Fermilab

CMS Trigger System

Christian Herwig
July 10, 2023

Who am I?

Grew up in small town Massachusetts.

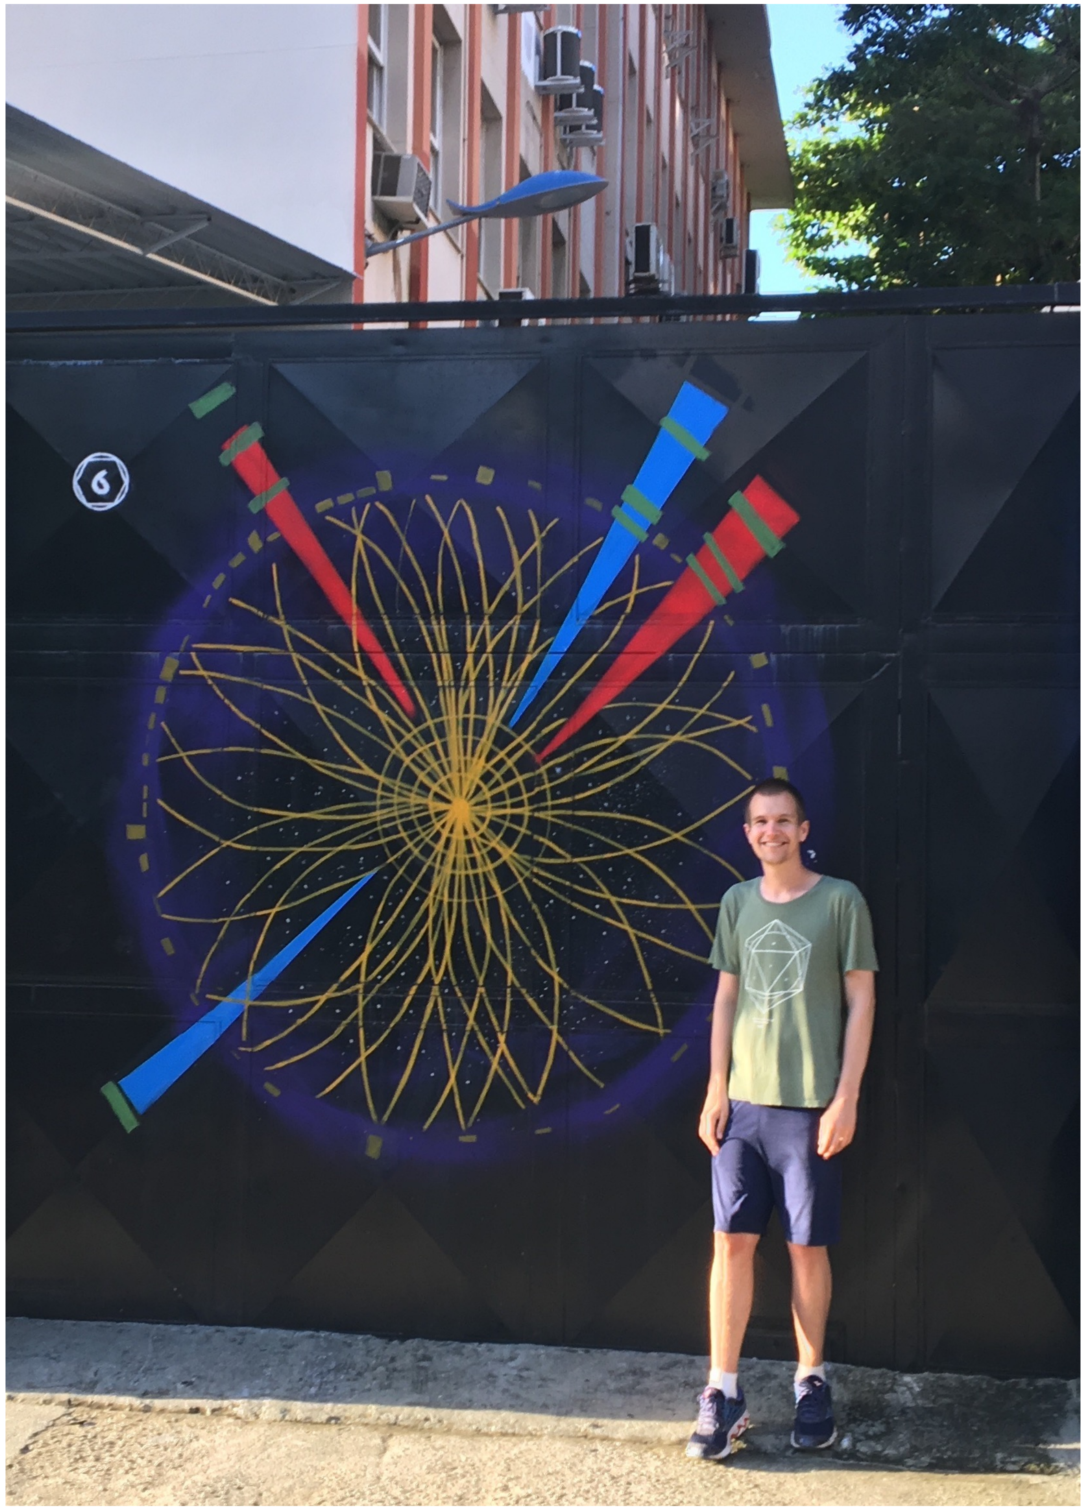

First came to Chicago for undergraduate, then found myself back again, @ Fermilab (Philadelphia↔Geneva→Chicago).

Main physics interests: Higgs (problems, solutions), Dark Matter, Trigger!

Can find me enjoying the outdoors while I can (hiking, swimming, cycling, sailing,…) or in a museum (aquarium) when not.

Why did we build the LHC?

Search for the Standard Model Higgs boson!

Tests the Electroweak theory at high precision

- Higher energies than before, large data-sets

Search for new particles & phenomena

- Explore the unknown: 2j, 2L resonances (Z’, gravitons, Higgs… TeV+)
    - Track record of discovery in “2X” final states
- Search for Dark Matter
    - WIMP miracle → DM with W/Z/h-like masses, couplings
- Heavy cousins of the top quark?
    - SM should have new “top-like” particles to solve hierarchy problems
- …and many more that address other deep questions
    - CP violation, matter/antimatter asymmetry, small neutrino masses,…

Target cross sections are small…

- At 13 TeV, the total cross section is σ_{pp} = 100 milli-barn
- Our target processes are many ***orders of magnitude*** more rare
    - Higgs boson: 49 pico-barn

      This is ($\sigma_{pp} / 2 \cdot 10^9$) !
    - Higgsino (150 GeV): 3.8 pb
    - Top squark (500 GeV): 0.6 pb
- Even the “most common” processes occur only in a fraction of events
    - Single-W production: 1 in 10^6
    - High-p_T jet (100 GeV): 1 in 10^5
- Most events: “soft, inelastic QCD”

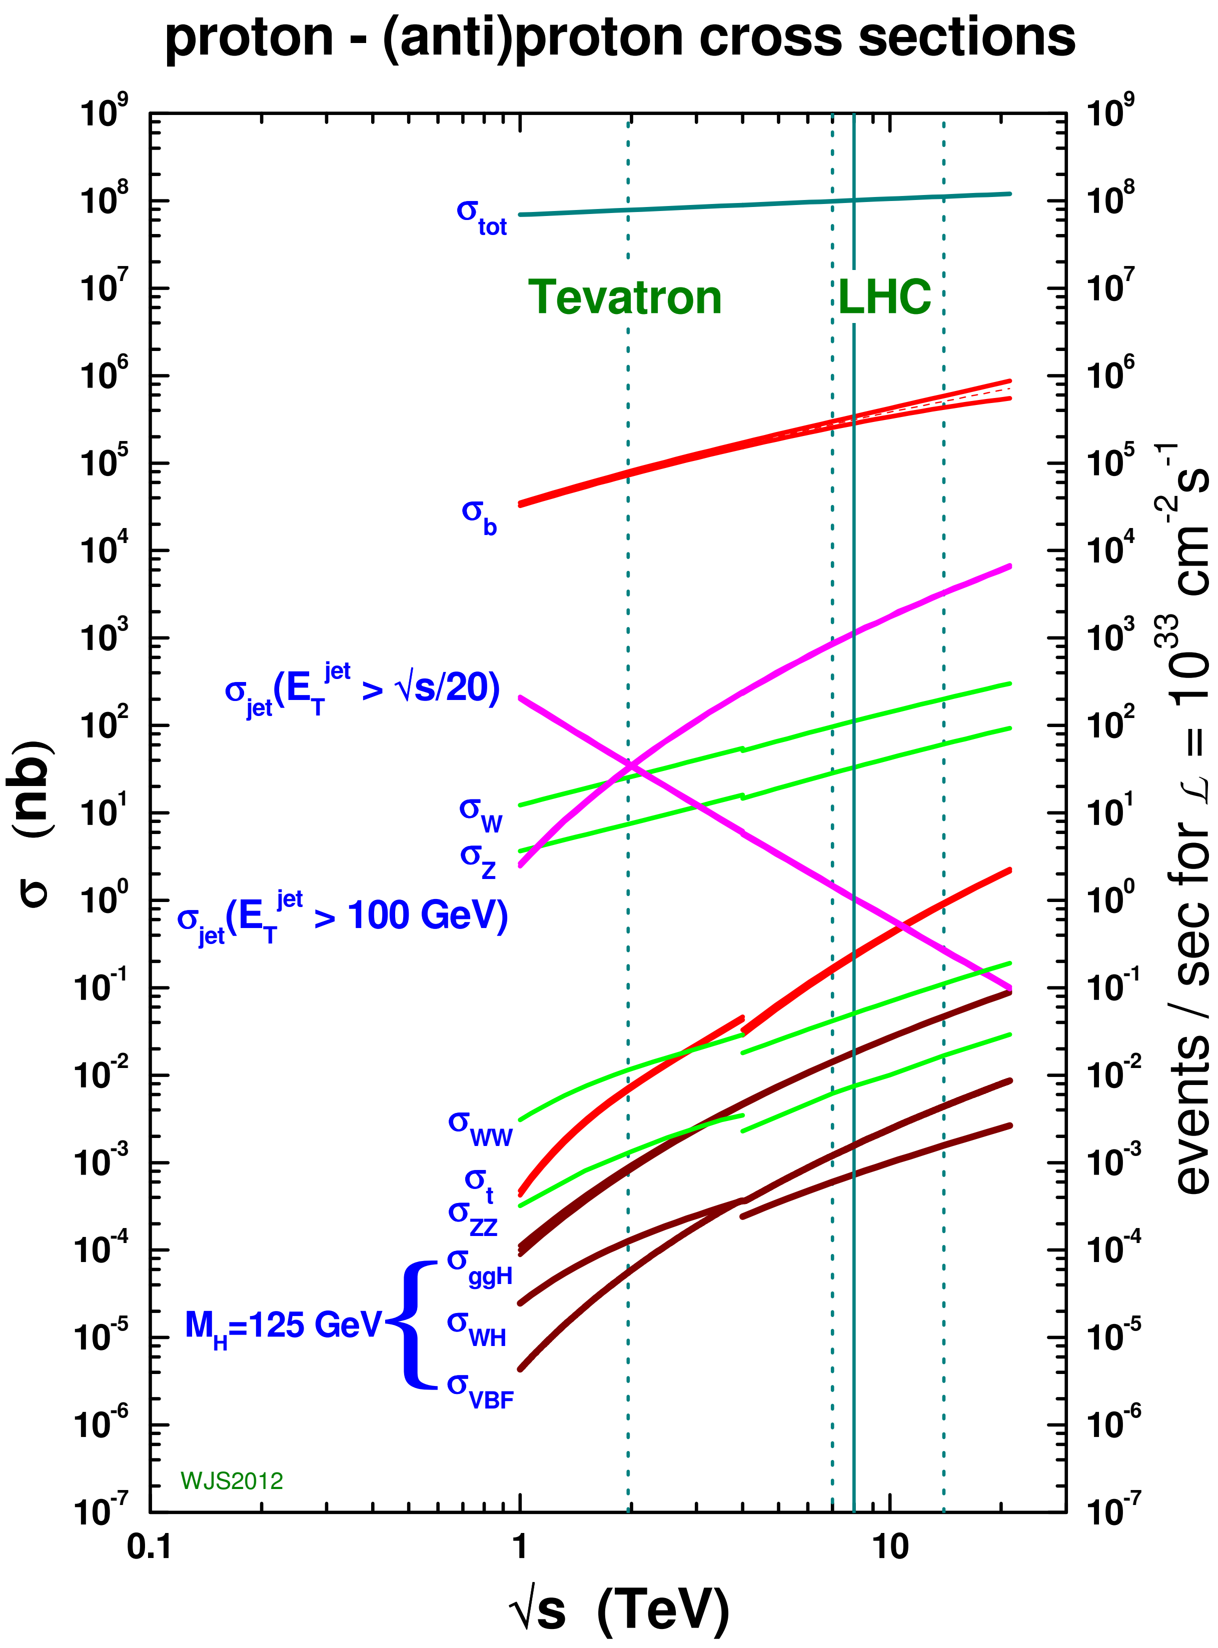

… so data sets must be large

Discovering rare phenomena requires many, many pp collisions.

- In 2012, Higgs discovery needed ~10/fb of 8 TeV data.
  - Produced 200k Higgs Bosons (Only 0.2% to ɣɣ, 0.01% to 4L).
  - This corresponds to 10^{15} total *pp* interactions!

The CMS detector takes high-fidelity images of each event, 40 million / sec.

- Each image is ~1 MegaByte → we generate 40 TeraBytes / second.
- Reconstructing each event takes ~1 second.

But reconstructing and storing 10^{15} collisions would take:

- 1 ZettaByte of data (Recall: Zetta > Exa > Peta > Tera > Giga > …)
- 32 Million years

How does CMS solve this problem? → the thrust of my talk today.

The Large Hadron Collider

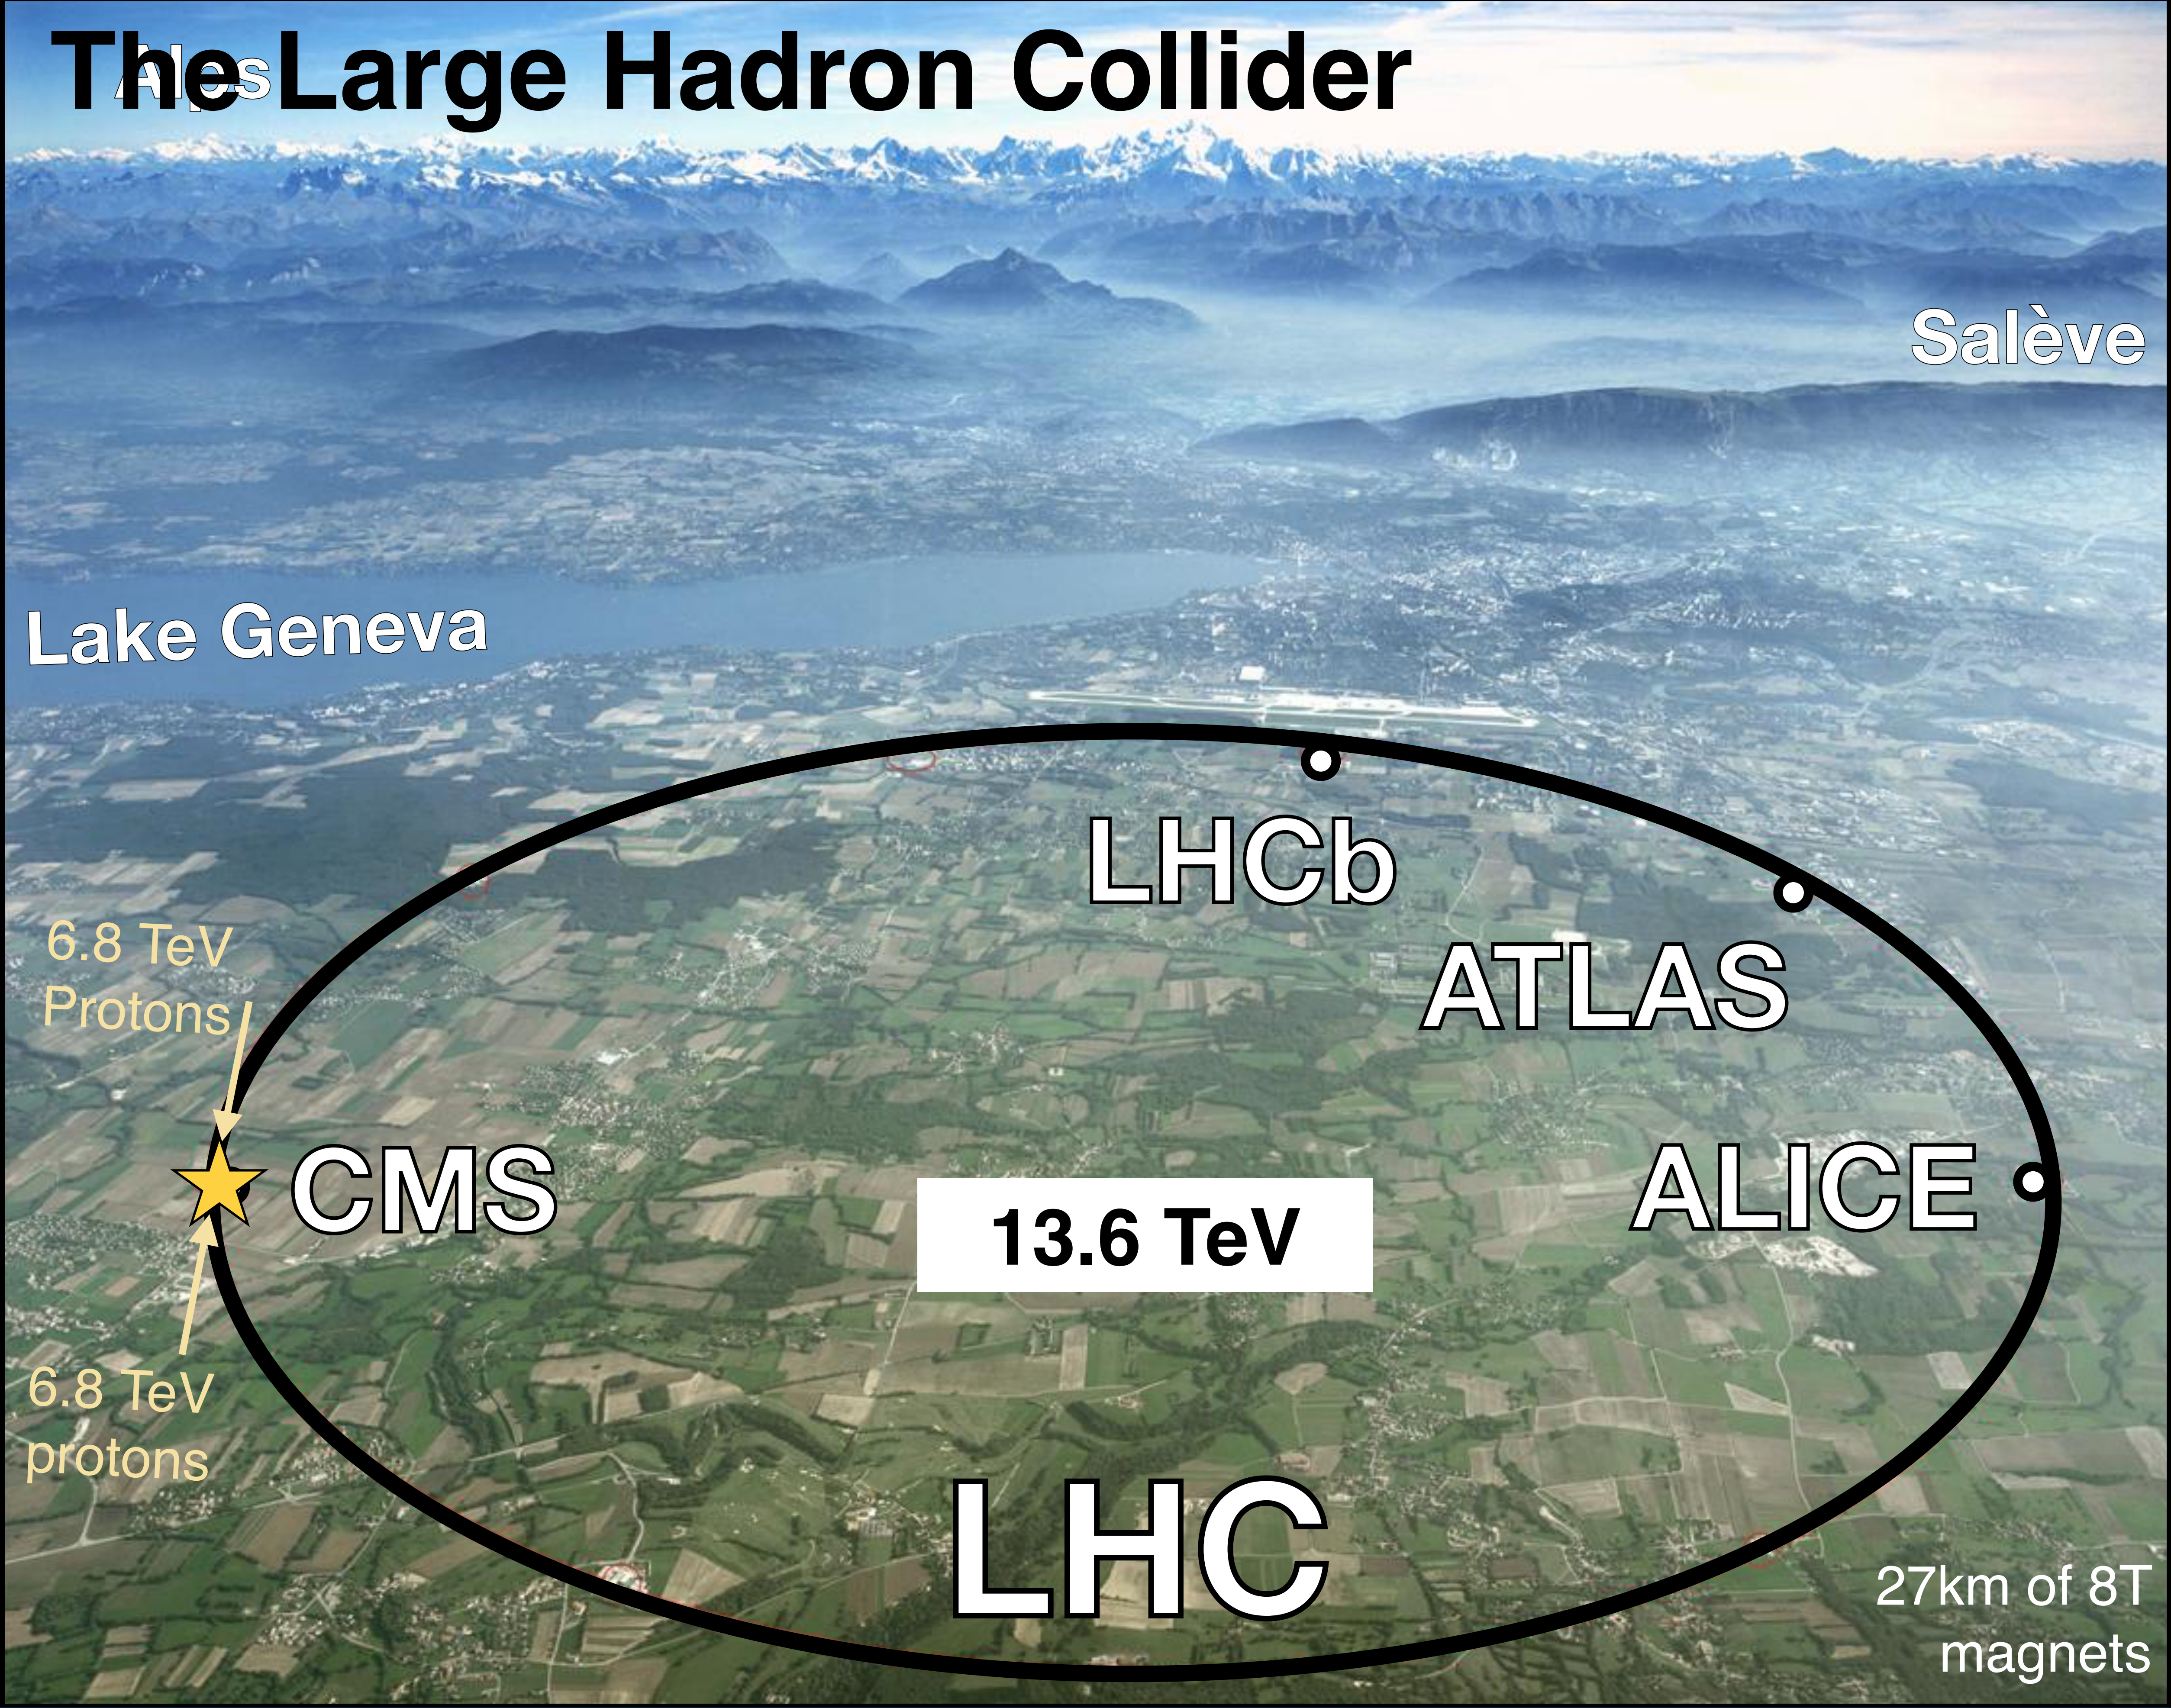

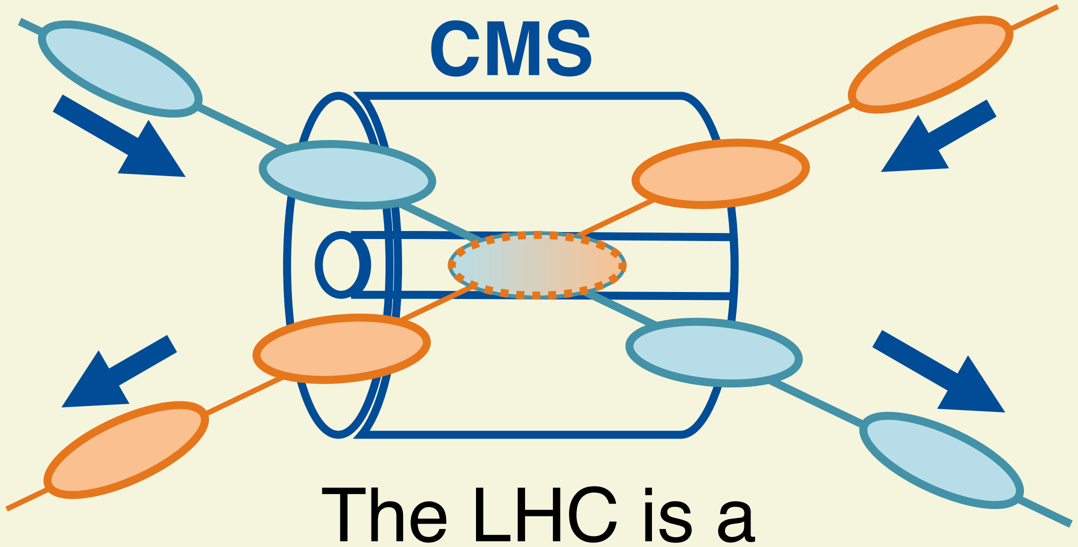

The LHC is a
proton **bunch** collider

Bunches collide every 25ns
this is *"one event"*.

Current average is about
~60 *pp* collisions / bunch

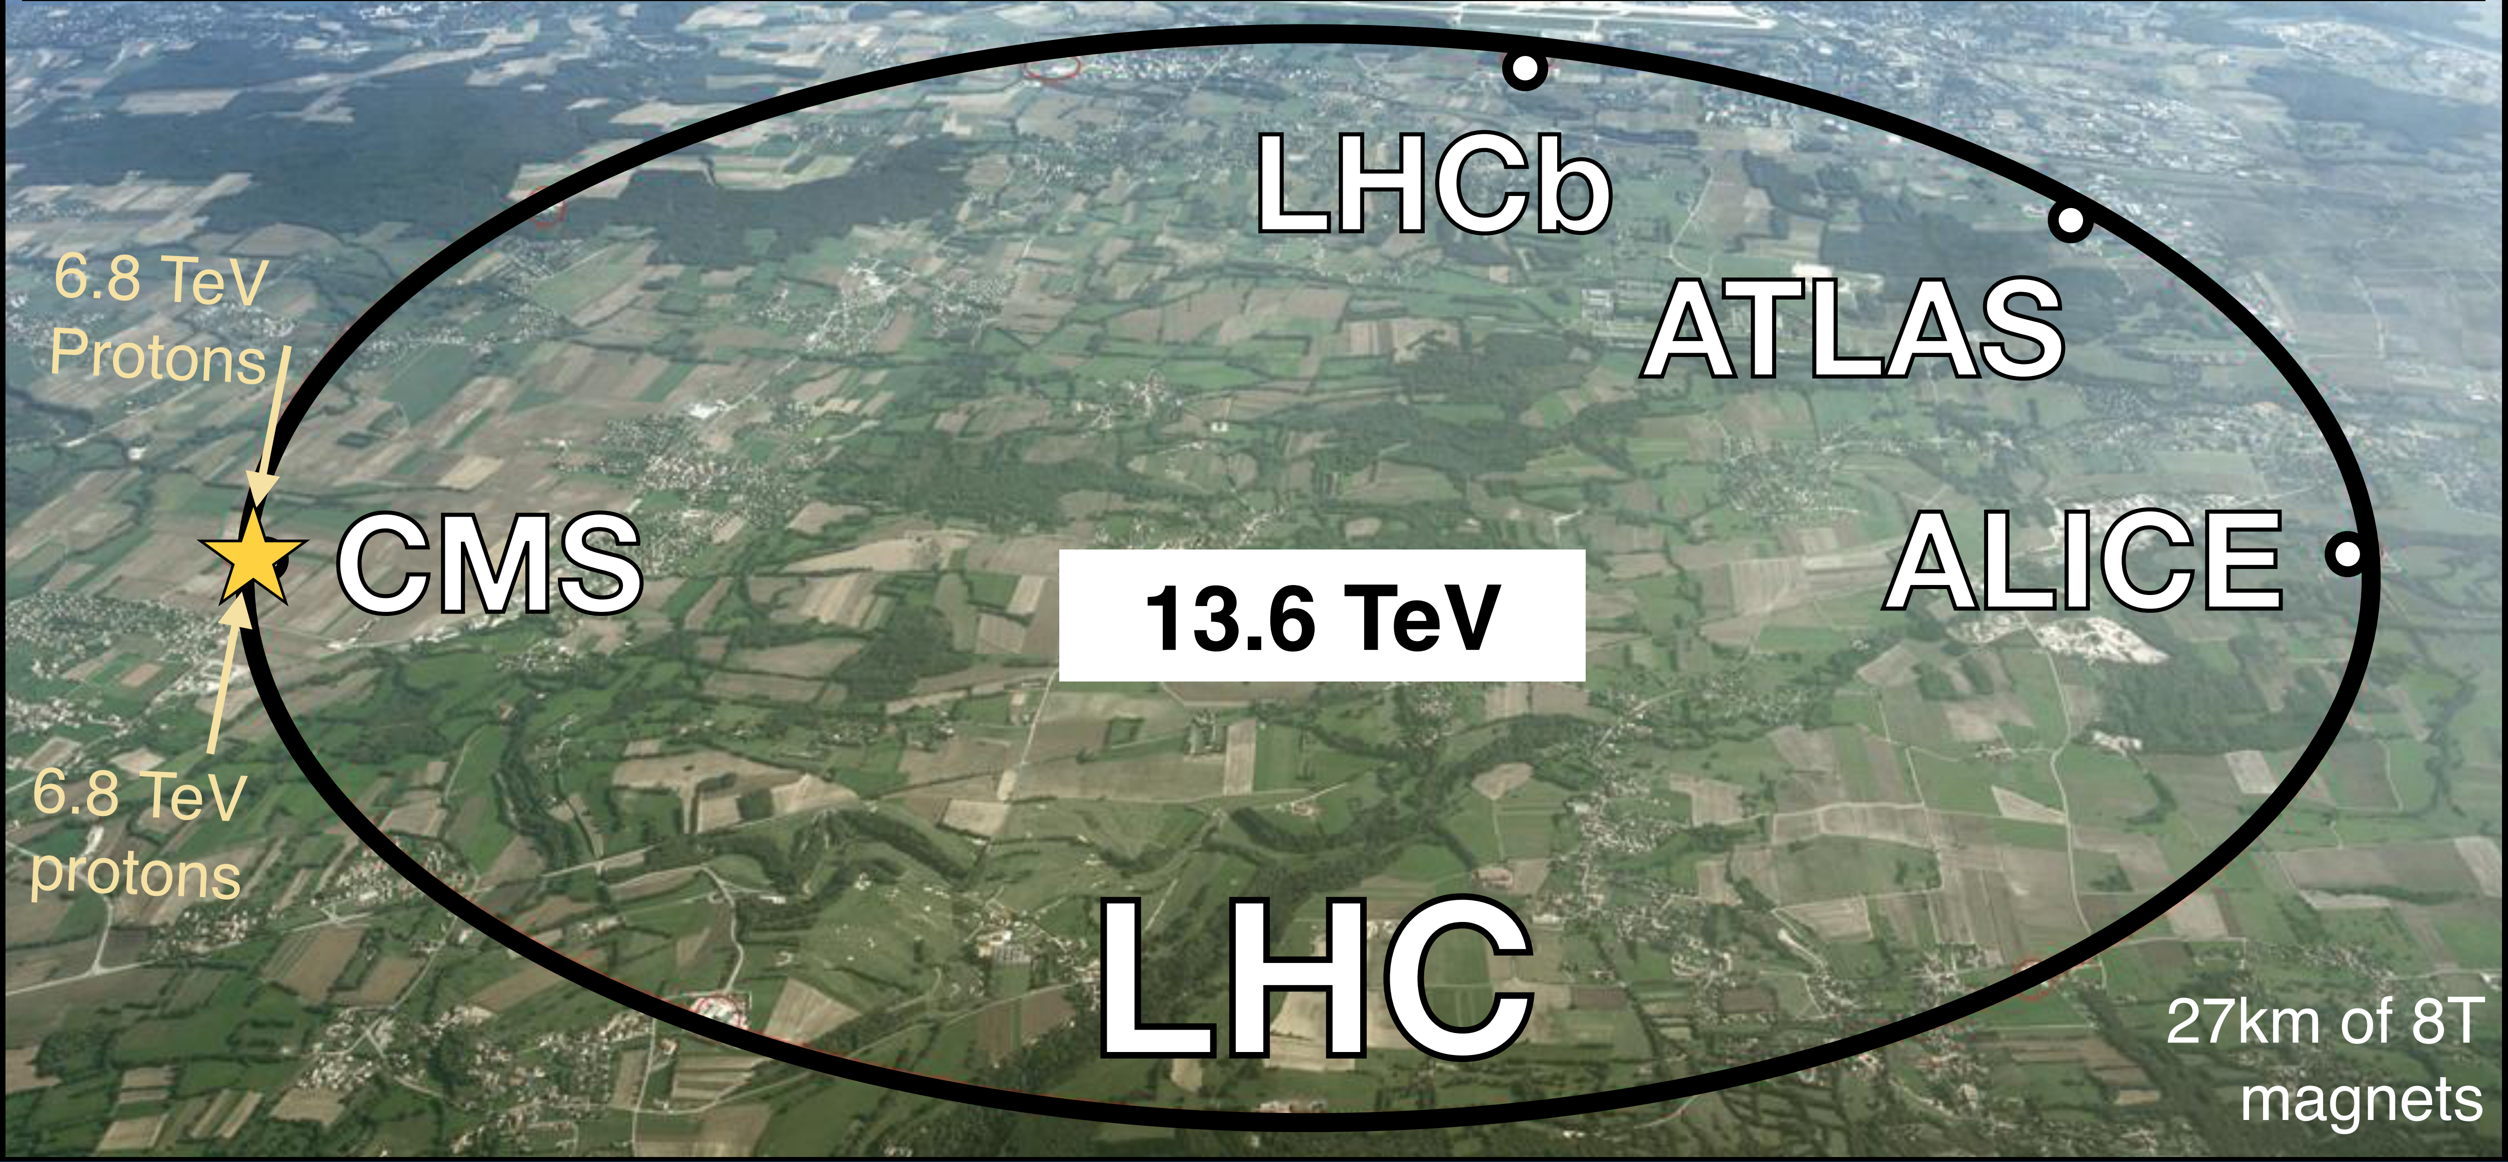

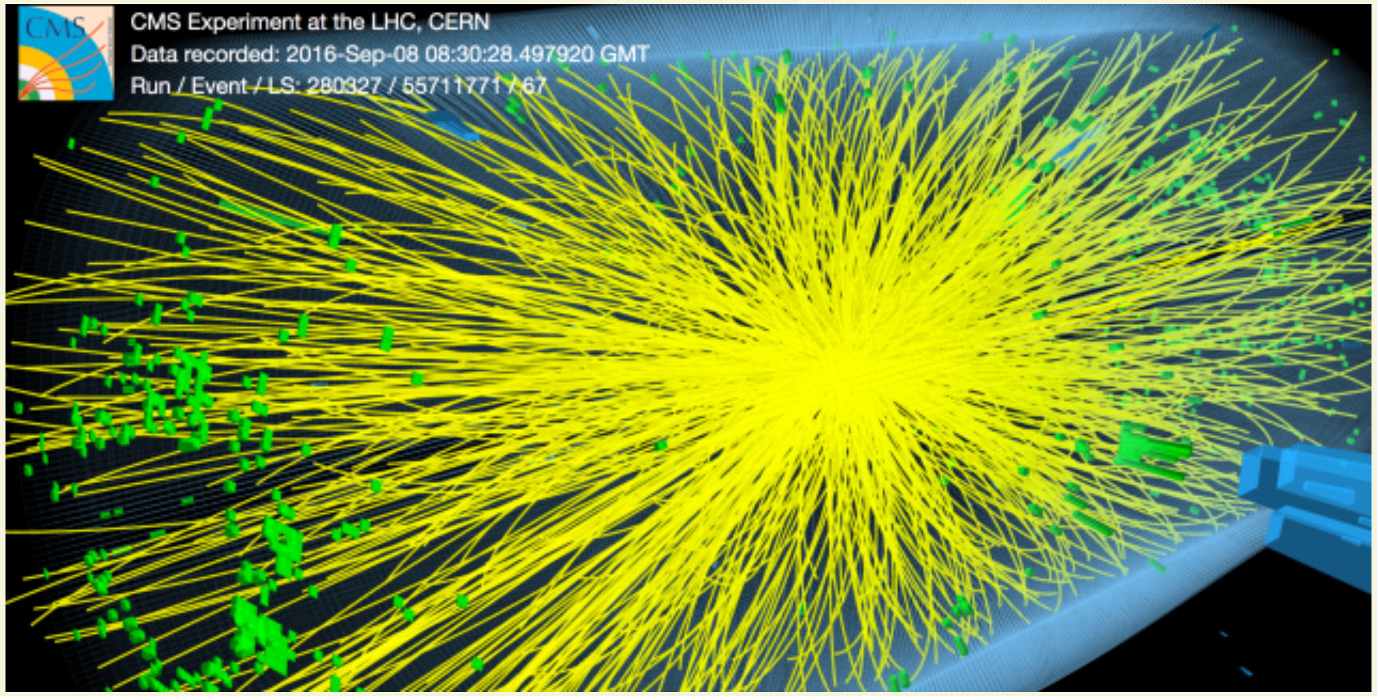

Bunches collide every 25ns this is *"one event"*.

Current average is about ~60 *pp* collisions / bunch

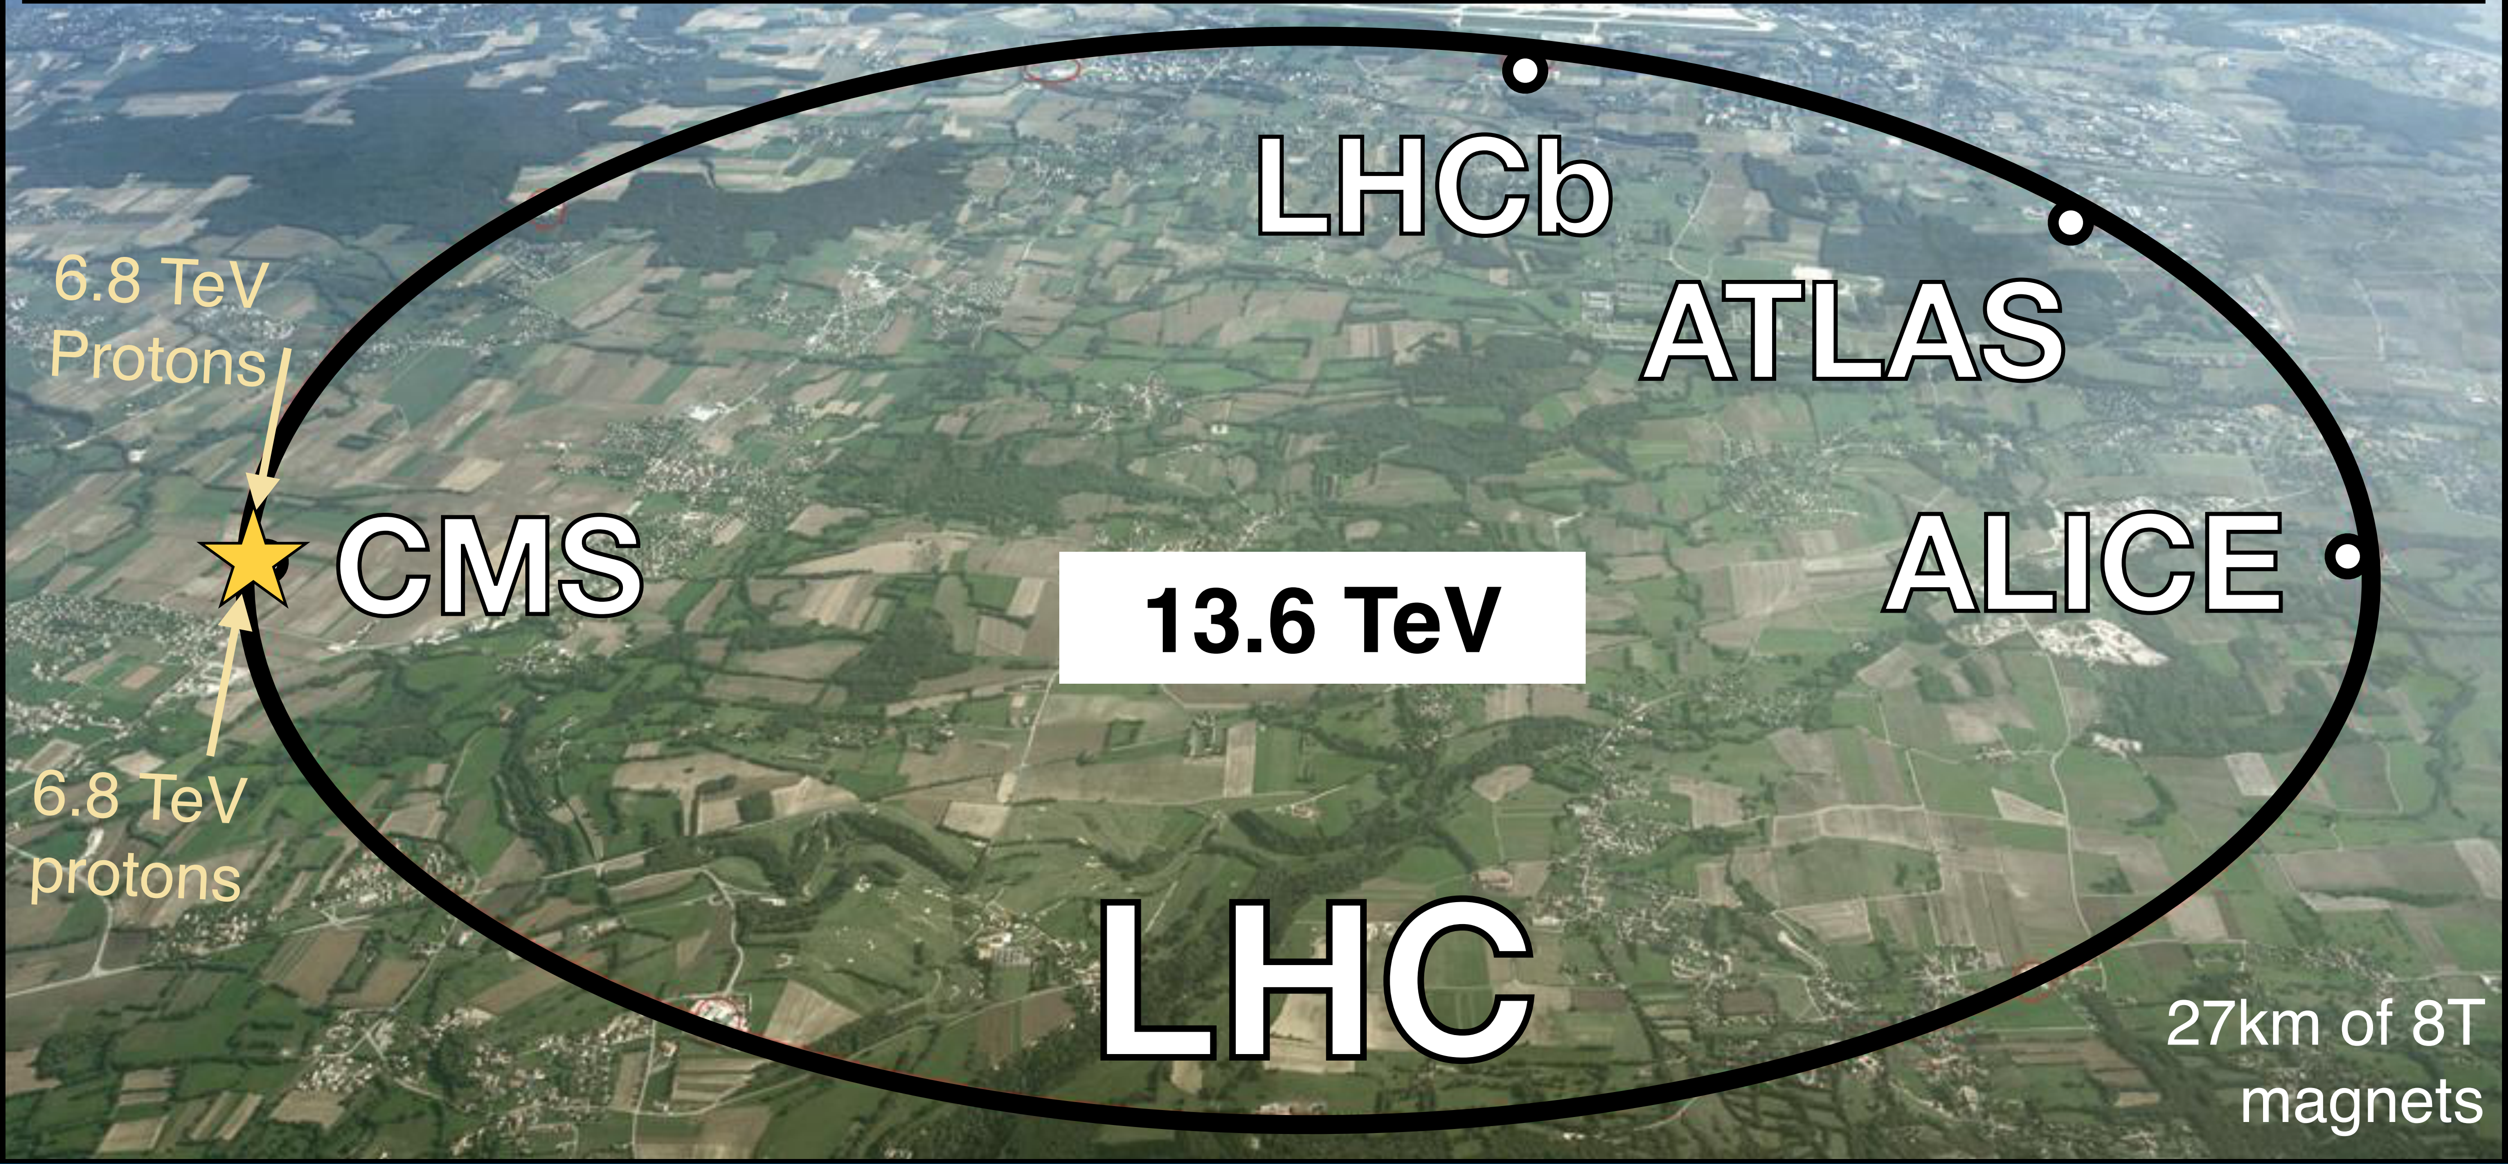

27km of 8T magnets

High-pileup data-taking

- Packing each bunch crossing full of more proton-proton collisions is an excellent way to accumulate data faster!
  - 60x collisions / event → 60x fewer events for a given dataset size
  - But are events “60x more complicated”? Yes and no…

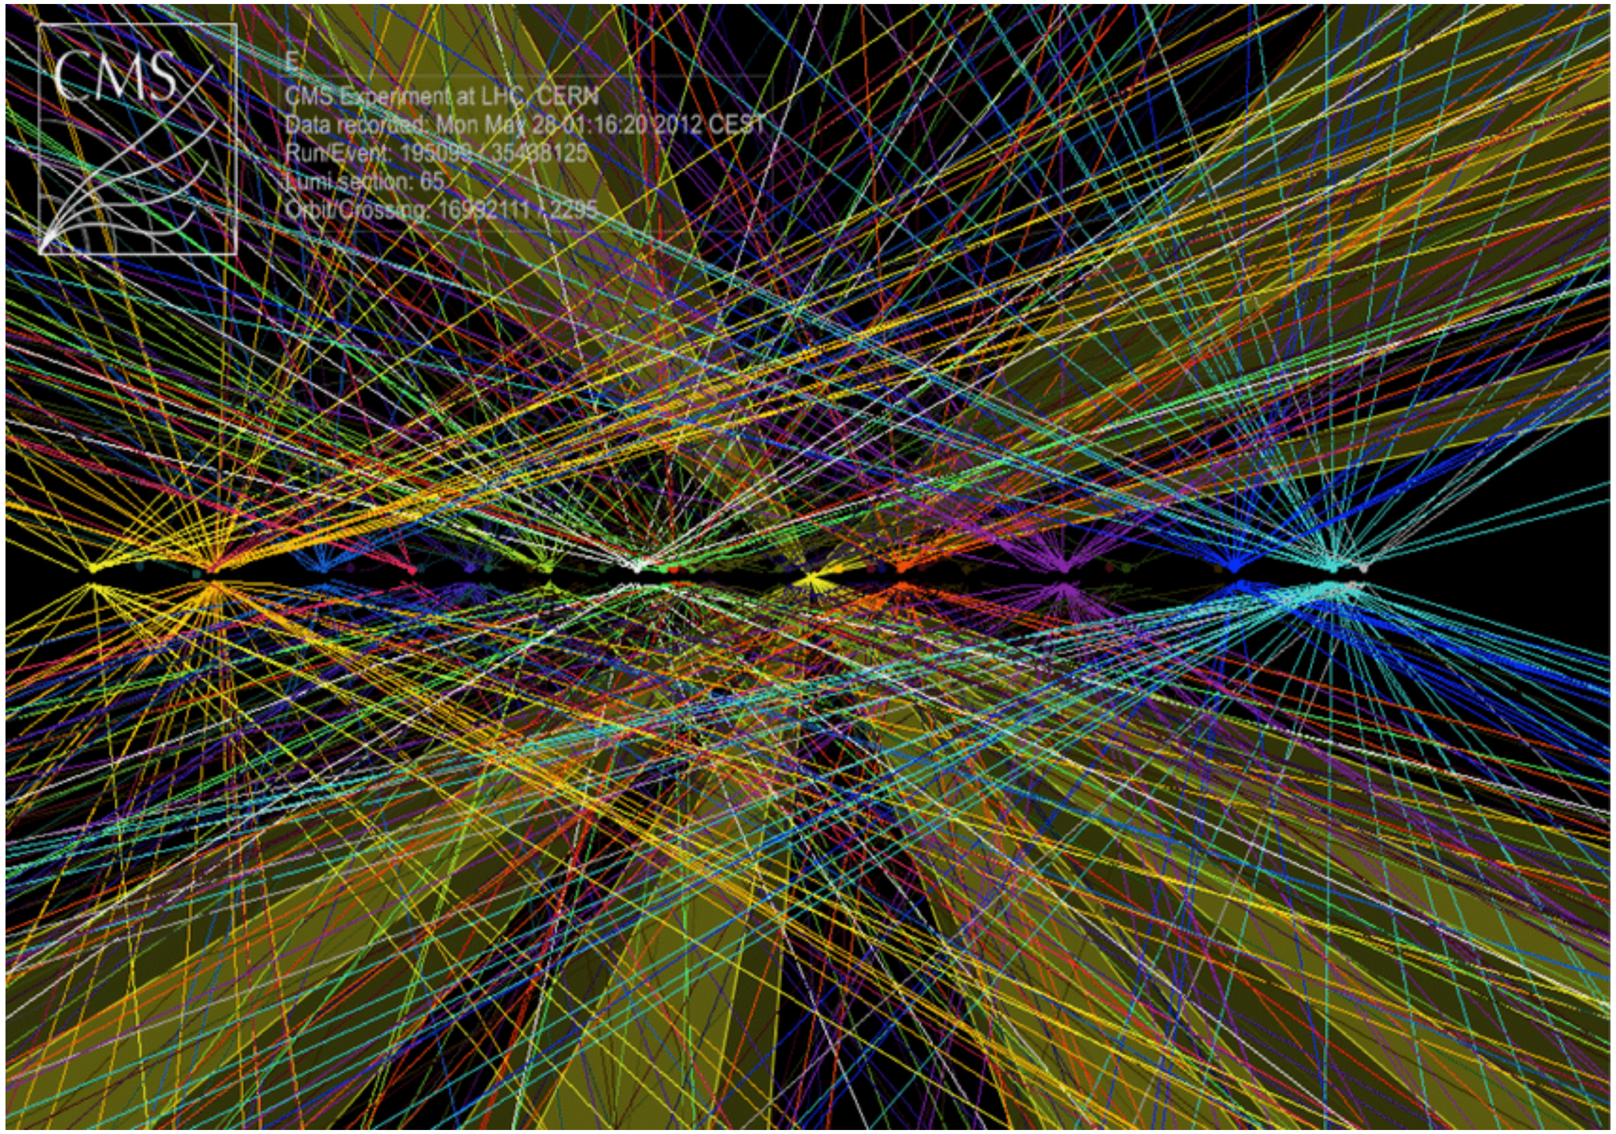

Challenge: can you tell the interactions apart?

- Charged particles: Silicon trackers provide excellent vertexing capability.
- Neutral particles are more difficult, but can be done “on average”.

But the tracking step is most computationally expensive!

Typical LHC collision: di-jets

- There are many, many ways to produce pairs of jets in LHC collisions
  - 8 gluons, and 5 x 3 light quarks!

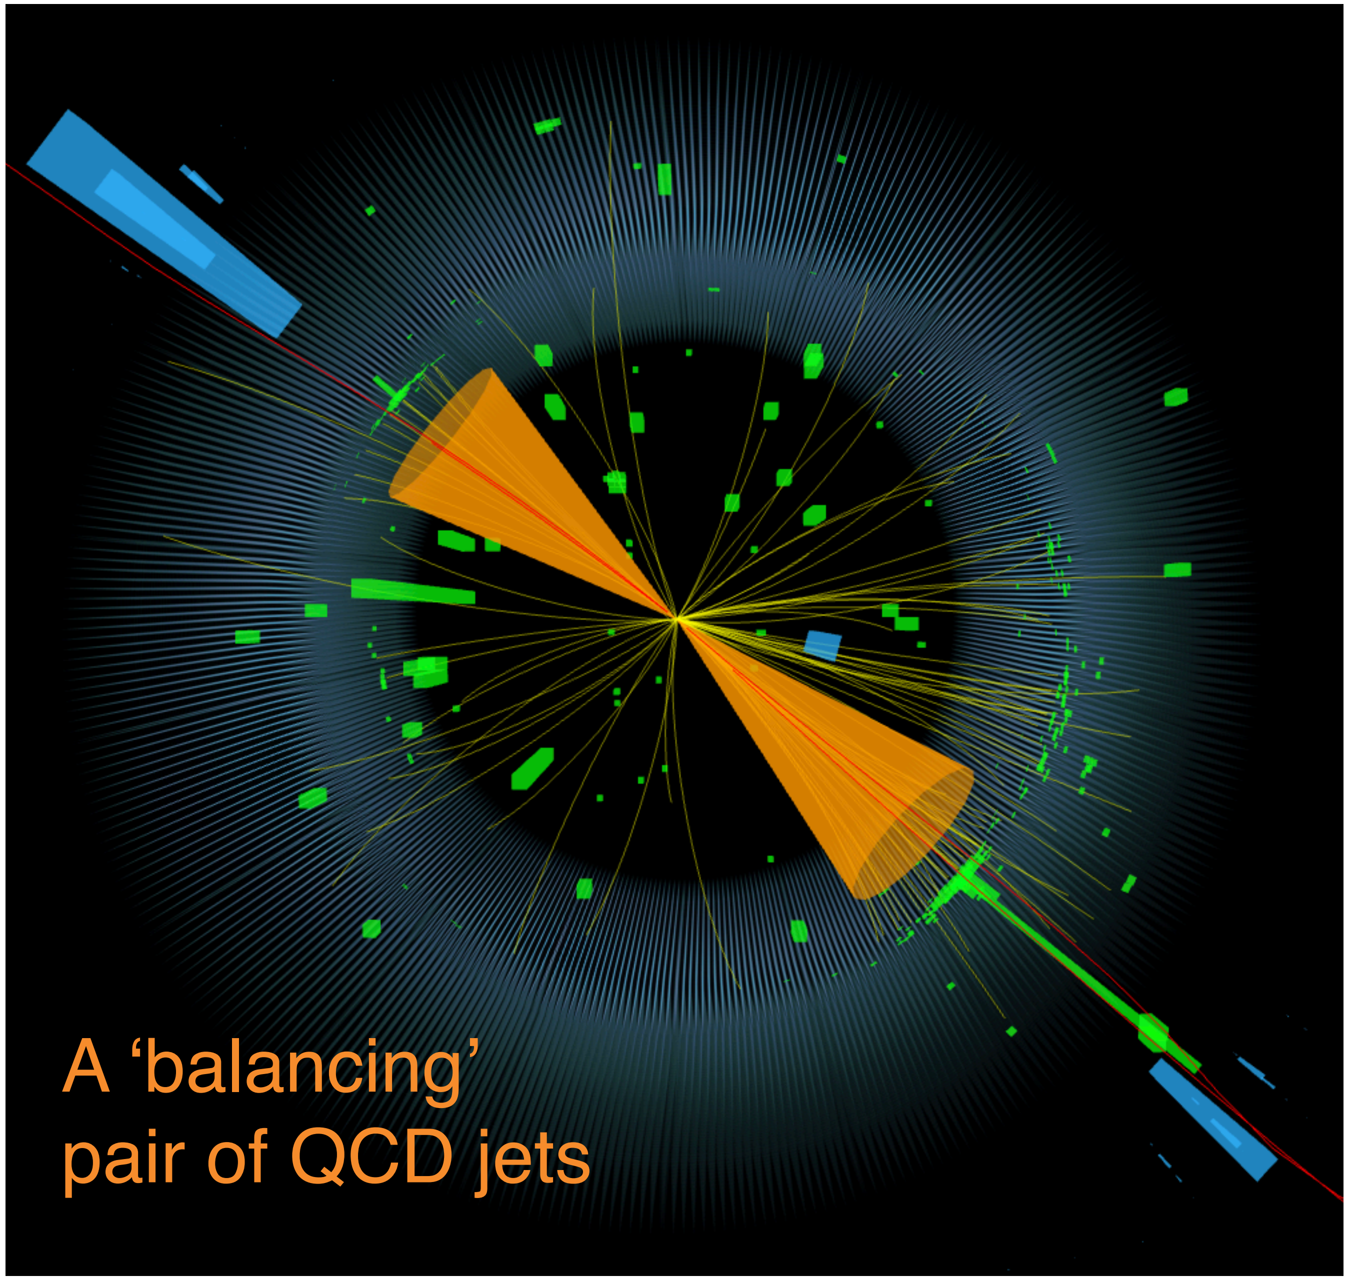

A ‘balancing’
pair of QCD jets

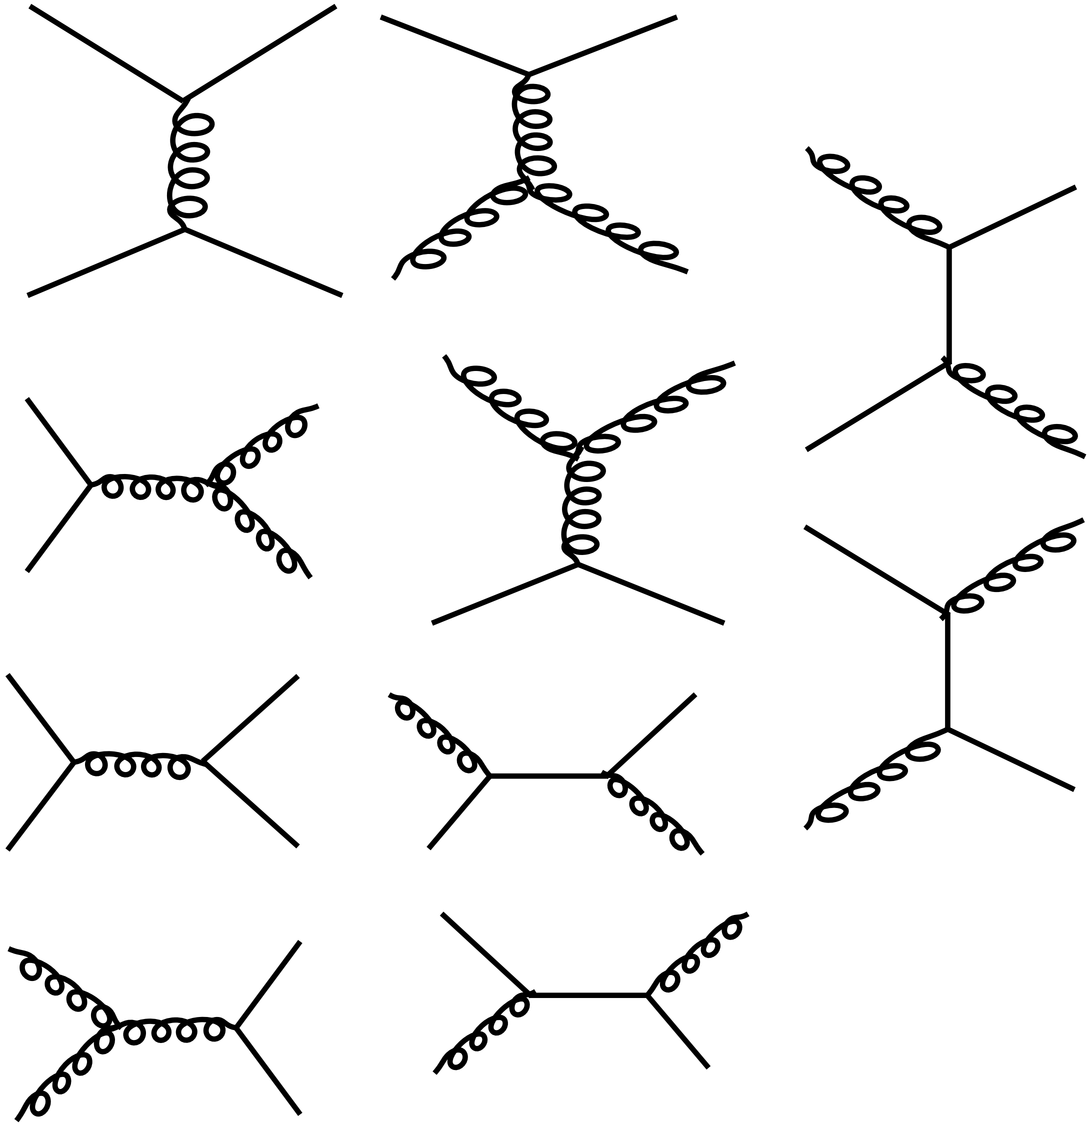

Interesting collisions: Higgs(→ZZ)

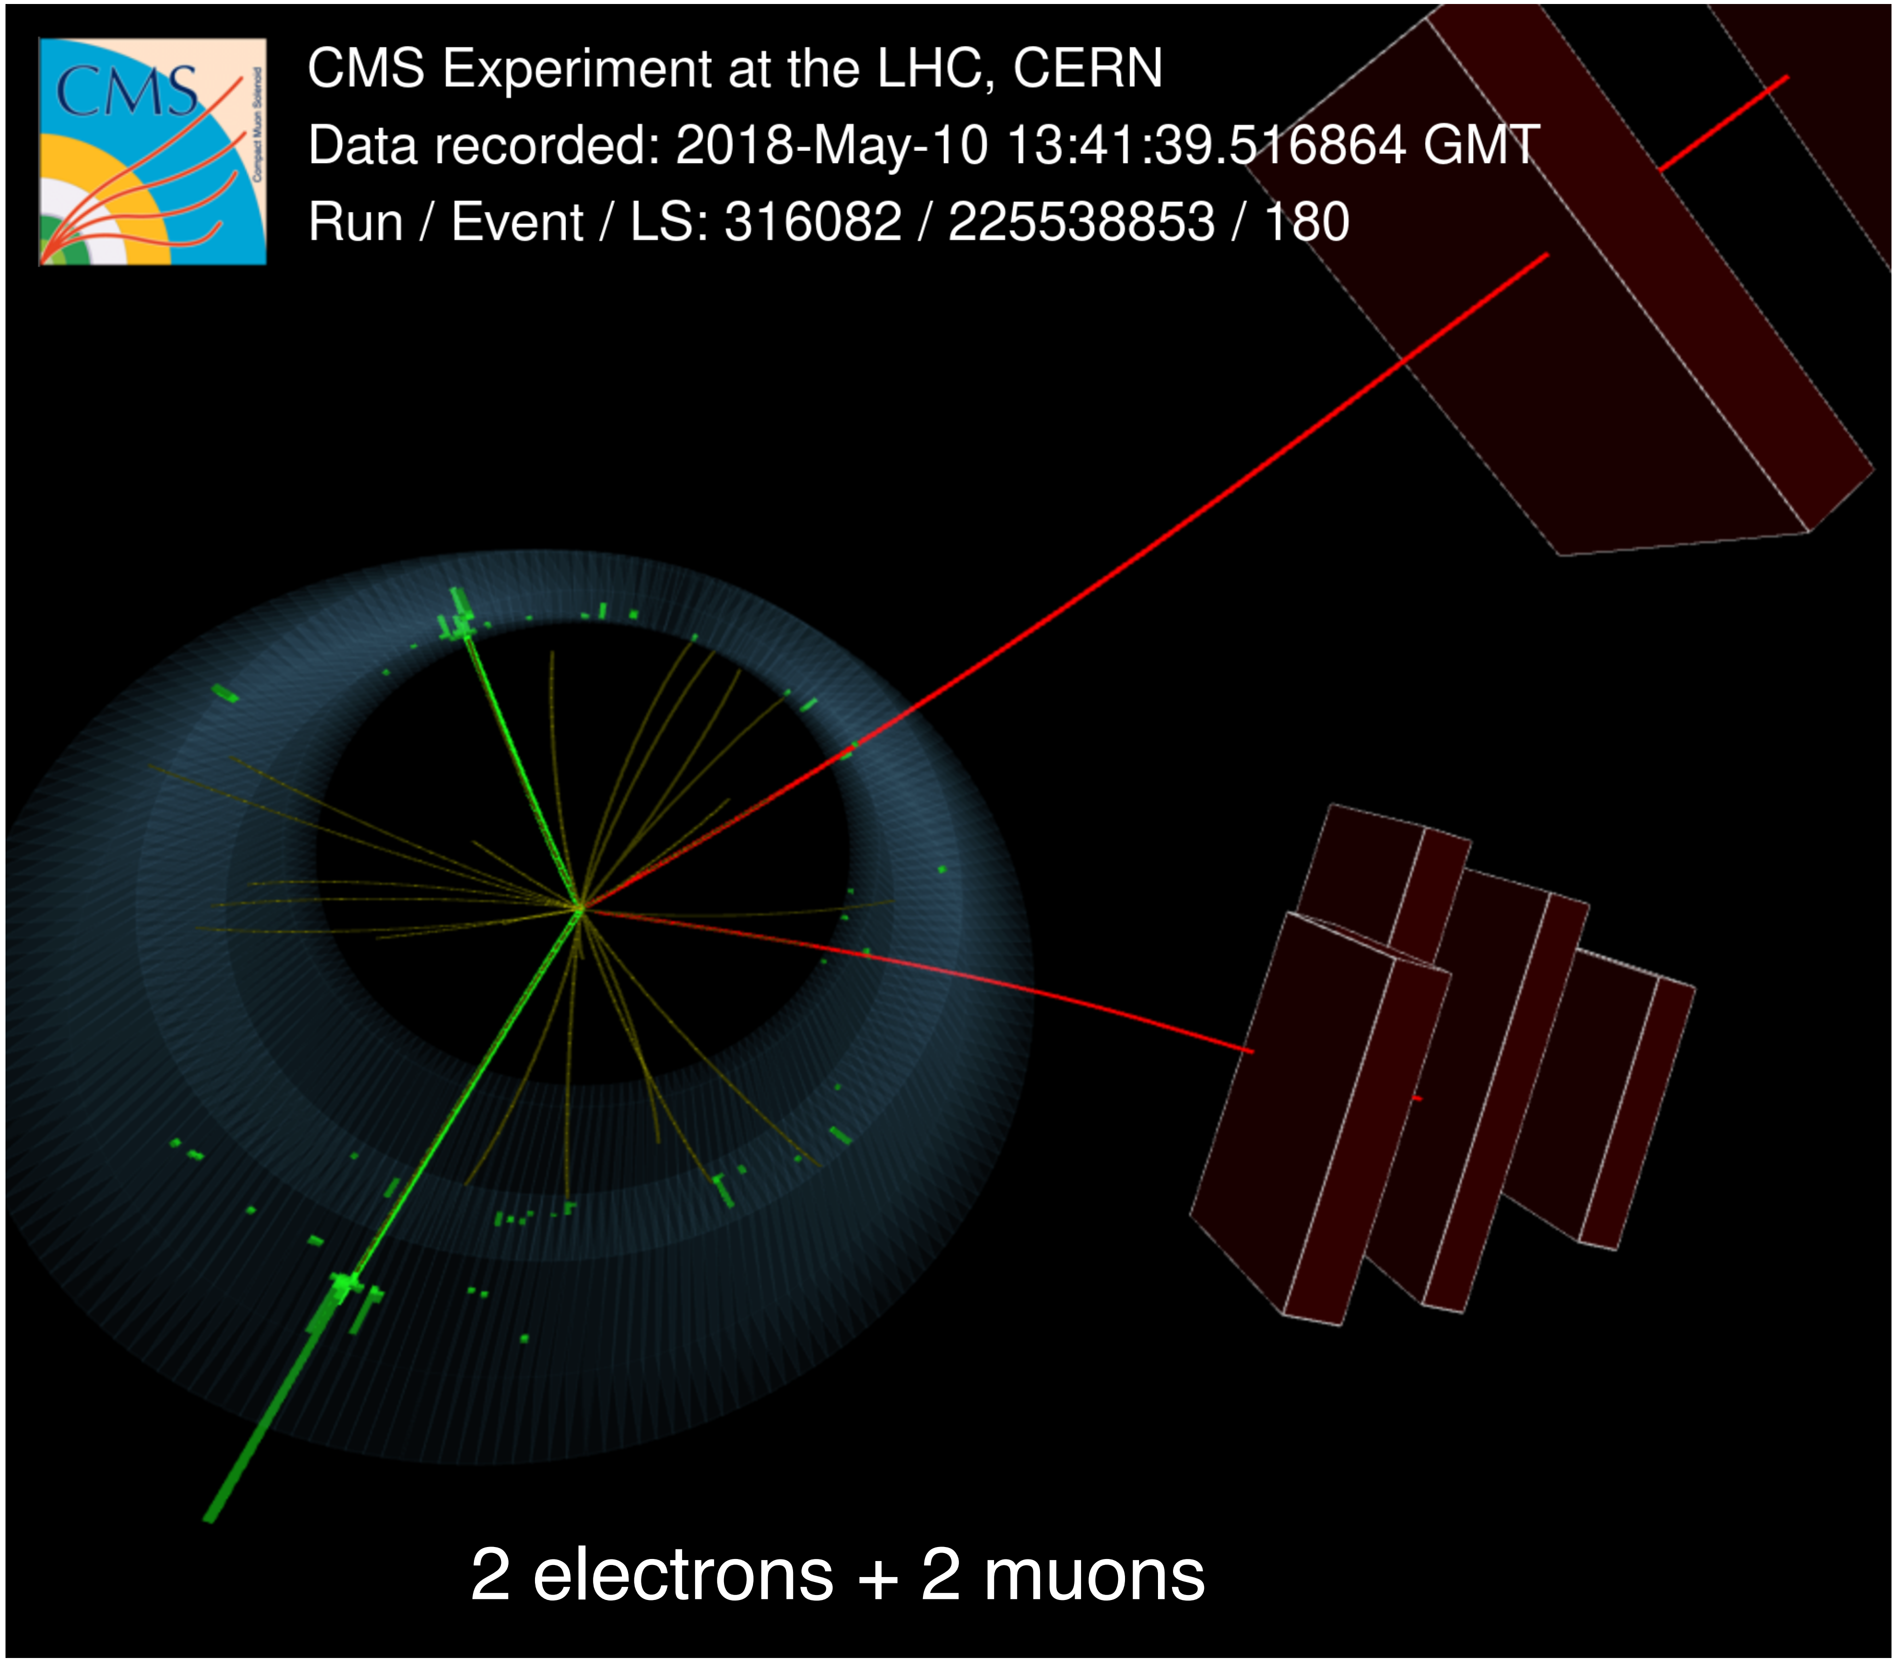

2 electrons + 2 muons

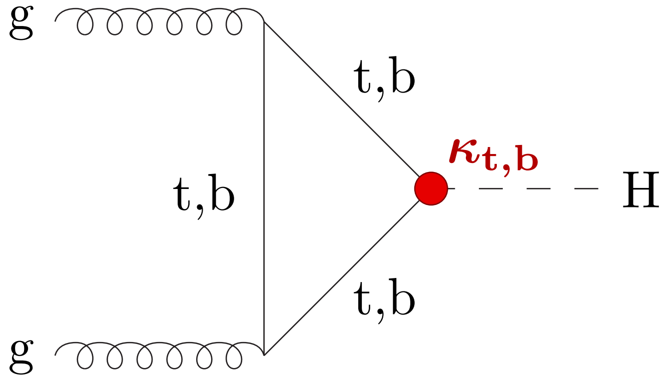

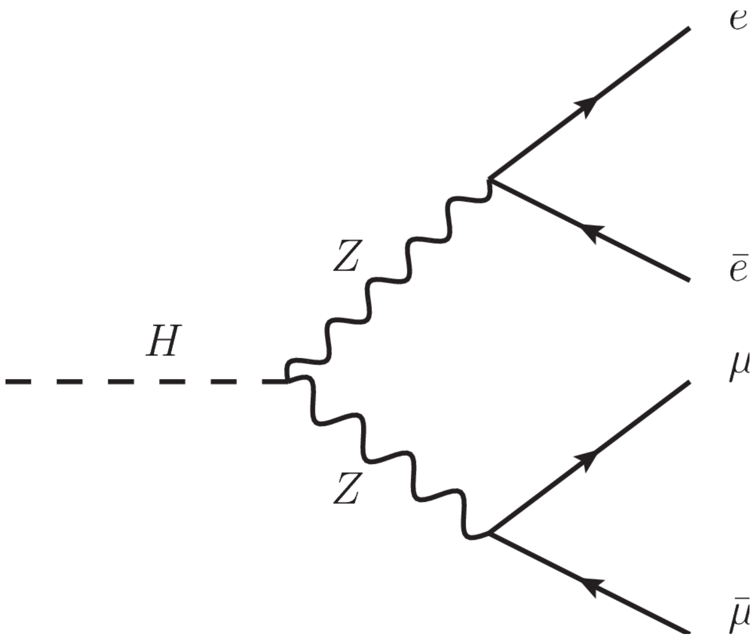

Interesting collisions: Higgs(→bb)

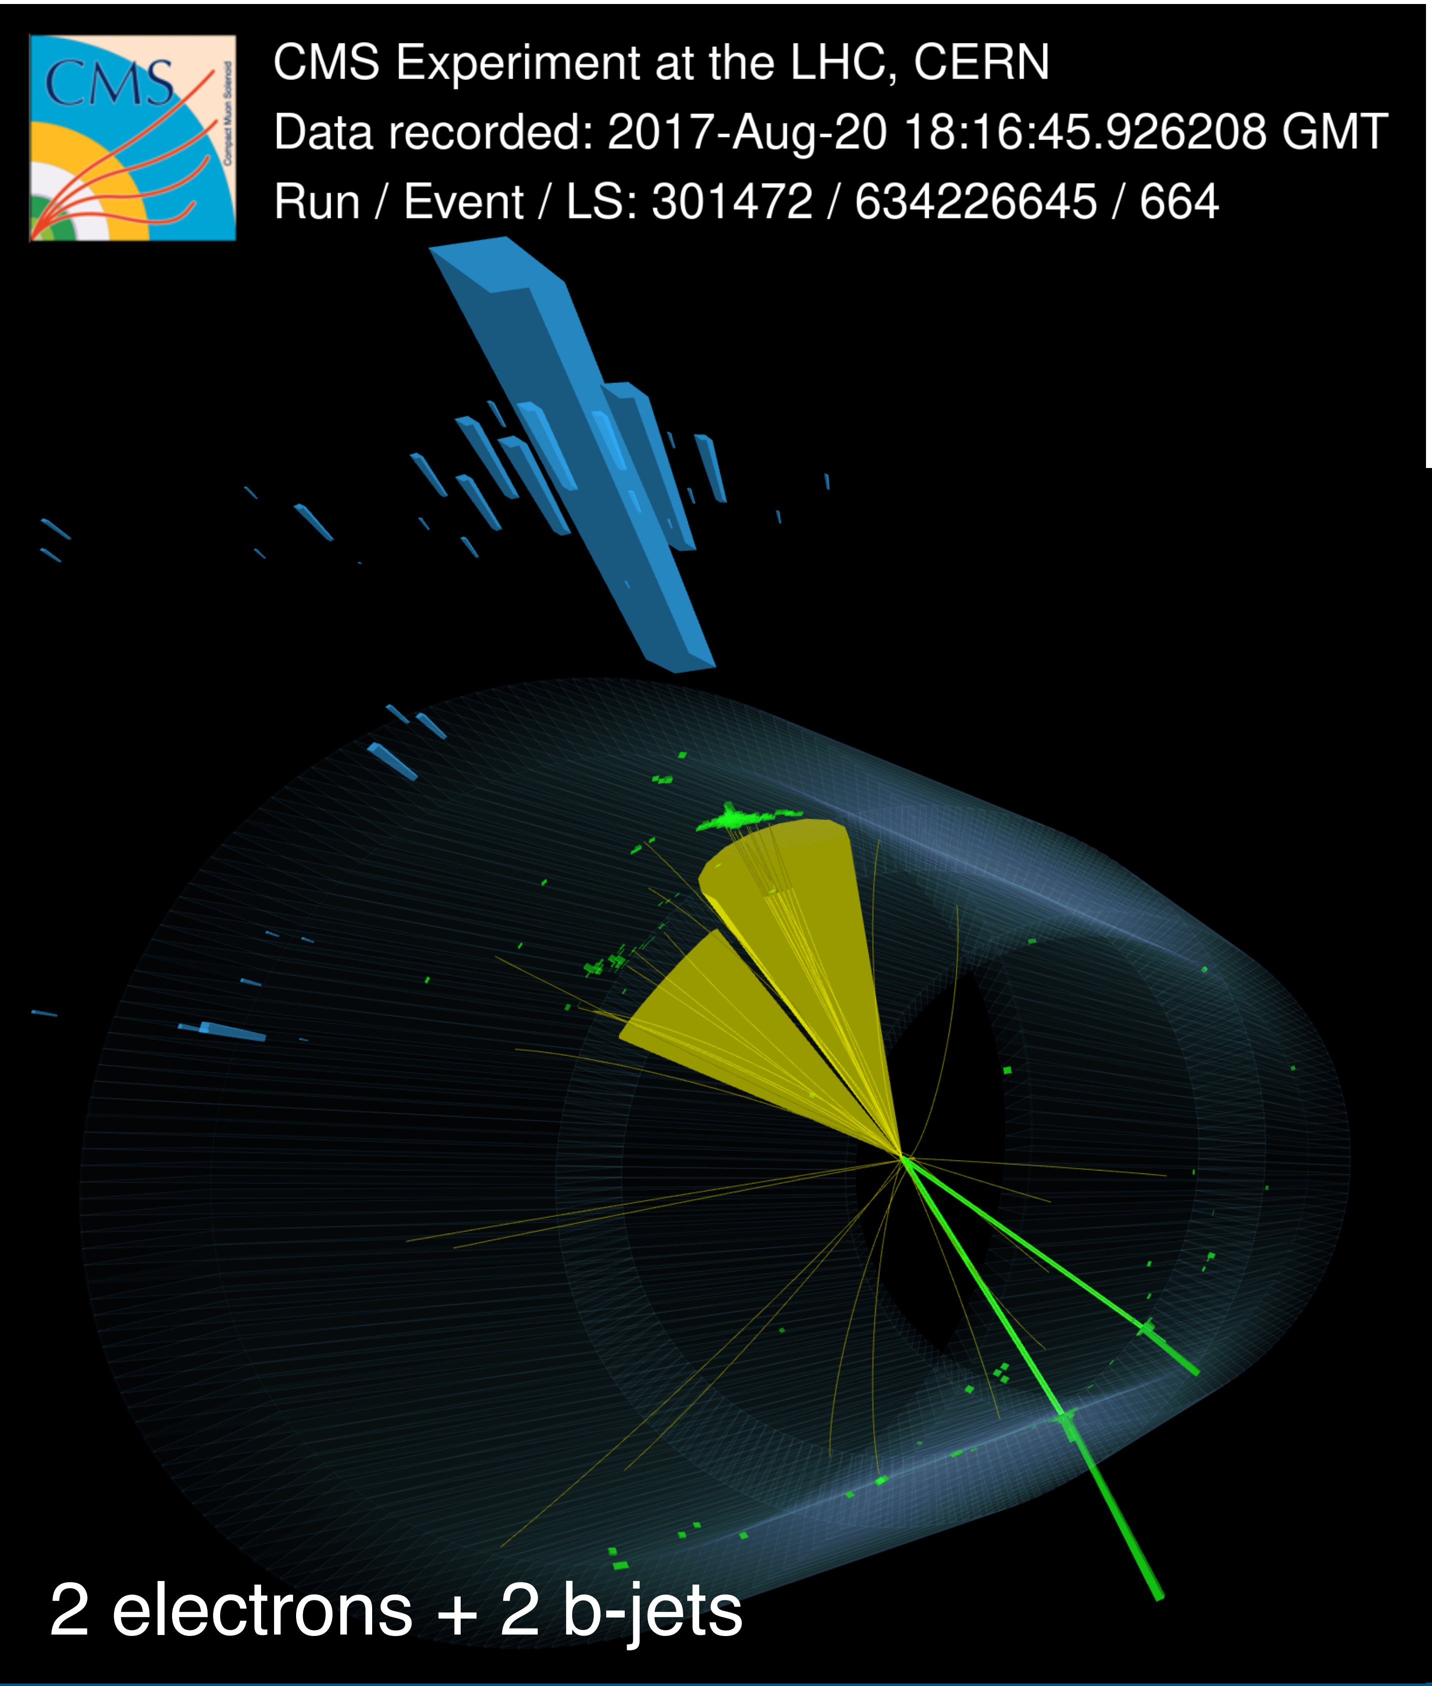
CMS Experiment at the LHC, CERN
Data recorded: 2017-Aug-20 18:16:45.926208 GMT
Run / Event / LS: 301472 / 634226645 / 664

2 electrons + 2 b-jets

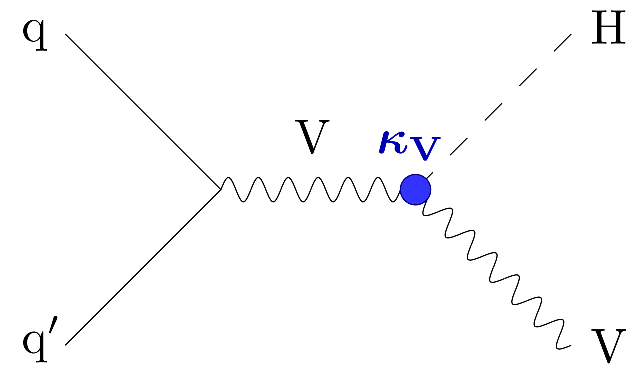

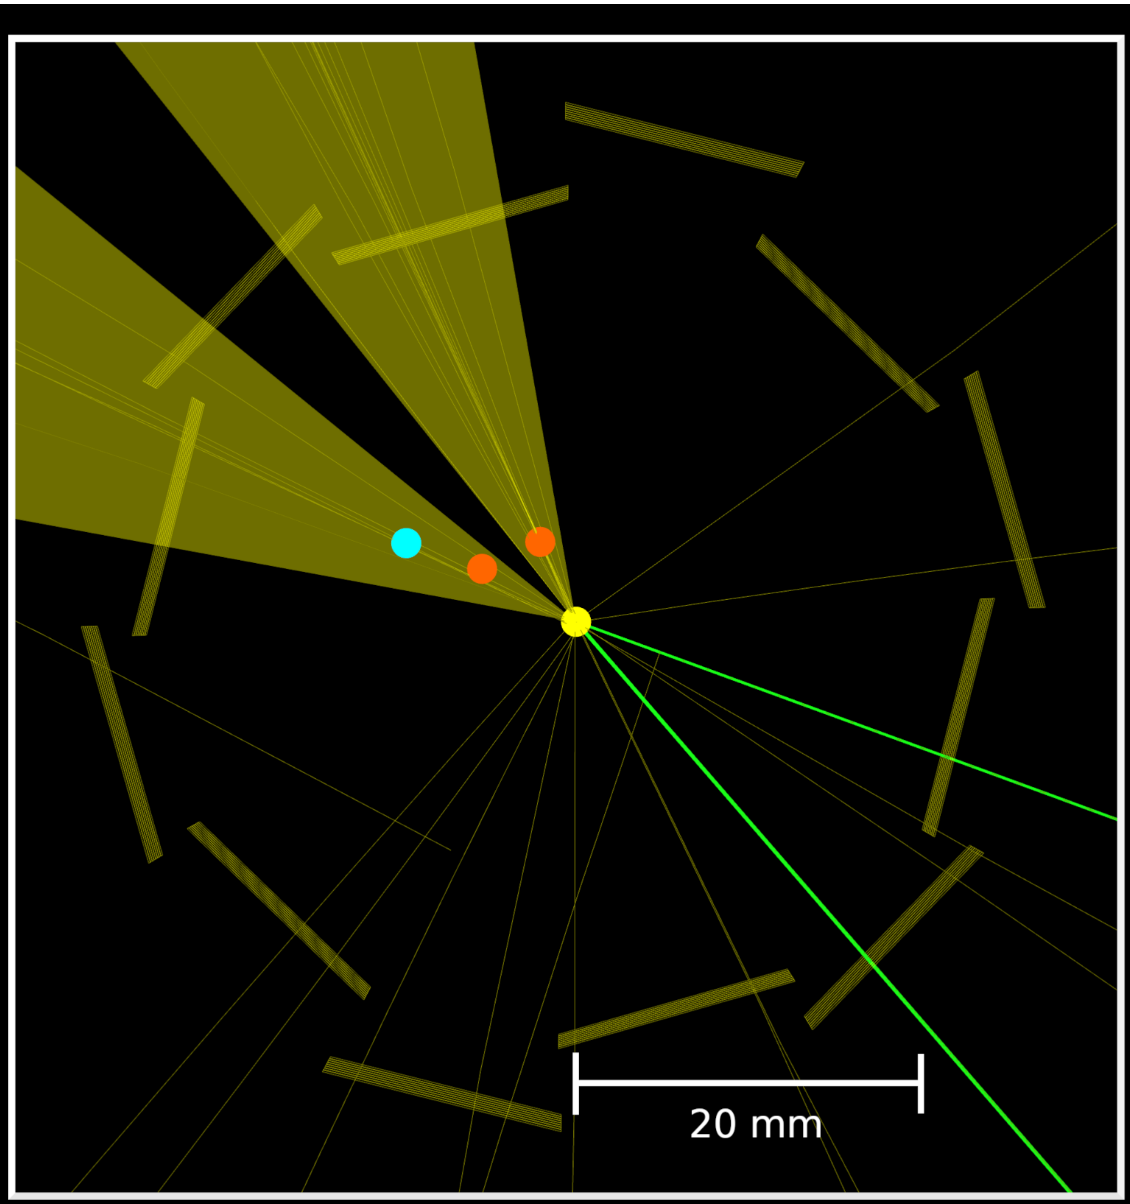

Interesting collisions: Dark Matter?

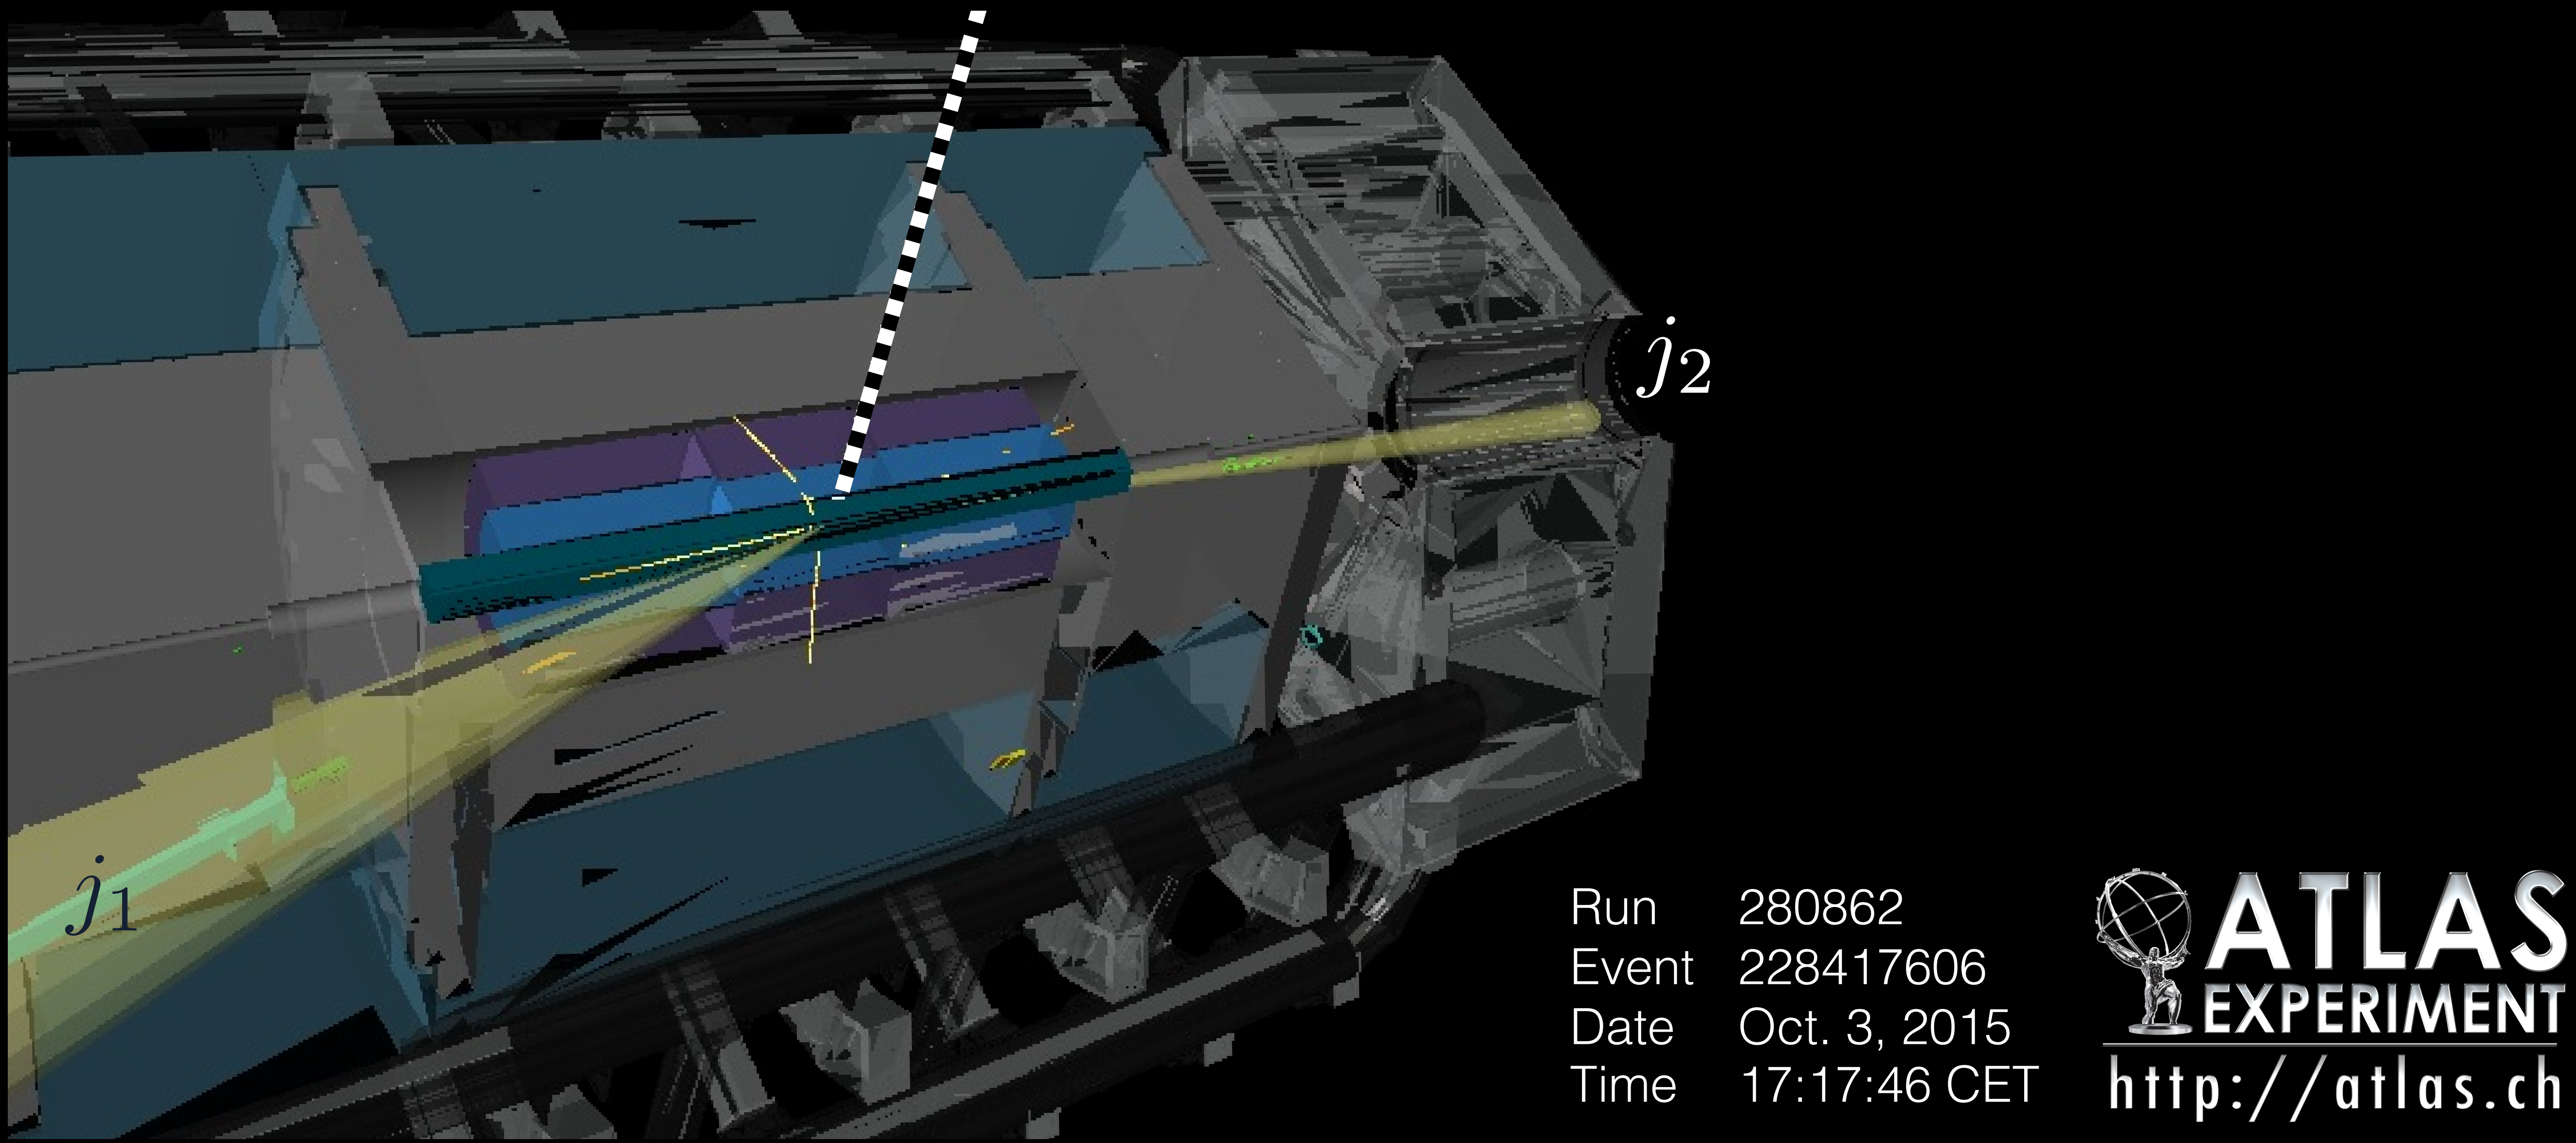

Interesting collisions: Dark Matter?

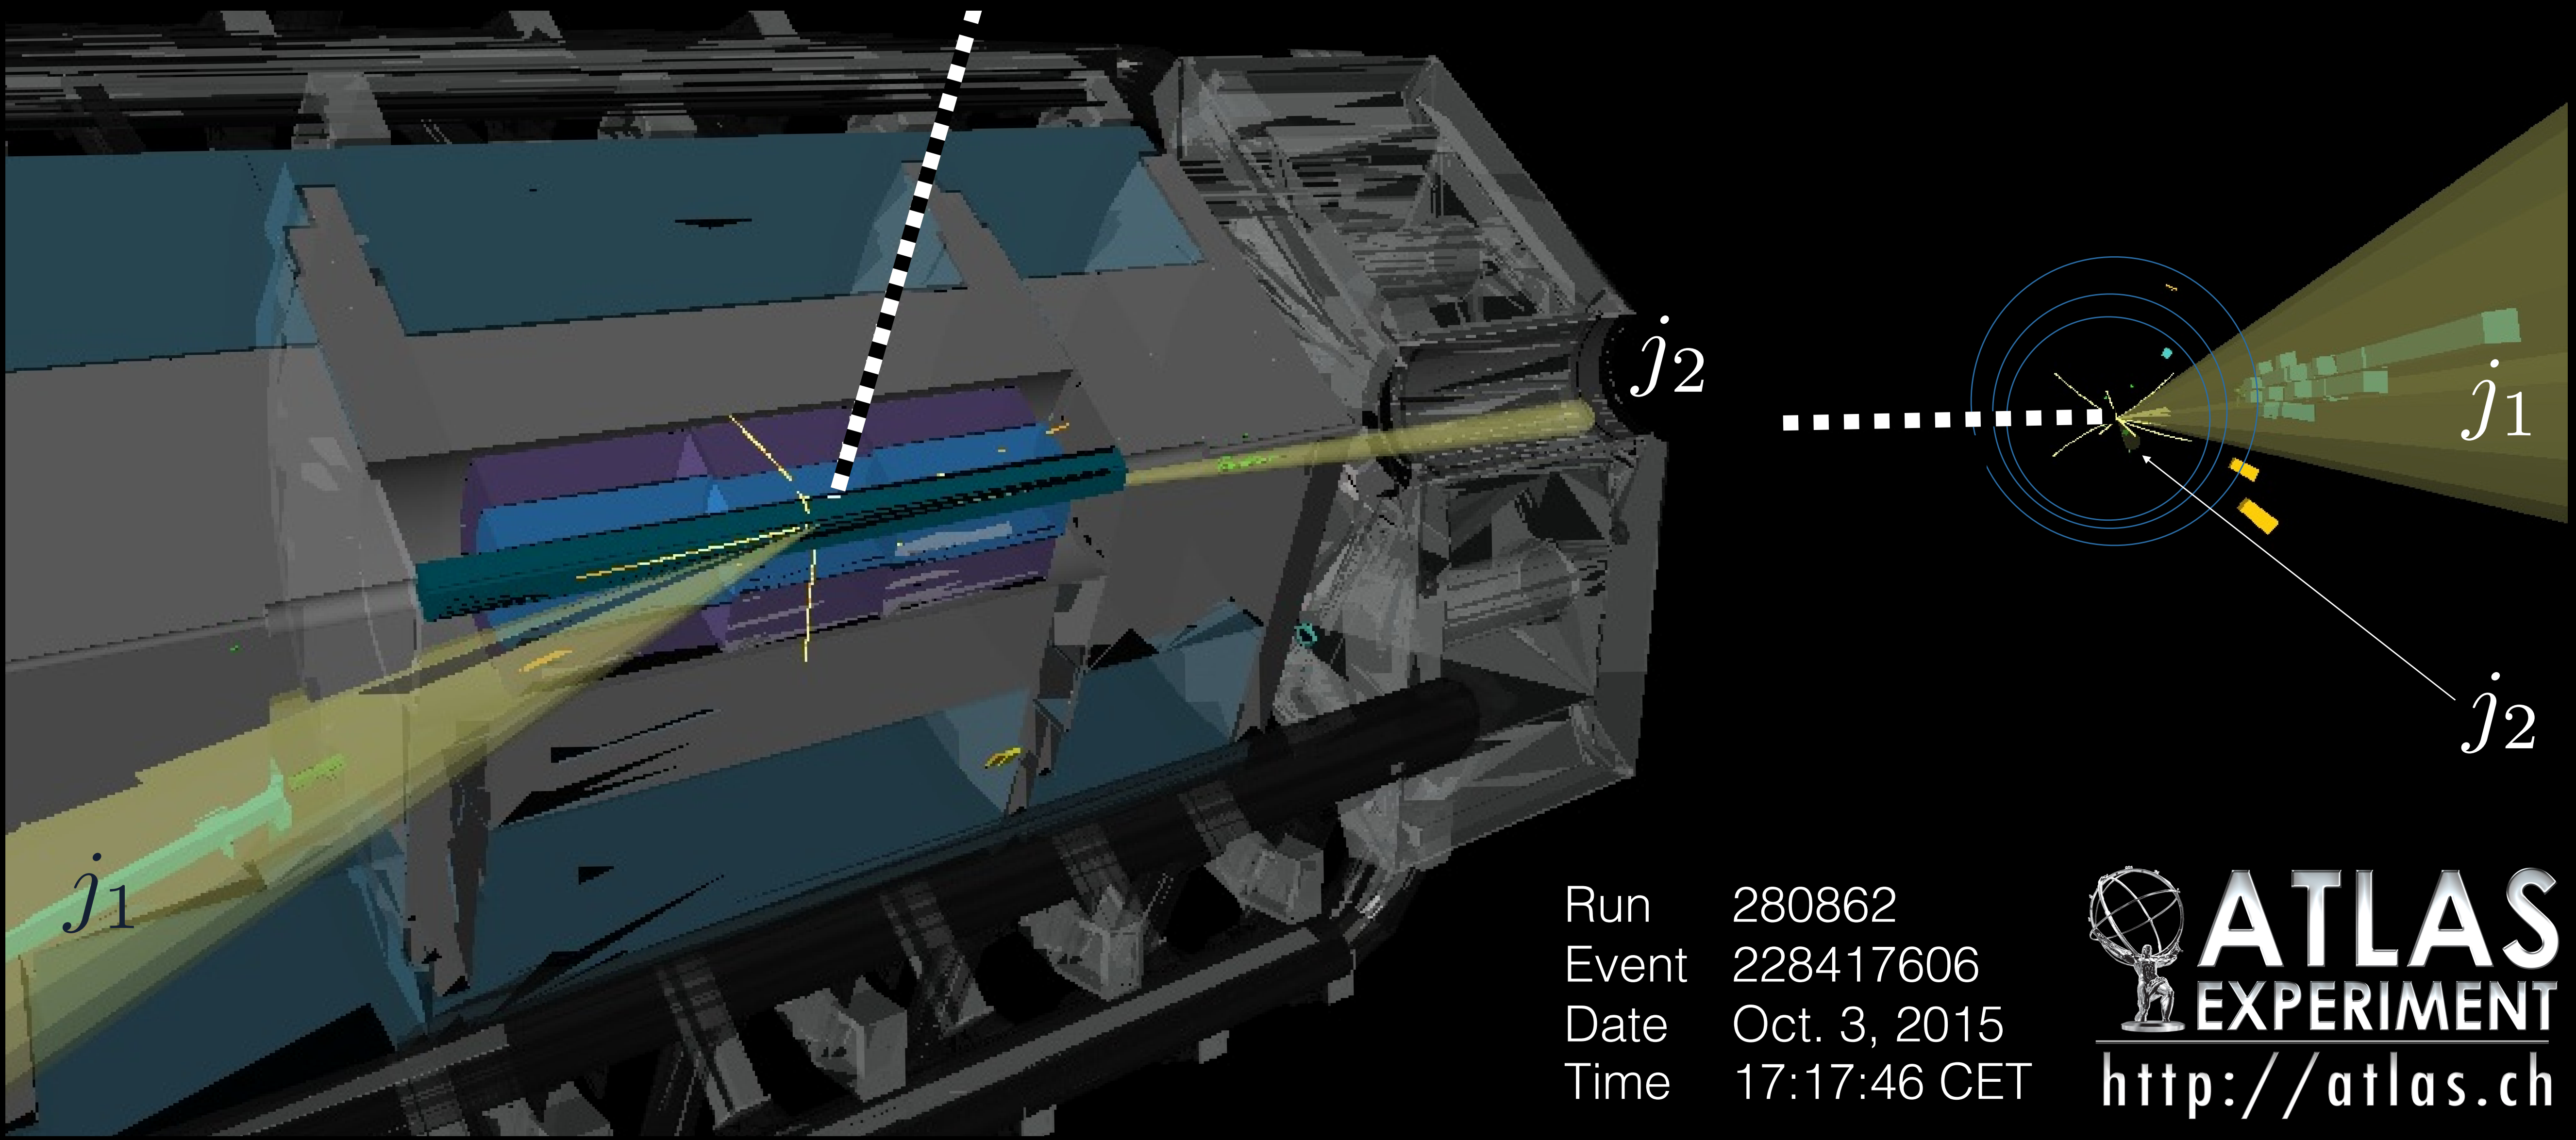

Run 280862
Event 228417606
Date Oct. 3, 2015
Time 17:17:46 CET

Interesting collisions: Dark Matter?

Candidate in signal region of $H \rightarrow \chi\bar{\chi}$ with two VBF jets (m_{jj} = 3.6 TeV)

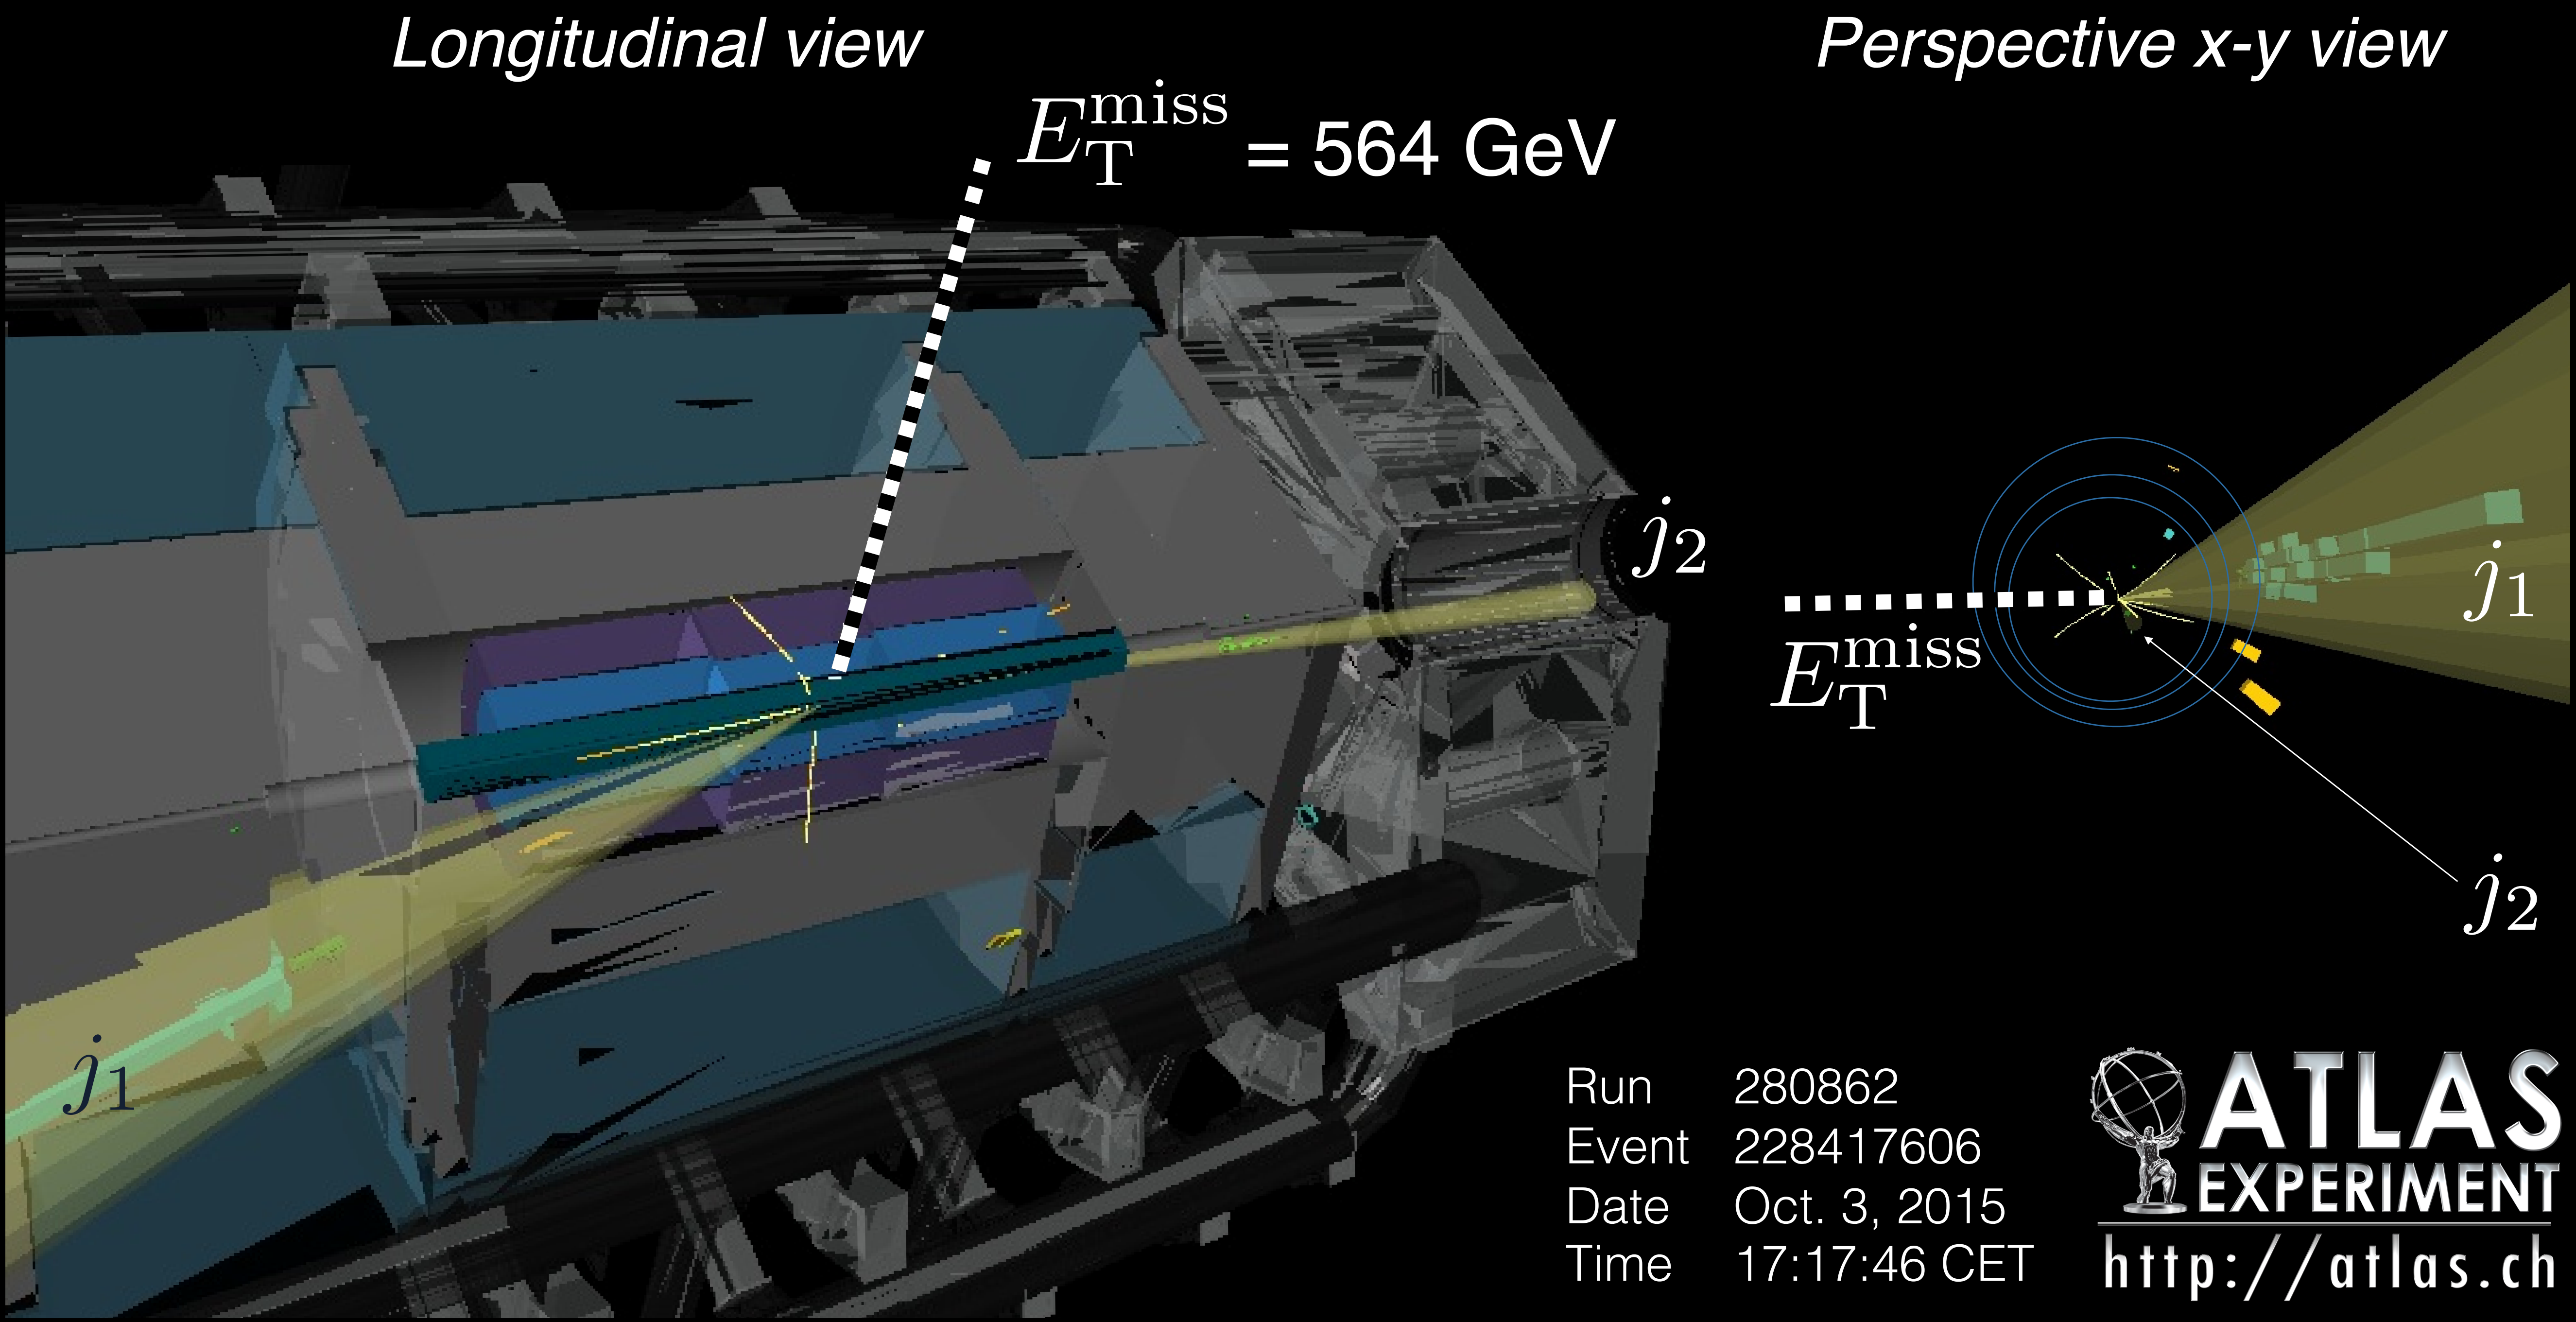

Run	280862
Event	228417606
Date	Oct. 3, 2015
Time	17:17:46 CET

Interesting collisions: top quark partner?

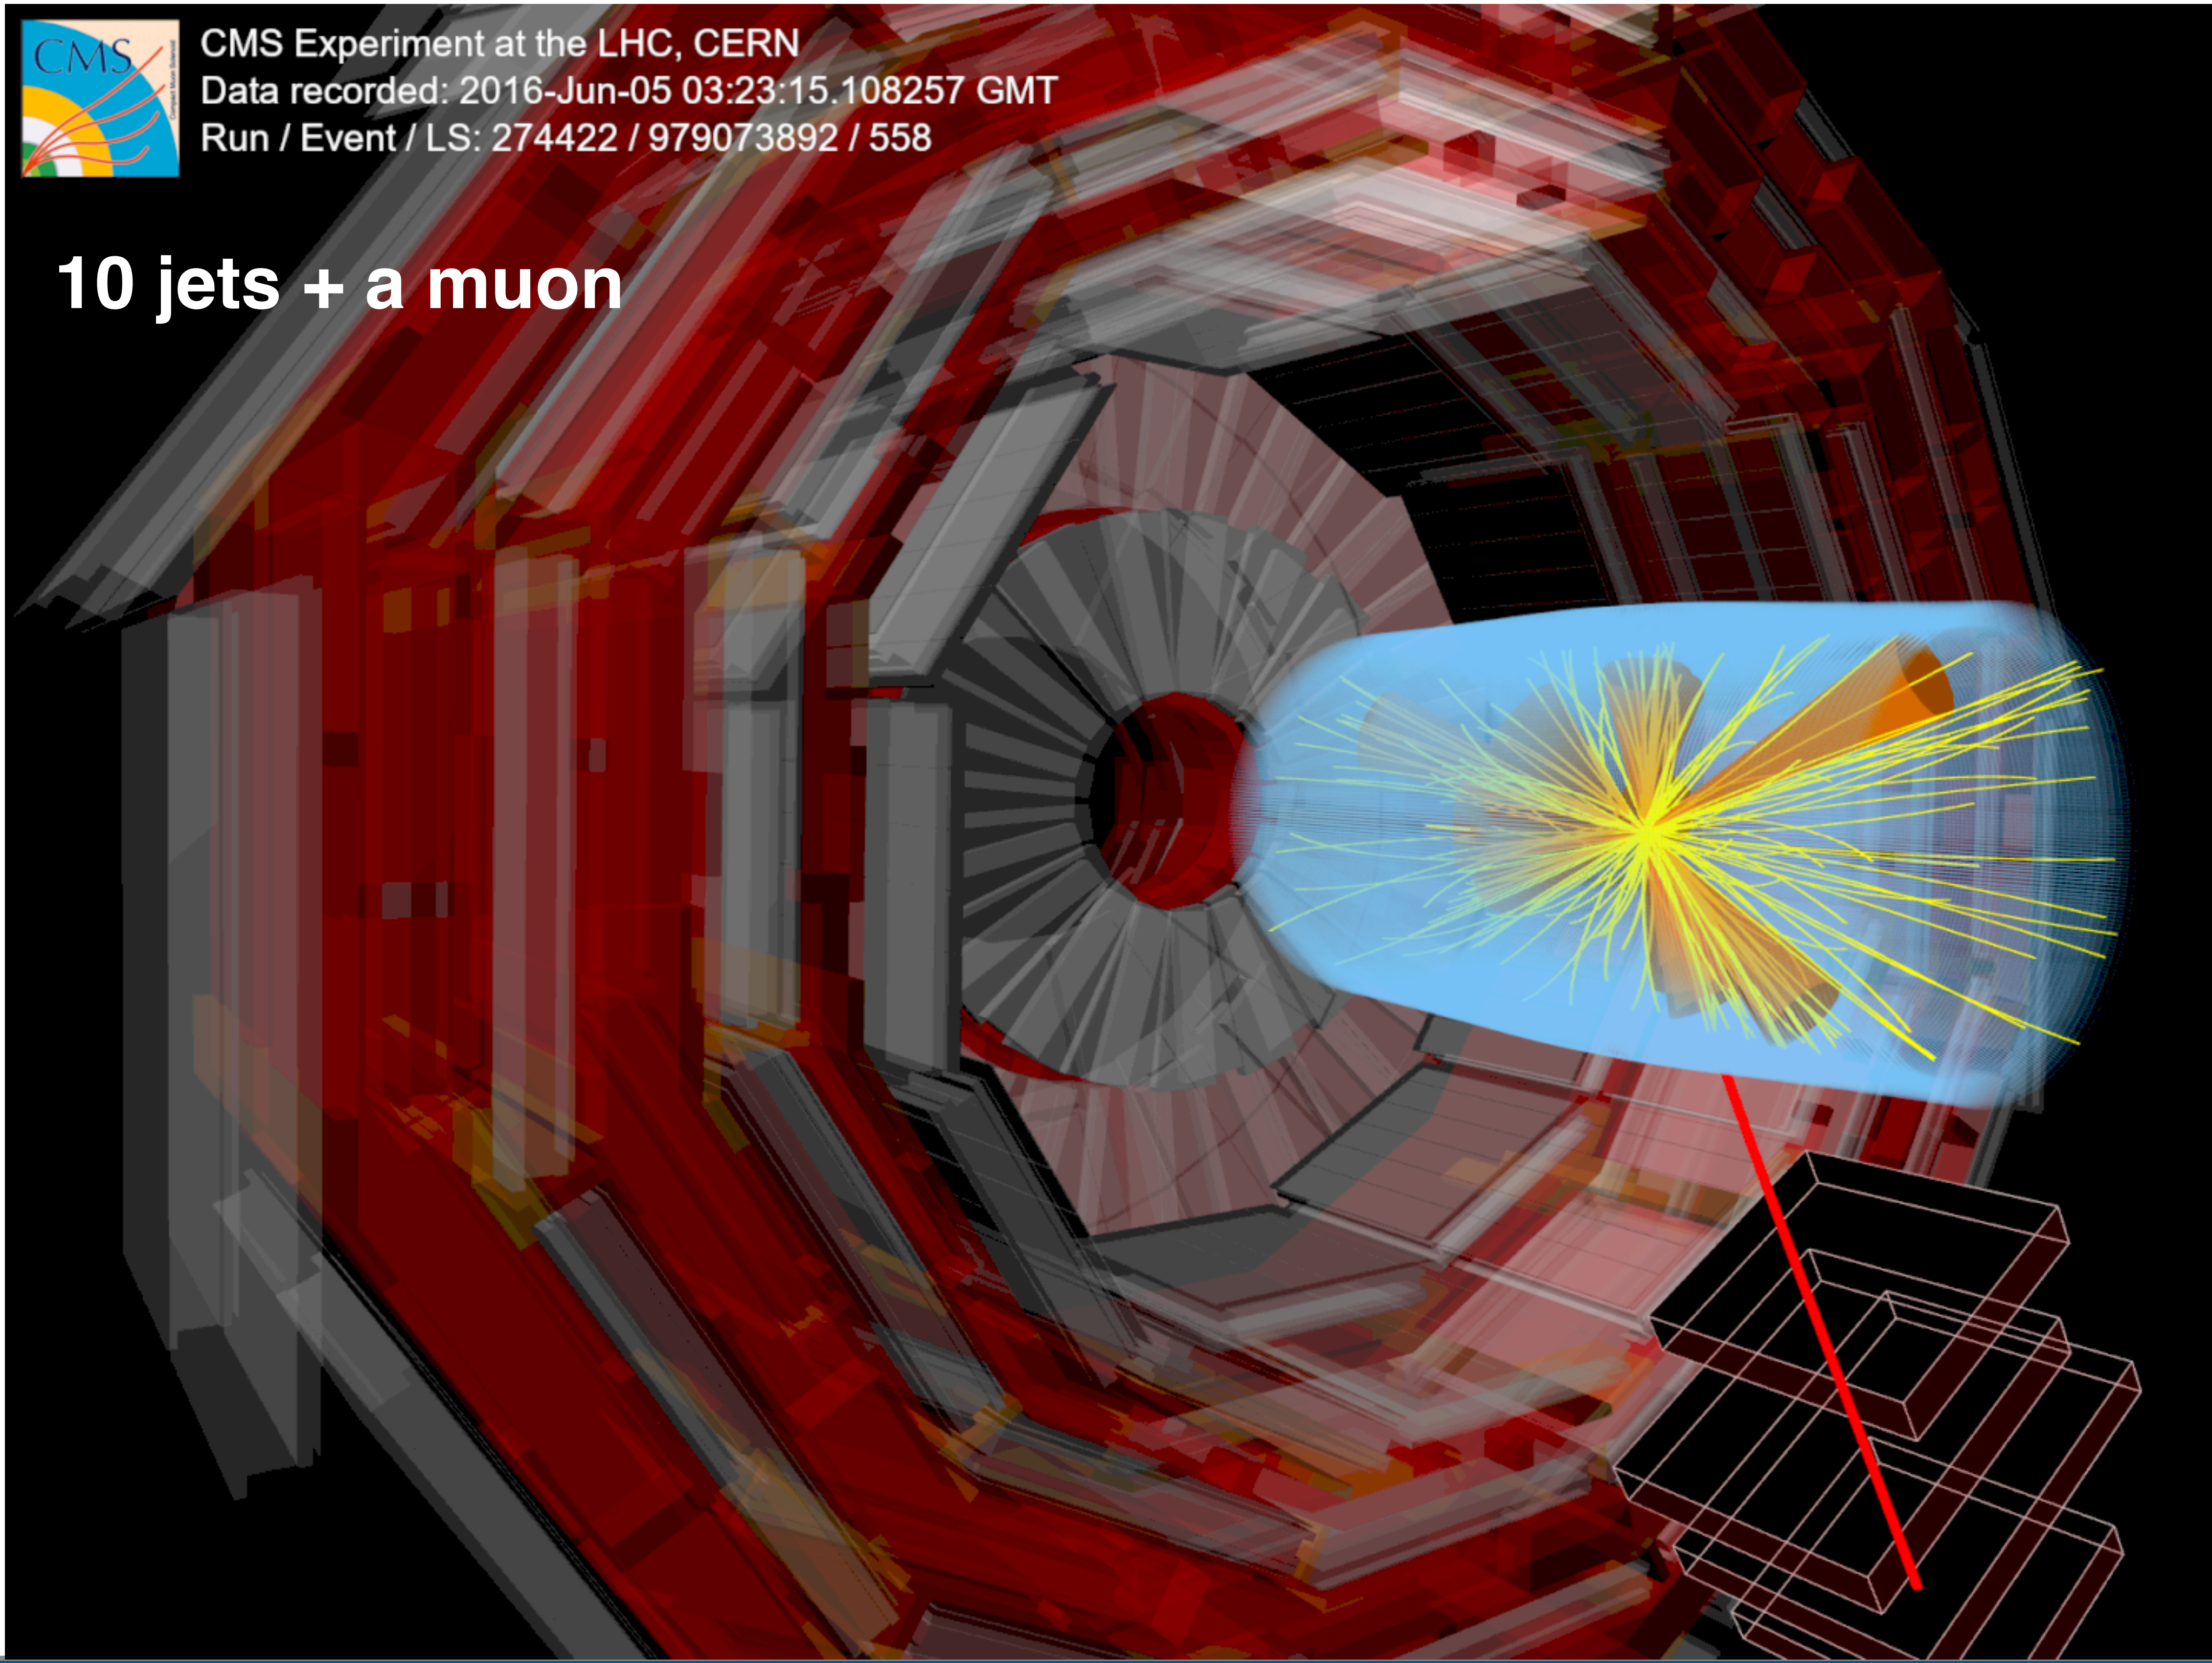

CMS Experiment at the LHC, CERN
Data recorded: 2016-Jun-05 03:23:15.108257 GMT
Run / Event / LS: 274422 / 979073892 / 558

10 jets + a muon

Interesting collisions: top quark partner?

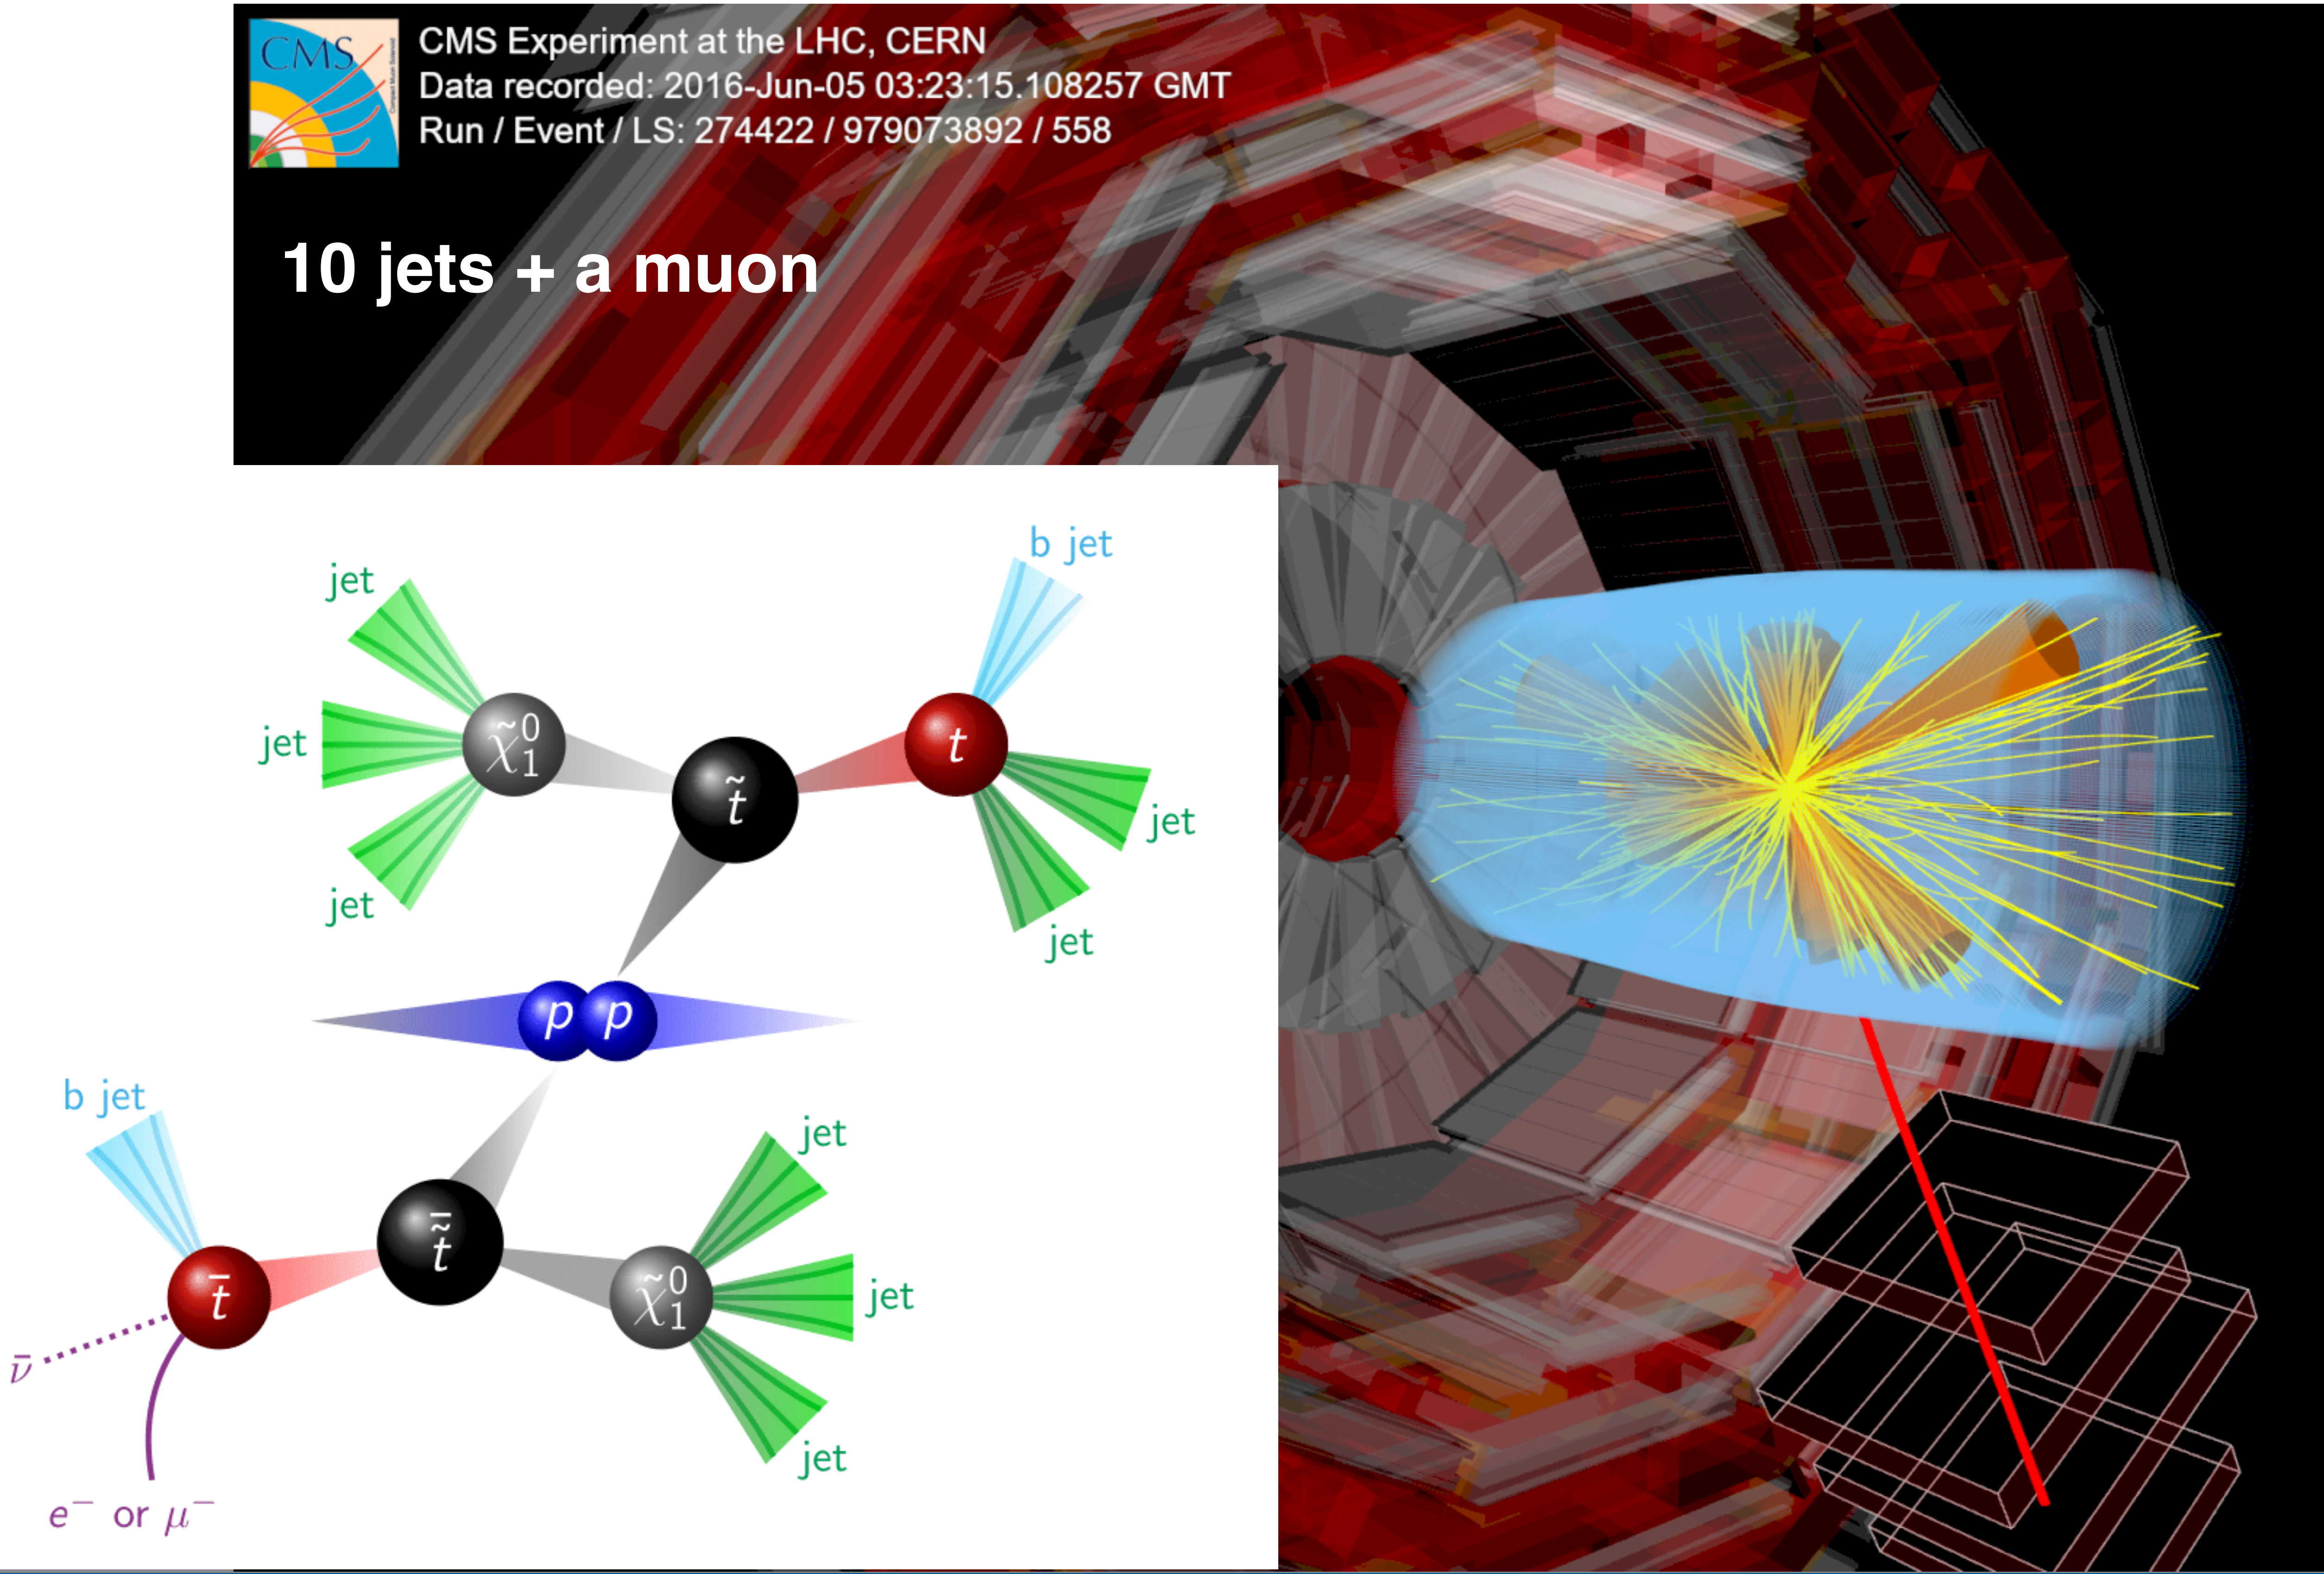

10 jets + a muon

The CMS Trigger Strategy

The events we are most interested in are largely “**spectacular**”

- They have striking features not found in most soft QCD interactions
    - Electrons, muons, tau leptons, photons, “missing-momentum”
    - High-energy objects
    - Multiple objects
    - Hadronic interactions too: e.g. many high-p_T jets, “missing-momentum”

These features can be *quickly identified* with a subset of the complete detector information.

→ This is the role of the CMS Trigger system!

CMS Trigger Design

System should efficiently select "interesting" events to save

Run 2+3
trigger
system

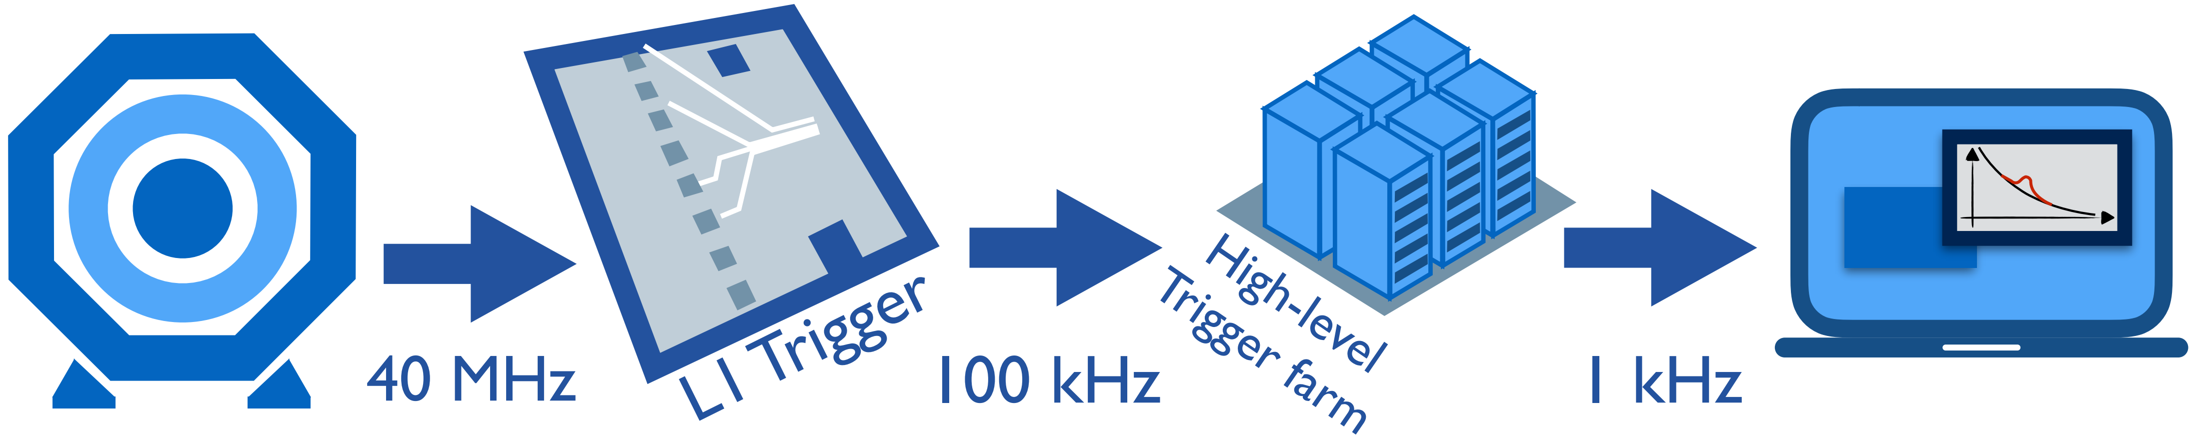

Level-1 Trigger:

Custom electronics

Simple algorithms

Limited data: calorimeters & muons

High-Level Trigger:

Commercial PCs, GPUs

Tracking, Advanced ML, …

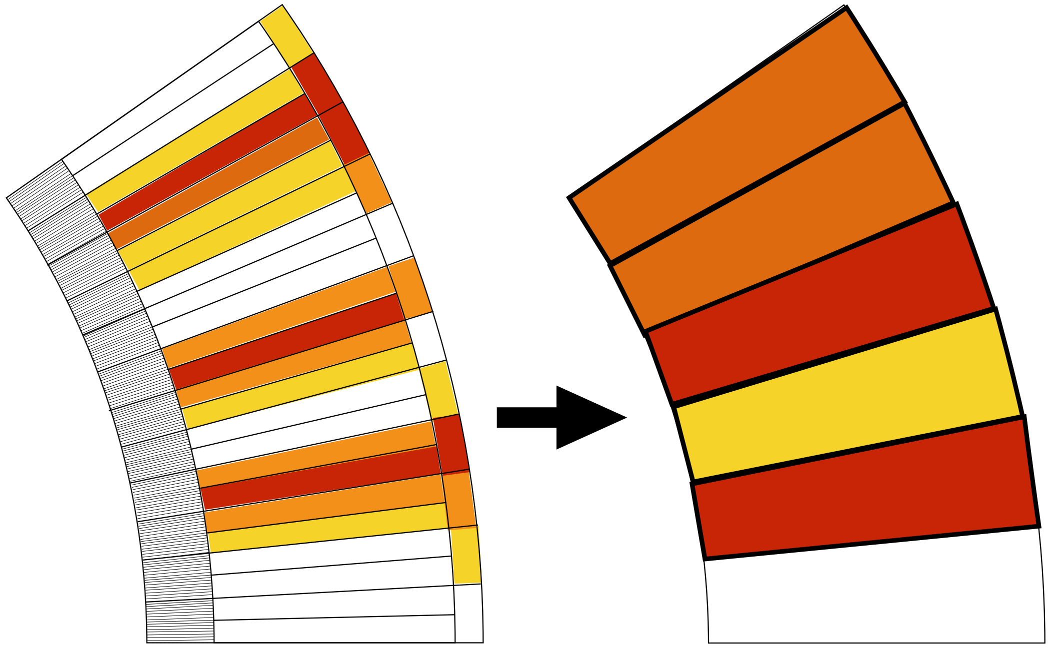

Calorimeter @ Level-1

Workhorse of the Level-1 Trigger

Field Programmable Gate Arrays "compute across space and time"

The underlying technology of LHC triggers

Fully re-programmable:

Build custom circuits by connecting:

Memories, Multipliers, and other configurable logic blocks

Enables:

- Highly parallel computation
- High-throughput (Tb/s)

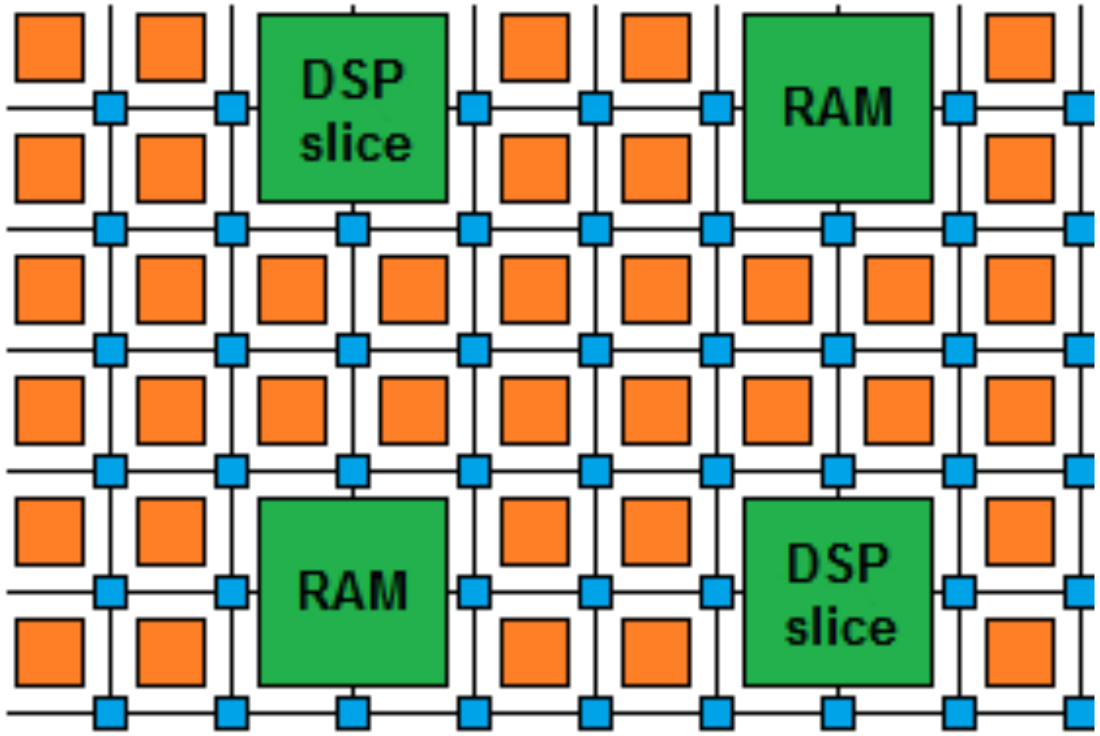

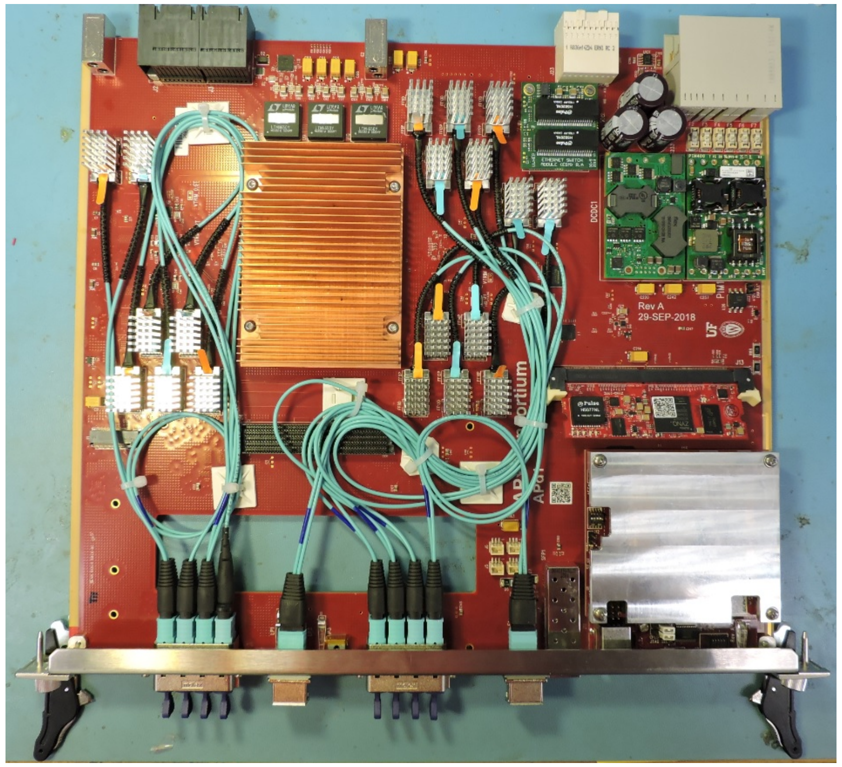

A prototype Trigger board for CMS

Unlocking FPGAs for physics

Recently, a number of advances have increased the suitability of FPGAs for physics applications.

Improving technologies:

Typical Run 1 FPGA:

~50 multiplier units

Target Phase-II FPGA:

~12k multiplier units

Accessible and automated design tools:

High-level synthesis (HLS) allows non-experts to create firmware designs with C code

→

You and me:

Growing community of physicists and engineers working to unlock our experiments full potential!

Triggering on electrons & photons

Electromagnetic showers are identified from their shapes in ECal+HCal

→ expect narrow, isolated clusters

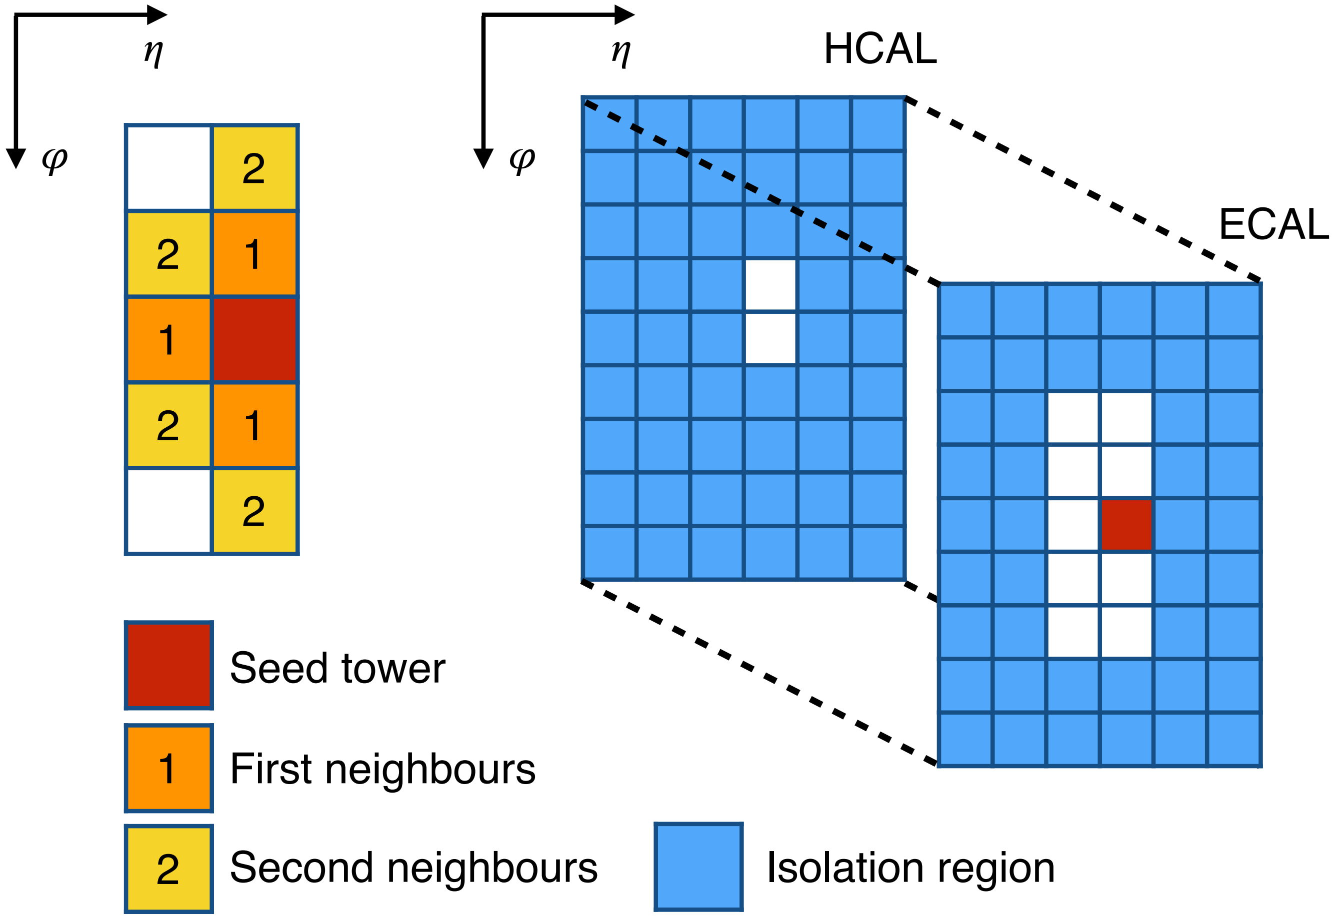

Typical thresholds @ 2E34

SingleIsoEG > 30 GeV
DoubleEG > 25,14 GeV
TripleEG > 18,17,8 GeV

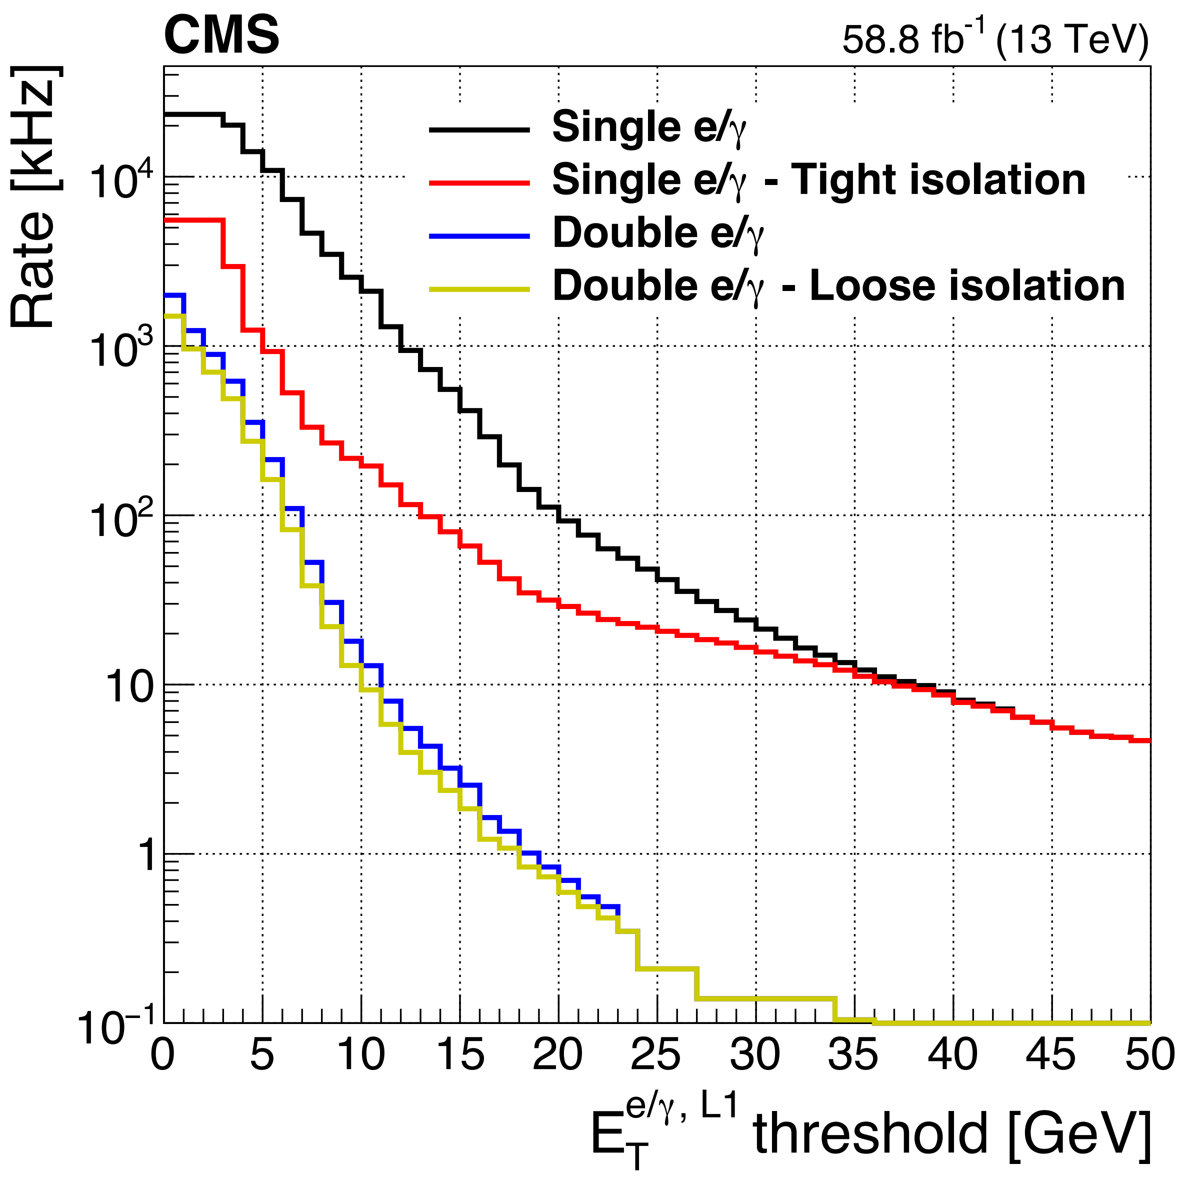

Triggering on electrons & photons

Electromagnetic showers are identified from their shapes in ECal+HCal

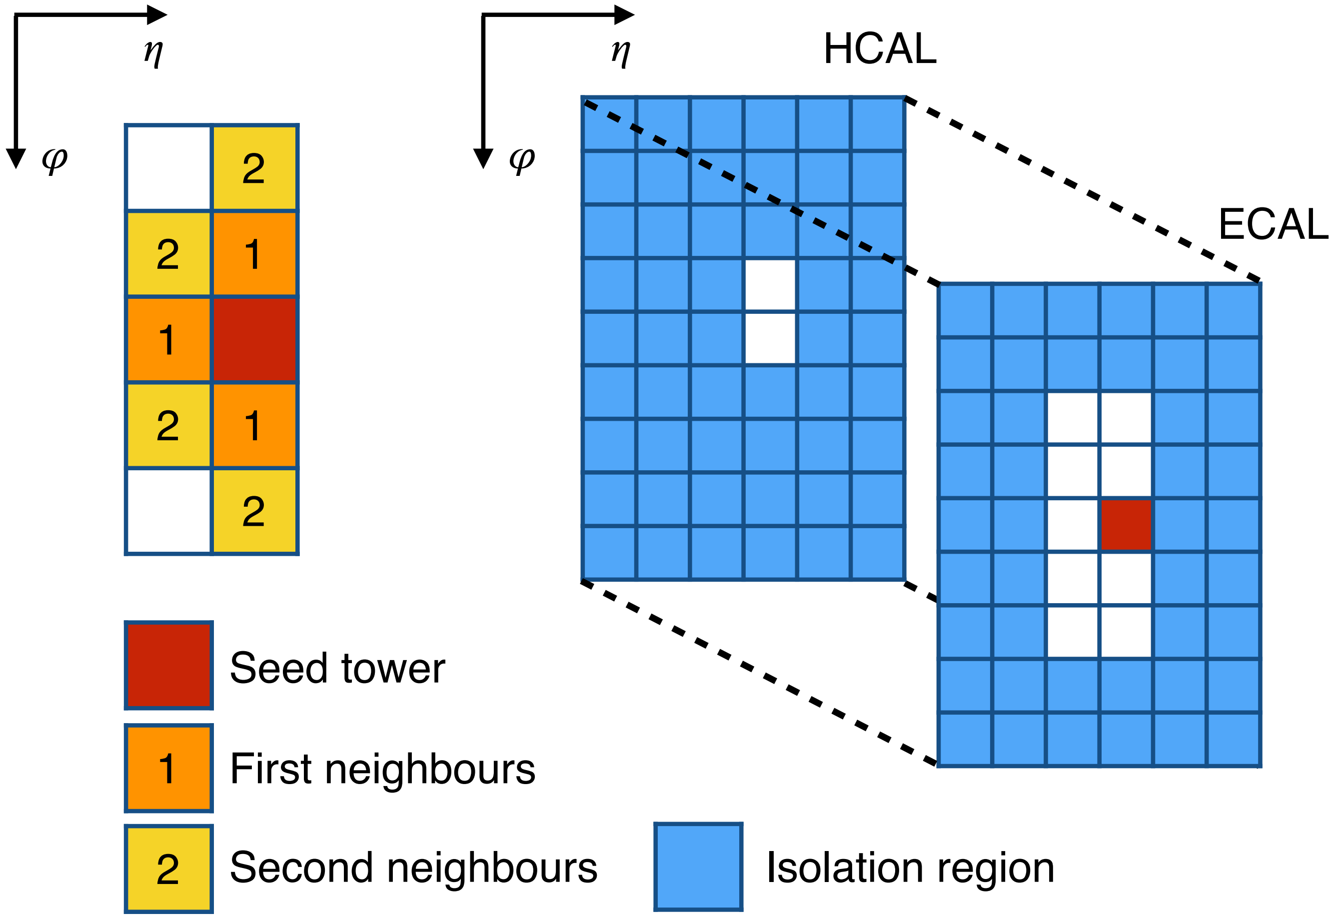

→ expect narrow, isolated clusters

Good correspondence between the trigger algorithm and full "offline reconstruction" →

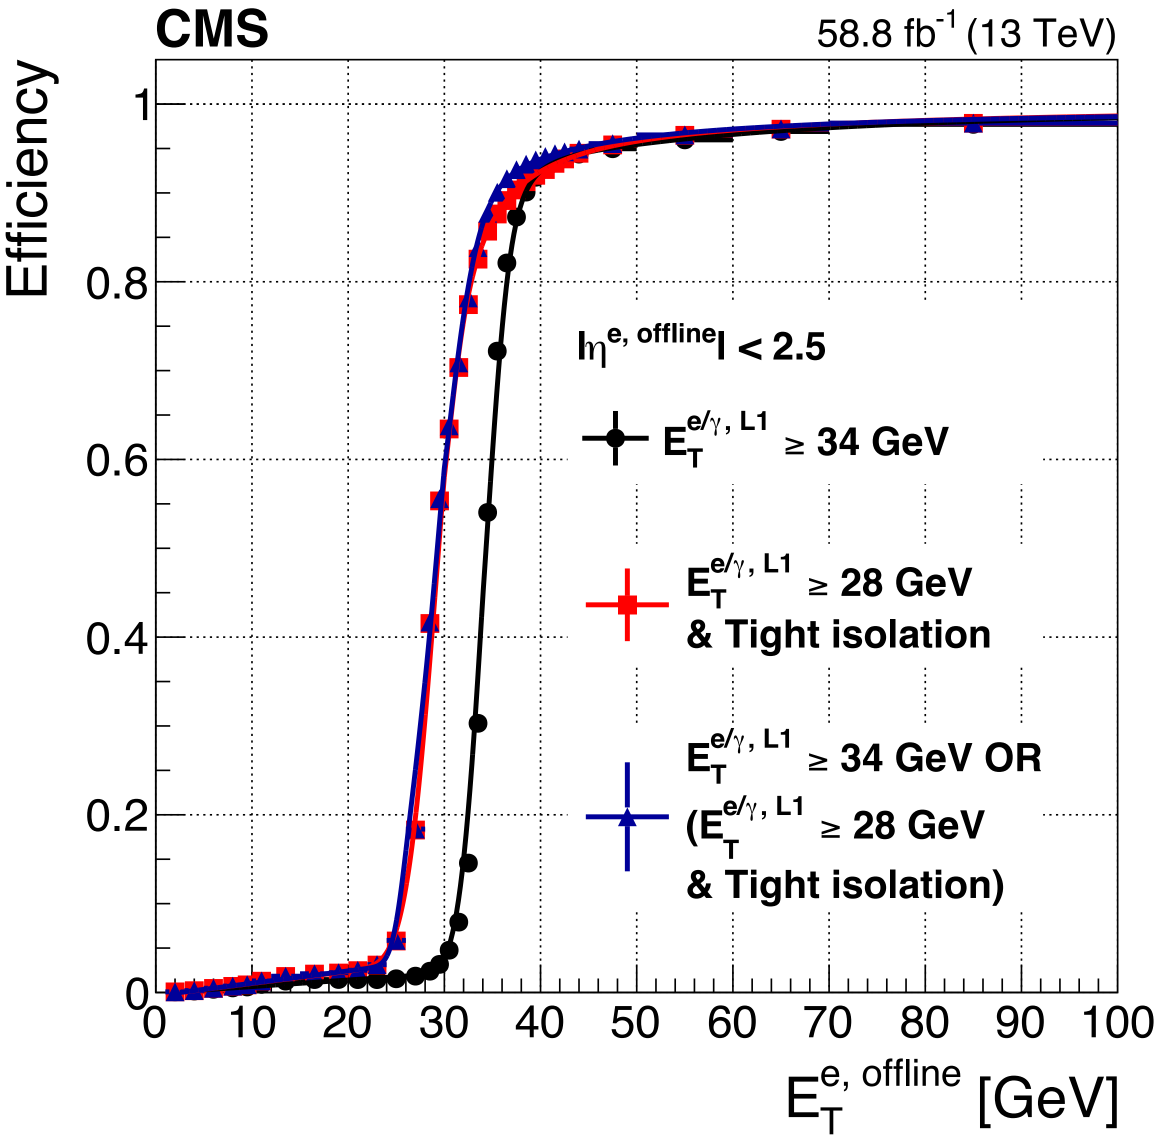

Other hadronic algorithms

Hadronic tau decays can be identified similar to e/ɣ

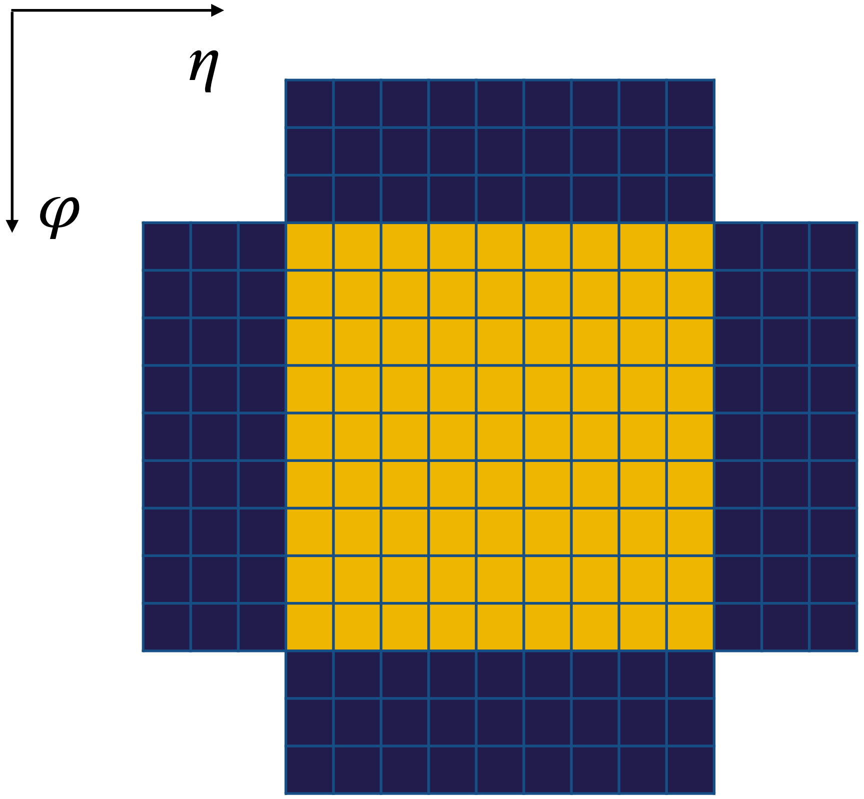

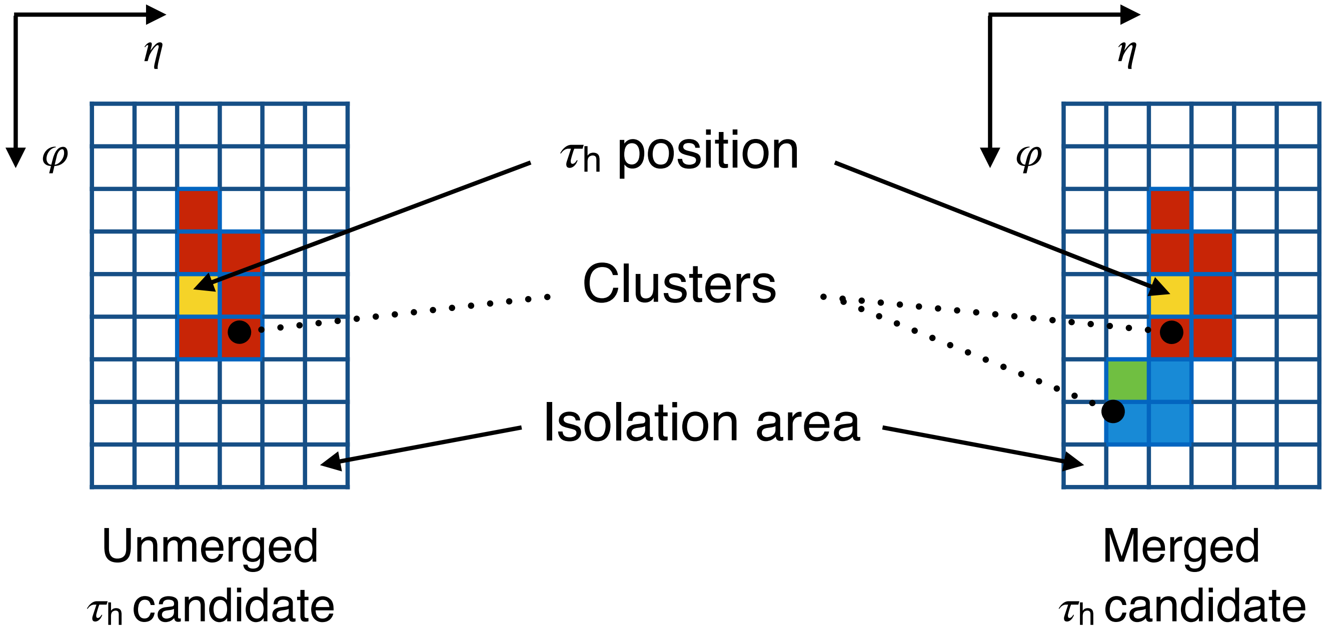

τ_h position

Clusters

Isolation area

Unmerged τ_h candidate

Merged τ_h candidate

Jet energies are especially sensitive to pileup collisions

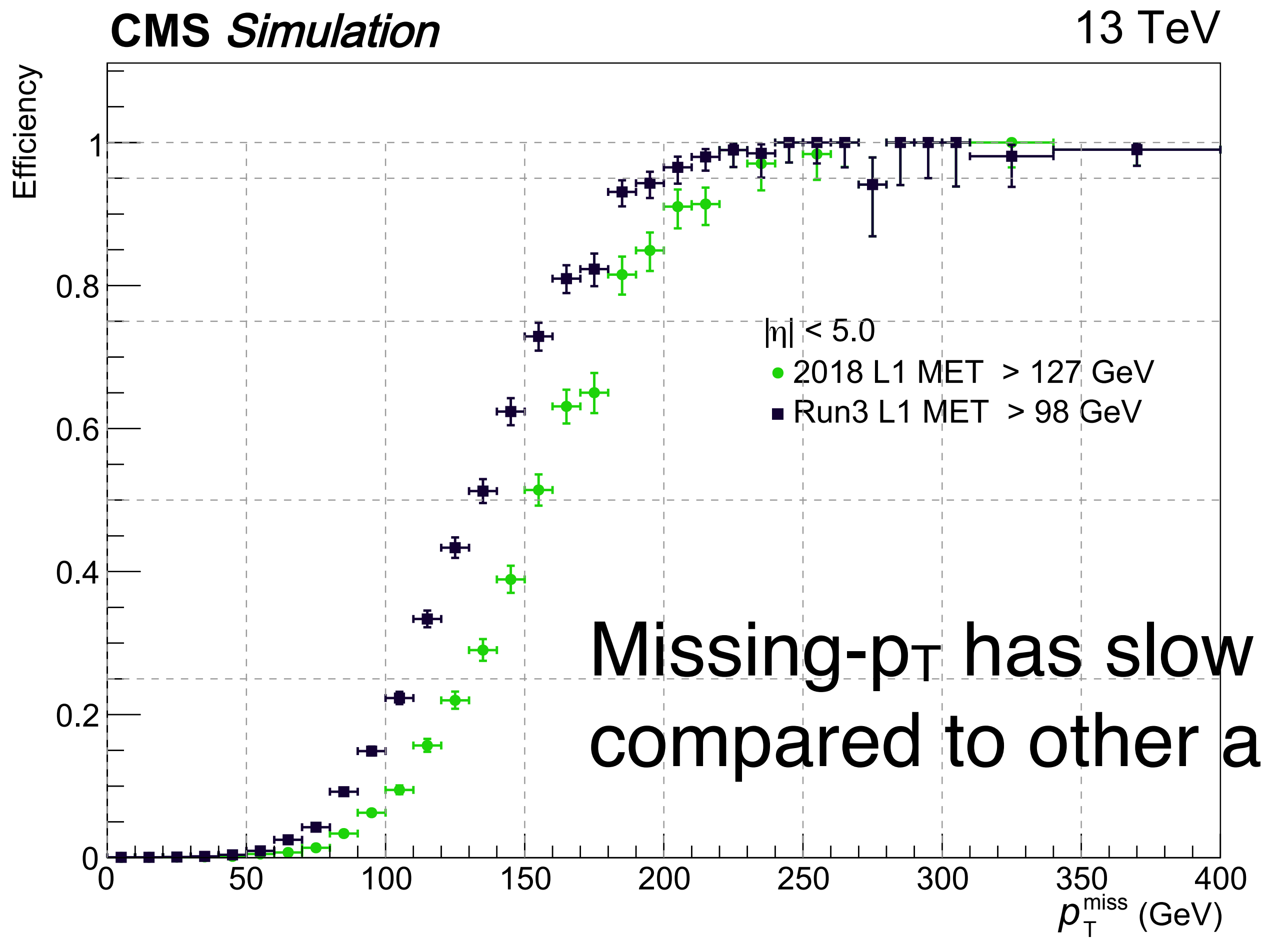

Missing-p_T has slow “turn-on” compared to other algorithms.

Muons

Similar to offline: low hit rates b/c few particles reach muon chambers. Missing the silicon tracker, which refines this p_T measurement.

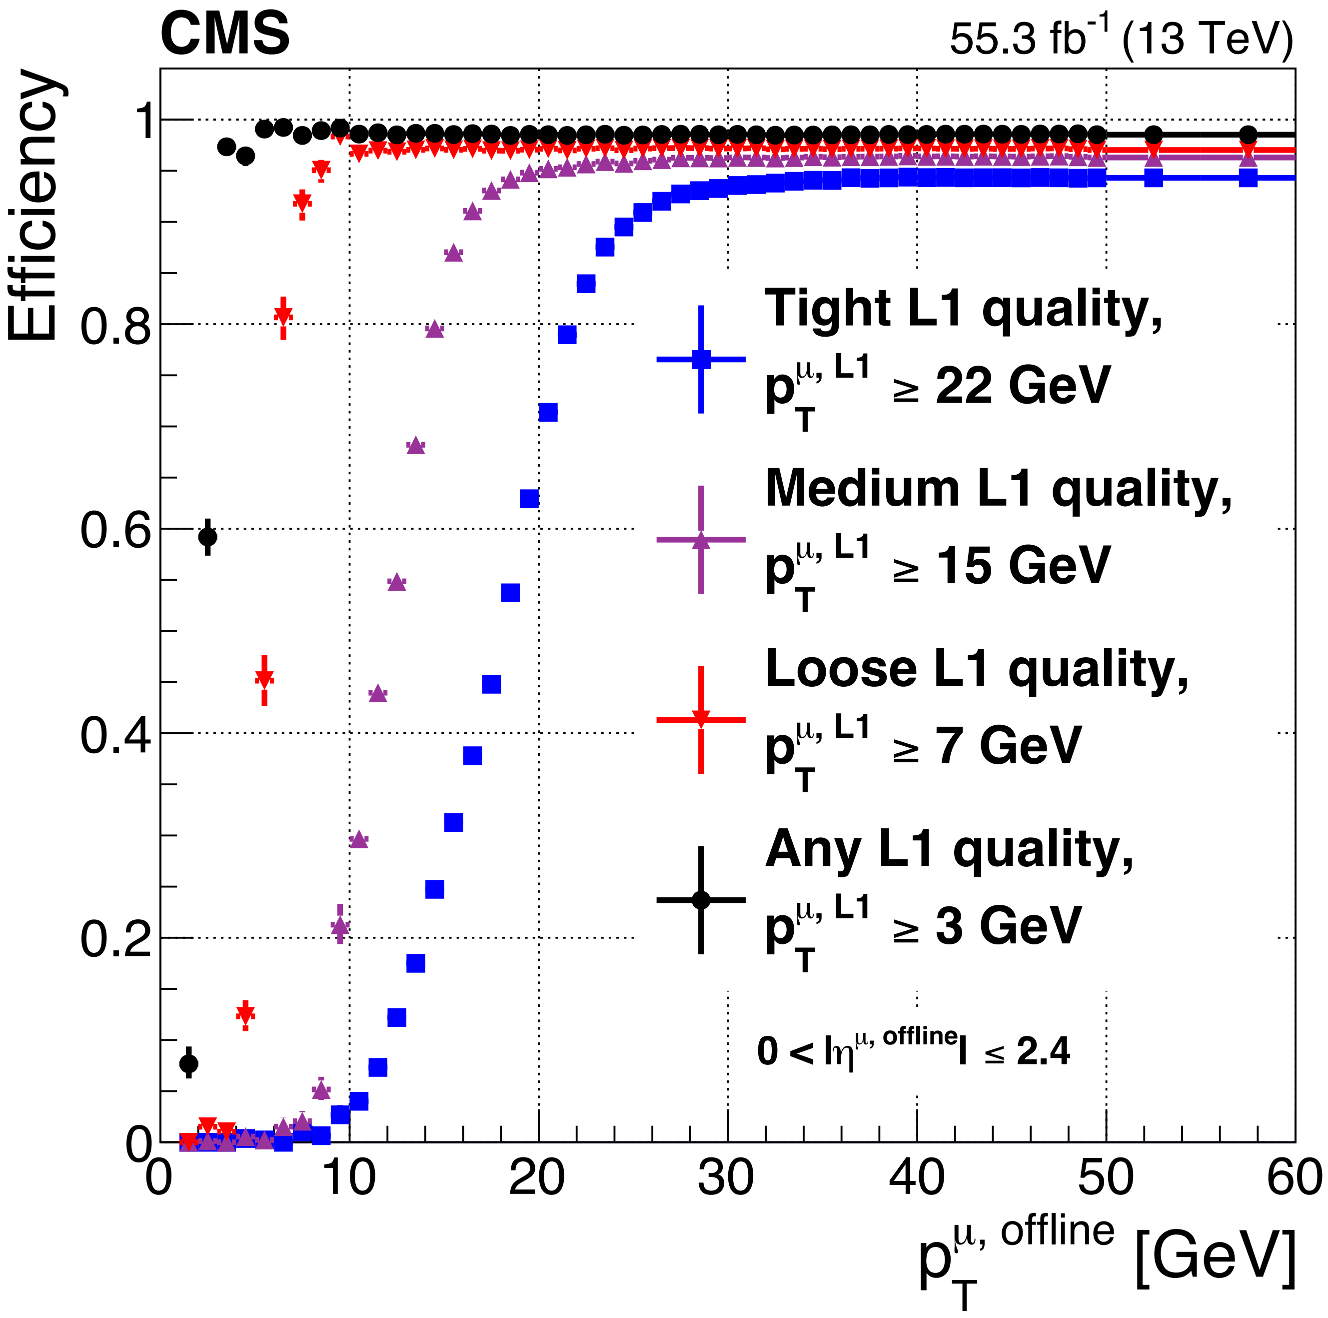

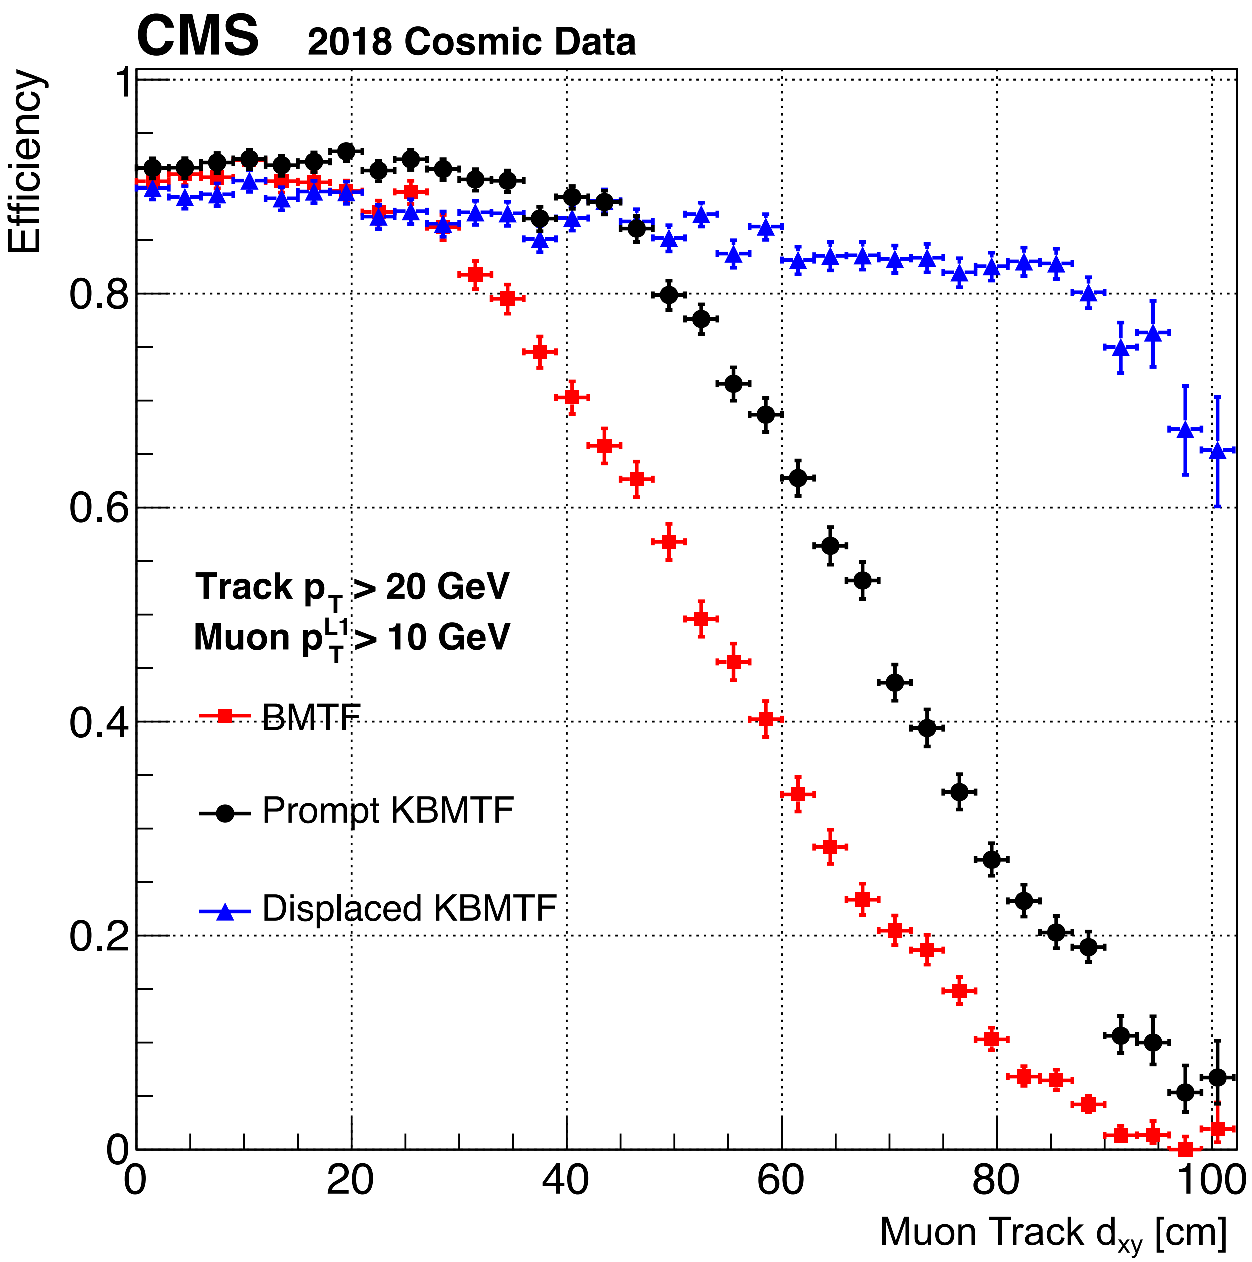

Breaking down the total L1 accept rate

The total 100kHz accept rate is split among many different algorithms.
Considerations: physics priorities, event purity, overlapping events,…

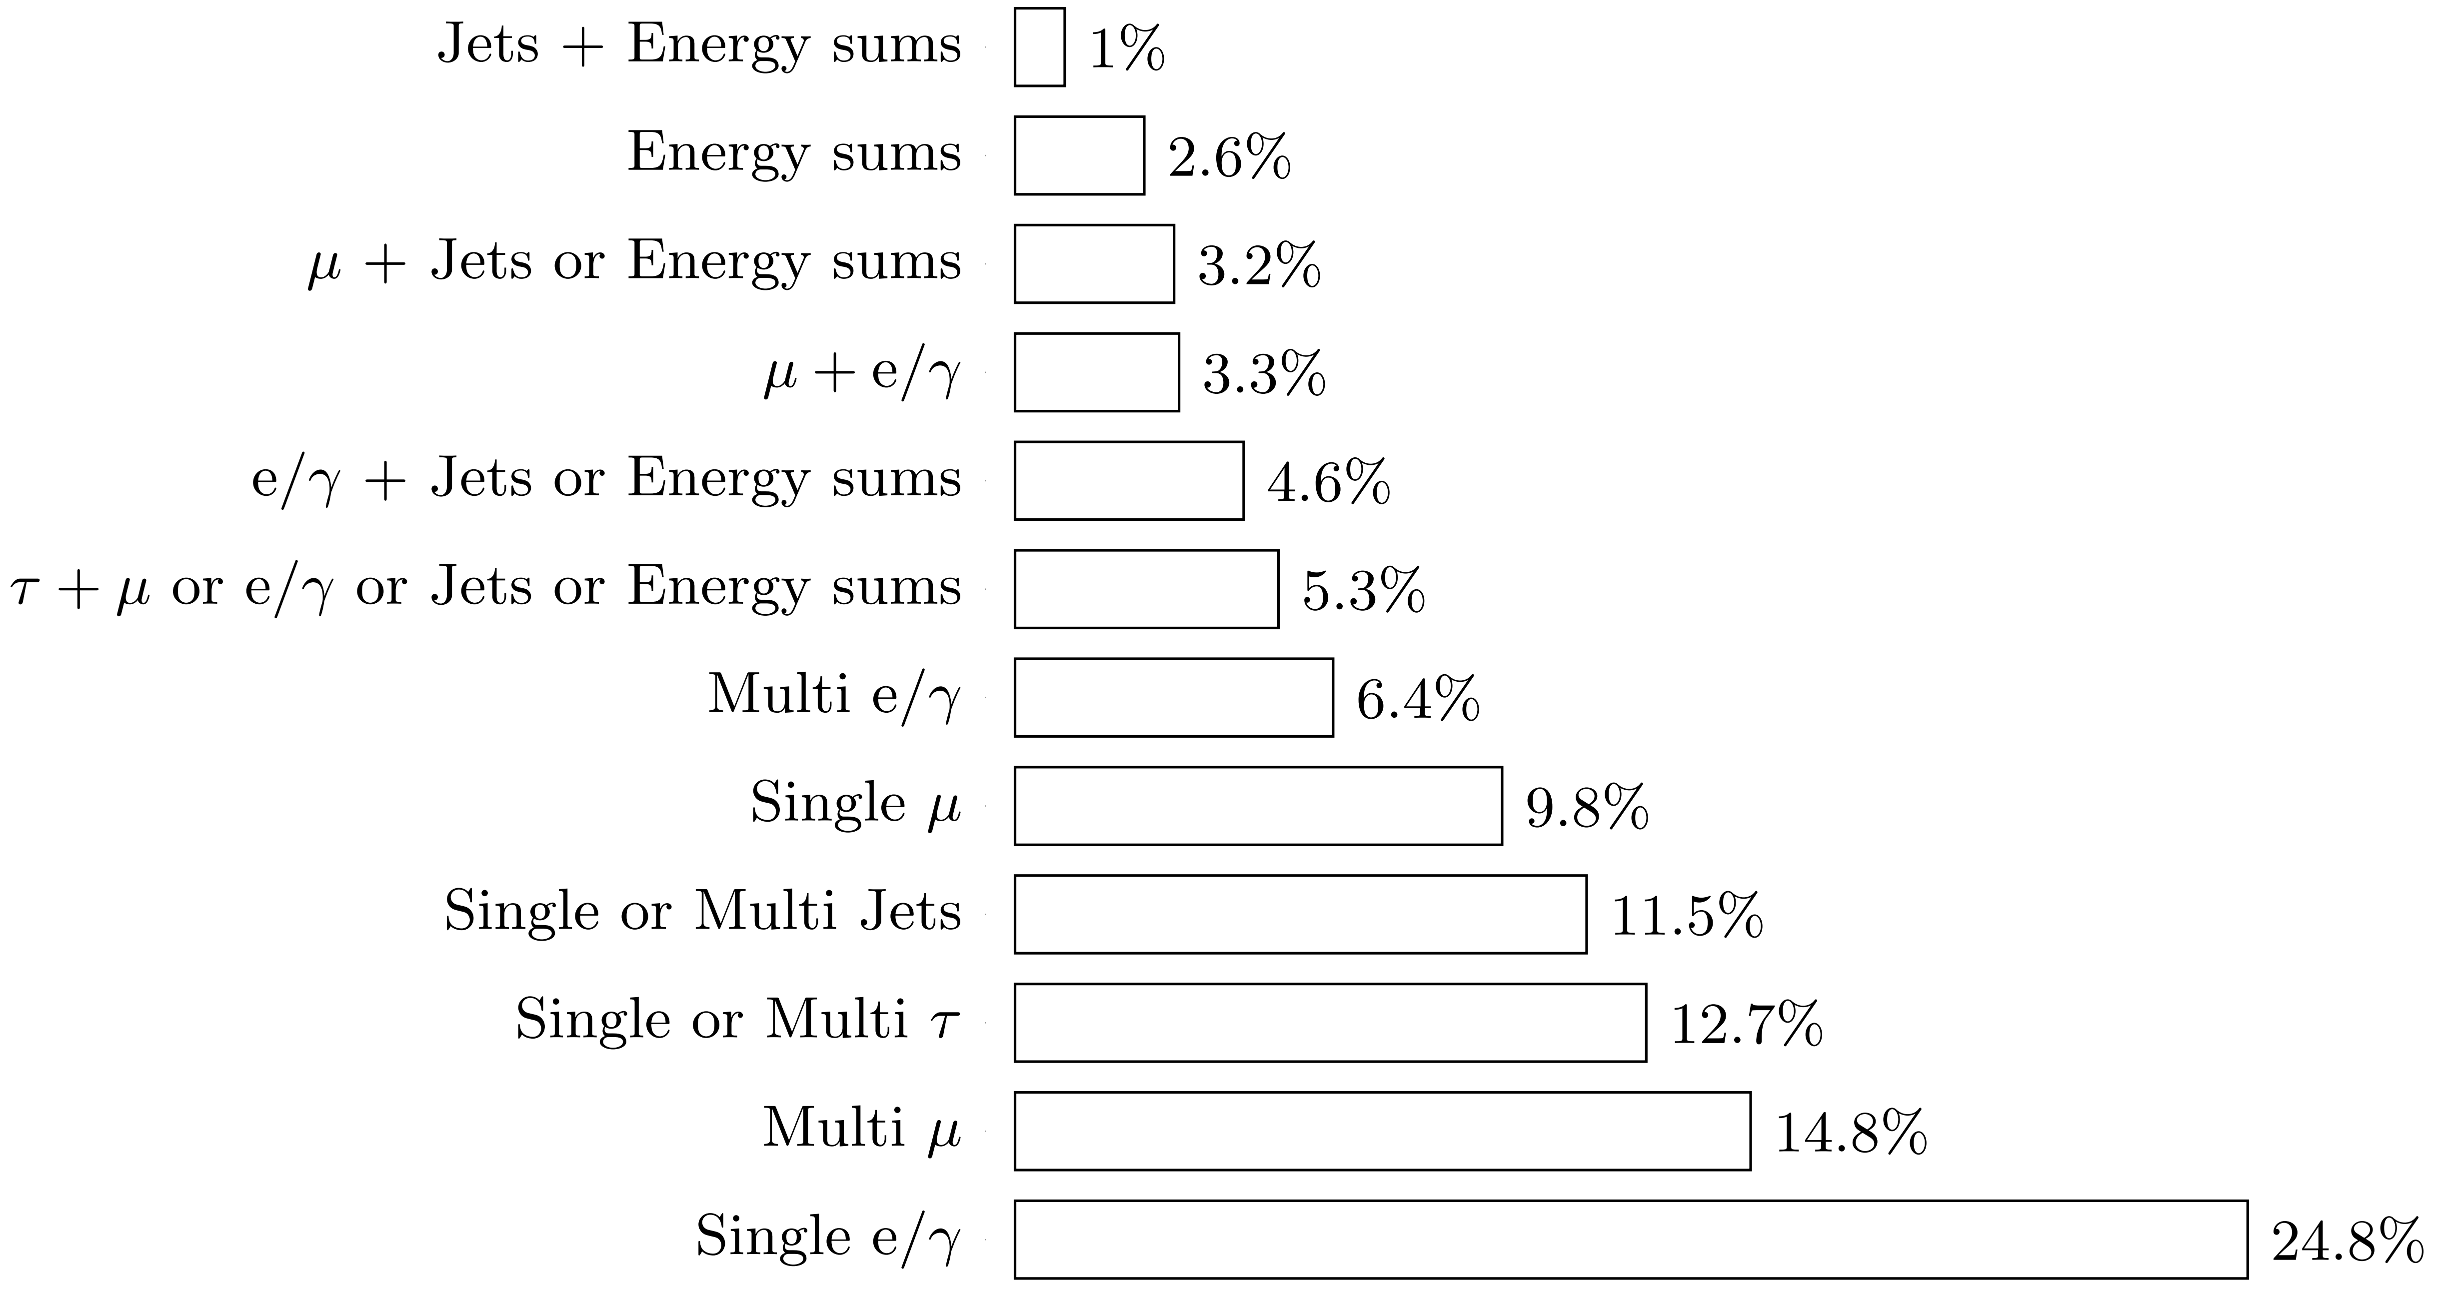

Level-1 trigger rates

Menus with different energy thresholds target different inst. luminosity.
Ideally the rate of accepted events scales linearly with pileup.

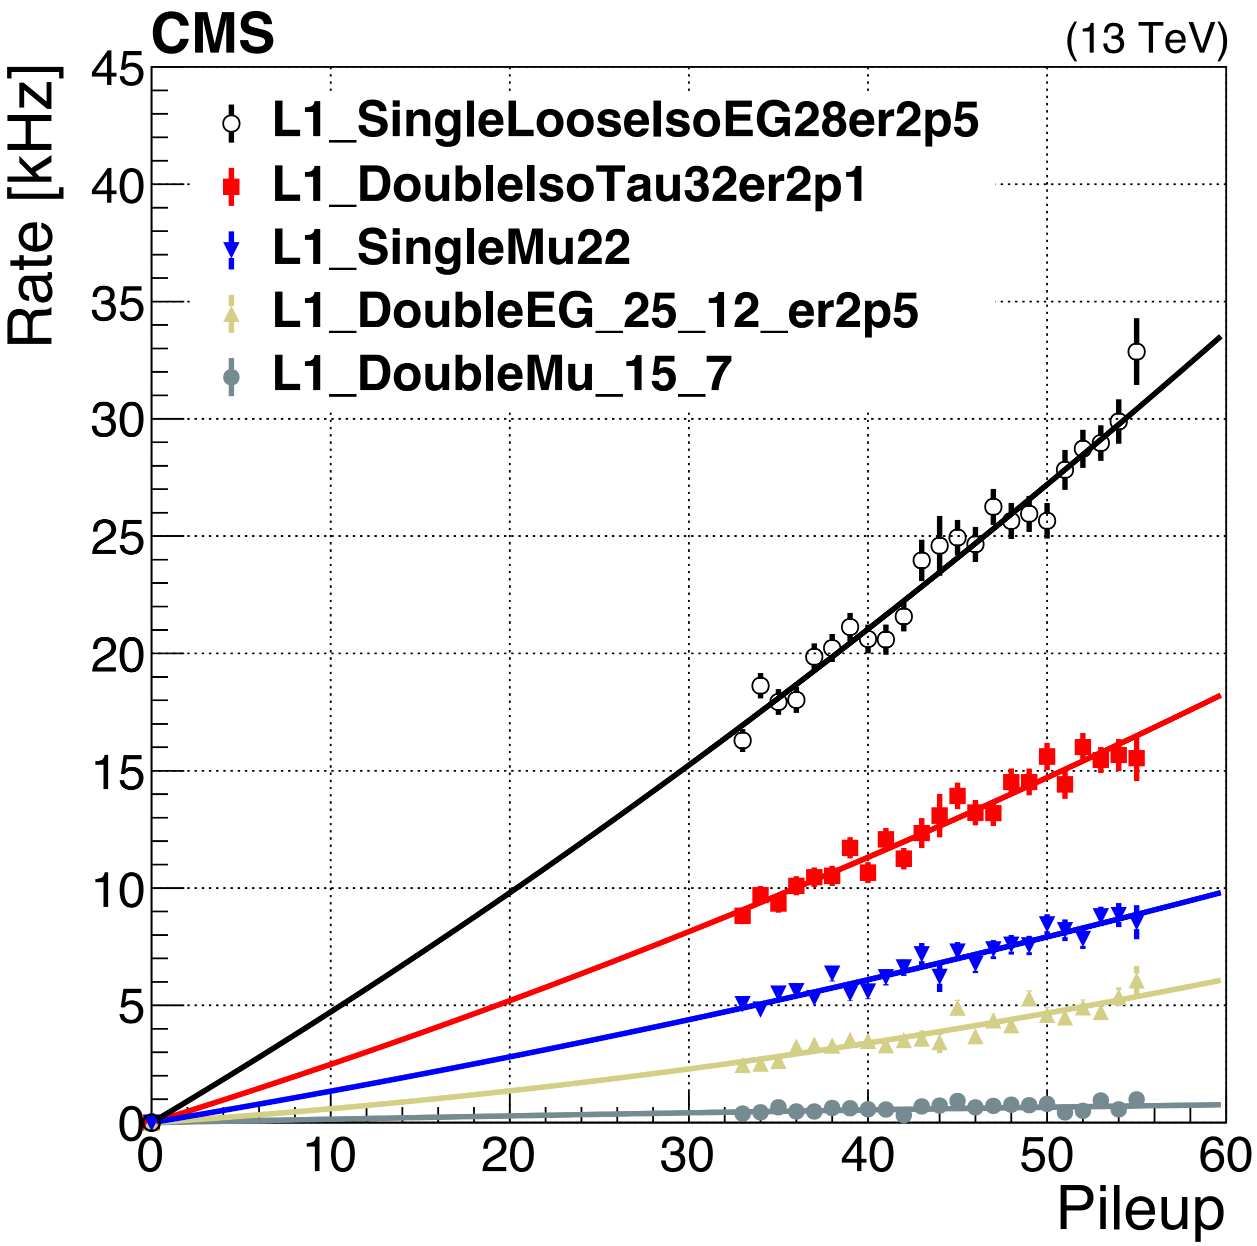

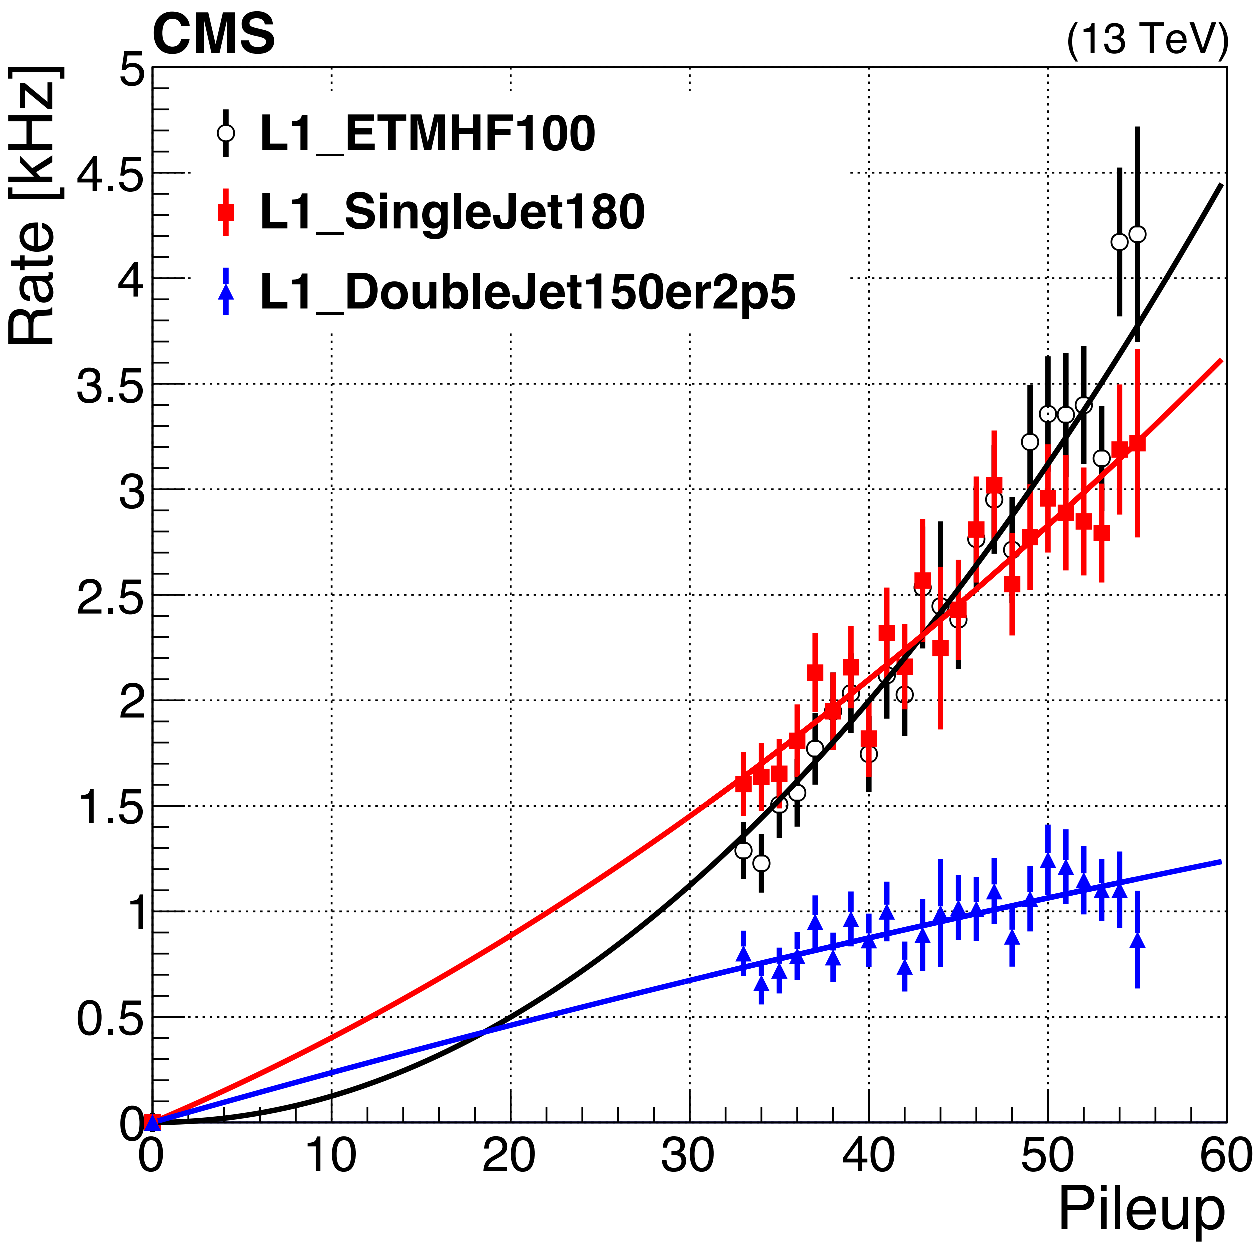

The high-level trigger

Software-based event filtering system (close-to-offline reconstruction).
Consists of 50k CPU cores and 400 GPUs.

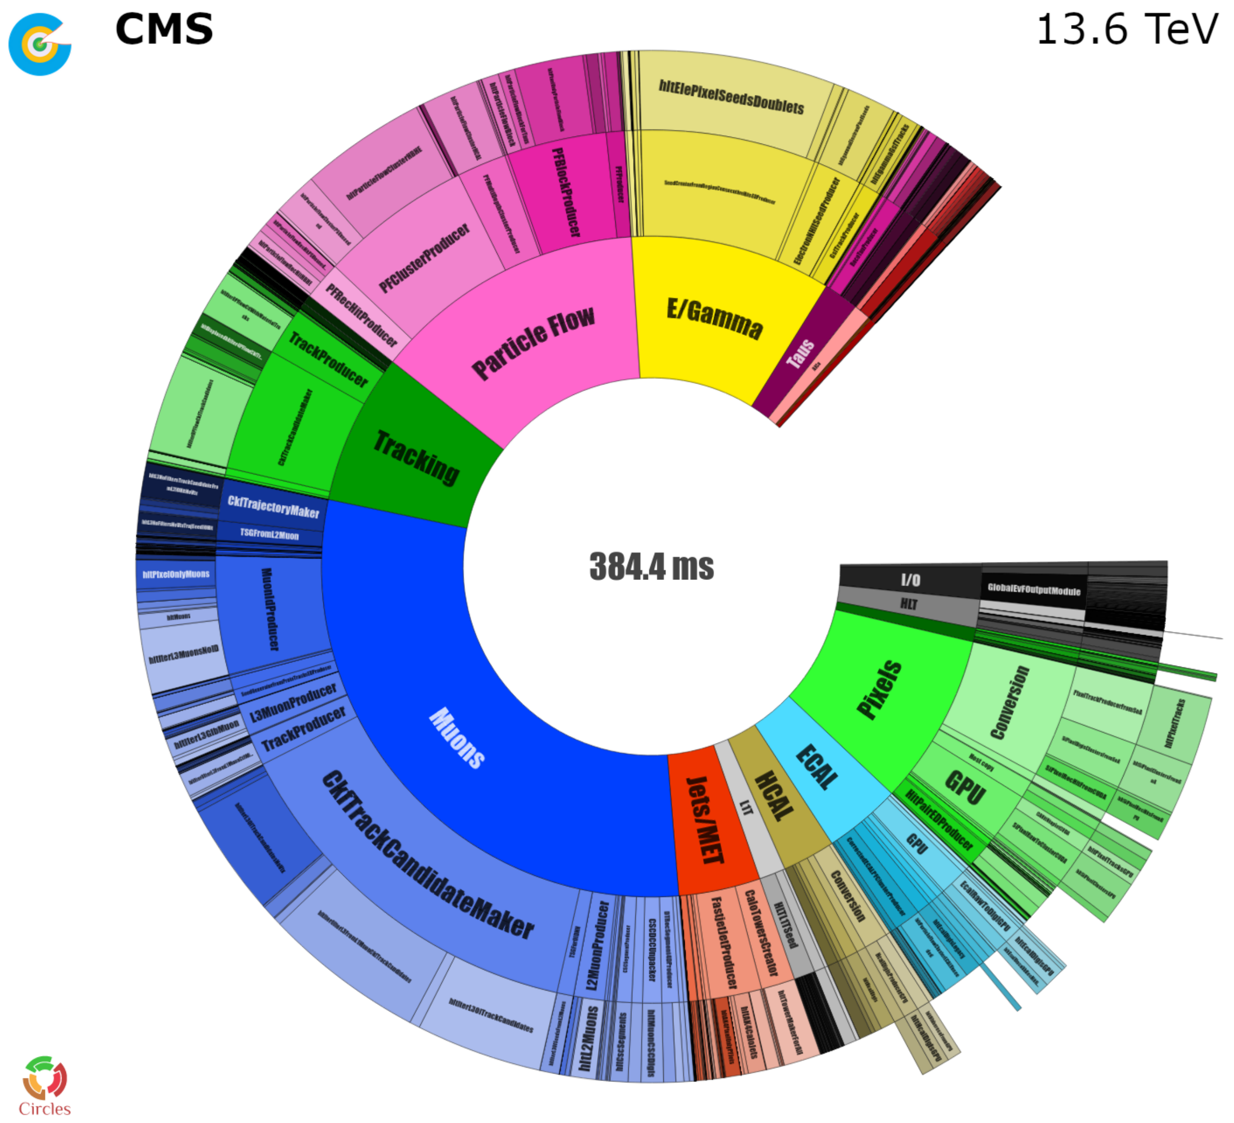

Average processing time: ~0.5 sec
Tracking offloaded to GPUs.

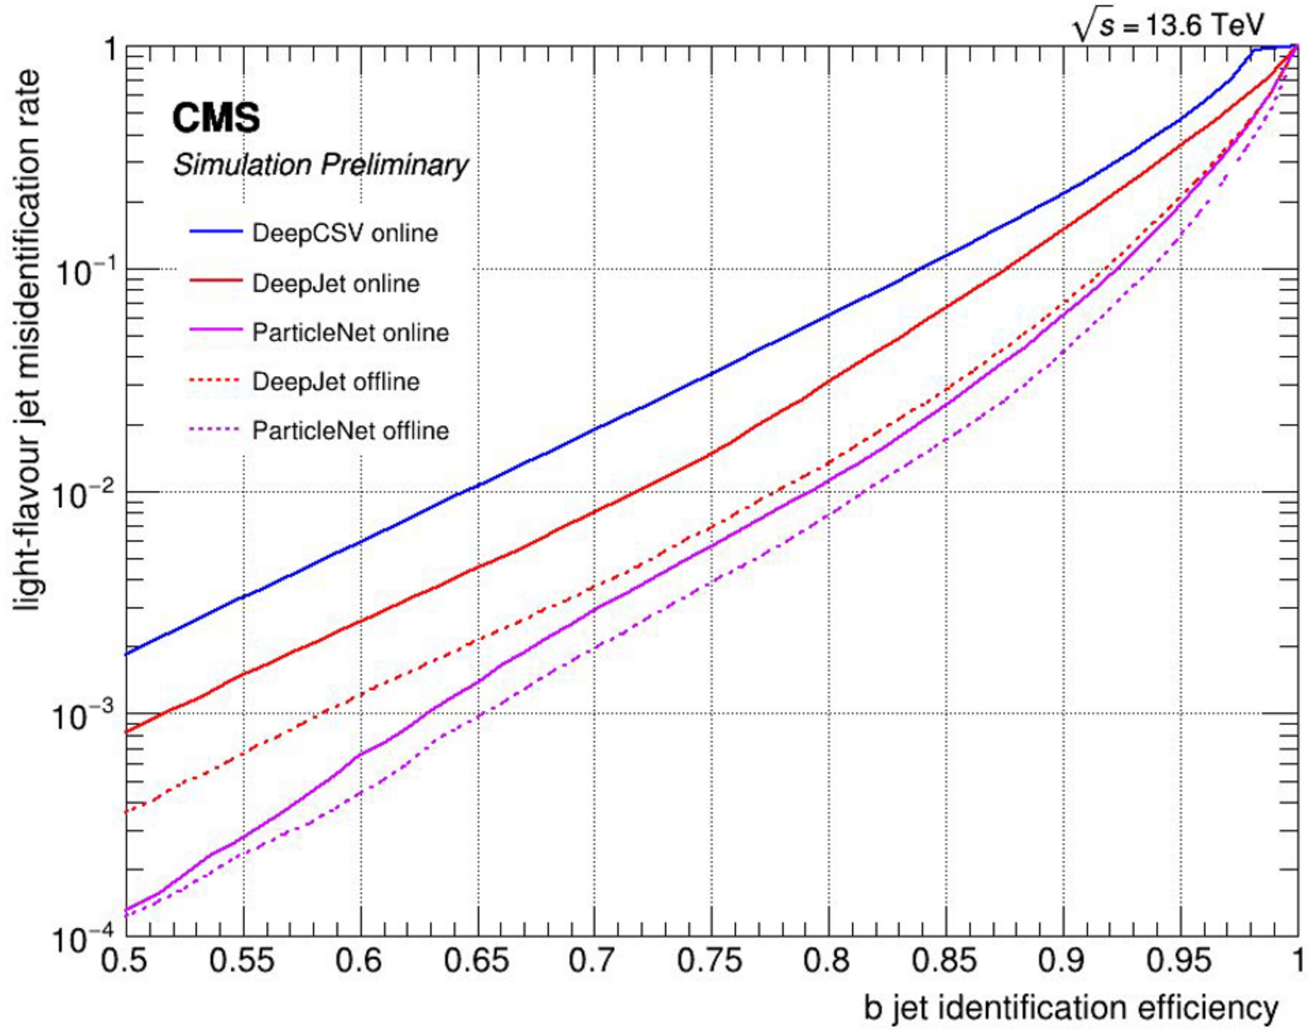

Secondary vertex reconstruction (B decays in jets) is near-offline.

The Trigger Menu

Each Level-1 trigger path can “seed” one or more HLT paths.
The full list of HLT paths is 500+ items long! Adds up to ~1000 Hz.

HLT algorithm	Rate
Isolated muon with $p_T > 24$ GeV	250 Hz
Isolated electron with $E_T > 32$ GeV	182 Hz
Particle-flow based $p_T^{miss} > 110$ GeV	81 Hz
4 PF jets with $p_T > 70/50/40/35$ GeV with 2 b tag	57 Hz
Two isolated tau leptons with $p_T > 35$ GeV	54 Hz
Muon with $p_T > 50$ GeV	51 Hz
Two electrons with $E_T > 25$ GeV	21 Hz
AK4 PF jet with $p_T > 500$ GeV	16 Hz
Two same-sign muons with $p_T > 18/9$ GeV	10 Hz

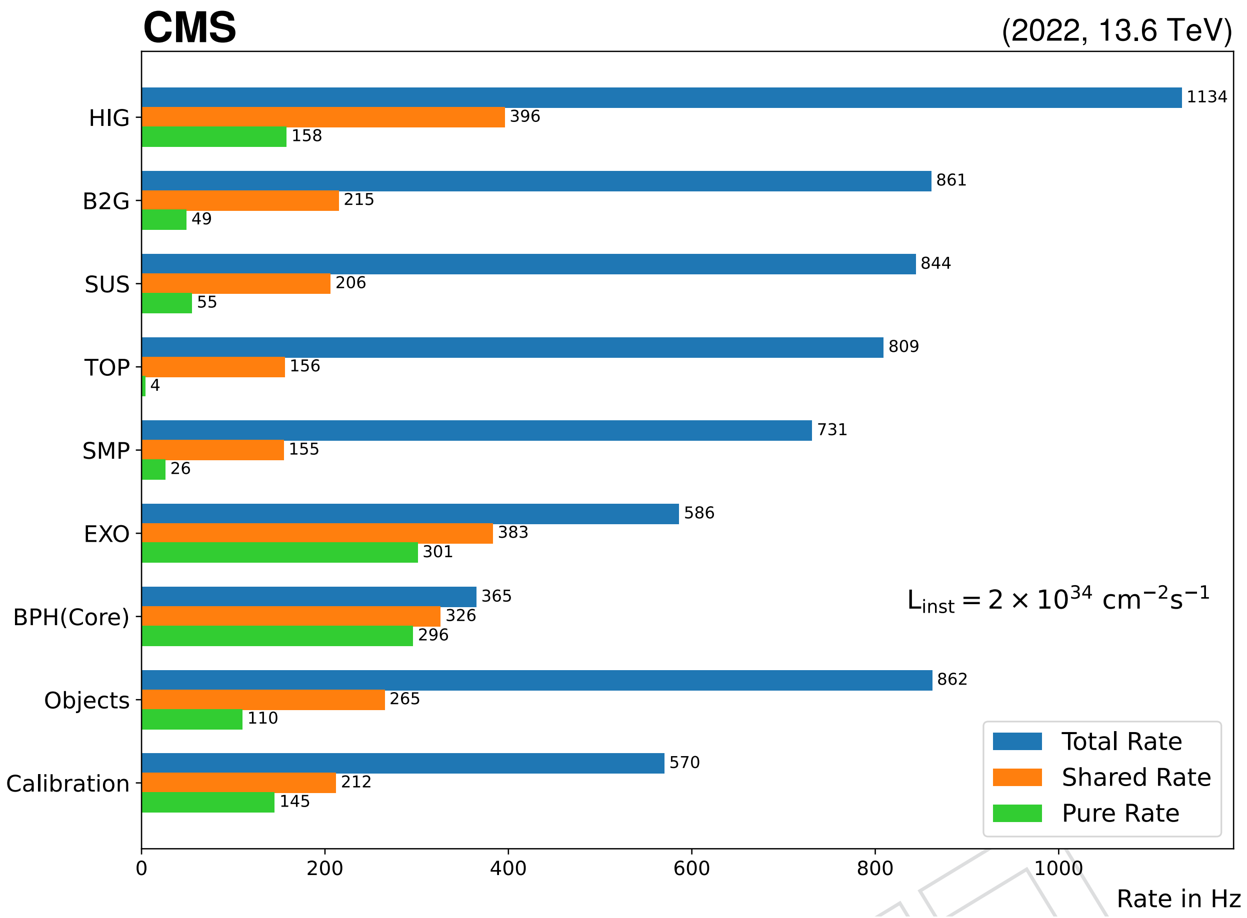

Each path can encode a long list of selection criteria!

Example: “**HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ_Mass3p8**”

Encodes most configurable information $p_T(1)$, $p_T(2)$, isolation, $|z_1\text{-}z_2|$, $m_{\mu\mu}$ but not all (e.g. muon η, dR, identification quality, …)

An analyzer’s perspective

All new searches & measurements must start with a critical question:
How am I going to record the data that contains my events??

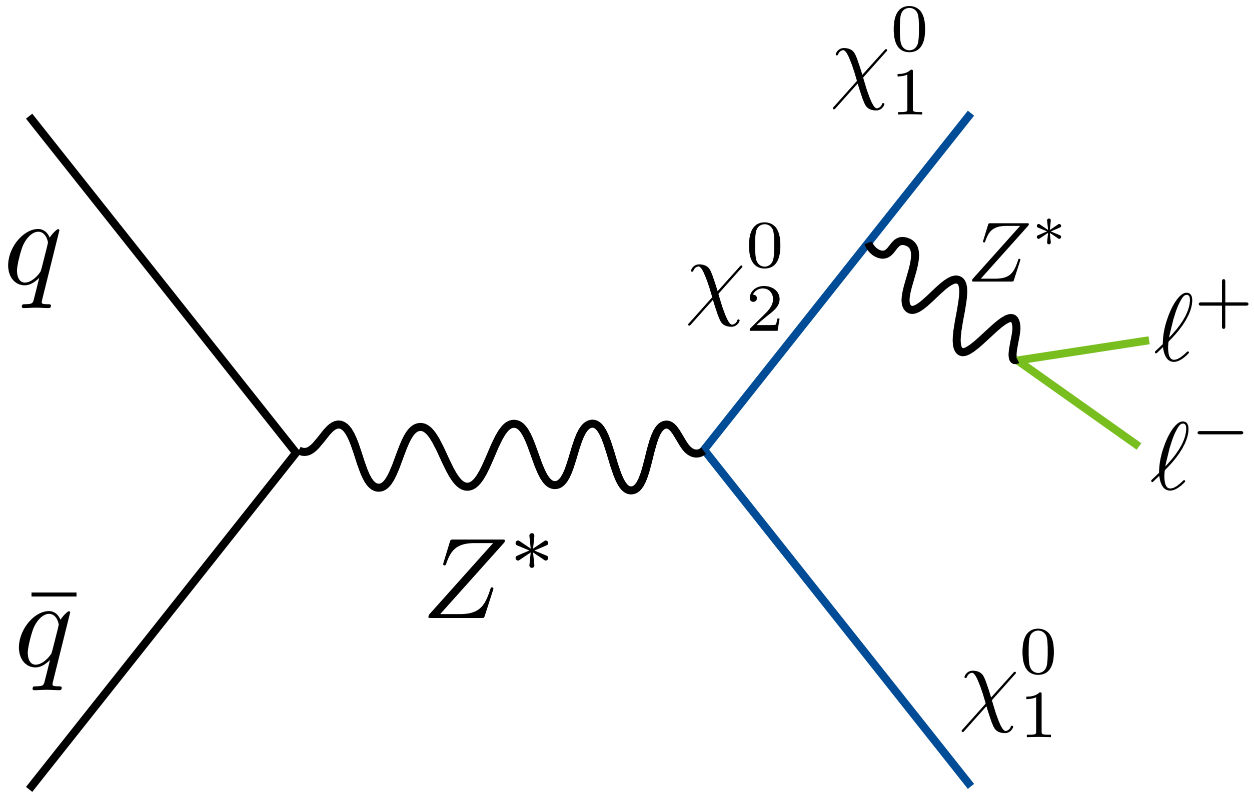

Feynman view

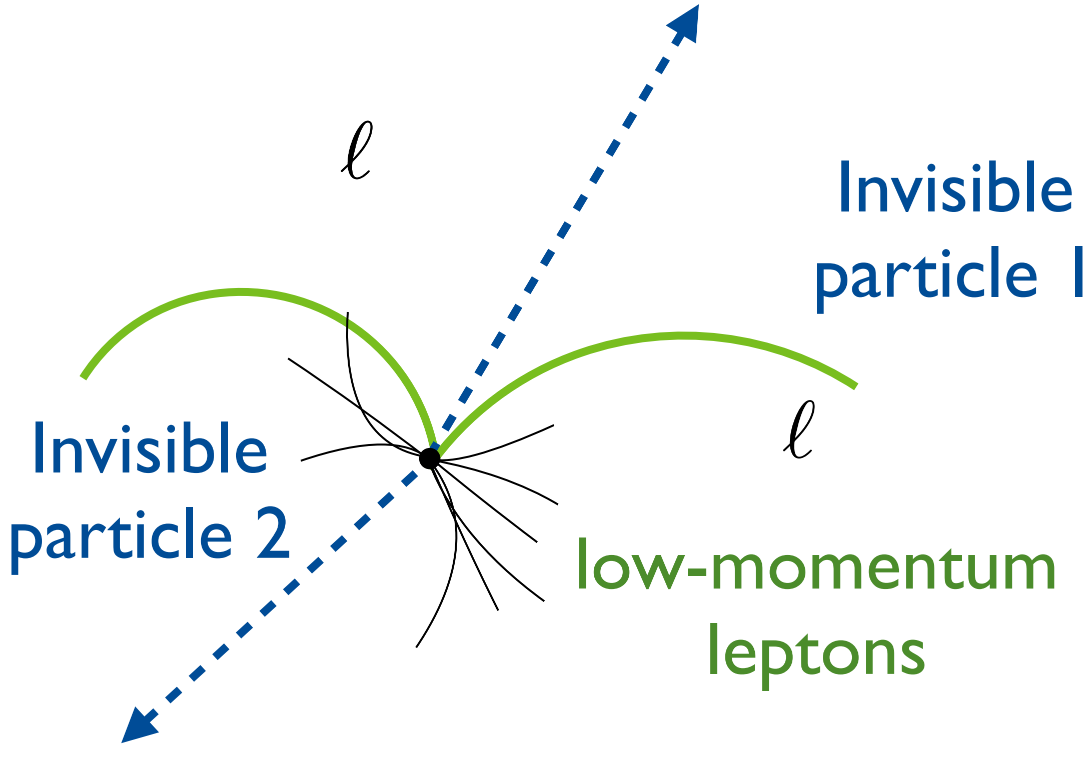

Detector view

An analyzer’s perspective

All new searches & measurements must start with a critical question:
How am I going to record the data that contains my events??

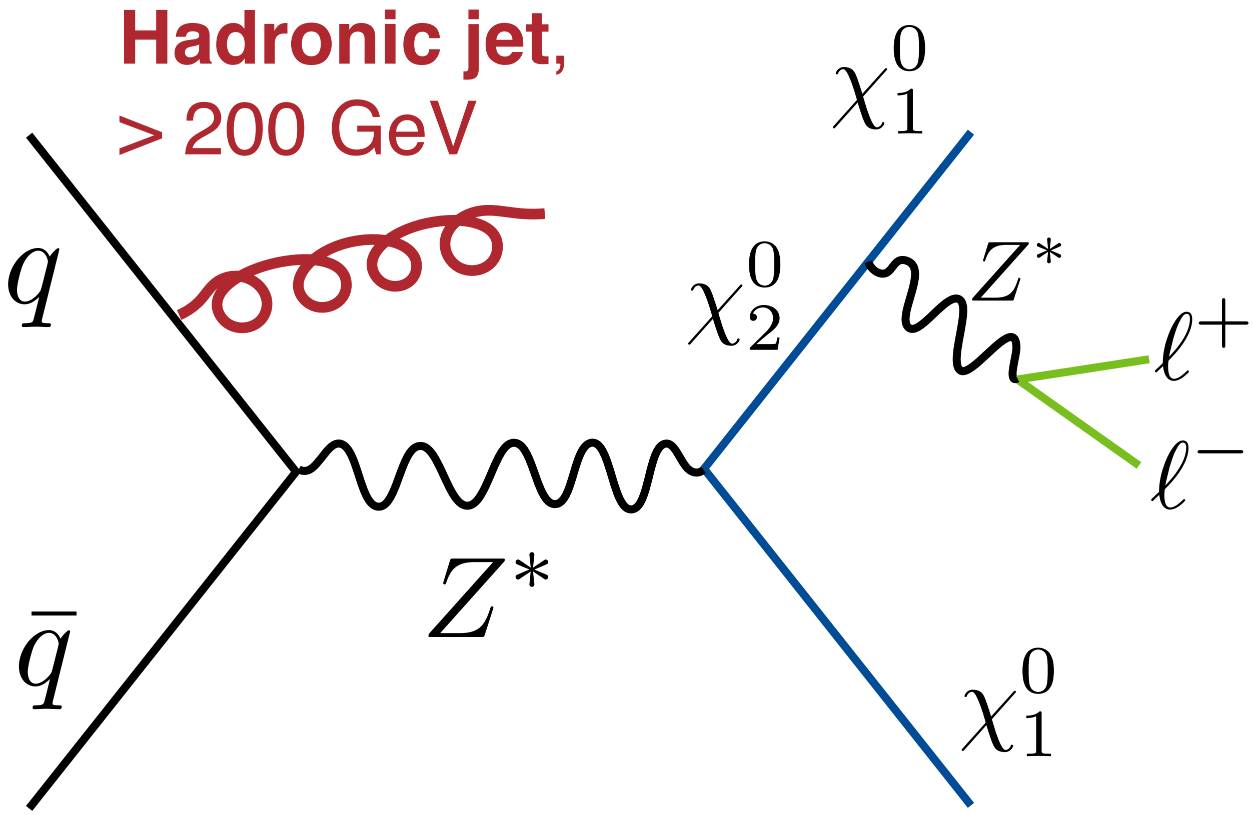

Feynman view

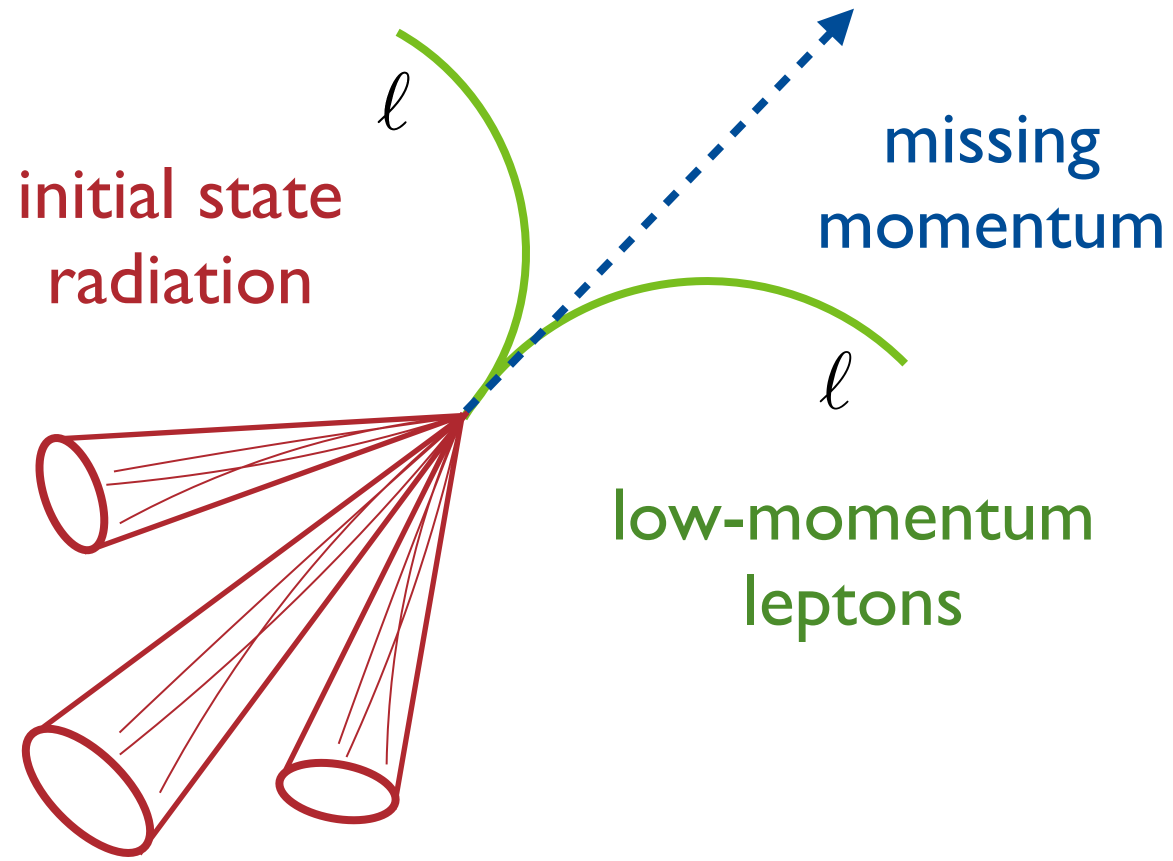

Detector view

Trigger shapes search: look at sub-set of events with a high-energy jet.
Boost the Dark Matter candidates to create “missing momentum”.

How do existing triggers perform?

Can test the standard missing-p_T trigger, measure performance in data.

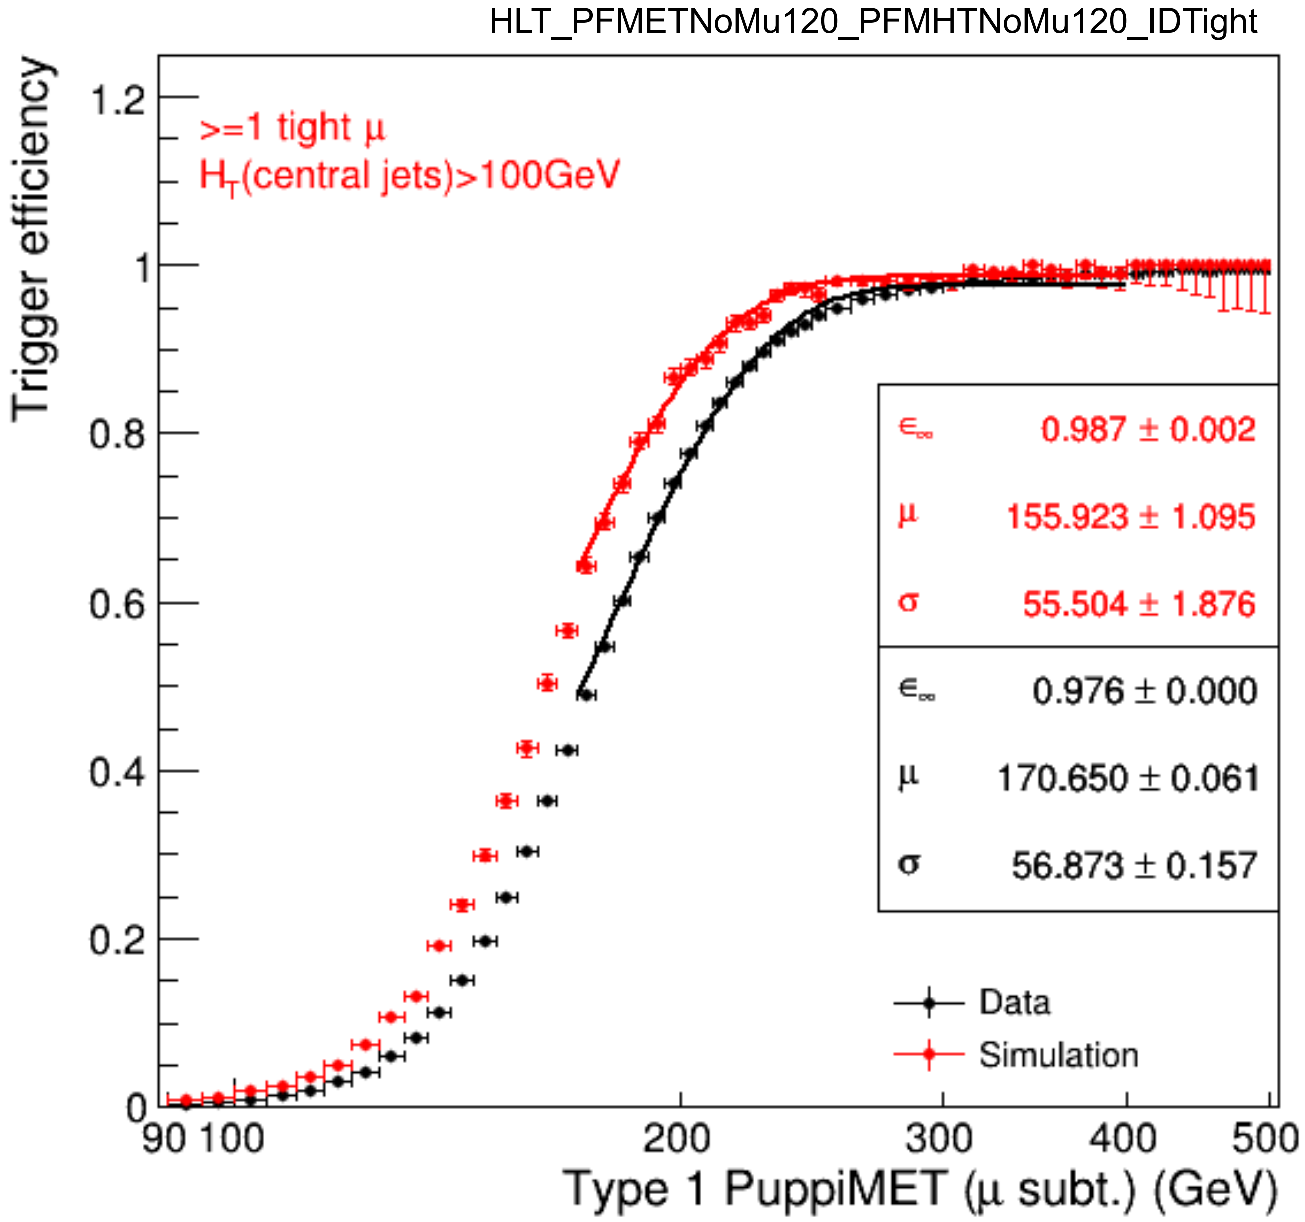

Make an unbiased measurement by selecting events with an orthogonal trigger (W→μν here).

Not perfectly modeled in simulation

Must be corrected to ensure we predict the expected number of events properly!

Still, the efficiency is quite low below 200 GeV. **Can we do better?**

Designing a new trigger path

Specific “cross-triggers” can exploit selections on multiple objects:

L1_DoubleMu3_SQ_ETMHF50_Jet60er2p5_OR_DoubleJet40er2p5

HLT_DoubleMu3_DZ_PFMET50_PFMHT60

Requires: 2 muons with $p_T > 3$ GeV and p_T-missing > 60 GeV.

Pro: recorded significantly more events for the analysis!

Con: trigger is COMPLICATED! Many parts of the efficiency to measure!

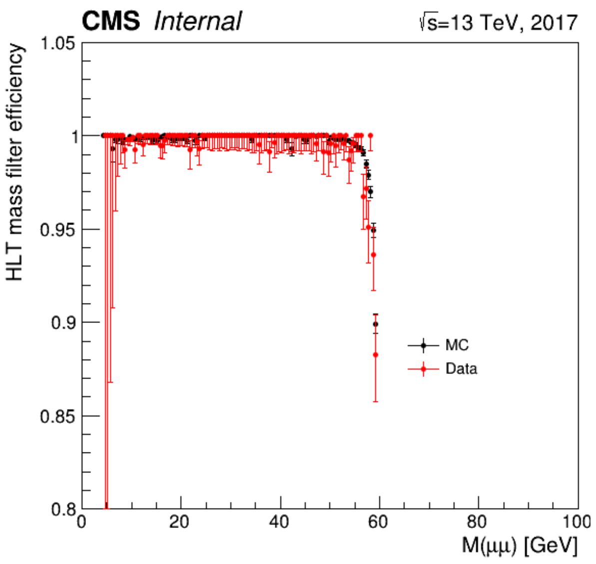

m(μμ)

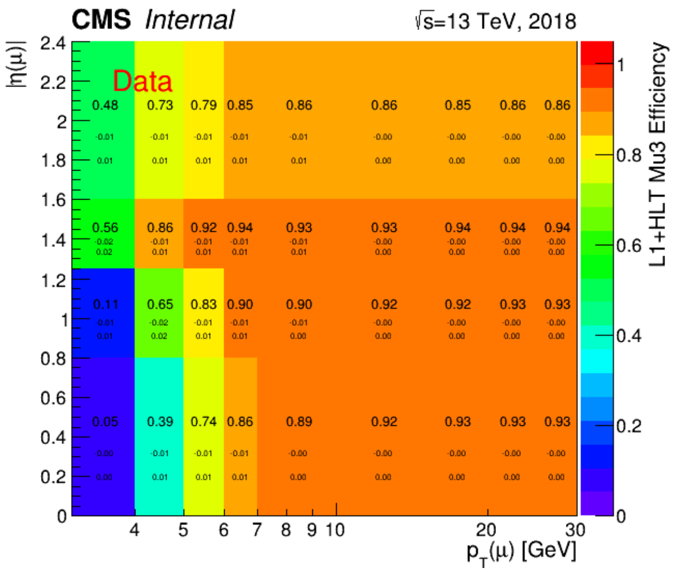

p_T, |η|(μ)

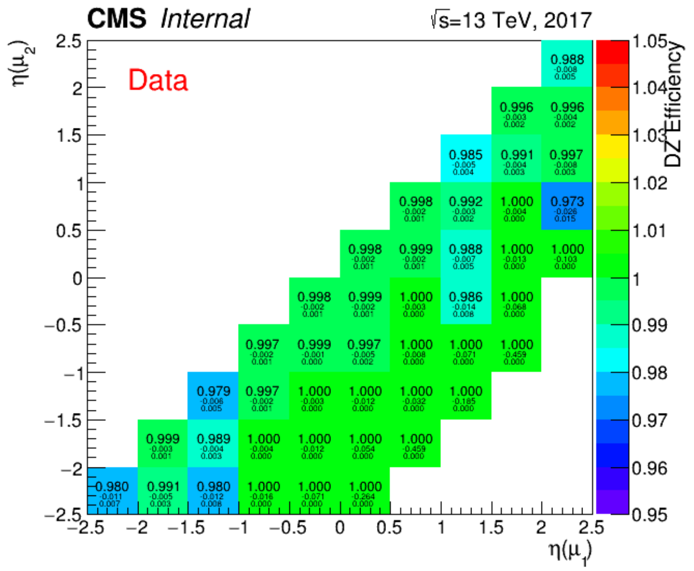

dz(μ_1,μ_2)

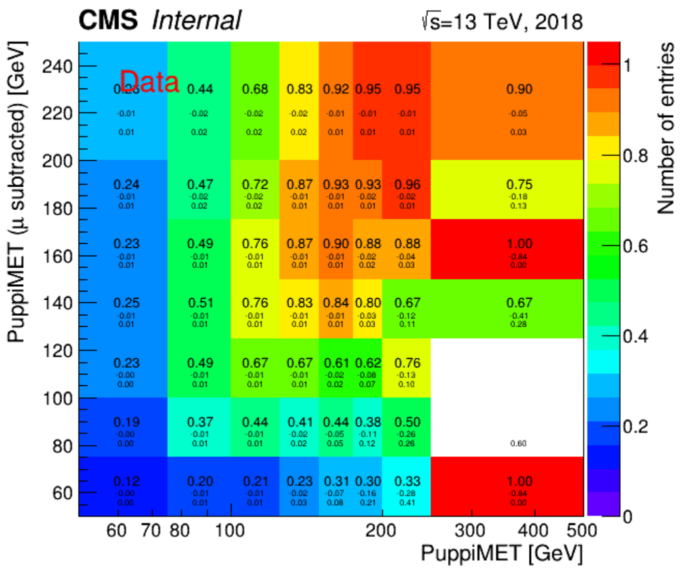

p_T-miss

So… how did we collect 10^{15} events?

Remember, we needed 10^{15} *pp* interactions to make 200k Higgs by 2012.

- Pack each bunch-crossing with 20 collisions / event
- Level-1 Trigger reduces rate from 40MHz to 100kHz.

Up to this point:

  - The total # of events is still 125 Billion
  - The total data rate is 100 GB/sec

- The High-Level Trigger farm (50k CPU cores) processes ≥2 events / second to “keep up” with the Level-1 trigger output.
  - HLT filters keeps 1 in 100 events, with the final 1kHz written to disk.

This means that “only” 1.25 Billion events are stored to disk.

  - An offline CPU farm of similar to HLT can reconstruct this in only $2.5 \cdot 10^4$ seconds (~7 hours)!

LHC / HL-LHC Plan

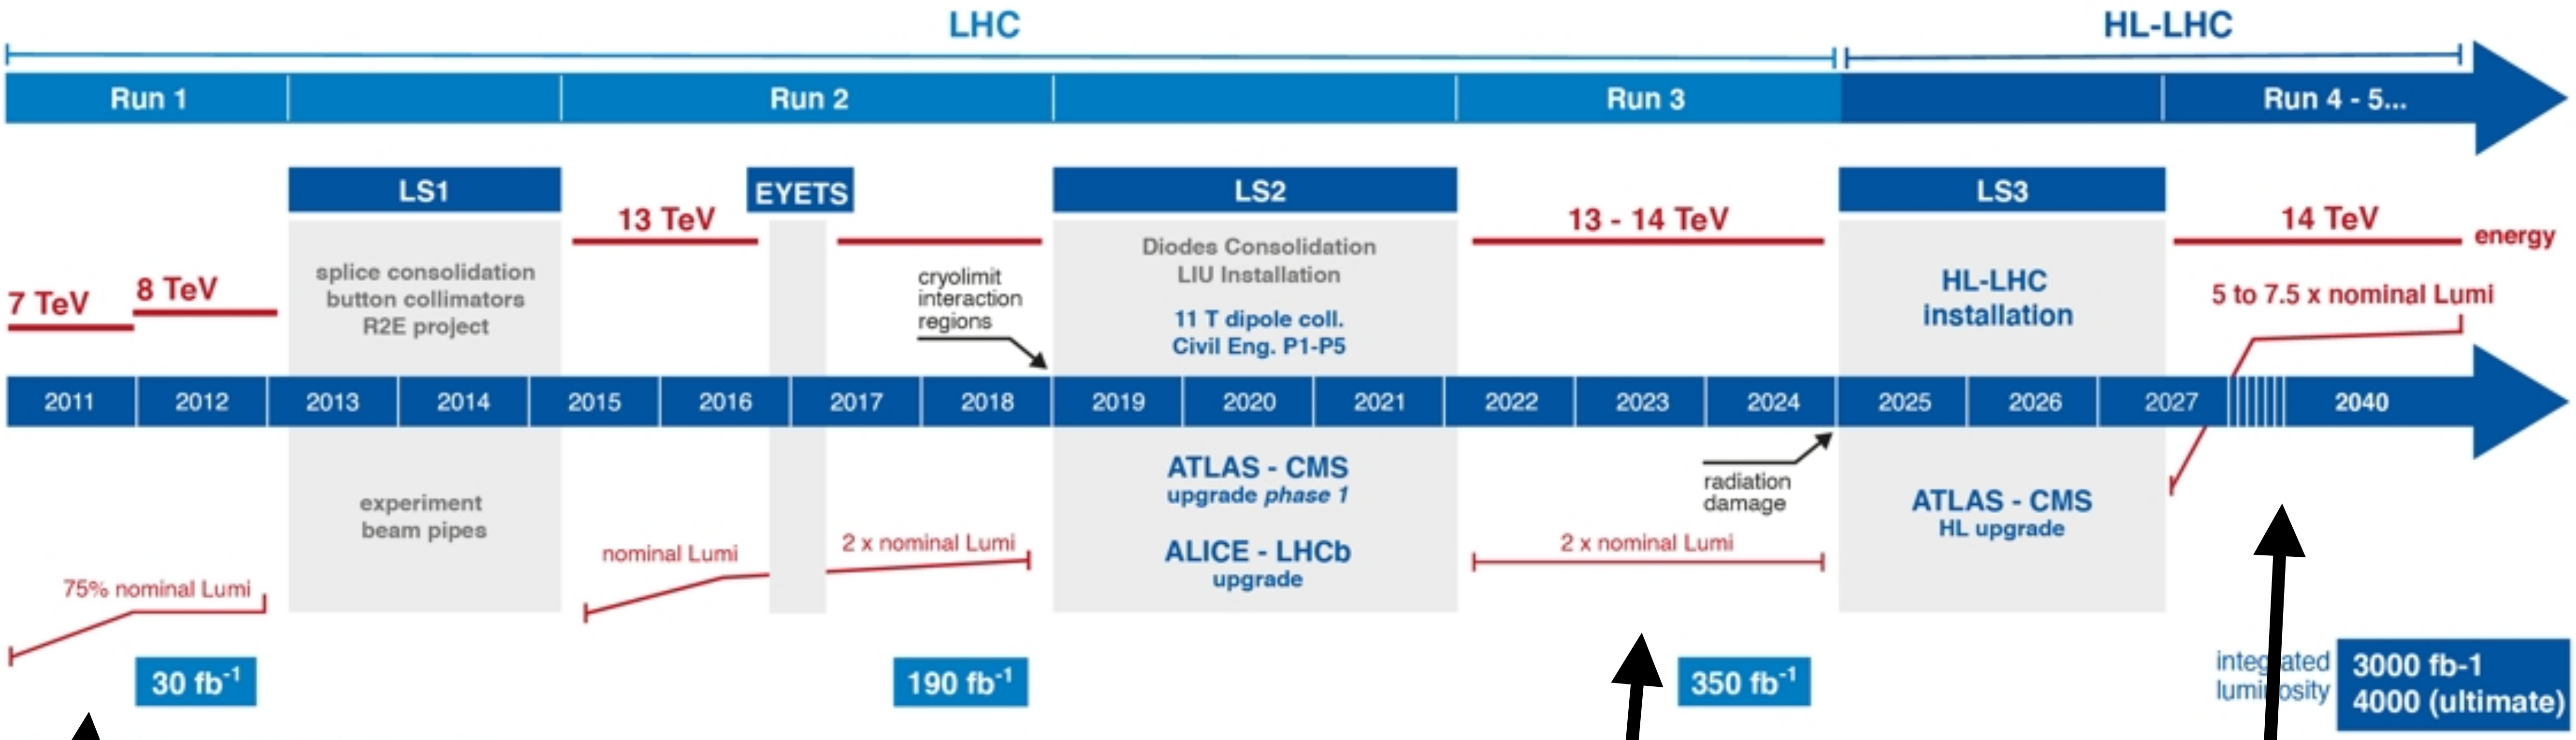
LHC

HL-LHC

Run 1

Run 2

Run 3

Run 4 - 5...

LS1

EYETS

LS2

LS3

7 TeV

8 TeV

13 TeV

13 - 14 TeV

14 TeV

energy

splice consolidation
button collimators
R2E project

cryolimit
interaction
regions

Diodes Consolidation
LIU Installation

11 T dipole coll.
Civil Eng. P1-P5

HL-LHC
installation

5 to 7.5 x nominal Lumi

2011 2012 2013 2014 2015 2016 2017 2018 2019 2020 2021 2022 2023 2024 2025 2026 2027 2040

experiment
beam pipes

ATLAS - CMS
upgrade *phase 1*

ALICE - LHCb
upgrade

radiation
damage

ATLAS - CMS
HL upgrade

75% nominal Lumi

nominal Lumi

2 x nominal Lumi

2 x nominal Lumi

30 fb^{-1}

190 fb^{-1}

350 fb^{-1}

integrated
luminosity

3000 fb-1
4000 (ultimate)

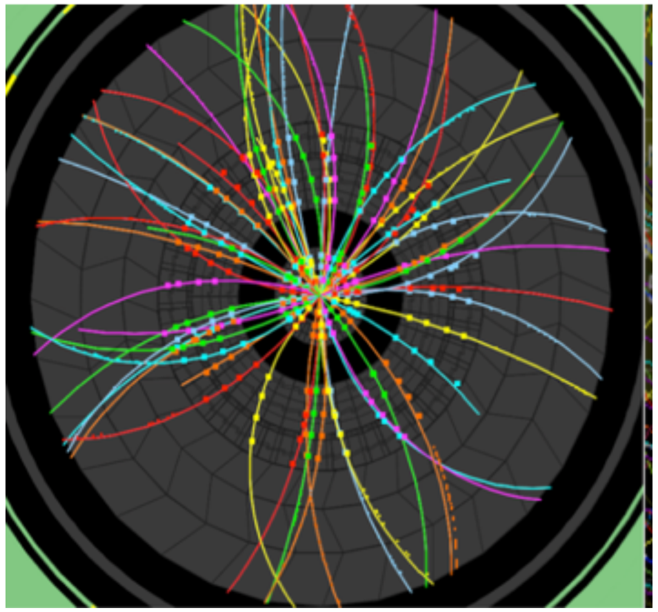

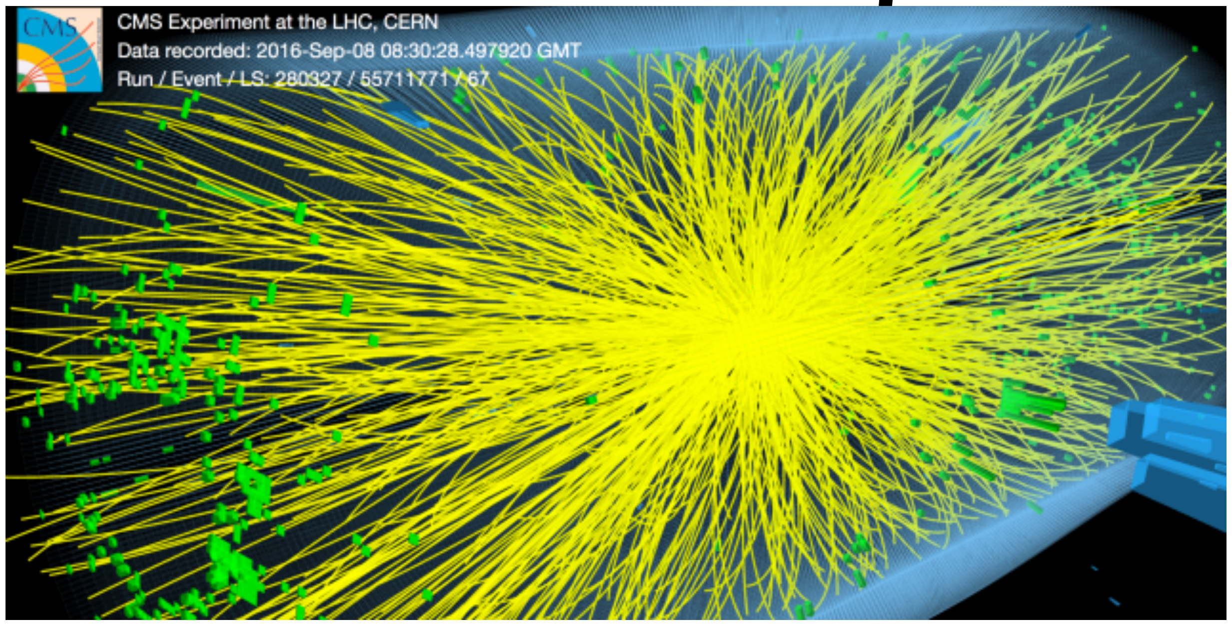

?

2 collisions

86 collisions

200 collisions!

Trigger-driven tracker design

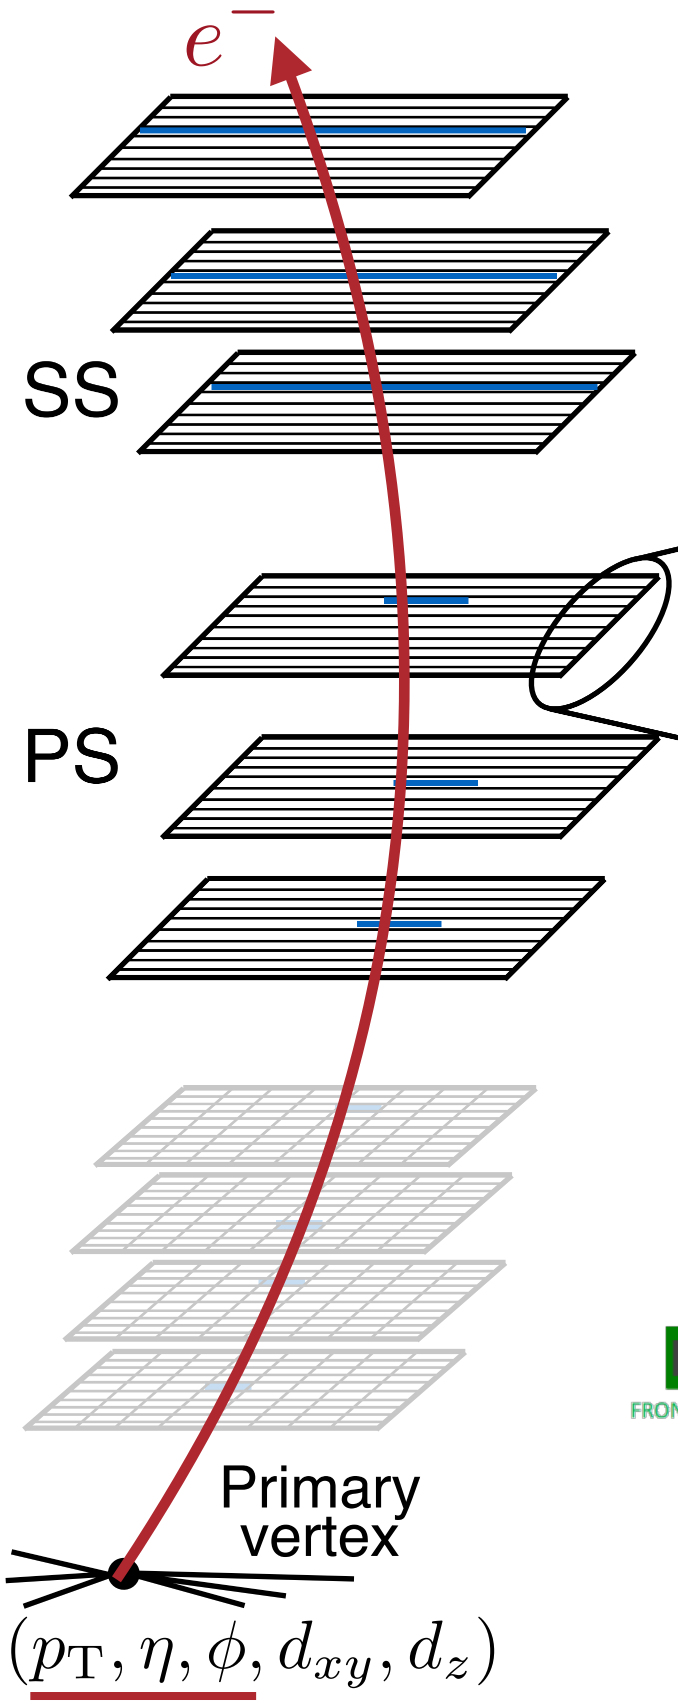

Large-radius sensors drive p_T measurement (lever arm).

Outer layers: 2 stacked sensors with 5cm strips “SS”.

Inner layers: strips (2.4cm) + macro-pixel (1.5mm) “PS”.

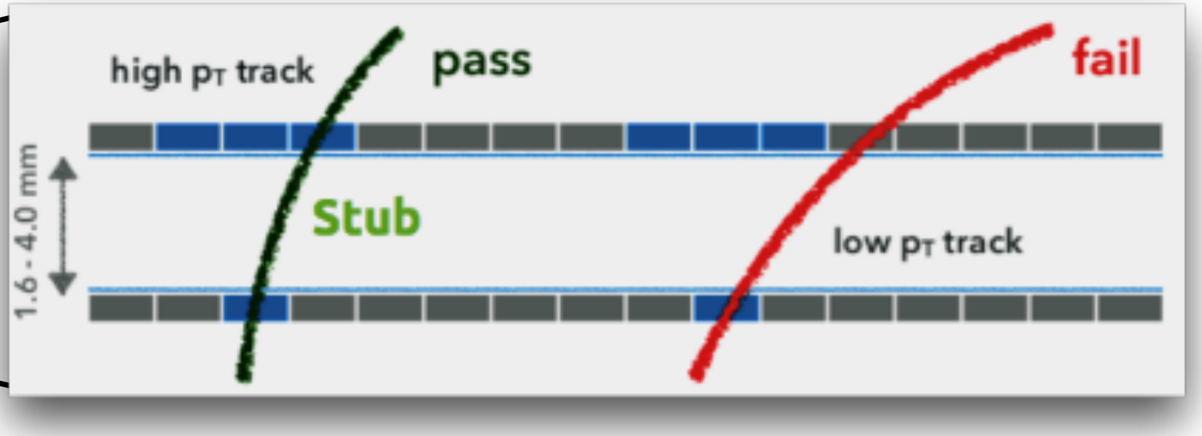

Double-layer strip modules provide local p_T measurement.

→ Intrinsic mechanism to filter hits from low-p_T tracks, allows high-p_T (2 GeV) track-finding in the trigger system!

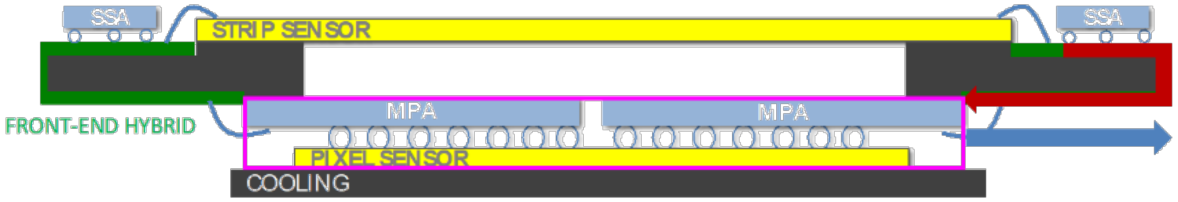

Schematic and prototype of a “MacroPixel+Strip” module

CMS Trigger Design (2023)

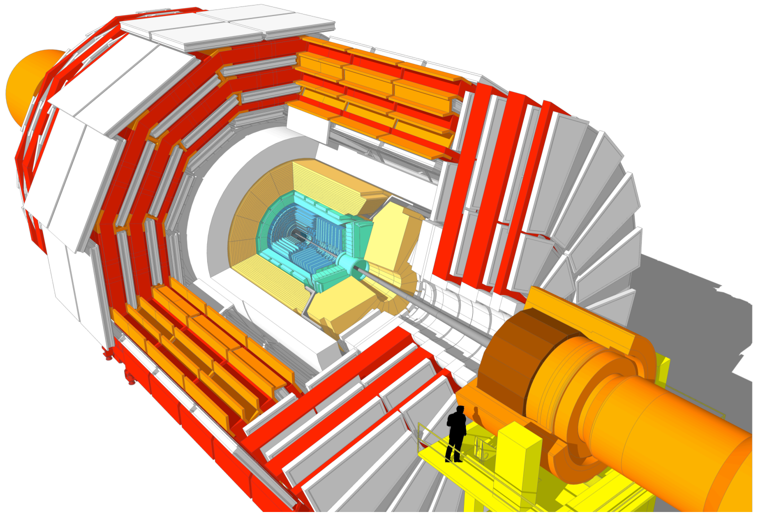

Data Acquisition path
(High-resolution data)

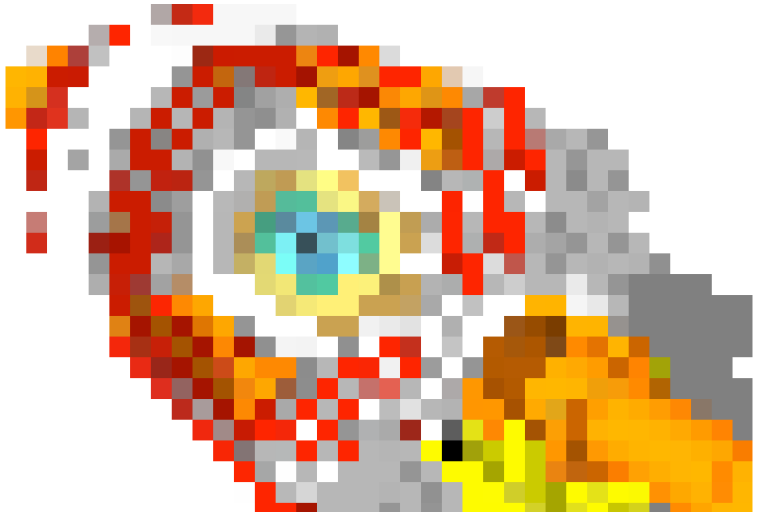

Trigger path
(Low-resolution data)
@ 40 MHz

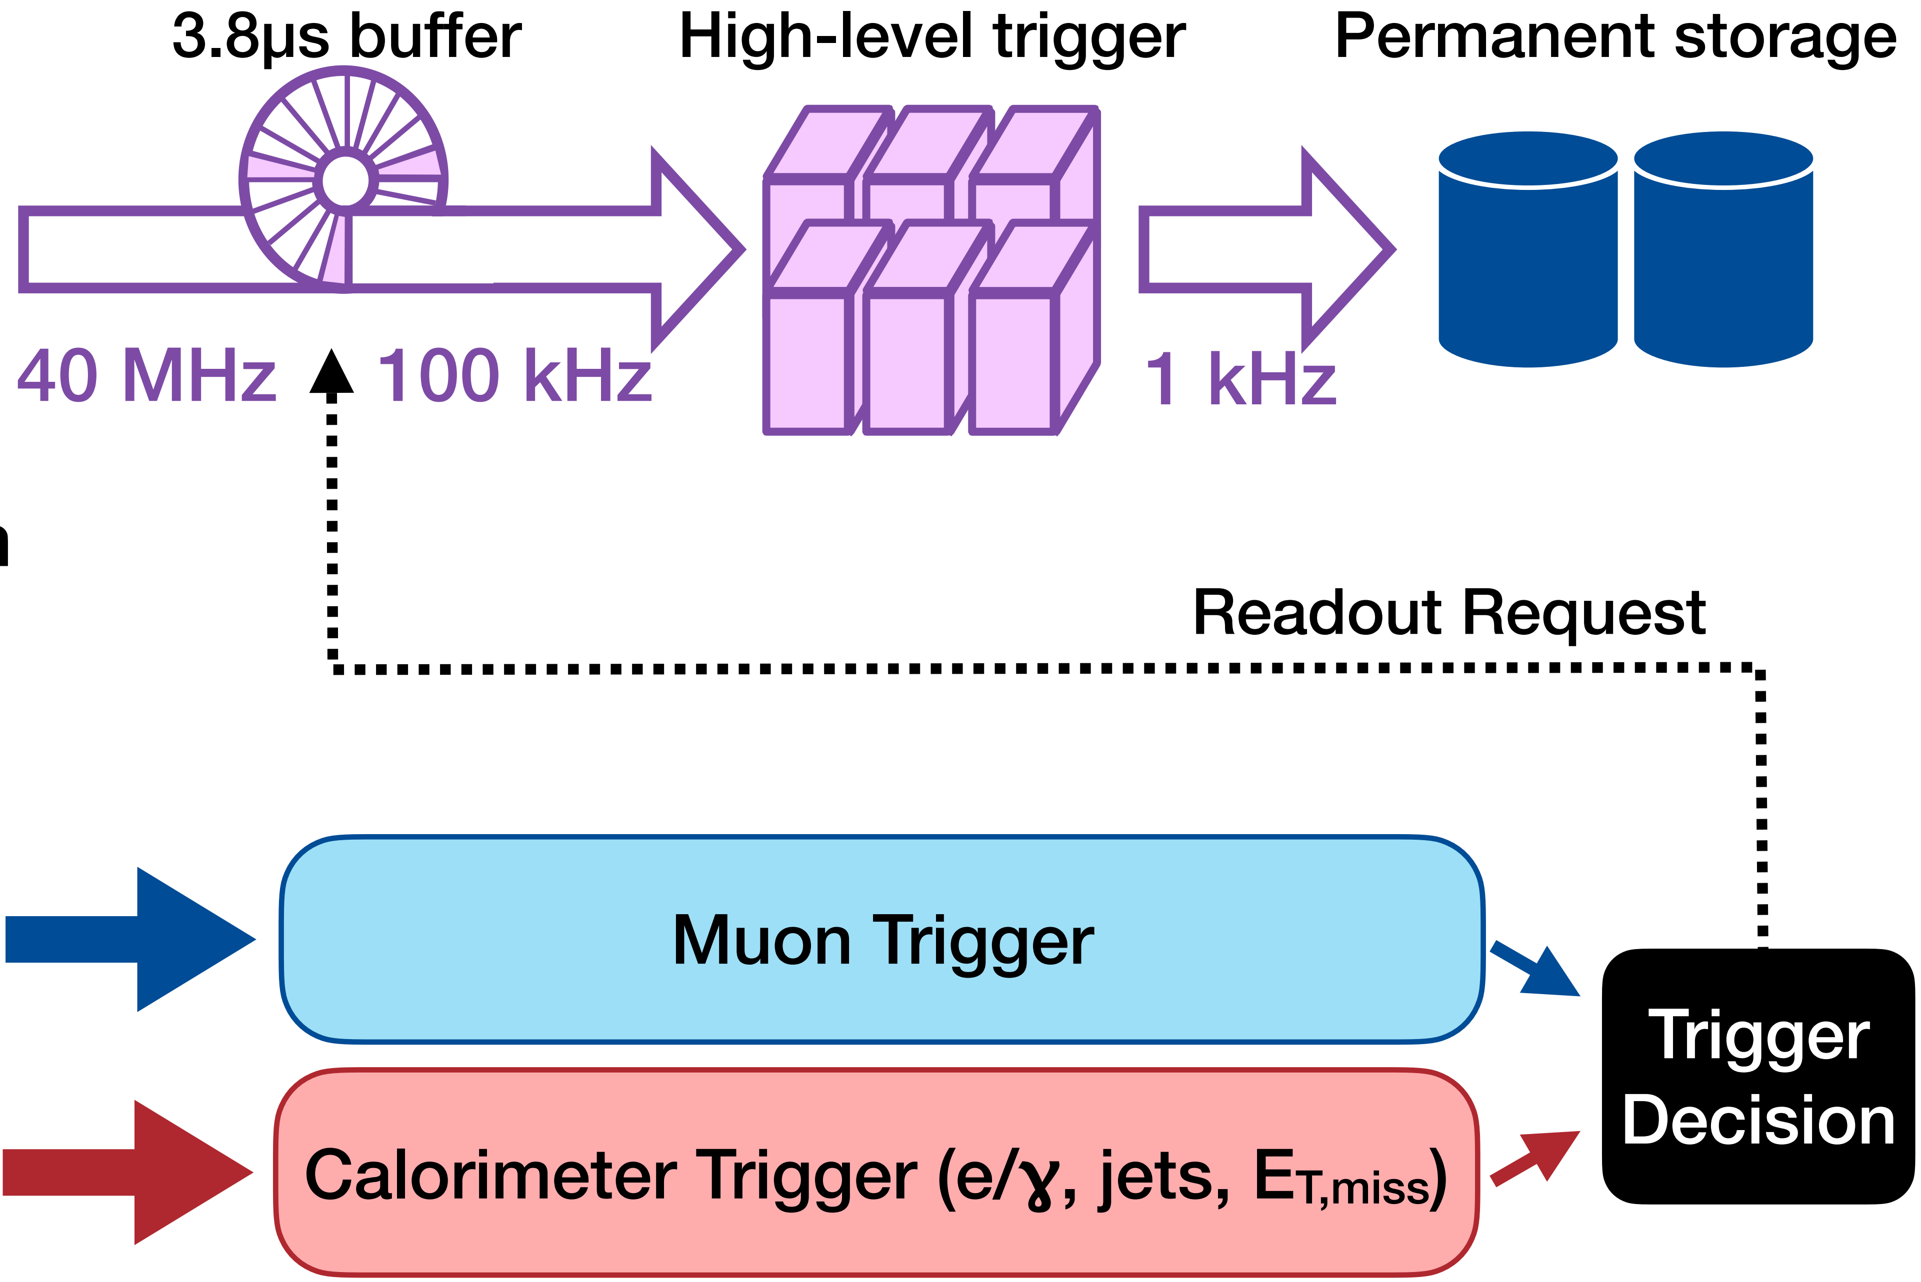

3.8μs buffer

High-level trigger

Permanent storage

40 MHz

100 kHz

1 kHz

Readout Request

Muon Trigger

Calorimeter Trigger (e/ɣ, jets, $E_{T,miss}$)

Trigger Decision

CMS Trigger Design (2023)

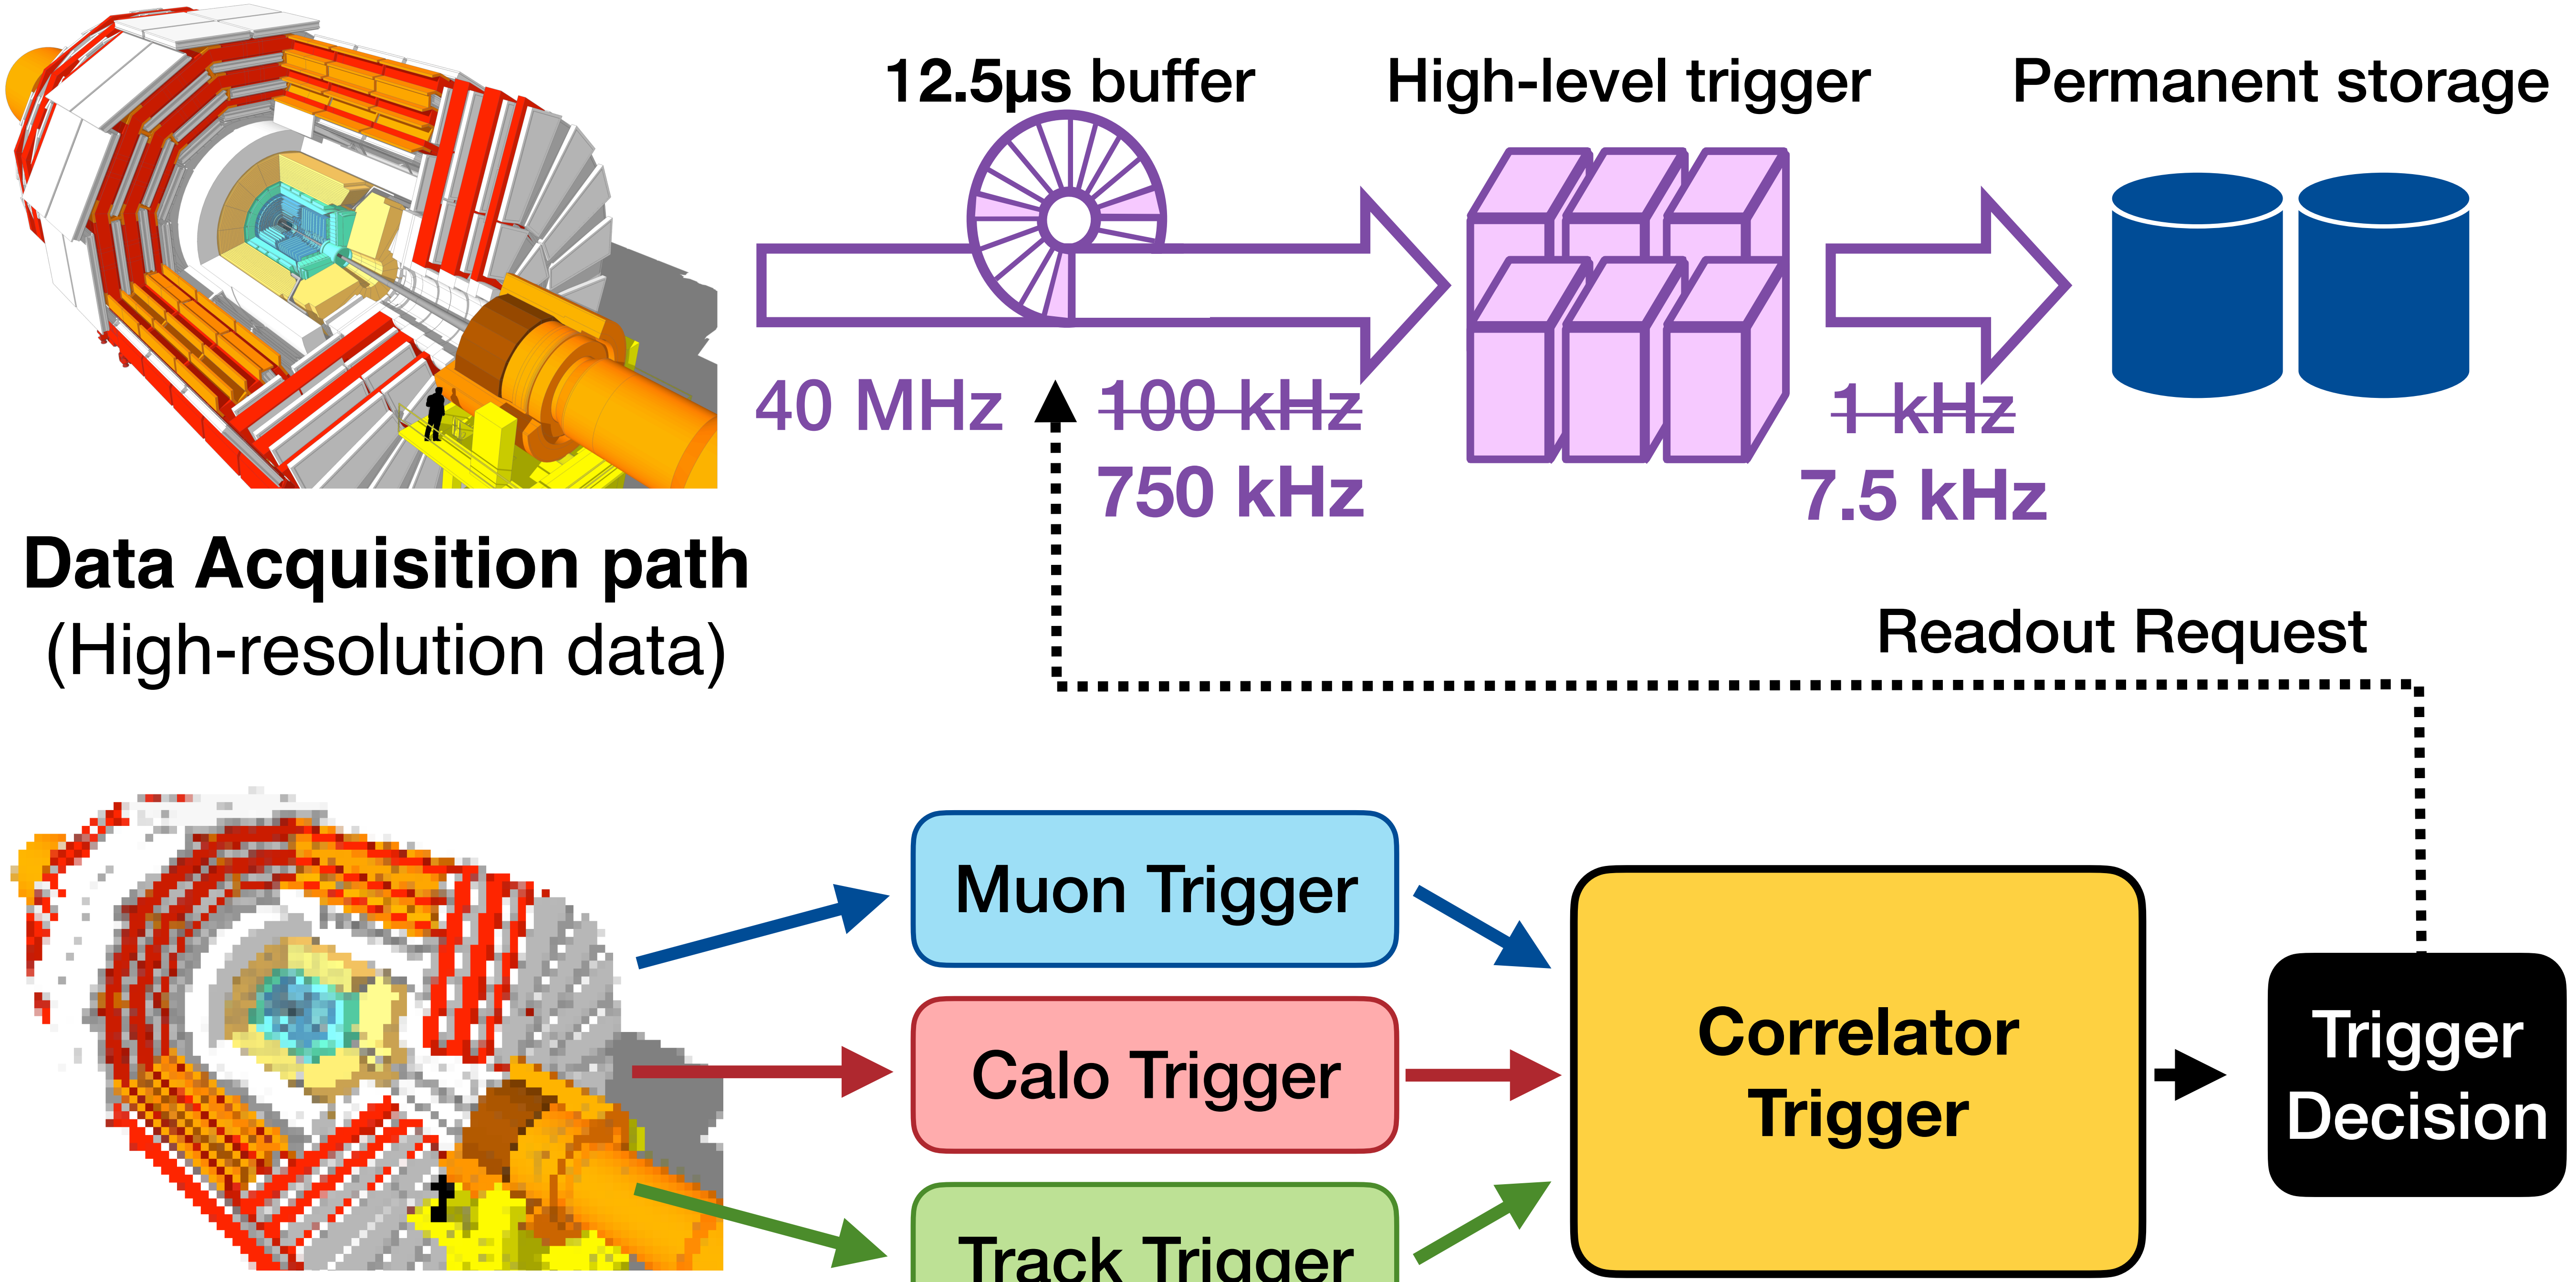

Data Acquisition path
(High-resolution data)

12.5μs buffer

High-level trigger

Permanent storage

40 MHz

~~100 kHz~~
750 kHz

~~1 kHz~~
7.5 kHz

Readout Request

Muon Trigger

Calo Trigger

Track Trigger

Correlator Trigger

Trigger Decision

Trigger path
(Low-resolution data)
@ 40 MHz

Complete Particle Reconstruction & Analysis

Level-1 Trigger system

One of many trigger processing boards:

- 120 x 28 Gbps optical links
- UltraScale+ FPGA (12k DSP slices)

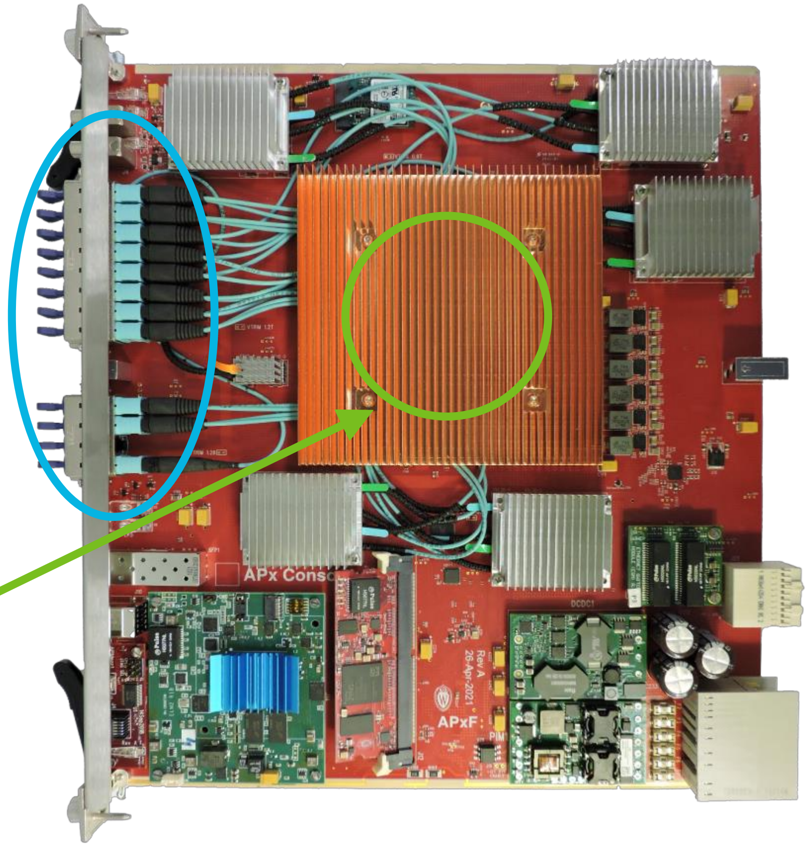

Example: Correlator Trigger

→ 66 processing boards

→ 6-18x time-multiplex factor

→ 1+1.5μs latency budget

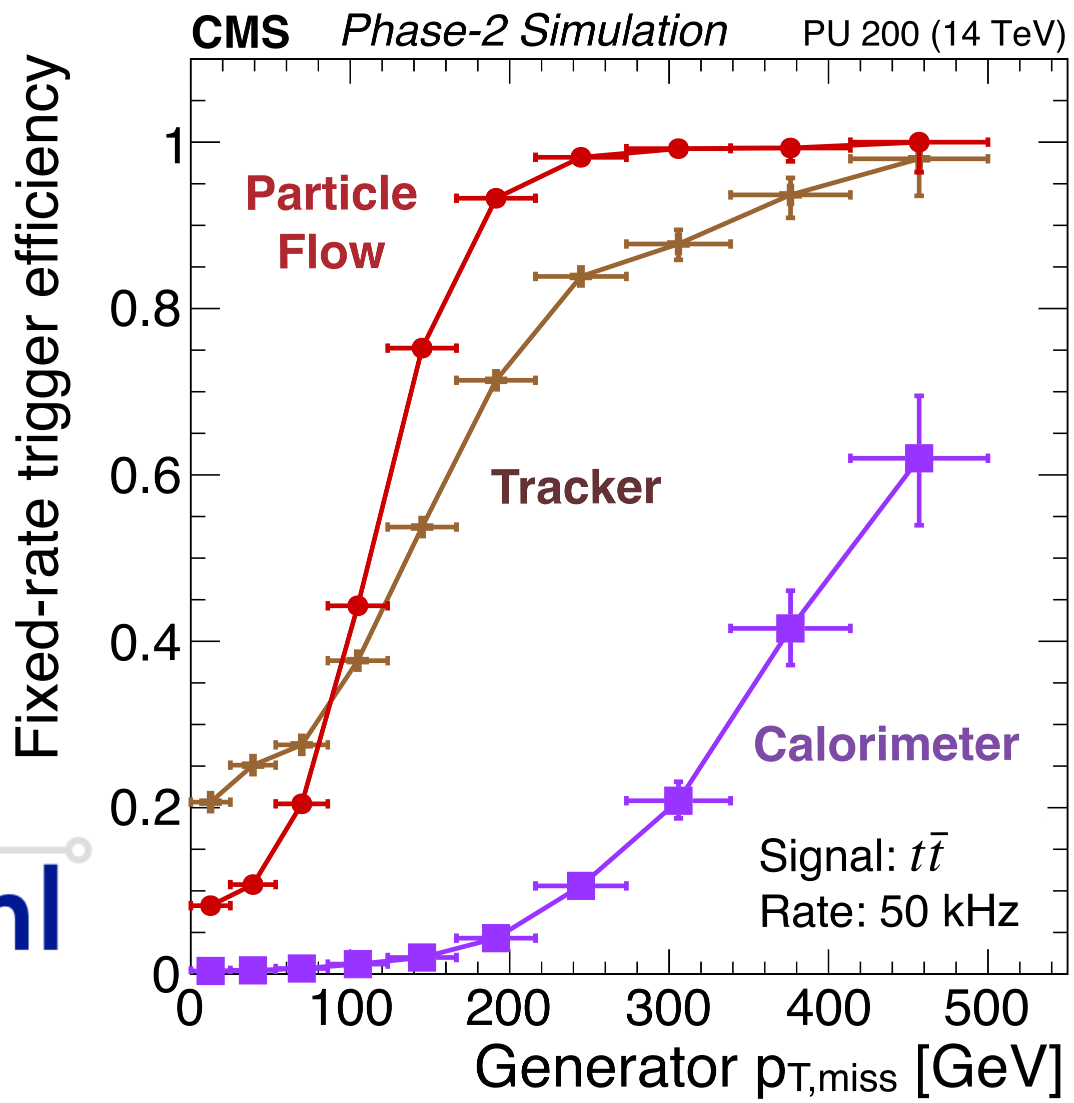

Building all particles allows for complex algos:

Jet-finding, τ ID, missing-p_T w/ neural nets,…

HLS: physicists can program FPGAs w/ C++

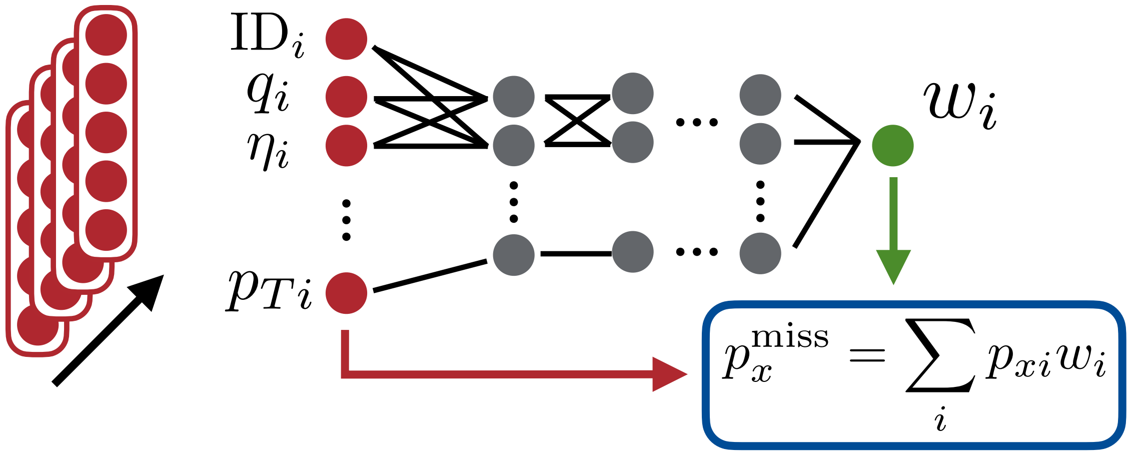

$$p_x^{\mathrm{miss}} = \sum_i p_{xi} w_i$$

hls4ml

Trigger will enable HL-LHC physics!

High-Luminosity data will dramatically enrich the physics potential of CMS.

Highlights: precision Higgs program & many opportunities for New Physics!

The upgraded Trigger System is critical to unlocking its power.

Trigger is the first step of your physics analysis!

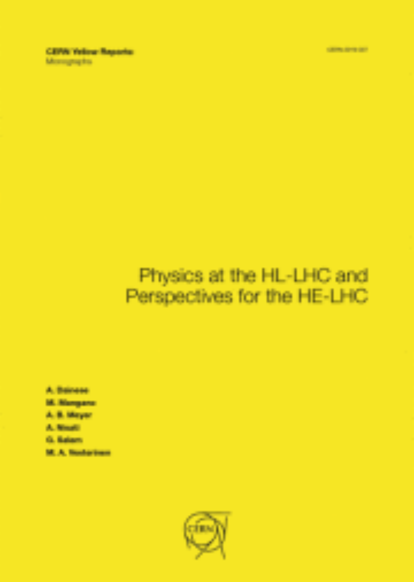

Physics at the HL-LHC and Perspectives for the HE-LHC

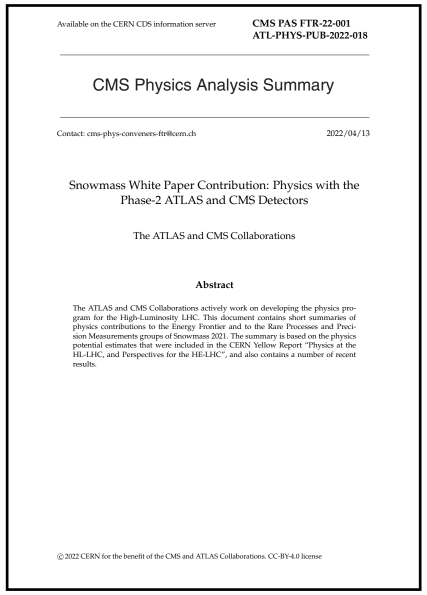

Available on the CERN CDS information server

CMS PAS FTR-22-001
ATL-PHYS-PUB-2022-018

CMS Physics Analysis Summary

Contact: cms-phys-conveners-ftr@cern.ch 2022/04/13

Snowmass White Paper Contribution: Physics with the Phase-2 ATLAS and CMS Detectors

The ATLAS and CMS Collaborations

Abstract

The ATLAS and CMS Collaborations actively work on developing the physics program for the High-Luminosity LHC. This document contains short summaries of physics contributions to the Energy Frontier and to the Rare Processes and Precision Measurements groups of Snowmass 2021. The summary is based on the physics potential estimates that were included in the CERN Yellow Report “Physics at the HL-LHC, and Perspectives for the HE-LHC”, and also contains a number of recent results.

© 2022 CERN for the benefit of the CMS and ATLAS Collaborations. CC-BY-4.0 license

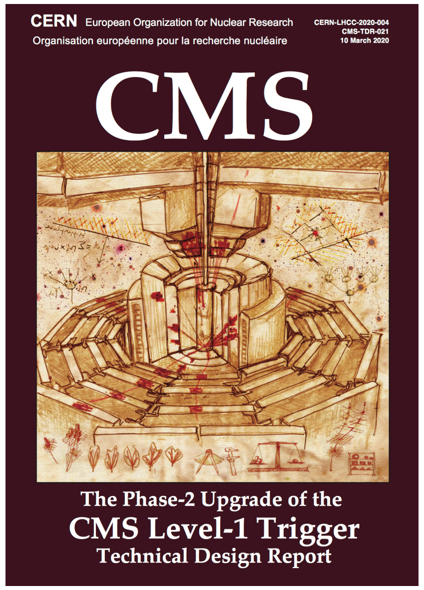

CERN European Organization for Nuclear Research
Organisation européenne pour la recherche nucléaire

CERN-LHCC-2020-004
CMS-TDR-021
10 March 2020

CMS

The Phase-2 Upgrade of the
CMS Level-1 Trigger
Technical Design Report

For more details:

→ CERN Yellow Report

→ ATLAS/CMS Snowmass reports

→ Trigger TDRs

Thanks for listening!