

Documentation of Changes with Respect to the  
Document 'Upgrade of the Readout & Trigger System  
TDR, CERN-LHCC-2013-019, LHCC-TDR-015,  
July 3, 2014, Version updated in June 2017'

for the TRD EDR

20th June 2018

# Contents

1	Addition to 8.2.2: Readout-out with modified data formats . . . . .	3
2	Addition to 8.4.1: Chamber optical transmit power and signal attenuation of optical fibers . . . . .	4
3	Addition I to 8.4.2: Triggering of the FED . . . . .	5
4	Addition II to 8.4.2: Busy handling . . . . .	6
5	Discussion on the location of the CRU (CR1 or cavern) . . . . .	7
6	Data format CRU $\rightarrow$ O2 . . . . .	7
7	TRD specific CRU Firmware . . . . .	8

# 1 Addition to 8.2.2: Readout-out with modified data formats

The final decision on the usage of the 128 bits that can be filled per MCM and per trigger has been taken.

The maximum number of tracklets per MCM will be reduced from 4 to 3 tracklets. Every tracklet consists of 32 bits and the tracklet data format itself will remain unchanged with respect to Run 2. This frees 32 bits which can be filled with additional data. We plan to use  $3 \times 10$  bits to encode additional digits of the PID values in the max. 3 measurable tracklets and thus enhance the PID precision. Effectively, this increases the PID encoding precision from 8 bit to 18 bit. 2 bits out of the total 128 bits will remain unused.

The changes in the data format are illustrated in Figure 1.

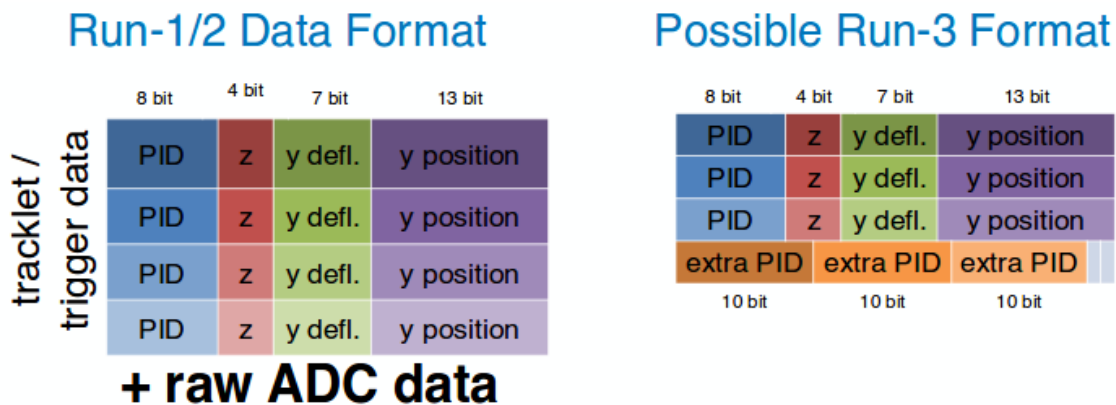


Figure 1: Modification of the data format from Run 2 to Run 3

The impact on data taking performance by reducing the maximum number of measurable tracklets from 4 to 3 per MCM and trigger has been studied in the meantime. Figure 2 shows the relative MCM multiplicity yields in different collision systems. The figure has been obtained by extrapolating TPC tracks to the position of the pad planes and counting the number of tracks that traverse a geometrical area equal to the size of readout pads connected to a single MCM. For central Pb-Pb collisions at  $\sqrt{s_{NN}} = 5.02$  TeV, we observe that the fraction of cases with 4 and more tracklets per MCM is within the order of 2 %. For peripheral Pb-Pb collisions and p-Pb collisions at  $\sqrt{s_{NN}} = 5.02$  TeV, the fraction is within the order of  $10^{-6}$ . We consider this impact on the tracking efficiency to be negligible.

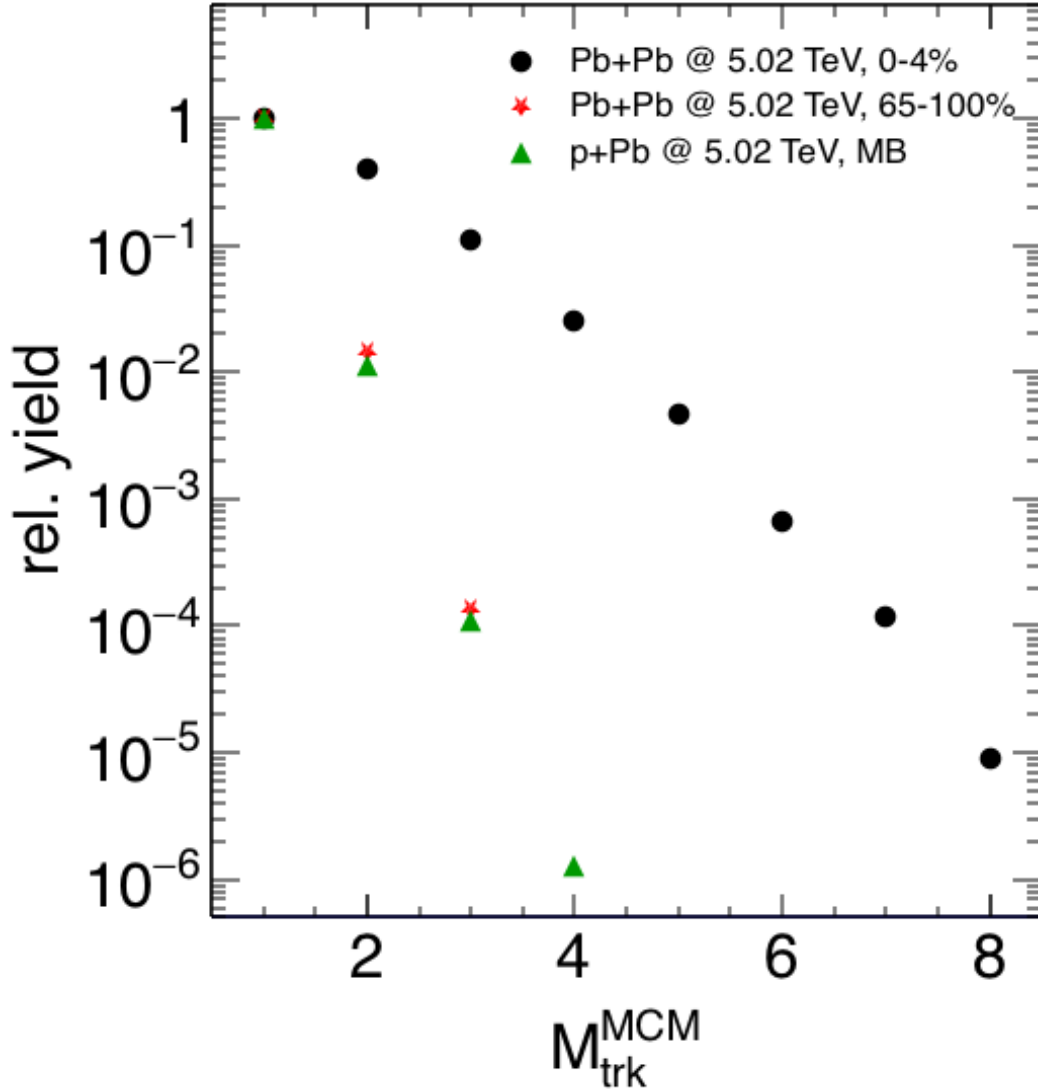


Figure 2: MCM multiplicity yields for different collision systems

## 2 Addition to 8.4.1: Chamber optical transmit power and signal attenuation of optical fibers

As announced in the TDR, additional tests concerning the chamber optical output power and the compatibility with a long optical link have been carried out.

In order to mimic the final setup, two long optical fibers were installed between the TRD patch panel in the UX25 cavern and CR1 for test purposes. This installation includes all necessary patch cables and connectors as they will be used in the final setup. The test fibers were strictly chosen to be of the same fiber quality (OM3 multi mode fiber) as they are ordered for Run 3.

Since the CRU and its final optical component were not available for the test, the signal attenuation on the long fiber was measured by connecting both test fibers in CR1 and looping the signal back to the current electronics. Comparing the optical powers seen by the receiver with and without loop yields a one-way loss of 30 % due to fiber length and connectors.

In addition, the data from a test chamber could be received successfully in CR1 using an Arria 10 development kit. The chamber used in the test delivered 578  $\mu\text{W}$  of optical power at the GTU without loop, translating to 405  $\mu\text{W}$  in CR1.

Unfortunately, data from a few chambers with optical powers as low as 100  $\mu\text{W}$  at the GTU level could not be read out in CR1 using the Arria 10 development kit.

The optical sender diodes of the chambers are mounted on ORI boards (Optical Read-out Interface) and are specified up to a maximum output power of 1000  $\mu\text{W}$ . However, as a precaution in order to avoid ageing effects or early diode failures, the output power was intended to be kept as low as possible. Therefore, the ORI boards allow a remote calibration of the power. After the installation of the first supermodules, a calibration procedure has been developed which automatically calibrates the output power of several chambers such that the optical power measured at the GTU level equals to 100  $\mu\text{W}$ . This calibration procedure needs to be initiated manually and it has only been done for the first installed supermodules. The supermodules installed later have never been calibrated because the optical component on the GTU turned out to be able to reliably receive data even from links with very low optical power. This is why the optical power currently varies across the chambers. We plan to calibrate all chambers during Run 3 commissioning to a common and sufficient value without unnecessarily expose the diodes to high currents. We have no reason at this stage to assume that the optical power of all chambers cannot be adjusted to the level required by the CRUs located in CR1.

Since the final version of the CRU card will come with a different optical component than the Arria 10 development kit, the precise minimum optical power threshold which the CRU requires, could not be measured yet. It will be measured, when the TRD receives the first CRU and will serve as a baseline for the final calibration of all chambers.

### **3 Addition I to 8.4.2: Triggering of the FED**

The requirement of the TDR to emulate the current L1 trigger pulse for the full TRD readout (for calibration purposes) comes with the disadvantage that the pulse needs to be precisely timed after the tracklet shipping is completed in order to make the FED handle it correctly. To overcome this issue, a special trigger sequence, known as test pretrigger sequence, will be used to trigger the CALIBRATION readout instead. For the test pretrigger sequence, the timing is not critical and it can be used similar to the PHYSICS trigger. It will signal to the FED that the next event should be read out fully, i.e. with full tracklet and raw data.

In case that the TRD runs in standalone mode without the CTP, the different trigger sequences won't be available from the CTP. Therefore the TRD will support a second method to change between PHYSICS and CALIBRATION readout. In this mode, all MCMs will use local but synchronized counters and a programmable number  $n$ . After every  $n$  PHYSICS events, one CALIBRATION event is sent.

## 4 Addition II to 8.4.2: Busy handling

The TRD team envisioned early on using a special busy merger unit that links together the individual CRU busy contributions before sending one common busy upstream to the CTP via a dedicated link. This option was put into question in the meantime since all CRUs are already connected to the CTP via TTC-PON. The TTC-PON link also supports upstream data and can therefore replace a dedicated busy upstream link to the CTP. In that case, the CTP receives 36 individual busy contributions and has to link them together by an 'or' operation on the CTP side. This solution also comes with the advantage that it only requires standardized hardware that is common to all detectors and eliminates the need of highly customized hard- or firmware mitigating the risk of future system failure and simplifying maintenance.

The TTC-PON busy upstream latency has been investigated. Latency and system costs highly depend on the number of CRUs connected to each optical TTC-PON network and thus on the total number of optical networks. The reason for this is that TTC-PON uses time multiplexing for the link arbitration to the individual CRUs. If  $n$  is the number of the CRUs per optical network, the upstream multiplexing latency is:

$$t_{\text{Latency,multiplex}} = 125 \text{ ns} \cdot n + 425 \text{ ns}$$

The additional constant 425 ns comes from the multiplexing between the CTP and the LTU. In case that the CRU is placed in CR1 and assuming a 120 m long optical TTC-PON fiber, additional 600 ns of optical propagation delay have to be added in order to obtain the total upstream latency.

Every optical network requires its own Optical Line Terminal (OLT) on the CTP side and its own optical splitter. Every OLT comes at the cost of 650 USD and splitter prices depend on the number of optical connectors. The cost of the optical networks is covered by the TRD upgrade CORE cost.

Table 1 shows example calculations for different numbers of CRUs per TTC-PON network and assuming the CRUs to be located CR1.

The last option of using 9 OLTs corresponds to the maximum what technically can be connected to the LTU. We will opt for it since it provides us the lowest possible latency.

Table 1: Example calculations: Latency vs. Cost

$n$	Latency	Number of OLTs	Splitters	Total cost
36	5.475 $\mu$ s	1	1 $\times$ 1:64	1 $\times$ 650 USD + 650 EUR
18	3.225 $\mu$ s	2	2 $\times$ 1:32	2 $\times$ 650 USD + 2x 350 EUR
9	2.1 $\mu$ s	4	4 $\times$ 1:16	4 $\times$ 650 USD + 4x132 EUR
4	1.475 $\mu$ s	9	9 $\times$ 1:4	9 $\times$ 650 USD + 9 splitters

## 5 Discussion on the location of the CRU (CR1 or cavern)

The busy upstream latency can be reduced by installing the CRU in the UX25 cavern instead of CR1 as suggested in the TDR. The FLP servers and CRUs would then be housed in C-racks which are now occupied by the GTU. Both options come with different advantages and disadvantages regarding various aspects. They are summarized below:

### CRU in CR1:

- Additional 600 ns busy upstream latency from the long fiber
- Easy maintainability and permanent accessibility
- GTU can stay in the C-racks and may be reused at a later stage
- No issues with radiation or magnetic field

### CRU in cavern UX25:

- Saves 600 ns of latency
- Access only possible with no beam in the LHC
- GTU has to be uninstalled
- FLP and CRU are not radiation tested, magnetic field could cause problems to fans in the FLP

## 6 Data format CRU $\rightarrow$ O2

The FED will send data to the CRU without any preceding header. Therefore the CRU has the task to packetize the data and equip every packet with a header before sending it to the O2 system. The header identifies to which trigger the data belongs and from which link the data were sent. For the header format, the TRD tries to comply as much as possible with the standard O2 header format, also called the RDHV2. For details, see Figure 3.

# RDH V2

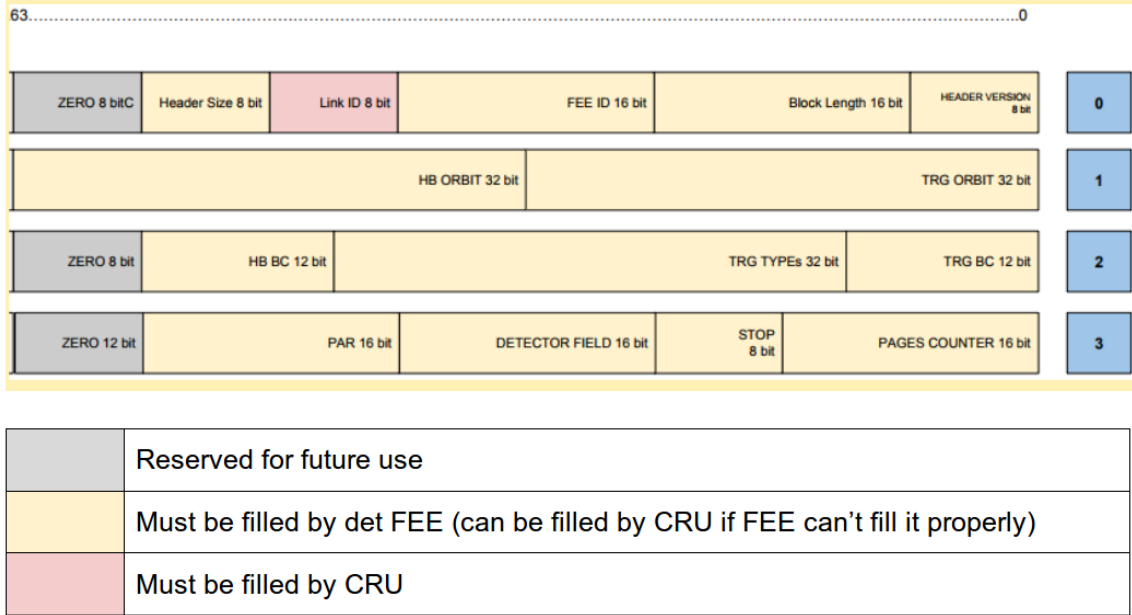


Figure 3: Raw Data Header, version 2

In principle, the raw data header specification requires the detector to chop the data such that for every trigger and for every link at least one packet is created. On the other hand, typical packet sizes should be close to the absolute maximum packet size of 8 kB in order to not deteriorate the performance by sending too many small packets and hence producing a large overhead. For the TRD in PHYSICS mode, the event size per link and trigger is limited to 1 kB for technical reasons, even in case of high multiplicity PbPb events. In case of low multiplicity events, the average event size per link is expected to be even smaller. Every CRU DMA channel has to handle the data from 15 links which could result in performance issues at high event rates.

As a tradeoff between RDHv2 compatibility and large packets, the CRU fills the field 'Link ID' with zeros and uses the packet payload for data from all 15 links. To distinguish data from different links offline, the data from a particular link are encapsulated into smaller units of a fixed 256 bit size. The most significant 16 bits out of every 256 bit word are used to encode the link id.

## 7 TRD specific CRU Firmware

Contrary to detectors using the standard GBT protocol, the TRD uses a custom link data format for the transmission of data from the FED to the CRU. The data format is based on an 8b10b encoding scheme. This means that the TRD CRU firmware requires not only a custom user logic, but also a custom link wrapper. To save resources, the TRD will use a custom bitfile, where unused elements of the common CRU firmware are removed



and replaced by the TRD specific logic. Figure 4 shows a block diagram of the TRD specific logic as it is currently implemented. All parts of the firmware have been tested in hardware simulations. Figure 5 shows an example of a simulation of the transmission of a test counter via the ORI link.

The link wrapper was in addition tested in the lab using an Arria 10 development kit and an ORI sender board from the detector. It also served as test equipment in the optical power tests, as described in the addition to chapter 8.4.1.

Further lab tests of the logic using a pseudo random data source are planned, but they require a test installation with a full TRD chamber and a CRU dedicated to the TRD. The test installation is prepared and we are currently awaiting the delivery of the CRU.

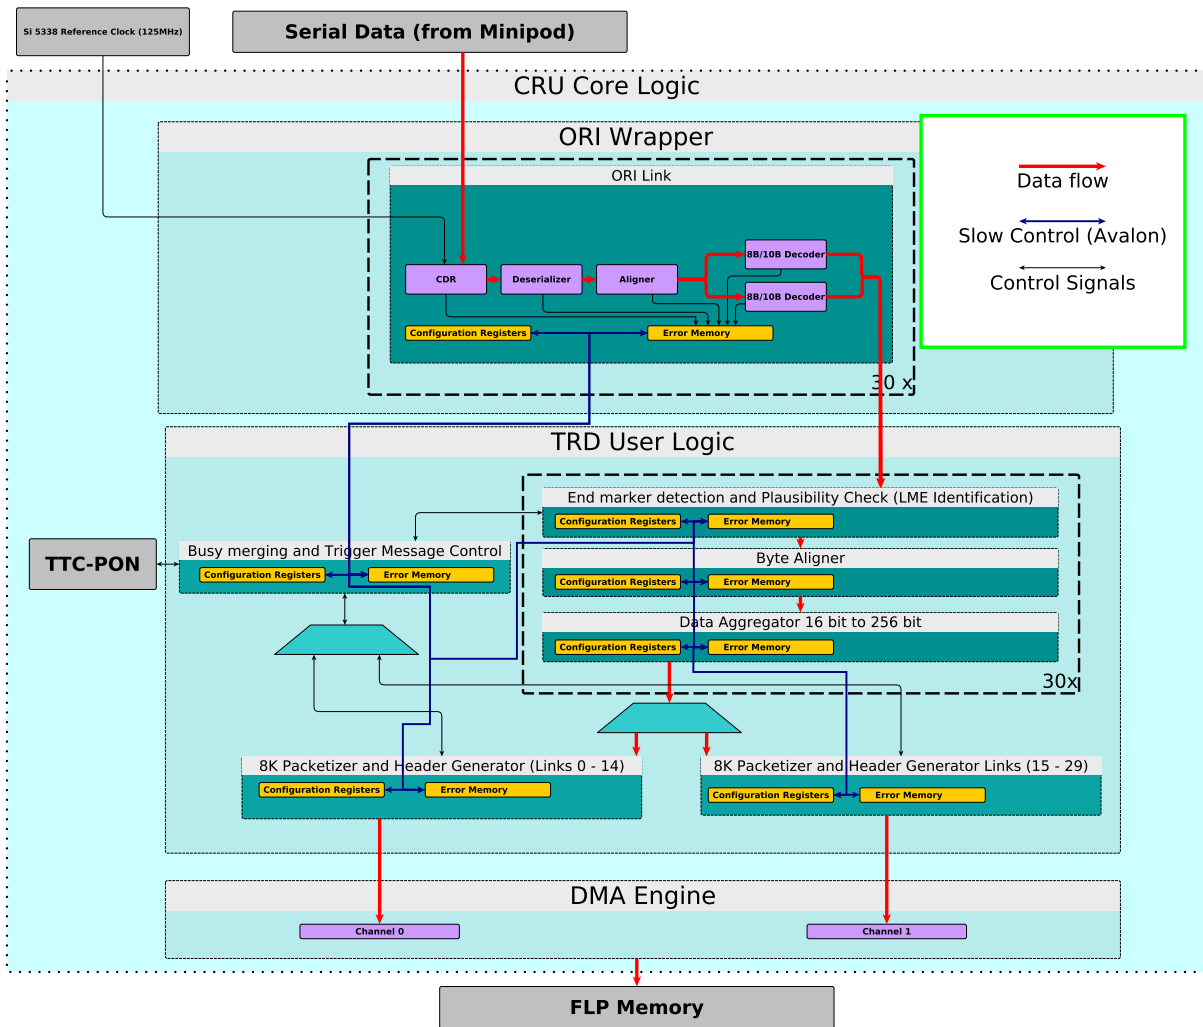


Figure 4: Block diagram of the TRD specific logic




Figure 5: Modelsim simulation of the transmission of a test counter via the ORI link. The counter data is encoded by an 8B10B encoder, interspersed with alignment characters and serialized (sender side). In the CRU, the sender clock is recovered and the data words are aligned to the word boundaries (receiver side)