

# ATLAS/TDAQ

---

## A trigger veto for the Bunch Counter Reset

Author: Thilo Pauly  
Status: Work in Progress

---

### Abstract

The Bunch Counter Reset (BCR) is a signal to synchronise bunch counters in the front- and back-end electronics of ATLAS. The value of the bunch counter, the bunch crossing identifier (BCID), is used to tag the data fragments in accordance with the bunch numbering scheme used to label bunches within the LHC bunch train. The BCID is also used as online and offline cross-check of the Level-1 Accept (L1A) synchronisation.

The TTC network transmits the L1A and the BCR to the sub-detector front-ends on two independent channels. The ID detectors (Pixels, SCT, TRT) use a different technology, where a single channel is shared for the transmission of the L1A and BCR. This has the effect that the two commands have to stay separate in time and no L1A can be sent during the transmission of the BCR. For the MDT, even though using the TTC network for transmission to the on-detector electronics, a restriction of the front-end TDC chip has the same effect.

In this note we describe in detail the restrictions we are facing and how they translate into a range of BCIDs, where no L1A triggers can be accepted. As this range covers BCIDs with proton collisions, we discuss how this forbidden range is moved into the abort gap of the LHC, where no collisions are expected. In order to guarantee stable data acquisition without loss of synchronisation, we introduce a protection mechanism in form of a trigger veto implemented in the Central Trigger Processor (CTP), which excludes triggers during the forbidden BCIDs.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	BCID and the LHC bunches . . . . .	3
1.2	BCID in ATLAS . . . . .	3
1.3	Global synchronisation of the BCID . . . . .	4
1.4	Purpose of BCID in ATLAS . . . . .	4
<b>2</b>	<b>Sub-detector restrictions</b>	<b>6</b>
2.1	General solution . . . . .	6
2.2	Pixels . . . . .	6
2.3	SCT . . . . .	7
2.4	TRT . . . . .	8
2.5	MDT . . . . .	9
2.6	Summary . . . . .	9
<b>3</b>	<b>Protection by a BCR trigger veto in the CTP</b>	<b>9</b>
<b>4</b>	<b>Conclusion</b>	<b>10</b>
<b>A</b>	<b>Tests for determining the range of forbidden BCIDs</b>	<b>11</b>
	<b>References</b>	<b>11</b>

# 1 Introduction

We start by revisiting the use and formation of the bunch-crossing identifier (BCID), first in the context of the LHC, then in the context of ATLAS.

## 1.1 BCID and the LHC bunches

The LHC radio-frequency (RF) system works at 400 MHz, dividing the LHC closed orbit by the harmonic number 35640 into 2.5 ns long RF buckets, potential wells that can hold particle bunches. In the nominal filling scheme with 25 ns bunch spacing, every 10th RF bucket is filled with particles, giving 3564 buckets where particles are allowed. These positions correspond to the different values of the BCID, a quantity which will be accurately defined later. Not all of the 3564 buckets are filled: 119 positions are kept free of particles for the extraction kicker to ramp up. This particle-free gap is also known as the **LHC abort gap**. Similarly, there are smaller gaps needed for the extraction and injection systems of the pre-accelerators. The exact bunch structure of the various filling schemes can be found in [1].

The **bunch labelling convention** for the LHC experiments is: all 25 ns nominal bunch positions (filled or not filled) are counted with the BCID starting from 0 for the last empty bunch of the LHC abort gap [2][3]. The first nominally filled bunch has therefore BCID=1. Since the harmonic number (with respect to 40 MHz) of the LHC is 3564, the range of BCID is 0 to 3563 (0x000 to 0xdeb).

In the LHC community, similar labels are used: The LHC beam instrumentation [4] uses BCID(BI) with the range 1 to 3564, with BCID(BI)=BCID(ATLAS) and the correspondence  $3564 \equiv 0$ . The LHC-RF convention [5] uses 2.5 ns bucket numbers reminiscent of the 400 MHz RF frequency, with BCID(RF) spanning from 1 to 35640, where BCID(RF)=1 coincides with BCID(ATLAS)=1.

It is worth noticing that when the LHC is operating in single-bunch mode, the single bunch will be called BCID=1. As we will see, the Pixel, SCT and MDT would not be able to receive triggers on BCID=1 if not special care is taken.

## 1.2 BCID in ATLAS

The BCID for ATLAS is a unsigned 12-bit integer with values spanning the range 0 to 3563. It is formed in every sub-detector read-out driver (ROD) and front-end with the help of local bunch counters that increment with the 40 MHz bunch clock [6]. The ROD-BCID is 12-bit long and enters as word in the ROD fragment header as specified in the data event format [7]. In the front-end, the BCID is formed as well, often a version of less than 12 bits, for example 8 bits for the Pixels and SCT, 4 bits for the TRT. This FE-BCID is sent as part of the data payload to the RODs. The RODs are comparing

the FE-BCID of the arriving data with the ROD-BCID of their corresponding header.

### 1.3 Global synchronisation of the BCID

In order to synchronise the BCID across all sub-detectors, a bunch counter reset (BCR) signal is used. This signal is derived from the 11 kHz revolution pulse (ORBIT signal) which is delivered by the LHC. The ORBIT signal is distributed to the sub-systems by the Central Trigger Processor (CTP) [8] and transformed into a BCR command in the sub-detectors' TTCvi [9]. The phase of the BCR with respect to the ORBIT signal can be adjusted in the TTCvi and downstream. The ORBIT signal is distributed in a tree-like network, with the following layers of nodes, starting from the top: RF, RF2TTC, CTP, LTP, TTCvi, sub-system specific layers, down to the leaves, which are RODs and front-ends. In most of the layers one has the possibility of delaying the ORBIT or BCR. Coarse delays in shallow layers typically act as global delays for the underlying sub-tree, whereas offsets in deeper layers are used for fine-tuning. The synchronisation procedure is done from the bottom up, that is, first the leaves of a node are synchronised, then the higher layers.

As the CTP (a single system) receives the ORBIT signal from the LHC and distributes it to all subdetectors, it can act as the BCID reference for ATLAS. The ROD of the CTP fixes the delay between ORBIT and BCID=0 to a value which is convenient for all sub-detectors to reach the same BCID as the CTP-ROD header. A value of 200 BC is currently chosen.<sup>1</sup>

After all sub-detectors have the same BCID as the CTP, one can align the BCID with respect to the LHC by adjusting the phase of the ORBIT signal in a single place: before the CTP, in the RF2TTC module. This way, the adjustment acts as a global offset for all sub-systems.

For the discussion on sub-detector restrictions, it is important to understand that for each sub-system, the time when the BCR is sent with respect to the L1A can be fully specified in terms of BCID.

### 1.4 Purpose of BCID in ATLAS

ATLAS uses the BCID in at least three different ways:

---

<sup>1</sup>The value of 200 BC was determined like this: In the TTCvi, one has to use an inhibit width of at least 51 BC to ensure synchronous transmission of the BCR. Allowing for a few more BCs until the BCR is sent out of the B-channel and allowing for differences in the sub-detector specific electronics, one arrives at around 70 BC. Taking into account the fact that the TRT sub-detector sends their BCRs roughly one Level-1 latency (currently 130 BC) earlier to the front-end than other sub-detector, one arrives at 200 BC. For illustration, the BCR delay parameter in the TTCvi (delay of the inhibit signal for B-Go channel 0) is currently set to: 113 for Pixels, 110 for SCT, 27 for TRT, 114 for MDT.

1. It labels the bunch-crossings within the LHC bunch train and thus is important in discussions with the LHC community and the other experiments, when individual bunches or bunch crossings need to be addressed.
2. It serves as a relative time-stamp of 25 ns precision relative to the LHC turn. This proves useful to independently **cross-check the L1ID synchronisation** of the dataflow and event building, in other words, that the L1A and the event counter reset (ECR) are properly transmitted.

The L1ID is a 32-bit identifier whose least significant 24-bits count the L1As (event counter, EC) and whose most-significant 8 bits count ECRs. The ECR resets the EC to 0xffffffff, so that the EC value of the first L1A after an ECR is 0x000000. The L1ID synchronisation is checked with the BCID at various stages: at the ROD, read-out system (ROS), and event builder (SFI). The principle is simple: at any of these stages, the BCID of the received data is compared to a reference BCID. If all the L1IDs are correctly formed and the BCR delays correctly timed in, the BCIDs will match. In case a L1A is missing (or too many) for certain parts of the system, the fragment of the next (or previous) event gets attributed with this L1ID. Since this other fragment comes from a different bunch crossing, it will have in general a different BCID, leading to a BCID mismatch.

The ROD compares its ROD-BCID with the BCID from the incoming data. In case of an inconsistency – a BCID mismatch – sub-detector specific actions are taken. Typical actions are: flagging the data as wrong, discarding the data payload for the fragment.

The ROS compares the BCID from the incoming readout links with respect to each other. On the SFIs, two independent checks are done: the event builder application can (enabled by default) check for all events for BCID mismatches between the RODs [10]. A separate application samples the data and does detailed ROD header BCID checks on a fraction of the events (in the Monalisa framework [11]).

3. It provides **a check of the sub-detector timing of the L1A with the bunch structure**: By triggering randomly in a window around BCID=1 one can check the timing of sub-detectors by looking at a suitable 'detector activity' observable per BCID. In case the detector is correctly timed in, a sharp edge should be visible with close to zero rate for BCID=0 and high rate at BCID=1. In order for this to be useful, a reasonably large window around BCID=1 should be available;  $\pm 6$  BC seems an appropriate choice, giving BCID=3559, 3560, 3561, 3562, 3563, 0, 1, 2, 3, 4, 5, 6, 7. Note, that even though any other edge in the bunch structure of the LHC could be used, the edge at BCID=0 is unique: it is present in all LHC filling schemes and the LHC is actively cleaning the abort gap of particles, which is not done for other gaps in the bunch structure.

## 2 Sub-detector restrictions

In this section, we describe restrictions from the ID detectors and the MDT, which do not allow L1A triggers for certain BCIDs, a forbidden range of BCIDs which is specific to their sub-system. As we will see, these ranges cover crossings of filled bunches, as well as a few empty bunch crossings of the LHC abort gap. Clearly, ATLAS wants to be able to trigger on all BCIDs that contain filled bunches, in particular BCID=1, which will host the collisions in single-bunch mode and is a reference bunch crossing for many LHC communities. In addition, the particle-free end of the abort gap is very useful for checking the sub-detector timing. For these reasons we propose to move the forbidden ranges of BCIDs sufficiently deep into the LHC abort gap so that L1A triggers can be accepted on  $\text{BCID} \geq 3559$  and  $\text{BCID}=0,1,2,3,4,\dots$ , and the full range of filled bunch crossings in any LHC filling scheme.

### 2.1 General solution

The forbidden BCID range can be moved deeper into the LHC abort gap by moving the BCR that is sent to the front-end. In principle, this could be achieved in two ways:

1. Move the BCR only for the sub-detectors in question, by changing the BCR-delay in the TTCvi and compensating for this shift in the ROD. This way, the BCID of the ROD header would stay correct and the ROD could perform BCID checking. This move would only introduce a re-definition of the FE-BCID, which is part of the data payload only. Alternatively, the BCR could be compensated for in the front-end, provided the front-end has this capability.
2. Move globally the BCR into the abort gap for all sub-detectors by changing the phase of the ORBIT signal with respect to the LHC, for example in the RF2TTC or in all TTCvis. This would lead to a global re-definition of the meaning of the BCID as bunch crossing label, inconsistent with the LHC machine and the other experiments. As we will see later, the TRT would have to delay its BCR, whereas the Pixels, SCT and MDT would have to advance it. Due to this opposite behaviour, one cannot solve the problem by a global shift of the BCR alone, but only in combination with a local shift as described before.

Based on discussions in previous ATLAS Technical Management Board meetings, the Executive Board decided to implement the first solution [12], which we will continue to discuss.

### 2.2 Pixels

The Pixel sub-detector uses the TTC network to provide the timing signals to the TIM modules. The TIM modules serve the RODs, which in turn provide the signals to the

back-of-crate cards (BOC), which transmit the signals optically to the front-ends. The BOCs encode the L1A into a 5-bit word, the BCR into a 9-bit word, and transmit them to the front-end over a single channel. Errors are created when the L1A and the BCR are issued too close in time. This translates into a range of 13 forbidden BCIDs[13], which was measured to be

3559, 3560, 3561, 3562, 3563, 0, 1, 2, 3, 4, 5, 6, 7.

It was recommended to add 1 BC (BCID=8) as safety margin. If L1A triggers are issued in this range, the following errors happen:

- BCID=3559..3562: The BCR will not be sent to the front-end. In this case the bunch counter will not be reset, which leads to BCID errors, i.e. correct data fragments with wrong BCIDs.
- BCID=3563: This error has not yet been sufficiently analysed
- BCID=0..7: The L1A will not be sent to the FEs and therefore no data fragments are sent to the ROD, which will time out while waiting for them to arrive. For a time-out value of 3.2 ms, for example, the dataflow frequency will drop down to about 300 Hz. At high L1A rates, the RODs will become busy and need to be re-configured.

This forbidden BCID range is now moved deeper into the abort gap by reducing the BCR-delay in the TTCvi by 14 BC, which effectively makes the BCR being sent to the front-end earlier. In the TIM, a BCID offset of  $-14$  BC gives the correct BCID for the ROD header and is used as a correction for the BCID checking in the RODs. After the correction, the forbidden BCID range reads:

3545, 3546, 3547, 3548, 3549, 3550, 3551, 3552, 3553, 3554, 3555, 3556, 3557, (3558).

## 2.3 SCT

The SCT uses a TTC partition to provide the timing signals to their TIM modules. As for the Pixels, the signals are distributed to the RODs and BOCs. In the BOCs, they are encoded and transmitted optically to the front-end. The L1A signal is encoded into a 3-bit word, the BCR into a 7-bit word. The forbidden region of BCIDs covers 13 BCs and is expected to be [14]

3558, 3559, 3560, 3561, 3562, 3563, 0, 1, 2, 3, 4, 5, 6

After introducing an offset of 16 BC in the TIM, TTCvi and ROD (in the same way as for the Pixels), the forbidden range translates into

3542, 3543, 3544, 3545, 3546, 3547, 3548, 3549, 3550, 3551, 3552, 3553, 3554.

## Open Issues for SCT

- Confirm this forbidden range experimentally
- Why is the range so large (13 BC)? Naively, one would expect  $L1A+BCR-1 = 9$ , or  $L1A+BCR+1 = 11$  if an additional BC is needed for a L1A.
- What are the error scenarios when triggering during the forbidden BCID range?

## 2.4 TRT

The TRT use the TTC to forward the timing signals to the TRT-TTC modules, which serve the RODs and the front-ends. While the ROD uses the full 12-bit BCID, the FE only uses the lower 4 bits. In the TRT-TTC, the L1A is encoded as a 3-bit command, the BCR as a 7-bit command, and transmitted to the front-ends and the RODs. The transmission shares a single channel, for which reason no L1As can be sent in a range of 11 BC around the BCR<sup>2</sup>.

In contrast to other sub-detectors, the TRT forms its FE-BCID at the beginning of the Level-1 pipeline, which is  $D$  bunch crossings deep. Thus, the BCR must be sent  $D$  bunch crossings earlier than the L1A for BCID=0, which is applied at the end of the Level-1 pipeline. It turns out that the current pipeline depth  $D = 130$  BC is longer than the LHC abort gap (119 BC), which means that the BCR should be sent a few bunch crossings before the abort gap begins, during time when collisions take place.

One can move the BCR into the abort gap by sending it to the front-end later. This can be done by increasing the inhibit delay of the BGo-0 channel in the TTCvi by the right amount, with the effect that the ROD- and FE-BCID become wrong. For the ROD, this can be compensated with a programmable delay in the TRT-TTC. The FE-BCID will stay wrong, but it can be taken into account during BCID checking in the ROD. Since the FE-BCID is only a 4-bit number, it is easiest to delay the BCR by multiples of 16 BC.

The current forbidden BCID range (for  $D = 130$  and no corrections applied) is unknown, but the corrected range was experimentally found to be [15]:

3542, 3543, 3544, 3545, 3546, 3547, 3548, 3549, 3550, 3551, 3552.

Currently, BCID checking in the ROD is switched off, but it is planned to switch it back on once the correction algorithm has been implemented in the ROD firmware.

## Open Issues for TRT

- For a given pipeline depth  $D$ , what is the 'raw' forbidden BCID range?

---

<sup>2</sup>Naively, one expects  $L1A+BCR-1 = 7 + 3 - 1 = 9$  BC. In case one additional BC is needed to read out the event, one would expect 11 BC.



- What are the error scenarios when triggering during the forbidden BCID range?
- Implementation of BCID checking in the ROD firmware
- The test of checking the forbidden BCID range should be repeated.

## 2.5 MDT

The MDT use the TTC system to feed the timing signals to the TIM modules, which serve the RODs, and to the on-detector service chamber modules (CSM), which use TTCrx chips to receive the signals. In the CSM, the L1A and the BCR are distributed to the TDC ASICs (AMT, ATLAS Muon TDC). Due to a limitation in the AMT, no triggers can be sent for FE-BCID=0, 1, 2.

In order to move the forbidden region deeper into the abort gap, the BCR is sent 12 BC earlier to the CSM, by subtracting 12 BC from the TTCvi BCR delay. The forbidden range then becomes:

$$3552, 3553, 3554.$$

This shift is compensated for the RODs by a BCID offset value of  $-12$  in the TIM. The same compensation is applied in the AMTs ( $-12$  for the parameter 'bunch-counter-offset-after-BCR').

In case triggers are sent during the forbidden region, the CSM will ignore them. For these events no data will arrive at the RODs, which eventually disables the corresponding inputs [16].

## 2.6 Summary

Fig. 1 summarises the range of forbidden BCIDs before and after the correction — the shift into the LHC abort gap.

## 3 Protection by a BCR trigger veto in the CTP

The CTP has the possibility to generate trigger conditions for 8 programmable sets of BCIDs through the so-called bunch group mechanism. These conditions can be combined with other trigger conditions to form trigger items. We propose that one bunch group (BGRP0) be reserved as protection for the BCR. This bunch group would contain all BCIDs but the forbidden ones and is linked with a logical AND to all trigger items. This way it is ensured that no L1A triggers can be issued for forbidden BCIDs. The list of

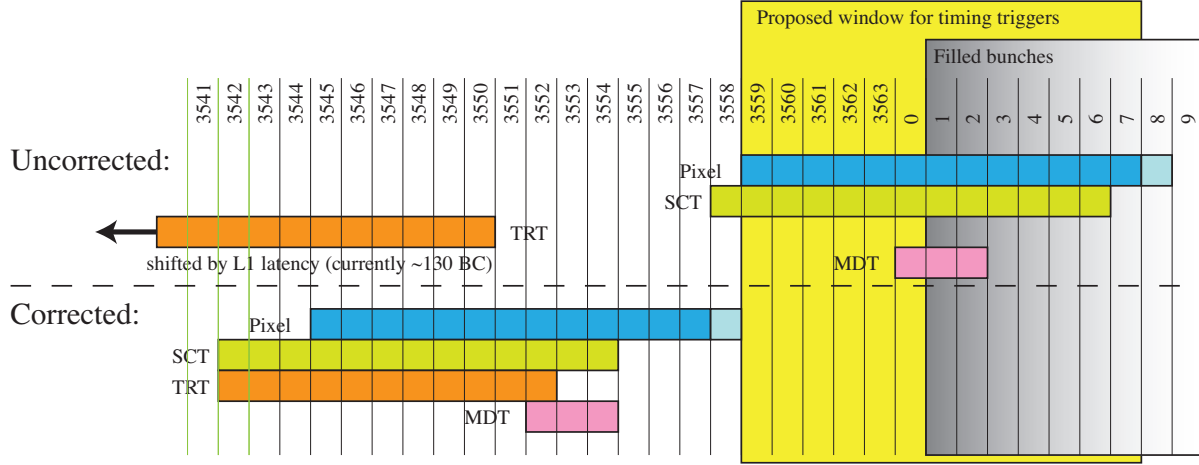


Figure 1: Range of forbidden BCIDs before and after the correction.

BCIDs is maintained in the trigger configuration database, which is populated with the trigger tool [17].

For this mechanism to work properly, it requires:

- the BCID to be aligned within the CTP and between CTP and sub-detector RODs.
- every trigger item to be combined with bunch group BGRP0 through a logic AND. The trigger tool will check that this requirement is being met while uploading a new trigger menu [18].

It is foreseen, at least temporarily, to monitor that the CTP respects the forbidden BCID range by monitoring the L1A per BCID in the CTPMON.

## 4 Conclusion

We have identified problematic issues for the Pixels, SCT, TRT and MDT, concerning the sending of the BCR and the impact on triggering on crossings of filled bunches. A solution is proposed that moves the BCR deeper into the LHC abort gap while keeping the BCID aligned with the rest of ATLAS at the ROD level. This allows that L1A triggers can be accepted for all relevant BCIDs: all potentially filled bunch crossings as well as a window of  $\pm 6$  BC around BCID=1. The CTP bunch group mechanism is used to protect these sub-detectors from receiving L1As during the time they send the BCR.

# Acknowledgements

We would like to thank our colleagues from the Pixel, SCT, TRT and MDT sub-detectors for developing and implementing the necessary corrections for the BCR: Iskander Ibragimov, Dave Robinson, Matt Warren, Bruce Gallop, Philippe Farthouat, Mike Hance, Dominick Olivito, Peter Lichard, Peter Wagner, Enrico Pasqualucci, Jay Chapman, Henk Boterenbrood, and Tiesheng Dai.

## A Tests for determining the range of forbidden BCIDs

The range of forbidden BCIDs for a particular sub-detector can be determined in a stand-alone setup: The pattern generator of the master LTP can be programmed to trigger only on one BCID, even though its value is not exactly known. By looking at the sub-detector ROD header, the BCID value corresponding to a given position in the pattern generator memory can be determined. This knowledge can be used to trigger on specific BCID values and check for proper data reception and synchronisation.

Using the CTP, the BCID range can be checked for more than one sub-detector at the same time: one can use a random trigger of 175 kHz in combination with a bunch group. If only one BCID is enabled, a rate of 50 Hz is obtained.

## References

- [1] Paul Collier, “*Standard Filling Schemes for Various LHC Operation Modes*,”  
<https://edms.cern.ch/document/489439>
- [2] “*Level-1 Trigger TDR: BCID synchronization*,” section 16.5.2.2 (page 441)  
<http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/TDR/tdr.html>
- [3] Wesley Smith (CMS), Richard Jacobson (LHCb), David Evans (ALICE), private communication.
- [4] Rhodri Jones, private communication.
- [5] Philippe Baudreghien, private communication.
- [6] ATLAS Trigger-DAQ Steering Group, “*Trigger & DAQ Interfaces with Front-End Systems: Requirement Document*,”  
<https://edms.cern.ch/document/384600/2>
- [7] A. dos Anjos, H.P. Beck, B. Gorini, “*The raw event format in the ATLAS Trigger & DAQ*,”  
<https://edms.cern.ch/document/445840>

- [8] S. Ask, et al., “*The ATLAS central level-1 trigger logic and TTC system*,” J. Instrum. 3 (2008) P08002.
- [9] Philippe Farthouat, Per Gällnö, “*TTC VME bus Interface TTCVI*,” <https://edms.cern.ch/document/110746/2>
- [10] Wainer Vandelli, private communication.
- [11] Roberto Ferrari, private communication.
- [12] Minutes of the ATLAS Executive Board meeting on Friday 23 May 2008.
- [13] Iskander Ibragimov, private communication.
- [14] Dave Robinson, Matt Warren, Bruce Gallop, private communication.
- [15] Mike Hance, Dominick Olivito, private communication.
- [16] Henk Boterenbrood, private communication.
- [17] A. dos Anjos, et al., “*The Configuration System of the ATLAS Trigger*,” ATL-COM-DAQ-2007-027, CERN, Geneva, 2007.
- [18] Jörg Stelzer, Takanori Kono, private communication.