Upgrade of the Readout & Trigger System TDR, CERN-LHCC-2013-019, LHCC-TDR-015, July 3, 2014, Chapter 8, Transition Radiation Detector - TRD

Version updated in June 2017

Contents

Acknowledgements 1

List of Figures 2

List of Tables 3

Chapter 8. Transition Radiation Detector - TRD 4

8.1 TRD upgrade strategy 4

8.2 Frontend operation and read-out 5

8.2.1 Current FEE read-out 5

8.2.2 Read-out with modified data formats 6

8.3 TRD Performance with new data formats 10

8.4 TRD read-out and trigger 12

8.4.1 TRD read-out unit 12

8.4.2 Trigger and busy handling 14

8.5 Schedule, funding and institutes 16

List of Figures

8.1 FEE event processing and read-out sequence in Run1. 6

8.2 TRD tracking performance with online tracklets. 11

8.3 TRD particle identification performance with online tracklets. 12

8.4 Block diagram of TRD read-out unit. 13

List of Tables

8.1 TRD read-out rates and data volume for different TRD data formats and event scenarios. 8

8.2 Event size and read-out input buffer capacity for various event types. 13

8.3 TRD institutes 16

8.4 TRD schedule 16

8.5 TRD funding 17

Chapter 8

Transition Radiation Detector - TRD

8.1 TRD upgrade strategy

The Transition Radiation Detector (TRD) has originally been designed for a Pb-Pb interaction rate of 8 kHz and for a significant event rejection from the level 1 (L1) trigger [?]. The existing processing and read-out of the front-end electronics (FEE) as well as the read-out and trigger functionality are optimised for these conditions and provide a fast L1 trigger contribution, implementing jet and electron triggers.

For the ALICE upgrade, the TRD detector must operate at much higher interaction rates, and the FEE and read-out system must accept the largest possible fraction of interactions without the need to provide a trigger.

Based on measurements in Pb-Pb collisions in Run1, it has been estimated that the chamber currents reach 6 μA at 50 kHz interaction rate. This leads to a total accumulated charge of 0.8 mC per cm of wire per year, assuming an average interaction rate of 50 kHz. As the chambers were validated for charges above 10 mC/cm, it is expected that no ageing effect will occur for the planned running time. The voltage drop at these currents, however, may result in significant gain variations in case of large variations of interaction rate, e.g. over the duration of a fill. No problems on detector stability or concerning space charge effects are expected.

An upgrade of the FEE hardware is not realistically feasible. Besides the design and production effort, it would require a complete disassembly and reconstruction of the 18 TRD supermodules and the FEE mounted on the 522 individual detector chambers. The chosen strategy is a reduction of event read-out time with the existing FEE by changing its mode of operation and limiting the amount of event data read from the FEE as detailed in Sec. 8.2. The impact on performance for tracking and electron identification has been extensively studied (see Sec. 8.3) to validate the proposed strategy.

The read-out of the optimised FEE data format at the full minimum bias event rate requires new hardware with increased bandwidth to the O^2 system as described in Sec. 8.4. The use of the proposed ALICE Common Read-Out Unit (CRU) is envisaged for this purpose.

8.2 Frontend operation and read-out

8.2.1 Current FEE read-out

The TRD FEE [?] is bound to operate in a triggered mode of single event read-out. An initial trigger level (LM, the functionality corresponds to the pretrigger in Run1) fixes the time reference for sampling and processing. A subsequent event can only be triggered after completion of the FEE event read-out or after the abort of the read-out sequence due a negative higher level trigger.

The front-end electronics comprises a hardware preprocessor for the calculation of quantities relevant for the finding of online tracklets, which are track segments in a single detector chamber. The preprocessor provides its results at a fixed time after the sampling has been started by an LM trigger. Further processing is done in CPUs in the FEE.

Figure 8.1 shows the timing sequence for a typical event. To recover the information before the arrival of the LM trigger, the digitised data are delayed in pipeline stages. With a drift time of 2.2 μs and a delay of 900 ns, the processing in the CPUs can start 3.1 μs after the interaction when all data have passed through the preprocessor and its results are available. The processing time depends on the complexity of the calculations. Finding tracklets using the preprocessor results takes about 1 μs.

The FEE read-out is organised in 60 trees (2 per chamber) per supermodule, each with 64 FEE devices (multi chip module - MCM) and equipped with one optical read-out interface (ORI). The read-out can operate in two modes: tracklet mode and raw read-out mode. The tracklet mode is implemented as pure push mechanism up the read-out tree without any handshaking. This avoids latency but is limited to the read-out of four 32-bit words for each MCM. The raw read-out mode has no practical limitation on the number of transmitted words but requires handshaking, which results in a total overhead of 8.32 μs in each read-out tree, in addition to the time for the actual data transfer with 8 bit at 120 MHz.

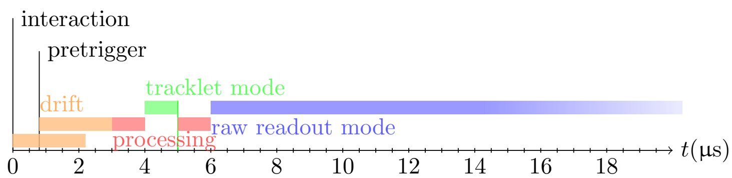

Figure 8.1: FEE event processing and read-out sequence as used in Run1. The event read-out timing is shown for an event with 25% of tracklet words.

The event read-out time, i.e. the time from an interaction until the FEE has finished shipping all data, depends on the FEE processing effort and the maximal data volume in a single read-out tree. With the raw read-out mode, the event read-out time will be of about 16 μs in addition to the transfer time for the actual data volume of a given event. The currently used read-out of full zero-suppressed ADC data results in event read-out time of several 10 μs and puts a severe limit on the maximum read-out rate.

8.2.2 Read-out with modified data formats

New data formats can be implemented within the capabilities of the existing FEE hardware, with the goal to minimise dead time by a reduction of data volume.

- Tracklet Read-out

  A significant read-out time reduction can only be achieved by avoiding the handshaking overhead in the raw read-out mode and by transferring most of the information via the four data words associated to each MCM.

  In this mode, the event read-out time is in the range of 4 μs up to an upper limit of around 8 μs, imposed by the maximum number of words available in this read-out mode.

  The most stringent constraint is the limitation to four 32-bit words (128 bits) per MCM which limits the acceptable local occupancy. Currently, in the tracklet mode one MCM can send up to four tracklet words, each with the following information: z-position (longitudinal) in units of padrow (4 bits), y-position (transverse) in units of 160 μm (13 bits), y-deflection (transverse) in units of 140 μm (7 bits) and PID information (8 bits).

  In order to extend the charge information used for PID, the read-out of 3 tracklets with 18 bits for PID information, or 2 tracklets with 40 bits for PID is foreseen. The bin widths for the position information would remain unchanged. Two charge slices are available directly from the preprocessor without additional delay. More slices could be calculated in the CPUs by looping over the data in the event buffers.

  Running with alternative data formats requires only a change of FEE configuration. Therefore, new formats can be tested and optimised with real data throughout Run2 without major disturbance for normal data taking.

- Raw data read-out

  The careful optimization of the payload of the 4 data words associated to each MCM, avoids the need for an extensive raw data read-out, as it was performed in run 1 and run 2 to a large extent. However, for some offline analyses (e.g. calibration), the availability of the full raw data is indispensable. Regarding that and the requirements on the read-out time, a tradeoff is required.

To keep the impact of the full raw data read-out on the busy time small, the full raw data read-out will only be performed for a small fraction of events. The data set for which full raw data will be available should not exhibit any bias, i.e. no data based online selection criteria will be applied. Instead, tracklet-only and tracklet+raw data read-out mode will alternate according to a fixed periodic sequence. The sequence will be steered by CTP and should be aligned with the LHC orbit signal. In the default configuration, the full raw data of one event per orbit will be read out.

Front-end read-out rates with new data formats

The read-out rate performance of the new data formats is shown in Tab. 8.1 for the case of Pb-Pb collisions which constitute the biggest challenge for the read-out given the large event sizes at comparably high interaction rates of 50 kHz or above.

Cases for the tracklet read-out are shown for the maximum event read-out time of 8 μs and another more typical value of 6 μs. Accepted event rates in the range of 60 kHz can be achieved for 100 kHz interaction rate, which is significantly higher than the accepted rates for any data format using the raw read-out mode of the FEE.

In all read-out scenarios with reduced or tracklet data, the data volume is below 14 Gb/s/sector.

	interaction rate [kHz]	Accepted rate [kHz]	Accepted fraction [%]	deadtime [%]	data volume [Gb/s/sector]
tracklet read-out only					
avg. deadtime 6 μs	50	38.5	76.9	23.1	4.73
	100	62.5	62.5	37.5	7.68
	200	90.9	45.5	54.5	11.17
avg. deadtime 8 μs	50	35.7	71.4	28.6	8.78
	100	55.6	55.6	44.4	13.65

Table 8.1: TRD read-out rates and data volume for different TRD data formats and event scenarios.

For the given read-out rates and data volumes, a Pb-Pb minimum bias raw event size of 210 kB/sector (28.3 kb/event/link) is assumed, derived from the experimental value of 170 kB/sector (2011 Pb-Pb data at $\sqrt{s_{NN}} = 2.76\ TeV$) and scaled to $\sqrt{s_{NN}} = 5.5\ TeV$. For the tracklet event size, occupancies of 25 % (6 μs case) and 50 % of the maximum number of tracklet words (8 μs case) are assumed. The numbers of accepted events are estimated based on FEE read-out time and interaction rate only.

Possible rate limits coming from the increased power consumption at the upgrade read-out rates were investigated in test runs where tracklets were produced artificially by adjusting the FEE baseline and cluster thresholds. Only the digital 1.8 V low voltage channels, which are used for components in the FEE chip that are clocked exclusively during event processing, show a significant dependence on read-out rate. Measured currents are below 150 A for all running scenarios up to 100 kHz read-out rate, well below the 200 A current limit of the LV supplies.

The TRD currently calibrates gain, drift velocity v_d, $E \times B$ effects and time-offset; about 30000 pp or 1500 Pb-Pb minimum bias events are needed to achieve a calibration point. The existing calibration procedures can be preserved with the new data formats by reading the full zero-suppressed ADC data instead of the tracklet words for a small subset of events, with negligible effects on deadtime and data volume. It is also conceivable that calibration could exclusively use tracklet words, doing gain calibration with the charge information available in the tracklet words and integrating other calibration parameters in a global alignment procedure.

As a conclusion, the tracklet read-out scenario would allow - with the existing FEE hardware - the reading out of more than 70 % of events at the envisaged Pb-Pb minimum-bias 50 kHz interaction rate, including also the TRD detector. A study on the impact on tracking and particle identification performance of the new format Pb-Pb is presented in the next section.

8.3 TRD Performance with new data formats

The performance for tracking and PID of the reduced information content of the tracklet read-out scenario described above is assessed by comparing it to the performance of the offline reconstruction based on full zero-suppressed ADC data (ZS) (for details see [**?**]) and TPC seeding. Results from pp data at 8 TeV (production LHC12f) are presented for two tracklet reconstruction scenarios:

- read-out tracklets obtained online as currently used for trigger purposes
- tracklets calculated offline from ZS data with an improved PID content.

Their matching was done with respect to corresponding global tracks by their azimuthal and polar positions at the radial distance of the anode wire of the corresponding TRD chamber. The offline residual misalignment is applied in both cases.

The tracklet reconstruction efficiency for online relative to offline scenarios is presented in Fig. 8.2 (left) for a single pp run. Due to systematic effects induced by drift being perpendicular to magnetic field deflection (E x B effects), positive and negative charged particles are influenced differently. They are therefore shown separately in order to assess the $p_T \sim 1.5\ GeV/c$ threshold above which reconstruction is not affected by particle charge. From Fig. 8.2 (left), we conclude that the TRD contribution to global tracks should remain unchanged within 4 % when using the tracklet read-out format.

The quality of TRD reconstruction for track position in the azimuthal plane with respect to global tracks is presented in Fig. 8.2. The residuals (Δy), obtained chamberwise, are characterised by Gaussian shapes with comparable sigmas, i.e. resolutions (Fig. 8.2 right) for both tracklet reconstruction scenarios. The TRD tracking performance remains unchanged for positive particles above $p_T \sim 1.5\ GeV/c$ and for negative particles above $p_T \sim 0.8\ GeV/c$.

The online tracklet performance at low p_T develops asymmetrically with particle charge due to the missing correction for the ion tails (Tail Cancellation - TC).

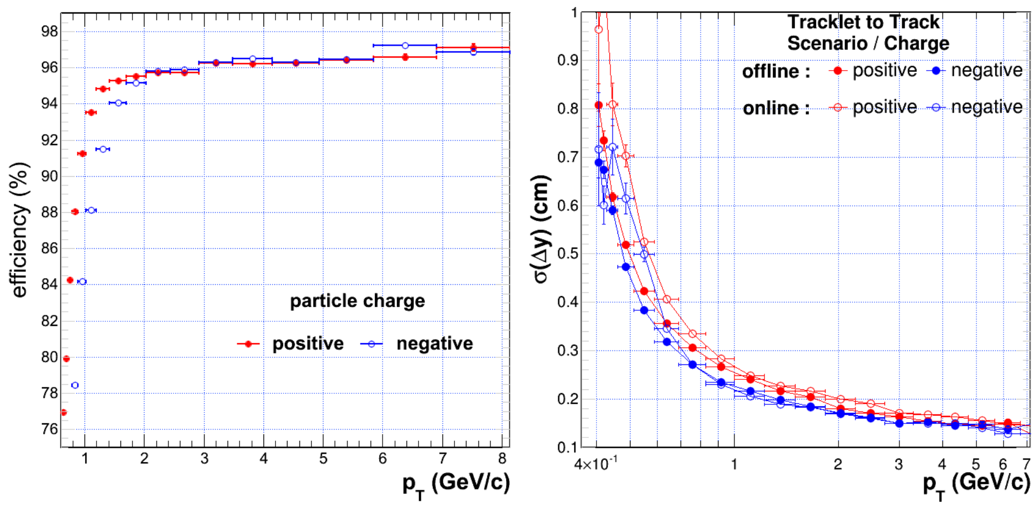

Figure 8.2: TRD reconstruction performance relevant for tracking. The reconstruction efficiency of online relative to offline (left) and the quality of azimuthal residuals (Δy) resolutions (right) for positive [red] and negative [blue] charged particles for the two tracklet reconstruction scenarios (right).

Characteristics relevant for particle identification of the estimation of the track angle, in a single TRD chamber, by the two tracklet scenario are presented in Fig. 8.3 (right). The Gaussian shaped residuals ($\Delta\phi$) are described by shifts with larger values obtained for the online tracklets and are mainly due to missing TC corrections and to limited calibration precision for the drift velocity and $E \times B$ effects.

In the left panel of Fig. 8.3, the particle identification (PID) performance is compared for online reconstructed tracklets optimised for triggering and normalised to global track inclination and offline tracklets, respectively. The pion efficiency at 75 % electron efficiencies for the online scenario is projected on the much higher statistics offline data set using a normalisation factor of 15 %. The target 1 % pion efficiency at 2 GeV/c momenta will be reached at 75 % electron efficiency.

The identification of particle species in this case is done offline based on reconstructed secondary vertices (V_0 candidates) due to photon conversion, K_0 and Λ decays. It is worth noting that the TRD 1-dimensional PID is formed out of two ingredients, the total charge and the track inclination. For online tracklets, PID and inclination (local momentum) cannot be optimised simultaneously. Rather, PID can be best calculated after global tracking is performed using the good online position

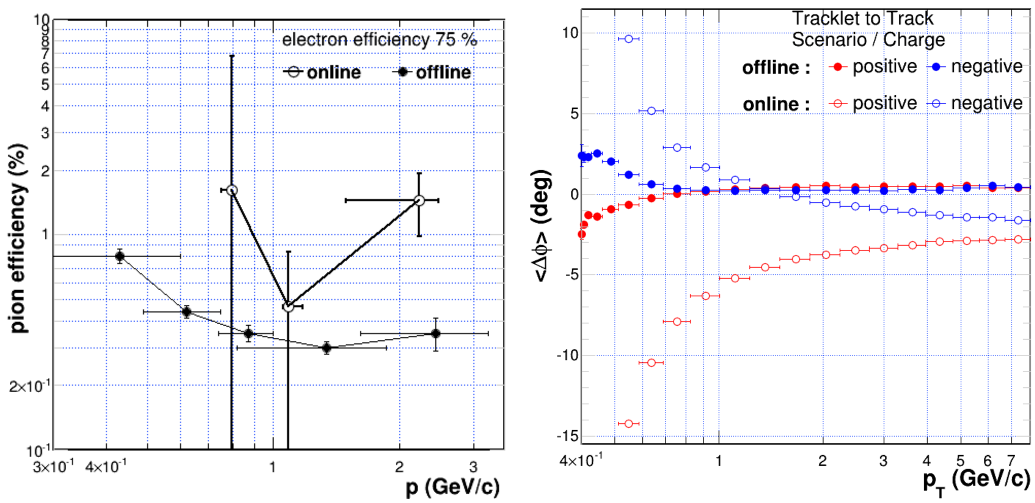

Figure 8.3: TRD reconstruction performance relevant for particle identification. The online [open symbols] tracklet scenario, trigger tuned and normalised, π efficiency for 75 % electron efficiencies and the corresponding ofline [filled symbols] performance for particles registered in 6 TRD layers (left) and the characteristic shifts of angular residuals ($\Delta\phi$) in the bending plane for positive [red] and negative [blue] charged particles for the two tracklet reconstruction scenarios (right).

information. For the upgrade data without TRD electron trigger, this poses no limitation.

8.4 TRD read-out and trigger

8.4.1 TRD read-out unit

In the following paragraphs the specifications for a new TRD read-out unit are presented (see also Fig. 8.4). The major upgrade is the higher bandwidth interface to the O^2 system: instead of the full zero-suppressed ADC data for a small subset of L1 accepted events, the full minimum bias triggered front-end electronics (FEE) data stream has to be transferred. Currently the GTU modules implement the TRD read-out functionality [?].

The expected data volume per sector is below 20 Gb/s (see Tab. 8.1). This translates into two read-out units (RU) per TRD sector, each with one DDL3 link to the O^2 system. The read-out has 30 optical input links, each transferring data from the FEE of one TRD half chamber at a net data rate of 2 Gb/s. In case of tracklet read-out, the data transfer from the FEE is active for less than 12 % of the time for

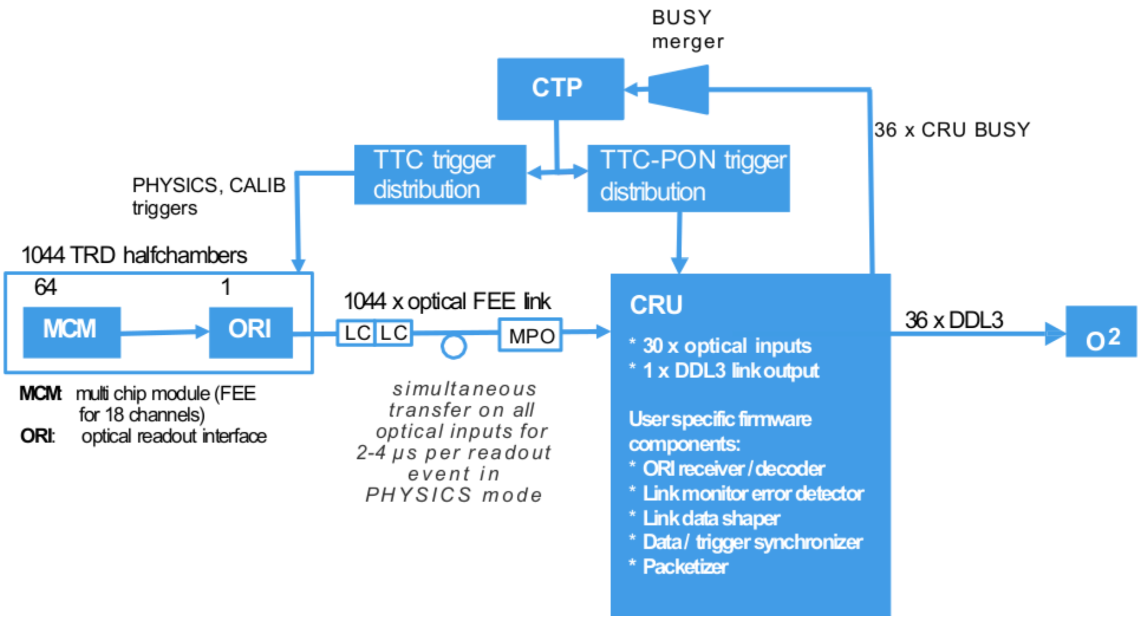

Figure 8.4: Block diagram of the TRD read-out unit.

all scenarios shown in Tab. 8.1. Together with the 5 times higher bandwidth of the DDL3, the 30:1 ratio of FEE input links and DAQ output link on one read-out unit is adequate.

The read-out must be able to handle different event types: (A) tracklet data and (B) full non zero-suppressed raw events (only in case of a calibration trigger). Table 8.2 shows event sizes and buffer requirements for typical cases of the various event types.

Event type	event size [kB/event/sector]	data volume [kb/evt/input link]	number of events in 512 kb link buffer
50 % tracklet words (A)	30.8	4	128
non zero-suppressed raw (B)	≈ 3000	400	1

Table 8.2: Event size and read-out input buffer capacity for various event types.

The input buffers on each FEE link act as multi-event buffer (MEB) for the read-out. The FEE data transfer to the read-out units uses a pure push mechanism without handshake or busy. To avoid data loss during this transfer, the read-out input buffer size and bandwidth have to be large enough to accept at least one event of maximum possible size simultaneously on all links. The largest possible event type, a non-zero suppressed raw event, requires a buffer size of 400 kb. This size allows the storage of a sufficient number of events for all event types in physics runs as shown in Tab. 8.2. With e.g. 512 kb link buffers, the single buffer can equally hold five central raw events or 64 events with maximum number of tracklets using a dynamic event buffer size.

For the interface to the O^2 system, there are no specific requirements from the TRD. Any chosen common ALICE DAQ link can be implemented for the TRD read-out with the choice of a suitable FPGA device. The goal for the TRD read-out is to use the ALICE common read-out unit (CRU). The scheme of 30 FEE optical links inputs and 1 DDL output maps to one AMC40 card.

The CRUs for the TRD will be located in the counting room CR1, allowing full accessibility during data taking. All 1044 fibers coming from the FEE and connected to the GTU in the C-area will be taken off from the GTU and extended to CR1. A preliminary test showed that the optical power of the laser diodes on the FEE is able to drive significantly longer link. However, this test didn't use the final setup, especially the patch panels which will be used to extend the fibers from the GTU racks as well as the multi-fiber push-on (MPO) connectors were not yet in place. An additional test with the final setup as in Fig. 8.4 will be performed.

8.4.2 Trigger and busy handling

For operating the TRD, trigger sequences need to be provided to the FEE and in parallel to the CRUs. The FEE mounted on the detector chambers will remain unchanged for the upgrade, employing a TTCrx device to receive and distribute trigger information to all FEE devices. Therefore, a TTC system for trigger distribution is needed for the TRD. The FEE requires a special trigger sequence on the TTC A-channel, which is not compatible with the standard ALICE TTC trigger sequence. It consists of individual pulses, one bunch crossing wide with a fixed timing for each provided trigger level. No TTC trigger messages are used on the FEE, thus no rate limit is imposed by using the TTC at high rates for the TRD.

For the TRD FEE, a single trigger level (LM) is sufficient to initiate the full processing and read-out sequence. For Run 2, a dedicated pretrigger system is in place, which sends the LM trigger to the FEE in case an L0 is expected to be issued by CTP. The employment of this system was necessary since the L0 from CTP is generated too late to sample the full signal shape including the early amplfication peak. The use of the pretrigger system will be discontinued in Run 3. Instead, the LM trigger will be delivered to the FEE via the TTC A-Channel by CTP directly. The timing of the LM signal has to be identical to the current TRD pretrigger with a maximum arrival time at the FEE of 900 ns after the collision. The latency of the TTC trigger distribution has to be short enough to fulfil this timing requirement.

The FEE supports up to two additional trigger levels: an L0 or L1 trigger, which, in case of tracklet read-out mode, is faster than 4 μs, aborts the FEE processing. The purpose of this function in Run 2 is the reduction of dead time in case of events which are rejected at the later L0 or L1 stage. For Run 3, this abortion must never occur and CTP must protect the FEE from receiving a new trigger pulse within the 4 μs timeframe after issueing an LM trigger.

CTP will also steer the alternation of tracklet-only and tracklet+raw data read out modes. For this purpose, an additional trigger type, the calibration trigger, will be introduced for Run 3. The calibration trigger has double the pulse width of the LM trigger and will initiate a full readout sequence (tracklet+raw data) on the FEE if it is encountered instead of an LM trigger. The FEE will be modified to support the reception of calibration triggers by a firmware update.

The CRUs will receive a full standard trigger sequence including trigger messages which are used for busy generation and event formatting. The fact whether a tracklet-only or tracklet+raw data readout is requested by CTP has to be encoded in the trigger message.

Contributing Institutes
University of Frankfurt, Germany
Gesellschaft für Schwerionenforschung, Darmstadt, Germany
University of Heidelberg, Germany
University of Münster, Germany
NIPNE Bucharest, Romania
Tokyo University, Japan

Table 8.3: TRD Institutes.

Schedule		
New data formats	2015-17	test runs with beam (pp, Pb-Pb); performance evaluation and optimisation of data content
	2016-18	modification of software (offline reconstruction, calibration, data quality monitoring)
Read-out unit	2015-17	firmware development and tests with prototypes
	2018	commissioning of all units

Table 8.4: TRD schedule.

The CRUs generate the TRD busy signal for the CTP. Each CRU asserts busy upon arrival of an LM trigger and releases the busy as soon as event end-markers are received on all FEE links or a time-out occurs. Moreover, busy is asserted in case of full buffers. Since 36 CRUs will process the incoming data from the FEE independently in parallel, each of them will generate an individual busy. All these busy contributions need to be ored without generating unnecessary latency in a dedicated busy merger unit. For this purpose, an additional CRU will be employed to fan in all individual busy signals.

8.5 Schedule, funding and institutes

Tables 8.3, 8.4 and 8.5 show the TRD institutes, schedule and funding.

		Funding
Read-out unit	37 CRUs, fiber connectors, crate, trigger & busy distribution	420 kCHF (CRUs in C-racks)
		additional cost for fiber routing to counting room
Manpower	firmware and software development	
	2015-16	1 FTE
	2017-19	2 FTE

Resources for funding will be requested from BMBF

Table 8.5: TRD funding.