Summer Student
Lecture Program 2022

Electronics, Trigger and Data Acquisition. 3/3

E. Meschi – CERN EP Dept. – CMS Experiment

- Trigger and DAQ system concepts
- From signal to physics through examples
- Timing
- Data transport, links, buses
- Queues and Event building
- On-line data processing

Credits:
SSLP ETD lecture series by R.Ferrari (2019)
EM: lectures on DAQ/Trigger at U.Padua 2018-2020
ISOTDAQ: International School of Trigger and DAQ
https://indico.cern.ch/event/928767/
Material from various papers and books (bibliography at the end)

- We have our ADC/TDC counts, but how do we correlate them with energy/position ?
- Need a calibration procedure

Calibration

- The experiments we discussed provide relative measurements. The values obtained via our system are in some (known) relation with the interesting quantity
  - Scintillator

$$Q \propto N_\gamma \propto E$$

  - XDWC

$$y = \alpha \cdot \Delta t + \beta = \alpha \cdot (t_{top} - t_{bottom}) + \beta$$

- Our instruments need to be in order to give us the answer we are looking for
  - We have to determine the **calibrated**parameters that transform the raw data into a physics quantity
  - The parameters normally depend on the experimental setup (e.g. cable length, delay settings, HV settings, …)
- In the design of our detector and DAQ we have to foresee calibration mechanisms/procedures

Ge crystal for isotope identification

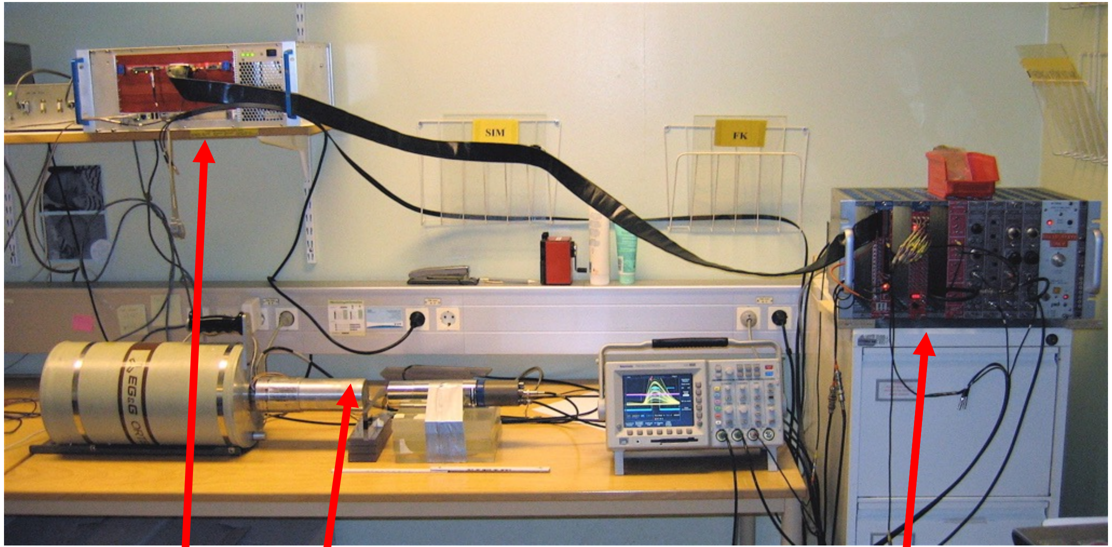

Crystal HPGe

Trigger & front-end

Readout (ADC)

Crystal calibration

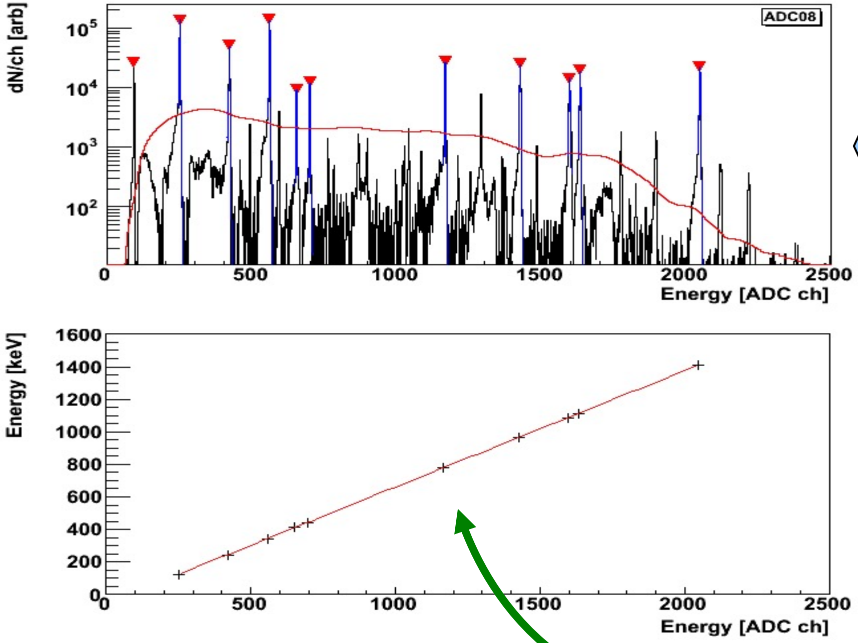

- ^{152}Europium reference source
  - Known emission lines
  - Peak find & fit
- Allow for definition of the parameters describing functional relation between ADC count and E

$$Q \propto N_\gamma \propto E$$

- Reality is not perfectly linear
  - Polynomial fits
  - Physical models

Isotope identification

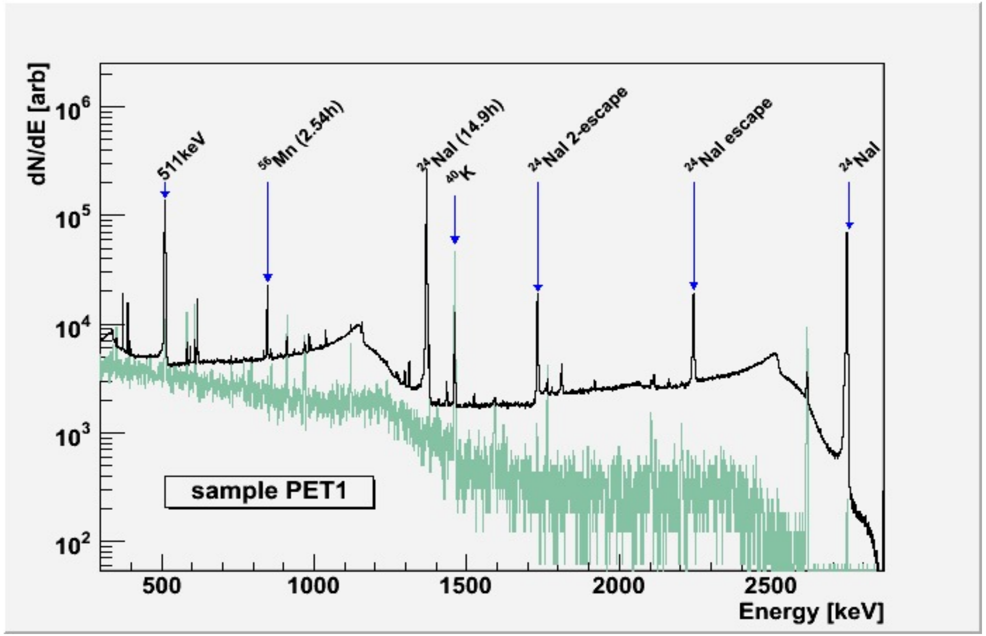

Calibrated crystal setup can be used to identify isotopes generated in irradiated samples

XDWC Calibration

- XDWC chamber have 3 calibration inputs that allow for independent calibrations of X and Y axis with only 3 different sets of data
- The calibration input simulate signals from particles respectively hitting
  - Right-top corner (X=Y=30mm)
  - Center (X=Y=0mm)
  - Left-bottom corner (X=Y=-30mm)
- The calibration data sets are collected with final setup and TDC
- Interpolating the three points in the *t*-*x* space, the parameters of the calibration equation can be measured

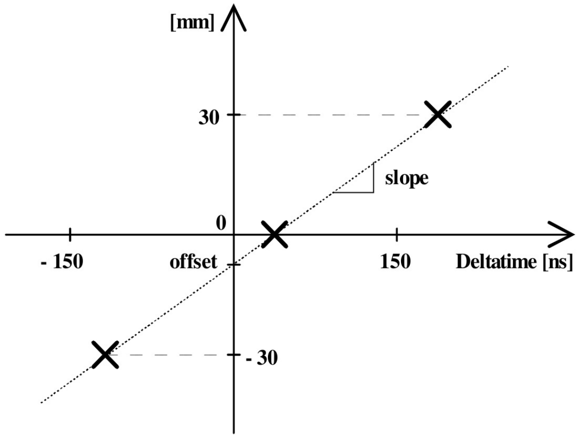

$$x = \alpha t^* + \beta$$

Calibrated XDWC

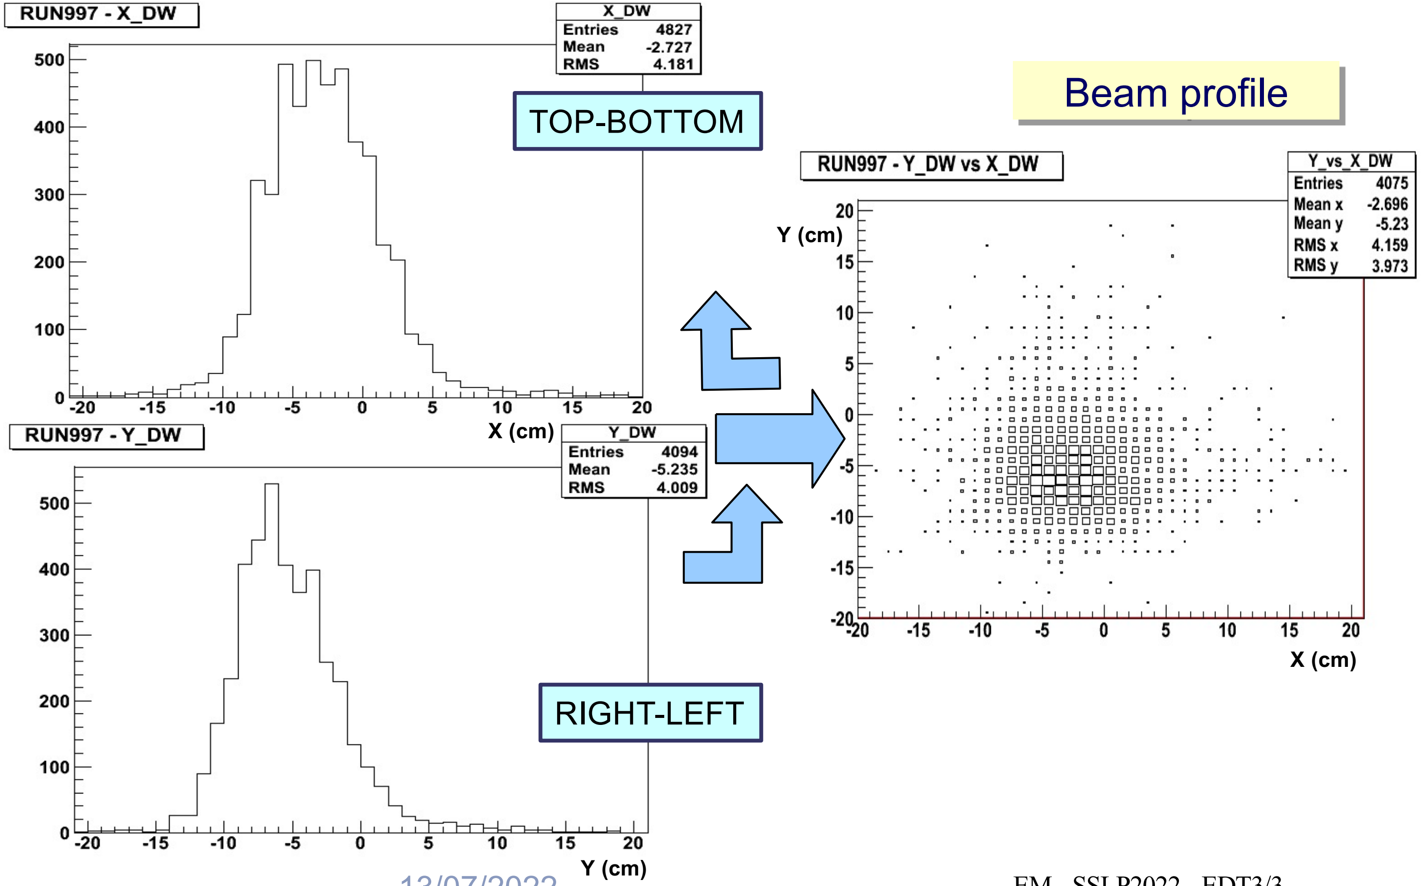

Calibration: summary

- Our DAQ chain provides us with numbers (raw data) but … what do they really mean ?
- Likely related with physics quantities but with relations (transfer functions) affected by several uncertainties:
- due to physical detection mechanisms
- due to signal processing
- Transfer functions usually parameterised, sometimes based on (look-up) numeric tables
- All system elements need to be calibrated to keep optimal knowledge of all parameters:
  - calibration procedures
  - calibration constants
- Calibration constants change with ageing (mainly due to radiation), beam conditions (electronics may have baseline drifting with pile-up), time … HV, LV, ...
- <u>The design of our detector and DAQ</u> has to foresee calibration mechanisms/procedures
- injection of known signals
- dedicated calibration *triggers* and data streams

Atlas tile calorimeter calibration system

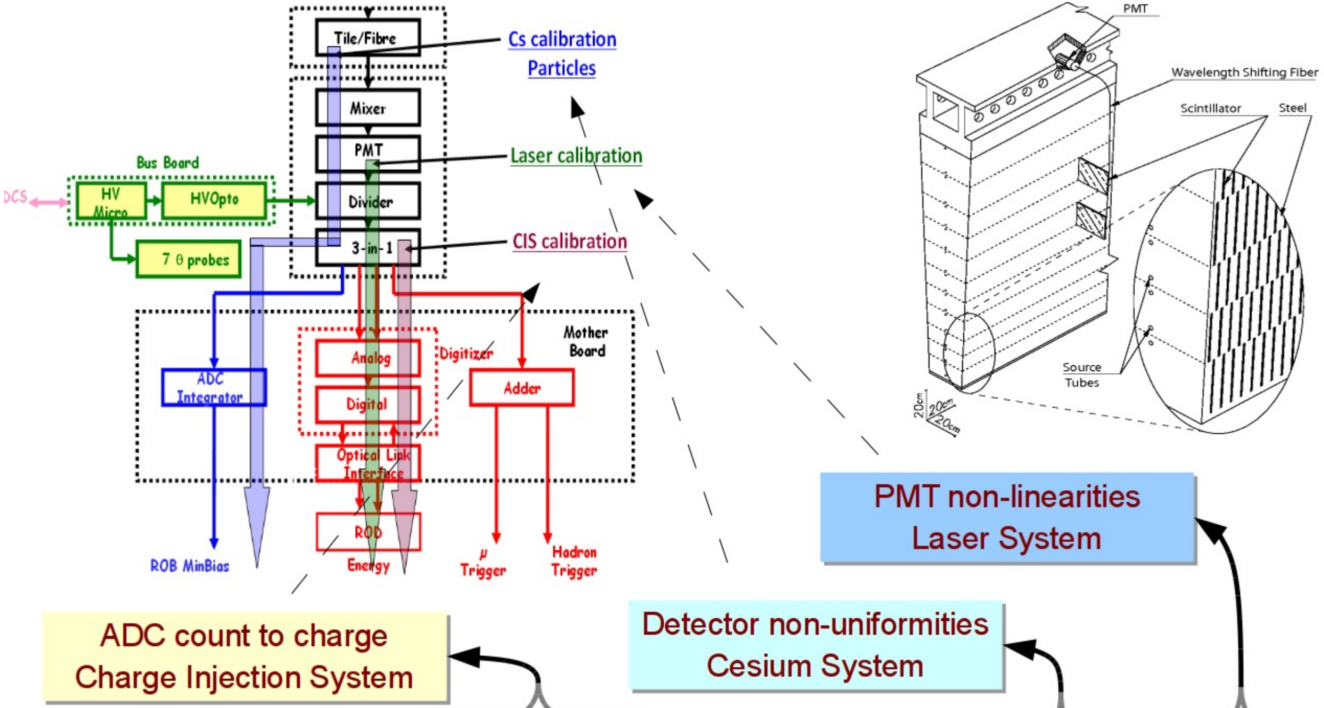

$$E_{channel} = A \cdot C_{ADC\rightarrow pC} \cdot C_{pC\rightarrow GeV} \cdot C_{Cs} \cdot C_{laser}$$

Time

Time

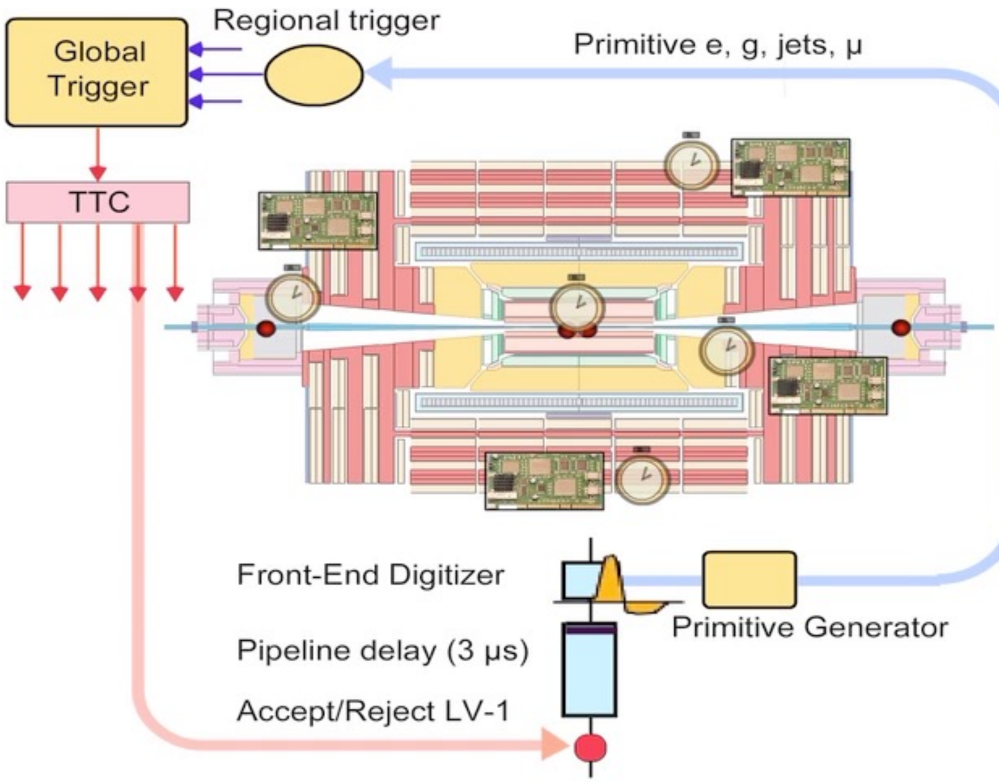
Global Trigger
Regional trigger
Primitive e, g, jets, μ
TTC
Front-End Digitizer
Pipeline delay (3 μs)
Accept/Reject LV-1
Primitive Generator

Clock distribution

- Let’s assume that your clock is received from an external source in “perfect” shape
  - For example the accelerator RF system
  - Or a time base for absolute time
- In order to have synchronized clocks in the detector, the signal has to be copied many times and distributed over thousands of connections
  - It is affected by
    - Length and quality of the transmission line
    - Environmental factors (temperature, humidity…)
  - It traverses a lot of intermediate electronics
  - In an environment full of noise sources
    - EM fields
    - Other electronics
    - Power supplies
    - …
- Distributing an accurate clock is far from trivial

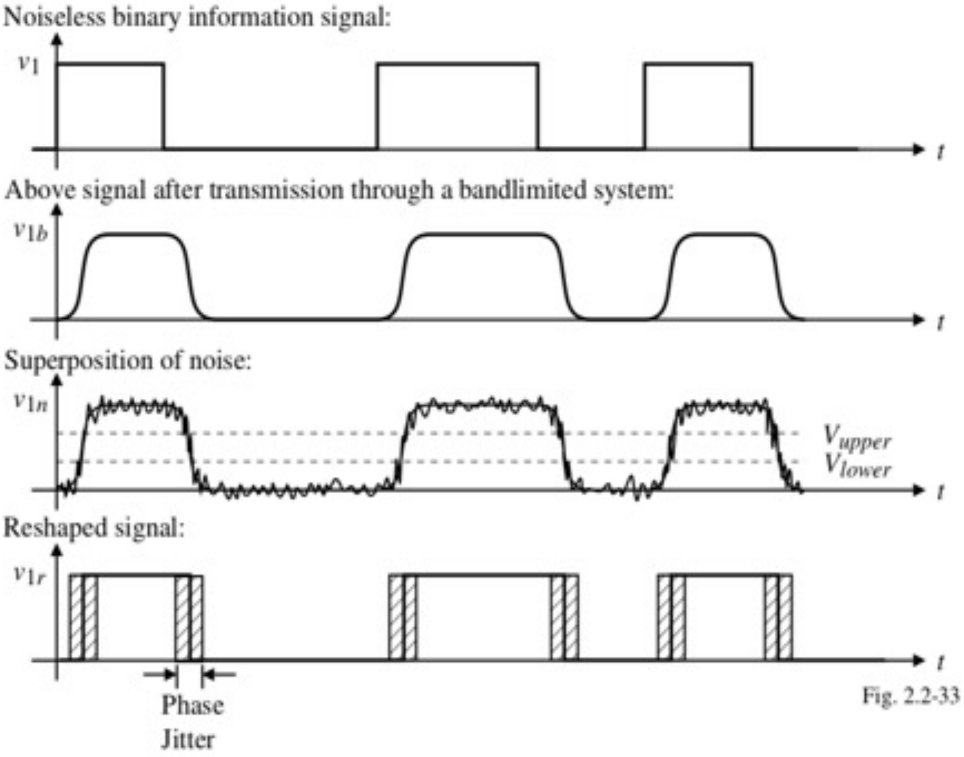

Fig. 2.2-33

Consequences of jitter

- The time references may be accurate in absolute terms but may have **phase differences at different places** in the detector
- The time reference may **not be stable and exhibit slow or fast phase variation**
- The jitter may affect your digitized data in different ways
  - **Distortion** of the digitized pulse
  - **Errors** in time measurement of a TDC
- On a trigger system the effect can be dramatic:
  - A trigger signal is **issued for the wrong time bucket** (corresponding to bunch-crossing or a specific instant in time)
    - As a result, the wrong information is stored (you take the picture too early, or too late)
- Mismeasurement due to jitter can also have serious effects
  - For example you can mis-measure a time of flight and take some wrong conclusion

Example: Augier

- **Detect air showers generated by cosmic rays above 10^{17} eV**
  - Expected rate $< 1/km^2$/century. Two large area detectors
- On each detector, a 3-level trigger operates at a wide range of primary energies, for both vertical and very inclined showers

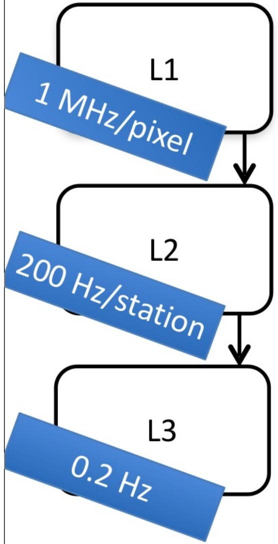

L1: (local) decides the pixel status (on/off)

- ADC counts > **threshold**
- ADC digitizes any 100 ns (**time** resolution)
- ADC values stored for 100 μs in **buffers**
- **Synchronized** with a signal from a GPS clock

L2: (local) identifies track segments

- Geometrical criteria with recognition algorithms on programmable patterns

L3: (central) makes spatial and temporal correlation between L2 triggers

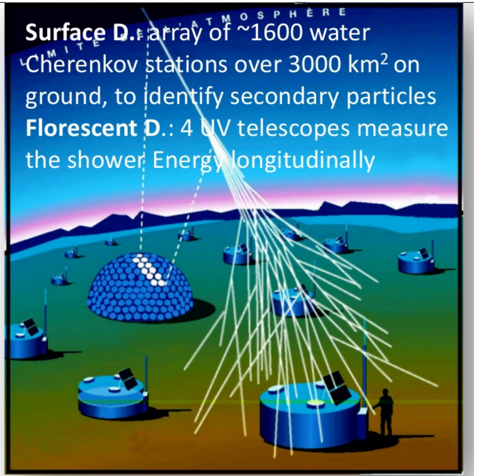

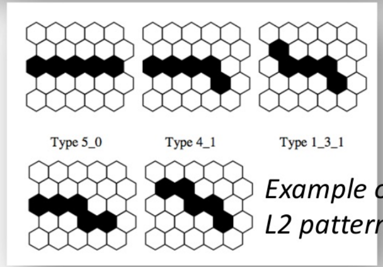

Example of L2 patterns

One event ~ 1MB → 0.2 MB/s bandwidth needed for the DAQ system

Where it's important ?

KM3NeT
Opens a new window on our universe

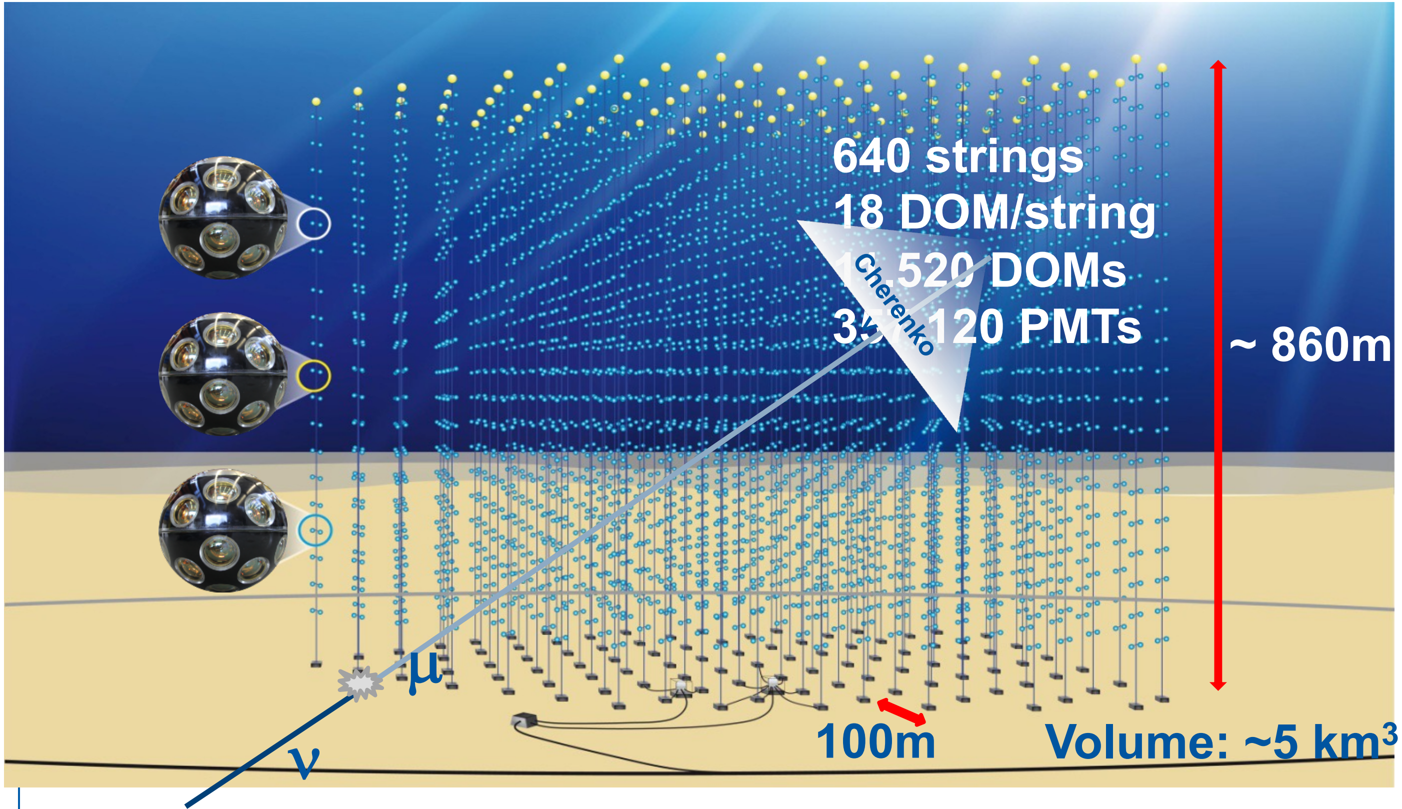

11.520 DOMs in the deep sea at 3-5 km depth

CERN

Aperture synthesis

SKA
SQUARE KILOMETRE ARRAY

SKA time distribution

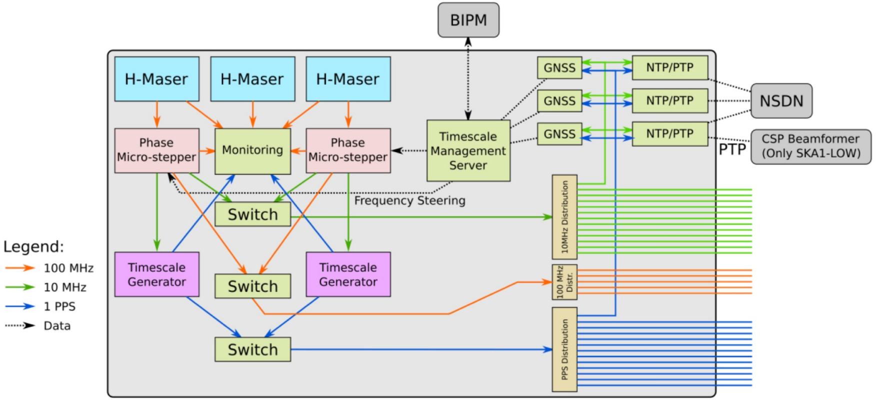

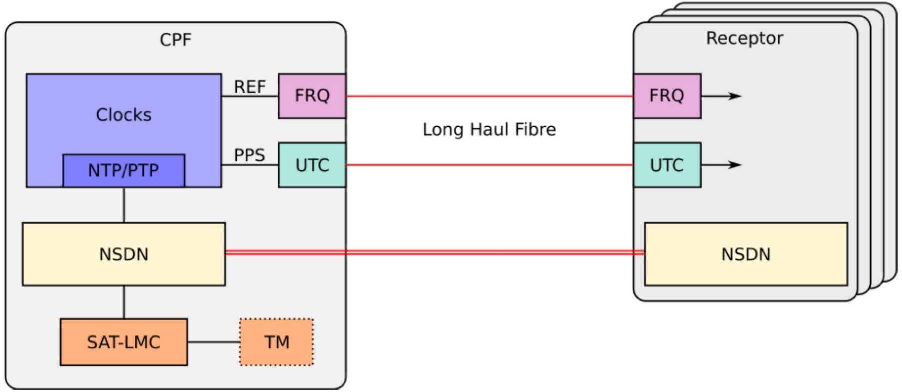


Optical fibres being deployed at MRO

Time at colliders

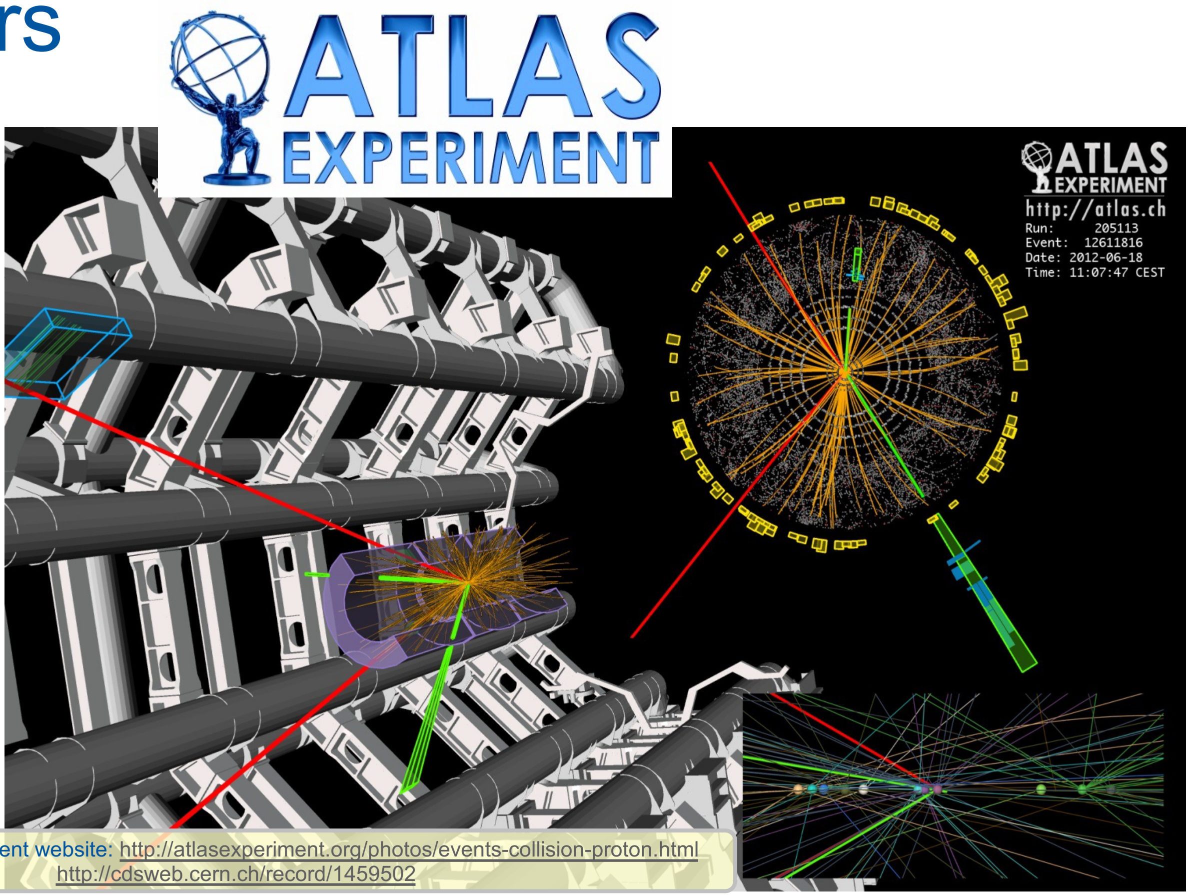

Picture ATLAS Experiment website: http://atlasexperiment.org/photos/events-collision-proton.html
http://cdsweb.cern.ch/record/1459502

Timestamping

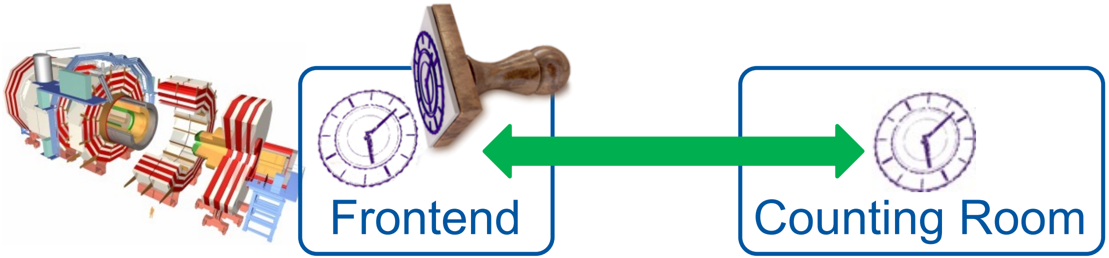

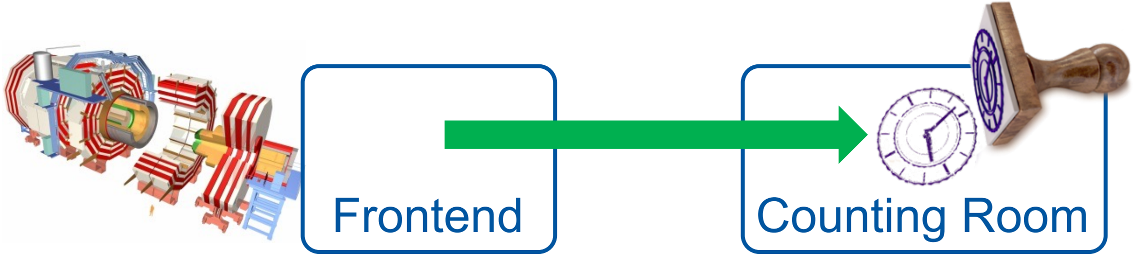

Syntonization

- Distribute frequency

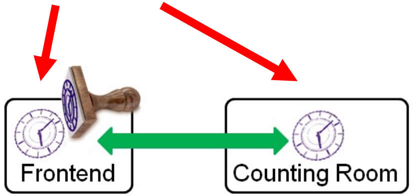

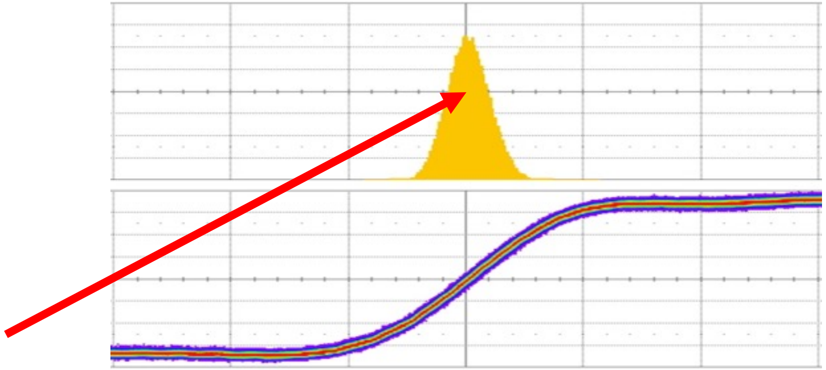

- Average clock edges

- Added single shot channel jitter

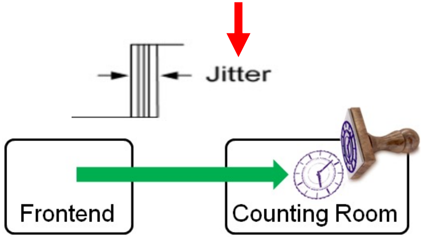

- Distributing frequency = Syntonization:
  - “The adjustment of two electronic circuits or devices in terms of frequency”

A timing system should distribute

- clock
  - with appropriate phase relative to some reference (e.g. the LHC bunch structure)

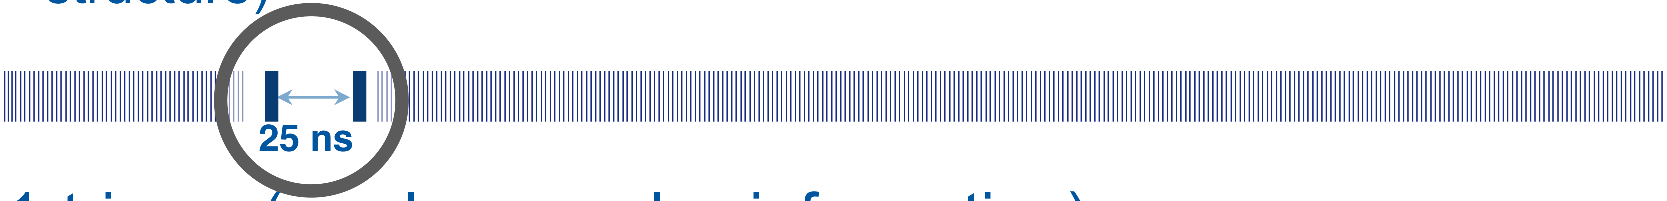

- L1-trigger (can be complex information)
- Bunch Counter (increment on LHC clock)
- Event Counter (increment on L1A)
- Additional synchronization commands for the front-end
- **Must be radiation and field tolerant (for colliders)**

One system

- Ideally, you want to distribute the clock AND the remaining information with a single system
- The clock could be embedded in a high-speed serial stream: can it be recovered with sufficient stability ?
  - What is “sufficient” ?

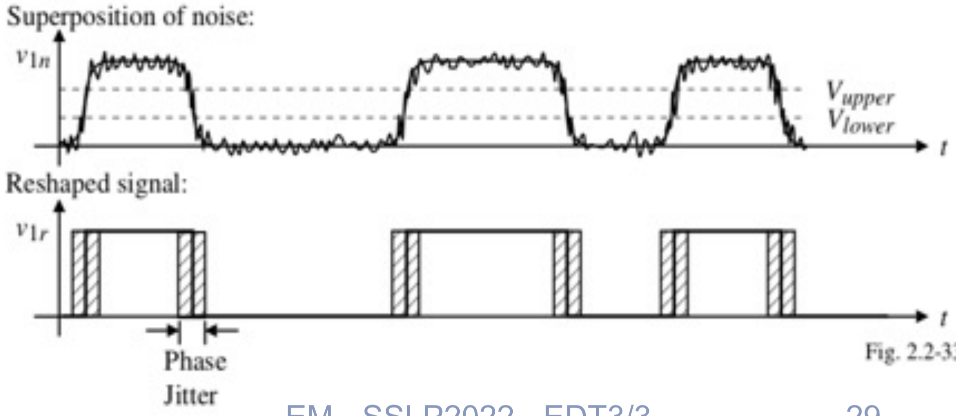

Fig. 2.2-33

Phase-locked loop

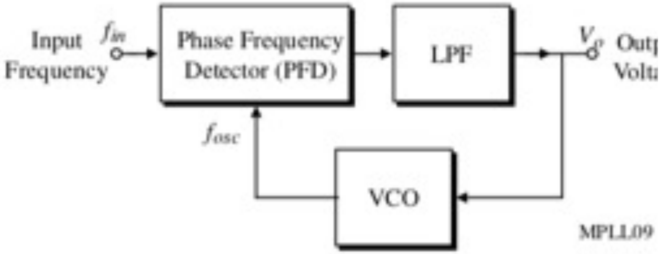

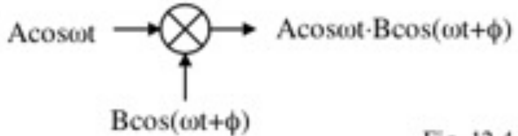

Fig. 12.4-2B

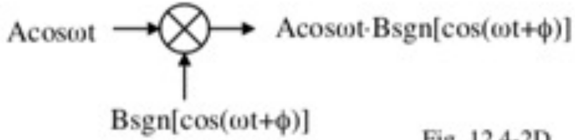

Fig. 12.4-2D

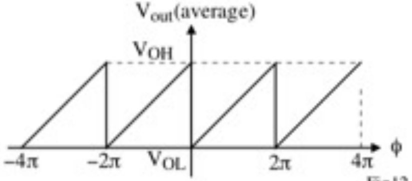

Fig 12.4-2F

$$K_D = \frac{V_{OH} - V_{OL}}{2\pi}$$

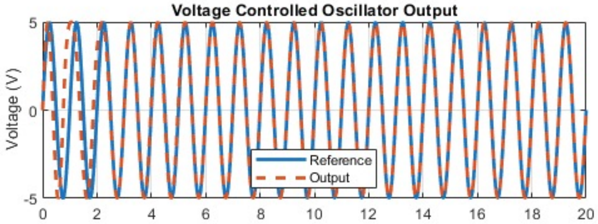

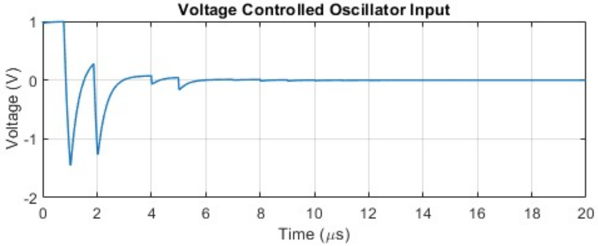

recover a noisy signal (jitter cleaning)
generate a stable frequency at multiples of an input frequency (clock synthesis)
distribute precisely timed clock pulses (clock distribution)

Problem: limited range of frequency
If input drifts -> unlock

Delay-lock loop:
VCO replaced by programmable delay line

Measuring phase noise

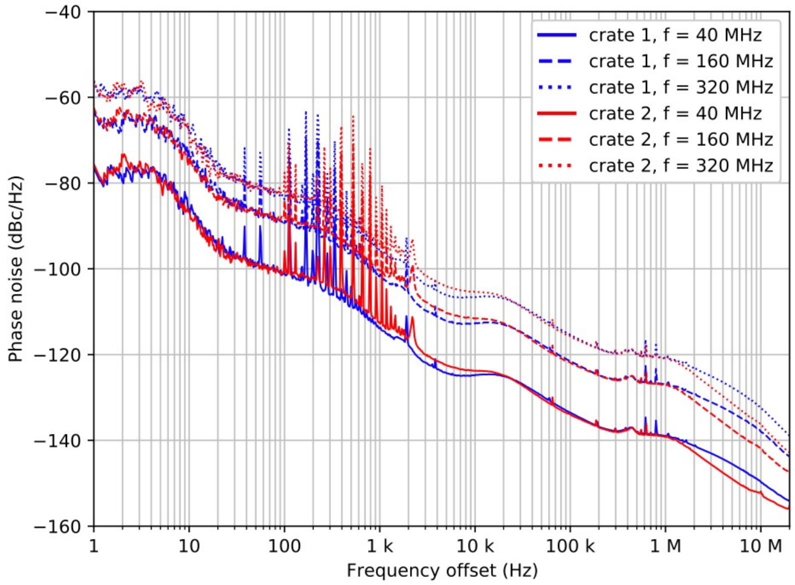

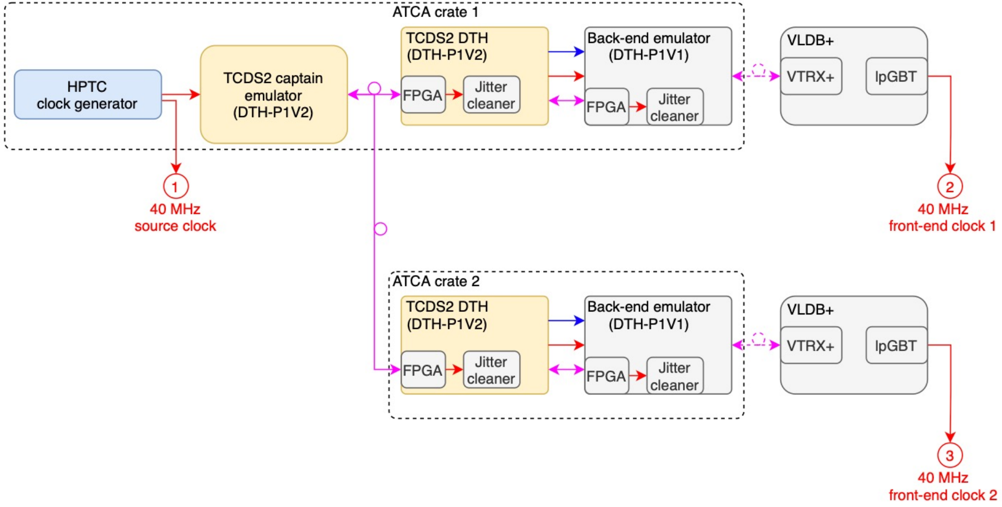

RMS jitter is obtained by integrating the phase noise power spectrum over the entire frequency band

More complex tools

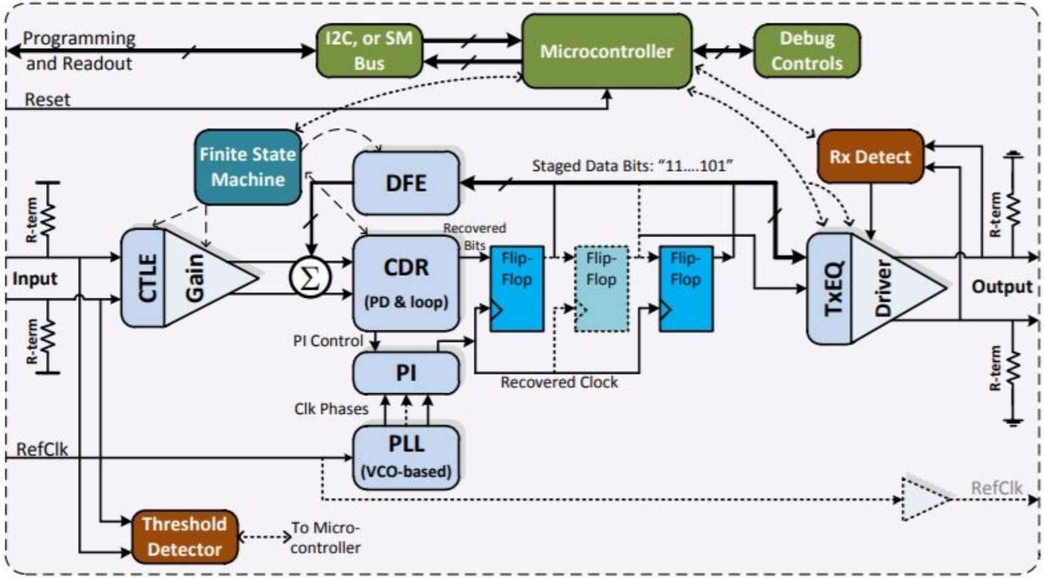

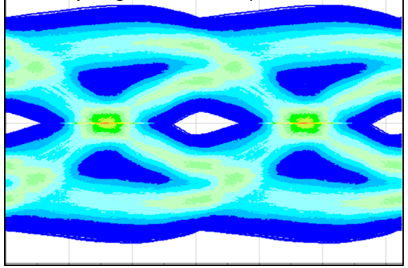
Input signal: Attenuated by channel loss

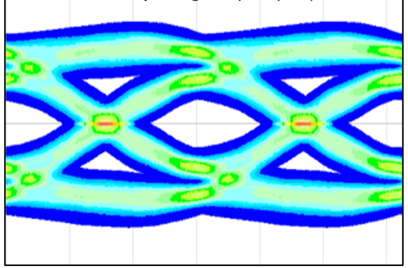
Redriver output: High frequency amplified

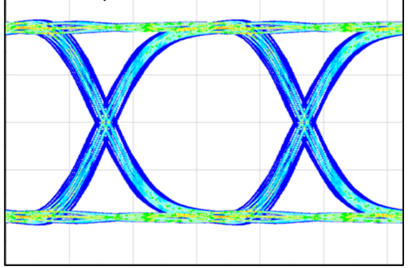
Retimer output: Data is recovered and retransmitted

https://pcisig.com/pci-express®-retimers-vs-redrivers-eye-popping-difference

What are the applications that demand most phase stability ?

What is pileup

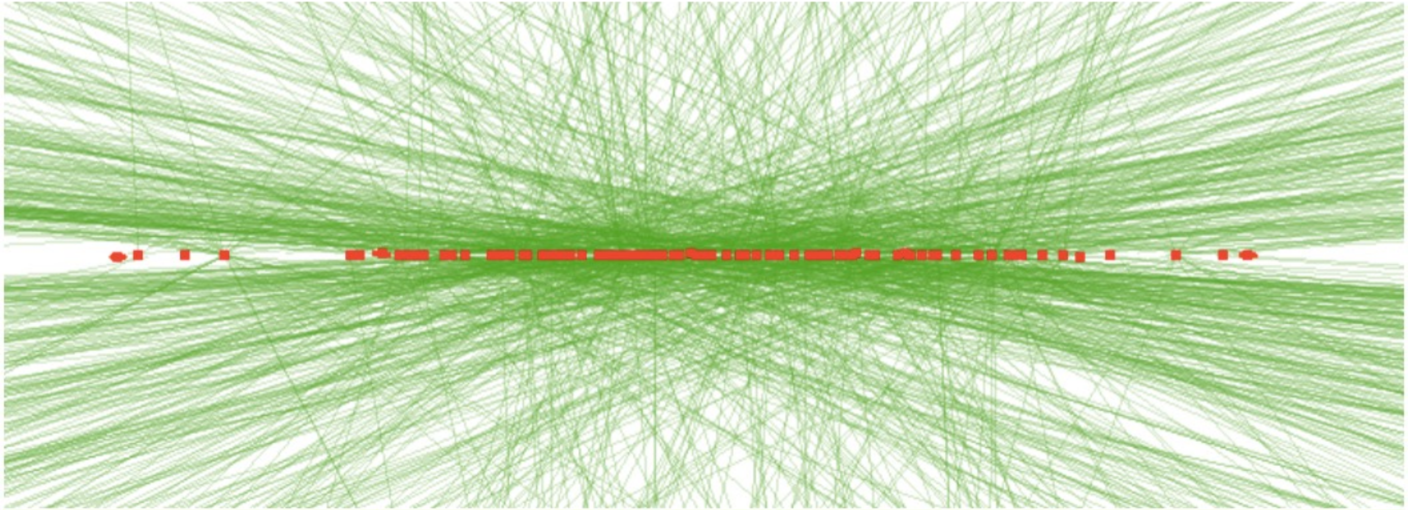

Top pair event + 140 additional low energy interactions
"Classical" spatial view of the vertices

Identifying vertex by time of flight

- With two time and position measurements eg. from two photons and with the constraint from the beam axis x and y location, the vertex x and t can be calculated.
- Equivalent to GPS with two sattelites.

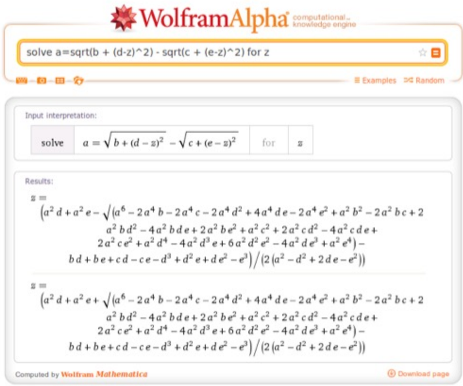

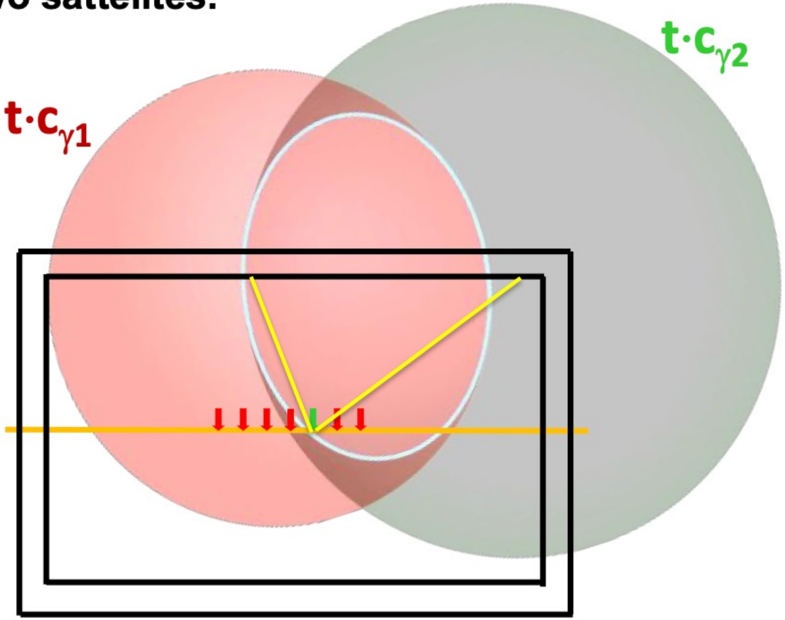

Pileup mitigation with timing detector

- **High pile-up has major impact on JetMET reconstruction.**
- **At low p_T (40 GeV) PU energy in the jet cone exceeds the jet energy by a factor 2 or more.**
- **Each PU event adds 3 GeV in quadrature to the MET resolution, resulting in about 40 GeV for 200 GeV.**

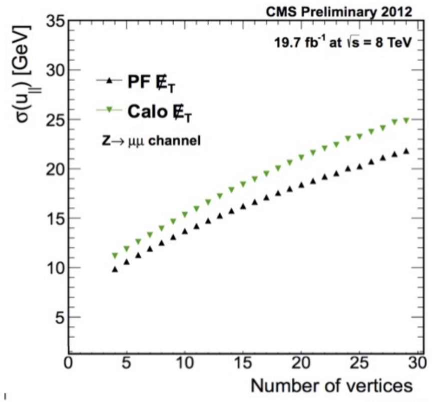

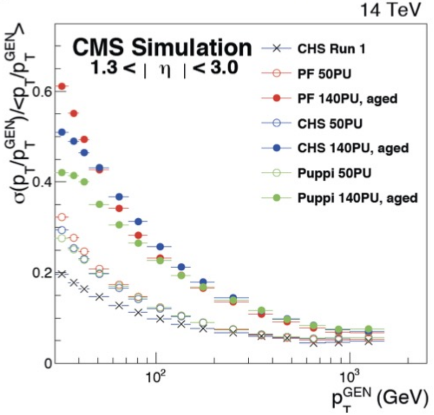

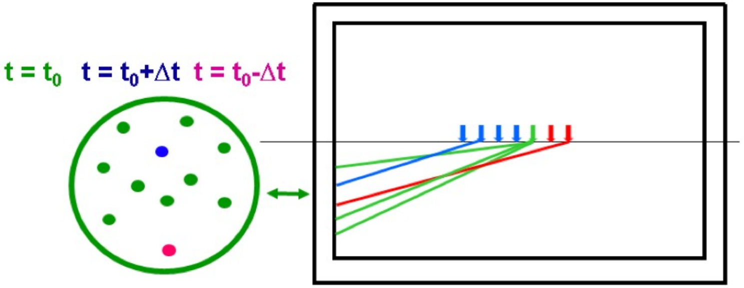

Correct vertex identification

- Synergies between Calorimetric and MIP timing : Increased efficiency for time vertexing.
- $H \rightarrow \gamma\gamma$ use case : Vertex ID efficiency 80% in Run I, 40% expected in HL-LHC, 55% with photon-only time vertex, 75% with photon+track time vertex.
- Results corresponds to a 10% and 30% increase in equivalent luminosity.
- Results to be presented at the ECFA workshop in Aix les Bains.

ECAL only timing

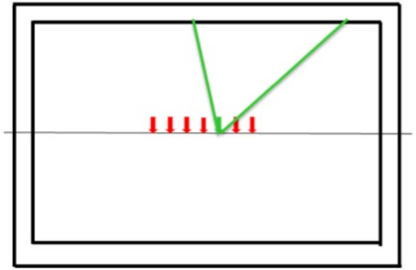

ECAL+MIP timing

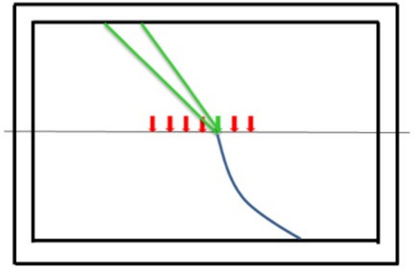

Requirements

The reference clock must have phase accuracy better than the target resolution

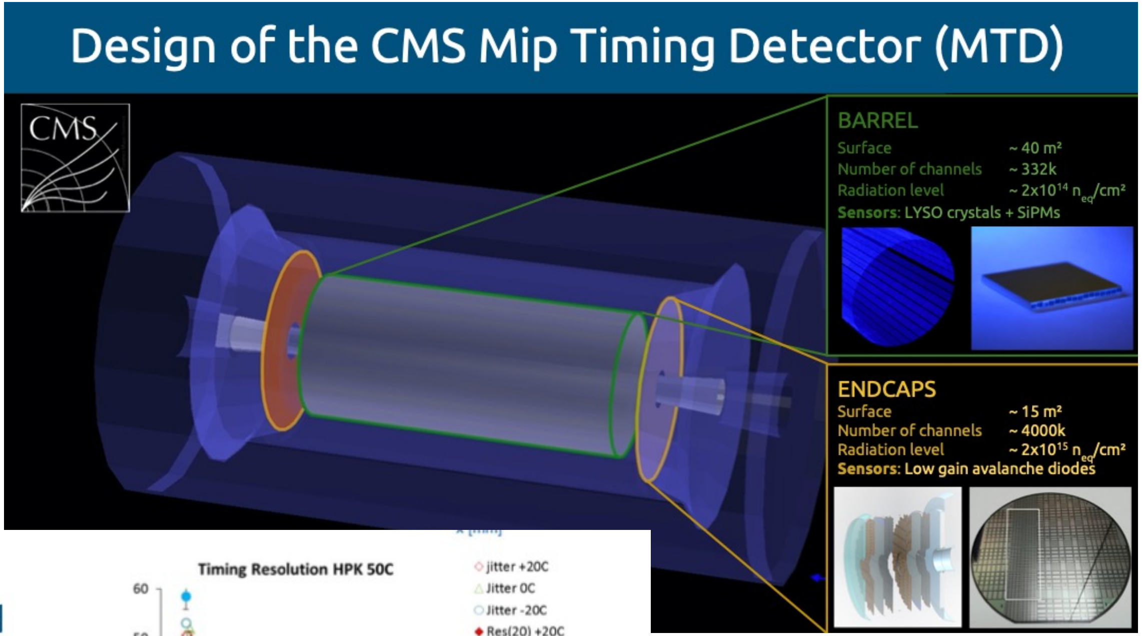

- Target time resolution of 30 ps achieved
  - Noise jitter term <25 ps for gain>15
    - ~N/(dV/dt)
  - Intrinsic limit from Landau fluctuations:
    - Spatially non-uniform energy deposits along the track cause event-by-event pulse distortions
    - Constant: ~25 ps

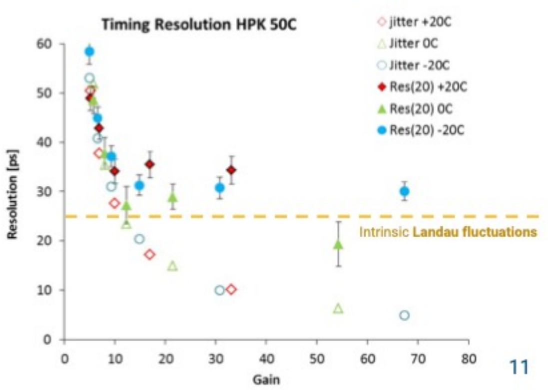

Wrap-up
Visionary part

DAQ: nice and easy

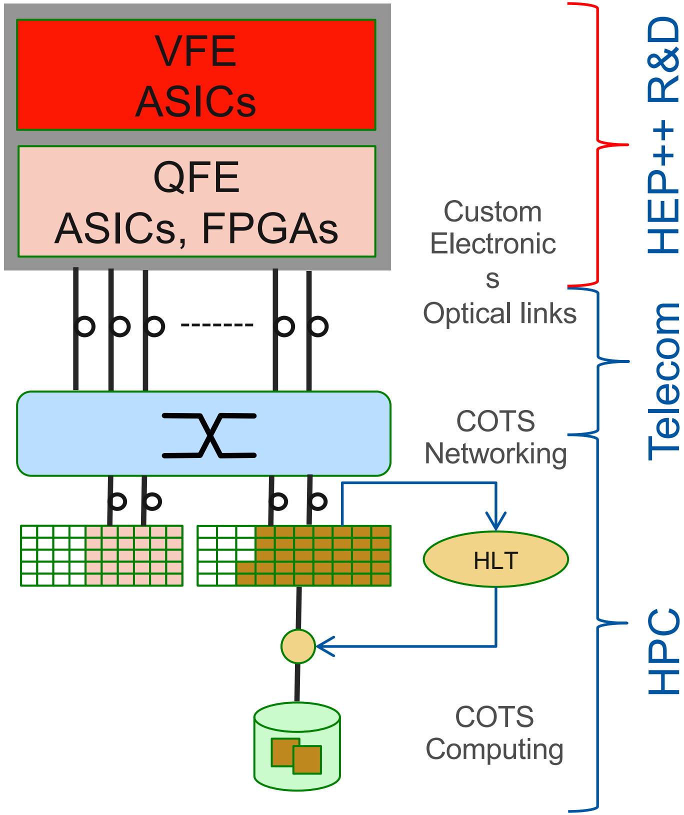

VFE does the analog part, ADC, low-level calibration, zero suppression, lossless compression, *optical links*

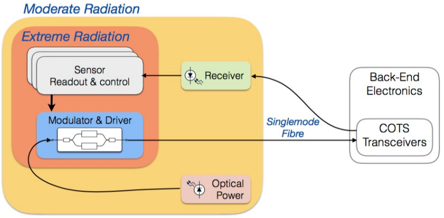

low-power, rad-hard *(rad-tolerant)*

QFE does medium scale aggregation, local reconstruction, ”lossy” compression, transition to standard protocol on optical links

asynchronous – precision clock - timestamping

COTS switched networks provide further aggregation, up to and including event building

COTS servers with co-processors do the final selection

One-size-fits-all type of problem

Some real-life problems

Beyond HL-LHC

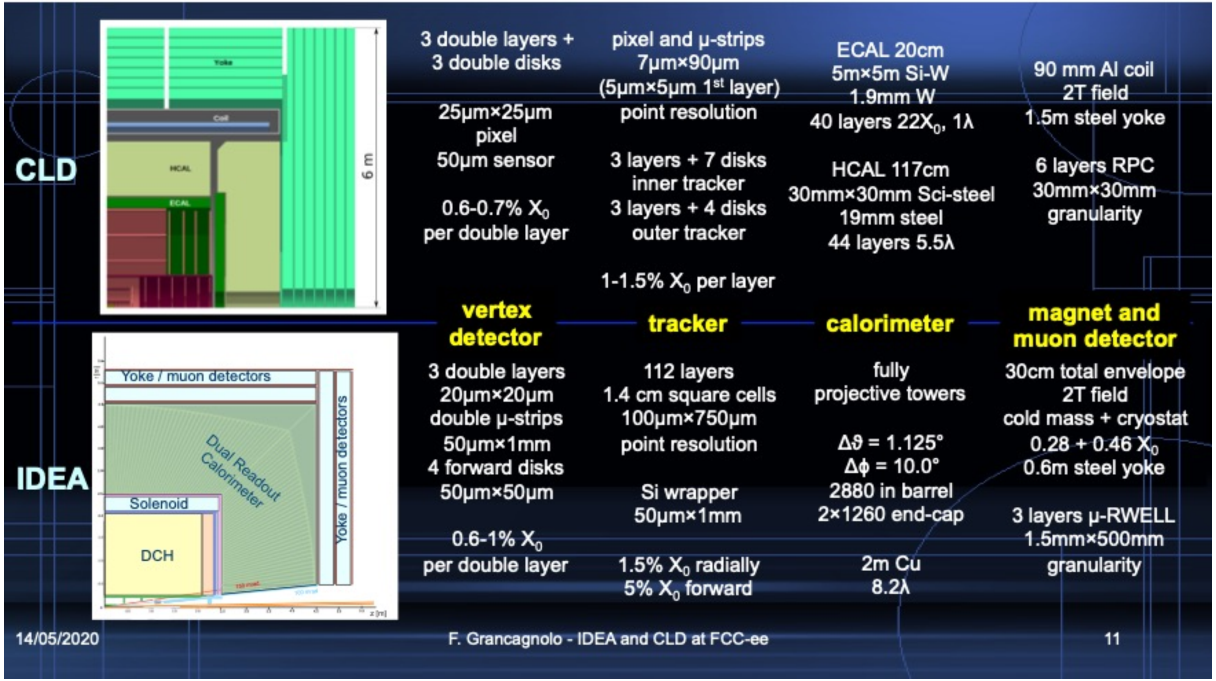

3.7 ns interbunch at Z pole

Rates are high but events small

Tracker with many channels but occupancy low (~kB??)

Trigger not challenging, but precision measurements benefit from multiple strategies

Trigger-less data rates similar to HL-LHC

Stepping stone for the hh detectors

Beyond HL-LHC

- Higher energy
  - larger B field, solenoid bore radius -> tracker
- 400 m^2 silicon, 10^{10} channels
- Higher to extreme fluences, less accessibility
- Luminosity: always as high as possible
  - Shorter interbunch to reduce PU ?

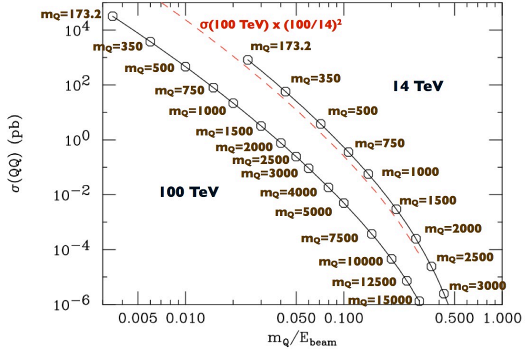

31 GHz of pp collisions
Pile-up 1000
4 THz of tracks

Comparison to ATLAS & CMS

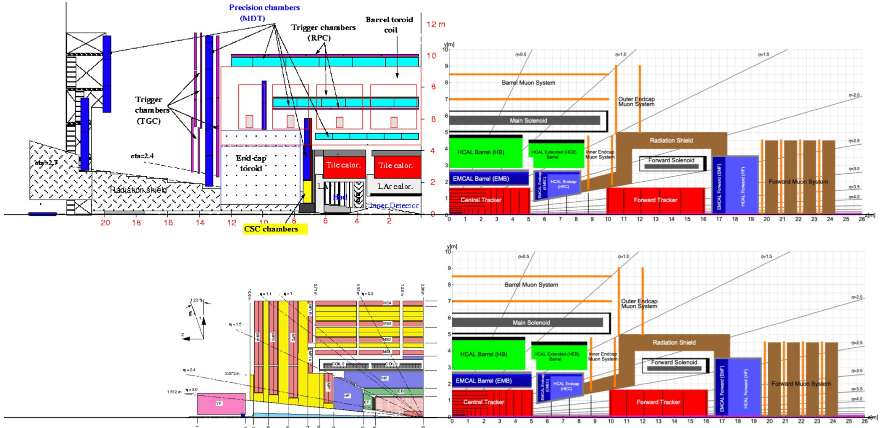

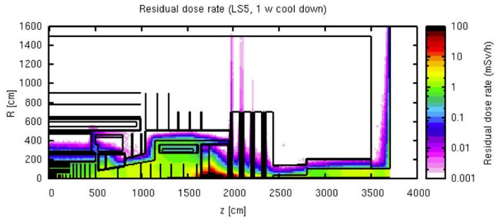

https://indico.cern.ch/event/727555/contributions/3461232/attachments/1869213/3075082/fcc_hh_detector_brussels_june_2019_riegler.pdf

Un-triggered readout at 40MHz **1000 - 1500TByte/s** over optical links to the underground service cavern and/or HLT

Not Just Colliders

- Neutrino@accelerator (DuNE…)
- Dark matter at BD or LL
- Next-generation specialized experiments and FT
- Neutrino (IceCube-2..)
- Astroparticle (CTA…)
- Radioastronomy (SKA…)
- GW (ET, CE…)

Multi-messenger astronomy:
network all the above in real time

Data reduction
Synchronization
Data rates
“Image” processing
AI applications
(CNN, GN, AE?)
Reliability

Precision clock distribution
Synchronization
High-bandwidth long-distance links
Reliability

Common needs for reliable systems

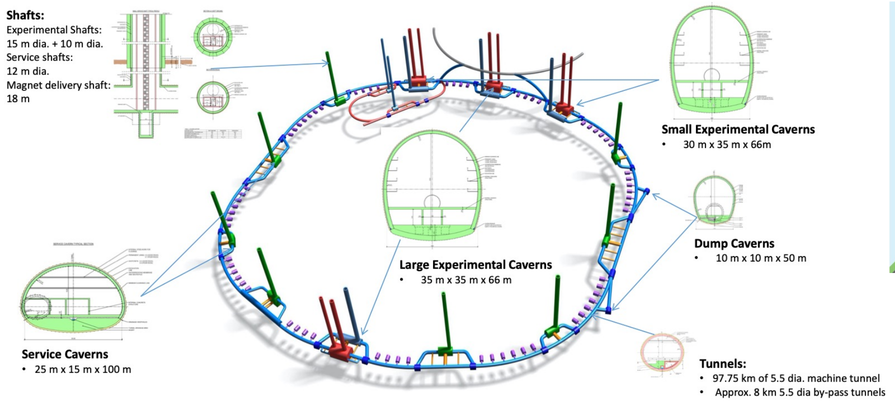

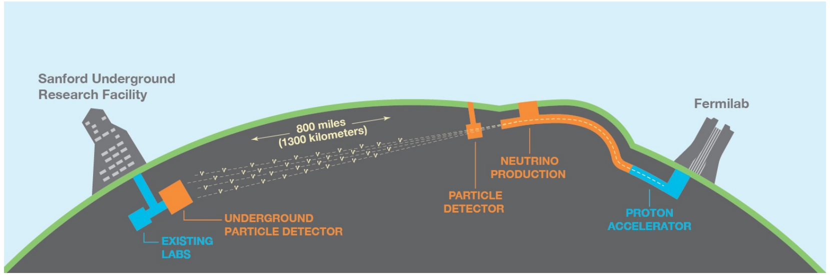

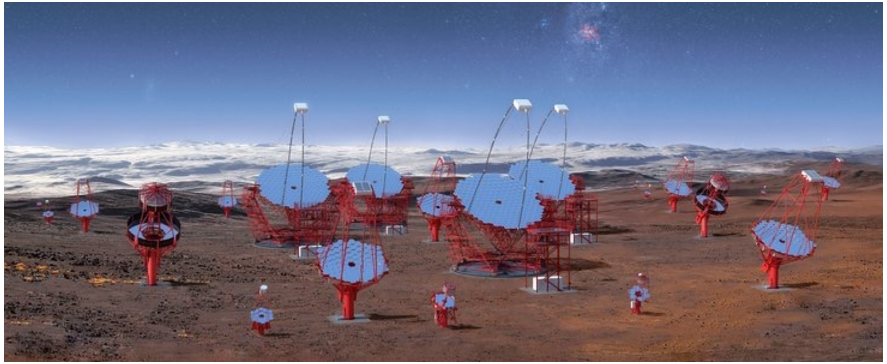

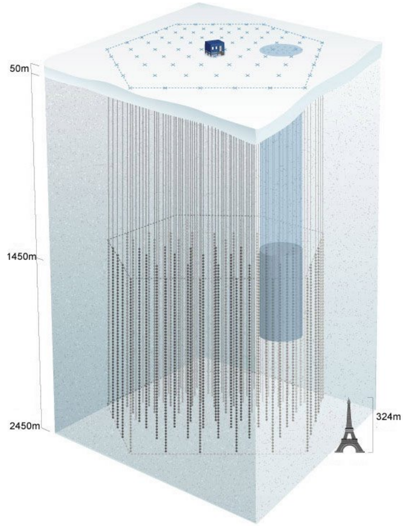

Optimistic outlook / lessons learned

- Move “as early as possible” to COTS networking
  - Requires buffer at sender (PC memory or something else)
  - Bare Ethernet not lossless unless sufficient buffer at switch (must keep usage below ~50%) - residual losses as deadtime (accounting) – need reliable delivery
- Could move “EVB” functionality to BE
  - Reliable protocol using HBM for congestion, aggregation
    - Integrated in back-end
  - Complete fw implementation of e.g. TCP
  - Delegate final EVB to HLT nodes (mild timestamping)
- Economic and sociological reasons (may) make this hard

More stuff (it time allows)

First stop: Readout at FE

- Optical Data Transmission technology is key
  - High Bandwidth **low mass, low power**
  - Immune to **electromagnetic interference** (+ isolation between power and readout)
  - Sufficiently **radiation tolerant**

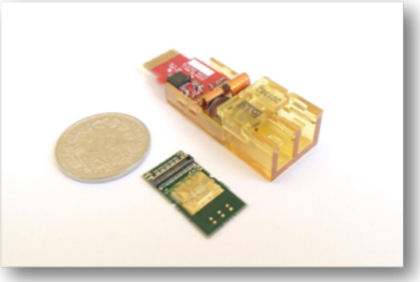

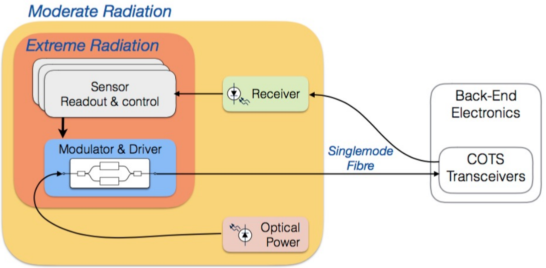

- Optical layer is only part of the story
  - **Acquisition, aggregation, and serialisation before** transmission over the link
- **ASICs** need to be specified and designed to meet system requirements
  - Increasingly complex over generations
  - Common developments are key
  - Reuse (of design, concept)

synchronous or asynchronous

- Using COTS at FE == profit of high bandwidth “for free”
  - Fixed speed and narrow locking range
- Asynchronous readout (with time stamping…)
  - less complex back-end electronics
  - In principle can use “standard” protocol (e.g. Ethernet) and connect “quasi-FE” to “commercial” equipment
    - Challenging clock distribution, disciplined clock, phase stability at FE
  - Additional bandwidth
- Synchronous
  - RF stability (ramps) and distribution vs. locking range
  - COTS not compatible with specific RF frequencies
    - need deterministic behavior (and if not, can losses be afforded/quantified)
      - Deterministic protocols hard
    - Reliable protocols require lots of buffer at the sender (and are not deterministic by construction)

“trigger-less” is good ?

~~Trigger~~ :== read out everything (or “just without” L1 trigger) ?

+ Reduce relevance of custom processors and special data paths
- More high speed links at or near the front end (with the usual problems…)
  - material budget (for power, cooling, fibres)
  - rad hardness (of optical transmitters, fibres)
    - Will get worse at “future” colliders
    - Optical on FE needs a lot of technology for which we are the only customers
- Do aggregation and optical links as close as possible to FE
  - “Lossy” data reduction, it looks a bit like trigger ? – but, at FE, level of aggregation is limited by
    - geometrical distribution of sources
    - space available
    - Environment to place connectivity / intelligence (the usual…rad,temp,field…)
- Transition to COTS
  - Link speeds are constrained by COTS standards
  - Speaks in favor of abandoning synchronous readout

Caveat

- Turns out that “doing everything in software” is not so easy even with a HW L1 trigger
  - size of the switched network, amount of buffer to allow for “non-deterministic” behaviour of software algorithms
- amount of computing needed vs. CPU evolution
  - With asynchronous time-stamped data, just assembling the right pieces would not be a trivial task
  - Lots of low-hanging fruit ends up being a bit too high
  - GPUs…help but transition is slow (but at least it is happening)

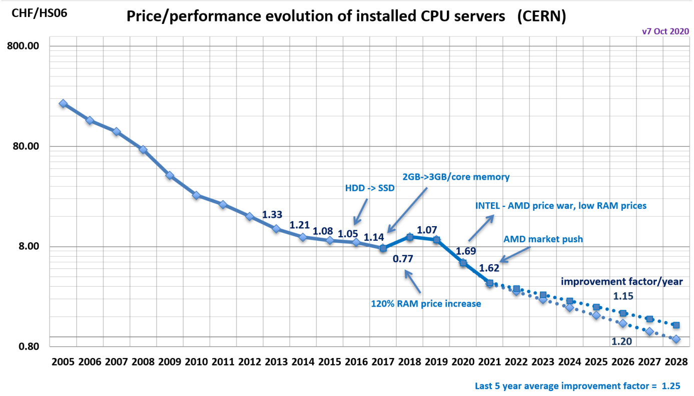

- …people end-up proposing FPGA custom co-processors…
- …or pre-processors (anyone ?) [while there is general consensus now that ASICs are hard a lot of people still think that custom boards with FPGAs are easy]…beats the purpose ?
  - Paying for CPU is not popular…people will want to develop their own boards
  - That’s how projects get funded at institutions: **not to develop software, or even firmware** (see later in “sociology”)

A look at the hardware triggers

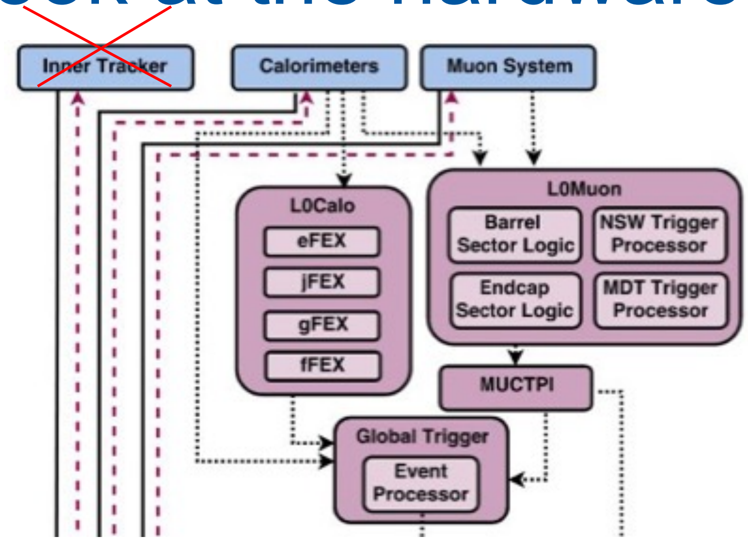

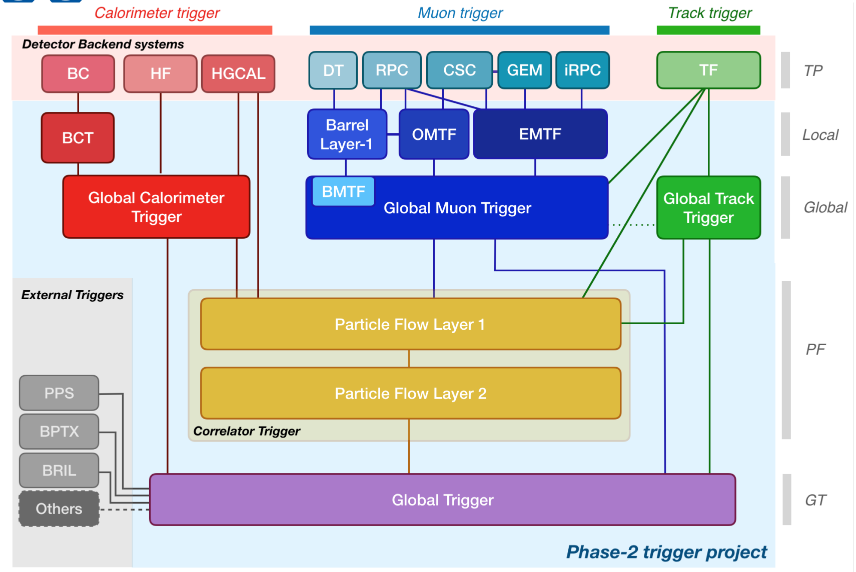

Some draconian **"lossy compression"** here

- Get rid of entire detectors, make "trigger cells", "trigger primitives"
- Final outcome is…well, one bit

Some "trigger-less" readout

- Trigger inputs ARE **read out at BX rate** (in some cases "primitives" are made at back-end, from streaming data)
- Trigger processors do **aggregation**
  - In some cases information from **multiple sub-detectors** are combined

There is a lot **of information in the intermediate layers** of the hw trigger…perhaps something to be learned there

Yet one wants to do fancy things

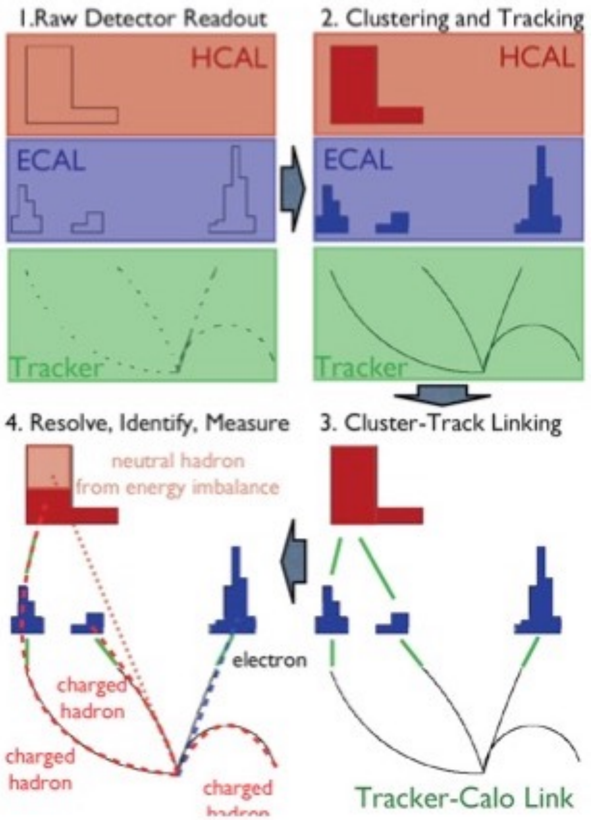

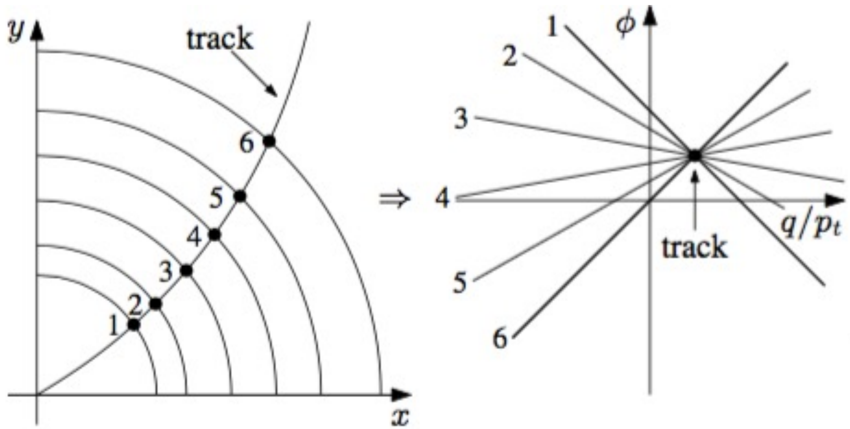

Several aspects to point out here:

- Enough information to “do the job”
- Aggregate it in the correct dimension
- Do all this within a short time
- Modular systems with custom interconnects
- Lots of complex algorithms to code in firmware

Some intelligence in the detector…

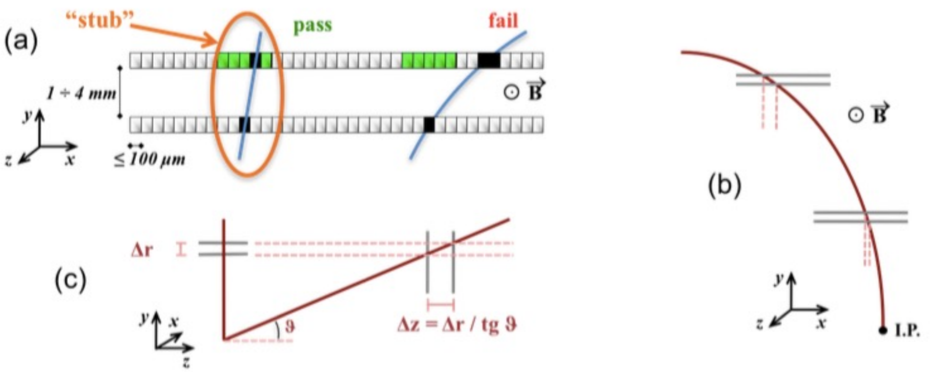

If one wants to do more Interconnects are needed

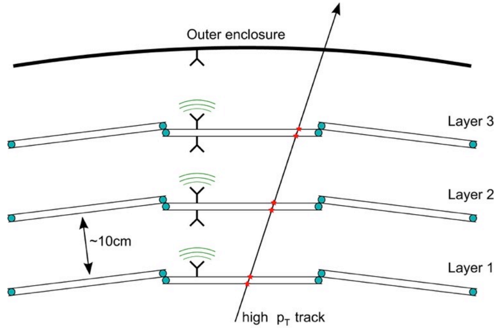

Inefficient at low Pt, large impact parameter, large η

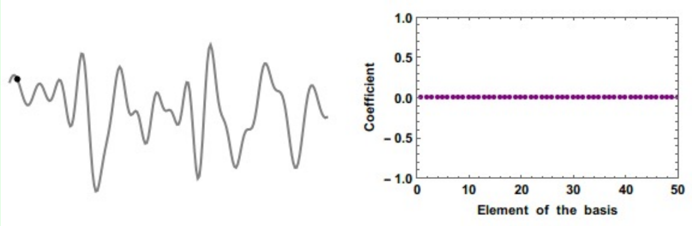

Lossy compression requires sparsity in some space
VFE connectivity insufficient (or in the wrong dimension)

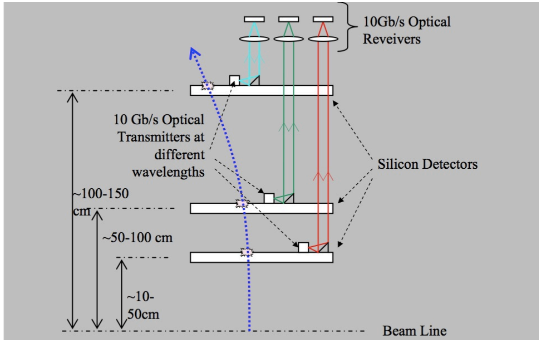

- Zero-th order problem: how to use complex, highly granular detectors in the trigger
- First order problem: how to make high-rate read-out possible (with or without trigger)
  - You can’t get rid of complex front-end (and back-end) if you can’t get rid of a hardware trigger
- Second order problem: build a tracker that spits out tracks, a calorimeter that spits out clusters ?

Some cursory conclusion

Hardware L1 is now a system of (generic) processors doing "local" event building, and processing events in parallel – **L1 and HLT are approaching**

- As an aside, this can and must be exploited to reduce the computation needs of the HLT
  - L1 objects much more useful
- Can capture L1 data for use in "physics at BX rate" (L1 "scouting")

Both L1 and HLT only require **"trivial" parallelism**: work on one BX (event) at a time.

- Driven by the almost-synchronous nature of the task – could be challenged by **asynchronous readout** and/or **multi-bx phenomena** (VLL)
- Information **only flows in one direction** – sub-optimal use of bi-directional links (but easy traffic) [on-demand event building…]

Some cursory conclusion

A hardware trigger can live with (very) lossy compression because it just needs enough information to classify events and be **“mostly correct”** – that L1 is only mostly correct is **an accepted fact of life** (see below)

Whatever readout scheme one chooses **needs to be “almost always correct”** – at the lowest level possible (i.e. not miss a track or know exactly what you missed):

- Ability to extract all relevant information from front-end electronics
  - Readout scheme
  - Lossy compression and traffic equalization (ML techniques…CNN, AE (?), compressive sensing (?))
- Ability to aggregate and process data **at sufficient scale** “on-detector”
  - Interconnect at “quasi” front-end
  - Powerful FPGA-based processors that can work in QFE environment
  - Still a lot of downstream connectivity, if possible on COTS
- Accept a redefinition of what “raw-data” means (notice that such a redefinition is mostly NOT yet accepted at the next level, i.e. HLT)

Future DAQ/trigger system will need **“holistic” system engineering at detector design level** that takes into account readout (and trigger if necessary)

Thank you