ATLAS Level-1Trigger The jFEX Module

– draft –

# **Project Specification**

# Version 0.0

Date: 08/08/2013 14:36:26

Bruno, Stefan, Uli Universität Mainz

# **Table of Contents**

1	Intro	Introduction				
	1.1	Related projects	3			
	1.2	The jFEX processor module	3			
	1.2.1	Baseline algorithms	4			
	1.2.2	2 Real-time data path	4			
Data reception			5			
Data duplication			5			
	Data processing					
	1.2.3 Clock distribution					
1.2.4 Configuration, monitoring, and control			6			
	Pı	e-configuration access / ATCA compliant slow control	6			
	E	mbedded controller	6			
	FPGA configuration					
	Environmental monitoring					
Module control			7			
	D	AQ and ROI	7			
2	Fund	ctional requirements	7			
	2.1	Signal paths	7			
	2.2	Real-time data reception	7			
	2.3	Real-time data processing	8			
	2.4	Clock distribution	8			
	2.5	Timing distribution	8			
	2.6	Configuration and JTAG	9			
	2.7	Module control	9			
	2.8	DAO and ROI	10			
	2.9	Extension mezzanine (X)	10			
	2.10	Zvng module (Z)	10			
	2.11	IPMC dimm	10			
	2.12	ATCA power brick	11			
3	Imp	ementation	11			
-	3.1	Real-time nath	12			
	3.2	Clock	13			
	3.3	Control	13			
	Mod	ule control	14			
	34	Board level issues: signal integrity power supplies and line impedances	14			
4	Firm	ware on-line software and tests	15			
•	41	Service- and infrastructure firmware	15			
	4.2	Algorithmic firmware	16			
5	Inter	faces : connectors pinouts data formats	17			
5	5 1	Internal interfaces	17			
	5.2	Front panel	17			
	53	Backplane connector layout	17			
	5.5	Interfaces to external systems	17			
	5.5	Data formats	1 x			
6	J.J Ann	endixwill move to separate document?!?!?	18			
0	- лрр 6 1	Checklist for detailed design	10			
	6.2	Closery	20			
	0.2 6.3	Change log	20 20			
	0.5	Change 10g	20			

# 1 Introduction

This document describes the specification for the Level-1 jet feature extractor module (jFEX). Section 1 of this document gives an overview of the module. Section 2 describes the functional requirements. Section 3 details the implementation. Section 5 summarizes the external interfaces.



Figure 1 : The jFEX system

## 1.1 Related projects

The jFEX system is a major processor system within the ATLAS Level-1 calorimeter trigger. The jFEX system will replace the current jet and energy sum trigger (JEP) in run 3. It will find jets and large taus and will calculate energy sums. The jFEX system consists of a single ATCA shelf of electronic modules (blades), eight jFEX modules, and two hub modules. The jFEX system also comprises the real-time input stage: optical fibre plant and rear-transition modules. The jFEX system is shown in Figure 1. It will share a large fraction of modules with the eFEX system. The jFEX module interfaces to a large number of modules and systems.

Phase-0 L1Calo modules <u>http://hepwww.rl.ac.uk/Atlas-L1/Modules/Modules.html</u> L1Topo Hub ROD eFEX

## 1.2 The jFEX processor module

The jFEX processor module is an AdvancedTCA (ATCA) module. The baseline design requires eight modules in a single ATCA shelf. Each jFEX module determines jets and taus within an octant of the ATLAS calorimeters, calculates partial energy sums within the octant and sends the results to the topology processor L1Topo. The jFEX module covers the whole eta range and the respective calorimeter systems: FCALs, end caps and barrel. Environment data required for the jet and tau algorithms is received from neighbouring octants on separate fibres, duplicated at source (see Figure 2, a section in eta is shown, cf. 1.2.2).



Figure 2 : jFEX core (green) and environment data duplicated at source (magenta) and sink (orange)

#### 1.2.1 Baseline algorithms

The jFEX is in charge of finding localized objects in the incoming calorimeter data, jets and large taus. Furthermore it calculates global quantities, the total and the missing transverse energy. The energy summation is performed in segments of the calorimeters separately, to generate data that can be used for pile-up corrections on the topology processor. The processors operate on separate electromagnetic and hadronic trigger towers at a granularity of  $0.1 \times 0.1$  in eta x phi.

The baseline jet and tau algorithms are sliding window algorithms operating on pre-summed electromagnetic + hadronic trigger towers within an environment of  $0.9 \times 0.9$  in eta and phi. That limits the maximum jet size to a radius of 0.45. An advanced option capable of doubling the jet cone size is discussed in a separate document. Within the limitations of environment available to each "core" cell, the definition of a "jet" is flexible. Figure 3a shows a jet definition with ROI (hatched) in use in run 1, though with higher granularity and slightly increased environment (blue). Taus are identified within smaller cones.



Figure 3 : jet algorithms

#### 1.2.2 Real-time data path

ATCA Backplane zone 3 of the jFEX is used for real-time data transmission. The input data enter the jFEX optically through the backplane. The fibres are fed via four optical backplane connectors that carry 72 fibres each. The optical signals are converted to electrical signals in 12-fibre receivers. The electrical high-speed signals (6.4 Gb/s line rate) are routed into six FPGAs, where they are de-serialised in MGT receivers; the parallel data are presented to the FPGA fabric. The processor FPGAs operate on their input data independently and in parallel. They process a calorimeter "core" area of  $1.2 \times 0.8$  in eta x phi each and require an additional processing environment of 0.4 to either side. Therefore a single FPGA will operate on a  $2.0 \times 1.6$  area worth of input data.

Figure 4 shows a conceptual drawing of the jFEX. The real-time processor FPGAs are labelled A to F. They are surrounded by the optical receivers. Non-real-time module control functionality is

implemented on FPGA K, as well as on mezzanine modules (IPMC module I, extension mezzanine). More detailed diagrams are shown in section 3.



Figure 4: the jFEX module

#### **Data reception**

The optical data arrive on the main board on 12-fibre ribbons. Since the backplane connectors support multiples of 12 fibres, the optical signals will be routed via "octopus" cable sets, splitting 72 fibres into groups of 12. The opto-electrical conversion will be performed in Avago MicroPOD 12-channel devices. The opto receivers exhibit programmable pre-emphasis so as to allow for improvement on signal integrity for given track length.

After just a few centimetres of electrical trace length, the multi-gigabit signals are de-serialised in the processor FPGAs. They allow for programmable signal equalization on their inputs. The exact data formats are as yet undefined, though standard 8b/10b encoding is envisaged for purpose of run length limitation and DC balance. The processors are supplied with required bias and termination voltages, as well as suitable reference clocks.

#### **Data duplication**

Data sharing between neighbouring FPGAs is accomplished by "PMA loopback", a scheme that splits a duplicated signal from the incoming data early on in the MGT data path and sends it out of the transmitter section of the same high-speed channel. There is a latency penalty of about one LHC bunch clock associated to this duplication scheme. The data duplicated by PMA loopback are shown in orange colour in Figure 2.

#### **Data processing**

The jet and tau algorithms, as well as energy summation are performed on the processor FPGAs for a core area of 1.2 x 0.8 each. All calculations for all "sliding" windows with an offset of 0.1 in eta and phi are computed concurrently. The jets are disambiguated within each FPGA, double counting of objects overlapping between neighbouring FPGAs are avoided by determination of relative maxima of energy deposition, including the environment data. Resulting jet candidates are reported on parallel

fabric I/O and fed into the merger FPGA which both finds objects within its own calorimeter area and consolidates results from the other processor FPGAs.

Energy sums (Et, Etmiss) are determined within a processor FPGA and are consolidated on the merger FPGA. In addition, eta projections of the total energy sums are calculated separately so as to allow for pile-up corrections downstream, on the L1Topo processor.

The merger FPGA reports the localized objects as well as global quantities to the L1Topo processor. It is able to supply data to up to 3 topology processors, on a total of 24 optical fibres.

#### 1.2.3 Clock distribution

The operation of the real-time data path requires low-jitter clocks throughout the system. For synchronous operation, data transmitters will have to be operated with clean multiples of the LHC bunch clock. Receiver reference clocks may as well be derived from local crystal oscillators, though tight limits on the frequency range will have to be observed. The jFEX module will be designed for 40.08 MHz operation of the real-time data path only.

The fabric clocks run at the LHC bunch clock frequency. The MGTs are clocked off multiples of the LHC clock. The jitter on the MGT bunch clock path is tightly controlled with help of a PLL device. The clock fan-out chips chosen are devices with multiple signal level input compatibility, and output levels suitable for their respective use, either LVDS or CML.

There are two independent clock trees for the fabric clocks into all FPGAs. There is one common MGT clock tree into all FPGAs. It carries the LHC based reference clock. Segmented crystal clock trees allow for different bitrates on the incoming MGT links and will in particular allow for an alternative line rate on tilecal inputs. The control FPGA receives additional local clocks since it handles DAQ, ROI, and control links as well.

The current L1Calo modules receive their LHC bunch clock and associated data through the TTCDec module, based on a TTCRx chip. Future LHC bunch clock distribution will differ from this scheme. The jFEX will receive its timing information through the hub modules, where the timing information will be re-formatted and thus shield the jFEX from any changes in clock definition. The hub clock interface is compatible to Xilinx MGT signals. The detailed protocol is defined by firmware only. The hub links are defined in table x.

#### 1.2.4 Configuration, monitoring, and control

#### Pre-configuration access / ATCA compliant slow control

The jFEX is a purely FPGA-based ATCA module. All communications channels are controlled by programmable logic devices and become functional only after successful device configuration. An initial step of module initialization is performed by an IPMC device. It communicates to the shelf via an I2C port (IPMB) available on all ATCA modules in zone 1. The prospective ATLAS standard (LAPP IPMC / mini-DIMM format) will be mounted on the jFEX. Additionally, the JTAG port in zone 1 will allow for limited access to the module at any time. The signals are routed to the FPGA components via an extension module.

#### **Embedded controller**

A small, auxiliary embedded controller will be available on the jFEX. It is Xilinx Zynq based. It will allow for pre-configuration access and will be able to configure the processor FPGAs. The device will boot off a microSD flash card.

#### **FPGA** configuration

The baseline FPGA configuration scheme on the jFEX is via the Zynq embedded controller. It will configure the FPGAs from data stored on the microSD card. A local SPI memory will allow for an alternative configuration method for the control FPGA. This FPGA does not contain any algorithmic firmware and is therefore not expected to be updated particularly often. The processor FPGAs will be configured off SD flash memory, which is sequenced by the control FPGA. In-situ configuration update will be possible via IP access to the control FPGA. For debug purposes USB/JTAG configuration will be available as well.

#### Environmental monitoring

Board level temperature and voltage monitoring will be a requirement. The default ATCA monitoring and control path is via I2C links in zone 1. The backplane I2C port (IPMB) is connected to the IPMC DIMM. On the jFEX configured FPGAs can be monitored for die temperature and internal supply voltage. This is achieved by the use of their built-in system monitor. These measurements will be complemented by board-level monitoring points connected via I2C. The IPMB based monitoring is meant to be used for basic functionality only. An additional path into the (IPbus based) module control circuitry will complement the ATCA standard monitoring scheme.

#### **Module control**

On standard ATCA modules, IP connectivity is mandatory for module control. This is generally provided by two 10/100/1000 Ethernet ports on the backplane in zone 2 (redundant base interface). The respective lines are wired to the extension module, where the required connectivity can be made to the IPMC, the embedded controller, or to an MGT link on the control FPGA via an SGMII Phy device.

#### **DAQ and ROI**

The jFEX will be able to capture incoming, outgoing and internal real-time data and send them into the readout stream. It will also be able to send region-of-interest data into the 2<sup>nd</sup> level trigger. The processor FPGAs carry latency buffers which are used to store readout and ROI data. Upon receipt of a level-1 accept signal the data are forwarded to de-randomizer buffers, located on the processor FPGAs as well. These data are then serialized onto the MGT backplane links (Zone 2 Fabric Interface) and routed to hub modules in charge of further formatting and compressing the data. Due to lack of handshake protocol the data volume to be read out of the processors will have to be kept well below the bandwidth limits of the backplane channel, so as to avoid any buffer overrun.

# 2 Functional requirements

This section describes the functional requirements only. Specific rules need to be followed with respect to PCB design, component placement and decoupling. These requirements are detailed elsewhere. For requirements on interfaces with other system components, see section 5.

## 2.1 Signal paths

The jFEX is designed for high speed differential signal transmission, both on external and internal interfaces. Two differential signalling standards are employed: LVDS and CML.

## 2.2 Real-time data reception

Real-time data are received optically from the back of the ATCA shelf; they are converted to electrical representation, transmitted to the processors and de-serialised in on-chip MGT circuits.

The requirements with respect to data reception and conditioning are:

- Provide four 72-way MPO/MTP compatible blind-mate fibre-optical backplane connectors in ATCA zone 3
- Route bare fibre bundles to 12-channel opto receivers
- Supply opto-electrical components with appropriately conditioned supply voltages
- Connect single ended CMOS level control lines to the control FPGA
- Provide suitable coupling capacitors for multi-Gigabit links

- Run the signal paths into the processors
- Design the data paths such that transmission rates up to 13Gb/s can be achieved

## 2.3 Real-time data processing

The jFEX processing resources are partitioned into six processors. This is a consequence of limitations on MGT link count on currently available FPGAs. The requirements on processing resources are not currently known. However, the envisaged algorithms are expected to use a small fraction of logic resources only.

The requirements with respect to the processors are:

- Process an area of 1.2 x 0.8 on each processor FPGA, along with +/- 0.4 environment
- Provide an input bandwidth and channel count sufficient to route an area 2.0 x 1.6 worth of calorimeter data into each processor FPGA
- Allow for additional connectivity in the calorimeter overlap region
- Duplicate incoming data to neighbouring FPGAs with PMA loopback scheme
- Process the real-time data : two sliding windows algorithms (jet, tau) plus energy sums
- Transmit the jet and tau candidate data to one of the processor FPGAs on parallel links
- Design the parallel data paths such that transmission rates up to 1Gb/s can be achieved
- Merge the per-FPGA results on one of the FPGAs
- Minimise latency on chip-to-chip data transmission
- Provide 24 output fibres to be routed to L1Topo. Line rate up to 13Gb/s per fibre.

## 2.4 Clock distribution

Both the FPGA fabric and the MGT links need to be supplied with suitable clock signals. Due to the synchronous, pipelined processing scheme most of the FPGA-internal clocks need to be derived from the LHC bunch clock or a device emulating it. Due to requirements on MGT reference clock frequency accuracy, a mixed 40.00/40.08 MHz operation is impossible. Therefore a requirement for 40.08 MHz operation has to be introduced.

The requirements with respect to the clock distribution on the main board are:

- Receive a clean LHC clock signal from the backplane from each of the hub modules on the fabric interface Zone 2
- Allow for selection of either clock as a system clock under firmware control (clock multiplexer)
- Allow for clock conditioning hardware in the clock path (jitter control, multiplication)
- The LHC-based clock must be used for reference on all real-time transmitters (to L1Topo)
- Provide the common LHC-based MGT clock to all FPGAs
- Provide additional crystal-based MGT reference clocks to all processor FPGAs for use on the receiver sections. Allow for segmentation of the clock trees to cope with possibly different line rates from various parts of the calorimetry.
- Connect the MGT clocks to the processor FPGAs such that the central quad of 3 quads is supplied.
- Provide two fabric clocks to all FPGAs (40.08 MHz crystal, LHC clock)
- Provide a separate crystal based MGT clock to the control FPGA for use on the control link (Ethernet)
- Provide a separate crystal based MGT clock to the control FPGA for use on the DAQ and ROI links

## 2.5 Timing command distribution

A small set of timing information is distributed to all detectors and the trigger system via the communications channels of the TTC system. This information comprises the Level-1 Accept signal, bunch count reset, broadcast commands and others. This information is received on the jFEX from the

hub modules via the backplane. The data format is not yet fixed, but the data will basically comprise the full TTC data stream, re-coded on a Xilinx MGT with 8b/10b encoding.

The requirements with respect to timing distribution are:

- Receive a multi-Gigabit data stream on a single signal pair from the backplane fabric interface from each of the hub modules
- Route the signals into the control FPGA
- Allow selection of either hub module as timing data source in firmware
- Extract the required data (L1A, BC reset, broadcasts) and route them on to the processor FPGAs on parallel I/O

## 2.6 Configuration and JTAG

JTAG is used for board level connectivity tests, pre-configuration access, and module configuration. During initial board tests and later hardware failure analysis, JTAG access will be required to connect an automatic boundary scan system, generally when the module is on the bench, not in an ATCA shelf. Also the initial configuration of non-volatile CPLDs will be performed at that stage.

The requirements with respect to boundary scan and CPLD configuration are:

- Allow for the connection of a boundary scan system to all scannable components of the jFEX: FPGAs, CPLDs, and mezzanine modules via standard JTAG headers, following the standard rules on pull-up, series termination, and level translation between supply voltage domains.
- Allow for CPLD (re)configuration, following the boundary scan tests, and occasionally on the completed system.
- There is currently no requirement nor possibility known regarding integration of devices sourcing or sinking MGT signals externally, into the boundary scan scheme

The requirements with respect to FPGA configuration are:

- Allow for static control of the FPGA configuration port settings and read-back of the status via the control CPLD.
- Provide SPI configuration of the control FPGA according to the KC705 configuration scheme
- Connect an SD card to both the control FPGA and the Zynq based controller
- Connect the serial configuration lines of the processor FPGAs to user I/O of the control FPGA and the Zynq device to allow for configuration off flash card data.

## 2.7 Module control

On ATCA modules serial protocols are used to achieve board level configuration and control. Typically Ethernet would be used to control module operation. On the jFEX IPbus is used to interface to the control FPGA, once it has been configured. IPbus is a small footprint IP stack, implemented in firmware, based on standard Ethernet. All board level control goes via the control FPGA. The control FPGA is in turn connected to a Zynq based processor and to an ATLAS standard IPMC for low level control and monitoring.

The requirements with respect to general board control are:

- Provide eight-lane access from the base interface in zone 2 to the extension mezzanine, compatible to 10/100/1000 Ethernet, so as to allow for an Ethernet Phy to be mounted on the extension mezzanine module
- Provide for Ethernet connectivity onwards to IPMC and Zynq
- Provide two-lane access from the mezzanine on to the control FPGA (one MGT, bidirectional)
- Provide bi-directional connectivity between processors and control FPGA
- Provide a control bus from the control FPGA to all opto-electrical transceivers (I2C and static control)
- Provide a single ended and a differential control bus from the control FPGA to the mezzanine module
- Provide an interconnect between control FPGA and control CPLD (via extension module)

The CPLD is in charge of mainly static controls that need to be asserted at power-up, before the FPGAs are configured.

The requirements with respect to the CPLD are:

- Communicate to the general board control system via a bus to the control FPGA
- Control the static FPGA configuration control lines

## 2.8 DAQ and ROI

A single lane for each DAQ and ROI transmission will be provided on each processor FPGA on the jFEX.

The requirements with respect to DAQ and ROI interface are:

- Connect each processor FPGA to the DAQ hub via fabric interface in zone 2
- Connect each processor FPGA to the ROI hub via fabric interface in zone 2
- Run DAQ and ROI Data on the backplane links at 6.4Gb/s, 8b/10b encoded
- Guarantee by choice of readout data that local on-processor buffers do not overflow

## 2.9 Extension mezzanine (X)

The extension mezzanine provides some connectivity and real estate for control purposes. It will be available on the jFEX prototype and might disappear on the production modules, with the respective functionality integrated on the mainboard. The requirements with respect to auxiliary controls on the mezzanine board are:

- Receive two 4-pair Ethernet signals from backplane zone 2 (base interface)
- Connect to two Ethernet ports on the IPMC
- Connect the mezzanine to the control FPGA via a single MGT for purpose of module control
- Connect the mezzanine to the control FPGA via an LVDS level bus
- Connect the mezzanine to the control FPGA via a CMOS bus for purpose of static and slow controls
- Connect the mezzanine to the control CPLD
- Connect the Zynq module via 32-way CMOS level control bus
- Drive the MicroPOD control buses (I2C)

## 2.10 Zynq module (Z)

The Zynq device might reside on a mezzanine or on the main board. Its main use is the SD card based configurator for the FPGAs. It will also be available for auxiliary controls and house keeping

The requirements with respect to the mezzanine board are:

- Wire the module to Ethernet via the extension module
- Connect the module to an SD flash card acting as processor boot device and configuration storage
- Connect a parallel buse to the control FPGA for purpose of module control
- Connect the Zynq processor to a RAM chip and some general I/O

#### 2.11 IPMC dimm

A "standard" ATLAS IPMC controller in mini-dimm format is used.

The requirements with respect to the IPMC mezzanine board are:

- Connect to the ATCA zone-1 control lines
- Connect to the zone-1 JTAG pins via a cable loop
- Connect to Ethernet via the extension module
- Connect power control / alarm on the ATCA power brick
- Supply the IPMC with the management voltage generated on the power bridge

## 2.12 ATCA power brick

A "standard" ATCA power brick is to be used to simplify ATCA compliant power handling.

The requirements with respect to the power brick are:

- 48/12V power brick
- ATCA power control / monitoring on-brick
- Internal handling of hot swap / in-rush control
- Internal management voltage generation

# 3 Implementation

This section describes some implementation details of the jFEX module. This section will be expanded before module final production.

The jFEX is built in ATCA form factor. The main board is built as a ~20 layer PCB. The PCB is approximately 2 mm thick and will fit ATCA standard rails. If required, the board edges would be milled down to the required thickness. An initial floor plan is shown in Figure 5. The real-time data path is indicated.



Figure 5 – The jFEX floor plan

## 3.1 Real-time path

The jFEX will require the largest FPGA devices available on the market. The design is heavily I/O bound, since a large environment will need to be received for each processor along with the "core" data.

Xilinx XC7VX690T will be used on the prototype. At the time of writing it cannot be ruled out that a larger device (XC7VX1140T) will be required to route additional data from the calorimeter overlap region into the FPGA. Since this device is not pin compatible there is a need to sort that issue out before detailed module design is started.

The back panel optical connectors provide a capacity of 288 fibres maximum, if four shrouds are mounted. The back panel connection scheme has been chosen to simplify maintenance on the module. However, in case of problems with the high density fibre connectors, fibre bundles could alternatively be routed to the front panel.

## 3.2 Clock

The clock circuitry comprises various crystal clocks, a jitter cleaner for LHC clock signal conditioning, and several stages of electrical fan-out. Various Micrel device types are used to drive and fan out clocks of LVDS and CML type at low distortion. All Micrel devices are sink terminated on-chip. The jitter cleaner used on the jFEX is a Silicon Labs 5326. It allows for jitter cleaning and frequency synthesis up to multiples of the bunch clock frequency.

A detailed block diagram of the clock path is shown in Figure 6.



Figure 6 – The jFEX clock distribution

## 3.3 Control

A detailed block diagram of control paths is shown in Figure 7.



Figure 7 – The jFEX control paths

#### Module control

The jFEX module control is done via Ethernet access. The required hardware components are available on the jFEX: an SGMII Phy (M88E1111 or similar) is connected between the base interface in Zone 2 and an MGT lane on the control FPGA. The firmware layer on top of the physical layer will be the IPbus suite maintained by the University of Bristol. This scheme is already in use on the topology processor L1Topo.

# 3.4 Board level issues: signal integrity, power supplies, and line impedances

The jFEX is a large, high-density module, carrying large numbers of high-speed signal lines of various signal levels. The module relies on single-ended CMOS, and differential (some or all of LVDS, PECL2.5, CML3.3, and CML2.5) signalling. System noise and signal integrity are crucial factors for successful operation of the module. Noise on operating voltages has to be tightly controlled. To that end, large numbers of decoupling capacitors are required near all active components. FPGAs are particularly prone to generating noise on the supply lines. Their internal SERDES circuitry is highly susceptible to noisy operating voltages, which tend to corrupt their high-speed input signals and compromise the operation of the on-chip PLL, resulting in increased bit error rates. To suppress all

spectral components of the supply noise, a combination of distributed capacitance (power planes) and discrete capacitors in the range of nF to hundreds of  $\mu$ F are required. On the FPGAs there are capacitors included in the packages for decoupling of noise up to highest frequencies.

The jFEX will receive its MGT supply voltages from switching regulators. Current drawn on the MGT supply lines was considered too high for linear regulators, if all links are operated at full speed. The converters are equivalent to the ones used on Xilinx evaluation modules, and the ripple, according to the specifications, is well below the specified limits for the MGTs.

The jFEX base frequency is 40.08 MHz. Parallel differential I/O operates at multiples up to 1Gb/s. Multi-Gigabit (MGT) links operate at 6.4 or 10+ Gb/s. This is only possible on matched-impedance lines. Differential sink termination is used throughout. All FPGA inputs are internally terminated to 100 $\Omega$  or to 50 $\Omega$ ||50 $\Omega$ , according to the manufacturer's guidelines. All lines carrying clock signals must be treated with particular care. There exists a checklist for the detailed module design.

# 4 Firmware, on-line software and tests

The jFEX is an entirely FPGA based module. For both hardware commissioning and operation a set of matching firmware and software will be required. These two phases are well separated and requirements differ considerably. Hardware debug and commissioning will require intimate knowledge of the hardware components and will therefore be in the hands of the hardware designers. Both firmware and software will be restricted to simple, non-OO code. Hardware language is VHDL, software is plain C. GUI based tools are not required and will not be supplied by the hardware designers. Module commissioning from a hardware perspective is considered complete once the external hardware interfaces, board level connectivity, and basic operation of the hardware devices have been verified. The hardware debug and commissioning will involve JTAG/boundary scan, IBERT/ChipScope tests, and firmware/software controlled playback/spy tests with any data source/sink device available. A multi-fibre output test module is currently being designed at Birmingham University. Initially an L1Topo prototype module will be available in the lab for test vector playback on a small number of 12-fibre ports.

IPbus based module control is being developed for the L1Topo processor and the required setup and expertise will be available in the test lab. Also IPMC control will have been explored on L1Topo by the time the jFEX prototype is available for first bench tests. Backplane outputs provided for DAQ and ROI data transmission will be tested for bit errors and link stability only. No specific data formats will be tested on these links.

In parallel to the hardware debug and commissioning, higher level software and firmware will be developed for later operation of the jFEX. It will be possible to re-use existent L1Topo firmware and software to a large extent.

The test environment available in the home lab will allow for simple real-time data path tests only. There is no hardware, software installation, nor expertise available to run any tests involving DAQ/RODs/ROSes. Therefore all system level tests will have to be done in an appropriately equipped L1(Calo) test lab. Currently the L1Calo CERN test rig seems to be the only available option, before the jFEX can be connected up at point 1.

## 4.1 Service- and infrastructure firmware

The correct operation of high-speed (MGT) links and FPGA-internal clock conditioning circuitry is prerequisite for any "physics" firmware running on the recovered data streams. This service- and infrastructure firmware is closely tied to the device family used. To a large extent the code is generated with IP core generator tools.

The bare MGT link instantiations are complemented with code for link error detection and monitoring, as well as playback and spy circuitry for diagnostics purposes. These functions are controlled by the crate controller CPU, via IPbus, as described above.

Real-time data are copied to the DAQ via ROD links upon reception of a Level-1 trigger. These data offer a further means of monitoring the correct functionality of the jFEX. There exists an additional data path to the LVL-2 ROI builders for the purpose of region-of-interest data transmission, as required.

All infrastructure described above has been successfully implemented and operated for all L1Calo / L1Topo processor modules. The detailed description is found in the collection of documents referenced in section 1. All functionality will be re-implemented on the jFEX. The JEM VHDL code is a purely behavioural description and can therefore be re-used with just minor modifications. Some optimization for the chosen FPGA devices might be desirable, but isn't strictly required.

## 4.2 Algorithmic firmware

There are several classes of physics algorithms currently being explored with help of physics simulations. The energy summation is expected to be simple and with low logic resource consumption. The jet sliding window baseline algorithm will be similar to the algorithm used in run1, though at slightly larger jet size and finer granularity.

Further algorithms will be coded and simulated in the participating institutes. It is envisaged to code and evaluate algorithms that are closer to current level-2 and offline jet analysis. There will be a separate document on software and on algorithmic firmware produced at a later date.

# 5 Interfaces : connectors, pinouts, data formats

## 5.1 Internal interfaces

The jFEX prototype module carries three daughter modules: an IPMC module, a Zynq processor, and an extension mezzanine.

The extension card is connected via a 400-way Samtec connector (SEAM on the mezzanine module, SEAF on the jFEX).

The Zynq based processor is a small mezzanine.

The IPMC module is a 244-pin mini-dimm, mechanically compatible to DDR3 memory modules. Documentation, including pinout is available from LAPP, Annecy.

## 5.2 Front panel

The front panel shows is carrying connectivity for both real-time output and non-real-time control:

Standard ATCA handle switches and LEDs (from IPMC) SD card slot 2 MPO connectors (from MicroPODs, to L1Topo)

## 5.3 Backplane connector layout

The backplane connector is made to standard ATCA layout in zones 1 and 2. Zone 3 is populated with four MTP/MPO connectors that connect onto a RTM with hermaphroditic blind-mate shrouds (MTP-CPI).

## 5.4 Interfaces to external systems

The jFEX interfaces to external systems mainly via optical fibres. The physical layer is defined by the fibre-optical components used. On the real-time data path fibre-bundles with multiples of 12 multimode fibres are used. The external connections are made by MTP/MPO connectors. The jFEX is equipped with male MTP connectors.

The calorimeter DPS (and legacy) processors are connected via 72-way fibre assemblies, for the L1Topo outputs 12-way assemblies will be used. Opto-electrical translation on the real-time path is made with MicroPOD 12-channel transceivers. These devices are fully compatible to MiniPOD devices, but offer higher design density. This comes at the expense of custom-designed heat exchange and retention systems. For the higher speed option (documented separately) a reduction in design density is expected such that standard MiniPOD devices might be employed.

Data rates and formats are defined by the FPGAs used for serialization and deserialization. For the jFEX 6.4 Gb/s or 10+ Gb/s input data rates will be supported. For the prototype module any rate up to 13 Gb/s will be supported by the processor FPGAs, though maximum data rate will depend on the exact device type being purchased in 2014. Mini/MicroPOD transceivers will (to current data sheets) limit the data rate to a maximum of 14 Gb/s. It should be noted that the Xilinx UltraSCALE device family is expected to be available in time for final module production. Therefore the production modules might slightly differ from the current design in terms of density and device count.

All real-time data streams are used with standard 8b/10b encoding. The data links are assumed to be active at all time. There are neither idle frames nor any packetizing of data. For reason of link stability and link recovery an option of transmitting zero value data words encoded into comma characters is being considered. This might also simplify initial link start-up. A first attempt has been made to define an encoding scheme for the L1Topo data links. The scheme will be used on the jFEX as well.

The baseline deserialization scheme is relying on a de-multiplexing the 6.4Gb/s data stream to 4 subticks (word slices) of 32 bits wide, transmitted on a 160MHz clock. It is anticipated that data will need to be fully de-multiplexed to the LHC base rate, since the algorithms require correlating all input data from all incoming sub-ticks. Further information might have to be transmitted to identify sub-ticks of the 40.08 MHz LHC bunch ticks. The de-multiplexing and word slice alignment scheme is a firmware option and will not affect the design of the interfacing processors. Compatibility to the data sources will be guaranteed.

I/O	From/to	bandwidth		
Real-time input	DPS/legacy	288 * up to 13 Gb/s	Opto fibre / MTP 72	MicroPOD, 8b/10b
Real-time output	L1Topo	24 * up to 13 Gb/s	Opto fibre / MTP	MicroPOD, 8b/10b
Control		2*1Gb/s	Ethernet electrical, zone 2	
DAQ	D-ROD/hub	6.4 Gb/s	Zone 2	
ROI	R-ROD/hub	6.4 Gb/s	Zone 2	
IPMB			Zone 1	
LHC clock	Hub/TTC		Zone 2	

**Table 1 : external interfaces** 

## 5.5 Data formats

The jFEX module is entirely based on programmable logic devices. Therefore the detailed data formats do not affect the hardware design of the module. At the time of writing this documents data formats on the real-time interfaces are being defined and written up. A preliminary input format ...

# 6 Appendix -- will move to separate document?!?!?

#### 6.1 Checklist for detailed design

Detailed rules regarding signal integrity are to be followed so as to make sure the high desity/high speed module can be built successfully. In addition a few details on signal wiring for FPGA control pins are listed. This list might be expanded for a detailed design review.

The rules with respect to power supply are:

- Use low-noise step-down converters on the module, placed far from susceptible components.
- Use local POL linear regulators for MGT link supplies
- According to the device specifications the following supply voltages need to be applied to the FPGAs: Vccint=1.0+/-0.05V, Vccaux=1.8V
- On all FPGA supply voltages observe the device specific ramp up requirement of 0.2ms to 50ms.
- Run all supply voltages on power planes, facing a ground plane where possible, to provide sufficient distributed capacitance
- Provide at least one local decoupling capacitor for each active component
- For FPGAs, follow the manufacturer's guidelines on staged decoupling capacitors (low ESR) in a range of nF to mF
- Observe the capacitance limitations imposed by the voltage convertors

- Minimise the number of different V<sub>CCO</sub> voltages per FPGA to avoid fragmentation of power planes
- avoid large numbers of vias perforating power and ground planes near critical components

The rules with respect to general I/O connectivity are:

- Tie Vccaux and most bank supplies to 1.8V. A given FPGA is supplied by only one 1.8 V plane.
- Use all processor FPGA banks for LVDS (1.8V) only
- Use HP banks on the control FPGA for LVDS connections to the processor FPGAs and mezzanine modules
- For the control FPGA only: wire a small number of banks for 3.3V single ended operation (HR banks)
- Neither reference voltages nor DCI termination are required on the processor FPGAs. Use respective dual-use pins for I/O purposes
- For the control FPGA HR banks allow for DCI termination on single ended lines

The rules with respect to single ended signalling are:

- Run FPGA configuration and FPGA JTAG clock lines on approximately 50  $\Omega$  point-to-point source terminated lines
- Observe the requirements on overshoot and undershoot limitation, in particular for System ACE and FPGA JTAG and configuration lines. Use slew rate limited or low current signals and/or series termination

The rules with respect to differential signalling are:

- For discrete components, use internally sink-terminated devices throughout. Any non-terminated high-speed devices need to be documented in a separate list.
- Use LVDS on all general-purpose FPGA-FPGA links
- Use LVDS on all GCK clock lines
- Use DC coupling on all LVDS lines
- Design all LVDS interconnect for 1Gb/s signalling rate
- Use CML signalling on all MGT lines, for both data and clocks
- Design all MGT data links for 10Gb/s signalling rate
- Generally use AC coupling on all MGT differential inputs and outputs, for both data and clocks
  - SFP devices might be internally decoupled, microPOD transmitters might have a sufficient common mode range to allow for direct connection
- Use CML on all common clock trees; rather than using AC coupling, observe the signalling voltage and input compatibility rules as outlined by the device manufacturers
- Use AC coupling or suitable receivers when crossing voltage or signal standard domains, except on LVDS
- Place coupling capacitors close to source
- Use bias networks on AC coupled inputs where required
- Route all differential signals on properly terminated,  $100 \Omega$  controlled-impedance lines
- Have all micro strip lines face a ground plane
- Have all strip lines face two ground planes or one ground plane and one non-segmented power plane
- avoid sharply bending signal tracks
- minimise cross talk by running buses as widely spread as possible
- Avoid in-pair skew, in particular for MGT links and clocks
- Make use of device built-in programmable signal inversion to improve routing
- Avoid impedance discontinuities and stubs, in particular on MGT links and clocks

The rules with respect to processor FPGA pre-configuration and configuration control pins are:

- Wire configuration lines for optional JTAG or slave serial configuration
- Allow mode lines M0, M2 to be jumpered to either Vcc or GND. Pre-wire to Vcc
- Connect M1 to the CPLD (GND=JTAG mode, Vcc=slave serial)
- Connect PROGRAM, INIT and DONE lines to the CPLD
- Pullup DONE 330Ω, INIT 4k7 PROGRAM 4k7
- Connect Vccbatt to GND

• Wire DIN, DOUT and CCLK (series terminated) configuration lines to the CPLD

The rules with respect to system monitor pins are:

- Connect DXN, DXP to I2C based monitoring circuits
- Decouple analog power and GND according to UG370 with ferrite beads and wire the system monitor for internal reference (both Vref pins to analog GND)
- Do not use analog sense lines Vn and Vp and connect to analog GND

## 6.2 Glossary

1000BASE-SX	Ethernet optical (multimode) physical layer			
8b/10b	Industry standard data encoding scheme for purpose of DC balance and run length			
	limitation (bit change enforcement)			
ATCA	Advanced TCA, Advanced Telecommunications Computing Architecture			
Avago	Manufacturer of 12-channel opto-electrical transceivers. The Avago transceivers used			
C	on THE JFEX are of MicroPOD type			
Base interface	ATCA compliant backplanes provide pre-defined redundant IP connectivity via			
	Ethernet 10/100/1000 from each slot to two modules in the crate (dual star)			
Fabric interface				
CML	Current Mode Logic, a high-speed differential signalling standard			
СТР	The Central Trigger Processor of ATLAS			
DAQ	Data Acquisition (link).			
L1Topo				
IBERT	Xilinx automated bit error rate test tool for MGTs			
IPMB	Intelligent Platform Management Bus (redundant, IPMB-A and IPMB-B), located in			
	ATCA zone 1			
JEM	Jet – and Energy module, being replaced by jFEX (continues to deliver tile signals)			
LVDS	Low-Voltage Differential Signaling standard			
MGT	Multi-Gigabit Transceiver			
MicroPOD	High density 12-channel opto-electric transceiver			
MPO/MTP	Industry standard optical connector for fibre bundles, here 12-72 fibres			
Phy	A device implementing the physical level (electrical representation) of a network			
	protocol			
Quad	The Virtex Serialiser/Deserialiser circuits are arranged in tiles of four MGT links each			
ROI	Region of Interest, as signaled to 2 <sup>nd</sup> level trigger			
RTDP	Real-time data path, i.e. the data path going to the CTP. Latency critical path.			
RTM	Rear Transition Module (note: ATCA RTMs mate with the front modules directly in			
	Zone 3, not via the backplane)			
SGMII	Serial Gigabit Media Independent Interface (for Ethernet Phy), 1.25Gb/s			
Zone	ATCA connector zones 1 (mainly power), 2 (backplane data links), 3 (RTM			
	connections)			

## 6.3 Change log

2013