Report of the Final Design Review

L1Calo Hub Module

Summary

The final design review of the L1Calo Hub module was held at CERN on March 20, 2016. This report summarises the outcome of the review.

Prepared by:	*Checked by:*	*Approved by:*
Stefan Haas (CERN)	**Ian Brawn (RAL)** **Philippe Farthouat (CERN)** **Ed Flaherty (Cambridge)** **Stefan Haas (CERN)** **Lorne Levinson (Weizmann)** **Marek Palka (Jagiellonian University)** **Elena Rocco (Mainz)** **David Sankey (RAL)** **Weiming Qian (RAL)** **David Strom (Oregon)** **Stefano Veneziano (Roma 1)**	**Kevin Einsweiler, LBNL**

for information, you can contact:	**Stefan Haas**	Tel. +41.22.767 9233	E-Mail stefan.haas@cern.ch

PURPOSE OF THIS REVIEW

The purpose of this review was to assess the status of the L1Calo Hub module in view of a proposed PRR in 2018.

MEMBERS OF THE REVIEW COMMITTEE

Review Committee

Ian Brawn (RAL)
Philippe Farthouat (CERN)
Ed Flaherty (Cambridge)
Stefan Haas (CERN)
Lorne Levinson (Weizmann)
Marek Palka (Jagiellonian University)
Elena Rocco (Mainz)
David Sankey (RAL)
Weiming Qian (RAL

Ex-Officio

Kevin Einsweiler, LBNL
David Strom, CERN
Stefano Veneziano (Rome)

AGENDA AND DOCUMENTATION

The agenda and documentation are available in Indico:

https://indico.cern.ch/event/713755

REVIEW OUTCOME

General remarks

The L1Calo hub module is part of the Phase-I upgrade of the trigger system. The module will be installed in each of the two hub slots of the ATCA shelves for the L1Calo eFEX, jFEX sub-systems. The L1topo shelf will host only one Hub module. Therefore the Hub module plays a central role in the shelf and needs to perform the following main functions:

- Receive TTC information from FELIX and distribute the recovered clock signal and the TTC commands to the FEX slots in the shelf.
- Act as an Ethernet switch for the ATCA base interface channels to enable IPbus communication with the FEX slots through the backplane as well as the Ethernet connection to the IPMCs on the node slots.
- Host the ROD mezzanine card. The ROD is designed by another team (Cambridge)
- Forward the backplane high-speed serial links carrying the readout data from the FEX modules to the ROD mezzanine.

An optional feature is the monitoring of the readout data received from the FEX modules through the backplane, however this functionality is not essential.

8 Hub blades have been assembled and tested. Tests have been performed at MSU with a single FEX Test Module (FTM) in a 14-slot ATCA shelf. Some tests were also performed at Cambridge with an eFEX prototype and a FEX Test Module (FTM), but these were done in a 6-slot shelf with a backplane that was not 40G qualified.

Clock distribution

The clock distribution circuit implemented on the Hub has a severe limitation, namely that only a single MGT reference clock frequency is foreseen. The Hub FPGA needs to extract the 40.08 MHz bunch clock from the GBT link coming from FELIX with fixed phase and distribute it to the FEX slots. The current version of the GBT-FPGA IP core developed at CERN, which would typically be used perform this function, requires a 120 MHz or 240 MHz MGT reference clock. In addition, the MGTs that forward the TTC information from the ROD FPGA require a reference clock that is phase and frequency locked to the extracted bunch clock. In order to ensure fixed and deterministic latency for these links operating at 6.4 Gb/s with 8B10B encoding, the MGT reference clock for these links has to be either 4- or 8-times the bunch clock frequency (160.32 MHz or 320.64 MHz respectively) depending on the internal transceiver interface width. Therefore at least two receiver reference clock frequencies are required. The design of the clock distribution should be reconsidered and an additional MGT reference clock frequency source added.

The bunch clock extracted from the FELIX GBT serial data stream also needs to be jitter cleaned and multiplied to be used as a reference clock. The phase of the jitter cleaned clock with respect to the extracted bunch clock has to be fixed and deterministic. All the L1Calo FEX designs use a device from Silicon Labs (Si5344/45) for this purpose. This device ensures a fixed input-to-output phase relationship by supporting a 0-delay mode and also meets the stringent phase noise requirements from Xilinx on the MGT reference clocks. This device was also extensively tested and qualified by the TTC team at CERN and is used in the upgrades of the back-end electronics of other parts of ATLAS as well as in ALICE and LHCb. However the Hub design team selected a different device (Connor Winfield SFX-524G). It was unclear if there is a good technical justification for this choice, or if it was lack of communication with other parts of the L1Calo Phase-I upgrade. No qualification measurements from this new device were presented. It is recommended to change the design to the proven Silicon Labs device, this will also enable the generation of different MGT reference clock frequencies in a programmable manner.

There were some questions concerning the logic levels for the fan-out of the bunch clock over the backplane: the Hub uses an LVDS fan-out chip, however it seems that the eFEX would prefer LVPECL levels. The compatibility of the signal level from the Hub with the eFEX, jFEX and L1Topo should be checked and interoperability tests performed.

Firmware

The firmware of the Hub is still under development. In particular the firmware that interfaces to FELIX in order to receive clock and TTC signals had not yet been implemented. Since this is an essential part of the Hub functionality, this needs to be added and fully tested with the highest priority.

It also appears that quite some effort went into implementing the Aurora 8B10B reception for the GTY transceivers on the Hub FPGA, since this IP core is not directly supported by Xilinx for these MGTs. This is unfortunate, given that monitoring the readout data received from the FEXes over the backplane is just an auxiliary feature of the Hub, It would have been preferable to focus the firmware development effort on essential features instead.

It was unclear if the IPbus functionality had already been fully implemented and tested, this also needs to be demonstrated as soon as possible.

In addition the firmware does not seem to be under revision control yet, the overall L1Calo guidelines on this should be followed.

Backplane Link Testing

Most of the tests appear to have been performed with a single FEX Test Module (FTM) in a 6-slot shelf. Some tests with an eFEX prototype, a FTM and ROD mezzanine have been performed at Cambridge. A Hub Test Module (HTM) is currently being developed in order to be able to populate a 14-slot shelf fully. This activity should be pursued in order to be able to perform tests in a 14-slot shelf with all the links operating concurrently. A comprehensive set of BER and eye-diagram measurements of all the backplane links should be made in this configuration in order to identify any

links that have a degraded performance. The test should be performed with a ROD mezzanine installed in order to also cover the backplane links to the ROD.

PCB

The Hub blade uses a 22-layer PCB. The dielectric material chosen is Isola FR408HR, which is not a dedicated high-speed laminate. For instance the dissipation factor (Df) of FR408HR is about 2.5 times higher than the one of Megtron-6 or Isola I-Tera, which are the dielectric materials used for the FEX blade PCBs. There was no technical justification given for this choice, in particular since the PCB production house is the same than the one used for the gFEX, Megtron-6 would clearly have been a possible choice. The choice of the dielectric should be reconsidered for the pre-production modules.

The routing in the area where the 74 MGT links enter the board from the ATCA backplane is very dense, and the design team itself addressed some concerns about possible crosstalk. A dedicated Hub test module is under development, which will allow populating all the FEX slots in the shelf and fully testing all the backplane links concurrently, which should show up any crosstalk issues.

Hardware monitoring

The I2C bus to access the on-board sensors is shared between two masters: the Hub FPGA and the IPMC. It is unclear how the arbitration will work and access from the IPMC to the sensor bus appears not to have been tested yet. The IPMC requires access to some critical temperature sensors (Hub FPGA and ROD FPGA on-chip temperature) in order for the ATCA shelf speed regulation through the shelf manager to work. This needs to be demonstrated.

CONCLUSIONS AND RECOMMENDATIONS

The recommendations given above should be addressed. The Hub design team is also strongly advised to re-consider some of the technical choices made and to adapt the design of the board for the pre-production modules. This concerns in particular the reference clock distribution scheme and the jitter cleaner device used as well the choice of PCB dielectric material.

In any case, the fixed phase of the clock extracted from the FELIX TTC stream needs to be demonstrated across modules resets, FPGA re-configuration and power-cycles. Since the standard GBT-FPGA implementation in a Xilinx Ultrascale FPGA can achieve a delay variation better than 100 ps (60 ps peak-to-peak over 1000 reset cycles), this is the performance figure to be met. The same applies to the latency for forwarding of the TTC information from the FELIX to the FEX slots, which also needs to be constant (to less than 1 ns) under these conditions.

In case the selected jitter cleaner device is kept, it has to be fully qualified in terms of input-to-output phase variation, including the effect of temperature. The input-to-output skew of the Silicon Labs Si5344 has been measured to be about 20 ps (over 10000 reset cycles) at room temperature. The performance of the selected jitter cleaner should match this figure. The phase-noise spectrum of the output clock used for the MGT reference also needs to be analysed against the requirements from the Xilinx specifications.

A signal integrity analysis of the most critical high-speed tracks (probably from the fan-out buffers to the ROD) should be done to assess the impact of the chosen dielectric with respect to a high-speed material such as Megtron-6.

It is recommend that a combined ROD/Hub intermediate design review be held prior to launching the pre-production of either module. The issues mentioned above have to be addressed at this time. In addition, a full chain slice test should have been demonstrated, including the transmission of a TTC stream from FELIX through the Hub to a FEX modules as well as readout data transfer from a FEX through the Hub and ROD to FELIX. The test should be performed in a full-size ATCA shelf loaded with the Hub/ROD and at least one FEX module prototypes.

In addition, a list of tests that need to be performed before a production readiness review (PRR) can be scheduled is given below:

- The TTC interface with FELIX need to be fully implemented and characterized.

- Firmware for the essential functions of the Hub needs to be ready and fully tested. This includes the interface to FELIX as well as to the ROD and the FEX modules with close to final data formats. IPbus should also be fully working.
- A scheme for asserting back-pressure to throttle the readout to prevent buffer overflows needs to be defined and documented. The function required on the Hub have to be implemented.
- A full-size ATCA shelf should be fully populated with Hub Test Modules and two Hub modules in order to run traffic on all the backplane links to and from the Hub concurrently. BER tests should be performed in this configuration to check for potential crosstalk issues.
- Temperature (FPGA and MiniPOD) and power dissipation tests with a Hub carrying a ROD mezzanine in a target ATCA shelf, if possible with all slots populated (e.g. with Hub Test Modules), have to be performed.
- Interoperability tests with close to final jFEX and eFEX prototype modules need to be performed. These tests have to done in a full size ATCA shelf with a backplane with the same characteristics as the one to be used in the experiment (ideally with the target shelf). Depending on the availability, interoperability tests with L1Topo should also be performed.
- Reception of readout data in the ROD mezzanine installed on the Hub has to be tested from all the FEX slots (could be one at a time). Tests should be performed using IBERT as wells with the Aurora protocol and full data format implemented in the firmware.
- Tests of the TTC clock reception on the each of the FEX slots as well as the TTC/control link to and from the FEX slots have to be performed under the same conditions as above.
- Monitoring on-board sensors (at the very least for the Hub and ROD FPGA temperatures) using the IPMC has to be tested. The sensor values have to be read out through the ATCA shelf-manager.

Since many of these tests involve the ROD and since the performance of the ROD heavily depends on various features of the Hub, it is strongly recommended to perform a common PRR of the Hub and the ROD.