Status on the TPC data sizes and the compression of the compressed TPC time frame

David Rohr
drohr@cern.ch, CERN

Offline Week

10.11.2017

CLUSTER FINDER COMPRESSION

David Rohr, drohr@cern.ch

Cluster finder compression status

- **In 2017, we have deployed a new version of the HLT hardware cluster finder and of the cluster compression.**
  - Improved noise filtering in the cluster finder.
  - Storing cluster charge and width with less precision.
  - → **23.9% less clusters, 13.7% less bits needed per cluster.**
- **On top, we have developed some more features to reduce the noise, which are not yet active in the HWCF:**
  - Can reduce the number of clusters by additional **1.5%** with better noise filtering.
  - We can reduce the clusters by **2.5%** by increasing the cut on qMax.
  - → No effect on tracking / dE/dx observed.

- **Baseline for run 3 compression:**
  - The run 3 compression prototype, achieves a factor of 9.1 on 2015 Pb-Pb data.
    - This already includes the higher qMax cut, and most but not all of the new HWCF features.
    - The prototype needs in average **36 bits** per cluster.
    - → Can expect another **2.5%** on top (but still to be verified).

David Rohr, drohr@cern.ch

DATA TYPES FOR TPC IN RUN 3

David Rohr, drohr@cern.ch

Draft of data types for TPC clustering / tracking / compression in O2

- **TPC produces the bulk of data, so a format that minimizes memory consumption is needed.**
  - **We should avoid data duplication, but some copy steps seem to be necessary.**

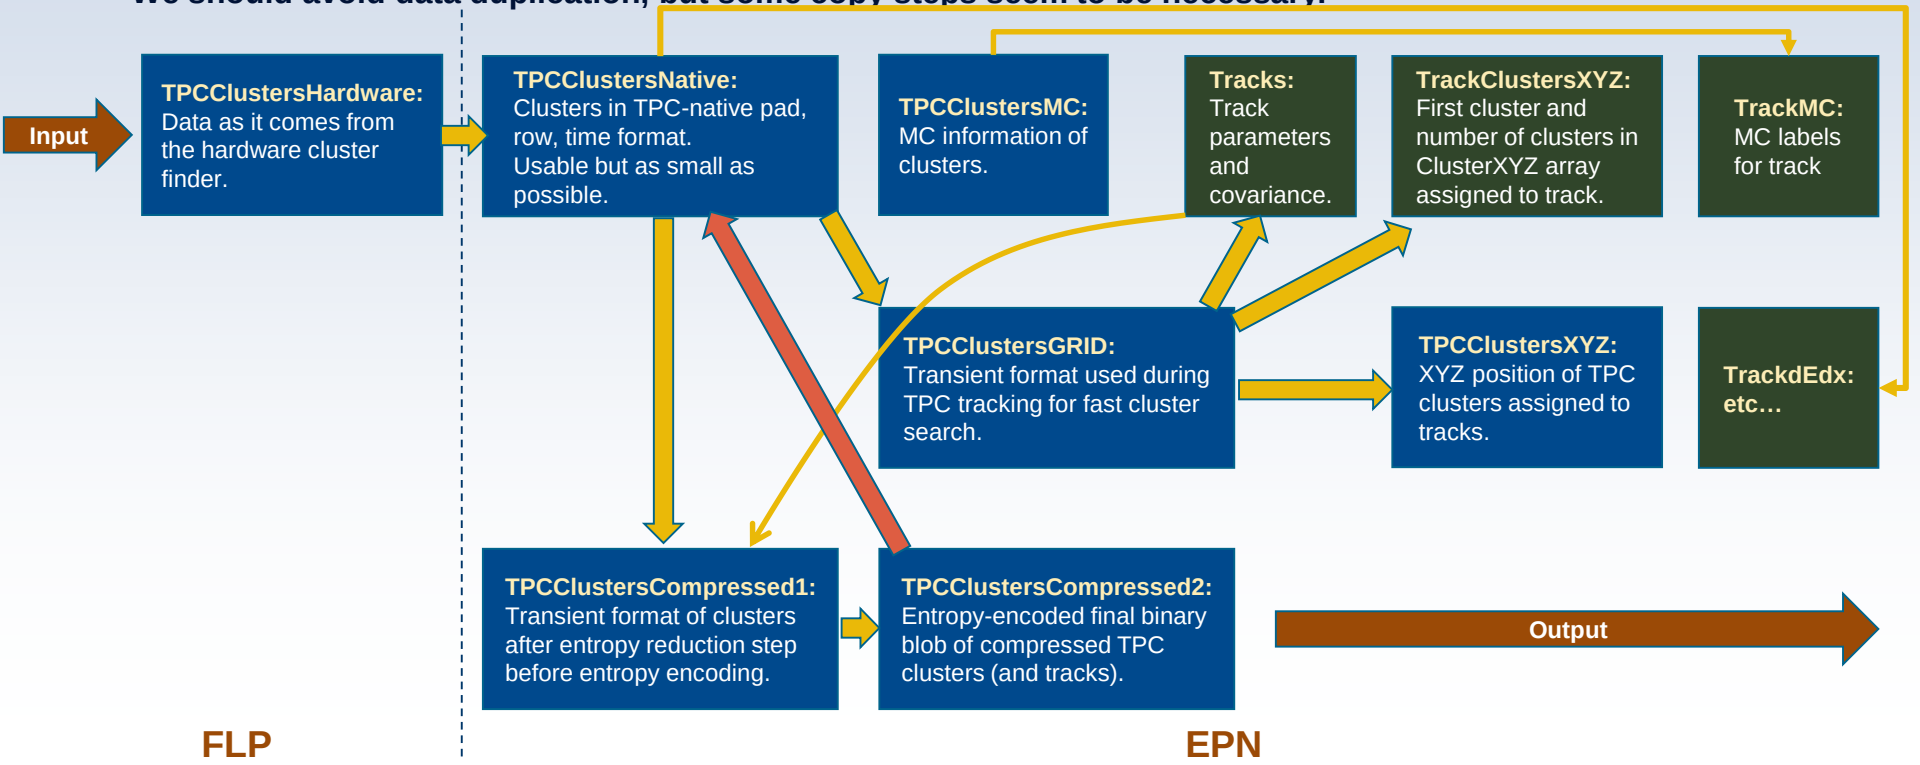

David Rohr, drohr@cern.ch

Draft of data types

- **Draft of members of data types.**
- **Sizes are sizes for unpadded SOA / padded AOS.**

TPCClusterHardware: (22 / 24 byte)

float	pad_pre
float	time_pre
float	sigmaPad2_pre
float	sigmaTime2_pre
uint16	qMax
uint16	qTot
uint8	row
uint8	flags

TPCClusterGrid: (10 / 12 byte)

uint16/uint32	z
uint16	y
uint32	clusterId

TPCClusterCompressed2: (~4.5 byte)

Binary blob

TPCClusterNative: (12 byte)

Uint32	time/flags
uint16	pad
uint8	sigmaTime2
uint8	sigmaPad2
uint16	qMax
uint16	qTot

TPCClusterCompressed1: (15 / 16 byte)

Uint32	time
uint16	pad
uint8	sigmaTime2
uint8	sigmaPad2
uint16	qMax
uint16	qTot
uint16	compressionMeta
uint8	flags

TPCClusterXYZ: (18 byte)

float	x
float	y
float	z
uint32	clusterId
uint8	sector
uint8	row

TPCTrack:

float	param[], cov[], x, alpha

TPCTrackClusters:

uint32	firstCluster
uint16	numClusters

David Rohr, drohr@cern.ch

Draft of data types

- **Draft of members of data types.**
- **Sizes are sizes for unpadded SOA / padded AOS.**

TPCClusterHardware: (22 / 24 byte)

float	pad_pre
float	time_pre
float	sigmaPad2_pre
float	sigmaTime2_pre
uint16	qMax
uint16	qTot
uint8	row
uint8	flags

TPCClusterGrid: (10 / 12 byte)

uint16/uint32	z
uint16	y
uint32	clusterId

TPCClusterCompressed2: (~4.5 byte)
Binary blob

TPCClusterNative: (1

Uint32	time/fl
uint16	pad
uint8	sigmaTi
uint8	sigmaPa
uint16	qMax
uint16	qTot

Just for reference…
For details, see
https://indico.cern.ch/event/678423/contributions/2778515/attachments/1553617/2442101/2017-11-07_HLT_TPC_Hands_on.pdf

ClusterXYZ: (18 byte)

at	x
at	y
at	z
t32	clusterId
t8	sector
t8	row

TPCTrack:

float	param[], cov[], x, alpha

TPCTrackClusters:

uint32	firstCluster
uint16	numClusters

David Rohr, drohr@cern.ch

Draft of data types for TPC clustering / tracking / compression in O2

- **TPC produces the bulk of data, so a format that minimizes memory consumption is needed.**
  - **We should avoid data duplication, but some copy steps seem to be necessary.**

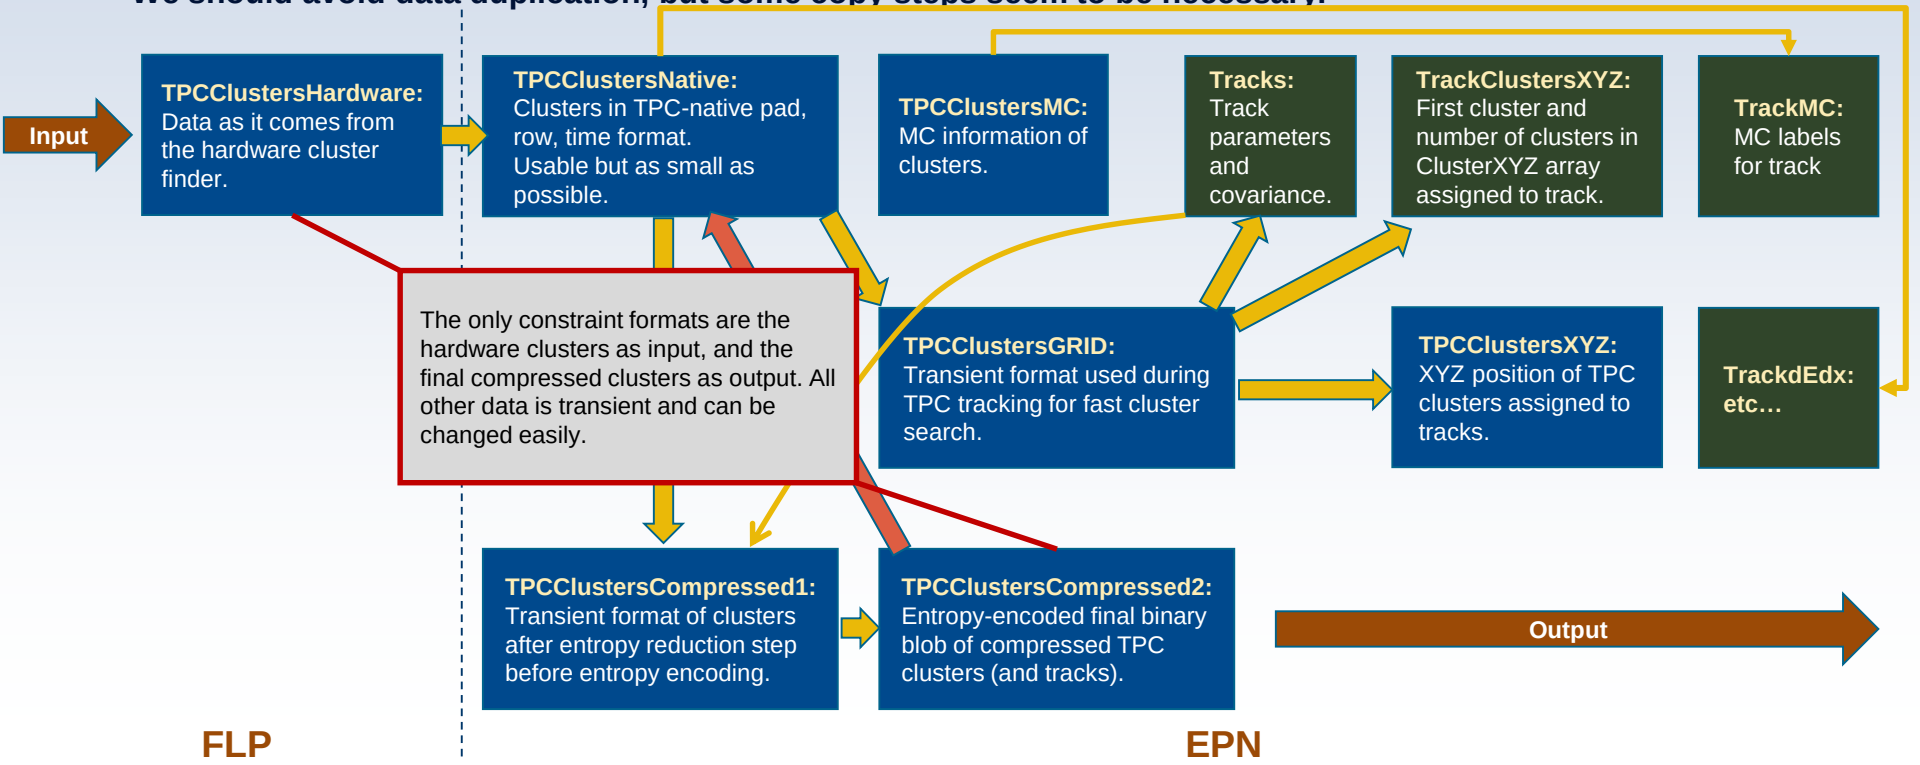

David Rohr, drohr@cern.ch

Draft of data types for TPC clustering / tracking / compression in O2

- **TPC produces the bulk of data, so a format that minimizes memory consumption is needed.**
  - **We should avoid data duplication, but some copy steps seem to be necessary.**

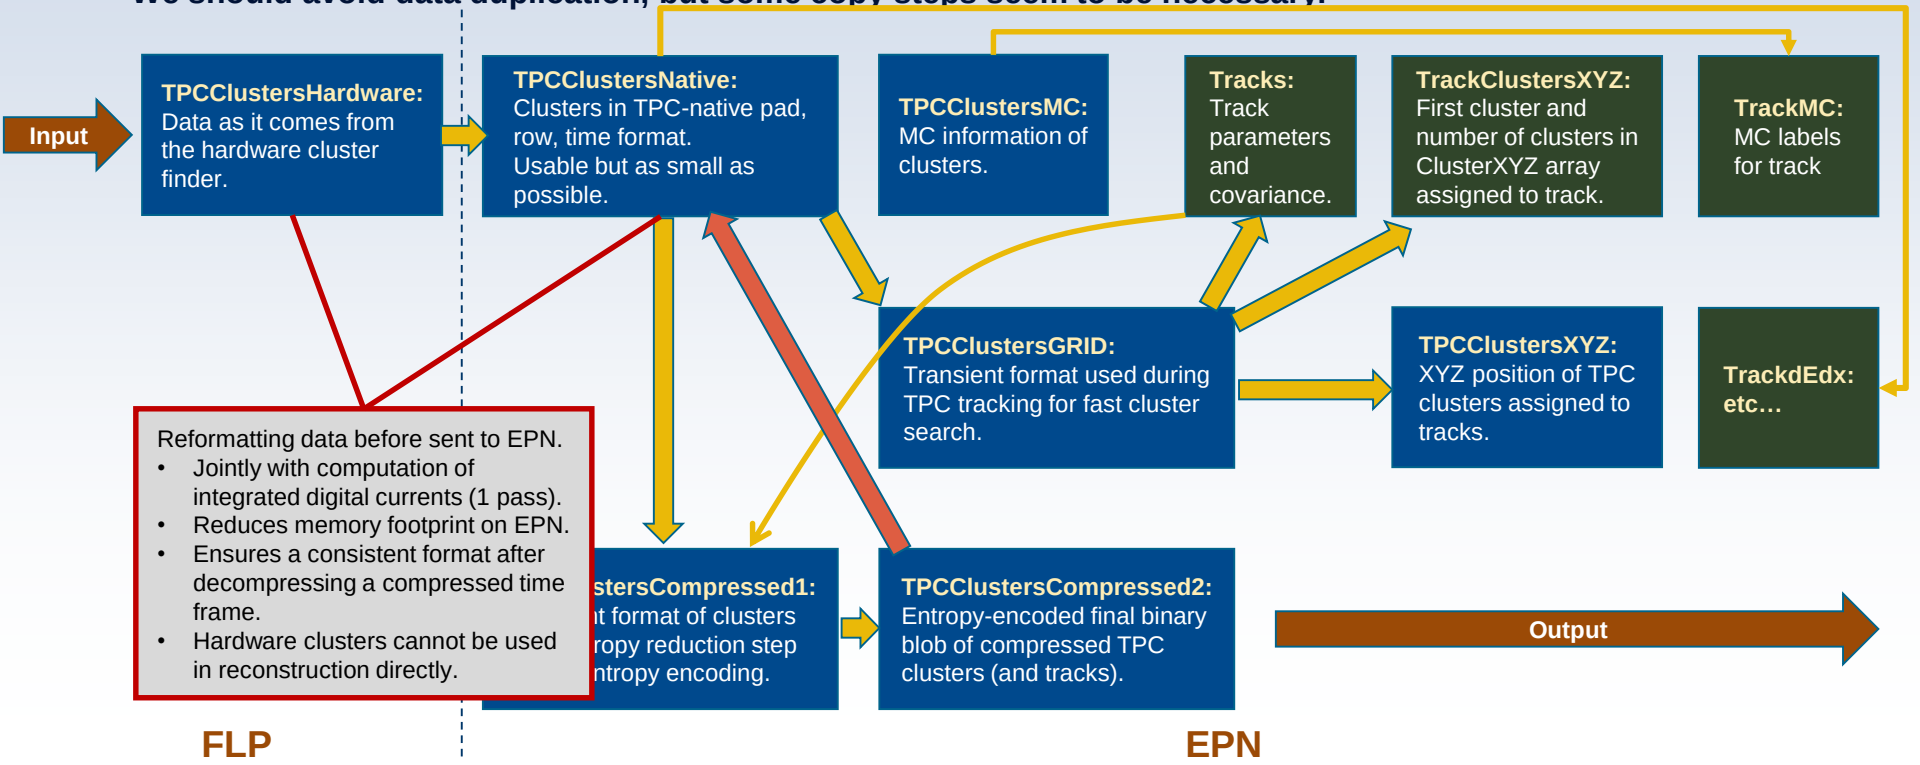

Draft of data types for TPC clustering / tracking / compression in O2

- **TPC produces the bulk of data, so a format that minimizes memory consumption is needed.**
  - **We should avoid data duplication, but some copy steps seem to be necessary.**

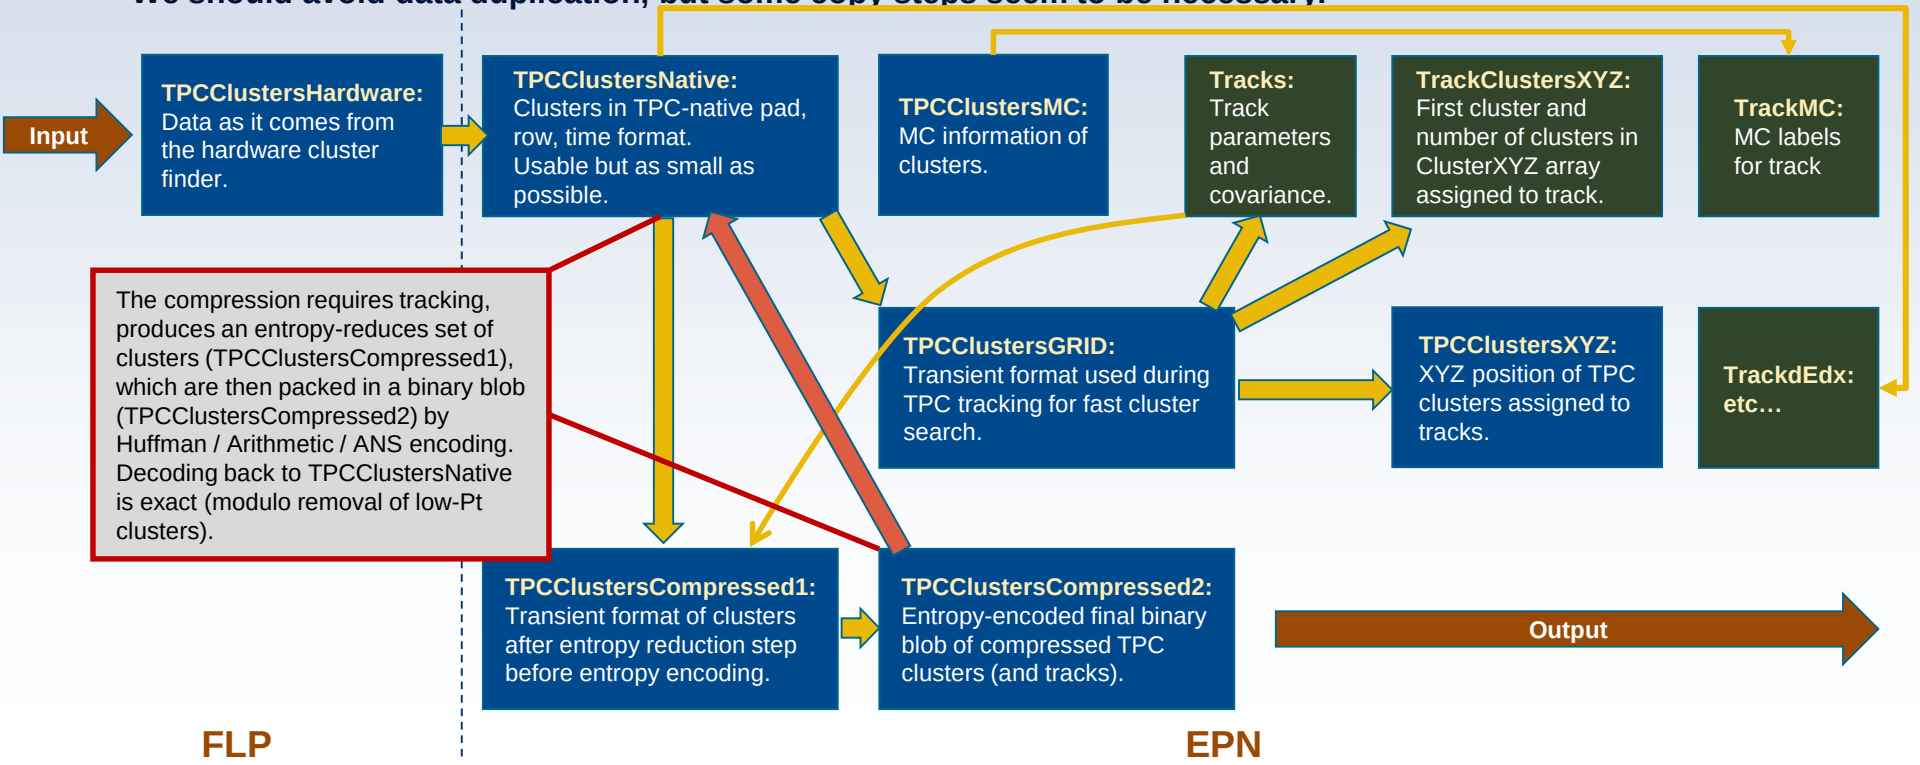

David Rohr, drohr@cern.ch

Draft of data types for TPC clustering / tracking / compression in O2

- TPC produces the bulk of data, so a format that minimizes memory consumption is needed.
  - We should avoid data duplication, but some copy steps seem to be necessary.

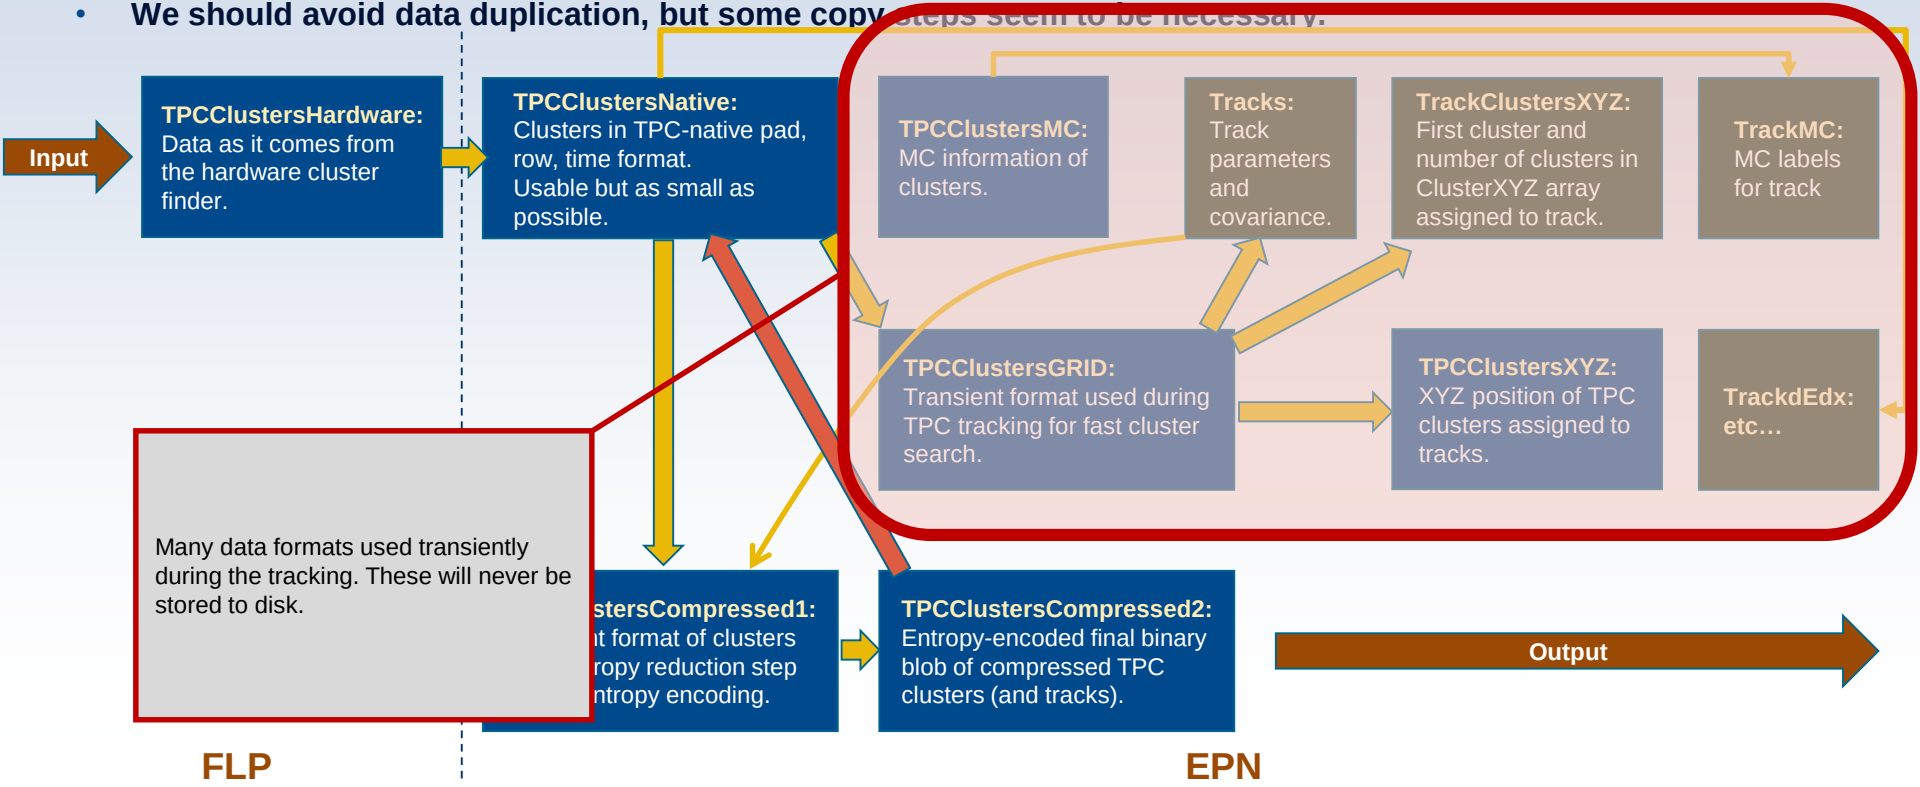

David Rohr, drohr@cern.ch

TPC DATA SIZES IN O2

David Rohr, drohr@cern.ch

Screenshot, pp, overlaid events, 50 kHz

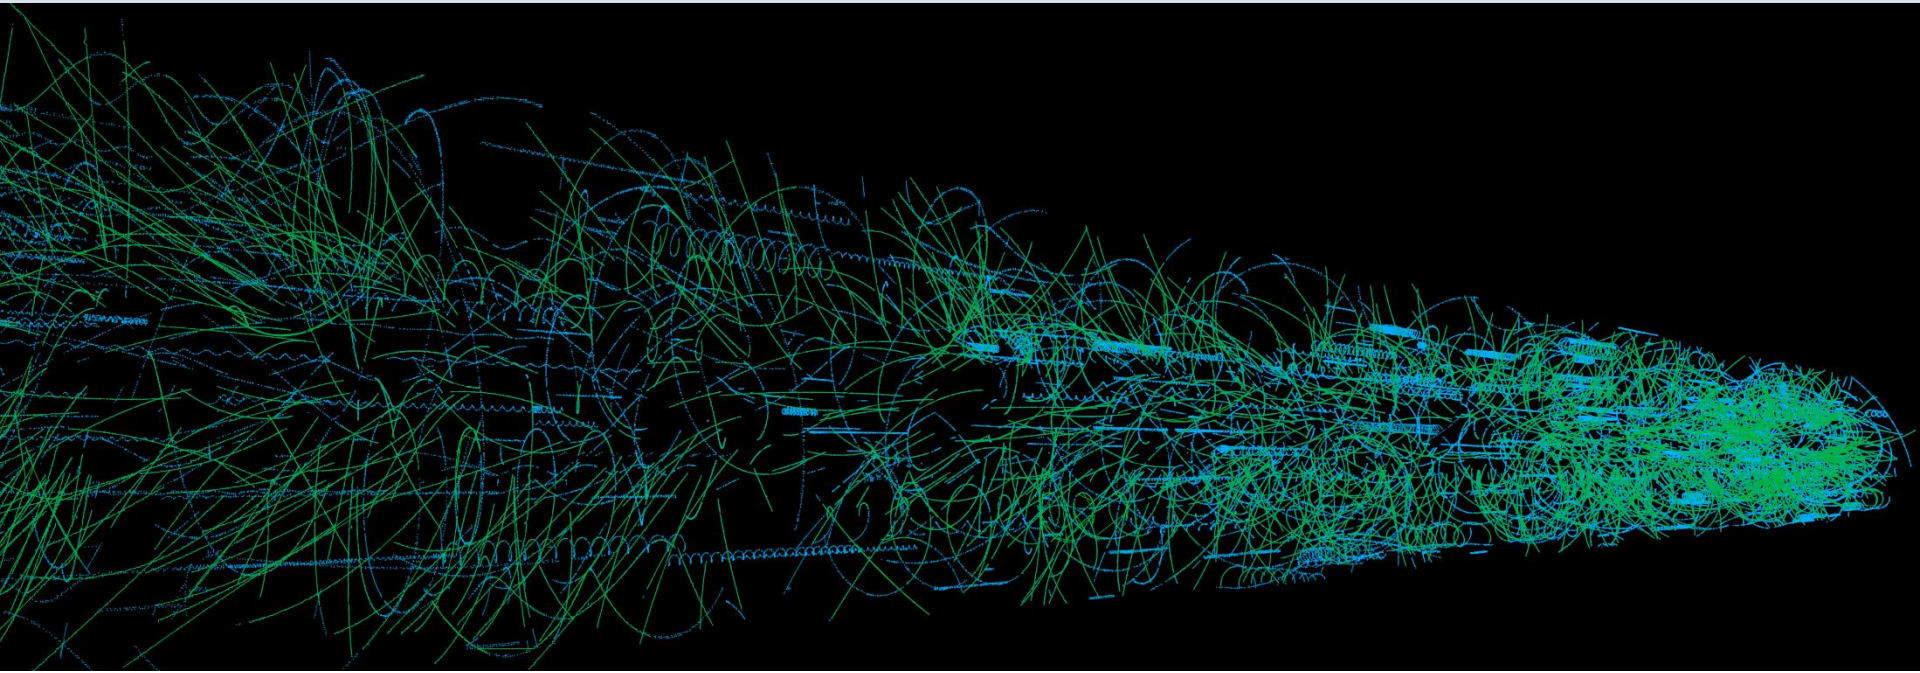

David Rohr, drohr@cern.ch

Screenshot, pp, overlaid events, 50 kHz

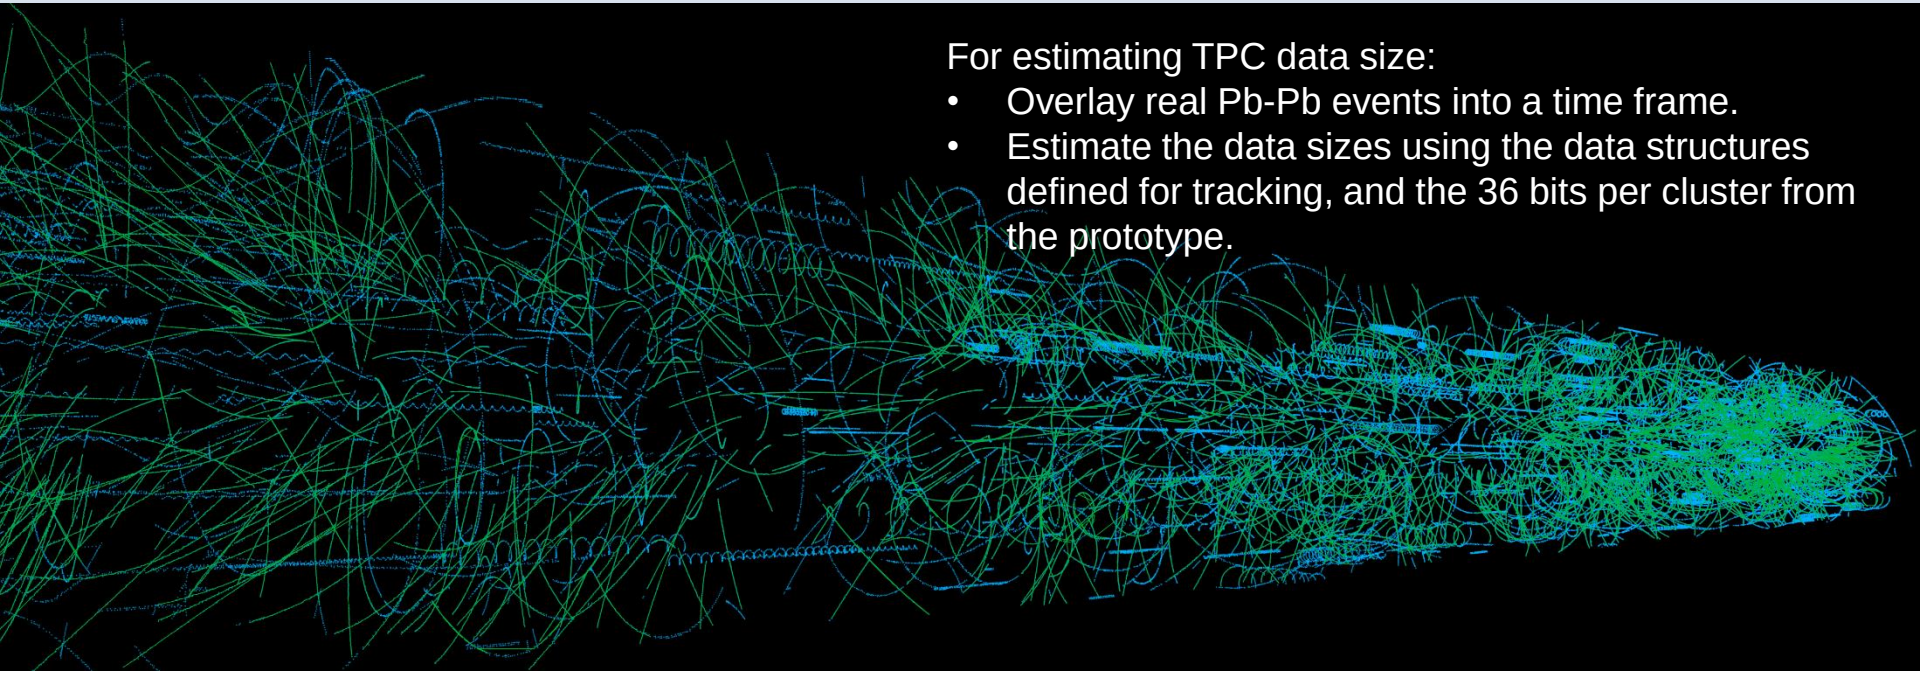

For estimating TPC data size:

- Overlay real Pb-Pb events into a time frame.
- Estimate the data sizes using the data structures defined for tracking, and the 36 bits per cluster from the prototype.

David Rohr, drohr@cern.ch

Data size estimations for run 3

- **Estimate the number of clusters per timeframe:**
  - Ran 2015 Pb-Pb (LHC15o) low-IR (no pile-up) data through new HLT hardware cluster finder.
  - Filtered for minimum-bias events only.
  - Exclude events with less than 1000 TPC hits (QED interactions, with come on top of the 50 kHz.)
  - → In average: 596011 hits per TPC event.
- **Assuming a time frame: (11245 orbits / sec, 1 TF = 256 HB frames (orbits), 22.8 ms).**
  - 678,429,907 hits per TF in average, 4,390,083 TPC tracks per TF
  - We assume 52% cluster attached to tracks, 40% of clusters can be removed because they are not used for physics.
- **Data sizes:**
  - New format yields 12 byte / cluster for the TPC input (most bits already stripped), 36 bit / cluster for the compressed TPC output.
  - → TPC Input data size per TF: **7.58 GB** (**2^{30} notation**).
  - → Memory needed during tracking: **12,01 GB** + additional temporary memory for track seeds.
  - → Size of compressed time frame (all clusters): **2.84 GB**.
  - → Size of compressed time frame (40% clusters removed): **1.71 GB**
- **Data rates (TPC only):**
  - Aggregate input (FLP to EPN): **333.05 GB/s** (**357,61** GB/s (**10^3** notation)).
  - Compressed output rate (all clusters): **124,89 GB/s** (**134,10 GB/s**).
  - Compressed output rate (40% clusters removed): **74,94** GB/s (**80,46 GB/s**).

David Rohr, drohr@cern.ch

What is missing / next steps

- **What was not included in this estimations?**
  - 2015 Pb-Pb data is with argon, run 3 will be neon → less clusters.
  - Run 3 will run a different cluster finder (similar number of clusters assumed as today).
  - Recent tracking improvements attach more clusters to tracks.
    - Track model compression should yield a slight improvement.
  - Run 3 will run at slightly higher LHC energy than Run 2.

- **Next steps:**
  - Develop a wrapper to run current prototype for run 3 cluster finder in AliRoot.
    - Reprocess the mode B data used for HWCF validation with run 3 cluster finder.
    - Run full AliRoot reconstruction and QA on run 3 clusters (as done for HLT HWCF validation).
    - This shall demonstrate the physics performance of the run 3 cluster finder, and will show how many clusters we can expect.
  - Rerun the run 3 prototype with the newest improvements to the cluster finder and the tracking, and see how much the factor 9.1 improves (few per cent at most).