PROCEEDINGS OF SCIENCE

Fast Readout Links

Francois Vasey[1]
CERN
1211 Geneva 23, Switzerland
E-mail: francois.vasey@cern.ch

Optical links have become a commodity in high energy physics experiments, as they have in telecom and datacom networks. However, the specific constraints found in detector front-ends often make it difficult to directly transfer the rapid progress observed in commercial optical links.

This paper compares the characteristics of optical links developed for tracking detectors with those found in typical datacom and telecom networks. From this comparison, we draw elements that might be characteristic of fast readout links developed for high energy physics in the coming five years.

The 20th Anniversary International Workshop on Vertex Detectors – Vertex 2011
Rust, Lake Neusiedl, Austria (June 19 – 24, 2011)

[1] Speaker

1. Introduction

To give a workshop presentation on a topic as broad as fast readout links and to elaborate on their status and evolution is a task which can be achieved, but only in a consciously biased and partial way. However, to write up a general paper that will resist time erosion is an impossible challenge as many of the forecasts made will inevitably be proven inaccurate within a few years. This paper will thus be concise, will focus on a few well established facts, and will avoid forecasts with too far reaching implications.

After defining in section 2 the metrics used to compare optical link characteristics, we will review in section 3 the performance of optical interconnects used in network applications surrounding our everyday life. In section 4, we will present the characteristics of optical links in operation at the Large Hadron Collider (LHC) as well as under development for the High Luminosity upgrade of the LHC (HL-LHC). Putting side by side the commercial and LHC applications will allow us to extract a few basic trends for future High Energy Physics (HEP) links.

2. Metrics

From a system designer perspective, the characteristics of an optical link can be narrowed down to the following five metrics:

- Bandwidth (in Gbps)
- Reach (in m)
- Energy (in pJ/bit) or Power (in mW/Gbps), 1pJ/bit= mW/Gbps
- Density (in mm^2/Gbps)
- Cost (in $/Gbps)

Plus of course the environmental resistance of the device which must match the system constraints.

These metrics are extremely useful to compare different link implementations, but those normalized by datarate obviously favour applications where the bandwidth can be concentrated to a few very high datarate channels. Applications (such as frequently found in HEP) where the bandwidth is distributed and cannot easily be aggregated, will use more physical channels for a given total bandwidth and will end up having a less favourable figure of merit. In the following sections, we will nevertheless use these five metrics to compare HEP and non-HEP links, and to extrapolate the characteristics one could expect from future developments.

3. Optical interconnects in commercial network applications

Historically, long-haul telecommunication networks (also referred to as Wide Area Networks (WAN)) have been the first large scale systems relying on fast optical data transfer. Due to the elevated costs associated to the deployment of such networks, WAN link performance has always been maximized to achieve ultimate bandwidth and reach. Nowadays, time as well as wavelength division multiplexing schemes are routinely used, together with advanced modulation formats and electronic mitigation techniques. Bandwidths as high as 10 Tbps per fibre, transported over hundreds of kilometers, have been reported [1]. In this context, optical component power, size and cost efficiencies are secondary issues and it therefore seems unlikely that WAN-optimized devices with all their complexity and cost will be much used in HEP networks. However, techniques developed to augment capacity and reach (advanced modulation formats and Forward Error Correction (FEC) for instance) will eventually become available to shorter reach networks. They should thus be monitored carefully and understood in order to be prepared for the long term evolution of the technology.

The access network is the branch of the network that reaches the end-user, be it a phone, a computer, a game-station or a TV-set. For many years, the fixed line access has been dominated by twisted pair and coax cable based technologies, but fibre is now (finally) reaching an increasing number of homes and buildings in Asia and America, with Europe slowly catching up. Most fibre to the home deployments rely on passive optical network (PON) architectures where tens of users share the same fibre in a time-multiplexed fashion. Reaches of several kilometers (up to 20km) at data rates of 2.5Gbps per fibre are typical (10Gbps are announced) [2]. Even though the PON star topology (point to multi-point) is not adapted to the high upstream bandwidth needs of HEP readout links, PON-optimized components (or subparts) produced in large volume and at low cost for the subscriber-end of the links are of high potential interest to HEP link developers.

Local area networks (LAN) are usually considered to closely match the HEP readout link requirements due to their similar reach (up to 300m) and data rate characteristics (1-10Gbps, quickly moving up to 100Gbps). They are implemented in very large quantities in campuses or data centers, featuring impressive power, density and cost figures. 10Gbps will become the baseline data rate of newly installed physical channels, while 40GBASE- and 100GBASE-Ethernet requirements will be met by aggregating 10Gbps lanes. In contrast to the long-haul and access networks which are both relying on single mode (SM) fibre, the vast majority of LANs in operation today are based on multi-mode (MM) technology. However, the trend towards longer reach and higher data rates, together with the convergence with the metropolitan area network, may change this situation in the future. It is worth noting that the ever increasing parallelism (and hence capacity) of LAN-optimized components may leave the HEP community behind at some point, as aggregating data from a detector front-end to a very high capacity node may prove more difficult than from a switch or router.

Table 1 presents a coarse comparison of the three link technologies reviewed above. Despite its oversimplification, it clearly shows the very large span covered by electro-optical technology today. If one only looks at the PON and LAN columns, and considers the absolute power, size and cost of a generic active component, typical figures of the order of 1W, 100-500mm^2 and 100$ per device seem to emerge.

Table 1: Optical link metrics for components used in WAN, PON and LAN networks: a simplistic comparison showing orders of magnitude only

	WAN	*PON*	*LAN*
Bandwidth [Gbps]	*1’000*	*1*	*100*
Reach [m]	*100’000*	*10’000*	*100*
Power [mW/Gbps]	*>1’000*	*1000*	*10*
Density [mm^2/Gbps]	*1’000*	*500*	*1*
Cost [$/Gbps]	*>100*	*100*	*1*

4. Optical interconnects at LHC, for HL-LHC and beyond

In the 2005-2008 timeframe, the LHC experiments have deployed optical links to an unprecendented scale for their readout systems, installing tens of thousands of point-to-point links between their front-ends and counting rooms. The lessons learned by ATLAS and CMS during the development, procurement, installation and commissioining of these systems have been summarized in [3].

In 2008, a common project was launched by ATLAS and CMS: the Versatile Link project [4]. It aims at developing a bi-directional 5Gbps optical link for HL-LHC, compatible with the installed cabling plants of the LHC experiments. It will thus be made available in SM as well as MM versions, and will feature a transceiver adapted to the constraints (environmental, mechanical, etc.) of detector front-ends.

Table 2 presents the link metrics achieved by the CMS-Tracker front-end optical components and expected from the Versatile Link front-end components.

Table 2: Optical link metrics for CMS Tracker and Versatile Link front-end components

	CMS Tracker	*Versatile Link (targets)*
Bandwidth [Gbps]	*0.32*	*5*
Reach [m]	*100*	*100*
Power [mW/Gbps]	*150*	*100*
Density [mm^2/Gbps]	*1000*	*100*
Cost [$/Gbps]	*900*	

When comparing Table 1 with Table 2, it is evident that apart from their reach, the components developed (or even under development) for HEP tracker front-ends do not match current LAN characteristics well. This counter-intuitive observation is linked to the fact that:

i) arrays have not (yet) penetrated most tracker detector front-ends and ii) radiation-tolerant electronics is generally based on more conservative technologies and dissipates more power than up to date commercial electronics. Based on these observations, one can infer that single-channel links developed for tracking detector front-ends during the next few years will probably not exceed 10Gbps in bandwidth, will dissipate of the order of 500mW and will have a footprint of a few $100mm^2$.

At their backends however, HEP links will be able to harness the full benefit of parallel optical engines developed for advanced LAN applications. This will result in more than one order of magnitude improvement in power efficiency and density compared to front-end components, and tight integration with field programmable electronics.

Over the past ten years, it has become clear that fast optical interconnects always outperform electrical interconnects for long and medium distances (>10m), but will they in the future also replace copper at the board-to-board, or even chip-to-chip level? Will we some day see optically coupled front-end chips, or backend boards connected to optical backplanes? Several studies show that supercomputers, in order to keep up with the steady increase of processing performance observed for the past 20 years (a 10-fold increase every 4 years), will need to rely heavily on intra-rack optical links [5]. Many development projects have demonstrated the feasibility of such dense and power efficient parallel interconnects. The IBM Terabus project for instance [6], has achieved an impressive 6.5mW/Gbps efficiency with a density of $0.1mm^2$/Gbps in an in-board 24-channel optical interconnect. Such results are state of the art. We believe they could be seen as an ultimate benchmark for future optical links based on arrays of directly modulated lasers. In the meantime however, thanks to advanced electronic processing, electrical links have managed to remain more efficient (3mW/Gbps) than their optical counterparts over short distances (<1m). In all likelyhood, electrical links will thus remain the established technology for intra-board, and possibly even board-to-board interconnects in the near future.

At the chip level, the input/output bandwidth limitations of current high density packages are a clear bottleneck to the scaling of their processing abilities. Chip-scale optical interconnects may offer a long term solution, though very low energy and novel optics will be needed to achieve the sub pJ/bit efficiencies (<1mW/Gbps) required to keep the chip power dissipation within a sustainable range (100W-200W) [7]. Modulator-based interconnects integrated with CMOS electronics may be one of the technologies allowing such a breakthrough, but more progress will be required before predictions can be claimed accurate, possibly by 2015.

5. Conclusions

Broadband communication has become a commodity, and the rapid growth in demand for capacity has traditionally been met by aggregating physical or wavelength lanes and by using advanced modulation techniques. In HEP, optical readout links have also become a commodity, but for tracker front-ends, it is not clear that any of the above mentioned aggregation techniques will be useful due to the distributed nature of the data sources and the strong preference to implement local electrical to optical conversion at relatively modest data rates (1-10Gbps).

In the short term, optical link developments for future tracking detectors will thus tend to concentrate, at the front-ends, on single-channel, 10Gbps radiation tolerant links based on directly modulated lasers. At the link back-ends, the need for high density will call for the use of commercially available parallel optical engines aggregating up to 12x10Gbps lanes.

In the second half of the decade, the Communication Technology Roadmap [8] points at the emergence of optical transceivers monolithically integrated with CMOS electronics. These will be very compact and low power devices coupled to single mode fibres, meeting the huge off-chip transmission bandwidth needs of advanced processors. Photonic integration with CMOS could be a paradigm changing breakthrough for designers of HEP tracking detectors also. It would give them access to massive bandwidth at a very low power dissipation cost.

For the next few years however, more conventional laser-based links will be what the HEP community will have to be content with.

References

[1] http://www.eetimes.com/electronics-news/4214401/ZTE-claims-10-Tbps-data-transmission-record

[2] M Daneshmand, C. Wang and W. Wei, *Advances in Passive Optical Networks*, IEEE Communications Magazine, Vol. 49, No. 2, Feb 2011. http://dl.comsoc.org/~ci

[3] K.K. Gan, F. Vasey and T. Weidberg, *Lessons Learned and to be Learned from LHC*, Joint ATLAS/CMS note, ATL-COM-ELEC-2007-001, CMS-IN-2007/066, https://edms.cern.ch/document/882775/3.8

[4] L. Amaral, S. Dris, A. Gerardin, T. Huffman, C. Issever, A.J. Pacheco, M. Jones, S. Kwan, S.-C. Lee, Z. Liang, T. Liu, Z. Meng, A. Prosser, S. Padadopoulos, I. Papakonstanstinou, C. Sigaud, S. Silva, C. Soos, P. Stejskal, J. Troska, F. Vasey, P. Vichoudis, T. Weidberg, A. Xiang and J. Ye, *The versatile link, a common project for super-LHC*, 2009 JINST 4 P12003

[5] M. B. Ritter, Y. Vlasov, J. A. Kash and A. Benner, *Optical technologies for data communication in large parallel systems*, 2011 JINST 6 C01012

[6] C. Schow, F. Doany and J. Kash, *Get on the Optical Bus*, IEEE Spectrum, Sept. 2010, p. 31

[7] D.A.B Miller, *Device Requirements for Optical Interconnects to Silicon Chips*, Proceedings of the IEEE, Vol. 97, No. 7, July 2009, p. 1166

[8] J. Michel, L.C. Kimerling, *Communications Technology Roadmap II*, presented at the 21st annual workshop on interconnections within high speed digital systems, Santa Fe, NM USA, 3-5 May 2010, http://mph-roadmap.mit.edu