CMS Upgrade MB Response to SLHC Document:

11.01: IP-Cores for Control and Readout in CMS Upgrades.

Contact Person: Guido Magazzu – UCSB, Physics Department

It is our intent to recommend this proposal for approval if revised to address the following requests (1-11), detailed below.

Please also see the comments from the referees below.

Specific requests for the revised proposal are:

1. Provide a plan to collaborate with the already approved GBT project and provide information on the complementarity of this project with the GBT project.

2. Provide results from the FF-TC1 chip evaluation.

3. Explain the difference between the FF-LYNX links and the GBT project E-links and why the E-links are not usable in place of the FF-LYNX links. Please also compare the technical performance.

4. Explain the choice of a 6-bit custom encoding scheme over use of standard encoding schemes.

5. Explain the plan to achieve the maximum data rate performance of at least 640 Mbps of the SERDES over all simulation corners.

6. Provide more detail on the impact of the VL frames with unbounded latency on the EMU and HCAL systems with results from simulation or a plan to do this.

7. Provide more detail on how this project would meet the needs of the track trigger upgrade, providing more details on the architecture proposed to use it.

8. Provide more details on how this project fits in with the schedules of the HCAL and EMU upgrades.

9. Provide a plan for radiation testing of the ASICs.

10. Provide details on the possibility of clock recovery to enable a “1-wire” scheme.

11. Provide milestones at points for subsequent reviews of the project tied to the ASIC production and testing.

In addition, we request you to address the detailed comments on the proposal from the referees below.

Referee #1:

I have reviewed SLHC R&D proposal 11.01 and I find that it is appropriate for CMS needs at SLHC and believe that it is not an excessively duplicated effort. My conclusions are based on the following observations and considerations:

1. Given that the FF-LYNX protocol is already being considered by some CMS sub detectors, it is clear that it is addressing the needs of CMS. In addition, it is a protocol that has already been tested and proven in silicon. However in its present interface implementations it requires a "2-wire" connection. If the FF-LYNX protocol were to be developed to be compatible with "1-wire" connections (both copper and optical), then it would be a much more attractive option for some CMS needs. But a "1-wire" scheme requires clock recovery capability. Although the scope of this proposal does not include clock recovery, it does address the necessary elements needed for developing clock recovery.
2. Since I am not aware of all the CMS SLHC projects currently being worked on, I cannot say with certainty that this proposal is not a duplication of effort. However I am aware of the GBT project and its apparent similarities with the FF-LYNX project. But I do not believe that they are in direct competition with each other but rather that the FF-LYNX project is potentially compatible with the GBT project. Also it may be possible that there are scenarios where an FF-LYNX solution in conjunction with a GBT solution is what is needed. So while there may be some overlap between the projects, I would not say that they are excessively duplicated.

General Comments:

While the challenges of developing high performance radiation tolerant designs are substantial, there seems to be a clear plan for overcoming the limitations of the current design. If the proposed implementations are successful, this will be a significant enhancement to the FF-LYNX protocol. In addition, it will form the basis for developing clock recovery capability.

Concerns:

The scope of this proposal is for two ASIC prototypes to be produced for testing the effectiveness of the IP-core implementation. One of the goals of the project is radiation tolerance. But the scope of the proposal does not include radiation testing. Clearly any evaluation of progress and plans for continuation will need to consider results of radiation testing even if funding for the testing is from another source.

Check Points:

I do not feel confident in setting a timescale for a milestone for this proposal, but an obvious check point would be after the first ASIC has been produced and tested and before the second ASIC is submitted.

Referee #2:

General:

This proposal describes the status of the FF-LYNX project, and highlights proposed activities for FY2012, namely two chip submissions evolving from the FF-TC1 chip under characterization and including rad-hard LVDS pads, 800Mbps SER/DES blocks, DC-balanced error detecting/correcting CODECS, etc...

A) Comments on FF-TC1 chip and its evolution to TC2-3

1. The TC1 chip is under evaluation, with August-11 stated as completion deadline. The results from this evaluation should be known before a decision on the proposal is taken
2. A custom 6bit encoding scheme is mentioned in section 2. Given the huge effort devoted to code-developments worldwide, the use of standard encoding schemes should be encouraged for TC2 and 3.
3. A maximum datarate of 640Mbps is claimed in section 2, while section 3 mentions that the SERDES does not operate up to this rate over all simulation corners. 640Mbs should be a minimum target rate for TC2 and TC3 designs. It is even doubtful that this rate will be sufficient for most SLHC detectors.
4. The FF-Lynx protocol has VL frames with unbounded latency. The effect of this on the target system(s) should be simulated with realistic FL frame rate and fluctuations (this may have been done, but is not reported in the proposal).
5. DC balanced encoding should be implemented in all channels (section 2 mentions only the 8xF channel).

B) Comments on target applications

The authors mention three applications:
- CMS Endcap Muon
- Redundant control system for CMS HCAL
- Readout system for CMS Track Trigger

It seems that the bandwidth requirements of a track trigger system will vastly exceed the capabilities of a 640Mbps-based system. Thus, the tracker application cannot in my opinion seriously be used as a justification for the TC2 and TC3 chip developments.
Leaving the track trigger aside, one should now consider the justification of developing an ASIC for the Endcap Muon system and for the HCAL redundant control system. How many chips are envisaged? The answer will probably speak for itself. CMS should resist the temptation of developing ASICs in advanced technologies for niche applications.

C) Comments on the innovative aspects of the proposal

After the global success of the GOL chip (1.6Gbps in 250nm technology), the proposal for a 640Mbps chip in 130nm technology cannot be claimed innovative. Moreover, at a time when all LHC collaborations (including CMS) endorse the GBT project (4.8Gbps in 130nm technology) and are represented in its specification group by their electronics coordinators, it is difficult for CMS to come up with an alternative protocol and chip that will result in different hardware,

firmware and software in the experiment. Finally, in many aspects, the FF-TC1 chip is already obsolete.

To conclude:
1. This project does not provide enough bandwidth for most SLHC applications and cannot be considered sufficiently future oriented or scalable. It may be appropriate for specific shorter-term upgrades, but cannot be claimed of general interest to CMS.
2. This project does duplicate the effort of the SLHC-wide GBT development. It distracts from the main objective of the collaborations to build common hardware and firmware for next generation links, and will result in the waste of badly needed resources (now when hardware needs to be built, and in the future when software needs to be developed and maintained). It targets a niche application and still faces many challenges (in particular the SERDES and CDR blocks) which have already been overcome by the GBT project.

I recommend not to support this project at the CMS level, and to encourage the initiators to join the common effort around the GBT project.

Referee #3:

1) It seems to me that the proposal actually IS focused to CMS and SLHC. It has a foreseen application in the muon system.

2) There is a possible application of the system in the HCAL to interconnect front end boards. We should encourage the group to extend the collaboration. Being in the HCAL community probably Drew Baden can comment more on this point. It would be nice to see this project used in more than one sub-detector.

3) I wonder if the project has been already discussed in the Electronics coordination and brought to the attention of the other sub-system upgrade projects. Magnus for sure knows this. May be more interested groups might join the project, or can give their opinion on the design. 4) Possible use in the track trigger was mentioned in the proposal. Being still in the far future I could imagine that the link-requirements for the track trigger will still change. Probably a link in that system will profit from further progress in technology expected up to the time it will be needed. But if the link will be deployed in the first shutdown for the muon system and possibly the hcal upgrade, a lot of experience will be collected, and it could be possible that a successor of the link will be used in projects like the tracker upgrade. Therefore the collaboration with the tracker upgrade team should be encouraged.

Referee #4:

Specifically the reviewers are asked to address:
1) The RD is appropriate for the needs of CMS at SLHC (ie focused)
This R&D proposal concentrates in 2012 on the "development of TX interfaces in the three speed options… compatible with optical links (…) and with single-wire links, i.e. clock and data encoded onto one serial line."
Comparing to CMS needs, possible applications of the FF-LYNX project mentioned are

- in the FF-EMU ASIC for the ME1/1 electronics upgrade (DCFEB and O-DMB boards)
- in VHDL modules as part of Actel devices in the HCAL control system for front-end boards
- in the Track Trigger project

It appears that the 2012 R&D proposed comes too late for the ME1/1 electronics upgrade, and is dedicated towards rad-hard LVDS pads and other ASIC developments that do not seem to be directly applicable to the HCAL control system project as it is based on Actel devices. So the justification for the 2012 R&D relies heavily on the Track Trigger project, involving collaboration with a FNAL group (R. Lipton).

The proposal claims that "a readout architecture for the Track Trigger detector based on the FF-LYNX protocol has been proposed…" but no references to the studies were included in this proposal, and so it was not possible to follow the argument in its details and the present reviewer is not an expert on the status of the Track Trigger designs.

_2) The RD is not excessively duplicated (ie we don't have too many people have working on the same topics)

The proposed links seem to be much slower than the proposed GBT links, and it would seem that the GBT links should become available on the time scale for development of a Track Trigger, so a direct comparison would also seem in order.

We would also like your input on suggested milestones or check-points where progress and plans for continuation should be reviewed.

As a one-year proposal, it would seem that if funded the progress and plans should be evaluated approximately one year from now.

Referee #5:

1 The RD is appropriate for the needs of CMS at SLHC (i.e. focussed)

The proposal appears to be focussed on potential, or perhaps probable, needs for planned and / or future CMS upgrades. Indeed a protocol compliant to FF-LYNX, which is only vaguely described in the Proposal, is adopted for the CMS EMU ME1/1 electronics upgrade readout and is being considered by others, e.g. the CMS HCAL RBX control redundancy scheme.

2 The RD is not excessively duplicated (i.e. we don't have too many people have working on the same topic)

The proposed protocol as presented appear as a complement to the already approved GBT (GigaBit Transceiver) development and its proposed lower speed interface links, the so-called e-links. The E-links appear to address the same issues as FF-LYNX in the Proposal.

3 Additional comments

There is no mention of long term support which would be required in case the FF-LYNX would replace or complement the current TTC.

4 Conclusion:

It seems like the FF-LYNX project would benefit from close collaboration with the GBT development project within which most of not all remaining difficult issues have been addressed or will have to be addressed. Resources could then be gathered rather than divided in order to deliver a solid set of ASICs and IP-Cores for CMS upgrades.

I would propose to approve with a rather strong request to seek coherence and collaboration with the already approved GBT project

Referee #6:

Generally speaking, this is a very good proposal, which addresses studies of new physics in areas The R&D is entirely appropriate, and I give this proposal a green light to proceed. I would suggest that they move forward in their efforts to instantiate the protocol into an ASIC, and continue their efforts to produce a protocol that has an embedded clock, something that would be extremely useful for all sub detectors of the upgrade. This is especially true given that the GBT project is our only alternative at this time to the TTC, which has to evolve. It is very risky to have a single alternative to this important issue. FX-LYNX is a very sensible and well thought out project and I give it a high rating to move forward.