WA105 Online Computing

Elisabetta Pennacchio
IPN Lyon

Software and Computing Bi-Weekly, December 15th 2015

- A short review of WA105 data acquisition system (already presented a few times at the general meetings and DAQ meetings)

- Data flow

- Online storage and processing requirements

The LBNO-DEMO/WA105 experiment at CERN

→ 1/20 of 10 kton DUNE detector

→ 6x6x6m^3 active volume, 300 tons , 7680 readout channels

→ Exposure to charged hadrons, muons and electrons beams (0.5-20(10) GeV/c)

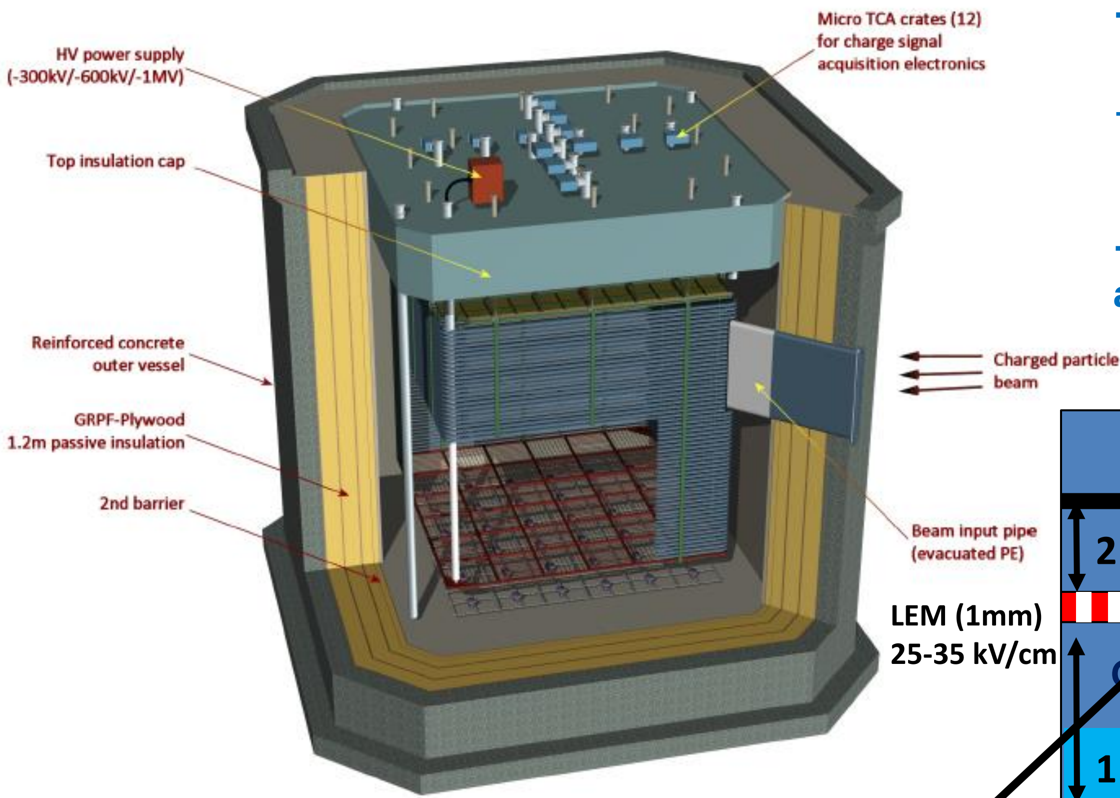

50x50 cm² LEM

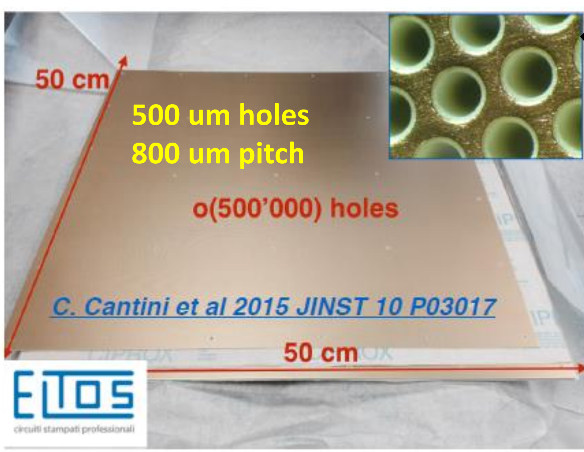

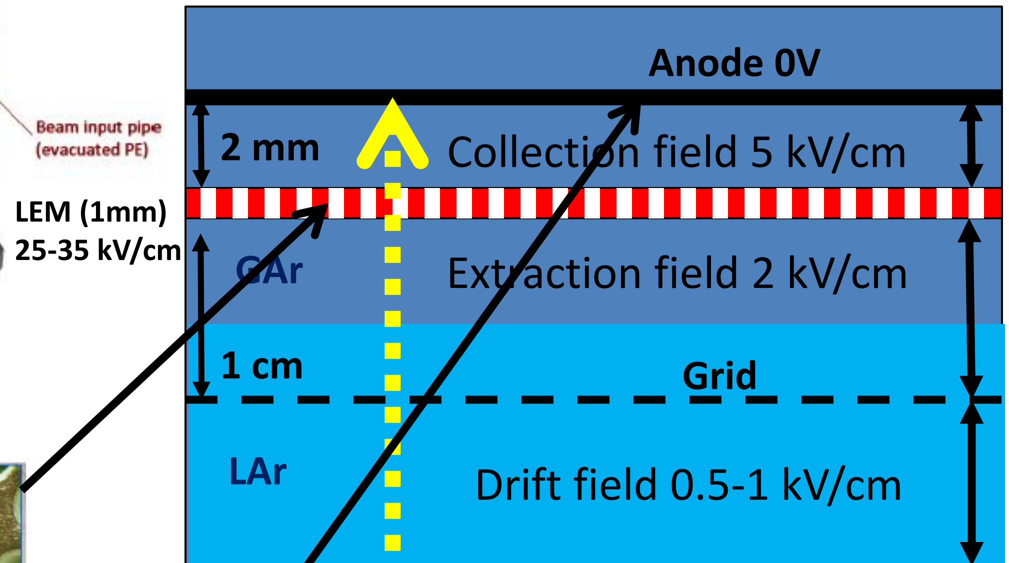

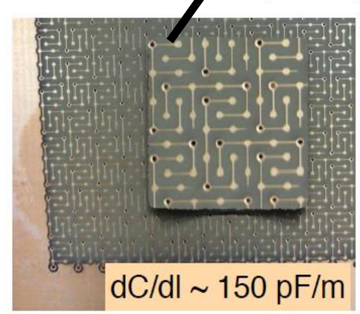

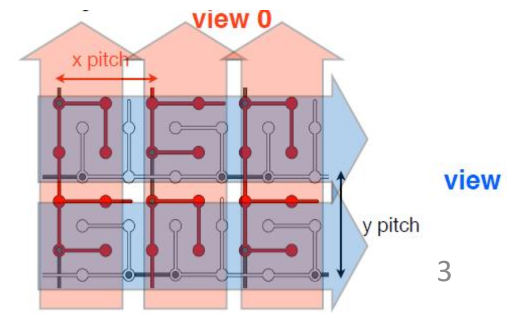

Double phase liquid argon TPC
6x6x6 m3 active volume

→ **Event size: drift window of 7680 channels x 10000 samples = 146.8 MB**

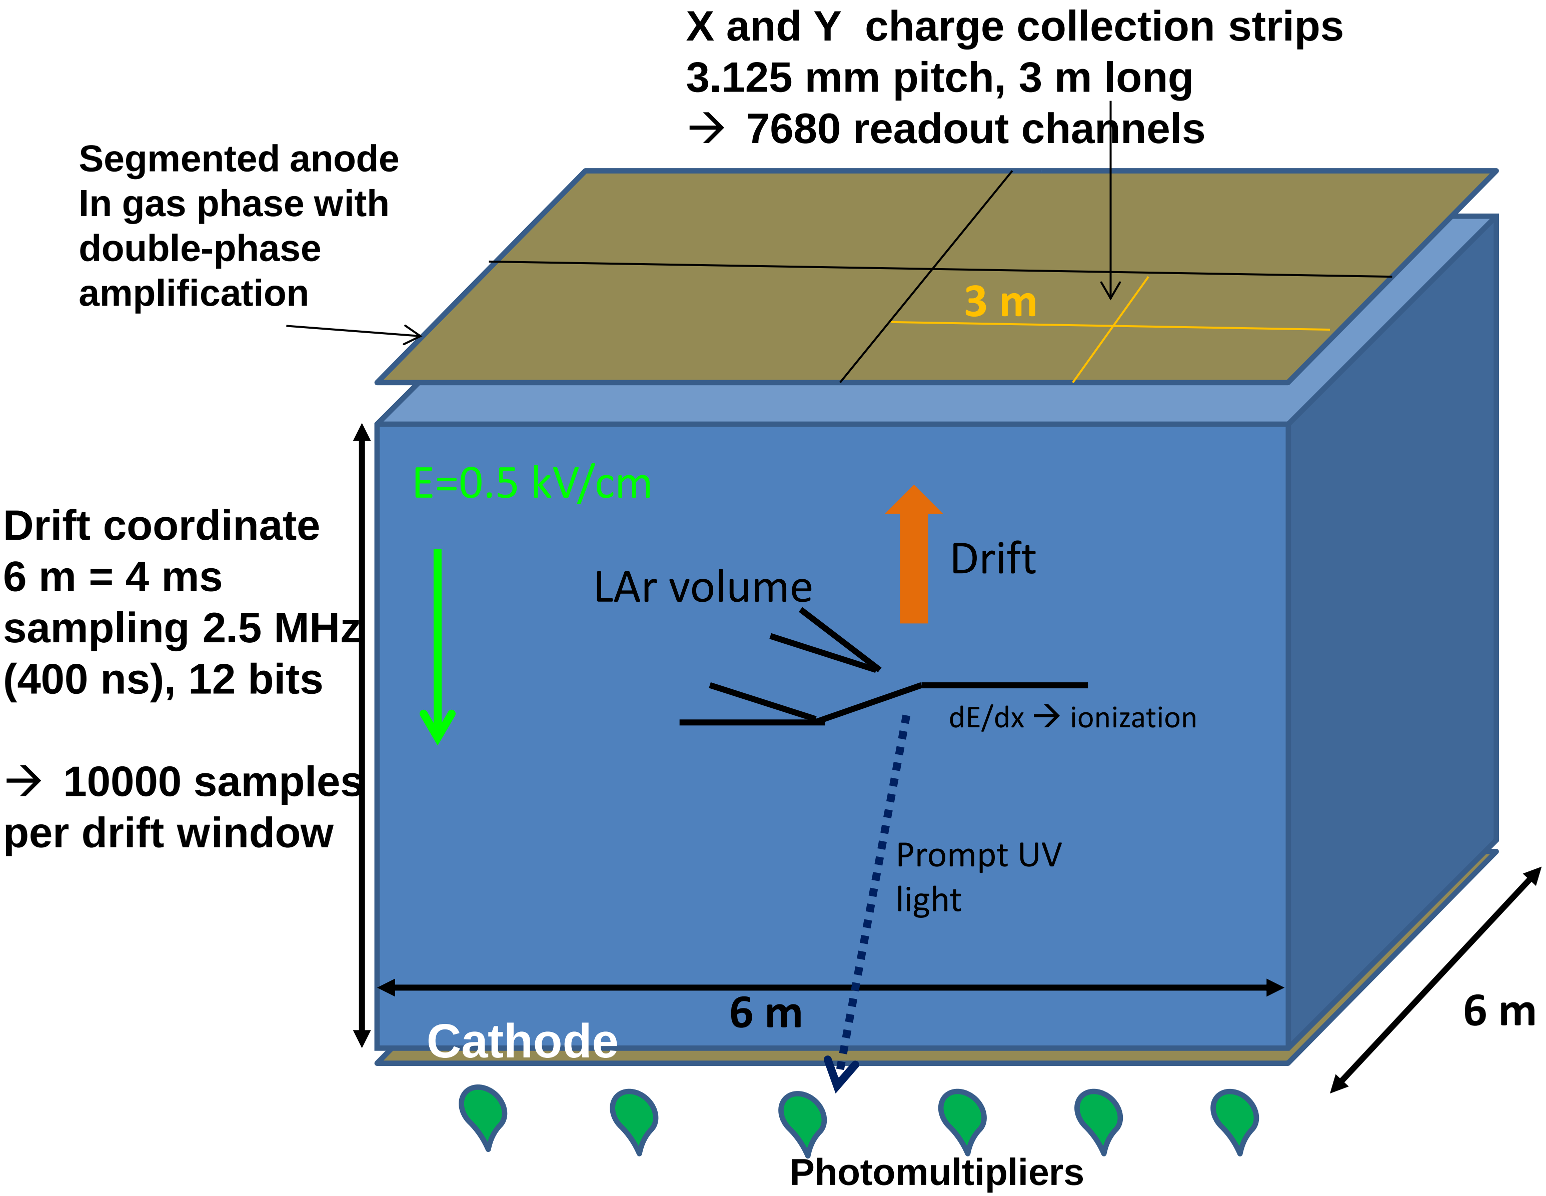

Accessible cold front-end electronics and uTCA DAQ system for WA105

Cryogenic ASIC amplifiers externally accessible at the bottom of the signal chimneys
→ short cables capacitance, low noise at low T, ongoing R&D since 2006 for LBNO/WA105/DUNE

ASIC (CMOS 0.35 um) 16 ch. amplifiers working at ~110 K to profit from minimal noise conditions:
- FE electronics inside chimneys, cards fixed to a plug accessible from outside
- Distance cards-CRP<50 cm
- Dynamic range 40 mips, (1200 fC)

(LEM gain =20)
- 1300 e- ENC @250 pF, <100 keV sensitivity
- Single and double-slope versions
- Power consumption <18 mW/ch

Digital electronics at warm on the tank deck:
- Architecture based on uTCA standard
- Local processors replaced by virtual processors emulated in low cost FPGAs (NIOS)
- Integration of the time distribution chain (improved PTP)
- Bittware S5-PCIe-HQ 10 Gbe backend with OPENCL and high computing power in FPGAs

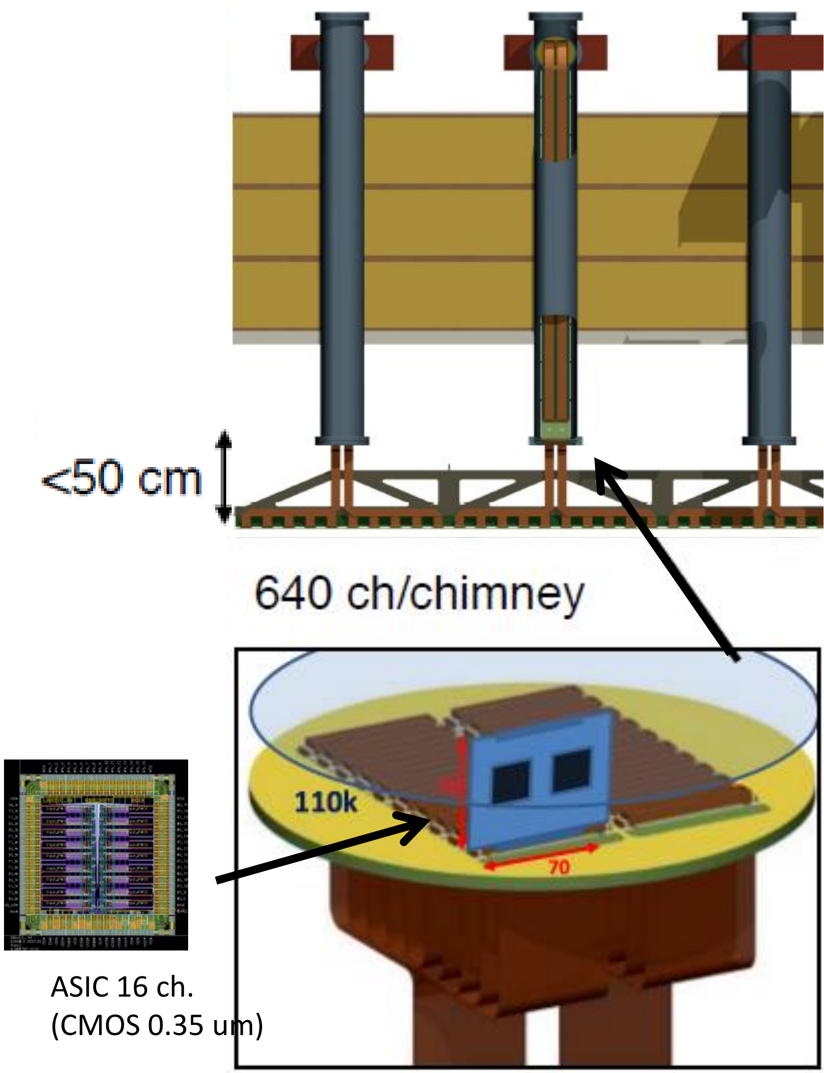

→ Large scalability 153600 channels for the DUNE 10 kton module at low cost

uTCA charge readout architecture

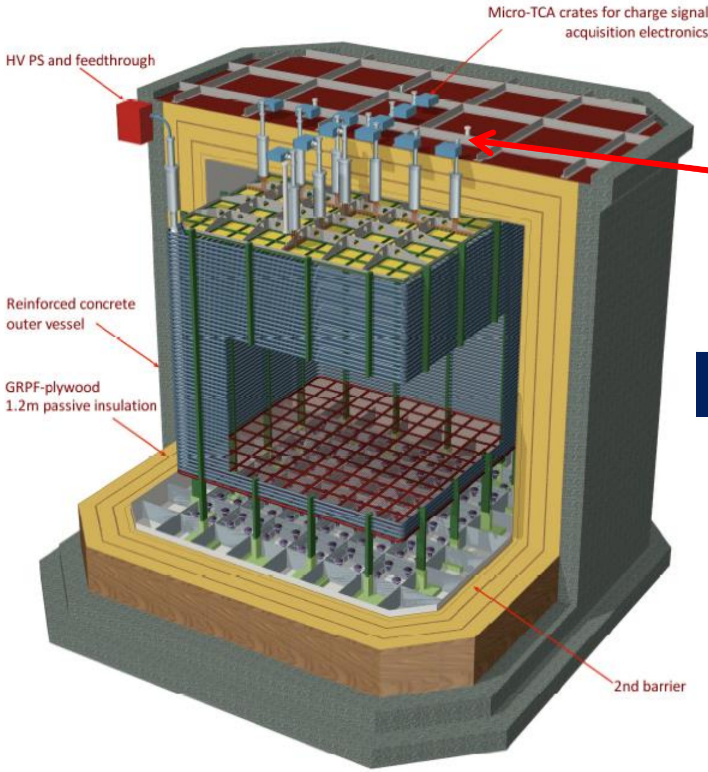

uTCA crates

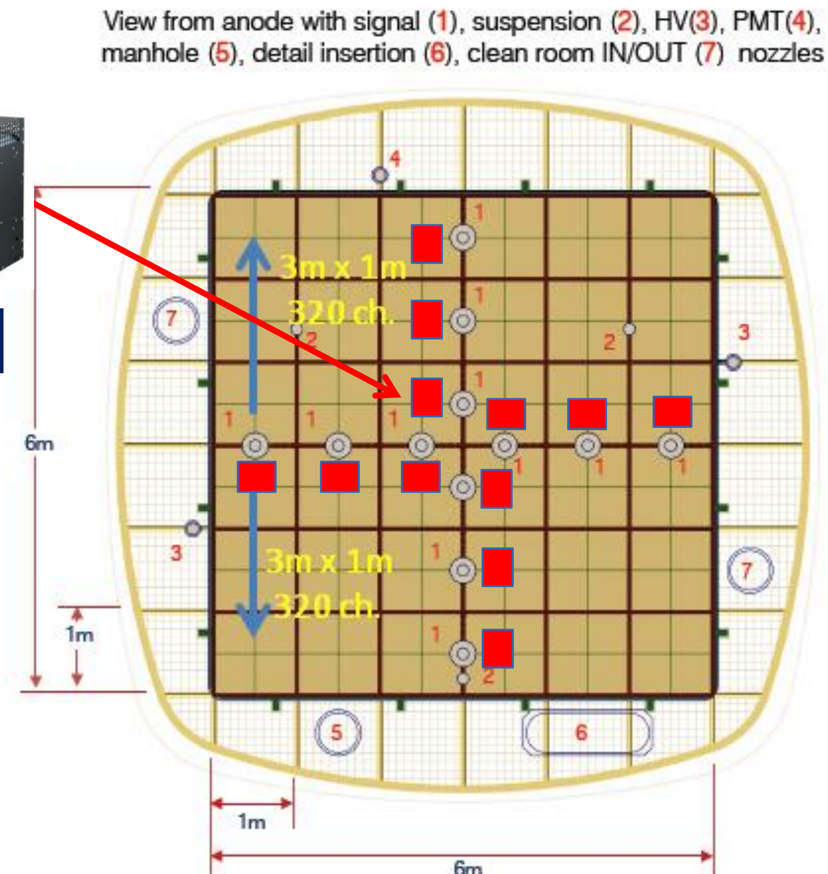

Readout in groups of 640 channels/chimney

top view

- 640 channels/chimney
- 1 uTCA crate/chimney, 10 Gb/s link
- 10 AMC digitization boards per uTCA crate, 64 readout channels per AMC board

→ 12 uTCA crates for charge readout + 1 uTCA crate for light readout

DAQ timing-trigger integration

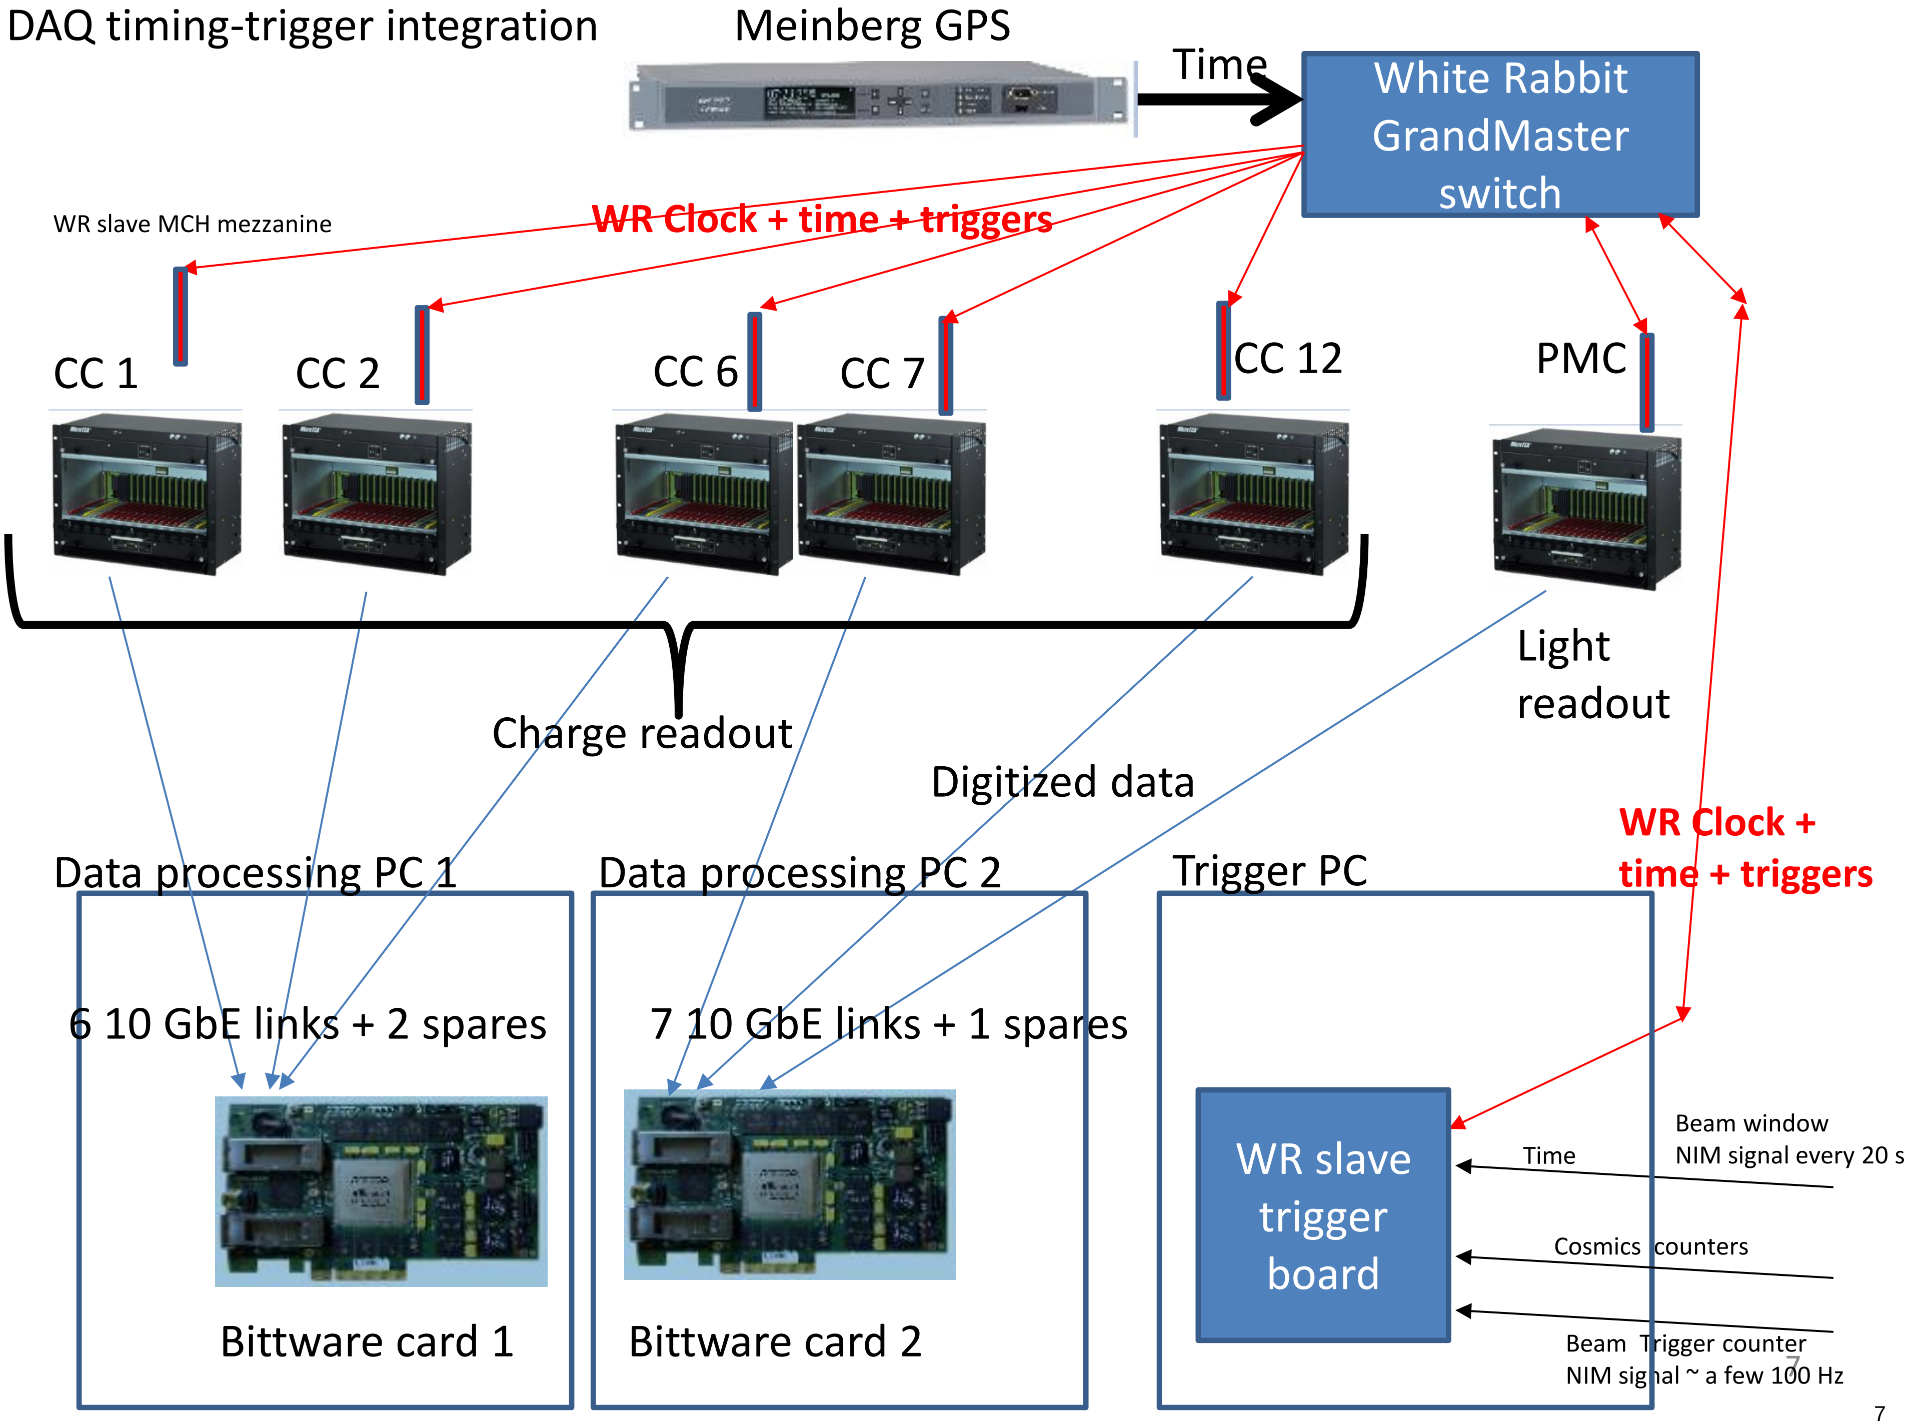

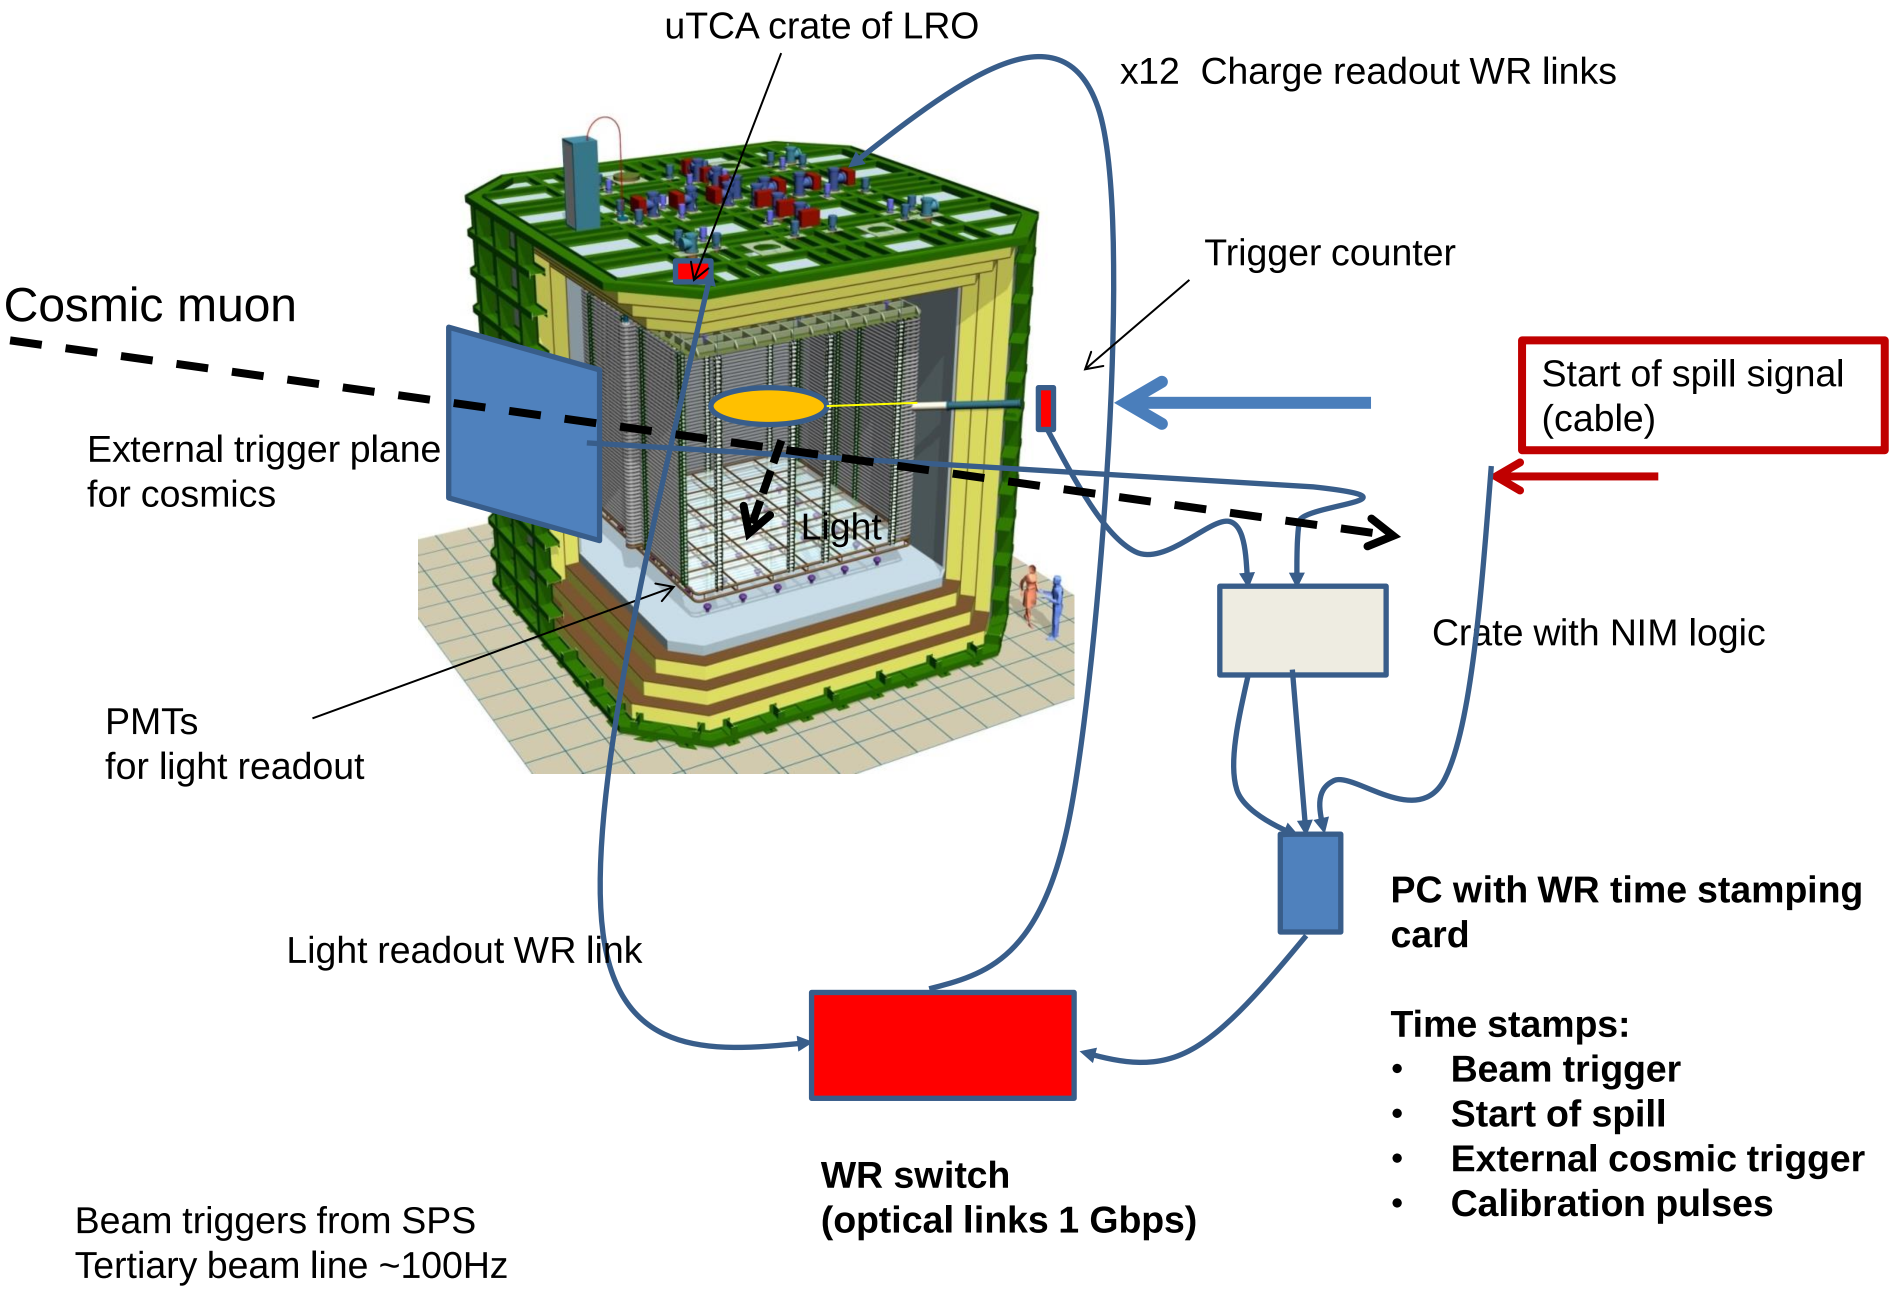
uTCA crate of LRO
x12 Charge readout WR links
Trigger counter
Cosmic muon
Start of spill signal
(cable)
External trigger plane
for cosmics
Light
Crate with NIM logic
PMTs
for light readout
PC with WR time stamping
card
Light readout WR link
Time stamps:
• Beam trigger
• Start of spill
• External cosmic trigger
• Calibration pulses
WR switch
(optical links 1 Gbps)
Beam triggers from SPS
Tertiary beam line ~100Hz

Beam time structure

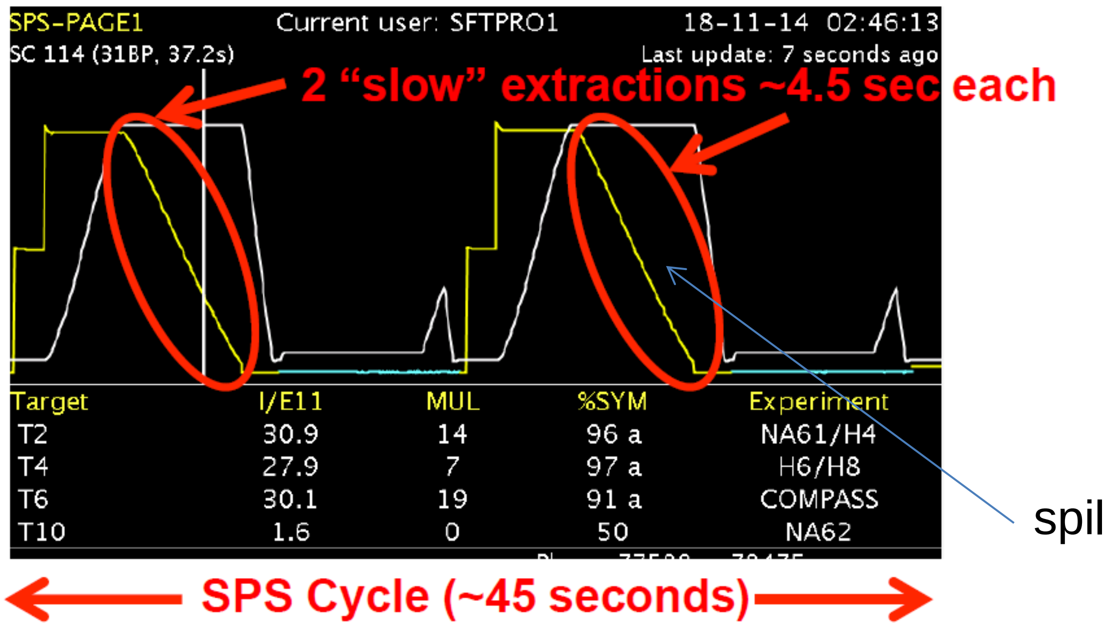

- 2 beam spills, of ~4.5 sec each, in a supercycle of ~ 45 sec (example, supercyle may vary with time depending on SPS operation mode)

- During spills: readout based on beam triggers from a beam trigger counter (charge+ continuous light readout in a window of +-4ms around the trigger for cosmic rays tagging)

- Out of spills: readout of cosmic rays (selected with external counters or internal light triggers) (charge+ self-triggering light readout)

For each beam trigger we can have on average 70 cosmics overlapped on the drift window after the trigger (these cosmics may have interacted with the detector in the 4 ms before the trigger and in the 4 ms after the trigger → chopped tracks, tapis roulant effect

In-spill cosmics in charge data

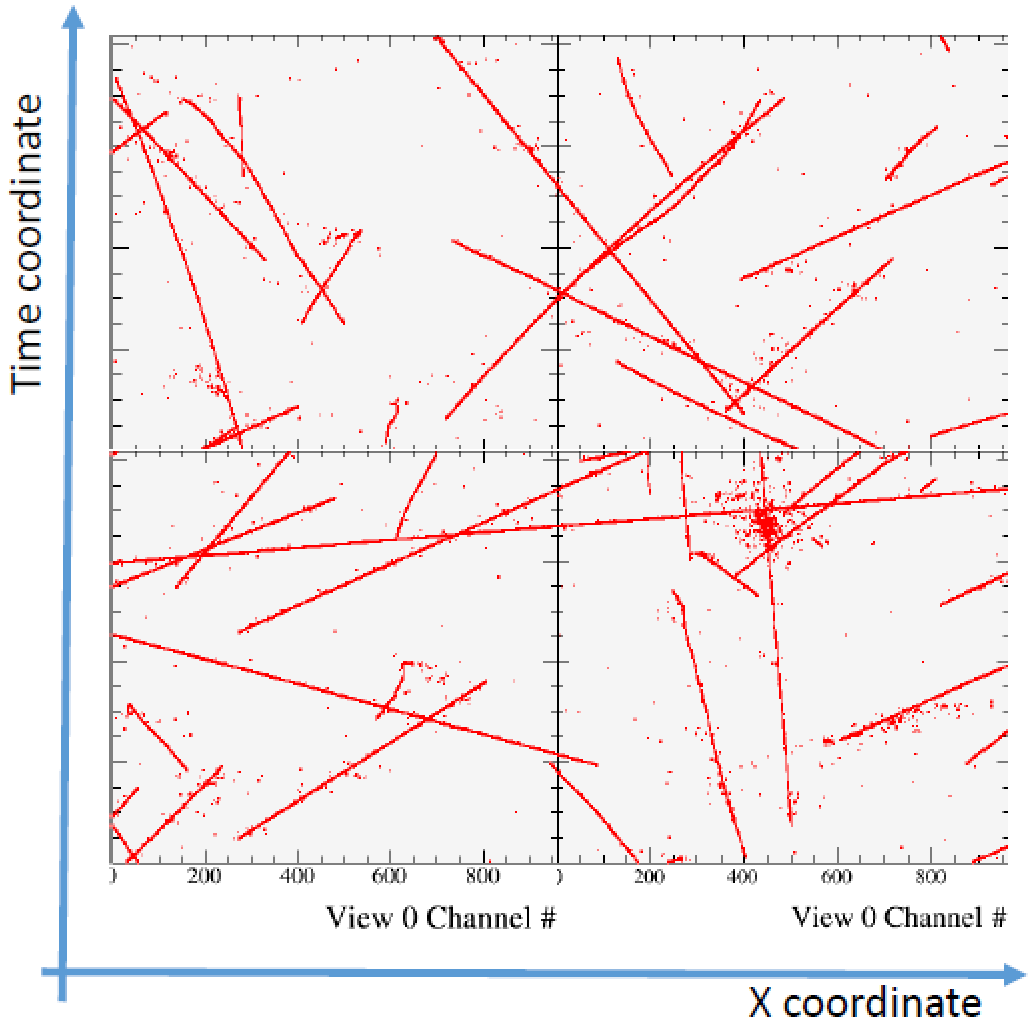

Example of cosmics only event in WA105 TPC (in one of the views)

- Red points are reconstructed hits
- TPC is readout in 4 3x3m^2 modules
- After track reconstruction:
  - Attempt to correlate found tracks with light data
  - Remove CR background from beam event
  - Select a subsample of long tracks for calibration purposes

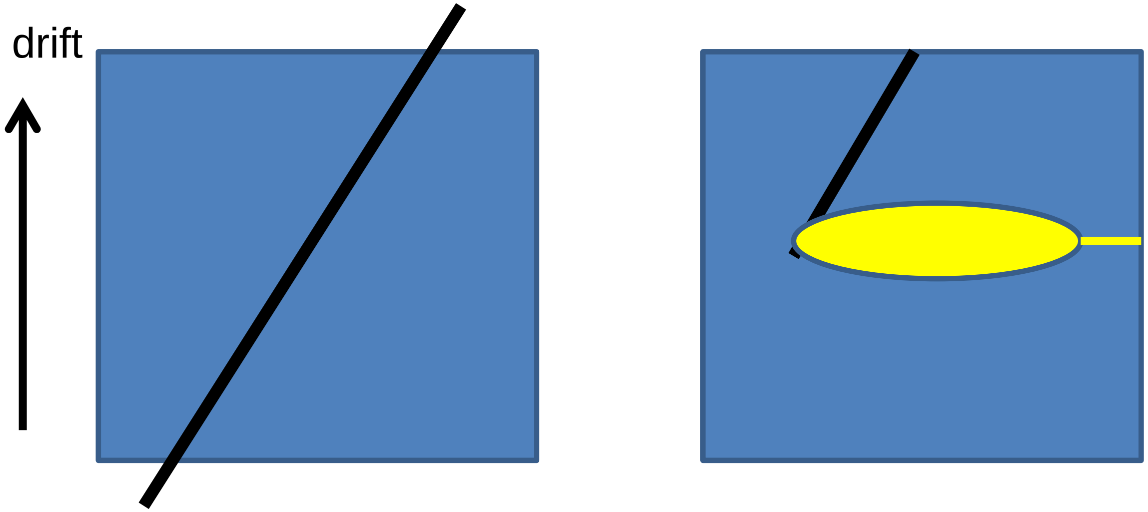

t=beam trigger - 2 ms

t=beam trigger → reconstructed event

The « tapis roulant » effect
+- 4 ms around the beam
trigger time

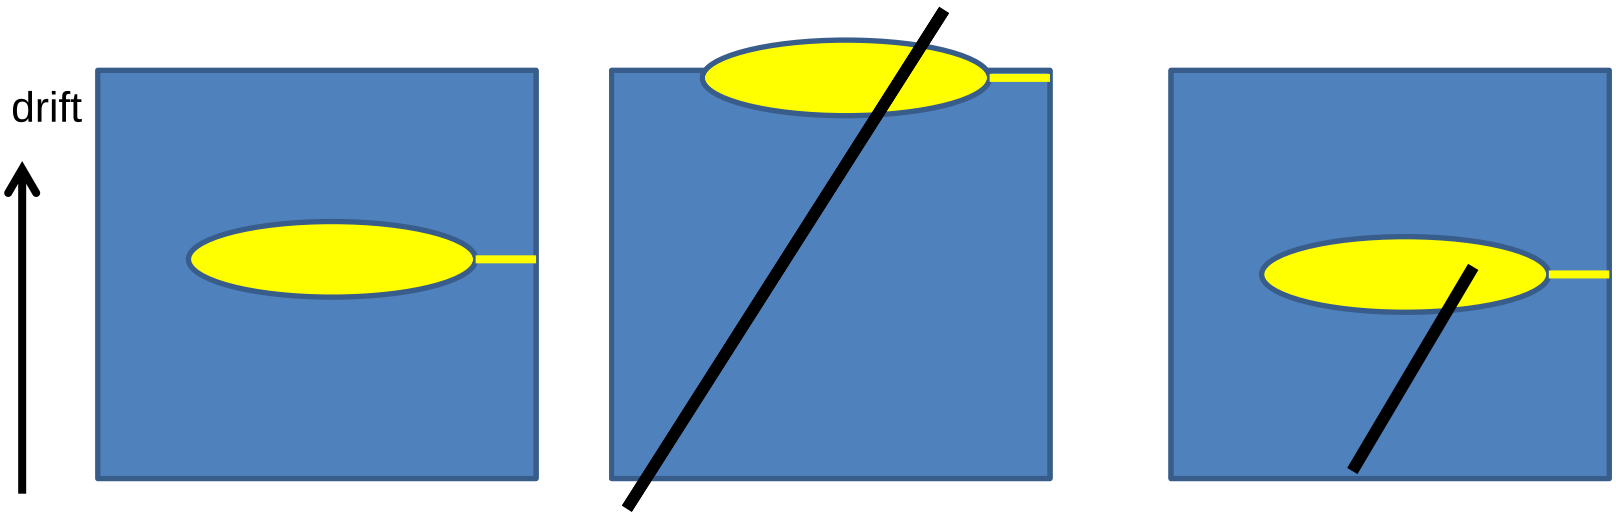

t=beam trigger

t=beam trigger + 2 ms

reconstructed event

→ During spills it is needed a continuous digitization of the light in the +-4 ms around the trigger time (the light signal is instantaneous and keeps memory of the real arrival time of the cosmics)

→ Sampling can be coarse up to 400 ns just to correlate to charge readout

Light readout electronics

Two modes of acquisition:

- External beam trigger to acquire ±4ms around the spill
- Internal trigger from PARISROC2 ASIC to acquire short time segments

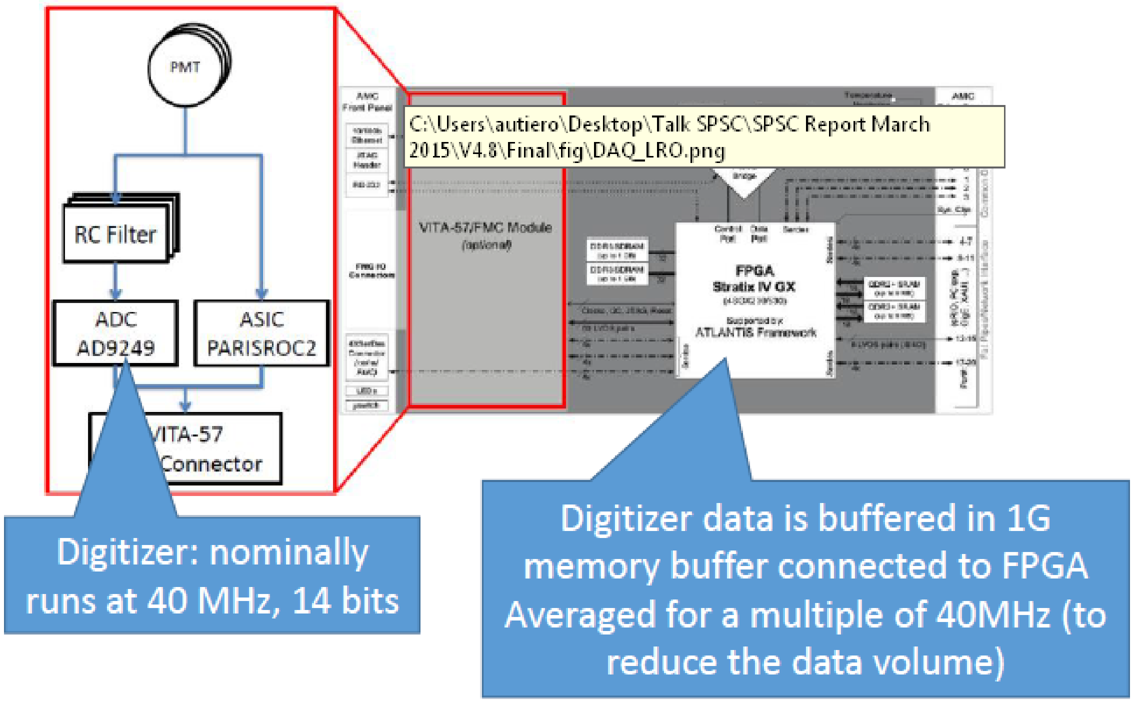

→ Sum 16 samples at 40MHz to get an effective 2.5 MHz sampling like for the charge readout

The LRO card has to know spill/out of spill
Out of spill it can define self-triggering light triggers when “n” PMTs are over a certain threshold and transmit its time-stamp over the WR

Data size

- A large fraction of the data are expected to be taken without zero skipping and exploiting loss-less compression and the system has been designed to support up to 100 Hz of beam triggers without zero-skipping and no compression

- 7680 channels, 10k samples in a drift windows of 4ms→ 146.8MB/events, No zero skipping

- Beam rate: 100Hz

- Data flow= 14.3 GB/s

- Light readout does not change in a significant way this picture (<0.5 GB/s)

⇩

→ Integrated local DAQ bandwidth on the “20 GB/s scale”

1 day of data ~ 1000TB (no zero skipping, no compression) (assuming 50% beam duty cycle + contingency)

Huffman lossless compression can reduce the non-zero-skipped charge readout data volume by at least a factor 10 (see Slavic’s presentation, S/N for double phase ~100:1, small noise fluctuations in absolute ADC counts)

DAQ architecture

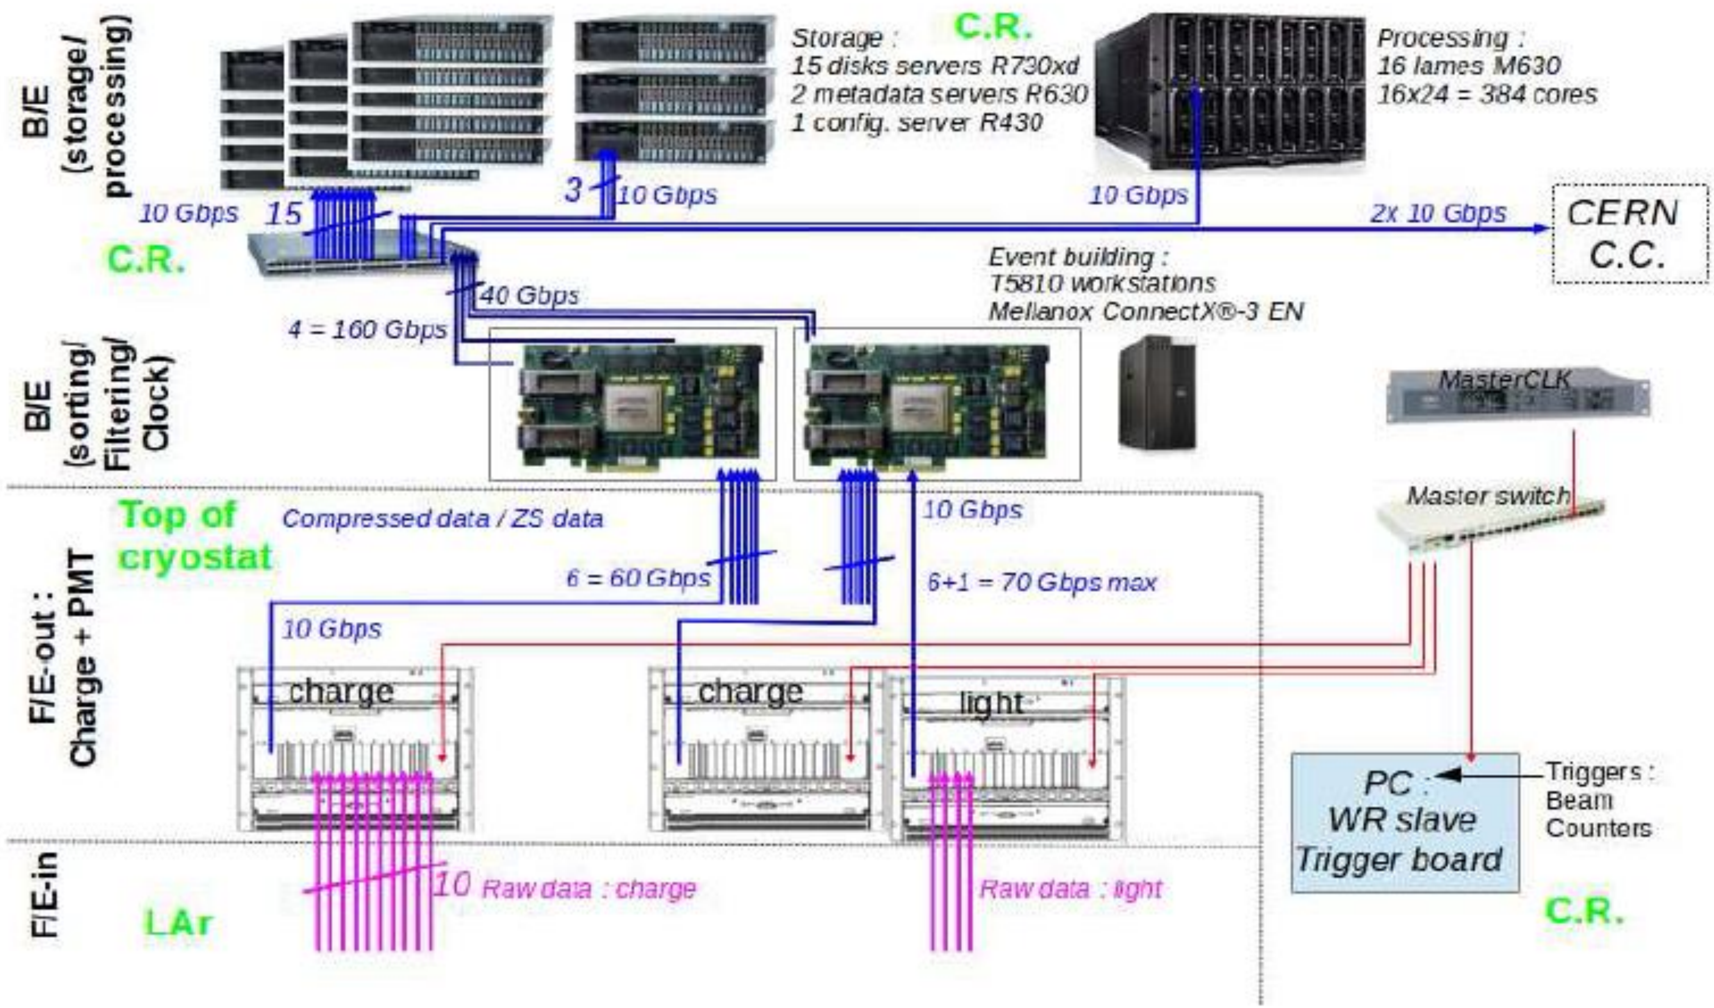

C.R. stands for Counting Room

Total data volume

TABLE XIX: *Requirements for particles and their momenta. The particle rate here is the rate within a spill, regardless of the spill length, slow extraction is assumed.*

Type	Momentum [GeV/c]	Rate [kHz]	Total	Time est. [hrs]
Muon tracks				
$\mu^{+/-}$	0.8, 1.0, 1.5, 2.0, 5.0, 10.0, 20.0	0.1	$5 \times 10^6 \times 14$	200
Shower reconstruction				
$\pi^{+/-}$	0.5, 0.7, 1.0, 2.0, 5.0, 10.0, 20.0	0.1	$5 \times 10^6 \times 14$	200
e	0.5, 0.7, 1.0, 2.0, 5.0, 10., 20.0	0.1	$5 \times 10^6 \times 7$	100

- WA105 TDR: 175 M triggers
- If totally stored in non-zero-skipped, lossless compression format
  (assuming Huffman, factor 10 compression: 15MB/event) → 2.4 PB

  (+ cosmic runs and technical tests)

Following discussions we had with the IT people at CERN this total data volume of 2.4 PB does not present any problem on the CERN scale

- Requested link from online-storage to CERN computing division at 20 Gbps, compatible with 100 Hz non-zero-skipped, Huffman compressed (factor 10) data flow

Online storage and processing

Online storage

- The B/E boards are hosted into Event Building Workstations transferring the data at 40 Gbps to the local storage/processing system (two links per WS)
- A standard 10/40 Gbps switch is used to interconnect all these elements and to send the data to be permanently stored to the CERN EOS/CASTOR system.
- The online storage/processing system fulfills the CERN requirements of a few days equivalent data storage local buffer for the experiment DAQ system.
- The online processing/storage system includes 15 disks servers (Object Storage Servers – OSS, DELL R730xd for the reference design) for a total capacity of the order of 1 PB, 2 redundant MetaData Servers (MDS, DELL R630), 1 configuration server (DELL R430) and 1 processing farm (DELL M1000e with 16 M630 blades leading to 384 cores).
- Final system software configuration under evaluation with data writing/access tests at the CCINP23 (LYON computing centre)
- Software for data transfer at CERN based on IT standards EOS/CASTOR scripts

Online processing

- Disentangling of cosmic ray overlaps ~70 tracks per drift window (overlapped to beam events in a +-4ms window around the beam triggers)

In-spill cosmics in charge data

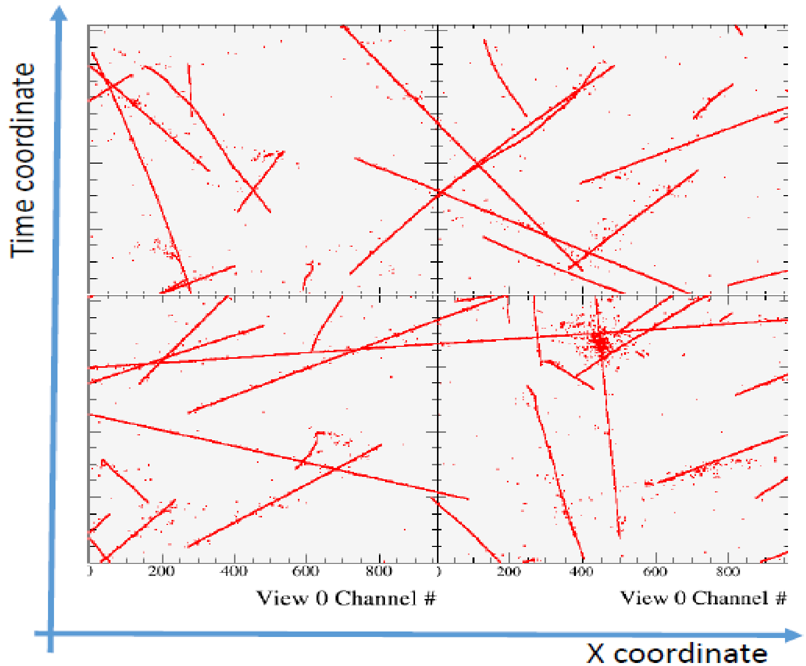

Example of cosmics only event in WA105 TPC (in one of the views)

- Red points are reconstructed hits
- TPC is readout in 4 3x3m^2 modules
- After track reconstruction:
  - Attempt to correlate found tracks with light data
  - Remove CR background from beam event
  - Select a subsample of long tracks for calibration purposes

8

- Use of cosmic rays for online purity analysis and gain monitoring (1-2 good tracks/event)

- Data reduction (exercise also for DUNE data online processing facility)

- The bulk of detector raw data processing will be done at CERN
- A copy of the processed raw data will be copied to LYON computing center (CCIN2P3), the transfer procedure has to be discussed and set-up with CCIN2P3 staff.
- Monte Carlo productions are centralized at the CCIN2P3
- <u>WA105 development environment is centralized at the CCIN2P3:</u>

→ all software is installed and running at the CCIN2P3

→ a complete working environment is set up for all WA105 users (environment variables are defined at login for different shells, common tools and development space)

→ a preinstalled, running version of WA105 code is available to users together with examples showing how to access simulation and reconstruction output root files

→ all source codes are also available on a SVN server (production version and ongoing developments)

→ the WA105 simulation and reconstruction environment is based on QSCAN (see Slavic's presentation)

→ At the CCIN2P3 is also available a complete LArSoft development environment mirrored from FNAL (updated every week)

Commissioning of the online processing scheme

- The configuration for the online storage and processing can be validated using the 3x1x1 prototype

- The 311 represents ¼ of the 666 as far as the charge readout is concerned : 3 µTCA crates, synchronized by WR, 1 B/E card, 1 switch 10/40 Gbps, 1 reduced version of the online storage/processing system (equivalent to ~100TB).

- The detector will only take cosmic data under different conditions for the detector debugging /commissioning

- The goal for the online storage system is to validate the configuration, the distribution of the data, the choice of the O/S, of the filesystem (NFS, beeGFS, lustre, xrootd etc), the redundancy of the MetaDataServer and the connection to the CERN EOS/CASTOR system.