HLT) so: commissioning and debugging also different

Trigger: A tricky business

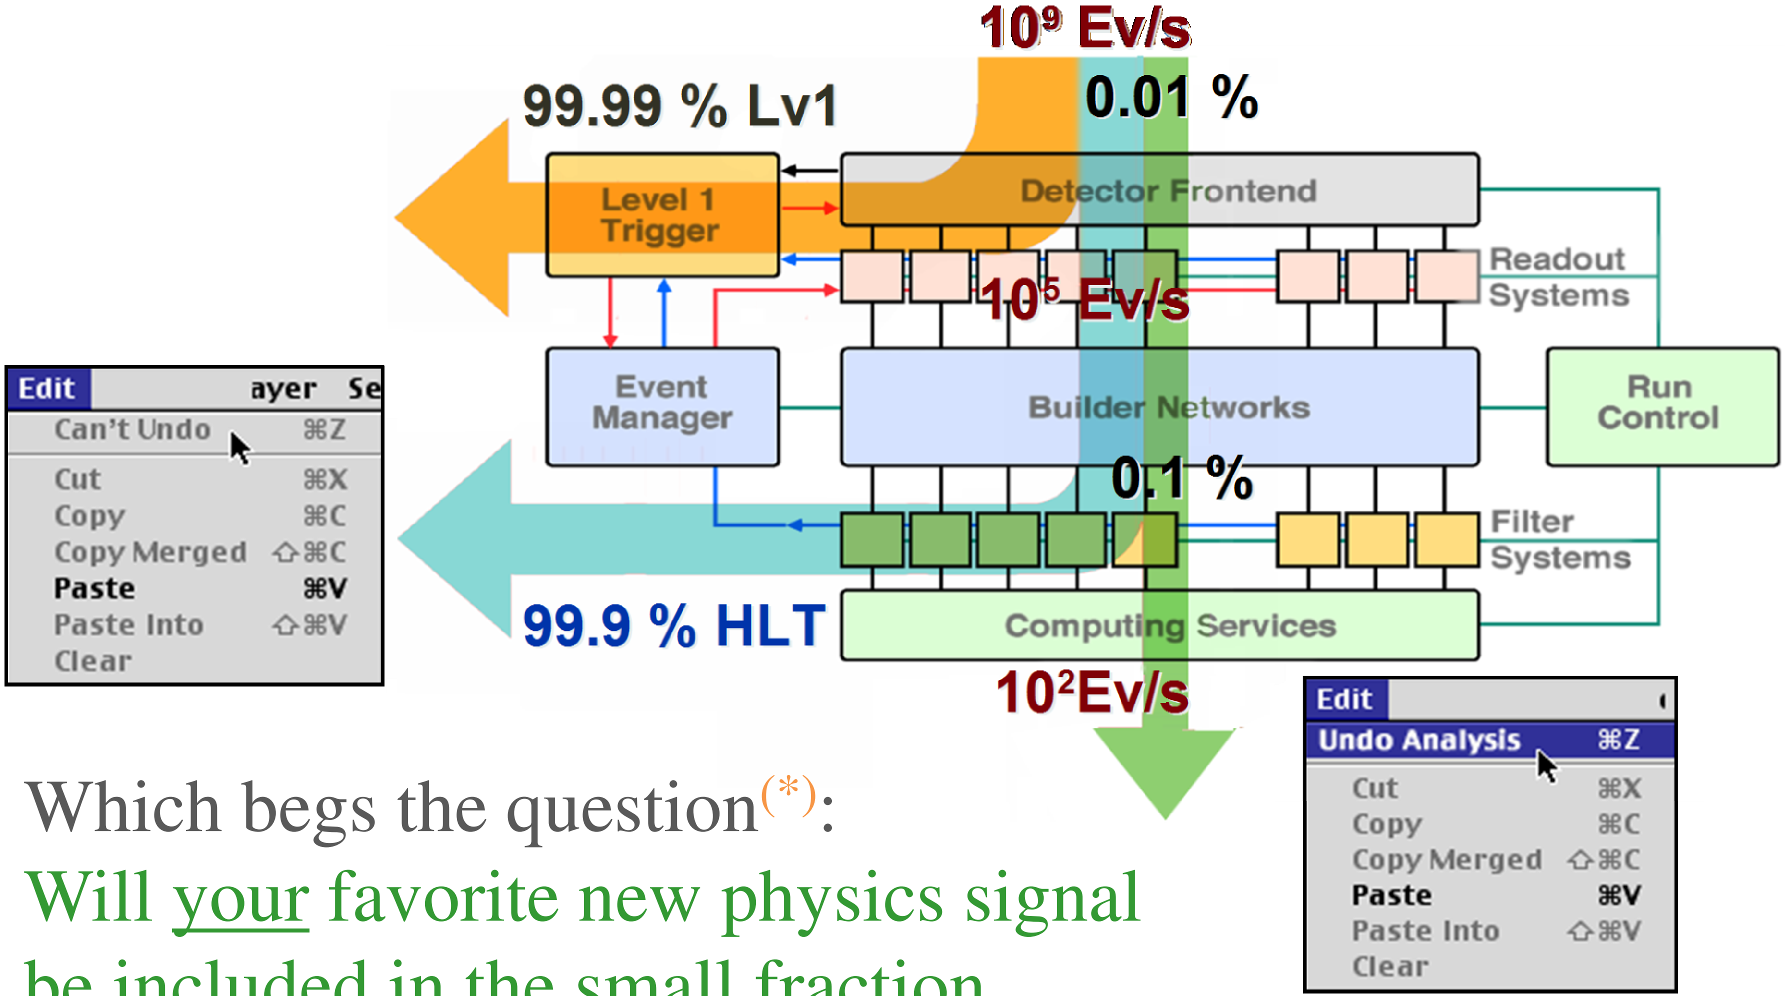

Which begs the question(*):

Will your favorite new physics signal be included in the small fraction of selected events? (unexpected signatures always a worry)

(*) LHC upgrade: 1B CHF, CMS+ATLAS detectors: 1.2B CHF

What are we trying to do? (v.3)

- Select the most interesting physics signals at LHC
  - ➢ 150-300 Hz out of ~ 1 GHz of “noise” (selection: 10^{-7})
- In real time
- Store them for off-line processing
- Don’t screw up

What to avoid

- Killing the interesting physics altogether
- Biasing the selected event samples:
  - Uncertainties in topologies of rejected events
  - Introduction of large systematic errors

arXiv:hep-ex/0502042v3 9 Feb 2007

Reduction of the Statistical Power Per Event Due

Measurements

Abstract

A cut on the max
the number of event
decreases the statist
parameter cut in the
to technical limitations, has the same effect. In this note we describe and
quantify the consequences of such a cut on lifetime measurements. We

The small loss of events due to a moderate upper lifetime cut is accompanied by a large loss of information, because not only a few events outside the allowed time window are lost, but also the information that there were only a few. This can have dramatic effects on the precision of the measurement. As shown

How to build good triggers

Ask old people

Learn from previous experiments

How to build good triggers

- No single silver bullet
- Using common sense (and trigger studies)

General strategies

- As simple as possible
- As inclusive as possible
- Robustness
- Redundancy

Simplicity

- Construct triggers with simple conditions
- Simple triggers easier to
  - ➢commission
  - ➢debug
  - ➢understand

General strategies

- As simple as possible
- As inclusive as possible
- Robustness
- Redundancy

Be inclusive

- Better to have one trigger covering similar analyses
- Even better: covering other, unrelated analyses
- Should be able to discover the unexpected as well

Strong social aspects, often ignored

- Competition inside experiment
  - One (wo)man's signal is another (wo)man's background
  - It's best for your analysis to rely on a popular trigger
- Inertia: people get used to "old" triggers
- Safety: people tend to ignore "new" triggers

General strategies

- As simple as possible
- As inclusive as possible
- Robustness
- Redundancy

Your favorite trigger should be deployed online as early as possible

Robustness

- Make sure your trigger can run for *many* events
  - ➢Including pathological events
  - ➢Including events with x10 more hits than MC predicts
- Make sure your trigger is immune
  - ➢To beam conditions, detector problems

Missing E_T : the popular trigger for

- SUSY particles
- Dark matter candidates
- But also: neutrinos (so: *W*s, Higgs, etc)

General strategies

- As simple as possible
- As inclusive as possible
- Robustness
- Redundancy

Missing E_T at DØ

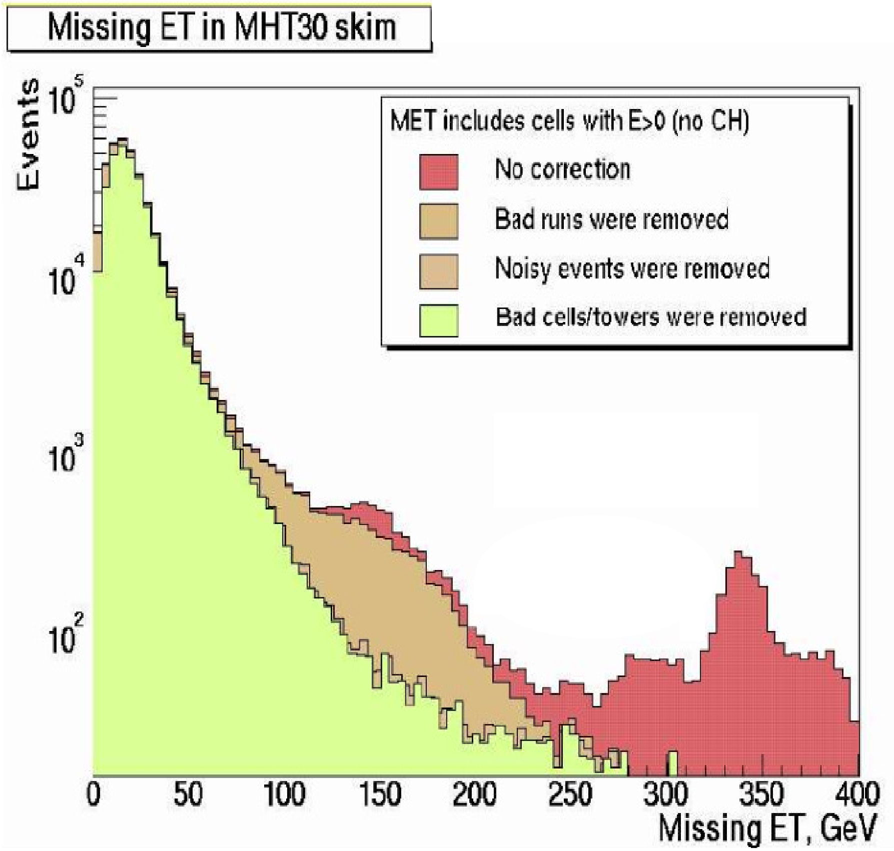

It takes a *long* time to

- Commission the detector for data-taking
- Remove all problematic runs
- Understand noisy environment
- Discover (and remove) problematic channels

Missing E_T :

- Not ideal for startup
- Typically the last trigger to be commissioned

Redundancy

- Make sure your signal can be selected by more than one triggers
  - ➢Helps to understand biases
  - ➢Ensures that if a trigger has problems (rates too high or instability) you still get your events

General strategies

- As simple as possible
- As inclusive as possible
- Robustness
- Redundancy

How is the trigger different at LHC?

Trigger trends

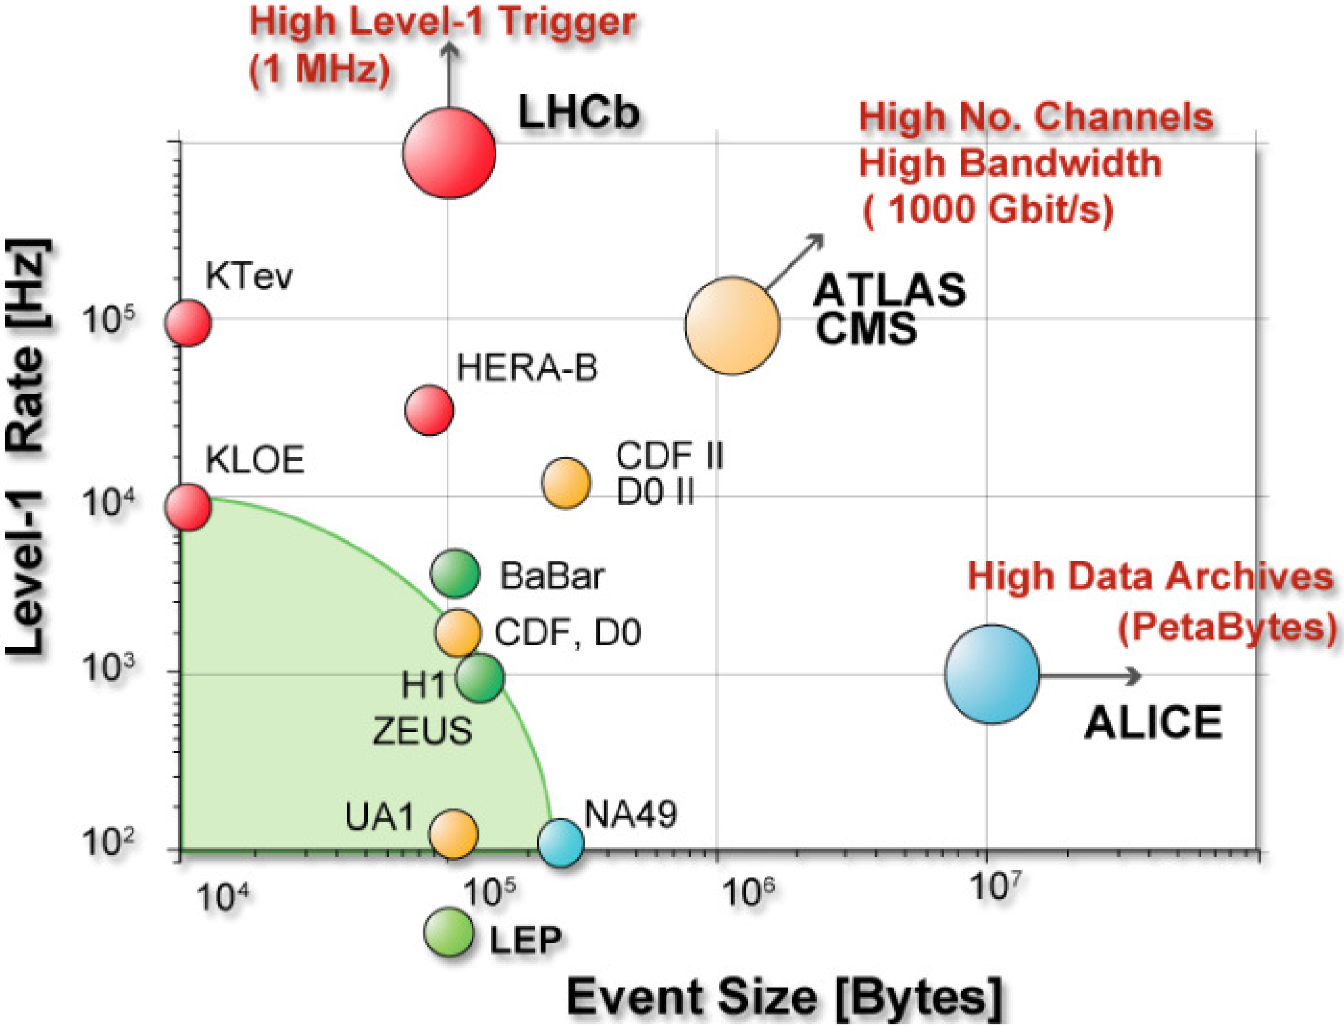

Luminosity, rates, event sizes:
all increased by ~an order of magnitude

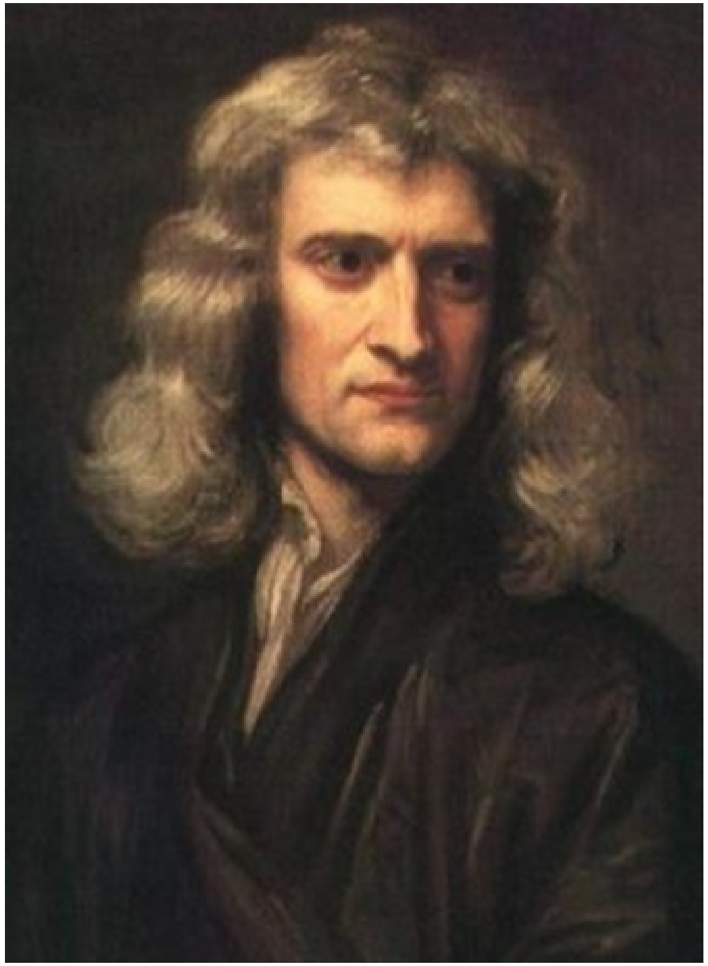

"If I have seen further it is by standing on the shoulders of Giants"

Evolution in computing

Advances in

- Networking (Ethernet, Terabit/s networks)
- PC industry (computing power and memory abundance)
- Software standards (Linux, http, XML, C++, Java)

offer affordable, modular, scalable, upgradable solutions

CMS trigger: scalable

The trigger at CMS evolves with luminosity

Adjusts to increases in:

- DAQ capacity (L1 rate)
- CPU-power needed at HLT
    - ➢By adding/upgrading PCs as necessary

CMS trigger at low luminosities

Lower luminosities allow us to trigger

- with lower thresholds, looser requirements
  - ➢e.g. no isolation on leptons
- on physics that we cannot trigger on later
  - ➢e.g. *B* physics, quarkonia or other low-p_T physics

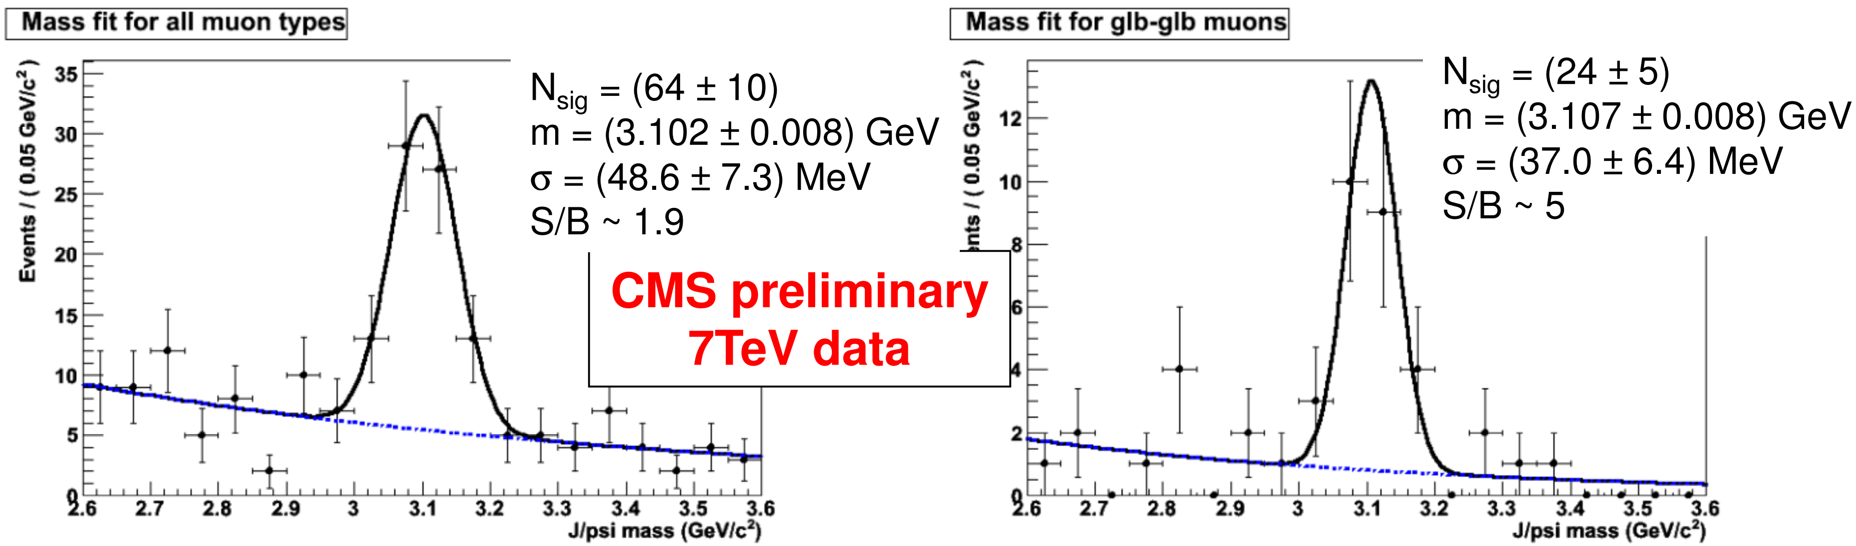

All di-muons

Only global-global muons

The
20 November – 16 December 2009
Revolution

The 2009 CMS Trigger Strategy

- Remove most requirements, put thresholds at lowest possible value
  - Calorimeter: low E_T electron ($\geq$ 5 GeV) or Jet ($\geq$ 10 GeV)
  - Muons without any threshold (effective kinematic limit: 3 GeV p_T cut)
  - Other candidate triggers for later running are active for diagnostic and study purposes (trigger efficiencies)

The 2009 CMS Trigger Strategy #2

- Bandwidth investment in “technical” triggers
  - Zero-bias triggers (using LHC machine clock information)
    - Filled-bunch coincidence
  - Beam-gas triggers
    - Unpaired filled-bunch + detector activity
  - Calibration, alignment, monitoring
  - Save small fraction of L1-accepted events irrespective of HLT decision (“Auto-accept” for CDF or “mark-n-pass” for DØ)
  - Suite of minimum bias triggers (any detector activity)

Trigger Commissioning at CMS

On the road to First Collisions

- Adapt to *rapidly* changing conditions
  - Beam “splashes” & circulating beam
  - Two (unstable) beams with magnetic field off
  - Two (stable) beams with magnetic field (and tracker) on

- Write out as many events as possible
  - Thou Shalt Not Unnecessarily Reject Events
  - Thou Shalt Capture as many bunches with protons as possible (“zero-bias” rate: from 11 to 88 kHz)
  - Thou Shalt Capture All Events with Any Detector Activity (“min.bias” rate: up to 600 Hz)

First collisions in CMS – Monday, 23 November 2009

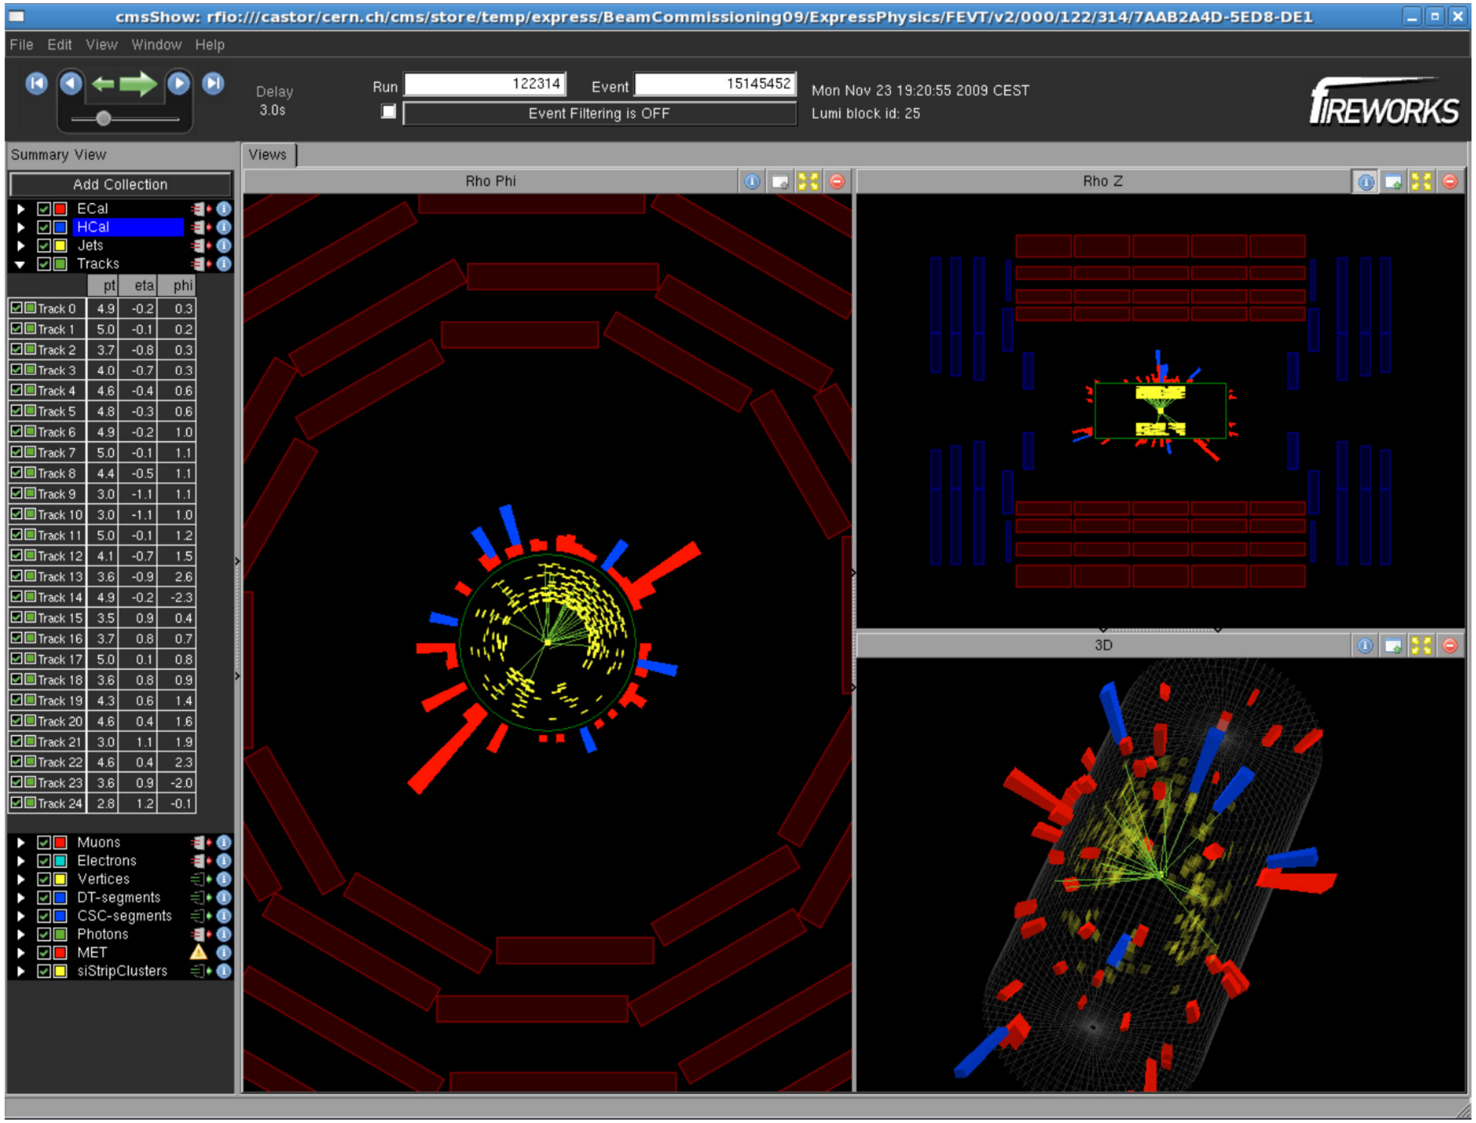

By the end of 2009

- Collision data taken at
  - 900 GeV (350 k min. bias events or 10 μb^{-1}), and
  - 2.36 TeV (20k min. bias events or $< 1\ \mu b^{-1}$)
    - Collider energy world record

- CMS took good quality data
  - > 99% of detector channels operational
  - High data-taking efficiency (> 80%)
  - Data can be analyzed very quickly

2010: Collisions at 7 TeV

Press Release

Search

this site All CERN

Releases | For Journalists | For CERN People

LHC research programme gets underway

PR07.10
30.03.2010

Geneva, 30 March 2010. Beams collided at 7 TeV in the LHC at 13:06 CEST, marking the start of the LHC research programme. Particle physicists around the world are looking forward to a potentially rich harvest of new physics as the LHC begins its first long run at an energy three and a half times higher than previously achieved at a particle accelerator.

"It's a great day to be a particle physicist," said CERN[1] Director General Rolf Heuer. "A lot of people have waited a long time for this moment, but their patience and dedication is starting to pay dividends."

First 7 TeV collisions in CMS – Tuesday, 30 March 2010

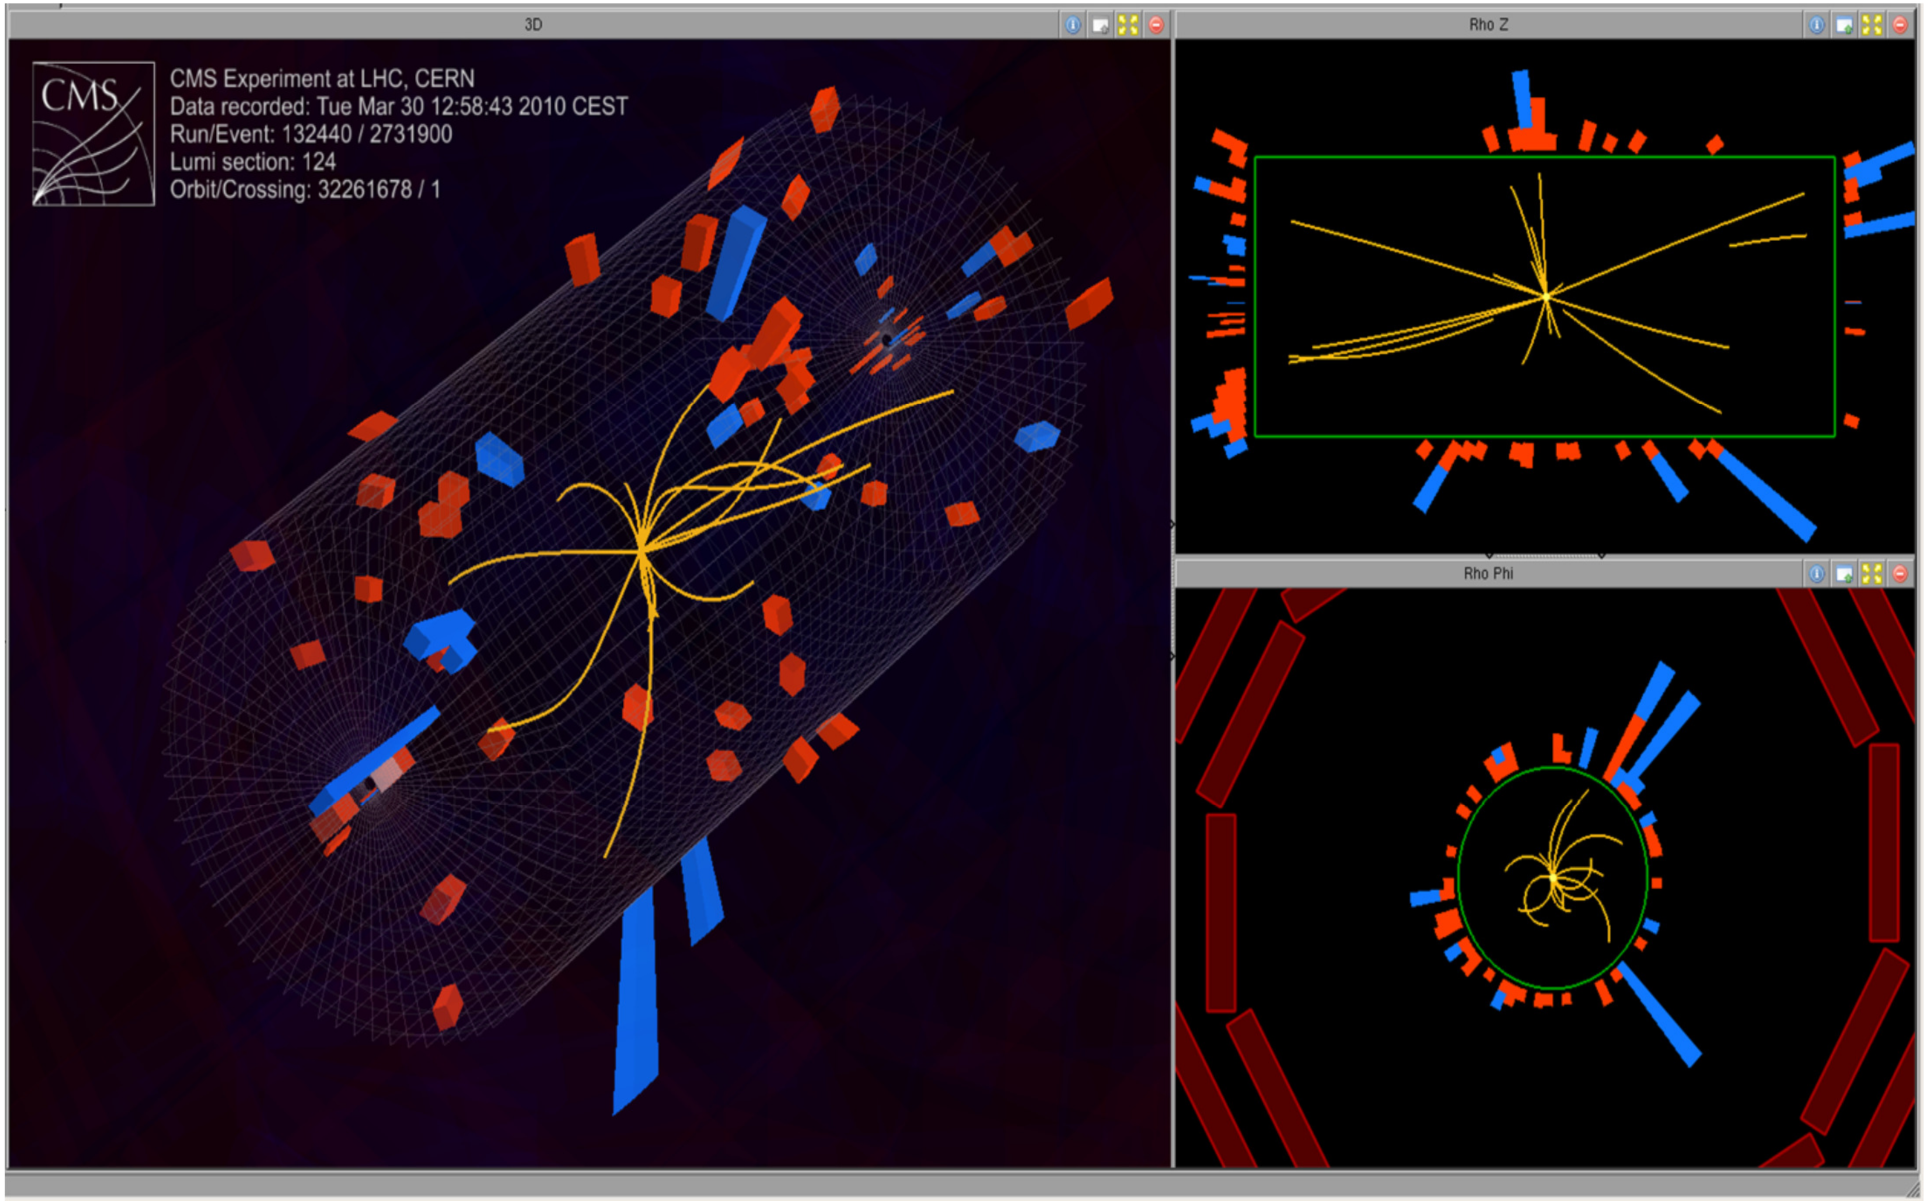

First 7 TeV collisions in CMS – Tuesday, 30 March 2010

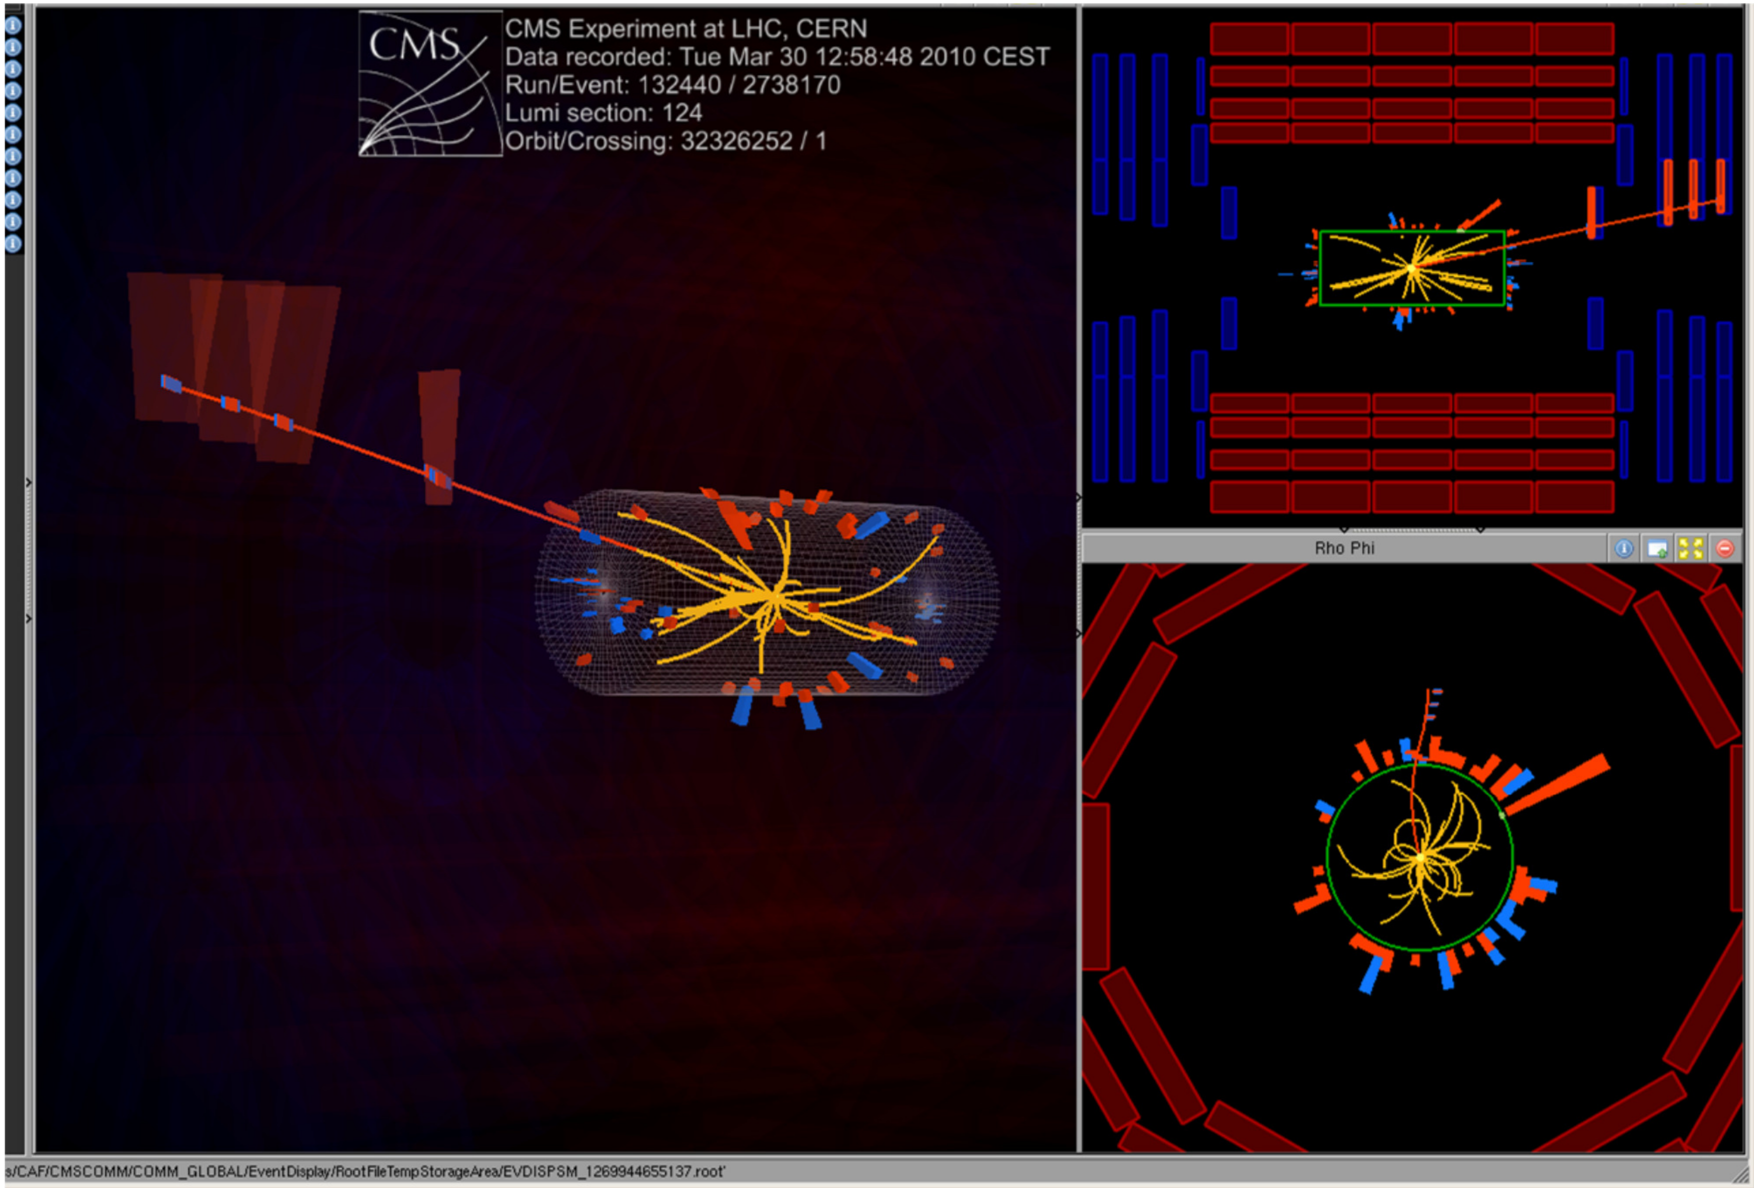

What is happening now

- LHC running at 7 TeV since March 30th
  - ➢ Have passed 1 nb^{-1} benchmark a month later
  - ➢ $L > 10^{28}$ cm^{-2} s^{-1} , total inelastic (min.bias) rate: ~ 1 kHz
  - ➢ Data-taking efficiency > 92% (96% in last two weeks)

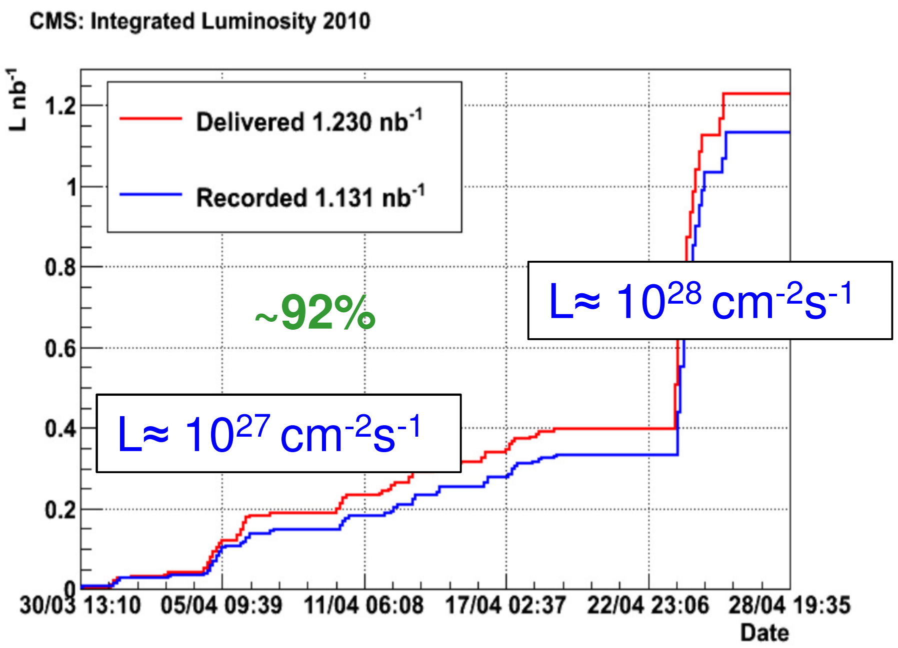

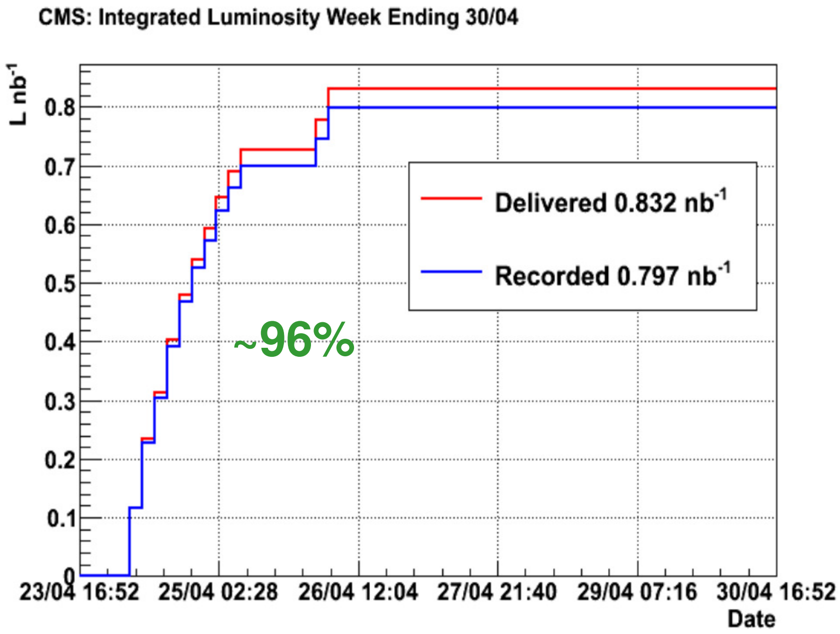

Trigger Performance

- Data collected from "express stream" are skimmed & provide a reference sample for fast-feedback trigger performance studies
- Monitoring is based on run-by-run & weekly basis
- Triggers demonstrate expected performance

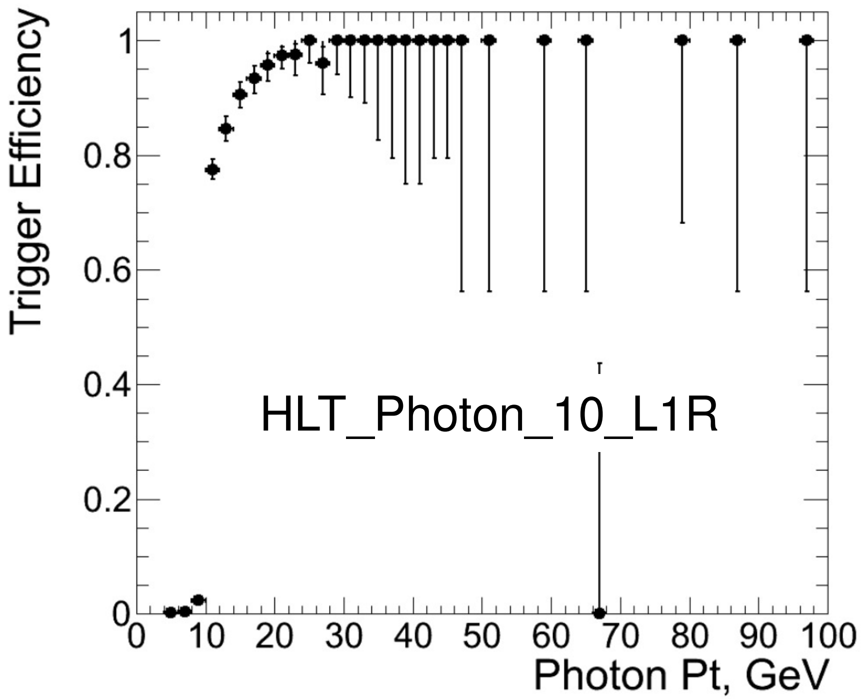

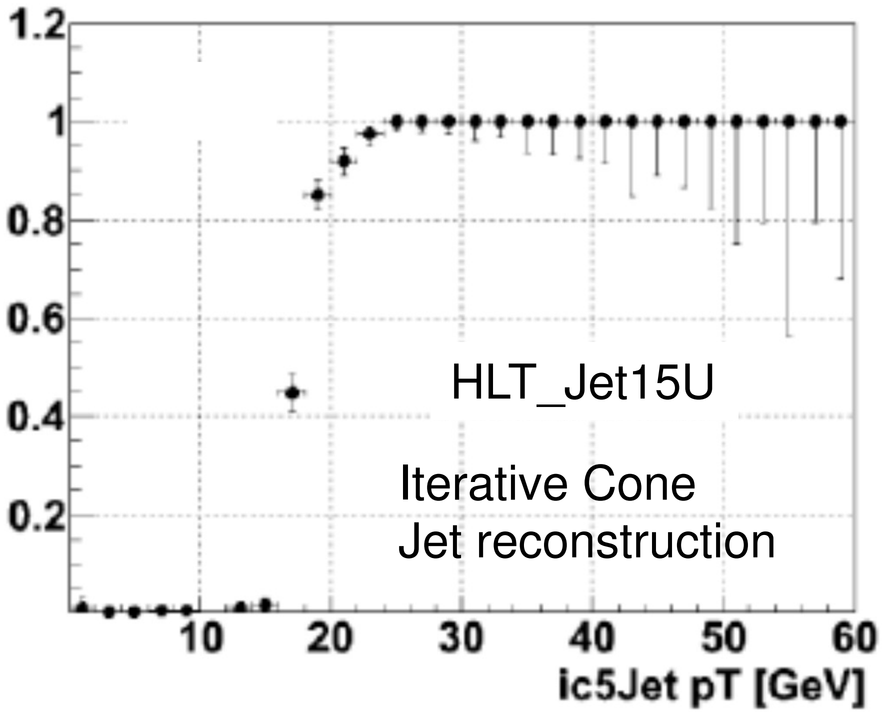

What is happening next

- Experiments preparing for ICHEP
  - Expecting 2 pb^{-1} by July
  - Expecting 100 pb^{-1} (?) by end of 2010

- Will skip 2010-11 Winter shutdown
  - Run "continuously" for ~1.5 years
  - Expecting > 1 fb^{-1} (?) by end of 2011

- Long shutdown afterwards to prepare for 14 TeV run

Trigger Menu Factory

- Running with "1E28" menu
- Have prepared the next series of menus
  - 1E29, 2E29, 4E29, 8E29…
- Re-optimization of (MC-tuned) 1E31 menu underway
- Start planning the 1E32 step (and beyond?)

Trigger Menu Factory #2

- Rate predictions based on MC & data
- Evolution strategy for higher luminosities
- Studying impact of pile-up on CPU-performance & rates
  - Dedicated "multiple vertex" trigger to capture pile-up events

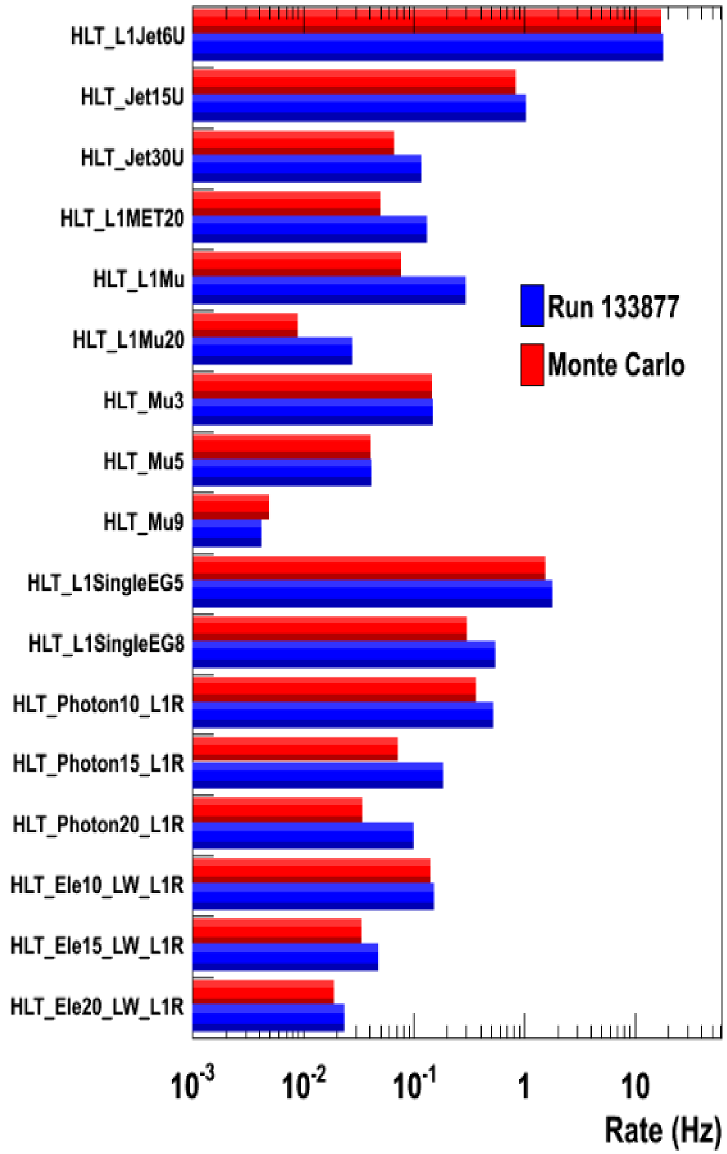

Luminosity effects

$H \rightarrow ZZ \rightarrow \mu\mu ee$ event with M_H= 300 GeV for different luminosities

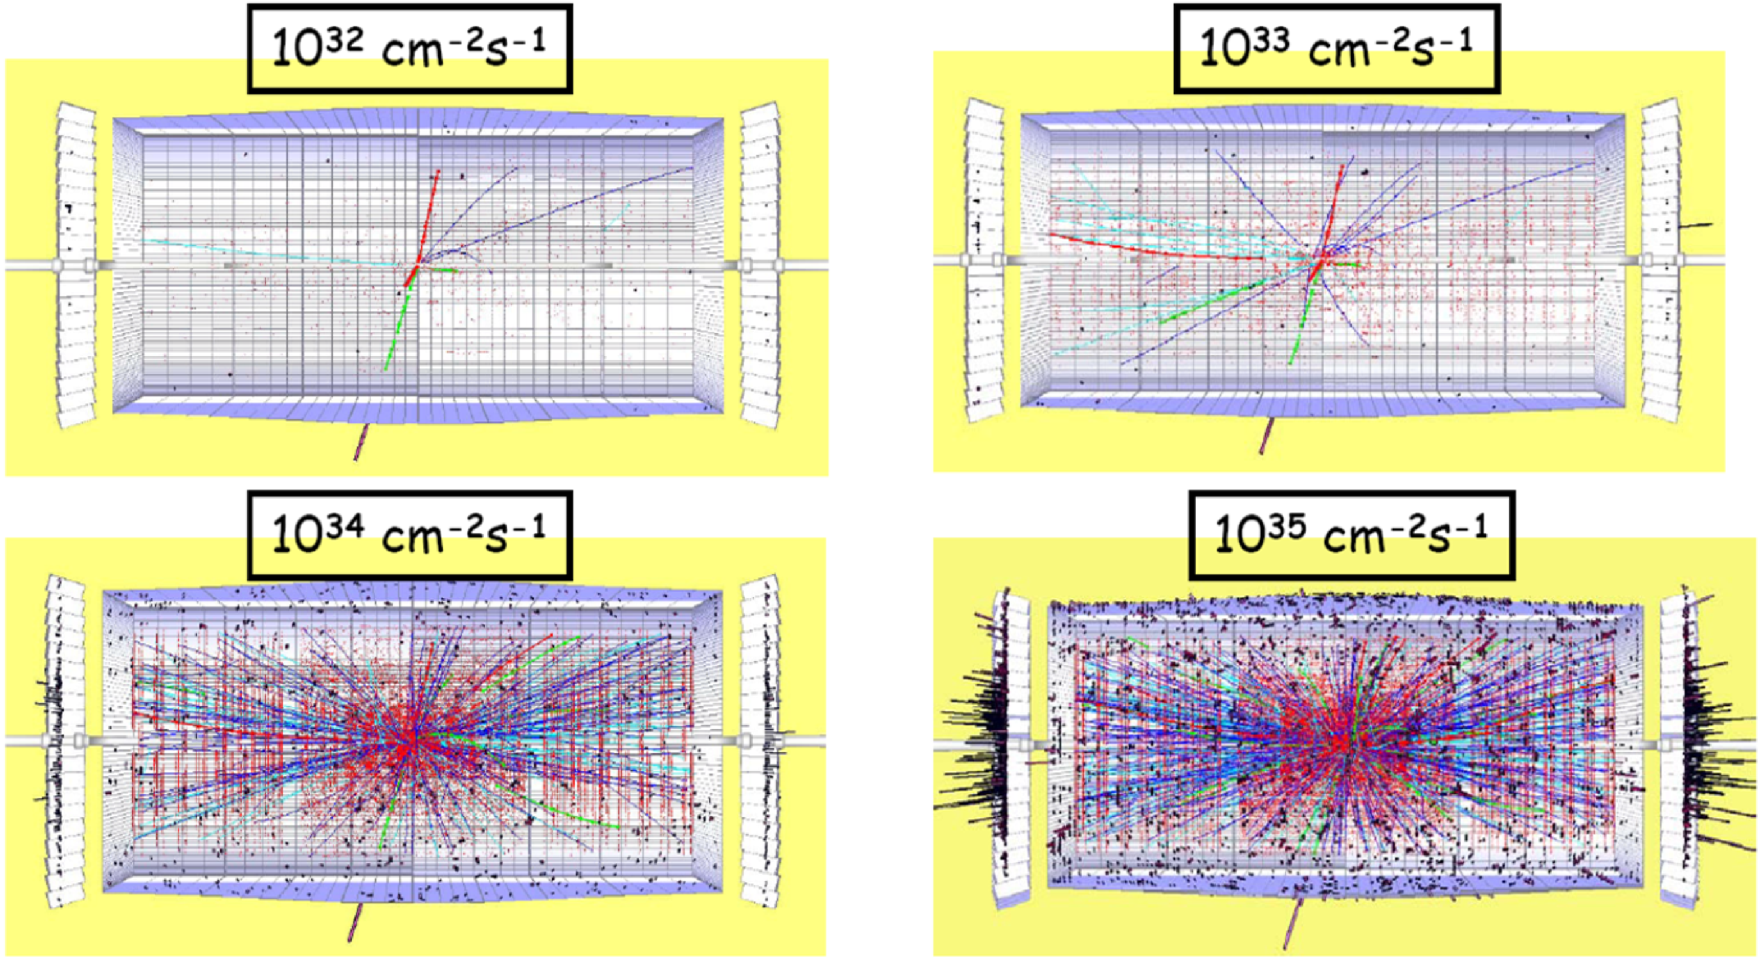

The Bottom Line

Fit everything into O(100) Hz

- How should the bandwidth be shared among the large number of available triggers?

  A difficult question – many things to consider:

  - Are triggers inclusive enough?
  - Which triggers are used by what physics analyses?
  - What are the experiment's priorities?

Example #1:
"Experiment X has a stronger chance of discovering the Higgs first"

Example #2:
"Rumors are that experiment Y is seeing a bump on channel Z.
We must increase bandwidth of corresponding trigger"

Epilogue

- The trigger is a dynamic creature, made by human beings
  - Bound to imperfections, common sense, inertia and strong personalities
  - Must evolve with time, luminosity increases and better detector understanding
  - It requires dedicated studies by analysis users

- But it remains the single most important item in hadron experiments: what makes the difference between discovering New Physics at LHC or not

The Beginning of The Journey

Credit for “Da Vinci” drawings: Sergio Cittolin

Backup

LHC Startup

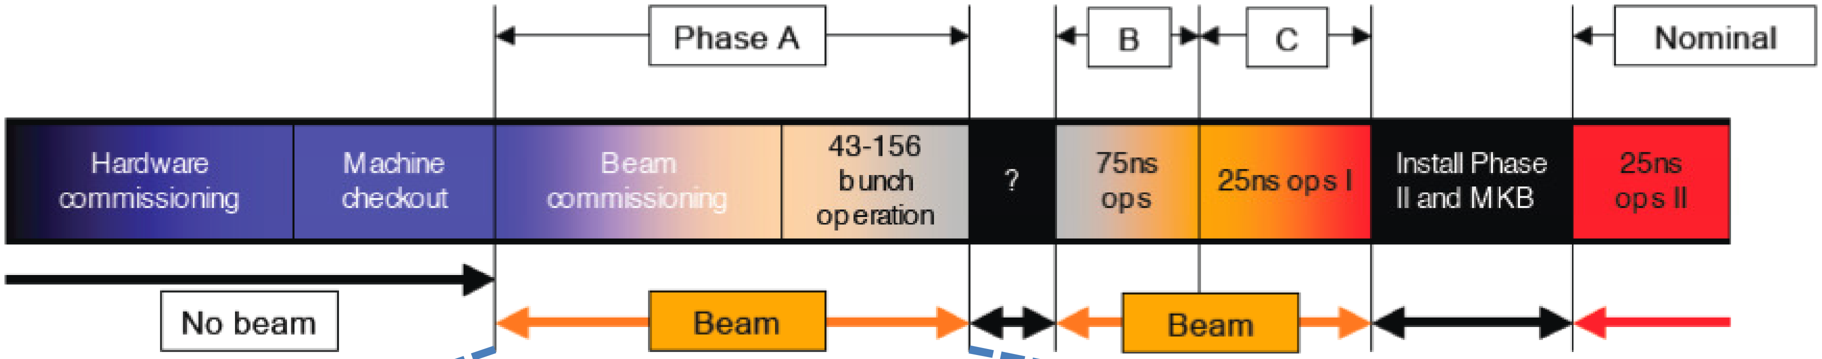

Bunches	β*	I_b	Luminosity	Event rate
1 x 1	18	10^{10}	10^{27}	Low
43 x 43	18	3×10^{10}	3.8×10^{29}	0.05
43 x 43	4	3×10^{10}	1.7×10^{30}	0.21
43 x 43	2	4×10^{10}	6.1×10^{30}	0.76
156 x 156	4	4×10^{10}	1.1×10^{31}	0.38
156 x 156	4	9×10^{10}	5.6×10^{31}	1.9
156 x 156	2	9×10^{10}	1.1×10^{32}	3.9

Maximum luminosity expected in 2010: 10^{32} cms^{-2} s^{-1}

Contingency & Early Menu Designs

- Safety factor of 3 in allocation of L1 bandwidth; only 17 kHz allocated to simulated channels – to account for:
  - Uncertainty in maximum DAQ bandwidth (especially at startup)
  - Input cross sections (especially QCD; Tevatron shows factors of ~2)
  - All that we have not simulated:
    - beam conditions, noise spikes, other electronics correlations…

- Safety factor of 2 in HLT accept rate; only 150 Hz allocated to simulated channels – to account for
  - Uncertainties in cross sections (e.g. heavy-flavor cross section)
  - Uncertainties in simulation (e.g. rate for a jet faking an electron: experience from Tevatron experiments shows Monte Carlo reliable to within a factor 2)

The CMS Detector

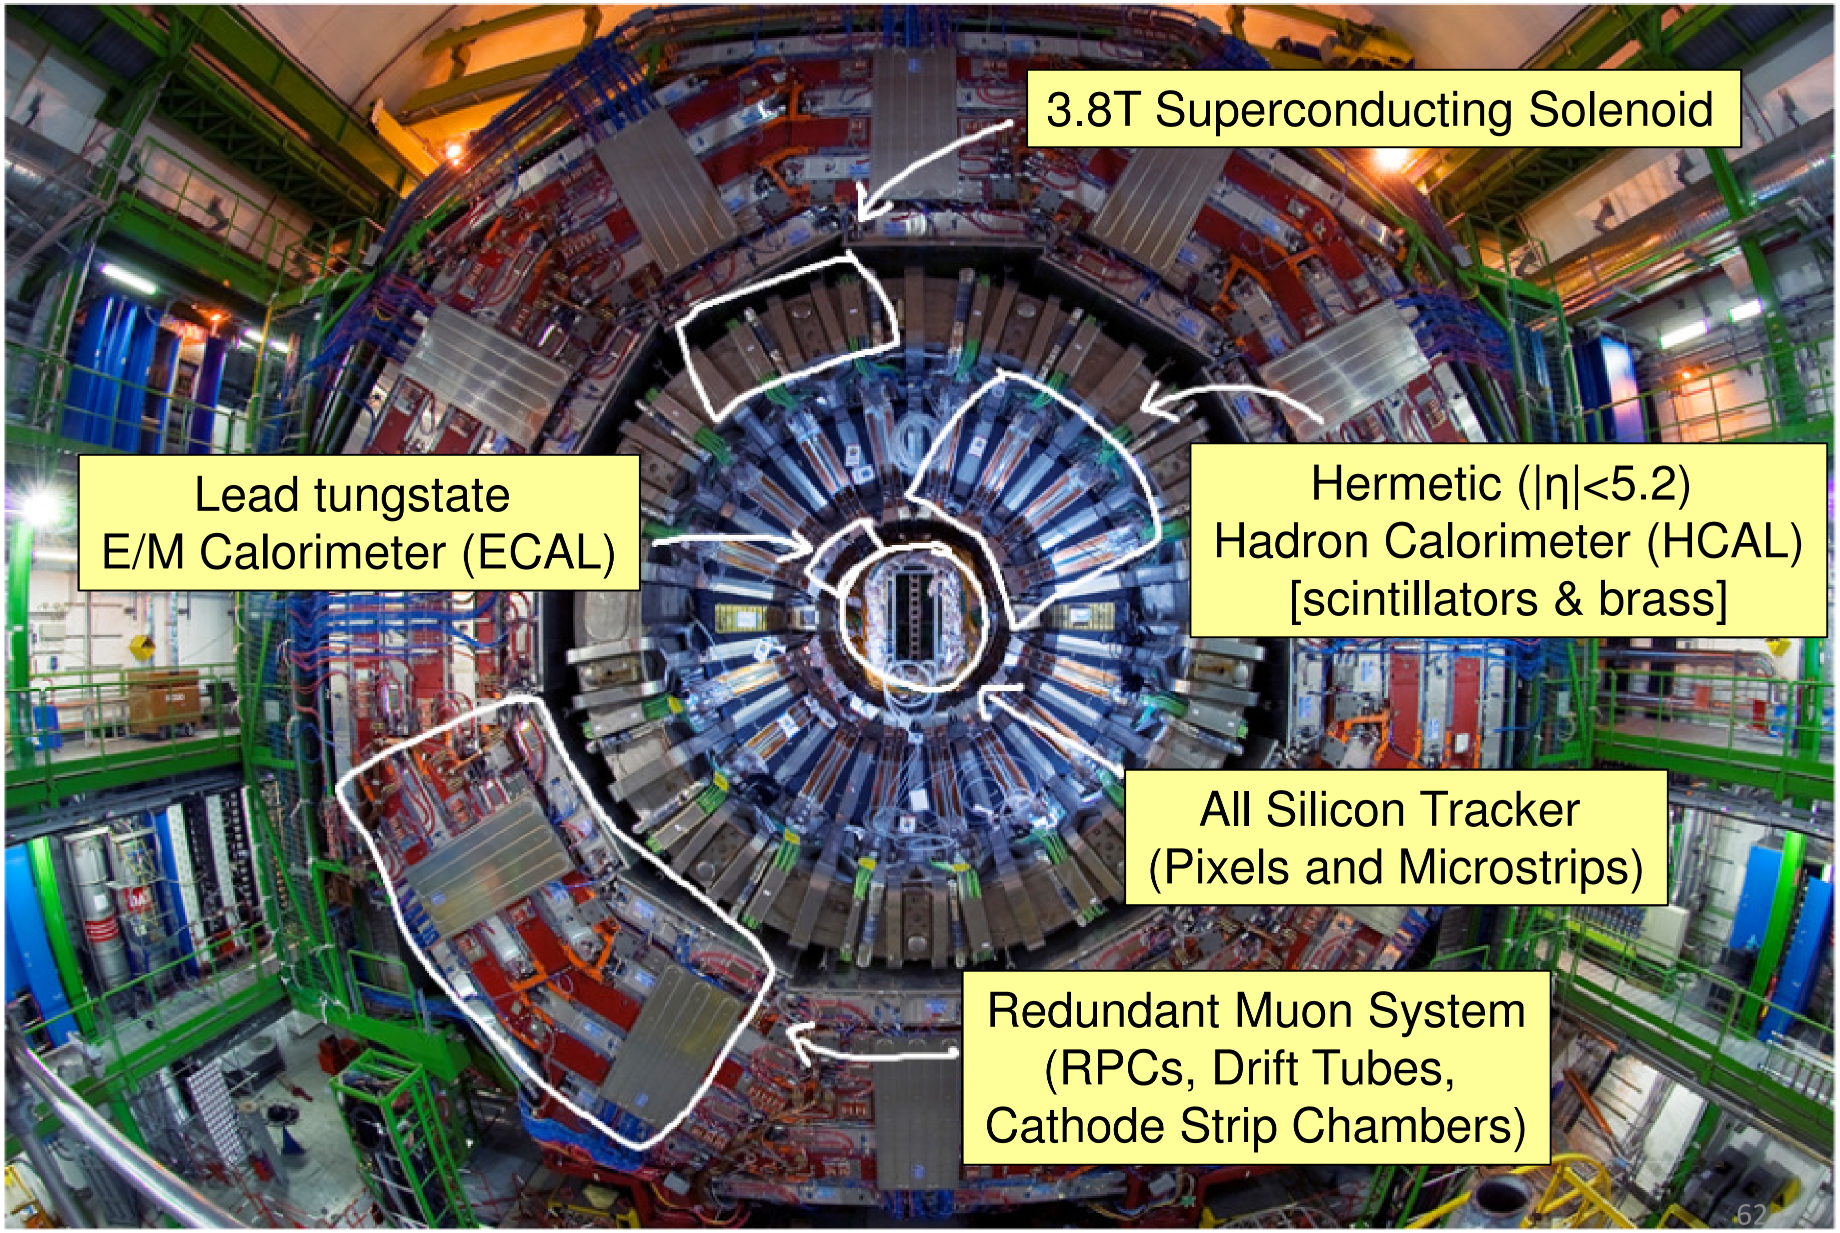

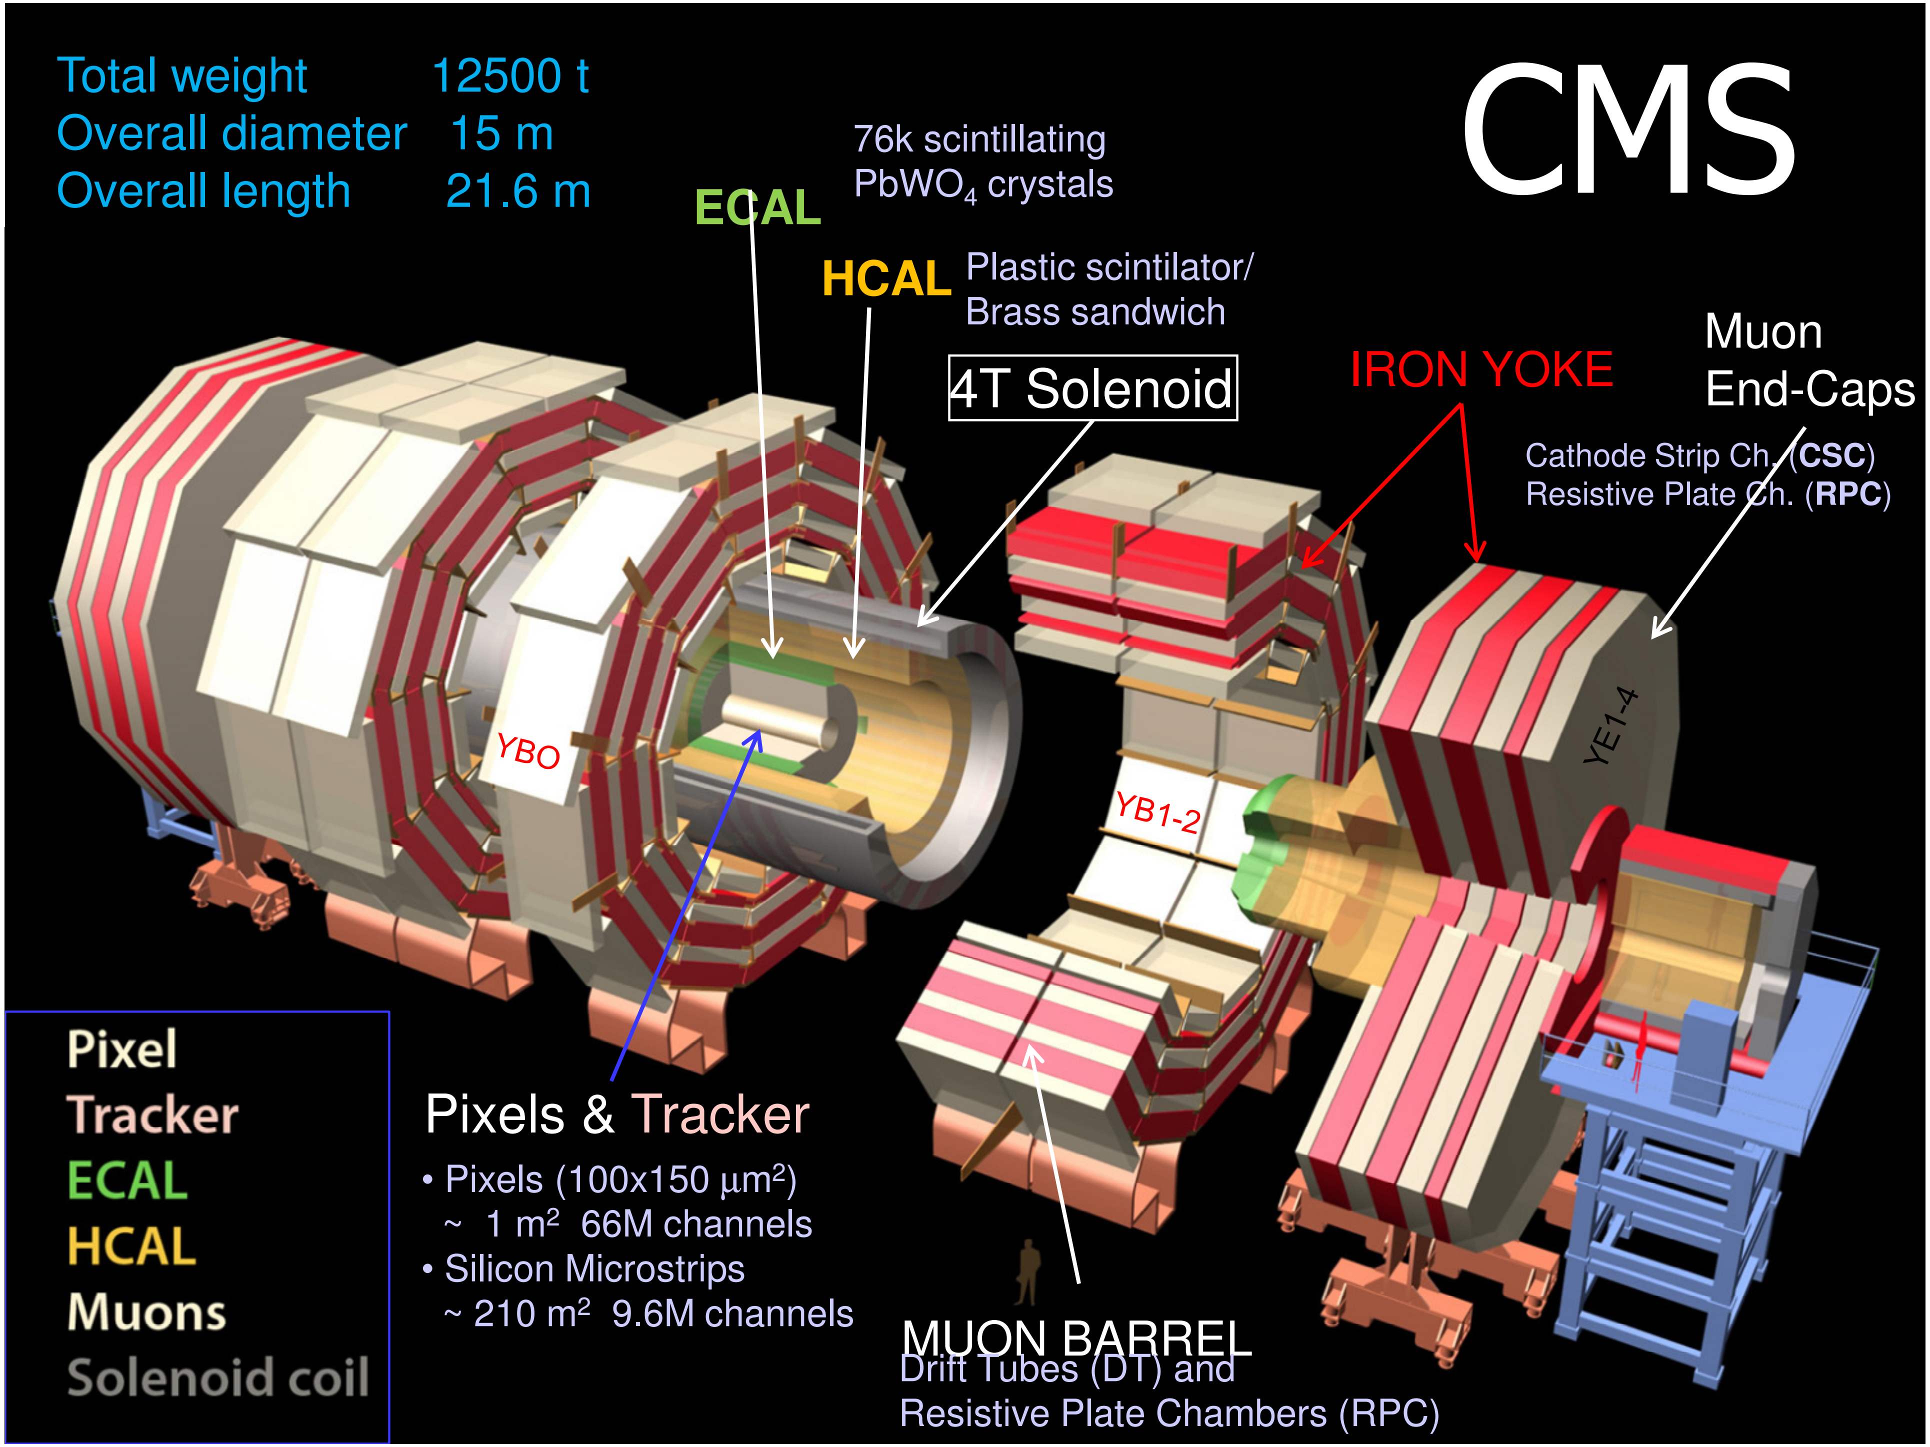
Total weight 12500 t
Overall diameter 15 m
Overall length 21.6 m
CMS
ECAL
76k scintillating $PbWO_4$ crystals
HCAL
Plastic scintilator/ Brass sandwich
4T Solenoid
IRON YOKE
Muon End-Caps
Cathode Strip Ch. (**CSC**)
Resistive Plate Ch. (**RPC**)
YBO
YB1-2
YE1-4
Pixel
Tracker
ECAL
HCAL
Muons
Solenoid coil
Pixels & Tracker
• Pixels (100x150 μm^2)
~ 1 m^2 66M channels
• Silicon Microstrips
~ 210 m^2 9.6M channels
MUON BARREL
Drift Tubes (DT) and
Resistive Plate Chambers (RPC)

Di-muon resonances: $J/\psi \to \mu^+\mu^-$

Pairs of global-global, global-tracker or tracker/tracker muons.

All muon tracks:

- Pixel layers with hits > 1
- Number of pixel+strip hits > 11
- |do| < 5 cm, dz < 20 cm

Global muons:

- global $\chi^2 < 20$

Tracker muons:

- track $\chi^2 < 5$
- TMLastStationAngTight bit on

Vertex probability > 0.1%

7TeV data

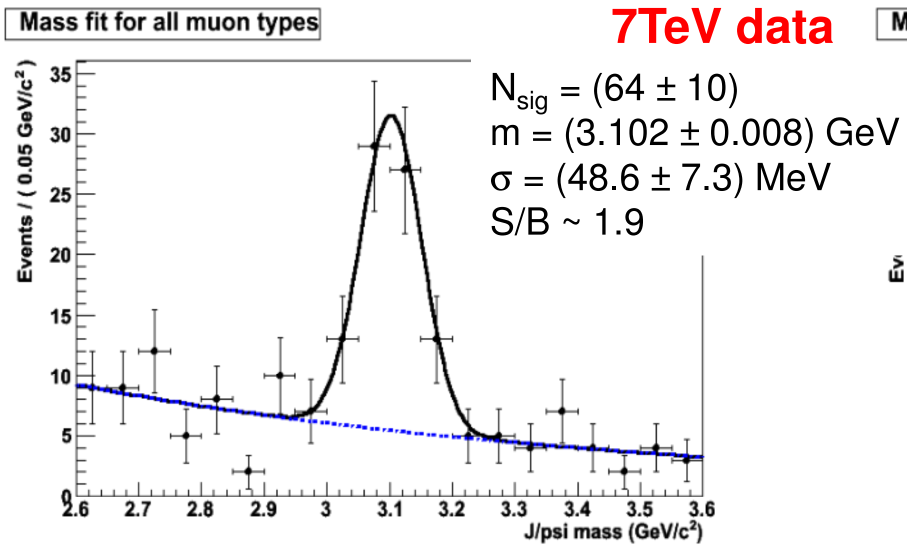

All di-muons

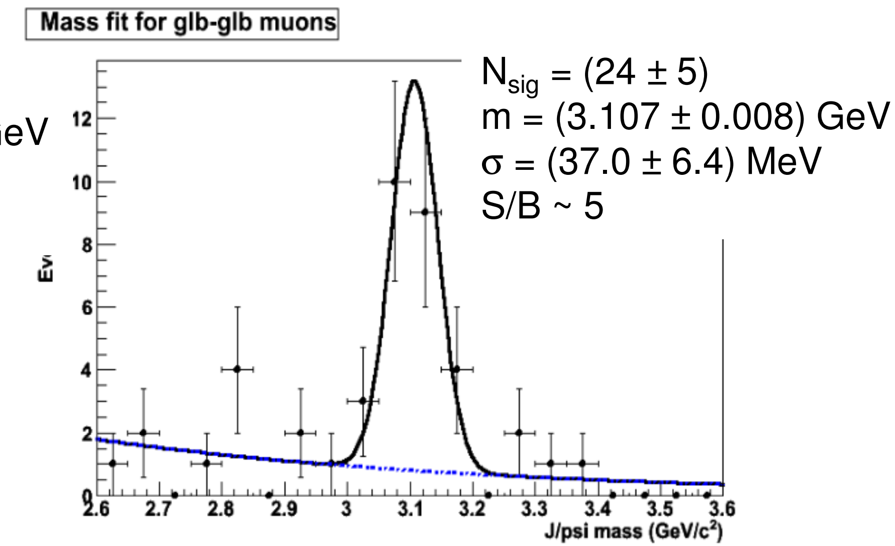

Only global-global muons