MAC Layer for Optically Enhanced Interconnectivity for Computing Platforms

G. Azmy[1], M. Abd El Meguid[1], M. Sowailem[1], R. Varano[2], D. Rolston[2], Prof. O. Liboiron-Ladouceur[1]
[1]Photonic Systems Group, Dept. of Electrical and Computer Engineering, McGill University; [2]Reflex Photonics
george.azmy@mail.mcgill.ca mohamed.mostafa2@mail.mcgill.ca

Motivation

The increasing demand of bandwidth from scientific modeling, simulations and video streaming has created a new bottleneck at the interconnect level between processors where traditionally copper is used.

Replacing copper with optics allows for a smaller footprint, longer cable runs and lower power consumption that fits in with the current top-of-rack architecture.

Bringing the optics closer to the processor:

- Limits RF interference
- Improves latency
- Opens the door to greater bandwidth.

Experimental Setup

The experiment is done using *Altera Stratix IV GX* FPGAs to emulate the servers interconnected using the *Reflex Photonics LightABLE* optical module. Each module has 12 optical channels running at 10Gb/s each, and configured as either transmitters (VCSEL arrays) or receivers (photodetector arrays). The project starts with deploying the project on a single FPGA for testing and our goal is to have 8 such optically interconnected servers in a chassis.

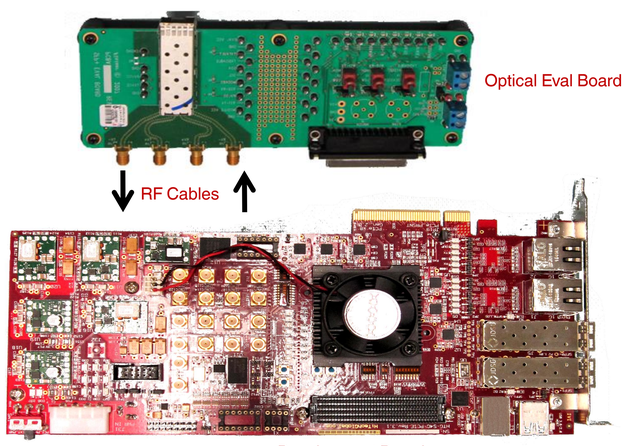

Development Board

Architecture Overview

Three optical modules (8 x 10Gbps channels/module) per server as follows:

- One TX block; 7 peer-to-peer (P2P) channels to the other servers, 1 broadcast channel
- One RX block with 7 P2P channels
- One RX block for receiving Broadcasts from the 7 other servers

The MAC layer supports, P2P, broadcast and multicast communication, utilizing all the channels available.

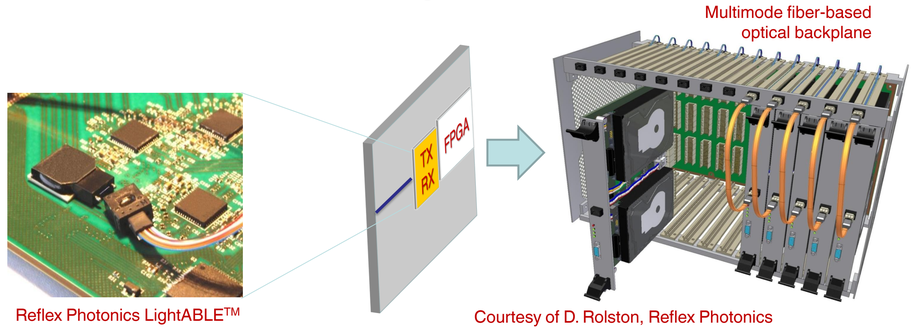

Reflex Photonics LightABLE™

Courtesy of D. Rolston, Reflex Photonics

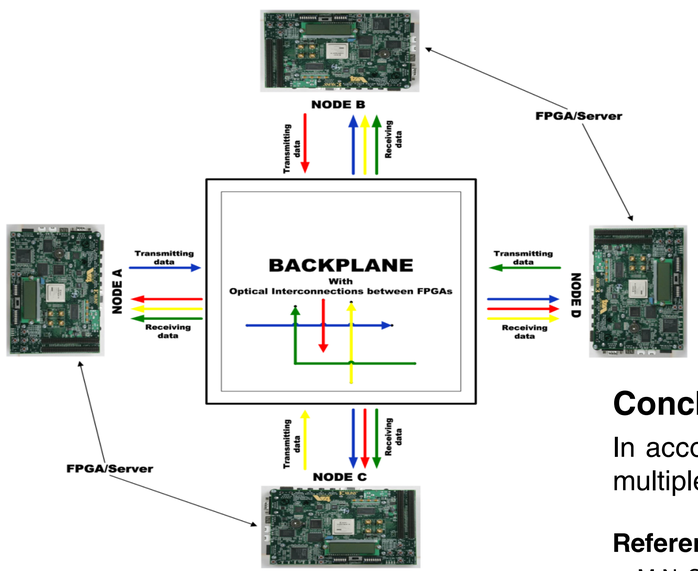

Overall Architecture

Networking (MAC)

The OSI model is implemented on the FPGAs. The MAC layer was developed in VHDL in compliance with the IEEE 802.3ae full duplex 10G Ethernet standard

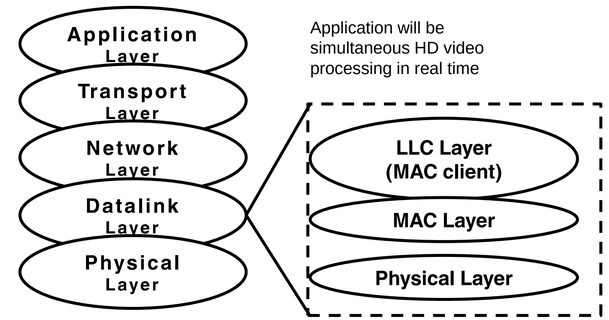

OSI Model

Ethernet Packet Format

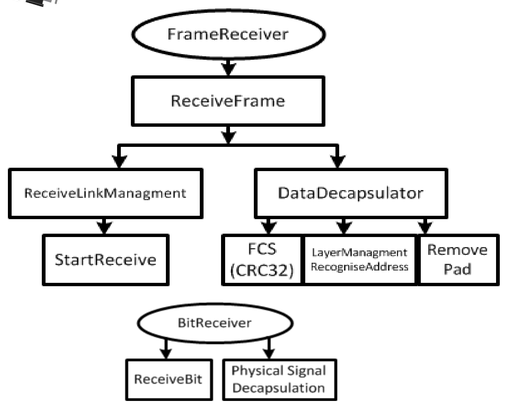

Receiver flow chart

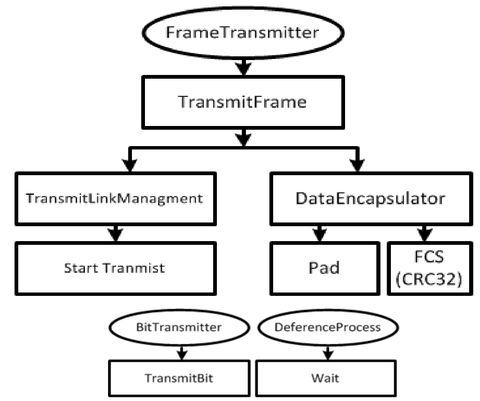

Transmitter flow chart

Conclusion

In accordance with the goals the mature test-bed will be replicated on additional FPGAs and multiple high-bandwidth HD video streams will be used to demonstrate the capabilities.

References

- M.N. Sakib, M. Sowailem, M.S. Hai, H. Abbas, G. Azmy, R. Varano, D. Rolston, and O. Liboiron-Ladouceur, “Development of Optically Enhanced Interconnectivity for Computing Platforms," CIPI Annual General Meeting, Ottawa, May 2011.
- IEEE 802.3ae 10GE http://standards.ieee.org/about/get/802/802.3.html