Online track summary

Vivian O'Dell, Andrew Norman, Tassos Belias, Niko Neufeld

Statistics

- 53 papers
- 25 posters, 28 – 2 talks (two were victims of the government shutdown)
- Good attendance in the parallel Online sessions

Very high quality throughout

Not mentioned != Not interesting

LHC status

- Excellent performance
- Better use of resources

Bring the data on!

Challenging days for Trigger

LHC had an extremely successful operation and luminosity ramp up!

- Luminosity increase
  -- Changes in trigger had to follow six orders of magnitude of changes in luminosity
- Pileup increase
  -- In-time (overlapping p-p collision events within a same BC) and out-of-time (from previous BC)
  -- Luminosity increase in the past years mostly by increase of bunch luminosity
  → Larger pileup than design
  -- Challenge of trigger to keep efficiency and rejection stable in high pileup conditions

Peak Luminosity

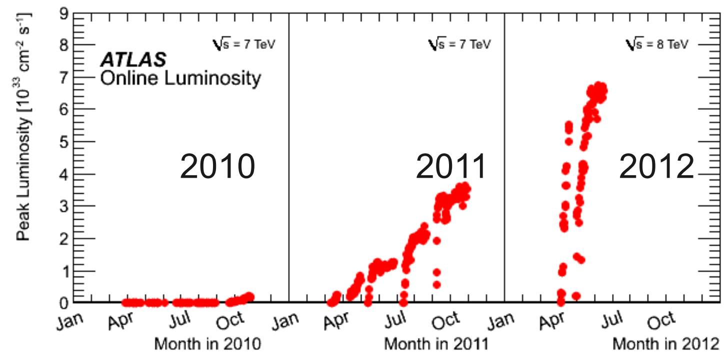

Peak Interactions per crossing

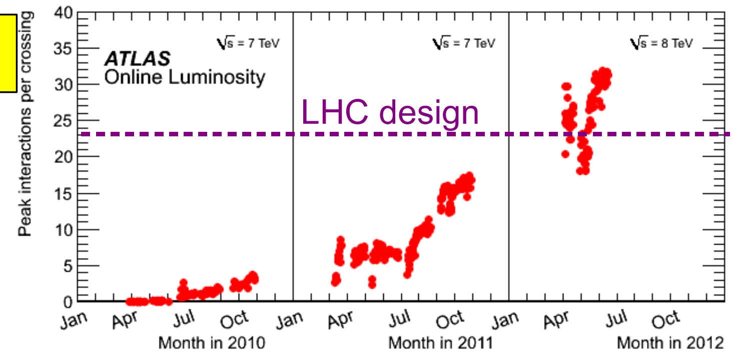

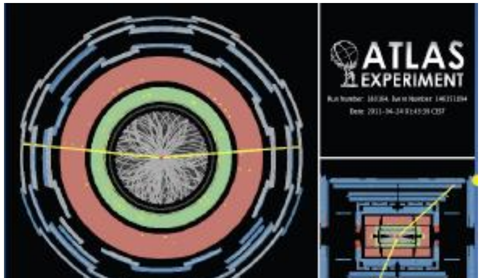

11 vertices with 1 Z→μμ

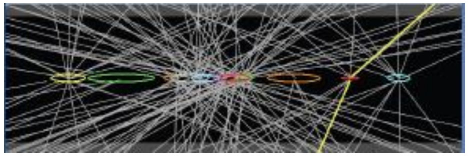

#503 Nagano-san

The menu, please!

Trigger Menu, Strategy and Evolution

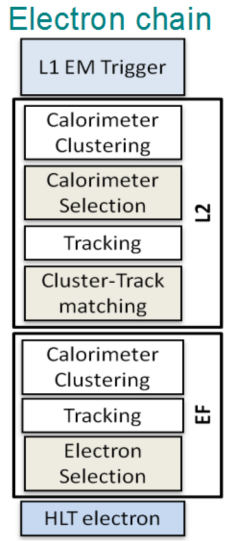

- Chain: one full L1→L2→EF selection sequence
  - -- E.g.: "e24i" chain = single electron trigger with $E_T > 24$ GeV
- Menu: full set of chains and prescale factors
  - -- Prescale factor = reduction factor to issue a trigger
  - -- A typical menu contains ~500 chains, to meet a large variety of physics goals at LHC, and also to contain sufficient supporting triggers (for background estimations, detector performances etc.)

- For 2011-12: managed to run with just 3 base menus for p-p
  - -- Some chains designed as extras in each menu dropped as luminosity increases to keep bandwidth under control
  - -- Avoid complication in physics analyses due to frequent trigger changes

Keep it simple, not stupid

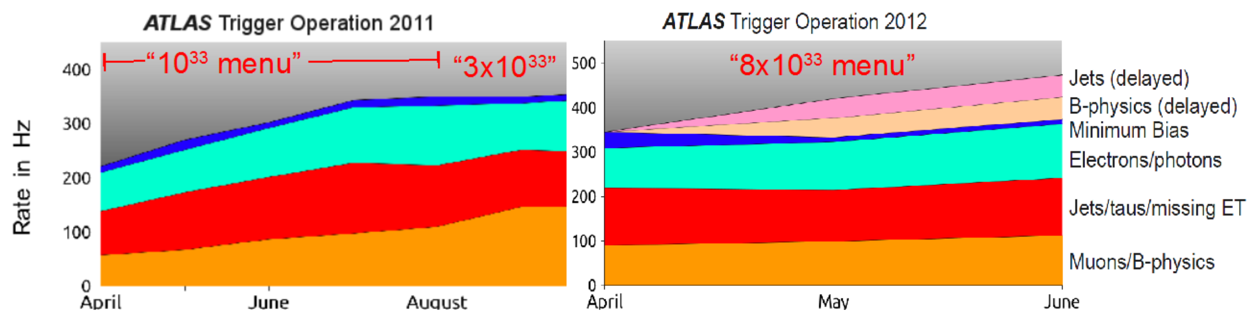

5

Towards 100% efficiency

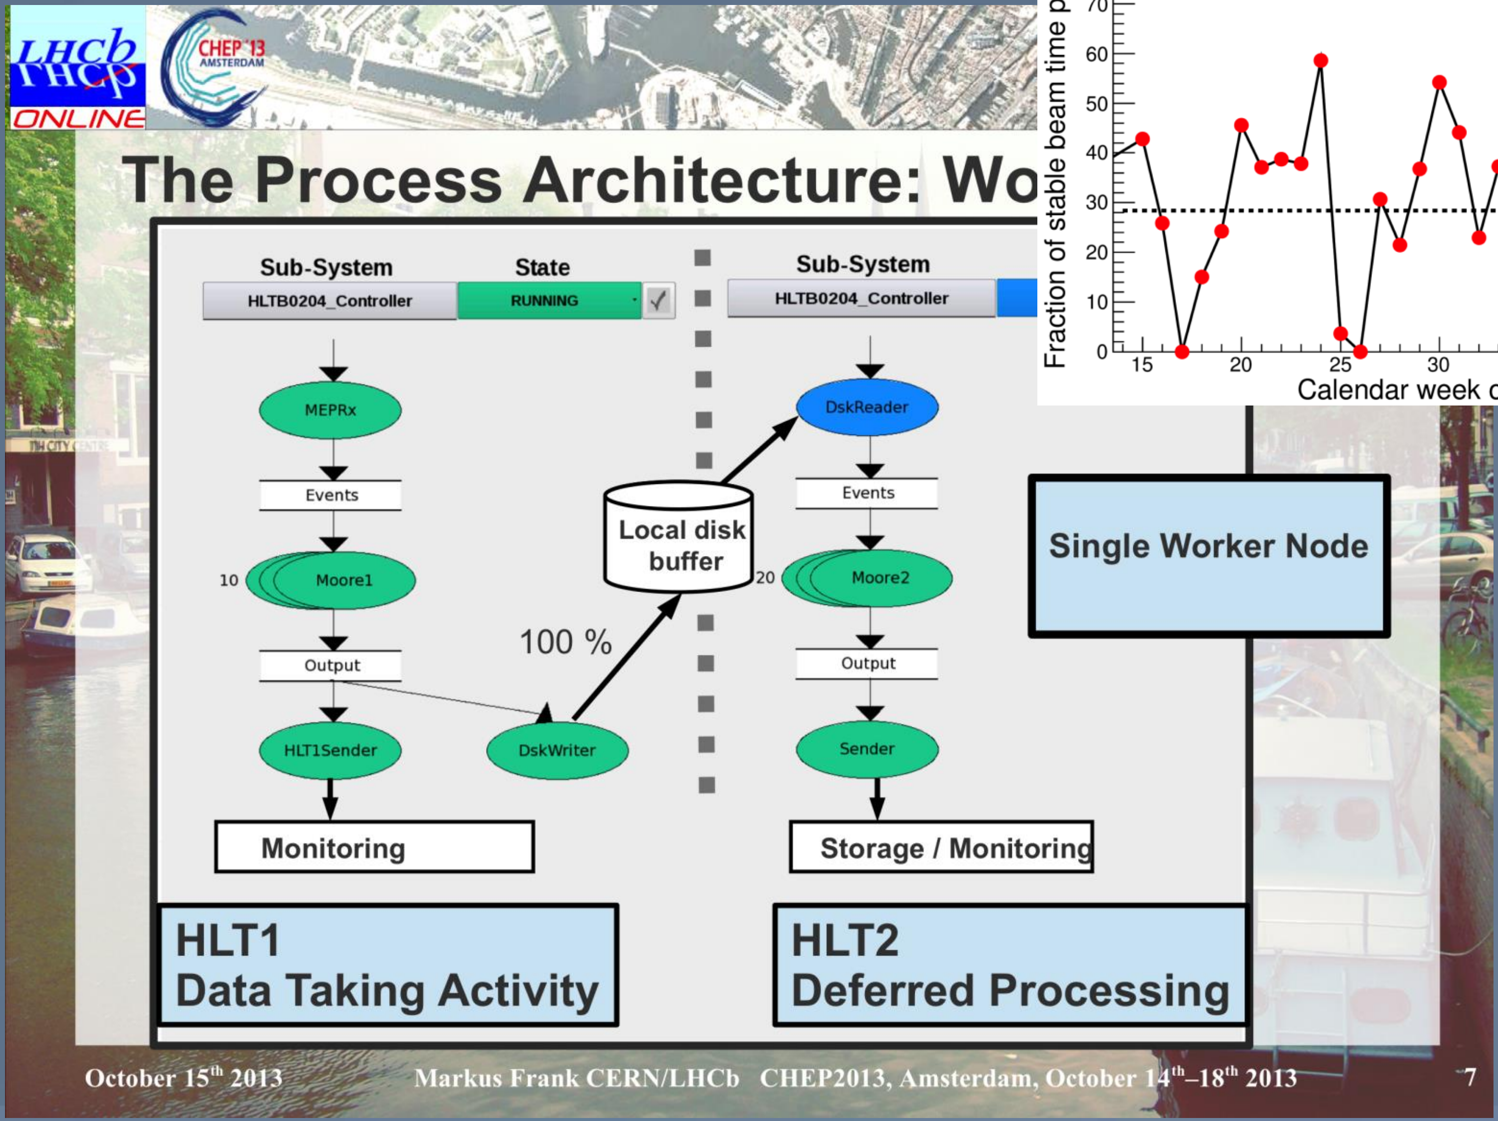

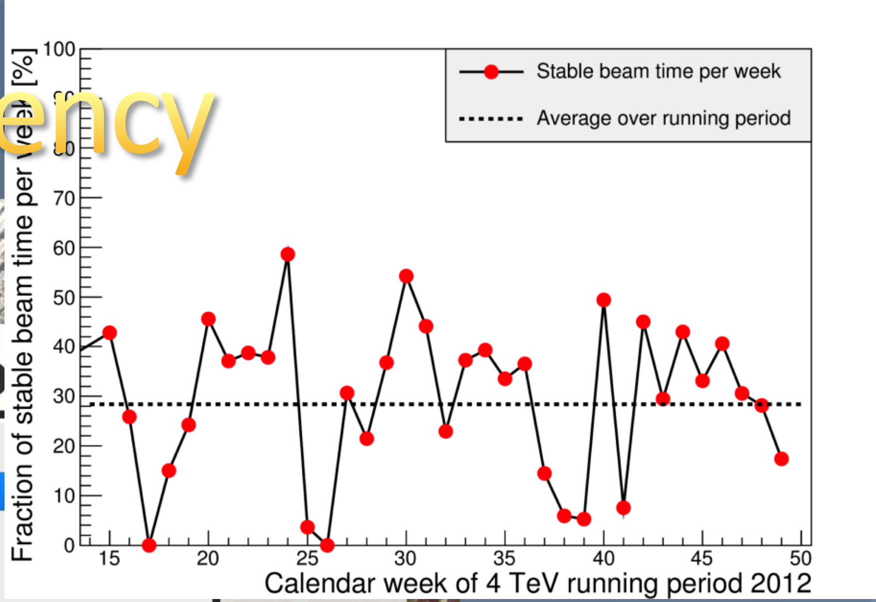

#35
M. Frank

ATLAS before … (run 1)

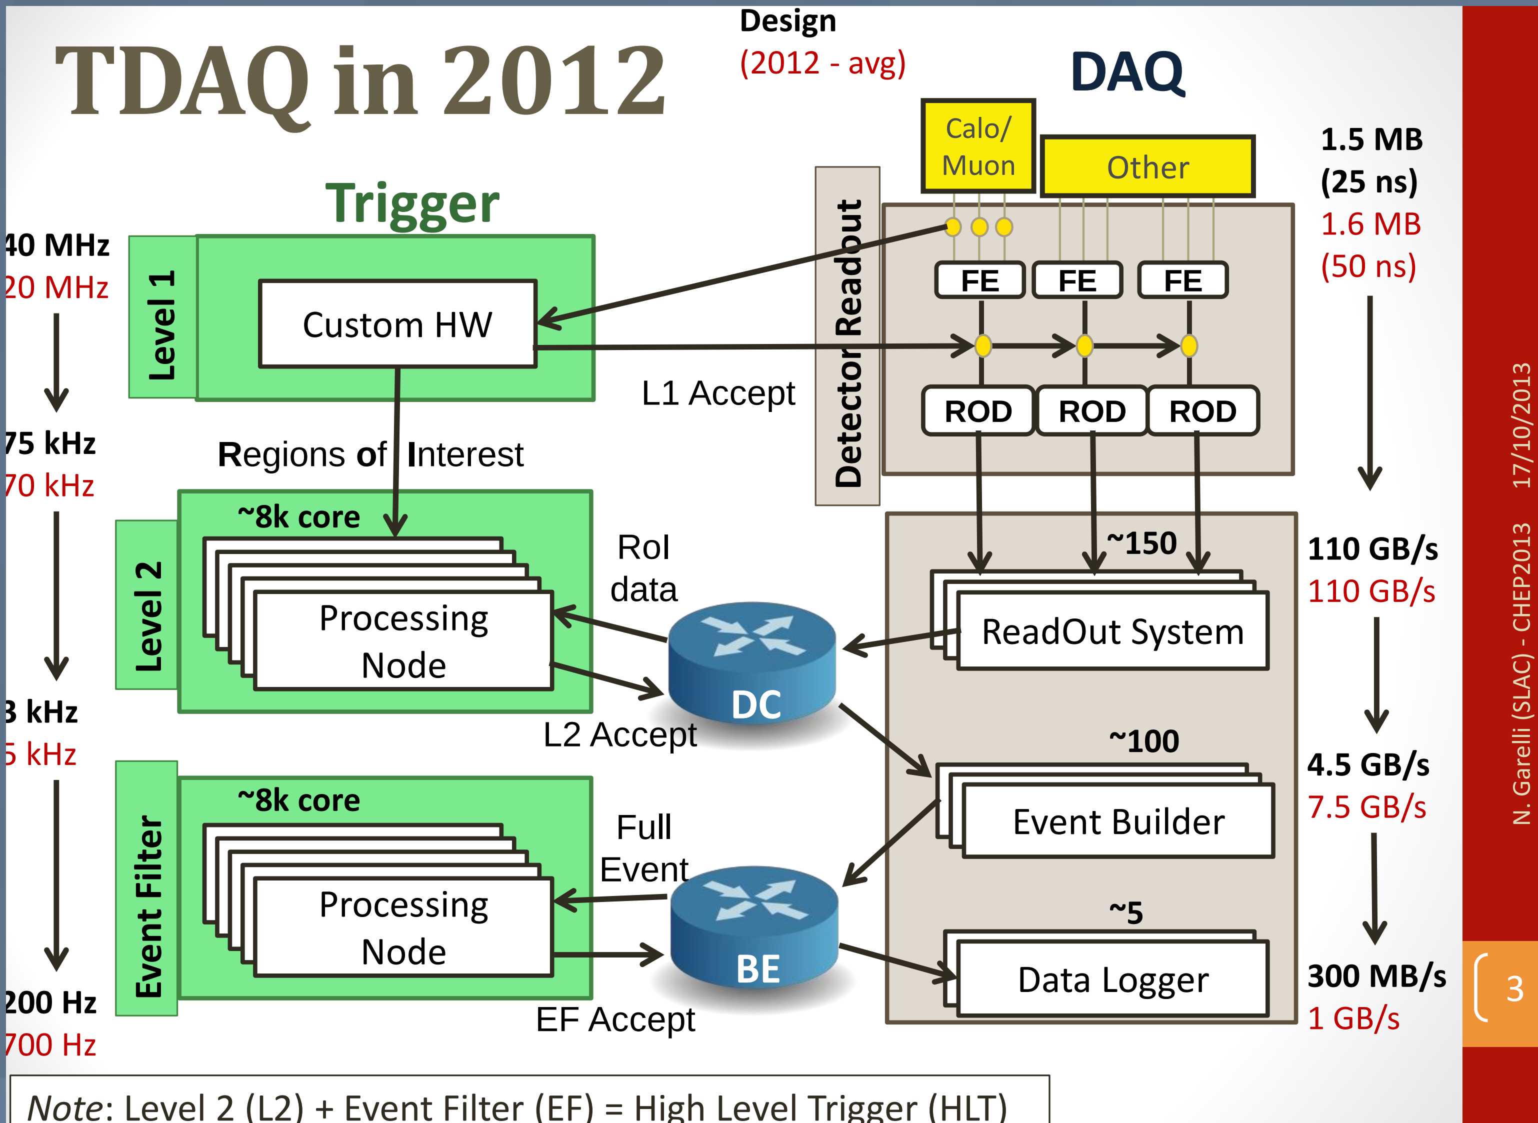

and after … (run 2)

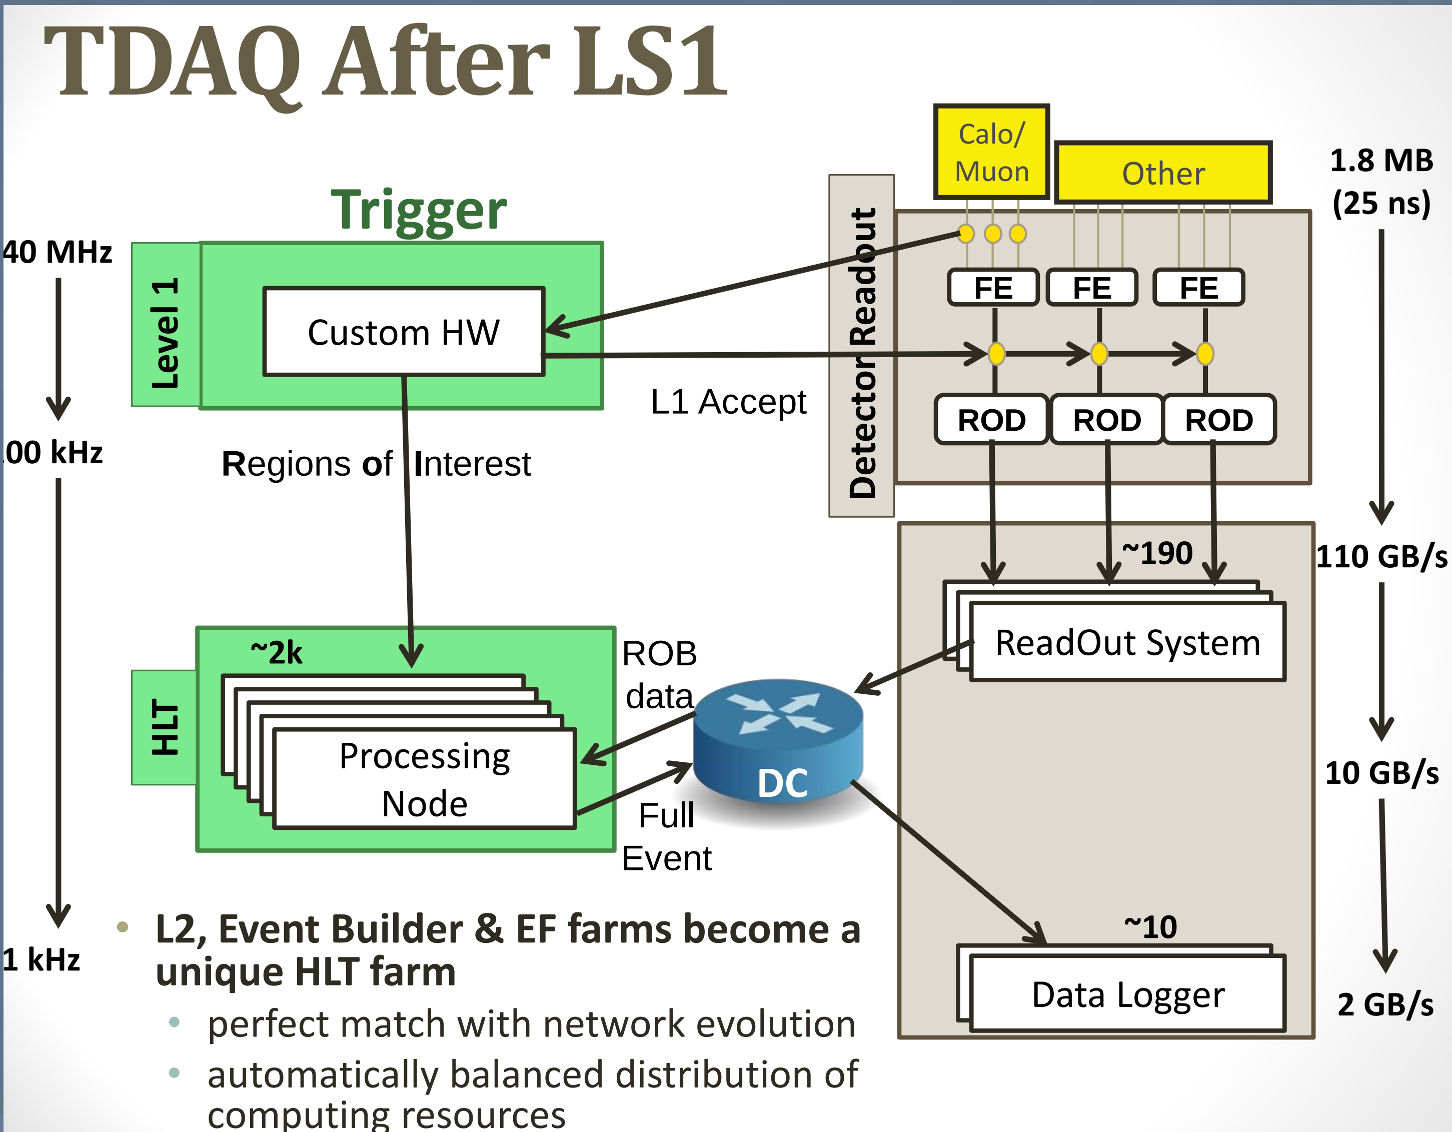

- More powerful networks and servers lead to simplified architectures
- Improved HLT software (fork-ing) mitigates memory pressure

LHC future

- LHC future is
  - immediate (Run 2)
  - soon (Run 3) → ALICE and LHCb go triggerless
  - coming (Run 4) → HL-LHC (upgrades for ATLAS and CMS) (no talks on this … yet – current work focusing on FE-electronics)

More and more bytes / second

- "Low-cost" 32 Tbit/s DAQ
- Based on InfiniBand & Ethernet
- FPGA receiver cards in standard server PCs

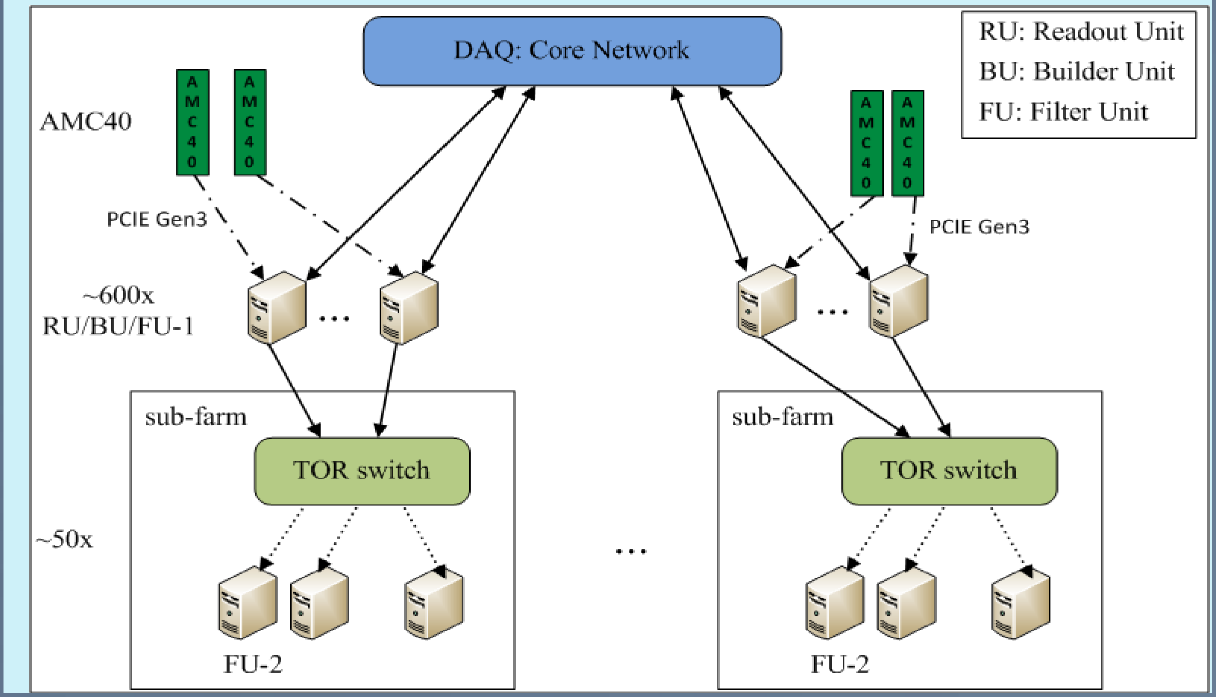

#54 G. Liu

O2: ALICE after LS2

- Continous readout of the TPC
- Event-building based on time-slices

Detector	Input to Online System (GByte/s)	Peak Output to Local Data Storage (GByte/s)	Avg. Output to Computing Center (GByte/s)
TPC	1000	50.0	8.0
TRD	81.5	10.0	1.6
ITS	40	10.0	1.6
Others	25	12.5	2.0
Total	**1146.5**	**82.5**	**13.2**

Full Online/Offline integration

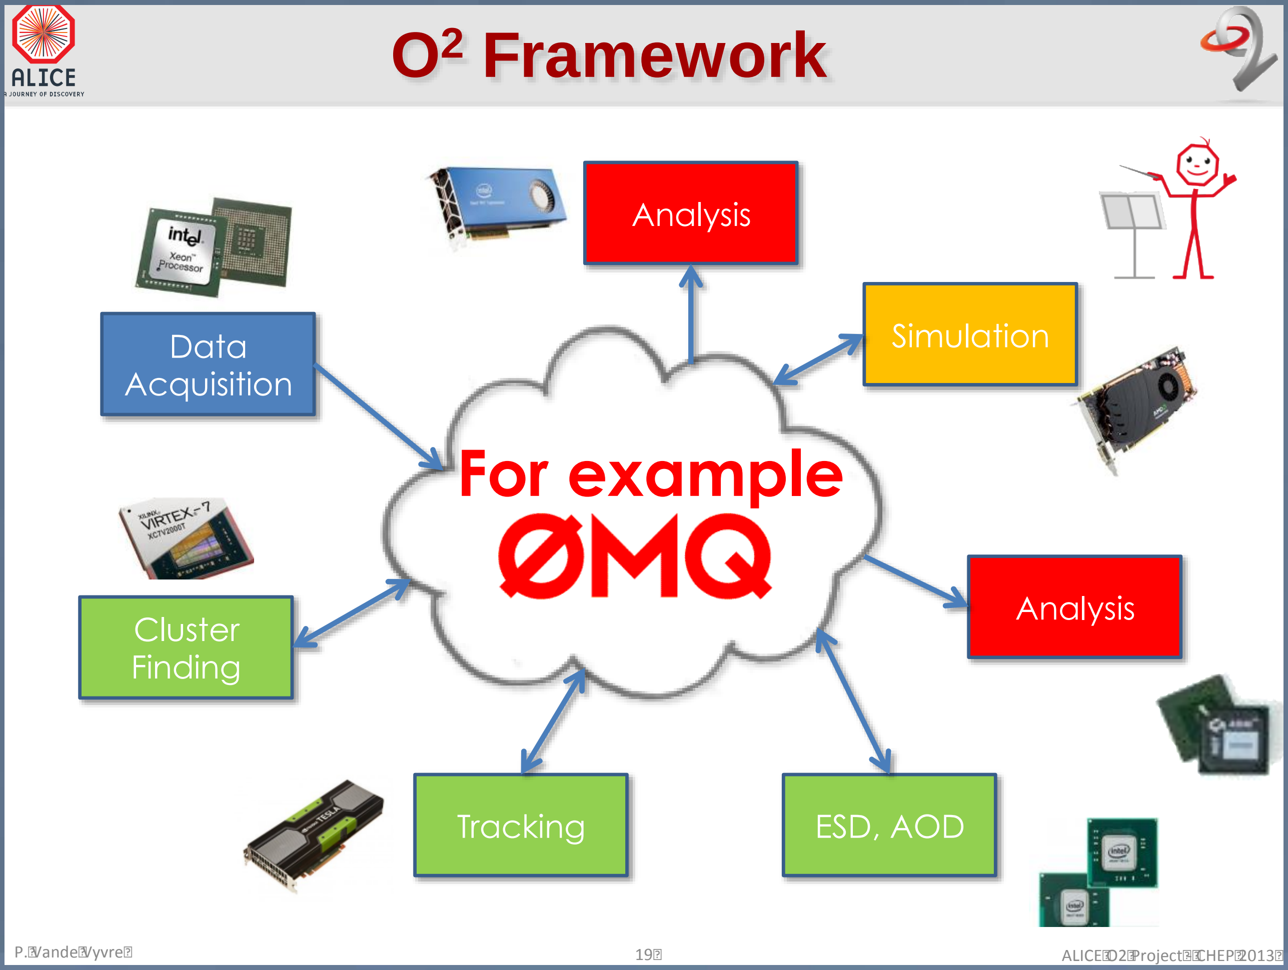

- Online resources very important
- Increase duty cycle (c.f. also LHCb deferral, similar initiatives in CMS, ATLAS

Technologies old and new

- FPGAs
- GPGPUs
- Xeon/Phi

Industry standards on FPGA

FEROL Introduction

- **Front-End Readout Optical Link (FEROL)**
  - Interface between custom and commercial hardware/network
  - Replace Myrinet NIC with custom FPGA based NIC card
- **Input:**
  - Legacy S-link input via FRL
  - SlinkXpress interface
    - 2x optical 6 Gbit/s interface
    - 1x optical 10 Gbit/s interface

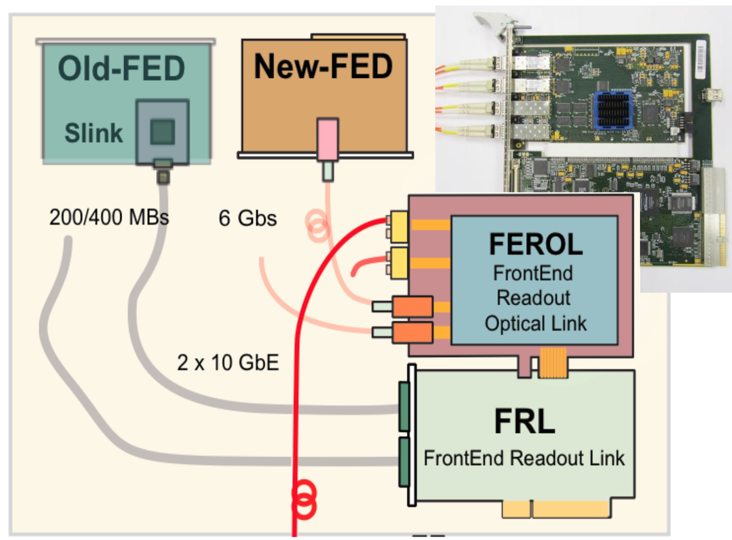

- **Output:**
  - Optical 10 Gbit/s Ethernet link
  - Optional second 10 Gbit/s Ethernet link
  - Runs a **standard protocol**: **TCP/IP** over 10Gbit/s Ethernet

#87
P. Zejdl

Ethernet, PCs – with FPGA support

#7 P. Gros 10-GigE for ATLAS Calo RO

A. Gianoli #33

Track#1 (Onli

L0TS: how to do it?

- use high performance PC to run selection algorithm
- use fpga board to handle fixed delay output to TTC (needs real-time)
- avoid memory-to-memory copy: use fpga board to collect primitives (udp packets) and put them into PC ram

HW used:

- core i7 920   2.67 GHz
- core i7 3930K 3.2 GHz
- Terasic DE4-230 board (Altera StratixIV, PCIe Gen2 x8, 4 eth ports)

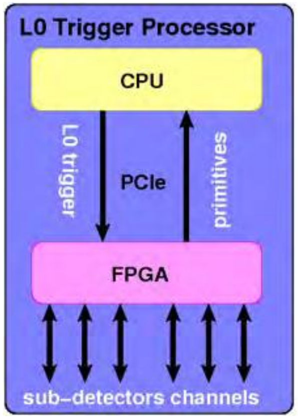

GPGPUs in L0, L1, L2, L3, …

#48 R. Ammendola
“How to get the data in/out quickly

also
#297 GPU Enhancement of the High Level Trigger to extend the Physics Reach at the LHC, H. Valerie

Many-core tracking (CDF example)

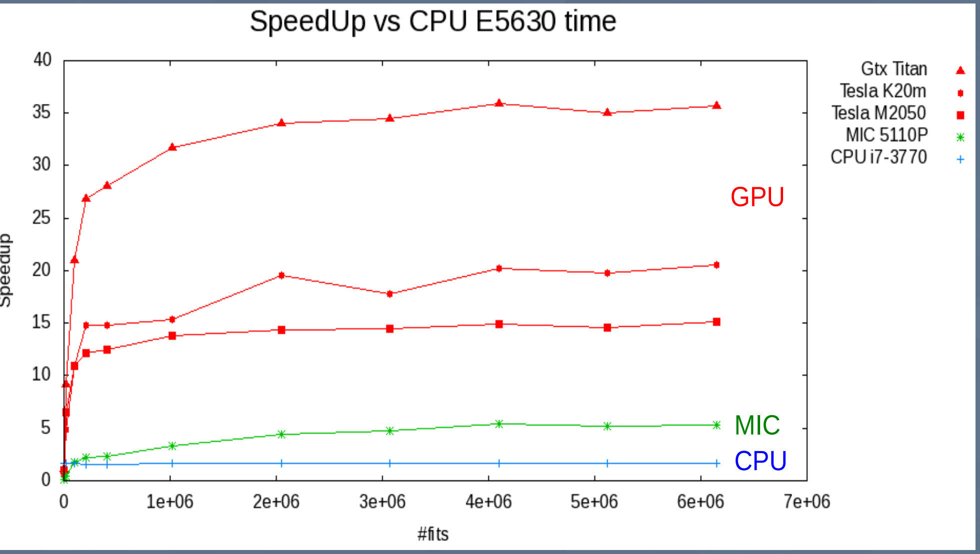

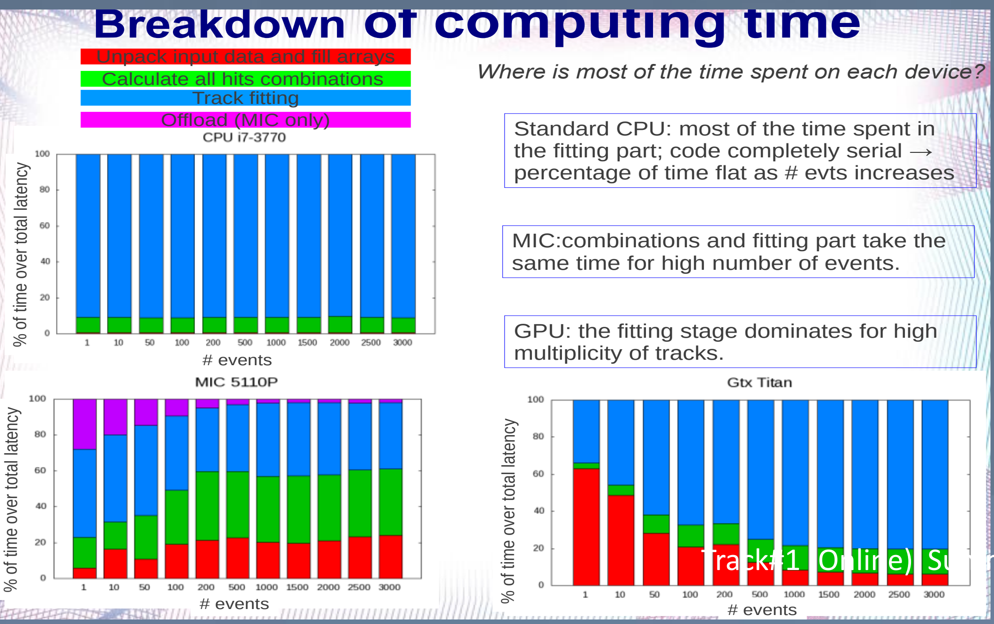

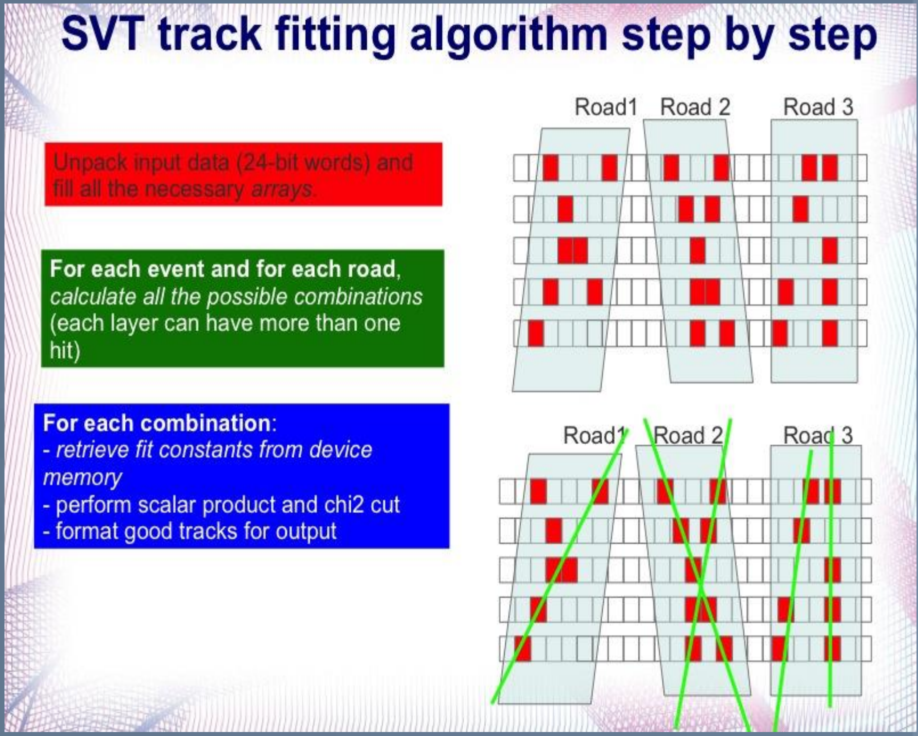

#78 S. Amerio

#430 A. Lonardo

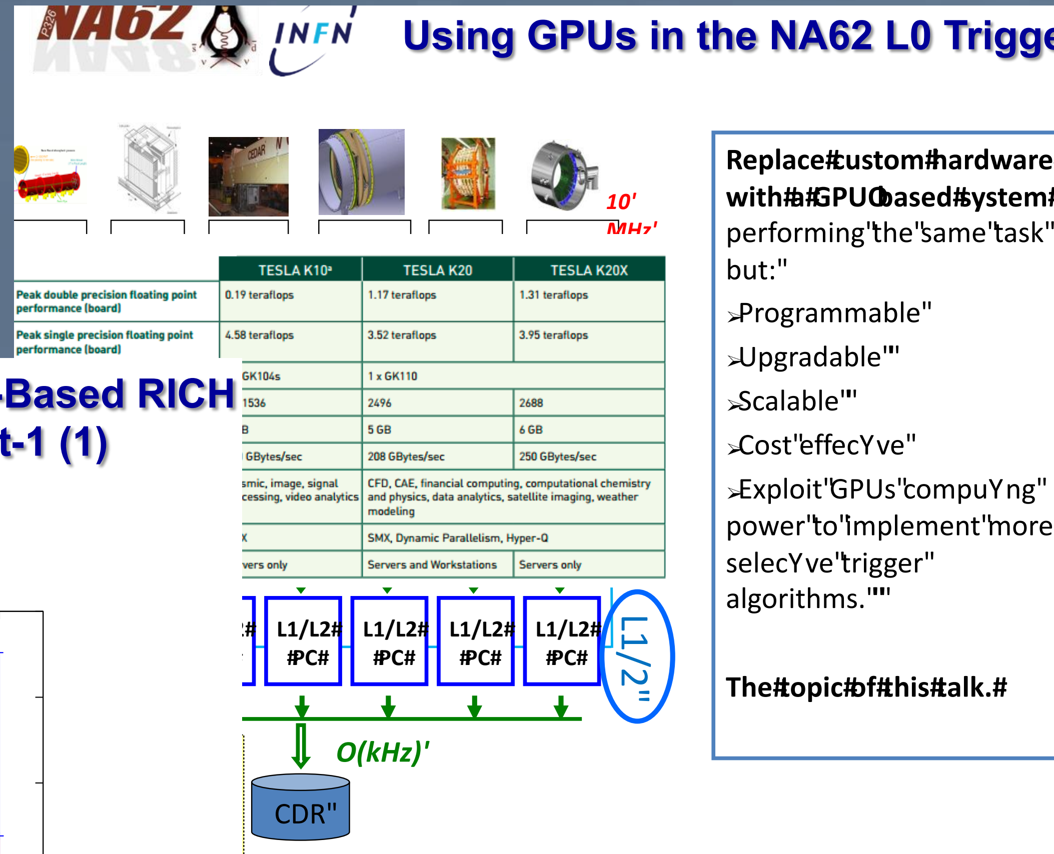

	TESLA K10ᵃ	TESLA K20	TESLA K20X
Peak double precision floating point performance (board)	0.19 teraflops	1.17 teraflops	1.31 teraflops
Peak single precision floating point performance (board)	4.58 teraflops	3.52 teraflops	3.95 teraflops
	GK104s	1 x GK110	
	1536	2496	2688
	B	5 GB	6 GB
	GBytes/sec	208 GBytes/sec	250 GBytes/sec
	smic, image, signal cessing, video analytics	CFD, CAE, financial computing, computational chemistry and physics, data analytics, satellite imaging, weather modeling	
	X	SMX, Dynamic Parallelism, Hyper-Q	
	vers only	Servers and Workstations	Servers only

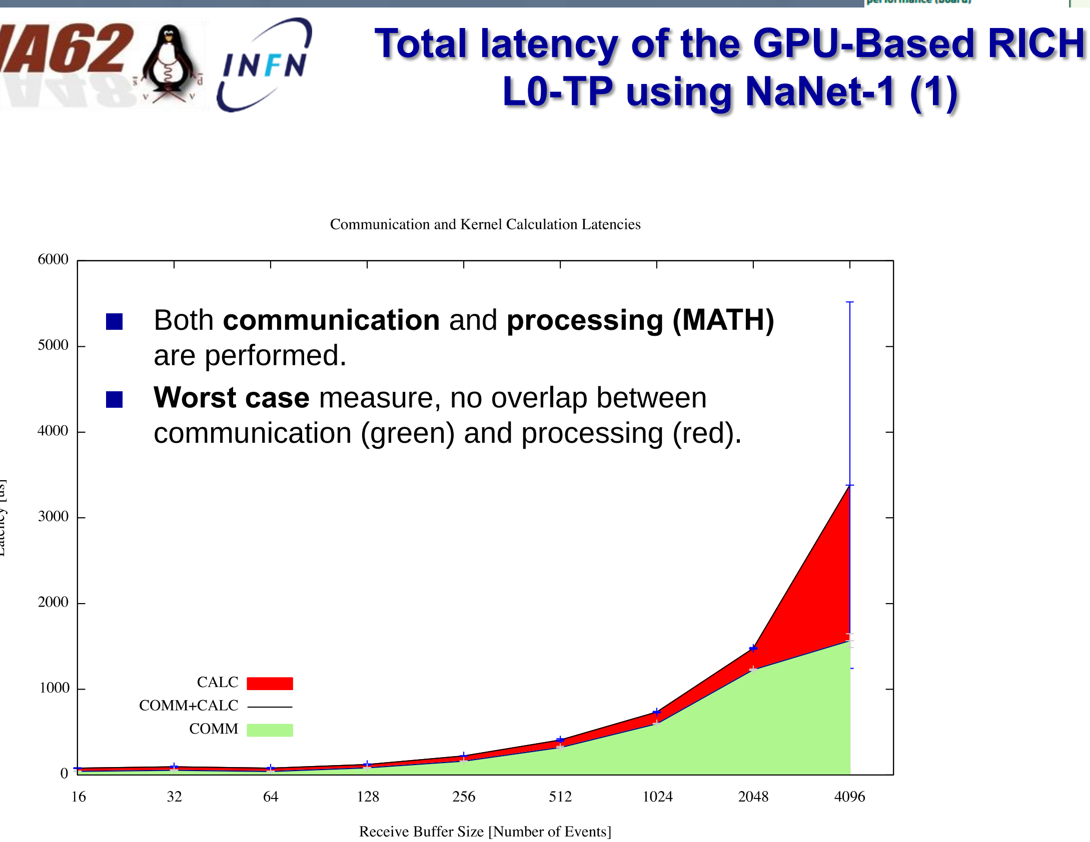

Controls (run & slow)

- Controls crucial for success of experiments
- Raw “size” requirements not growing as fast as in DAQ systems, BUT
  - 24/7/365 (often)
  - crucial for overall expt. efficiency
  - large, sophisticated software systems

Controls for Telescopes

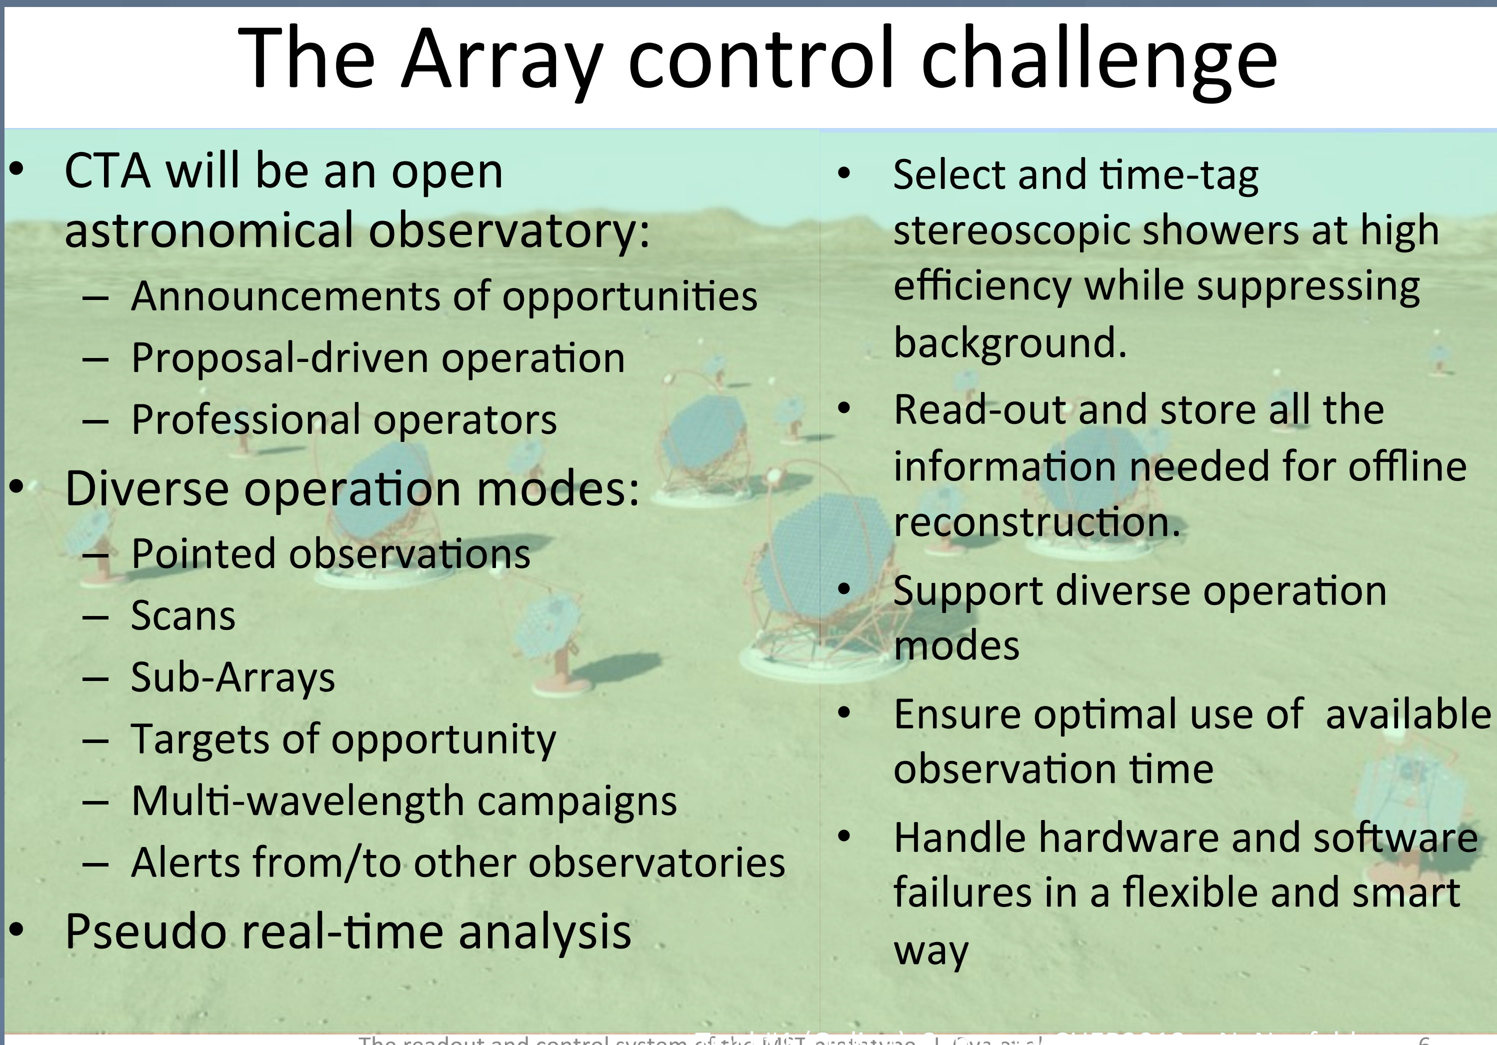

The Array control challenge

- CTA will be an open astronomical observatory:
  - Announcements of opportunities
  - Proposal-driven operation
  - Professional operators
- Diverse operation modes:
  - Pointed observations
  - Scans
  - Sub-Arrays
  - Targets of opportunity
  - Multi-wavelength campaigns
  - Alerts from/to other observatories
- Pseudo real-time analysis

- Select and time-tag stereoscopic showers at high efficiency while suppressing background.
- Read-out and store all the information needed for offline reconstruction.
- Support diverse operation modes
- Ensure optimal use of available observation time
- Handle hardware and software failures in a flexible and smart way

The readout and control system of the MST prototype I. Oya et al. 6

#21 I. Oya

LHC controls – ready for 5 more years

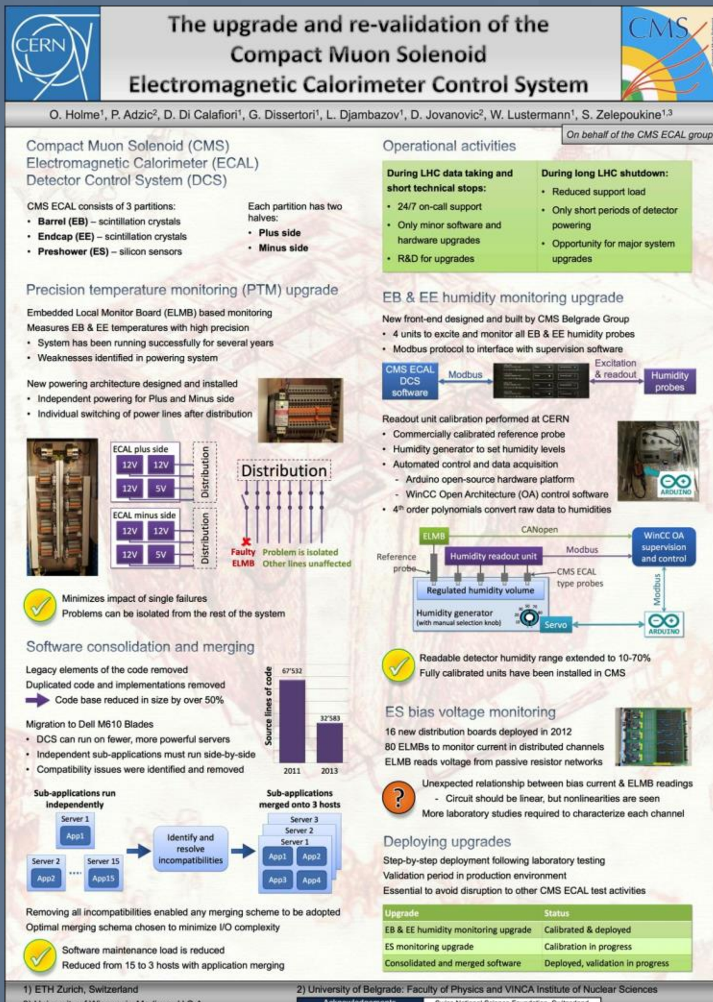

#227 O. Holme

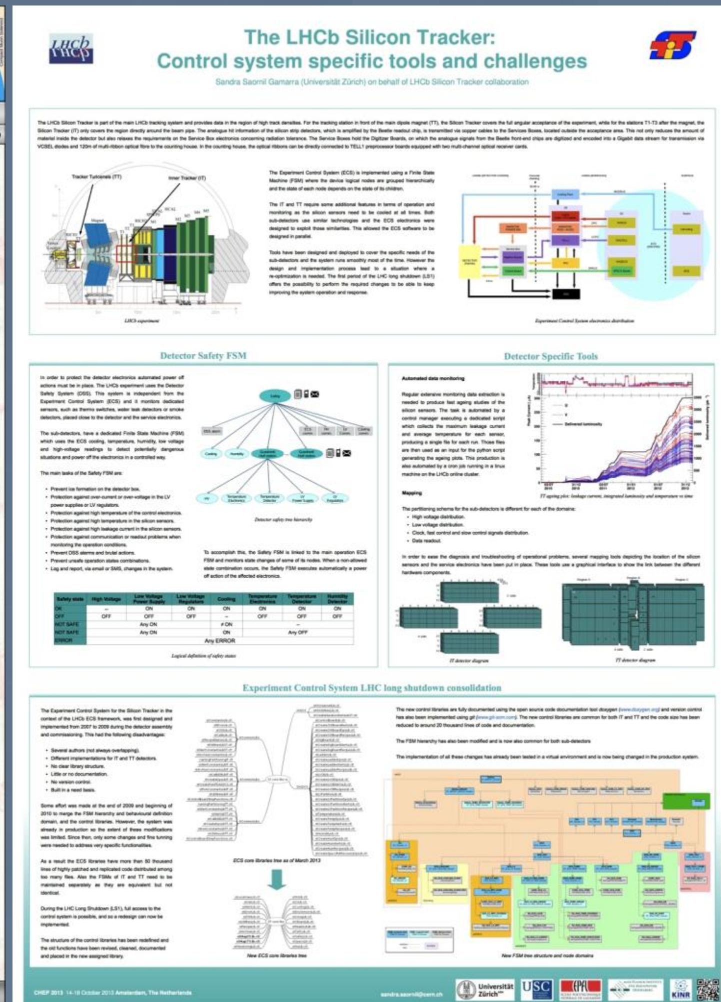

#190 S. Saornil

- 24/7/365
- Consolidation of software
- New hardware: redundancy
- all upgrades on production systems
- Sophisticates detectors with 1000s of parameters
- Flawless operation

DAQ around the world

- Accelerator based experiments
  - Belle
  - Nova
  - MICE
  - Panda
- Astronomy
  - CTA
  - H.E.S.S
  - Icecube

H.E.S.S.

#12 A. Balzer

Data Flow

- Own Receiver for each hardware Controller
- Every Receiver may have a Displayer
- Data storage classes are experiment specific
- Data is stored using ROOT object serialization mechanism
- Data calibration and reconstruction algorithms can be used online and offline

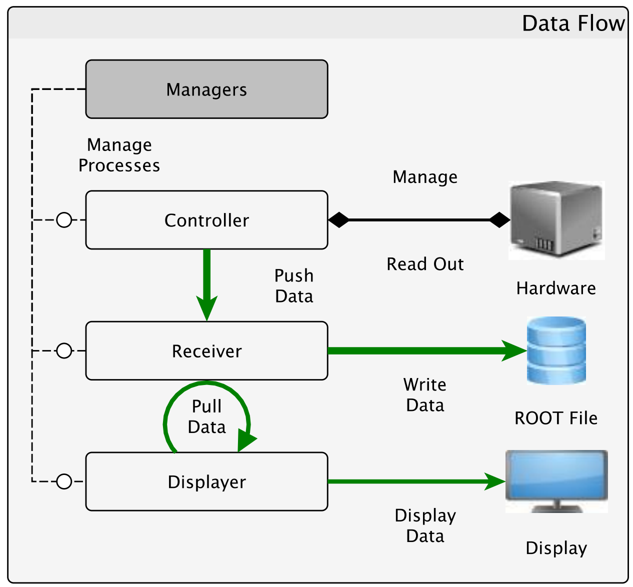

H.E.S.S.

Arnim Balzer | H.E.S.S. DAQ | 14.10.2013 | Page 11

DESY

- Have to serve multiple “experiments”→ flexibility
- > 240 processes for calibration, control, configuration, alarms, etc…
- Using CORBA, C++ and python
- Very high efficiency >99%(!)

IceCube: DAQ @ -40 °C

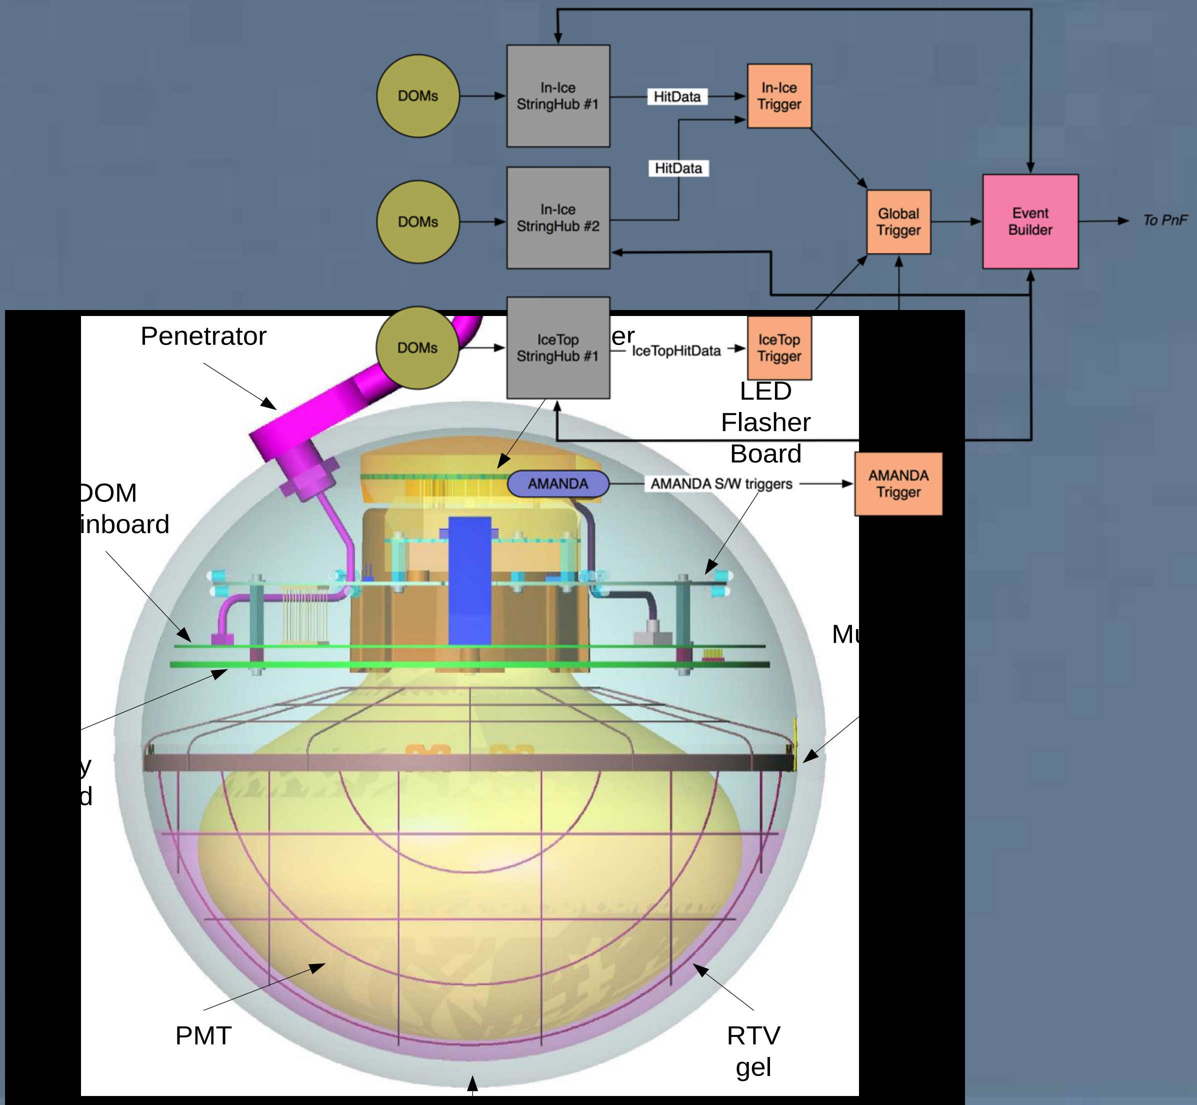

DOMs triggering on average 500 Hz in ice and 2 kHz IceTop → **3 MHz** aggregate hit rate.
LC tagged hits sent to triggers in-ice triggers select hit clusters: only these sent to triggers currently : approx **50 kHz rate**.
Global Trigger Rate 8 kHz but many overlapping triggers which get merged down to ultimate rate of triggers to Event Builder of **2.7 kHz**.
Event Builder requests data in ± 10 µs windows (typ. … up to 10 ms in extreme cases); resultant data rate to disk is **10 MB/sec**.
HitSpooling (not shown) all hits from 1st stage (DOM triggers) written to disk buffer on hubs at (aggregate) rate of **250 MB/sec**. Stored in circular buffer for several hours. Useful for external triggers such as SNDAQ and GRBs.

#390 K. Haenson

Ethernet/IB and a modern DAQ framework

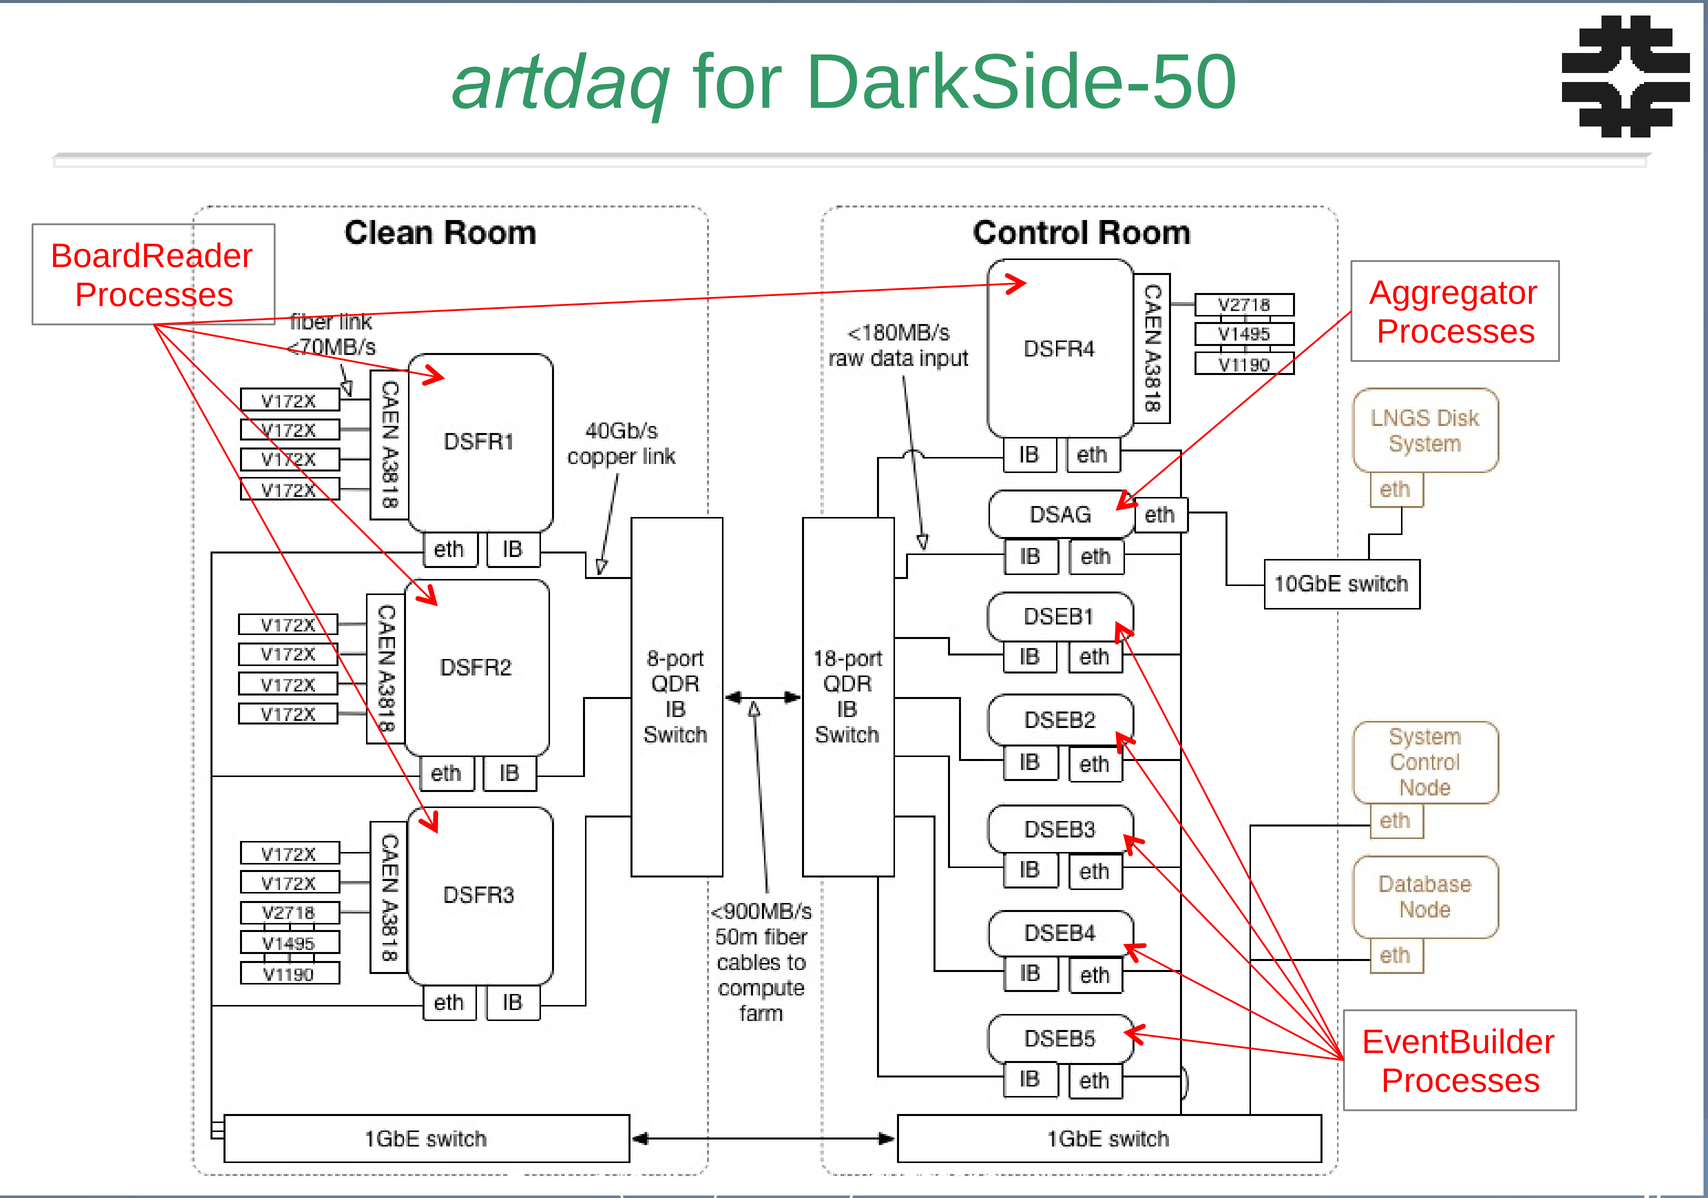

#466 K. Bierty

Time Synchronisation: NOVA

- Need to sync 350k channels with 10 ns precision
- Do not know exact arrival time of spill
- Every channel has a TDC
- Distributed time from GPS
- Continuous read-out with high precision time-stamp 56/64 bits
- Pure software trigger

GPS distributed timing

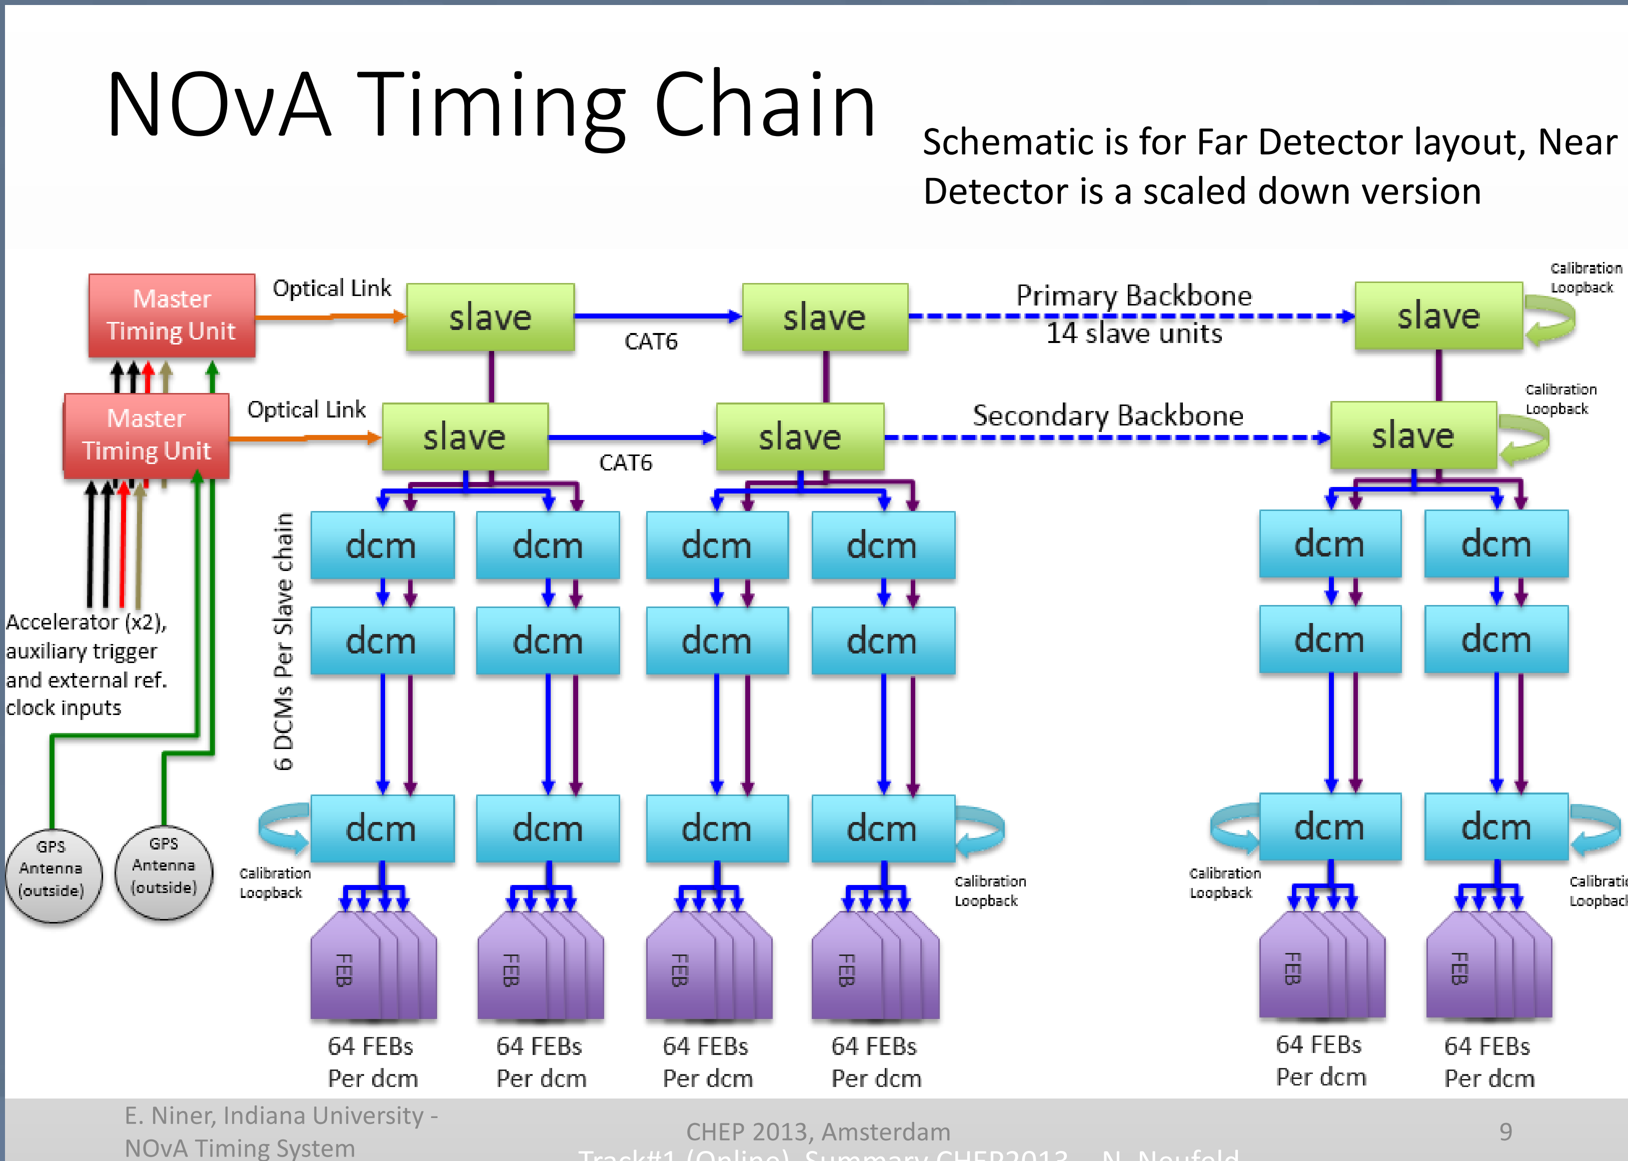

Resisting NIH

Software technologies for CTA

ACS: ALMA Common Software

- Framework for distributed applications used for the control of the ALMA array, one of the largest astronomical projects, similar in many aspects to CTA
- C++, Java and Python implementations
- Framework for a distributed system
- Environment for implementing CORBA applications
- Uses Container/Component Model

OPC UA: OPC Unified Architecture

- Open Platform Communications specification from the OPC Foundation
- A cross-platform service-oriented architecture for process control
- Multi-platform implementation, including portable ANSI C, C++, C#, Java and .NET implementations
- High scalability: from smart sensors and actuators to mainframes and servers on embedded systems

ALMA

__DevIO__: is an ACS simple and generic abstraction of hardware monitor and control point, based on the Bridge design pattern.

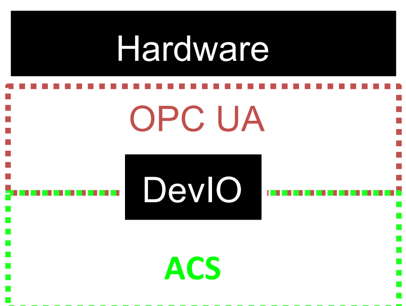

OPC FOUNDATION

The readout and control system of the MST prototype I. Oya et al. 8

- Re-use and adapt existing controls frame-work
- Use open standards (OPC UA, Java, Ethernet)

Away with the trigger!

Trigger-less PANDA

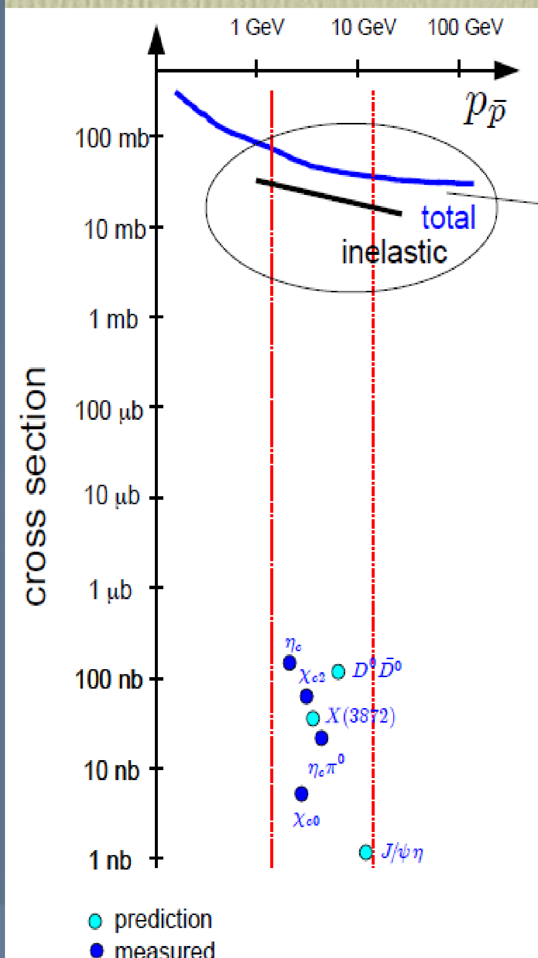

- 20 MHz event rate (peaking at 50 MHz).
- Strong event overlap.
- Lack of simple features distinguishing the interesting events from the background.

Physics Book criteria for triggering of:

- J/psi (→ base for many charmonia)
  - Invariant Mass: Tracking/Momentum
  - Electron ID: Tracking, cluster energy, track/cluster match
  - Muon ID: Tracking, Muon detector information
  - Vertex: Tracking
- D/Ds Mesons
  - Pi0s: EMC clusters
  - Inv. Mass: Tracking
  - Kaon, Pion ID: dE/dx, DIRC info (w/ track match), ToF (track match)
  - Vertex: Tracking
- Baryons
  - Inv. Mass: Tracking
  - proton, pion ID: DIRC info (w/ track match)
  - Vertex: Tracking

- No hardware trigger possible.
- Full event reconstruction, with track finding & fitting, as well as particle identification, needed to extract interesting events.

#4259 R. Karabowicz

BELLE2 – Trigger & DAQ at 10^{35} (!!)

Vital statistics at the design luminosity
of $L = 8 \times 10^{34}$ cm^{-2}s^{-1}

- ✓ Beam collision frequency ~ 508 MHz
  - ⇒ bunch separation ~ 2 ns
- ✓ Separation between physics events ~ 25000 bunches ~ 50 μs
- ✓ Good physics rate ~ 20 kHz *(see next slide)*
- ✓ L1 trigger peak rate ~ 30kHz *(not much junk!)*
- ✓ L1 latency = 5 μs
- ✓ Minimum time between triggers ~ 2 ns
- ✓ Data reduction: 80 TByte/s → 60 GByte/s
- ✓ L1 efficiency for $\Upsilon(4S)$ events ~ 100%

3

KLM trigger in Belle2 → using a common module

Barrel K_L–Muon (KLM) data flow for RPCs

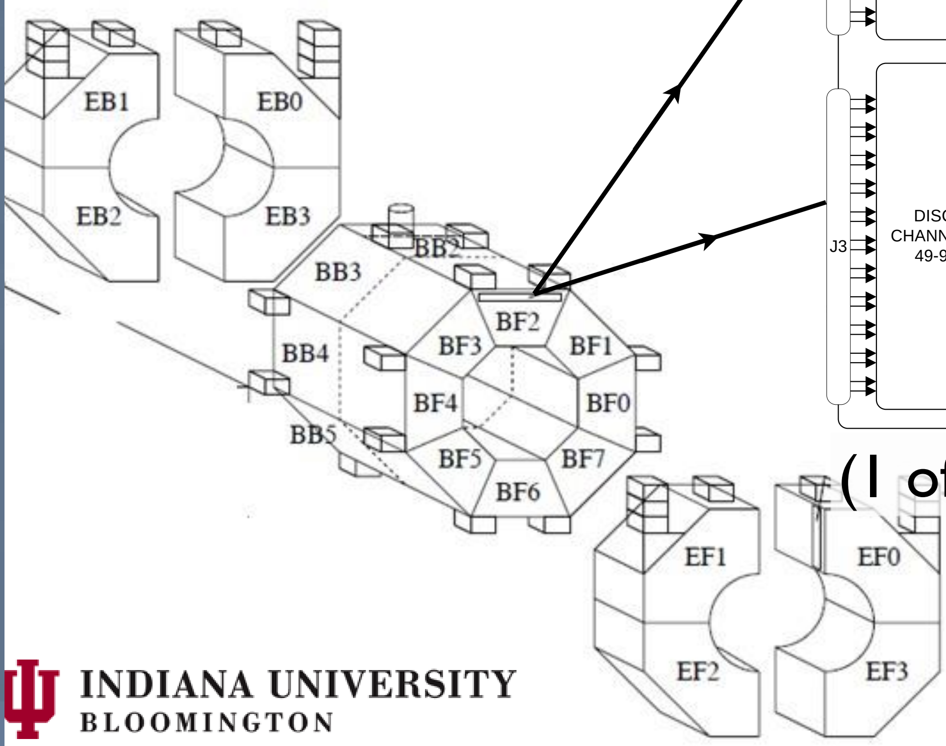

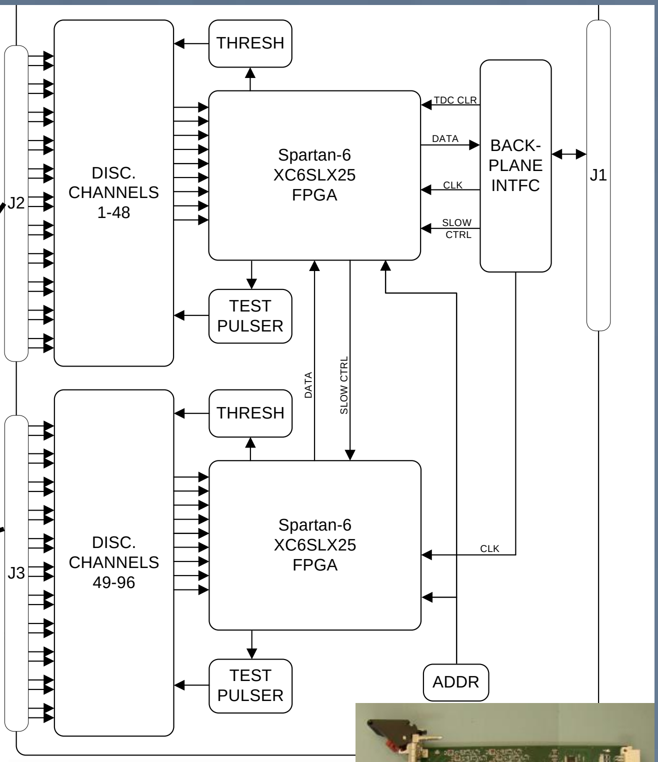

(1 of 13x8x2)

MICE → towards the far future

#447 P. Hanlet

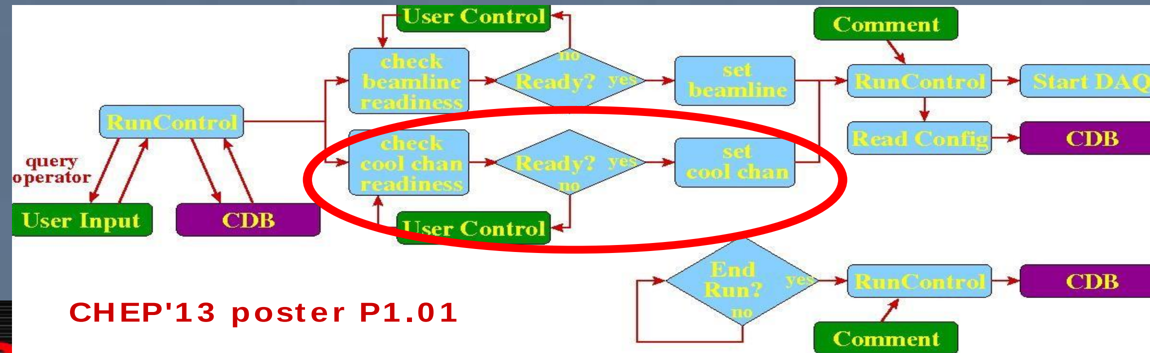

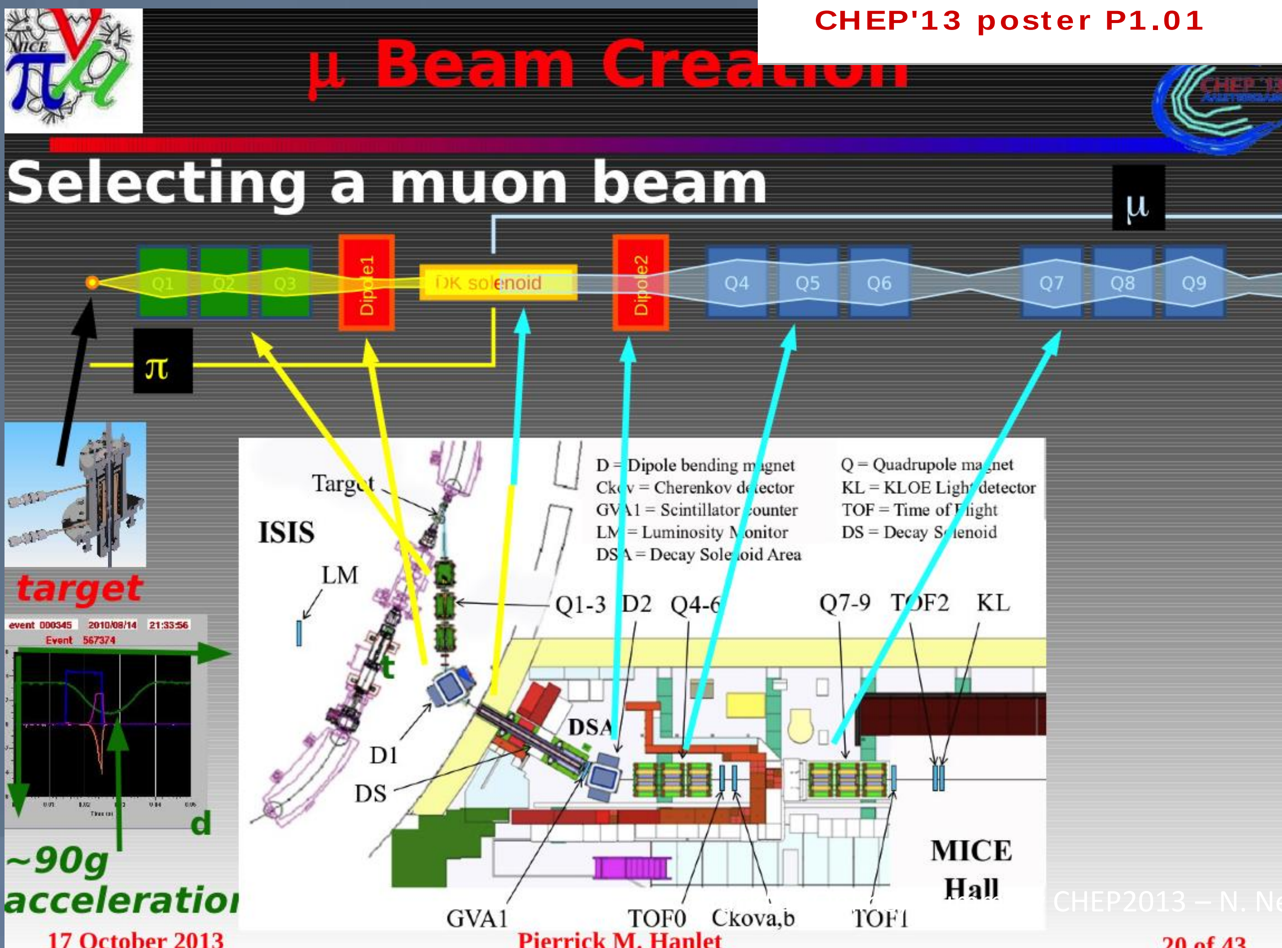

- Precision experiment 0.1%
- Tight integration of machine and experiment devices
- Highly automated control

Summary of Summary

- == input for the Summary of the Summaries
- Trend to use more and more COTS equipment and all-software based solutions continues
- DAQ systems outside HEP have been growing a lot → challenges comparable, similar ideas & synergy coming on