Technical Specification

ATLAS Level-1 Calorimeter Trigger Upgrade

Jet Feature Extractor (jFEX) Prototype

Sebastian Artz, Stefan Rave, Ulrich Schäfer,
Esteban Torregrosa

Draft

Version: 0.2

14 July 2014

Contents

1 Related Documents 3
2 Conventions 3
3 Introduction 4
3.1 Overview of Phase-1 System 5
3.2 Overview of Phase-2 system 6
4 Functionality 7
4.1 Real-Time Data Path 8
4.1.1 Input Data Granularity 8
4.1.2 Feature Identification Algorithms 8
4.1.3 Processing Area 9
4.1.4 Result Merging 10
4.1.5 Output Bandwidth 11
4.2 Readout Data Path 11
4.3 Latency 14
4.4 Error Handling 14
4.5 Interface to TTC 14
4.6 Slow Control 15
4.7 Environment Monitoring 15
4.8 Commissioning and Diagnostic Facilities 15
4.9 ATCA form factor 16
5 Data formats 16
Input data 16
5.1 16
5.1.1 Calorimeter Data Format 17
5.2 Real-Time Output Data 17
Readout data 19
5.3 19
6 Implementation 20
6.1 Input Data Reception and Fan Out 20
6.2 Processor FPGA 21
6.3 Merger FPGA 22
6.3.1 Result merging 22
6.3.2 Readout 23
6.3.3 Control 23
6.4 Clocking 24
6.5 High-Speed signals on the PCB 24
6.6 FPGA configuration 25
6.7 The IPM Controller 25
6.8 Power Management 26
6.9 Front-panel Inputs and Outputs 26
6.10 Rear-panel Inputs and Outputs 27
6.10.1 ATCA Zone 1 27
6.10.2 ATCA Zone 2 27
6.10.3 ATCA Zone 3 28
6.11 LEDs 28
6.12 Instrument Access Points 29
6.12.1 Set-Up and Control Points 29
6.12.2 Signal Test Points 29
6.12.3 Ground Points 29
6.13 Floor plan 29
7 Front-Panel Layout 31
8 Programming model 31

8.1 Guidelines 31
8.2 Register Map 31
8.3 Register Descriptions 32
9 Glossary 32
10 Document History 34

1 Related Documents

[1.1] ATLAS TDAQ System Phase-I Upgrade Technical Design Report, CERN-LHCC-2013-018, http://cds.cern.ch/record/1602235/files/ATLAS-TDR-023.pdf

[1.2] L1Calo Phase-I Hub Specification *(not yet available)*

[1.3] L1Calo Phase-I ROD specification *(https://twiki.cern.ch/twiki/pub/Atlas/LevelOneCaloUpgradeModules/Hub-ROD_spec_v0_9.pdf)*

[1.4] L1Calo Phase-I eFEX Specification *(https://twiki.cern.ch/twiki/pub/Atlas/LevelOneCaloUpgradeModules/eFEX_spec_v0.2.pdf)*

[1.5] L1Calo Phase-I Optical plant Specification *(not yet available)*

[1.6] ATCA Short Form Specification, http://www.picmg.org/pdf/picmg_3_0_shortform.pdf

[1.7] PICMG 3.0 Revision 3.0 AdvancedTCA Base Specification, *access controlled*, http://www.picmg.com/

[1.8] L1Calo High-Speed Demonstrator report *(https://twiki.cern.ch/twiki/pub/Atlas/LevelOneCaloUpgradeModules/HSD_report_v1.02.pdf)*

[1.9] Development of an ATCA IPMI controller mezzanine board to be used in the ATCA developments for the ATLAS Liquid Argon upgrade, http://cds.cern.ch/record/1395495/files/ATL-LARG-PROC-2011-008.pdf

2 Conventions

The following conventions are used in this document.

A programmable parameter is defined as one that can be altered by slow control, for example, between runs, not on an event by event basis. Changing such a parameter does not require a re-configuration of any firmware.

Where multiple options are given for a link speed, for example, the readout links of the jFEX are specified as running up to 10 Gb/s, this indicates that the link speed has not yet been fully defined. Once it is defined, that link will use a single speed. No links on the jFEX will support more than one speed in operation.

In accordance with the ATCA convention, a crate of electronics here is referred to as a shelf.

Where the term jFEX is used here, without qualification, it refers to the jFEX module. The jFEX subsystem is always referred to explicitly by that term.

3 Introduction

This document describes the jet Feature Extractor (jFEX) module of the ATLAS Level-1 Calorimeter Trigger Processor (L1Calo) [1.1] . The jFEX is one of several modules being designed to upgrade L1Calo, providing the increased discriminatory power necessary to maintain trigger efficiency as the LHC luminosity is increased beyond that for which ATLAS was originally designed.

The function of the jFEX module is to identify, in the data received from the electromagnetic and hadronic calorimeters, large energy deposits indicative of jets and τ particles. Global parameters like total transverse energy and missing transverse energy are also calculated. The jFEX does this, using algorithms of greater complexity and data of finer granularity than those processed by the current L1Calo system. At the same time larger jets can be calculated than those possible in the Phase-0 system.

The jFEX will be installed in L1Calo during the long shutdown LS2, as part of the Phase-1 upgrade, and it will operate during Run 3. It will remain in the system after the Phase-2 upgrade in LS3, and will operate during Run 4, at which time it will form part of L0Calo. The following sections provide overviews of L1Calo in Run 3 and L0Calo in Run 4. Where the requirements on the jFEX are different in these two runs, it is the most demanding requirements that are used to specify the jFEX. Where these arise from Run 4, this is stated explicitly in the text.

This is a specification for a prototype jFEX. This prototype is intended to exhibit the full functionality of the final module, plus some additional functionality for research purposes, such as the ability to handle different speeds of multi-Gb/s link. The other differences between prototype and final functionality are noted in the appropriate sections. Excepting these, the functionality described here can be regarded as that of the final jFEX.

3.1 Overview of Phase-1 System

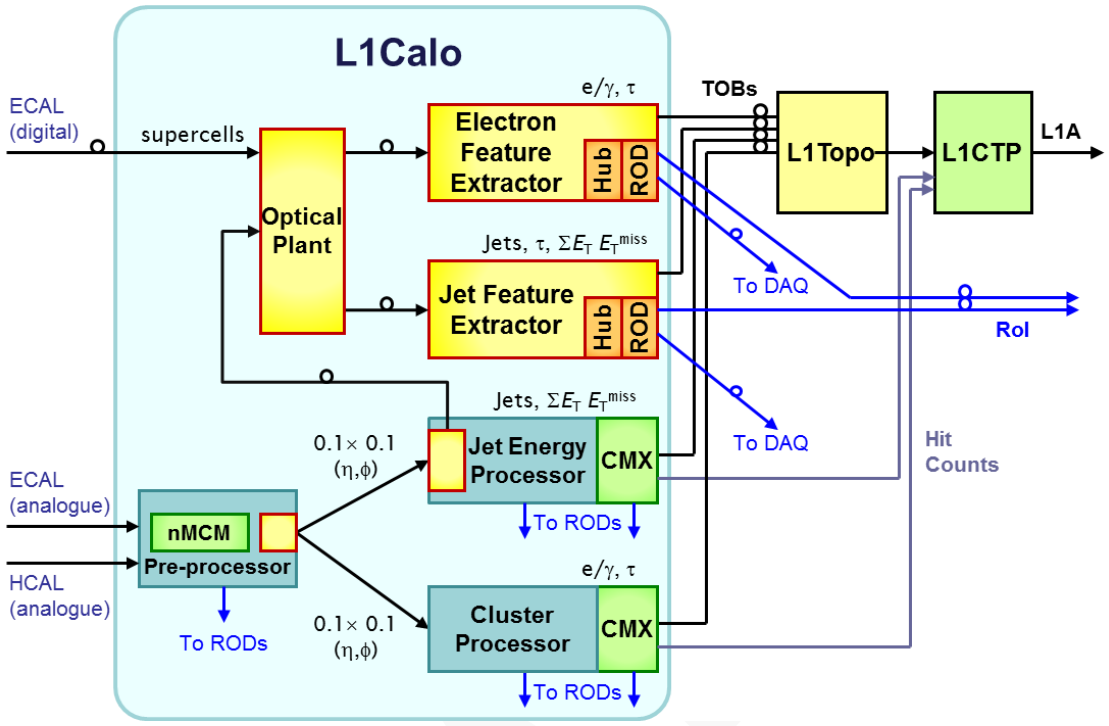

Figure 1. The L1Calo system in Run 3. Components installed during LS2 are shown in yellow /orange.

In Run 3, L1Calo contains three subsystems installed prior to LS2 (see Figure 1):

- the Pre-Processor, which receives shaped analogue pulses from the ATLAS calorimeters, digitises and synchronises them, identifies the bunch-crossing from which each pulse originated, scales the digital values to yield transverse energy (E_T), and prepares and transmits the data to the following stages;
- the Cluster Processor (CP) subsystem, comprising Cluster Processor Modules (CPMs) and Common Merger Extended Modules (CMXs), which identifies isolated e/γ and τ candidates;
- the Jet/Energy Processor (JEP) subsystem, comprising Jet/Energy Modules (JEMs) and Common Merger Extended Modules (CMXs), which identifies energetic jets and computes various energy sums.

Also installed prior to LS2 is the Topological Processor (L1Topo), which receives data from L1Calo and L1Muon and applies topological algorithms and kinematic cuts.

Additionally, it contains the following two subsystems, installed as part of the Phase-1 upgrade in LS2:

- the eFEX subsystem, comprising eFEXs, and Hub modules [1.2] with Readout Driver (ROD) daughter cards [1.3] , which identifies and counts isolated e/γ and τ candidates, using data of finer-granularity than does the CP subsystem;
- the jet Feature Extractor (jFEX) subsystem [1.4] , comprising jFEX modules and Hub modules with ROD daughter cards, which identifies energetic jets, large τ candidates and computes various energy sums, using data of finer-granularity while handling larger jet

windows than does the JEP subsystem. The increased granularity also allows for more flexible and more complex jet algorithms.

In Run 3, the electromagnetic calorimeter (ECAL) electronics provide L1Calo with both analogue signals (for the CP and JEP) and digitised data (for the eFEX and jFEX subsystems). From the hadronic calorimeters (HCAL), only analogue signals are received. These are digitised on the JEP and output optically towards the eFEX and jFEX subsystems. Initially at least, the eFEX and jFEX subsystems operate in parallel with the CP and JEP subsystems. Once the outputs of the eFEX and jFEX have been validated, removing the CP and JEP systems from L1Calo is an option, excepting that section of the JEP used to provide hadronic data to the FEX subsystems.

The optical signals from the JEP and ECAL electronics are sent to the FEX subsystems via an optical plant [1.5] . It breaks apart the fibre bundles and regroups them, changing the mapping from that employed by the ECAL and JEP electronics to that required by the eFEX and jFEX subsystems.

The outputs of the eFEX and jFEX (plus CP and JEP) subsystems comprise Trigger Objects (TOBs): words which describe the location and characteristics of any candidate trigger objects found. These words are transmitted to L1Topo via optical fibres. There, they are merged over the system and topological algorithms are run, the results of which are transmitted to the Level-1 Central Trigger Processor (CTP).

The eFEX and jFEX subsystems are both implemented using the ATCA standard [1.6] [1.7] . The jFEX comprises one shelf of 7 jFEX modules. L1Topo comprises up to four identical modules, each of which receives a copy of all data from all the jFEX modules.

As for the other L1Calo processing modules, on receipt of an L1A the jFEX subsystem provides Region of Interest (RoI) and Readout data to Level-2 and the DAQ system respectively. Each jFEX module outputs these data on to the shelf backplane, and two Hub modules in each shelf aggregate the data and implement the required ROD functionality (via a daughter board). Additionally, the Hub modules provide hubs on the TTC, control and monitoring networks.

3.2 Overview of Phase-2 system

In the Phase-2 upgrade, the entire calorimeter input to L1Calo (from both the ECAL and the HCAL) will migrate to the digital path, rendering the Pre-Processor, CP and JEP subsystems obsolete. The eFEX and jFEX subsystems will thus form the front-end of the L1Calo system at Phase-2. Hence, not only must they be designed for adequate performance at Phase-1, they must also be compatible with the Phase-2 upgrade.

The Phase-2 upgrade will be installed in ATLAS during LS3. At this point, substantial changes will be made to the trigger electronics: the entire calorimeter input to the trigger (from both the ECAL and the HCAL) migrates to the digital path; the latency available to the Level-1 trigger is greatly increased; L1Track is introduced and requires seeding. To meet these new opportunities and demands, L1Calo is split into L0Calo (a Level-0 calorimeter trigger) and L1Calo.

The remaining parts of the previous L1Calo system, the eFEX and jFEX subsystems remain, but they now constitute L0Calo (see Figure 2). They supply L0Topo (and thence L0CTP) with real-time TOB data and, on receipt of an L0A, they supply readout data to the DAQ system, plus RoI data to L1Calo and L1Track.

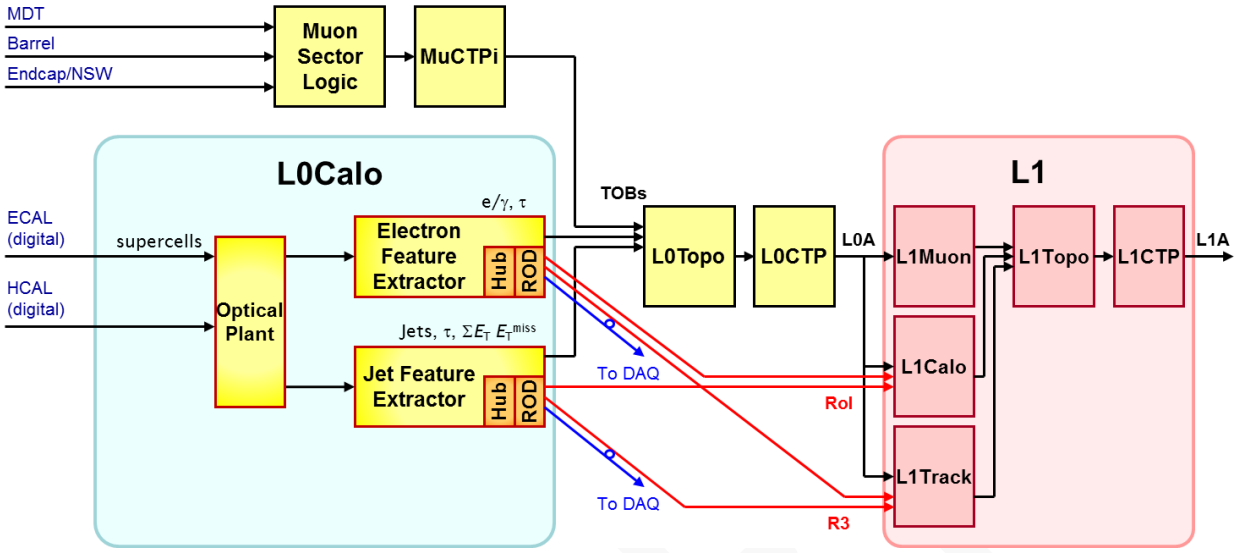

Figure 2. The L0/L1Calo system in Run 4. The new L1 system is shown in red /pink. Other modules (yellow /orange) are adapted from the previous system to form the new L0Calo.

4 Functionality

Figure 3 shows a block diagram of the jFEX. The various aspects of jFEX functionality are described in detail below. Implementation details are given in section 6.

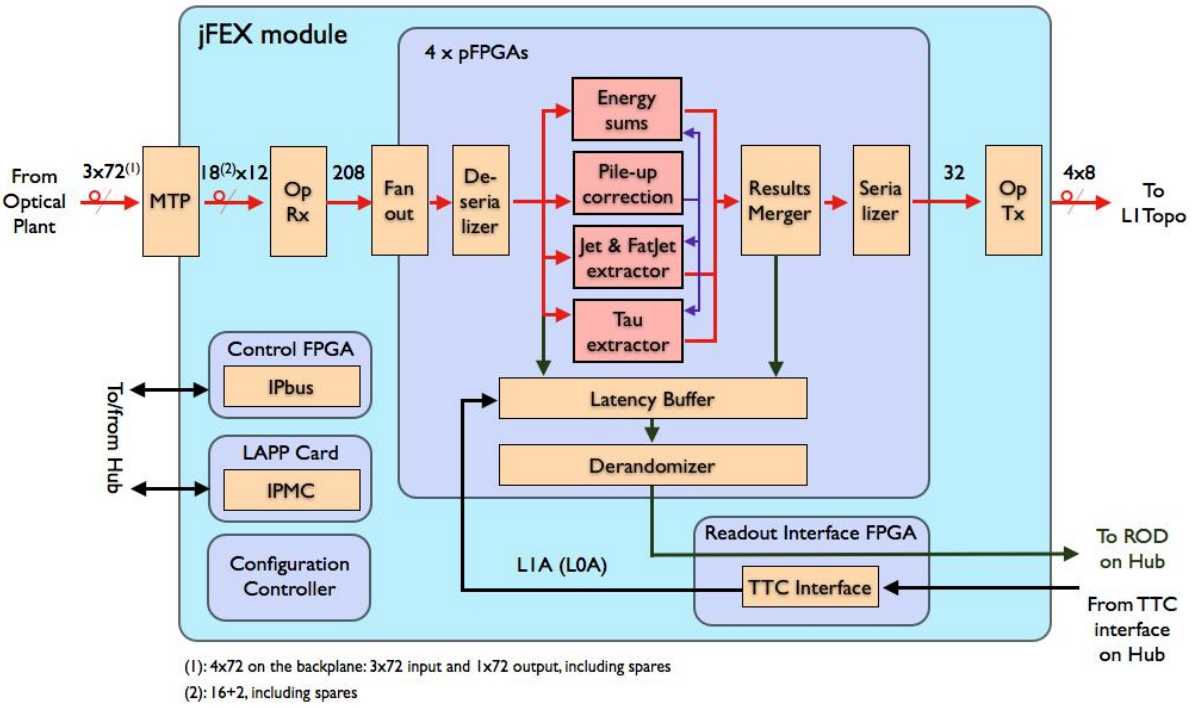

Figure 3. A block diagram of the jFEX module. Shown are the real-time and readout data paths. With the exception of the L1A, control and monitoring signals are omitted for simplicity.

4.1 Real-Time Data Path

4.1.1 Input Data Granularity

The jFEX receives data from all calorimeter parts. From ECAL and HCAL data is received as single E_T values from each 0.1 × 0.1 (η, ϕ) trigger tower, summed in depth, covering the central region $|\eta| < 2.5$. The region $2.5 < |\eta| < 3.2$ is covered in 0.2 × 0.2. The FCAL region of $3.2 < |\eta| < 4.9$ is divided in up to 12 η-bins, depending on the FCAL layer, with a binning width of 0.4 in ϕ.

4.1.2 Feature Identification Algorithms

The jFEX system examines data from the Electromagnetic and Hadronic Calorimeters, within the range $|\eta| < 4.9$, to identify energy deposits characteristic of hadronic jets. The data from the range $|\eta| < 2.5$ is also used to identify energy deposits from τ particles, which are larger than the windows used by the τ identification algorithms on the eFEX system. Combining the data from all jFEX modules, the global parameters E_T and E_T^{miss} are calculated.

Multiple versions of the sliding window algorithm are implemented in parallel, to achieve the best result in identifying jets from energy deposits in both calorimeters.

Physics studies to determine the optimal algorithms are on-going. The following descriptions are indicative; whilst the precise details are undefined, the overall structure and complexity of the algorithms are understood.

The algorithm employed for the purpose of identifying large energy deposits from τ particles, measures the diffuseness of the deposits, thus distinguishing those produced by τ particles from the more diffuse deposits typical of jets interacting in the ECAL. The algorithm can be divided into two steps: the first step is to seed the process of finding τ candidates by searching for characteristic shower shapes and applying a hadronic veto, followed by measuring the energy of the candidate particles.

The process of finding jet candidates is based on the sliding window algorithm. Energy deposits are summed around a central trigger tower over a small area. If this sum is a local maximum, compared to its 8 neighbours, the central trigger tower is considered as the position of a jet candidate. Once a candidate is found, a bigger area around the position is summed, to calculate the transverse energy. In the JEP system the size of the seeding area is 0.4 × 0.4 (η × ϕ) with up to 0.8 × 0.8 for the energy calculation. The jFEX system can calculate candidates with up to 1.7 × 1.7. As a further improvement non-square jet windows are feasible, due to the finer granularity. Also weights can be applied to each individual trigger tower within a jet window. The values that provide the best trigger performance are yet to be determined by studies.

In parallel to the regular jet finding an additional algorithm is used to identify heavily boosted objects, like top quarks, which decay into several, separate jets. For this purpose the sliding window algorithm can be extended to a maximum size of 2.2 × 2.2, using a granularity of 0.2 × 0.2. For the inner region of 1.8 × 1.8 the finer granularity can be used to identify a substructure. The detailed mechanism for finding these "fat jets", as well as the identification of its sub structure, is a matter of ongoing studies.

In addition to the algorithms mentioned above recent development on triggering algorithms, like the so called jet-without-jet algorithms [ref], are being explored to be implemented inside the jFEX.

For those trigger candidates that pass the tests above, Trigger Object words (TOBs) are generated and passed to the merging logic (see below). These TOBs describe the location and type of the candidate and its energy.

At the edges of the η range of the calorimeter, data for the full search windows are not available. In these instances, modified algorithms are applied (the details of which are to be defined).

Beside these TOBs, the parameters E_T and $E_T{}^{miss}$ are estimated on each jFEX module for a subset of the detector and sent to L1Topo to be merged to a global value.

The increased instantaneous luminosity of the LHC in Run 3 and Run 4 will produce unprecedented levels of pile-up. To compensate for this, event-by-event corrections, which cope with the fluctuations of pile-up energy deposits, are determined and applied on each jFEX module. These corrections are viable for jet finding, as well as for $E_T{}^{miss}$ calculations.

4.1.3 Processing Area

The feature identification algorithms used by the jFEX define a core area, over which TOB candidates are found, and a larger environment area, from which the algorithms need to receive the data. As the core areas examined by neighbouring instances of each algorithm are contiguous, the environment areas overlap. In the central region each module covers the whole ϕ-ring over 0.8 in η as the core area. Five modules are required for $|\eta| \leq 2.0$. The region within $2.0 < |\eta| < 4.9$ is covered by a single module for each side of the detector. This is possible due to the coarser granularity and missing environment at the outer end. Contiguous to its core area any module receives data in 0.1 × 0.1 granularity for 0.8 in η in each direction and an additional row in 0.2 × 0.2 granularity beyond this, to support the identification algorithm for fat jets. Hence a single module in the central region receives data from a region of 2.8 × 6.4 (η × ϕ).

In total, seven jFEX modules are required to process all of the data from the calorimeters within the range $|\eta| < 4.9$. The hardware on all jFEX modules is the same; the differences in core area processed are implemented via firmware.

On each jFEX module there are 4 Processor FPGAs. Each of them covers a quarter of the modules core area in ϕ. Combined with the environment of 0.8 in both, η and ϕ, a single FPGA receives data from a total of 2.4 × 3.2 (η × ϕ) in a granularity of 0.1 × 0.1. The extended environment increases the covered area to 2.8 × 3.6 (η × ϕ), using a coarser granularity in the outer region.

Each jFEX module receives only one copy of the data from every trigger tower. The required duplication for overlapping areas is handled internally as shown in Figure 4. The core area (displayed in red) and the environment within the same ϕ range (shown in violet) are received directly from the source. This information is duplicated early in the MGT data path and retransmitted via PMA "loopback" to the neighbouring FPGAs. The duplicated data is shown light blue. The extended environment shown in green is transmitted via the same links that carry the data for the fine granularity. Each input contains 16 trigger towers in fine granularity, arranged in 0.4×0.4, and 2 additional cells covering 0.2×0.4. The extended environment shown in orange is transmitted from the neighbouring FPGAs via low latency 1Gb/s differential links.

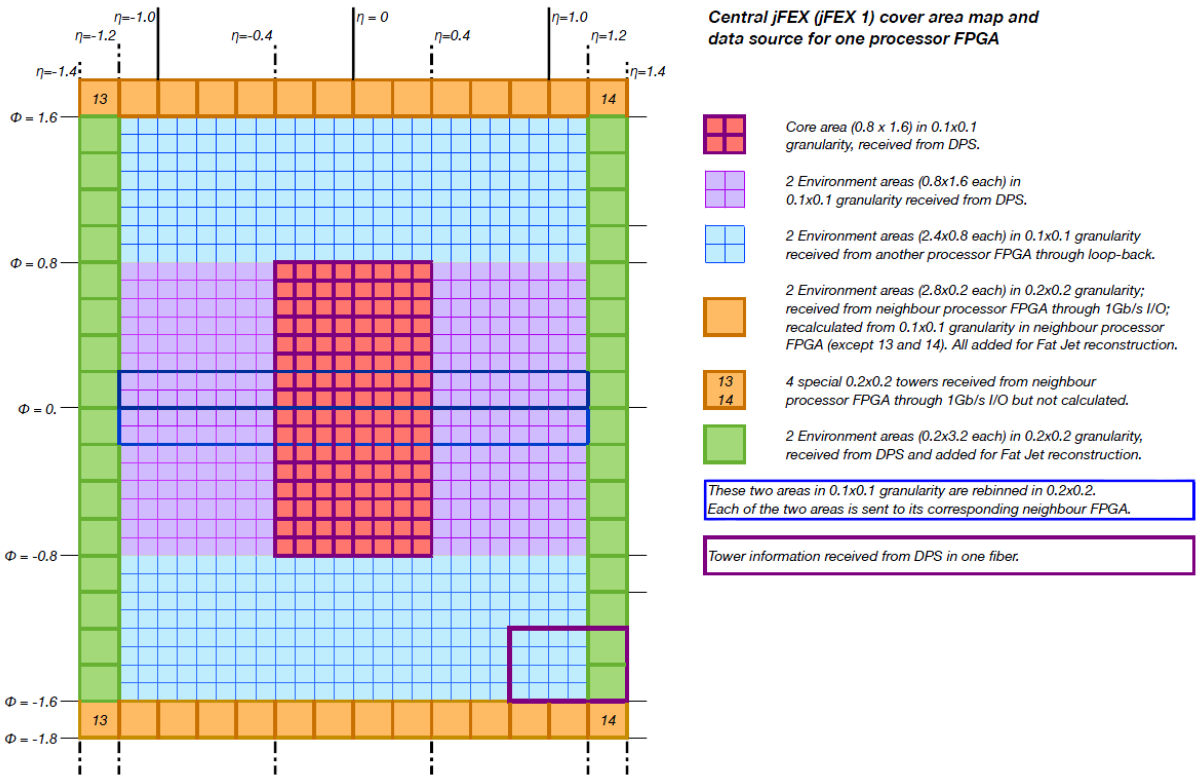

Figure 4. **The input data for a single Processor FPGA, including granularity and source. The red area shows the core area, while the blue and violet shows the environment area. The orange and green shows the extra environment area added for fat jets feature extraction.**

4.1.4 Result Merging

The feature-identification algorithms described above are implemented in four Processor FPGAs, each of which processes data from its own core area of the calorimeters. The outputs from these algorithms are TOBs as well as subtotals for E_T and $E_T{}^{miss}$. In order to make the best use of the limited bandwidth available for transmitting data to L1Topo (see below), these TOBs are 'merged' over each module. That is, they are prioritised and, potentially, sorted. The exact algorithms to be employed for this purpose are to be determined. Their complexity may be limited by the latency they add to the Level-1trigger.

Merging of the output is performed on a separate Merger FPGA. This device receives the results from each of the Processor FPGAs on 48 differential signal pairs operating at 1 Gb/s. This number of links at the given speed allows any one of the Processor FPGAs to supply the entire real-time output of the jFEX concerning TOBs. The exact number of links required for this purpose depends on the bandwidth from the Merger FPGA to L1Topo.

4.1.5 Output Bandwidth

The real-time output data of the jFEX are transmitted to the L1Topo system. This system comprises multiple, identical modules, and the jFEX is assumed to transmit the same data to each module. Furthermore, each L1Topo module houses two FPGAs, and, ideally, the jFEX should transmit the same data to each L1Topo FPGA, to maximise the flexibility and minimise the latency of L1Topo.

Due to bandwidth constraints at the L1Topo input, the real-time output of the jFEX must be carried on a maximum of four fibres per L1Topo FPGA. The exact format of the jFEX TOBs is yet to be defined, but they are estimated to be approximate 40 bits in size. If each L1Topo FPGA must receive its own copy of the jFEX data, using four fibres at 12.8 Gb/s, with 8b/10b encoding, this means no more than 24 TOBs can be output from a jFEX module in any bunch crossing. Studies are underway to confirm that this is sufficient.

If the available bandwidth is not sufficient, there is the option for increasing the real-time output bandwidth of the jFEX. The requirement to transmit a copy of the data to each L1Topo FPGA could be dropped, allowing the jFEX to transmit its data on eight fibres to each L1Topo module. In this scenario, data is transferred between FPGAs on L1Topo. This option also requires additional parallel links running between the Processor FPGAs and the Merger FPGA.

The jFEX has upwards of 40 fibre outputs, allowing copies of the output data to be provided to up to 10 L1Topo FPGAs (on four fibres each), or 5 L1Topo modules (on eight fibres each). Presently, it is not foreseen that there will be more than four L1Topo modules used in L1Calo. The extra output capacity of the jFEX is spare.

4.2 Readout Data Path

On receipt of an L1A signal, the jFEX provides data to a number of systems: in Run 3, it provides RoI data to Level-2; in Run 4, it provides RoI data to L1Track and L1Calo (the jFEX being part of L0Calo in Run 4); in both Run 3 and Run 4, it provides data to the DAQ system. Collectively, these data are referred to here as readout data.

The jFEX outputs a single stream of readout data, which contains the superset of the data required by all of the downstream systems. In Run 3, these data are transmitted across the crate backplane to a ROD. In Run 4, there are two RODs per crate and the jFEX transmits identical readout data to both RODs. It is the RODs that are responsible for formatting the data as required by the downstream systems, and handling the multiple interfaces.

For each event that is accepted by the Level-1 trigger, the jFEX can send three types of data to the readout path: final TOBs, expanded TOBs (XTOBs) and input data. The final TOBs are copies of those transmitted to L1Topo. In normal running mode these are the only data read out. The XTOBs are words that contain more information about trigger candidates than

can be transmitted on the real-time data path (see section 5.2). They are extracted from the real-time path before the merging process and therefore, as merging may reduce the number of TOBs, the number of XTOBs may be larger than the number of TOBs. To minimise the amount of readout data generated, XTOBs are not normally read out. However, this functionality can be enabled via the slow control interface. This cannot be done dynamically for individual events.

The input data comprise all data received from the calorimeters. They are copied from the real-time path after serial-to-parallel conversion and after the CRC word has been checked. There are a number of programmable parameters, set via slow control, that determine which input data are read out. These are as follows.

- The Input Readout mode: by default, only input data from fibres that have generated an error are read out. However, the readout of data received without error can also be enabled.
- The Input Channel Mask: the read out of individual channels of input data, from individual FPGAs, can be disabled. A channel here means the data received at an FPGA from one fibre. In total, a Processor FPGA on the jFEX receives up to 104 channels of input data. However, many of these data are redundant copies, created because of the need to fan out data between the FPGAs. The Input Channel Mask provides a way of stripping redundant channels from the jFEX readout. It also allows data from permanently broken links to be excluded from the readout process.
- The Input Readout Veto: this veto is asserted for a programmable period (0-256 ticks) after the read out of any Input Data. It provides a means of pre-scaling the amount of Input data read out, preventing it from overwhelming the readout path.

The mechanism for capturing readout data is illustrated in Figure 5. For every bunch crossing all input data, intermediate and final TOB data are copied from the real-time path and written to scrolling, dual-port memories. They are read from these memories after a programmable period, of up to 3 μs. At this point they are selected for readout if they meet both of the following criteria: an L1A pertaining to them is received, and they are enabled for readout by the control parameters described above. Otherwise, they are discarded.

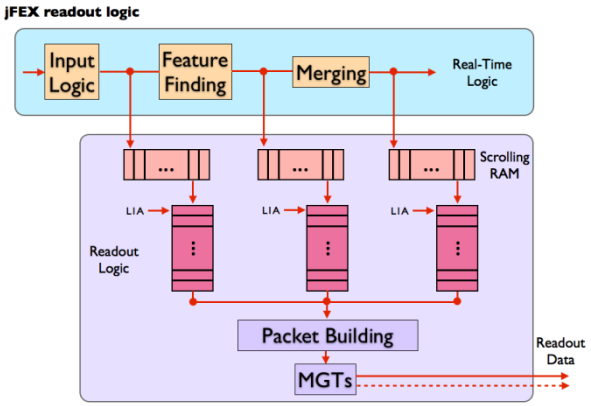

Figure 5. A functional representation of the jFEX readout logic.

For each L1A, data from a time frame, programmable via control parameters, can be read out. The selection of data for read out is a synchronous process with a fixed latency, and it is the

period for which data are held in the scrolling memories that determines the start point of this time frame. The correct value must be determined when commissioning L1Calo (it should correspond to the period from when the data are copied into the scrolling memories, to when an L1A pertaining to those data is received at the jFEX, plus or minus any desired offset in the time frame). The jFEX hardware allows the read out of overlapping time frames. At low rates (including everything before Run 4) the jFEX expects never to read out overlapping time frames. At high rates, the read out of overlapping time frames will be possible, but the frame length and trigger rate will need to be controlled carefully to prevent buffer overflow.

It is possible that for a BC there will be no TOB data (or XTOB data when enabled) to be captured. In such cases a control word is inserted into the readout path to indicate this. This word, which is used for flow-control, is internal to the jFEX; it is not passed to the ROD.

Data that are selected for readout are written to FIFOs, where they are stored before transmission to the ROD. This storage is necessary for two purposes: first, data are sent to the ROD in formatted packets, requiring some data to be stored as the packet is built; secondly, the peak rate at which readout data are captured by the jFEX exceeds that at which they can be transferred to the ROD.

All of the readout logic described above is implemented in the Processor FPGAs. Downstream of the FIFOs, the readout logic is implemented in the Merger FPGA. The data from the Processor FPGAs to the Merger FPGA are transferred via multi-Gb/s transmitter-receiver (MGT) links.

In the Merger FPGA, the data are built into packets and transmitted, via the shelf backplane, to a ROD (in Run 3) or two RODs (in Run 4, each ROD receiving a copy of the same data). Six links, each running at up to 10 Gb/s, carry the data to a ROD.

The transfer of data from the FIFOs, via the Merger FPGA, to the ROD(s), is initiated whenever the FIFOs are not empty. There is no backpressure asserted from the ROD(s) to pause transmission.

Table 1 shows an estimate of the maximum readout bandwidth required of the jFEX. This maximum case occurs in Run 4, where readout is initiated by the L0A signal at a rate of 1 MHz. However, only the TOB data need to be read out at this rate. In normal operation, the input data are only read out at a pre-scaled rate for monitoring purposes. They are also read out if an error is detected, but if that error is persistent, those data are also pre-scaled. Thus, for the input data, a maximum readout rate of 50 KHz is acceptable. The number of bunch crossings from which data is read out after an L1A (L0A) can be set via control parameters. For normal operation a window size of three bunch crossings is assumed in this calculation.

Readout of the XTOB data is optional. The calculation shown in Table 1 assumes a pre-scaled rate of 500 KHz. The number of XTOBs in this calculation is based on the number of TOBs that can be sent in the real time data path from a single Processor FPGA. The maximum number of generated XTOBs depends on the exact implementation of the algorithms and is thus not yet known. Based on the available bandwidth, the maximum rate for readout of XTOBs can be calculated, once details for the algorithms are known.

Data	No. Chan.	Bits / Chan. / BC (post 8b/10b)	BC / Event	Trigger Rate / KHz	Bandwidth / Gb/s
Input data	416	320	3	50	19.97
XTOBs	96	80	3	500	11.52
TOBs	4	320	3	1000	3.84
Total					**35.33**

Table 1. An estimate of the maximum readout bandwidth required for a jFEX module. For XTOBs, a channel equals an XTOB; for the other types of data, a channel equals a fibre.

4.3 Latency

A breakdown of the estimated latency of the real-time path of the jFEX is given in the ATLAS TDAQ System Phase-1 Upgrade Technical Design Report [1.1] .

4.4 Error Handling

The data received by the jFEX from the Calorimeters are accompanied by a CRC code. This is checked in the Processor FPGAs, immediately after the data are converted from serial, multi-Gb/s streams into parallel data. If an error is detected, the following actions are performed:

- All data to which a detected error pertains are suppressed (i.e. set to zero) on the real-time path. They are passed to the readout path as received.
- The Error Check Result for the current clock cycle is formed from the 'OR' of all error checks for the current bunch crossing.
- The Input Error Count is incremented for any clock cycle where there is at least one error in any input channel.
- A bit is set in the Input Error Latch for any channel that has seen an error. These bits remains set until cleared by an IPBus command.
- The global Input Error bit is formed from the 'OR' of all bits in the Input Error Latch.

The Error Check Result, Input Error Count, Input Error Latch and Input Error bit can all be read via IPBus. A single IPBus command is provided to clear all of these registers. The Error Check Result and Input Error Count are included in the readout data for the current bunch crossing. The jFEX does not generate any other external error signal, so data monitoring or regular hardware scanning must detect an error condition.

4.5 Interface to TTC

TTC signals are received in the jFEX shelf in the Hub-ROD module. There, the clock is recovered and commands are decoded, before being re-encoded using a local protocol (to be

defined). This use of a local protocol allows the TTC interface of the shelf to be upgraded without any modification of the jFEX modules.

The jFEX module receives the clock and TTC commands from the Hub-ROD via the ATCA backplane. It receives the clock on one signal pair and the commands on a second (see section 6.10 for details).

4.6 Slow Control

An IPBus interface is provided for high-level, functional control of the jFEX. This allows, for example, algorithmic parameters to be set, modes of operation to be controlled and spy memories to be read.

IPBus is a protocol that runs over Ethernet to provide register-level access to hardware. Here, it is run over a 1000BASE-T Ethernet port, which occupies one channel of the ATCA Base Interface. On the jFEX there is a local IPBus interface in every FPGA, plus the IPMC. These interfaces contain those registers that pertain to that device. The Merger FPGA implements the interface between the jFEX and the shelf backplane, routing IPBus packets to and from the other devices as required. The Merger FPGA also contains those registers which control or describe the state of the module as a whole. For those devices such as MiniPODs, which have an I^2C control interface, an IPBus-I2C bridge is provided.

4.7 Environment Monitoring

The jFEX monitors the voltage and current of every power rail on the board. It also monitors the temperatures of all the FPGAs, of the MiniPOD receivers and transmitters, and of other areas of dense logic. Where possible, this is done using sensors embedded in the relevant devices themselves. Where this is not possible, discrete sensors are used.

The voltage and temperature data are collected by the jFEX IPMC, via an I^2C bus. From there, they are transmitted via IPBus to the ATLAS DCS system. The jFEX hardware also allows these data to be transmitted to the DCS via IPMB and the ATCA Shelf Controller, but it is not foreseen that ATLAS will support this route.

If any board temperature exceeds a programmable threshold set for that device, IPMC powers down the board payload (that is, everything not on the management power supply). The thresholds at which this function is activated should be set above the levels at which the DCS will power down the module. Thus, this mechanism should activate only if the DCS fails. This might happen, for example, if there is a sudden, rapid rise in temperature to which the DCS cannot respond in time.

4.8 Commissioning and Diagnostic Facilities

To aid in module and system commissioning, and help diagnose errors, the jFEX can be placed in Playback Mode (via an IPBus command). In this mode, real-time input data to the jFEX are ignored and, instead, data are supplied from internal scrolling memories. These data are fed into the real-time path at the input to the feature-extracting logic, where they replace the input data from the calorimeters.

Optionally, the real-time output of the jFEX can also be supplied by a scrolling memory. It should be noted that, in this mode, the jFEX will process data from one set of memories, but the real-time output will be supplied by a second set of memories. Depending on the content of these memories, this may result in a discrepancy between the real-time and readout data transmitted from the jFEX.

In Playback Mode the use of the input scrolling memories is mandatory, the use of the output scrolling memories is optional, and it is not possible to enable Playback Mode for some channels but not others. Playback Mode is selected, and the scrolling memories loaded, via the slow control interface. The scrolling memories are 256 words in depth.

In addition to the above facility, numerous flags describing the status of the jFEX can be read via the slow control interface (see section 8). Access points are also provided for signal monitoring, boundary scanning and the use of proprietary FPGA tools such as ChipScope and IBERT.

4.9 ATCA form factor

The jFEX is an ATCA module, conforming to the PICMG® 3.0 Revision 3.0 specifications.

5 Data formats

The formats of the data received and generated by the jFEX have yet to be finalised. Those defined here are working assumptions only.

5.1 Input data

The jFEX modules receive data from the calorimeters on optical fibres. For the region $|\eta| < 2.5$ each fibre carries the data for 16 adjacent triggers towers, i.e. an area of 0.4×0.4 ($\eta \times \phi$). The data from the calorimeters within $2.4 < |\eta| < 3.2$ can be sent using one fibre per 0.4 in ϕ, due to the coarser granularity of only 12 trigger towers. This also leaves spare capacity to include the additional information from the Tile-HEC overlap, which cannot be included in the fibres of the central region. The Processor FPGAs covering $0.4 < |\eta| < 1.2$ do not receive these fibres. The corresponding modules receive data from the overlap region on separate fibres. The complete data from the FCAL, covering all three layers, is carried on three fibres per 0.8 in ϕ. These mappings are independent of the line rate of the optical fibres. The data format and content, however, are dependent upon the line rate. For the baseline line rate of 12.8 Gb/s, the data are encoded as specified below. For lower line rates, no format is specified here. Further study is required to establish the optimal bandwidth.

Data from the calorimeters are transmitted to the jFEX as continuous, serial streams. To convert these streams into parallel data, the jFEX logic must be aligned with the word boundaries in the serial data. The scheme for achieving this is yet to be defined, but there are a number of possible mechanisms. For example, boundary markers can be transmitted during gaps in the LHC bunch structure. These markers are substituted for zero data and are interpreted as such by the jFEX trigger-processing logic. Periodic insertion of such markers allows links to recover from temporary losses of synchronisation automatically.

5.1.1 Calorimeter Data Format

- Up to 13-bit data are provided for each of the 16 trigger tower.
- Up to 13-bit data are provided for each of the two 0.2 × 0.2 (η × ϕ) cells.
- A 10-bit cyclic redundancy check is used to monitor transmission errors.
- 8b/10b encoding is used to maintain the DC balance of the link and ensure there are sufficient transitions in the data to allow the clock recovery.
- Word-alignment markers (8b/10b control words) are inserted periodically, as substitutes for zero data.

Using 8b/10b encoding, the available payload of a 12.8 Gb/s link is 256 bits per bunch crossing (BC). The above scheme uses up to 244 bits (data from 18 cells, plus a 10-bit CRC). The remaining 12 bits/BC are spare. The order of the data in the payload is not yet defined.

5.2 Real-Time Output Data

The Real-time output of the jFEX comprises TOBs, each of which contains information about a jet or τ candidate, such as its location and the deposited energy. Figure 6, Figure 7 and Figure 8 show the draft format of the jet TOB, fat jet TOB and τ TOB respectively. Besides these candidates the global values, E_T and $E_T{}^{miss}$, are transferred to L1TOPO. They are sent separately as E_T, $E_T{}^{X}$ and $E_T{}^{Y}$ as 13-bit energy values.

Due to multiple jet finding algorithms, the jet TOBs include two energies. The results from two algorithms, which are based on the same seeding procedure, can be compressed into one TOB. The sizes of the remaining TOBs are adjusted to match the jet TOBs.

The TOBs are transmitted to L1Topo on optical fibres. The line rate and protocol used for this transmission is the same as that used to transmit data from the calorimeters to the jFEX. The baseline specification is thus as follows.

- The data are transferred across the optical link at a line rate 12.8 Gb/s.
- 8b/10b encoding is used to maintain the DC balance of the link and ensure there are sufficient transitions in the data to allow the clock recovery.
- Word-alignment markers (8b/10b control words) are inserted periodically, as substitutes for zero data.

For TOBs of 40 bits, four links at the given specifications allow a maximum of 24 TOBs and the global values to be transmitted to L1Topo per bunch crossing.

Should the specification of the jFEX inputs change, the specification of the real-time outputs will be updated to match. (Using a common line rate and encoding scheme enables the output data to be looped back to the inputs for diagnostic purposes.)

Hadronic Jet TOB

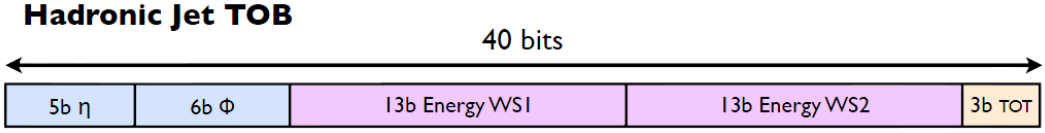

TOB contents	No. Bits	Comments
η	5	Local η coordinate, within core area of jFEX
Φ	6	Global Φ coordinate
Energy WS1	13	Energy computed for window size 1
Energy WS2	13	Energy computed for window size 2
Trigger Object Type (TOT)	3	Trigger object : Had. Jet, Fat Jet, tau
TOTAL	**40**	

Figure 6. Draft jet TOB Format.

Fat Jet TOB

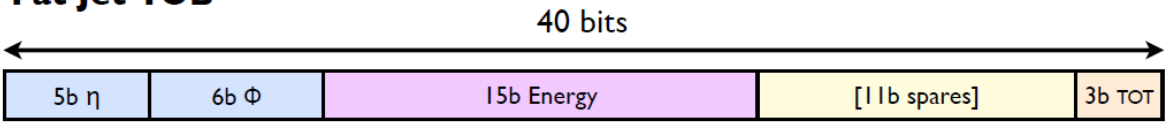

TOB contents	No. Bits	Comments
η	5	Local η coordinate, within core area of jFEX
Φ	6	Global Φ coordinate
Energy	15	Energy
[spares]	11	Spare bits TBD final use
Trigger Object Type (TOT)	3	Trigger object : Had. Jet, Fat Jet, tau
TOTAL	**40**	

Figure 7. Draft fat jet TOB Format.

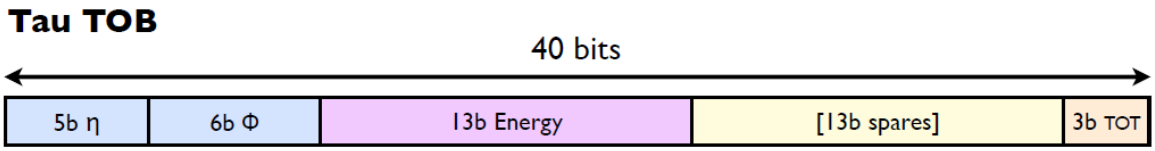

TOB contents	No. Bits	Comments
η	5	Local η coordinate, within core area of jFEX
Φ	6	Global Φ coordinate
Energy	13	Energy
[spares]	13	Spare bits TBD final use
Trigger Object Type (TOT)	3	Trigger object : Had. Jet, Fat Jet, tau
TOTAL	**40**	

Figure 8. Draft τ TOB format.

5.3 Readout data

On receipt of an L1A, the jFEX transmits to the ROD a packet of data of the format shown in Figure 9. This packet contains up to three types of data: TOBs, XTOBs and input data (see below). The TOBs are exact copies of those output from the jFEX on the real-time path. The input data are copies of the calorimeter data as received in the Processor FPGAs. The XTOBs are words that contain more information about trigger candidates than can be transmitted on the real-time data path. If the readout of XTOBs is enabled, any TOB in the readout data will have a corresponding XTOB. The readout data may also contain XTOBs for which there is no corresponding TOB. Such XTOBs describe trigger candidates for which TOBs have not been transmitted to L1Topo because of the input bandwidth limit of that module. The exact format of the XTOBs is yet to be determined. Preliminary assumptions introduce a width of up to 64 bits.

The data in the readout packet are from a programmable window of bunch crossings. The size of this window is the same for all types of data and is limited by the available memory in the Processor FPGAs. The size can be set via control parameters.

Within the packet, the data are organised first according to type, and then according to bunch crossing. Headers mark the boundaries between data types and bunch crossings. Not every type of data is necessarily present in a packet. If a data type is absent, then the headers for that data are also absent. In the extreme, the packet may contain no data, in which case just the packet header and footer are transmitted.

The jFEX readout packets are transmitted to the ROD via six links at up to 10 Gb/s per link, using a link-layer protocol that is to be defined.

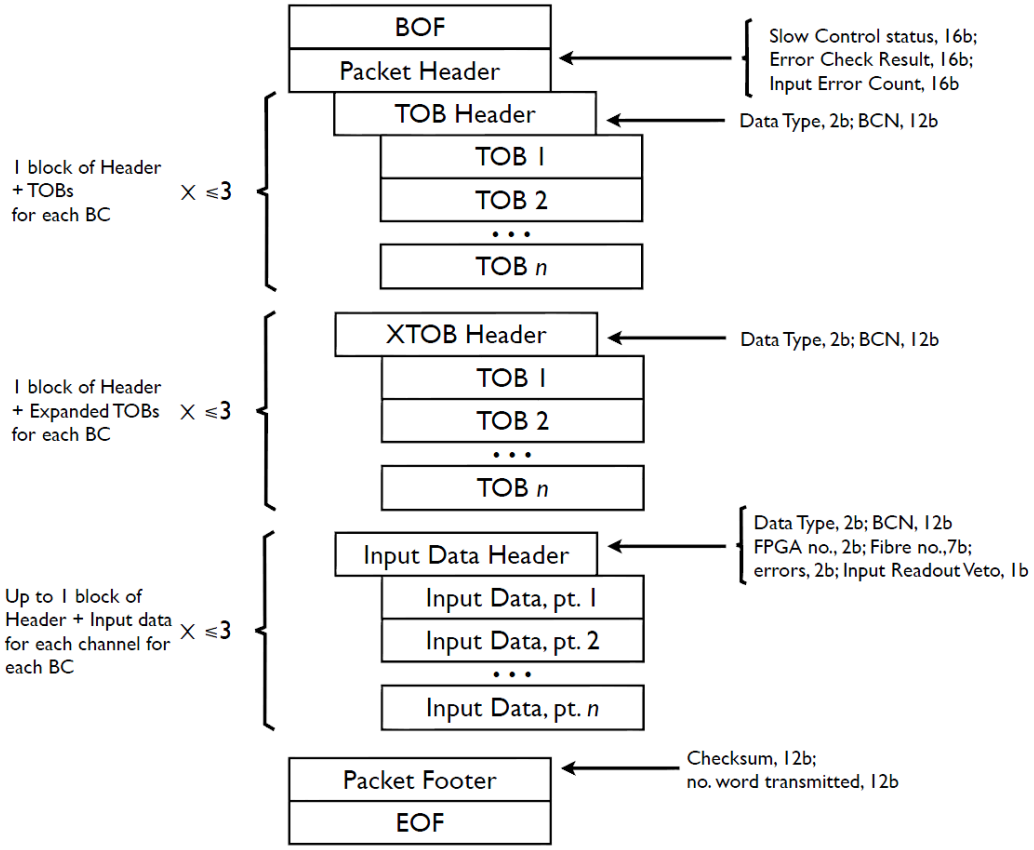

Figure 9. A provisional format for a readout data packet.

6 Implementation

The description of the implementation is based on jFEX modules in the central region. Details of the implementation differ on modules covering the outer regions, due to changes in the input data granularity and the η coverage.

6.1 Input Data Reception and Fan Out

The jFEX receives data from the calorimeters via optical fibres. Each fibre carries data from an area of 0.4×0.4 (η, ϕ). In order to cover an area as described in section 5.1, a single jFEX module must receive data on up to 192 fibres. Two modules require up to 16 additional links, carrying the data from the Tile-HEC overlap, making a total of 208.

The input fibres to the jFEX are organised into 18 ribbons of 12 fibres each. They are routed to the jFEX via the rear of the ATCA shelf, where a rear transition module provides mechanical support. Optical connections between the fibres and the jFEX are made by up to four 72-way Multi-fibre Push-On/Pull-Off (MPO) connectors, mounted in Zone 3 of the ATCA backplane. These connectors allow the jFEX to be inserted into, and extracted from, the shelf without the need to handle individual ribbon connections.

On the jFEX side of the MPO connectors, 18 optical ribbons (each comprising 12 fibres) carry the signals to 18 Avago MiniPOD receivers. These perform optical to electric conversion. They are mounted on board, around the Processor FPGAs, to minimise the length

of the multi-Gb/s PCB tracks required to transmit their output. If the positioning on the PCB does not allow using MiniPODs , the smaller MicroPODs are used instead.

All of the signals received from the calorimeters are transmitted to two Processor FPGAs. The Processor FPGA, which has a core region that covers the region in ϕ from which the data on a fibre is originated, receives the incoming signal. The data is retransmitted to one of the neighbouring Processor FPGAs via "PMA loopback". Once the signal has been received by the FPGA and equalisation has been performed, but before the signal has been decoded, it is sent from the high-speed receiver to the paired high-speed transmitter. There is a latency penalty of >25 ns and some degradation of signal quality associated with this method. The jFEX must therefore handle upwards of 416 multi-Gb/s signals.

6.2 Processor FPGA

There are four Processor FPGAs on the jFEX. The functionality they implement can be grouped into real-time, readout and slow-control functions. All Processor FPGAs on a jFEX module have the same functionality. The differences between the Processor FPGAs on different modules are caused by the varying core areas covered by a certain module and are implemented via different firmwares.

Every Processor FPGA performs the following real-time functions.

- It receives, from MiniPOD optical receivers, up to 104 inputs of serial data at 12.8 Gb/s and additional data from neighbouring FPGAs on 1Gb/s differential links. These carry data from the calorimeters, from an environment of 2.8×3.6 (3.9×3.6 in forward region).
- It applies the feature-identification algorithms described in section 4.1.2 to the calorimeter data, to identify and characterise jet and τ objects and calculate global values.
- For each jet and τ object found, it produces a TOB, as described in section 5.2.
- It prioritises the TOBs, and if the number it has found exceeds the number that can be transmitted to the Merger FPGA in one BC, the excess TOBs are supressed.
- It transmits its TOB results to the Merger FPGA via 48 differential signal pairs at a bandwidth of 1Gb/s per pair.

Each Processor FPGA can process a core area of calorimeter data of 0.8×1.6 (2.9×1.6 in the forward regions). The differences between the modules depending on their covered η range are implemented via firmware. The hardware is the same for all modules.

On the readout path (described in section 4.2), each Processor FPGA performs the following functions.

- The Processor FPGA records the input data and the TOBs generated on the real-time path in scrolling memories, for a programmable duration of up to 3 μs.
- On receipt of an L1A, it writes data from the scrolling memories to the FIFOs, for a programmable time frame. This is only done for those data enabled for readout by the control parameters.

- The Processor FPGA transmits data from the readout FIFOs to the Merger FPGA, via a 12.8GB/s MGT link.

For slow control and monitoring, each Processor FPGA contains a local IPBus interface, which provides access to registers and RAM space within the FPGAs.

The Processor FPGA is a Xilinx XCVU190. The dominant factor in the choice of device is the available number of multi-Gb/s receivers, low latency parallel links and logic resources.

Of the 120 high speed links available in the XCVU190, depending on the covered η range, 80 to 104 are used. The spare resources can be used for slow control functions, to further reduce the required number of parallel links.

Regarding general-purpose I/O, of the 520 pins available, a maximum of 96 are required to transport real-time output data, 84 are required to transmit/receive pile-up sums, 56 are required to transmit/receive additional data from the extended environment, 50 are planned for slow control functions. The remaining 258 pins are left for spare resources. In case of a latency penalty at high bit rates, these spare resources can be used to decrease the transmission speed on latency critical connections. Up to 96 additional pins are required if different data are to be sent to the two FPGAs on an L1Topo module.

6.3 Merger FPGA

The fifth FPGA has multiple tasks, which can be divided into the following categories

- Result merging: It merges the results from the from Processor FPGAs and sends the real-time output to L1Topo
- Readout: It implements the final stages of the jFEX readout path and it implements the interface between the jFEX and the local TTC network on the shelf backplane.
- Control: It provides the interface between the jFEX and the IPBus network on the backplane.

The exact device used for the Merger FPGA has still to be evaluated, but preliminary investigations suggest a (compared to the Processor FPGAs) smaller Xilinx UltraScale FPGA is suitable.

6.3.1 Result merging

The Merger FPGA implements the final stages of the jFEX real-time output path.

- It receives TOBs and global parameters from the Processor FPGAs, via a total of 384 differential signal pairs, at a bandwidth of 1Gb/s per pair.
- It prioritises the TOBs from the Processor FPGAs. If the number of TOBs found by the jFEX exceeds the number that can be transmitted to L1Topo in one BC, it supresses the excess TOBs. The algorithms to be used here are to be defined.
- It calculates module based results from the Processor FPGA based values of parameters E_T and $E_T{}^{miss}$.
- It duplicates the final set of TOBs internally.

- It transmits all copies of the set of TOBs to MiniPOD electrical–optical transmitters. Each set of TOBs is transmitted over its own high-speed link at 12.8 Gb/s per link.

6.3.2 Readout

With regard to the readout path (which is described in section 4.2), the functions of the Merger FPGA are as follows.

- It monitors status flags from the readout FIFOs in the Processor FPGAs.
- If the FIFOs contain the data for at least one complete event, the Merger FPGA initiates the building of a readout packet.
- It reads from the FIFOs all data required for that event, as indicated by the control parameters (Input Readout Mode, Input Readout Veto, etc.).
- It builds the readout event packet.
- It transmits packets of readout data to the ROD, via the shelf backplane, from six multi-Gb/s transceivers.

On the prototype jFEX, data are received by the Merger FPGA from the Processor FPGAs via a total of four 12.8-Gb/s links (one from each Processor FPGA). The maximum, total input bandwidth is thus 51.2 Gb/s. Data are output from the Merger FPGA on six multi-Gb/s transceivers, at a maximum line rate of up to 10 Gb/s per transceiver. The maximum data rate out of the Merger FPGA on the readout path, aggregated over the six transmitters and ignoring framing information, is 60 Gb/s. Assuming 8B/10B encoding, this gives a payload of 48 Gb/s.

With regard to the TTC interface, TTC commands and information are transmitted within the jFEX shelf using a local protocol. This isolates the jFEX module from any changes in the ATLAS TTC system between Run 3 and Run 4. The local TTC protocol is yet to be defined, but it is estimated to require much less bandwidth than the 1 Gb/s allocated. The functionality of the Merger FPGA with respect to the TTC interface is as follows.

- Control commands, such as Event Counter Reset, Bunch Counter Reset and L1A, are received from the shelf backplane via a single link of 1 Gb/s.
- These commands are decoded and passed to the relevant hardware on the jFEX via individual control lines.
- Any information requested from the jFEX hardware is received by the Merger FPGA on individual status lines. There it is built into a packet and transmitted over the shelf TTC network.

In addition to the above, the Merger FPGA contains a local IPBus interface, which provides control and monitoring access to registers and RAM space within the FPGA.

6.3.3 Control

The Merger FPGA provides the interface between the jFEX and the IPBus network on the backplane. It contains those control and status registers that concern the operation of the jFEX as whole (i.e., all registers not specific to one particular FPGA), plus a switch to direct

IPBus traffic to and from those registers and IPBus-accessible RAM blocks that are implemented in the other FPGAs.

In addition to high-level control, IPBus provides a pathway for transmitting environment data (on temperatures and voltages) from the jFEX to the ATLAS DCS system. This traffic is also handled by the Merger FPGA, which routes packets between the DCS system and the IPM Controller on the jFEX.

6.4 Clocking

There are two types of clock sources on jFEX: on-board crystal clocks and the LHC TTC clock, received from the ATCA backplane. These clock sources are fed via the clocking circuitry to five FPGAs. The 40.08MHz TTC clock has too much jitter to drive multi-Gb/s links directly. A PLL chip is therefore used to clean up the jitter on this clock. From the input of 40.08 MHz the PLL chip can generate clocks of frequency $n \times 40.08$ MHz within a certain range. This flexibility allows the multi-Gb/s links on the jFEX to be driven at a large range of different rates. The TI CDCE62005 has been tested and verified on the High-Speed Demonstrator [1.8] and thus is considered an option for the jFEX. Another option is the Si5326 which is currently used on the L1Topo.

To facilitate standalone tests of the high-speed serial links on the jFEX, an on-board crystal clock of 40.08MHz is also provided.

The reference clocks for the MGTs on both the Processor FPGAs and the Merger FPGA are driven by PLL chips. The readout links will probably run at a slower speed than the real-time links, as they are copper links over an ATCA backplane. Therefore, the readout links and real-time links are driven with separate PLL chips.

A 125MHz crystal clock is provided for the Merger FPGA for its Gigabit Ethernet interface to the shelf IPBus network. On the jFEX, the protocol between the IPBus master (in the Merger FPGA) and the IPBus slaves (in the other FPGAs) runs using this clock. Hence, the jFEX module control function over IPBus is independent from the TTC clock domain.

The 40.08MHz clock and its multiples (e.g. 160.32MHz or 320.64MHz) from a PLL chip are also connected to the global clock inputs of all the Processor FPGAs and the Merger FPGA.

6.5 High-Speed signals on the PCB

The jFEX is a very high-speed and very high-density ATCA module, which has about 450 optical fibre links running at a speed of 12.8Gb/s, and many copper links running up to 12.8Gbps over the backplane. In addition, the tight ATLAS L1Calo latency margin requires hundreds of parallel links running at up to 1Gb/s between FPGAs for results merging and data sharing on the jFEX.

Signal integrity is a big challenge for the jFEX design. The designing will be accompanied by detailed PCB simulations.

6.6 FPGA configuration

The jFEX houses five big FPGAs: four Processor FPGAs and the Merger FPGA. The configuration of these FPGAs is done using a small device (microcontroller or another FPGA). This device contains an integral flash memory from which it loads its configuration data on power up. (These data are downloaded to this memory during commissioning of the jFEX, via the JTAG Boundary Scan port.) In case an additional FPGA is used as the Configurator, the firmware loaded into the Configurator is Xilinx System ACE *SD Controller* IP. Once configured it becomes a System ACE controller, responsible for the configuration process of the other FPGAs on the jFEX. It initiates this process as soon as it, itself, is configured.

The configuration data for the five FPGAs are stored on the jFEX in a micro SD flash card. They are stored as collections of firmware, where one collection comprises one firmware load for each FPGA on the jFEX, excluding the Configurator. Up to eight firmware collections can be stored on the jFEX and handled by the Configurator. The collections are enumerated and by default collection zero is loaded into the FPGAs. This choice can be over-written by IPBus. Currently, only two firmware collections are foreseen for the jFEX: the normal, running-mode firmware and a diagnostic collection. Extra capacity for a further six collections is spare. The configuration data stored in the micro flash SD card can be updated via IPBus.

Re-configuration of the FPGAs can be initiated via IPBus and via the low-level management IPMI bus. The Configurator must be re-configured separately from the other FPGAs, which must be re-configured as a group. As the IPBus interface is implemented in the Merger FPGA, and the firmware of the Merger FPGA can be updated over IPBus, it is possible, by uploading bad firmware, to place the jFEX in a non-working state from which it cannot be recovered via IPBus. For this reason a firmware collection that is known to work should always be kept in the micro flash SD. This ensures it is always possible to restore the jFEX to a working state via the IPMI bus.

6.7 The IPM Controller

For the purposes of monitoring and controlling the power, cooling and interconnections of a module, the ATCA specification defines a low-level hardware management service based on the Intelligent Platform Management Interface standard (IPMI). The Intelligent Platform Management (IPM) Controller is that portion of a module (in this case, the jFEX) that provides the local interface to the shelf manager via the IPMI bus. It is responsible for the following functions:

- interfacing to the shelf manager via dual, redundant Intelligent Platform Management Buses (IPMBs), it receives messages on all enabled IPMBs;
- negotiating the jFEX power budget with the shelf manager and powering the Payload hardware only once this is completed (see section 6.8);
- managing the operational state of the jFEX, handling activations and deactivations, hot-swap events and failure modes;
- implementing electronic keying, enabling only those backplane interconnects that are compatible with other modules in shelf, as directed by the shelf manager;

- providing to the Shelf Manager hardware information, such as the module serial number and the capabilities of each port on backplane;
- collecting, via an I^2C bus, data on voltages and temperatures from sensors on the jFEX, and sending these data, via IPBus, to the Merger FPGA;
- driving the ATCA-defined LEDs.

The jFEX uses the IPMC mezzanine produced by LAPP as the IPM Controller [1.9] . The form factor of this mezzanine is DDR3 VLP Mini-DIMM.

6.8 Power Management

With regard to power, the hardware on the jFEX is split into two domains: Management hardware and Payload hardware. The Management hardware comprises the IPM Controller plus the DC-DC converters and the non-volatile storage that this requires. By default, on power up, only the Management hardware of the jFEX is powered (drawing no more than 10 W), until the IPM Controller has negotiated power-up rights for the Payload hardware with the shelf manager. This is in accordance with the ATCA specification. However, via a hardware switch it is also possible to place the jFEX in a mode where the Payload logic is powered without waiting for any negotiation with the shelf controller. This feature, which is in violation of the ATCA specification, is provided for diagnostic and commissioning purposes.

On power-up of the Payload hardware, the sequence and timing with which the multiple power rails are turned on can be controlled by the IPM Controller. Alternatively, by setting hardware switches, these rails can be brought up in a default sequence defined by resistor-capacitor networks on the module.

Excluding the optional exception noted above, the jFEX conforms to the full ATCA PICMG® specification (issue 3.0, revision 3.0), with regard to power and power management. This includes implementing hot swap functionality, although this is not expected to be used in the trigger system.

Power is supplied to the jFEX on dual, redundant -48V DC feeds. Two Emerson ATC250 (or similar) convertors accept these feeds and provide a power supply of 3.3 V to the Management hardware, and a supply of 12V to the Payload hardware. This 12V supply is stepped down further, by multiple switch-mode regulators, to supply the multiplicity of voltages required by the payload hardware.

For the power supplies to the FPGA multi-Gigabit transceivers, the PCB design guidelines and noise requirements specified in the UltraScale Series FPGAs GTH Transceiver User Guide (UG576) and GTY Transceiver User Guide (not yet available) will be observed.

6.9 Front-panel Inputs and Outputs

The following signals are, or can be, input to the jFEX via the front panel.

- Auxiliary clock. This input allows the jFEX to be driven by an external 40 MHz clock, in the absence of a suitable clock on the backplane. The optimum physical form factor for the signal is to be identified.

The following bi-directional control interfaces are available on the front panel. See section 6.12 for the use of these interfaces.

- JTAG Boundary Scan. The optimum physical form factor for this interface is to be identified.
- 1G Ethernet socket.

6.10 Rear-panel Inputs and Outputs

6.10.1 ATCA Zone 1

This interface is configured according to the ATCA standard. The connections include

- dual, redundant -48V power supplies,
- hardware address,
- IPMB ports A and B (to the Hub module),
- shelf ground,
- logic ground.

Figure 10 shows the backplane connections between the jFEX and the Hub module, which are located in Zones 1 and 2 of the ATCA backplane. See the ATCA specification for further details.

6.10.2 ATCA Zone 2

6.10.2.1 Base Interface

The Base Interface comprises eight differential pairs. Four of these are connected to hub slot one and are used for module control, the other four are connected to hub slot two and are used to carry DCS traffic. Both of these functions are implemented using IPBus, running over 1G Ethernet links.

6.10.2.2 Fabric Interface

The Fabric Interface comprises 16 differential signal pairs, eight of which are connected to hub slot one, and eight of which are connected to hub slot two. Those signal pairs connected to hub slot one are used as follows:

- One signal pair is used to receive the TTC clock.
- One signal pair is used to receive decoded TTC commands, plus near real-time signals such as ROD busy. The protocol is to be defined. The link speed does not exceed 10 Gb/s.
- Six signal pairs are used to transmit readout data. The link speed does not exceed 10 Gb/s.

Those signal pairs connected to hub slot two are reserved for the same functions as above. Potentially, this allows redundant connections to be made to this hub slot. However, the

firmware necessary to drive and receive data to and from the Fabric Interface of hub slot two is undeveloped.

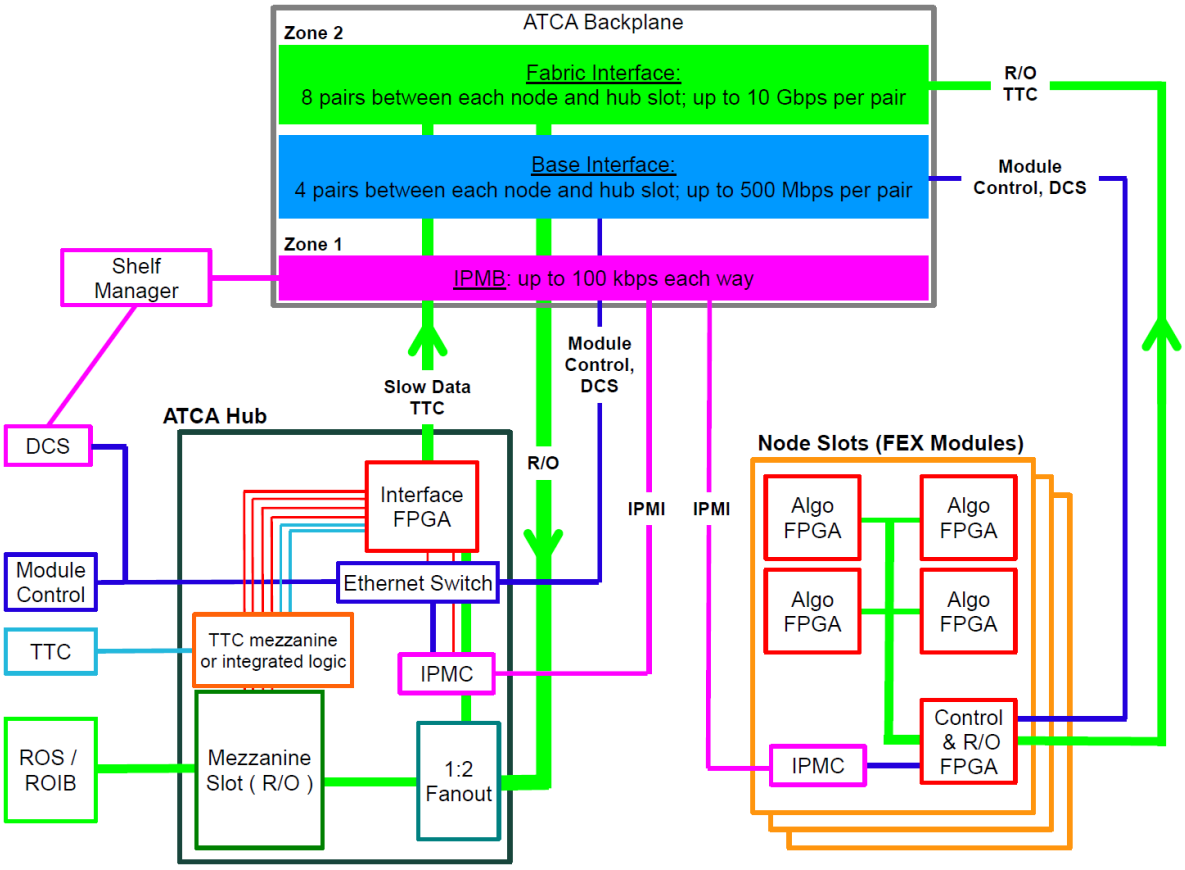

Figure 10. The ATCA backplane connections between the jFEX and the Hub module.

6.10.3 ATCA Zone 3

ATCA zone houses four 72-way optical MPO connectors. Three of these house a total of up to 208 fibres, carrying data from the calorimeters to the jFEX (see section 6.1). At the rear of the MPO connectors, optical fibres carry data from the calorimeters to the jFEX via the L1Calo Optical Plant. These fibres are supported in the jFEX shelf by a (passive, mechanical) rear transition module (RTM). On the jFEX side of the connectors, fibre ribbons carry the calorimeter data to MiniPOD receivers, mounted in board. The optical connections are made on the insertion of the jFEX into the shelf, and broken on its extraction. The fourth MPO connector houses fibres, carrying TOB data from the Merger FPGA to the L1Topo modules.

6.11 LEDs

All LEDs defined in the ATCA specifications are located on the jFEX front panel. In addition, further status LEDs are provided on either the front panel or the top side. These indicate functions like power, Done signals, L1A receipt und further LEDs for diagnostic purposes for all FPGAs.

6.12 Instrument Access Points

6.12.1 Set-Up and Control Points

The following interfaces are provided for the set-up, control and monitoring of the jFEX. They are intended for commissioning and diagnostic use only. During normal operation it should not be necessary to access the jFEX via these interfaces.

- The JTAG Boundary Scan port: via this port a boundary scan test can be conducted, all FPGAs on the jFEX can be configured, the configuration memory of the Configurator can be loaded and the FPGA diagnostic/evaluation tool ChipScope can be run, including for IBERT tests. This port is on the front panel.
- The 1G Ethernet port: this port provides an auxiliary control interface to the jFEX, over which IPBus can be run, should there be a problem with, or in the absence of, an IPBus connection over the shelf backplane. It is on the front panel and connected to the Merger FPGA.
- The RS232 port: this port provides a control interface of last resort, available if all others fail. It is mounted on the top side of the module and connects to the Merger FPGA. Firmware to implement this interface will only be developed if needed.

6.12.2 Signal Test Points

Due to the sensitive nature of multi-Gb/s signals, no test points are provided on PCB tracks intended to carry multi-Gb/s data. If such signals need to be examined, this must be done via firmware. Test points are placed on a selection of those data and control tracks that are not operating at multi-Gb/s.

For each FPGA, spare, general-purpose IO pins are routed to headers. Furthermore, spare multi-Gb/s transmitters and receivers are routed to SMA sockets. With appropriate firmware these connections allow internal signals, or copies of data received, to be fed to an oscilloscope, for example, or driven from external hardware.

The exact number of test connections, and those signals on which a test point can be placed most usefully, are to be determined during schematic entry.

6.12.3 Ground Points

At least six ground points are provided, in exposed areas on the top side of the module, to allow oscilloscope probes to be grounded.

6.13 Floor plan

Figure 11 shows a preliminary floor plan of the jFEX module. This will be used as a guide for the layout process; the exact location of components may change as the physical constraints on the layout are better understood.

The routing of over 400 signals at multi-Gb/s presents a significant challenge for the design of the jFEX PCB. In order to minimise track lengths and routing complexity for these signals, the Avago MiniPOD receivers are placed around the Processor FPGAs. However, this creates an additional constraint on the layout: the need to accommodate routing paths for the fibre-optic ribbons carrying the data to these receivers. To connect the MPO connectors to the receivers the ribbons need to twist, curve and bypass large heat sinks on the FPGAs. It can be seen in Figure 11 that large components have been excluded from some areas of the floor plan, to allow space for the routing of the fibre-optic ribbons.

In addition to those components shown in Figure 11, glue logic is placed on the underside of the module.

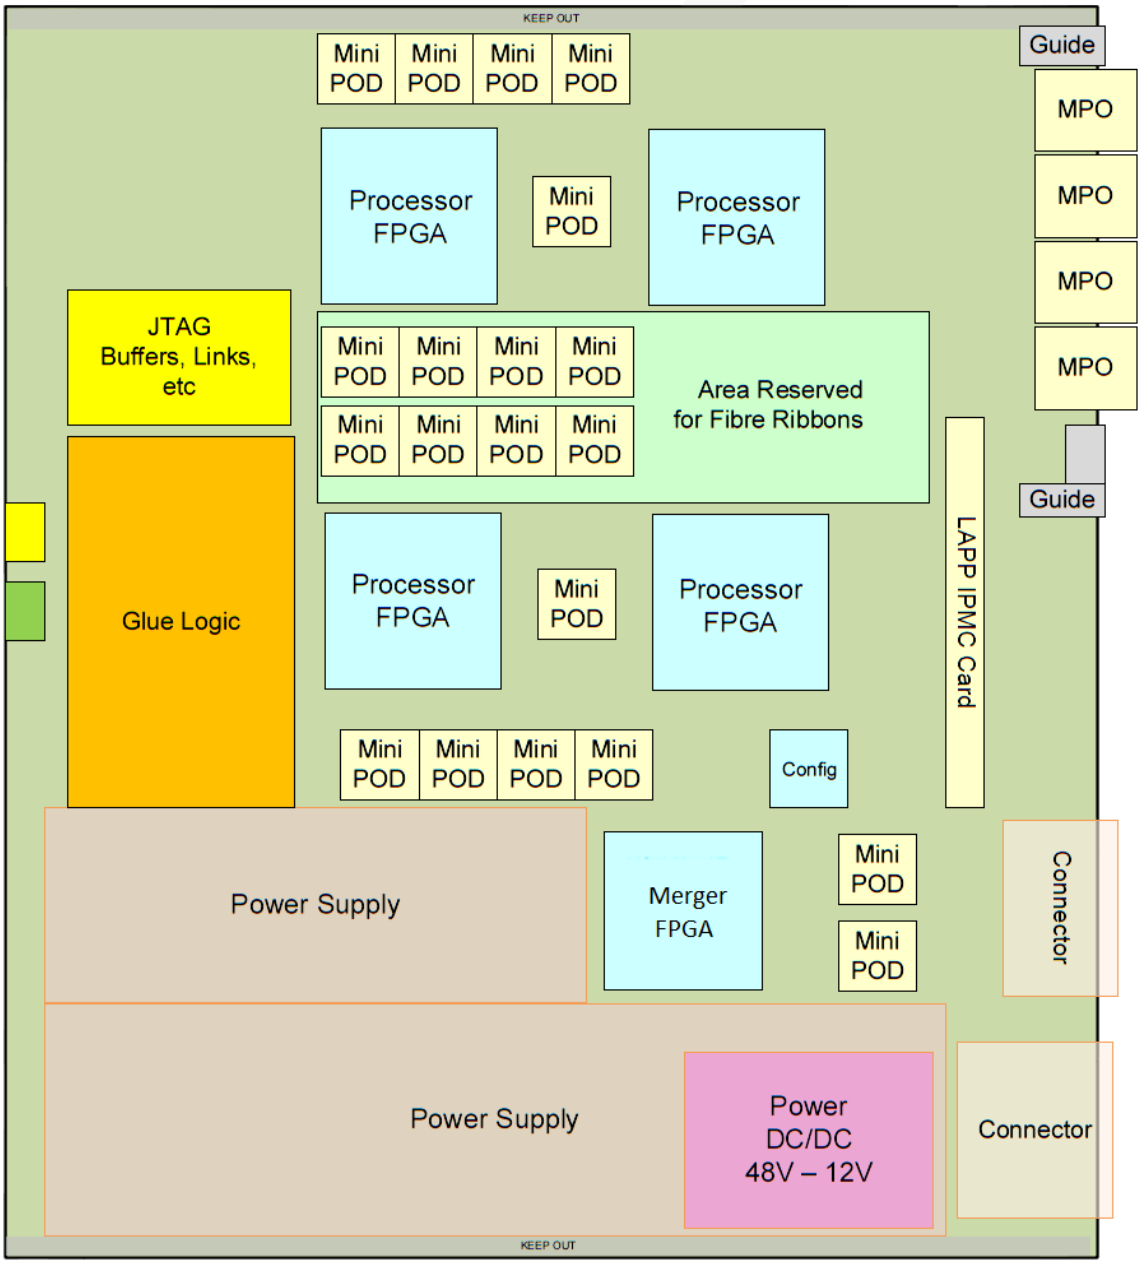

Figure 11. A floor plan of the jFEX, showing a preliminary placement guide.

7 Front-Panel Layout

Figure 12. Preliminary front panel layout (not to scale).

Figure 12 shows a preliminary template for the front panel layout of the jFEX. Shown are the JTAG port for boundary scanning and FPGA access, an auxiliary Ethernet control port, status LEDs and the ATCA extraction/insertion handles. These components are not drawn to scale.

8 Programming model

8.1 Guidelines

The slow-control interface of the jFEX obeys the following rules.

- The system controller can read all registers; there are no ‘write only’ registers.
- Three types of register are defined: Status Registers, Control Registers and Pulse Registers.
- All Status Registers are read-only registers. Their contents can be modified only by the jFEX hardware.
- All Control Registers are read/write registers. Their contents can be modified only by system controller. Reading a Control Register returns the last value written to that register.
- All Pulse Registers are read/write registers. Writing to them generates a pulse for those bits asserted. Reading them returns all bits as zero.
- Attempts to write to read-only registers, or undefined portions of registers, result in the non-modifiable fields being left unchanged.
- If the computer reads a register (e.g. a counter) which the jFEX is modifying, a well-defined value is returned.
- The power-up condition of all registers bits is zero, unless otherwise stated.

8.2 Register Map

The full register map will be developed during the design process and documented here. The following is an incomplete list of the requirements that have been identified thus far.

- Programmable parameters:
  - Enable intermediate TOB readout
  - Input error registers
  - Input mask bits

- Input Readout Mode
- Input Readout Mask
- Input Readout dead time length
- Programmable input delay
- Readout frame length
- Readout Offset, Input data
- Readout Offset, Intermediate Data
- Readout Offset, Final TOBs

- Status words:
  - The Error Check Result,
  - Input Error Count,
  - Input Error Latch
  - Input Error Bit Mask
  - Readout parameters

- Memory access:
  - All dual-port RAM on readout path
  - All FIFOs on readout path (non-destructive, random access).
  - Playback and spy buffers, memory mapped

8.3 Register Descriptions

This section is a place holder, to be completed during the design process.

9 Glossary

ATCA	Advanced Telecommunications Computing Architecture (industry standard).
BC	Bunch Crossing: the period of bunch crossings in the LHC and of the clock provided to ATLAS by the TTC, 24.95 ns.
BCMUX	Bunch-crossing multiplexing: used at the input to the CPM, JEM (from Phase 1) and eFEX, this is a method of time-multiplexing calorimeter data, doubling the number of trigger towers per serial link.
CMX	Common Merger Extended Module.
CP	Cluster Processor: the L1Calo subsystem comprising the CPMs.
CPM	Cluster Processor Module.
DAQ	Data Acquisition.
DCS	Detector Control System: the ATLAS system that monitors and controls physical parameters of the sub-systems of the experiment, such as gas pressure, flow-rate, high voltage settings, low-voltage power supplies, temperatures, leakage currents, etc.
ECAL	The electromagnetic calorimeters of ATLAS, considered as a single system.
eFEX	Electromagnetic Feature Extractor.

FEX	Feature Extractor, referring to either an eFEX or jFEX module or subsystem.
FIFO	A first-in, first-out memory buffer.
FPGA	Field-Programmable Gate Array.
HCAL	The hadronic calorimeters of ATLAS, considered as a single system.
IPBus	An IP-based protocol implementing register-level access over Ethernet for module control and monitoring.
IPMB	Intelligent Platform Management Bus: a standard protocol used in ATCA shelves to implement the lowest-level hardware management bus.
IPM Controller	Intelligent Platform Management Controller: in ATCA systems, that portion of a module (or other intelligent component of the system) that interfaces to the IPMB.
IPMI	Intelligent Platform Management Interface: a specification and mechanism for providing inventory management, monitoring, logging, and control for elements of a computer system. A component of, but not exclusive to, the ATCA standard.
JEM	Jet/Energy Module.
JEP	Jet/Energy Processor: the L1Calo subsystem comprising the JEMs.
jFEX	Jet Feature Extractor.
JTAG	A technique, defined by IEEE 1149.1, for transferring data to/from a device using a serial line that connects all relevant registers sequentially. JTAG stands for Joint Technology Assessment Group.
L0A	In Run 4, the Level-0 trigger accept signal.
L0Calo	In Run 4, the ATLAS Level-0 Calorimeter Trigger.
L1A	The Level-1 trigger accept signal.
L1Calo	The ATLAS Level-1 Calorimeter Trigger.
LHC	Large Hadron Collider.
MGT	As defined by Xilinx, this acronym stands for Multi-Gigabit Transceiver. However, it should be noted that it denotes a multi-gigabit transmitter–receiver pair.
MiniPOD	An embedded, 12-channel optical transmitter or receiver.
MicroPOD	An embedded, 12-channel optical transmitter or receiver, smaller compared to the MiniPOD.
MPO	Multi-fibre Push-On/Pull-Off: a connector for mating two optical fibres.
PMA	Physical Media Attachment: a sub-layer of the physical layer of a network protocol.
ROD	Readout Driver.
RoI	Region of Interest: a geographical region of the experiment, limited in η and ϕ, identified by the Level-1 trigger (during Run 3) as containing candidates for Level-2 trigger objects requiring further information. In Run 4, RoIs are

	used in the same between the Level-0 and Level-1 triggers.
Shelf	A crate of ATCA modules.
SMA	Sub-Miniature version A: a small, coaxial RF connector.
Supercell	LAr calorimeter region formed by combining E_T from a number of cells adjacent in η and ϕ.
TOB	Trigger Object.
TTC	The LHC Timing, Trigger and Control system.
XTOB	Extended Trigger Object. A data packet passed to the readout path, contained more information than can be accommodated on the real-time path.

10 Document History

Version	Comments
0.0	Internal circulation
0.1	L1Calo circulation
0.2	PDR draft