Demonstration of Intra-Data Center Link Based on 1x4 Multicore Fiber (MCF) Edge-Coupled to Silicon Photonics

Doug Butler(1), Norbert Schlepple(2), Sergey Ten(1), Cliff Sutton(1), Kevin Bennett(1), Matthew Tuggle(1), Marc Lorenzo(1), Jeff Clark(1), Pushkar Tandon(1), Aramais Zakharian(1), Anthony Artuso(1), Vasanta Rao(3), Dingchuan Chen(3), Nikolai Kalnin(3), Ravi Tummidi(2), Nadine Tsai(2), Bob Bullock(2), Alberto Cervasio(4), Joyce Peternel(3) and Matt Traverso(3)

(1) Corning Incorporated, Corning, NY, butlerdl@corning.com
(2) Cisco Systems, Allentown, PA, nschlepp@cisco.com
(3) Cisco Systems, San Jose, CA
(4) Cisco Systems, Vimercate, Italy

Abstract *We demonstrate 100 Gb/s single-wavelength as well as four wavelength CWDM4 error-free performance of a 2 km data center link consisting of 1x4 MCF, MTP connectors and effective coupling into silicon nitride (SiN) waveguides on Cisco's standard photonic integrated circuit platform.*

Introduction

Power consumption in hyperscale data centers (HSDCs) becomes a serious problem since high SerDes data rates increase the share of power consumed by networking, thus taking away power from revenue generating servers[1]. Moving transceivers closer to the switch ASIC using on-board and ultimately co-packaged optics (CPO) relaxes I/O power requirements. The transition to CPO is expected to start with 100 Gb/s lane rates when SerDes start consuming significant amount of power and at the same time the number of lanes reaches 512 and 1024 (i.e. ASIC capacity of 51.2 Tb/s and 102.4 Tb/s). The large number of lanes particularly in case of parallel (DR) interfaces leads to necessity to manage 1000+ fibers in the switch box and between switches. Equally challenging is handling many optical lasers and signal I/Os along the limited physical width of an individual optical engine surrounding the switch ASIC. Several solutions were proposed to address the linear core density requirement e.g., packing regular 125 µm cladding fibers in 127 µm pitch fiber array unit (FAU) or even reducing the cladding diameter to 80 µm. The most radical proposal is to use multicore fiber (MCF) that would increase the linear core density beyond what reduced cladding fibers can achieve[2]. For example, the use of 1x4 MCF will reduce by 75% the number of fibers in the FAU, will reduce the number of alignments to silicon photonics (SiPh) chip and the number of cables between switches.

In this paper we demonstrate the feasibility of an intra-DC switch to switch link based on 1x4 MCF with standard 125 µm cladding and MCF connectors. In addition, coupling loss and optical on-chip crosstalk of the 1x4 MCF aligned to low-loss silicon nitride (SiN) edge couplers[3] on Cisco's standard photonic integrated circuit (PIC) platform is demonstrated. Results are in line with single mode fiber (SMF)

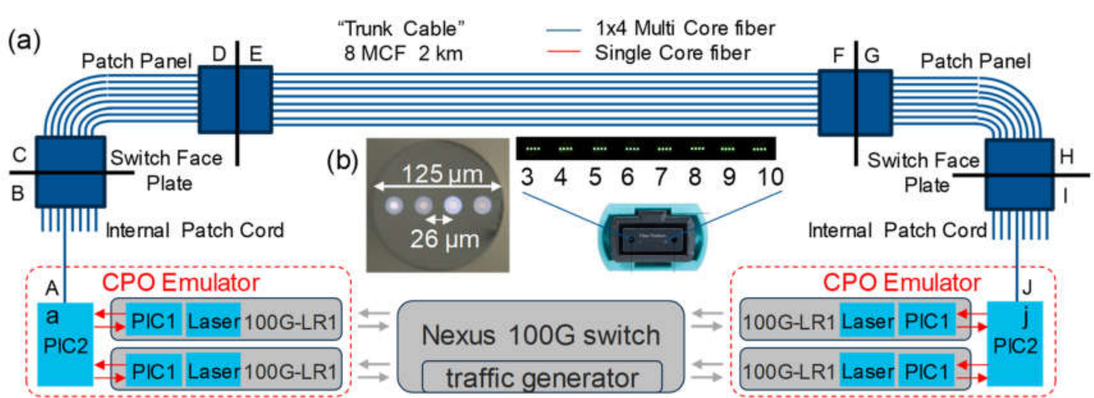

Fig. 1: (a) Schematic of the MCF link with 2 km trunk and shorter patchcords. The link is terminated on either end with our CPO emulator comprising of passive PIC2 and active PIC1+Laser. (b) Cross-section of the 1x4 MCF with standard 125 mm cladding and front facet of the MTP connector with MCFs in Positions 3-10 are shown.

978-1-6654-3868-1/21/$31.00 ©2021 IEEE

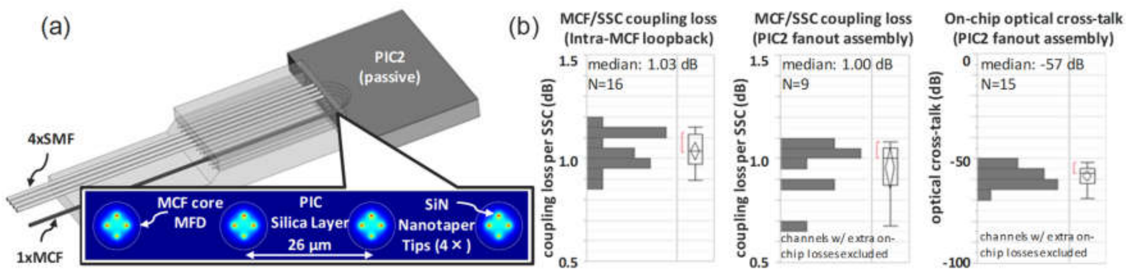

Fig. 2: (a) Illustration of FAU attached to PIC2 with MCF aligned to four Cisco SSC at a 26 µm pitch, TE near field of four SiN nanotaper SSC is shown. (b) Coupling losses between MCF and SSC are shown for intra-MCF loopback as well as MCF embedded in the five-fiber FAU. No significant on-chip optical crosstalk including among SSC is observed.

coupling performance. Finally, error-free 100 Gb/s PAM4 single-wavelength as well as CWDM4 bidirectional point to point data traffic is demonstrated. This is the first to our knowledge demonstration of a broadband MCF link with standard 125 µm cladding and MTP connectors.

Multicore Fiber Design

The link demonstrated in this experiment is shown in Figure 1a and it contains segments typical for the data center structured cabling. In the middle, there is a trunk section (E-F) that is 2 km long connected using MTP connectors (see Figure 1b) to the shorter 3 m patch cords (C-D and G-H) representing connections from the patch panels for structured cabling to the front panels of the switch. Sections A-B and I-J represent in-switch cable assemblies from the front panel to the SiPh CPO. There are 8 MCF fibers in the link, placed in Positions 3-10 in the MTP connector. We chose 8x4 optical cores in the link because that would correspond to the capacity of 1.6 Tb/s and 6.4 Tb/s with 100 Gb/s DR and FR interfaces, potentially routable (as a duplex fiber pair) in granularities of 400 Gb/s and 1.6 Tb/s respectively. One of the eight MCFs in the link is connected to a CPO emulator on each end for live traffic testing. A 1x4 MCF cross-section is shown in Figure 1b. The most important design parameter is the distance between optical cores (core pitch) that becomes critically important for standard 125 µm cladding fiber. The criticality of the core pitch is dictated by a delicate balance between inter-core cross talk and radiative losses of the outer cores due to their proximity to the cladding interfaces. Our design chooses a core pitch of 26 µm that minimizes radiative loss of outer cores while keeping cross talk at sufficiently low level. Core design is like that of bend insensitive fiber with a mode field diameter (MFD) of 8.5 µm. The chromatic dispersion of the fabricated MCFs has zero wavelength in the range of 1300-1320 nm with the average at 1310 nm and the dispersion slope of 0.093 ps/(nm^2×km). We use standard MTP connector with MCFs placed in the positions 3-10 and angularly aligned (horizontally) with the accuracy of ≤0.3 deg (see Figure 1b).

Multicore Fiber Coupling to PIC

We use Cisco's previously presented[3] and well characterized spot size converter (SSC) based on four SiN nanotapers for integrating the MCF interface into our standard silicon photonics platform. The multi-prong edge coupler allows for 0.5 dB or less of coupling loss between a single standard SMF fiber over the entire O-Band with <0.3 dB polarization dependent loss (PDL) for TE/TM light under ideal conditions. In order to provide for the significantly reduced channel pitch on the utilized MCF, we had to bring the group of four SSC closer together accordingly. However, no design changes to the SSC itself were made. Figure 2a shows the FAU with one embedded MCF and four SMF in contact with our custom fanout chip. Four loopback waveguides on the chip connect one MCF core with its respective SMF core for breakout of the optical channels. To ensure negligible on-chip propagation and polarization-dependent losses, respectively, we chose to stay entirely in SiN with the signal paths, however, transitioning into silicon waveguides is possible at any point along the way as was the case here with power taps on the outermost channel leading to on-chip monitor photodiodes used for assembly purpose. The magnified insert in Figure 2a illustrates the MCF with its four cores as it aligns with the four SiN nanotapers on the SSC. Rotational clocking of the MCF within the FAU body is critical to avoid additional excess losses from lateral fiber core to SSC offsets, most significantly on the two outer cores. Figure 2b (left) shows minimum coupling losses obtained from looping core 1 to core 4 and core 2 to core 3, respectively, on an individual MCF and dividing equally the difference between MCF input powers and output powers after compensating for on-chip losses. Coupling loss distribution with median of 1 dB per core is in line with standard SMF coupling to the same SSC considering process tolerances in addition to the previously

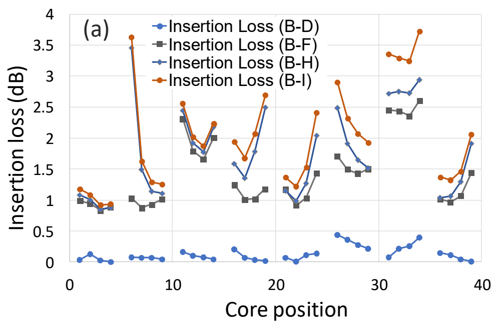

(b)

	Counter-Propagating		Co-Propagating	
MCF core number	1	2	3	4
Data Stream Direction	Up	Down	Up	Up
Line Rate (Gb/s/λ)	106.25			
Wavelength (nm)	1310			
Link Loss Budget (dB)	11.1	9.9	8.7	7.6
Receive OMA (dBm)	-7.1	-5.7	-4.4	-3.5
Pre-FEC BER (over 190 seconds)	7.7E-6	3.1E-6	5.1E-7	6.7E-8
Uncorrectable FEC Errors (>20 hours)	0	0	0	0
Frame Input/Output Errors (>20 hours)	0	0	0	0
Frame CRC Errors (>20 hours)	0	0	0	0

Fig. 3: (a) Insertion loss of the link (point B-I) for each of the 32 cores, including 2 km of MCF. (b) Link performance for counter-propagating and co-propagating channels, respectively. No post-FEC errors are detected.

reported[3] ideal single core performance. Having confirmed this coupling loss parity between MCF and SMF, we can equally distribute coupling losses between SMF input and MCF output on PIC2 using the same method. Figure 2b (center) indicates no additional lateral/rotational offsets on the FAU in our PIC2 assembly during peak alignment, while figure 2b (right) shows no on-chip optical crosstalk on nearest neighbor channels because of the reduced pitch between our SSC. No polarization control was applied in these measurements.

Link Performance

Figure 3a shows insertion loss of the MCF link from B to I incrementally as link segments are added i.e. from B to D, from B to F etc. with the total link loss not exceeding FR4 link loss budget of 4 dB. In fact, the demonstrated loss includes the loss of connector at the front panels which is normally included in the power budget of the transceiver i.e. normally not counted in the link loss budget. We also show that the insertion loss of the fiber is <0.83 dB for the 2 km length and it demonstrates that the design objective of the fiber was fulfilled – the insertion loss of the outer cores is only 0.07 dB higher than inner core over the link length. Connector loss on average is less than 0.8 dB but some variability is observed, especially for outer cores, thus providing a direction for improvement.

A CPO emulator on either end of the link consists of two Cisco Silicon Photonics QSFP28 100G-LR1 transceivers (LR1 transmitter delivering additional optical power to provide for PIC2 coupling losses; no forward error correction (FEC) difference versus FR1) representing PIC1+Laser[4], as well as an external PIC2, see Figure 1a. In this configuration PIC1 integrates active transmit and receive circuitry while PIC2 is passive and solely to provide the MCF coupling interface. PIC1 and PIC2 are otherwise built on the same PIC platform and can be monolithically integrated on a future CPO design iteration. To test link performance, we use an external 100G PAM4 traffic generator with 90% line rate utilization and a Nexus 9316D switch. Figure 3b summarizes link performance, including pre-FEC bit error rate (BER) on all four lanes well below the 2.4E-4 limit ensuring error-free post-FEC operation as indicated by the absence of uncorrectable FEC errors on transceiver level as well as no input/output nor CRC errors on frame level. Compared to the received OMA, pre-FEC BER tracks well with the reported[4] receiver sensitivity of our SiPh transceiver and suggests no further MCF impairments versus standard SMF. No BER degradation on co-propagating channels 3 and 4, both at 1310 nm optical carrier wavelength, implies insignificant optical crosstalk over the entire link. To expand on the optical performance, we confirmed more than 20 hours of error-free operation of four co-propagating 100G CWDM4 streams in four cores, demonstrating broadband operation (insertion loss and crosstalk) of the link.

Conclusions

We demonstrated 100 Gigabit Ethernet compliant error-free performance on all lanes over a 2 km data center link on a 1x4 multicore fiber with MTP connectors and effective coupling into SiN waveguides on Cisco's standard PIC platform. The loss, faceplate-to-faceplate, was below 4 dB and MCF coupling loss to our SSC was in line with standard SMF coupling performance. Running CWDM4 error free over the link indicates the low crosstalk and insertion loss demonstrated with 100Gb/s/λ can be achieved across the O-band.

Acknowledgements

Corning authors would like to acknowledge Song Yu from Facebook's Global Data Center Connectivity Group and Ruby Chen from Sourcing and Operation Engineering Team for guidance and fruitful discussions.

References

[1] M. Filer "Opportunities for co-packaging in future Microsoft data center networks," presentation at ARPA-e ENLITENET meeting, January 13, 2021.

[2] H. C. Nguyen "Multi-Core Fibers: Considerations for CPO" OIF presentation oif2021.098.01, 2021

[3] R.S. Tummidi, M. Webster, "Multilayer Silicon Nitride-based Coupler Integrated into a Silicon Photonics Platform with <1.0 dB Coupling Loss to a Standard SMF over O, S, C and L optical bands" OFC 2020.

[4] Matt Traverso, Marco Mazzini, Kadaba R. Lakshmikumar, Sanjay Sunder, Alex Kurylak, Craig Appel, Cristiana Muzio, Ravi Tummidi, Romesh Kumar Nandwana, Alberto Cervasio, Mary Nadeau, Weizhuo Li, Jarrett Neiman, Jock Bovington, Mark Webster, "Integrated Silicon Photonics Transceiver Module for 100Gbit/s 20km Transmission," in Optical Fiber Communication Conference 2021