



Low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber

ZICHUAN ZHOU,^{1,4} HUBERT DZIECIOL,¹  KARI CLARK,¹  YUAN LUO,^{2,5} DAVID RICHARDSON,³ FRANCESCO POLETTI,³ POLINA BAYVEL,¹  RADAN SLAVIK,^{3,6}  AND ZHIXIN LIU^{1,*} 

¹*Optical Networks Group, Dept. of Electronic and Electrical Engineering, University College London, Torrington Place, London WC1E 7JE, UK*

²*School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, Guangdong 518172, China*

³*Optoelectronics Research Centre, University of Southampton, Southampton SO17 1BJ, UK*

⁴*zczlz0@ucl.ac.uk*

⁵*luoyuan@cuhk.edu.cn*

⁶*r.slavik@soton.ac.uk*

**zhixin.liu@ucl.ac.uk*

Abstract: Fast (nanoseconds) optical wavelength switching is emerging as a viable solution to scaling the size and capacity of intra-data center interconnection. A key enabling technology for such systems is low-jitter optical clock synchronization, which enables sub-nanosecond clock and data recovery for optically switched frames using low-cost methods such as clock phase caching. We propose and demonstrate real-time low-latency wavelength-switched clock-synchronized intra-data center interconnection at 51.2 Gbd using a fast tunable laser (with ns scale switching time) and ultra-stable-latency hollow core fiber (HCF) for optically-switched data center networks. For wavelength-switched systems, we achieve a physical layer latency below 46 ns, consisting of 28 ns transceiver latency and a 18 ns inter-packet gap. Finally, we show that by exploiting the low chromatic dispersion and thermally-stable latency features of HCF, active clock phase tracking can be entirely eliminated.

Published by Optica Publishing Group under the terms of the [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

1. Introduction

The drastic increase in the size and information rate in data center networks has motivated research on fast optical switching to address the scalability and capacity challenges in intra-data center networks [1–3]. Compared to the current multi-layer electrical switching architecture, optical switches allow for a flat data center topology, which significantly lowers power consumption by substantially reducing the number of optical transceivers [4–7]. Optical switches for data center networks need to provide fast switching speed because production cloud data-center traffic is dominated by small size data packets, with 98% of the packets comprising <576 bytes (which corresponds to 90 ns assuming a 51.2 Gbps line rate) [4–6]. Among fast optical switching techniques, wavelength switching simultaneously offers the potential merits of higher scalability, lower power consumption, and lower end-to-end latency [8–12]. The aspect of low end-to-end latency is of particular interest to time-sensitive applications such as high-frequency trading, as well as real-time processing of virtual reality and autonomous car signals. End-to-end latency is fundamentally limited by signal processing and transmission in the physical layer. Therefore, it

is beneficial to understand the physical layer end-to-end latency of an optical switching system and how to minimize it for a low-latency system design.

In an optically switched system, the total physical layer end-to-end latency is the sum of the fiber propagation time, the time gap between consecutive optical packets, and transceiver latency. The transceiver latency consists of the time for clock and data recovery (CDR) to lock, signal processing (e.g., filtering and equalization), data sampling and detection, and forward-error correction (FEC) [9–11]. Generally, the transceiver latency for an optically-switched system is higher than that of a continuous data transmission system due to the additional time needed for burst-mode CDR, signal conditioning, adaptive equalization for optical bursts coming from different nodes, and the inter-packet gap required for optical switching [10]. The transceiver latency of standard optical transceivers (e.g. 100 Gb/s) is usually in the order of microseconds, dominated by CDR (90 μ s) [14] and FEC (230 ns) [15], which is comparable to 100s meters of propagation latency in standard single mode fiber (SSMF). Considering short-reach intra-data center interconnect distances of less than 100 m, low latency optical packet transmission system optimization requires mainly transceiver latency and the inter-packet gap to be minimized.

To this end, many techniques that individually address the above-mentioned challenges have been demonstrated. For instance, in [9,10,13], we and others have demonstrated optical clock synchronization and clock phase caching to enable sub-nanosecond CDR, which minimizes the CDR-induced transceiver latency. Hollow core fiber (HCF) allows a propagation speed that is 99% of the speed of light in vacuum, a 34% reduction of light propagation latency compared to SSMF [16,17]. Besides the low light propagation latency, the HCF low sensitivity of light propagation latency to ambient temperature (low thermal sensitivity) has also been explored to eliminate the CDR locking time as well as any network overhead for the clock phase caching [18]. To minimize the inter-packet gap, researchers have optimized driving waveforms for wavelength tunable lasers and demonstrated sub-10 ns wavelength-switching time across 30 nm wavelength bandwidth using commercially available digital-supermode distributed Bragg reflector (DS-DBR) tunable lasers [7], permitting a small gap between consecutive packets. A further reduction of the inter-packet gap to <100 ps was reported using a laser array or an optical frequency comb followed by a semiconductor optical amplifiers (SOA) array as switches [12]. Although these techniques have been individually demonstrated in previous studies, to the best of our knowledge, there has not been a study that jointly exploits the above-mentioned techniques to investigate the physical-layer end-to-end latency of a fast optical switching system.

In this paper, we design and demonstrate a real-time fast wavelength-switched transmission system with a minimum transceiver latency of 28 ns and an inter-packet gap of 18 ns. The use of a tunable laser only for wavelength-switching ensures lower power consumption and higher OSNR which is crucial for high baud rate signaling compared to the SOA array-based approach. This ultra-low transceiver latency and inter-packet gap are achieved by jointly harnessing techniques including optical clock synchronization [9,10,13], fast-wavelength switching [7] and several advantages of hollow core nested anti-resonant nodeless hollow core fiber (HC-NANF) versus SSMF [16–18]. Specifically, we jointly exploit the ultra-low thermal sensitivity and the low and flat dispersion of HCF to achieve a minimum clock phase shift during wavelength switching, which allows any signal processing latency associated with clock tracking and recovery to be eliminated [19]. By carefully characterizing the wavelength switching dynamics using frequency discriminators of different resolutions, we optimize the driver waveform to achieve fast wavelength switching with minimal overshooting, ensuring low inter-packet gap and change of group delay (GD). Our system design allows sub- 10^{-9} bit error ratio (BER) optical packet transmission, thus also eliminating any latency associated with FEC. In the demonstrated 100 m link (333 ns transmission time in HCF), compatible with an optically switched data center network, we achieved an overall HCF latency-dominated end-to-end latency of only 379 ns (transceiver + gap + fiber propagation latencies), which is mainly dominated by the fiber

propagation latency. An SSMF-based system would require 480 ns for light propagation in conjunction with 100s of nanoseconds transceiver time for CDR.

2. Proof-of-concept experiment for studying the latency and performance of optically-switched system

2.1. System architecture

The architecture of the envisaged fast wavelength-switching system is shown in Fig. 1. The system concept is the same as the previously demonstrated systems [9–13], except that a tunable laser is employed for wavelength switching. A centrally synchronized optical clock is sent to the top-of-the-rack (TOR) switches for clock frequency synchronization. Here we use intensity modulated optical clock signals that are transmitted over control plane fibers for clock synchronization. Other clock synchronization methods, such as radio over fiber [20], could also be exploited. The tunable lasers in the TOR are externally modulated by data signals. The tunable lasers are calibrated and driven with optimized control waveforms to achieve fast wavelength switching for the wavelength channels aligning to the array waveguide grating (AWG) for switching the physical connections to different TORs [7,8,12]. N AWGs are used for N racks to achieve meshed connections with non-blocking switching. The inter-packet gap, which depends on the switching time, is minimized to achieve a low bit error ratio (BER), as detailed in the experimental setup.




Fig. 1. Conceptual diagram of a fast wavelength-switching system with optical clock synchronization. DS-DBR lasers are employed in this work but other fast tunable lasers can be used for this architecture.

2.2. Proof-of-concept experiment

Due to the availability of components, we used a two-node system (see Fig. 2) with fast-wavelength switching for the proof-of-concept experiment. This simplification does not compromise the study of the physical latency as it contains all the latency elements detailed in the introduction. As shown in Fig. 2, the optical clock signal was generated by externally modulating a continuous wave (CW) laser emitting 12 dBm at 1550.8 nm using a quadrature-biased Mach-Zehnder modulator (MZM1) driven with a 400-MHz clock signal. This resulted in a 400-MHz intensity-modulated optical clock signal of 3 dBm. The optical clock signal was split by a 1:4 splitter and sent to the transmitter (Tx) and receiver (Rx) nodes via 2-m SSMF patch-cords and a 400-m HCF, respectively. The HCF was pigtailed on both ends with SSMF with total loss of 5.5 dB, including 4 dB connection loss and 1.5 dB transmission loss. To compensate for the fiber loss, an erbium-doped fiber pre-amplifier (EDFA1) was used to amplify the clock signal to 12 dBm

to ensure low clock jitter after transmitting over the 400-m of HCF. Recent advances in HCF technology have achieved 0.3 dB connection loss (both ends pigtailed) and 0.17 dB/km fiber loss, permitting practical use of HCF for large-scale clock synchronization [21,22]. At both the Tx and Rx nodes, the optical clock signals were detected by a PIN photodiode followed by a transimpedance amplifier (TIA). The input optical power of the photodetector was adjusted to -5.5 dBm for both Tx and Rx nodes and the corresponding 400 MHz electrical clock signal power was -7.5 dBm.

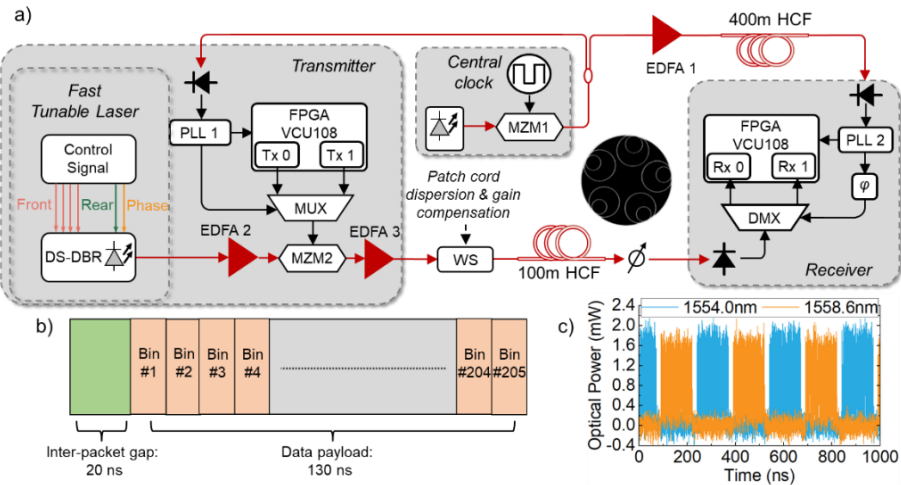


Fig. 2. (a) Experimental setup of wavelength switched clock synchronized data center interconnection; (b) packet structure; (c) received optical waveforms (orange: 1558.6 nm packet, blue: 1554 nm packet). MZM: Mach Zehnder modulator; PLL: Phase locked loop; MUX/DMX: RF multiplexer/demultiplexer; WS: waveshaper; PD: Photodetector; φ : Phase shifter.

The detected 400 MHz clock signals were fed into electronic phase lock loops (PLL1&2). Both PLL1 and PLL2 (TI LMX2595) generate a 400 MHz reference clock for the field programmable gate arrays (FPGAs, Xilinx VCU108) and a 25.6 GHz clock signal that drove electronic multiplexer and demultiplexer at the Tx and Rx, respectively. The Tx FPGA generated two streams of synchronized 25.6-GBd non-return-to-zero (NRZ) packets that were multiplexed to 51.2 GBd NRZ using an external RF multiplexer driven with the 25.6 GHz clock signal generated from PLL1. At the Rx node, the optical 51.2 GBd OOK packets were firstly detected by a 70 GHz photodetector then the electrical packets were amplified and then demultiplexed by an RF demultiplexer, which output two streams of 25.6 GBd OOK data, connecting to two GTX receivers (Rx1 and Rx2) on the receiver FPGA for packet detection and BER calculation. Similar to the Tx node, the 25.6 GHz clock generated by PLL2 was connected to RF demultiplexer for sampling. An RF phase shifter was inserted between the RF demultiplexer and PLL2 output to optimize the sampling phase.

The structure of the optical packets is shown in Fig. 2(b), which consisted of a 20 ns inter-packet gap (shown in orange) followed by a 130 ns payload (832 bytes) using a pseudorandom binary sequence (PRBS) of 2^9 length. The short 130 ns payload length (832 bytes) was used to emulate the small-packet dominated traffic patterns typical within production cloud data centers [10]. The duration of the inter-packet gap was determined by the time required to stabilize the power and wavelength of the DS-DBR laser (switching time), which was around 20 ns in our experiment. The DS-DBR laser switching dynamics characterization setup is described later in the sub-section 2.3. To ensure time synchronization between wavelength-switching and packet transmission,

buffers were implemented at the Tx FPGA to align the inter-packet gap with the beginning of wavelength switching event. At the receiver side, the payload was divided into 205×16 -bit bins for BER calculation [10]. The number of bit errors falling in each bin across all packets arriving at the receiver was recorded in real time.

For fast-wavelength switching, we used a digital supermode distributed Bragg reflector (DS-DBR) laser outputting 7 dBm CW light. The DS-DBR laser consists of one front grating and one rear grating for wavelength selection and a phase section for fine wavelength tuning [23]. As shown in Fig.2a, six channels of synchronized digital-to-analog converters (DAC) operating at 1 GSa/s were employed to configure the DS-DBR laser, including 4 channels for front grating selection and two channels for the rear grating and the phase section. The DACs were triggered by the distributed clock for time synchronization.

The wavelength of the DS-DBR laser can be tuned by over 40 nm by adjusting the current injection to the rear and front gratings [7]. We chose two representative switching pairs in this study. In the first case, the wavelength was switched between 1554 nm and 1558.6 nm by tuning the rear grating only. In the second case, both rear and front gratings were tuned to achieve a wavelength switching between 1544 nm and 1558.6 nm. For both cases, the driving waveforms to the rear and phase section current were pre-equalized to minimize switching time using a linear regression algorithm studied in [7]. The output of the DS-DBR laser was amplified by a polarization maintaining EDFA (EDFA2) before being modulated by 51.2 GBd electronic packet using a 40-GHz MZM biased at the quadrature point. The 51.2 GBd electronic packets were generated by time-interleaving two 25.6 GBd digital NRZ data packet streams using an external data multiplexer driven by the distributed clock. The optical data packets were subsequently amplified from 9 to 17 dBm using EDFA3 before transmitting through a waveshaper and a 100-m HCF. To study the impact of the HCF on system performance, the waveshaper was configured to compensate for the residual dispersion of the SSMF patch-cords as well as the limited gain-flatness due to the EDFAs. The EDFAs were employed due to the relatively low output power of the laser and the high insertion loss of waveshaper (9 dB). They could potentially be removed by using a tunable laser with a high-power booster semiconductor optical amplifier (SOA) and low-loss AWG.

Figure 2(c) shows the optical waveforms when the DS-DBR laser is configured to switch between 1554 and 1558.6 nm. The blue and orange curve corresponds to the 1554 nm and 1558.6 nm packets, respectively. Active CDR clock phase tracking was not needed due to the low thermal sensitivity of the HCFs. By reducing the electronic processing to sampling and detection only, we achieved a minimal transceiver latency of 28 ns [14,19,24].

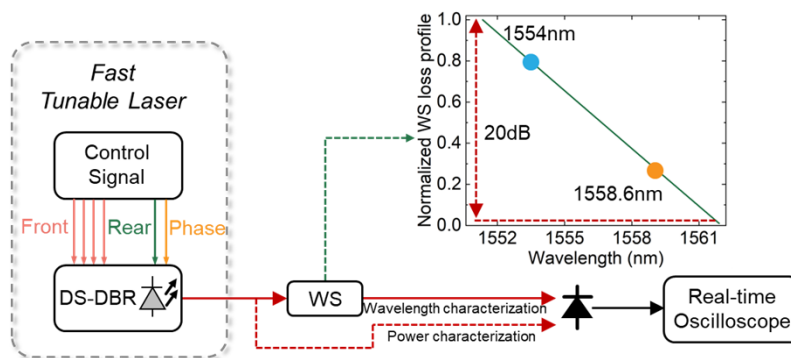


Fig. 3. Experimental setup for DS-DBR wavelength switching characterization. Inset shows the normalized waveshaper loss profile (in linear scale), using 1554 nm ↔ 1558.6 nm as an example.

2.3. Experimental setup for studying wavelength-switching dynamics

The switching dynamics has previously only been studied using optical power [7]. To study the wavelength switching dynamics, we developed an experimental setup shown in Fig. 3. The laser was configured to continuously switch between two wavelengths λ_1 (1544 or 1554 nm) and λ_2 (1558.6 nm) with 150-ns time interval. The output of the laser passed through a waveshaper configured as a linear wavelength discriminator for the two switching pairs described above. A 20-GHz bandwidth photodetector was used to detect output of the waveshaper and captured by a 100 GSa/s (20-GHz analog bandwidth) 10-bit real-time oscilloscope for analysis. The dynamic range of the waveshaper was 20 dB and the wavelength range was set to be 10 nm or 20 nm to cover the whole range of wavelength deviation during switching event.

3. Experimental results and discussion

3.1. DS-DBR wavelength switching characterization

Figures 4(a) and 4(b) show the measured DS-DBR lasing wavelength change during the switching events: 1554 nm \leftrightarrow 1558.6 nm and 1544 nm \leftrightarrow 1558.6 nm, respectively, over 150 ns time interval (a \leftrightarrow b means continuously switching between a and b). In both cases, the laser's wavelength reached the target wavelength (within <0.1 nm accuracy) within 18 ns. However, the wavelength could vary more than 3 nm during switching. In practice, this would lead to crosstalk to other wavelength channels during the switching. Figure 4(c) and 4(d) show the change of optical power during the switching events, showing about 2.5 mW power variation during switching and the power stabilizes 18 ns after switching. We further characterized the switching dynamics after passing through a 100-GHz-bandwidth AWG as the passive switching core in data center networks. The AWG was emulated by configuring the waveshaper to generate two 100-GHz-bandwidth super-Gaussian filter centered at two target wavelengths with an out-of-band suppression of 40 dB. Figure 4(e) and 4(f) show that when DS-DBR lasing wavelength drifts outside the passband of AWG, the optical power drops to zero, indicating burst-mode optical packet transmission requiring a minimum 18 ns inter-packet gap for switching. Receivers could ignore incoming information during the wavelength switching period to avoid wavelength crosstalk from other channels impacting data transmission, a process that is aided by central clock synchronization. This wavelength crosstalk could also be addressed by adding a shutter (e.g., reverse-biased semiconductor optical amplifier) at the output of the DS-DBR laser as a gate which absorbs the light during wavelength tuning as well as providing optical gain [25].

3.2. Impact of DS-DBR switching dynamics on BER

To study the impact of DS-DBR laser switching dynamics on data transmission performance, we measured the BER across the data packet using the setup shown in Fig. 2. To evaluate the BER at different time slots, we count the errors for each bin using 2.7×10^{12} packets, corresponding to 4.4×10^{13} bits for each bin, as shown in Fig. 5(a) (1554 \leftrightarrow 1558.6 nm) and Fig. 5(b) (1544 \leftrightarrow 1558.6 nm).

For both switching cases, the measured BER remained above 10^{-9} between 5 ns and 18 ns. This is due to both wavelength and power variation of the DS-DBR laser during switching events. From 18 ns to the end of packet, the BER remains below 10^{-9} because the DS-DBR laser becomes stabilized. These results, together with Fig. 4, characterize the impact of switching dynamics on transmission performance of optically switched DCI, concluding that at least an 18 ns inter-packet gap is required to ensure sub- 10^{-9} BER.

3.3. Impact of GD change on wavelength-switched data transmission

In conventional systems using SSMF or dispersion shifted fiber (DSF), the CDR clock phase has to be tracked due to wavelength dispersion at the telecom C band and the relatively high thermal




Fig. 4. DS-DBR laser wavelength switching characterization. DS-DBR laser wavelength variation during (a) 1554 nm ↔ 1558.6 nm switch, (b) 1544 nm ↔ 1558.6 nm switch. DS-DBR power variation (without 100 G AWG) during (c) 1554 nm ↔ 1558.6 nm switch, (d) 1544 nm ↔ 1558.6 nm switch. DS-DBR power variation (with 100 GHz AWG) during (e) 1554 nm ↔ 1558.6 nm switch, (f) 1544 nm ↔ 1558.6 nm switch

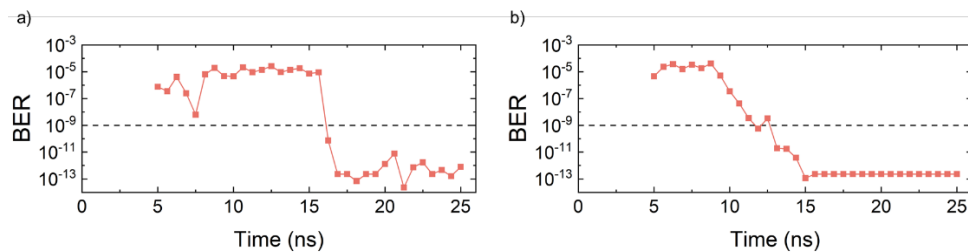


Fig. 5. BER distribution of data packet. (a) 1554 nm ↔ 1558.6 nm switch, (b) 1544 nm ↔ 1558.6 nm switch.

sensitivity of SSMF. Our HCF addresses both these issues by providing low and flat dispersion as well as 20 times lower thermal sensitivity compared to SSMF. Our system simultaneously leverages both features to achieve data transmission without tracking clock phase for data sampling, eliminating any additional latency associated with CDR.

We first evaluate the sampling clock tolerance of the 51.2 GBd OOK packets by manually tuning the sampling clock phase and record the BER of the packets at 1558.6 nm for the 1554 ↔ 1558.6 nm switching pair. We record a tolerance to sampling clock misalignment of ± 3.7 ps, indicating that any wavelength change or thermally-induced latency should be less than 7.4 ps. Figure 6(b) shows the measured GD of the 100 m HCF used in this experiment, showing a GD slope of about 0.19 ps/nm. Thus, the 1554 nm ↔ 1558.6 nm (4.6 nm spacing) and the 1544 nm ↔ 1558.6 nm (14.6 nm spacing) switches only resulted in a change of propagation

latency by about 0.9 ps and 2.5 ps, respectively. This is within the 7.4 ps tolerance of the clock phase offset. Considering a 2.5 ps maximum clock phase offset and a thermal coefficient of delay of 2 ps/(km·°C) [17,18], the 500-m HCF link (100 m data transmission and 400 m clock transmission) indicates a tolerance to temperature variation of about 9.8 °C, which is higher than the recommended data center temperature operation range requirement (18-27 °C) [26]. To operate beyond this temperature range or with increased fiber length, one could apply clock phase caching [10] that synchronizes the sampling clock or exploit advanced coating or fiber designs that further reduce the TCD [27,28].

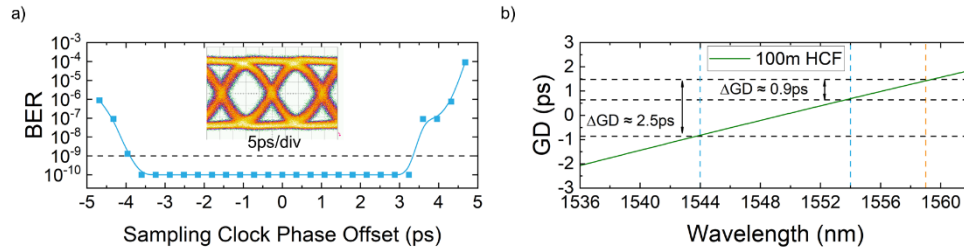


Fig. 6. (a) BER tolerable sampling clock phase offsets, (b) measured GD of the 100-m HCF.

Figure 7(a) and 7(b) show the average BER of the data packets using HCF (closed markers connected with solid lines) and SSMF (open marker with dashed lines). In this case, the DS-DBR laser was configured to switch between 1554 nm and 1558.6 nm (4.6 nm spacing). Each BER point was calculated by recording the error bits over 10 seconds (a total of 1×10^{14} bits). The optimal sampling clock phase was calibrated once for the 1558.6-nm data packets at the beginning of the measurement. When transmitting through 100-m HCF, the BER for the 1558.6-nm packets was less than 10^{-11} whilst the BER for the 1554-nm packets was slightly higher due to the offset of the sampling clock phase. Both wavelengths achieved sub- 10^{-9} BER without any clock phase tracking when using HCF. On the other hand, the BER for the 1554-nm packets when transmitting through 100 m SSMF was significantly degraded to about 5×10^{-4} due to larger dispersion induced clock phase offset, which was about 7.6 ps when assuming CD of 16.5 ps/(nm·km) for SSMF.




Fig. 7. (a) Packet BER over 15 minutes for 1554 nm ↔ 1558.6 nm and 1544 nm ↔ 1558.6 nm switch, (b) packet BER comparison between 100 m HCF and SSMF.

4. Conclusion

We combine technologies including optical clock synchronization, fast-tunable lasers and HCF and demonstrate the first CDR-free fast-wavelength-switched packet transmission for low-latency data center interconnection. We systematically study wavelength and power dynamics during switching, showing that the wavelength of DS-DBR laser changes more than 3 nm during this process. Our-proof-of-concept experiment achieved 18 ns switching time for 51.2 Gb/s

burst-mode transmission, limited by the power and wavelength variation during laser switching in an AWG based optical switching system. We further show that our CDR-free system can operate with sub- 10^{-9} BER without any active clock phase tracking, which is enabled by the low and flat dispersion as well as the low thermal sensitivity of HCF. This permits eliminating any transceiver latency due to clock synchronization, equalization, and FEC, leading to an ultimate low transceiver latency of 28 ns due to sampling and detection only. The results and studies reported here can serve as a basis for studying the latency and performance of future optically switched networks. The cost drivers of the optically switched system under investigation predominantly reside in the tunable lasers and the Hollow Core Fibre (HCF). While the present cost of individual tunable lasers is high for DC application, successful demonstrations of laser array integration promise a compact footprint and a potential significant reduction in cost at volume production. Similarly, HCF, despite its inherently lower material cost compared to its solid core counterparts, is undergoing a similar evolution. With the recent breakthrough in achieving record low loss HCF, the emphasis has shifted towards mass production, thereby attracting substantial investment. The benefits showcased in this study provide an exemplar for the adoption and mass production of HCF in future data center networks.

Funding. National Natural Science Foundation of China (62102343); Engineering and Physical Sciences Research Council (EP/P030181/1, EP/R035342/1, EP/V051377/1).

Disclosures. The authors have no conflicts of interest to declare.

Data availability. Data underlying the results presented in this paper are available in Ref. [29].

References

1. H. J. Dorren, E. H. Wittebol, R. D. Kluijver, G. G. D. Villota, P. Duan, and O. Raz, "Challenges for optically enabled high-radix switches for data center networks," *J. Lightwave Technol.* **33**(5), 1117–1125 (2015).
2. X. Zhou, R. Urata, and H. Liu, "Beyond 1 Tb/s Intra-Data Center Interconnect Technology: IM-DD OR Coherent?" *J. Lightwave Technol.* **38**(2), 475–484 (2020).
3. Cisco Visual Networking, "Cisco annual internet report (2018–2023)," White paper, Cisco Public (2018).
4. I. White, E. Aw, K. Williams, H. Wang, A. Wonfor, and R. Penty, "Scalable optical switches for computing applications [Invited]," *J. Opt. Netw.* **8**(2), 215–224 (2009).
5. L. Qiao, W. Tang, and T. Chu, " 32×32 silicon electro-optic switch with built-in monitors and balanced-status units," *Sci. Rep.* **7**(1), 42306 (2017).
6. Q. Cheng, S. Rumley, M. Bahadori, and K. Bergman, "Photonic switching in high performance datacenters [Invited]," *Opt. Express* **26**(12), 16022–16043 (2018).
7. T. Gerard, H. Dziencial, J. Benjamin, K. Clark, H. Williams, B. Thomsen, D. Lavery, and P. Bayvel, "Packet Timescale Wavelength Switching Enabled by Regression Optimisation," *IEEE Photon. Technol. Lett.* **32**(8), 477–480 (2020).
8. T. Verolet, S. Almonacil, M. Szczerban, J. Estaran, R. Boddeda, F. Boitier, J. Provost, H. Mardoyan, F. Blache, J. Decobert, S. Olivier, A. Shen, and S. Bigo, "Ultra-Fast Tunable Laser Enabling 4 ns Coherent Slot Switching Beyond 100 Gbit/s," *European Conference on Optical Communications* (2020).
9. H. Ballani, P. Costa, R. Behrendt, D. Cletheroe, I. Haller, K. Jozwik, F. Karinou, S. Lange, K. Shi, B. Thomsen, and H. Williams, "Sirius: A flat datacenter network with nanosecond optical switching," *Association for Computing Machinery*, 782–797 (2020).
10. K. A. Clark, D. Cletheroe, T. Gerard, I. Haller, K. Jozwik, K. Shi, B. Thomsen, H. Williams, G. Zervas, H. Ballani, P. Bayvel, P. Costa, and Z. Liu, "Synchronous subnanosecond clock and data recovery for optically switched data centers using clock phase caching," *Nat. Electron.* **3**(7), 426–433 (2020).
11. Z. Hu, Z. Zhou, C. C. -K. Chan, and Z. Liu, "Equalizer State Caching for Fast Data Recovery in Optically-Switched Data Center Networks," *J. Lightwave Technol.* **39**(17), 5362–5370 (2021).
12. A. S. Raja, S. Lange, M. Karpov, K. Shi, X. Fu, R. Behrendt, D. Cletheroe, A. Lukashchuk, I. Haller, F. Karinou, B. Thomsen, K. Jozwik, J. Liu, P. Costa, T. J. Kippenberg, and H. Ballani, "Ultrafast optical circuit switching for data centers using integrated soliton microcombs," *Nat. Commun.* **12**(1), 5867 (2021).
13. Z. Zhou, K. A. Clark, C. Deakin, and Z. Liu, "Clock Synchronized Transmission of 51.2 GbD Optical Packets for Optically Switched Data Center Interconnects," *J. Lightwave Technol.* **40**(6), 1735–1741 (2022).
14. Xilinx, DS893 (v1.12) Virtex UltraScale Architecture Data Sheet: DC and AC Switching Characteristics (2019); https://www.xilinx.com/support/documentation/data_sheets/ds893-virtex-ultrascale-data-sheet.pdf.
15. O. S. Sella, A. W. Moore, and N. Zilberman 2018. "FEC Killed The Cut-Through Switch," *Workshop on Networking for Emerging Applications and Technologies* (2018).
16. F. Poletti, N. Wheeler, M. Petrovich, N. Baddela, E. Numkam Fokoua, J. R. Hayes, D. R. Gray, Z. Li, R. Slavík, and D. J. Richardson, "Towards high-capacity fibre-optic communications at the speed of light in vacuum," *Nat. Photonics* **7**(4), 279–284 (2013).

17. R. Slavík, G. Marra, E. Fokoua, N. Baddela, N. Wheeler, M. Petrovich, F. Poletti, and D. J. Richardson, "Ultralow thermal sensitivity of phase and propagation delay in hollow core optical fibres," *Sci. Rep.* **5**(1), 15447 (2015).
18. K. A. Clark, Y. Chen, E. R. Fokua, T. Bradley, F. Poletti, D. J. Richardson, P. Bayvel, R. Slavik, and Z. Liu, "Low thermal sensitivity hollow core fiber for optically-switched data centers," *J. Lightwave Technol.* **38**(9), 2703–2709 (2020).
19. Z. Zhou, H. Dziecil, K. Clark, R. Slavík, H. Sakr, K. Harrington, D. J. Richardson, F. Poletti, and Z. Liu, "Low-latency Optically-switched Data Centre Interconnects enabled by Hollow Core Anti-resonant Fiber," *Conference on Lasers and Electro-Optics* (2022).
20. Z. Zhou, D. Nopchinda, M. -C. Lo, I. Darwazeh, and Z. Liu, "Simultaneous Clock and RF Carrier Distribution for Beyond 5 G Networks Using Optical Frequency Comb," *European Conference on Optical Communication* (2022).
21. D. Suslov, E. Fokoua, D. Dousek, A. Zhong, S. Zvánovec, T. Bradley, F. Poletti, D. J. Richardson, M. Komanec, and R. Slavík, "Low loss and broadband low back-reflection interconnection between a hollow-core and standard single-mode fiber," *Opt. Express* **30**(20), 37006–37014 (2022).
22. G. T. Jasion, H. Sakr, J. R. Hayes, S. R. Sandoghchi, L. Hooper, E. Fokoua, A. Saljoghei, H. C. Mulvad, M. Alonso, A. Taranta, T. D. Bradley, I. A. Davidson, Y. Chen, D. J. Richardson, and F. Poletti, "0.174 dB/km Hollow Core Double Nested Antiresonant Nodeless Fiber (DNANF)," *Optical Fiber Communications Conference and Exhibition* (2022).
23. A. J. Ward, D. J. Robbins, G. Busico, E. Barton, L. Ponnampalam, J. P. Duck, N. D. Whitbread, P. J. Williams, D. C. J. Reid, A. C. Carter, and M. J. Wale, "Widely tunable DS-DBR laser with monolithically integrated SOA: design and performance," *IEEE J. Select. Topics Quantum Electron.* **11**(1), 149–156 (2005).
24. Xilinx GTY TRx latency: <https://support.xilinx.com/s/article/66341>, (2021).
25. T. Gerard, C. Parsonson, Z. Shabka, B. Thomsen, P. Bayvel, D. Lavery, and G. Zervas, "AI-optimised tuneable sources for bandwidth-scalable, sub-nanosecond wavelength switching," *Opt. Express* **29**(7), 11221–11242 (2021).
26. Data Center Networking Equipment—Issues and Best Practices TC9.9 (ASHRAE, 2016); https://tc0909.ashraetcs.org/documents/ASHRAE_TC0909_Power_White_Paper_22_June_2016_REVISED.pdf.
27. E. N. Fokoua, M. N. Petrovich, T. Bradley, and R. Slavík, "How to make the propagation time through an optical fiber fully insensitive to temperature variation," *Optica* **4**(6), 659–668 (2017).
28. M. Bousonville, M. K. Bock, M. Felber, T. Ladwig, T. Lamb, H. Schlarb, S. Schulz, and C. Sydlo, "New phase stable optical fiber," *Beam Instrumentation Workshop* 101–103 (2012).
29. Z. Zhou, H. Dzieciol, K. Clark, Y. Luo, D. Richardson, F. Poletti, P. Bayvel, R. Slavik, and Z. Liu, "Dataset for "Low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber"," *UCL RPS* (2023), <https://doi.org/10.5522/04/21789161>.