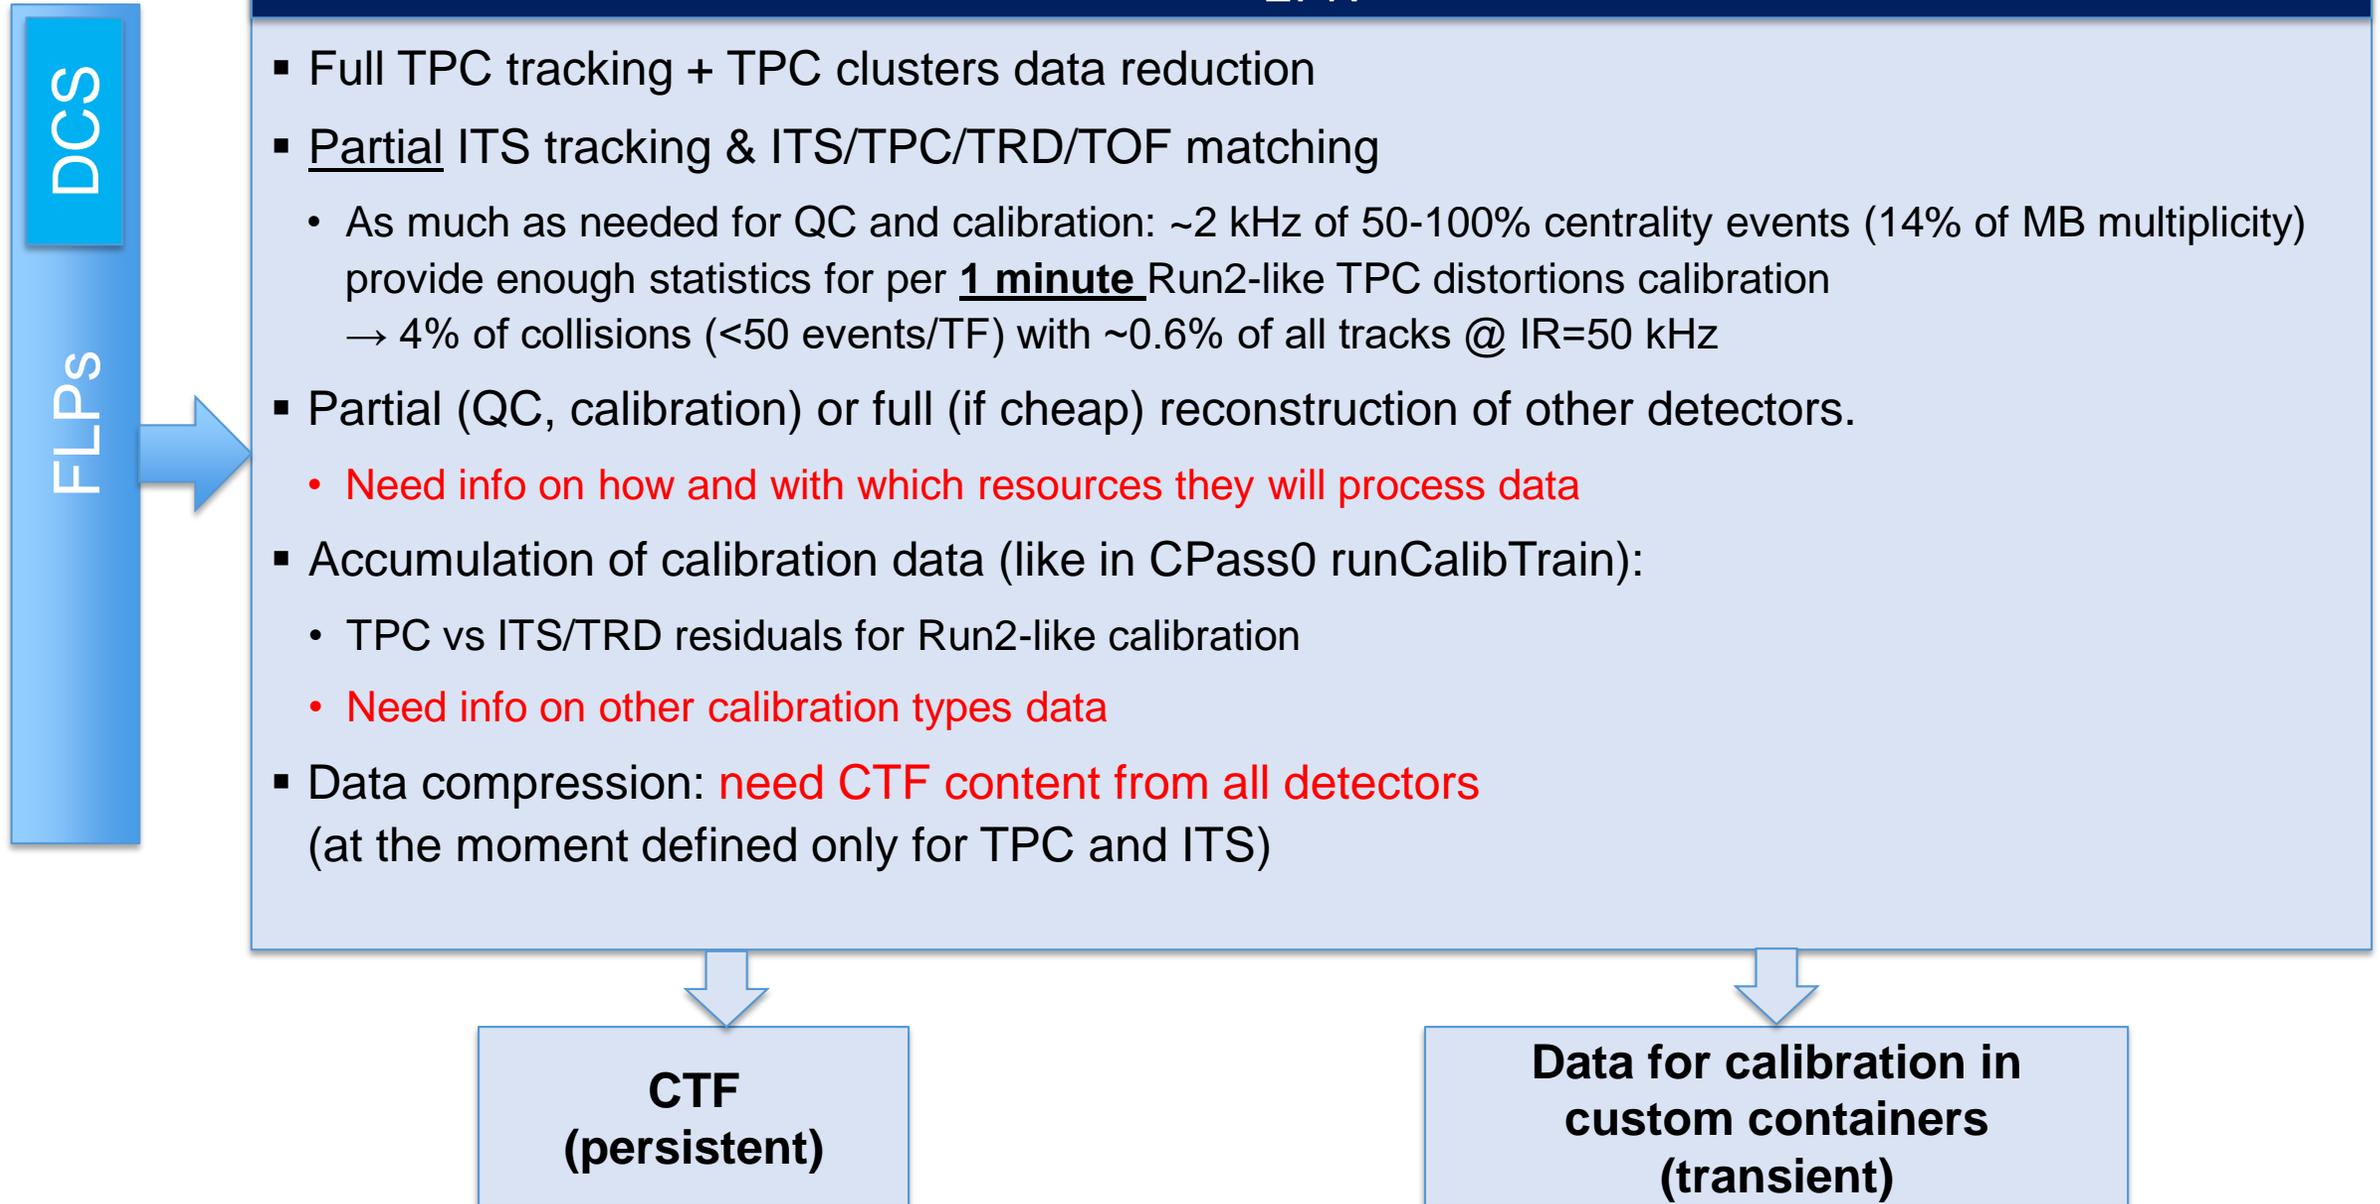


Reconstruction and Calibration in Run3

Synchronous stage



Synchronous stage

- **TPC tracking:** current status already good enough for this stage (dE/dX missing but not needed), GPU version fits well to ~30s time slot allocated for ~23 ms TF reconstruction.
- **ITS tracking:** calibration needs only tracks $p_T > 0.5$ GeV/c, no particular requirements on the efficiency → single pass for straight tracks in 50 events will take $\ll 1$ s on single CPU core, tracking on GPU is progressing.
 - Depends on the tracklet-vertexing component which needs speed up.
- **ITS/TPC matching:** operational on CPU
- **ITS-TPC / TRD matching:** OK on CPU, in progress on GPU (tested in HLT framework), see TRD report
- **ITS-TPC / TOF matching:** CPU version in validation (with some limitations due to the time/length integration missing on the framework level), see TOF report.
- All these components are supposed to run both on the CPU and GPU, see report by David.
- O2 farm can be used for truly synchronous processing with ~50% efficiency during the beam (tailored for 50kHz IR while $\langle IR \rangle$ is ~25kHz) and idle between the fills → 25% average efficiency: Only TPC compression can depends on tracking → if needed, data of other detectors can be buffered and processed at inter-fill or low-IR periods

Data compression

▪ TPC data compression

▪ **clusters reduction**: two alternative scenarios:

- Plan A (TDR): keep whatever is not tagged as junk (loopers, noise) after protecting clusters attached to useful tracks and those in the proximity of the track:

Status:

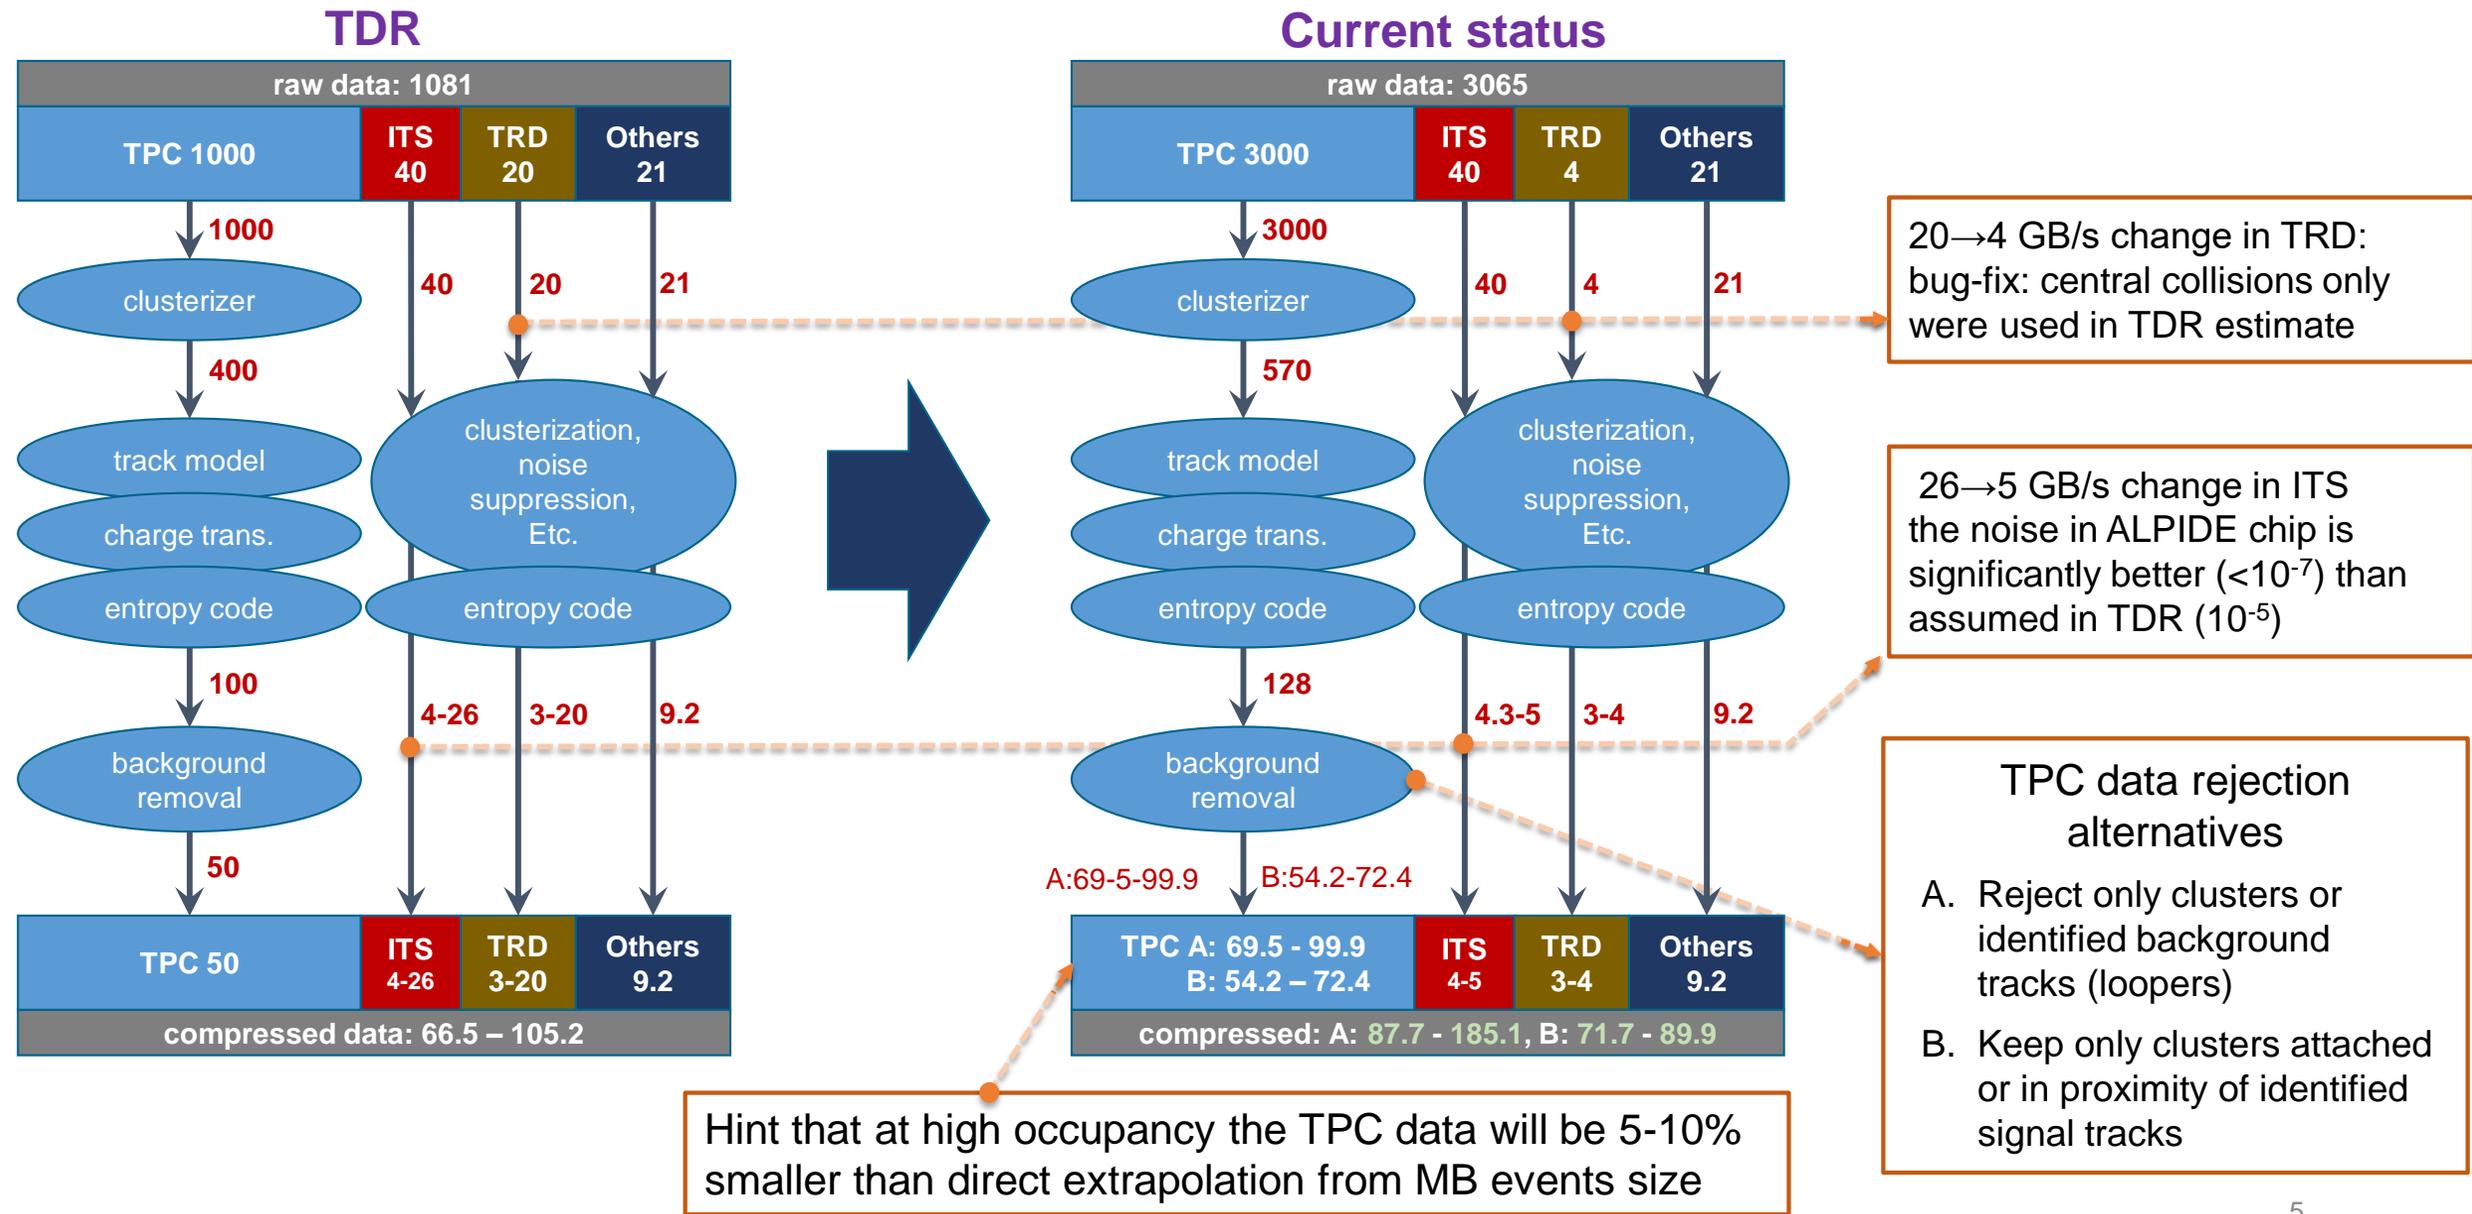
- removable clusters from tracking, merging ($p_T > 10$ MeV): **current 12.5%**, upper limit: 22.5%
- removable clusters below 10 MeV (no good progress yet): upper limit 13.6%
- with improved definition of proximity to track: 8% less fake protected
→ **current: 12.5%, upper limit : 39.1% = ((22.5 + 13.6) / 0.92)**
- Plan B: keep only protected clusters (currently 62.7%):
→ **current: 37.3%, upper limit : 52.5% = (100 - 51.7) / 0.92**

▪ **Track-based entropy compression** being implemented

▪ ITS data compression

- **cluster → pattern ID reduction** : in validation
- **entropy compression**: to be implemented

CTF Data Rates (in GB/s, 10^9 bytes) for Pb-Pb @ 50 kHz



Asynchronous stage

EPN (+GRID?)

Calibration

- Extraction of calibration from the data accumulated at synchronous stage, populating the CCDB (equivalent of the MakeOCDB + TPC SP calibration of Run2)
- If needed: consider reconstruction pass over tiny fraction of data for fine calibration (a la CPass1) and validation: ~0.1% of data will provide MB events statistics / time like in Run2 PbPb CPass1 (~25 Hz)

PPass reconstruction

CCDB

CTF

- High-frequency fine-grained TPC SP calibration (before TF processing)
- Full tracking/matching, PID in barrel detectors, global fits, primary vertex finding
- EMCAL/PHOS partial (cells selection) or full (clustering, matching to tracks) reconstruction
- MUON/MFT tracking, matching
- V0-finding, physical event building

AOD (~15% of CTF size)

Major reconstruction tasks

		CPU	GPU
TPC	Tracking (*)	done	done
	dE/dX	Q1/2019	Q2/2019
	Compression	Q1/2019	Q3/2019(*)
ITS	Tracking finding	done / extra passes: Q2/2019	done
	Track fitting	done	Q2/2019
	ITS-TPC matching	done / afterburner: Q2/2019	Q3/2019
	Compression	Q1/2019	Q3/2019(*)
TRD	Matching to ITS-TPC	done	Q4/2018
TOF	Matching to ITS-TPC	done (in validation)	Q2/2019
EMCAL	Clustering	Q2/2019	-
PHOS	Clustering	Q2/2019 ?	-
MUON	MCH clustering, tracking	Q4/2019 ?	-
	MID	done (in validation)	
MFT	Tracking (standalone)	Q1/2019	-
	Matching to MCH	Depends on MCH schedule	
FIT	T0+ reconstruction	done (in validation)	-
	V0+ reconstruction	?	
HMPID	Clustering, matching	?	-

(*) TPC reconstruction is operational as DPL device, others still need to be interfaced to DPL.

While wrapping a task to DPL device is relatively straightforward (after gaining some experience), to minimize changes and facilitate benchmarking it is better to make them DPL ready from the very beginning.

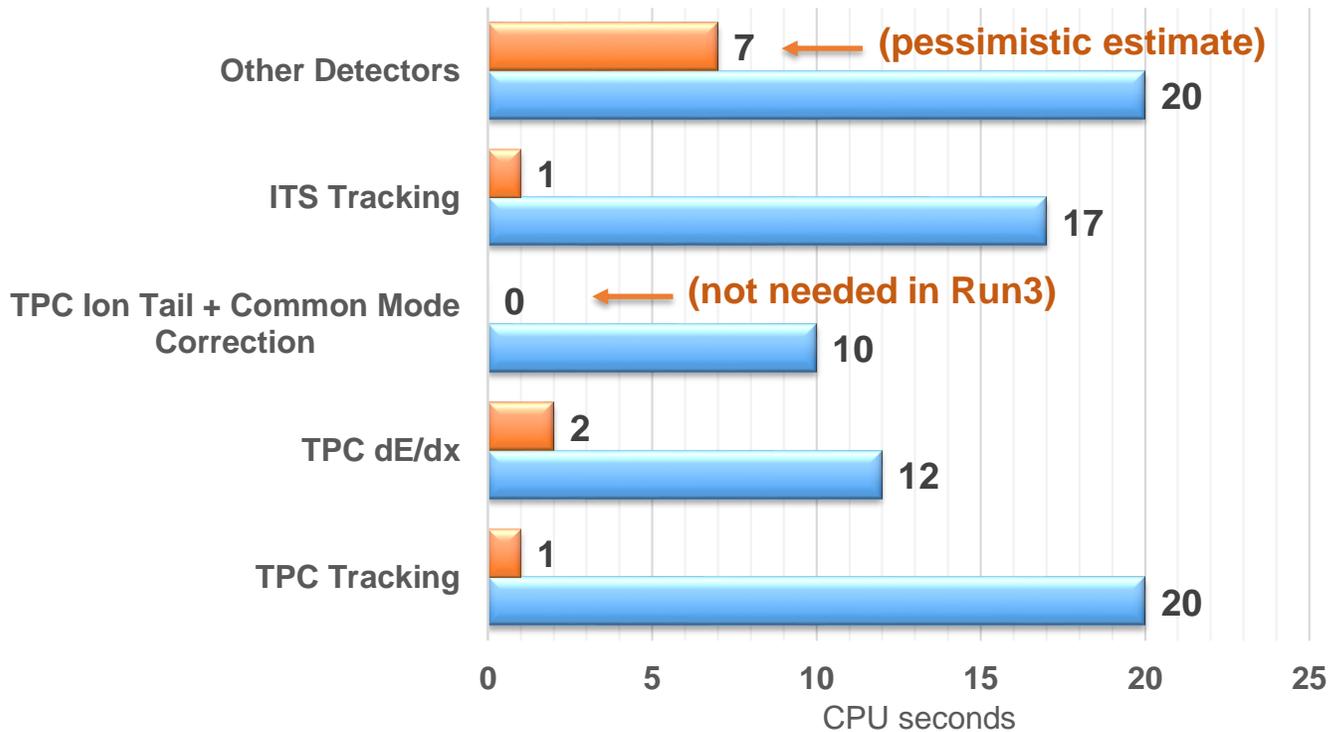
(*) Feasibility of entropy compression on GPU is under study

See details of developments foreseen on GPUs in David's talk

Reconstruction Performance Improvements for O2

Time for 1 minimum bias Pb-Pb events, 1 CPU core

Run2: 79s → Run3: 11s



Time estimate dominated by “Other detectors”, which assume only factor 3 speed up wrt AliRoot → most probably the real time will be a fraction of that

Expected reconstruction output (AOD)

- To keep the output AOD size ~15% of CTF (rather than >150% with ESD format), we need to sacrifice:
 - option of drawing the trees interactively (need data bit-wise packaging and containers optimization, will be handled by WP1)
 - significant part of current ESD/AOD content (multiple ExternalTrackParam parameterization, large bitmasks of TPC padrows, lot of redundant information).

- Very [preliminary format](#) leads to
 - ~80 B (after compression) per barrel track (>80% of AOD size) containing essentially
 - Parameterization + errors at the vertex (can be propagated) + at most 1 extra (outer) parametrization w/o errors or extrapolated coordinates
 - Ncl counters (pattern for ITS), χ^2 for fits, dE/dx and TOF info, flags and indices (instead of pointers).
 - No V0s stored. Even if V0-finding done at reconstruction stage, only track indices can be stored
 - Calorimeters (only EMCal responded) will store only cells which have chance to end up in a cluster (possibly also an index for cluster attachment)
 - Muon will store ~190 B/track and 20 B/cluster (~3% of total data)

Calibration components

In preparation

- TPC distortions calibration with track residuals, ETA: Q4/2018-Q1/2019
 - Extraction from filtered residuals is being ported from AliRoot, then the ITS/TRD+TOF interpolation and residuals extraction will be implemented
 - Container for parameterized corrections is implemented (to be validated)
- TPC VDrift calibration in synchronous stage: ETA: Q4/2018-Q1/2019
(fine calibration produced offline in TPC distortions extraction)
- High-frequency (~5ms) TPC distortions calibration with digital currents, ETA: Q2/2019
 - Both CPU and GPU versions in development
 - Current understanding: 4 transient maps per TF produced from currents on-the fly immediately before the processing of the TF (good for CCDB!)
 - Rough estimates: 100 CPU s per map
- TOF channels calibration: Q1/2019

DCS data usage survey

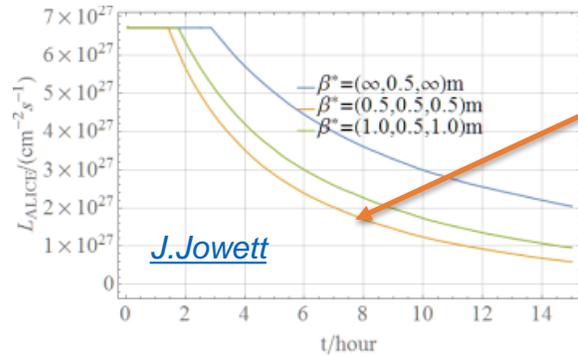
was started but filled by very few detectors

<u>Detector</u>	Type of data (configuration, HV, LV, P...)	Amount of data per timeframe	Amount of data per run (if not per timeframe)	Multiple (M)/ Single (S) value in a timeframe	Processing level (FLP, EPN-sync, EPN-async, Post-processing of sync. stage data, None)	Type of process that will use the data (calibration/simulation/reconstruction/QA)	Acceptable latency in case of use in synchronous mode	Acceptable resolution on the timestamp associated to the measurement	
Det0	Temperature	10 kB		Multiple	FLP	calibration	5 s	0.1 C (or 1%)	
	Current	20 kB		Multiple	EPN	calibration	1 ms	2%	
	...								
TOF	Configuration	N/A	~2 MB	None - SOR	MC	simulation	EOR	N/A	should go to CCDB
	Thresholds	N/A	40 kB	None - EOR	MC	simulation	EOR	N/A	should go to CCDB
	Currents	N/A	6 kB/10 mins	None - EOR	MC	simulation	EOR	1 s	should go to CCDB; 1 v
	HV	N/A	6 kB/10 mins	None - EOR	MC	simulation	EOR	1 s	should go to CCDB; 1 v
ZDC	Configuration	N/A	~few tenths of kb	S	EPN	calibration/reconstruction	EOR	100 ms	to CCDB
	Thresholds	N/A	~few tenths of kb	S	MC	simulation	EOR	N/A	to CCDB
	Detector table positions	N/A	~few tenths of kb	S	MC	(simulation)	EOR	N/A	to CCDB
	HV	N/A	~few tenths of kb	S	MC	(simulation)	EOR	N/A	to CCDB
EMC	Temperature	N/A	~3.5kB/min	S	postprocessing	calibration/analysis	EOR	1s	to CCDB
	Configuration	N/A	2kB	S	postprocessing	simulation/reconstruction	EOR	N/A	to CCDB
	STU-TRU error counter	N/A	1 kB/min	S	postprocessing	reconstruction	EOR	1s	to CCDB
HMPID	HV	N/A	~ few KB	None - EOR	EPN	reconstruction	N/A	N/A	to OCDB
	Pressure	N/A	~ few KB	None - EOR	EPN	reconstruction	N/A	N/A	to OCDB
	Freon transparency	N/A	~ few KB	~one per month	EPN	reconstruction	N/A	N/A	to OCDB

Summary

- Good progress on heaviest reconstruction tasks but the status/benchmarks for most of small (in AliRoot) tasks are not clear yet
- By the mid-2019 we should have the reconstruction devices “mostly ready”: functional and providing realistic benchmarks (exceptions?)
- Calibration components are almost untouched yet (except the heaviest ones, requiring R&D)
 - Detectors are asked to provide details of the calibration components requirements (using DCS data, where they are running, dependencies on other detectors) and start implementing them asap (to have them ready by the end 2019 to allow benchmarking, fine-tuning etc)

Computing resources requirements / usage efficiency



$\beta^* = (0.5, 0.5, 0.5)\text{m}$ scenario
 $L_{\text{int,annual}} = \eta \langle L \rangle T_{\text{run}}$
 $= (50\%)(3.0 \times 10^{27} \text{ cm}^{-2}\text{s}^{-1})(24 \text{ day})$
 $\approx 3.1 \text{ nb}^{-1}$ (c.f. target of 2.85 nb^{-1})

Operation efficiency in last Pb-Pb run in 2015 was 62%. Even higher in 2016 p-Pb.

8-35% more statistics than expected in TDR

Requirements for O2 farm @ P2 are defined by synchronous processing of peak IR:

- At 50 kHz leveling $L_{\text{max}} / \langle L \rangle \sim 0.25 \rightarrow \frac{1}{4}$ of O2 farm power is used in average
- At 40 kHz leveling: lose 10% data (for 5h SB fill, less for longer) but would have better quality with 20% smaller farm used with higher efficiency

Computing power needed:

- Synchronous mode: will need 1500 modern GPUs to cover sync. processing of TPC, ITS, TRD and TOF. Assuming 12 CPU cores per GPU for the tasks which cannot be run on GPU. Other detectors have no yet estimate for required power: \rightarrow assume (worse case) can be buffered and processed at low-IR/no beam time?
- Asynchronous mode: with 10x more time/event and 4x more efficient use of O2, the PPass should take 2.5x data-taking time on O2 alone (~ 60 days), w/o need of involving the GRID.

