



**ATLAS Note**  
ATLAS-INT-2019-XXX  
1st May 2020



Draft version 2.0

1

2

## **New Small Wheel Trigger Processor**

3

4

5

## **Firmware Requirements & Implementation**

6

NSW Trigger Processor Group

7 © 2020 CERN for the benefit of the ATLAS Collaboration.

8 Reproduction of this article or parts of it is allowed as specified in the CC-BY-4.0 license.

## 9 Contents

10	<b>1 Introduction</b>	<b>6</b>
11	1.1 Overview of the NSW trigger	6
12	1.2 System granularity and terminology	6
13	1.3 Use Cases	7
14	1.4 Requirements and Limitations of the trigger	8
15	1.5 Description of the Trigger Processor hardware	9
16	<b>2 Algorithms</b>	<b>12</b>
17	2.1 sTGC algorithm	12
18	2.1.1 sTGC input capture	14
19	2.1.2 Distributor	16
20	2.1.3 FPGA resources	16
21	2.2 Micromegas algorithm	17
22	2.2.1 Description	17
23	2.2.2 MM input format	19
24	2.2.3 ART Data Protocol	19
25	2.2.4 Input data alignment and clock domain crossing	20
26	2.2.5 FPGA Resources	21
27	2.2.6 Simulation	21
28	2.2.7 Hardware Testing	21
29	2.2.8 Micromegas trigger alignment constants	22
30	2.3 Clocking	23
31	<b>3 Output to the Sector Logic</b>	<b>27</b>
32	3.1 Transfer of track candidate segments from MM FPGA to sTGC	28
33	3.2 Removal of “duplicate” segments	28
34	3.3 Fixed latency output to the Sector Logic	30
35	<b>4 Configuration, control and status of algorithm and operating parameters</b>	<b>31</b>
36	<b>5 Latency estimate</b>	<b>33</b>
37	<b>6 Monitoring</b>	<b>34</b>
38	6.1 Monitoring Data	36
39	6.2 Correlated MM– sTGC monitoring	36
40	<b>7 TTC and Level-1 Accept</b>	<b>37</b>
41	7.1 TTC data	38
42	7.2 Level-1 Accept Algorithm	38
43	7.3 Level-1 Accept Implementation	38
44	7.4 Read out data format	40
45	7.4.1 Micromegas read out data	40
46	7.4.2 sTGC read out data	42
47	7.4.3 Read out bandwidths	43
48	<b>8 Configuration, control and monitoring of the hardware via I2C</b>	<b>43</b>

49	8.1 IPMI monitoring	44
50	<b>9 FPGA configuration</b>	<b>46</b>
51	<b>10 Playback of test data</b>	<b>47</b>
52	<b>11 Firmware organization and continuous integration</b>	<b>48</b>

53 **Revision History**

<b>Revision</b>	<b>Date</b>	<b>Author(s)</b>	<b>Description</b>
0.0	10-08-2017	L. Levinson	Created
0.1	02-10-2017	L. Levinson	Released to reviewers
0.2	12-06-2018	L. Levinson	Base for June 2018 IDR: corrected Pad Trigger data, Table 5.
0.3	20-04-2019	L. Levinson	Base for April 2019 review. SCX protocol is now defined so that no changes are needed to any OPC UA software. The firmware was moved from an SVN to a Git repository. sTGC input changed from $3 \times 40$ bits to $4 \times 30$ bits, Section 2.1.1. Changed format of the Pad Trigger data sent, Table 5. The bitmap ART output format is used; references to list format were removed from Section 2.2.2. Extensively updated L1A readout, Section 7. Added details on the transfer of segments from MM to sTGC, the removal of “duplicate” segments, and fixed latency, to Section 5. Extensively updated Monitoring, Section 6. Correlated monitoring by exchanging data replaced by just signaling to send monitoring data for the current BC via FELIX. Latency table updated. Table 12. FPGA configuration, Section 9 updated. Added Playback, Section 10. Removed Manpower and Phase-2 compatibility sections.
1.0	27-04-2019	L. Levinson	Released to reviewers
1.1	30-04-2019	L. Levinson	Updates to Figures 6, 22. Added link to Excel file listing status and configuration registers. Updated MM Monitoring section. Formatting, English and spelling corrections.
1.2	22-05-2019	L. Levinson	Corrected Table 12, Latency. Still not final.
1.3	10-10-2019	L. Levinson	Added description of BC clock recovery to Section 3.3. Added use of the swROD’s round-robin E-links feature to attain needed L1A output bandwidth in Section 7. Replace table for MM readout data format in Section 7.4. Updated Monitoring, Section 6.

- 2.0      30-04-2020    L. Levinson, *et al.*    For May 2020 PRR:
- Added a brief description of the carrier card to the Introduction, Section 1, including reference.
  - Updated diagram of carrier card (with latest information from Samway), Figure 5.
  - Added Section 2.2.4, Input data alignment and clock domain crossing.
  - Correct position of Monitor bit in the output to the Sector Logic to be high bit (Table 9).
  - Added comment on bands that span two TDS ASIC's in Section 2.1.
  - Added overview of clocking as Section 2.3.
  - Revise Section 3.3 on fixed latency of the output to the Sector Logic.
  - Added description of the high-level user interface that exports all the needed formats needed to define the configuration and status registers, Section 4.
  - Replaced L1A data format diagram, Figure 27.
  - Added L1A read out bandwidth, Section 7.4.3.
  - Updated Micromegas latency in Section 5.
  - Add text for IPMI monitoring for the new carrier, Section ??.
  - Improved text for FPGA configuration and IPMI monitoring, Section ??.
  - A graphical tool is no longer used to create the top level VHDL of the sTGC design (Section 11).
  - Newly instituted Continuous Integration and Continuous Deployment are described in Section 11.

## 1 Introduction

This document presents the requirements for firmware and its implementation for the NSW Trigger Processors. Information about the whole New Small Wheel detector system can be found in the New Small Wheel Technical Design Report [1]. Documentation of the Trigger Processor can be found in the Preliminary Design Review documents [2–4] and the NSW Trigger Processor documents available in the NSW Electronics SharePoint <http://www.cern.ch/ATLAS-NSW-ELX>. The documentation includes a description of the interfaces to the two NSW trigger detectors (Micromegas, MM, and small Thin Gap Chambers, sTGC), of the interface to the muon end-cap Sector Logic and of the ancillary functionality of the Trigger Processor. This document presents the requirements and the implementation of the firmware for the Trigger Processor for both detector technologies (sTGC and MM), including the trigger algorithms, monitoring and other ancillary functions.

### 1.1 Overview of the NSW trigger

The main goal of the NSW trigger is to provide additional information to the muon Level-1 trigger (L1) in the endcap region  $1.0 < |\eta| < 2.4$ , in order to dramatically reduce fake triggers arising from particles that are not high- $p_T$  muons originating in the interaction point (IP). A major source of fake triggers is low energy particles, mainly protons, generated in the material located between the Small Wheel and the end-cap middle station, the Big Wheel (BW). As shown in Figure 1, these particles can cross the end-cap trigger chambers at an angle similar to that of real high  $p_T$  muons. The NSW trigger signal is based on track segments produced online by the small Thin Gap Chambers (sTGC) and Micromegas chambers (MM) comprising the NSW detectors. These candidate track segments are sent to the new Sector Logic that uses them to corroborate trigger candidates from the BW TGC chambers. The Sector Logic sends Level-1 trigger candidates to the ATLAS Muon Central Trigger system.

### 1.2 System granularity and terminology

The endcap trigger system operates independently on each detector endcap (side A and side C). Several other factors define the granularity of the endcap trigger system. The following recalls the terminology and boundary conditions of the overall endcap trigger system:

- A **New Small Wheel sector** comprises  $1/16^{\text{th}}$  sector of a wheel (i.e. one endcap). Each wheel has eight large and eight small sectors. The detector and trigger sectors are coincidental.
- The **Big Wheel trigger detector (BW-TGC)** has 12 detector sectors per wheel. Each sector is further divided into trigger sectors, four trigger sections per detector sector in the region at larger radius (endcap region) and two trigger sectors in the region at smaller radius (forward region). This segmentation results in a total of 48 trigger sectors in the endcap (larger radius) region and 24 trigger sectors in the forward region.
- A **Sector Logic board** serves two adjacent BW-TGC trigger sectors. For each side (A/C), there are 24 Sector Logic boards for the BW-TGC endcap region and 12 Sector Logic boards for the BW-TGC forward region.
- A Sector Logic board receives data from up to six NSW sectors.

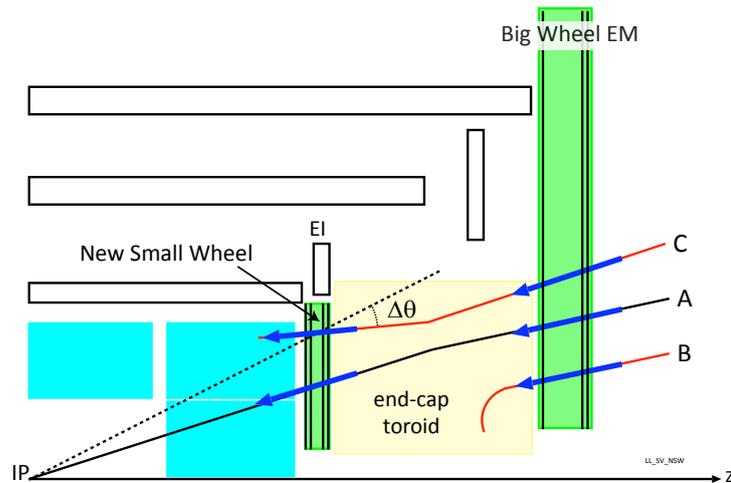


Figure 1: Schematic of the muon endcap trigger. The existing Big Wheel trigger accepts all three tracks shown. With the addition of the NSW to the muon end-cap trigger, only track 'A', the desired track, which is confirmed by both the Big Wheel and the NSW, will be accepted. Track 'B' will be rejected because the NSW does not find a track coming from the interaction that matches the Big Wheel candidate. Track 'C' will be rejected because the NSW track does not point to the interaction point (IP). The NSW logic restricts  $\Delta\theta$  to a value consistent with the track to have originated from the IP.

- 91 • A **NSW Trigger Processor ATCA carrier board** (corresponding to two NSW sectors) holds two  
92 mezzanine cards to serve a NSW octant.
- 93 • In the NSW Trigger Processor, there is one **FPGA** per detector technology (MM or sTGC) per NSW  
94 sector. (MM) and sTGC information is processed by separate algorithms in separate FPGA's. The  
95 two FPGA's are located on the same mezzanine card. Low latency LVDS lines allow communication  
96 between the MM and sTGC FPGA's. The track segments found by each algorithm must be merged,  
97 i.e. duplicates removed and sent as a single stream to the Sector Logic. The MM algorithm finishes  
98 first, so it sends its results to the sTGC FPGA for merging.
- 99 • One NSW trigger sector needs to deliver data to up to seven Sector Logic boards. The maximum  
100 fan-out of seven is needed for large NSW sectors, due to the overlap with the Big Wheel trigger  
101 sectors when multiple scattering, misalignments and magnetic field dispersion are taken into account.

### 102 1.3 Use Cases

103 The Use Cases for the Trigger Processor are:

- 104 1. Search for MM and sTGC segments on every bunch crossing.
- 105 2. Merge segments into one list for the New Sector Logic.
- 106 3. On Level-1 Accept, send Trigger Processor inputs and outputs for the relevant Bunch Crossings to  
107 the ROD via FELIX.
- 108 4. Download simulated input data to the FPGA and perform the trigger algorithm.

- 109 5. Configure the algorithm and readout parameters, and the FPGA's themselves, from the DAQ  
110 configuration database.
- 111 6. Accumulate and send performance statistics to the monitoring system, e.g. segments generated per  
112 BC, segments resulting in L1A,  $R$ - $\phi$ ,  $\theta$  distribution of segments, number of valid planes per BC. . .
- 113 7. Monitor the efficiency and purity of the Trigger Processor's segment finding and matching with the  
114 Big Wheel by forcing a Level-1 Accept at a low rate (e.g. a few 10's of Hz) for bunch crossings in  
115 which a segment has been found. This is done by setting an "NSW monitor" bit in the packet sent to  
116 the Sector Logic on every bunch crossing.
- 117 8. Send sector data from any bunch crossing that has a "potential" segment to a monitoring point. This  
118 enables understanding algorithms, input integrity or BC alignment. For sTGC "potential" can mean,  
119 for example, a Pad Trigger for the sector. For MM, this corresponds to strip data in a roads.
- 120 9. Monitor board and crate voltages and temperatures and send them to DCS. This may done by the  
121 IPMI shelf manager.

#### 122 1.4 Requirements and Limitations of the trigger

123 The main requirement for the NSW trigger is to provide track segments from the NSW detectors to be  
124 matched to track segments from the Big Wheel. The angular resolution on the NSW track segments  
125 should be 1 mrad. In the data-taking period immediately after the installation of the NSW, during the Long  
126 Shutdown 2 (LS2), the Big Wheel trigger granularity is limited to an angular resolution of 3 mrad or larger.  
127 During the Long Shutdown 3 (LS3), new BW trigger electronics and a new MDT Level-1 trigger will be  
128 deployed, allowing for 1 mrad angular resolution in the BW.

129 **Matching resolution requirements** The NSW measures the radial coordinates in two planes, the azimuthal  
130 coordinate,  $\phi$ , and the angle,  $\Delta\theta$ , of track segments inside the wheel, i.e. before the end-cap toroid.  $\Delta\theta$  is the  
131 angle of the segment with respect to an 'infinite momentum track', i.e. a line from the IP to the segment's  
132 radial position in the NSW. The radial coordinate is measured by high-precision strips in both detectors.  
133 For the sTGC,  $\phi$  is determined by the triggering tower of sTGC pads, and for the MM, by small-angle  
134 stereo strips. The angle  $\Delta\theta$  is to be measured to an accuracy close to 1 mrad. The NSW trigger logic will  
135 reject track segments with  $\Delta\theta > \pm(7 - 15)$  mrad (the final, configurable, value will be determined from  
136 future studies). The angle  $\Delta\theta$  is passed to the Sector Logic, but is not used in the Phase-1 trigger decision.  
137 The corroboration with the BW trigger is done by projecting the 'infinite momentum track' through the  
138  $R - \phi$  point of the segment in the NSW onto the plane at the nominal Z of the most downstream sTGC  
139 detector's wire plane. This matching requires the NSW trigger segments to have:

- 140 •  $\phi$ -resolution of 20 mrad
- 141 •  $\eta$ -resolution of 0.005

142 The sTGC Trigger Processor can find up to four segments. This limit arises from the fact that the Router  
143 has four output fibers for transferring strip charges from up to four TDS ASICs. In one bunch crossing,  
144 at 4.8 Gb/s, a TDS ASIC can transmit data for only one band of strips. In the case where a band spans  
145 the strips handled by two TDS ASICs (about 10% of the bands), two fibers are needed to transmit the  
146 band data. Simulations for Phase-2 indicate that bunch crossings with more than three active bands will  
147 be less than one per mil. The MM is limited to sending up to eight segments. The Sector Logic requires

148 the removal of duplicates and can accept at most eight segments from the NSW Trigger Processor. The  
 149 algorithm to remove duplicates from the sTGC and MM is described in Section 3.2.

150 **1.5 Description of the Trigger Processor hardware**

151 Figure 2 shows an overview of the New Small Wheel trigger and readout paths. Figure 3 shows a block  
 152 diagram of the Trigger Processor mezzanine card where most of the firmware resides. This includes inputs  
 153 form the sTGC and MM detectors, the Trigger Processor algorithms, and the output to the Sector Logic.  
 Figure 4 shows the input and output connections of the Trigger Processor.

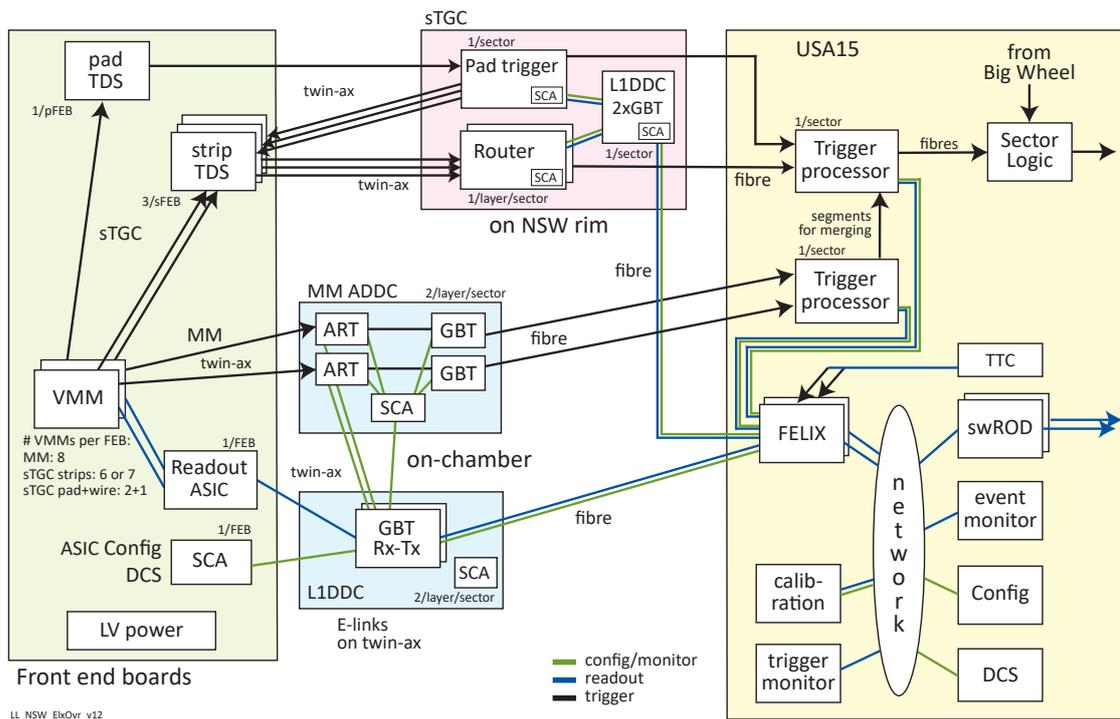


Figure 2: Block diagram of the NSW electronics showing both trigger and readout paths from the VMM front end ASIC. The readout path passes through the L1DDC's to FELIX. The sTGC trigger path passes through the Pad TDS, Pad Trigger, strip TDS and Router to the Trigger Processor. The MM trigger path passes through the ADDC to the Trigger Processor.

154

155 The Trigger Processor hardware consists of an AdvancedTCA (ATCA) carrier board and two mezzanine  
 156 cards. The boards are an evolution of the SRS boards designed within the RD51 collaboration [3, 4]. Each  
 157 mezzanine has two FPGA's, each with 36 bi-directional fiber optic links (microPODs). The connector to  
 158 the carrier card is **not** the ATCA standard Advanced Mezzanine Card (AMC) connector. Each mezzanine  
 159 is intended to serve one sector – one FPGA for MM and one for sTGC. There are 68 fast, low latency  
 160 LVDS connections between the two FPGA's on a mezzanine. Within one bunch crossing, eight segments  
 161 can be communicated between the two algorithms. The segments found by the detector that finishes first  
 162 are sent via these low latency links to the other for merging; up to eight segments can be sent to the Sector  
 163 Logic. The FPGA's are Xilinx Virtex 7 XC7VX690T-FFG1158 FPGA's with GTH transceivers.

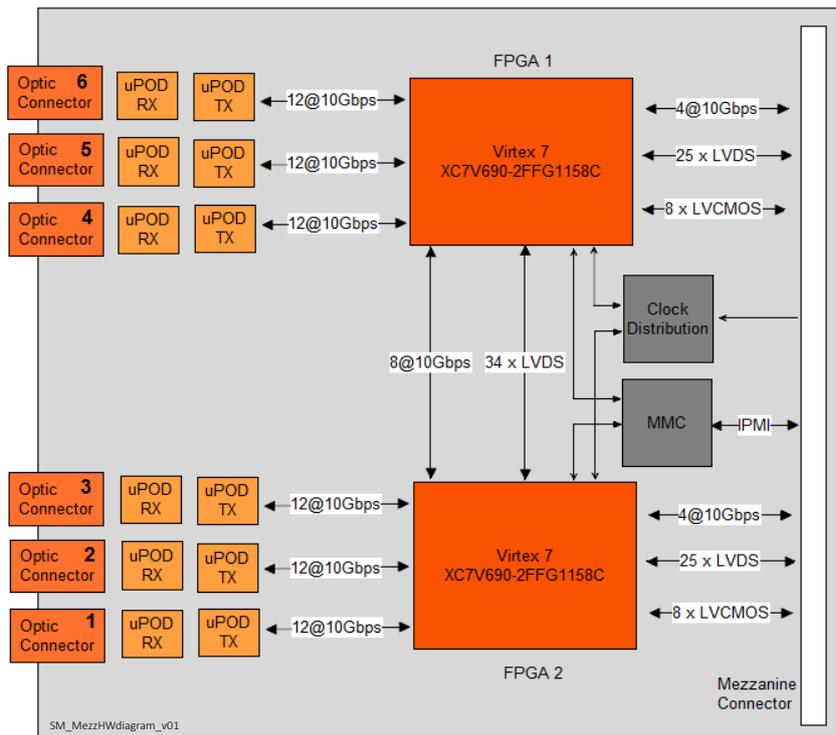


Figure 3: Trigger Processor mezzanine card. Note all LVDS, transceiver and CMOS links are bidirectional and so include twice as many signals as labeled in the diagram.

NSW Trigger Processor context

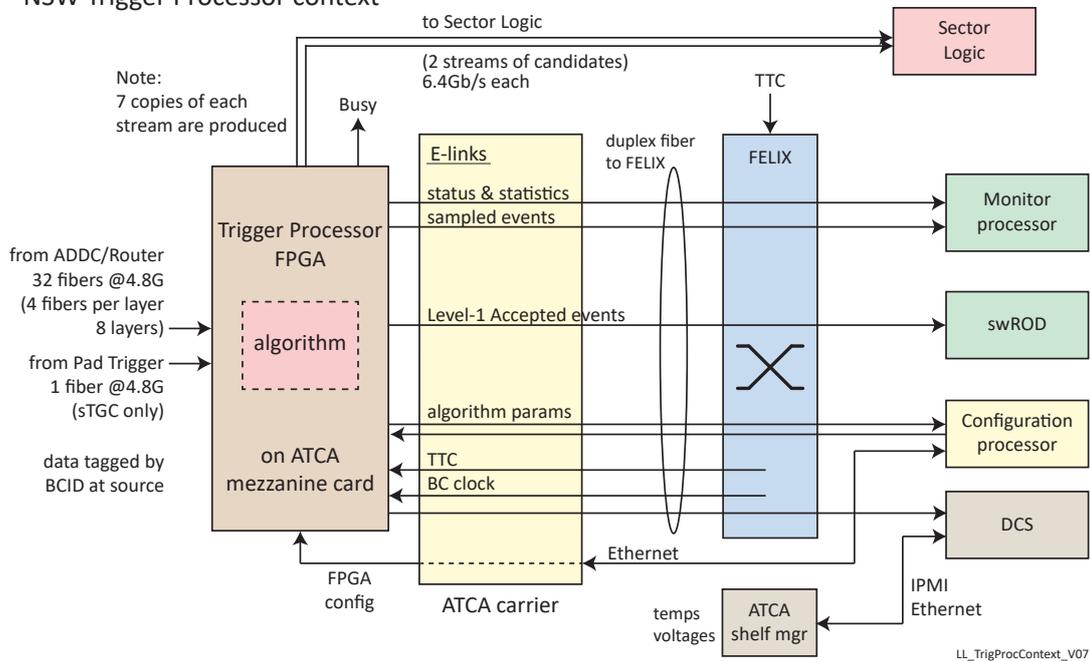


Figure 4: Trigger Processor context diagram

164 Figure 5 shows a block diagram of the Trigger Processor carrier board [5]. The FPGA's are Xilinx Kintex  
 165 Ultrascale XCKU060, the board manager is a Zynq XC7Z015 and the gigabit Ethernet switch used has  
 166 external connections to the ATCA Base interface and the Rear Transition Module (RTM), as well as  
 167 internal connections to the IPMC, Zynq and FPGA's. The FPGA's on the carrier board support the trigger  
 168 algorithms running on the mezzanine FPGA's. It can distribute either a Bunch Crossing (BC) clock  
 169 (40.079 MHz) recovered from FELIX or a local oscillator, and monitor and configure the trigger algorithm.  
 170 In addition, data from the mezzanine FPGA's sent to FELIX pass through the carrier FPGA's (Level-1  
 171 Accept and monitoring).

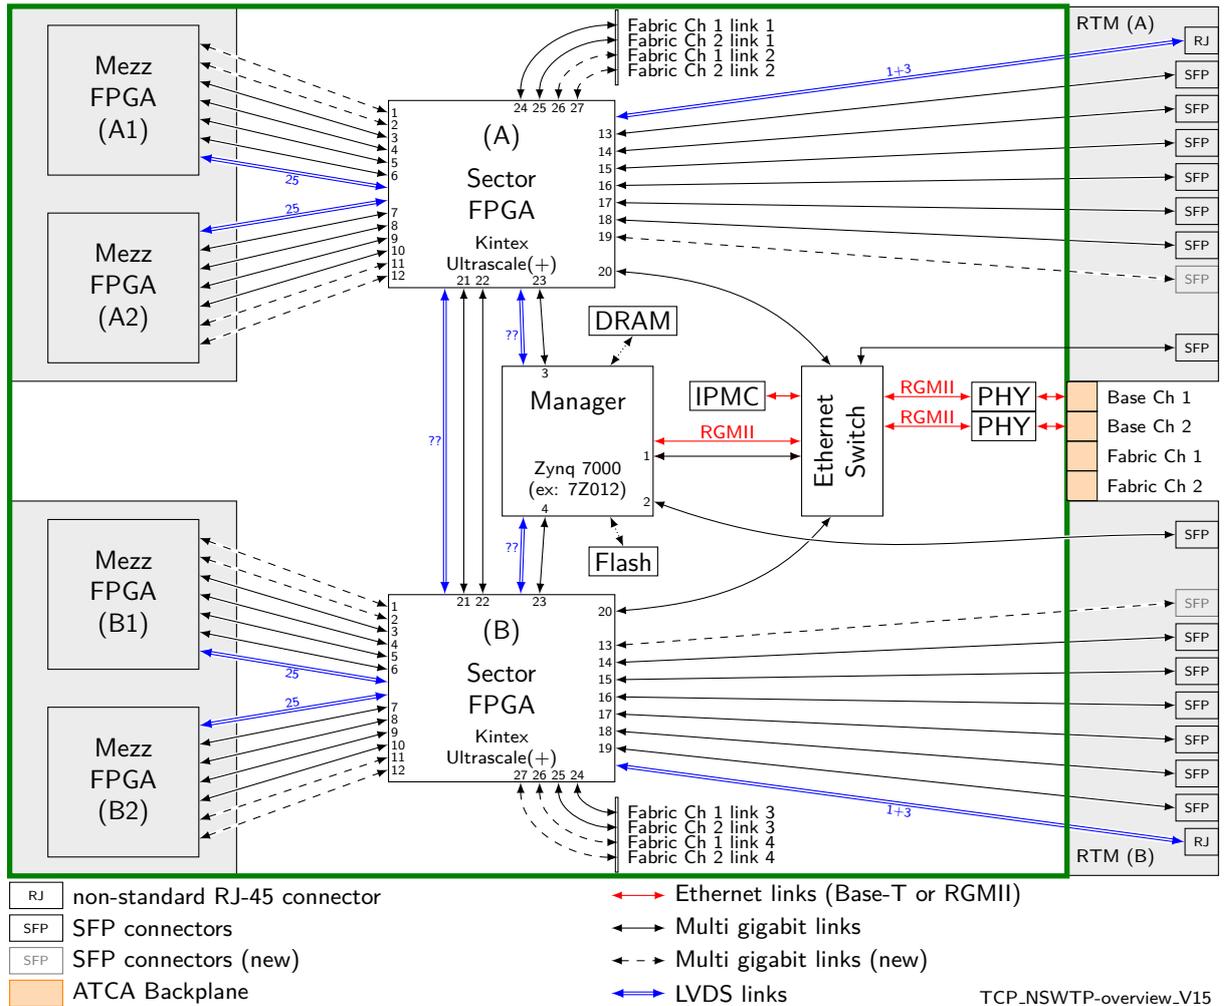


Figure 5: Trigger Processor carrier card highlighting the interconnections.

## 172 2 Algorithms

173 The output of the Trigger Processors to the Sector Logic is up to eight unique segments that point to Big  
 174 Wheel (BW) Regions-of-Interest (RoIs). Table 2 shows the  $\eta$  boundaries of the RoIs. The  $\phi$  RoIs are  $15^\circ/4$   
 175 in size in the Endcap ( $|\eta| < 1.92$ ) and  $7.5^\circ/4$  in the Forward part ( $|\eta| > 1.92$ ). The list of segments must be  
 176 transmitted every bunch crossing. To relax the precision requirements for construction and installation  
 177 of the NSW, first order alignment corrections must be done by the Trigger Processor. Since R and  $\phi$  RoI  
 178 indices and not R- $\phi$  coordinates in the NSW coordinate system are sent to the Sector Logic, the final  
 179 look-up table from coordinates to indices should take into account misalignments of the Big Wheel.

Table 2:  $\eta$  boundaries of the Big Wheel RoIs

Forward		Endcap			
$\eta$	tan	$\eta$	tan	$\eta$	tan
2.4148	0.18022	1.9172	0.30053	1.4622	0.48973
2.3777	0.18714	1.8932	0.30817	1.4428	0.50049
2.3439	0.19370	1.8679	0.31643	1.4227	0.51187
2.3100	0.20050	1.8420	0.32518	1.4015	0.52424
2.2755	0.20768	1.8153	0.33443	1.3783	0.53822
2.2417	0.21497	1.7882	0.34417	1.3528	0.55405
2.2070	0.22277	1.7608	0.35427	1.3257	0.57157
2.1734	0.23056	1.7328	0.36500	1.2972	0.59068
2.1411	0.23835	1.7050	0.37597	1.2681	0.61114
2.1088	0.24639	1.6781	0.38694	1.2384	0.63284
2.0781	0.25431	1.6524	0.39779	1.2086	0.65570
2.0476	0.26247	1.6274	0.40864	1.1798	0.67883
2.0166	0.27101	1.6038	0.41923	1.1528	0.70146
1.9855	0.27991	1.5810	0.42974	1.1275	0.72361
1.9537	0.28931	1.5591	0.44013	1.1036	0.74534
1.9214	0.29921	1.5385	0.45015	1.0807	0.76704
1.9163	0.30082	1.5187	0.46005	1.0587	0.78873
		1.4997	0.46982	1.0375	0.81043
		1.4811	0.47959	1.0334	0.81474

### 180 2.1 sTGC algorithm

181 The sTGC algorithm is based on computing the radial position and segment direction from centroids of  
 182 strips in two groups of four detector layers. Section 3.2 of the Preliminary Design Review [2] details the  
 183 sTGC trigger algorithm, including that for alignment corrections. Figure 6 shows a block diagram of the  
 184 sTGC Trigger Processor algorithm firmware.

185 The requirements for the sTGC algorithm are:

- 186 • Up to four bands of strips in each of eight layers must be capable of being processed to find up to  
 187 four segments. This limit arises from the four fibers from each Router to the Trigger Processor, one  
 188 band of strips per fiber per bunch crossing.

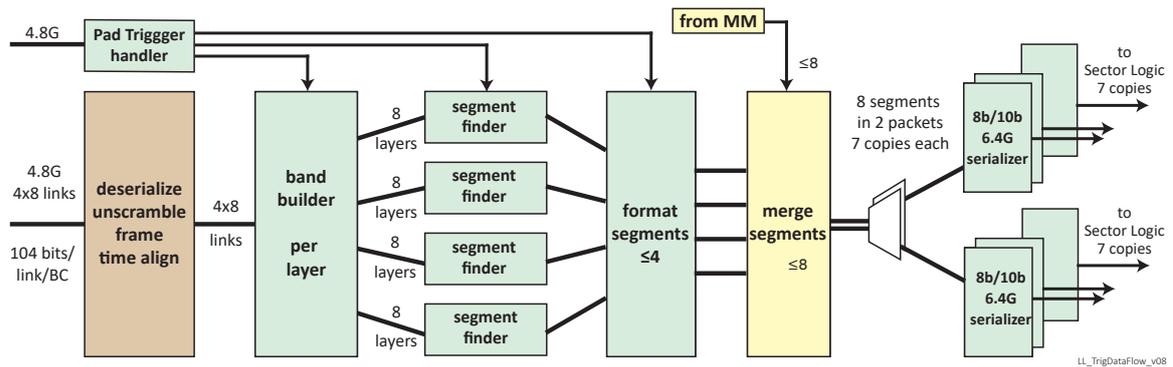


Figure 6: Block diagram of the sTGC trigger algorithm firmware. The data is captured and frame aligned with a 240 MHz clock from the deserializer. The distributor, segment finders and downstream blocks run with a 320 MHz clock. FIFO's between the clock domains must be avoided because of their long latency.

- 189 • The  $\phi$  resolution is determined by the pad geometry. See Table 3. A  $\phi$ -id is passed from the Pad
- 190 Trigger to the Trigger Processor, unchanged, via the strip TDS and Router as well as directly from
- 191 the Pad Trigger.
- 192 • Cuts and thresholds are required to reject  $\delta$ -rays (a wide cluster) and neutrons (very large charges per
- 193 strip)
- 194 • Due to dead areas (frames) some tracks can pass through at most four active detector regions. A
- 195 4-layer algorithm, with reduced resolution, must be provided.
- 196 • First order alignment corrections must be made to reduce the needed precision in construction and
- 197 mounting the sTGC wedges.

198 The requirements for a valid centroid for a cluster are:

- 199 • From two up to seven strips can form a cluster. (configurable)
- 200 • The strips must be adjacent.
- 201 • One hole (a strip below threshold) will be allowed. (Not allowed in the current algorithm)
- 202 • The strips of a cluster must be within the data sent by the TDS in one BC. (The TDS can include
- 203 strips in a configurable window up to 25 nsec beyond the current BC.)
- 204 • Some ( $\approx 10\%$ ) bands of strips span the boundary between two TDS ASIC's. In this case, the two
- 205 parts of a band are sent independently from the two TDS ASIC's through the Router and then via
- 206 two fibers to the Trigger Processor where they must be concatenated.

Table 3:  $\phi$  segmentation of the sTGC

Small Sector							
	$\eta$ max	$\eta$ min	One pad		One tower		
			Width	Nmax	Width	Nmax	
Quadruplet 1	2.414	1.917	5°	4	1.25° <sup>(1)</sup>	16	
Transition	1.917	1.853		(2)		(3)	
Quadruplets 2&3	1.853	1.281	7.5°	3	1.875° <sup>(1)</sup>	12	
Large Sector							
	$\eta$ max	$\eta$ min	One pad		One tower		
			Width	Nmax	Width	Nmax	
Quadruplet 1	2.452	1.961	5°	7	1.25° <sup>(1)</sup>	28	
Transition	1.961	1.900		(2)		(3)	
Quadruplets 2&3	1.900	1.298	7.5°	5	1.875° <sup>(1)</sup>	20	

<sup>1</sup> 1.25° = 22 mrad; 1.875° = 33 mrad

<sup>2</sup> Either Quadruplet 1 or Quadruplets 2 & 3

<sup>3</sup> The handling in the Pad Trigger of the transition between Quadruplet 1 and Quadruplets 2 & 3 is not totally defined yet.

### 207 2.1.1 sTGC input capture

208 The trigger input data arrives from the 32 Routers and the Pad Trigger at 4.8 Gb/s, i.e. 120 bits per bunch  
 209 crossing. Table 4 shows the format of the sTGC input data from the Router. Figure 7 is a block diagram  
 210 of the sTGC signal capture logic. From the Router, the data arrives either as one 30-bit “null” frame or,  
 211 for trigger data, as a four 30-bit data frame sequence. 26 bits of each frame are scrambled. 30-bit frame  
 212 alignment is done by searching for a stable header pattern, either 0b1010 or 0b0110. Should unexpected  
 213 data be found in the header position, the searching is repeated. Note that the output from the serial  
 214 transceivers is 20 bits at 240 MHz, whereas the descrambler works with 30-bit words, with output at  
 215 160 MHz. Therefore three 20-bit words must be captured to provide two 30-bit words for the descrambler.  
 216 In fact, a fourth 20-bit word must be captured to allow for the arbitrary bit alignment of the incoming data.  
 217 If the 30-bit frame is a null, its data is written to output registers. A 30-bit data frame which immediately  
 218 follows a null frame is assumed to be the first of a 4-frame data sequence. This 120-bit alignment is  
 219 assumed until another null frame is received. Different links may arrive in different 240 MHz clock periods.  
 220 They are aligned to a shared 240 MHz clock. The BCID is used to confirm that all 32 inputs belong to the  
 221 same bunch crossing.

222 From the Pad Trigger, data arrives from the transceiver, already 8b/10b decoded, as 16-bits at 240 MHz.  
 223 The K28.5 comma character is used to align the 120-bit data frames. The 16-bit data is written to a  
 224 width-matching FIFO at 240 MHz. The 96-bit data word is read from the FIFO by the Trigger Processor  
 225 Distributor with a 320 MHz clock. The data format is shown in Table 5.

Table 4: Format of the sTGC 120-bit input data from the Router. The data payload is 104 bits.

Data	Frame 0		Frame 1		Frame 2		Frame 3	
bits	header	data	header	data	header	data	header	data
	4	26	4	26	4	26	4	26
	1010	xxxx	1010	xxxx	1010	xxxx	1010	xxxx

Null	Frame 0		
bits	header	ID	undef
	4	10	16
	0110		xxxx

data	Band-ID	Hi/Lo	Strip charges	$\phi$ -ID	BCID
bits	8	1	14x 6= 84	5	6

ID	Fibre	Sector	Layer	Endcap
bits	2	4	3	1

Table 5: The data record sent every bunch crossing by the Pad Trigger to the Trigger Processor. Each Band-id is 8 bits; Each phi-id, 6 bits. Missing bands have band-id =0xAA. The high flag bit indicates a “monitor this BC” request to be passed to the CTP via the Sector Logic. The BCID is 12 bits. The 16-bit coincidence field shows, for each band for each wedge, the of number of layers in coincidence:  $4 \times 2 \times 2$ -bits. The 12 bytes are 8b/10b encoded; the first is a K28.5 comma symbol.

-	0-1	2-5	6-8	9-10
comma	flags   BCID	band-id[3..0]	$\phi$ -id[3..0]	coincidence

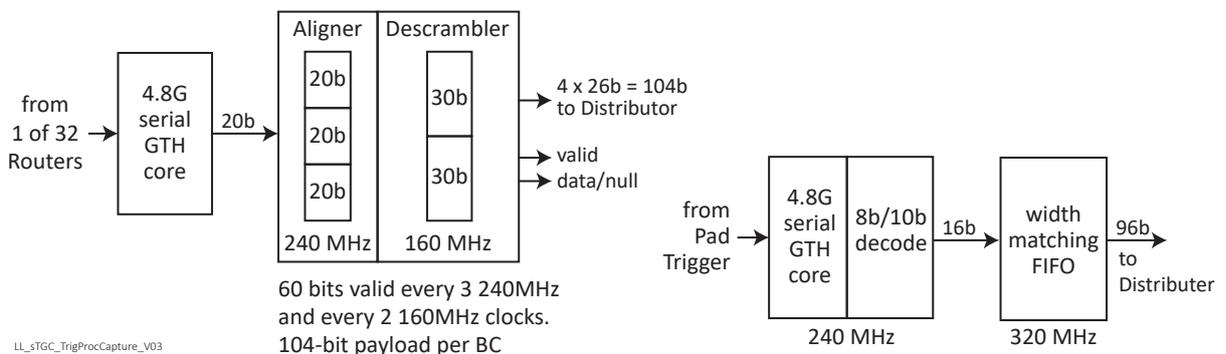


Figure 7: Block diagram of the sTGC signal capture logic. The clocks shown are derived from the 240 MHz deserializer RX clock via a MMCM.

## 2.1.2 Distributor

The Distributor of the Trigger Processor must pre-process the information sent by the Routers to the Trigger Processor to assemble a message per strip-band (containing the information from all detector layers), and route the assembled messages to the centroid finders. The assembled messages must contain: the strip-band ID, the high/low indicator for the group of strips, the strip data, the  $\phi$ -ID, radial offset of the starting point of the band, and the BCID. The Distributor must arrange strip information from the detector layers in increasing Z, according to a LUT mapping the serializer channel to Z position. The Router channel ID, which identifies the detector layer, is obtained from the ID in the NULL frames sent by the Router when it has no active data, e.g. before the Pad Trigger is enabled.

Table 6: Format of the Distributor output, for one layer, to the centroid finders. A band is actually 17 strips; the Hi/Lo bit indicates if the 14 6-bit charges transmitted are for the upper or lower 14 strips in the band. The centroid algorithm must take the implied offset into account.

Field	Strip data (layers 1,...,8)	band-ID	Hi/Lo	$\phi$ -ID	Data valid
bits	$(6 \times 14 \times 8)$	8	1	5	1

The strip information is delivered to the Trigger Processor through the 32 fibers connecting the eight Routers (one per detector layer) to the Trigger Processor. In each of these fibres, the Router sends information from a single TDS ASIC. A complication arises from the fact that strip bands may span two TDS ASICs. In this case, each TDS ASIC transmits the strip data that it has, and the Router then receives two active inputs for a single band. For simplicity, and to reduce latency, the Router does not merge the two streams to a single output fiber, but just passes them on two fibers to the Trigger Processor. In this case, the most significant bit within the 8-bit band-ID information is to be used to indicate which of part of the strip band is transmitted from each of the two TDS ASICs. The Distributor then must route the appropriate strip charges from the two strip records to their correct location in the 14-strip record for the detector layer that corresponds to these TDS ASICs. To make this routing possible, a LUT describing the correspondence between band-IDs and the correct location within the 14-strip record is to be provided per detector layer. The merging is implemented using a barrel shifter (6-bit x 14) to route the strips to their correct location. The barrel shifter is controlled by a 4-bit value extracted from a LUT addressed by the 8-bit band-ID. Following the barrel shifter, a 1-to-2 de-mux (84-bit) is used to route the strip data to the complementary strip info if merging is necessary. The de-mux is controlled based on the Pad Trigger data as explained below. The de-mux is followed by a 84-bit x 2 OR that merges the hi and lo strip in the band. Figure 8 illustrates this merging.

Because it does not pass via the TDS ASICs and Router, the information sent by the Pad Trigger arrives  $\sim 5$  BCs before the TDS data. It is used for two purposes. First, it is used to configure the routing, so complementary TDS messages (belonging to the same band and layer) are concatenated. Second, it is used to configure the routing of the strip information so that for a given band, all its layers are sent to the same sTGC algorithm instance. The Pad Trigger data consists of four flag bits, the BCID, and for each of four bands, the band-id,  $\phi$ -id, and for each wedge of the band, the number of active layers. The 11 bytes of data are shown in Table 5.

## 2.1.3 FPGA resources

The current mezzanine FPGA design includes the blocks shown in Figure 6. It excludes the calculation of the alignment corrections, but includes the I2C block for configuring the external jitter cleaners and a

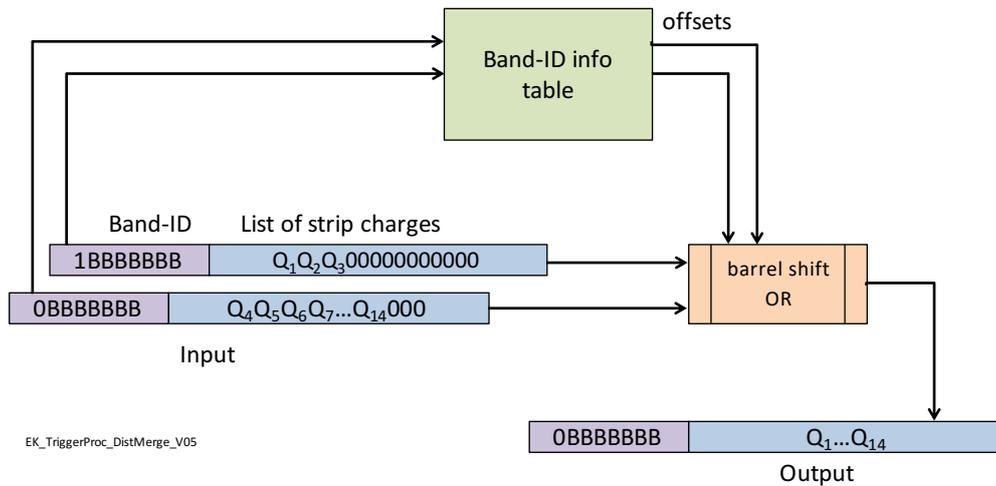


Figure 8: Merging of strip information from a strip band spanning two TDS ASICs – strips 1 to 3 in one, 4 to 14 in the second. There is a separate look-up table for each layer.

261 sniffer for debugging that was used to output UDP packets to Ethernet via the carrier. The FPGA resource  
 262 consumption is as follows:

263 LUT's: 20%; Block RAM: 42%; DSP blocks: 7%, MMCMS (clock blocks): 20%. **To be updated**

## 264 2.2 Micromegas algorithm

265 This section describes the Micromegas Trigger Processor (MMTP) introduced in [6] and its implementation,  
 266 as well as necessary modifications to correct for detector misalignments. It also includes details of simulation,  
 267 and testing on hardware. Further details of the algorithm performance and alignment corrections can be  
 268 found in [7].

### 269 2.2.1 Description

270 The MMTP algorithm has been described in detail in [6]. Here, its main features are summarized. The main  
 271 components of the algorithm are shown in Figure 9. In order to handle the detector hit data throughput,  
 272 there will be 16 copies of the algorithm operating in parallel. Each algorithm will share information with  
 273 its neighbor to avoid boundary issues. The clock alignment stage is handled in deserializer block, all logic  
 274 following this stage is clocked using the same 320 MHz clock. This clock is derived from the 40 MHz  
 275 bunch crossing clock at the FPGA input with a 0-phase offset to maintain a known phase relationship with  
 276 other components in the system.

277 **Decoder** Incoming strip hit addresses are decoded into global slope values using a multiplication with a  
 278 constant. A strip's stored slope value is defined as the orthogonal distance between a given strip and the  
 279 beam line divided by the z location of the relevant detection plane. It is pre-computed taking into account a  
 280 strip offset and a z position stored for each of the 8 planes and 16 radial segments of each wedge. The  
 281 slope range is divided into 512 slope-roads which will be used to form a track candidate.

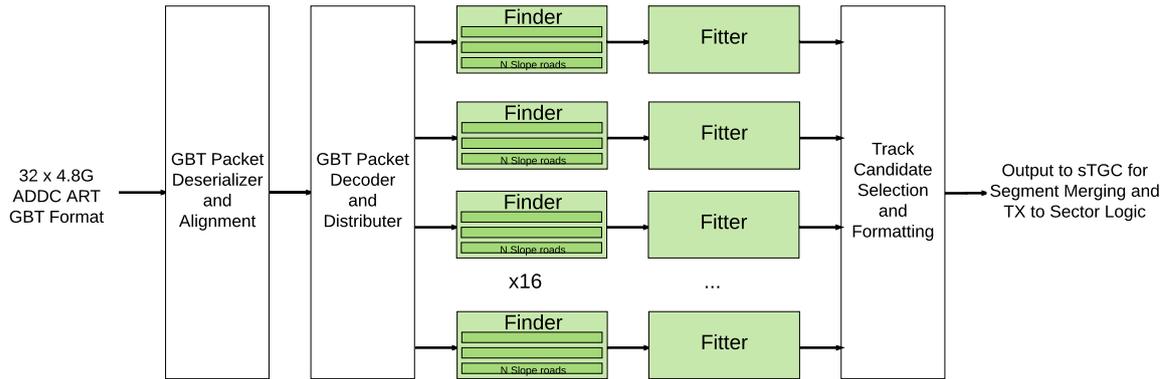


Figure 9: Block diagram of the MM algorithm

282 **Finder** The track finding algorithm is slope-road based with the wedge divided from bottom to top  
 283 into approximately 1000 roads and corresponds to about eight VMM strips. Hit data from the decoder  
 284 is routed to the corresponding road and that road is marked as being hit. Details of the hit that can't be  
 285 calculated from the road itself is stored. Each slope-road can hold a single hit for each plane and this hit  
 286 expires after the hit integration time. The current implementation uses a hit integration time of eight bunch  
 287 crossings. The slope roads will overlap each other by 1/2 to accommodate tracks on the road boundaries.  
 288 Each slope-road is checked once per bunch crossing to determine if a coincidence threshold has been met.  
 289 Coincidence requires a minimum number of planes to be hit and the oldest hit of the track to be expiring.  
 290 Coincidence identification is accomplished using combinatorial logic and a priority encoder. The strip  
 291 number and slope for each hit are calculated and passed to the track fitting algorithm.

292 In an effort to reduce the impact of background hits, the Finder algorithm has been rewritten to separate the  
 293 X-plane and UV-plane coincidences. The fitter will first identify a coincidence trigger on the X planes  
 294 and then scan from left to right all of the UV-roads that overlap the triggered X-road. The overlapping  
 295 area of UV-roads are “diamonds”. For each X trigger there can be as many as 57 UV-road diamonds that  
 296 are searched. The number of diamonds is configured for each algorithm region to match the length of  
 297 the Micromegas detector strips. Every bunch crossing each of the 16 algorithm regions can process two  
 298 independent X-road coincidences. For every X-road coincidence the Finder can process three UV-road  
 299 coincidences.

300 The Finder implementation focuses on minimizing the resources used in each slope road. This component  
 301 is multiplied by the number of slope roads and number of planes, thus, uses the most resources in the total  
 302 design.

303 **Fitter** In the fit, individual hit slopes in a slope-road are used to calculate global slopes associated  
 304 with each plane type, which are averages (e.g.  $M_X^g$  for the average slope of horizontal planes). These  
 305 in turn are used to calculate the three composite slopes: slopes associated with the horizontal ( $m_x$ ) and  
 306 vertical coordinates ( $m_y$ ) and the local slope of hits in the horizontal planes ( $M_X^l$ ), all of which are shown  
 307 in Equation 1. In Equation 1,  $\theta_{st}$  is the stereo angle of  $1.5^\circ$ ; the sums are over relevant planes;  $\bar{z}$  is the  
 308 average position in  $z$  of the horizontal planes; and  $y_i$  and  $z_i$  in the local slope expression refer to the  $y$  and

309  $z$  coordinates of hits in X planes.

$$m_x = \frac{1}{2} \cot \theta_{st} (M_U^g - M_V^g), m_y = M_X^g, M_X^l = \frac{\bar{z}}{\sum_i z_i^2 - 1/n (\sum_i z_i)^2} \sum_i y_i \left( \frac{z_i}{\bar{z}} - 1 \right) \quad (1)$$

310 From these composite slopes, the familiar expressions for the fit quantities  $\theta$  (the zenith),  $\phi$  (the azimuth<sup>1</sup>),  
 311 and  $\Delta\theta$  (the difference in  $\theta$  between the direction of the segment extrapolated back to the interaction point  
 312 and its direction when entering the detector region; the following is an approximation) may be calculated,  
 313 as noted in [6]:

$$\theta = \arctan \left( \sqrt{m_x^2 + m_y^2} \right), \phi = \arctan \left( \frac{m_x}{m_y} \right), \Delta\theta = \frac{M_X^l - M_X^g}{1 + M_X^l M_X^g} \quad (2)$$

314 The fitter implementation is accomplished primarily by taking advantage of cascading DSP48E1 slices to  
 315 minimize resources and latency. The fitter is implemented as a pipeline and can accept one new track from  
 316 the finder every 320 MHz clock cycle.

### 317 2.2.2 MM input format

318 The ART (Address in Real Time) data from an entire sector will be transmitted to a single Trigger Processor  
 319 via 32 ADDCs. Each ADDC will transmit its data on a single fibre optic link. The Trigger Processor  
 320 therefore uses 32 fibres to receive the ART data from one sector. The MM and sTGC Trigger Processor  
 321 will share the same ATCA carrier card and each carrier will support two sectors.

### 322 2.2.3 ART Data Protocol

323 The ART Data from the ADDC will be transmitted using the GigaBit Transceiver (GBT) architecture and  
 324 transmission protocol in a low-latency widebus mode at a rate of 4.8 Gb/s. The Trigger Processor will take  
 325 advantage of the GBT firmware developed by the GBT Project and optimized for latency by Kai Chen, to  
 326 implement the receivers.

327 The GBT packet in widebus mode will provide 112 data bits and arrives once every bunch crossing. One  
 328 ADDC will service 32 VMMs and each packet can contain ART data from a maximum of eight triggered  
 329 VMMs. Each VMM will be uniquely identified to determine which MM strip on the sector is hit.

330 The GBT data will be encoded in the “hitmap” format where the triggered VMMs are represented by  
 331 asserting a bit in a 32-bit hit list as shown in Figure 10. The triggered VMM strip number is provided in an  
 332 encoded list as shown in Table 7.

---

<sup>1</sup> Defined with respect to the center ( $y$ ) axis and *not* the axis of the strips ( $x$ ) as is sometimes typical, so a hit along the center of the wedge has  $\phi = 0$

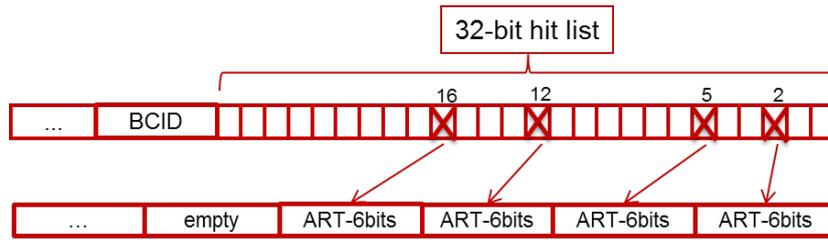


Figure 10: Option 1 VMM ID encoding using 32-bit hit list.

0b1010	BCID(12)	ERR_FLAGS(8)	HIT_LIST(32)	ARTDATA_PARITY(8)	8xART_DATA(6)
--------	----------	--------------	--------------	-------------------	---------------

Table 7: Option 1 ADDC GBT packet format.

333 Where:

334 BCID: 12-bit bunch crossing ID

335 ERR\_FLAGS: ADDC error indicators from the eight binary encoders attached to each priority  
336 encoder in the hit selection logic

337 HIT\_LIST: 32-bit list of flags corresponding to each of the 32 VMMs. 0: no hit, 1: hit. A register  
338 controls if this is a filtered (i.e. 8 hits max) or an unfiltered copy of the VMM flags registered in a  
339 particular BC.

340 ARTDATA\_PARITY: 8-bit parity the ART data computed by each of the 32 ART de-serializer units.  
341 Each bit corresponds to one of the ART data field selected by the priority unit

342 ART\_DATA: 6-bit triggered VMM strip number

### 343 2.2.4 Input data alignment and clock domain crossing

344 The ADDC data from 32 fibers arrives via GBT FPGA on a 40 MHz clock that is aligned to each receiver's  
345 RX out-clock. In order to align each fiber to a single recovered BC Clock, the data is registered twice using  
346 a clock that is phase aligned to the recovered BC clock. The first register uses a 40 MHz clock with settable  
347 phase adjustment in 45° increments specific to each fiber. (See Figure 11.) This register will be used to  
348 account for any individual fiber length differences. The second register will also be a 40 MHz clock with  
349 settable phase adjustment but the phase will be a single constant set for all fibers. This register will align  
the data from all fibers and provide a fixed latency.

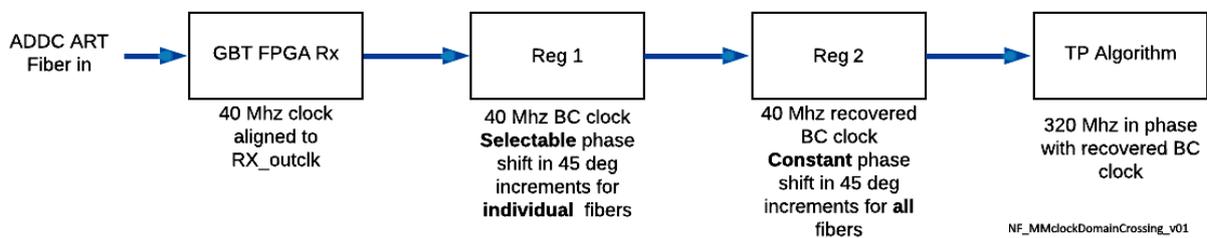


Figure 11: Block diagram of the Micromegas clock domain crossing

350

### 351 2.2.5 FPGA Resources

352 A full wedge design has been implemented. This includes 16 parallel algorithms, Finder and Fitter,  
353 along with the logic to handle tracks on the algorithm boundaries. A block diagram of the algorithm  
354 implementation is shown in Figure 9. The implementation uses a nominal configuration based on algorithm  
355 performance studies, Ref. [7] with the addition of a new Finder design using fine slope-roads and UV-road  
356 diamond triggering. The full algorithm resources required are ~65% of the Xilinx XC7VX690T FPGA. The  
357 basic alignment correction functionality involving plane shifts and tilt have been implemented. Alignment  
358 corrections involving twist and rotate have not yet been implemented. The implementation for the full  
359 wedge includes monitoring, L1A packet generation, and the lateral transfer of track segments to the sTGC  
360 FPGA. The removal of duplicate track segments has not yet been implemented.

361 Vivado 2018.3 is used in non-project mode to enable scripts to control the implementation process. This  
362 has greatly aided in timing closure. The implementation requires the use of placement blocks to meet  
363 timing.

### 364 2.2.6 Simulation

365 A software model of the algorithm has been developed to evaluate its performance using Athena generated  
366 VMM ART hit data. The firmware design is simulated using the same Athena data and the results are  
367 compared between the software model and firmware simulation for validation. Early testing of the algorithm  
368 used both behavioral and timing simulations. Recent algorithm validation focuses on hardware testing with  
369 behavioral simulations as a diagnostic tool. Modelsim is used for firmware simulations.

### 370 2.2.7 Hardware Testing

371 The MMTP design is being tested using two separate hardware platforms, the Xilinx VC707 Development  
372 board, and the ATCA HORX hardware. The HORX hardware is being tested with an ADDC emulator  
373 to generate the GBT packets containing Athena generated ART data. These packets are transmitted and  
374 received with a fiber in a loop-back scheme allowing the MMTP to receive data as if it were sent from an  
375 actual ADDC. The MMTP hardware results are compared with the results of the software simulation for  
376 verification.

377 The primary method of algorithm development has been in a cosmic ray test stand. This system includes  
378 an entire ART data chain consisting of a Micromegas octuplet, eight MMFE8s, two ADDC prototypes,  
379 and one MMTP running on a Xilinx VC707 development board. This system allows us to move beyond  
380 simulations and exercise the MMTP using muon tracks from a Micromegas detector. Over the past months,  
381 MMTP diagnostic data from millions of cosmic ray muon triggered events has been collected. A full  
382 description with test results can be found in [8].

383 During past integration weeks at CERN the HORX MMTP was successfully integrated with the ADDC V2  
384 in the ART data chain. The ART data from pulsed MMFE8s were successfully received by the MMTP  
385 and generating track candidates. All of the boards were synchronized to a clock that was sourced by  
386 FELIX. FELIX via RTM was used for all of the communication during tests utilizing SCAX for algorithm  
387 configuration, monitoring and L1A functionality to receive algorithm and segment data.

388 **2.2.8 Micromegas trigger alignment constants**

389 The expected tolerance for installation of the Micromegas chambers is roughly 2 mm, although it possible  
 390 to have larger excursions. In order to retain optimal performance of the trigger, alignment corrections  
 391 need to be performed. The corrections need to be performed rapidly to minimize latency, and also to have  
 392 a small number of constants to avoid large look-up tables. Figure 12 shows the categories of possible  
 393 misalignments.

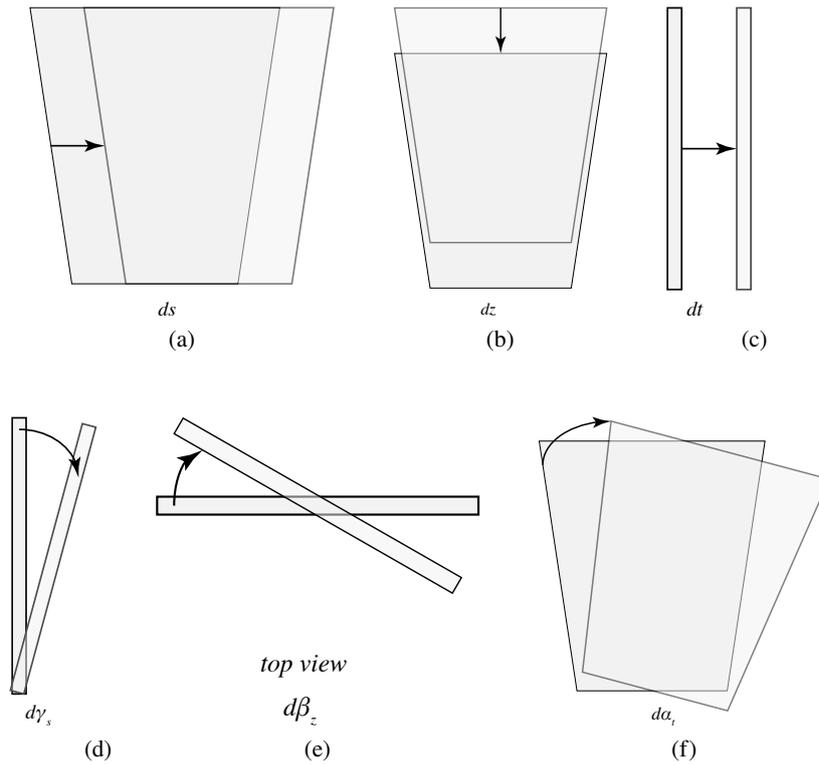


Figure 12: The different misalignment cases as defined in the AMDB manual.

394 **Translations** Translations were studied up to 5 mm out of perfect alignment, and showed that the  
 395 corrections gave acceptable performance.

- 396 a Translation along horizontal strip direction (s): This only affects small angle stereo. It is calculated  
 397 ahead prior to a run, stored in a small look up table that has 11 constants in memory, one lookup,  
 398 one addition – two operations.
- 399 b Translation perpendicular to beam and strip direction (z): This only requires a change in the calculated  
 400 constant ybase at the start of the run with no additional parameters.
- 401 c Translation along beamline (t): As with the translation in z, this requires only a change in an  
 402 initializing constant at the start of the run with no additional parameters.

403 **Rotations** Rotations of modules were studied up to  $\pm 1.5$  mrad, and it was found that acceptable corrections  
 404 can be made with a manageable number of parameters.

405 d Chamber rotations toward and away IP (gamma): A two-dimension lookup table is used to  
 406 parameterize the wedge into an 8 x 8 grid, and with one offset, this yields 56 parameters (64-8) and  
 407 one operation.

408 e Chamber rotations about wedge vertical (beta): This rotation creates a negligible effect for any  
 409 conceivable misalignment. No correction is performed.

410 f Chamber rotations in the azimuthal direction (alpha): This rotation can be corrected for with a  
 411 look-up table that partitions the wedge into discrete subsections. It requires two operations per hit in  
 412 the x strip (non-stereo strip), and requires 400 constants in memory.

413 A summary of correction methods, including resources necessary for the individual analytic cases, is shown  
 414 in Table 8.

Table 8: A summary of corrections with additional constants/operations (written as  $n_{const}c/n_{ops}op$ ;  $n_X$  is the number of X hits in a fit) necessary for analytic corrections. Yes means a correction exists but might not entirely remove misalignment effects, while yes+ means a quality of correction is only limited by knowledge of misalignment and memory.

	$\Delta s$	$\Delta z$	$\Delta t$	$\gamma_s$	$\beta_z$	$\alpha_t$
Analytic Resources	yes+ 11c/2op	yes+ 0c/0op	yes+ 0c/0op	yes 56c/1op	no —	yes 400c/2 $n_X$ op, 32c/12 $n_X$ op
Simulation	yes+	no	no	no	yes+	yes+

### 415 2.3 Clocking

416 The Trigger Processor has no direct connection to the TTC. The BC clock and TTC bits are received via  
 417 the FELIX fibre connection. The BC clock is recovered with fixed latency by the fixed latency version of  
 418 the CERN GBT-FPGA firmware [9, 10] running in the Carrier. Figures 13 and 14 show the details of the  
 419 fixed latency BC clock recovery from the 240 MHz user bus clock from the 4.8 Gb/s deserializer. The  
 420 120-bit frame transferred every BC clock consists of a 4-bit fixed header (Idle=0b0110 or Data=0b0101)  
 421 and a 116-bit scrambled payload. The fixed latency of the recovered BC clock is attained by shifting its  
 422 rising edge so that it is aligned to the start of the fixed header pattern.

423 All the needed design clocks, 40, 80, 160, 240 and 320 MHz, are generated from the recovered bunch  
 424 crossing clock, each with an edge aligned to the BC clock edge. The deserializer and serializer reference  
 425 clocks are generated from the BC clock by “zero-delay” jitter cleaners that align their output to the input  
 426 BC clock. Since the various design clocks preserve fixed phase with respect to the BC clock, clock domains  
 427 can be crossed without FIFO’s. This avoids the latency inherent in FIFO’s. Figures 14 and 16 show the  
 428 Micromegas and sTGC, respectively, clock schemes.

429 Note that Micromegas Trigger Processor does not need to transmit its segment output to the sTGC Trigger  
 430 Processor with fixed latency. The Micromegas segment output is passed to the sTGC clock domain where  
 431 it is buffered in a FIFO until the sTGC algorithm completes.

432 See also Section 3, Output to the Sector Logic.

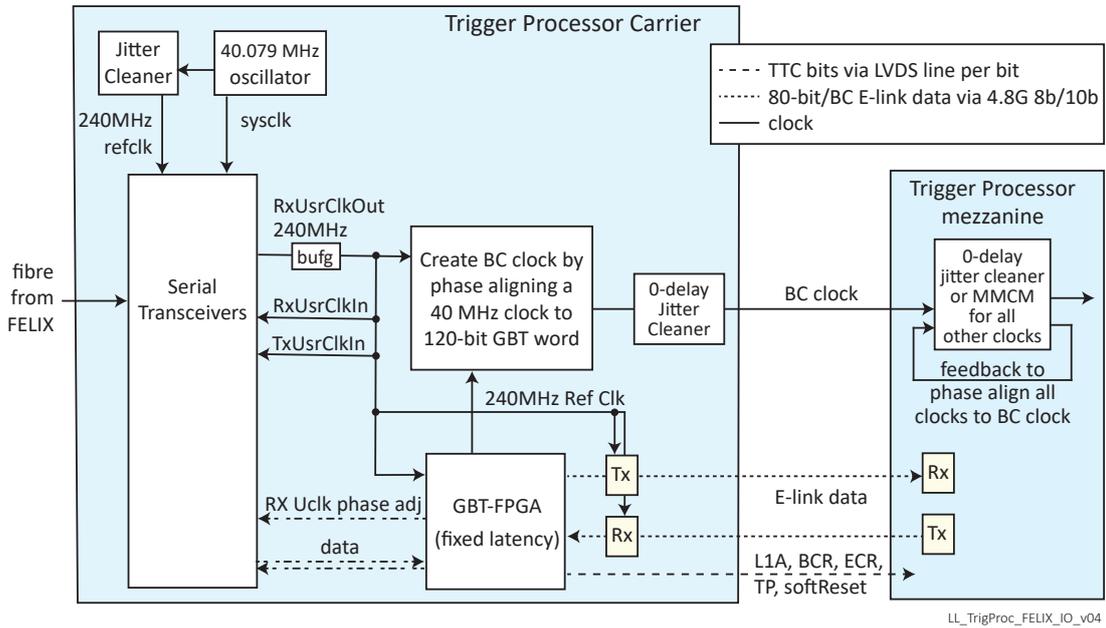


Figure 13: Details of the fixed latency BC clock recovery from FELIX. A 120-bit packet is received every bunch crossing.

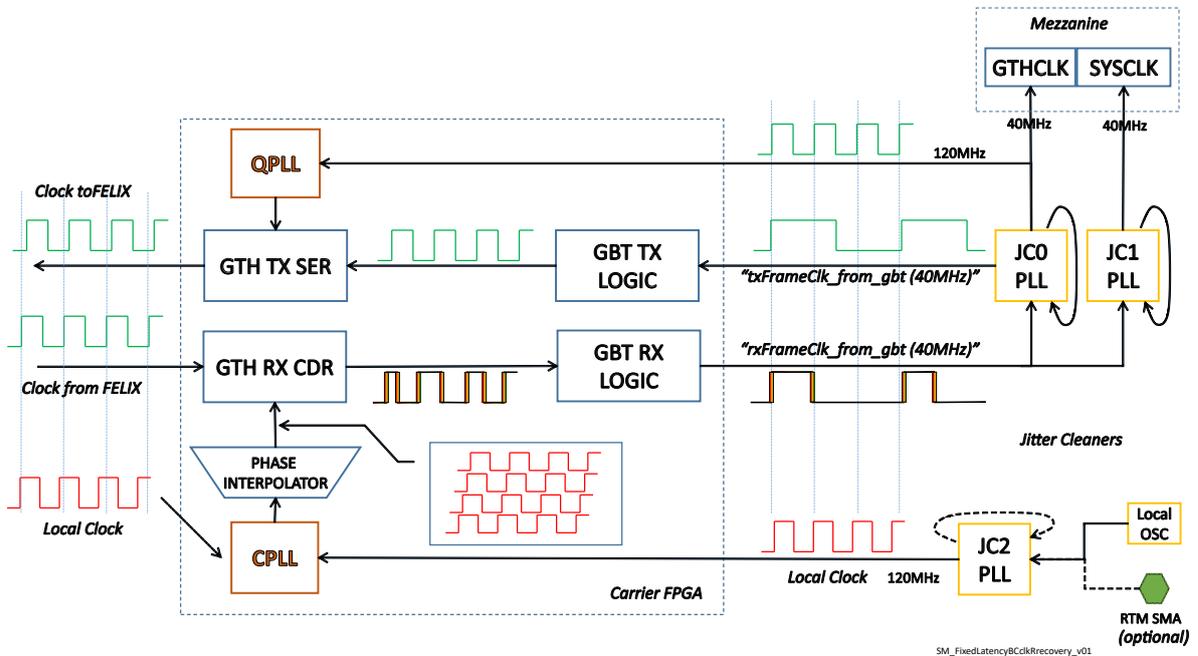


Figure 14: Fixed latency BC clock recovery – details

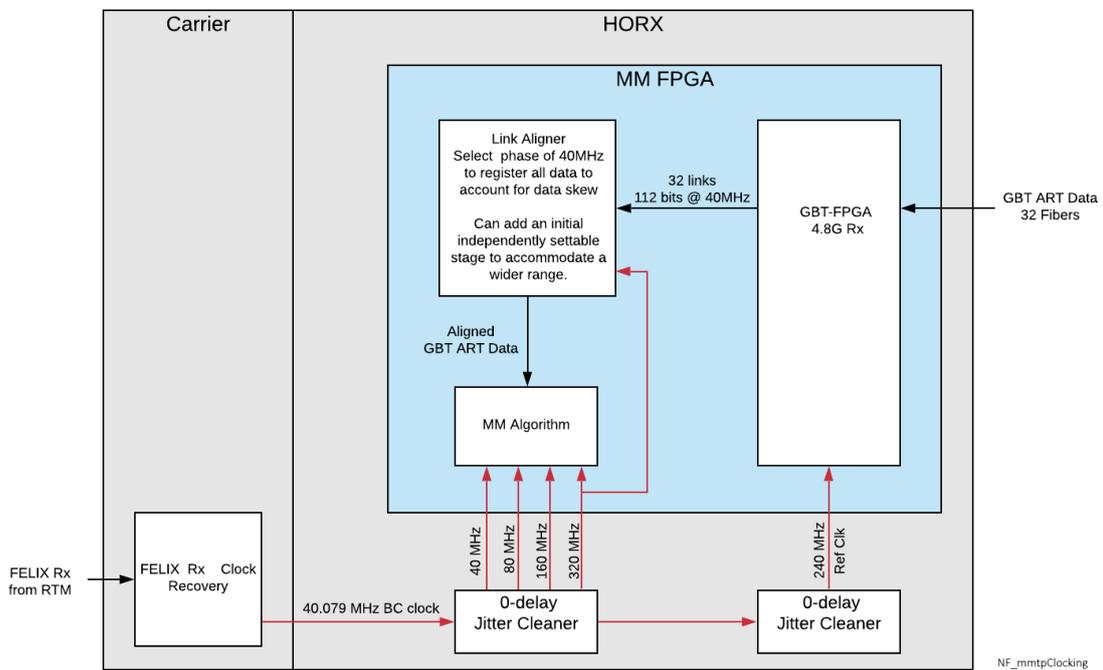
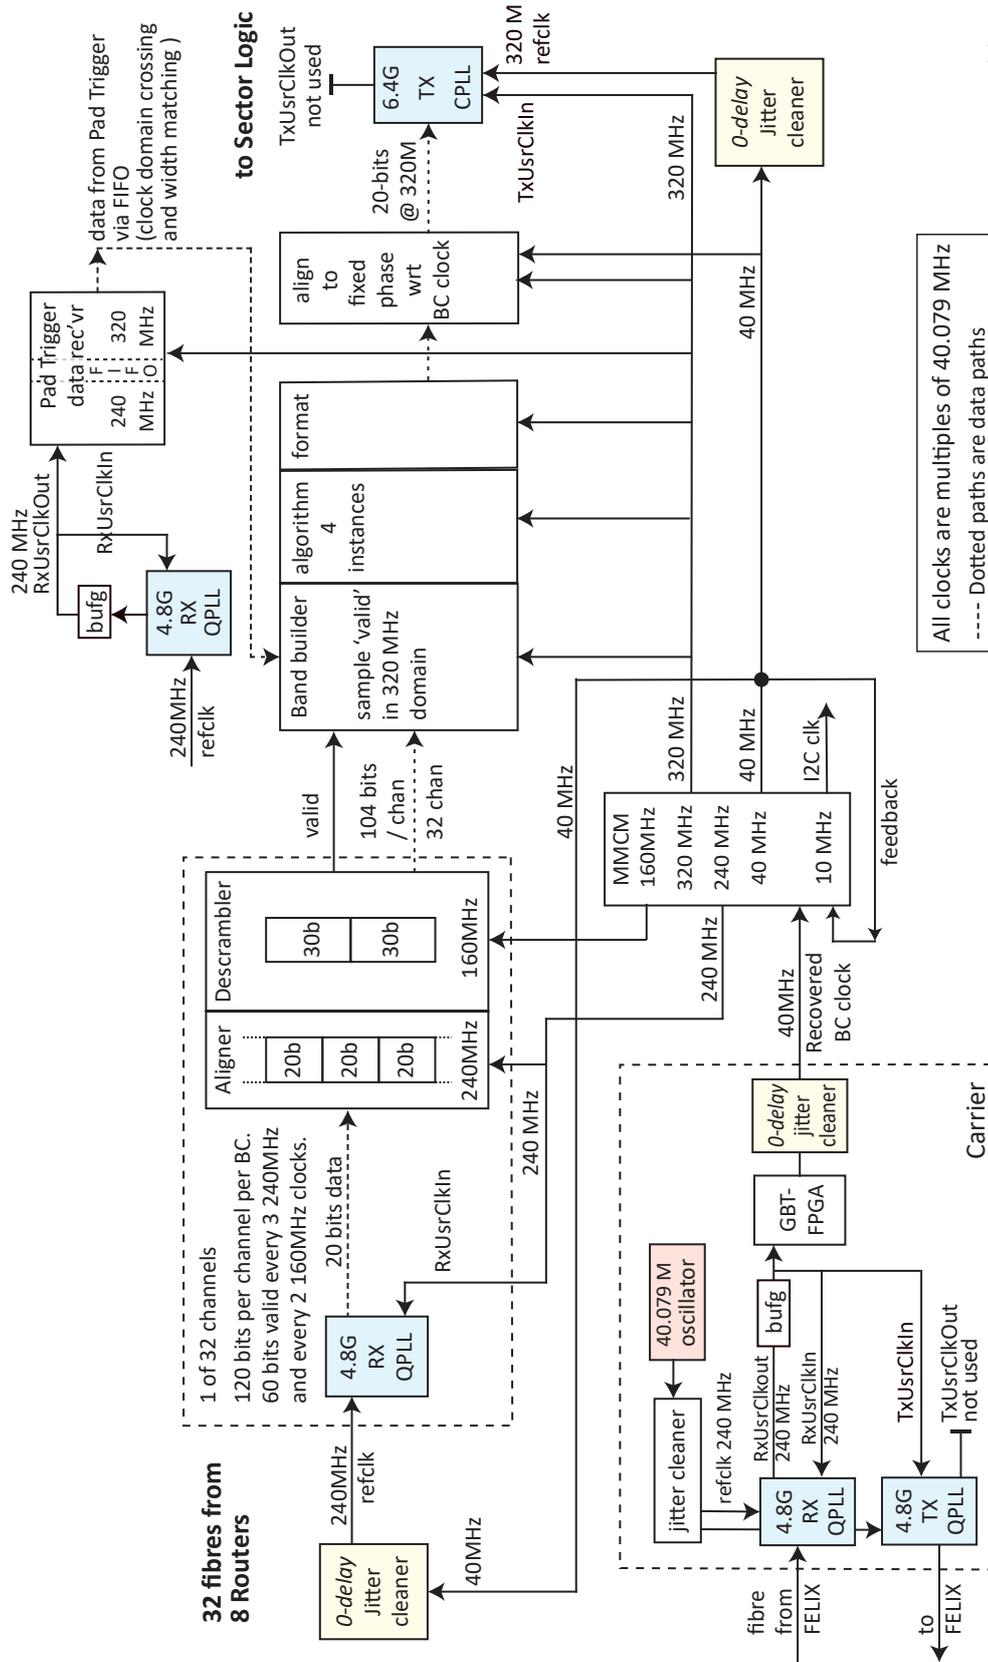


Figure 15: Micromegas trigger clocking



LL\_sTGCTriggerClocking\_v06

Figure 16: sTGC trigger clocking

### 3 Output to the Sector Logic

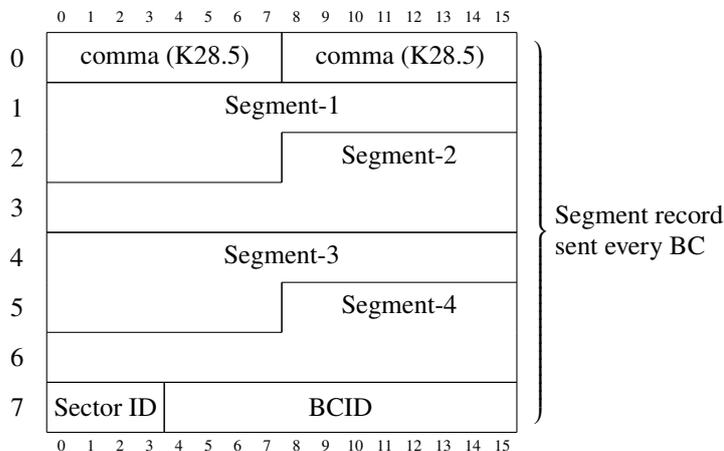
The output data format of a track segment from the NSW Trigger Processor is shown in Table 9. One track segment is represented in a 24-bit field, that includes  $\Delta\theta$ ,  $\phi$  and  $R$  position information, and three flags. Required resolutions (1 bit) are approximately 1 mrad ( $\pm 15$  mrad full scale) for  $\Delta\theta$ , 20 mrad for  $\phi$  and 0.005 in pseudo-rapidity  $\eta$ .  $\Delta\theta$  is stored as a 1's-complement signed value, with "-0" (0x1F) indicating that there is no segment present in that segment field. The  $\phi$  resolution flag when set, indicates that the segment is from the sTGC whose  $\phi$  resolution is one bit less than the MM. In this case, the low bit of the  $\phi$  index is zero. The low resolution flag is set when only one quadruplet had a 3-out-of-4 coincidence and therefore  $\Delta\theta$  has poor resolution. If data only from bunch crossings that are accepted by the Level-1 trigger are recorded, we cannot know if missing triggers are due to improper functioning of the NSW trigger. The Monitor flag is set when this bunch crossing has been sampled as "interesting" by the NSW Trigger, e.g. a 3-out-of-4 coincidence. When the Monitor flag is set, the Sector Logic forwards it to the MUCTPI and similarly the MUCTPI to the CTP. The CTP generates a Level-1 Accept with trigger type "NSWMON trigger". This allows a monitoring process to see all the NSW front end, Sector Logic and MUCTPI data for that bunch crossing. The CTP ensures that the rate of these events does not exceed a few 10's of Hertz.

Table 9: Format of the 24-bit track segment sent to the Sector Logic ("Monitor" is the *hi-order* bit.)

Field:	Monitor	Spare	Low res.	$\phi$ resol.	$\Delta\theta$ (mrad)	$\phi$ index	$R$ index
Num of bits:	1	2	1	1	5	6	8

Up to eight track segments are transmitted to the Sector Logic in each LHC bunch crossing via two optical fibres, four candidates per fiber. Table 10 shows the format of the transmitted data. 8b/10b encoding is used. The segment frame is comprised of two comma codes for frame alignment, up to four track segments (three bytes of data each), the NSW Sector ID (4 bits) and a BCID number (12 bits). A missing segment is indicated by a  $\Delta\theta$  of "-0" (0b10000). Note that the End cap is not identified. The 16 bytes, 160 bits after 8b/10b encoding, are sent in a single bunch crossing at 6.4 Gb/s. The segments found by MM and sTGC must be merged into a single list. Duplicates (i.e. equal  $\Delta\theta$ ,  $\phi$  index and  $R$  index) must be removed, giving priority to the higher quality segment.

Table 10: Data format from the NSW Trigger Processor to the Sector Logic. The data is transmitted with 8b/10b encoding in one bunch crossing (16 bytes at 6.4 Gb/s). The format of the 24-bit segment is shown in Table 9. The comma character (K28.5) is an 8b/10b out-of-band control code for symbol alignment and packet framing. The NSW Sector ID is the *low-order* 4 bits and the BCID is 12 bits.



### 3.1 Transfer of track candidate segments from MM FPGA to sTGC

Figure 17, shows a block diagram of how the track candidate segments are sent from the MM FPGA to the sTGC FPGA where they will be merged. The process starts by collecting the segments from the 16 algorithm regions. Each region can generate eight segments every bunch crossing. A priority encoder will be used to select eight total segments each bunch crossing from all regions. The selected segments will be transferred to the sTGC FPGA using 16 640 Mb/s LVDS signals. The receiver in the sTGC uses independently delayed versions of the same signal to set the sampling point to the center of the data eye and provide segment data to the sTGC logic that will be aligned to the sTGC clock with no additional clock domain crossing necessary.

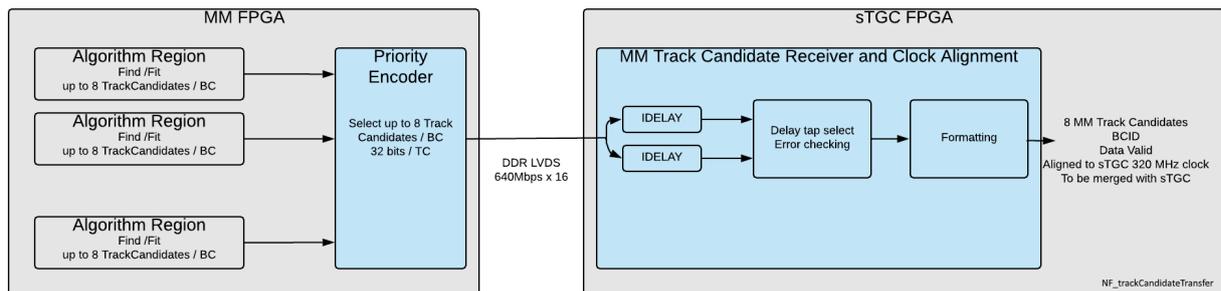


Figure 17: Transfer of MM track segments from the MM FPGA to the sTGC FPGA on the HORX mezzanine card

### 3.2 Removal of “duplicate” segments

Figures 18, 19 illustrate the two step process of removing duplicate segments so that at most eight segments will be passed to the Sector Logic from at most four sTGC and at most eight MM segments. The first step is to mark all “duplicate” segments so that only one is potentially chosen by the second step. Duplicates are segments that have equal  $\Delta\theta$ ,  $\phi$  index and R index. The fuzziness of the match, i.e. the number of significant bits to be compared, is configurable and will be determined from experience. Optionally, the  $\phi$ -index of the MM, which has better  $\phi$  resolution, can be taken instead of from the sTGC. In addition, lower quality segments from four-layer versus eight-layer coincidences will be disfavoured.

The second step is to gather up to eight segments to the output registers. This is done by using the vector of 12 chosen bits to address a look-up table whose outputs control multiplexers that route the segment data to the eight output segment registers. The look-up-table was prepared by software that chooses, for each of the  $2^{12}$  possible vectors of chosen valid segments, which of the chosen segments (or null segment) is to be routed to each of the eight output segment registers. If there are more than eight segments after duplicate removal, it is currently undefined which eight will be chosen.

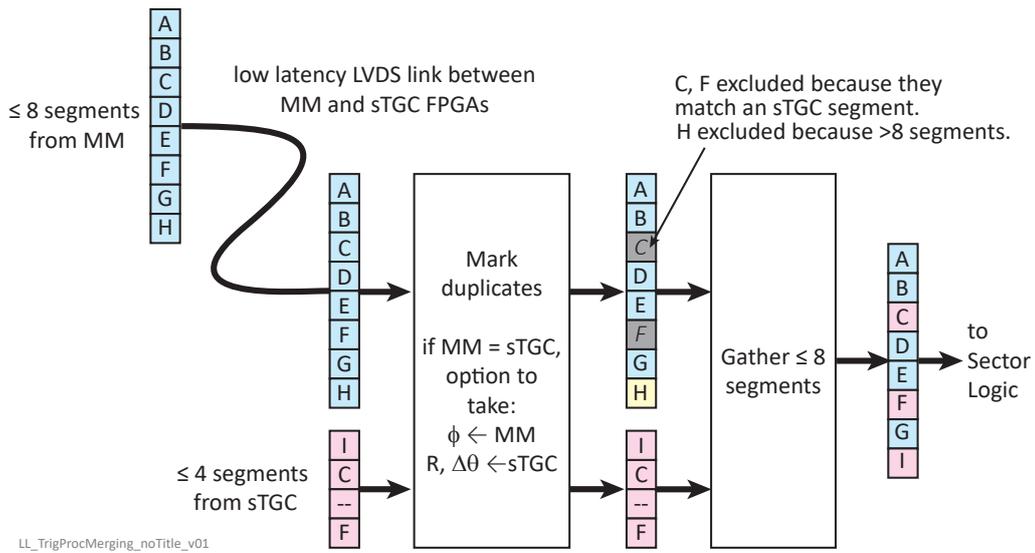


Figure 18: Merging of MM and sTGC segments for the Sector Logic by removal of duplicate segments

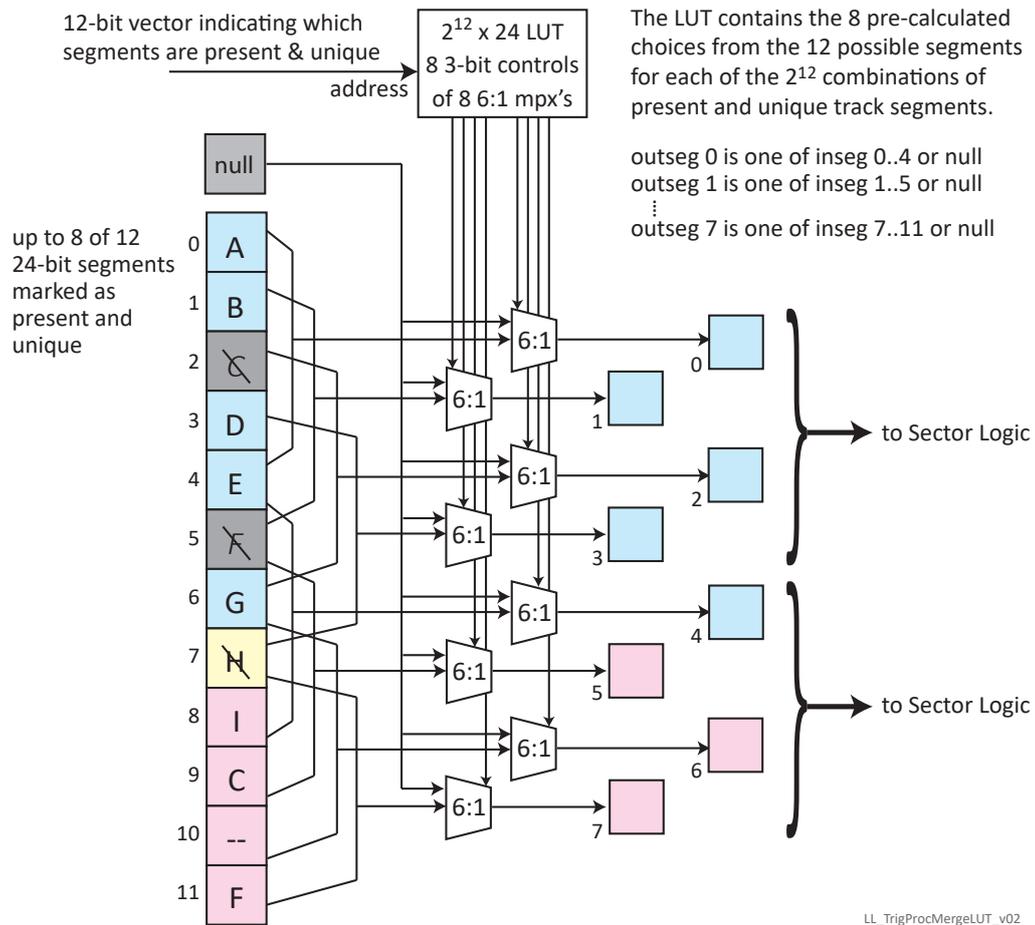


Figure 19: Gathering the chosen segments for output to the Sector Logic

### 479 3.3 Fixed latency output to the Sector Logic

480 The sTGC Trigger Processor runs with clocks that are multiples of the BC clock recovered with fixed  
 481 latency by the carrier FPGA. The fixed latency version of the CERN GBT-FPGA firmware [9, 10] is used.  
 482 The 320 MHz FPGA design clock is edge-aligned to the 40 MHz recovered BC clock by an MMCM. A  
 483 zero-delay jitter cleaner guarantees that the output serializer reference clock is also edge-aligned to the BC  
 484 clock.

485 The start of the output from the GTH serializer transmitting to the Sector Logic that was initiated from the  
 486 320 MHz design domain is therefore guaranteed to have fixed phase with respect to the 320 MHz design  
 487 clock. It remains to start at a defined one of the eight 320 MHz clock edges in a BC to guarantee fixed  
 488 phase with respect to the BC clock.

489 Although the Trigger Processor algorithm pipeline has fixed latency, its start time is not fixed across power  
 490 cycles and resets. The delay between the arrival of a bit to the deserializer input and the write to the parallel  
 491 deserializer output userbus varies by  $\pm 1$  receiver userbus clock between power cycles and resets. Since the  
 492 pipeline starts as soon as all the input data is ready (after configurable delays compensate for differences in  
 493 fibre length), there will therefore be a variation (between power cycles) in the time the output segments are  
 494 ready to transmit to the Sector Logic.

495 Fixed latency is attained by measuring, over several power cycles of the whole trigger path, the latest  
 496 320 MHz phase within a BC at which the segment data is ready and saving that value as a configuration  
 497 parameter,  $N$ . The output to the Sector Logic is then always begun at the 320 MHz clock with phase  
 498  $N+1$  (modulo 8), that is after the longest possible variation in the latency. The phase counter is reset by  
 499 BCR. See Figure 20.

500 In this way, the latency from collision at the interaction point to the start of output to the Sector Logic is  
 501 fixed across power cycle and resets.

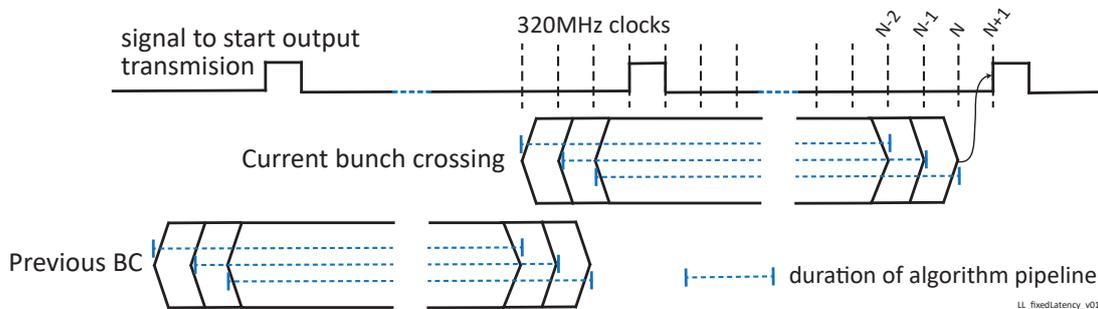


Figure 20: Timing diagram showing the (fixed) duration of the algorithm pipeline for three different power-ups (dotted lines). To guarantee fixed latency of the output, transmission to the Sector Logic is begun one clock after the latest possible completion of the pipeline. “ $N$ ”, between 0 and 7, is one of the eight phase 320 MHz phases of the BC clock at which the latest completion of the pipeline occurs. “ $N$ ” is found by observing several power cycles of the complete trigger path.

502 The fixed latency across power cycles of the Trigger Processor of the output transmission has been verified  
 503 by measurements with an oscilloscope. The time delay between the edge of BCR from the TTC module  
 504 (used as a stand-in for the collision) and a signal when the segment data was sent to the output serializer  
 505 was measured and found not to change more than 1 ns after power cycles and resets. A test point on the  
 506 Trigger Processor PCB was used to extract this signal.

## 4 Configuration, control and status of algorithm and operating parameters

Table 11 shows the status and control registers common to both sTGC and MM, specific to the sTGC and specific to the MM.

For configuration and status, there is no logical distinction between the Trigger Processor and the Front end devices which are configured via FELIX and the SCA ASIC. For the back-end software, uniformity with the SCA path is desirable. The Trigger Processor already requires a FELIX connection for readout on Level-1 Accept. These considerations point to using the SCA protocol between the NSW configuration process and the Trigger Processor FPGA's.

The SCA protocol uses an OPC client to interface to both the DCS and the NSW configuration process, and an OPC server to communicate to FELIX. On the Front end boards, an SCA is connected to an E-link of the same GBTx ASIC that handles the TTC and Level-1 Accept readout E-links. The Trigger Processor implements firmware using the SCA protocol so that its configuration register banks emulate SCA I2C devices. Each of up to 16 I2C masters can address up to 1024 32-bit registers. The SCA protocol is a very simple command-response, stateless protocol. Figure 21 shows the data formats and the communication path from the NSW configuration process to the status and control blocks of each FPGA on each mezzanine card of the Trigger Processors. The implementation is called SCA eXtension (SCAX) and uses 8b/10b encoding instead of the HDLC bit-stuffing used by the SCA ASIC. Detailed documentation can be found at [12–14].

Each of the 16 SCAX I2C masters accesses the FPGA logic's registers via an associated register file, which acts as a demultiplexer for write commands, or as a multiplexer for read commands. Upon a user's write request, the corresponding I2C master drives its register file's demultiplexer with the register address and propagates the data-to-be-written to the desired register. Upon a read request, the register address is used to drive the corresponding register contents to the I2C master. The data that were read are then forwarded to the back-end software.

Important extensions are: 1) A register address can point to a RAM and the FPGA generates read or write enables and automatically increments the RAM address on every read or write; 2) A register address can point to a FIFO and the FPGA generates read or write enables on every read or write. This is used to configure look-up tables and to read out debug info captured in a FIFO.

A high-level user interface has been developed called `NSWSCAXRegisters` which creates vhd, C++ header files, xml for OPC UA, JSON for the configuration software, and documentation tables. The required registers and their properties are described in a CSV format. Scripts then build a YAML database of these registers and the package `wuppercodegen` is used in conjunction with a series of code templates written using the JINJA2 standard. In addition, a Continuous Integration setup runs such that user interaction is confined to CSV updates, and running of the scripts and production of the various sources are automatically done for the user. This integration also ensures full versioning of the register configuration.

Table 11: Status and control registers. For a complete list, see the CSV files [11]. The NSW-CAXRegisters package generates vhd, C++ header files, xml for OPC UA, JSON for configuration, and documentation based on the CSV input.

<b>Status registers</b>	<b>Control registers</b>
FELIX RX link OK	Endcap/Sector ID
Jitter cleaner status (2)	$\Delta\theta$ cut
Firmware version / time stamp	L1A BC window
	Z (IP to layer distance)
	1/Z
	monitoring BC window
	BCID offset
	Level-1 ID
	XVC IP low byte
	HORX PLL reset
<hr/>	
<b>sTGC status registers</b>	<b>sTGC control registers</b>
Input QPLL lock (3)	enable MM segments
Input CPLL lock (36)	enable output to Sector Logic
input link OK (36)	max cluster size
SL output CPLL lock (14)	min cluster size
	charge threshold for determining cluster size
	charge threshold for determining centroid
	strip direction (+/-) list
	radial offset (per gas gap)
	params for alignment corrections
	for each 1/4 pad in each layer of a sector
	which of the 8 320 MHz phases of BC clock to start output to SL
<hr/>	
<b>MM status registers</b>	<b>MM control registers</b>
GBT link status	MM strip zero offset
PLL lock status	List of inverted MMFE8s
GBT packet alignment error	Cotangent $1.5^\circ$
FIFO overflow errors	Cosecant $1.5^\circ$
Communication errors	Hit integration BC window width
	Physical arrangement of x u v planes
	Cartesian ROI map
	parameters for $M_X^l$ calculation
	parameters for $\Delta\theta$ calculation
	List of connected ADDCs
	GBT link reset
	GBT PLL reset
	GBT pkt alignment control
	Diagnostic FIFO enable
	ART data enable

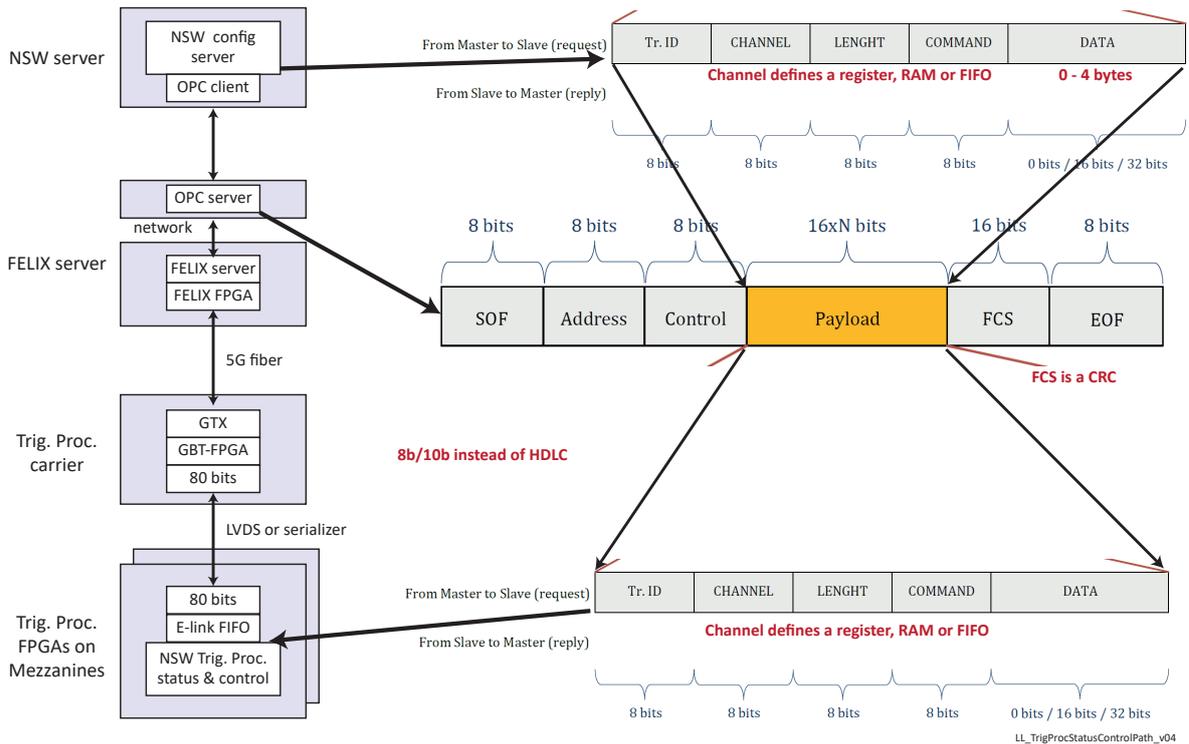


Figure 21: Configuration and control path using the SCA OPC-UA protocol.

## 5 Latency estimate

542

543 The latency budget for the NSW from bunch crossing to input of the segment information to the Sector  
 544 Logic is 1075 ns. See [15] for a detailed breakdown of the latencies along sTGC trigger path and fiber  
 545 routing. The latency of the current design of the NSW trigger path leaves little contingency. It is therefore  
 546 imperative to keep the latency as low as possible at every step of the chain, including in the Trigger  
 547 Processor. The Trigger Processor latency estimation is given in Table 12 for both detector technologies.  
 548 The accounting includes the input/output serializers, the trigger algorithm, the algorithm that merges the  
 549 MM and sTGC segments and the transmission time to the Sector Logic. The Trigger Processor start  
 550 times are based on individual latencies of the upstream elements which are a mixture of measurements,  
 551 simulations and estimates. Work is ongoing to improve the estimates and some measurements will be made  
 552 in the NSW Vertical Slice. Consequently the total latency from bunch crossing to input to the Sector Logic  
 553 is still uncertain. Since the sTGC is the critical path, segment merging is planned to be done in the sTGC  
 554 FPGA.

555 **Micromegas Trigger Processor latency** The latency from VMM output to Trigger Processor track  
 556 segment coincidence trigger (finder) have been measured during a hardware integration workshop. Timing  
 557 measurements for the Micromegas track segment fitter components are taken from behavioral simulations.  
 558 The total algorithm latency from GBT fiber input to segments arriving in the sTGC is 267 ns. This estimate  
 559 assumes a three bunch crossing hit integration window. Recent studies suggest the hit integration window  
 560 will need to be increased to four or more bunch crossings. Each additional bunch crossing will increase the  
 561 latency by 25 ns.

Table 12: Trigger Processor latency for the MM and sTGC streams. The time required to merge the segments from the two streams is assigned only to the sTGC, since the merging is done in the sTGC FPGA.

	sTGC (ns)	MM (ns)	
Start time	795	637	Arrival of input data
Input deserializer (Rx)	77	44	
Clock alignment & distribution	31	44	
BC integration window		75	3 BCs for MM
Trigger algorithm	56	25	320 MHz clock
MM Track Segment Fitter		41	dTheta calculation critical path
Format & Tx MM segments to sTGC		38	up to 8 pipelined segments
End	959	904	MM finishes first
Segment merging algorithm driving output serializer	13		Assigned to sTGC
Output to Sector Logic serializer (Tx only)	21		to 1st bit out (Deserialization is in Sector Logic latency budget)
Fiber to Sector Logic	25		5 m fiber @ 5 ns/m
End	1018		

562 **sTGC Trigger Processor latency** The sTGC Trigger Processor latency has not yet been measured. This  
563 section will be updated with results from the simulation.

## 564 6 Monitoring

565 Monitoring enables validation of the data processing in the Trigger Processor by providing data for  
566 debugging and fast online analysis. Monitoring covers L1A events, a sampling of bunch crossings of  
567 interest, as well as additional dedicated monitoring data to validate the trigger algorithms and data flow.  
568 The monitoring data is intended to identify algorithm failures that may arise from occurrences such as:

- 569 • Misalignments
- 570 • Fake events from backgrounds and noise
- 571 • True events that are missed due a range of effects including non-functional detector elements or  
572 timing errors

573 Figure 22 shows a block diagram of the NSW Trigger Processor monitoring path. While the details of the  
574 data used for monitoring depend on the specific detector technology, sTGC or MM, the common aspects  
575 and infrastructure will be described first here.

576 Monitoring data is buffered on the Trigger Processor pending a Send Monitoring Data request for monitoring,  
577 which could correspond to a L1A, a random signal or an internally generated signal. The monitoring data  
578 is sent from the Mezzanine cards through the carrier card FPGAs to the Rear Transition Module (RTM)  
579 behind the carrier card. The baseline, as shown in Figure 22, is to buffer on the Mezzanine card FPGAs,

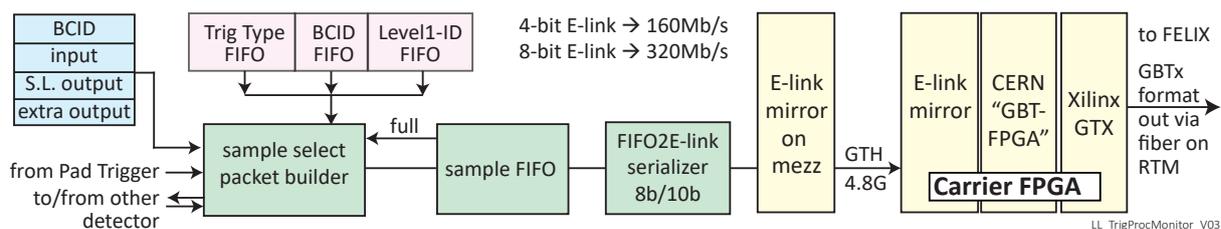


Figure 22: Block diagram of the NSW Trigger Processor monitoring output. The “from Pad Trigger” and “to/from other detector” signals enable choosing the same BCID to be monitored by all three components.

580 though an alternative is to buffer on the carrier card FPGAs (for example to reduce resources used in the  
581 Mezzanine FPGAs).

582 Complete data from a sample of bunch crossings are collected for sending to a monitoring process. Example  
583 selection criteria are: a random BC, a BC with at least one segment found, a BC with segments outside the  
584  $\Delta\theta$  cut, a Level-1 accept BC, etc. The data buffered includes all the input data and the output segment data  
585 that is sent to the Sector Logic for a configurable window around that BC.

586 Example diagnostics that the monitoring can provide include: the number and distribution in  $R - \phi$  of  
587 bunch-crossings that have candidates but are not accepted by Level-1, multiplicity and distribution of input  
588 VMM ART signals and sTGC bands per BC, multiplicity and distribution of slope coincidences,  $\theta$ ,  $\phi$  and  
589  $\Delta\theta$  distributions, types of track candidates (number of coincidence layers, etc.)

590 In addition to sending data to a monitoring process, the SCAX system allows readout of counters and  
591 statistics accumulated in the FPGA that can be transferred to the DAQ Information System (IS). For  
592 example: the number of segments accepted and rejected per bunch-crossing, the count of link transmission  
593 errors and the volume of monitor data.

594 The monitoring data streams are transferred via E-links on GBT links connected to FELIX as part of  
595 the ATLAS TDAQ system. The GBTx ASIC is emulated by the CERN GBT-FPGA firmware and the  
596 Xilinx GTH serializer/deserializer. Interfacing the data streams to the GBT protocol as E-links is the  
597 baseline choice due to its simplicity and because this link already is required for the Level-1 Accept data.  
598 The monitoring data is then transferred to FELIX for communication via Ethernet to processes on PCs.  
599 FELIX routes the various logical data streams to the appropriate network end-points, including Monitor  
600 processes.

601 Two LVDS lines between the sTGC and MM FPGAs exchange the Send Monitoring Data signal across  
602 the two technologies so that the same BCID from both detectors can be monitored. The sTGC Trigger  
603 Processor also receives requests to monitor signals from the Pad Trigger based on 3-out-of-4 coincidences  
604 in a pad tower.

605 Monitoring data is buffered at several stages:

- 606 • The monitoring data is first stored in a circular buffer that continuously keeps a moving window  
607 of data samples (typically comprised of the raw hit information and the corresponding trigger  
608 information available for each BC). There is a dedicated buffer for each readout element (typically  
609 strip or pad, or layer). When a Send Monitor Data signal is received, a contiguous block of data  
610 samples for the BC range desired (typically 5 to 10, depending on the type of monitoring data that is  
611 being sent).

- 612 • The monitoring data is then packaged into a second stage FIFO in preparation for transmission.
- 613 • Once the second stage FIFO has received all the data following the Send Monitoring Data signal,
- 614 the results are merged, the header information included and data are sent out sequentially in E-link
- 615 packets.

## 616 6.1 Monitoring Data

617 The data available to be included in the MM monitoring data stream is collected at various points in  
 618 the algorithm pipeline. Data that is included in the monitoring stream will be selected as needed due  
 619 to the throughput limitations. Monitoring data can contain two general types of data. This is shown in  
 620 Figure 23.

- 621 1. ART GBT data packet as sent from the ADDC. Only a small subset of the total ART data being  
 622 received can be selected to be included in the monitoring data. ART Data, is selectable at the fiber  
 623 level. All non-zero GBT packets on the enabled fibers will be output through the monitoring path.
- 624 2. Coincidence Trigger Data. The Coincidence trigger data will contain ART hit data that forms a  
 625 coincidence trigger along with algorithm debug values and track fit results.

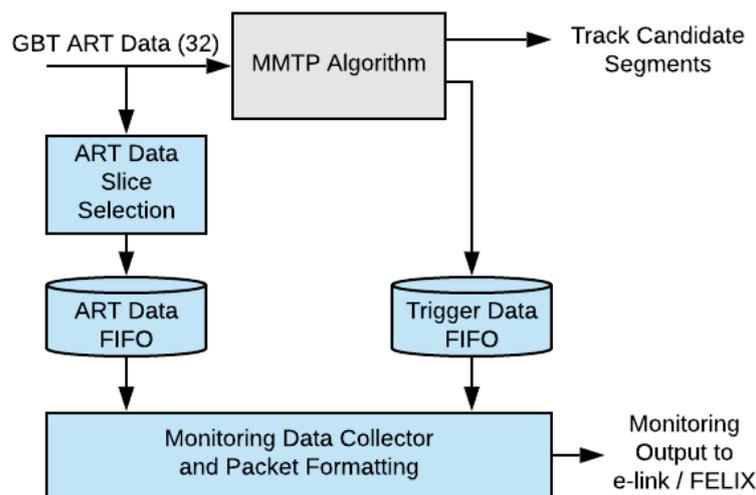


Figure 23: Block diagram showing the monitoring data flow for the Micromegas algorithm.

626 Figure 24 shows the current debug monitoring interfaces for the sTGC. The output is via E-links to FELIX  
 627 using the FPGA-GBT framework shown in Figure 22.

## 628 6.2 Correlated MM– sTGC monitoring

629 In order to perform correlated MM – sTGC monitoring, each Trigger Processor can signal the other that the  
 630 current bunch crossing should be monitored by both. In this case both will send monitoring data via FELIX  
 631 to the NSW monitor processor. The Pad Trigger can also signal the sTGC Trigger Processor to monitor the



657 **Sector-ID register:** The sector id of this instance of the Trigger Processor.

658 **Output bandwidth** The bandwidth needed is expected to be larger than that provided by a single 320 Mb/s  
659 E-link, partly because more than three bunch crossings are likely to be read out for MM. The software  
660 ROD supports transmitting L1A events in round-robin over several E-links. Ten such E-links can be  
661 carried by a GBT normal-mode fibre, of which at least eight would be available for readout of L1A  
662 data.

663 For sTGC, the case with zero-suppression of the input by only a factor of two and no suppression  
664 of the output requires  $1/2 \times 32 \text{ inputs} \times 104 \text{ bits} = 208 \text{ bytes}$  for the input data and  $8 \text{ segments} \times$   
665  $16 \text{ bytes} = 128 \text{ bytes}$  for the output data, for a total of 336 bytes per Level-1 Accept. At 100 kHz and  
666 including 8b/10b encoding, two 320 Mb/s E-links are sufficient.

## 667 7.1 TTC data

668 The eight TTC bits that are received on every bunch crossing are:

669 L1A, L0A, BCR, ECR, ECOR, test\_pulse, soft\_rst, SCA\_rst

670 L0A, ECOR are required only in a two-level trigger. In a single level trigger, L0A is sent along with L1A.  
671 The current Trigger Processor firmware does not support a two-level trigger. Soft reset is used to flush all  
672 data in the data path, synchronously over the whole NSW data path.

## 673 7.2 Level-1 Accept Algorithm

674 The Level-1 Accept algorithm concept may be described as a matrix, as shown in Figure 25. Each row  
675 except the first one handles one data stream and each column handles an L1A. The first row stores the  
676 header information that is common to all processing cells in a column. Note each data stream may contain  
677 different types of information (raw data, segments, etc.).

678 L1A requests are buffered until a processing column is available. Once all the data has been gathered for  
679 one L1A request, the information is sent to the packet builder module and the column is released for use by  
680 another L1A. The packet builder sweeps the columns for data to be sent and the data is sent in the same  
681 order that the L1A were received.

## 682 7.3 Level-1 Accept Implementation

683 Figure 26 shows a block diagram for the data path for the Level-1 Accept data. The input data is buffered  
684 for the duration of the trigger latency (latency pipeline). Once the L1A signal is issued, the data associated  
685 to the L1A (e.g. BCID, Orbit count, Level-1 ID) is collected (derandomizer) in a processing column. Once  
686 all the data associated to the L1A is collected, the L1A packet builder retrieves all the data in the desired  
687 format and sends it to FELIX. The Level-1 Accept implementation is located on the mezzanine FPGA's.

688 The clock used is 320 MHz ( $8 \times$  BC clock). The data input interface is a variable length payload associated  
689 with a BCID and a write strobe to validate the payload. Each processing row has a data input interface.  
690 The payload length is a multiple of 16 bits in order to interface to the FELIX E-link module.

691 In this implementation, the data flows continuously from the algorithms to the processing cells, where the  
692 data is matched to the L1A BCID. The matching is performed in a configurable (and possibly asymmetric)

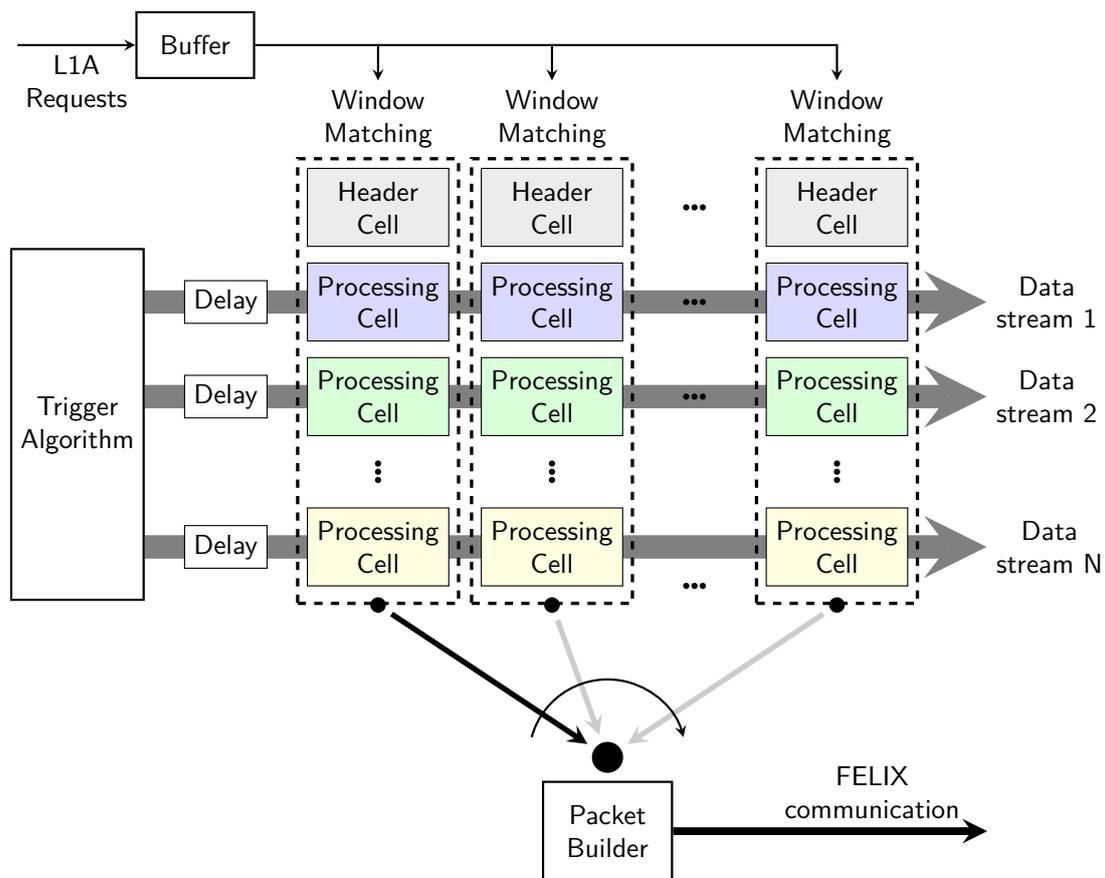


Figure 25: Conceptual design of the Level-1 Accept algorithm

693 window around the BCID of interest. The data formats for the data input to the L1A firmware are defined  
 694 below. The input buffers are connected to all concurrent processing cells.

695 The L1A Manager receives a single boolean signal indicating a L1A for a BC and calculates a range of  
 696 BCID's associated to the L1A. The range is defined according to a configurable offset. The L1A manager  
 697 then configures and activates the first processing column that is available with the BCID of interest. Note  
 698 that multiple columns can be active simultaneously. The L1A manager watches for the end of the process  
 699 and instructs the L1A packet builder to drain the corresponding column.

700 Each processing cell stores the data in its FIFO. The data format used is the same as the input information, with  
 701 the exception of headers that are added for subsequent packet handling downstream. This implementation  
 702 allows for flexible input data streams, accommodating both sTGC and MM, as well as different data types  
 703 (raw data, segments, etc.). The header is included by the header cells and is common for all packet types.

704 The output of the L1A packet builder is a "FIFO to E-link" module provided by FELIX. The L1A packets  
 705 are sent via FELIX to the swROD and monitoring processes for further processing.

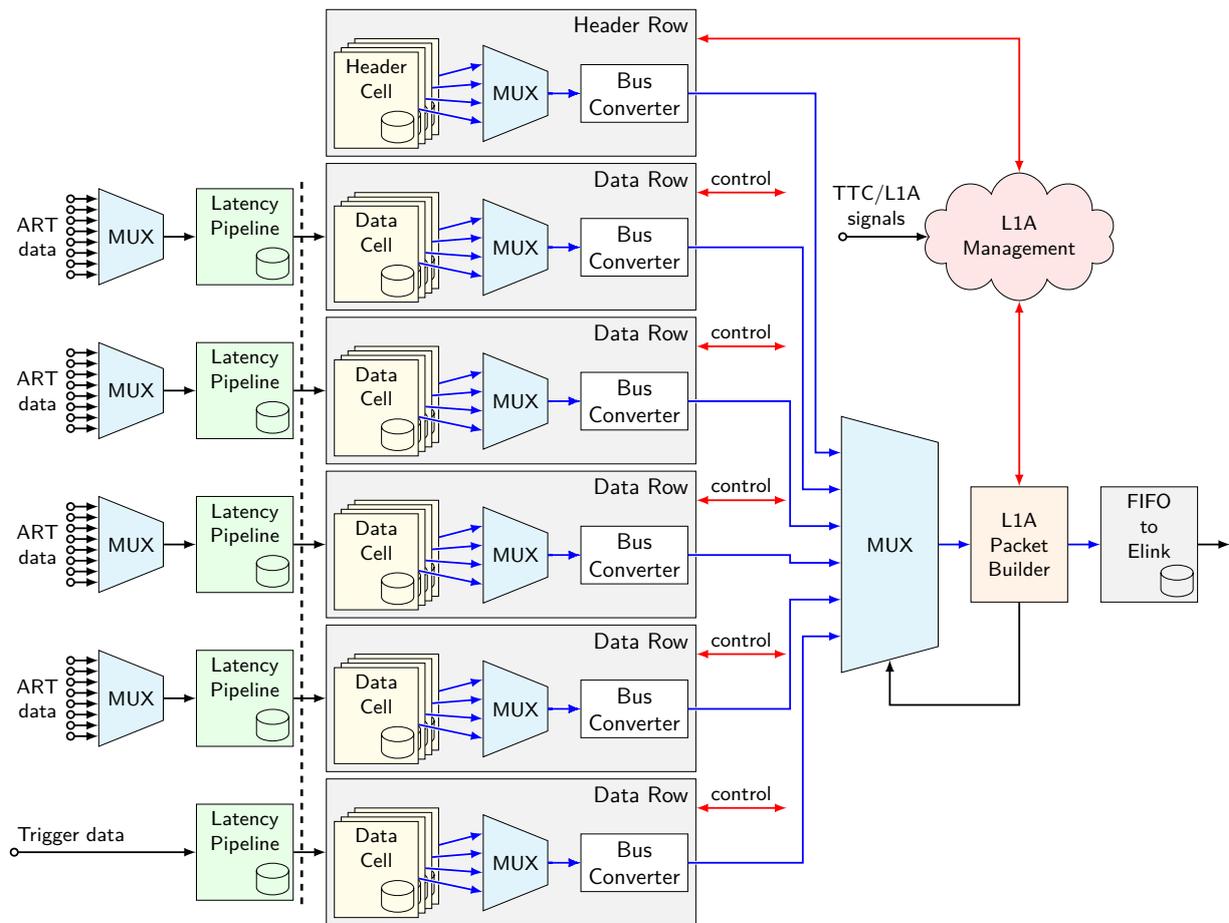


Figure 26: Level-1 Accept data flow

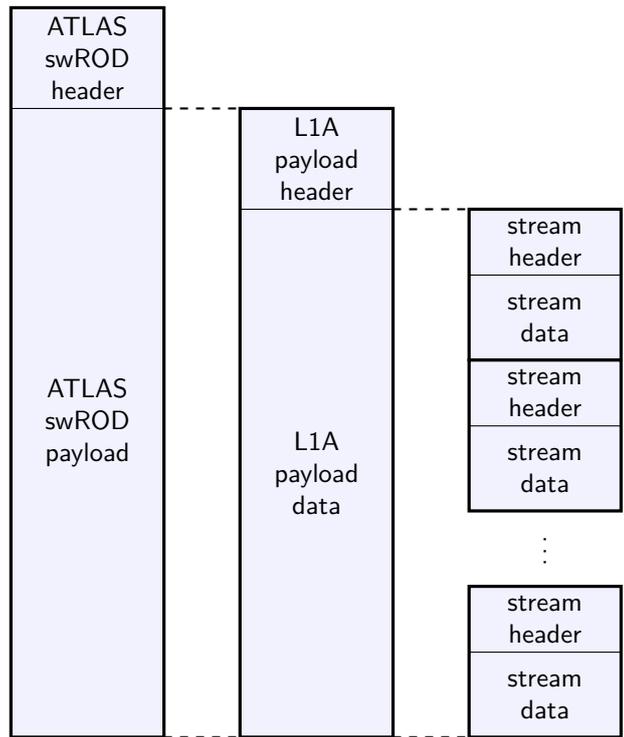
## 706 7.4 Read out data format

707 The format of the Trigger Processor data fragments sent to the swROD include a packet header followed by  
 708 the specific information provided by the corresponding trigger algorithm (sTGC or MM, raw data, trigger  
 709 segments, etc.). The overall format is shown in Figure 27 and the specific header is shown in Figure 28.  
 710 The data formats listed below for sTGC and MM correspond to the detector-specific data streams included  
 711 in the full packet.

### 712 7.4.1 Micromegas read out data

713 For the MM, three kinds of information are foreseen to be output by the L1A processor: raw data, hit data  
 714 and fitter data. The data formats are: (See also Figure 29.)

- 715 • *MM Raw data*: comprised of 256-bit packets sent over 16 links  $\rightarrow$  16 links  $\times$  1 packet/BC corresponds  
 716 to 4k bits. If the data extended over a  $\pm 8$  BC window, this corresponds to 64k bits. The proposed  
 717 implementation uses a 512-bit bus with the 320 MHz clock, such that one BC of raw data can be  
 718 read per BC.



TCP\_NSWTP\_L1A\_data\_format\_V01.pdf

Figure 27: Level-1 Accept data format

Standard Header for input to the swROD and monitor																															
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
spare		orbit		BCID								flags[6..0]				EC	sector id				frag id										
Level-1 ID																															
<b>Flags for Trigger Processor output to swROD</b>												MM VMM hit data						frag ID = 0													
6	5	4	3	2	1	0						sTGC VMM hit data						frag ID = 1													
Mon						#segments				Pad Trigger input data						frag ID = 2															
										Pad Trigger output data						frag ID = 3															
Mon:		Monitoring of this BC was forced																MM trigger processor input data		frag ID = 4											
										MM trigger processor segment data						frag ID = 5															
										sTGC trigger processor input data						frag ID = 6															
										sTGC trigger processor segment data						frag ID = 7															
										output data to the Sector Logic						frag ID = 8															

Figure 28: Level-1 Accept data header

719  
720  
721  
722

- *MM Hit data* comprised of 416-bit packets. It is potentially larger than raw data if there are multiple packets per BC for a given region. The proposed implementation uses a 576-bit bus with the 320 MHz clock, such that one BC of hit data can be read per BC.
- *MM Fitter data*: comprised of 416-bit packets. This data is only generated in case of a coincidence.

Figure 29: Format of the MM data sent for a Level-1 Accept to the swROD, excluding raw data

	<b>name</b>	<b>bits</b>	<b>description</b>
Segment 0			
	... Fitter output		
	ROI	16	
	DTheta	16	
	mx	16	
	MX local	16	
	MV Global	16	
	MU Global	16	
	my/MX Global	16	
	Event Counter	16	
		16	bytes / segment header
	... for 8 layers		
	Hit information on track	16	
		16	bytes / 8 layers/ segment
	for one segment	32	bytes / segment
... for up to 8 segments		256	bytes for max 8 segments

723 **7.4.2 sTGC read out data**

724 The format of the sTGC Trigger Processor output segment data fragment sent to the swROD is shown in  
 725 Figure 30. Also output is a fragment containing the input data with format shown in Table 4.

Figure 30: Format of the sTGC segment data sent for a Level-1 Accept to the swROD

	<b>name</b>	<b>bits</b>	<b>description</b>
Segment 0	band-id	8	
	Phi-id	8	
	Delta theta	8	
	R index	8	
	map of valid layers	8	
	number of valid layers	8	
	spare	16	
		8	bytes / segment header
	... for 8 layers		
	layer Centroid	16	
	layer num hits in centroid	4	
	layer error flags	12	why layer was rejected
		32	bytes / 8 layers/ segment
	for one segment	40	bytes / segment
... for up to 4 segments		160	bytes for max 4 segments

726 **7.4.3 Read out bandwidths**

727 **Micromegas Level-1 Accept data payload** The L1A data packet consists of RAW ADDC hit data and  
 728 coincidence triggered segment data. The ADDC data is in the form of raw un-decoded ADDC GBT packet,  
 729 rather than the individual hits. This means that each GBT packet containing one or more hits will contribute  
 730 96 bits to the L1A packet. Packets with no hits are suppressed.

731 Using an average strip hit rate of 8 kHz for  $\mathcal{L} = 2.8 \times 10^{34}$  (from a plot produced by Alex Tuna at  $\mathcal{L} = 7 \times 10^{34}$   
 732 giving 20 kHz), the probability of an ART GBT packet containing at least one hit is

733 
$$\lambda = 8 \text{ kHz} \times 2048 \text{ strips/GBT} \times 25 \text{ ns} = 0.41$$
  
 734 
$$1 - e^{-\lambda} = 34\%$$

735 The ART data contribution for a 5 BC window for 32 ART fibers would be

736 
$$34\% \times 96 \text{ bits/BC} \times 5 \text{ BC's} \times 32 \text{ ART's} = 5.2 \text{ Kb}$$

737 For the coincidence triggered segment data, we currently collect 304 bits. Assuming we have five segments  
 738 in a L1A window this will contribute 1.5 Kb. The total data size is then 6.7 Kb; a L1A rate of 100 kHz yields  
 739 a required data rate capacity of 670 Mb/s. Six E-links, each providing 250 Mb/s after 8b/10b encoding, are  
 740 available, providing a total of 1500 Mb/s. Of the 96 bits in each ART GBT packet, many will be zeroes.  
 741 Data compression should be able to significantly reduce the bandwidth.

742 **sTGC Level-1 Accept data payload** For  $\mathcal{L} = 5 \times 10^{34}$ , the Table 13 shows that the expected number of  
 743 32 Mb/s E-links needed for the sTGC Trigger Processor is 1.36. Scaling to the expected Run 3 luminosity  
 of  $3 \times 10^{34}$ , we require 80% of a 320 Mb/s E-link. Note that six such E-links are available.

Table 13: Expected number of 32 Mb/s E-links needed for the sTGC Trigger Processor

16	bits for header	
104	bits per input TDS	104 means no compression
8	layers	
1.33	average no of bands per BC. This is an exaggeration.	
3	BC's per L1A	
0.1	L1A rate, MHz	
348	Mb/s	(435 bytes/L1A)
435	Mb/s with 8b/10b	
1.36	# E-links	at $\mathcal{L} = 5 \times 10^{34}$

744

745 **8 Configuration, control and monitoring of the hardware via I2C**

746 Figure 31 shows a simplified diagram of the I2C buses available in the NSW TP mezzanine board. Each  
 747 FPGA is connected to related set of microPODs and clock and jitter cleaner devices, which allows not only  
 748 sensor reading, but also parameters configuration. There is also a bus connecting the local microcontroller

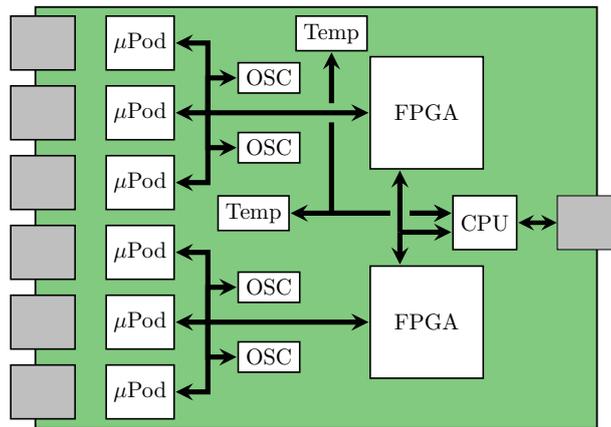


Figure 31: Simplified diagram for the I2C buses in the NSW TP mezzanine board.

749 to the FPGA devices, making read and write operations possible. There is another bus connecting the  
 750 microcontroller to discrete temperature sensors for environment sensing.

751 Figure 32 shows the path from the FPGA internal parameter table to the DCS. The MMC microcontroller  
 752 has a user I2C port connected to the FPGA's on the mezzanine. This port provides IPMI access to the FPGA  
 753 internal parameter table. The parameters include the FPGA parameters (temperature, internal voltage) and  
 754 the optical link status registers from the micropods (RX power, link-loss, TX fault, temperature, etc.). This  
 755 path can also be used to monitor application flags (Trigger Processor status, error flags, etc.).

756 Alternatively, instead of the I2C slave, the SCAX machinery can be used to read out the parameter memory  
 757 in Figure 32. This path goes to DCS via FELIX instead of via the IPMC and shelf manager.

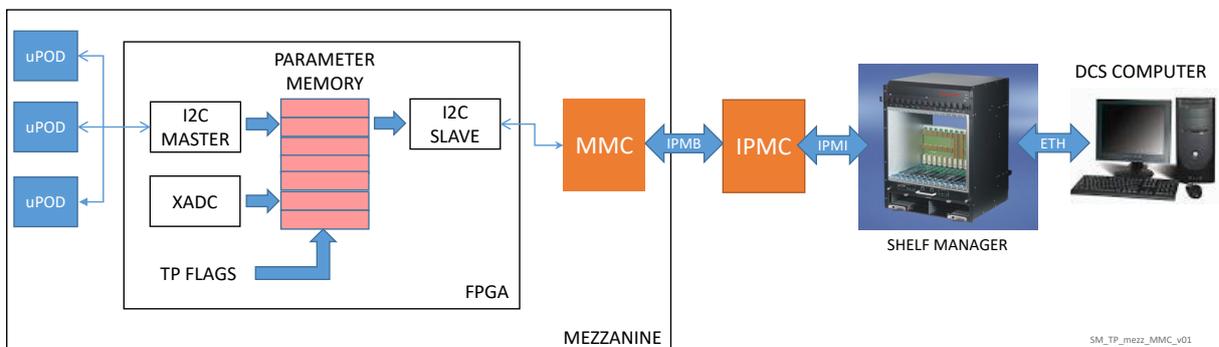


Figure 32: Simplified diagram for the access of the sensor parameters sent to the DCS.

## 758 8.1 IPMI monitoring

759 The management of the hardware temperatures is accomplished by the Intelligent Platform Management  
 760 Interface (IPMI) protocol, which changes the speed of the fans in the ATCA shelf according to sensor  
 761 readings and the configurable recommended ranges. Figure 33 shows some levels of managing resources.  
 762 The Shelf Manager (ShM) in the ATCA crate gathers information of the devices inside the crate and  
 763 makes decisions. The Intelligent Platform Management Controller (IPMC) is the standard interface for all

764 the possible sensors in the board and responds directly to the ShM. In addition, Modular Management  
 765 Controllers (MMC) are used for specific units and purposes (e.g. the Mezzanine card, the RTM, etc.). All  
 766 the communication ShM-IPMC and IPMC-MMC relies on the I2C protocol. The ranges and descriptions  
 767 of sensors are stored in the MMC or in the IPMC as Sensor Data Records (SDR) and are promptly available  
 768 upon system startup.

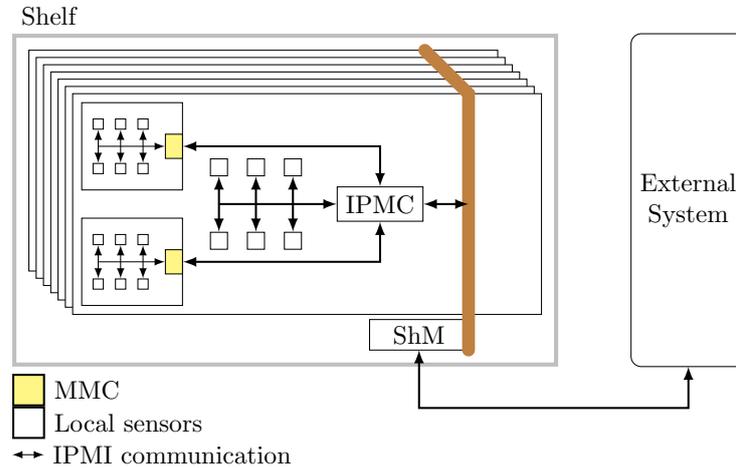


Figure 33: Simplified diagram for the I2C buses in the NSW TP mezzanine board.

769 Each NSW TP mezzanine board is equipped with an MMC, as shown in Figure 31 (labeled as “CPU”),  
 770 which is able to read information from sensors for temperature, from the micropods and FPGA’s. The IPMC  
 771 then collects the information from the MMC’s (mezzanine and RTM) and also for the many components  
 772 on the carrier board, like FPGA’s, SoC, power supply, etc, as detailed in the diagram in Figure 34. Only the  
 773 Zynq requires some programmable logic design for this communication but relies on IP Cores from Xilinx  
 774 and software drivers. None of the additional items involves FPGA designs and the monitoring happens  
 775 directly at the hardware level.

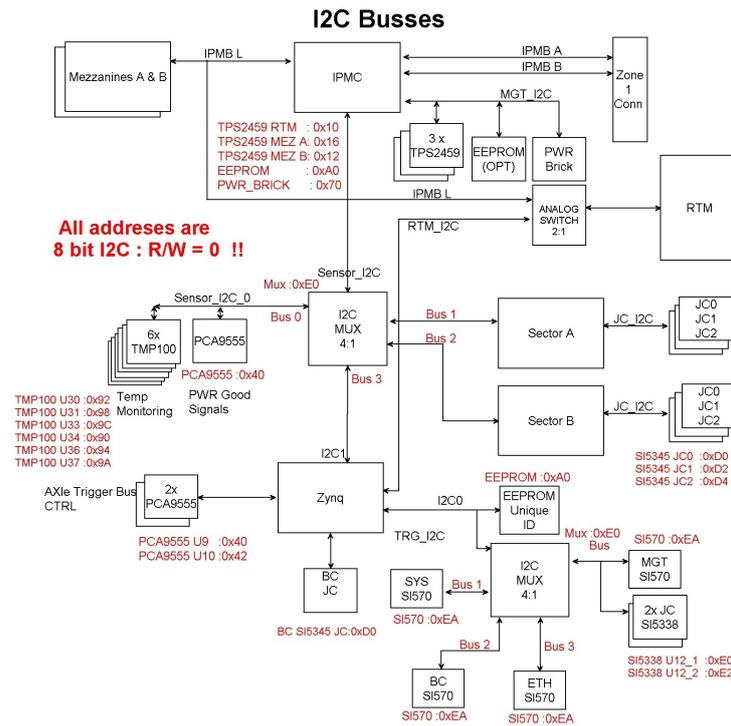


Figure 34: Detailed I2C map for the NSW Trigger Processor carrier board, involving access to mezzanines and RTM buses. Among other functions, the IPMC monitors FPGA's and SoC, transceivers, power supplies, and other sensors (temperature, current and voltages).

## 776 9 FPGA configuration

777 Each Trigger Processor ATCA carrier has seven FPGA's: two Virtex 7's on each mezzanine, two Ultrascale  
 778 Kintex and the Zynq SoC FPGA on the carrier. In order to program all the components on the board, a  
 779 scheme to expose the JTAG chains over the network is used. The Xilinx Virtual Cable (XVC) protocol  
 780 is used to translate between network protocol (TCP/IP) and JTAG signals. This solution, supported by  
 781 Xilinx development tools, provides access to the reconfigurable devices inside the crate, enabling remote  
 782 connections.

783 The NSW Trigger Processor carrier board includes a Zynq System on a Chip (SoC), comprising a Xilinx  
 784 FPGA and an ARM processor in the same chip. Figure 35 shows the implementation of the TCP/IP stack  
 785 and XVC service. The Zynq device is in the JTAG chain and provides the FPGA programming capability.  
 786 However, since the processor itself cannot be reprogrammed once it is running the XVC service, an external  
 787 flash and boot loader is used, which are directly controlled by the processor and outside the JTAG chain.

788 In this scheme, the Zynq SoC on the blade provides configuration resources for its own FPGA, as well  
 789 as for the JTAG chains in its mezzanines. It also provides network access using the fabric interface of  
 790 the backplane or via RTM, ensuring that all the configuration resources for every blade in the shelf are  
 791 available through network switches. If the Zynq SoC configuration is disrupted, it can be reprogrammed  
 792 through the XVC service provided by the CERN IPMC as an alternative path, also available via network  
 793 connection. As yet another recovery procedure, the carrier Zynq is also equipped with two flash memories.  
 794 A golden firmware is stored in the secondary one, which can be activated by the IPMC via shelf manager.

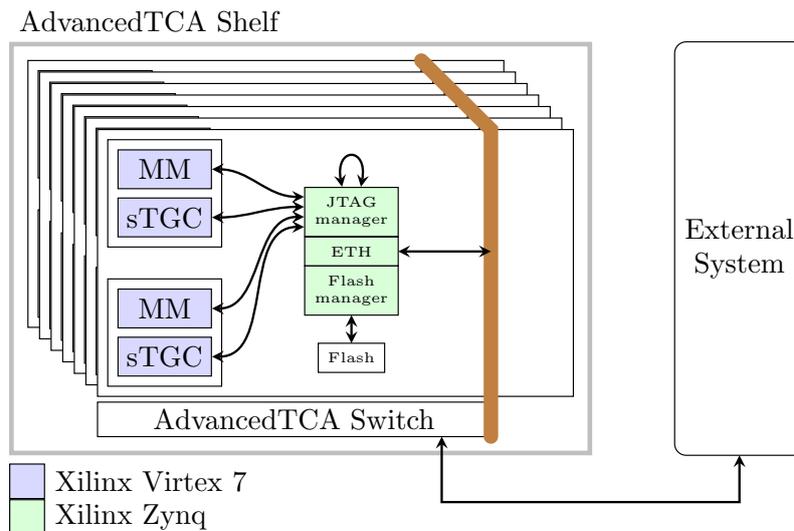


Figure 35: Example of programming scheme using the XVC service via Zynq SoC and network through an ATCA network switch. A similar approach can be done with a regular network switch since each carrier provides network access on the RTM.

## 10 Playback of test data

In order to test the Trigger Processor it is possible to *play-back* simulated events (Athena). To produce the expected inputs to the Trigger Processor with an approximate description of their time wise features (in particular to emulate the correct time relationship between the MM and sTGC Trigger Processors), events obtained from the Athena simulation are used as inputs to a time-based simulation of the NSW trigger chain. The expected inputs are loaded to the Trigger Processor via the SCAX and then played back (event  $\rightarrow$  Tx  $\rightarrow$  Rx  $\rightarrow$  Trigger Processor algorithms).

The time-based simulation is driven by an event-based scheduler. The first version of the simulation is to be released in the upcoming weeks. Currently the simulation includes approximate models for the following items in the NSW trigger chain.

- MM: MMFE8
- MM: Cabling from MMFE8 to ADDC
- MM: ADDC
- MM: Fibres from ADDC to Trigger processor
- sTGC: sTGC strip adapter boards (to describe which strips are connected to which VMM channel)
- sTGC: sTGC front end boards
  - VMM direct output 6-bit ADC
  - TDS
  - VMM to TDS connections
- sTGC: Pad trigger (using an external coincidence table)
- sTGC: Cabling from Pad Trigger to sTGC strip FE-board
- sTGC: Cabling from strip FE-board to Router
- sTGC: Router (includes mis-match between fiber input ordering and serializer block order)

- 818 • sTGC: Fibres from Router to Trigger Processor

## 819 **11 Firmware organization and continuous integration**

820 The firmware language used in the Trigger Processor collaboration is VHDL. Xilinx Vivado is used for  
821 both synthesis and place & route of the firmware. All firmware is stored in the CERN Gitlab group  
822 [https://gitlab.cern.ch/nsw\\_trigger\\_processor\\_elx](https://gitlab.cern.ch/nsw_trigger_processor_elx).

823 Multiple institutes are involved with the development of the Trigger Processor firmware. There are firmware  
824 blocks specific for sTGC, specific for MM and blocks that are shared. The blocks need to be integrated  
825 with each other into the two top-level designs: sTGC and MM. The standard tool flows from Xilinx Vivado  
826 or Mentor Graphics were not efficient for such a collaboration. We have adopted the design flow developed  
827 by the FELIX team that makes it possible to work together.

828 By default, Vivado creates a multi-level directory structure containing all the source files. The Vivado  
829 project is not included in the source tree under version control, but is generated by sourcing a TCL script  
830 from within the Vivado GUI. This script creates a fresh Vivado project by copying all the needed source  
831 files from the repositories; no source files are kept inside the Vivado project tree. An additional benefit of  
832 this approach is that the huge Vivado project is not included in version control. The project contains many  
833 binary and generated files; these files would create many changes in the version control system and would  
834 complicate the manual browsing through the repository.

835 A complete set of firmware needed for the Trigger Processor is versioned together in a so-called  
836 *release* repository at [https://gitlab.cern.ch/nsw\\_trigger\\_processor\\_elx/nsw-trig-proc](https://gitlab.cern.ch/nsw_trigger_processor_elx/nsw-trig-proc).  
837 This release repository represents an ensemble of git commits (as git submodules) that represent compatible  
838 Trigger Processor firmware, including all dependencies. Major versions of the entire Trigger Processor  
839 firmware set are git tagged from this release repository creating a uniform release structure for all relevant  
840 firmware. These tags take the form NSWTrigProc-XX-XX-XX.

841 Additionally, the health of all firmware is ensured using a Continuous Integration and Continuous  
842 Deployment (CI/CD) setup in CERN GitLab. Using custom CI/CD “runners” (or optionally shared CERN  
843 runner resources with the use of a Docker image-based setup), every push event launches a pipeline of  
844 jobs on the runners to build the projects and obtain bit files, storing them as artifacts for testing. Merge  
845 requests into the primary development branches are required to pass these build tests to ensure the primary  
846 development branch always functions. Tagging git commits also triggers a CD job to deploy the Vivado  
847 output files to a tag area in the EOS file system for preservation of tagged bit files as well as quick  
848 deployment of stable firmware builds to test sites.

## References

- 849
- 850 [1] ATLAS New Small Wheel collaboration, *New Small Wheel Technical Design Report*, (2013), URL:  
851 <https://cds.cern.ch/record/1552862> (cit. on p. 6).
- 852 [2] ATLAS New Small Wheel Trigger Processor Working Group, *Trigger Processor Design Review*  
853 *Report*, (2015), URL: [https://edms.cern.ch/file/1470527/1/NSW\\_TriggerProcessor\\_](https://edms.cern.ch/file/1470527/1/NSW_TriggerProcessor_PDR_20150210.pdf)  
854 [PDR\\_20150210.pdf](https://edms.cern.ch/file/1470527/1/NSW_TriggerProcessor_PDR_20150210.pdf) (cit. on pp. 6, 12).
- 855 [3] eicSys GmbH, *EATCA-100, FPGA based ATCA blade for RD51 SRS system, User Manual*, (2014),  
856 URL: [https://edms.cern.ch/file/1476256/1/NSW\\_TriggerProcessor\\_SRS\\_ATCA\\_PDR\\_](https://edms.cern.ch/file/1476256/1/NSW_TriggerProcessor_SRS_ATCA_PDR_20150209.pdf)  
857 [20150209.pdf](https://edms.cern.ch/file/1476256/1/NSW_TriggerProcessor_SRS_ATCA_PDR_20150209.pdf) (cit. on pp. 6, 9).
- 858 [4] Sorin Martoiu, *High-Density Optical Receiver Mezzanine board for ATCA-SRS*, (2015), URL:  
859 [https://edms.cern.ch/file/1476256/1/NSW\\_TriggerProcessor\\_SRS\\_Mezzanine\\_](https://edms.cern.ch/file/1476256/1/NSW_TriggerProcessor_SRS_Mezzanine_PDR_20150209.pdf)  
860 [PDR\\_20150209.pdf](https://edms.cern.ch/file/1476256/1/NSW_TriggerProcessor_SRS_Mezzanine_PDR_20150209.pdf) (cit. on pp. 6, 9).
- 861 [5] ATLAS New Small Wheel Trigger Processor Group, *NSW Trigger Processor Carrier Board*  
862 *Version 2 Specification*, (2018), URL: [https://espace.cern.ch/ATLAS-NSW-ELX/Shared%](https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%20Processor/NSWTP_NewCarrierBoardRequirements.pdf)  
863 [20Documents/NSW%20Trigger%20Processor/NSWTP\\_NewCarrierBoardRequirements.](https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%20Processor/NSWTP_NewCarrierBoardRequirements.pdf)  
864 [pdf](https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%20Processor/NSWTP_NewCarrierBoardRequirements.pdf) (cit. on p. 11).
- 865 [6] B Clark, D Lopez Mateos, N Felt, J Huth, and J Oliver, *An Algorithm for Micromegas Segment*  
866 *Reconstruction in the Level-1 Trigger of the New Small Wheel*, (2014), URL: [https://cds.cern.](https://cds.cern.ch/record/1753329/files/ATL-UPGRADE-INT-2014-001.pdf)  
867 [ch/record/1753329/files/ATL-UPGRADE-INT-2014-001.pdf](https://cds.cern.ch/record/1753329/files/ATL-UPGRADE-INT-2014-001.pdf) (cit. on pp. 17, 19).
- 868 [7] S Chan and et al., *Micromegas Trigger Processor Algorithm Performance in Nominal, Misaligned, and*  
869 *Corrected Misalignment Conditions*, (2015), URL: [https://cds.cern.ch/record/2113121/](https://cds.cern.ch/record/2113121/files/ATL-COM-UPGRADE-2015-033.pdf)  
870 [files/ATL-COM-UPGRADE-2015-033.pdf](https://cds.cern.ch/record/2113121/files/ATL-COM-UPGRADE-2015-033.pdf) (cit. on pp. 17, 21).
- 871 [8] J. Farah et al., *Test of the Micromegas Trigger Processor with Cosmic Ray Muons*, (2015), URL:  
872 <https://cds.cern.ch/record/2285496/files/ATL-COM-MUON-2017-049.pdf> (cit. on  
873 p. 21).
- 874 [9] M Barros Marin et al., *The GBT-FPGA core: features and challenges*, (2014), URL: [https :](https://cds.cern.ch/record/2158963/files/10.1088/1748-0221/10/03/C03021.pdf)  
875 [//cds.cern.ch/record/2158963/files/10.1088\\_1748-0221\\_10\\_03\\_C03021.pdf](https://cds.cern.ch/record/2158963/files/10.1088/1748-0221/10/03/C03021.pdf) (cit. on  
876 pp. 23, 30).
- 877 [10] J Mendez, S Baron, P Leitao, and A Caratelli, *An Algorithm for Micromegas Segment Reconstruction*  
878 *in the Level-1 Trigger of the New Small Wheel*, (2017), URL: [https://cds.cern.ch/record/](https://cds.cern.ch/record/2312287/files/PoS(TWEPP-17)083.pdf)  
879 [2312287/files/PoS\(TWEPP-17\)083.pdf](https://cds.cern.ch/record/2312287/files/PoS(TWEPP-17)083.pdf) (cit. on pp. 23, 30).
- 880 [11] ATLAS New Small Wheel collaboration, *NSW Trigger Processor status and control registers*, (2019),  
881 URL: [https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%](https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%20Processor/Registers.xlsx)  
882 [20Processor/Registers.xlsx](https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%20Processor/Registers.xlsx) (cit. on p. 32).
- 883 [12] ATLAS New Small Wheel collaboration, *SCX - Requirements and Implementation of Access to Con-*  
884 *trol and Status Registers in FPGAs via FELIX*, (2018), URL: [https://espace.cern.ch/ATLAS-](https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%20Processor/SCX_requirements.pdf)  
885 [NSW-ELX/Shared%20Documents/NSW%20Trigger%20Processor/SCX\\_requirements.pdf](https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%20Processor/SCX_requirements.pdf)  
886 (cit. on p. 31).
- 887 [13] ATLAS New Small Wheel collaboration, *SCA eXtension Deployment User Guide*, (2019), URL:  
888 [https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%](https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%20Processor/scx_userGuide.pdf)  
889 [20Processor/scx\\_userGuide.pdf](https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/NSW%20Trigger%20Processor/scx_userGuide.pdf) (cit. on p. 31).

- 890 [14] C Bakalis et al., *SCA eXtension: a Design for FPGA Parameter Configuration within the ATLAS*  
891 *DAQ Scheme*, (2019), URL: [https://cds.cern.ch/record/2711577/files/ATL-COM-DAQ-](https://cds.cern.ch/record/2711577/files/ATL-COM-DAQ-2020-016.pdf)  
892 [2020-016.pdf](https://cds.cern.ch/record/2711577/files/ATL-COM-DAQ-2020-016.pdf) (cit. on p. 31).
- 893 [15] S. Bressler, A. Koulouris, A. Lanza, L. Levinson, and S. Zimmermann, *Trigger latency for the*  
894 *sTGC in the New Small Wheel*, (2016), URL: [https://edms.cern.ch/file/1604270/1/sTGC\\_](https://edms.cern.ch/file/1604270/1/sTGC_Latency_fiber_routing_final_20160314.pdf)  
895 [Latency\\_fiber\\_routing\\_final\\_20160314.pdf](https://edms.cern.ch/file/1604270/1/sTGC_Latency_fiber_routing_final_20160314.pdf) (cit. on p. 33).