



O^2 Technical Design Report: reconstruction, calibration, compression, simulation.

LHCC review, 2015-06-01

Mikołaj Krzewicki (FIAS)

Introduction

- Change of paradigm for data taking and reconstruction:
 - Triggerless, continuous data readout.
 - Basic computing unit changes:
 - Event (single interaction) $\rightarrow$ Time frame (time window of readout).
- Vast amounts of data ($\times 100$ wrt. Run1):
 - online (synchronous with data taking) compression needed.
- Reconstruction & calibration performed on-the-fly.
 - Necessary to achieve the desired compression ratio.
 - Lossy compression: part of information will be discarded.

Table 8.1: Summary of the different detector calibrations for TPC and ITS.


Detector	Calibration	Processing type	Beam [y/n]	Freq.	N. of events
TPC	Pedestal/Noise	FLP	n	fill	O(100)
	Pad Gain	Standalone runs	n	year	-
	Drift Velocity	EPN (sync)	y	15 min ¹	O(1k)
	Space Charge	EPN (sync)	y	15 min ¹	O(5k)
	Drift Velocity	EPN, (async ²)	y	15 min	O(1k)
	Gain	EPN (async ³)	y	15 min	O(1k)
	Space Charge	EPN (async ⁴)	y	5 ms	all
ITS	Noisy/dead/faulty channels	Standalone runs	n	run	O(1k)

Table 8.2: Summary of the different detector calibrations for TRD.

Detector	Calibration	Processing type	Beam [y/n]	Freq.	N. of events
TRD	Pedestal/Noise	FLP	n	fill	100
	t_0	FLP (monitoring on EPNs)	y	run	O(10k)
	Chamber Status	EPN	y	run	O(35k)
	Drift Velocity	EPN	y	run	O(3.5k)
	ExB	EPN	y	run	O(3.5k)
	Chamber Gain	EPN	y	run	O(35k)
	Pad by pad gain	standalone runs	n	year	O(2G)
	PID reference	Asynchronous	n	period	all

- Reconstruction & calibration performed in 2 stages:
 - synchronous to data taking (up to level needed for compression).
 - asynchronous to data taking (final “physics grade” calibration and data)
- Per-fill calibrations w/o beam mean a short standalone run (~ few minutes) before the fill start
- Calibrations requiring accumulation of large statistics are used at physics analysis stage
- Long-time special standalone runs (Kr)


Data flow




Data flow, First Level Processors

- Preparation of sub-timeframes.
- Calibration: per-channel, pedestal, etc...
- Local reconstruction:
 - FPGA based cluster finding: TPC
 - CPU based cluster finding: ITS, MFT, HMP
 - for ITS and MFT possible FPGA implementation.
 - pre-clustering for MCH and MID

Figure 8.1: Schematic outline of the reconstruction and calibration data flow.




Data flow: Event Processing Nodes




- Sub-timeframe assembly into timeframes.
- Calibration:
 - TPC space charge distortions:
average map rescaled by multiplicity (TPC currents)
integrated over last 160ms (full ion drift time).
- Reconstruction:
 - TPC: Track finding, junk removal.
 - ITS: Track finding (full or partial, depending on noise levels).
 - MCH: Cluster finding.
 - FIT: Interaction type, preliminary vertex, multiplicity
- Compression.

Data flow: Event Processing Nodes

- Calibration:
 - TPC: average space charge map (SCD, integrated) periodically updated using ITS/TRD track interpolations ($O(15\text{min})$).
 - Drift velocity from TPC-ITS matching ($O(15\text{min})$).
 - Histograms for high frequency SCD correction (TRD-ITS interpolation)
- Reconstruction:
 - TRD: Matching of TPC tracks to online trackless.
 - ITS/TPC: standalone tracks matching.




Data flow: asynchronous processing



- After data taking, using the CTFs stored in the O₂ cache.
- Calibration:
 - TPC: final calibration.
 - Drift velocity ($O(15\text{min})$).
 - Gain ($O(15\text{min})$).
 - Space charge (high freq. $O(5\text{ms})$).
 - TRD: final calibration (per run).
 - Chamber status.
 - Drift velocity.
 - ExB.
 - Chamber gain.
- Reconstruction:
 - TPC,ITS: final tracking.
 - ITS/TPC/TRD: final matching, refit.
 - MFT/MCH: track finding, matching.

Data flow: asynchronous processing

- Calibration:
 - Calorimetry, CPV, TOF, HMPID.
- Reconstruction:
 - TOF, HMPID, EMCAL, PHOS/CPV: matching to global tracks.
 - global refit inward.
- Event extraction.
 - Primary vertex finding.
 - secondaries: Vos, cascades.
 - Track-to-vertex association.



Event extraction

- Time resolution of global track (ITS+TPC) is determined by $\sigma_Z \sim 0.25\text{cm}$ of TPC track at outer ITS radius:
 $\sigma_t \approx \sigma_Z^{TPC} / V_{drift}^{TPC} \approx 0.1\mu\text{s} \rightarrow$ vertex built with global tracks is >99% pile-up free
- Time resolution of ITS standalone track (no TPC match) determined by ITS readout integration frame (5-20 μs)
 - full train seen ($\mu\text{Tr} \sim 0.38$)
 - tracks which cannot be unambiguously attributed to global vertex will see $\sim 18\%$ pile-up (>100% in 200kHz pp)
 - problem mostly for secondaries: their attribution in the analysis to interaction vertex will depend on distance to closest pile-up vertex




Figure 8.4: Probability of at least one pile-up collision happening in proximity (Y-axis) for selected vertex as a function of the pile-up rate (X-axis). Luminous region $\sigma_Z = 6\text{cm}$ is assumed.

- Store primary vertices with unambiguously attributed tracks, Vo's;
- separate container of ITS standalone tracks (secondaries) with candidate primary vertices (to be attributed in the analysis)

Data reduction

Chapter 9

Data reduction

Data reduction is the driving requirement of the synchronous reconstruction since the peak data rate from the read-out of the detectors is expected to be ~ 1.1 TB, hence an overall reduction by factor ~ 14 . Listed below are the data reduction procedures envisaged by different detectors, excluding those that have a comparatively negligible data rate.

TPC data reduction

- Uncompressed data rate: ~1TB/s.
- Cluster finding (factor 2.5).
 - Compression ~1.25 in FPGA cluster finding.
 - Data format optimisation ~factor 2.
- Cluster compression.
 - Currently: entropy coding (Huffman) - factor ~1.6.
 - O2: compression applied to more cluster parameters.
 - O2: additional compression using tracking information.
 - Estimated total: factor ~3
- Background identification: factor ~2.
- Charge transformations for clusters associated to tracks: factor ~1.35.

TPC cluster compression




Figure 9.1: Data compression factor versus raw data size achieved using cluster finder and data-format optimizations in the High-Level Trigger during the Pb–Pb data taking in RUN 1. Adopted from [3].

- Current HLT compression performance:
 - about a factor ~ 4.2 .
 - cluster finding + format optimisation + entropy coding.

Fast cluster finder

- FPGA based algorithm
- used in HLT on the H-RORC (run 1) and C-RORC (run 2).
- Data flow architecture
 - highly pipelined and parallel.
- DDL2 version: 25X faster than CPU.
- each C-RORC has 6 instances:
 - 150 core equivalent.




Figure 6.4: Performance of the FPGA-based FastClusterFinder algorithm for DDL1 and DDL2 compared to the software implementation on a recent server PC. The plot on the top of the figure shows the required processing time relative to event size. The one on the bottom of the figure shows the same measurements as rate equivalent.

HLT TPC Track Finder




Figure 6.5: Tracking time of HLT TPC CA tracker on Nehalem CPU (6 Cores) and NVIDIA Fermi GPU.

- Cellular automaton track finder algorithm (GPU).
- CPU based merger + Kalman track fit (currently ~3 supporting CPUs).
- 1 GPU (+3CPU's) replaces ~28 CPU's.

TPC background identification




Figure 9.2: Projection of a pp event on the transverse plane of the TPC, showing clusters attached to physics tracks (red) and not assigned clusters (black).

- Identified junk removal:
 - clusters from low-momentum tracks (delta electrons).
 - upstream ion-gas/beampipe interactions -> tracks parallel to the beam pipe.
- Prototype shows rejection factors from 2.5 (PbPb) to 4 (high lumi pp).

TPC compression: summary

Table 9.1: Data reduction steps for the TPC.

Data Reduction Step	Data Reduction Factor	Data Rate at 50kHz Pb–Pb interactions (GB/s)
Detector Input	—	1000
Cluster Finding	2.5	400
Cluster Compression (Cluster+Track Info.)	3	135
Background Identification	2	68
Charge Transformation & Compression	1.35	50

- factor 20 compression feasible.

ITS data reduction

- Uncompressed: @50kHz Pb-Pb $\rightarrow$ ~8.2GB/s related to physics.
- Noise contribution:
 - up to ~28GB/s (10^{-5} random noise probability per pixel).
- + 10% protocol overhead: 40 GB/s.
- Huffman encoded clusters.
 - Multiplicity dependent tables ~23bits or ~19 bits (if noise dominates) per cluster.
 - output: from ~4.2GB/s to ~26GB/s (with full noise).
- If noise dominates: additional factor 3 compression will be necessary.
- If synchronous reco feasible - discard clusters not attached to tracks.
 - Resulting in peak data to storage rate: ~3GB/s.
- Otherwise: use partial reco and partial cluster removal from outer layers.
 - Performance under study.

TRD compression

- Uncompressed data rate: 20GB/s
- only 10% of online TRD trackless can be matched to TPC tracks.
- work in progress to check this for run 3.
- compression strategy:
 - store only trackless which are in the vicinity of the TPC extrapolation.
 - this should lead to rate reduction to ~3GB/s

Data reduction summary

Table 9.2: Data rates for input to O² system and output to permanent storage for routine data taking with Pb–Pb at 50kHz interaction rate.

Detector	Initial raw data rate [GB/s]	Compressed data rate [GB/s]	Data reduction factor
TPC	1000.0	50.0	20
ITS	40.0	26.0(8.0)	1.5(5)
TRD	20.0	3.0(?)	6.6
TOF	2.5	2.0	1.25
MCH	2.2	0.7	2.9
MFT	10.0	5.0	2
EMCAL	4.0	1.0	4
PHOS	2.0	0.5	4
Total	1080.7	88.2(70.2)	12.3(15.4)

- Target compression factor of 14 feasible.

Simulation

- Driving requirement: simulation needs grow by factor ~ 20 .
- Most of this increase will be handled by fast/parametrised Monte Carlo.
- Need for full simulation still exists,
 - largest resource consumer: particle transport.
 - needs to take advantage of multi-core architectures.
 - needs to scale in a heterogeneous environment (GRID/Cloud, opportunistic supercomputing resources).
 - $\rightarrow$ Geant4
- Better simulation strategy:
 - Merging techniques: at kinematics/MC level (signal into simulated background).
 - Embedding of simulated signals in real data.

Simulation

- Currently used: Geant3, Geant4 and FLUKA
 - as plugins to the Virtual Monte Carlo (VMC) layer.
- Multithreaded Geant4 (via VMC) already tested and shown to scale.
- Fast/parametrised simulation
 - within the VMC framework.
 - plugins would apply parametrised resolution/efficiency effects (instead of running the full particle transport).
 - shares similarities with ATLAS Integrated Simulation Facility Approach.

backup

TOF compression

- Uncompressed data rate: 2.5GB/s
- Optimisation of the data format: ~20% reduction.
 - Performed on the FLP.
 - Remove slow control information (temperatures, thresholds).
 - Redundant, available in the DCS.
- Suppression of headers/trailers from TRM (readout) for empty modules.
 - Continuous mode in pp data compression factor lower:
 - 10 μ s windows implies a lower fraction of empty modules.

MUON chamber and identifier

- Uncompressed data rate: $\sim 2.2 \text{ GB/s}$
- If pad information can be safely dropped:
 - Can be compressed to $\sim 1.5 \text{ GB/s}$ on the FLPs (pre-clusterisation).
 - clusterisation (EPN): total 0.7 GB/s to storage.
- If pad information cannot be dropped: $\sim 2.1 \text{ GB/s}$ to storage.
 - no gain in performing clusterisation in synchronous mode.
- under study.

MFT

- Uncompressed data rate $\sim 10\text{GB/s}$.
- Same technology as the ITS.
 - Worst case scenario: noise level of $\sim 10^{-5}$.
- Similar compression approach as with the ITS.
 - Huffman coding $\rightarrow \sim 5\text{GB/s}$

EMCAL/PHOS


- EMCAL uncompressed data rate: ~2GB/s.
- PHOS uncompressed data rate: ~4GB/s.
- Data reduction on the FLP.
 - compress 30 10-bit samples on two floats:
 1. Cell energy.
 2. Time information.
- Data reduction factor 4.

Event extraction

- 50kHz PbPb: assume 12 equally-spaced trains of 48 bunches each (50ns filling), like for TPC TDR $\rightarrow \mu\text{BC} \sim 0.0079$
- Time resolution of global track (ITS+TPC) is determined by $\sigma_Z \sim 0.25\text{cm}$ of TPC track at outer ITS radius:

$$\sigma_t \approx \sigma_Z^{TPC} / V_{drift}^{TPC} \approx 0.1 \mu\text{s} \rightarrow \text{vertex built with global tracks is } >99\% \text{ pile-up free}$$


- Time resolution of ITS standalone track (no TPC match) determined by ITS RO integration frame (5-20 μs ?)
 - $\rightarrow$ full train seen ($\mu\text{Tr} \sim 0.38$)
 - $\rightarrow$ tracks which cannot be unambiguously attributed to global vertex will see $\sim 18\%$ pile-up ($>100\%$ in 200kHz pp)
 - $\rightarrow$ problem mostly for secondaries: their attribution in the analysis to interaction vertex will depend on distance to closest pile-up vertex



- Depending on filling scheme and ITS integration time, it might be preferable to avoid strict collision definition in AOD, instead storing primary vertices with unambiguously list attributed tracks, Vo's + separate container of ITS standalone tracks (secondaries) with candidate primary vertices (to be attributed in the analysis)

TPC data compression

gx1:gx0 {eventNr==674&&cl.fUsed==0}



Clusterization & Huffman encoding

Routinely used in data taking
suppression factor of > 4 (5) achieved

Identifying loopers

pp, p-Pb real data

70% (80% in pp w/ beam-gas background)
of clusters not assigned to track $\rightarrow$
identify low momentum (< 50 MeV/c)
tracks looping in magnetic field and
remove their clusters

assume $> 80(90)\%$ efficiency

$1 - 0.7 \times 0.8(0.9) = 0.44(0.33)$

$\rightarrow$ factor $> 2(3)$

Hierarchical

cluster – tracklet – track model

store params w.r.t. resolution, saves bits
+ makes Huffman encoding more efficient

25 instead of 77 bits per cluster

$\rightarrow$ factor > 3 (to be demonstrated)

Huffman x loopers x hierarchical track model

= 4 (5) x 2 (3) x 3 (?)

= 24 (45?)