CMS Note xxxx/96

A method for synchronization of the trigger data
DRAFT

J. Varela
CERN and LIP-Lisbon

30 September 1996

1. Introduction

The CMS trigger system can be viewed as a massive parallel processor that computes local trigger objects, followed by a tree like structure that selects the highest rank objects in the detector or performs global energy sums. The entire system works in synchronous pipeline mode, processing events at the 40 MHz LHC rate.

The system is based on the assumption that, at every processing stage, the data are synchronized and belong to the same bunch crossing. This goal has to be achieved without making use of bunch crossing identifiers attached to the data, which would imply a non acceptable complication of the trigger system design.

In this note we propose a method for the synchronization of the trigger data that builds on the ideas presented in reference [1]. The method takes profit of the gaps in the LHC bunch structure (sequence of clock cycles without particles). During the largest gap (127 clock periods) occurring every 89 μs, each transmitter sends synchronization data which is used by the receivers to adjust their phase. The start and the end of the synchronization gap is identified by commands send by the TTC System [2]. To implement the method, two specialized circuits, a Synchronization Transmitter (Sync TX) and a Synchronization Receiver (Sync RX), are needed in the trigger data links .

2. Characterization of the problem

The synchronization is a general issue for the trigger system and we shall be able to identify the points in the trigger data flow path where special synchronization circuitry is needed. The method we are proposing is sufficiently general and can be used widely in the trigger system. However, to facilitate the presentation, we will take here the example of the calorimeter trigger data synchronization at the input of the regional processors.

In figure 1 we show a timing diagram of the calorimeter trigger. We assume that, at the output of the front-end electronics, the data corresponding to the 36 crystals of one trigger tower are synchronous. With this assumption there is no need for a synchronization step at the input of the Trigger Primitive Generators (one per trigger tower).

However, due to different flight paths to different regions of the detector, the various TPGs are not necessarily synchronized with each other. In addition, the lengths of the fibers from the TPGs to the Regional Trigger Processors (RTP) could be different. As a result, data corresponding to the same crossing arrive at different times at the RTPs inputs .

The proposed synchronization method uses FIFOs at the end of the transmission lines to re-synchronize the data. The key points in the FIFOs operation are the following: i) after a clear, the first data entering every FIFOs correspond to the same crossing ; ii) the readout of the FIFOs is controlled by a common signal arriving simultaneously to every channel.

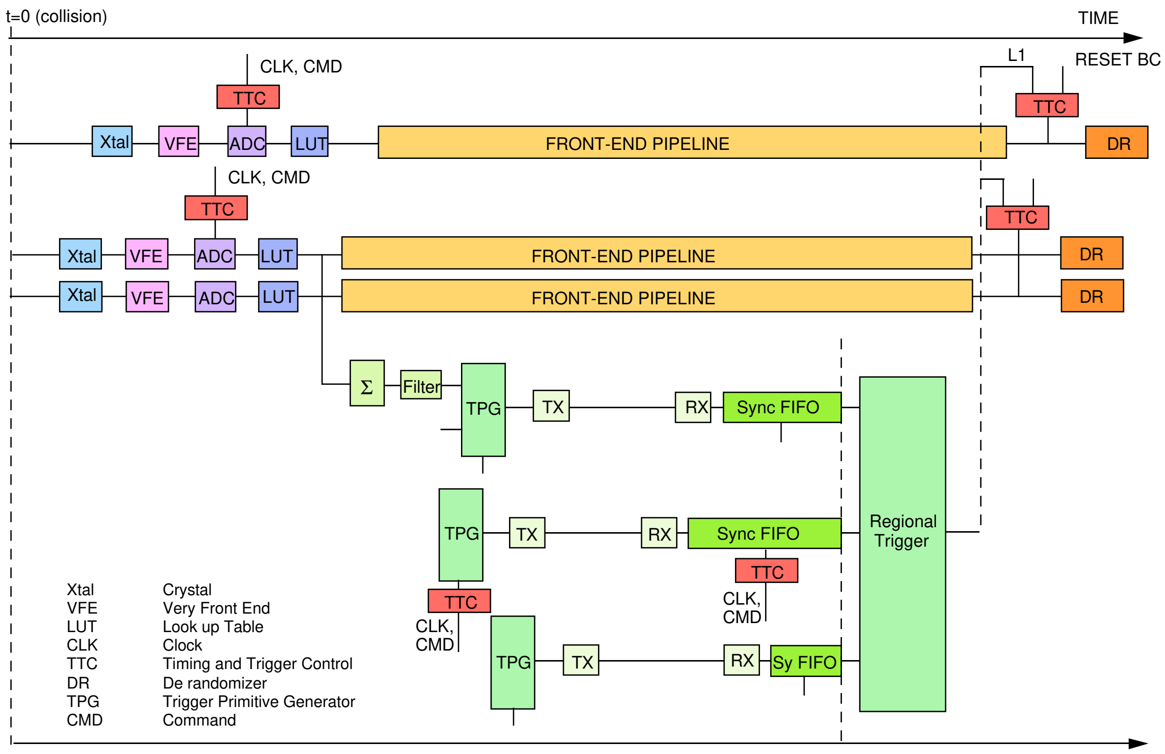

Figure 1 - Timing diagram of the calorimeter trigger

3. LHC bunch structure and TTC signals

The LHC bunch structure is represented in figure 2. One LHC orbit corresponds to 3564 periods of the 40.08 MHz clock. The structure has several gaps, that is sequences of missing bunches, of predefined length and location. We will concentrate in the LHC extraction gap (127 missing bunches), that for simplicity we will call the gap. By definition, the first bunch after the gap is called the 'bunch Ø'. In addition to the 40 MHz clock, the LHC machine can provide a 'bunch Ø clock' at a frequency of 11.25 kHz, the frequency of the LHC orbits.

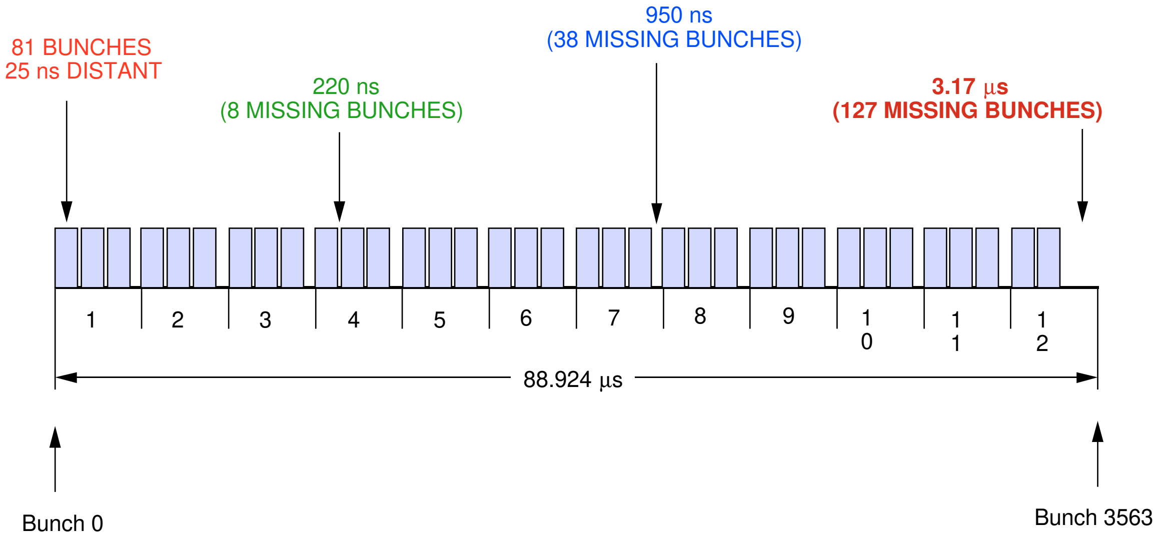

Figure 2 - The LHC bunch structure

The TTC system distributes to the front-end, trigger and readout electronics, through optical fibers, the Timing and Trigger control signals. The following signals are available from the TTC receiver circuit:

- LHC clock (CLK).
- Level 1 Trigger Accept (L1A).
- Bunch Counter Reset (BCR).
- Event Counter Reset (ECR).
- Synchronous Commands Signals: broadcast fast control signals, deskewed by a separate delay from that of L1A and BCR.
- Asynchronous commands and data: individually-addressed slow controls and data.

The L1A and BCR timings are adjusted to match the Level 1 Accept and the Bunch Counter phase to the arrival of the corresponding sub-detector data at the output of the front-end readout pipelines.

In order to be able to perform the synchronization operations described in the next sections, the TTC system is programmed to distribute to the trigger synchronization circuits, at both ends of the trigger links, the following Synchronous Commands:

- Bunch Crossing Ø (BCØ): signal synchronous with the first clock signal after the gap (bunch Ø). The frequency of the BCØ signal is 11.25 kHz.
- Start Of Gap (SOG): signal synchronous with the first clock signal in the gap. The frequency of the SOG and BCØ signals are the same.

In the transmitter side, the BCØ timing is adjusted to match the arrival of the bunch Ø data at the trigger synchronization circuits. The phase of the SOG signal is adjusted so that it arrives 127 clock cycles (or the actual gap length) before the BCØ signal. The BCØ and BCR signals have the same frequency but their phase differs by an amount equivalent to the trigger latency.

4. Trigger synchronization actions

In each trigger data link, the synchronization actions are performed by two circuits, the Sync Tx and the Sync Rx. At the front-end, the Sync Tx receives data from the trigger primitive generators, receives the TTC signals and controls the Link Transmitter. In particular, we assume that the Link accepts, together with the data frame, one DATA/CONTROL FLAG which distinguishes real data from control information (this facility is available in the HP G-Link, for example).

At the receiver side, the Sync Rx receives from the Link Rx, the data frame, the data/control flag and the transmitter clock (TX CLK). The Sync RX receives also the TTC signals, in particular the clock at the receiver side (RX CLK). Figure 3 shows the global scheme of a synchronization link.

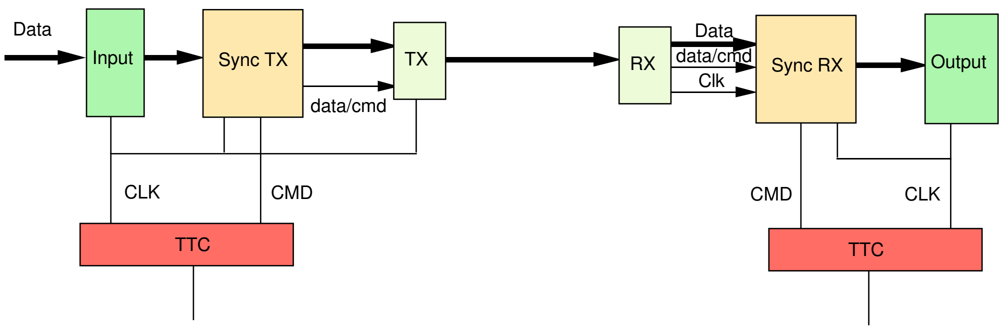

Figure 3 - Global scheme of a synchronization link

4.1. Actions at the Synchronization Transmitter

The main blocks of the Sync Tx circuit, shown in figure 4, are the following:

- Synchronization Data Generator (generates synchronization data to be sent during the gap);
- Multiplexer, which allows to switch between real data and synchronization data.
- TTC Command Decoder
- Accumulator, which allows to build the synchronization statistics (see §5).

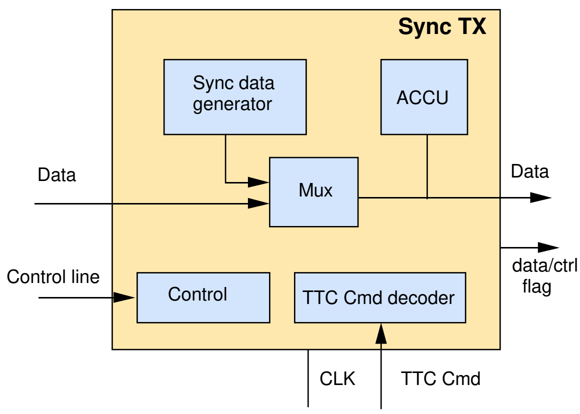

Figure 4 - Block diagram of the Synchronization Transmitter Circuit.

The TTCcmd BCØ, if properly adjusted in time, indicates that the data incoming in the next clock (or eventually in the second next clock) corresponds to the bunch Ø. When this command is decoded, the circuit switches the Multiplexer to real data and sets the data/control flag in data mode.

On the other hand, the TTCcmd SOG indicates that the next clock cycle (or eventually in the second next clock) corresponds to the first clock in the gap. When this command is decoded, the circuit starts the Synchronization Data Generator, switches the Multiplexer to synchronization data and sets the data/control flag in control mode.

The Synchronization Data Generator can be a simple counter. The counter data allows to maintain the activity in the link during the gap, and is used on the receiver end to trigger some of the Sync Rx synchronization actions. Among the 127 synchronization words, the receiver should be able to recognize three words, for example words 1, 64 and 127, as synchronization control information.

4.2. Actions at the Synchronization Receiver

The main blocks of the Sync Rx circuit, shown in figure 5, are the following:

- The Synchronization Data Decoder (SDD)
- The Synchronization FIFO
- The TTC Command Decoder
- The Accumulator (see §5)

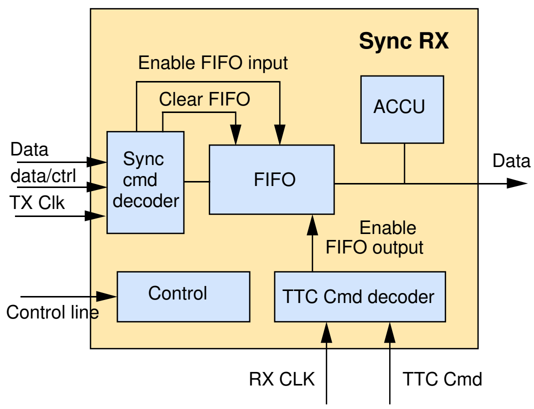

Figure 5 - Block diagram of the Synchronization Receiver Circuit.

When the link is in control mode, the SDD interprets the synchronization data words, recognizing the following commands:

- Synchronization word #1 - Disable FIFO input

  This command indicates that one LHC orbit is completed and in consequence the input of data in the FIFO should be inhibited. Note that this event happens at different times in different channels.

- Synchronization word #64: Clear FIFO

  This command indicates that the FIFO should be prepared to receive data from the next LHC orbit. The timing of this command shall be such that all the FIFO data from the previous orbit was already read. The suggested interval between the FIFO input disable and the FIFO reset (64 clock periods) is more than enough to fulfill that condition.
- Synchronization word #127: Enable FIFO input

  This command indicates that the next LHC orbit starts in the next clock and in consequence the FIFO input shall be enabled. Again, this event happens at different times in different channels.

When the link is in data mode, the input data is simply written in the Synchronization FIFO at the transmitter clock (TX CLK) frequency.

The clock given by the TTCrx ASIC at the reception (RX CLK) drives the FIFO readout. The TTCrx ASICs are housed in the Clock Control Cards (CCC) in the trigger crates (one CCC per crate). It is the responsibility of the CCC to distribute the clock (as well as other synchronous TTC commands) to all Sync Rx circuits in the crate with exactly the same phase. This guarantees that all the synchronization FIFOs in the crate are read at the same moment and that the trigger computations in that crate are in phase.

The read access to the Synchronization FIFO is controlled by the signals originated in the TTC Command Decoder. When the TTC cmd BCØ is identified, the FIFO output is enabled and the next data extracted from the FIFO corresponds to the bunch Ø. The timing of the BCØ command must be adjusted in such a way that all the Sync FIFOs in the trigger crate have already received data from the bunch Ø. On the other hand it can not be issued much latter than the latest data to not affect significantly the trigger latency. Again, it is the responsibility of the CCC to distribute the BCØ command to all Sync Rx circuits in the crate with exactly the same phase. In this way we guarantee that the trigger data in one trigger crate is synchronous and corresponds to the same bunch crossing.

Similarly, when the TTC cmd SOG is identified the circuit disables the FIFO output. During the synchronization gap the circuit will output zeros in the data bus.

5. Timing adjustments

In order to have the trigger data well synchronized, the timing of the TTC commands BCØ and SOG, at the transmitter and at the receiver sides, must be adjusted. Because the phase between BCØ and SOG is fixed (equal to the synchronization gap width), there are only two timing parameters to be determined, the deskewings of Tx BCØ and Rx BCØ.

The method used to compute these time adjustments is based on the ideas pioneered by the RD12 collaboration [2]. The distribution of the number of events as a function of the bunch crossing number (event time spectrum) is compared (correlated) with the expected LHC bunch structure to extract the time adjustment. The function of the Accumulator in the synchronization circuits is to build the event time spectra in-situ.

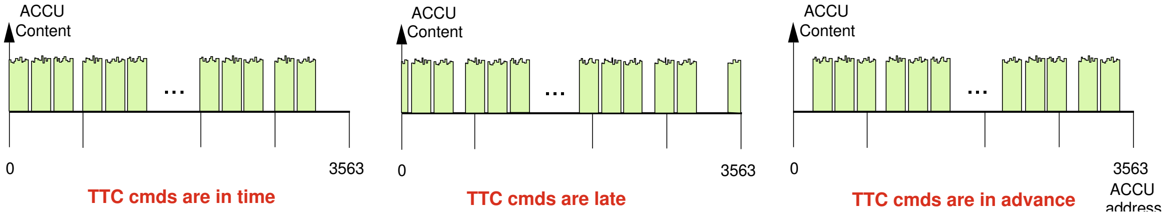

Figure 6 - Examples of event time spectra

5.1. Timing adjustment in the transmitter

The operation of the Accumulator (ACCU) in the Sync TX circuit is the following:

- When the command BCØ is received, the ACCU address is set to zero.
- At each clock cycle,

  the ACCU content (at the current address) is incremented if the input data is above the noise threshold (programmable),

  and the ACCU address is incremented.
- When the SOG is received, the circuit stops updating the Accumulator content.

After a certain number of LHC orbits the ACCU contains the event time spectrum with enough statistics to be compared (in the DAQ controllers) with the expected LHC structure. In figure 6 we show typical examples of the event time spectrum. The first spectrum reproduces exactly the LHC bunch structure, which means that the command BCØ arrives to the Sync TX circuit in phase with the data that effectively is produced by the collisions occurring at bunch Ø. The other two examples correspond to situations where the BCØ signal is either late or in advance relative to the data.

The time interval needed to accumulate the time spectra depends on the channel occupancy and on the statistics required. For example, a precise spectrum with about 1000 events per bin for a channel with occupancy 10^{-4} takes about 15 mn to accumulate.

Ideally, the DAQ controllers would read the ACCU content without disturbing the link operation. This facility allows to monitor the synchronization during the data taking.

5.2. Timing adjustments in the receiver

The operation of the Accumulator in the Sync RX circuit is similar. When the command BCØ is received, the ACCU address is set to zero, and at each clock cycle the ACCU content is updated and its address is incremented. The event time spectrum stored in the ACCU allows to compute the amount of deskewing needed on the Rx BCØ signal.

This time adjustment is less critical than in the transmission side. It is intended to verify that the BCØ doesn't arrive too soon, before the data from bunch Ø has arrived to the Rx FIFO, neither too late, in which case the FIFO would reach the full state or the clear FIFO signal would happen before the end of the complete FIFO readout.

6. Compatibility with calibration and monitoring activities

The trigger synchronization actions do not require the use of the entire 127 clock period gap. In fact, part of the gap will certainly be needed for monitoring or calibration triggers. We suggest to allocate for these purposes a certain number of clock periods just after the last bunch (bunch 3436). The TTCcmd SOG used by the trigger synchronization would then correspond to the first clock after the calibration/monitoring gap (see figure 7).

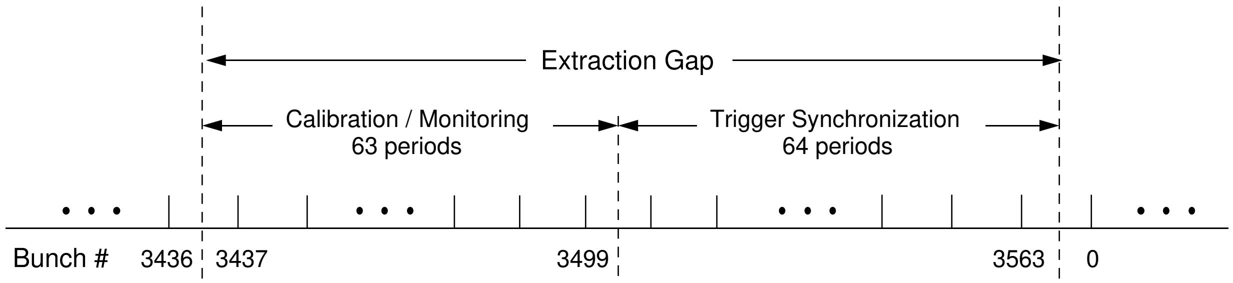

Figure 7 - Allocation of the gap region to trigger synchronization and monitoring/calibration activities.

7. Conclusions

We have presented the principles of the synchronization scheme of the trigger data. The method makes use of the LHC bunch gap and allows a resynchronization of the data every LHC orbit (88 μs). It relies on the distribution by the TTC of two signals (commands) synchronous with the bunch Ø and the start-of-gap events.

References

[1] M. A. Thompson, 'SDC Link Word Synchronization', SDC Note SDC-93-536, June 23, 1993.

[2] RD12 Status Report, CERN/LHCC 95-26, April 24, 1995.