Fermilab

Trigger and DAQ challenges at the LHC

Christian Herwig

July 22, 2024

About me

Grew up in coastal Massachusetts

About me

Toronto
MICHIGAN
Chicago
NEW HAMPSHIRE
NEW YORK
MASSACHUSETTS
CT
RI
New York
PENNSYLVANIA
Philadelphia
OHIO
DIANA
Cincinnati
NJ
MARYLAND
DE
WEST VIRGINIA
THE UNIVERSITY OF CHICAGO
1890

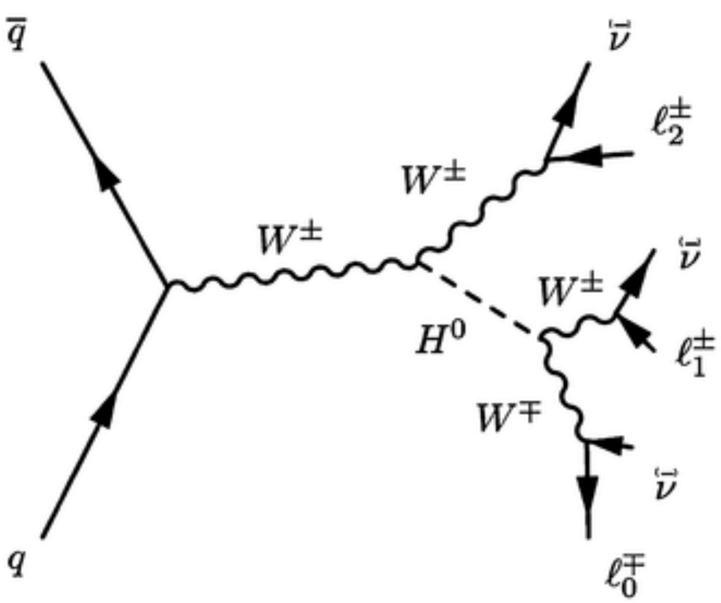

$\bar{q}$
q
$W^{\pm}$
$W^{\pm}$
$W^{\pm}$
$W^{\mp}$
H^0
$\overset{(-)}{\nu}$
$\ell_2^{\pm}$
$\overset{(-)}{\nu}$
$\ell_1^{\pm}$
$\overset{(-)}{\nu}$
$\ell_0^{\mp}$

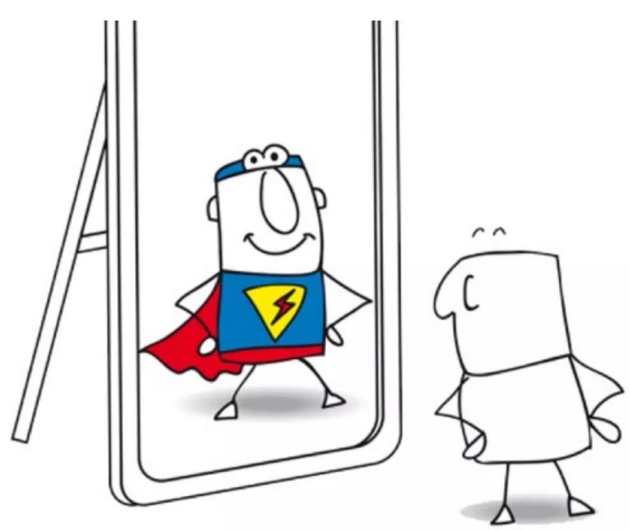

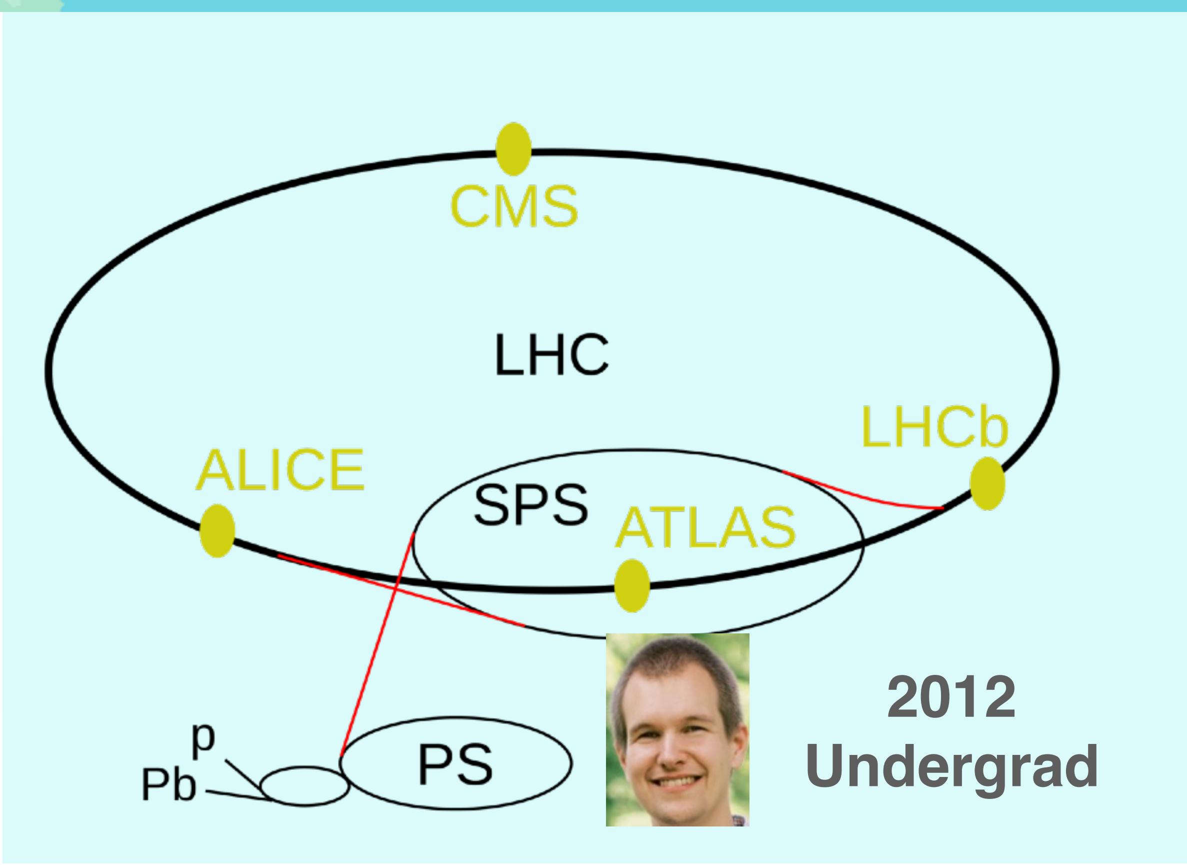

CMS
LHC
LHCb
ALICE
SPS
ATLAS
p
Pb
PS
2012
Undergrad

About me


MICHIGAN
Toronto
NEW HAMPSHIRE
NEW YORK
MASSACHUSETTS
Chicago
CT
RI
New York
PENNSYLVANIA
ILLINOIS
INDIANA
OHIO
Philadelphia
NJ
Indianapolis
Cincinnati
MARYLAND
DE
WEST VIRGINIA
CERN

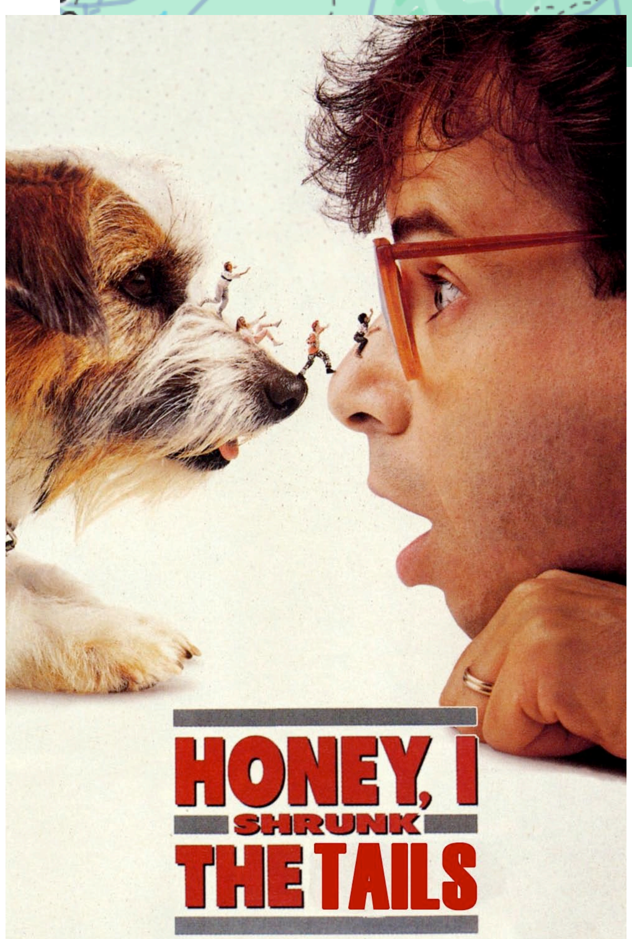
HONEY, I
SHRUNK
THE TAILS

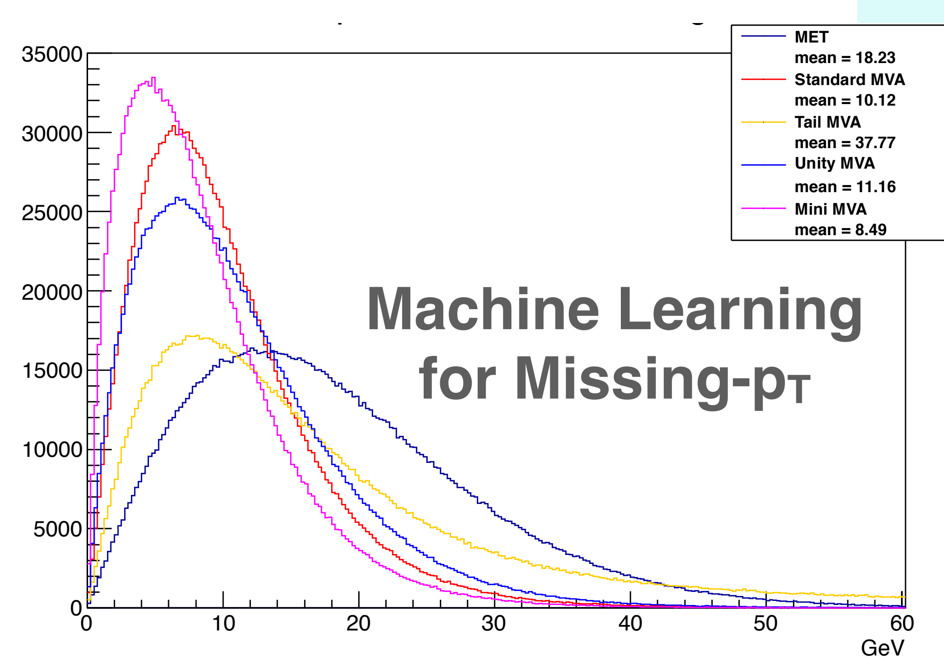
MET mean = 18.23
Standard MVA mean = 10.12
Tail MVA mean = 37.77
Unity MVA mean = 11.16
Mini MVA mean = 8.49
Machine Learning for Missing-p_T
GeV

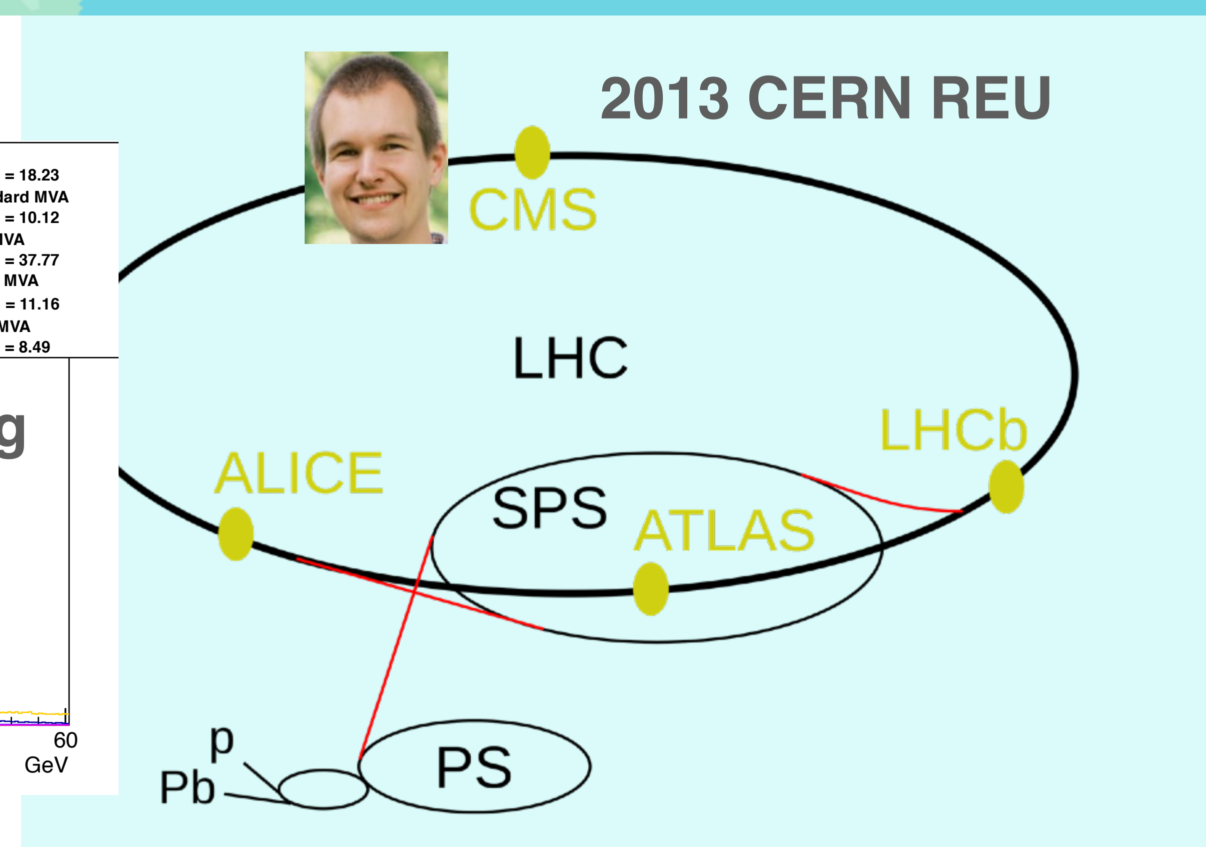
2013 CERN REU
CMS
LHC
ALICE
LHCb
SPS
ATLAS
p
Pb
PS

About me

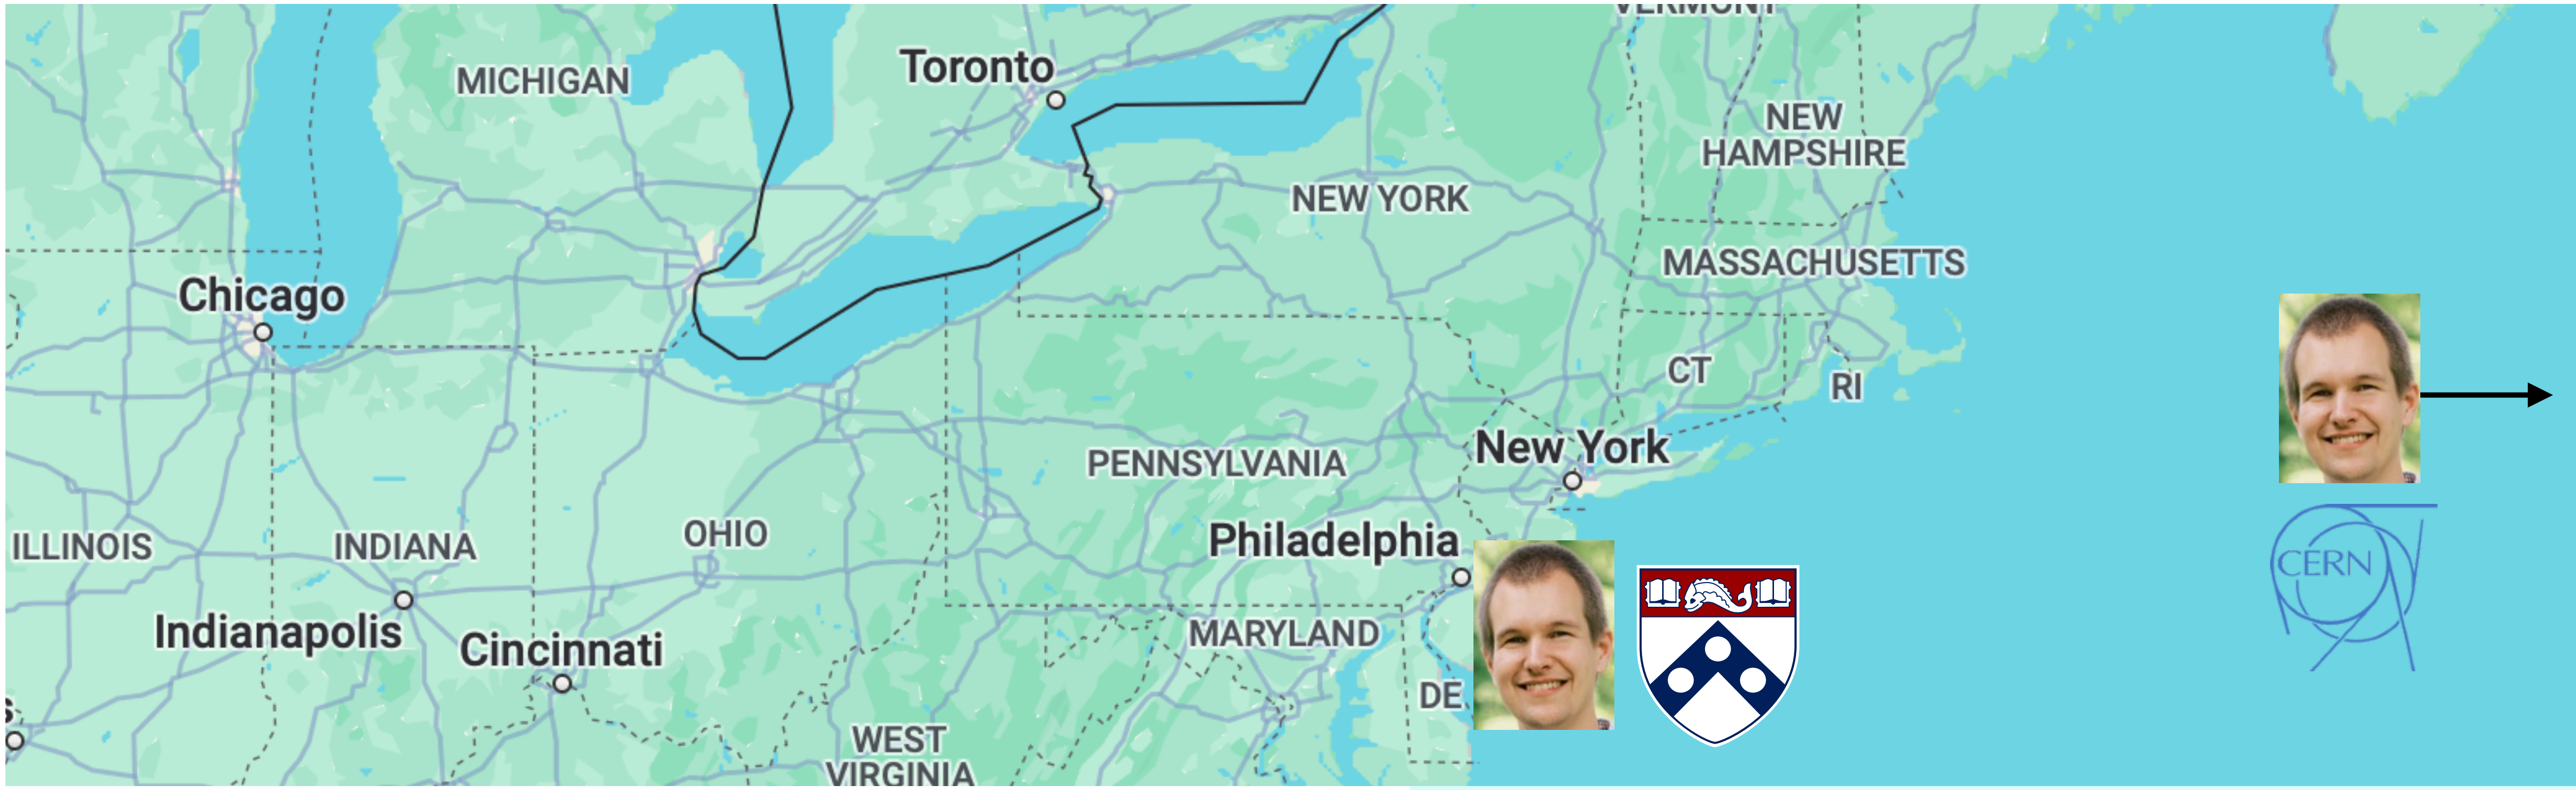
MICHIGAN
Toronto
NEW HAMPSHIRE
NEW YORK
MASSACHUSETTS
Chicago
CT
RI
New York
PENNSYLVANIA
ILLINOIS
INDIANA
OHIO
Philadelphia
Indianapolis
Cincinnati
MARYLAND
DE
WEST VIRGINIA
CERN

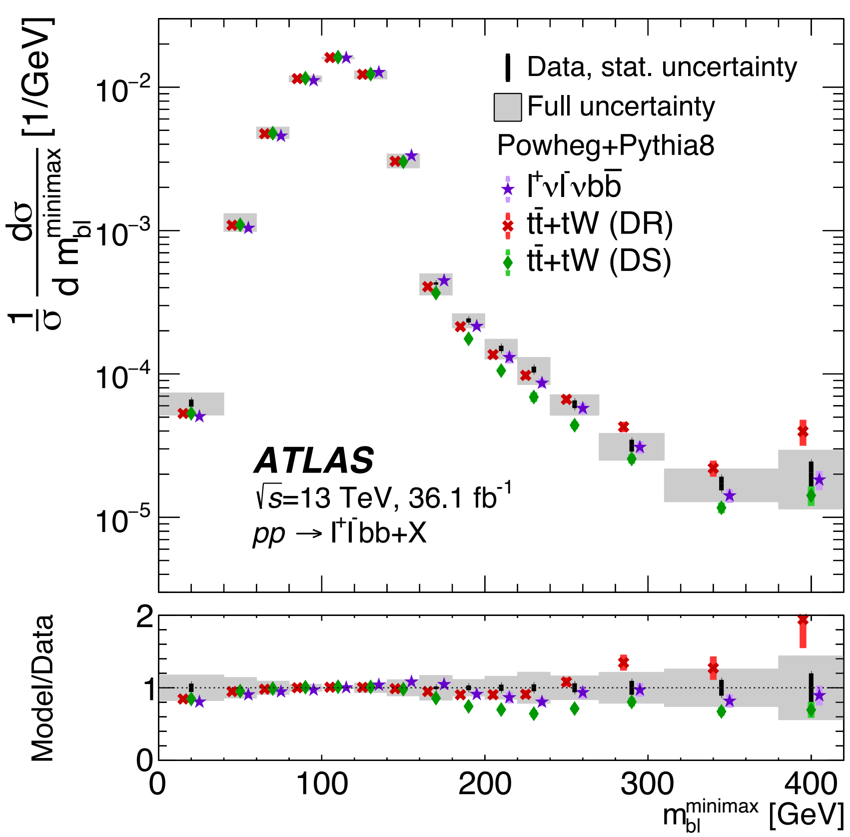
Data, stat. uncertainty
Full uncertainty
Powheg+Pythia8
$l^+\nu l^-\nu b\bar{b}$
$t\bar{t}$+tW (DR)
$t\bar{t}$+tW (DS)
ATLAS
$\sqrt{s}$=13 TeV, 36.1 fb^{-1}
$pp \to l^+l^-bb+X$
$\frac{1}{\sigma}\frac{d\sigma}{d\,m_{bl}^{minimax}}$ [1/GeV]
Model/Data
$m_{bl}^{minimax}$ [GeV]

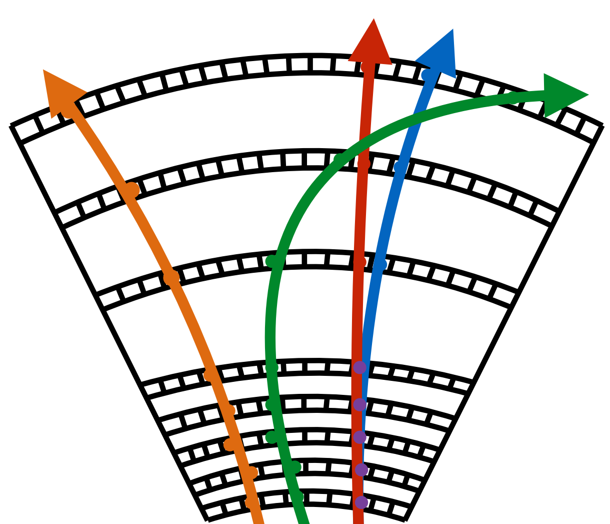

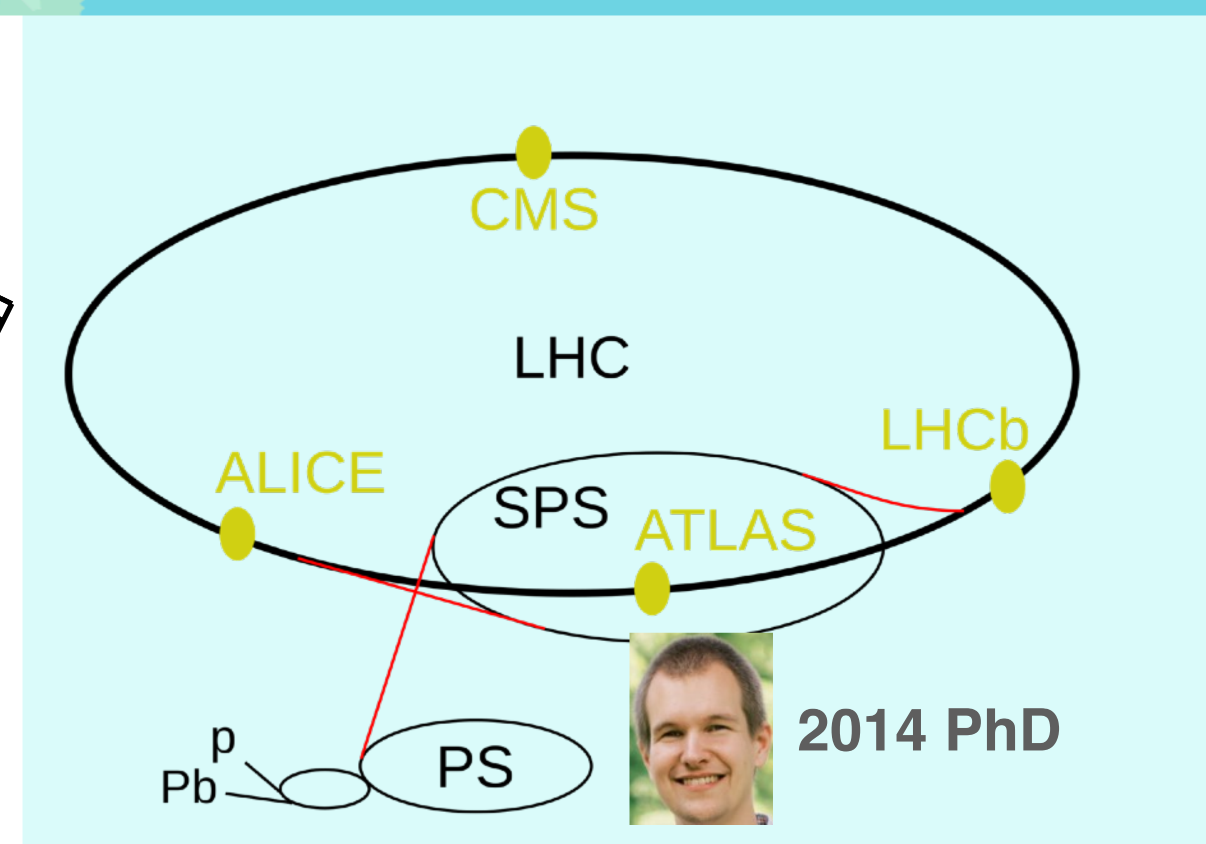
CMS
LHC
LHCb
ALICE
SPS
ATLAS
p
Pb
PS
2014 PhD

About me

MICHIGAN
Toronto
NEW HAMPSHIRE
NEW YORK
MASSACHUSETTS
Chicago
CT
RI
New York
PENNSYLVANIA
INDIANA
OHIO
Philadelphia
NJ
Indianapolis
Cincinnati
MARYLAND
DE
WEST VIRGINIA

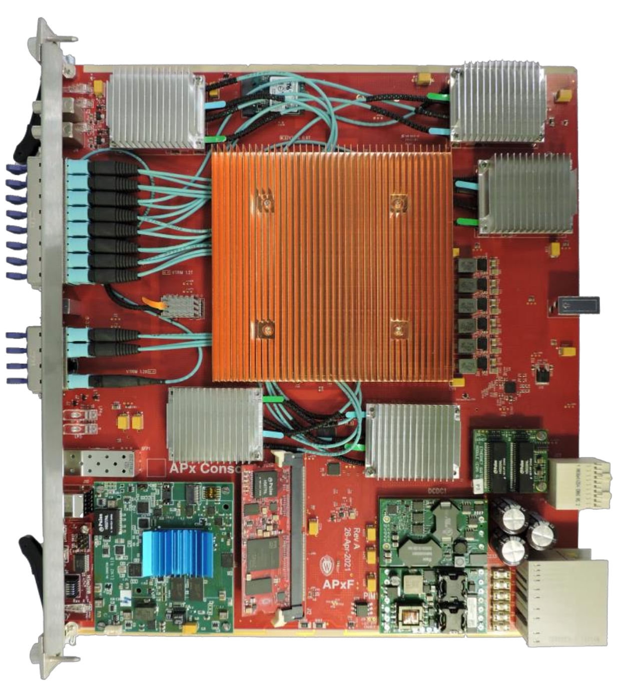

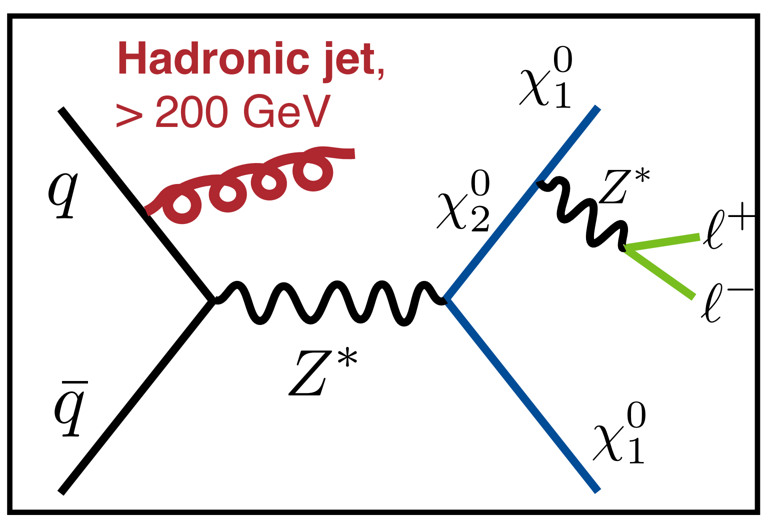

Hadronic jet,
> 200 GeV
q
$\bar{q}$
Z^*
χ_2^0
χ_1^0
Z^*
ℓ^+
ℓ^-
χ_1^0

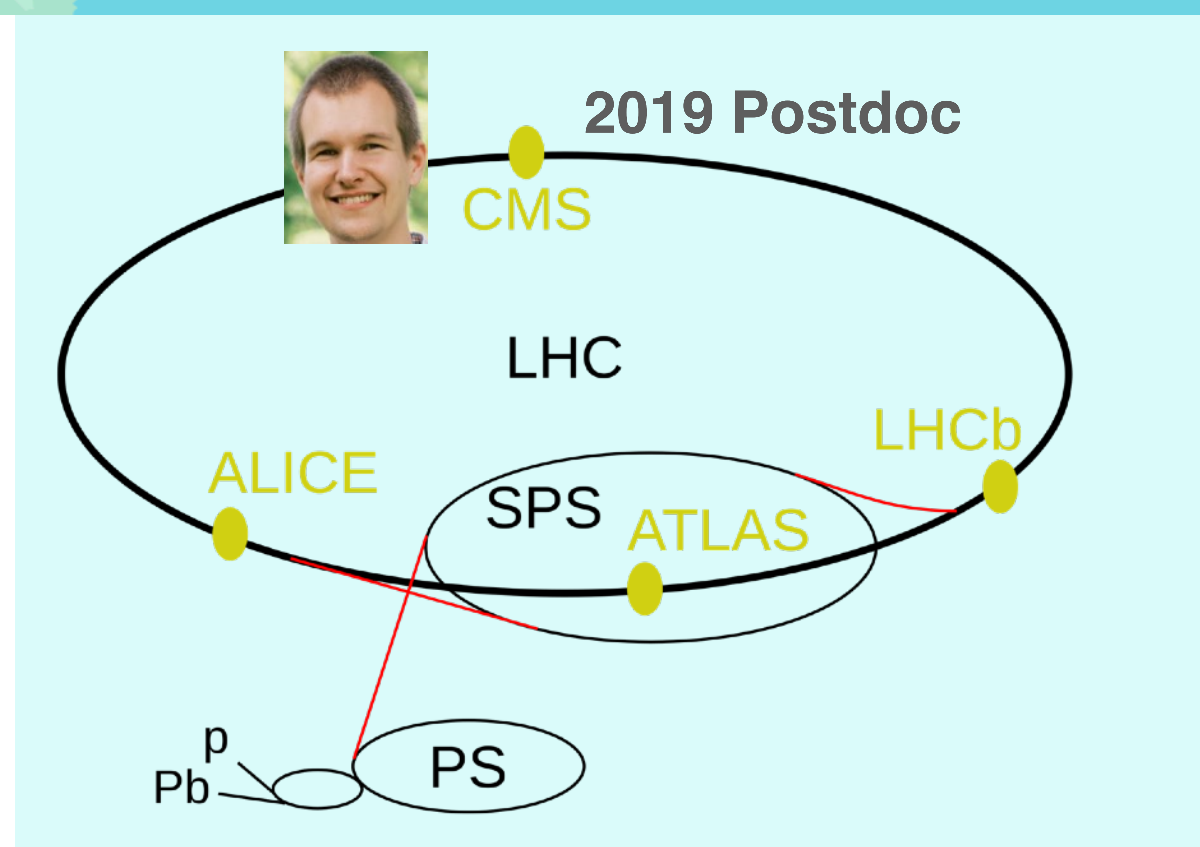

2019 Postdoc
CMS
LHC
LHCb
ALICE
SPS
ATLAS
p
Pb
PS

About me

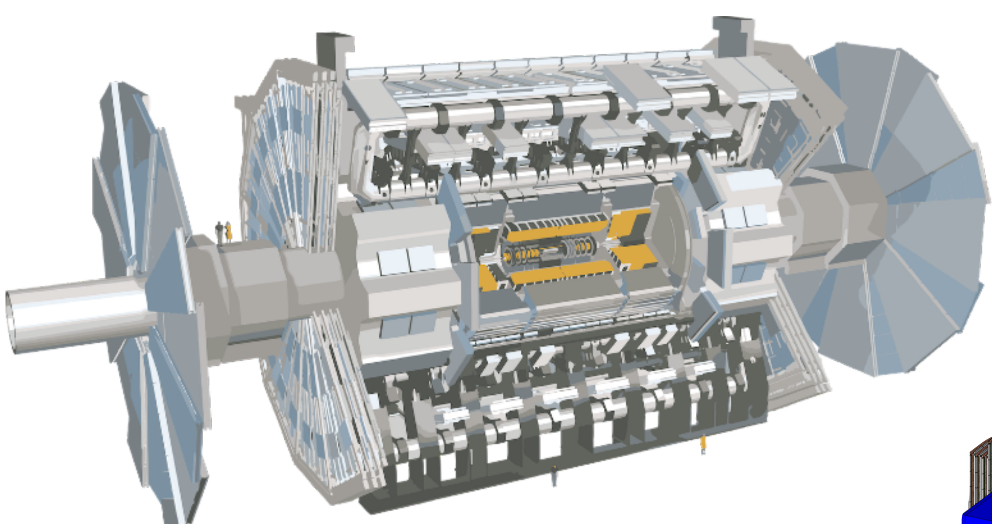

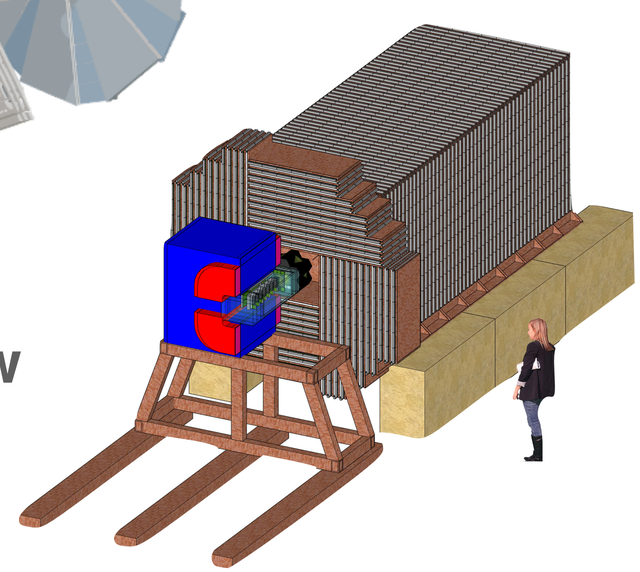

Continuing the hunt for New Physics with accelerators!

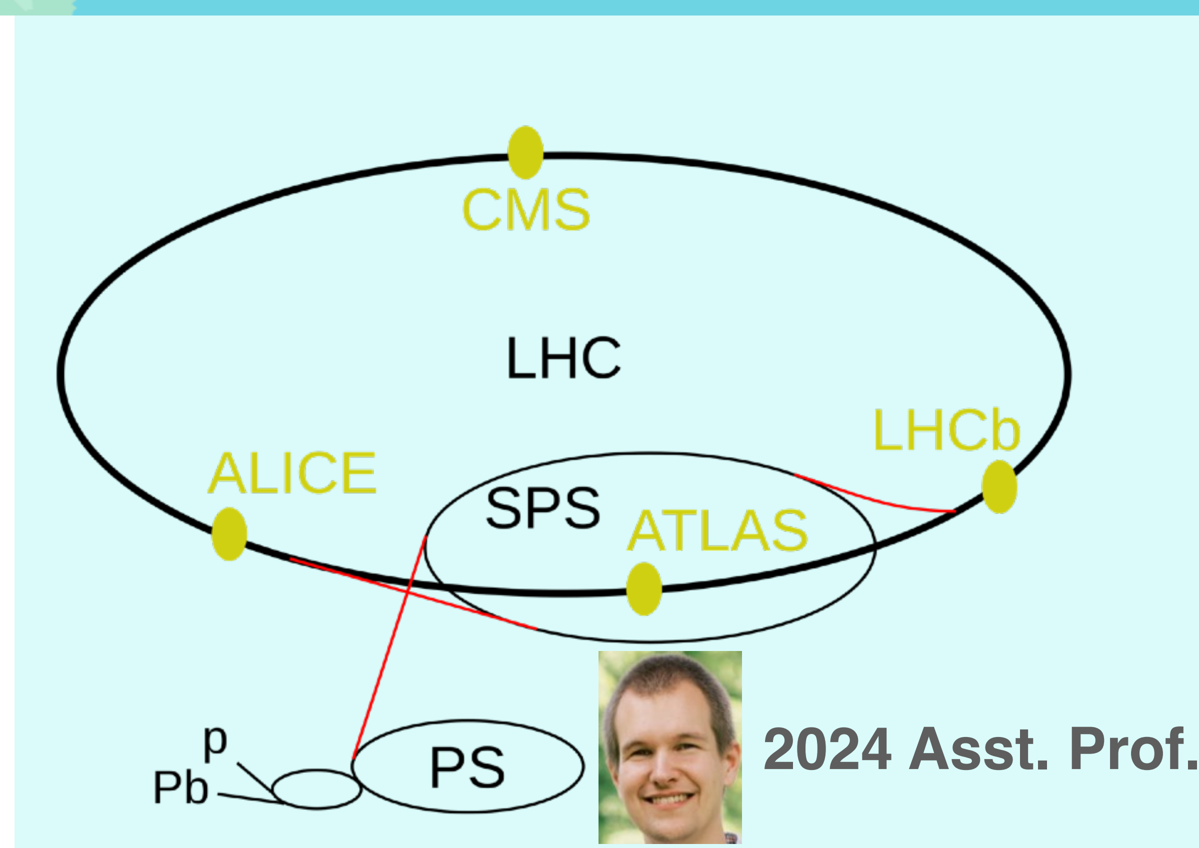

Plan

- **This afternoon: Concepts**
  - Introduction and physics goals of the LHC Experiments
  - The role of the trigger and DAQ systems
  - DAQ concepts: simple toy model → complex systems
  - Trigger concepts: from hardware to architecting a system
- **Tomorrow morning: Details and Implementations**
  - The challenge of high luminosity, and the upgrades
  - ATLAS, CMS hardware triggers, and their evolution
  - Software triggers, and new paradigms for analysis
  - Trigger menus, and you
  - Conclusions, and looking farther ahead

Finding more information

- These lectures borrow inspiration and material form many excellent schools in the past on this and similar topics.
    - Thanks to Sergo Jindariani, Darin Acosta, Lauren Tompkins
- If interested, can consider other schools on dedicated topics! e.g.
    - International School of Trigger and DAQ (ISOTDAQ)
    - Excellence in Detectors and Instrumentation
- And of course the primary material itself!

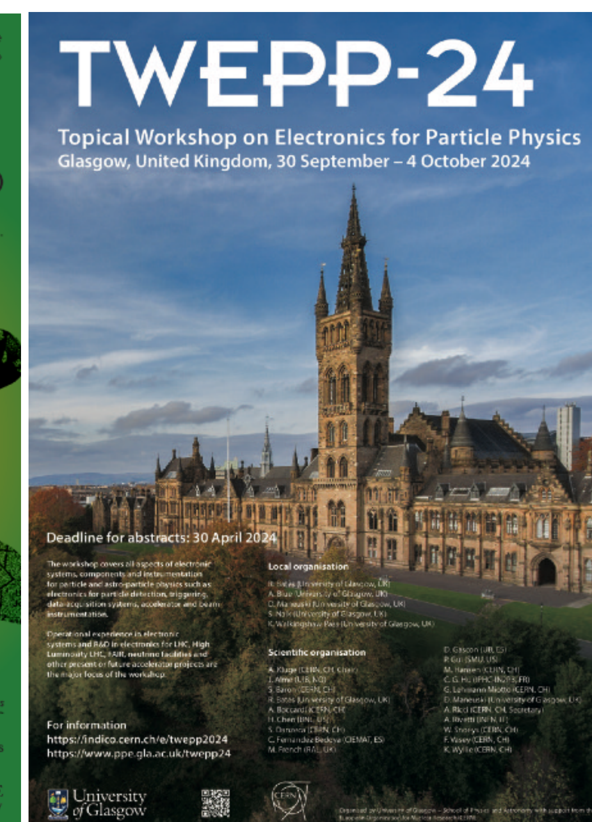

Physics goals (I)

The high energy and large luminosity available at the LHC enables a diverse physics program (ATLAS, CMS, LHCb, ALICE, …)

Test properties of copiously-produced particles to high precision (B, W, Z, …)

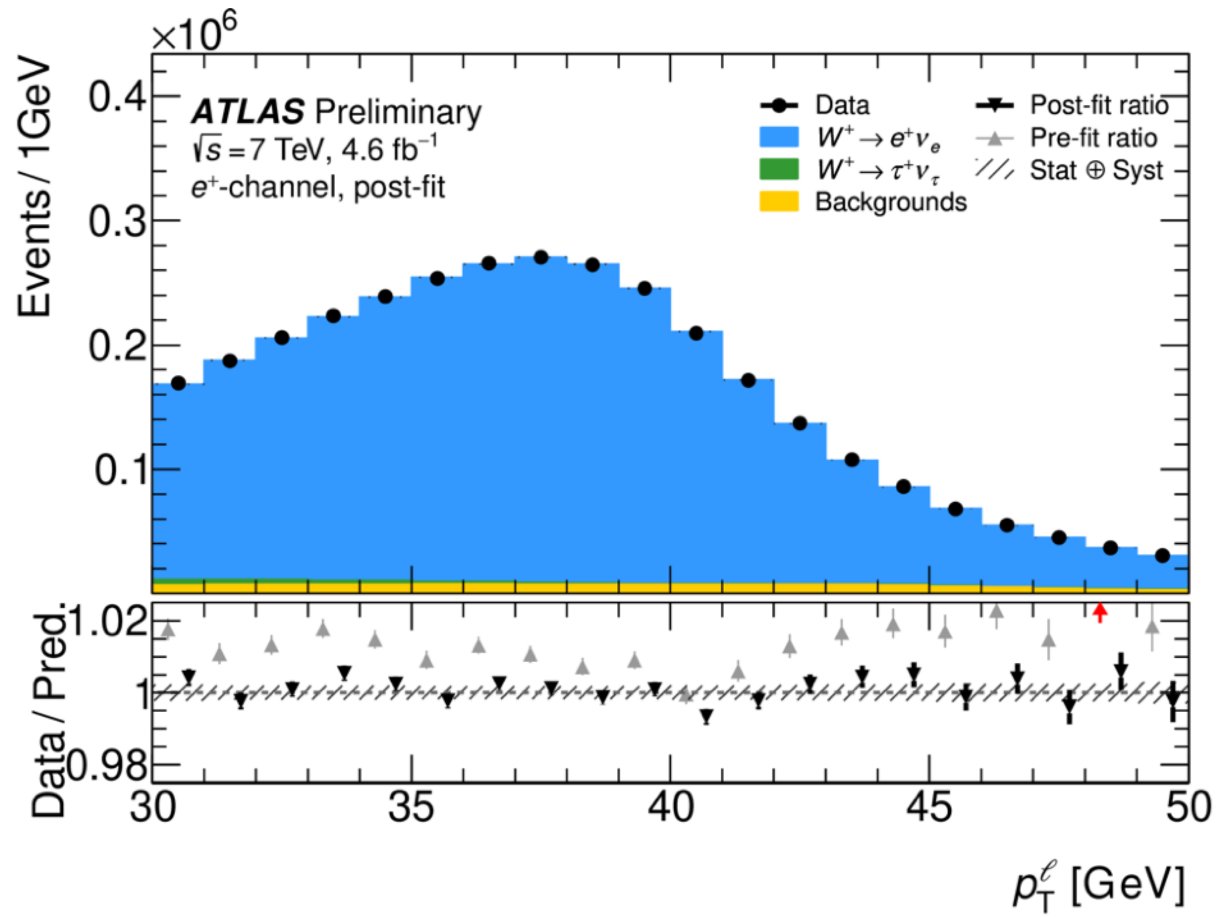

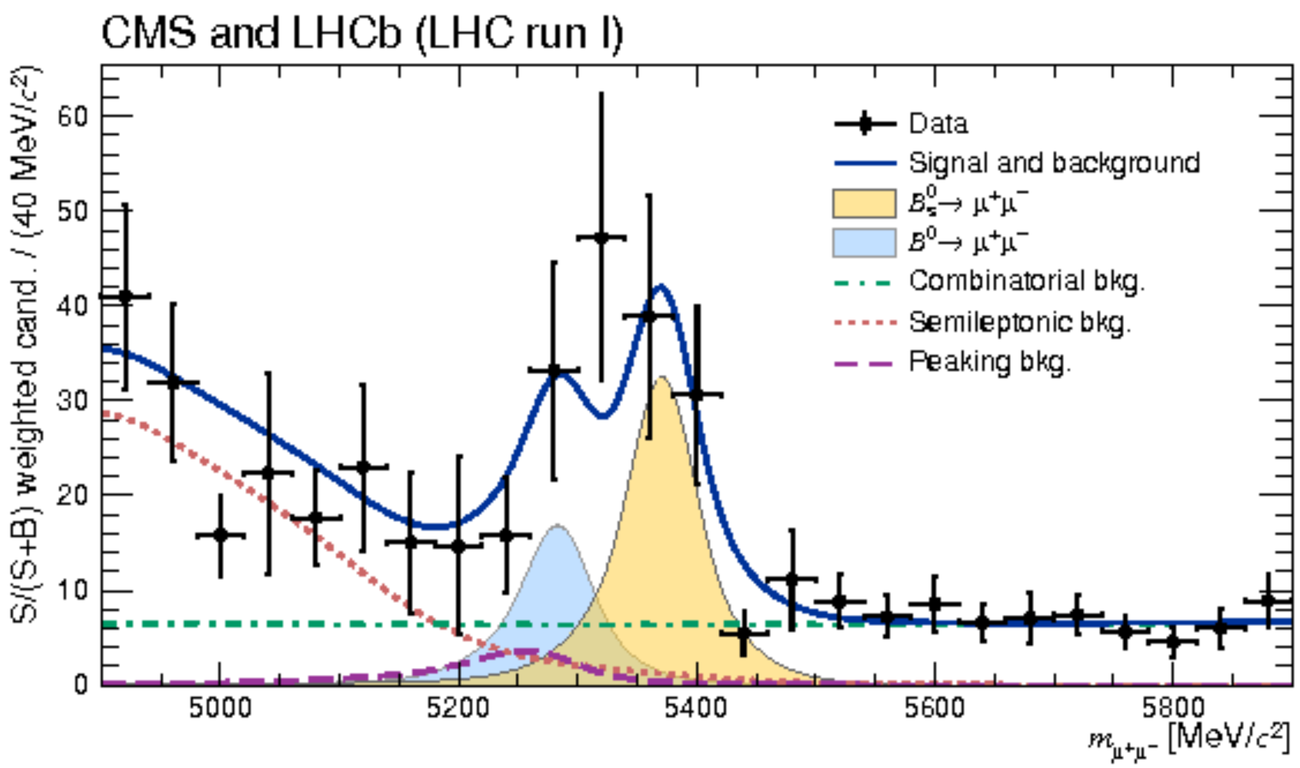

Hunt for **rare SM** phenomena

Physics goals (II)

The high energy and large luminosity available at the LHC enables a diverse physics program (ATLAS, CMS, LHCb, ALICE, …)

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**?
  - New "top-like" particles key to hierarchy problem?
- …and **many more** that address other deep questions
  - CP violation, matter/antimatter asymmetry, small neutrino masses,…

The ATLAS detector

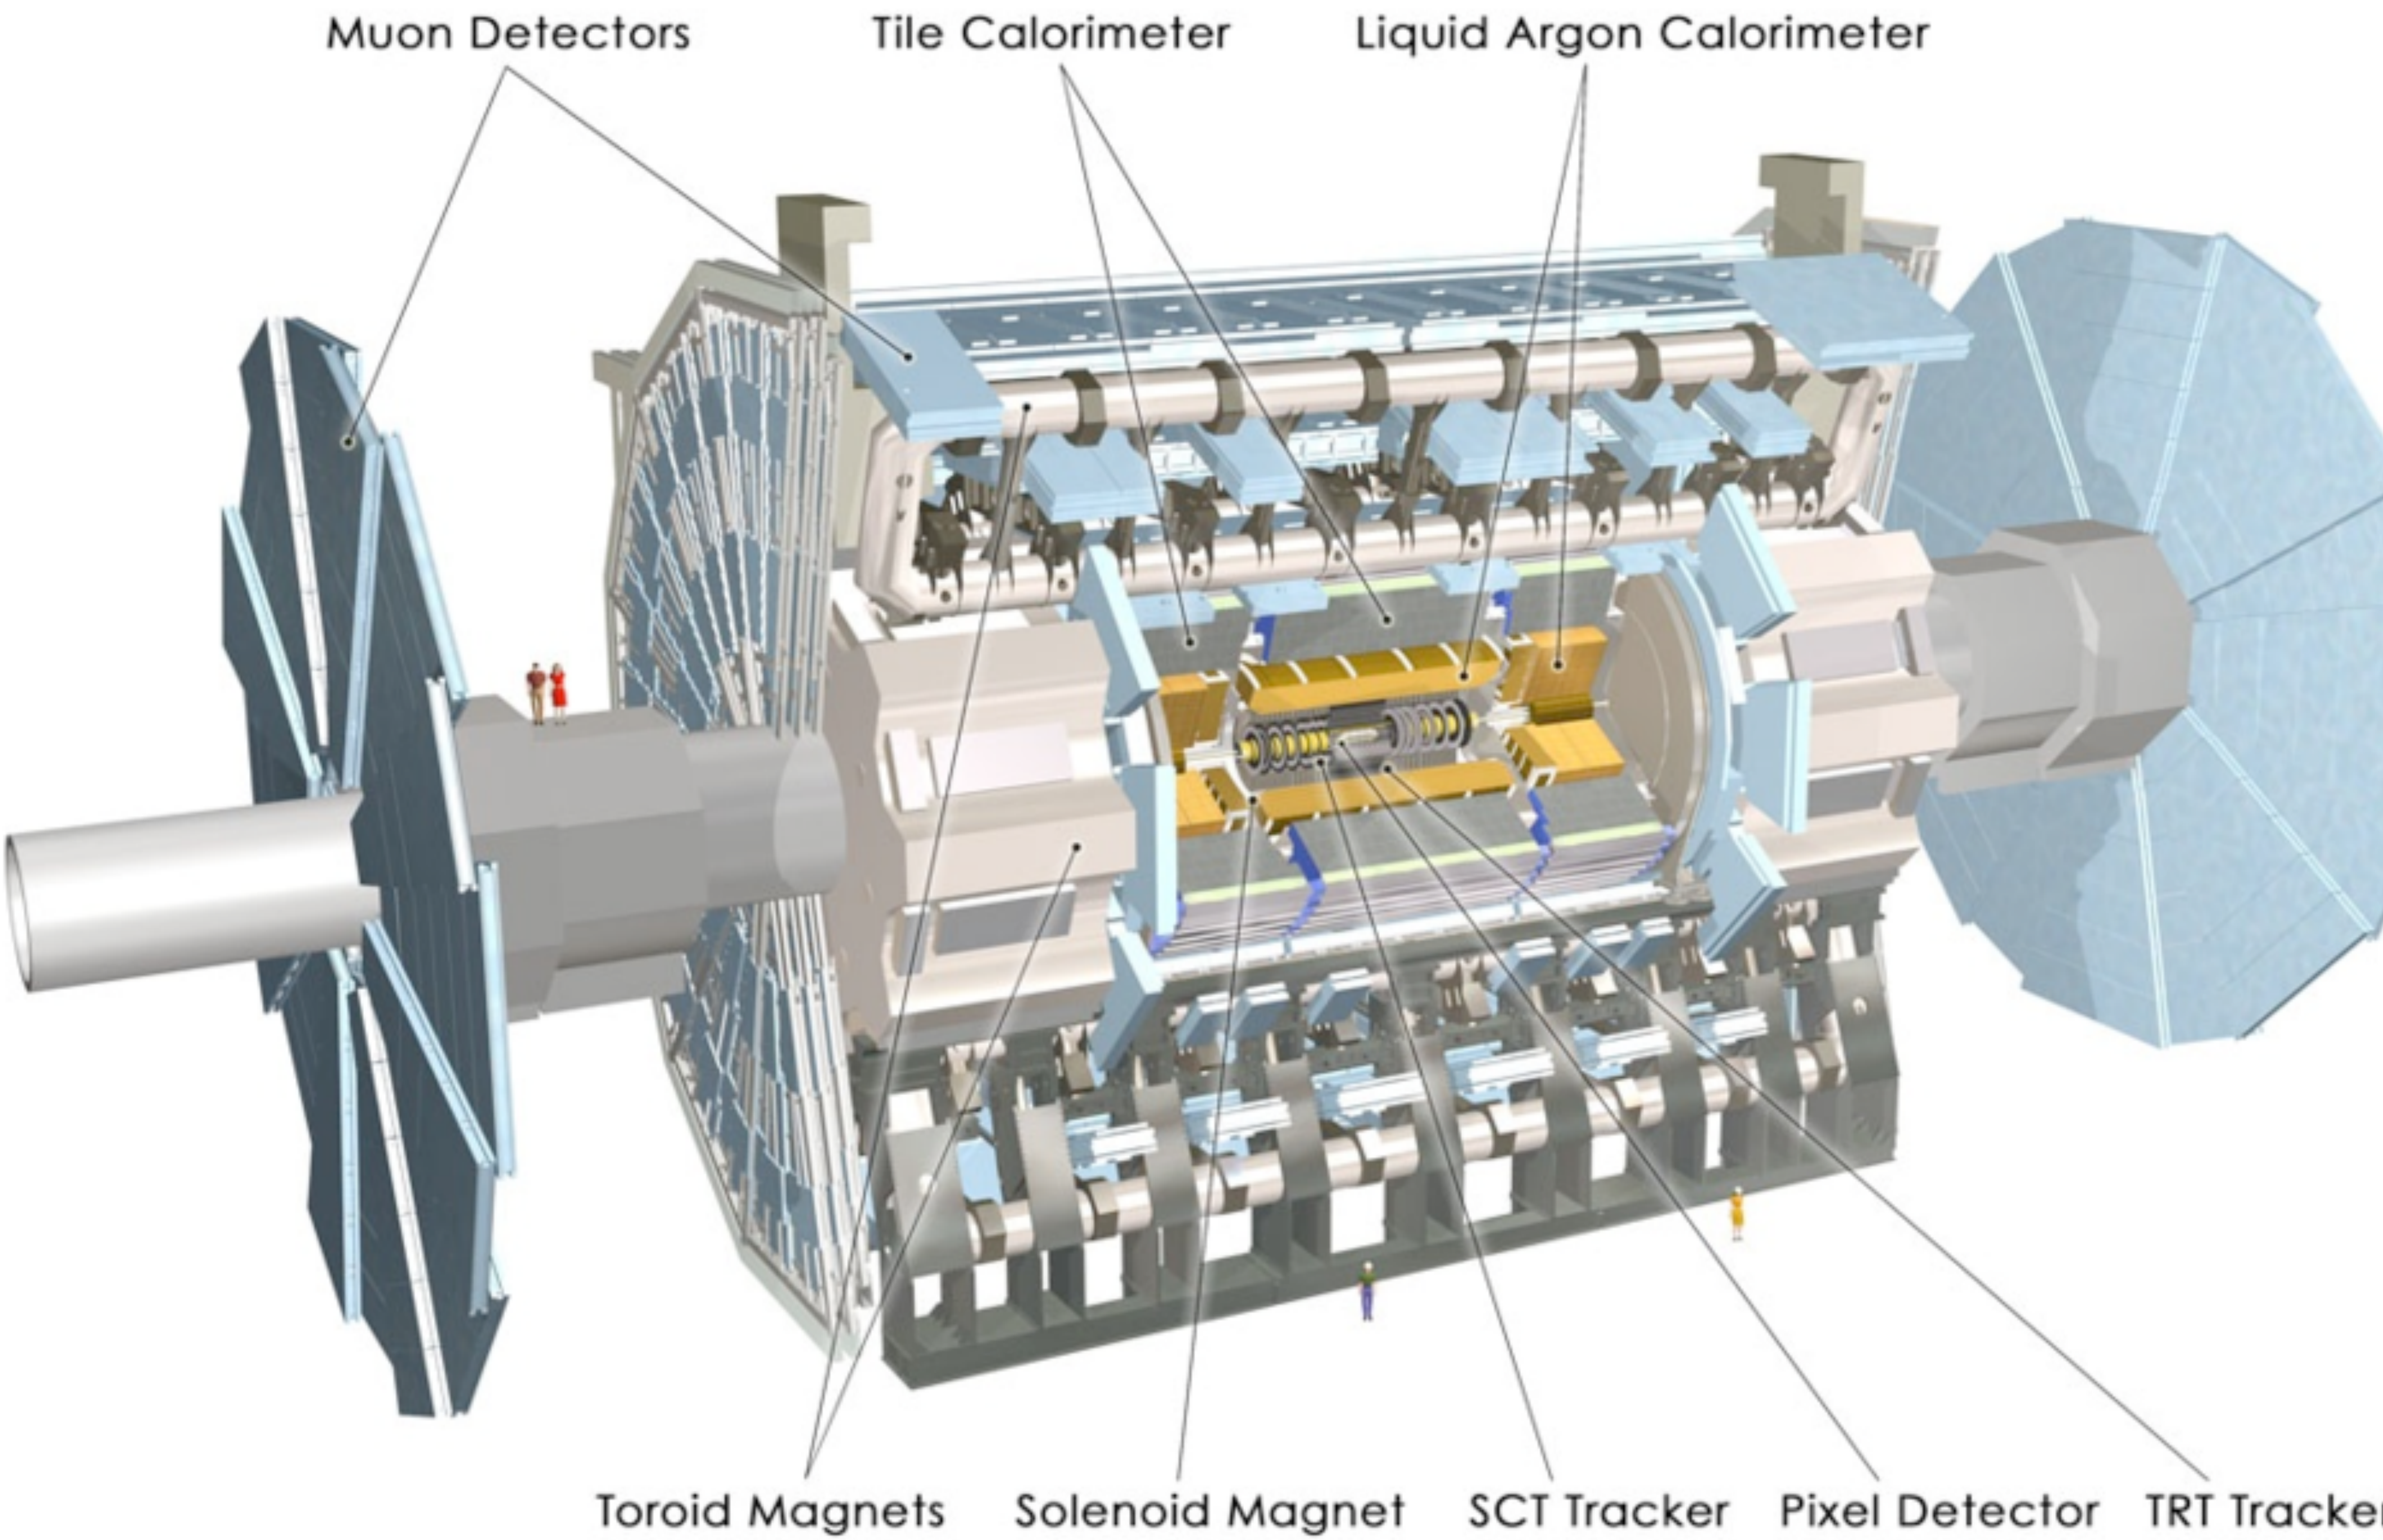

CMS Detector

Total weight	: 14,000 tonnes
Overall diameter	: 15.0 m
Overall length	: 28.7 m
Magnetic field	: 3.8 T

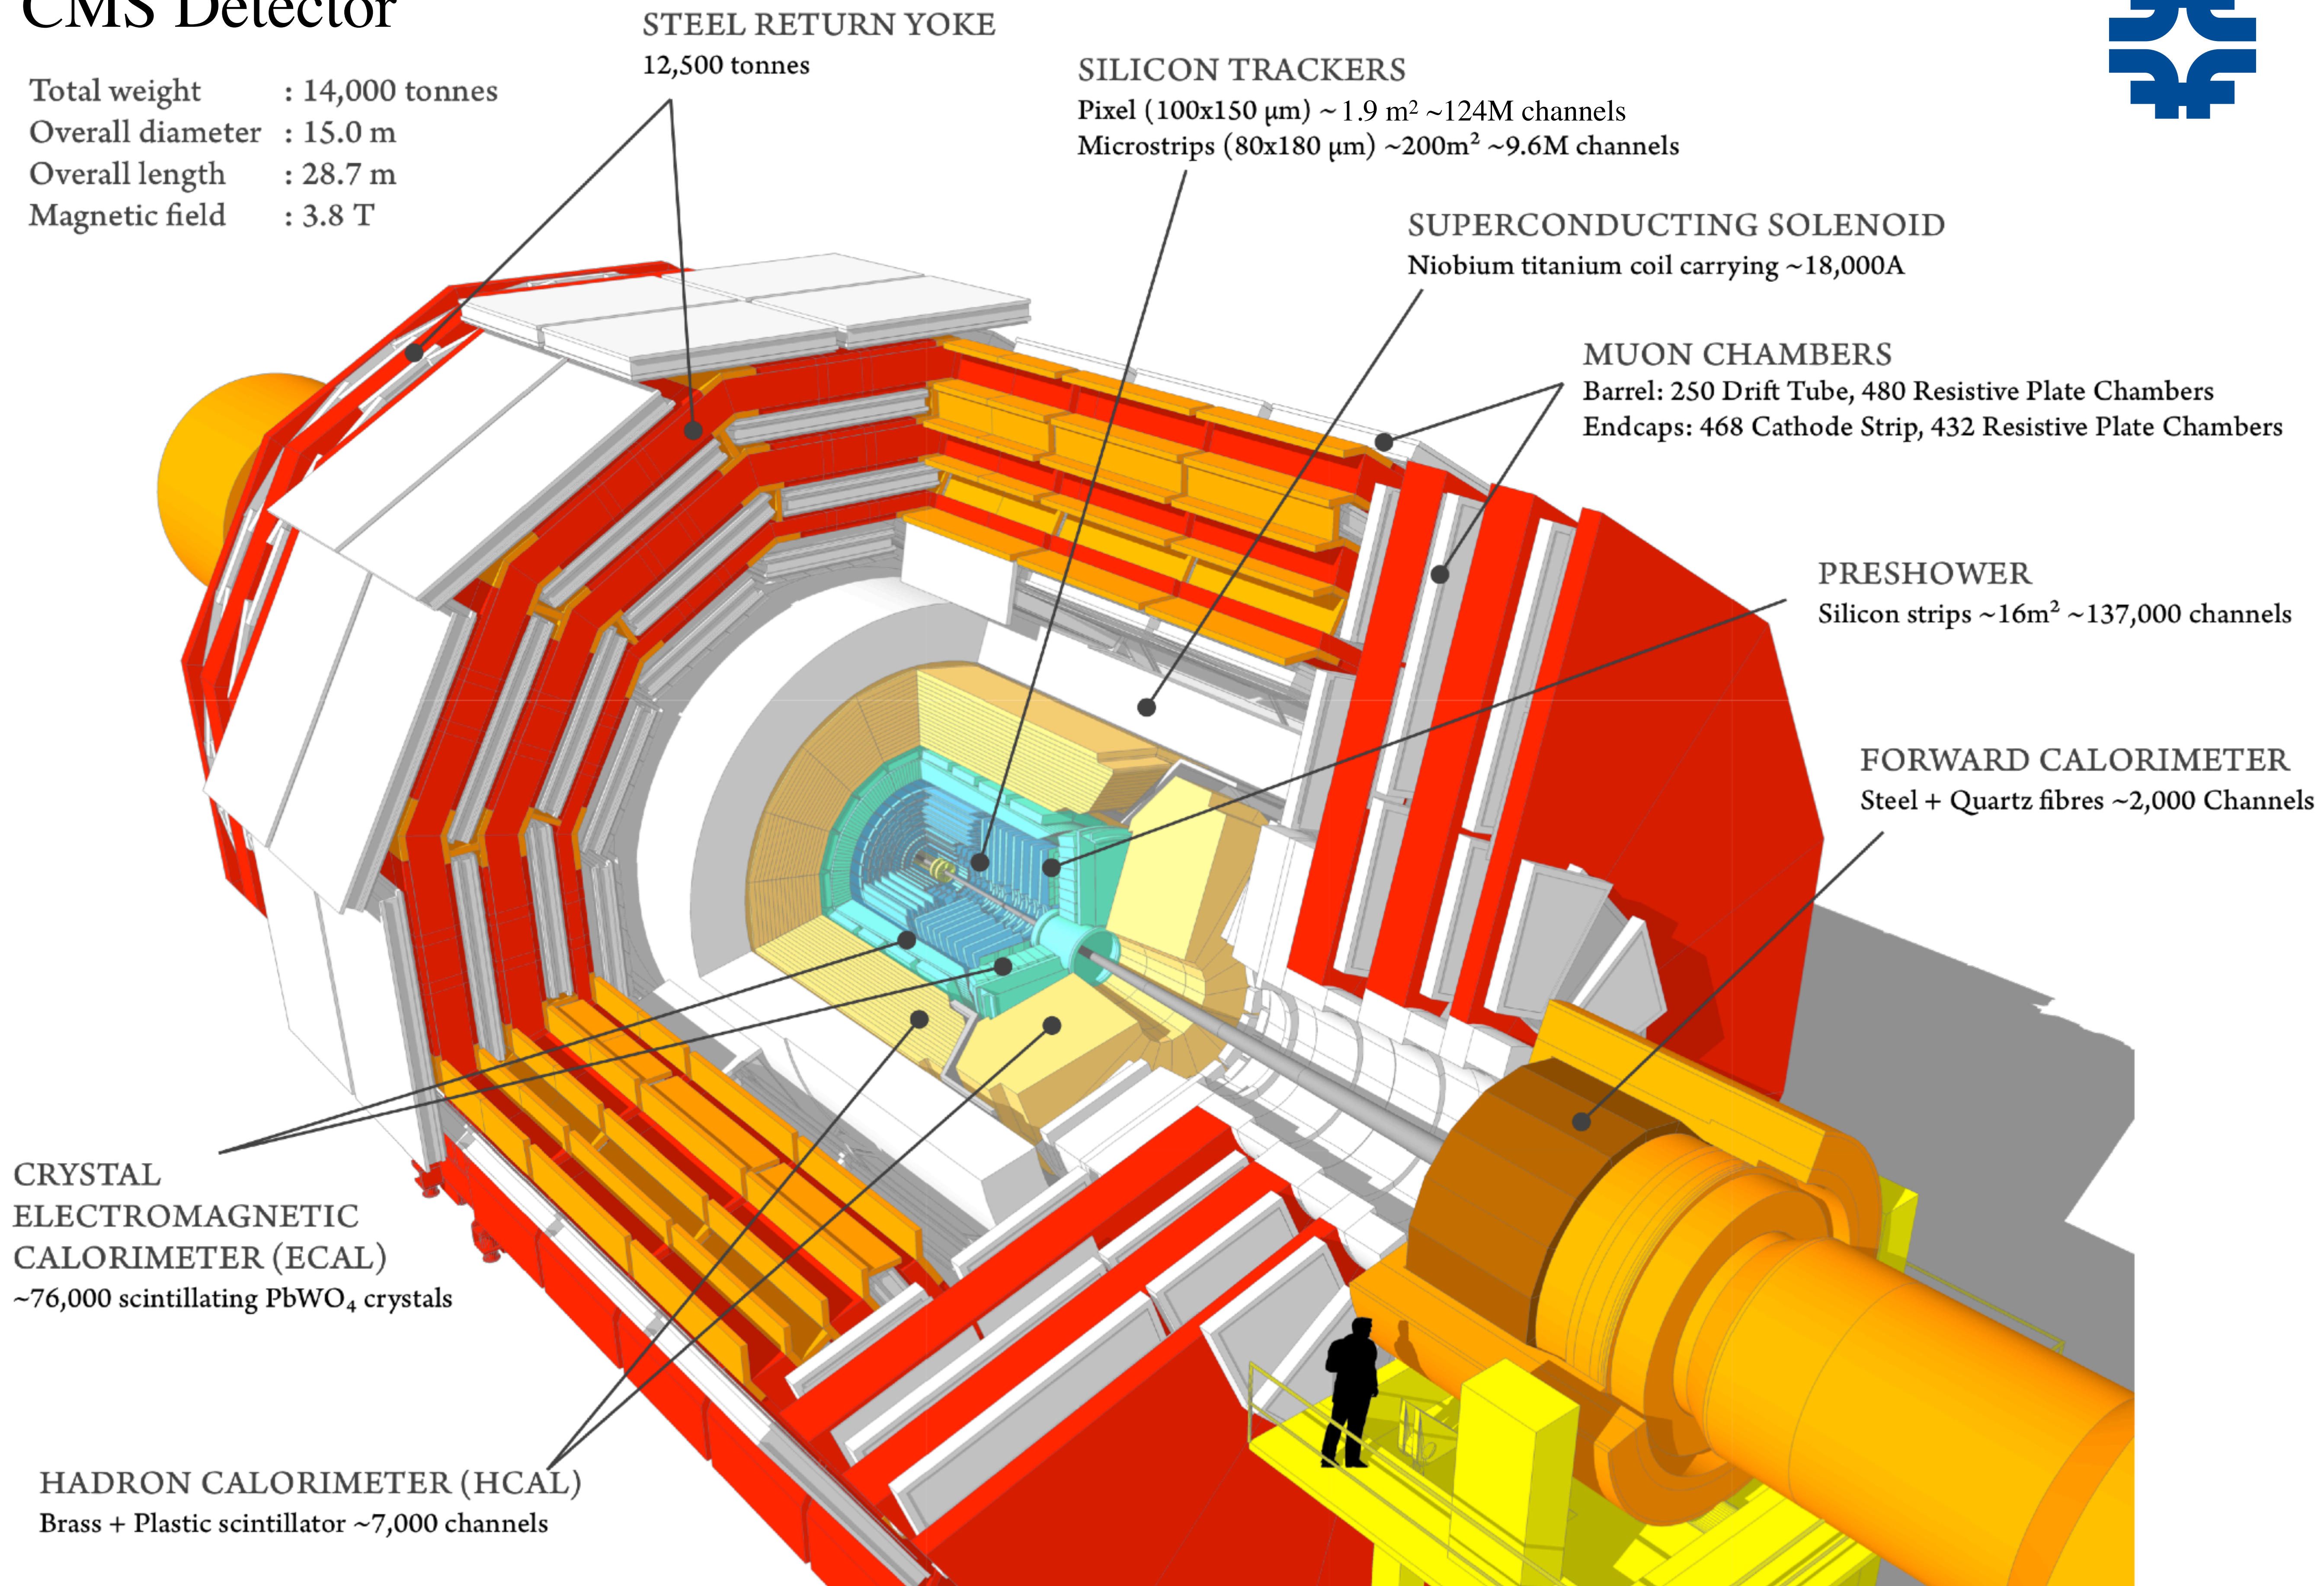

LHCb Detector

Weight: 5,600 tonnes
Height: 10 m
Length: 20 m

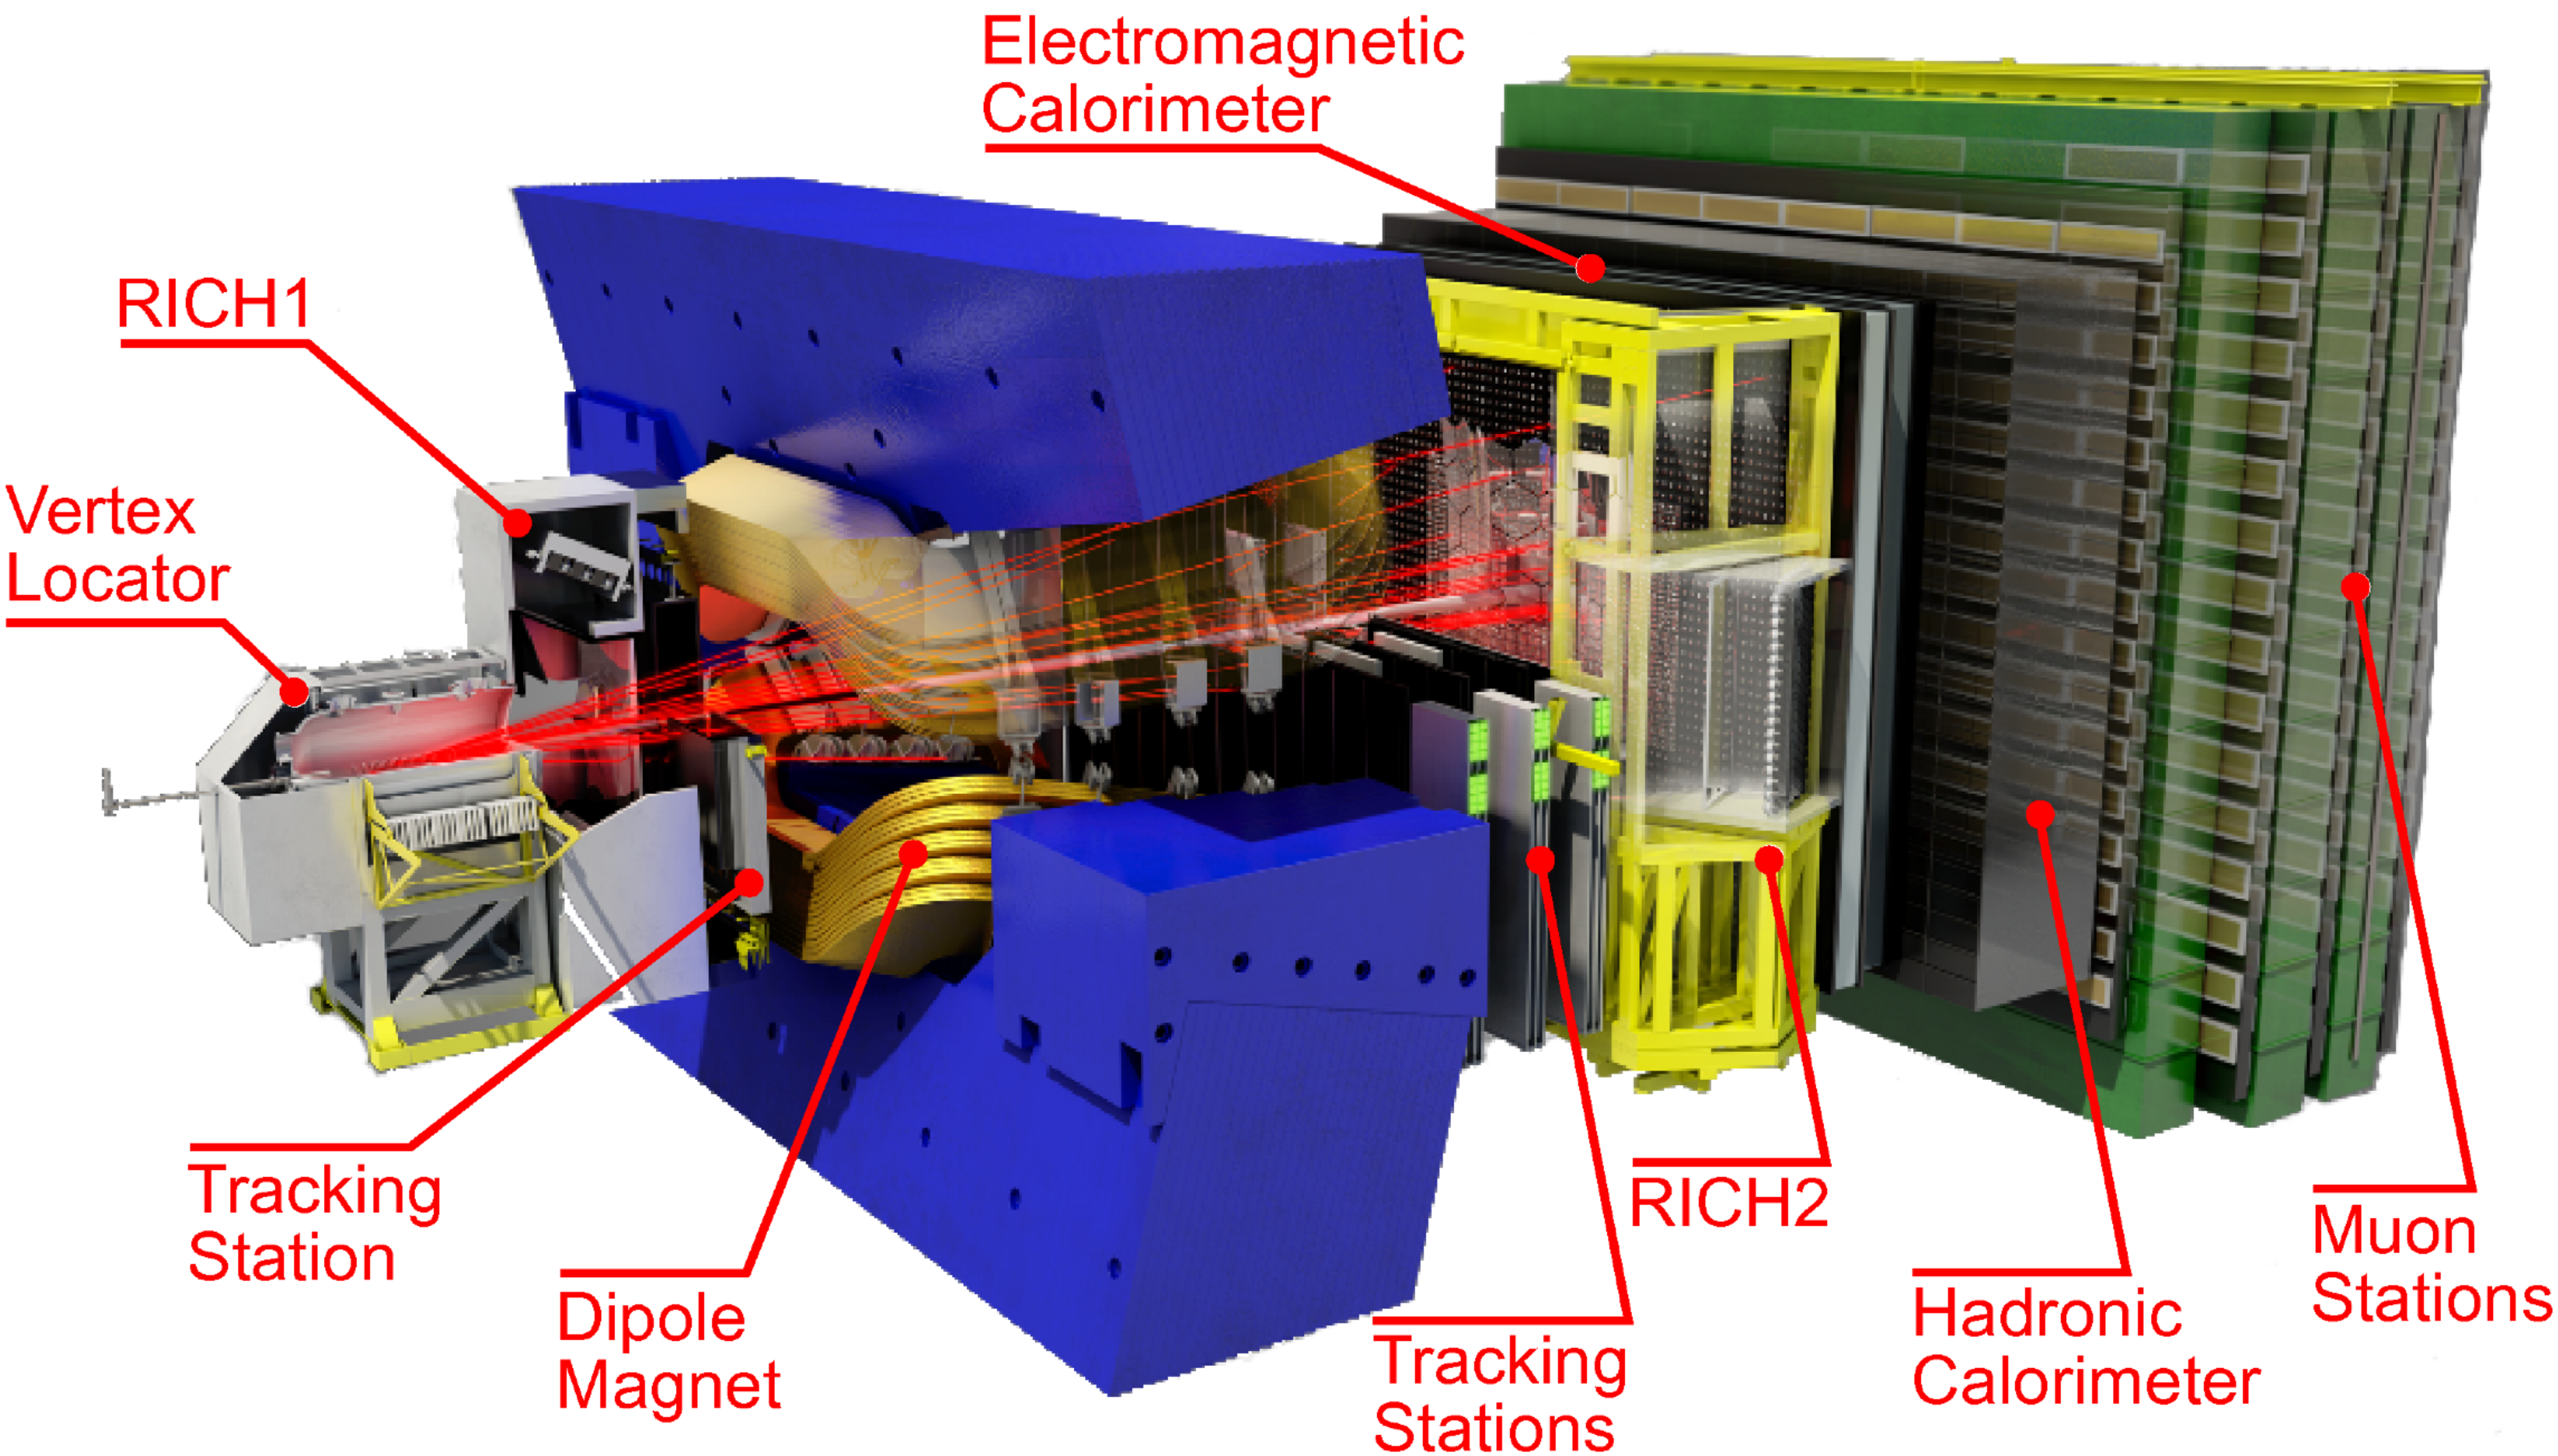

Interesting collisions: Higgs(→ZZ)

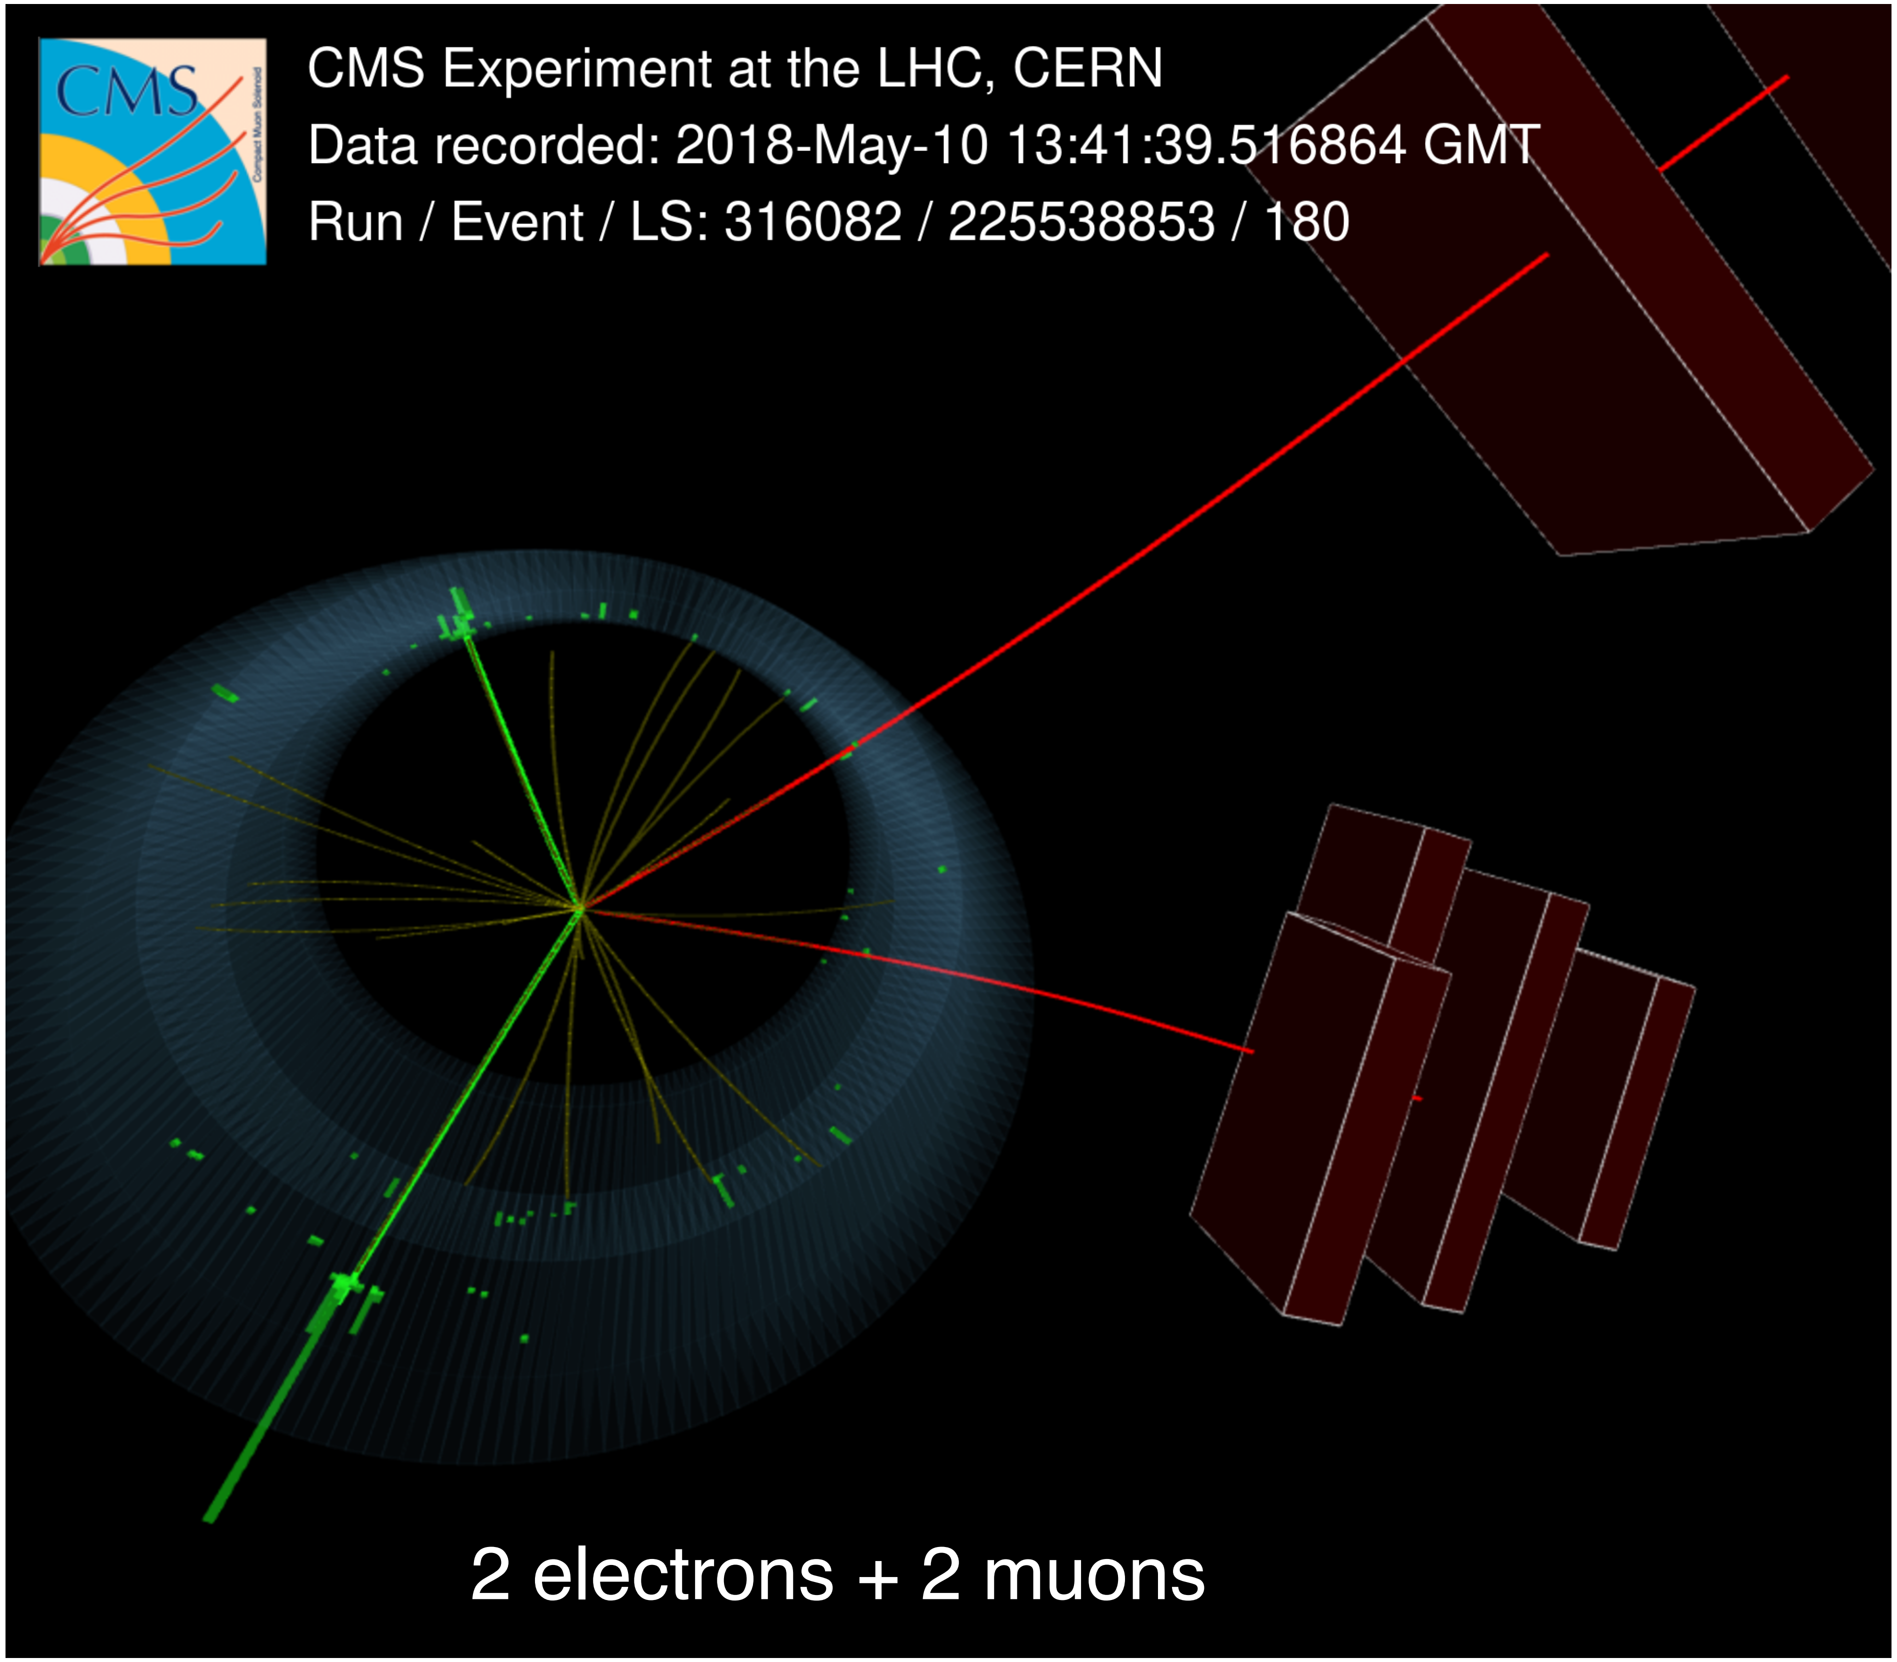

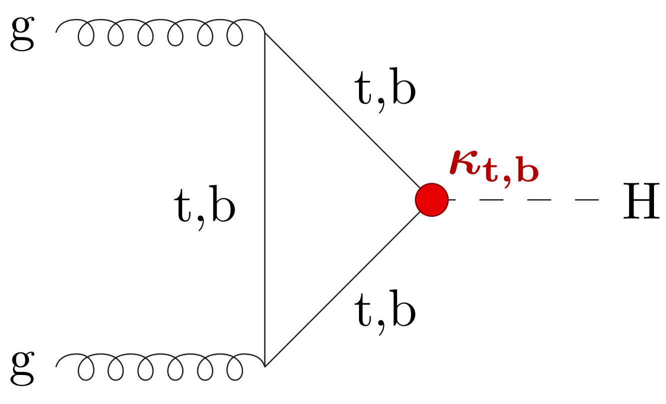

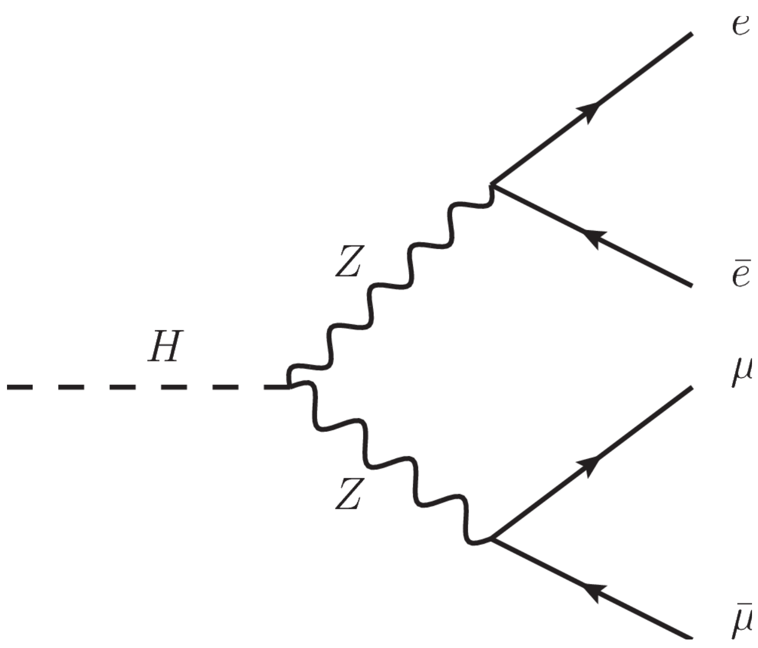

Interesting collisions: Higgs(→bb)

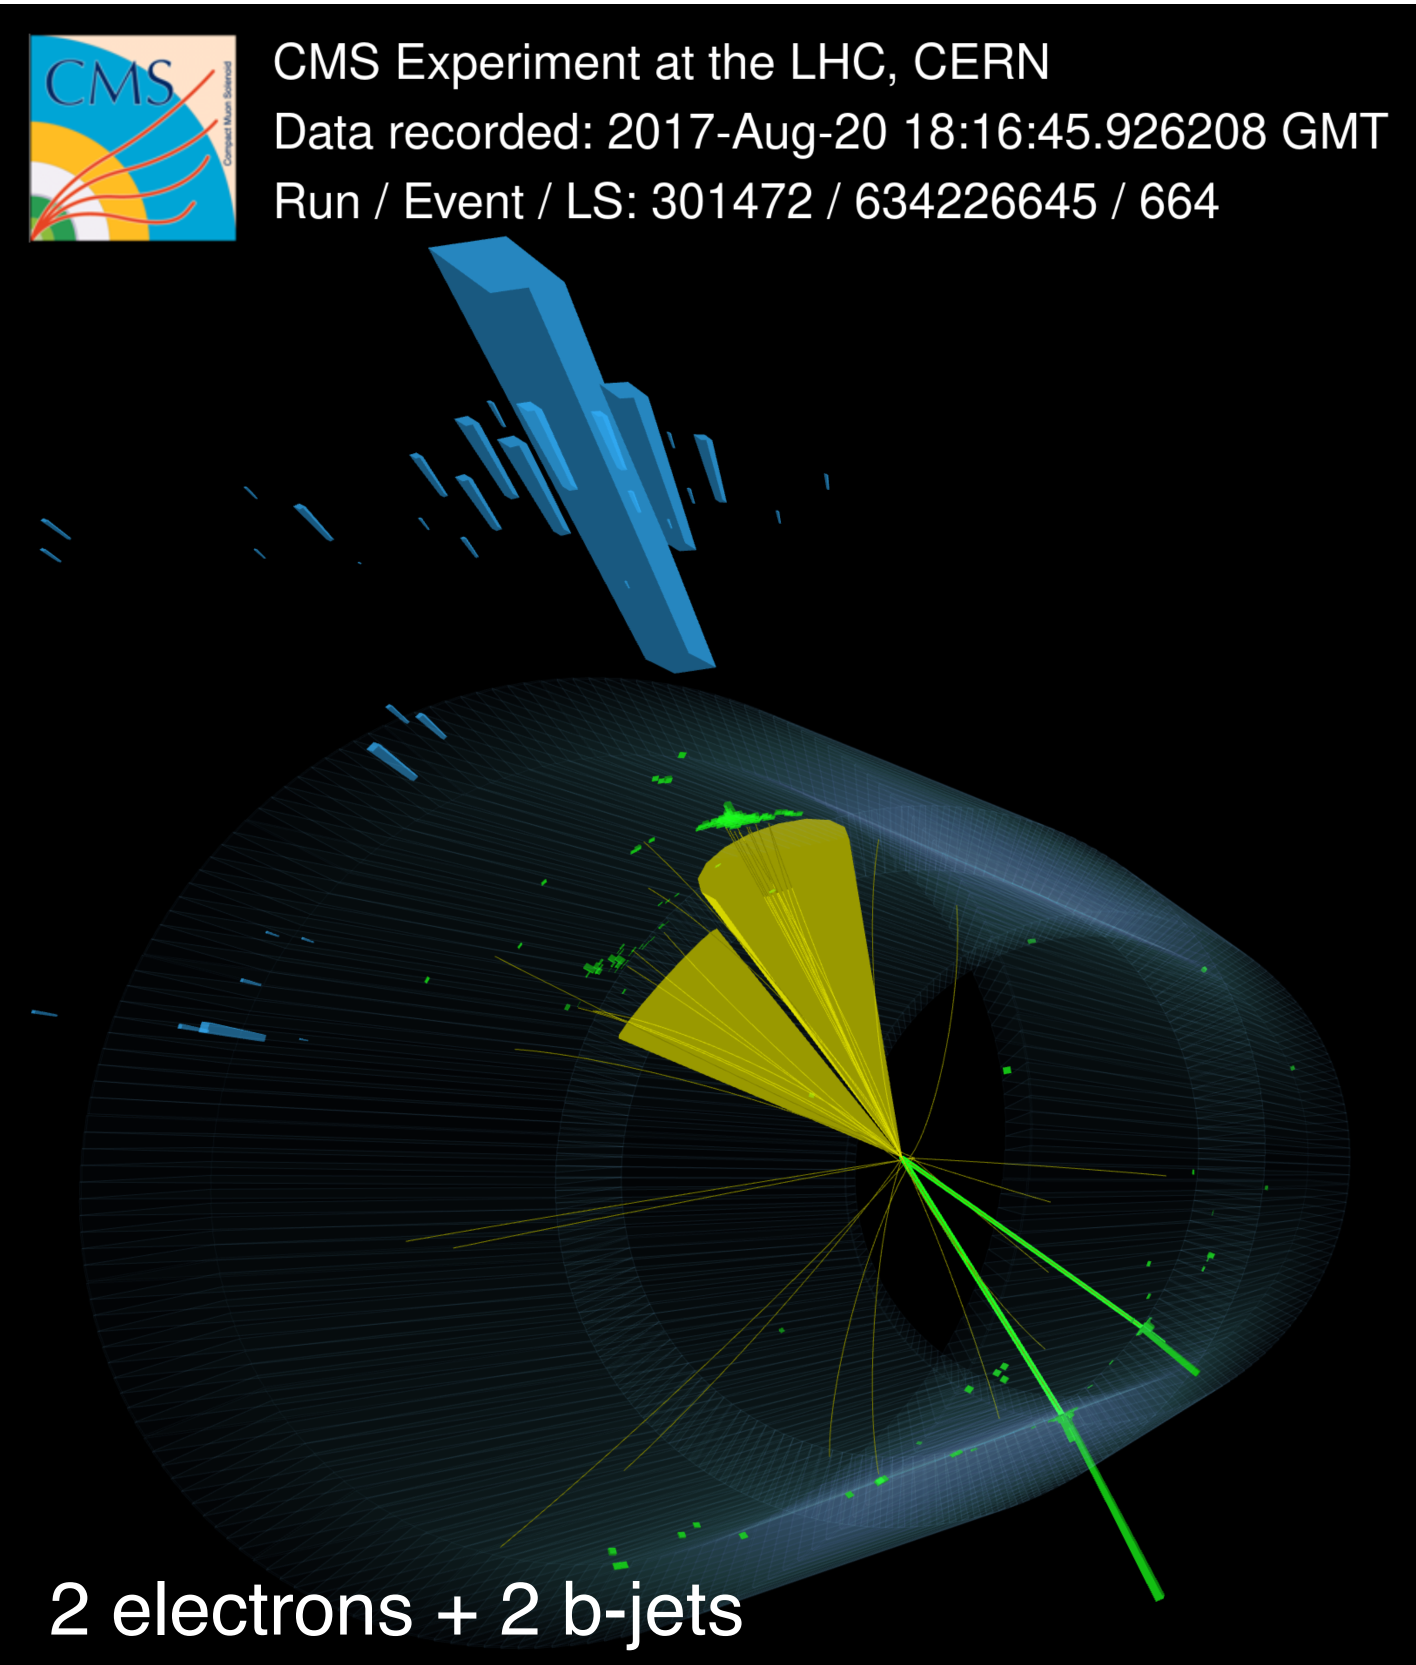

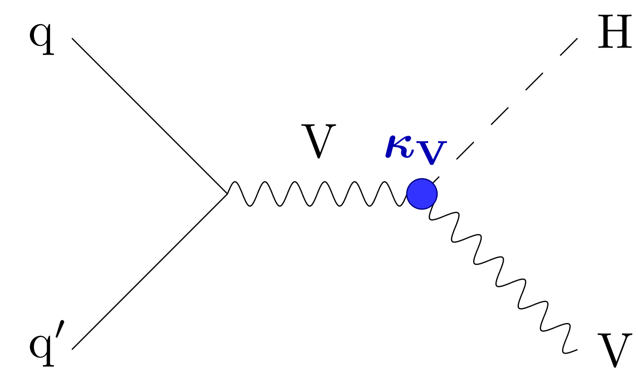

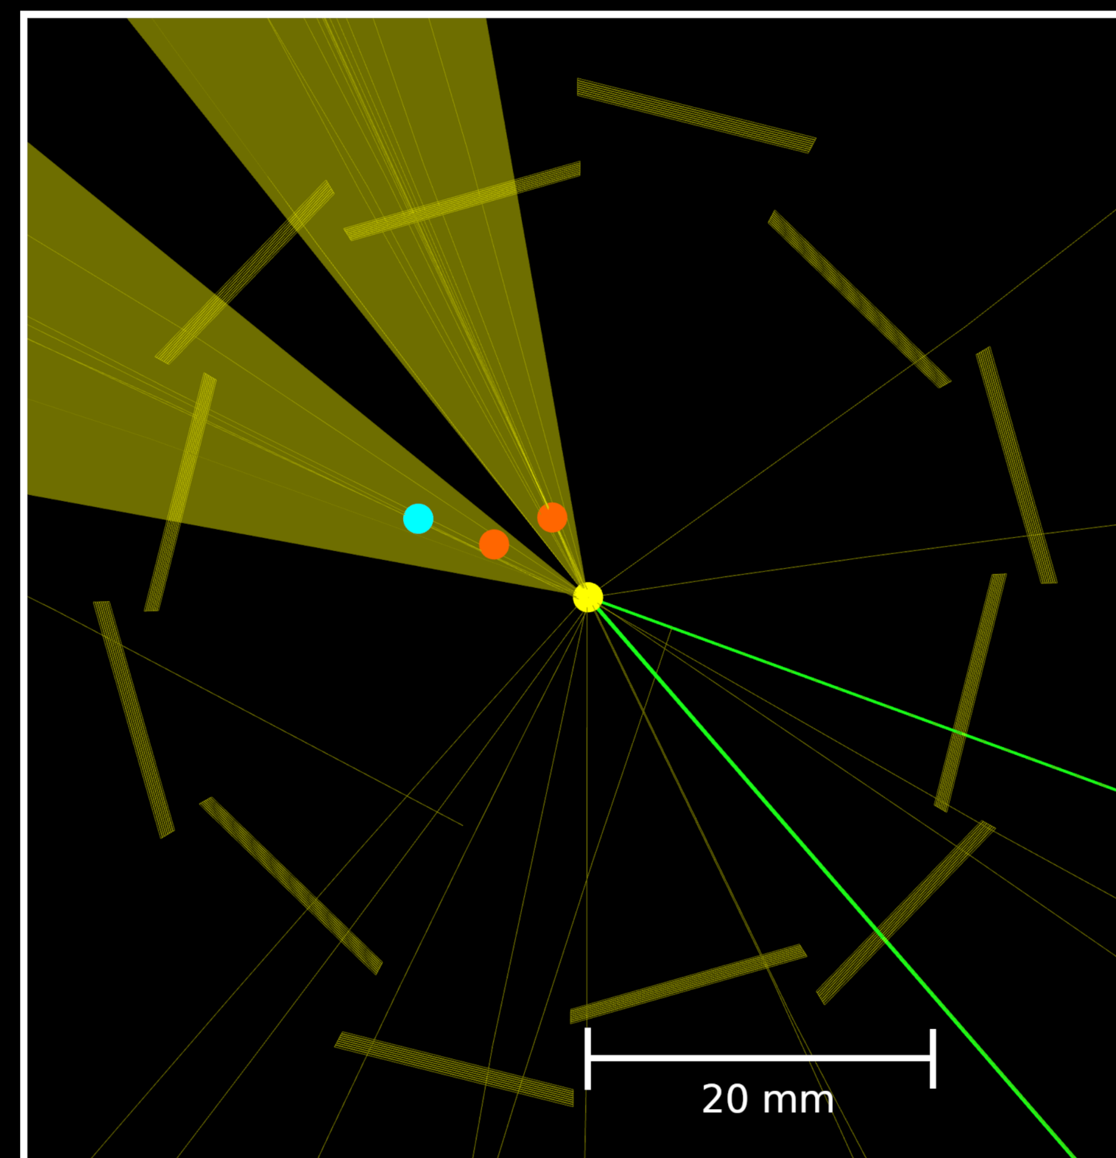

Interesting collisions: Dark Matter?

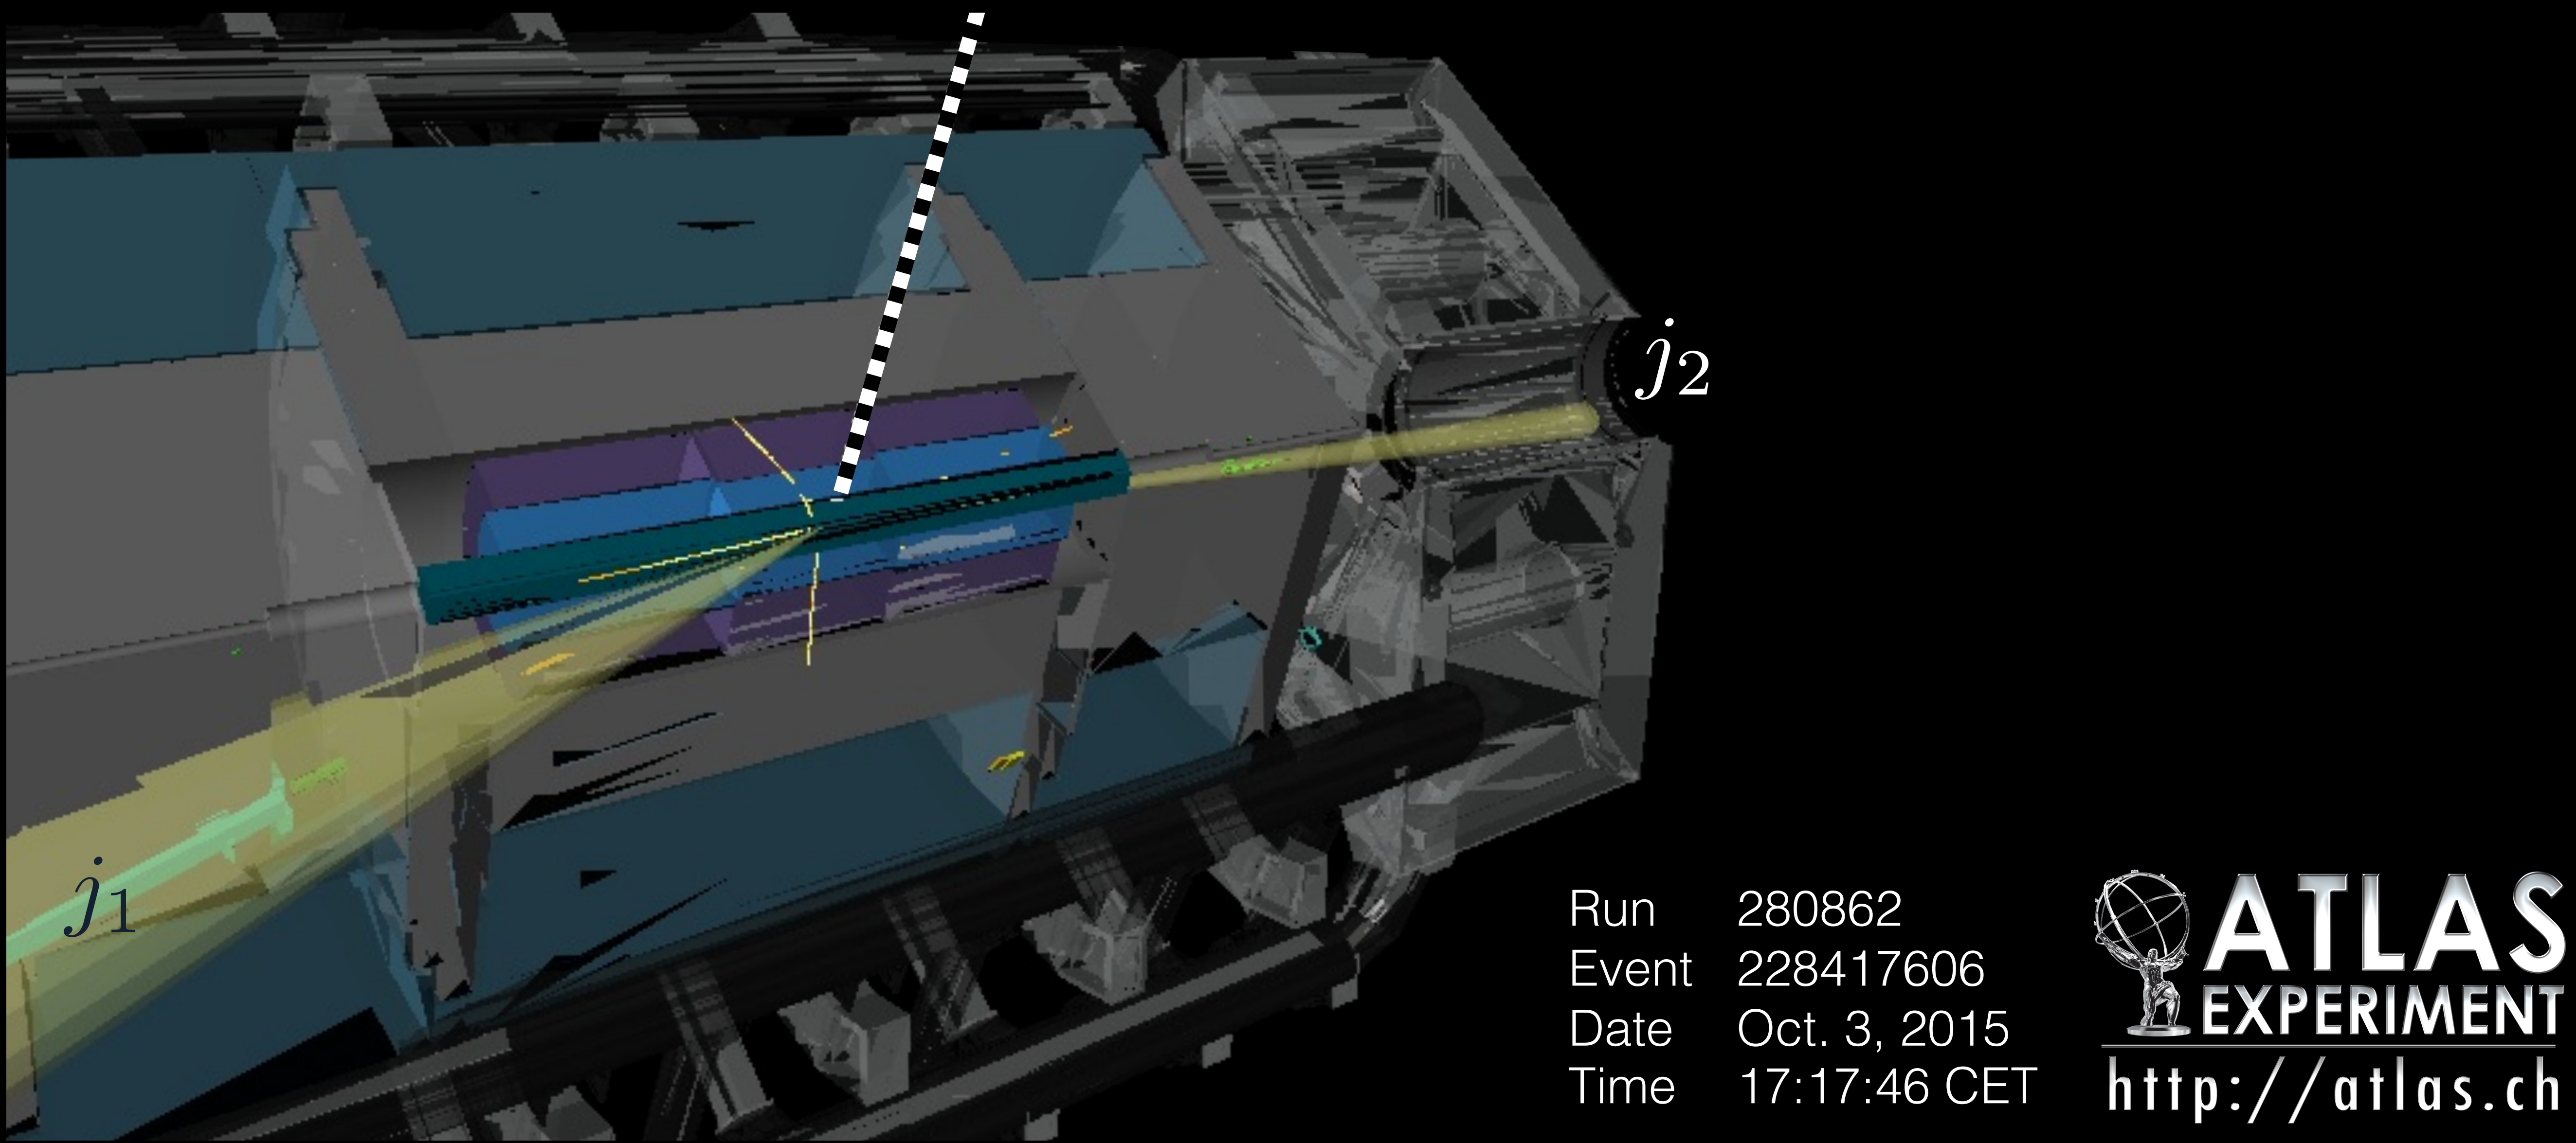

Interesting collisions: Dark Matter?

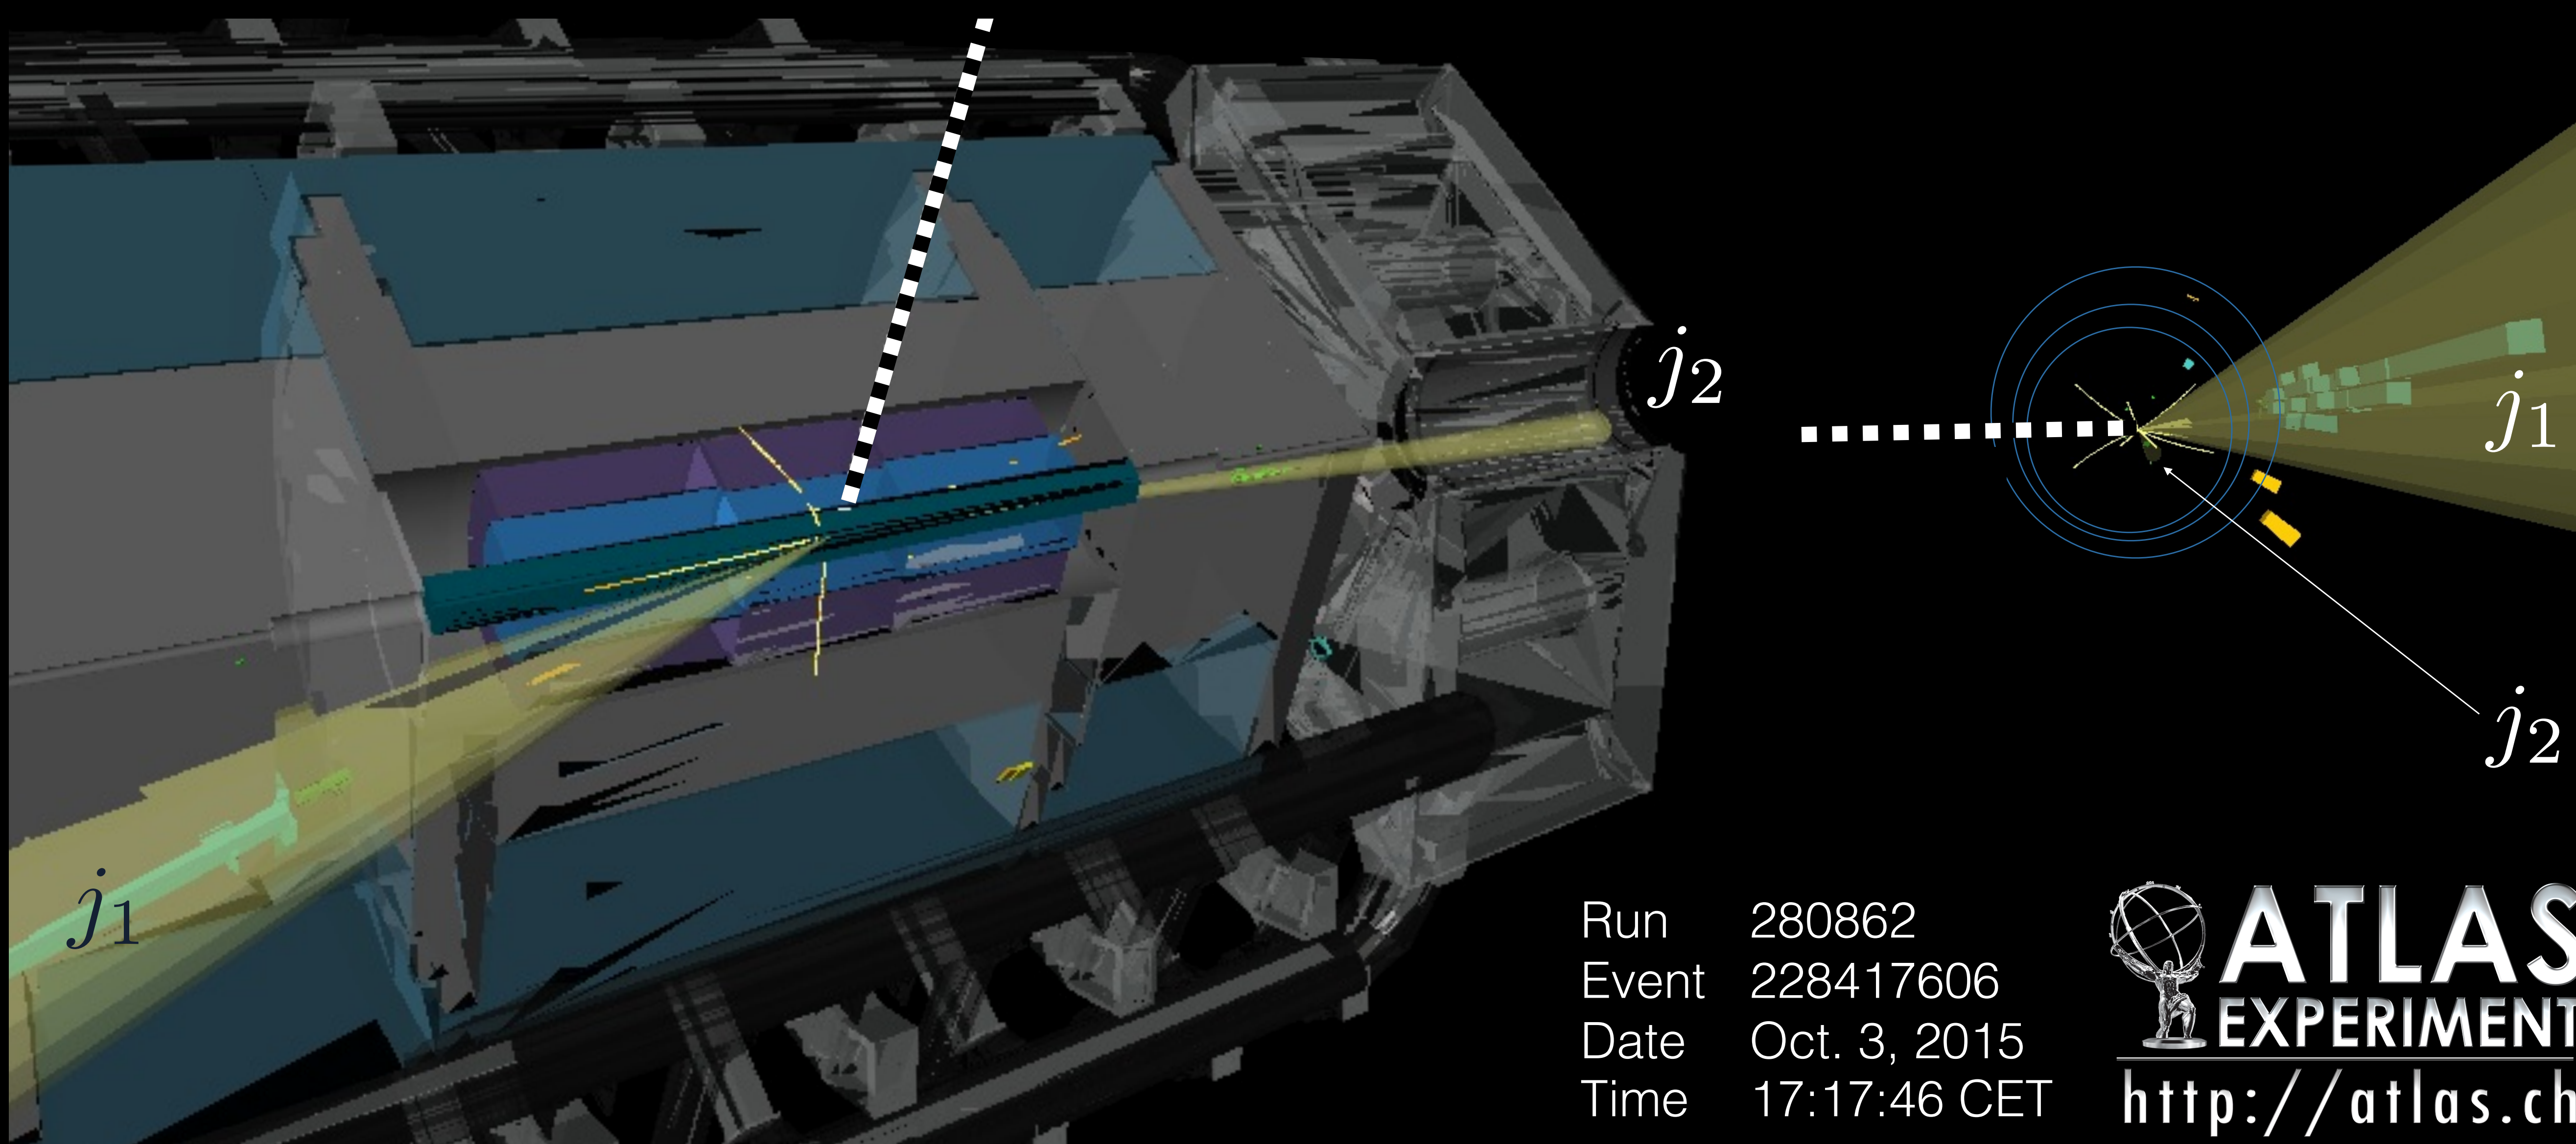

Interesting collisions: Dark Matter?

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

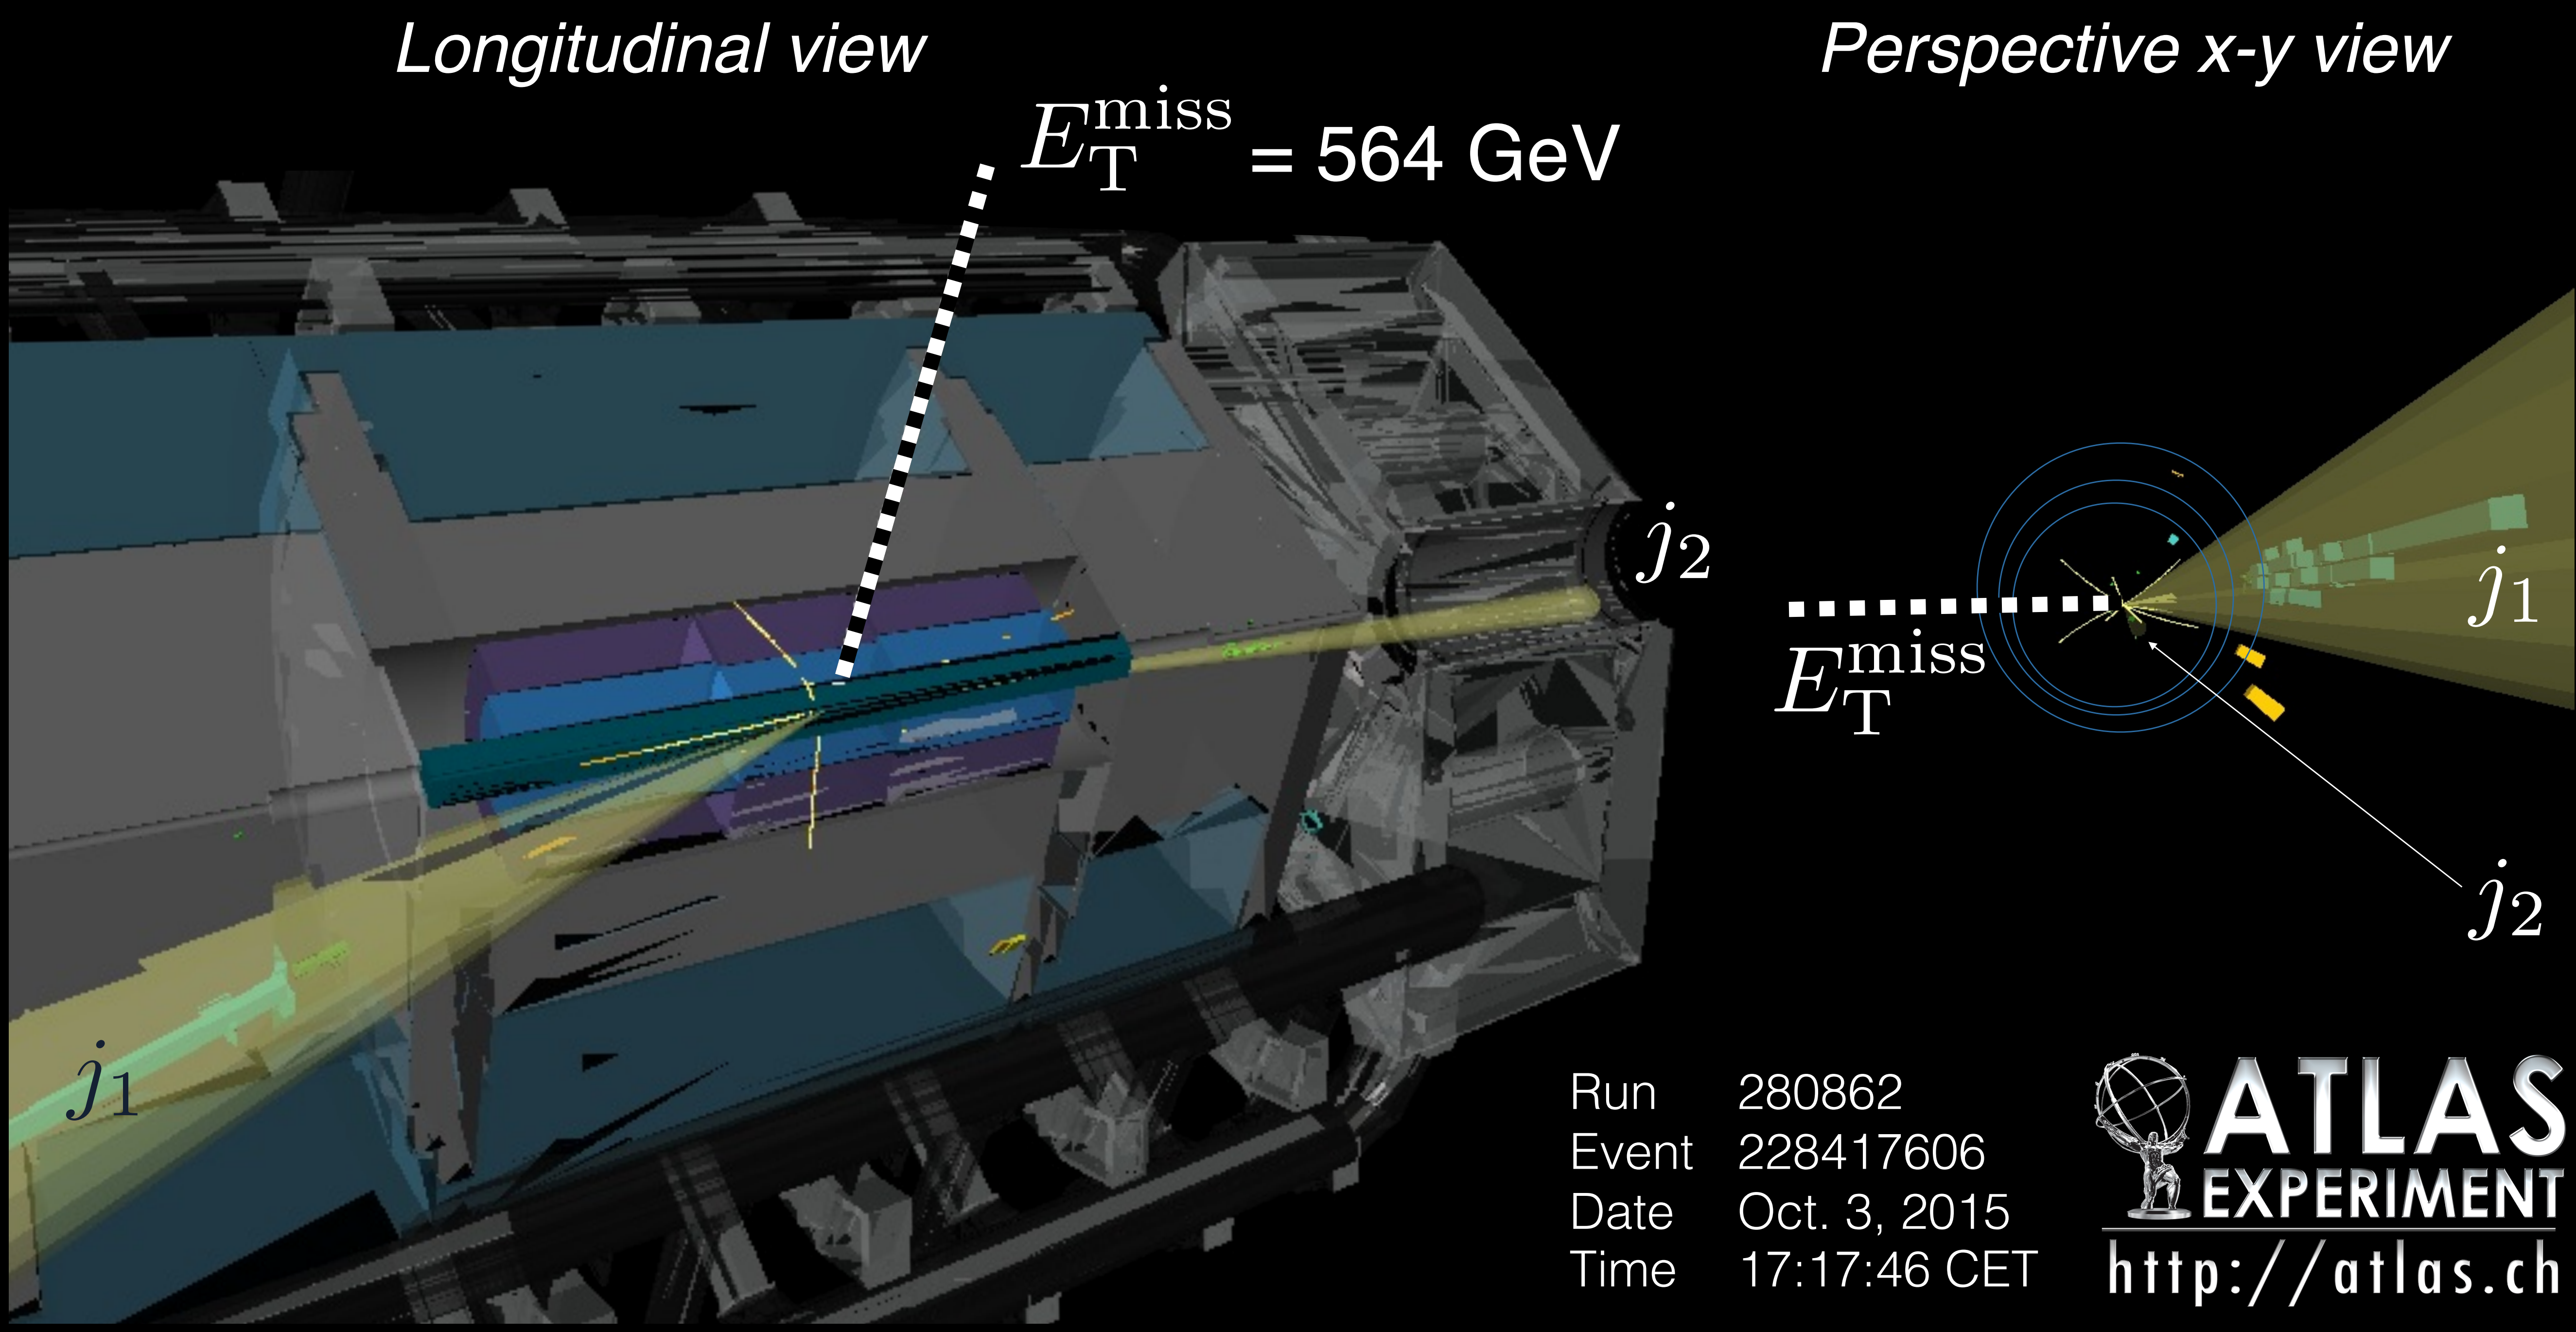

Interesting collisions: top quark partner?

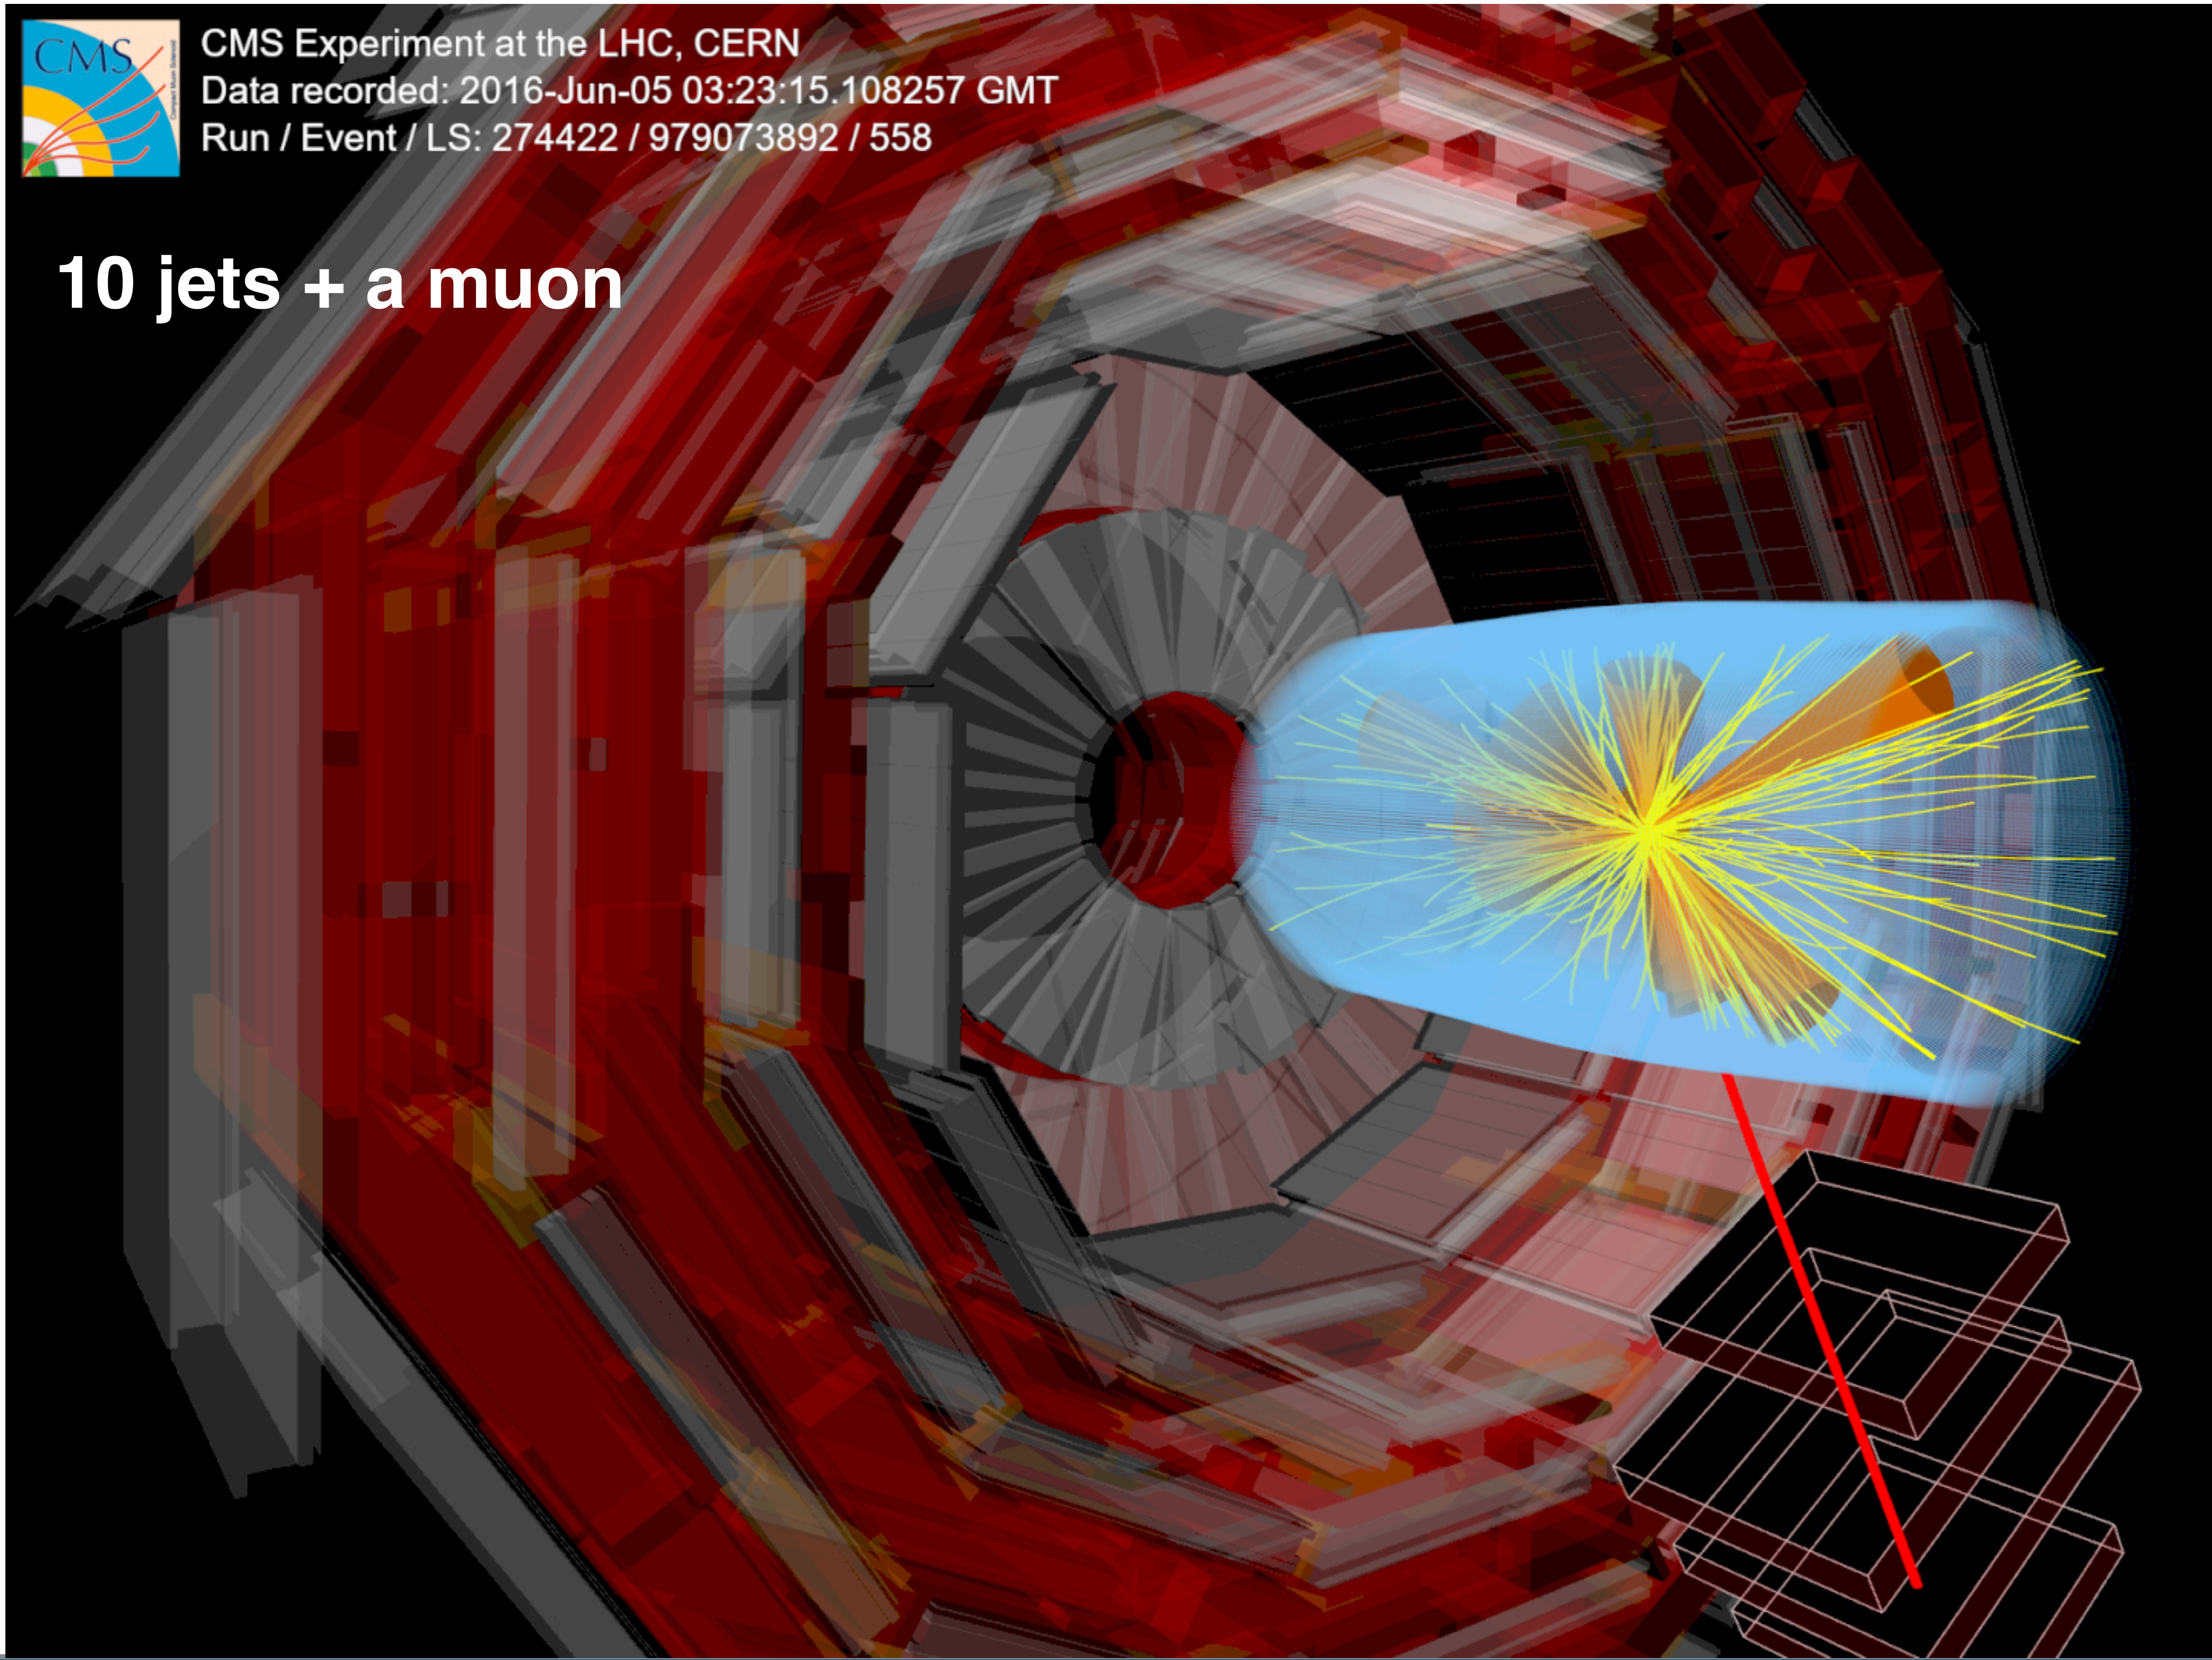

Interesting collisions: top quark partner?

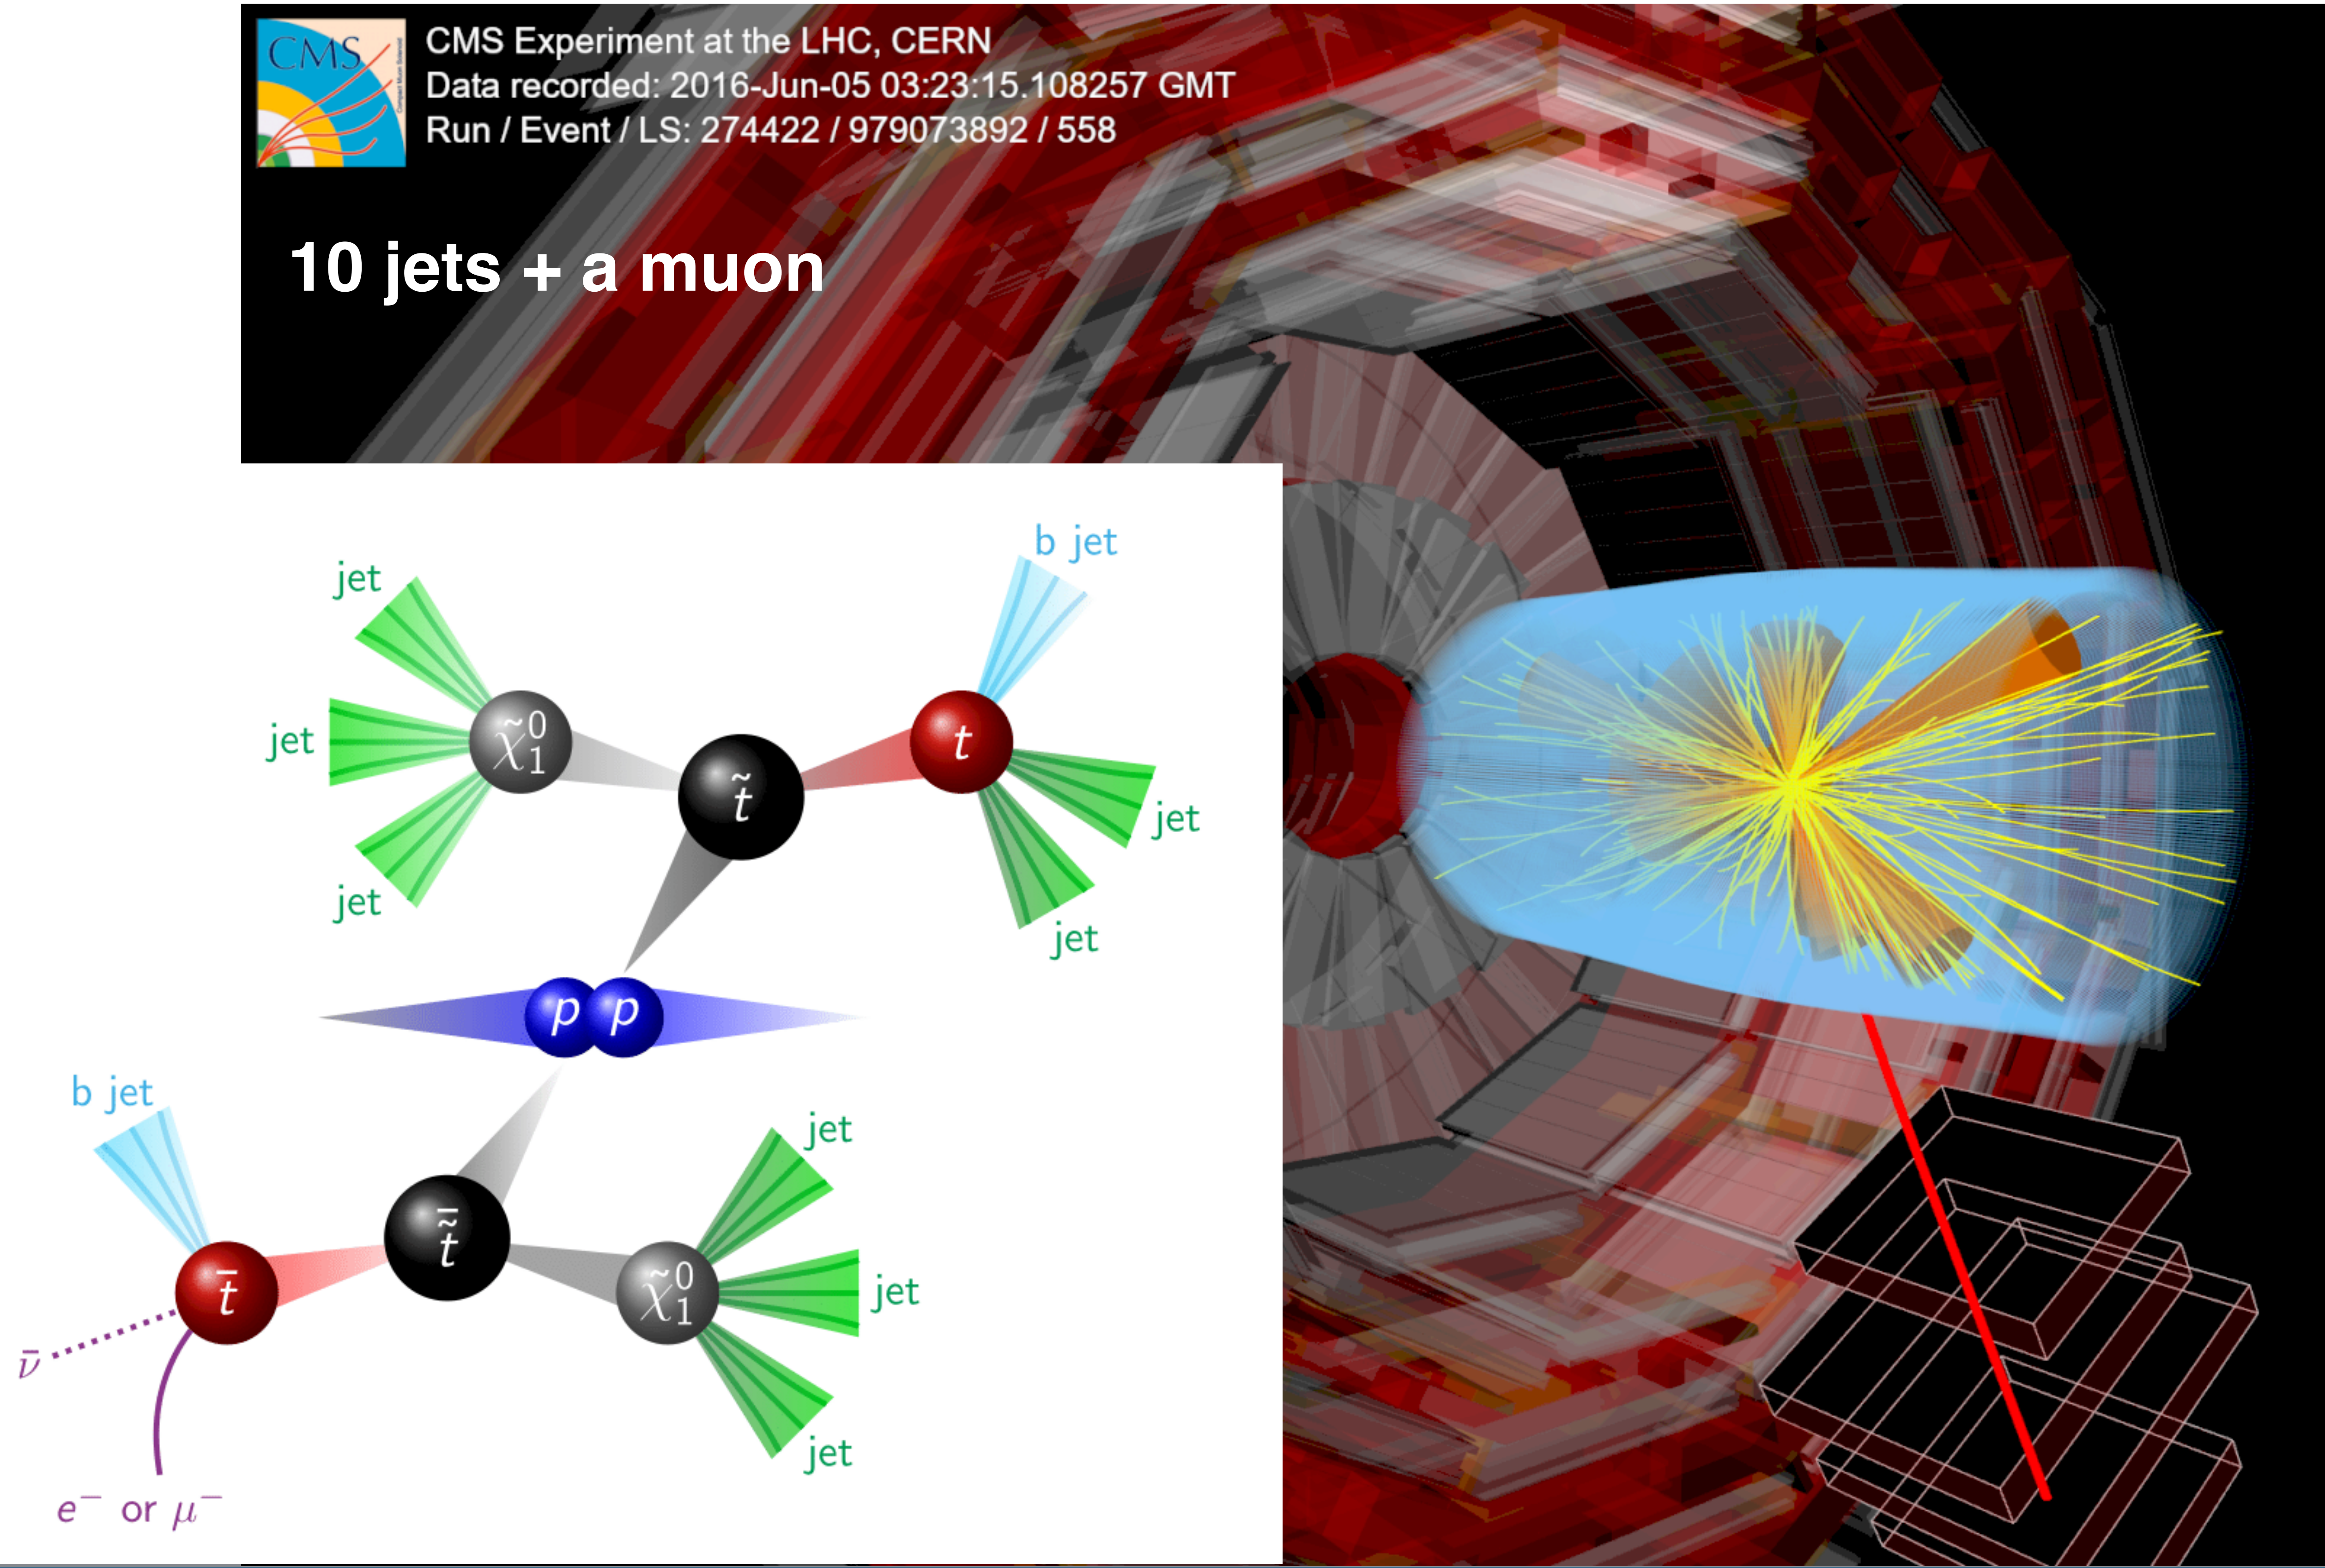

Interesting collisions: b decays?

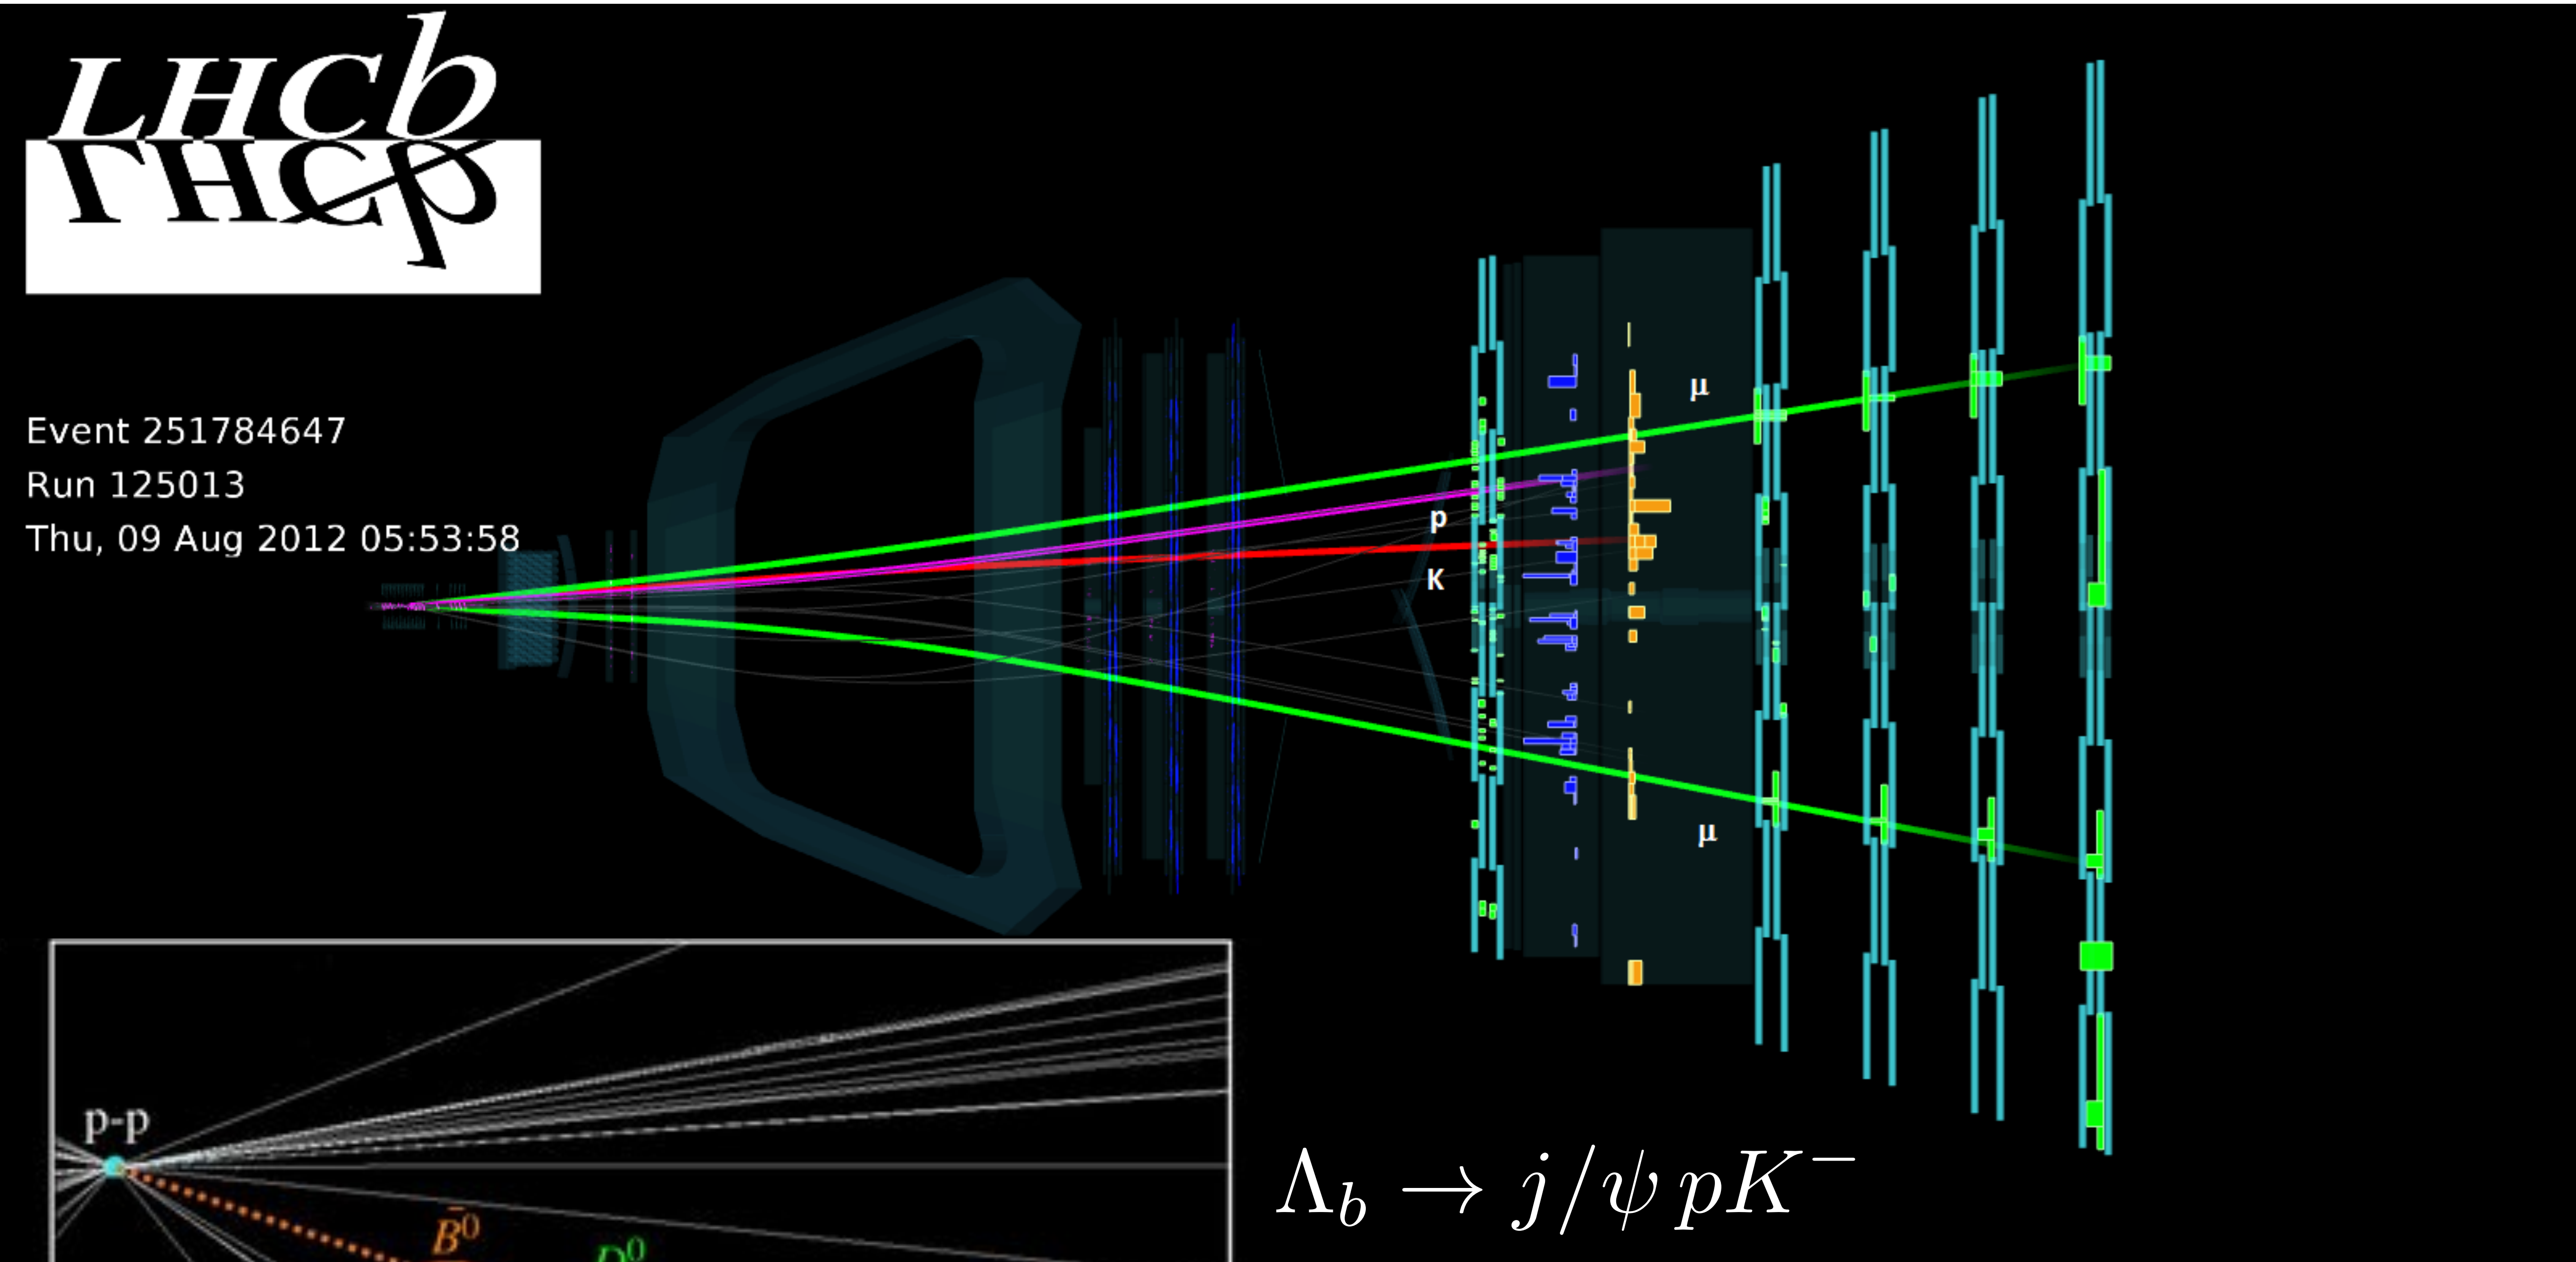

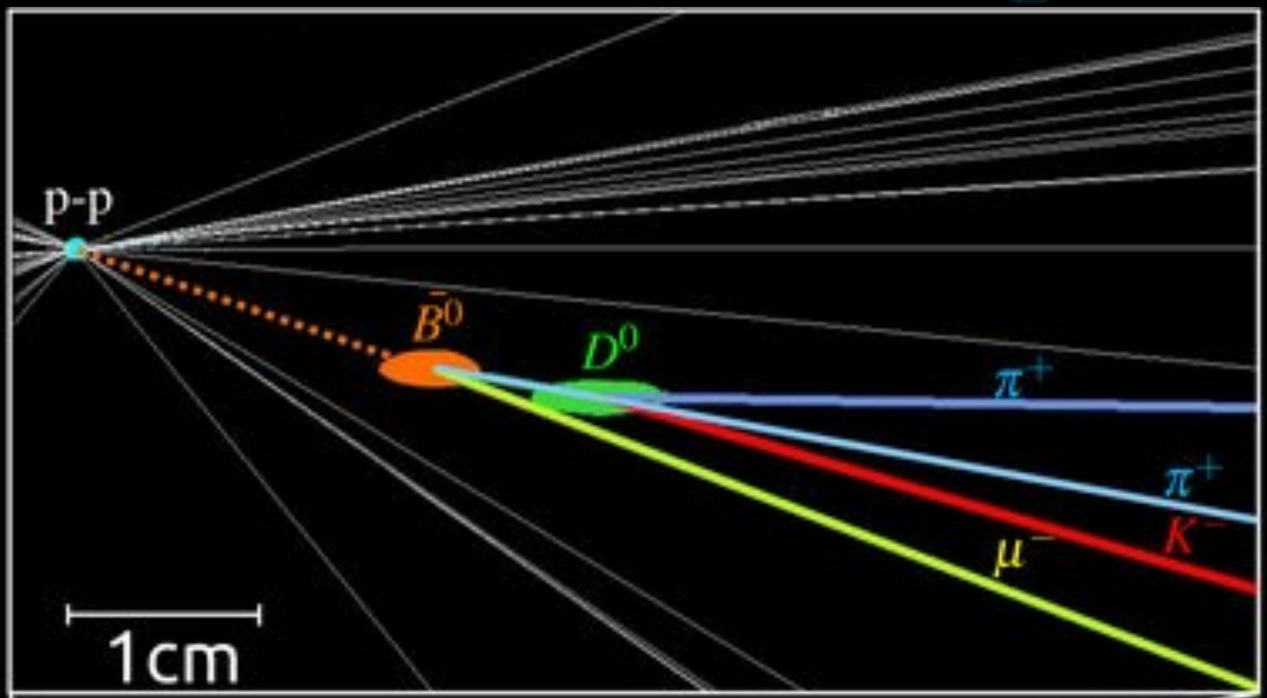

$$\Lambda_b \to j/\psi \, p K^-$$

b hadron decays:

A few low-energy leptons & hadrons

Typical LHC collision: di-jets

- There are many, many ways to produce pairs of jets in LHC collisions
  - 8 gluons + 5x3 light quark combinatorics, and α_S is large!

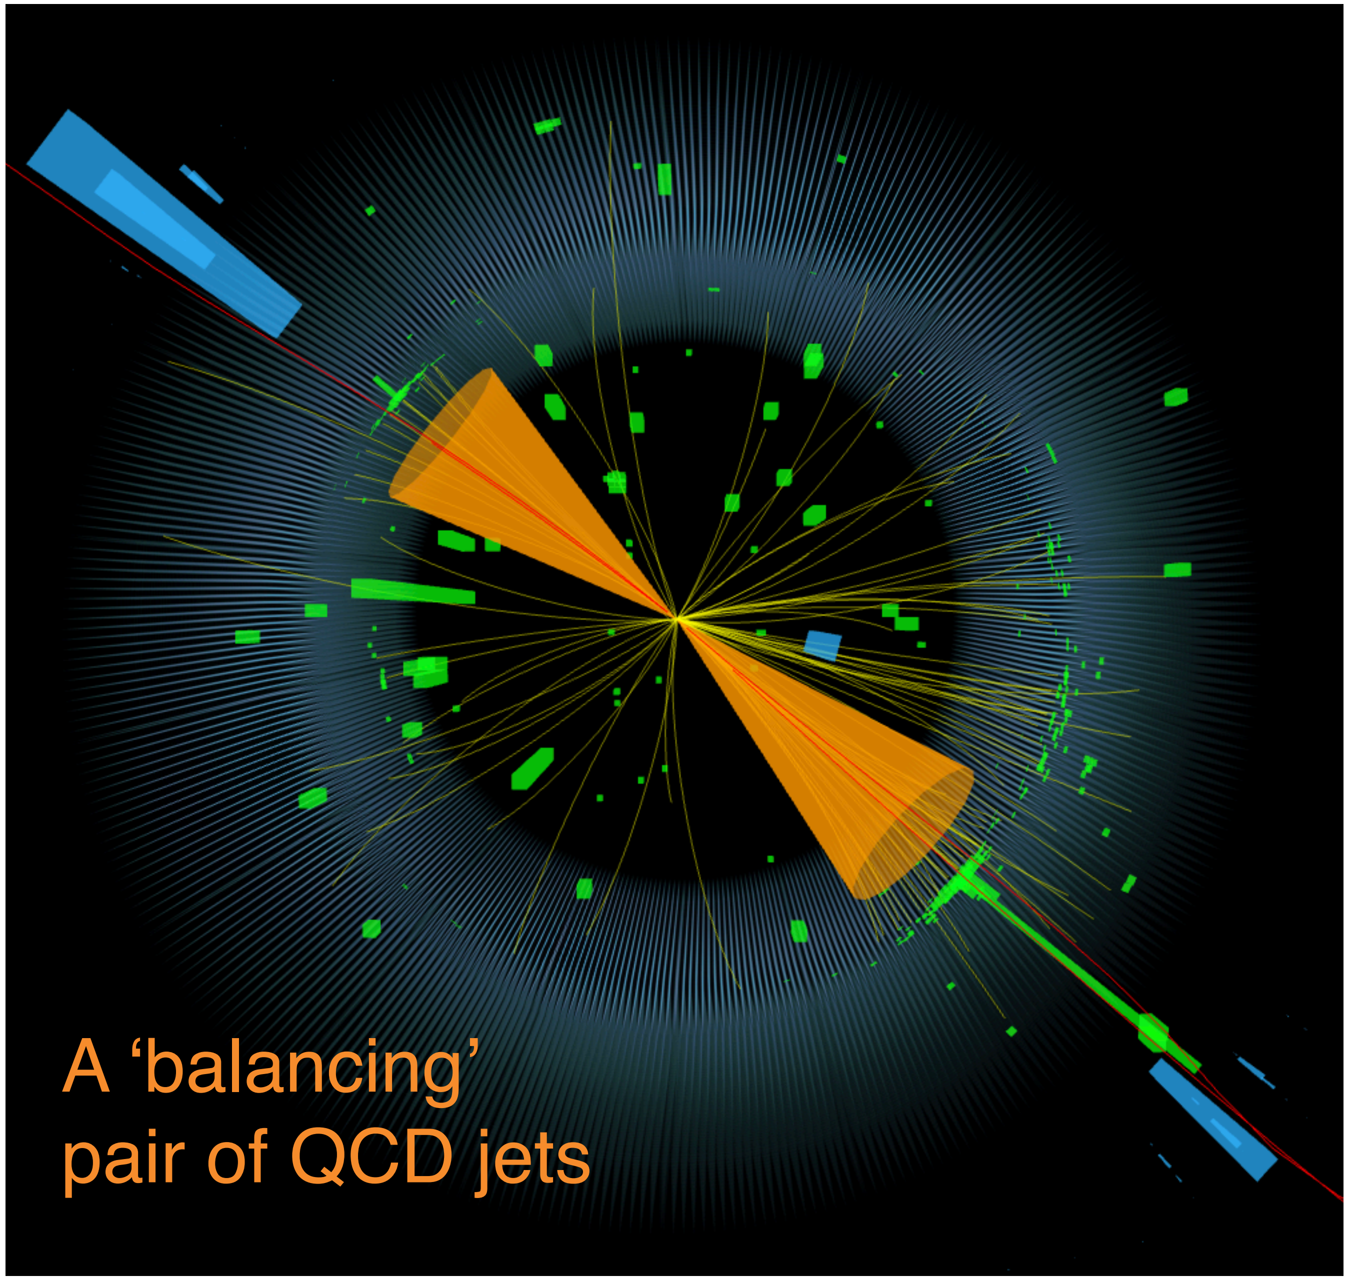

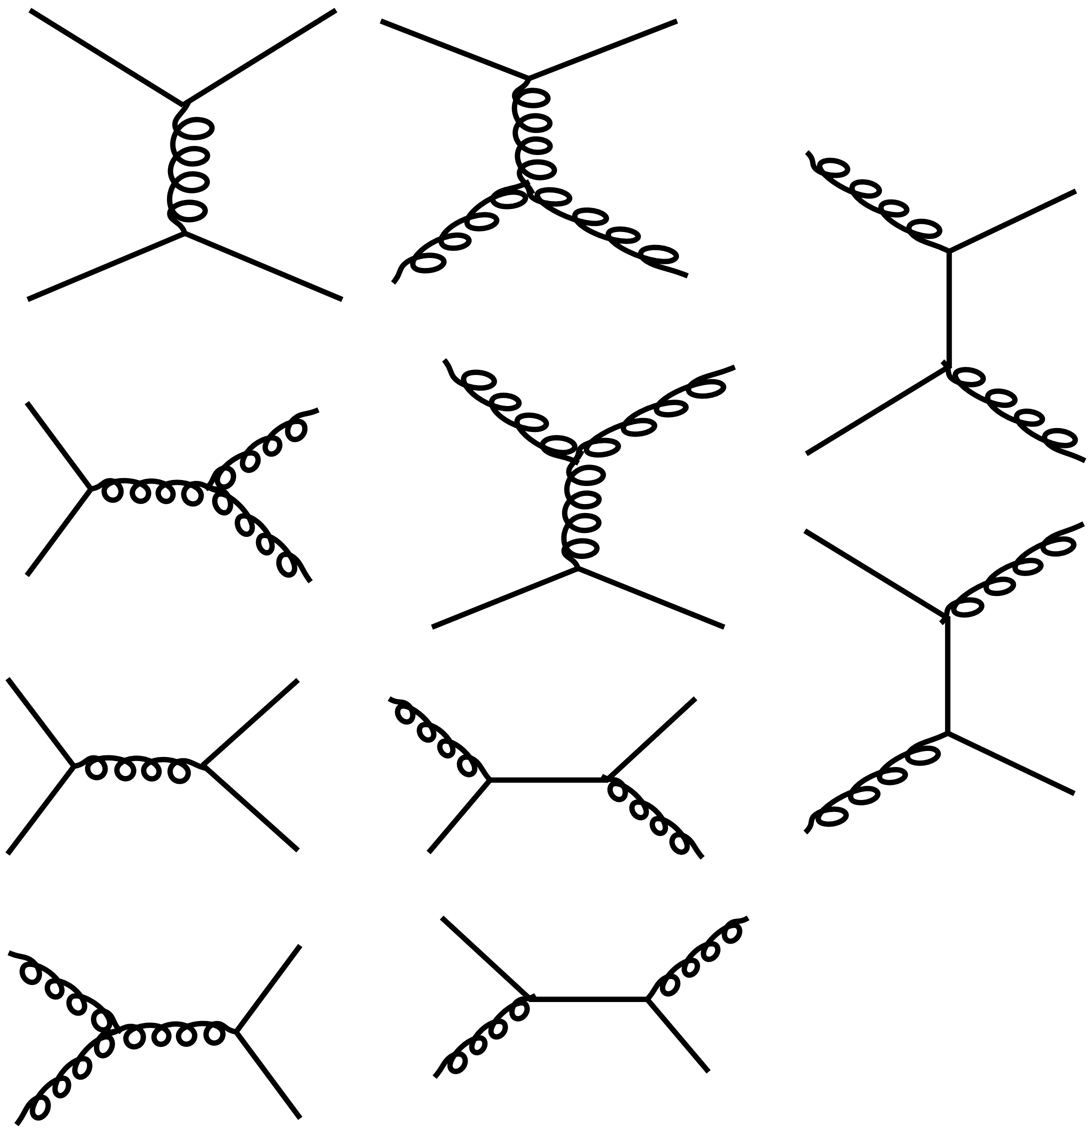

Trade-offs of a hadron collider

Pro: big **σ** & **L** @ HC usually mean the largest samples of target particles

Con: Soft QCD interactions dwarf the typical processes of interest

→ High-p_T jet (100 GeV): 1 in 10^5

→ Single W production: 1 in 10^6

→ Higgs boson: 1 in a Billion

→ Top partner (500 GeV): 1 / 100 Billion

Would be more convenient if our collider ONLY made B, W, Z, Higgs, BSM, ….

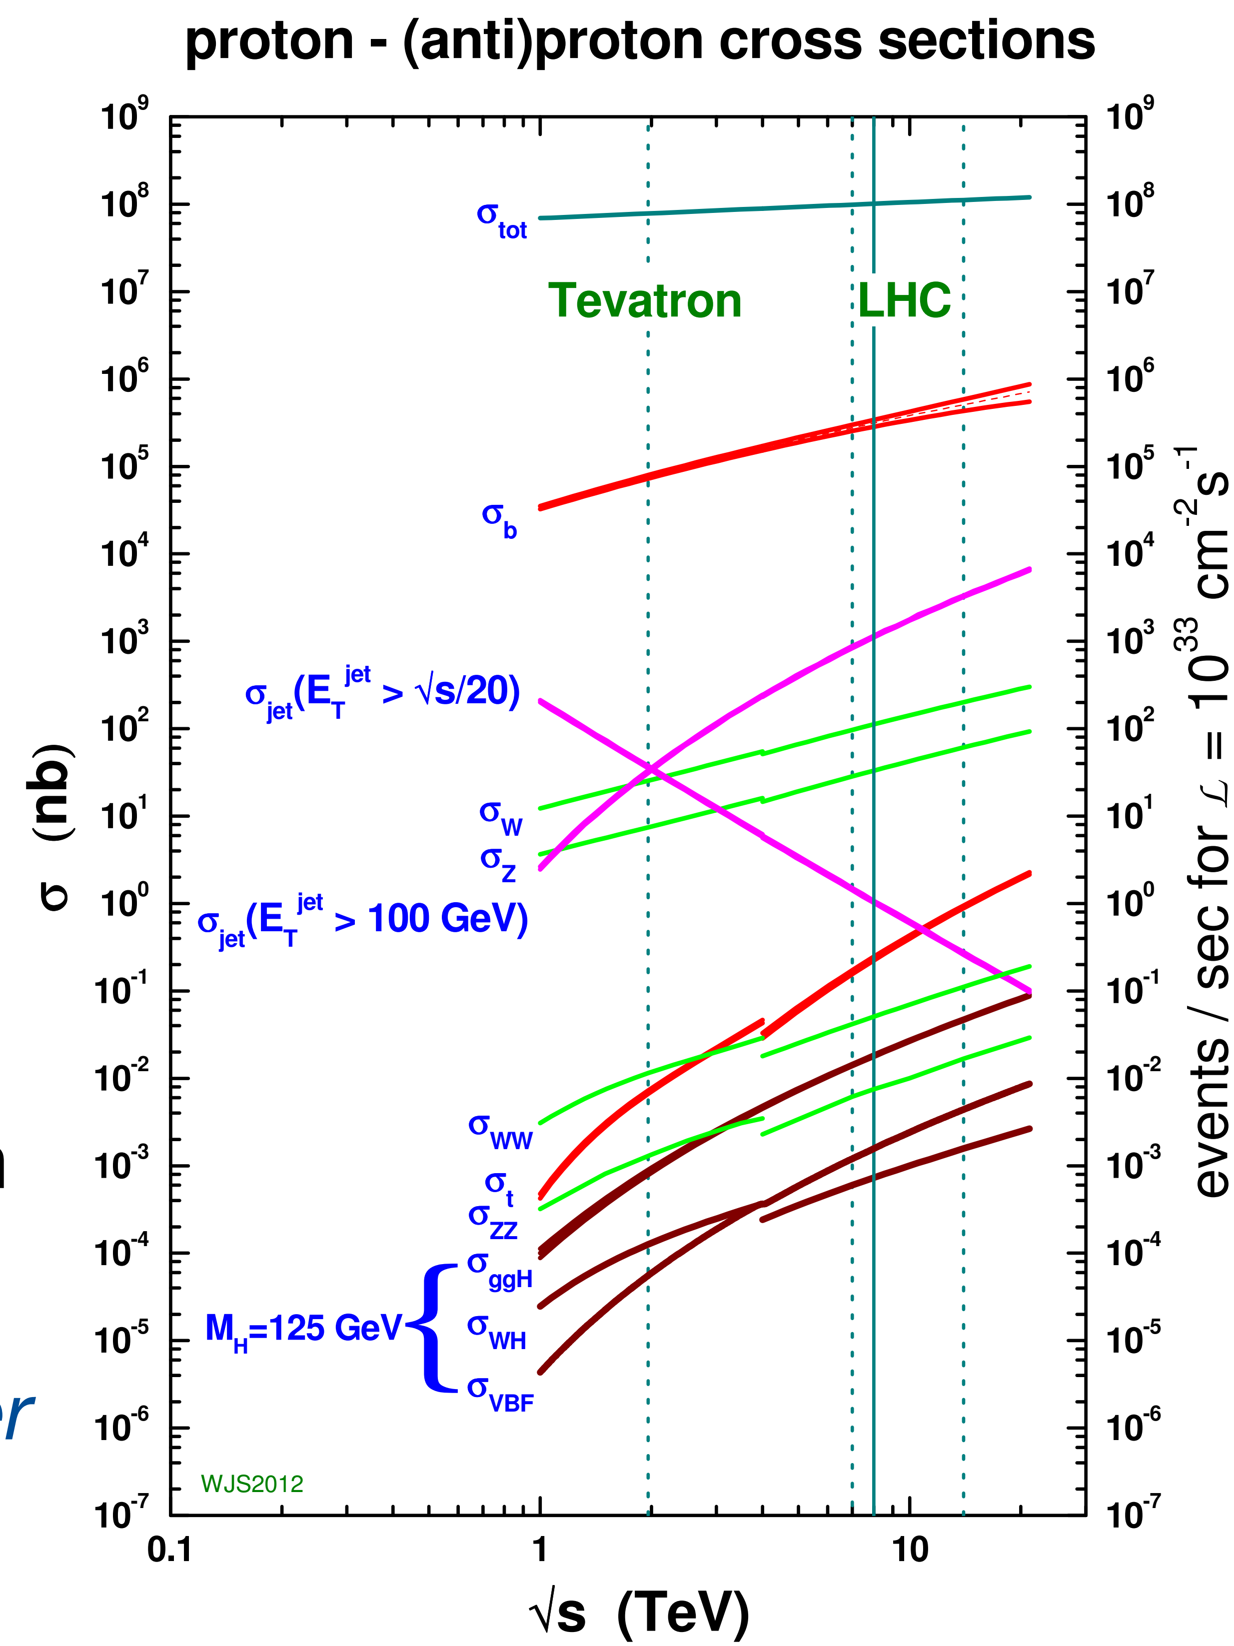

Trigger system: the next best thing

A system to filter events, recording only the "signatures of interest"

(Ideally $\varepsilon_B \sim 10^{-5}$ and $\varepsilon_S \sim 1$)

Without a trigger

Trigger system: the next best thing

A system to filter events, recording only the "signatures of interest"

(Ideally $\varepsilon_B \sim 10^{-5}$ and $\varepsilon_S \sim 1$)

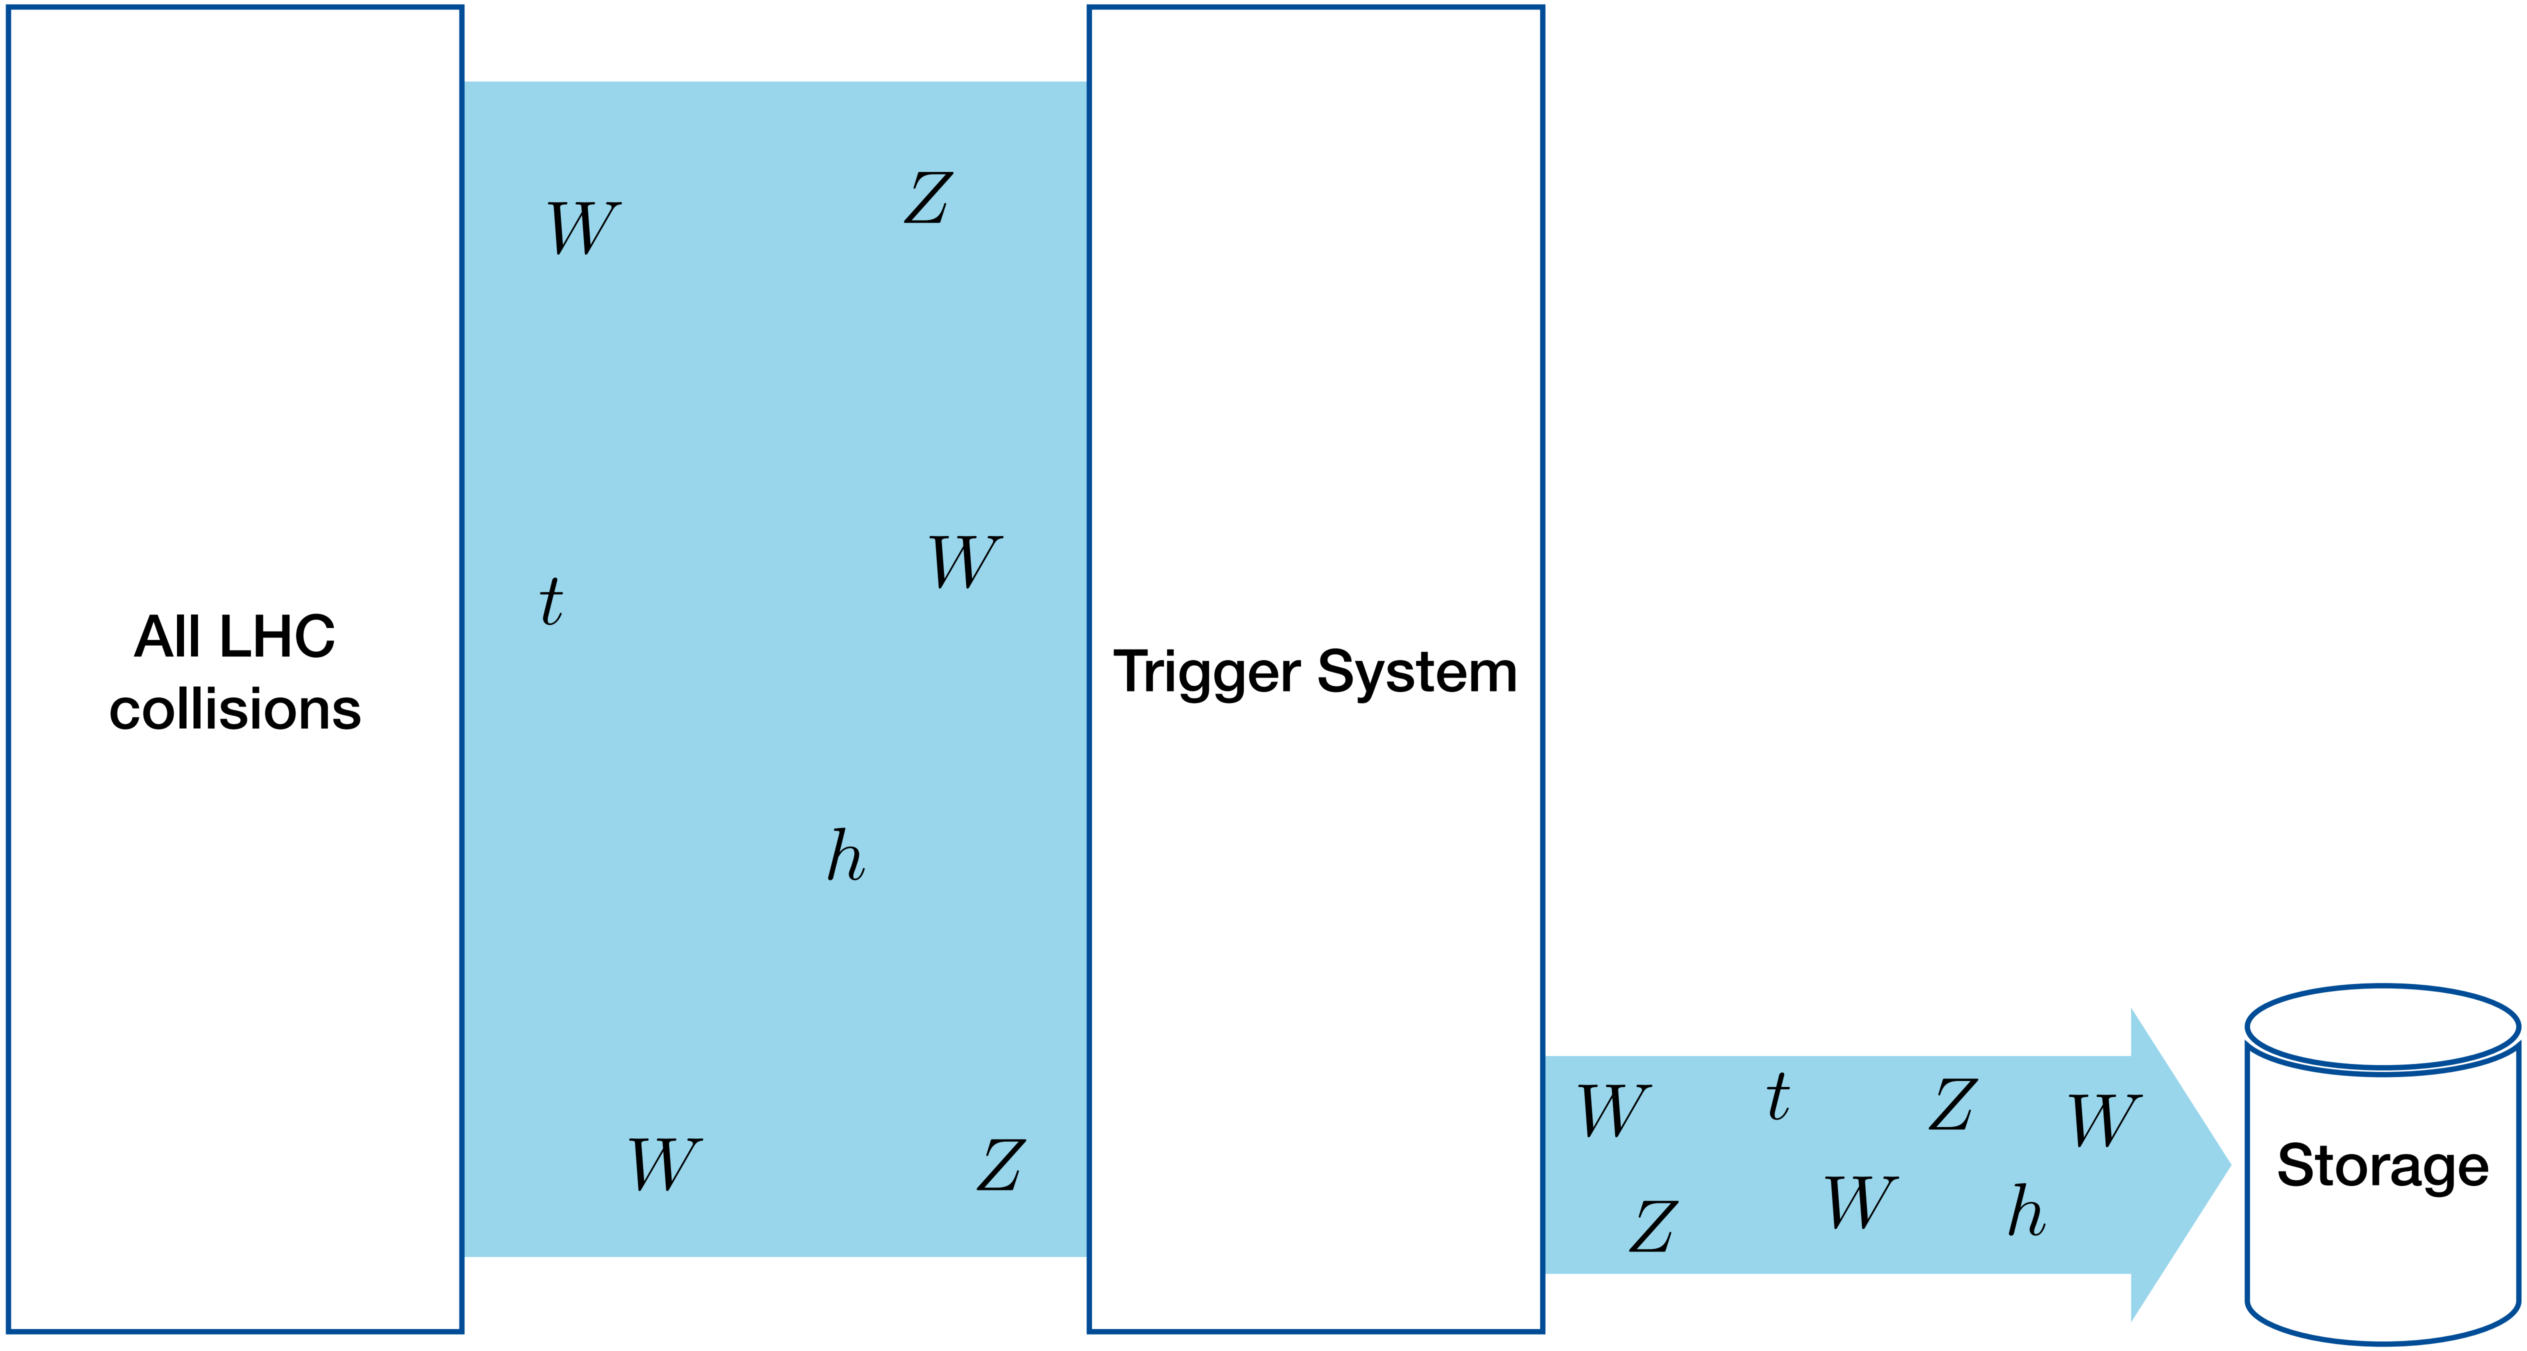

Trigger system: the next best thing

A system to filter events, recording only the "signatures of interest"

(Ideally $\varepsilon_B \sim 10^{-5}$ and $\varepsilon_S \sim 1$)

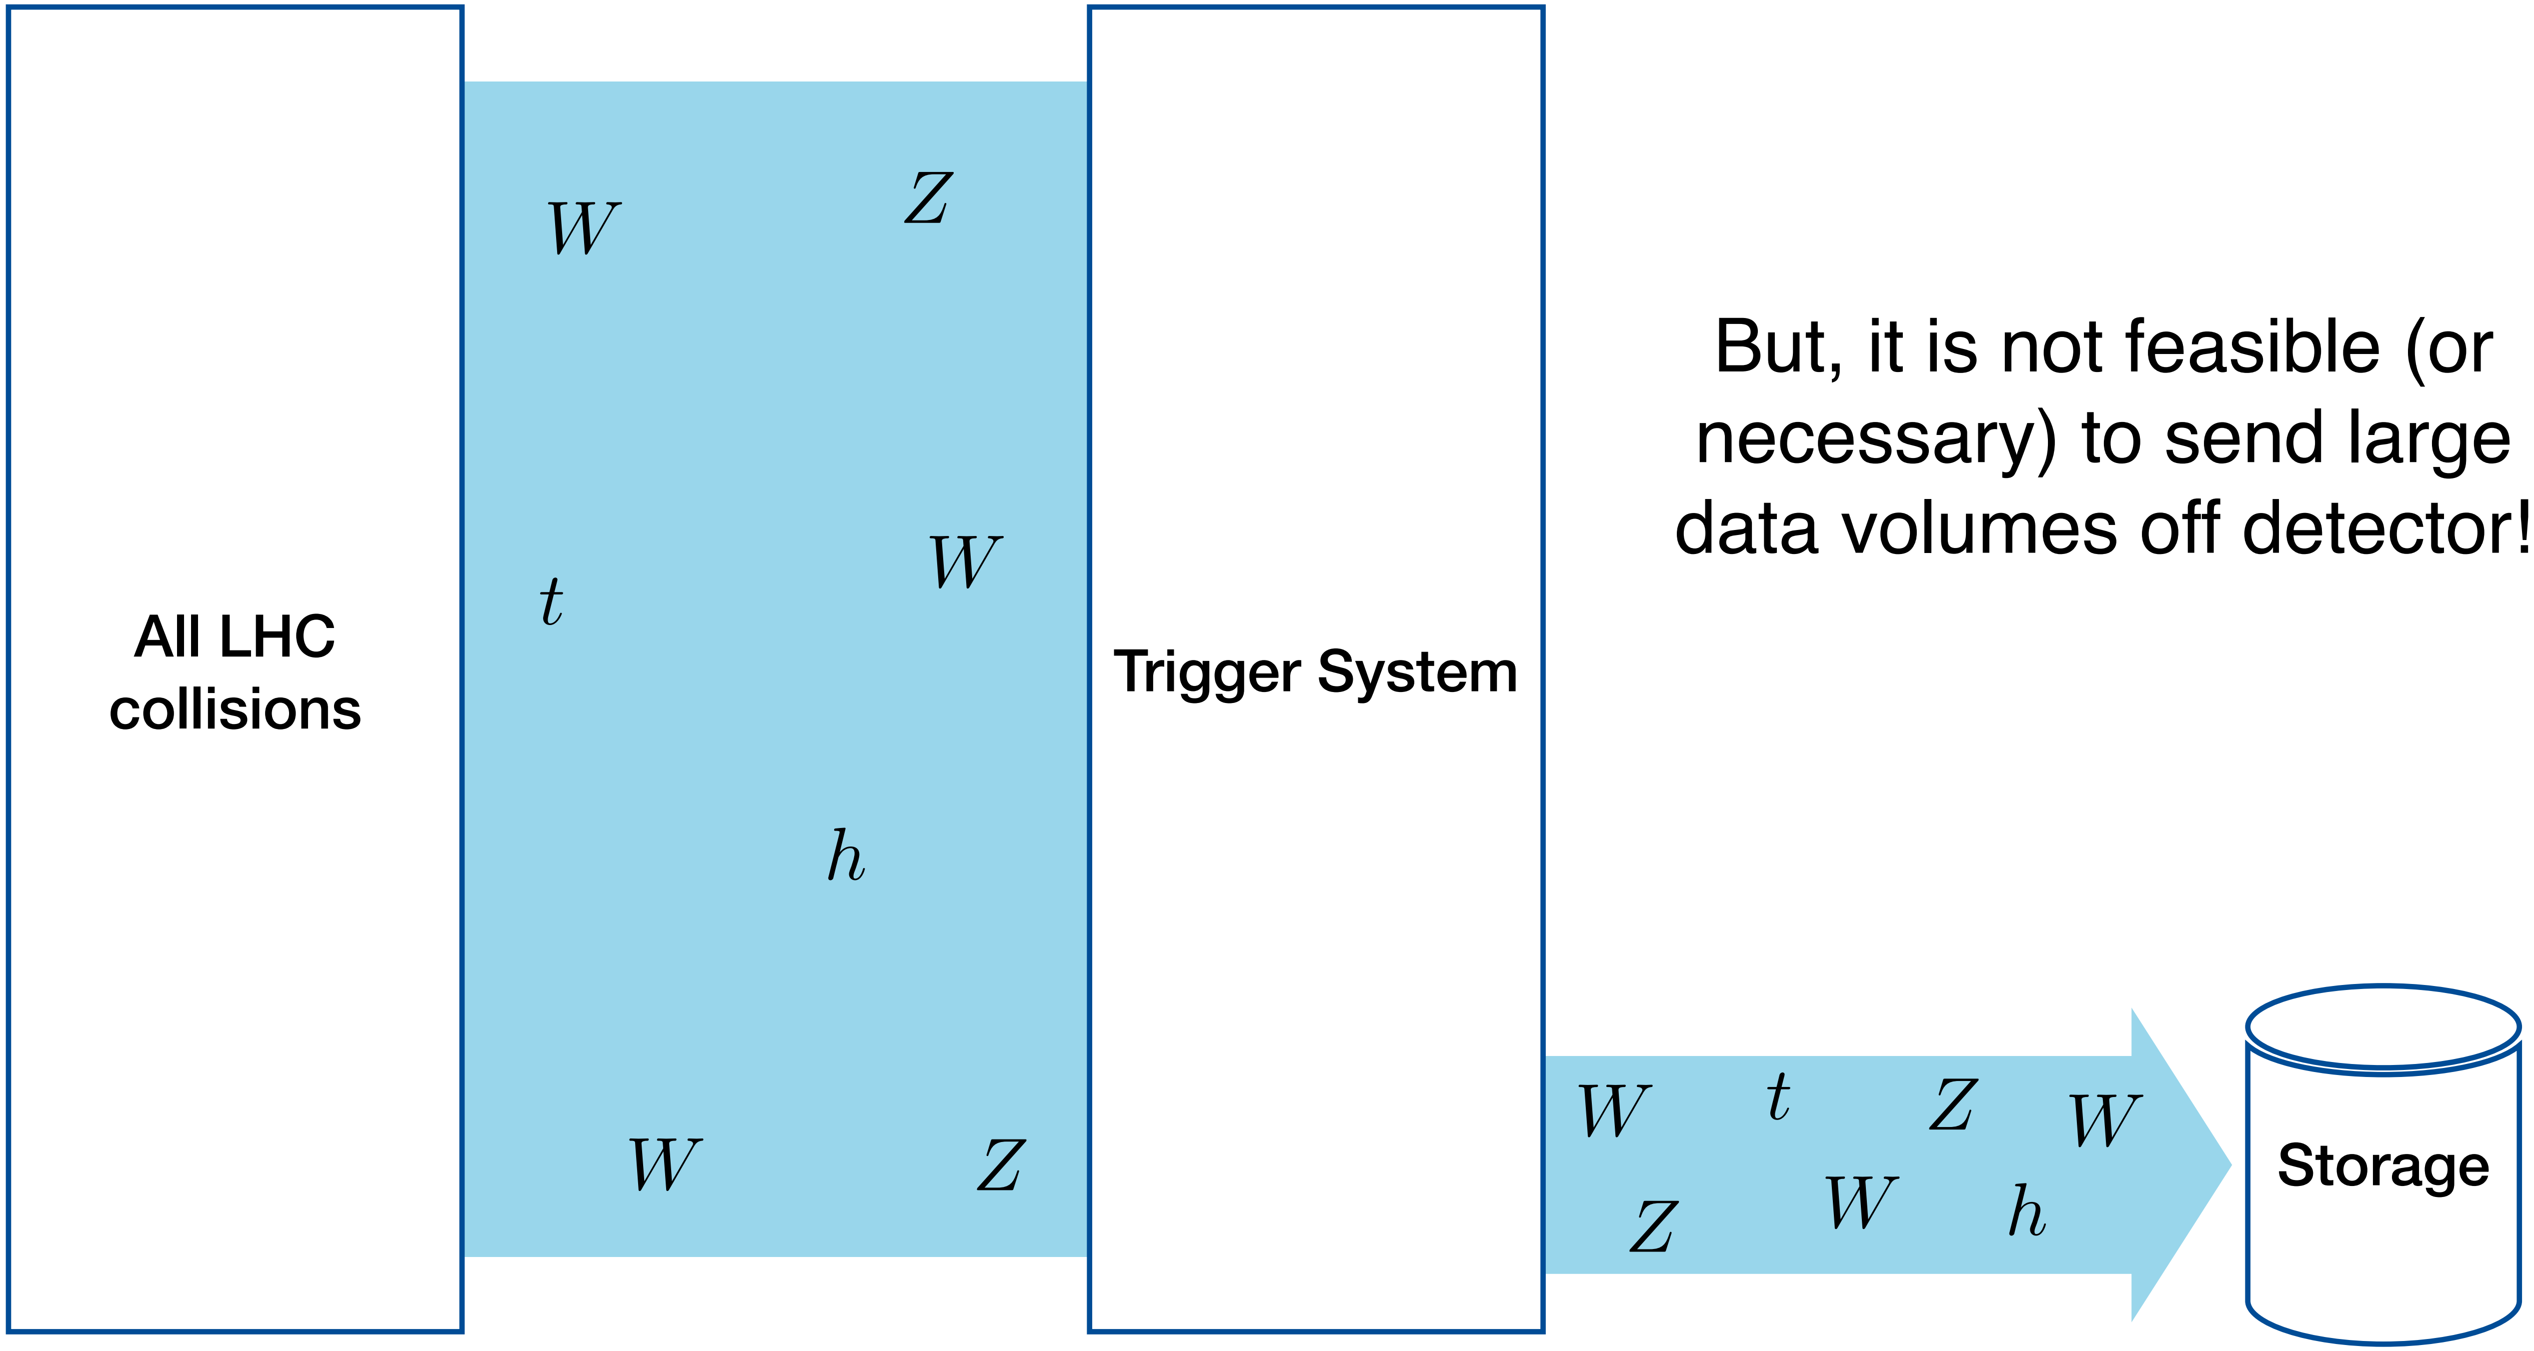

But, it is not feasible (or necessary) to send large data volumes off detector!

Trigger system: the next best thing

A system to filter events, recording only the “signatures of interest”

(Ideally $\varepsilon_B \sim 10^{-5}$ and $\varepsilon_S \sim 1$)

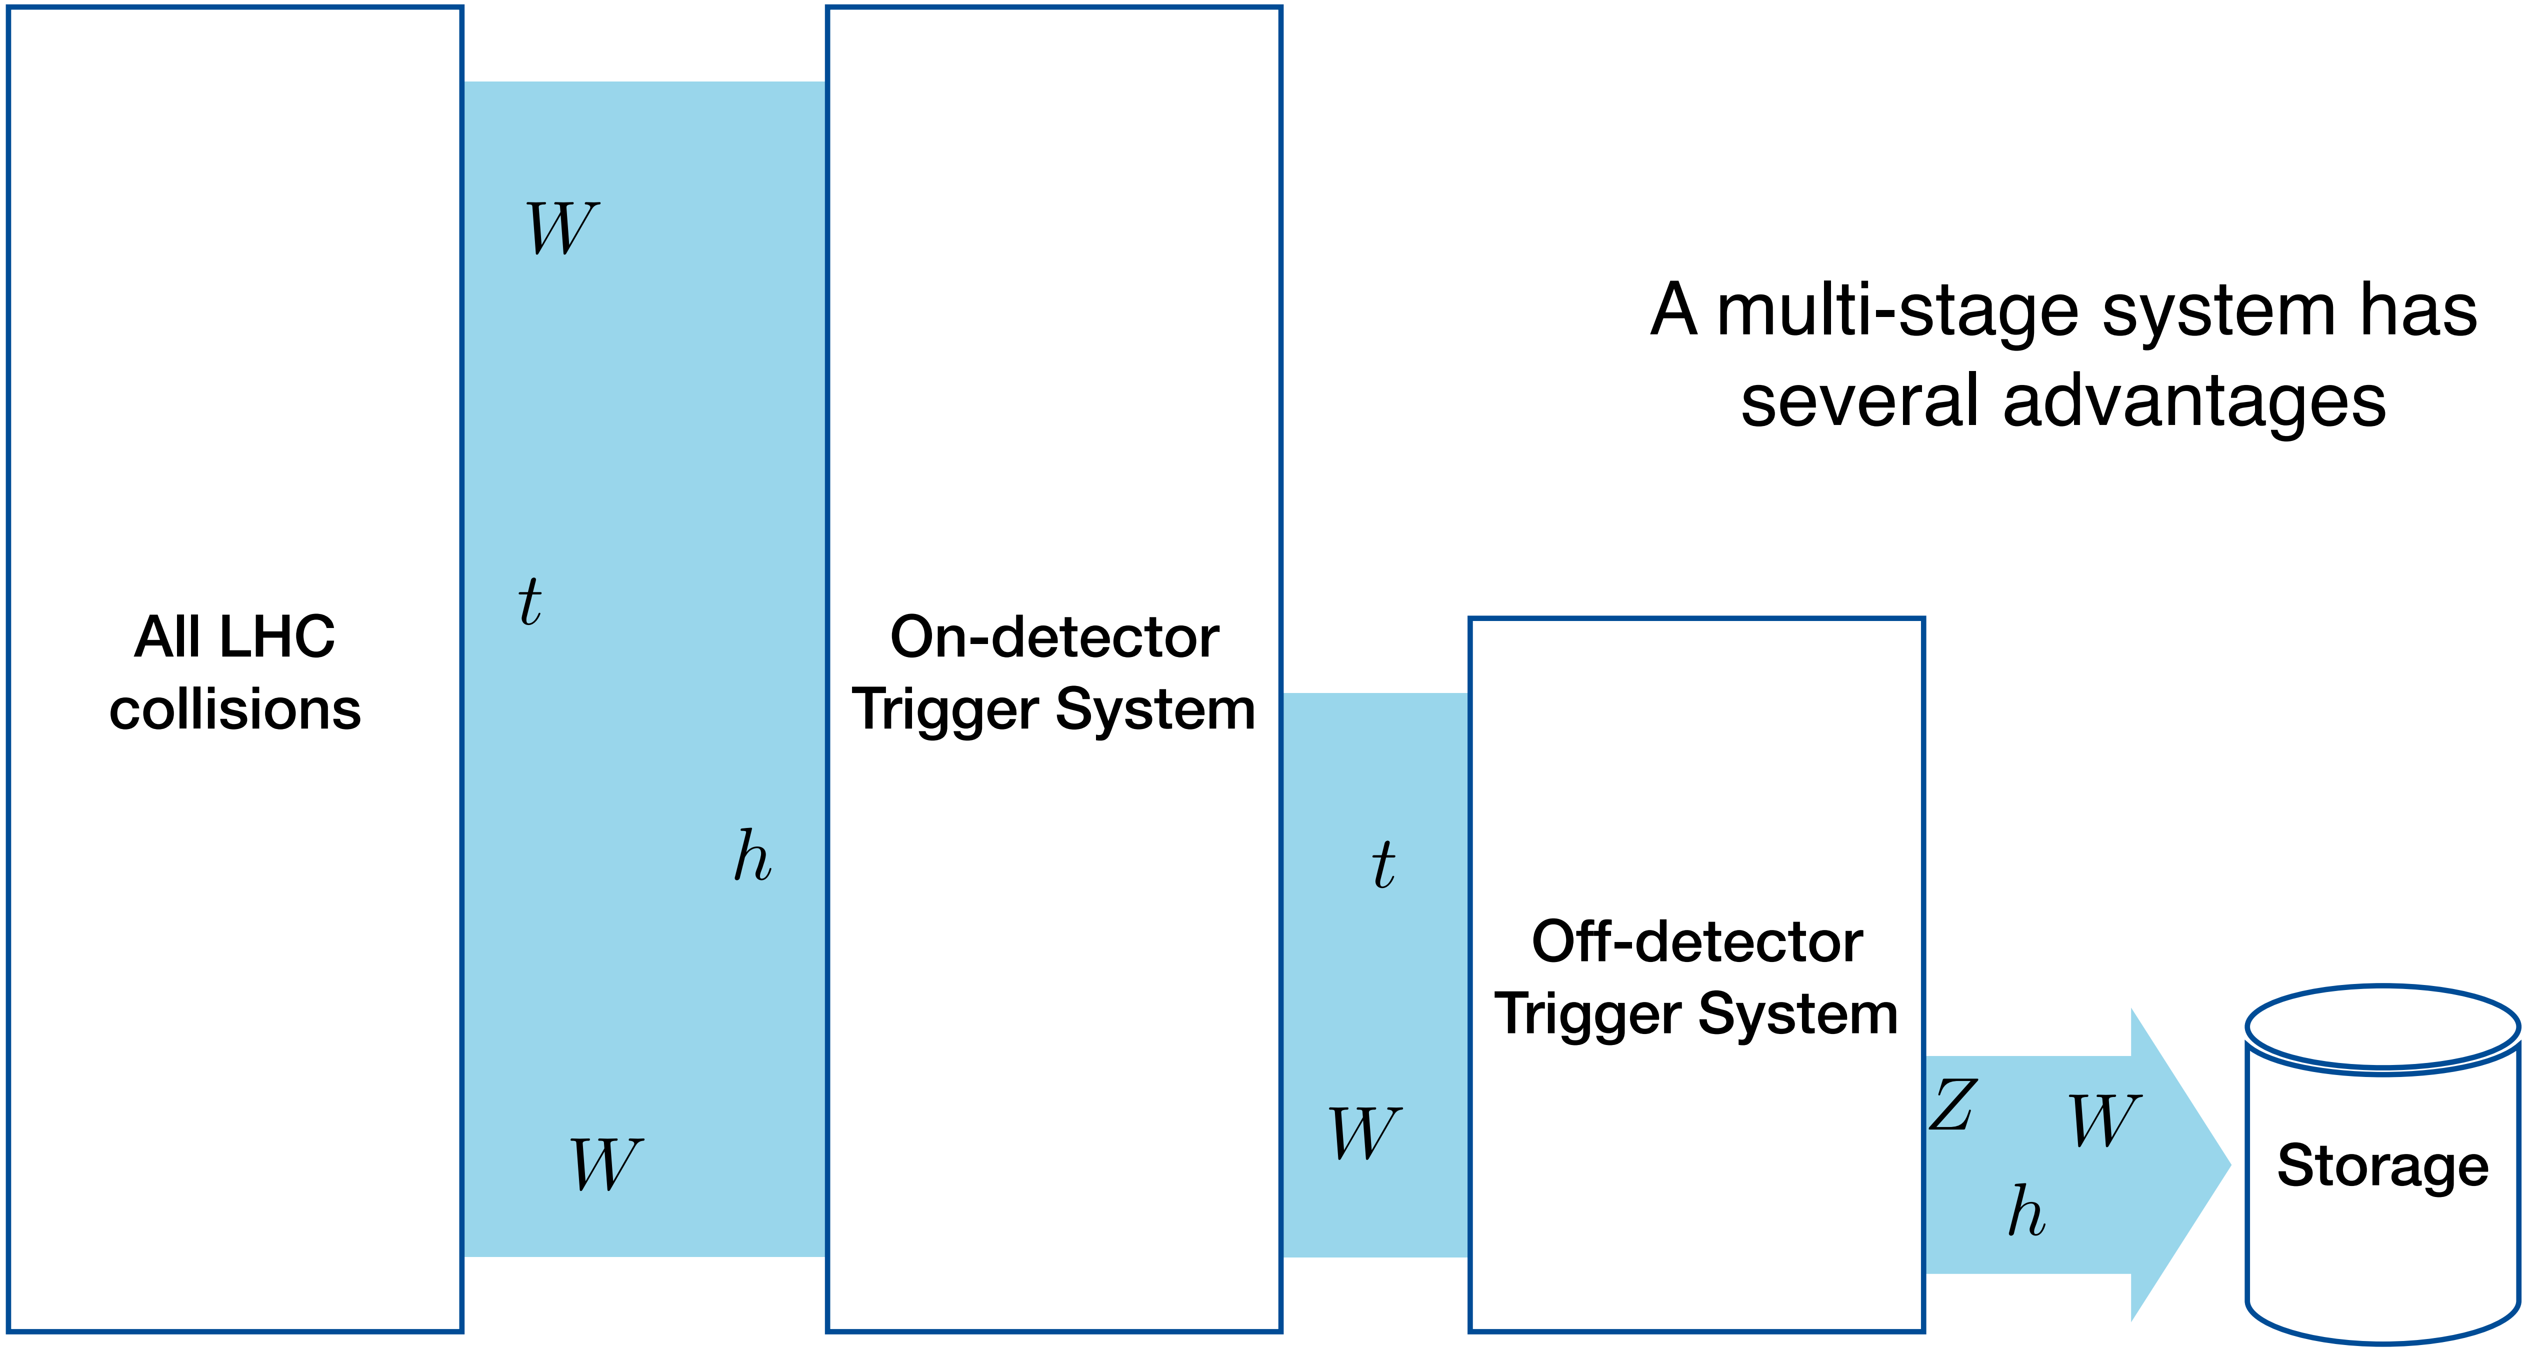

A multi-stage system has several advantages

Adding a little more realism

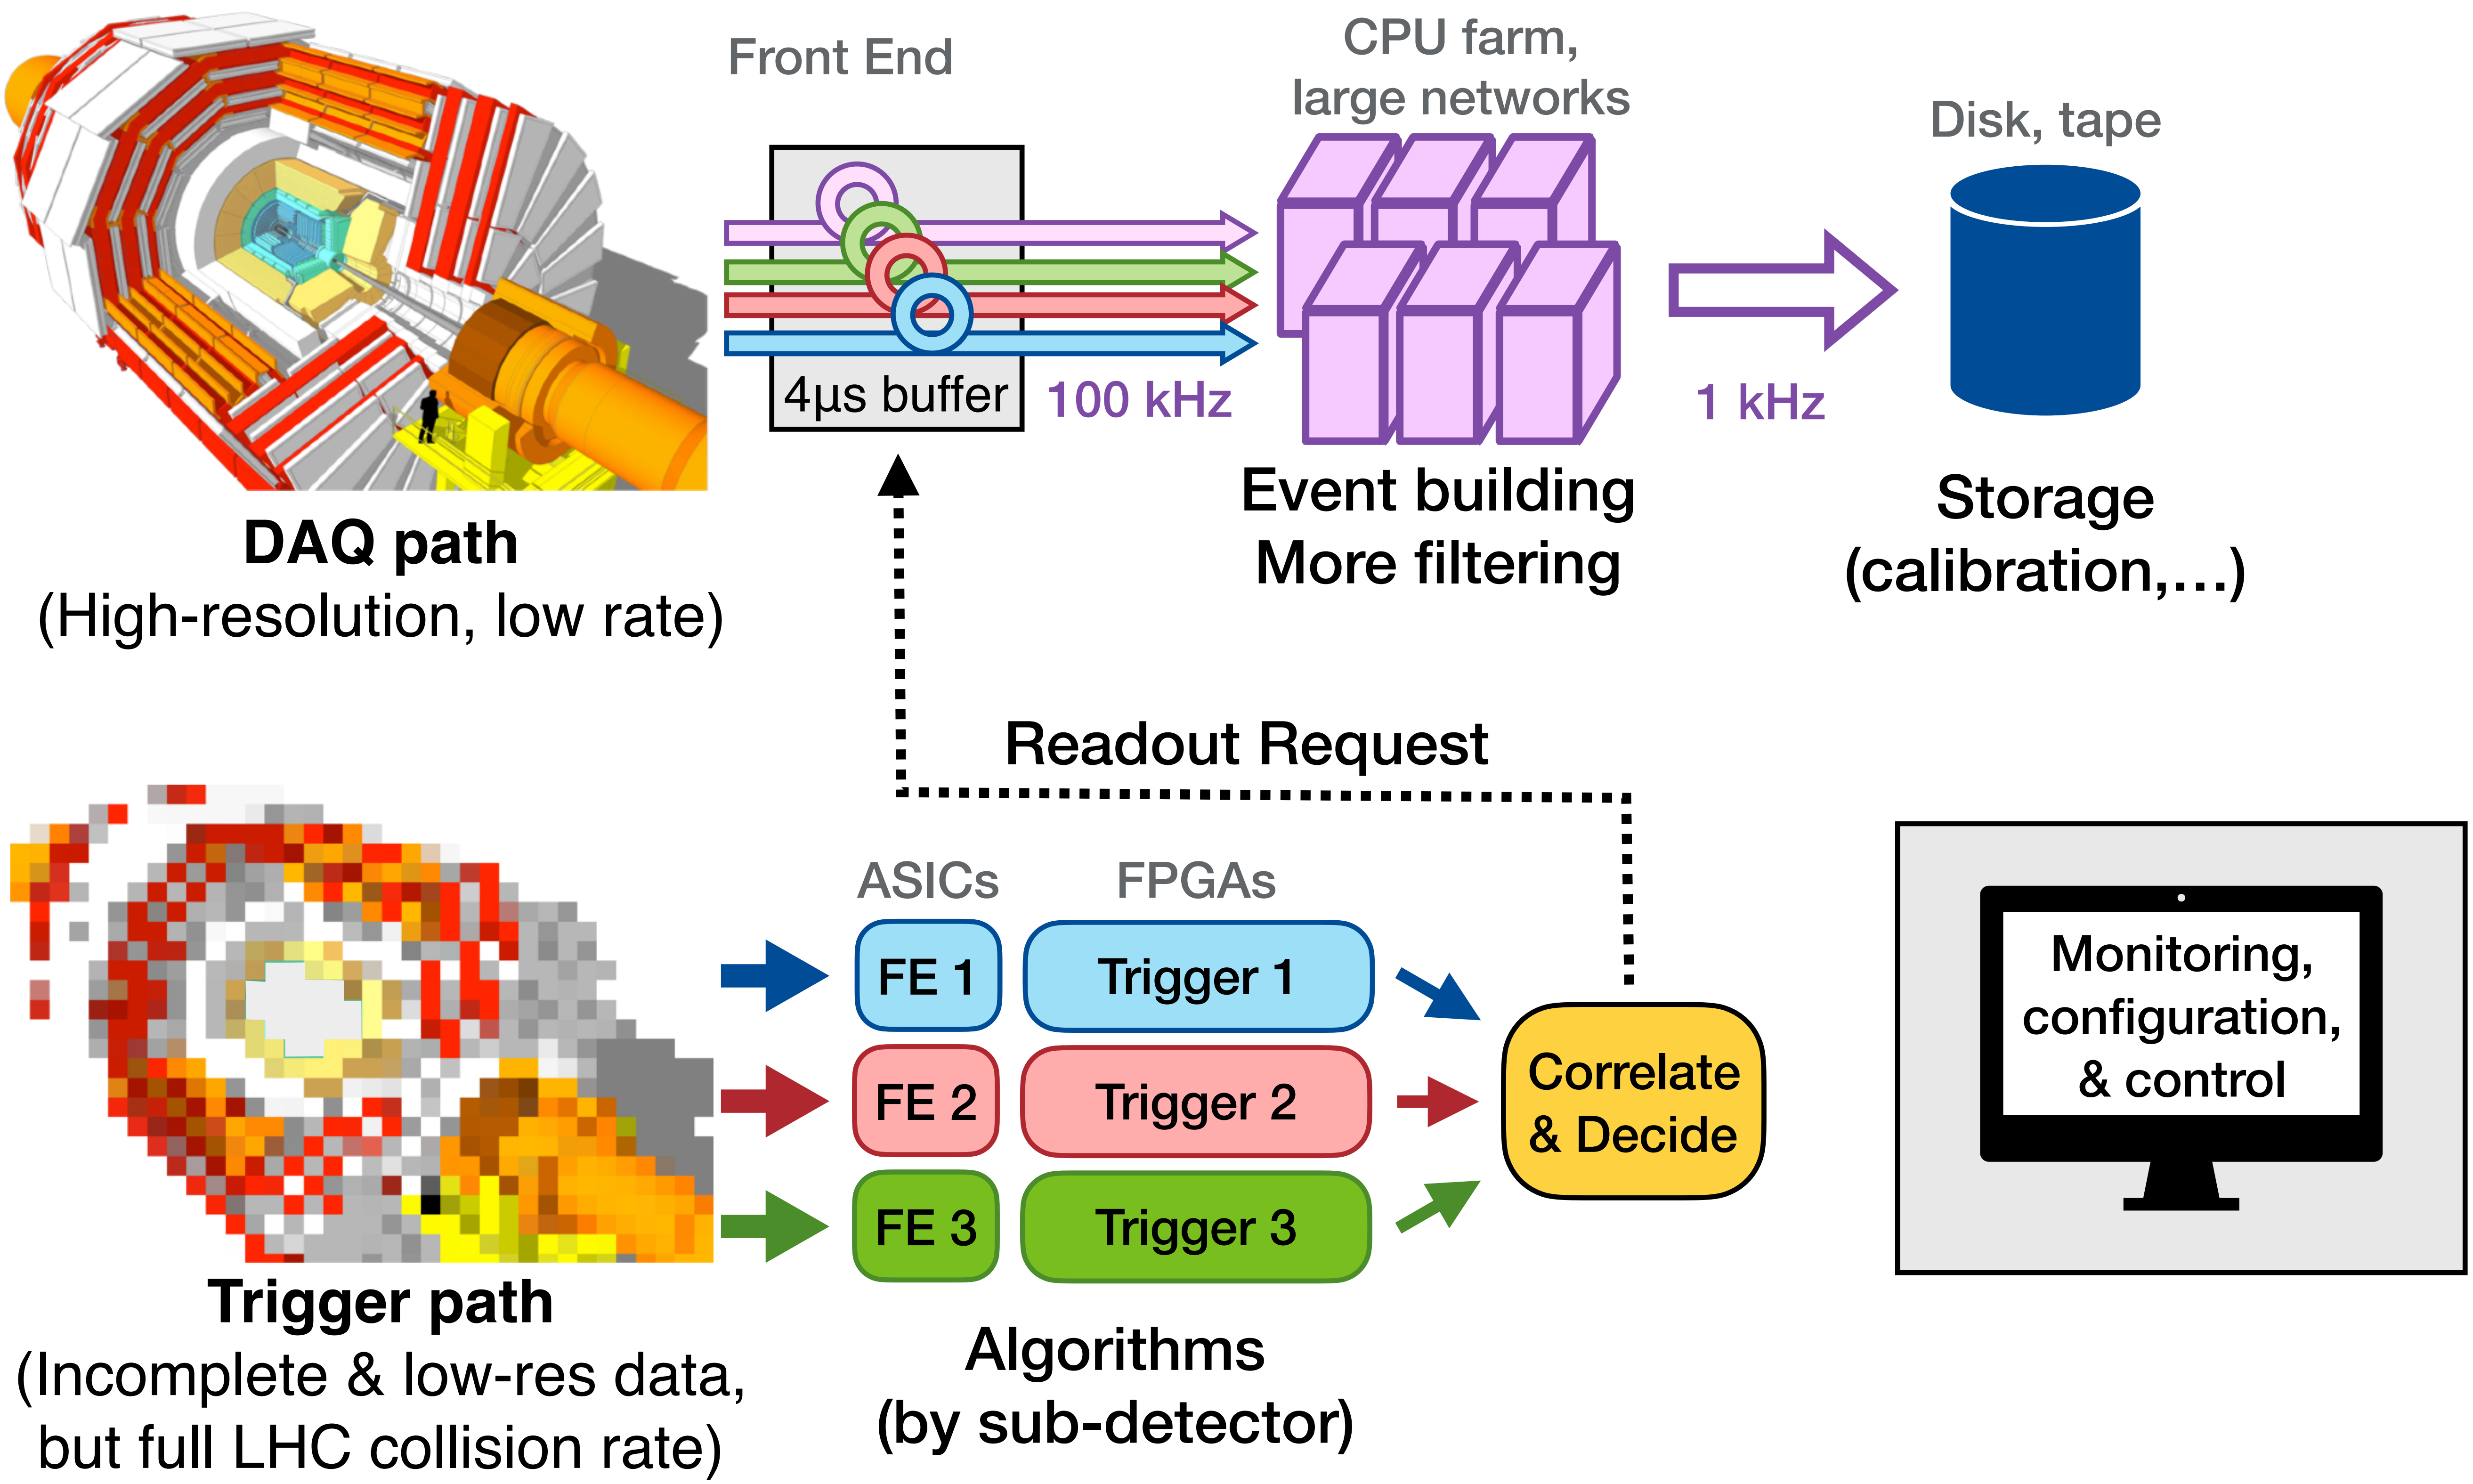

Some approximate numbers from ATLAS/CMS

Plan

- **This afternoon: Concepts**
  - Introduction and physics goals of the LHC Experiments
  - The role of the trigger and DAQ systems
  - **DAQ concepts: simple toy model → complex systems**
  - Trigger concepts: from hardware to architecting a system
- **Tomorrow morning: Details and Implementations**
  - The challenge of high luminosity, and the upgrades
  - ATLAS, CMS hardware triggers, and their evolution
  - Software triggers, and new paradigms for analysis
  - Trigger menus, and you
  - Conclusions, and looking farther ahead

The simplest DAQ

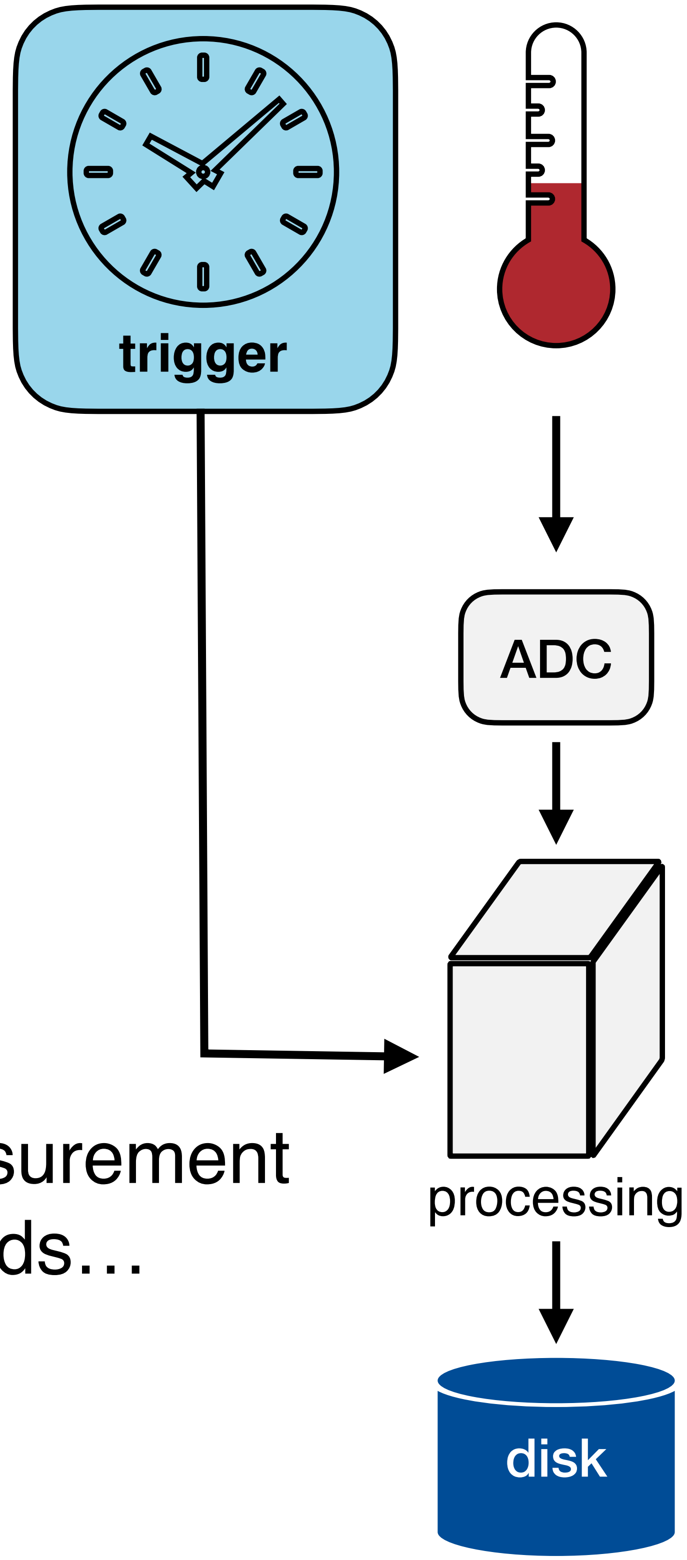

Measure a temperature (e.g.) at fixed frequency

Analogue to digital conversion, digitize signal ("front end electronics")

CPU for readout and processing

Write results to storage

Trigger a measurement every N seconds…

The simplest DAQ

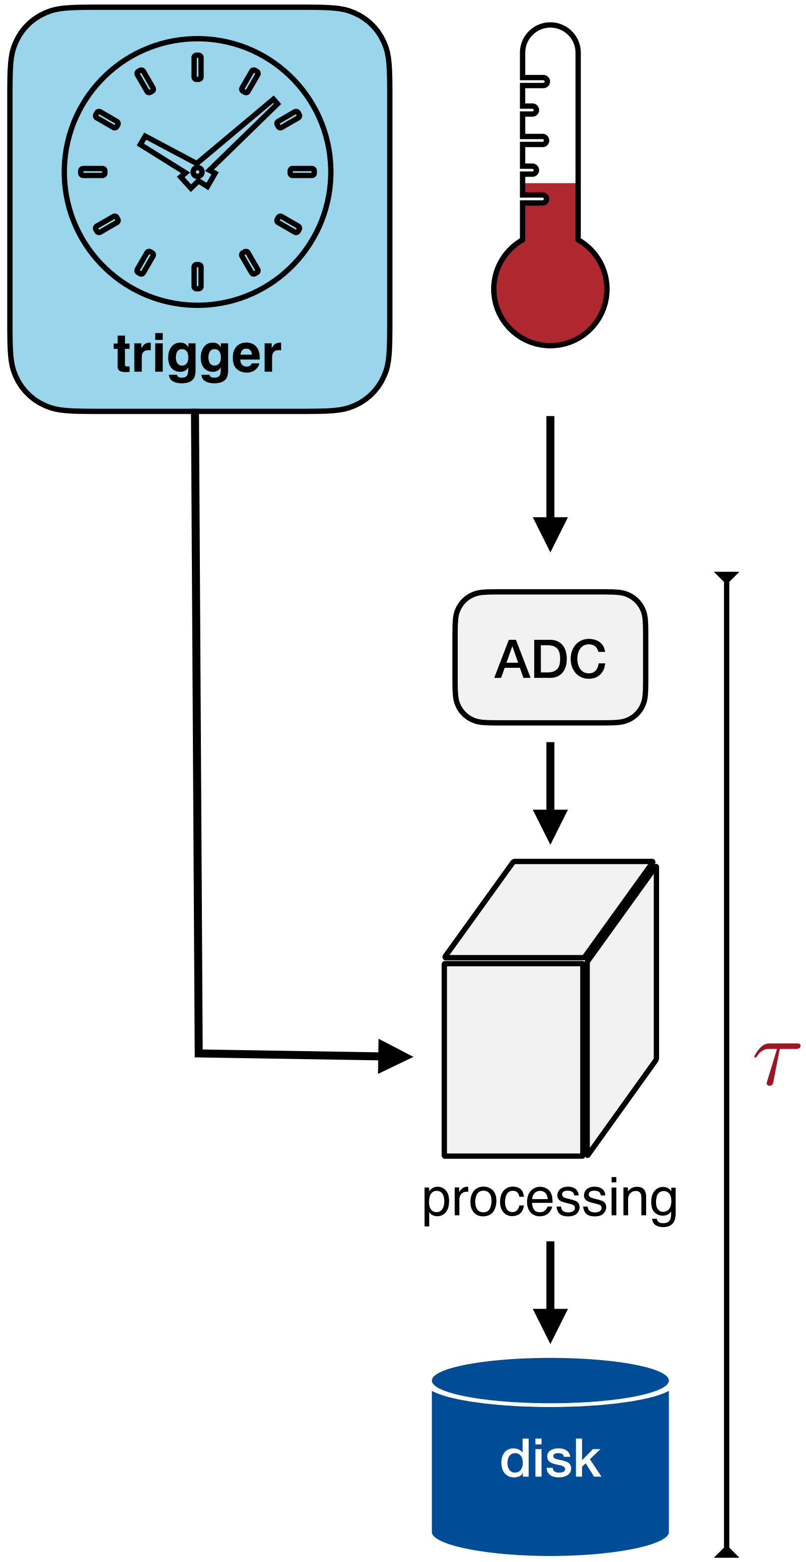

How often can we make a measurement?

Limited by the total time to process and record the "event":

$\tau = \tau(\text{ADC}) + \tau(\text{process}) + \tau(\text{write})$

Triggering a measurement each τ, leads to a maximum rate of $R=1/\tau$

E.g. $\tau=1\text{ms} \rightarrow R = 1\text{kHz}$

Simple DAQ, with a ‘real’ trigger

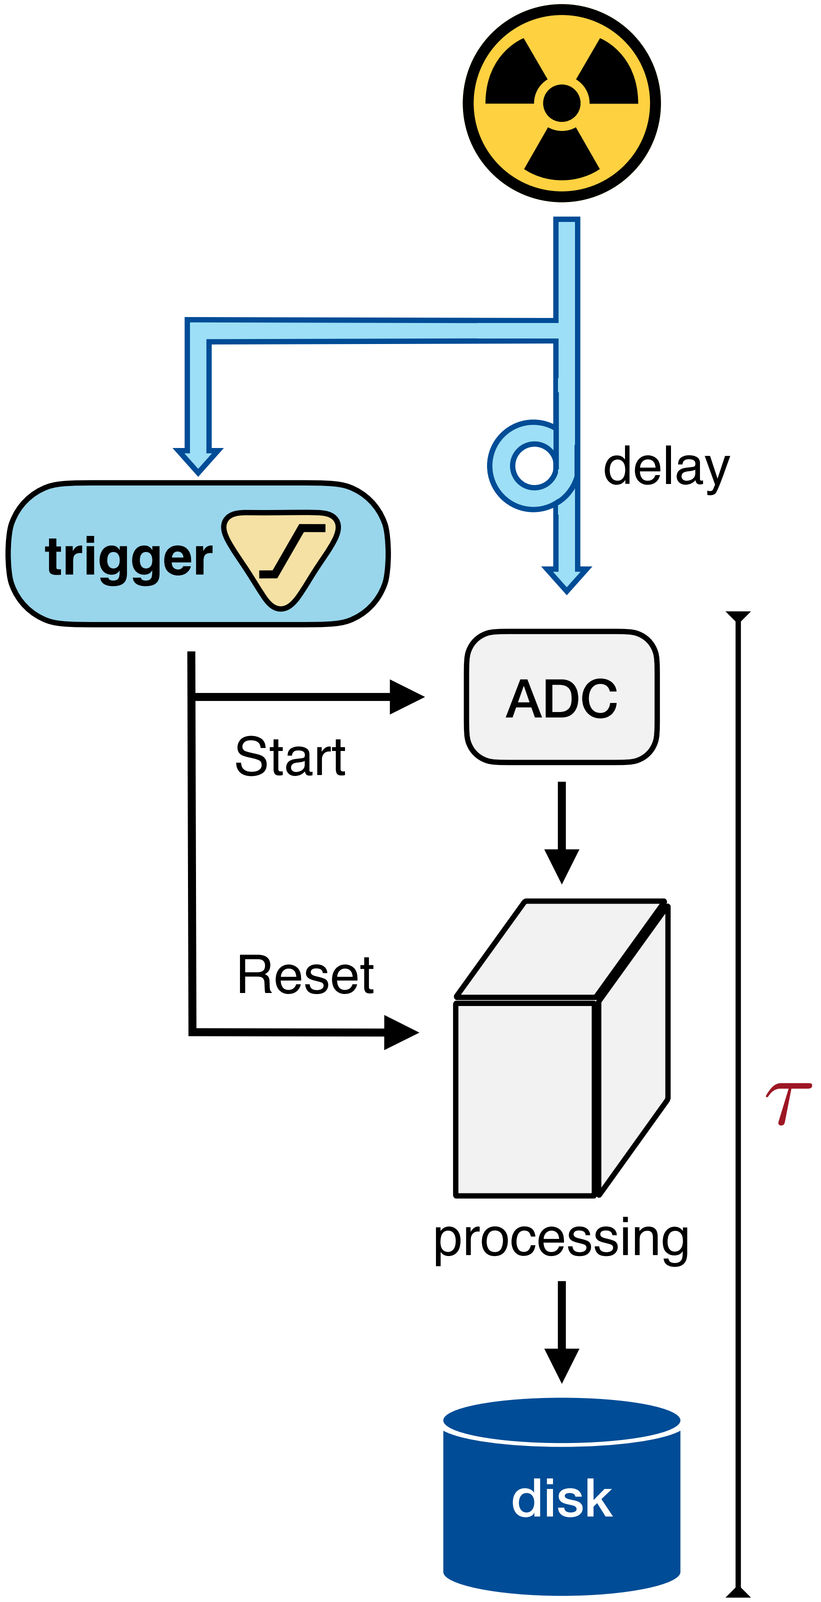

Consider instead events that occur asynchronously, and unpredictably (e.g. radioactive decay)

Delay signal, to account for trigger latency

Discriminator triggers data collection on a pulse’s rising edge

→ Start ADC readout

→ Ensure CPU is ready to receive

How often can we make a measurement?

Simple DAQ, with a ‘real’ trigger

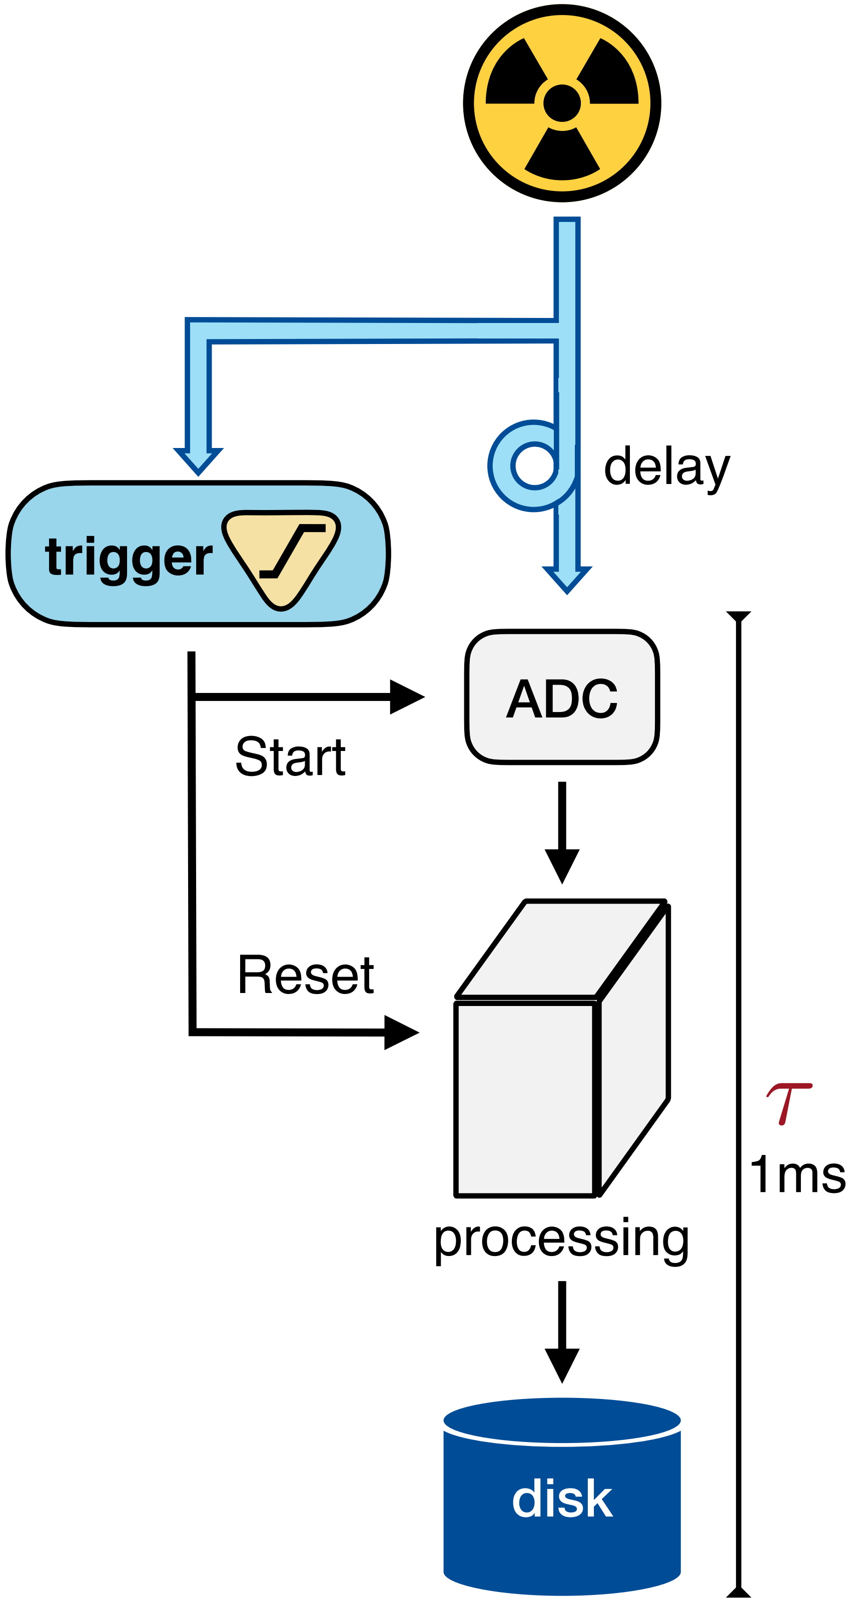

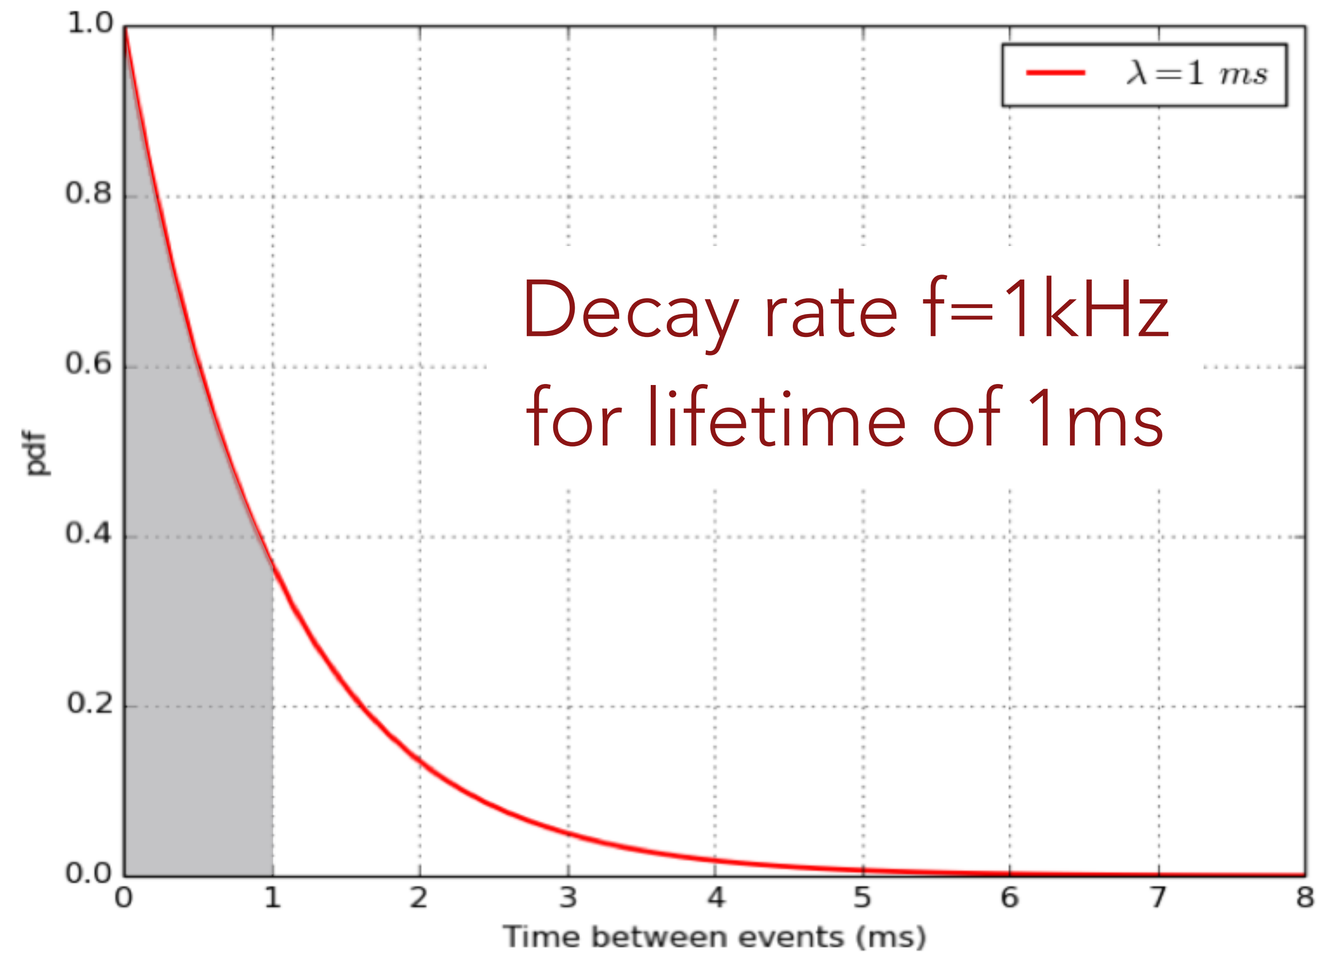

What if a new signal arrives with $\Delta t < \tau$?

Simple DAQ, with a ‘real’ trigger

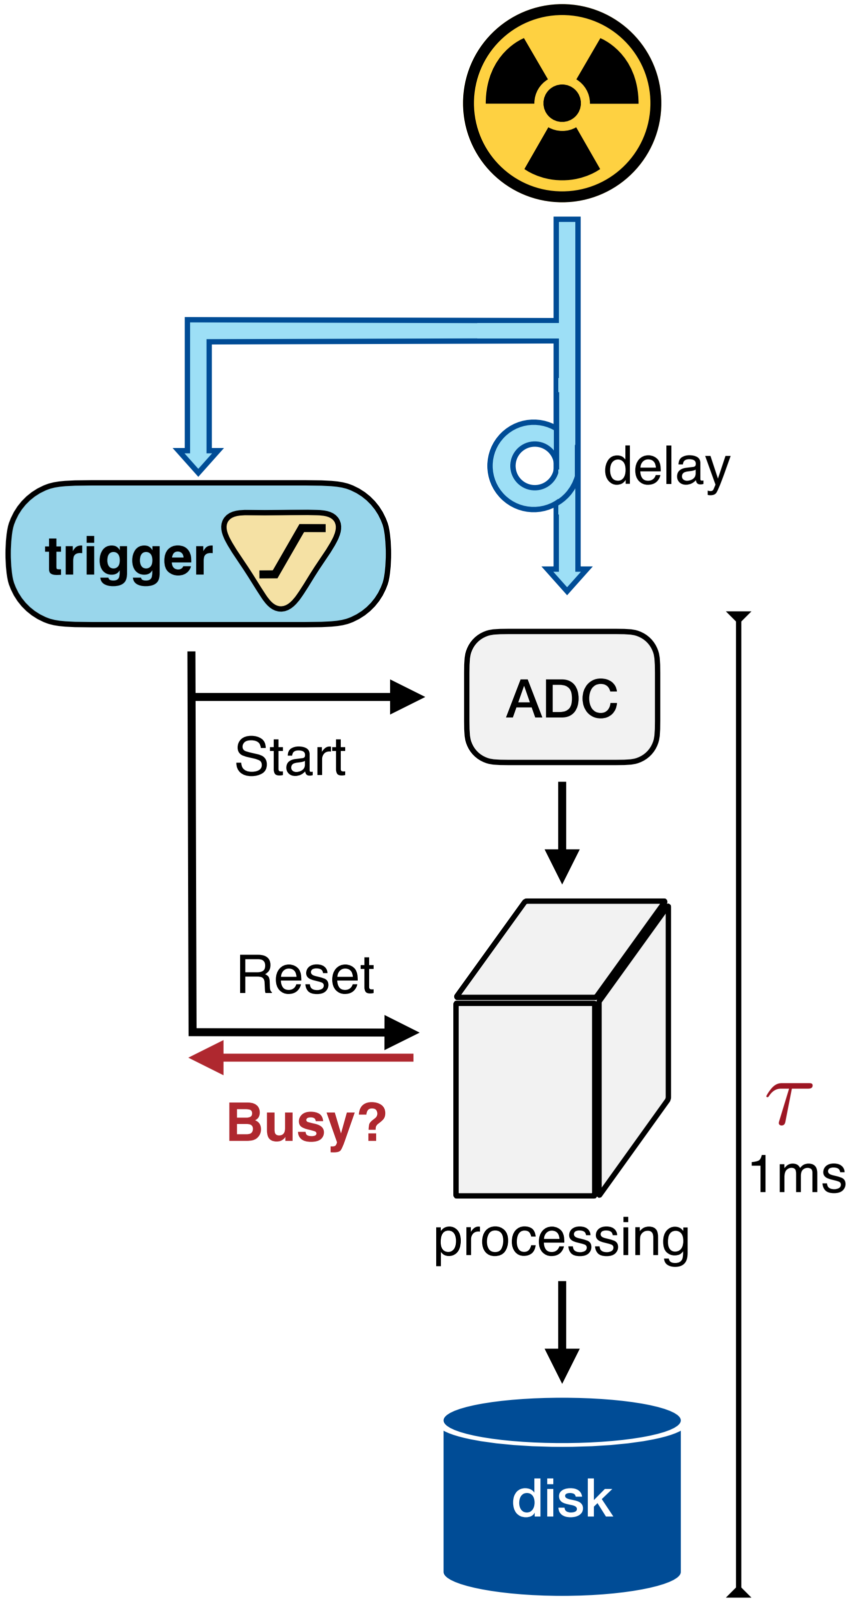

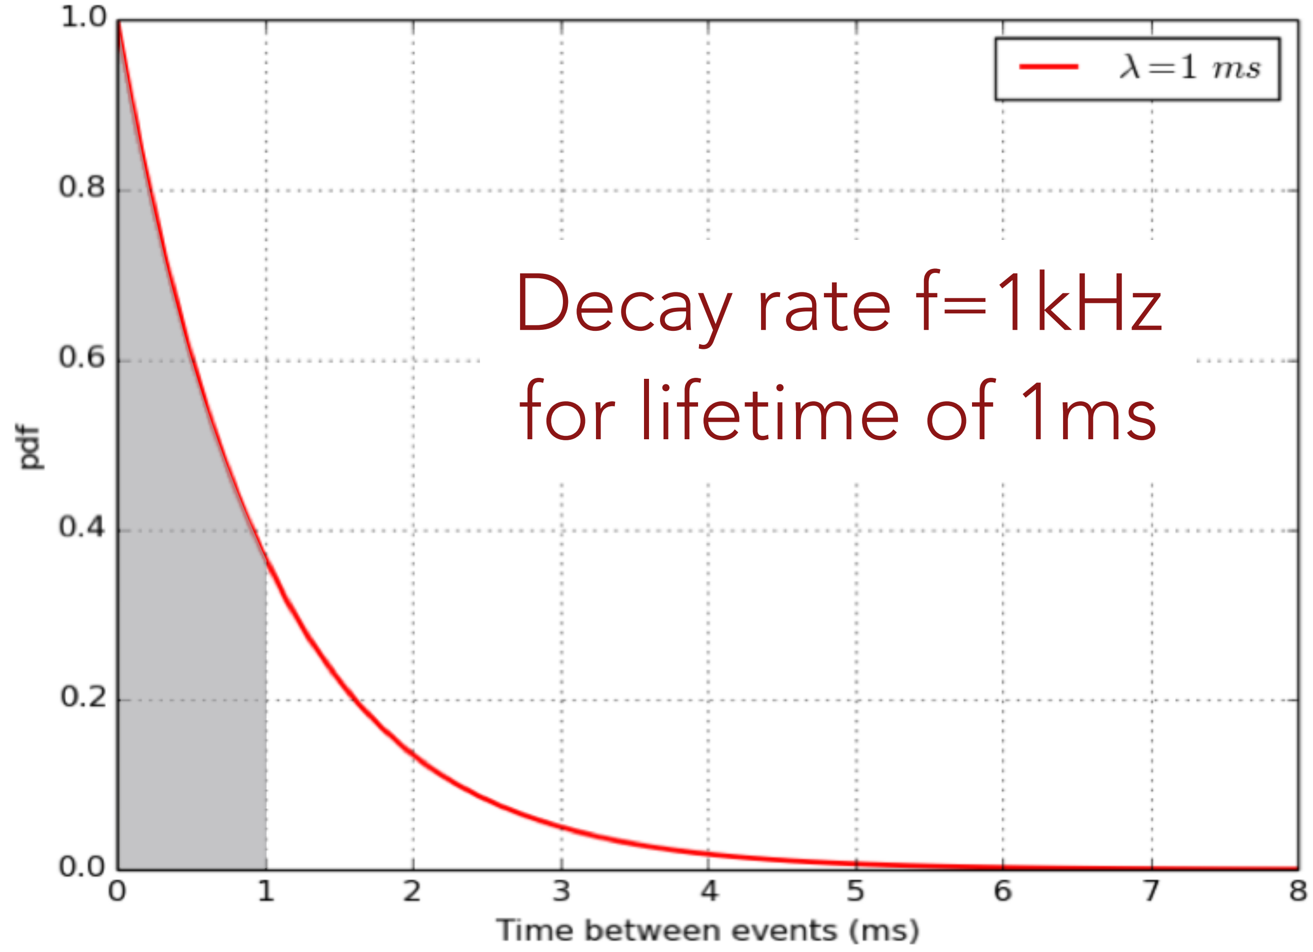

What if a new signal arrives with Δt < τ ?

“Dead time” τ: time to processes an event, before system can handle new triggers

Simple DAQ, with a 'real' trigger

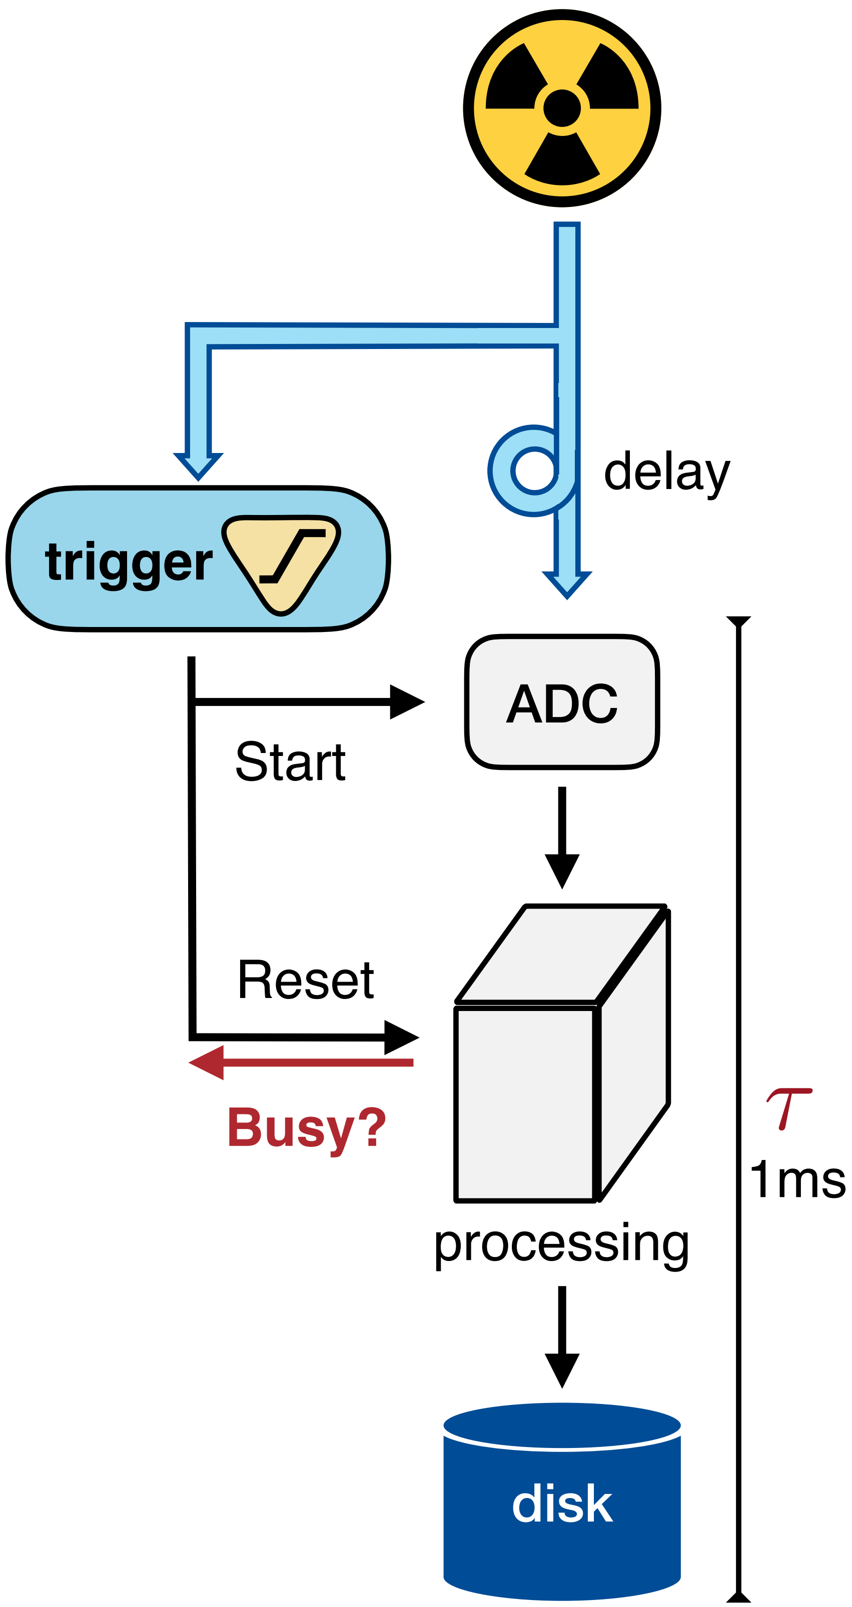

What is the *actual DAQ rate* **v** = f * P(free)?

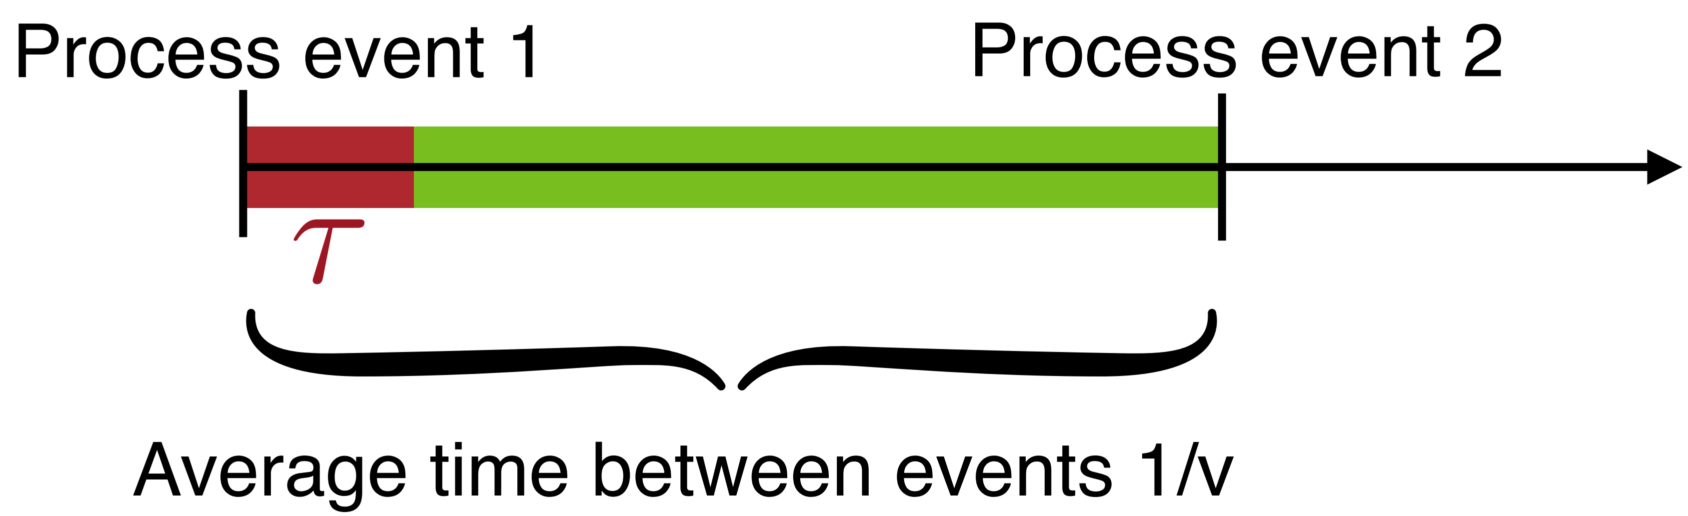

Simple DAQ, with a 'real' trigger

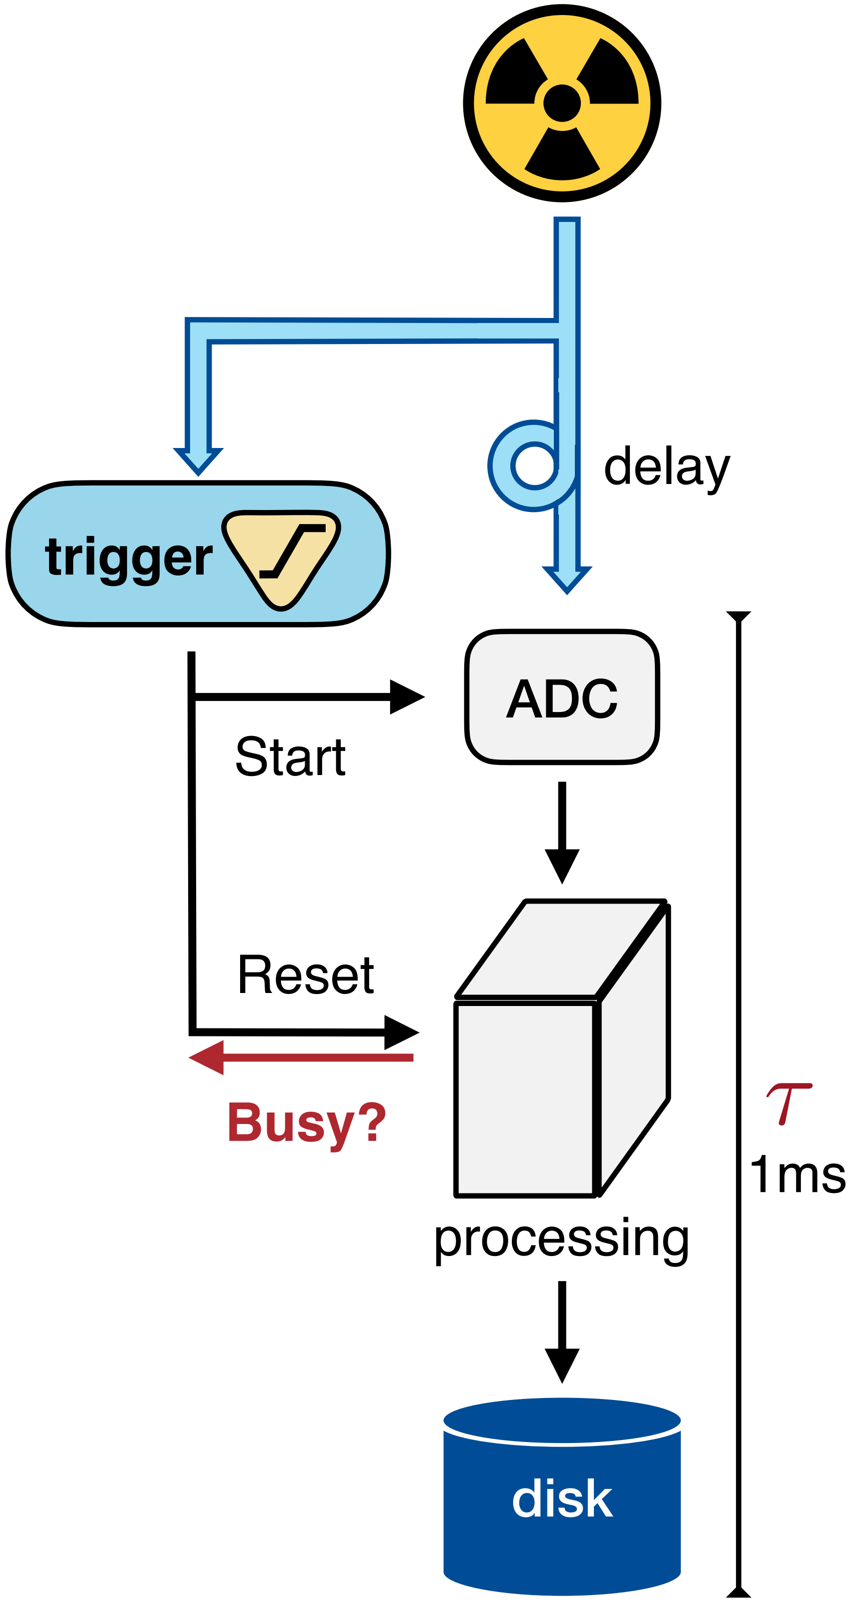

What is the *actual DAQ rate* **v** = f * P(free)?

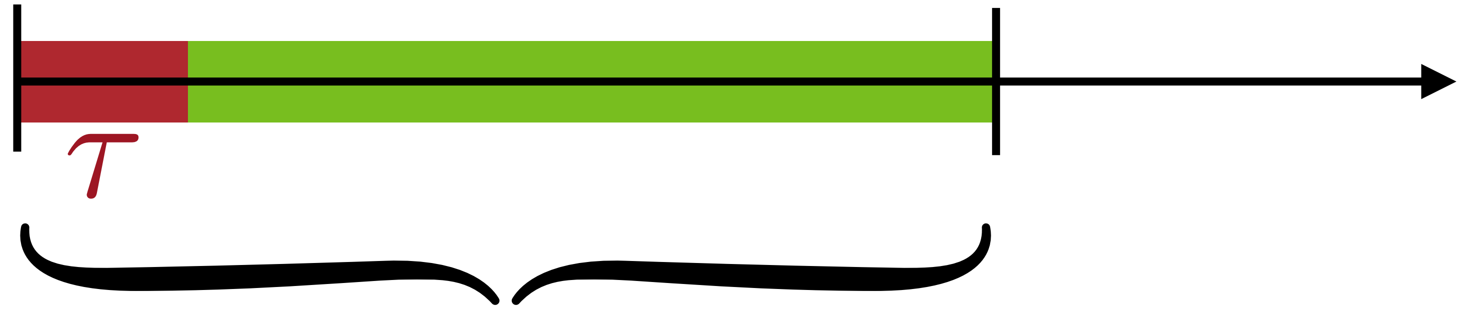

Average time between events 1/v

P(busy) = τ/(1/v) = τv, and P(free) = 1-τv

Simple DAQ, with a ‘real’ trigger

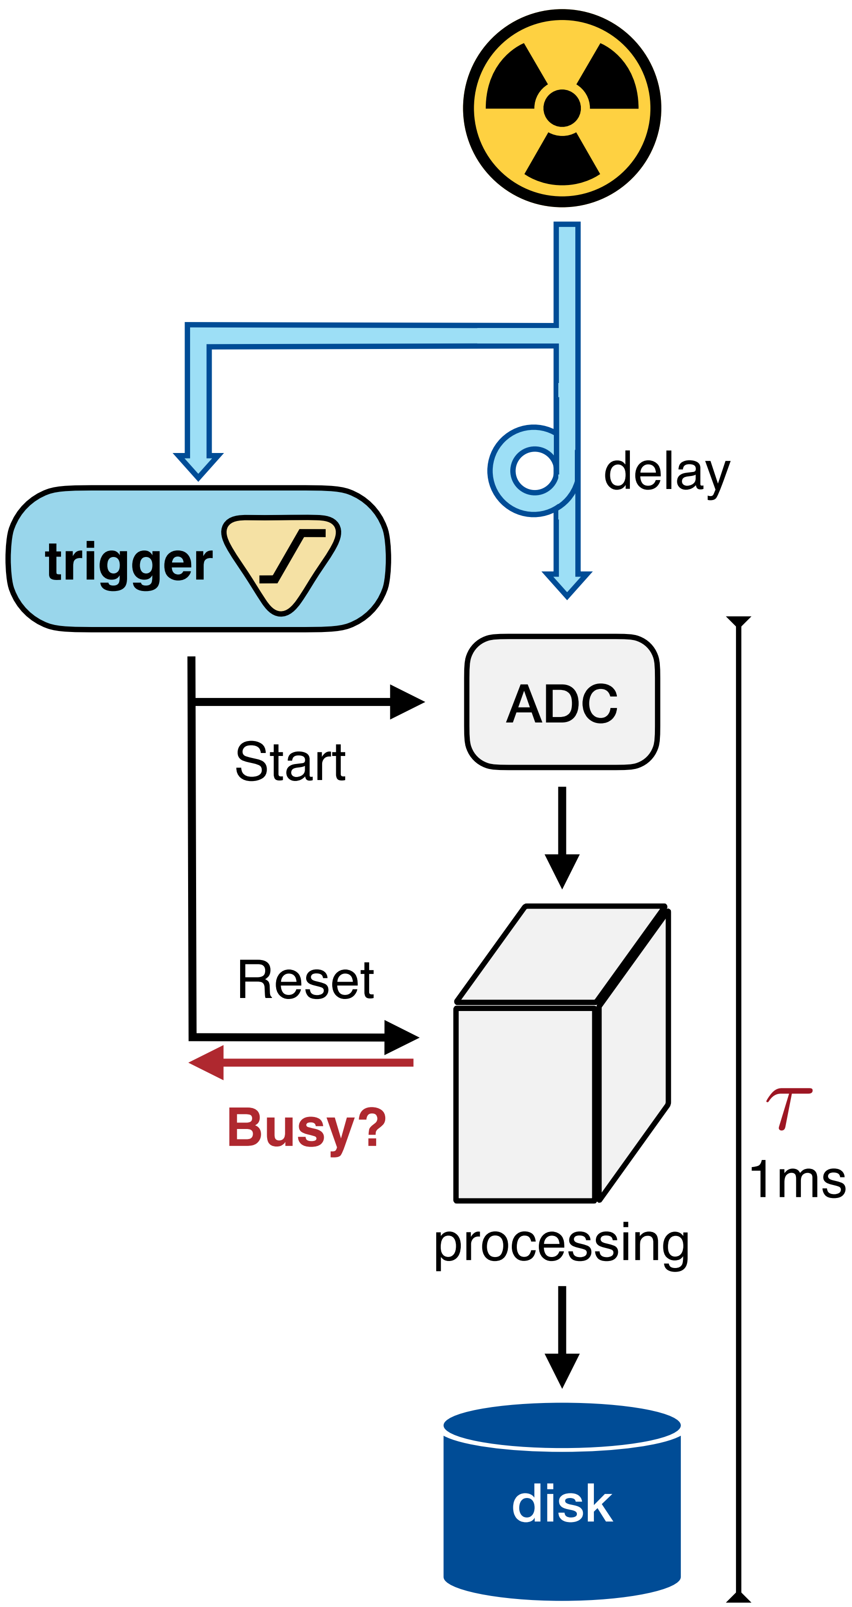

What is the *actual DAQ rate* **v** = f * P(free)?

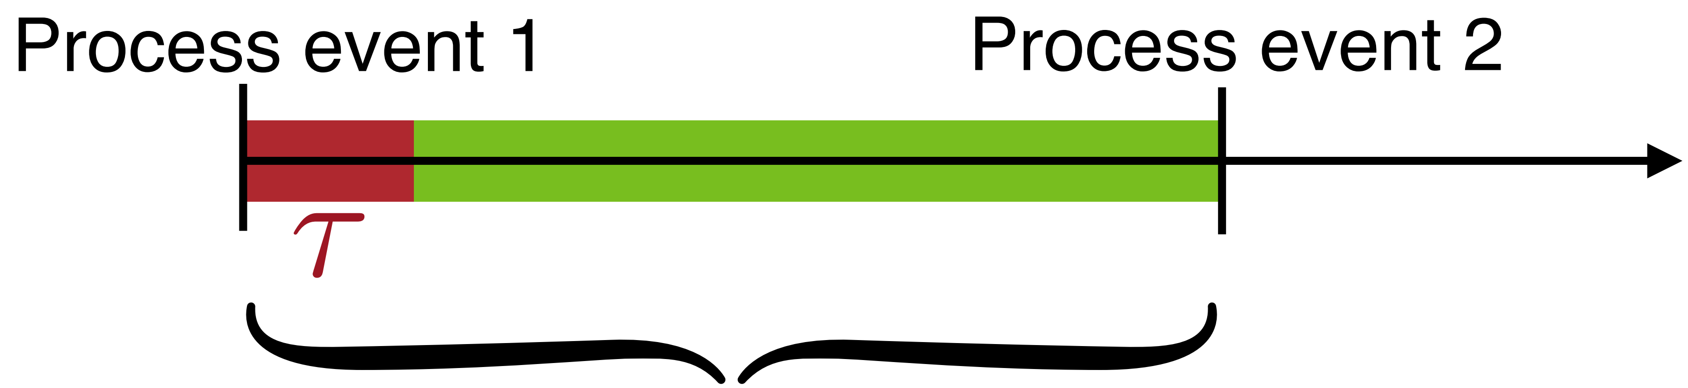

Average time between events 1/v

P(busy) = τ/(1/v) = τv, and P(free) = 1-τv

So the rate is given by **v** = f(1-τv)→ v=f/(1+fτ)

Simple DAQ, with a ‘real’ trigger

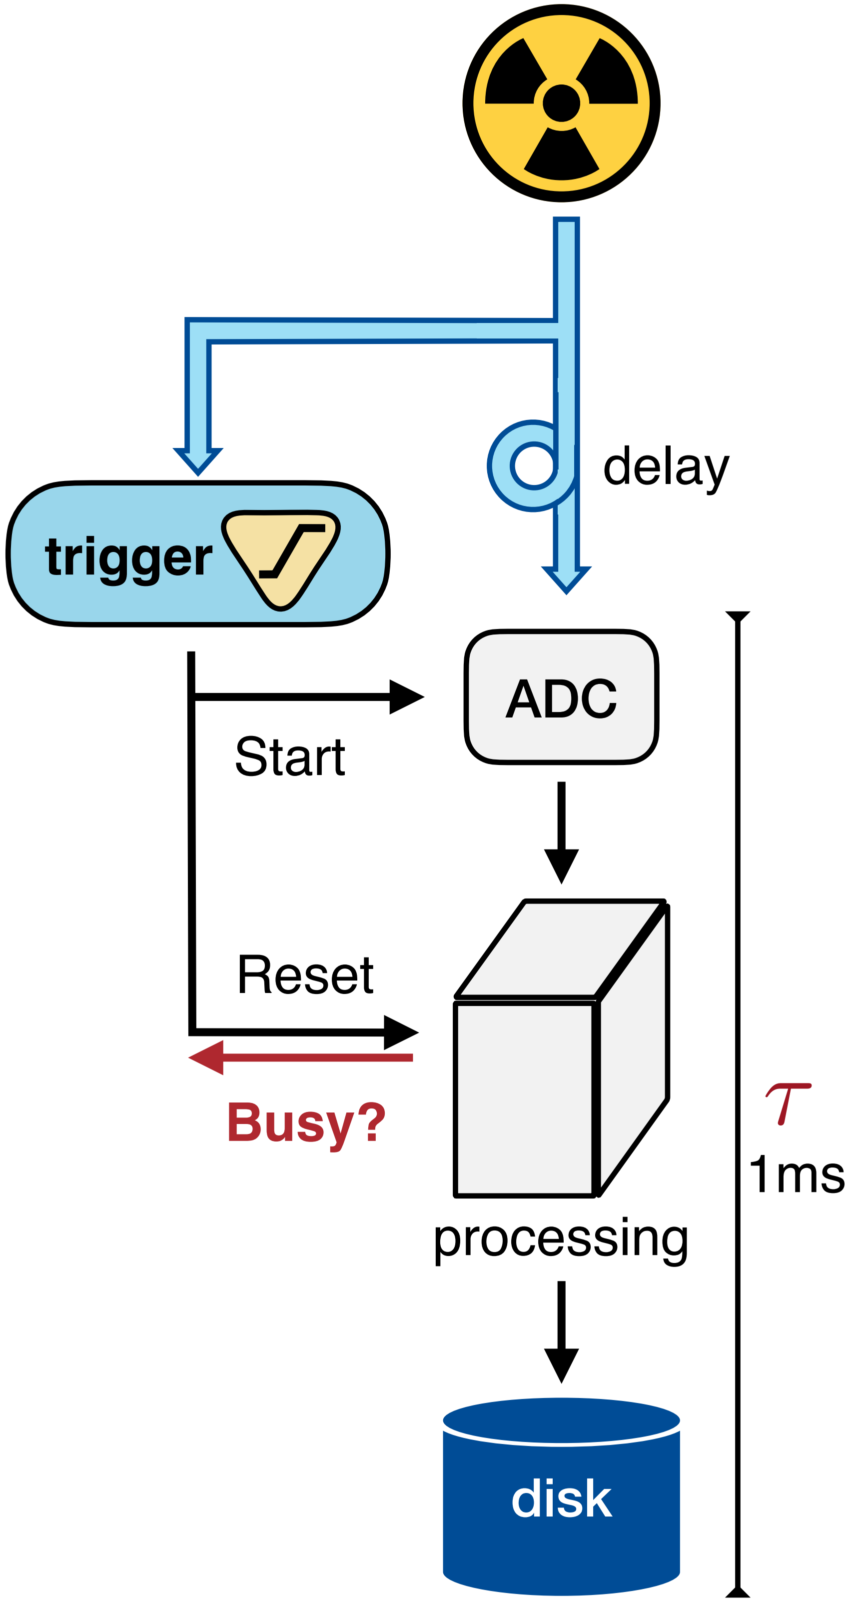

What is the *actual DAQ rate* **v** = f * P(free)?

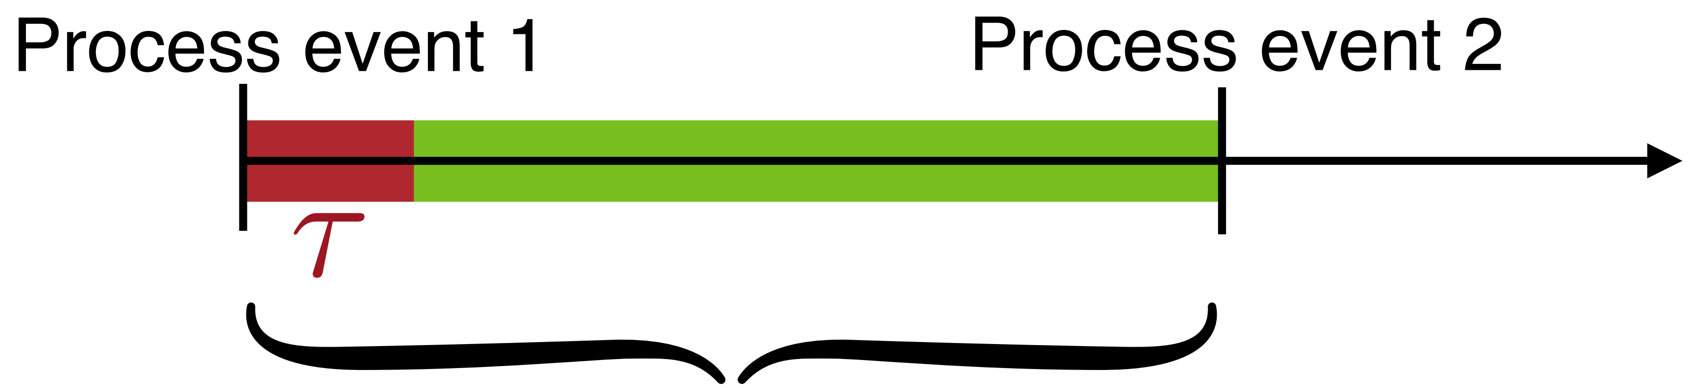

Average time between events 1/v

P(busy) = τ/(1/v) = τv, and P(free) = 1-τv

So the rate is given by **v** = f(1-τv)→ v=f/(1+fτ)

Our example: f = 1kHz, τ = 1 ms → v = 0.5 kHz

This means a data-taking efficiency ε of 50%!!

Simple DAQ, with a ‘real’ trigger

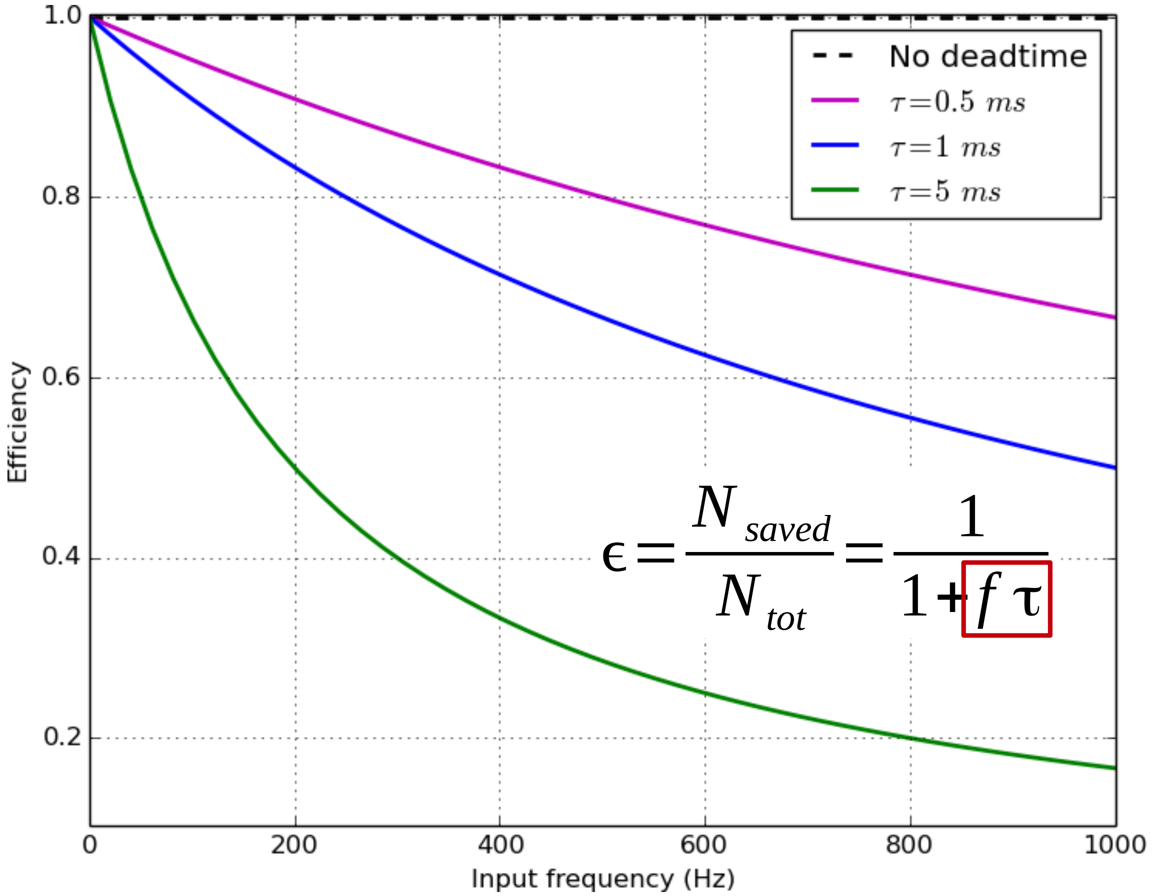

Maximize ε by minimizing the dead time.

To achieve $\varepsilon > 99\%$ requires $\tau < 0.01$ms

This is *100x more stringent* than what’s required for a fixed-frequency experiment

→ The burden of random inputs!

Simple DAQ, with a 'real' trigger

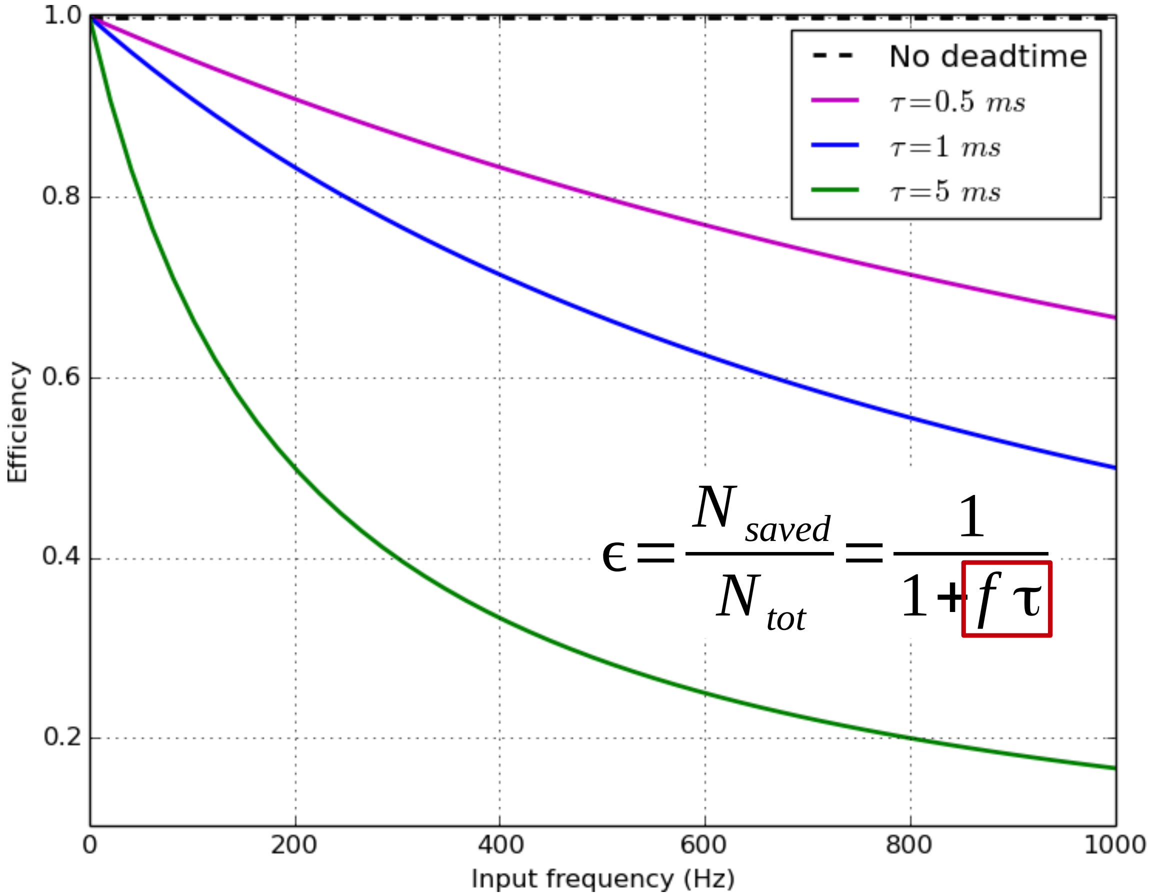

$$\epsilon = \frac{N_{saved}}{N_{tot}} = \frac{1}{1 + f\,\tau}$$

Maximize ε by minimizing the dead time.

To achieve ε > 99% requires τ < 0.01ms

This is *100x more stringent* than what's required for a fixed-frequency experiment

→ The burden of random inputs!

How to cope? → **De-randomization**

Queuing the data averages out the time to access

e.g. a First-In First-Out (FIFO) buffer

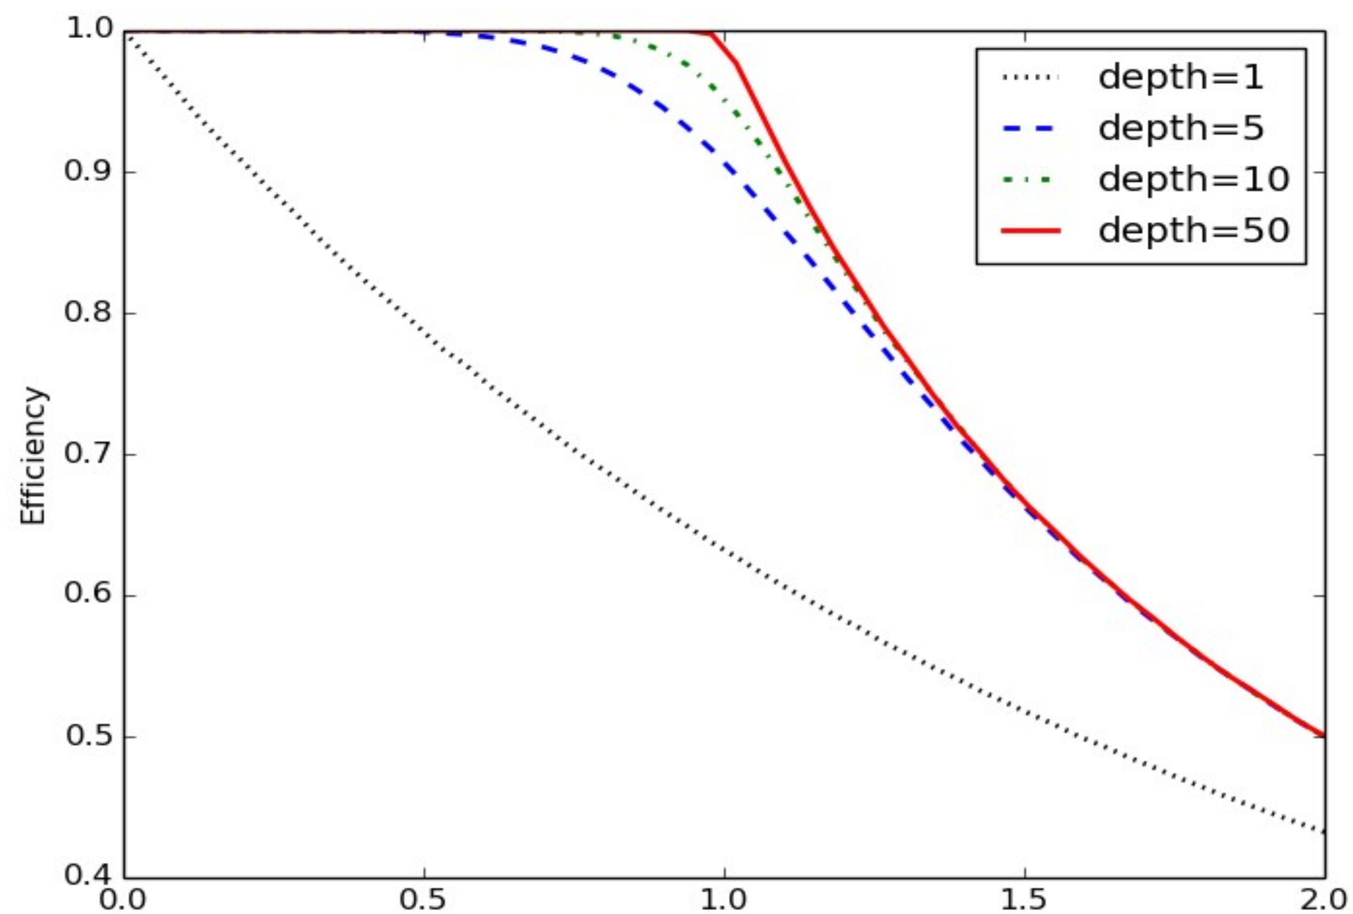

Fill rate / empty rate

Simple DAQ, with a ‘real’ trigger

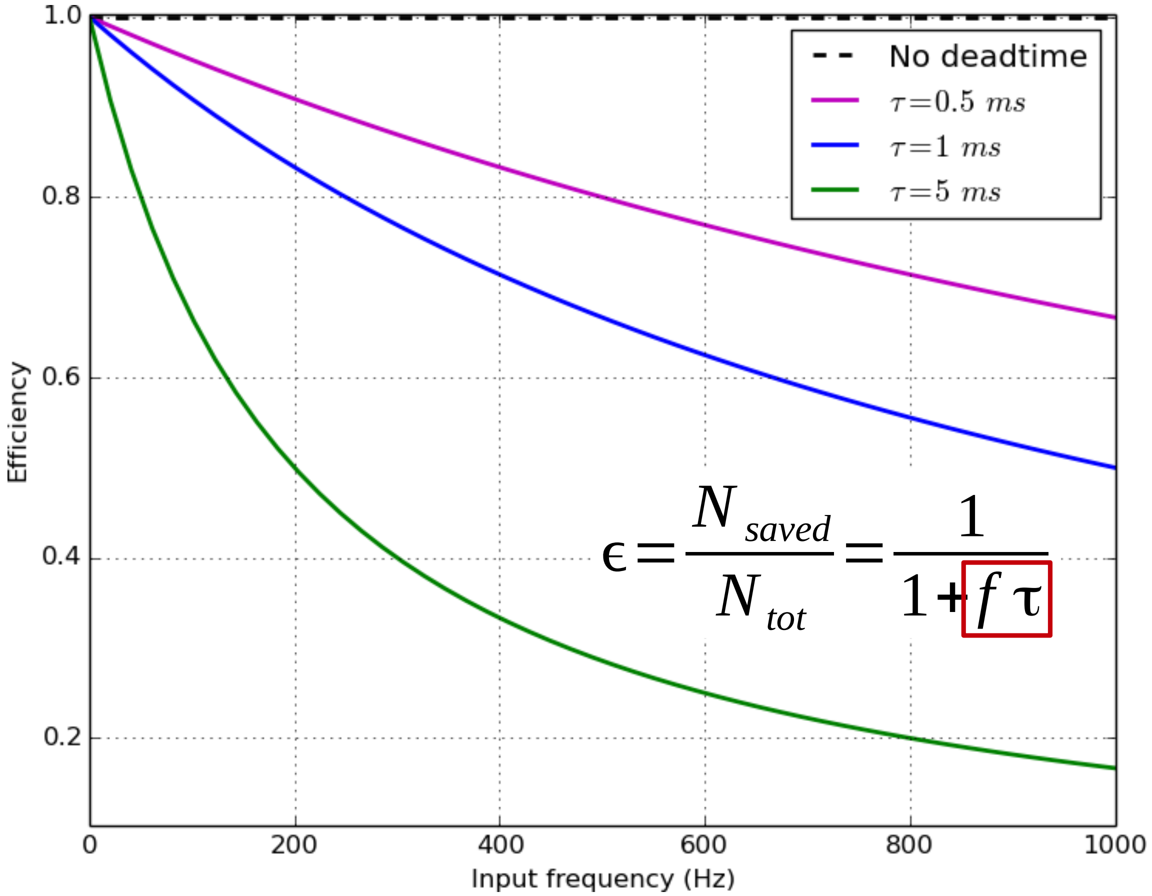

Maximize ε by minimizing the dead time.

To achieve ε > 99% requires τ < 0.01ms

This is *100x more stringent* than what’s required for a fixed-frequency experiment

→ The burden of random inputs!

How to cope? → **De-randomization**

Queuing the data averages out the time to access

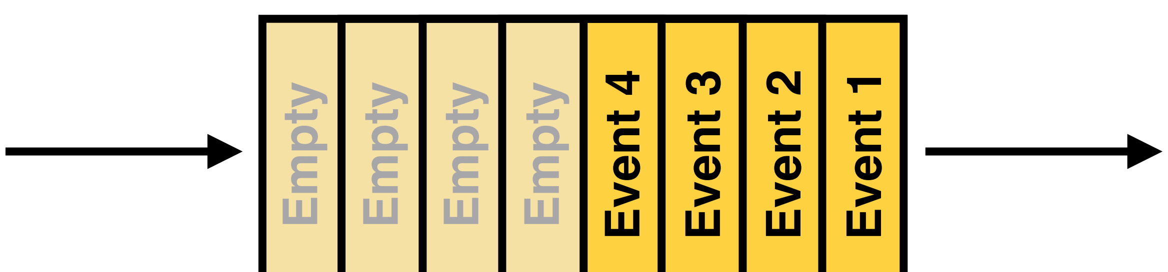

e.g. a First-In First-Out (FIFO) buffer

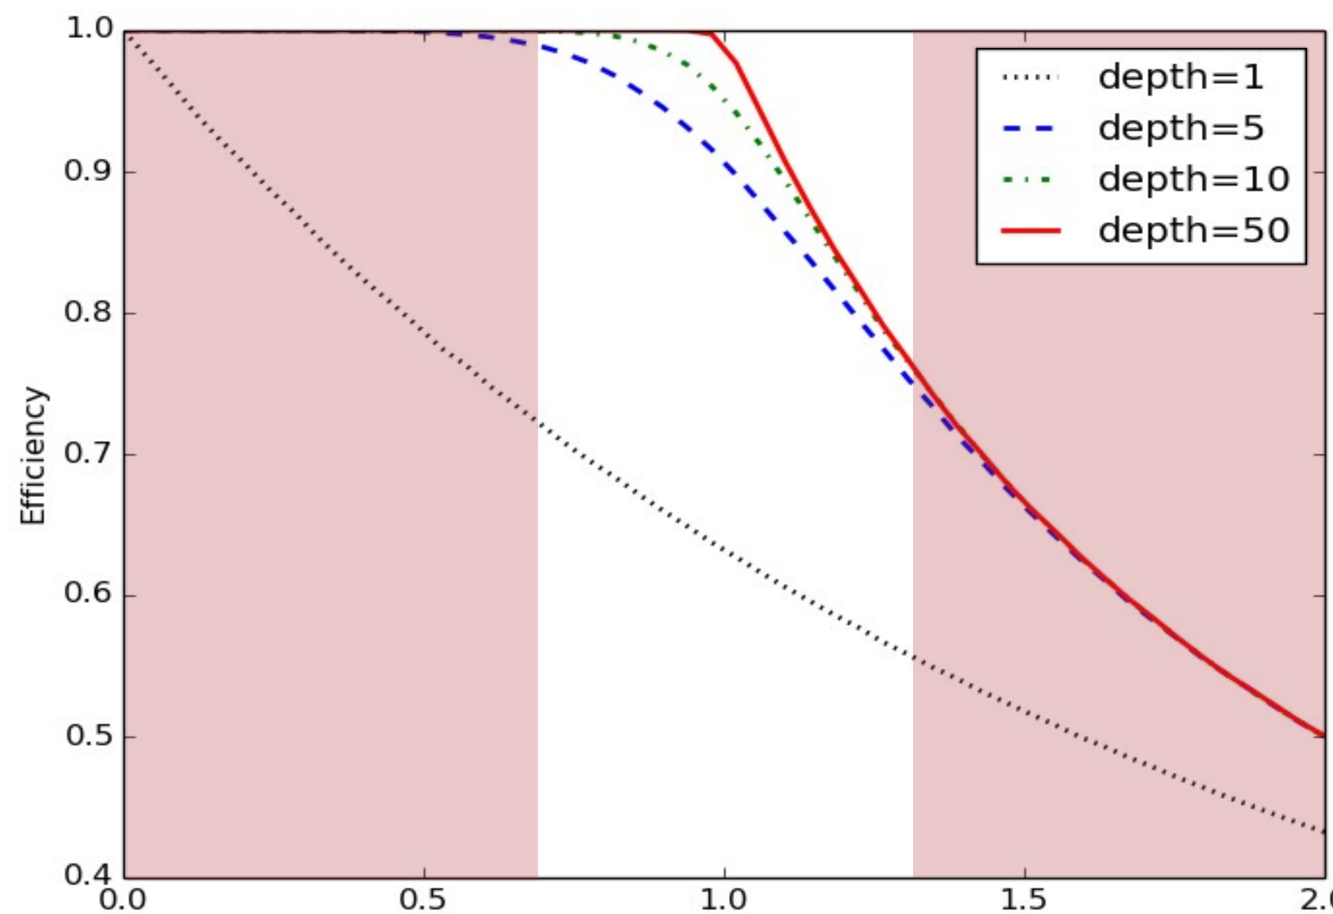

Fill rate / empty rate

Simple DAQ, made efficient

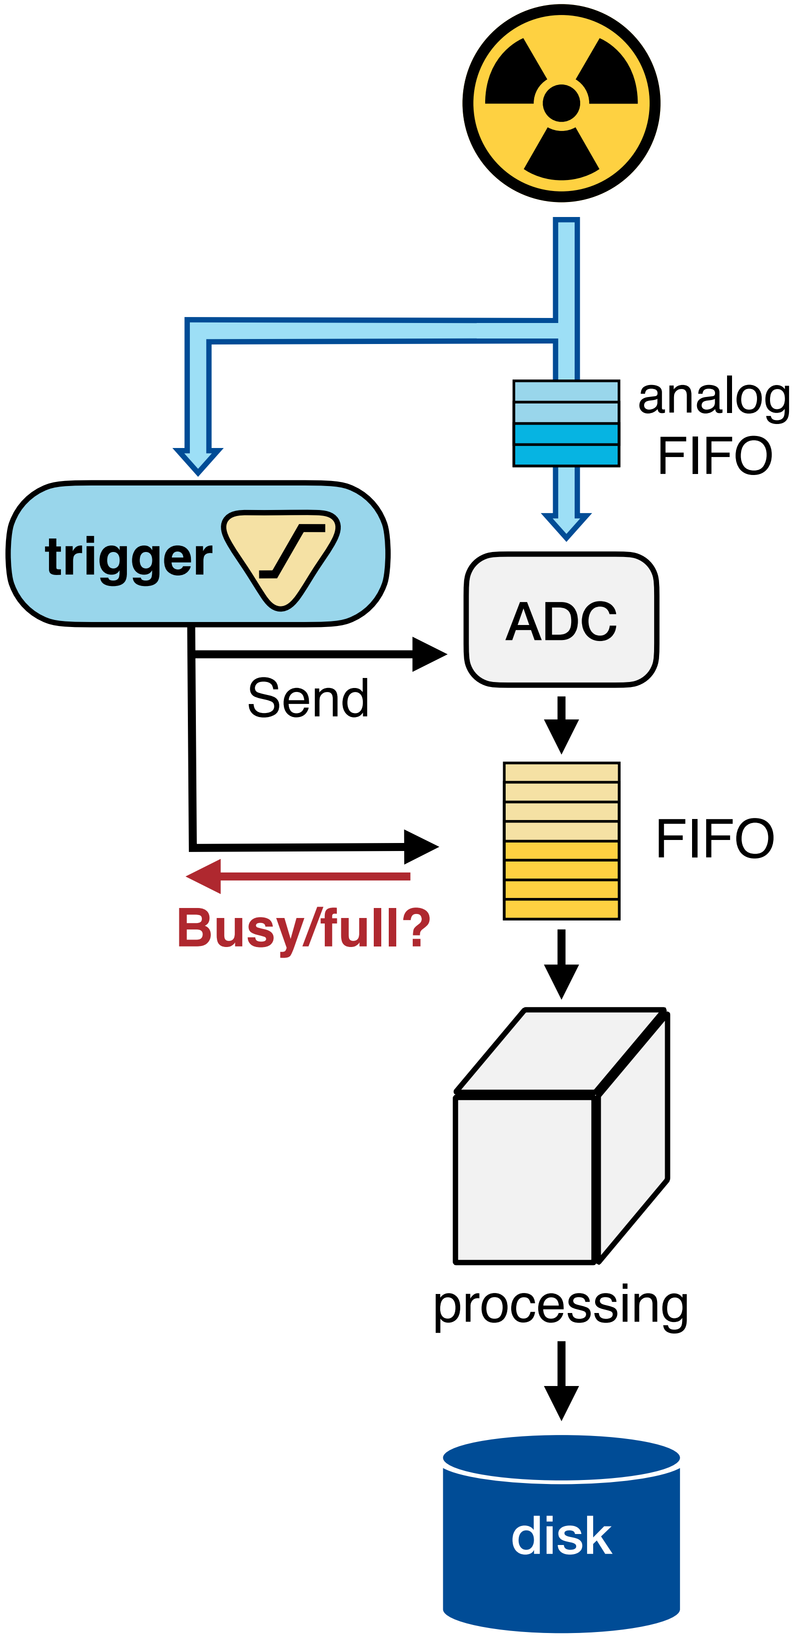

With sufficient size, buffers rarely fill

Processors can now reliably read at fixed frequency

Simple DAQ, made efficient

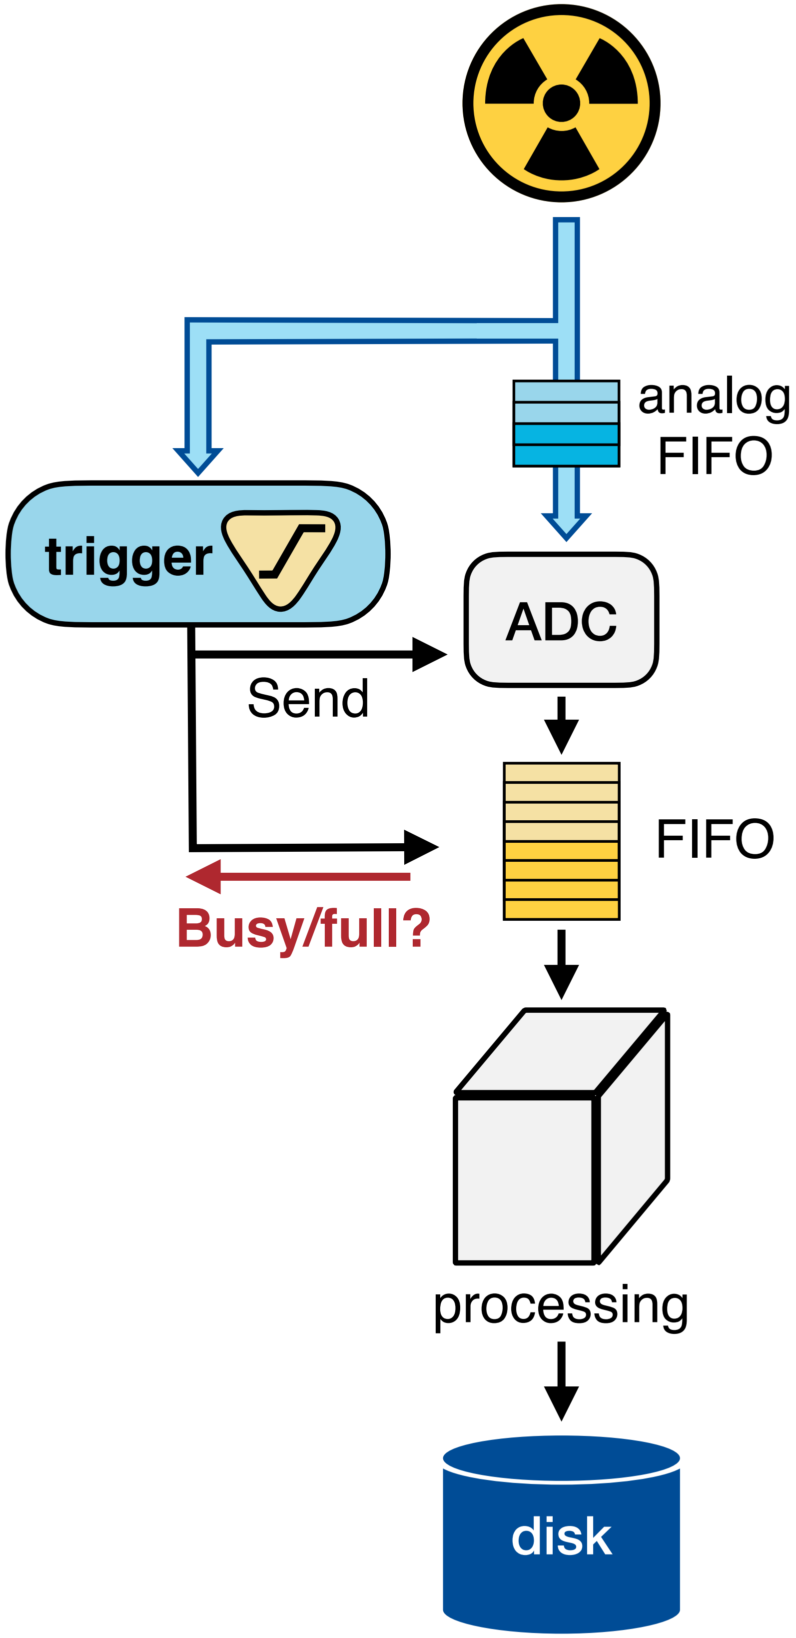

FIFOs may also be added in analog logic

With sufficient size, buffers rarely fill

Processors can now reliably read at fixed frequency

A many-channel experiment

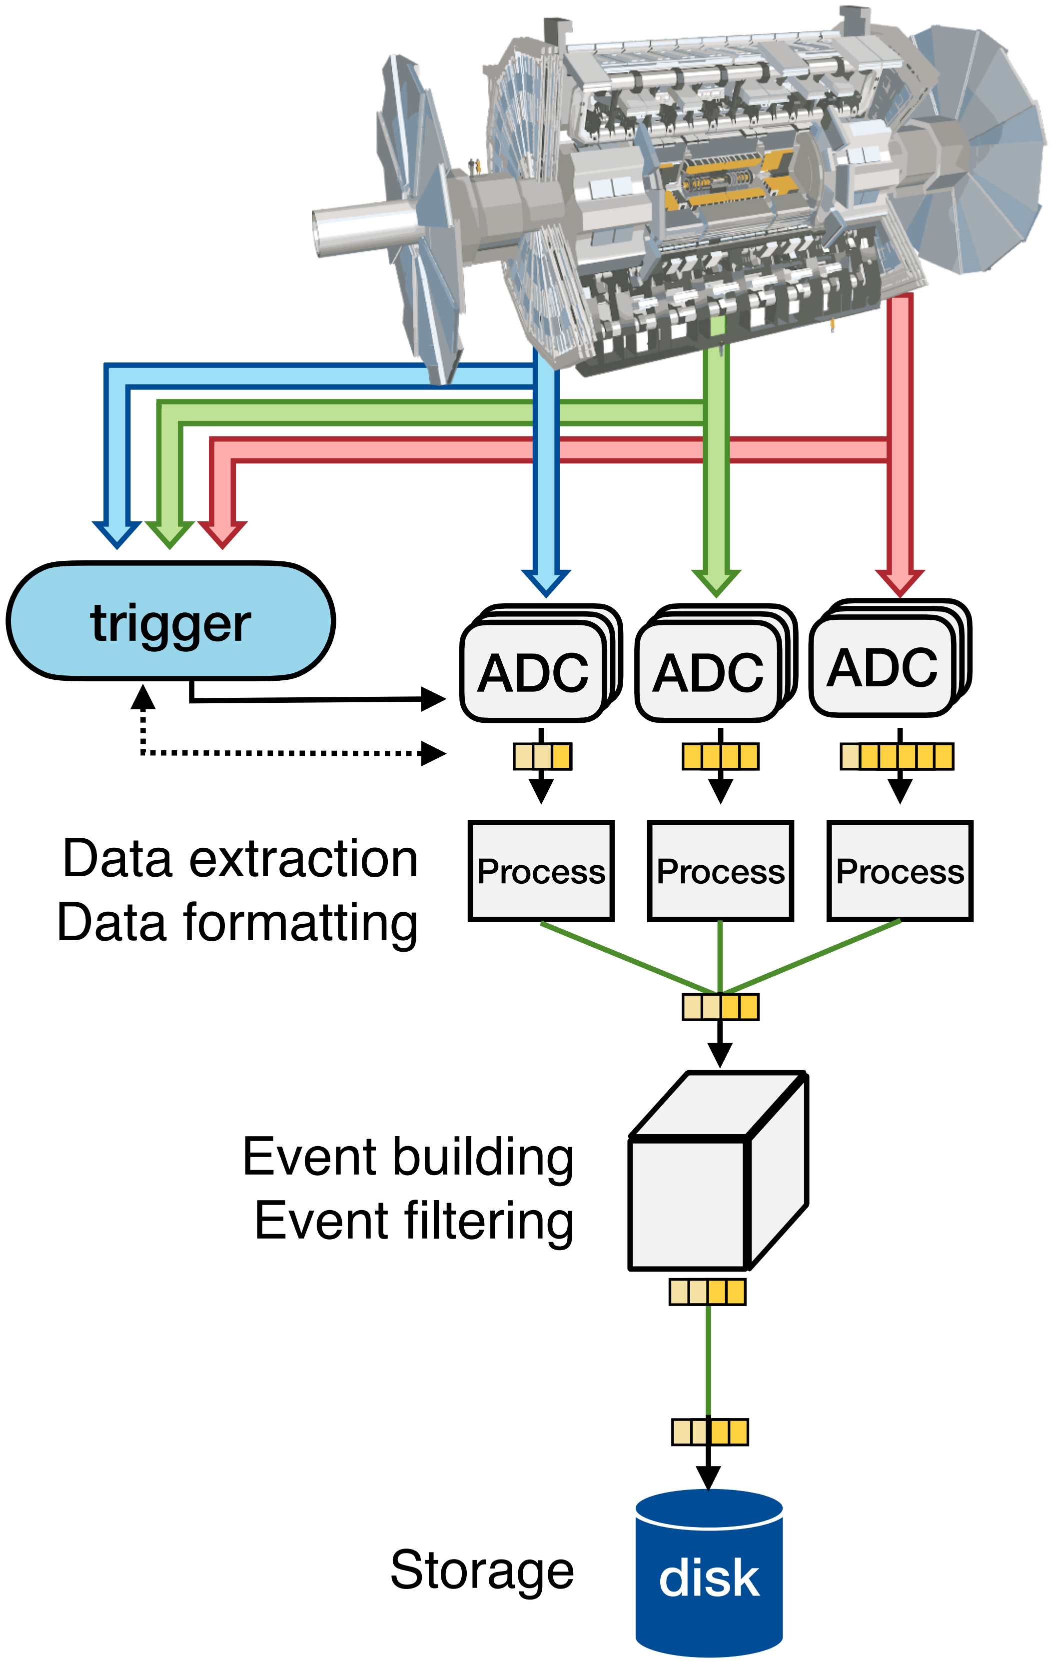

Many front-ends with buffers of various depths

FIFOs throughout the DAQ system

Back-pressure may propagate up the readout chain

(Where did it originate? →Monitoring!)

Ex. dead time at ATLAS

Simple veto:

Skip N bunches after accept

Complex veto: "Leaky bucket"

Bucket fills on each accept

Assume a fixed drain rate

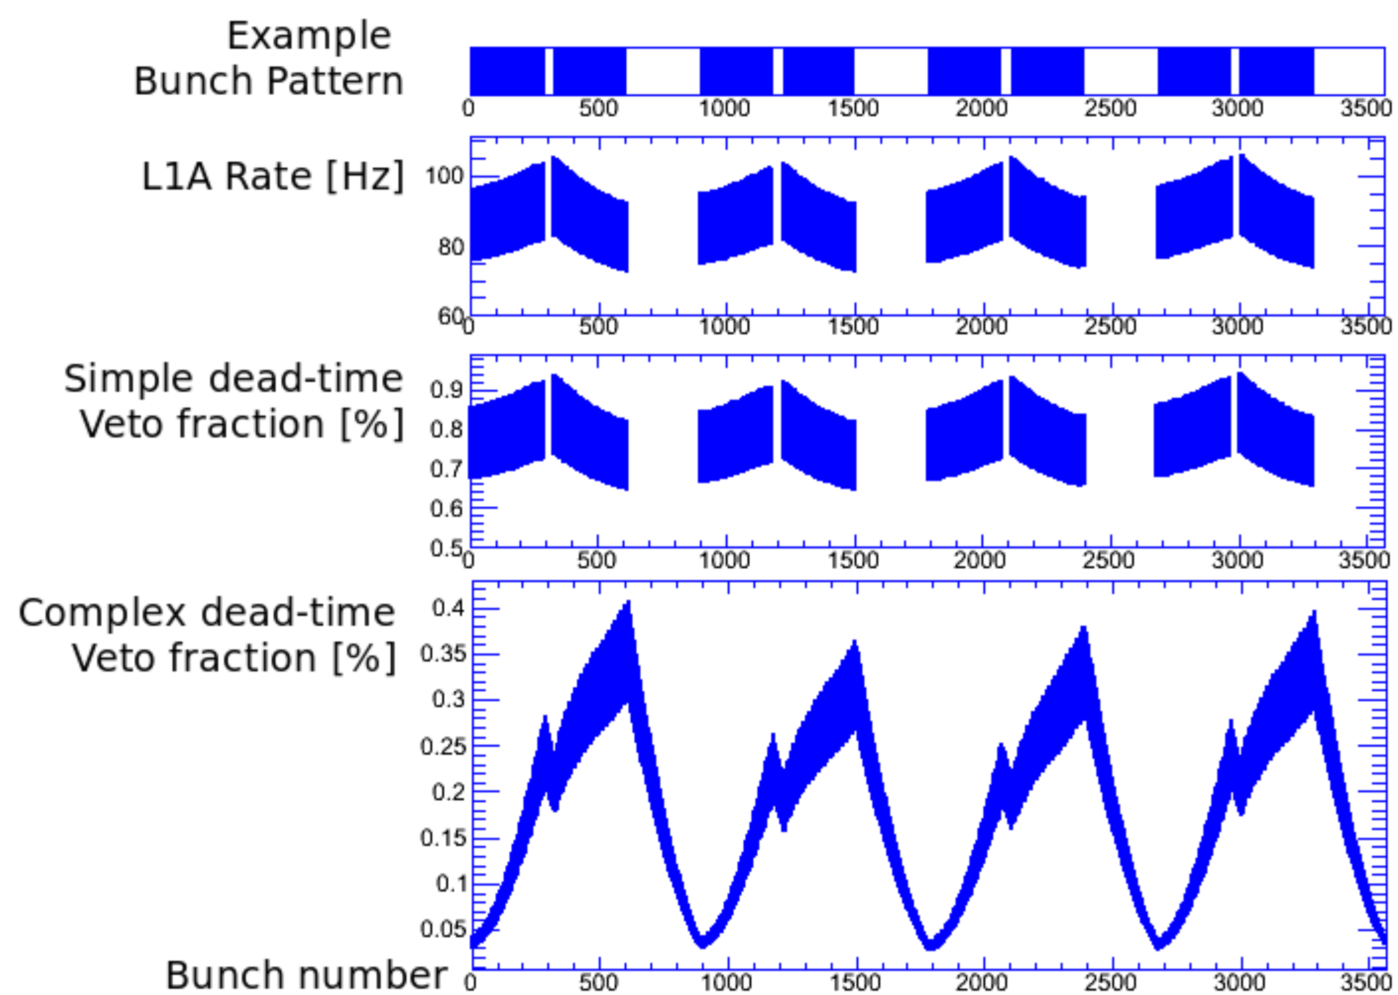

Many possible origin sources

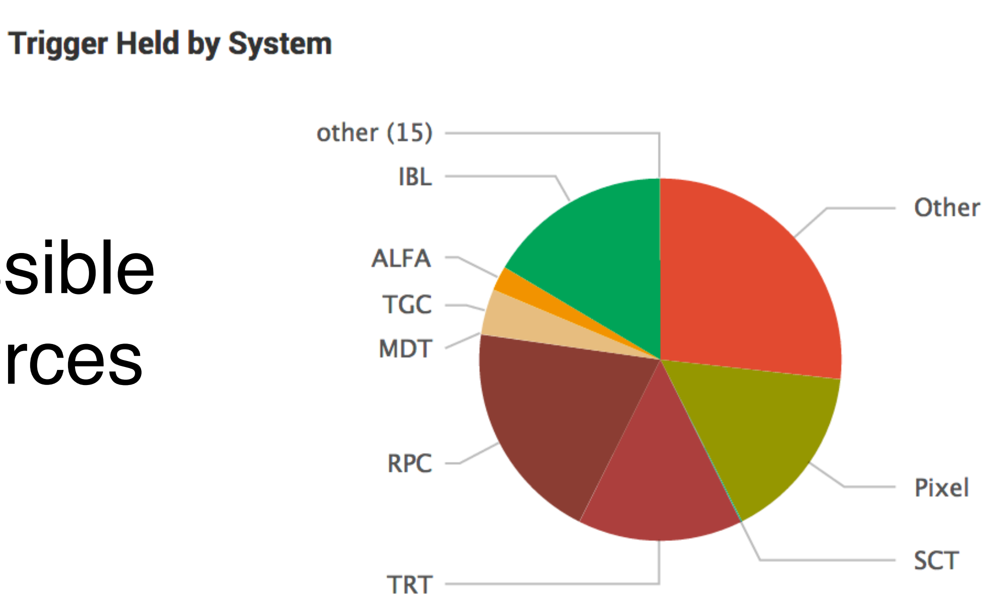

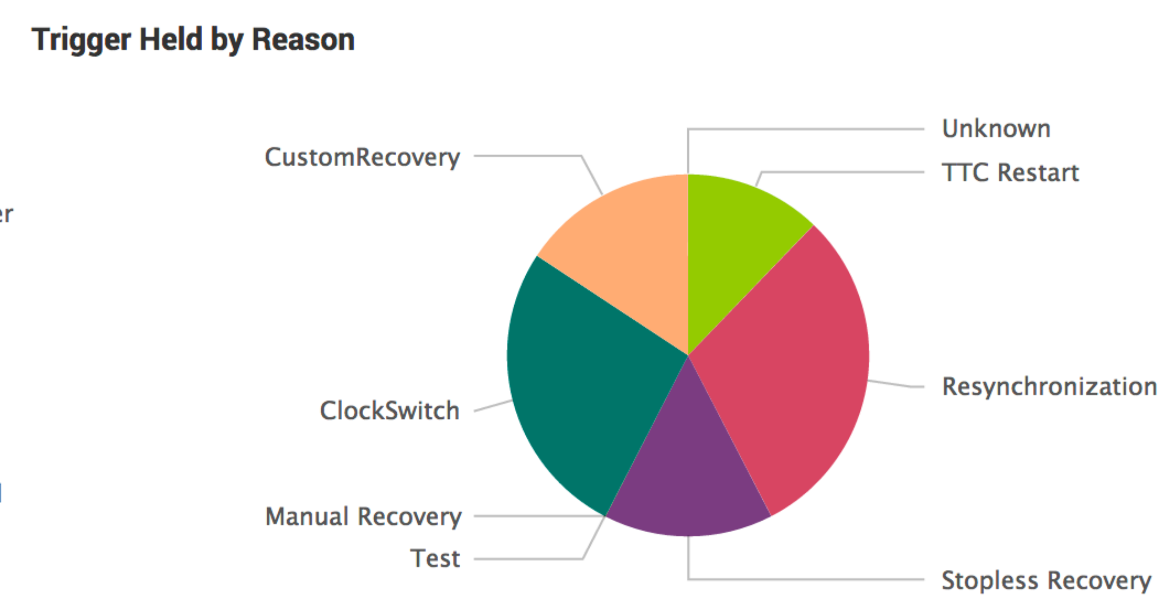

Synchronizing signals

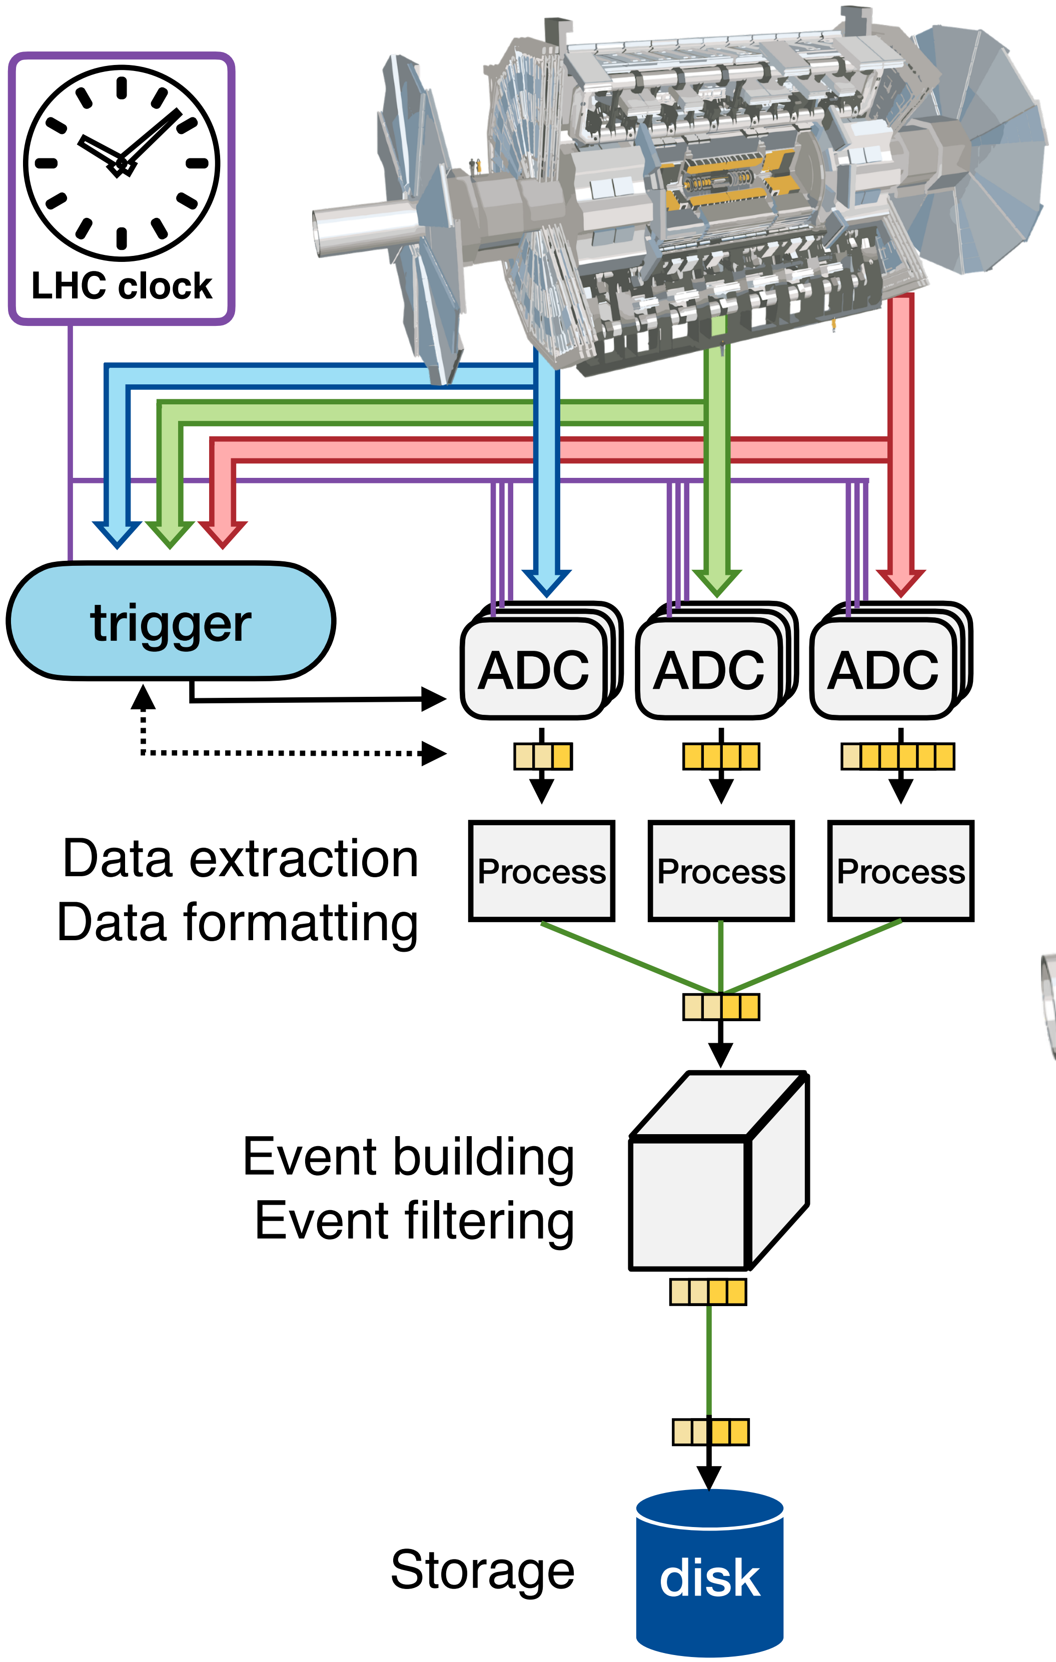

Many front-ends with buffers of various depths

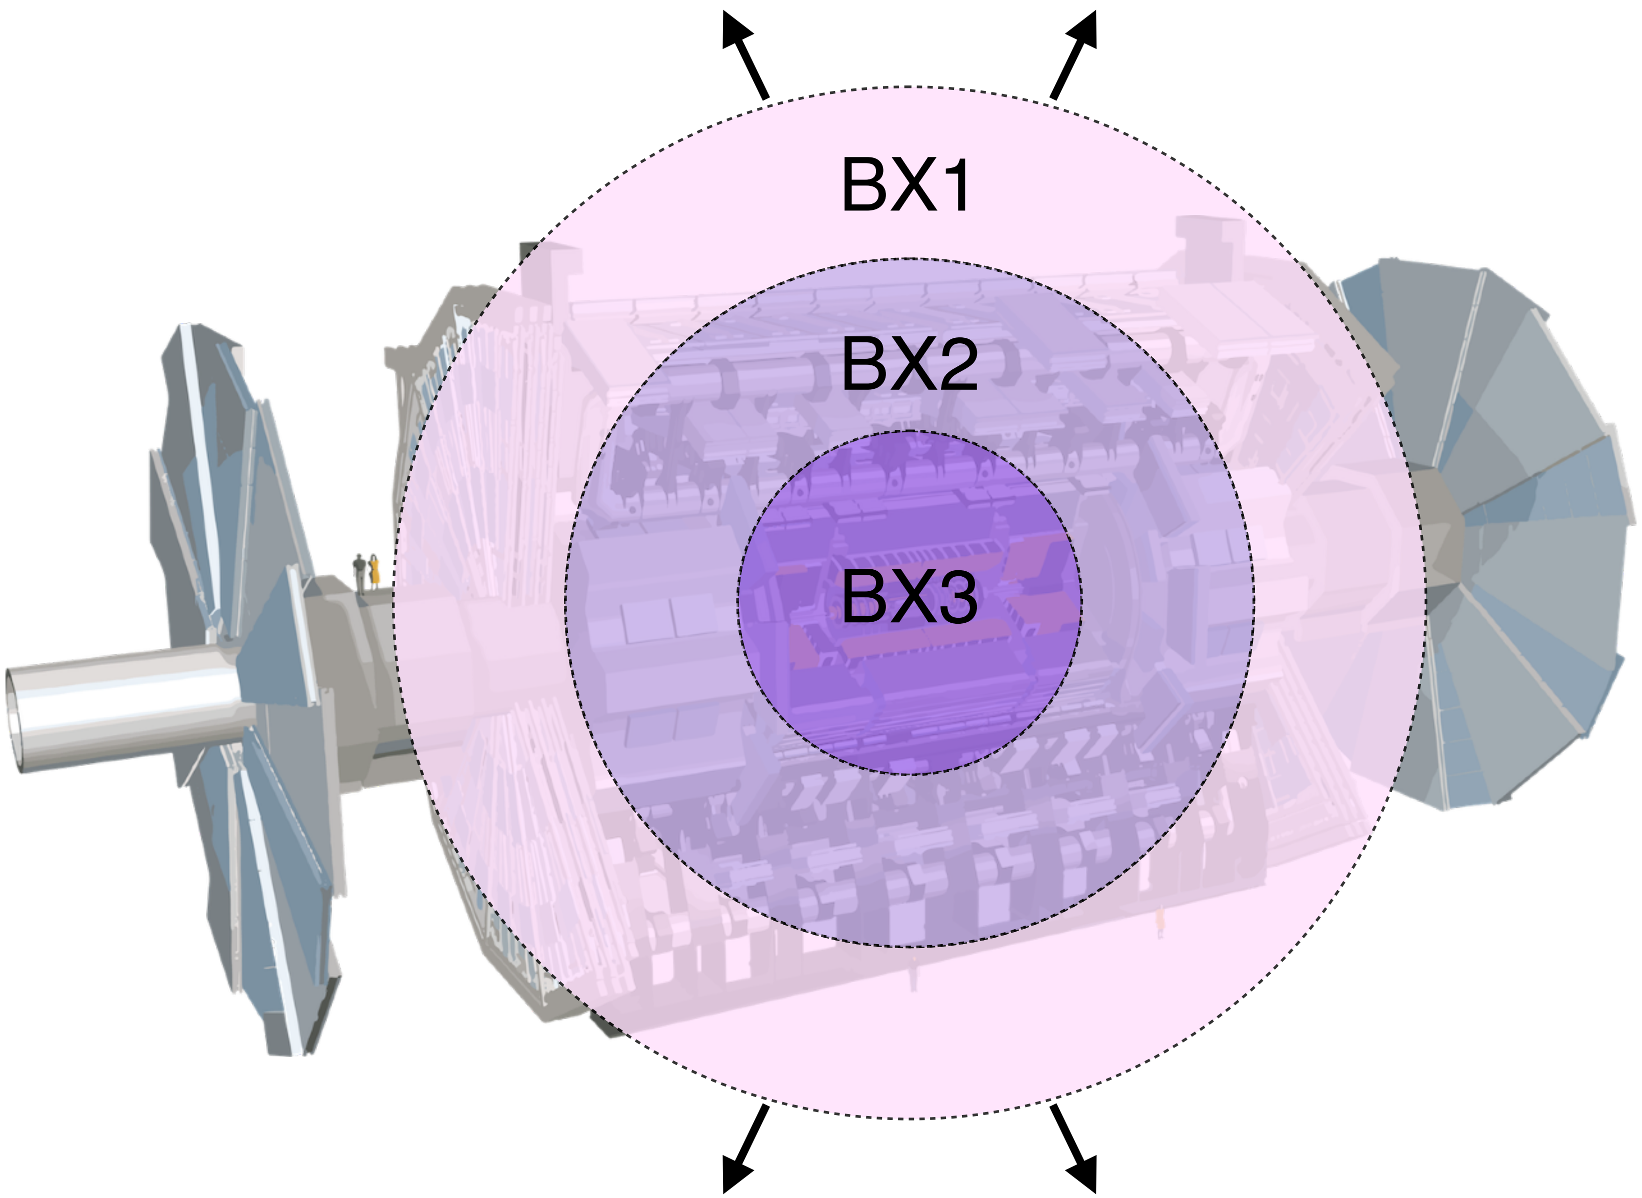

Also critical for aligning signals in time!

Synchronizing signals?

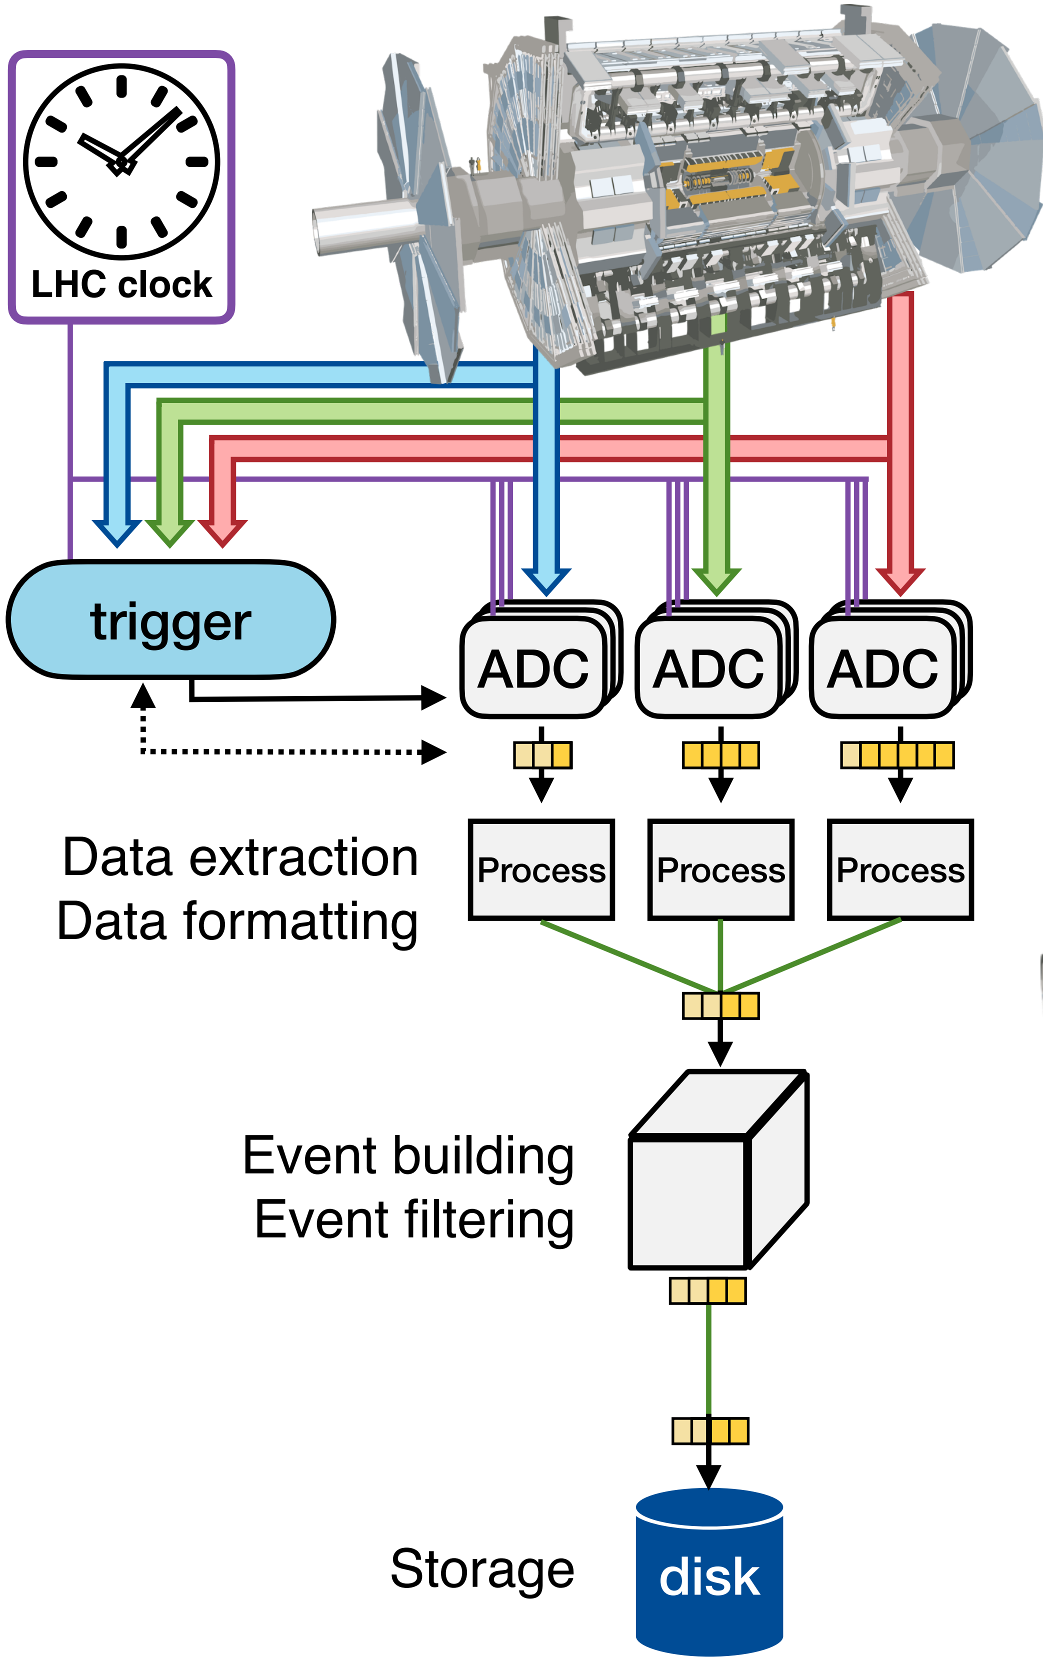

Many front-ends with buffers of various depths

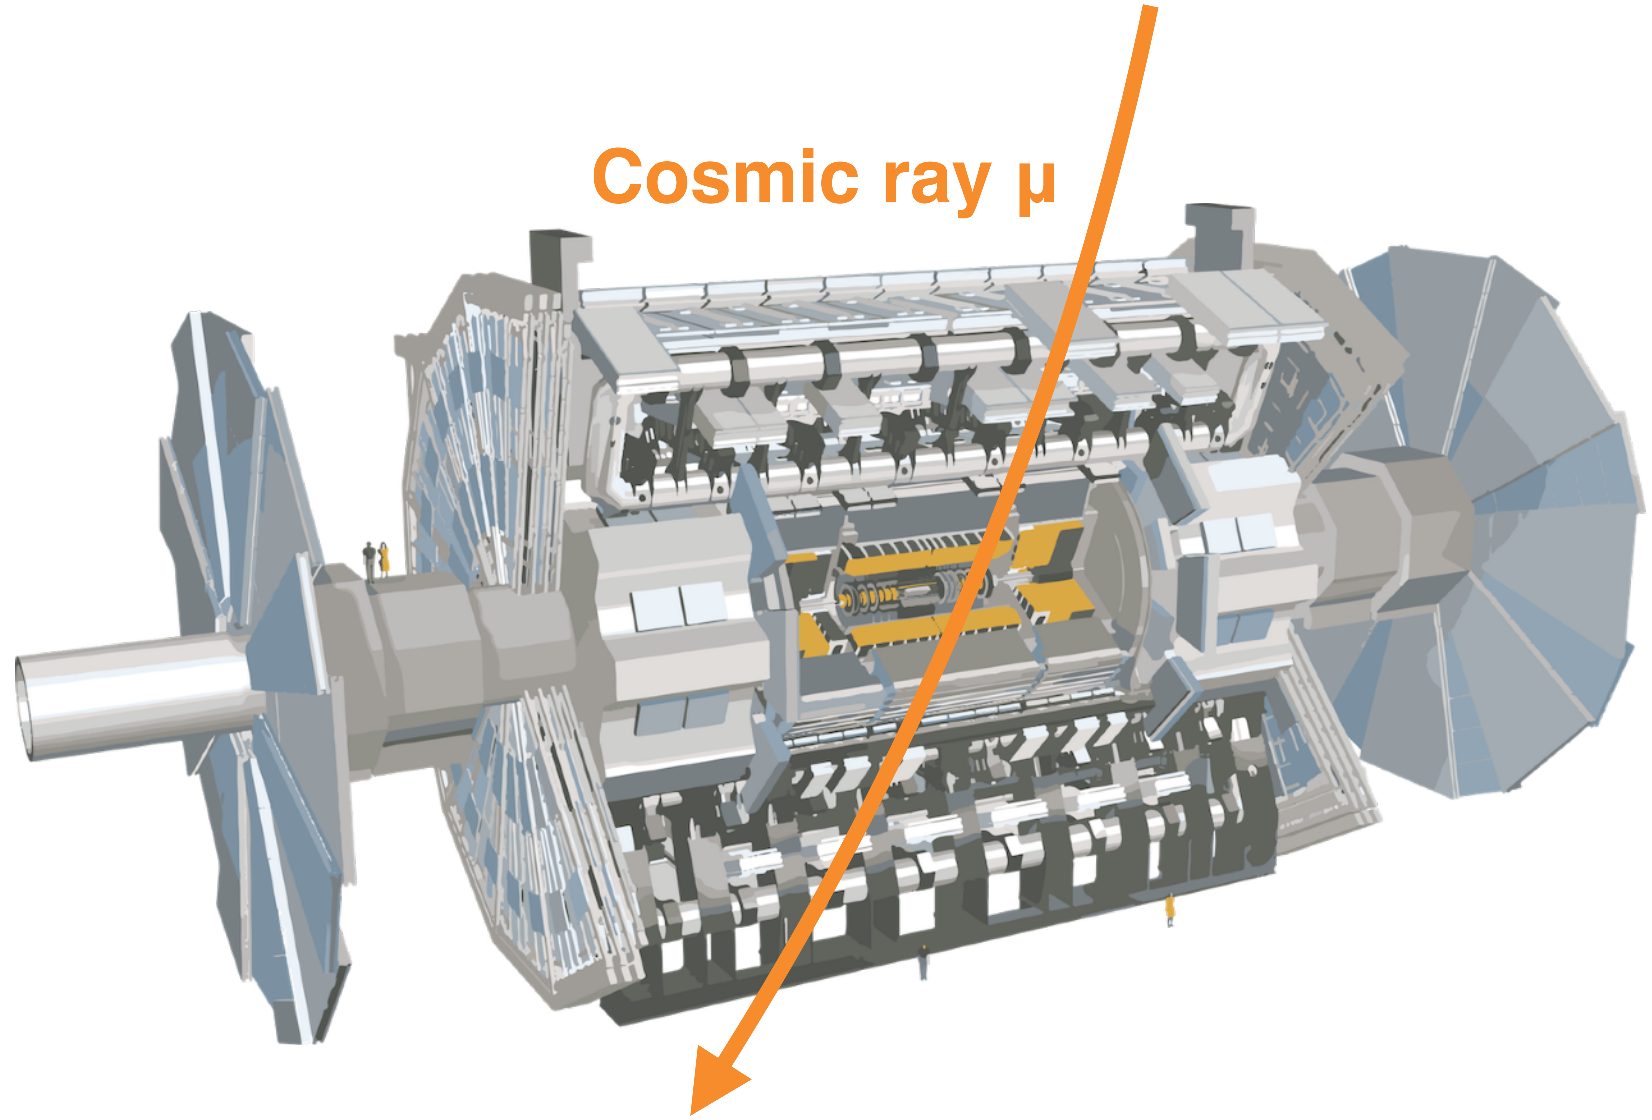

Side Q: then how can we reconstruct particles that don't originate from the beam spot?

Scaling up to the full DAQ

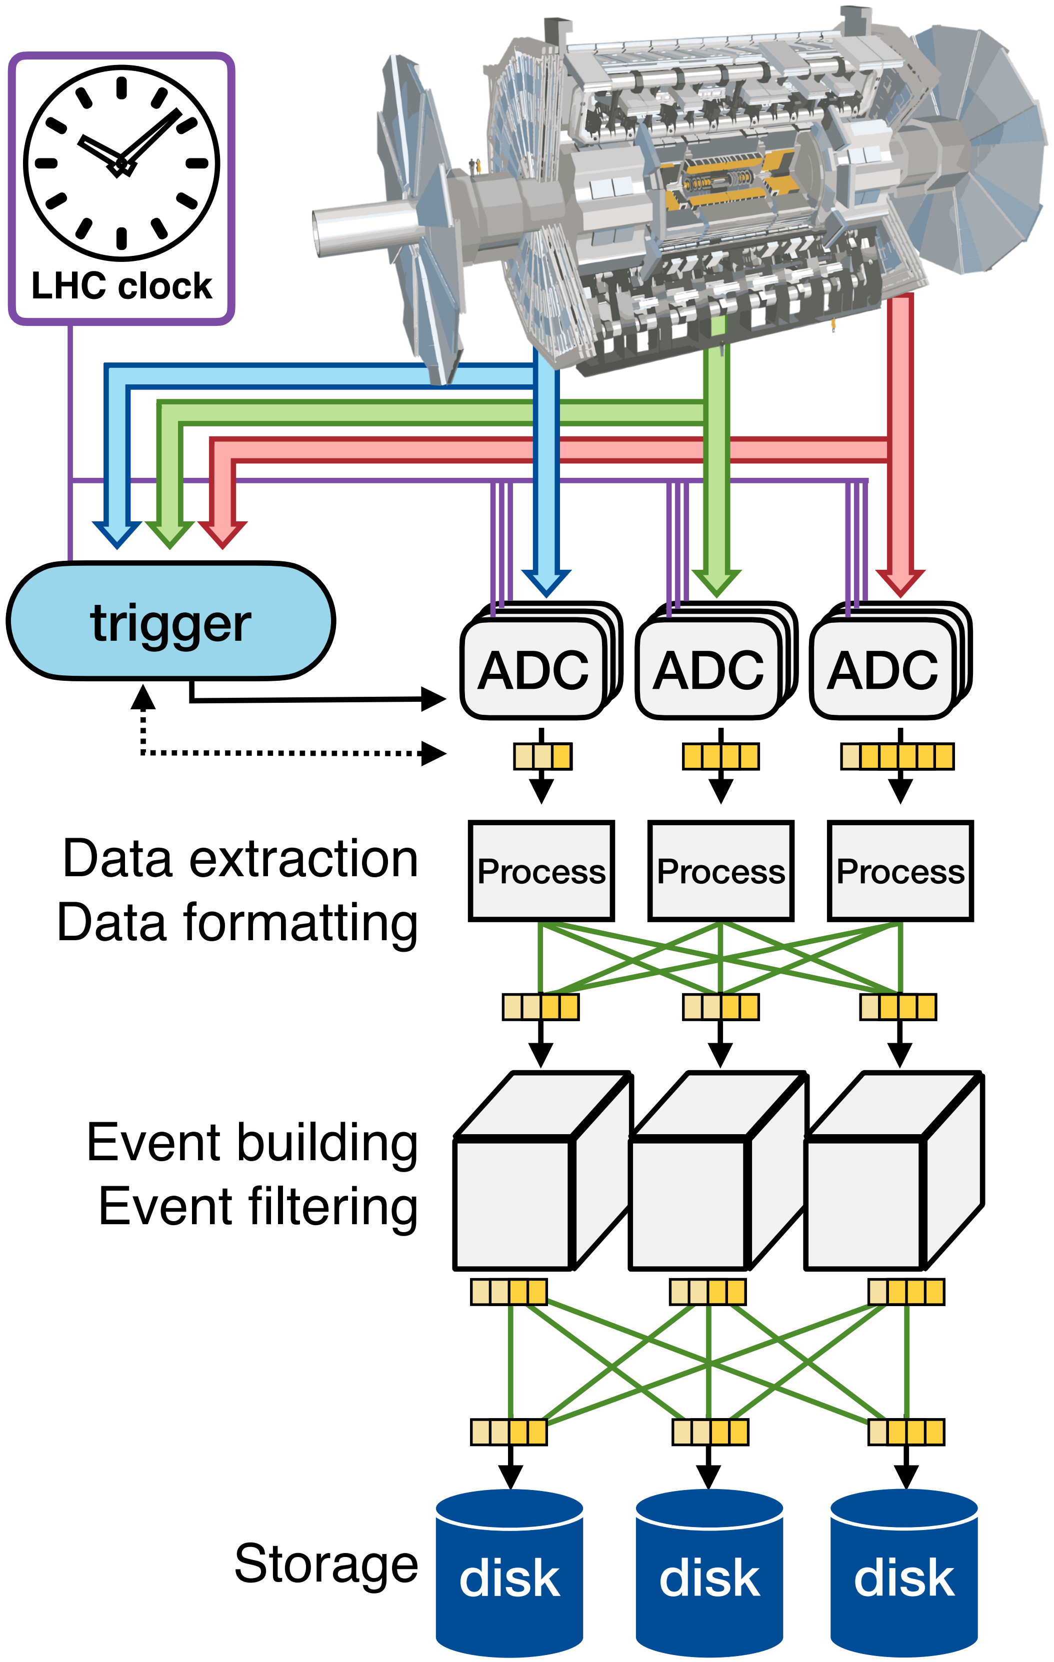

System of many processors, event builders, filters, disks, … requires some choice of **readout topology**.

Buses

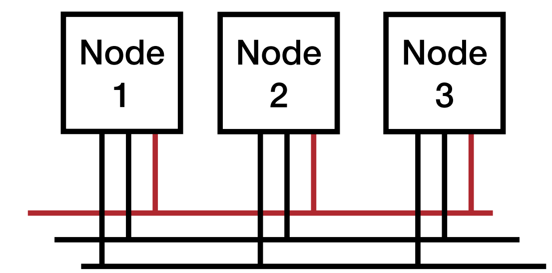

control & data lanes
(simplicity)

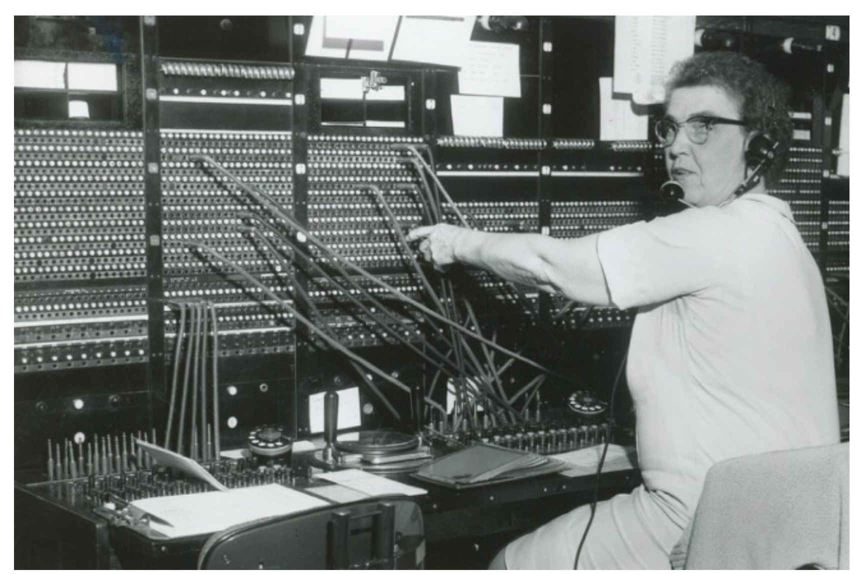

Networks

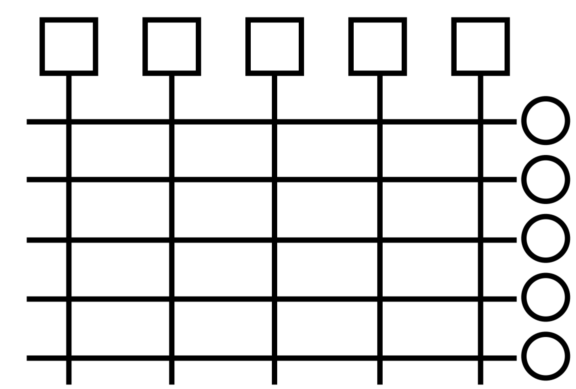

point to point
(scalability)

Scaling up to the full DAQ

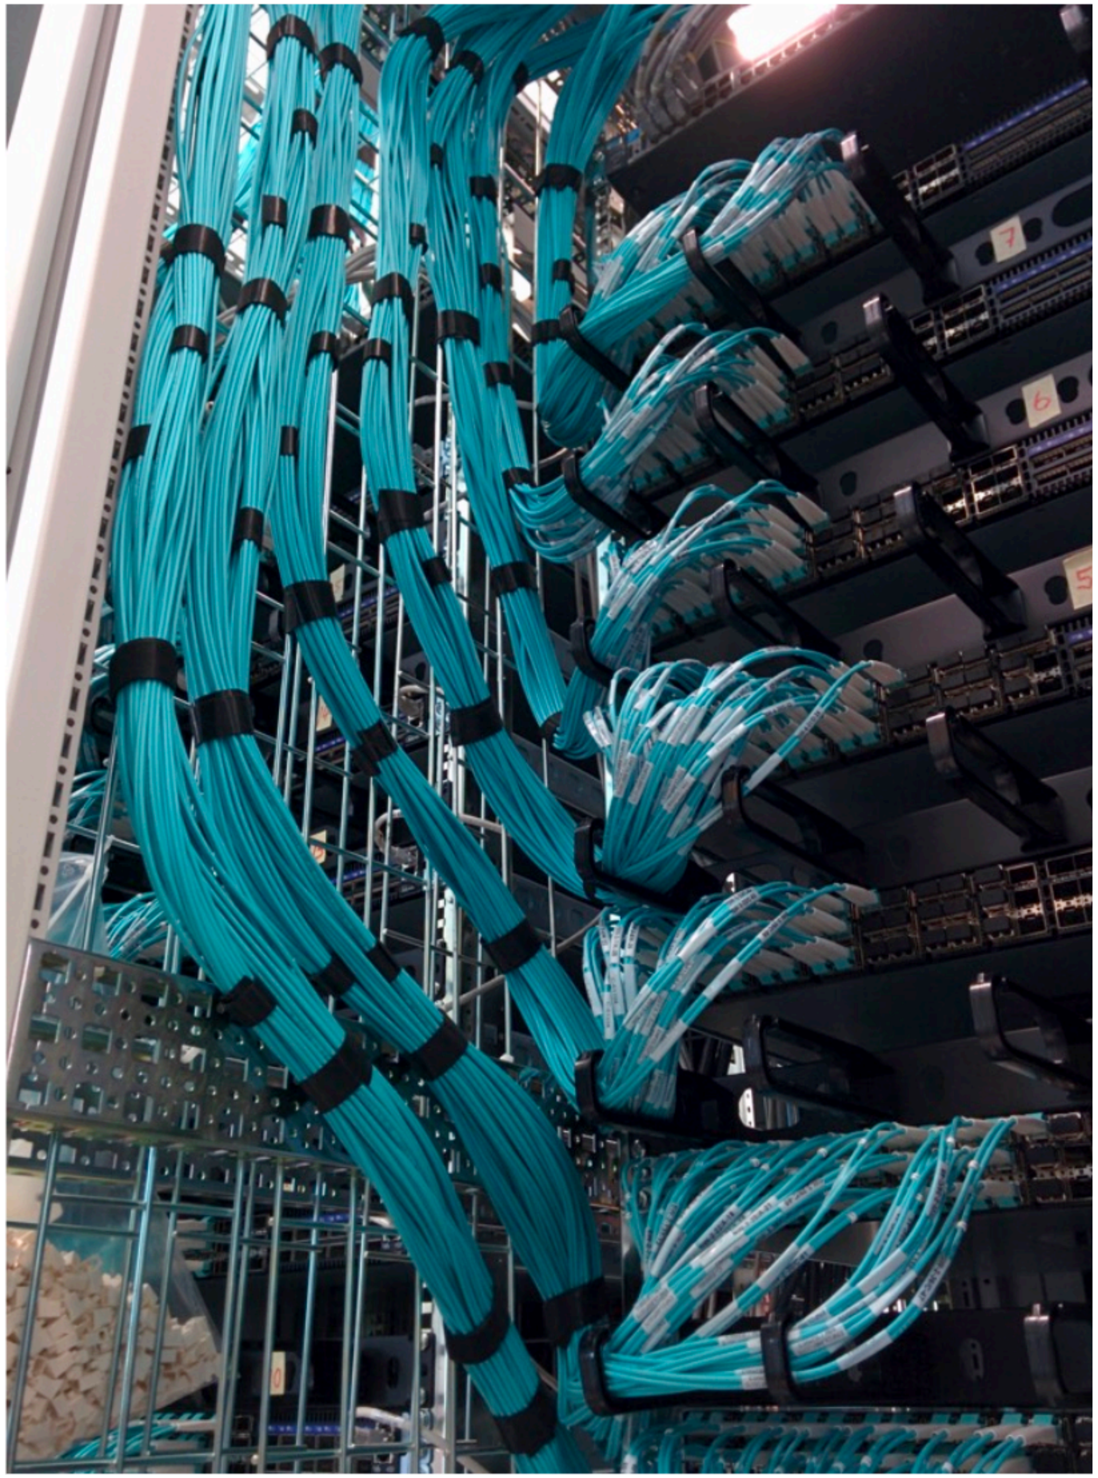

CMS Data concentrator ‘patch panels’ and switches

A heterogenous system

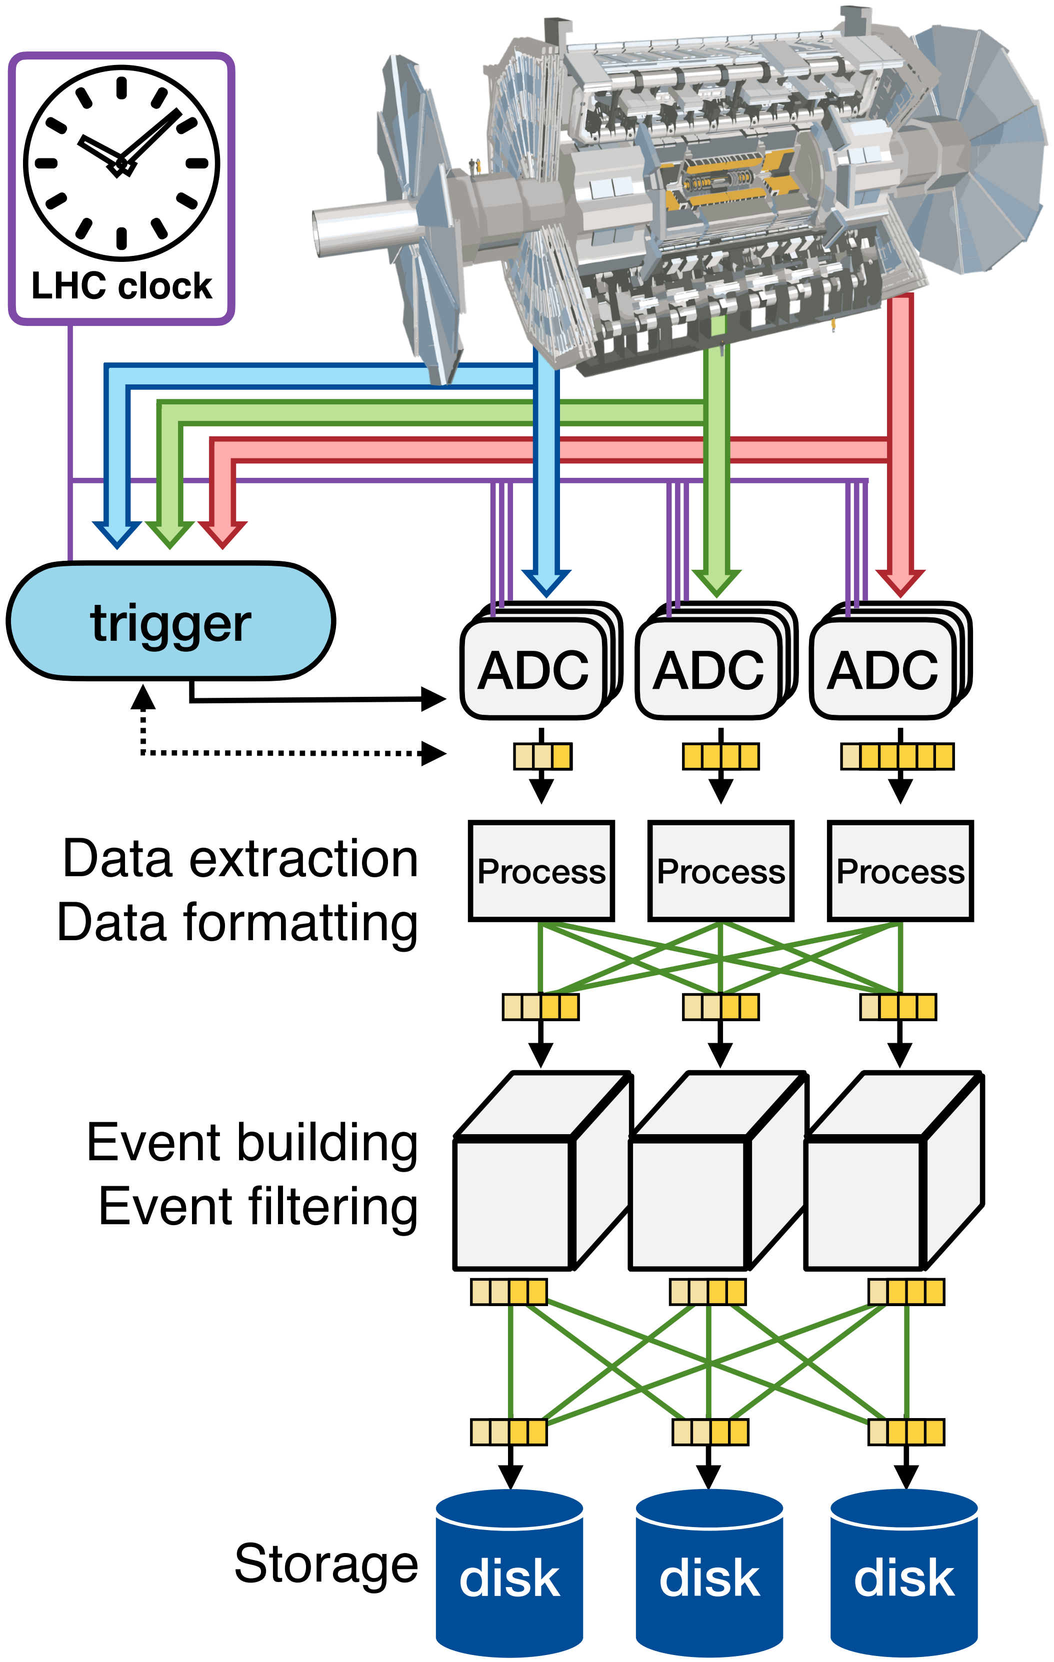

ASIC

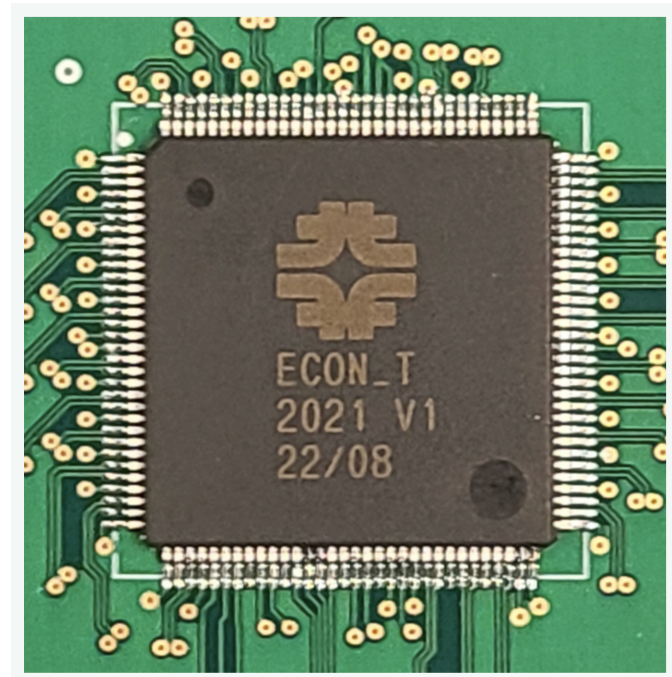

FPGA

CPU

GPU

A heterogenous system

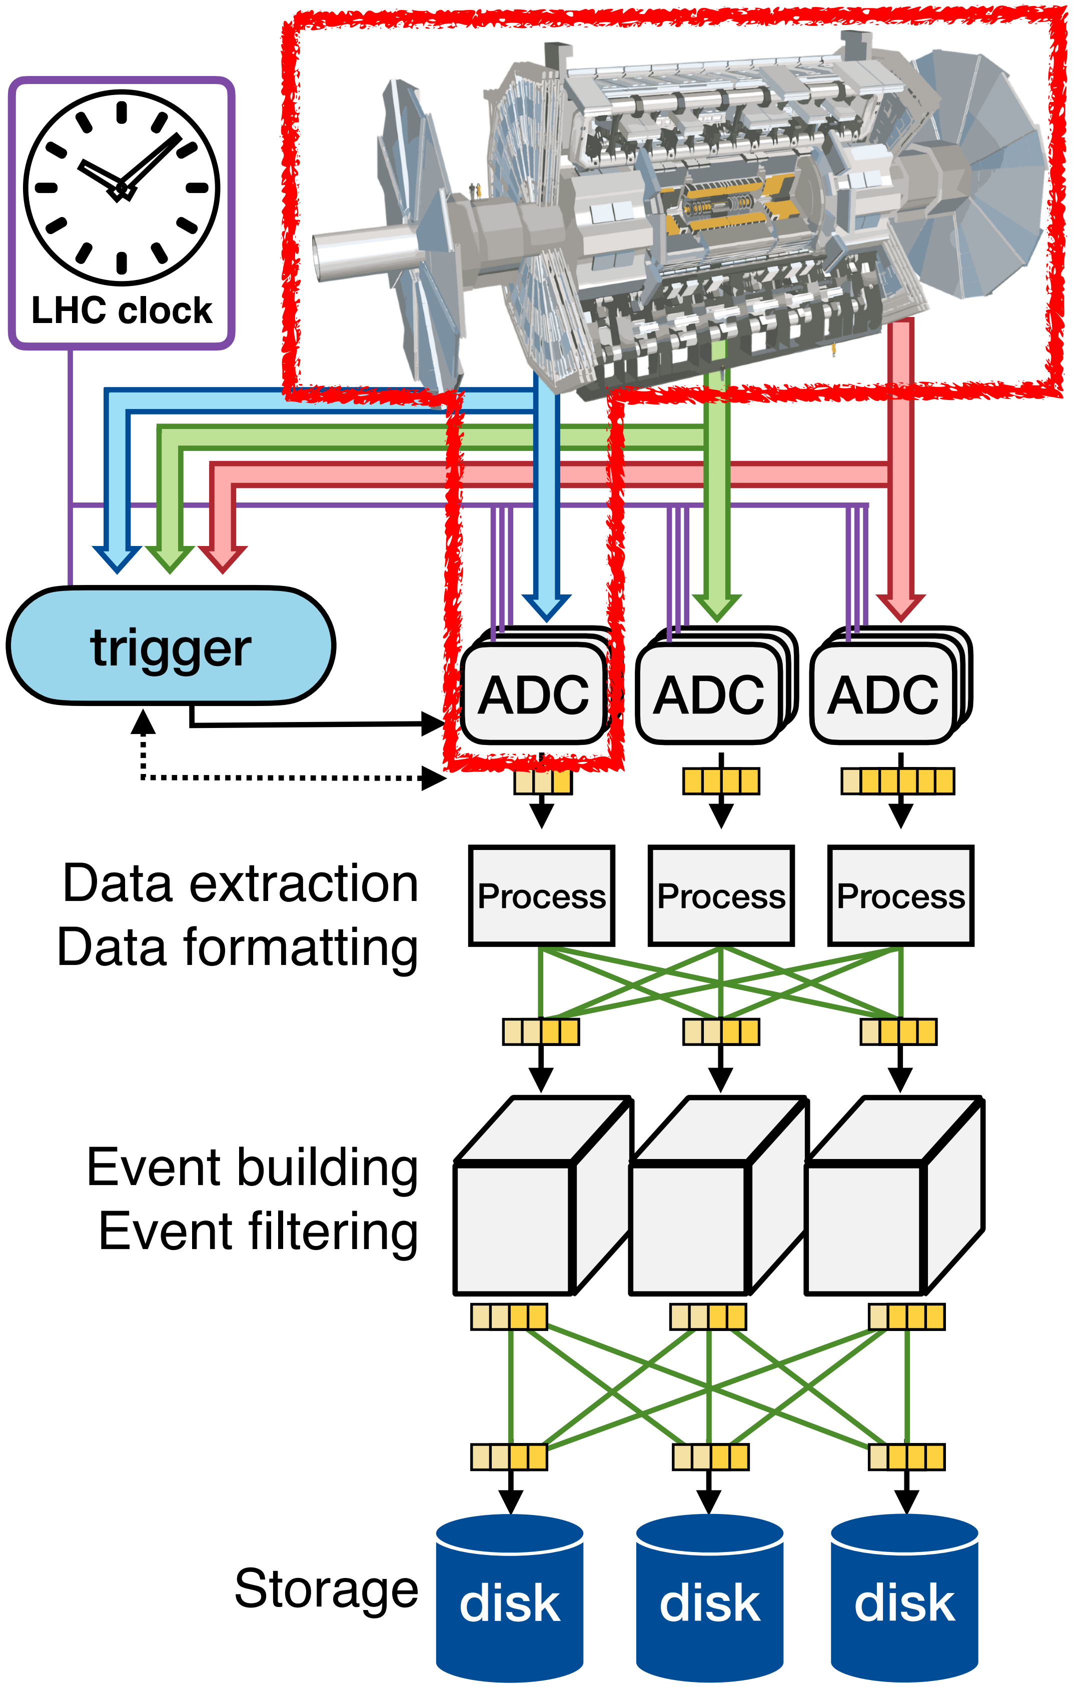

ASIC

FPGA

CPU

GPU

A heterogenous system

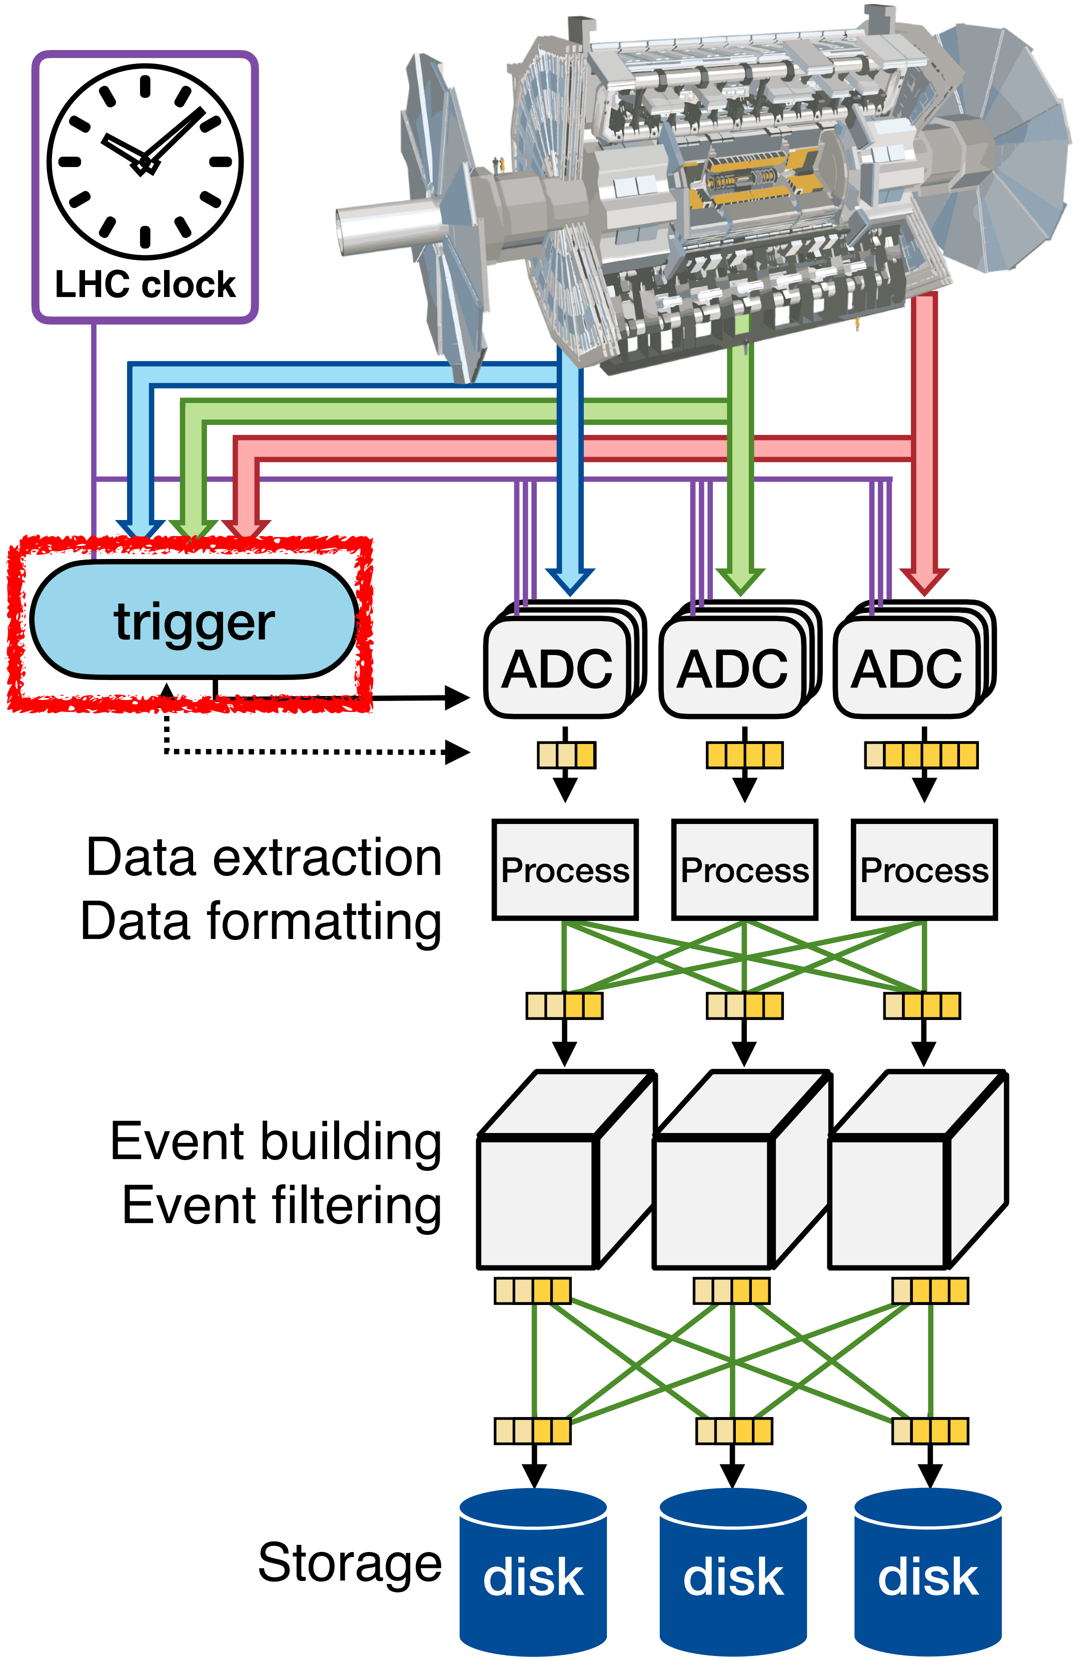

ASIC

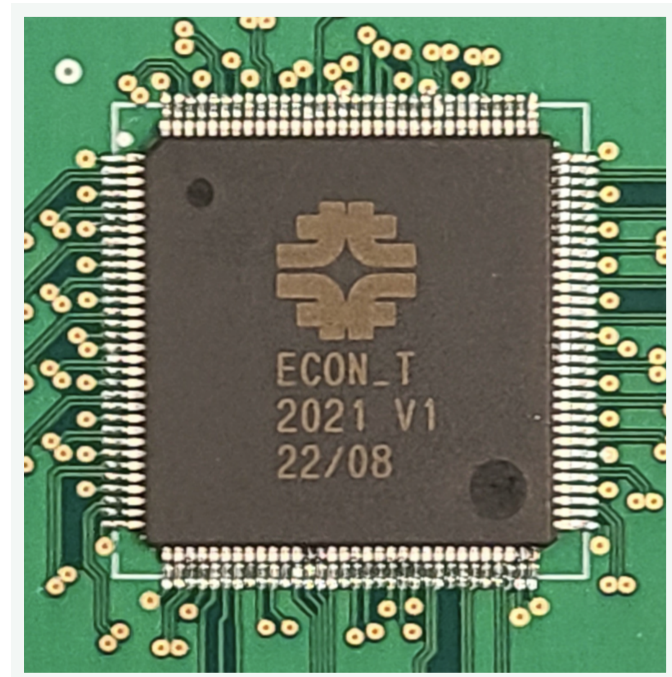

FPGA

CPU

GPU

A heterogenous system

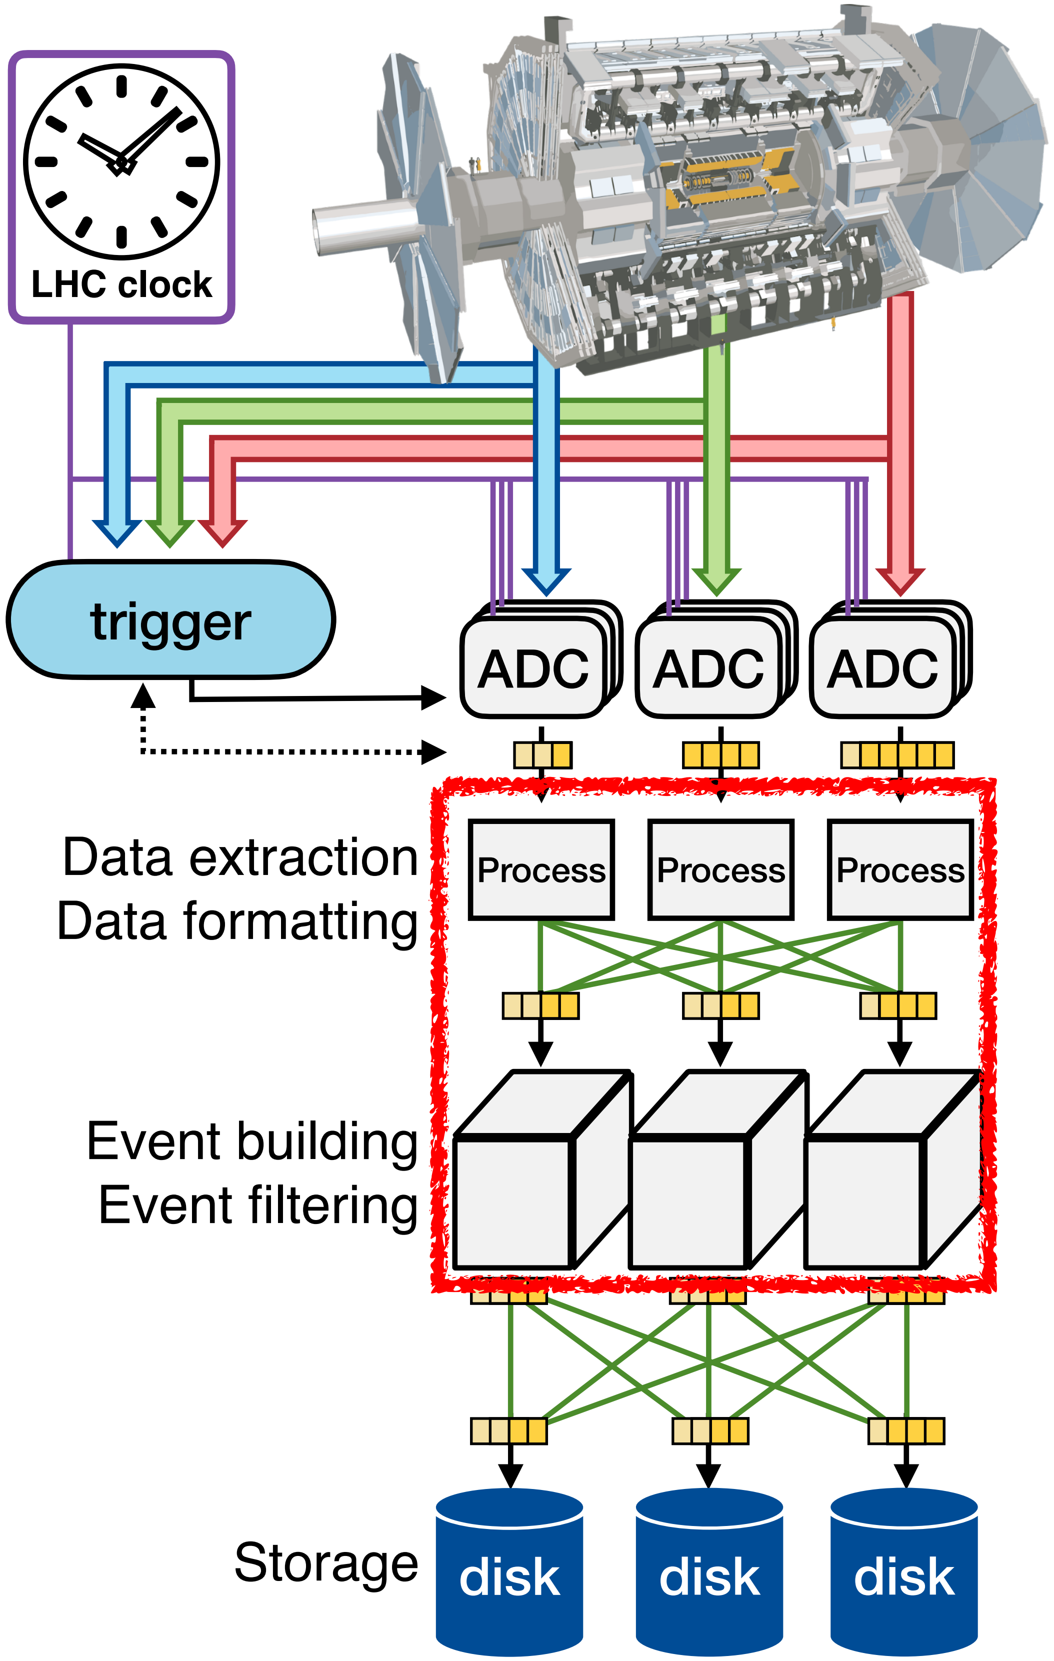

ASIC

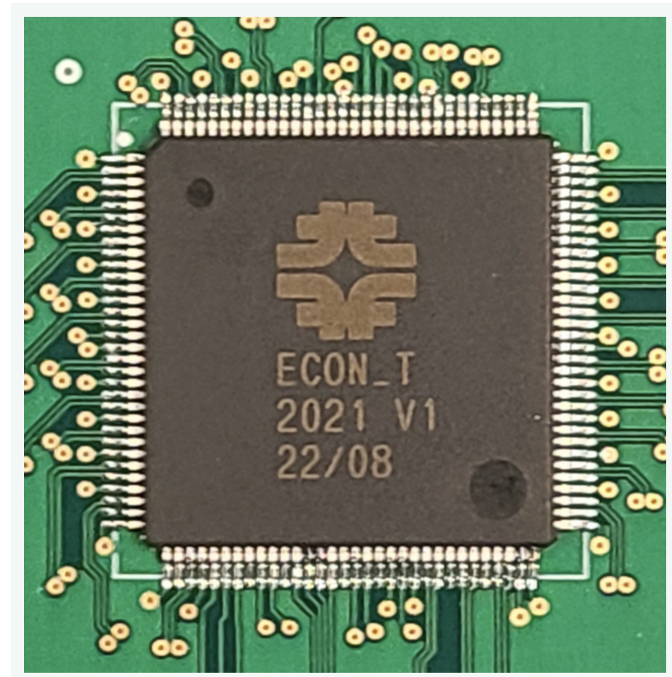

FPGA

CPU

GPU

A heterogenous system

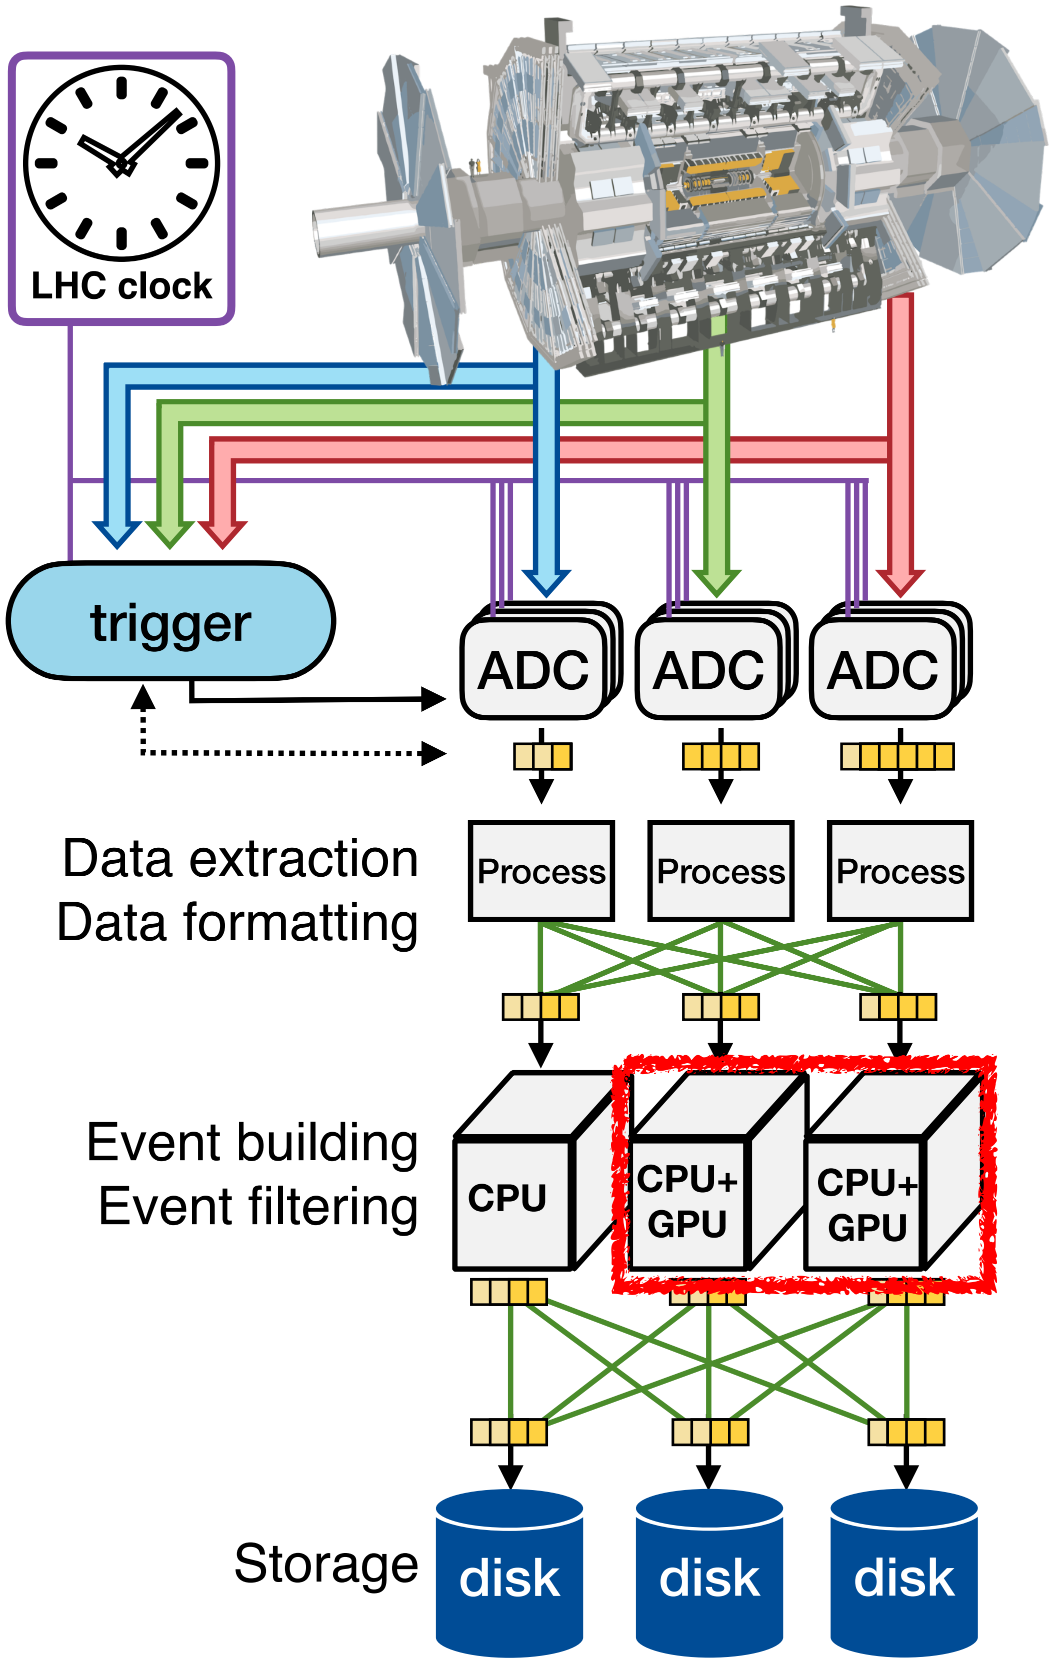

ASIC

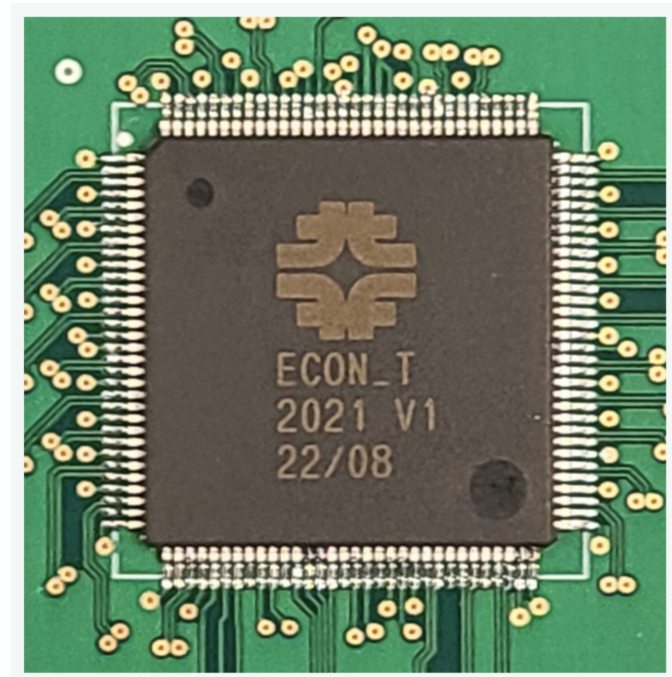

FPGA

CPU

GPU

Inside an FPGA (I)

Field Programmable Gate Arrays "compute across space and time"

The workhorse technology of LHC hardware triggers

Fully re-programmable:

Build custom circuits by connecting:

Memories, Multipliers, and other configurable logic blocks

Enables:

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

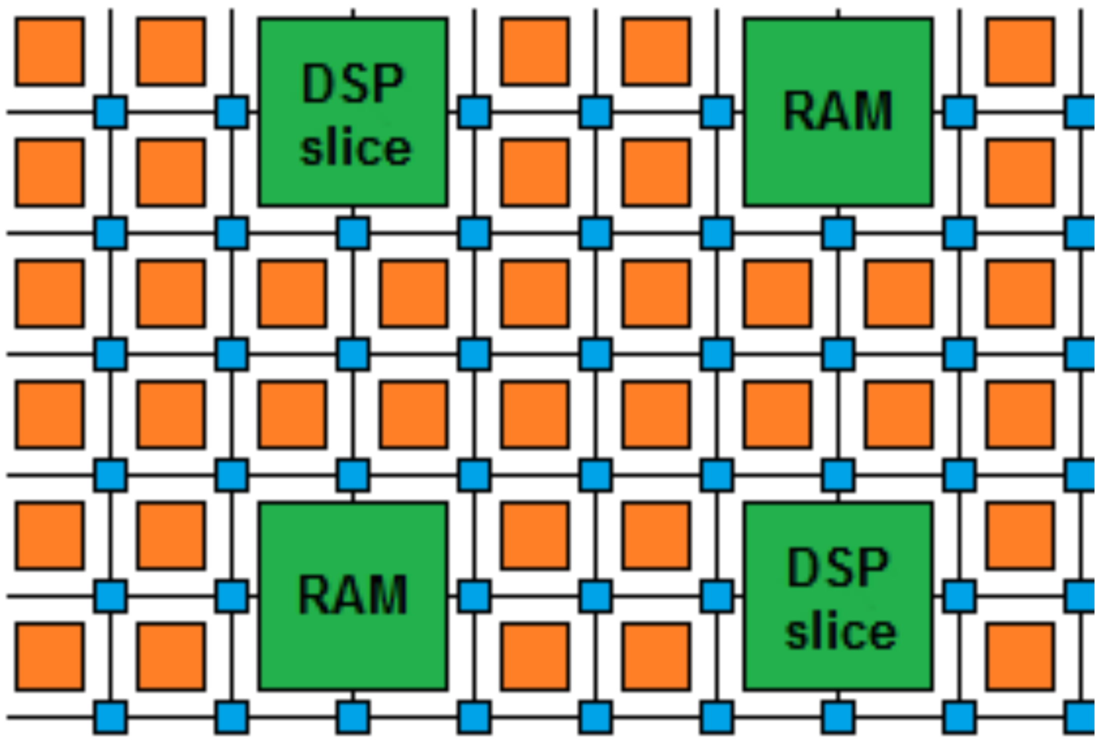

Inside an FPGA (II)

Combination of elements, including LUTs, memories and DSP slices.

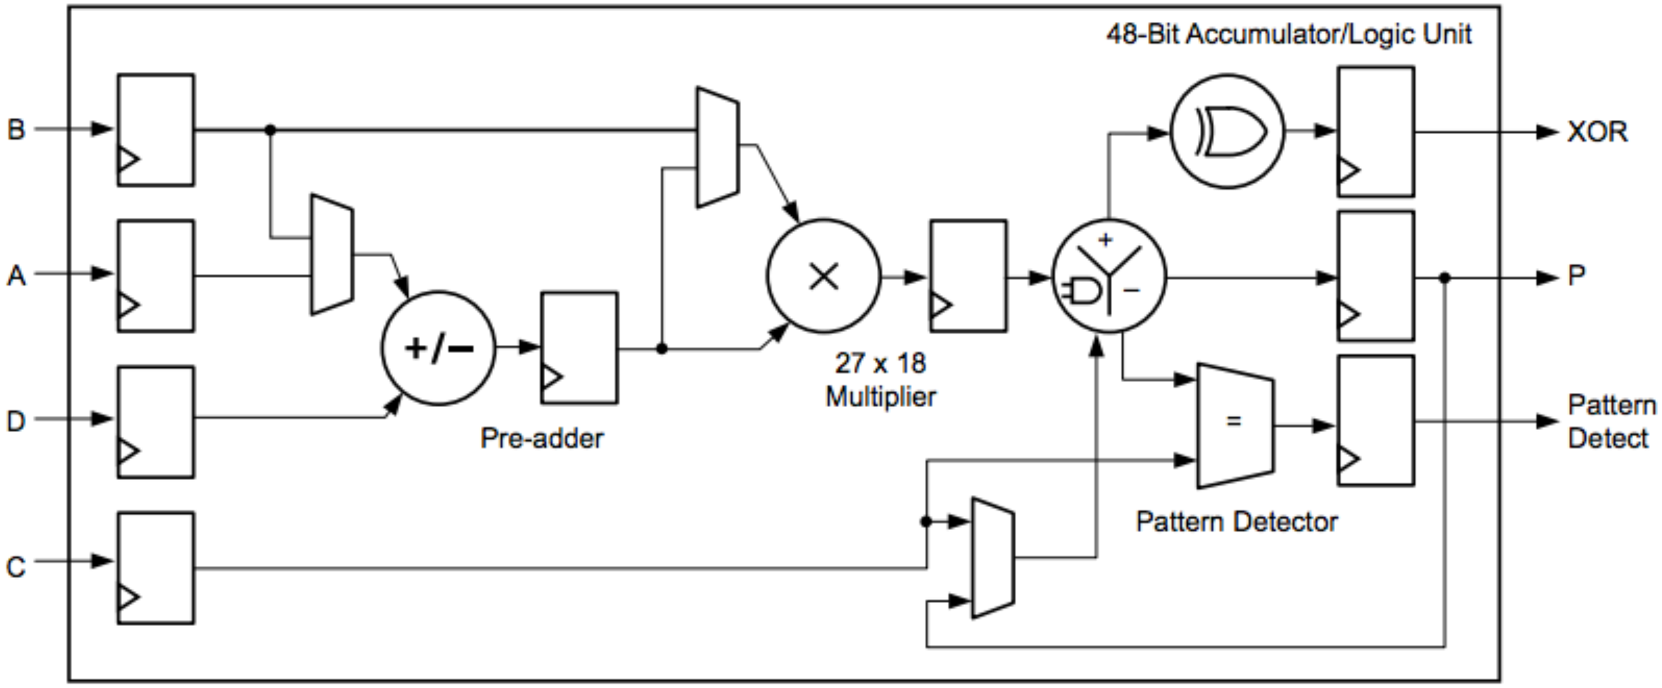

Xilinx DSP slice 48b accumulator

Rapidly improving tech:

Typical Run 1:

~ 100 DSP / FPGA

Typical today:

~ 5k DSP / FPGA

Many DSPs can make "on the fly" calculations feasible.
(Especially tasks with many/large matrix multiplications)

Aside: but trigger is not just an FPGA!

The complete trigger system includes **hundreds of similar boards**.

High speed optical links
(120 x 28 Gbps)

UltraScale+ FPGA
(12k DSP slices)

System-on-module

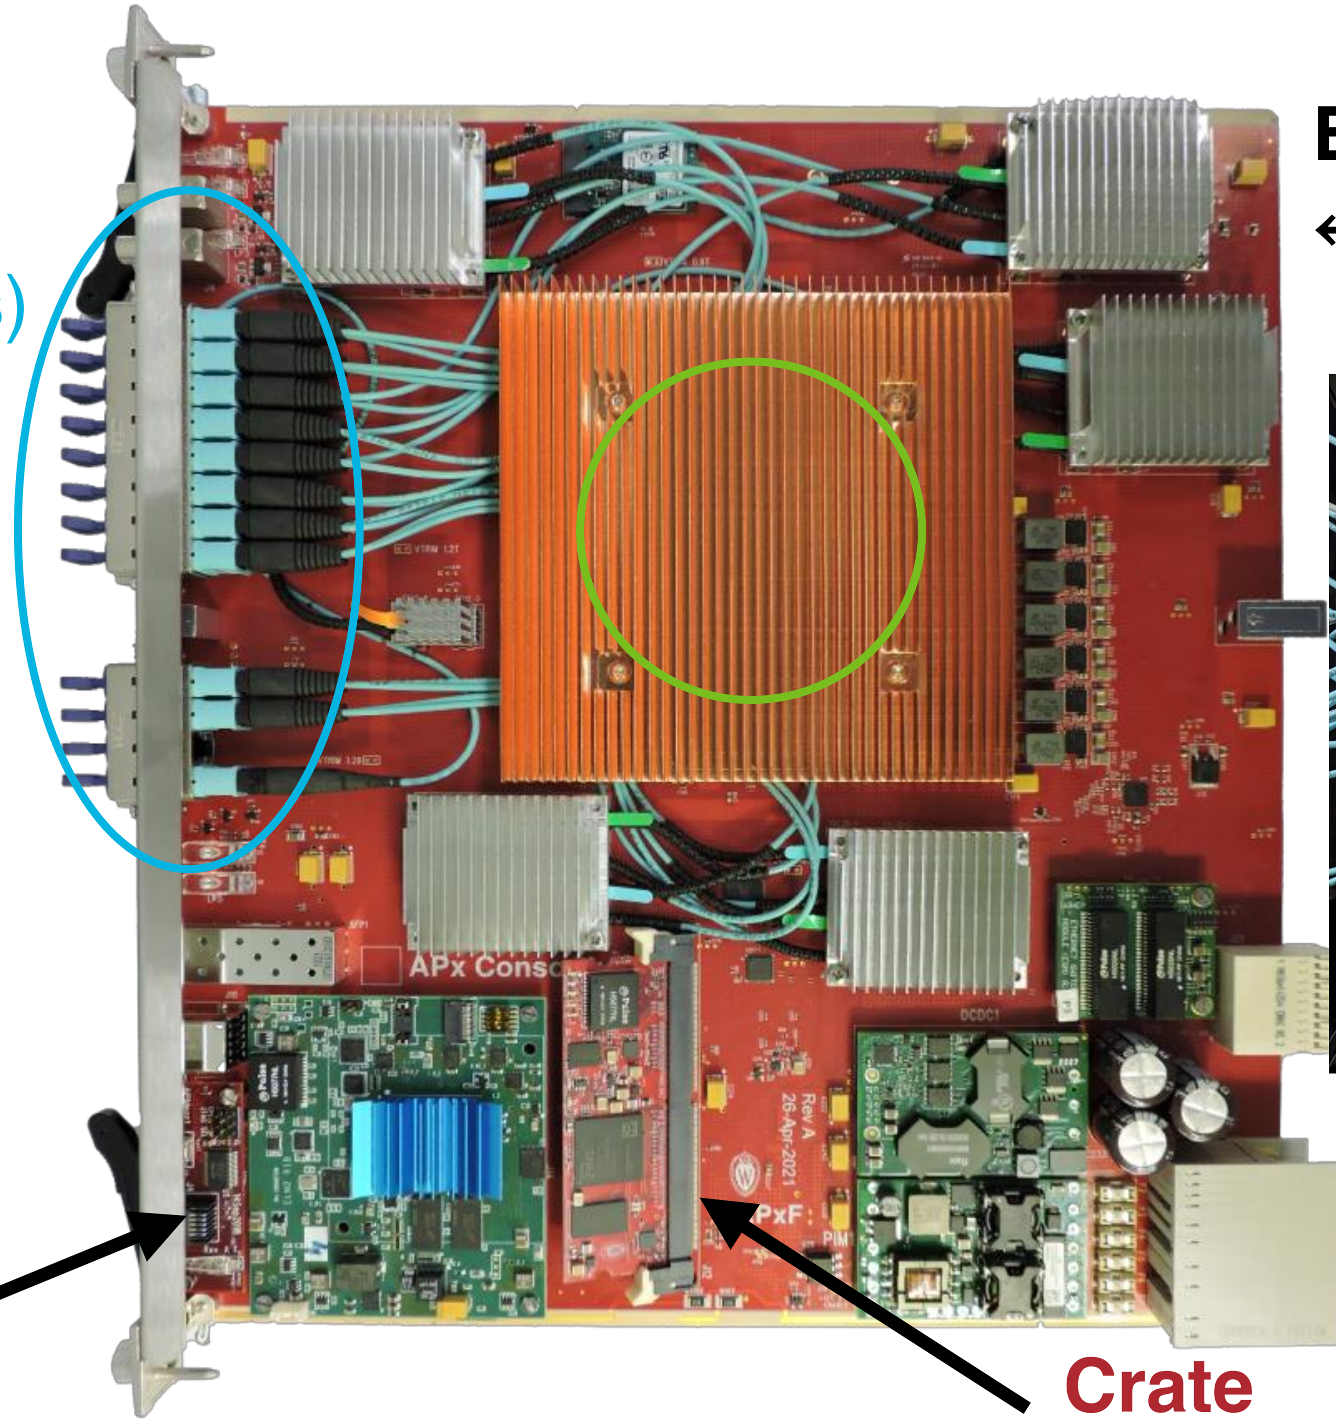

Electrical ↔optical

heatsinks

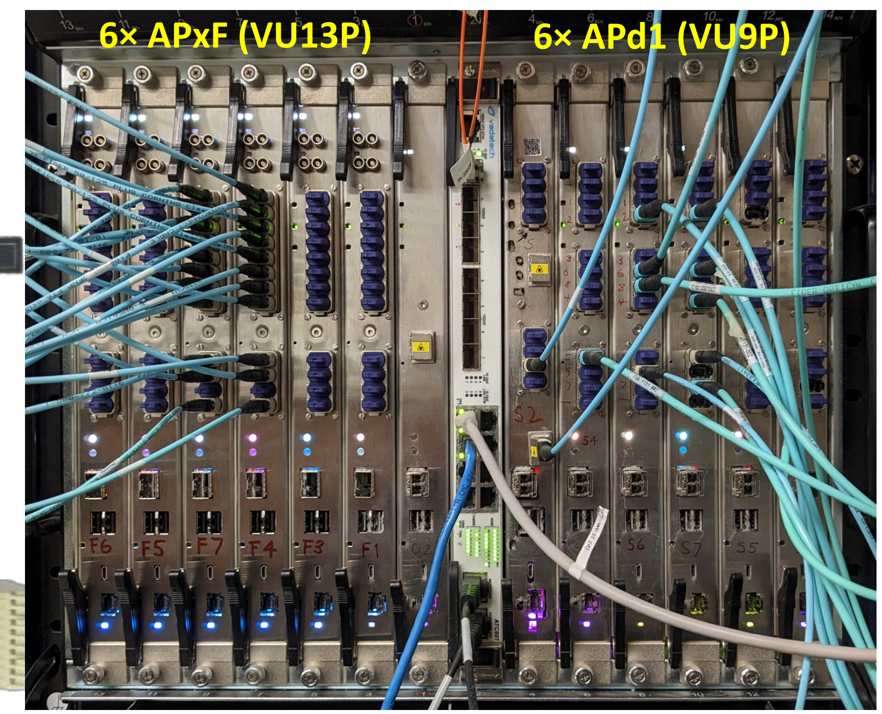

Shelf of 12 'blades'

Crate management

"APx" prototype trigger board for CMS Phase-II

Plan

- **This afternoon: Concepts**
  - Introduction and physics goals of the LHC Experiments
  - The role of the trigger and DAQ systems
  - DAQ concepts: simple toy model → complex systems
  - **Trigger concepts: from hardware to architecting a system**
- **Tomorrow morning: Details and Implementations**
  - The challenge of high luminosity, and the upgrades
  - ATLAS, CMS hardware triggers, and their evolution
  - Software triggers, and new paradigms for analysis
  - Trigger menus, and you
  - Conclusions, and looking farther ahead

Designing a rapid-reconstruction system

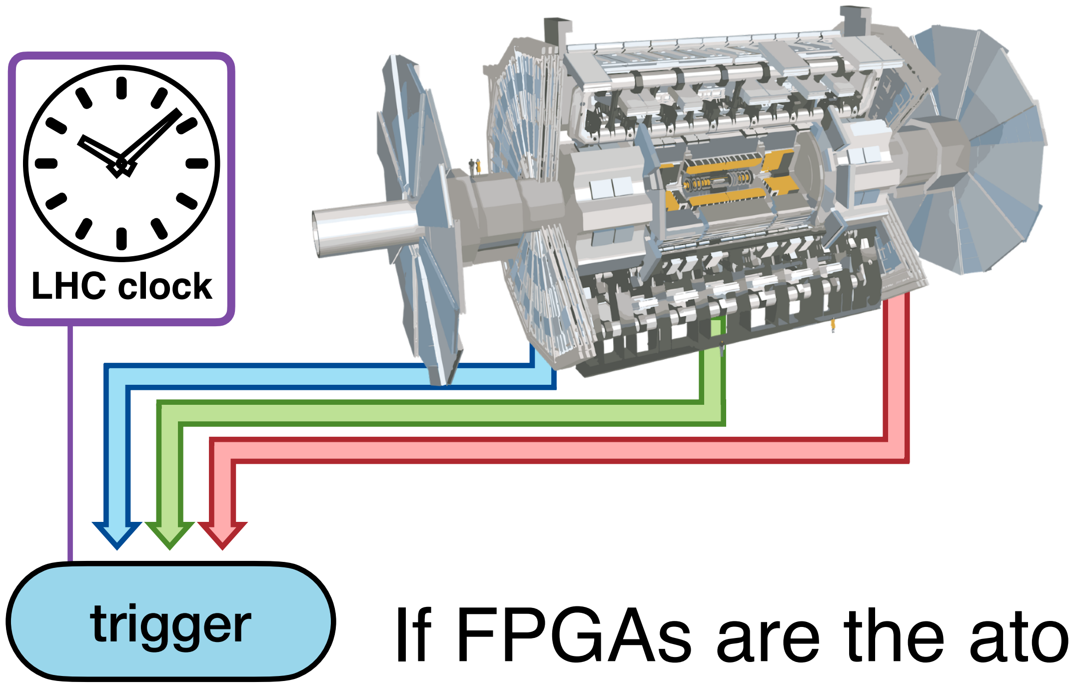

Trigger MUST process new events at the experimental frequency — quickly!

@LHC this is 40 MHz, 1/(25 nanosec)

If FPGAs are the atomic units, how to architect a complete system?

Designing a rapid-reconstruction system

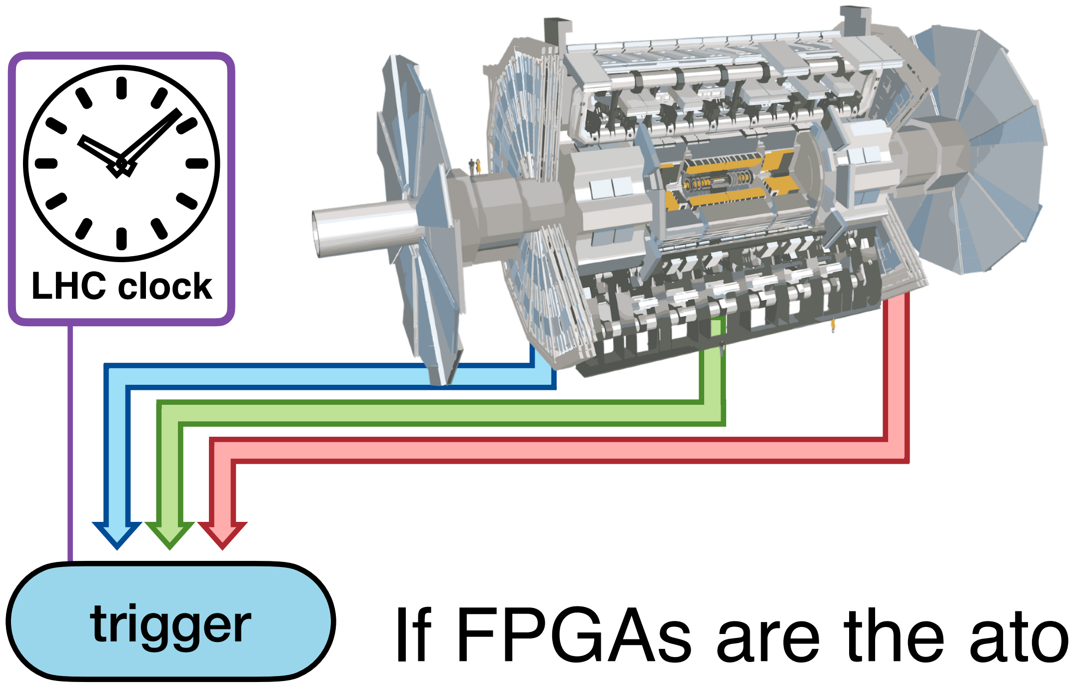

Trigger MUST process new events at the experimental frequency — quickly!

@LHC this is 40 MHz, 1/(25 nanosec)

If FPGAs are the atomic units, how to architect a complete system?

Pipelined processing

Time multiplexing

Latency bugdets

Simplified algorithms

Regional reconstruction

Pipelining

Trigger systems rely on **pipelines** to perform *complicated reconstruction* tasks while handling a stream of *continuous inputs*.

They enable systems with *long latency*, *short initiation interval*.

???

???

???

A real life example

Trigger systems rely on **pipelines** to perform *complicated reconstruction* tasks while handling a stream of *continuous inputs*.

They enable systems with *long latency*, *short initiation interval.*

Example: laundry in my apartment building:

1 washer (runs 30min) and 1 dryer (runs 1hr).

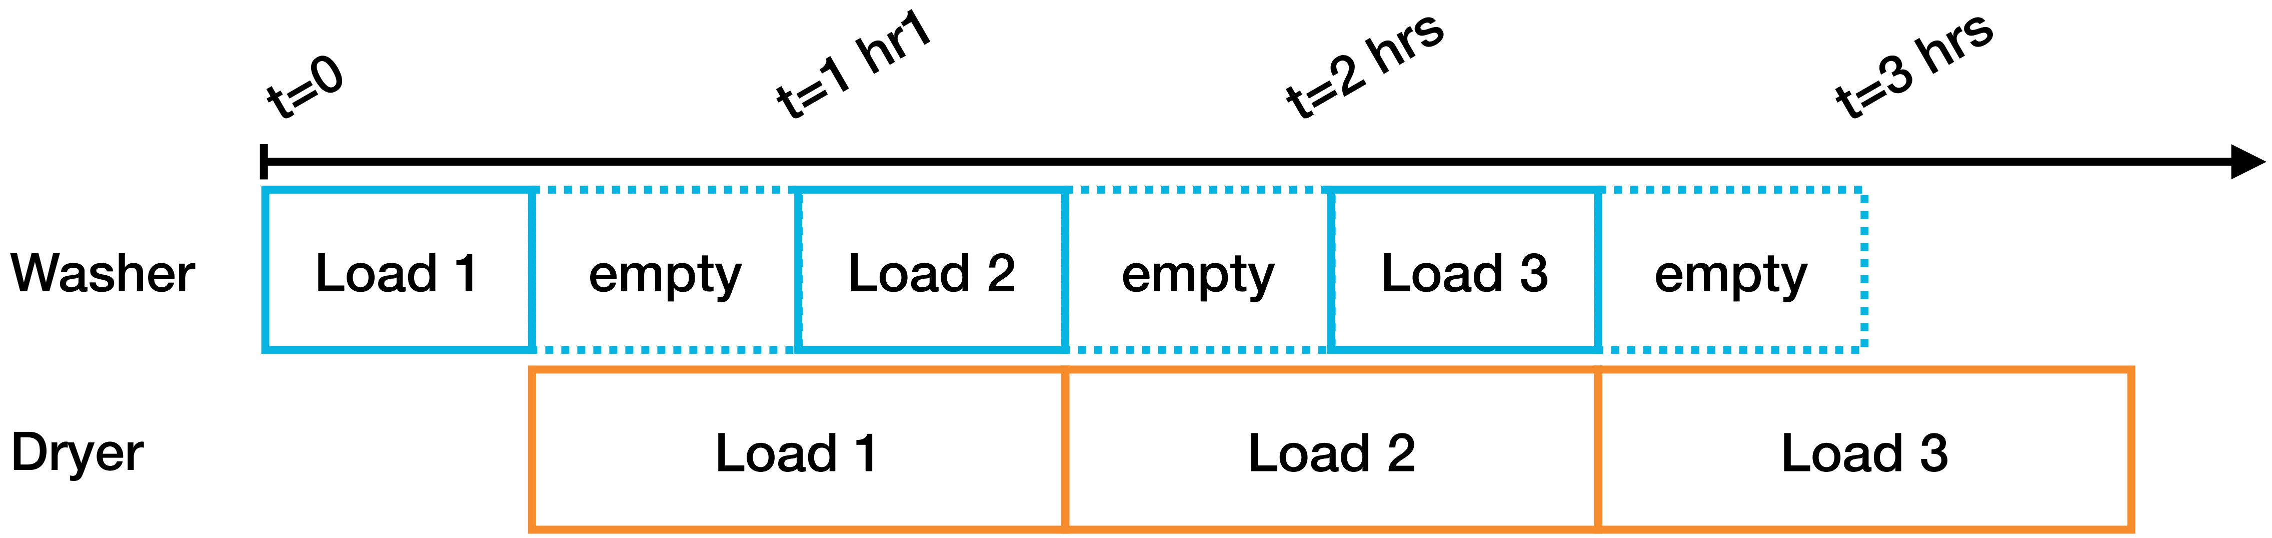

A real life example

Trigger systems rely on **pipelines** to perform *complicated reconstruction* tasks while handling a stream of *continuous inputs*.

They enable systems with *long latency*, *short initiation interval.*

Example: laundry in my apartment building:

1 washer (runs 30min) and 1 dryer (runs 1hr).

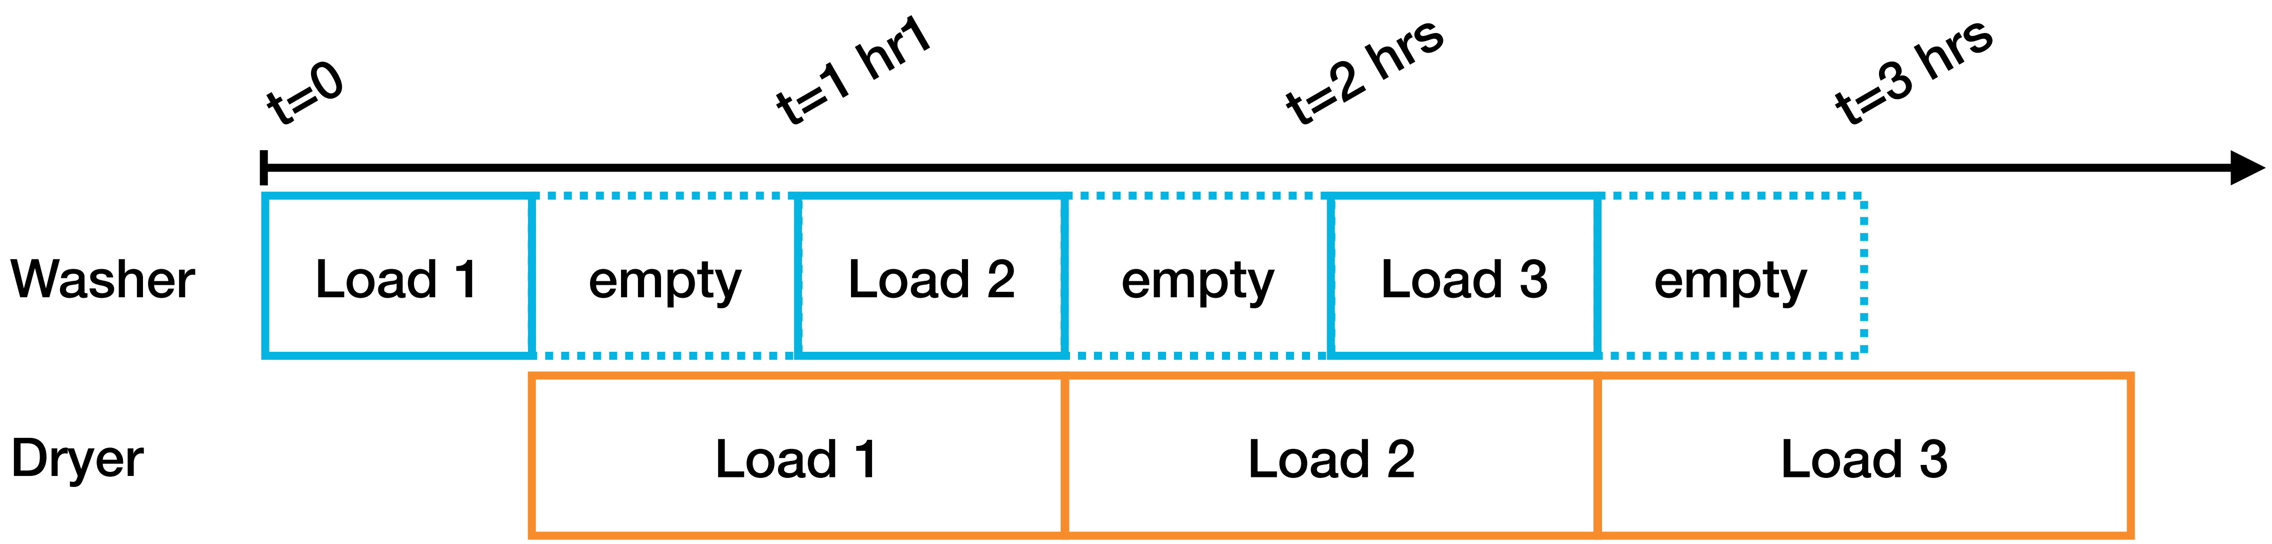

What are the ii, latency of the washer? The dryer? The entire system?

A real life example

Trigger systems rely on **pipelines** to perform *complicated reconstruction* tasks while handling a stream of *continuous inputs*.

They enable systems with *long latency*, *short initiation interval.*

Example: laundry in my apartment building:

1 washer (runs 30min) and 1 dryer (runs 1hr).

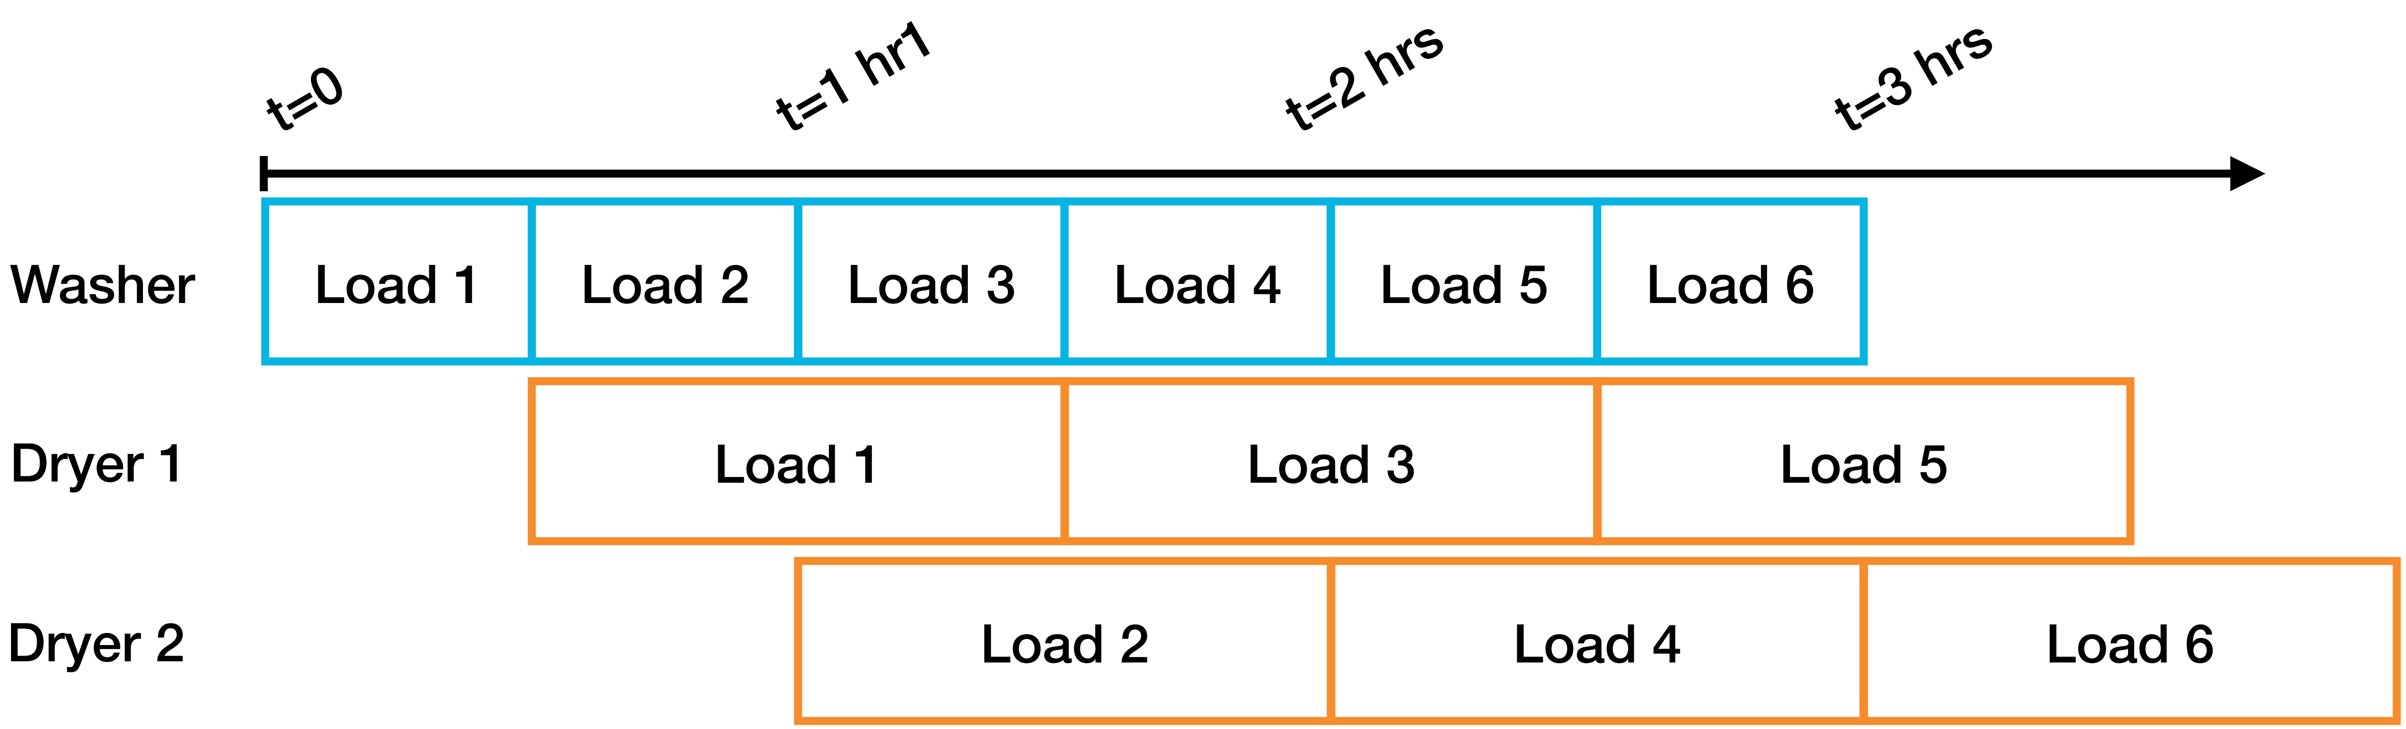

What are the ii, latency of the washer? The dryer? The entire system?

Pipelining

Trigger systems rely on **pipelines** to perform *complicated reconstruction* tasks while handling a stream of *continuous inputs*.

They enable systems with *long latency*, *short initiation interval.*

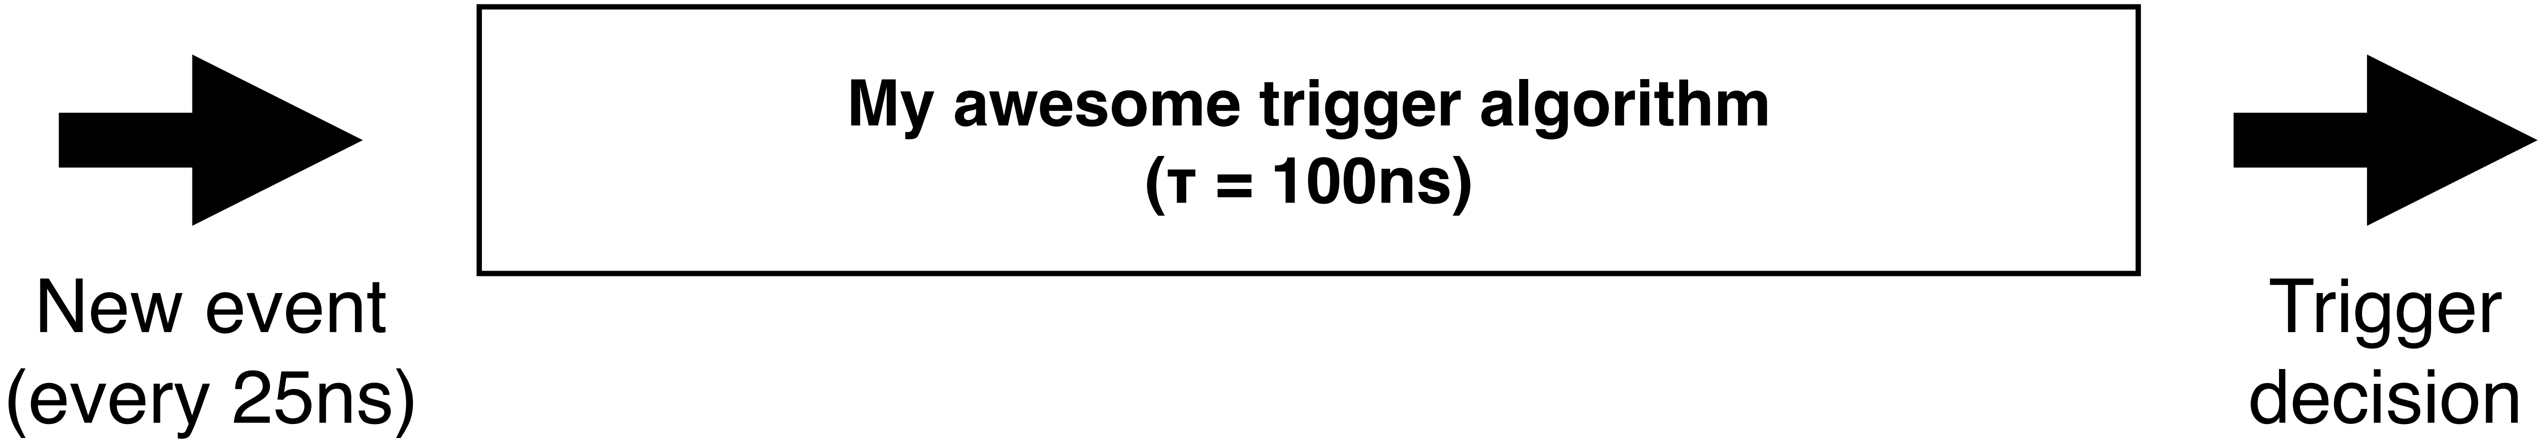

Pipelining

Trigger systems rely on **pipelines** to perform *complicated reconstruction* tasks while handling a stream of *continuous inputs*.

They enable systems with *long latency*, *short initiation interval.*

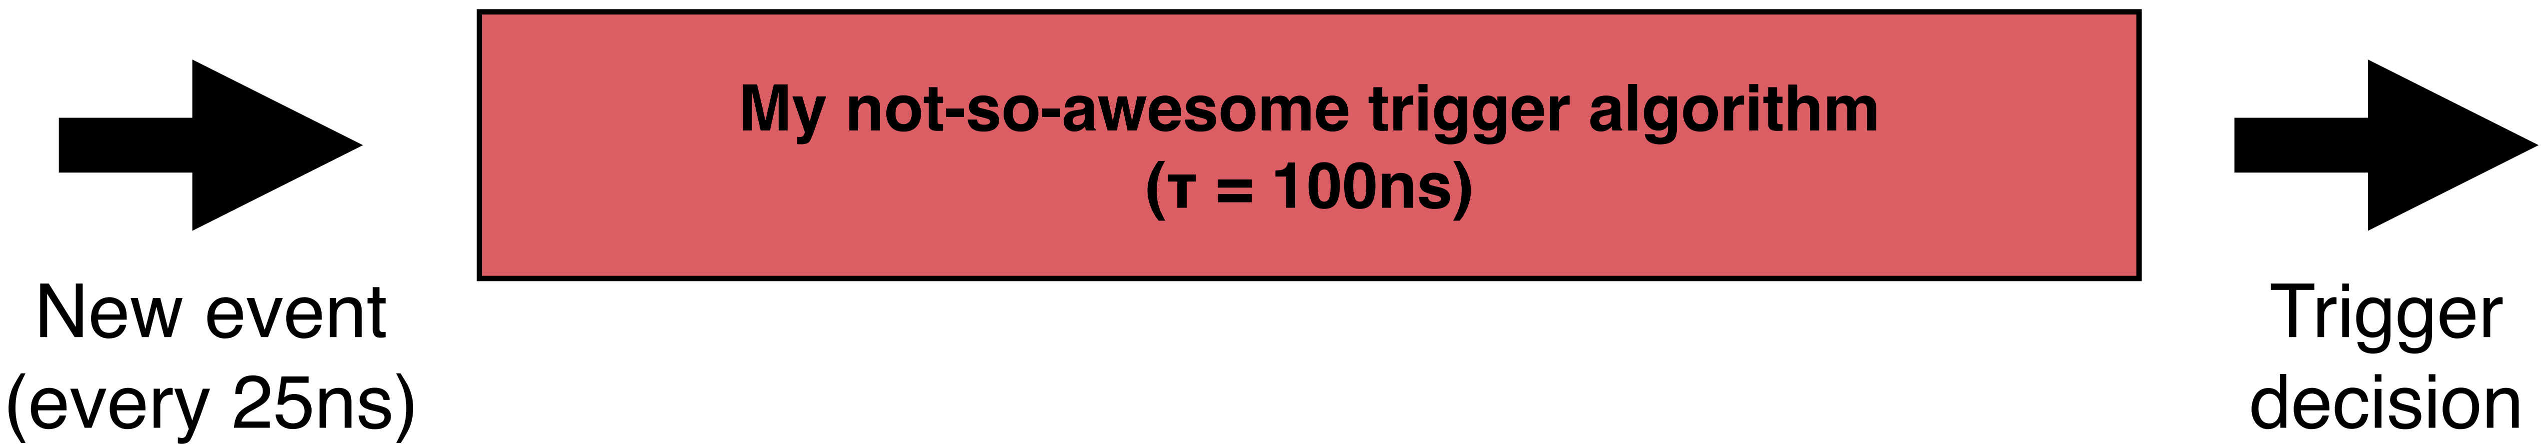

The block is still processing when a new event arrives!

Pipelining

Trigger systems rely on **pipelines** to perform *complicated reconstruction tasks* while handling a stream of *continuous inputs*.

They enable systems with *long latency*, *short initiation interval.*

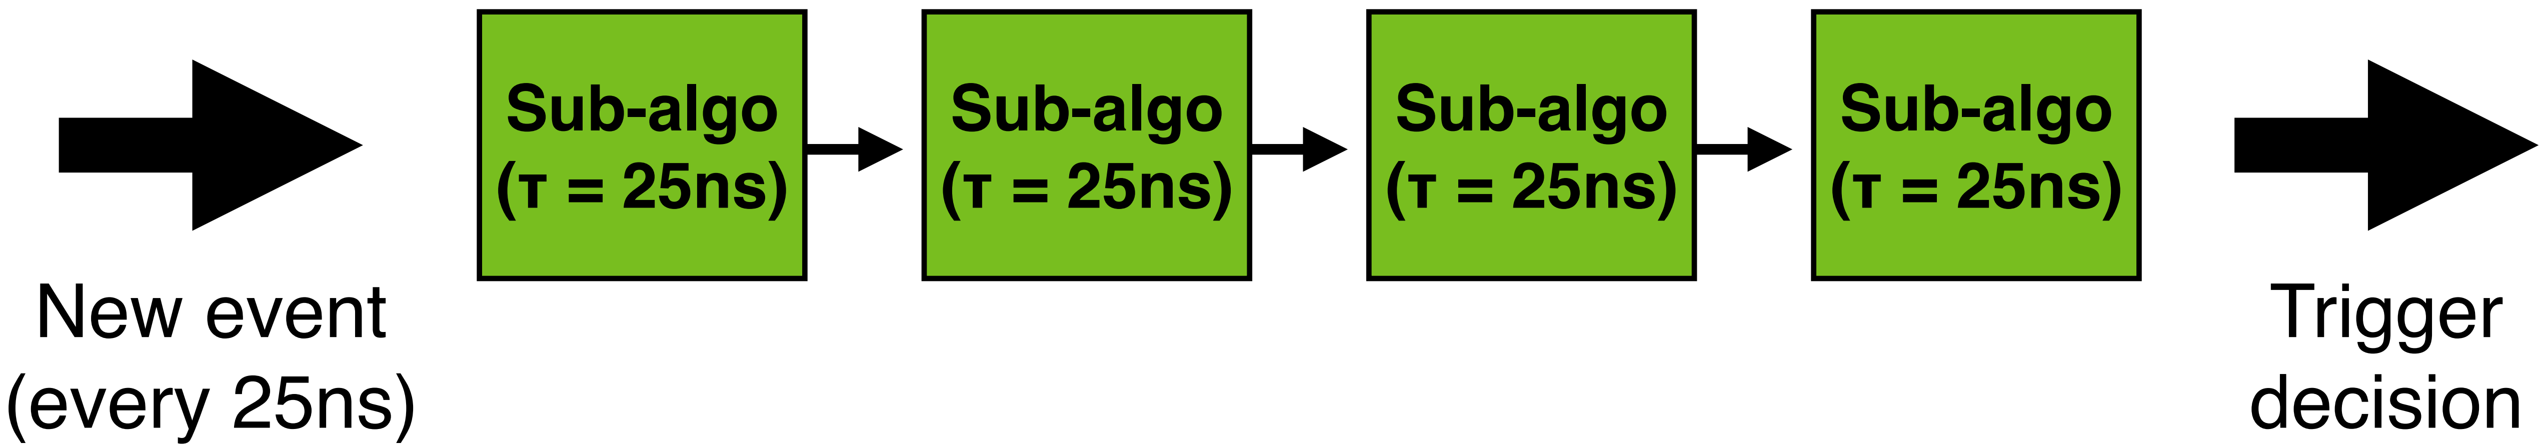

Modular logic can process many events at once

(and again, intermediate buffers ease synchronization)

Pipelining

Trigger systems rely on **pipelines** to perform *complicated reconstruction* tasks while handling a stream of *continuous inputs*.

They enable systems with *long latency*, *short initiation interval.*

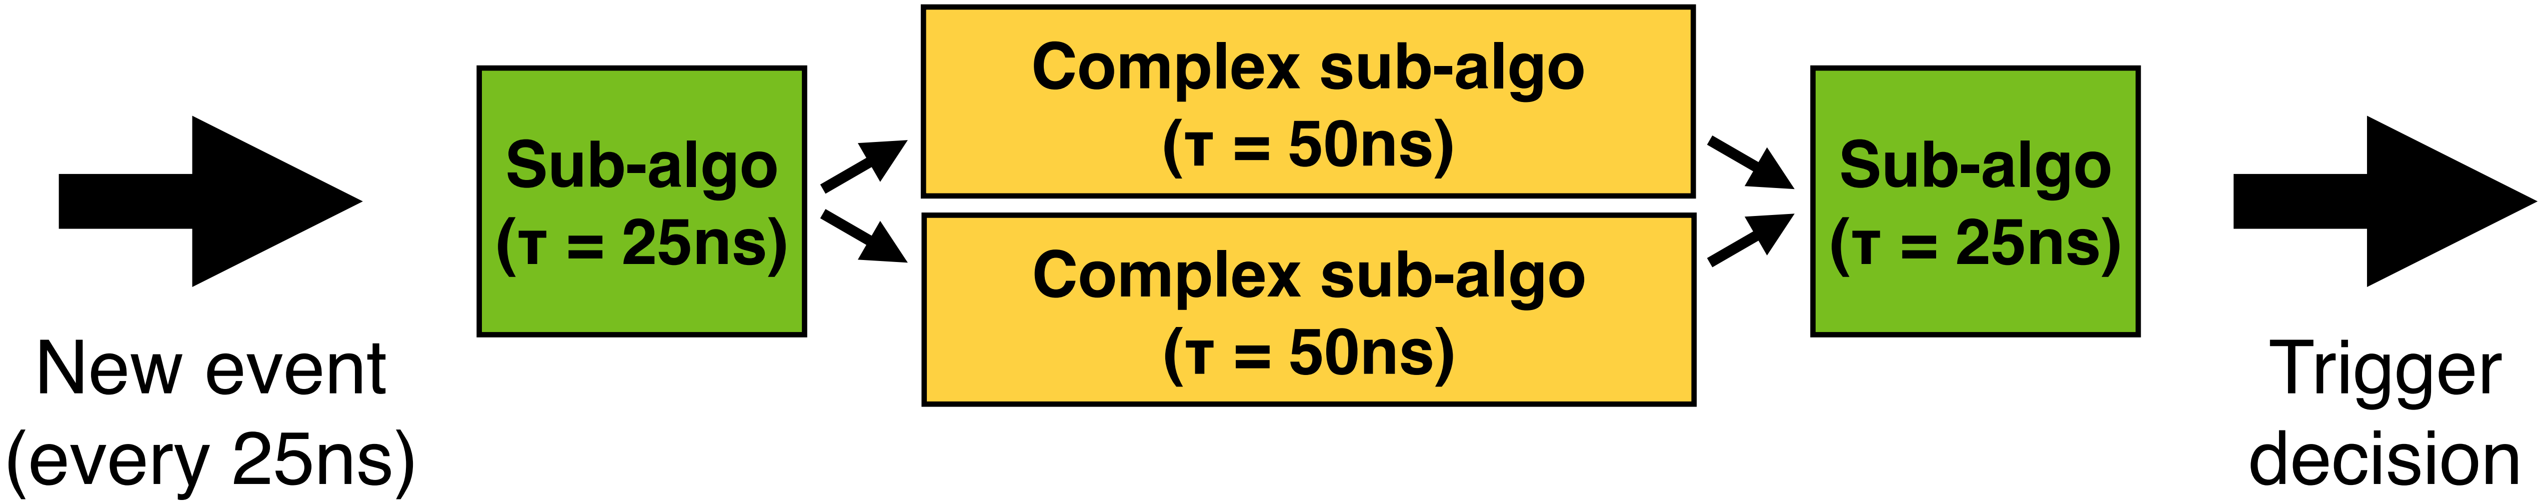

Sometimes this is inconvenient or impossible → parallel processors

(More complex building blocks, at the expense of more resources)

Regional processing

Particle reconstruction is **an** **inherently local** **task**
→ process parallel regions

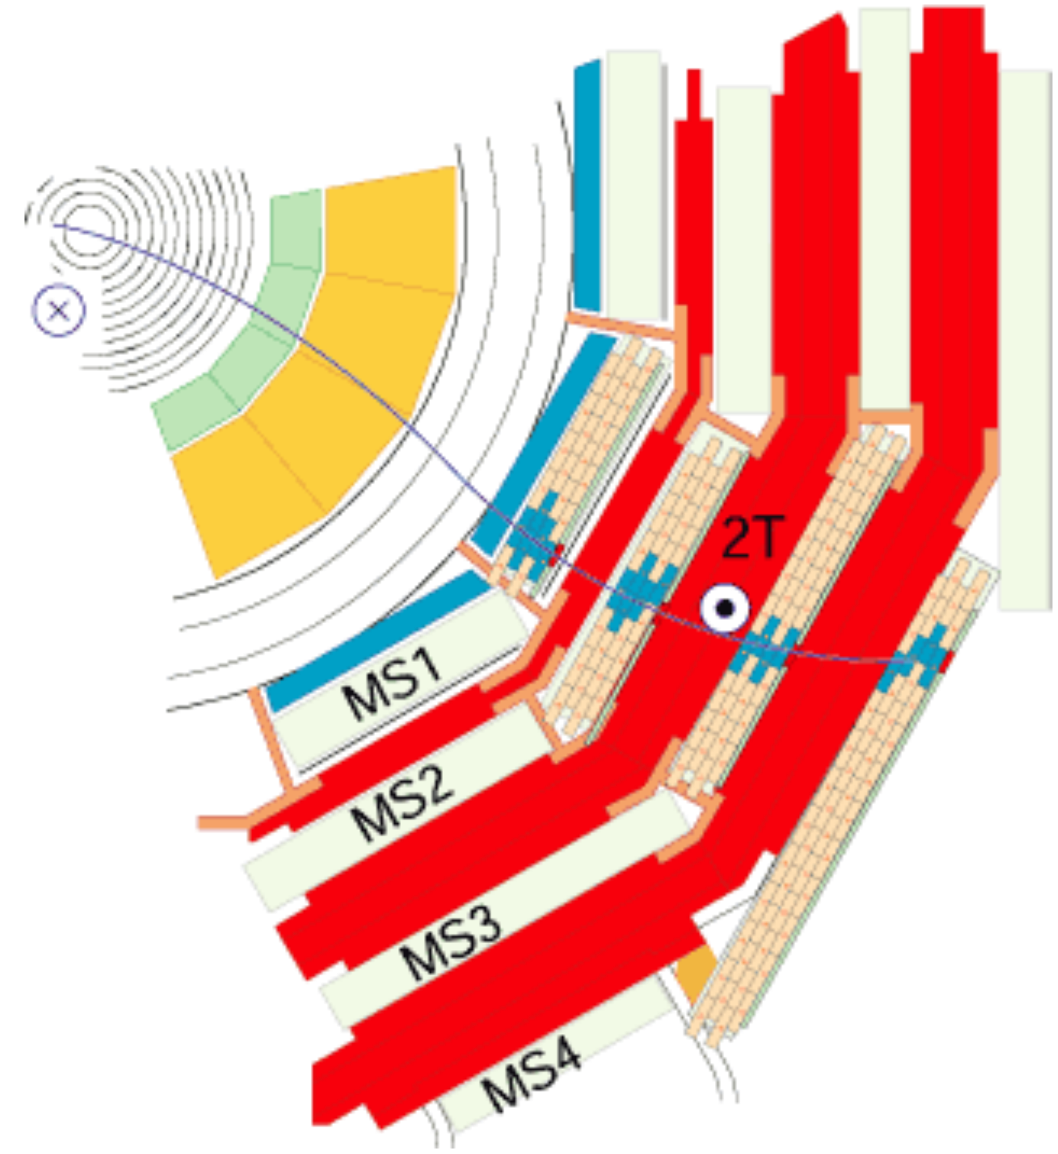

Regional processing

Particle reconstruction is **an** **inherently local** **task**
→ process parallel regions

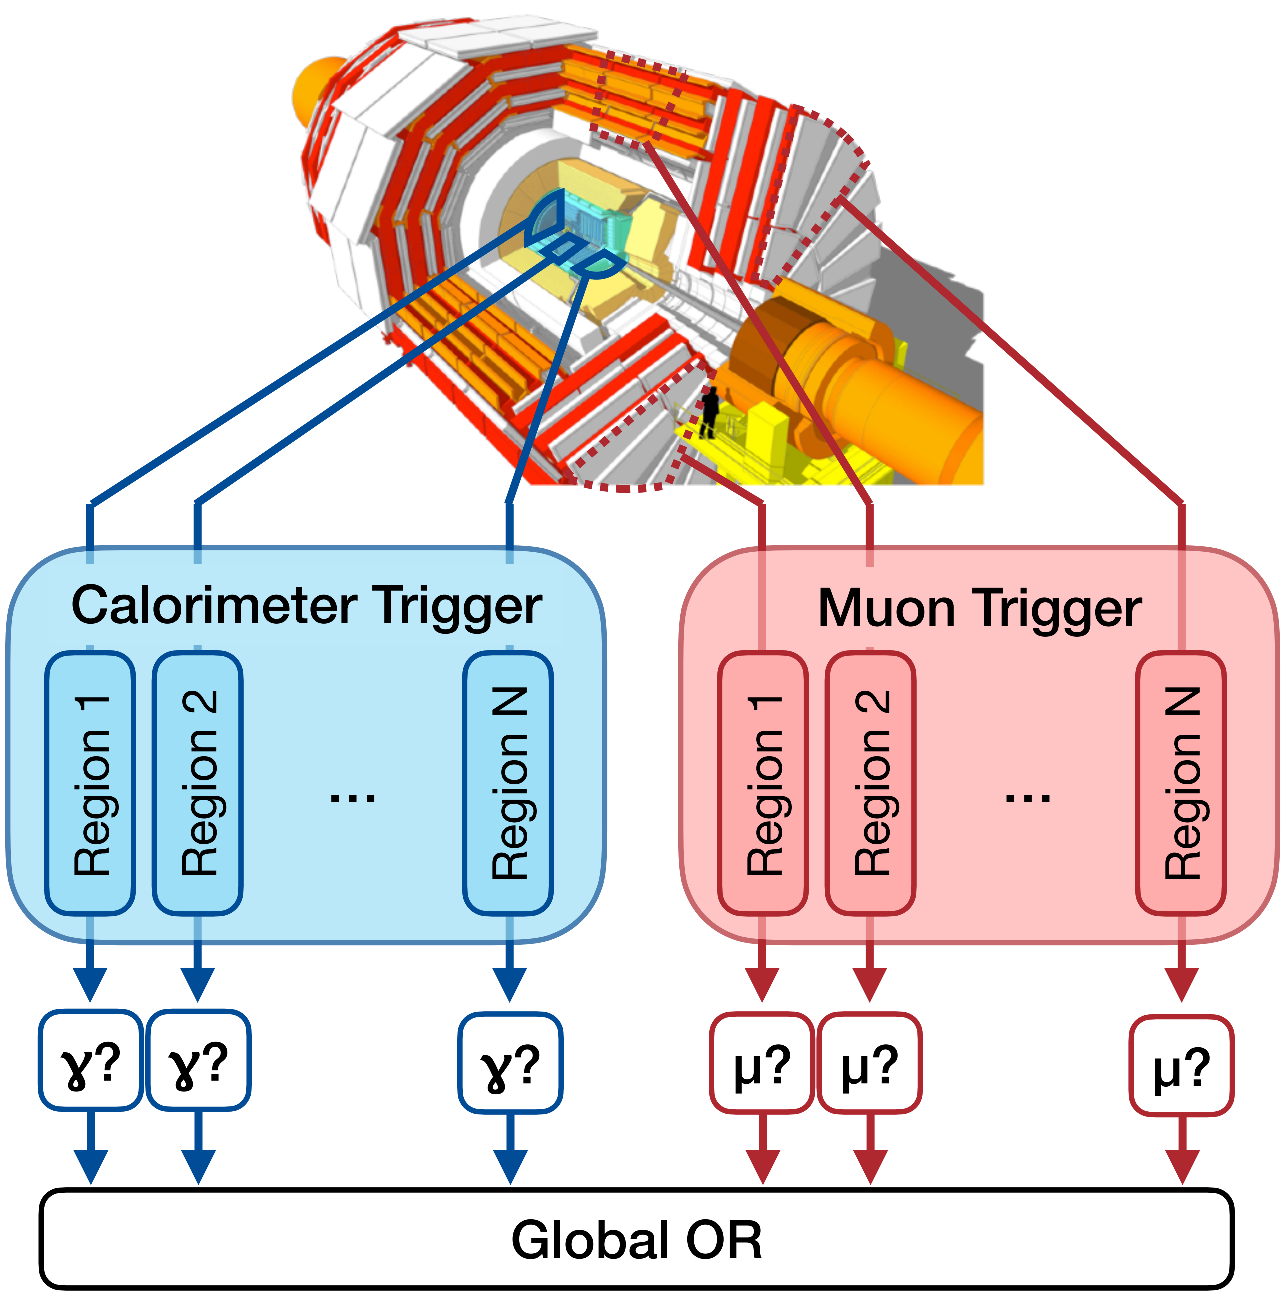

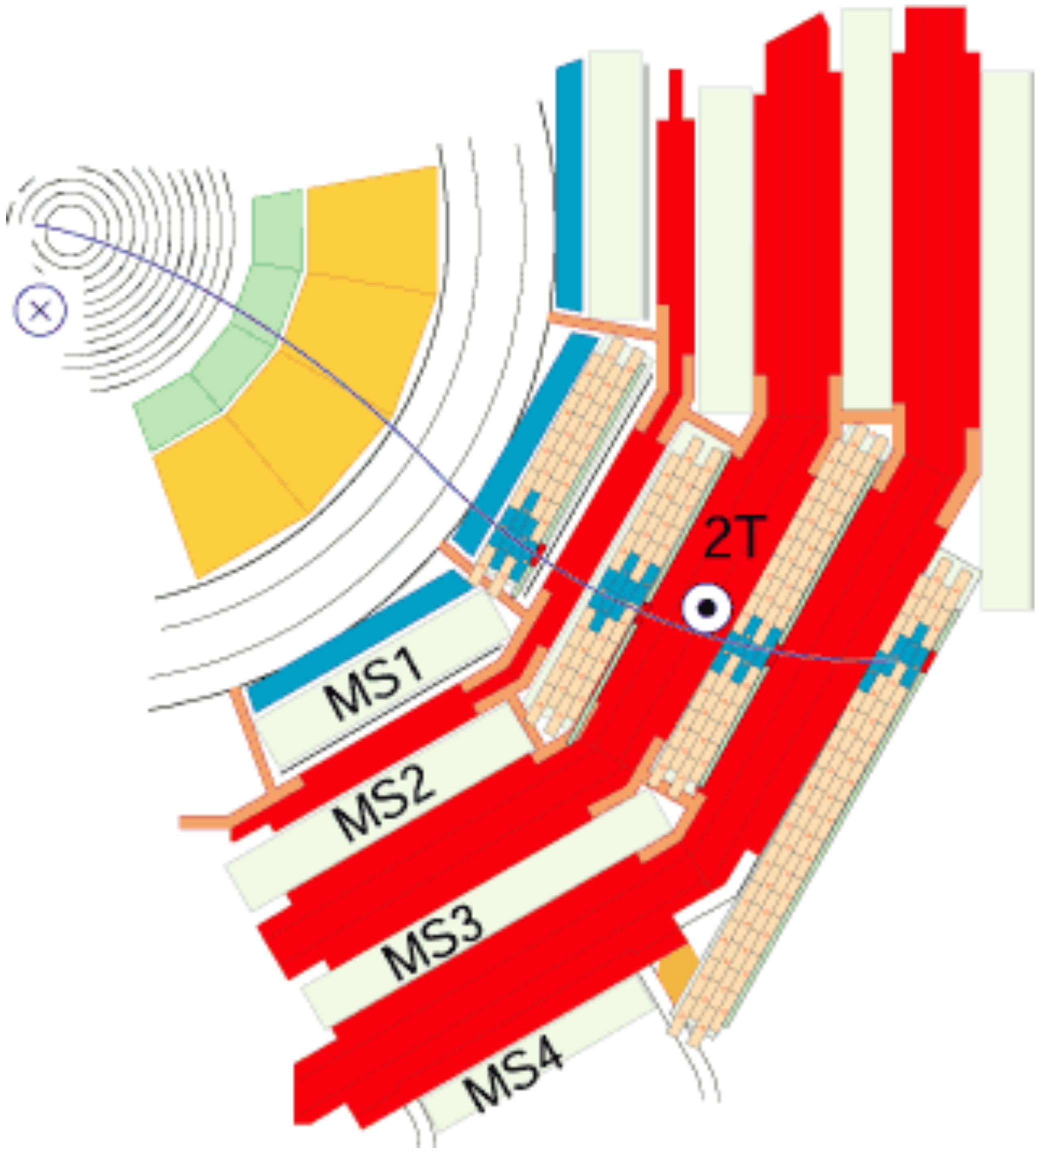

No “event building” necessary!

Regional processing

Particle reconstruction is **an inherently local task**
→ process parallel regions

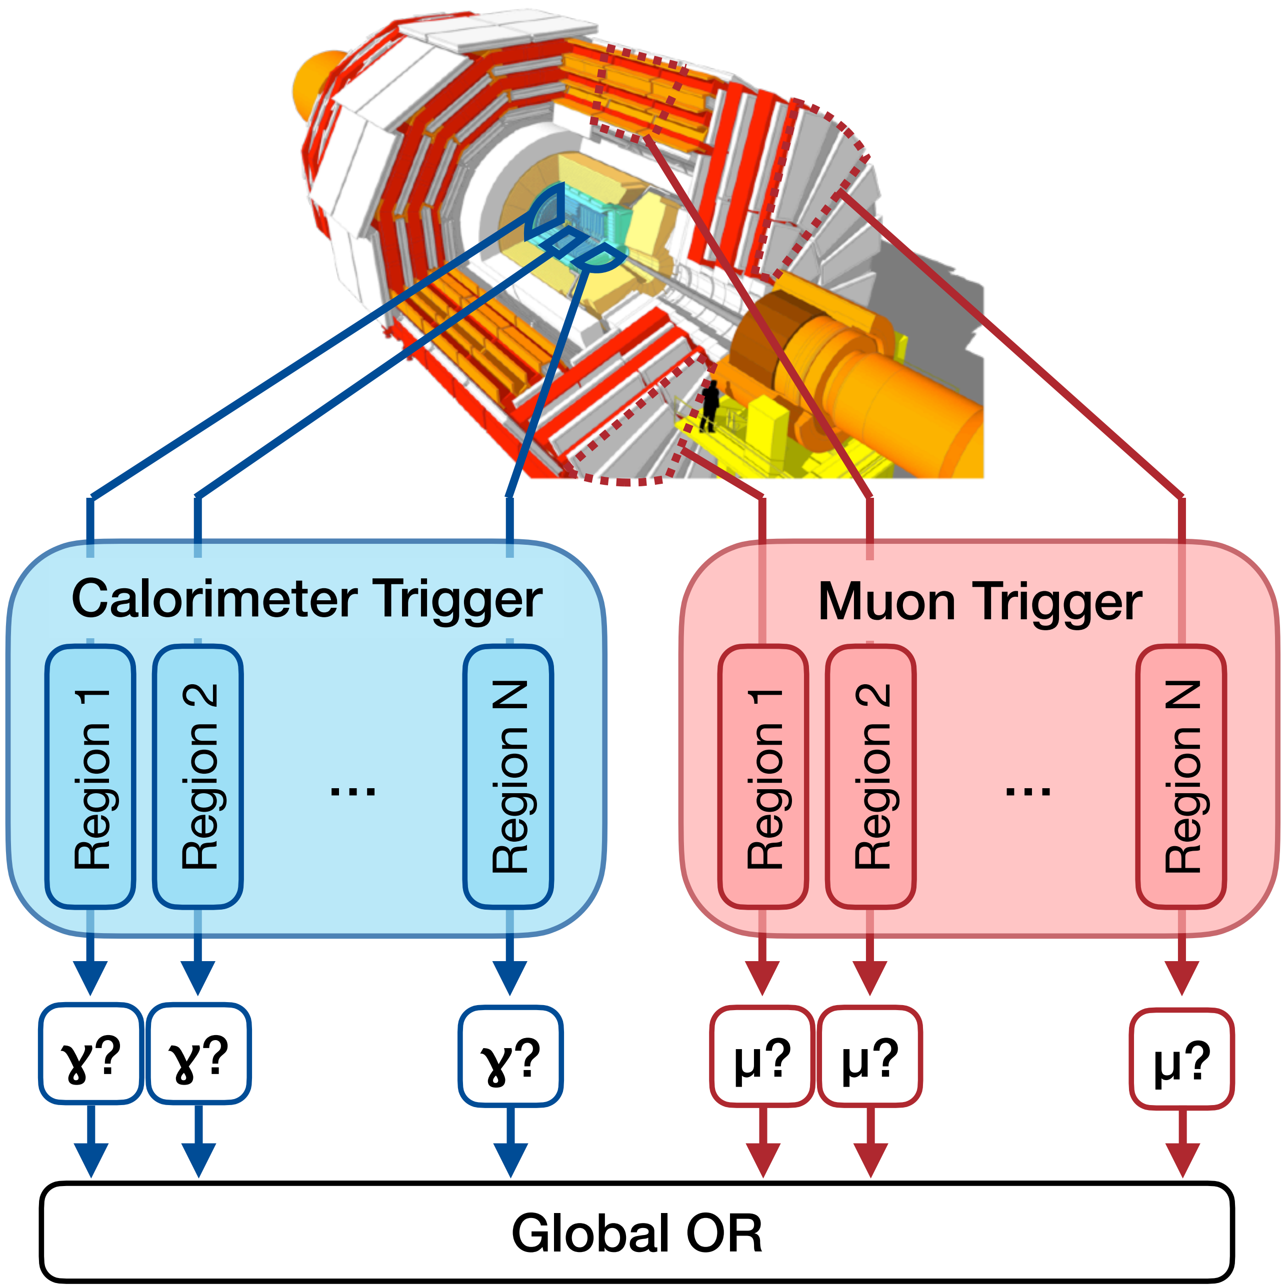

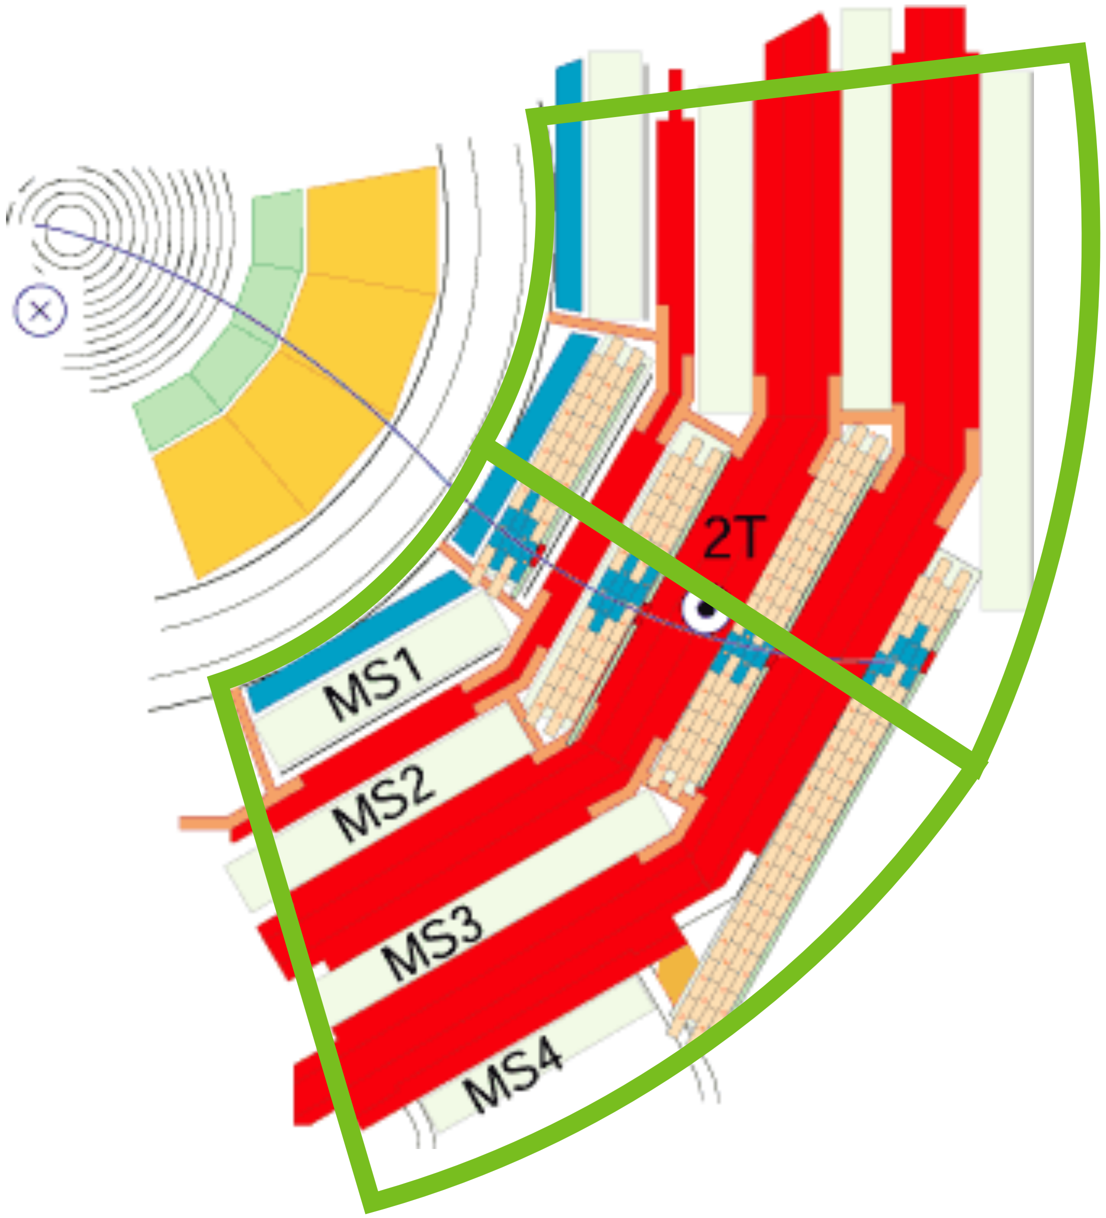

Caveat: some data-sharing always needed. How to deal with overlaps?

Regional processing

Particle reconstruction is **an inherently local task**
→ process parallel regions

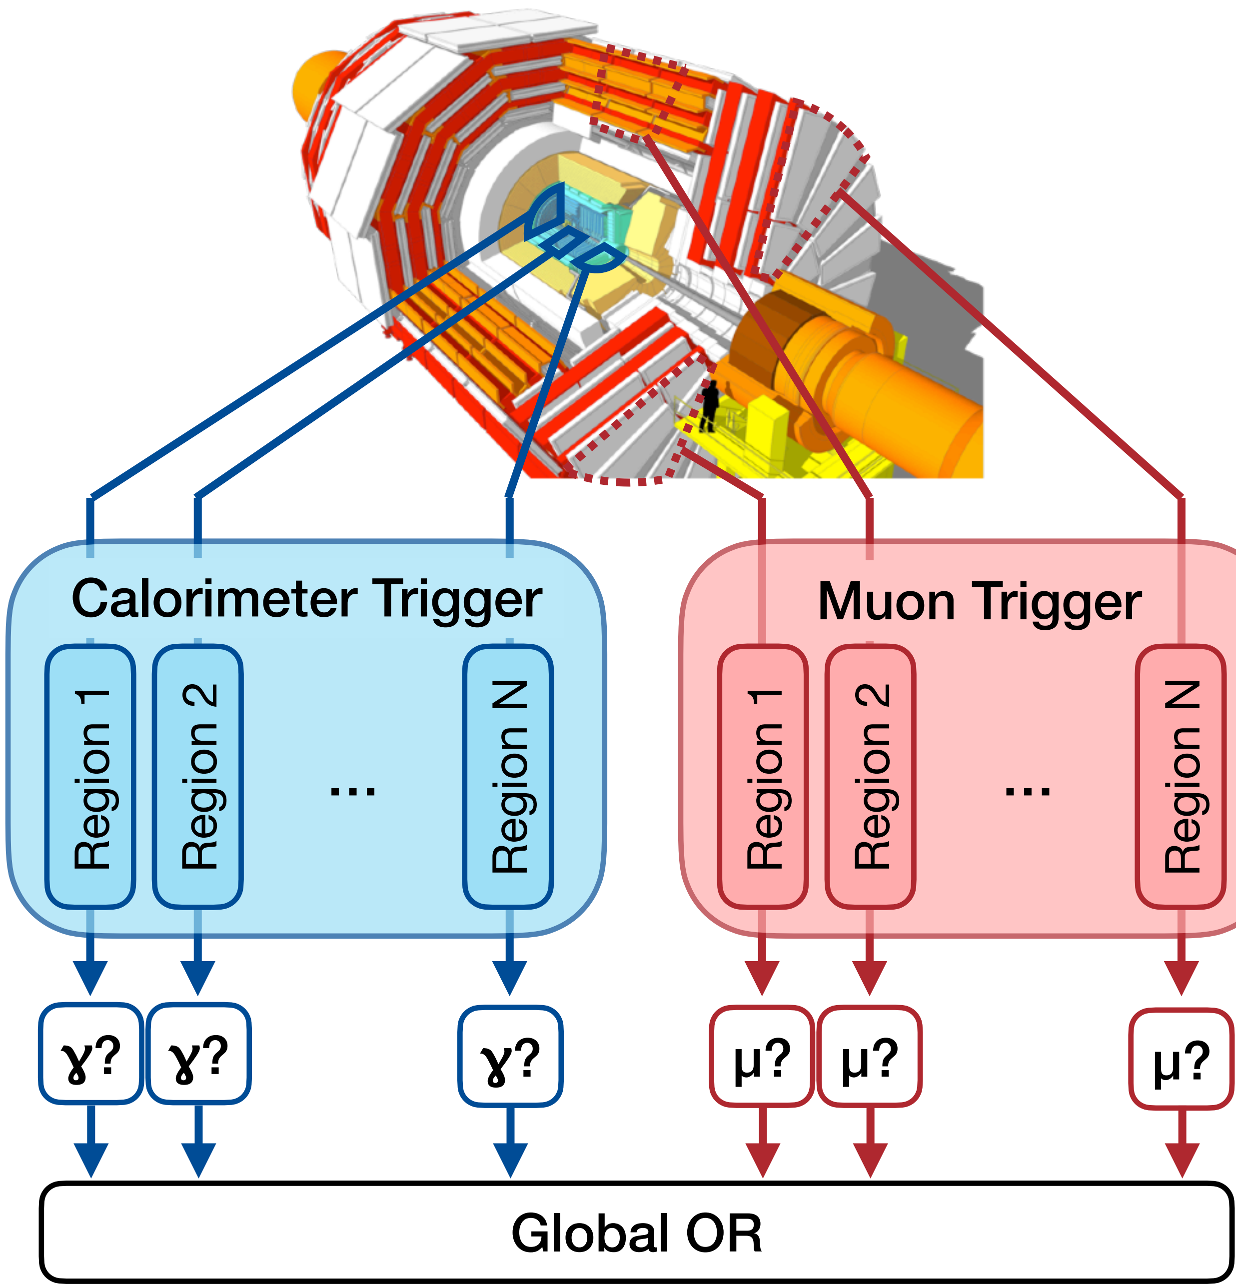

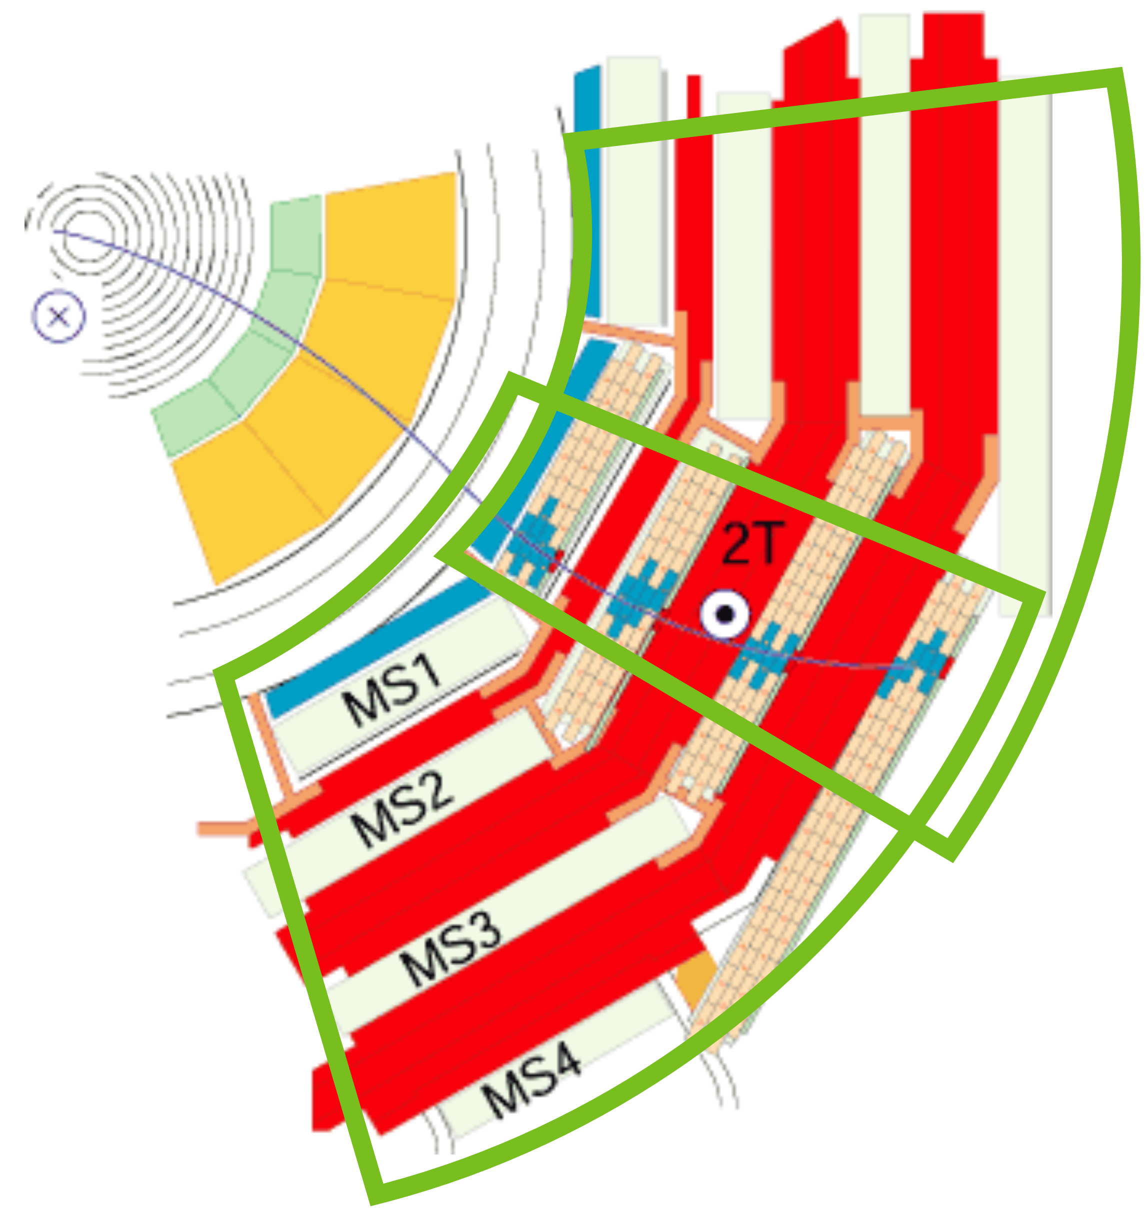

Caveat: some data-sharing always needed. How to deal with overlaps?

Regional processing

Particle reconstruction is an inherently local task
→ process parallel regions

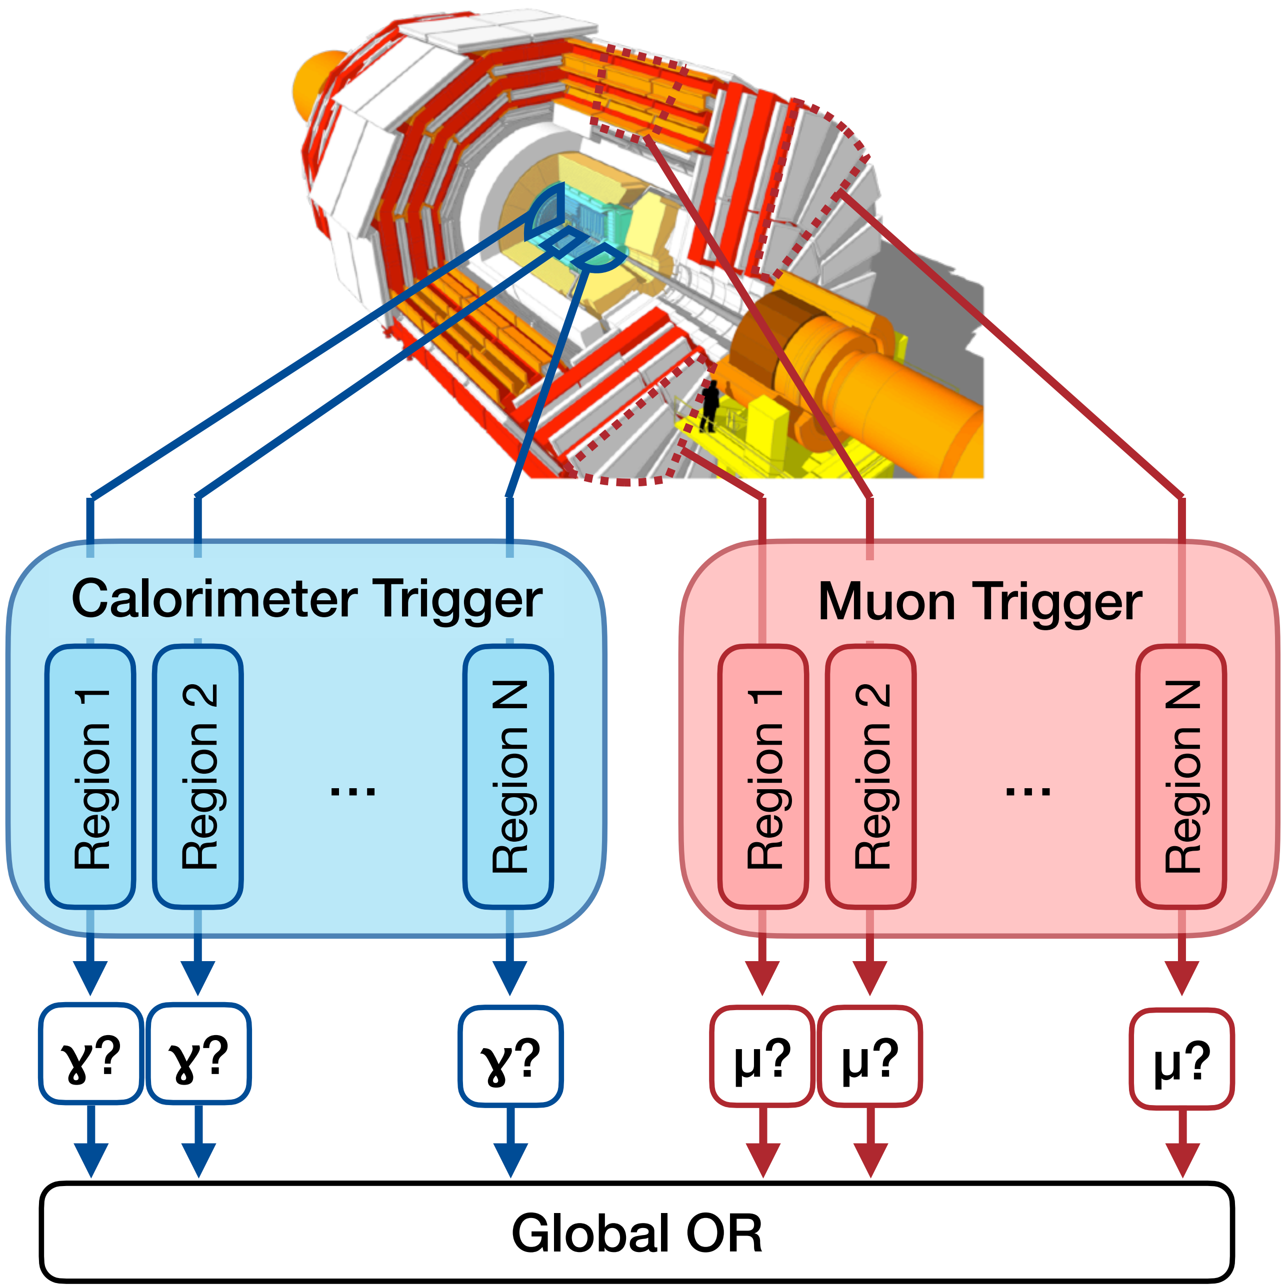

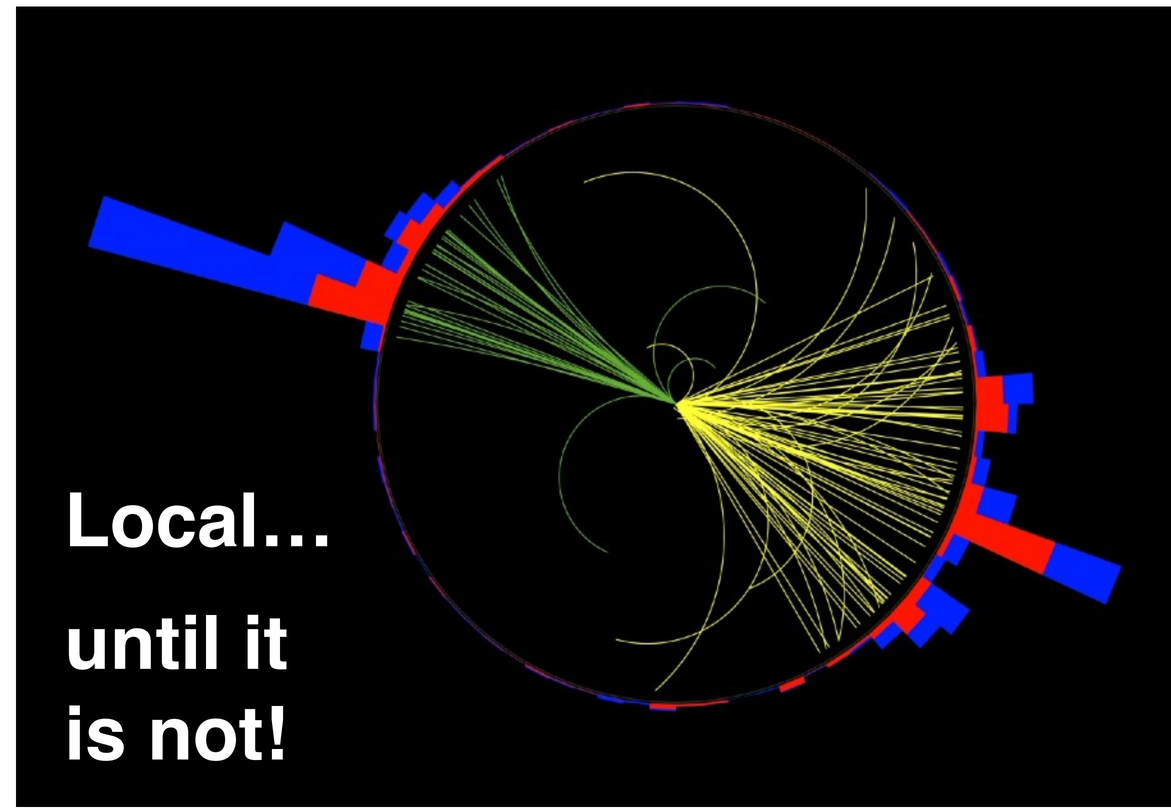

Particle jets occupy 1/2 detector!

E.g. a decaying top quark with large Lorentz boost

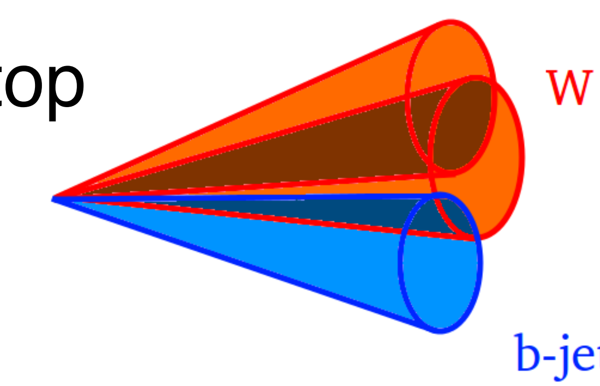

Regional processing

Aside: data-sharing a **far larger concern** for detectors near the collision point.

Example: find charged particles in a **patch of angular coordinates**.

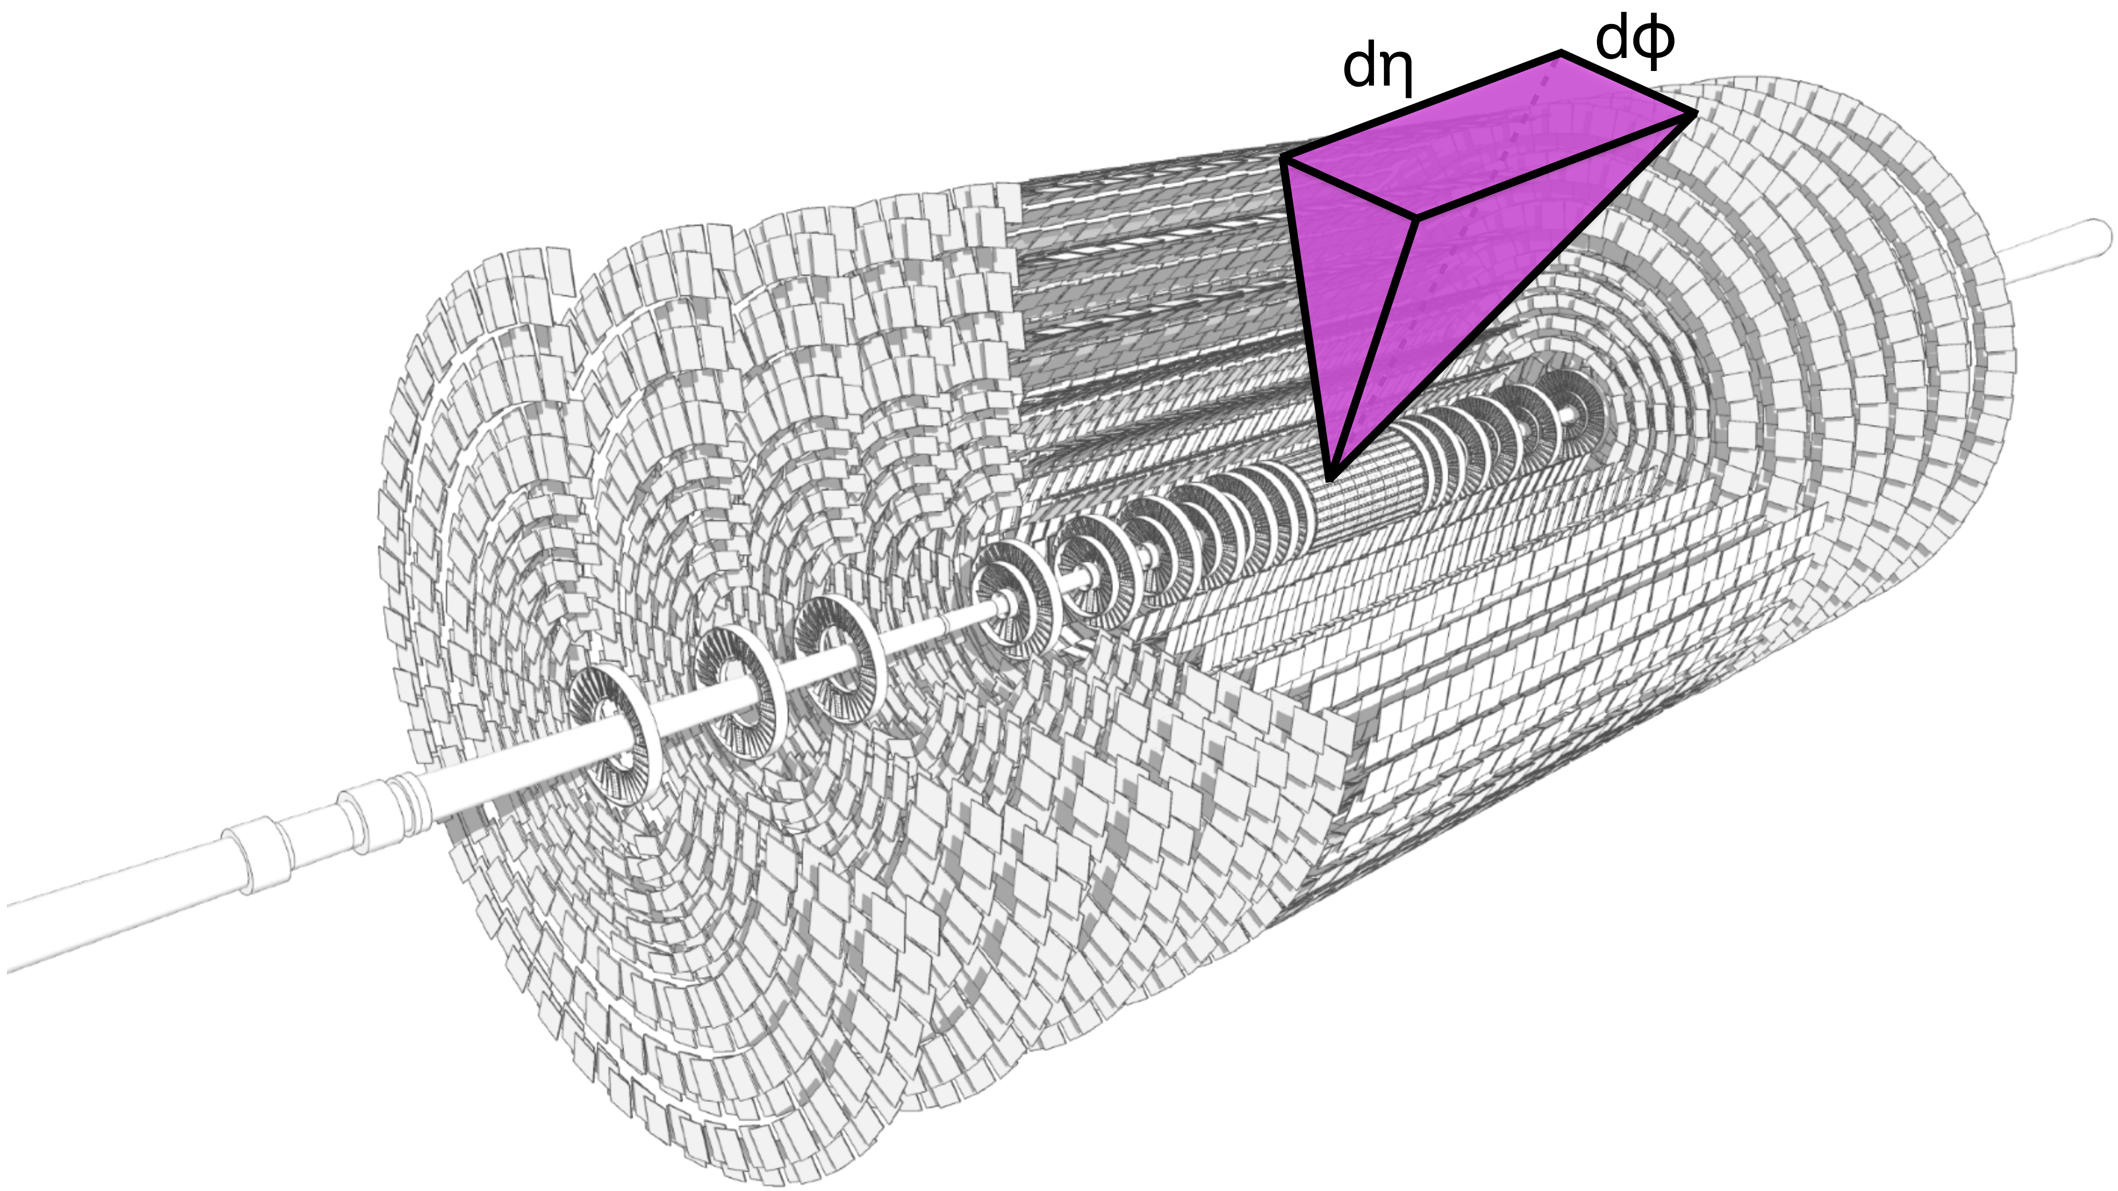

Regional processing

Aside: data-sharing a **far larger concern** for detectors near the collision point.

Example: find charged particles in a **patch of angular coordinates**.

Total data required to find tracks can be several times the **fiducial region**!

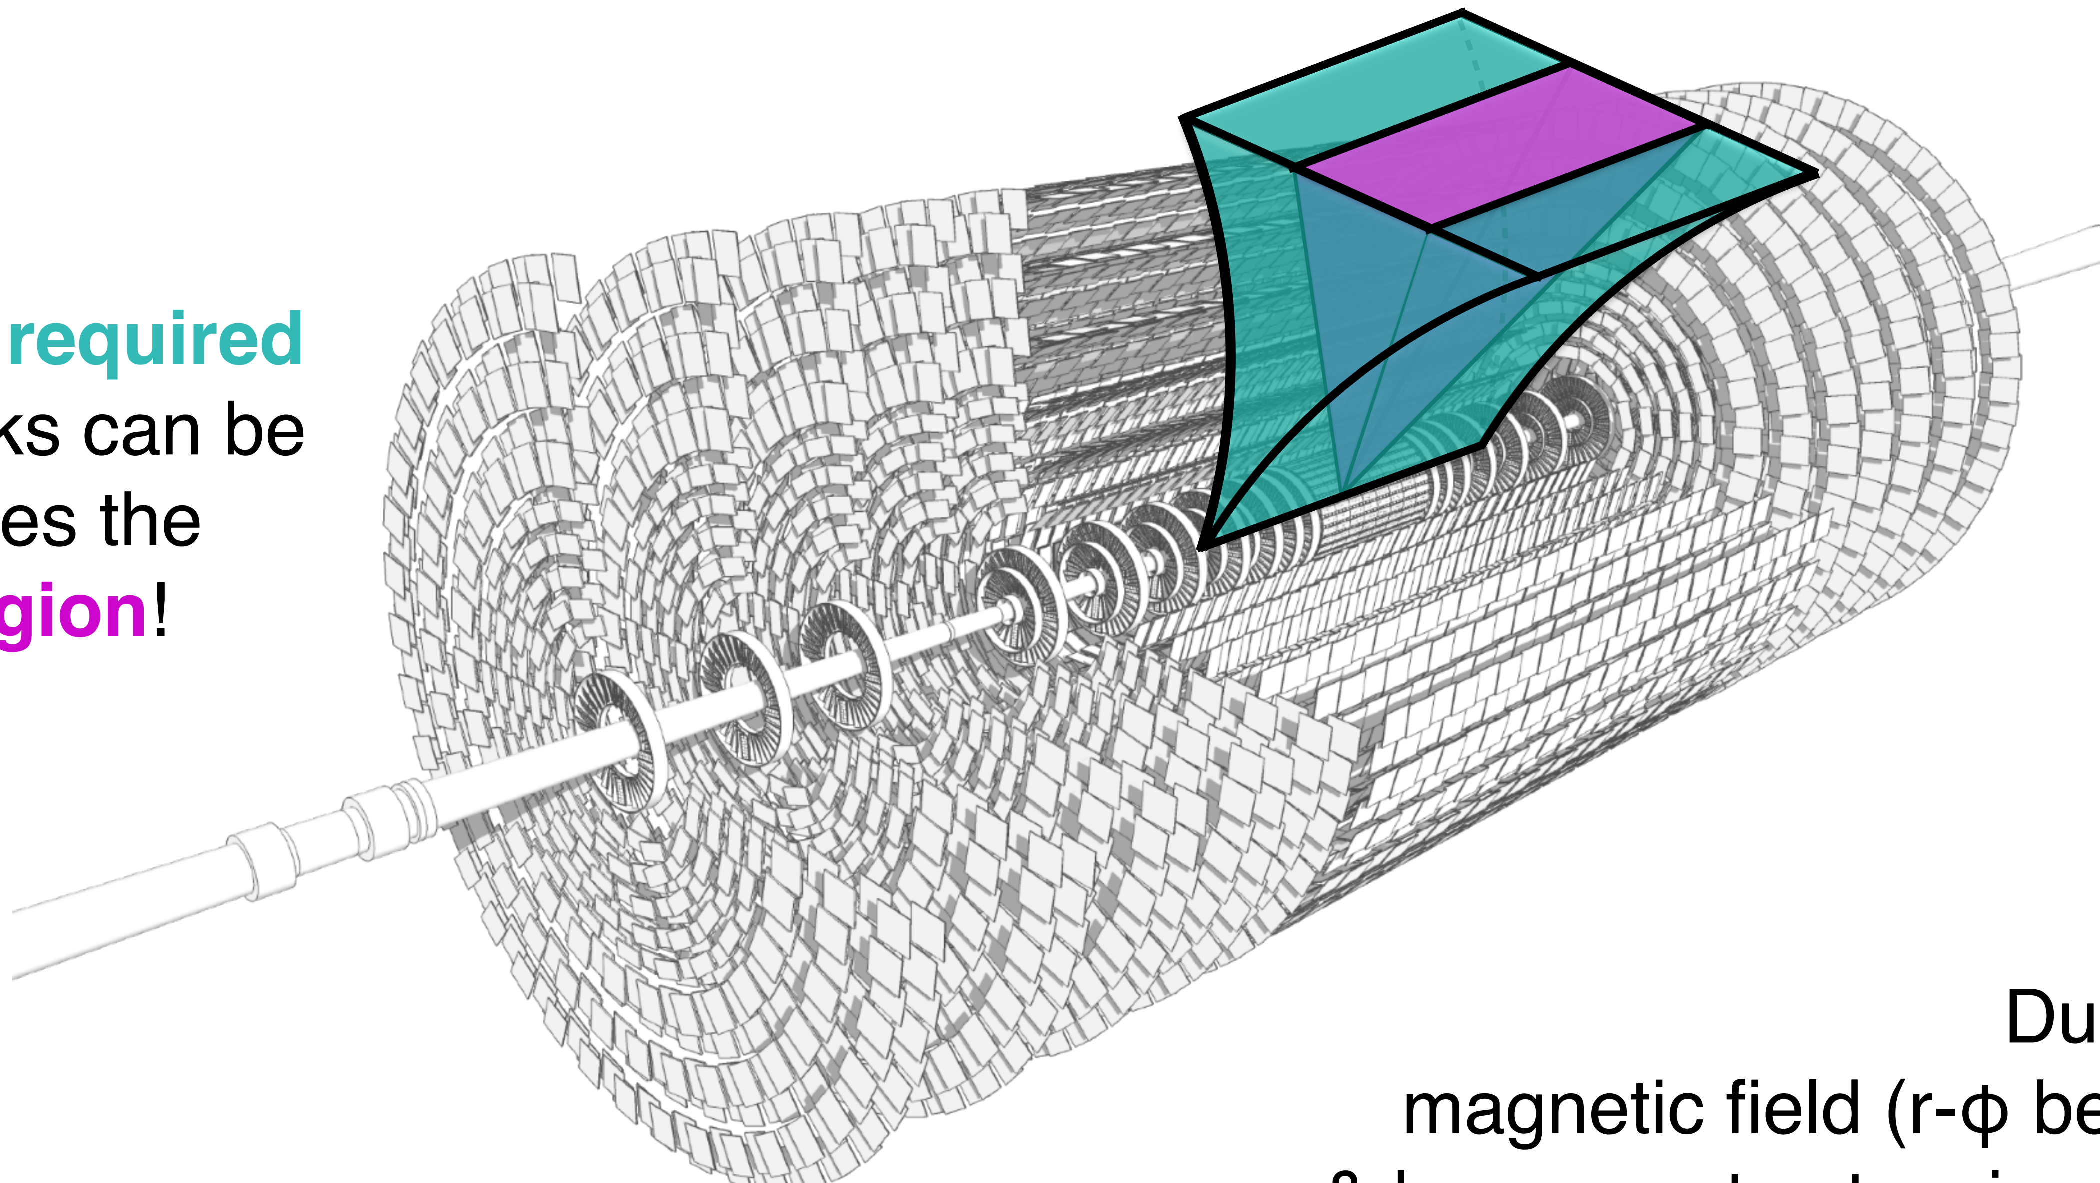

Due to magnetic field (r-ϕ bend) & beam-spot extension in z

Plan

- **This afternoon: Concepts**
  - Introduction and physics goals of the LHC Experiments
  - The role of the trigger and DAQ systems
  - DAQ concepts: simple toy model → complex systems
  - Trigger concepts: from hardware to architecting a system
- **Tomorrow: Details and Implementations**
  - **Some context: challenges of high luminosity & the upgrades**
  - ATLAS, CMS hardware triggers, and their evolution
  - Software triggers, and new paradigms for analysis
  - Trigger menus, and you
  - Conclusions, and looking farther ahead

Break