ALICE

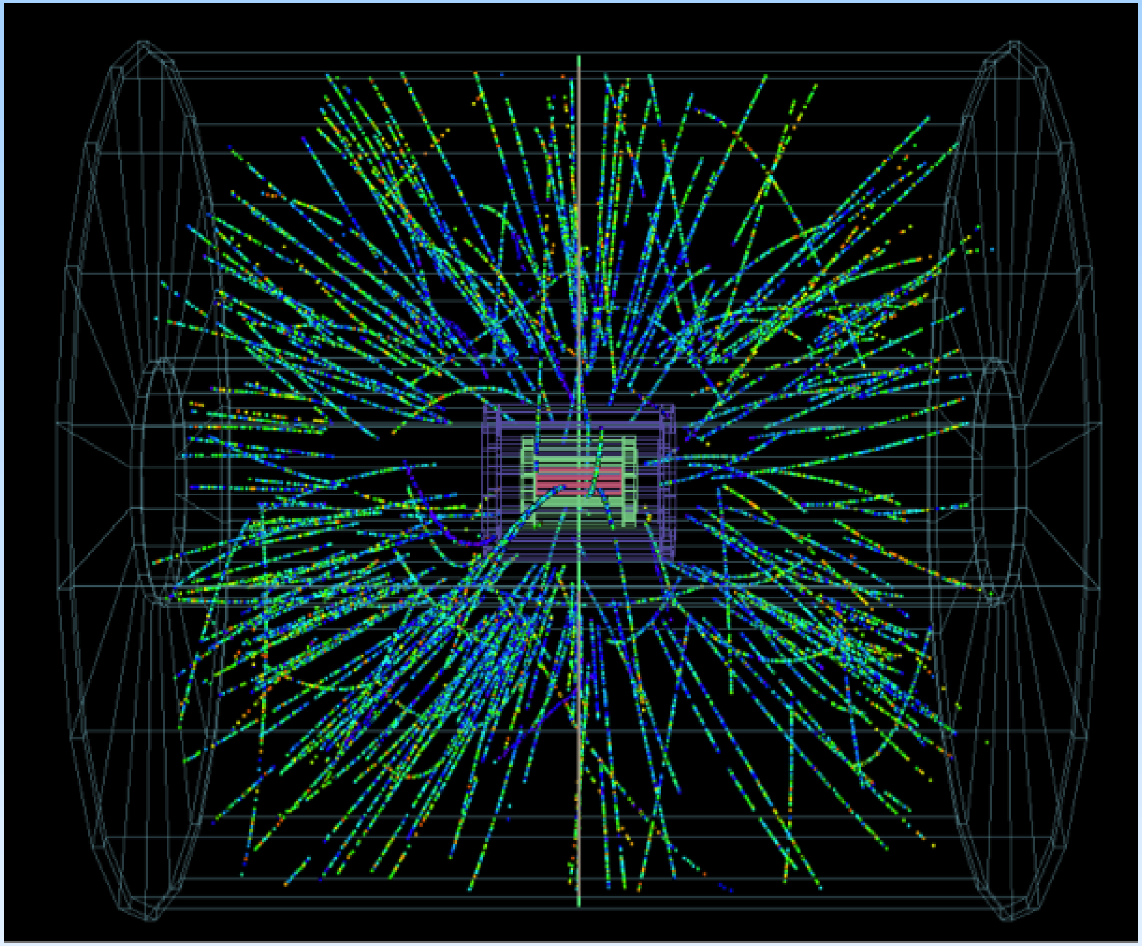

Silvia Masciocchi
s.masciocchi@gsi.de

Event reconstruction workshop, GSI, October 29, 2012

Executive summary - 1

The ALICE readout is not yet continuous now (well organized in events, processed partly online, mostly offline, etc ...)

But it is planned to be continuous after the upgraded version of ALICE, from 2018 on.

Part of that needs to be prepared already for the data taking 2015-2017, after the first long shutdown (March 2013-2015).

ALICE joins the club: we start now to think about a continuous readout and online data processing

(although I am not the right person to present this 😉)

HUGE THANKS TO THORSTEN KOLLEGGER !!!!!

Executive summary - 2

We believe ALICE (*) can contribute very importantly to the development of online data processing

BECAUSE we are a running experiment:

- Hot and hard experience of real data taking conditions
- Calibration needs
- Changing background conditions (machine)
- Real detector (some dead parts, status of detector changing in time, misalignment, …)
- Interdependence of detectors

Most of this stuff is extremely hard to simulate!!!

(*) same for HADES

ALICE today

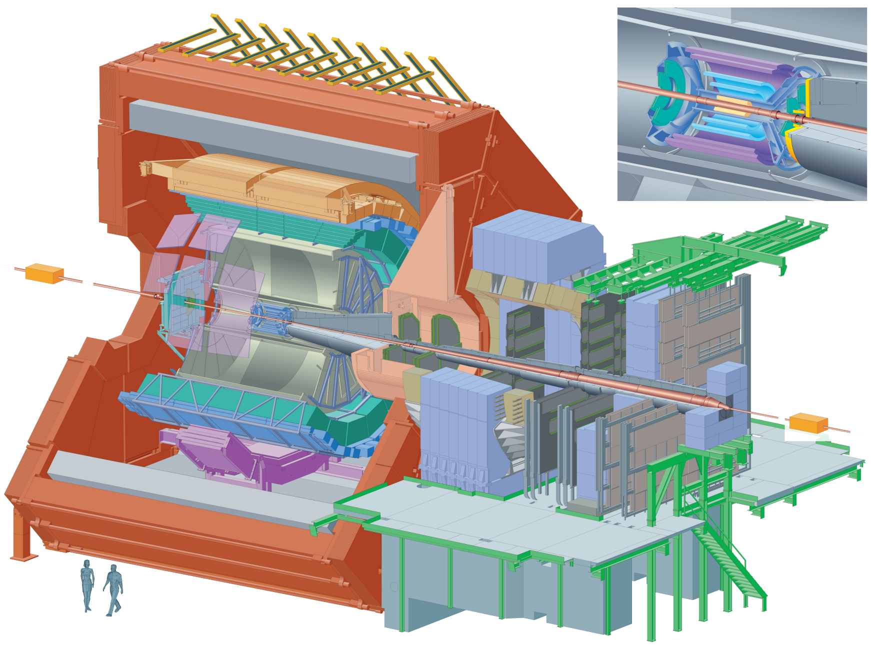

High Level Trigger
- Partial online event reconstruction
- "Latest" calibration available
- TPC data compression
  (factor ~5-6)

Data Acquisition
- Detector algorithms

Shuttle

OCDB: Offline Condition DataBase

ALICE today: HLT

Central barrel
Data structure: 18 sectors x 2 sides

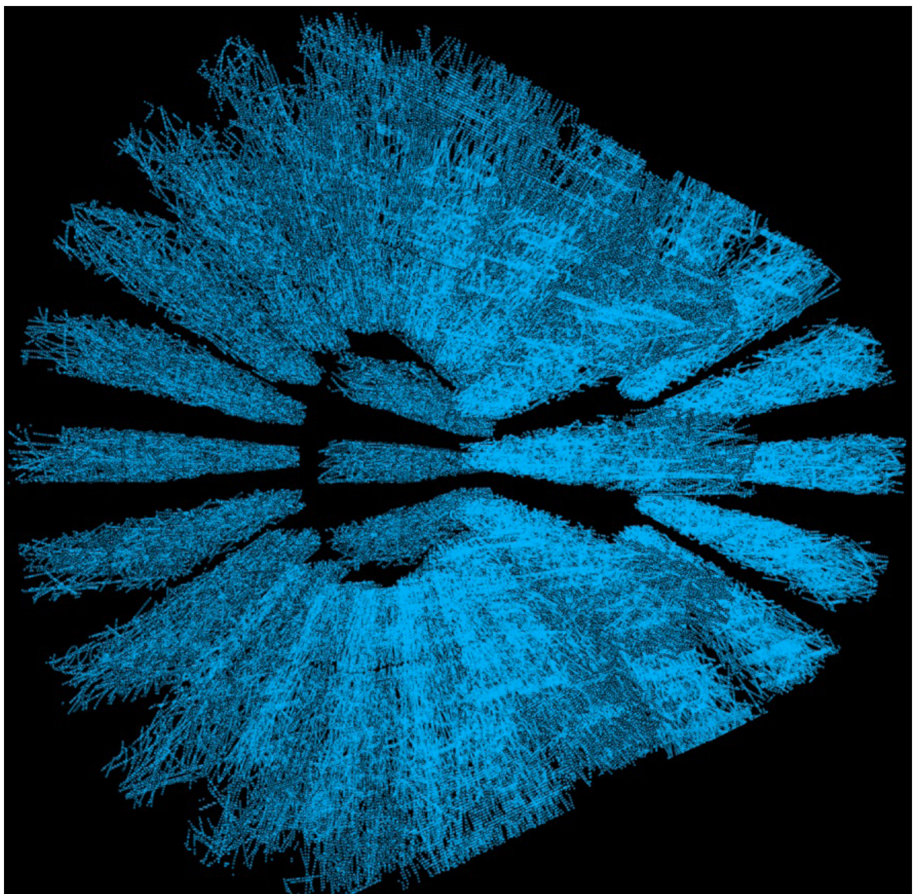

David Rohr (FIAS)

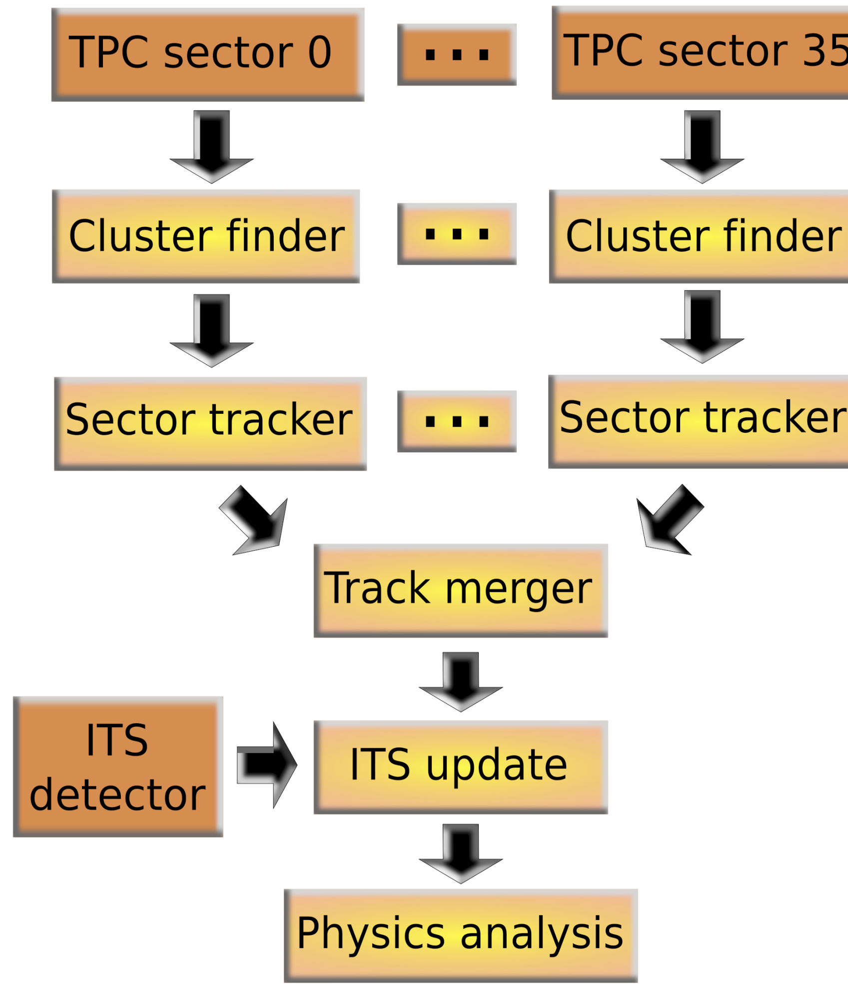

Sergey Gorbunov (FIAS)

HLT TPC Sector Tracker: The Cellular Automaton method

1. Neighbours finder:

■ For each TPC cluster it finds two (up&down) neighbours which compose the best line

2. Evolution

■ non-reciprocal links removed

■ one-to-one linked clusters are compose track segments

3. Other steps

■ fit, search for missed hits, and the final track selection

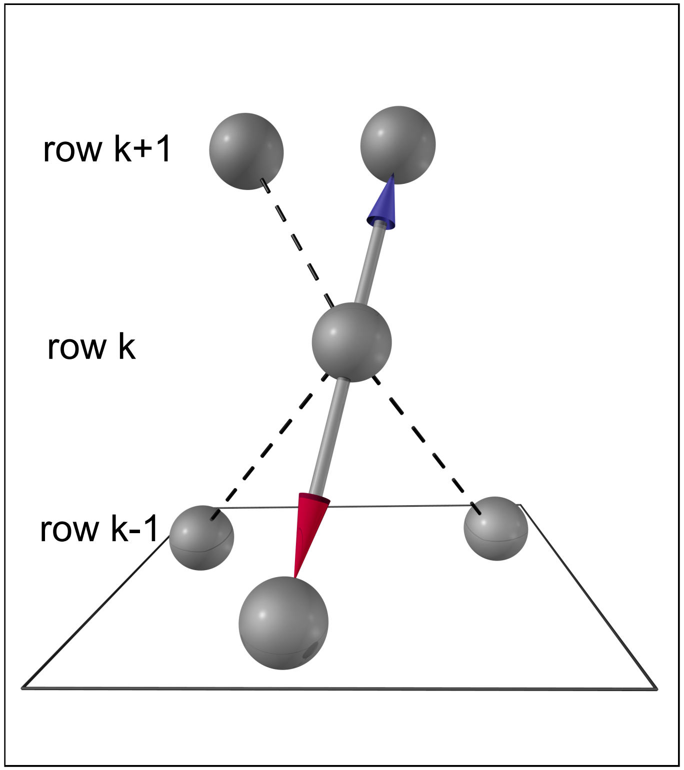

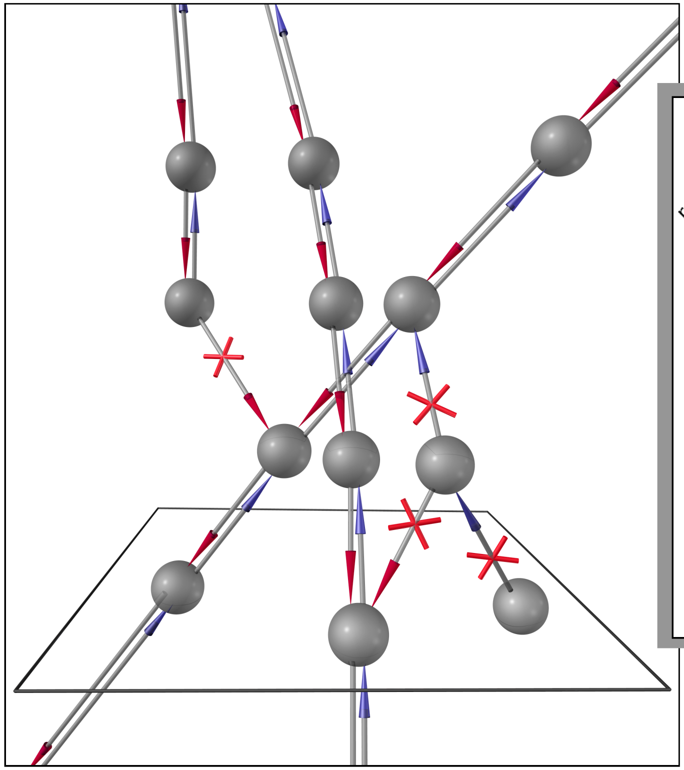

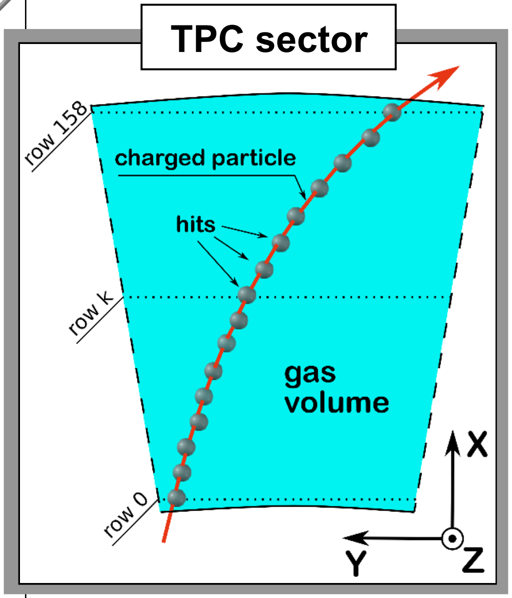

Sergey Gorbunov (FIAS)

ALICE today: HLT

Supported detectors: ITS (SPD & SSD), TPC, MUON, TRD, EMCal, PHOS

TPC

- Clusters

  Fast-Cluster-Finder implemented on FPGA coprocessor (~40 times faster than CPU implementation)

  Since 2011: online clusters are stored instead of raw data, and used by the offline reconstruction

- Tracks

  Tracker runs on CPU and GPU

  Next: HLT tracks will be used as seeds for the offline reconstruction

ALICE today: HLT

Silicon Pixel Detector (SPD) vertexer

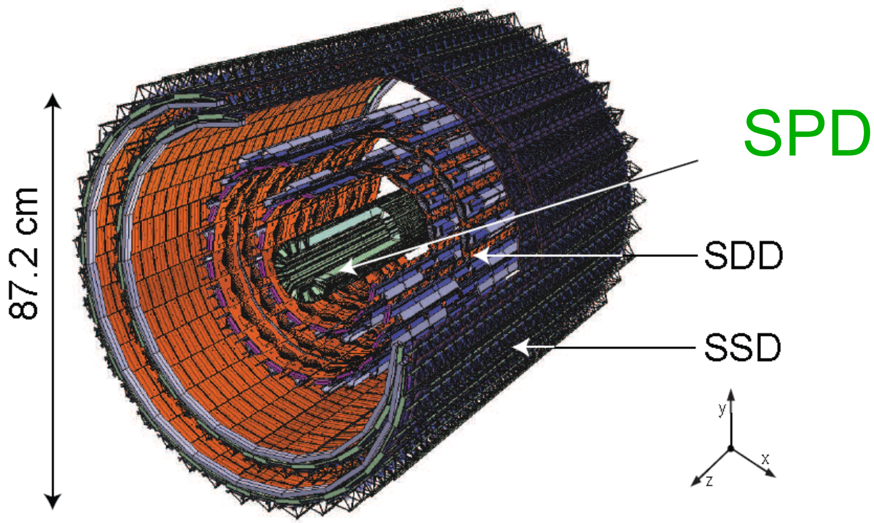

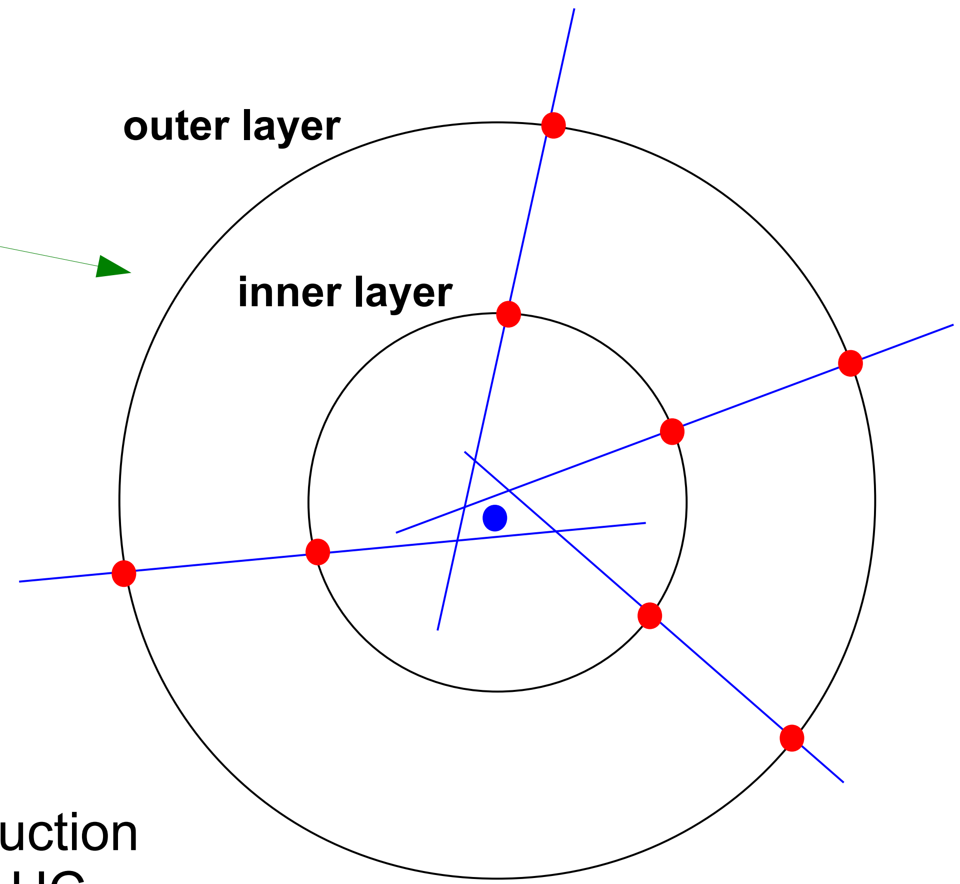

Standalone primary vertex reconstruction
Useful for online monitoring of the LHC beams interaction point in ALICE

ALICE today: offline reconstruction

- Uses HLT clusters for TPC + raw data for the rest
- Extremely involved calibration schema
  CPass0 – CPass1 – VPass - PPass
  - Fine granularity in time needed (e.g. 15 minutes for TPC!)
  - Interdependence of detectors
  - High order calibration to provide data ready for physics analysis
- All calibration information written in the OCDB
- Writes ESDs (Event Summary Data)

- ESDs are later reduced to AODs (Analysis Object Data)
- “calibration” related information needed for analysis are stored in an OADB (Offline Analysis Data Base)

ALICE upgrade

LoI and ITS CDR submitted to the LHCC on September 6, 2012

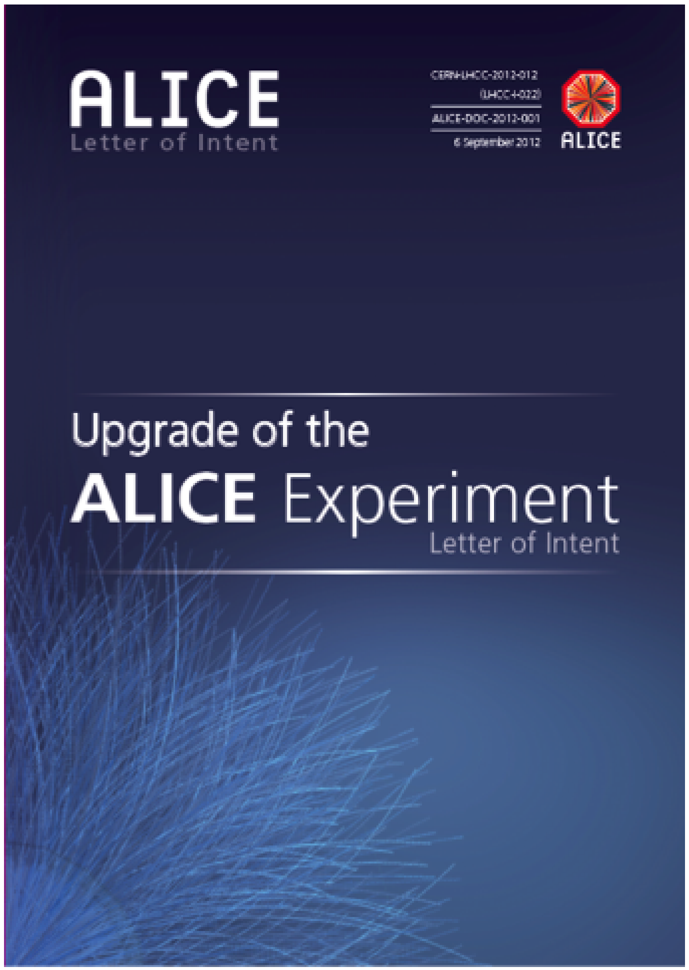

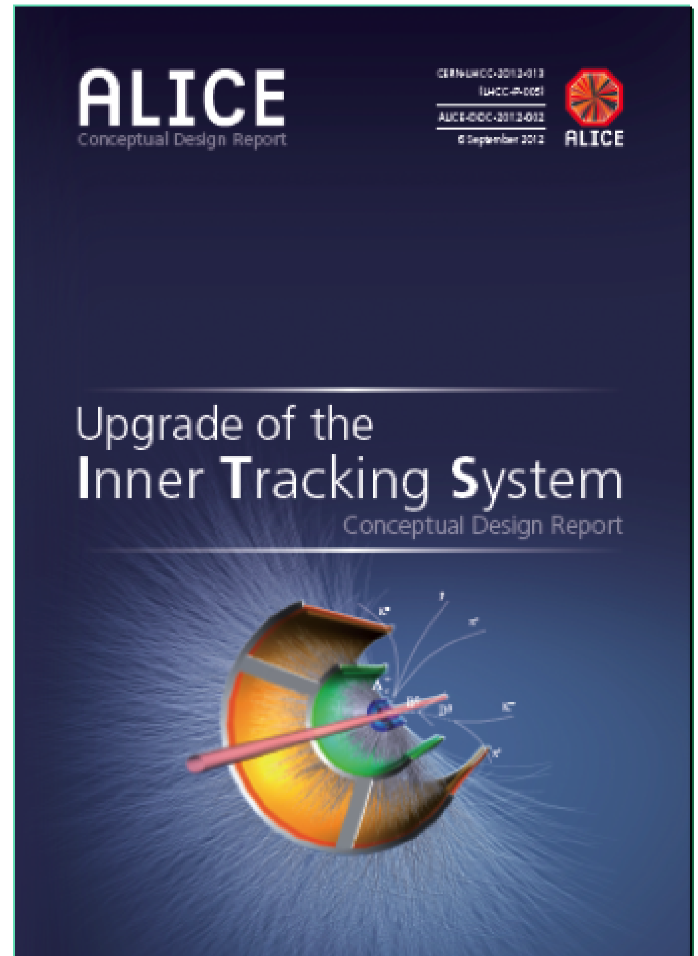

Endorsed by the LHCC on September 27, 2012

ALICE upgrade (2018) and before (2015)

- **Physics:**
  - the focus is on rare probes (heavy flavours, quarkonia, dielectrons, heavy nuclei, etc) which require very high statistics
    → aim at 10 nb^{-1} integrated luminosity
  - Most probes have low S/B: classical trigger or event filter approach is not efficient

- **Online system requirements:**
  - Sample the full 50 kHz Pb-Pb interaction rate
    (current limit at ~ 500 Hz → factor 100 increase!!)
  - Detector readout: ~ 1.1 TByte/s
    (but storage bandwidth limited to ~ 20 GByte/s)

Pb-Pb expected rates in 2015: up to 20 kHz !!

ALICE tomorrow: strategy

ALICE

Data reduction by (partial) online reconstruction and compression

Store only reconstruction results, discard raw data

- Demonstrated with TPC clustering since Pb-Pb 2011
- Optimized data structures for lossless compression
- Algorithms designed to allow for offline reconstruction passes with improved calibrations

⇨ Implies much tighter coupling between online and offline reconstruction software

ALICE tomorrow: detector readout

Combination of continuous and triggered readout

Continuous readout for TPC and ITS

- At 50 kHz, ~5 events in TPC during drift time of 92 μs
- Continuous readout minimizes needed bandwidth
- Implies event building only after partial reconstruction

Fast Trigger Processor (FTP) complementing CTP

- Provides clock/L0/L1 to triggered detectors and TPC/ITS for data tagging and test purposes

ALICE tomorrow: processing power

Online systems

Estimate for online systems based on current HLT processing power
- ~2500 cores distributed over 200 nodes
- 108 FPGAs on H-RORCs for cluster finding
    - 1 FPGA equivalent to ~80 CPU cores
- 64 GPGPUs for tracking (NVIDIA GTX480 + GTX580)

Scaling to 50 kHz rate to estimate requirements
- ~ 250.000 cores
- additional processing power by FPGAs + GPGPUs

⇨ **1250-1500 nodes in 2018 with multicores**

ALICE tomorrow: processing power

Offline system

Estimate for offline processing power
- Today: 2 month with 10^4 cores for 1 month Pb-Pb run
- 1 month Pb-Pb run after upgrade: ~2 x 10^{10} events, two orders of magnitude more than today

⇨ 10^6 cores required after upgrade

Expected performance increase per node until 2018: factor 16
- Additional gain by code optimization, use of online reconstruction results and farm

Offline raw storage requirement increases by factor ~10

ALICE tomorrow: parallel reconstruction

Change in computing paradigm: single core clock speed stagnating, instead multi/many-core
- several R&D projects underway,
  some already in production use

HLT CPU/GPU tracking

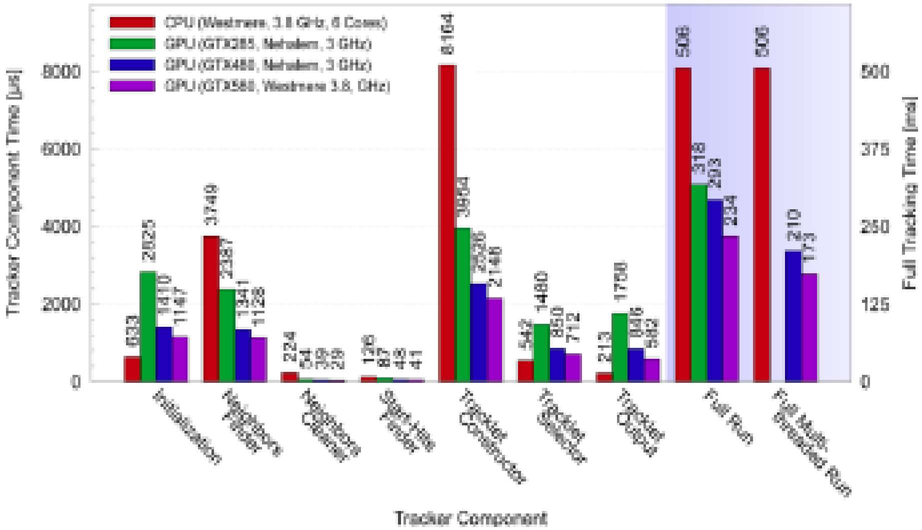

Parallelized transport in simulation

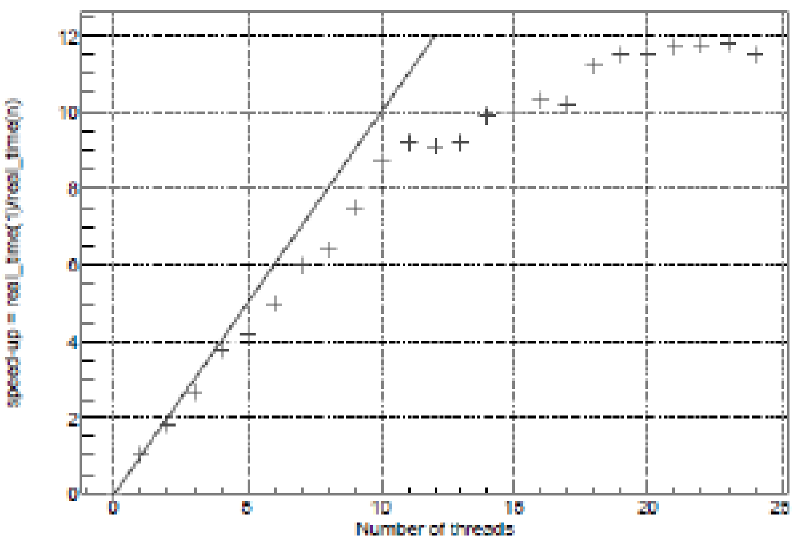

ALICE tomorrow: summary

- Current DAQ – HLT – Offline will become ONE online, central system
- Calibrations have to be brought online
- Raw data are compressed/partly reconstructed and never leave the online farm (except for tape storage)
- All raw data reconstruction is done online and produces output of the quality of PPass1 now
- Further PPass's can be done when there is no data-taking (which assumes that we store the data in a format which allows for recalibration)


Plenty of fun ahead of us
Let's do it together !
Courtesy of
David Rohr