Design and testing of the high speed signal densely populated ATLAS calorimeter trigger board dedicate to jet identification

B. Bauss, V. Bücher, R. Degele, H. Herr, S. Rave, E. Rocco, U. Schäfer, J. Souza, S. Tapprogge, M. Weirich
Institut für Physik, Johannes Gutenberg Universität, Mainz Germany
A. Brogna, C. Kahra
Detector Lab, PRISMA cluster of Excellence, Mainz Germany

ATL-DAQ-PROC-2017-009
13 March 2017

***Abstract*—The ATLAS experiment has planned a major upgrade in view of the enhanced luminosity of the beam delivered by the Large Hadron Collider (LHC) in 2021. As part of this, the trigger at Level-1 based on calorimeter data will be upgraded to exploit fine-granularity readout using a new system of Feature Extractors (three in total), which each uses different physics objects for the trigger selection. The contribution focusses on the jet Feature EXtractor (jFEX) prototype. Up to a data volume of 2 TB/s has to be processed to provide jet identification (including large area jets) and measurements of global variables within few hundred nanoseconds latency budget. Such requirements translate into the use of large Field Programmable Gate Array (FPGA) with the largest number of Multi Gigabit Transceivers (MGTs) available on the market. The jFEX board prototype hosts four large FPGAs from the Xilinx Ultrascale family with 120 MGTs each, connected to 24 opto-electrical devices, resulting in a densely populated high speed signal board. MEGTRON6 was chosen as the material for the 24 layers jFEX board stack-up because of its property of low transmission loss for high frequency signals (GHz range) and to further preserve the signal integrity special care has been put into the design accompanied by simulation to optimise the voltage drop and minimise the current density over the power planes. The jFEX prototype was delivered at the beginning of December and the preliminary results on the design validation and board characterisation will be reported.**

I. Introduction

After the Long Shutdown 2 (LS2) of the accelerator complex at CERN, the Large Hadron Collider (LHC) will restart with an increased luminosity of 2.5 $\times$ 10^{34} cm^{-2}s^{-1}. The ATLAS experiment [1] has planned a detector upgrade for the Phase-I period (from 2021 to 2026) to address the new challenging accelerator machine conditions and to maintain the sensitivity to electroweak physics without being affected by the increased number of pile-up events. In this contribution, the focus will be on the upgrade of the ATLAS Level-1 Calorimeter system [2]. Fig. 1 (taken from [2]) shows the pipelined Level-1 Calorimeter system (L1Calo) as it will look like at the beginning of Run 3. It will be composed of the current legacy system (in green), which will stay until the new system (in yellow) is fully commissioned. The new system is composed of an optical plant (FOX) that distributes digitised data from the calorimeters to the Feature EXtractors (FEXs). The FEXs differ in the physics objects used for the trigger selection and they are the: electron FEX (eFEX) [3], the jet FEX (jFEX) and the global FEX (gFEX) [4].

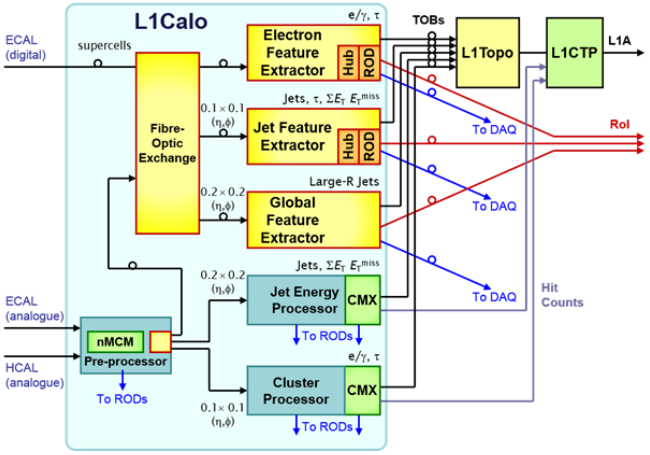

Fig. 1. Level-1 Calorimeter trigger system at the beginning of LHC Run 3. New components are shown in yellow. The new system will run in parallel with the current legacy system until the new system will be fully commissioned.

II. jFEX requirements and features

The jFEX board will identify jet and large-area tau candidates and calculate ΣE_T and E_T^{miss}. It will receive data from the central and forward calorimeter with various granularities over the full η range of $\pm$ 4.9. The input bandwidth has been maximised so that the jFEX can exploit the calorimeter information at the granularity of 0.1 $\times$ 0.1 in $\eta \times \phi$ whilst maintaining the capability of large area jet clustering. Several algorithms for the jet identification will run on the jFEX, with the current baseline algorithm being the "*Sliding window jet algorithm with Gaussian weighting*", which identifies the local maximum with the sliding window algorithm and for each local maximum summing the surrounding cells with a Gaussian weighting distribution. Each FPGA covers an area of 2.4 $\times$ 3.2 in $\eta \times \phi$, with the core area of 0.8 $\times$ 1.6. The data are duplicated on the module within the FPGA using the PMA (Physical Medium Attachment) loop-back (see Section

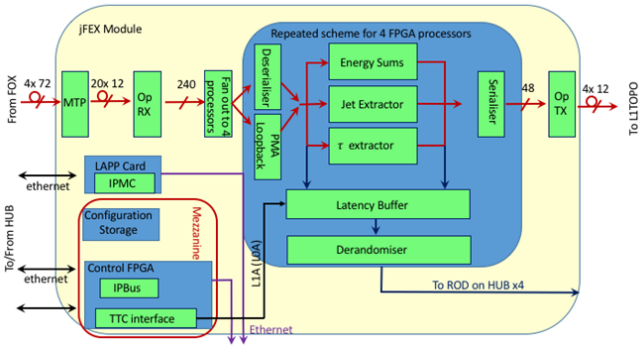

Fig. 2. jFEX block diagram. The sketch shows the functionality of one of the four FPGAs and on bottom left the functionalities of the extension mezzanine. Data and module monitoring are performed through the connection to the HUB and the ROD via backplane.

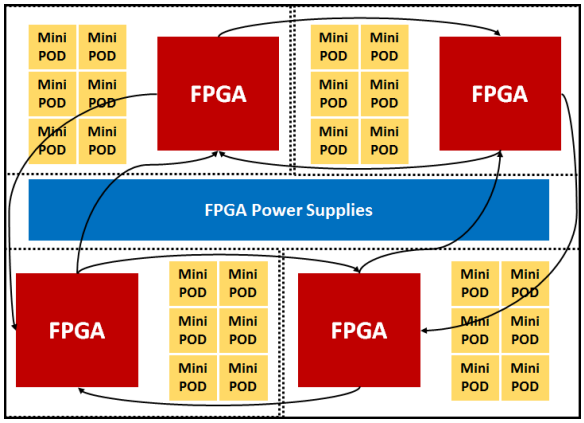

Fig. 3. The jFEX simplified block diagram. The board is divided in four identical blocks, being each consisted of one FPGA and six miniPODs.

IV-A). The need of large processing power is expected due to the complexity of the algorithms that will be implemented.

Fig. 2 shows the jFEX processing block diagram, where all the jFEX functionalities are sketched. The digitised calorimeter data are fed into the jFEX through the Fibre Optical Plant (FOX) and then received on the board through the electro-opto devices, miniPOD [5]. Each module covers the whole ring in ϕ with overlapping regions between processors in η and ϕ duplicated on board using PMA loop-back. The results of jet identification and tau algorithms and the calculation of transverse energy sum and missing transverse energy are sent to the Level-1 Topological Processor board [6] (L1Topo in short) as Trigger Objects (TOBs).

The board control is located on an extension mezzanine (see Section II-B) to allow flexibility in upgrading control functions and components without affecting the main board. This philosophy guarantees smooth and reliable board operation and compatibility with the surrounding trigger system for the next twenty years of data taking.

The data readout checks and external communications are performed via the ROD [2] and the HUB [2].

A. jFEX design challenges

The high input bandwidth and the large processing power required drove the FPGA choice. The Xilinx Ultrascale XCUV190-2FLGA2577 device [8] meets the requirement of the large number of Multi Gigabit Transcievers (MGT) and of the processing power. This device has 120 MGTs, capable to handle up to 3.6 Tbps input bandwidth. The jFEX board hosts four Xilinx Ultrascale and twenty-four miniPODs, twenty receivers and four transmitters. The jFEX board is an ATCA [7] board therefore the routing space available is quite limited taking into account all the components on it.

The challenges faced during the design of the board are summarised as follows:

- signal integrity
- FPGA power estimate
- FPGA power dissipation (not discussed in this contribution)

In order to cope with all requirements and challenges, the jFEX has a modular design which consists of four identical blocks, being defined by one FPGA and six miniPOD. This approach allows a symmetric board design and consequently, an equalization of the length of the traces between processors and miniPOD and also for the FPGA inter-connections. Additionally, the bus lengths are minimised and there are no crossing busses, what improves the signal integrity.

Fig. 3 shows the jFEX design block diagram, with the four FPGA/miniPOD groups highlighted with dashed black squares.

Concerning the FPGA power estimation, a first step was the usage of the Xilinx Power Tool and the performance of the FPGA power plane and thermal simulation (see Section III) taking into account the increase in resource usage due to the jet identification algorithm.

B. Mezzanines

The jFEX board control resides on the extension mezzanine, allowing modifications, implementation of new functionalities and future upgrades without affecting the design of the main board. The current first version of the extension mezzanine hosts several functionalities, such as the possibility to configure the FPGA through the JTAG chain and via the *selectmap* mode, the IPBus master, the reception and the decoding of the Trigger, Timing and Control (TTC) data, spare clocking circuitry that has been used for testing two different types of jitter cleaner (Si5338 and Si5345), I2C buses and the PicoZed [10] plugged into an FCI connector.

To gain more flexibility during the prototype phase, the FPGA power supplies were also mounted on small power cards where the tuning and the optimisation could be done without affecting the design of the main board. After the optimisation of the power modules for ripple and current stability, the power module design will be directly implemented on the main board for the pre-production phase. The requirements on the power module are low ripple and current stability even with sudden changes in the FPGA drawn current so as to not affect the FPGA performance. According to our estimates, the drawn

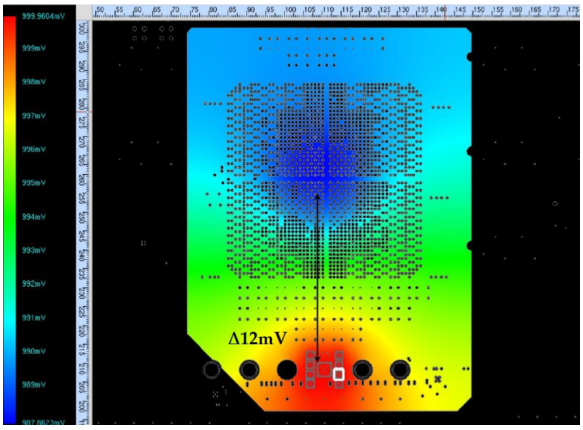

Fig. 4. Simulation of the VCC_int power plane for one single processor. The reported voltage drop between the current pin and the centre of the FPGA is 12 mV.

current on the VCC_int could reach 60A and this requires a DC/DC converter able to deliver such an amount of current.

III. Simulation

The simulation activity was carried out in parallel to the jFEX design to validate the design choices.

Fig. 4 shows the power plane for the VCC_int of one processor. The final shape of the plane has been reached after several simulation iterations, with a special effort on minimising the voltage drop across the FPGA. The maximum voltage drop allowed in the Ultrascale specification is only $\pm$ 20 mV and the expectation with the current power plane design is 12 mV.

The current distribution was also investigated because of high current density spots between 50 - 60 A/mm^2 seen in the simulation. Such spots typically occur on power planes close to the distribution pins. With the help of thermal simulation it was possible to quantify the accumulated power in these points. In all the cases, the accumulated power was of the order of μW or in the worst situations, mW, therefore harmless for the board. The calculated total heat caused by the current flow is, in the worst case, 6.4 °C.

To assess the quality of the high-speed signal, post-layout simulations were performed on the whole board. The single ended and the differential pair tracks for reflected and transmitted signals lay on the regions of *S parameter (dB)* versus the frequency range (GHz) according to the SFP+ specifications.

IV. Measurements

A. PMA loop-back measurement

The jFEX algorithms require data duplication between the processors. The Xilinx evaluation board, VCU110 [12], was used to assess the quality of the PMA loop-back with a similar FPGA but different package compared to that which will be mounted on the jFEX board. Fig. 5 shows the corresponding eye diagram of the IBERT [13] test for a link speed of 28 Gbps with the limit of Bit Error Rate (BER) set at 10^{-13}.

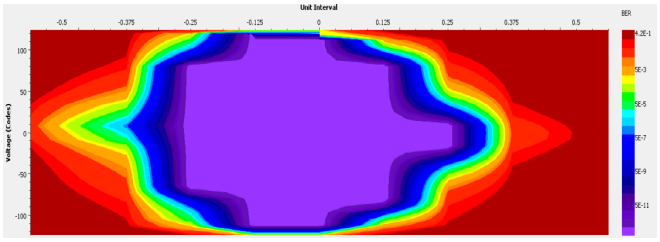

Fig. 5. Eye diagram of the IBERT test set at 10^{-13} BER limit for one link running at 28 Gbps.

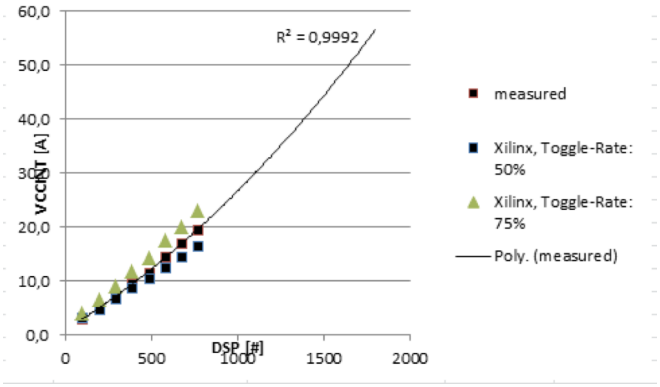

Fig. 6. Plot of the FPGA current consumption on the VCC_int as function of the number of DSPs used. The measured values on the evaluation board are in dark red, while in blue and green are the calculated points by the Vivado tool for two different data toggle rate.

In further tests, the BER limit of 2.25 $\times$ 10^{-16} was reached without error.

The PMA loopback measurement proves the feasibility of the data duplication on the jFEX board and at higher link speed compared to the current baseline link speed for all the three FEXs of 11.2 GBps.

B. FPGA power consumption tests

One of the challenges faced during the jFEX design was to properly assess the FPGA power consumption given that the algorithm might be modified and further optimised in the future. The strategy adopted was to implement different algorithm configurations that foresee different amount of DSP usage in the FPGA of the Xilinx evaluation board, VCU110 to map out the current behaviour as a function of the DSP usage. Fig. 6 shows the comparison between the measurement performed on the evaluation board, namely the dark red dots, and the expected current values for different data toggle rate calculated by the Vivado tool for the implemented firmware.

The data fit of the measured points behaves polynomially and, assuming a possible maximum FPGA usage around 60% corresponding to 1200 DSPs for the implementation of different algorithms on the processors, the expected maximum current is around 35A, validating the choice of the high current DC/DC converter for the VCC_int.

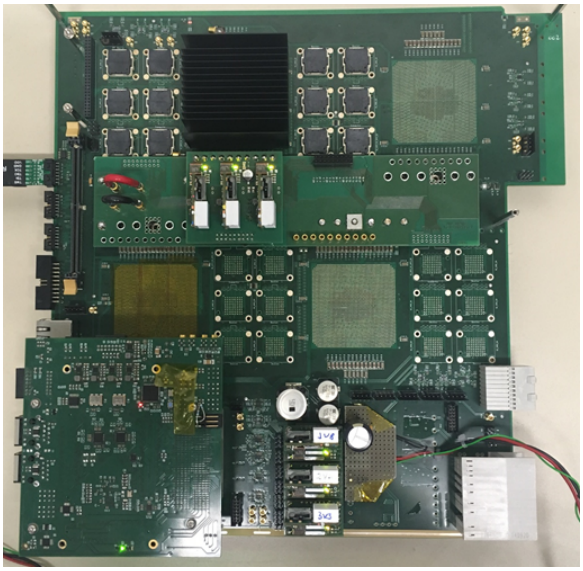

Fig. 7. Picture of the jFEX prototype.

V. jFEX Prototype

The jFEX prototype, see Fig. 7 was delivered at beginning of December 2016. The board hosts one Xilinx Ultrascale and twelve opto-electrical devices, ten receivers and two transmitters. Although the board is not fully assembled, the design validation and prototype characterisation is possible and the signal integrity, the maximum link speed possible (>11.2 Gb/s) and the data duplication on board via PMA loop-back will be assessed. Firmware tests are also planned.

VI. Conclusion

The ATLAS Level-1 Calorimeter trigger will be upgraded for the start of the LHC Run 3 to cope with the higher luminosity of the LHC accelerator whilst preserving the sensitivity to electroweak processes. The jFEX is one of the three new feature extractors that will be installed and it is an ATCA board hosting four Xilinx Ultrascale that are able to handle up to 3.6 Tbps input bandwidth.

Parallel to the design activity there has been quite an extensive and complementary work on power and thermal simulation carried out. As a result no issues with the power plane and thermal performance are expected with the current design. Similarly the post layout simulation of the reflected and transmitted signals show to be in agreement with the SFP+ specification.

To gain flexibility during the prototype phase, the FPGA power supplies are located on four daughter modules and, in turn, on a bigger daughter module plugged on the main board. With the same philosophy the board control is located on an extension mezzanine, where on the first version some functionalities have been added for debugging purpose only. Unlike the FPGA power supply modules, the board control will be always located on the extension mezzanine, allowing possible modifications required for a smooth operation over the future twenty years of data taking.

The sliding window jet algorithm with Gaussian weighting, the baseline algorithm for jet identification, has been already implemented in order to assess the impact of the firmware on the FPGA resource usage. Currently there are two versions available: the first one is meant to be used to measure the current load on the FPGA as function of the DSP consumption to better assess the FPGA power consumption while the second one is an optimised version meant to be used in the final implementation to extract physics results.

The jFEX prototype board has been delivered at the beginning of December. Preliminary test results of design validation and prototype characterisation will be reported.

Acknowledgment

This research has been supported by German Federal Ministry of Education and Research (BMBF) and by the Cluster of Excellence, Precision Physics, Fundamental Interactions and Structure of Matters (PRISMA) funded by the German Research Council (DFG).

The authors would like to thank all the collaborators of ATLAS L1Calo collaboration for the fruitful discussions and suggestions.

References

[1] ATLAS Collaboration, The ATLAS Detector, JINST 3 (2008) S08003
[2] "Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System", CERN-LHCC-2013-018
[3] "Design and test performance of the ATLAS Feature Extractor trigger boards for the Phase-I Upgrade", W. Qian *et al.*, under press in JINST Proceeding section for the TWEPP 2016 conference.
[4] "The Prototype of Global Calorimetric Hardware Trigger for ATLAS at High Luminosity LHC", H. Chen *et al.* contribution to the IEEE NSS and MIC 2016 conference Strasbourg.
[5] www.avagotech.com
[6] "Upgrade of the ATLAS Level-1 Trigger with event topology information", E. Simioni *et al.* J.Phys.: Conf. Ser. (JPCS), Volume 664, 2015
[7] https://www.picmg.org/openstandards/advancedtca/
[8] www.xilinx.com
[9] www.ge.com
[10] https://www.xilinx.com/products/boards-and-kits/1-58nuel.html
[11] "SFF-84331 Specification for Enhanced Small Form Factor Pluggable Module SFP+", ftp://ftp.seagate.com/sff
[12] https://www.xilinx.com/products/boards-and-kits/dk-u1-vcu110-g.html
[13] https://www.xilinx.com/support/documentation/sw_manuals/xilinx2016_3/ug973-vivado-release-notes-install-license.pdf