

# LOW THERMAL SENSITIVITY HOLLOW CORE FIBRE FOR OPTICALLY-SWITCHED DATA CENTRE APPLICATIONS

*Kari A Clark<sup>1\*</sup>, Yong Chen<sup>2</sup>, Eric R Numkam Fokoua<sup>2</sup>, Tom Bradley<sup>2</sup>, Francesco Poletti<sup>2</sup>, David J. Richardson<sup>2</sup>, Polina Bayvel<sup>1</sup>, Radan Slavík<sup>2</sup>, and Zhixin Liu<sup>1</sup>*

<sup>1</sup>*Department of Electronic and Electrical Engineering, University College London, London, United Kingdom*

<sup>2</sup>*Optoelectronics Research Centre, University of Southampton, Southampton, United Kingdom*

*\*kari.clark.14@ucl.ac.uk*

**Keywords:** Optical fiber communication, Optical fibers, Optical packet switching, Data Centers, Clock recovery

## Abstract

We demonstrate 20-times greater tolerance to temperature variation using hollow core fibres compared to SMF-28 in a fast optical switching system. With frequency and once-only phase synchronisation, we obtained error-free transmission of short packets with <625ps clock recovery time in 25.6 Gb/s real-time systems.

## 1 Introduction

In today's data centres, data traffic switching is performed in the electronic domain, which has several drawbacks including limited scalability, high power consumption, and high latency [1]. Optical switching has the potential to achieve future hyper-scale intra-data centre interconnections, offering large number of input/output ports, high bandwidth, and a significantly reduced number of opto/electronic conversions, reducing power consumption, cost, and latency [1]. Recent research has highlighted that one of the major challenges to implementing optical switching in a data centre is the need for fast clock and data recovery (CDR). When two network nodes communicate via a dynamic optical path through an optical switch, the CDR locking time can be much longer than the time required to transmit short data packets, greatly reducing the overall data throughput [1]. Sub-nanosecond CDR locking time is therefore required to overcome the throughput bottleneck [2].

CDR requires recovery of both the frequency and phase of the clock embedded in received data signals. The synchronisation of transceiver clock frequencies can be achieved by distributing an optical reference clock to all network nodes. Nevertheless, as shown in [2], clock phase recovery can still take up to 10 s of nanoseconds because the CDR module must search through a range of clock phases for each new incoming data signal. This search for the correct sampling clock phase is necessary because of the change of the signal propagation delay due to environmental temperature variation. In a practical data centre environment, the operating temperature range could vary from 18 to 27 °C [3] and the rate of change in a typical production cloud data centre could be up to 0.03 K/s [2]. For standard single mode fibres (SMF-28), the typical rate of temperature-induced propagation delay change is about 40 ps/(km·K) [4]. Assuming a typical 1-km short-reach data centre link, a 0.5 °C temperature change could lead to a data delay change of half a symbol period for a 25 Gbaud signal, causing bit errors if the clock phase remains constant.

Previous research proposed a clock phase caching method, which configures the transceivers to memorise and then track the relative clock phase offset between different transceiver pairs, enabling sub-nanosecond CDR locking time [2]. An alternative to the phase caching method is to use a transmission fibre insensitive to the temperature change, such that a fixed clock phase can be kept without any update.

A very promising low-thermal sensitivity optical fibre is hollow core fibre (HCF), which has been demonstrated to have about 20 times lower thermal coefficient of delay than SMF-28. In SMF-28, the 40 ps/(km·K) thermal coefficient is predominantly caused by the thermo-optic effect of silica glass, i.e. the change of refractive index with temperature, which accounts for about 38 ps/(km·K). The remaining 2 ps/(km·K) is due to thermally-induced fibre elongation. As light propagates in HCFs mostly through the air (rather than silica glass), the thermo-optic contribution is negligible, making the thermal coefficient of HCF ~2 ps/(km·K), caused primarily by the fibre elongation. A specific design of HCF was even demonstrated to have a thermal coefficient of zero [5]. Although this was achieved only at a specific wavelength, it opens a new opportunity for ultra-fast clock recovery that is immune to the impact of temperature variation.

In this paper, we experimentally investigate the performance of HCF in clock-synchronised optical transmission systems, with a focus on the tolerance to temperature variation. We first studied the temperature tolerance in a straight-line transmission system without any active clock phase tracking. We then measured and compared the clock phase variation of HCF with SMF-28 in a lab environment, followed by a 2×1 optically-switched system demonstration with sub-nanosecond clock recovery achieved without any clock phase update. Our experiments show that HCF offers 20 times higher resilience to temperature variation than conventional SMF-28 in a clock-synchronised optical switching system with sub-nanosecond clock recovery.

## 2 Experimental set-up

### 2.1 Point-to-point clock synchronised transmission via HCF

We first study the effect of the HCF low thermal sensitivity on data transmission in a point-to-point transmission experiment as shown in Fig. 1. An optical clock signal was generated by modulating an 800 MHz reference clock onto a 1550 nm optical carrier, via a LiNbO<sub>3</sub> Mach-Zehnder modulator (MZM). The optical clock signal was split into two components, with one component distributed to the transmitter (Tx FPGA, Xilinx VCU108) and the other to the receiver. The component of the optical clock signal sent to the transmitter was passed through a first 1 km (fully connectorised) length of HCF to emulate the worst case scenario where the clock source is adjacent to the receiver and far from the transmitter, causing the change of temperature to result in clock phase drift in the clock and data paths that is entirely additive. The optical clock signals were converted to 800 MHz electrical clock signals for frequency synchronisation of the two FPGAs through on-chip digital phase lock loop (PLL) modules. We programmed the Rx CDR to record the detected clock phase at the receiver. At the transmitter side, the externally modulated laser (EML) emitted a 13 dBm continuous wave (CW) signal at 1555 nm which was modulated by 25.6 Gb/s NRZ data of PRBS length of  $2^{31}$  by a 35-GHz electro-absorption modulator, outputting a 3 dBm modulated signal that was launched into a second 1 km length of (fully connectorised) HCF. The two HCFs used were 19-cell hollow core photonic bandgap fibre (HC-PBGF). The end-to-end loss of the data and clock paths were 6.5 dB and 7.0 dB, respectively, consisting of about 3 dB connector loss and 4 dB fibre loss. At the receiver side, the signal was detected by a 20-GHz bandwidth photodiode followed by a trans-impedance amplifier before the real-time FPGA receiver. The optical power was intentionally attenuated to -10.5 dBm to match intra-data centre transmission standards [6]. Bit errors were counted in real-time and used to calculate the bit error ratio (BER).

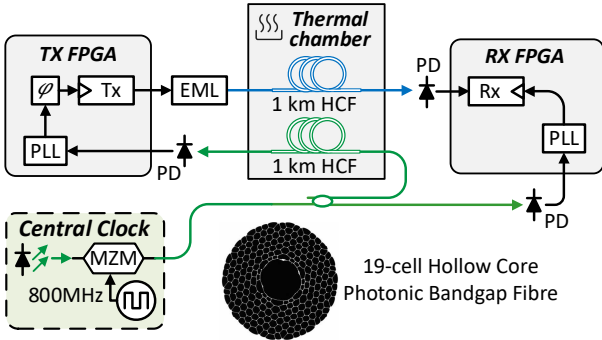


Fig. 1: Point-to-point HCF synchronisation experiment. PLL – Phase locked loop, EML – externally modulated laser, MZM – Mach Zehnder Modulator, HCF – hollow core fibre, PD - photodiode.

To investigate the impact of temperature change on our system, both HCF spools were contained within a thermally controlled chamber with less than 0.05 °C temperature fluctuation. The chamber was initially stabilised at 33.5 °C and the receiver data sampling clock phase was set to be half a symbol period away from the optimum value (the highest Q factor). The initial temperature of 33.5 °C was chosen to ensure a significant offset from ambient temperature, maximising the thermal chamber temperature stability at each

point. We then increased the chamber temperature and recorded the BER after the temperature became stabilised (within  $\pm 0.05$  °C). Each BER point was measured from  $7.68 \times 10^{12}$  bits.

### 2.2 Optically-switched prototype system with HCF

We next investigated the effect of the low thermal sensitivity of HCF on BER and CDR locking time with a 2x1 optical switch between two transmitters, Tx0 and Tx1 and a receiver, Rx, based on FPGAs, as shown in Fig. 2. In this experiment, the optical clock was sent to Tx0 through 1 km of HCF, and directly to Tx1 and the Rx. The packets were formed by PRBS signals with a length of  $2^9$  to emulate the 64-byte short packets of data centre traffic. The optical packets were generated by driving EML0 and EML1 with the 25.6 G/s NRZ signals produced by Tx0 and Tx1. Their output optical signals were at 1555 nm and 1553 nm, respectively. The optical packets from both transmitters had 60 ns duration and were interleaved by a 2x2 LiNbO<sub>3</sub> MZM switch with 5 ns gap. The interleaved packets leaving the switch were amplified by an erbium doped fibre amplifier (EDFA) to compensate for the loss of the LiNbO<sub>3</sub> switch (about 6 dB). The interleaved packets then passed through 1 km of HCF before being received by the Rx node. Like the point-to-point experiment, both HCF fibre spools were contained within the thermally controlled chamber that was initially stabilised at 32.5 °C. This initial temperature was again chosen to maximise thermally controlled chamber stability. The clock phases of the two transmitters were set to be a half symbol period apart from their respective optimum values. We then measured the BER and CDR locking time at different temperature values. In this experiment, the CDR in the receiver was running constantly. As a result, only the clock phase offset between the two transmitters contributed to the BER and CDR locking time, which is entirely due to the 1 km length of HCF between the central clock and the clock input to Tx0.

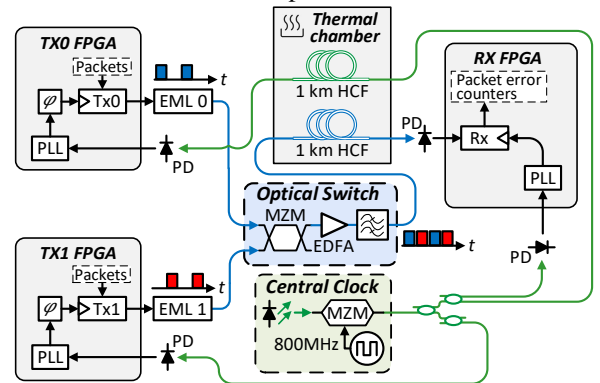


Fig. 2: Synchronous HCF 2x1 optically switch experimental setup. EML – externally modulated laser, EDFA – erbium doped fibre amplifier.

To calculate the overall BER and CDR locking time, the first and second sequences within each packet are divided into  $64 \times 16$ -bit bins, and the number of bit errors falling in each bin is recorded in real-time at each temperature setpoint until  $10^{13}$  total bits (summed across all bins) have been received. CDR locking time was calculated as the first bin in the packet with a BER of  $< 10^{-10}$ , if all following bins also have BER of  $< 10^{-10}$ . The overall receiver BER was calculated by summing the errors falling in all 16-bit bins.

### 3 Results and discussion

Fig. 3 shows the impact of temperature change on the BER of the point-to-point clock-synchronised transmission system. The blue markers show the measured BERs using the HCFs. The inset shows the eye diagram captured from the receiver FPGA. When temperature increases, the counter propagated clock path and data path led to a total change of the receiver clock phase shift of about 4 ps/K (total HCF length of 2 km). We observed an error free (BER  $<10^{-12}$ ) temperature range of about 2 K (35 to 37 °C). To compare with SMF-28, we assumed a thermal coefficient of 40 ps/(km·K) and calculated the tolerance to temperature variation based on the measured BER at different clock phase offset values from our previous work [2]. Shown as cross markers in Fig. 3, the error free temperature range using SMF-28 was only about 0.1 °C, a factor of approximately 20 smaller than HCF.

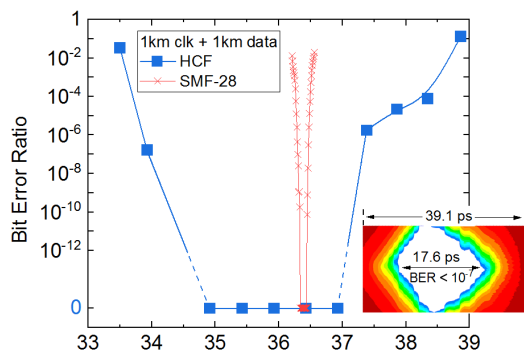


Fig. 3: Impact of temperature variation on the performance of a point-to-point clock-synchronised transmission system. Blue marker: BER using HCF; Red cross marker: using SMF-28.

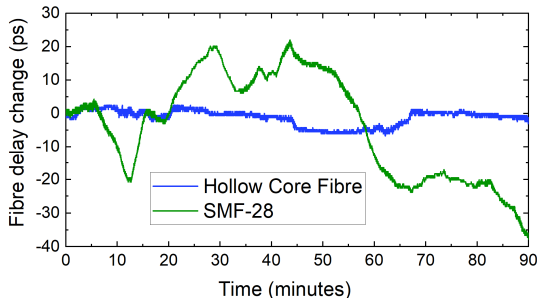


Fig. 4: Comparison of clock phase stability. Blue: HCF; Green: SMF-28.

Fig. 4 compares the fibre delay variation of the 2 km of HCF vs 2 km of SMF-28 under ambient lab temperature conditions. Over the one and a half hours observation period, the ambient temperature fluctuation was about 1 °C. The SMF-28 link caused significantly more change in fibre delay shift, consistent with our BER results in Fig. 3.

Fig. 5 shows the BER and CDR locking time of the 2×1 optical switch system experiment. The CDR locking time was calculated from BER as discussed in section 2.2. In contrast to [2], the phase of the transmitters is not updated to adapt to changing temperature, causing clock phase mismatch between the two transmitters. Using HCF, we observed a window of 4 K temperature change over which the received data has a BER  $<10^{-12}$  and a CDR locking time of  $<625$  ps. The tolerance to temperature variation is consistent with the clock phase offset caused by the 1 km HCF length difference between the

two clock to receiver via transmitter paths, as shown in the experimental setup in Fig. 2.

The 20 times lower thermal coefficient makes HCF a robust transmission medium for clock-synchronised data centre networks. In addition, HCF has the natural advantage of 1.45 times faster transmission speed than SMF-28, benefiting the need for low latency in data centre networks. Although the current loss of HCF is higher than SMF-28 (lowest demonstration of 1.3 dB/km [7]), the total loss of 2.6 dB for a 2 km fibre link is already under the allowable channel insertion loss budgets for intra-data centre interconnection standards [6]. We estimate that a system operating with HCF with a round-trip data + clock path of 200 m, sufficient to interconnect a cluster, would be able to withstand a proportionally larger temperature range of 20 K. This would potentially allow the possibility for a centrally-synchronised optical switch to operate indefinitely with calibration of the clock phases only at power up.

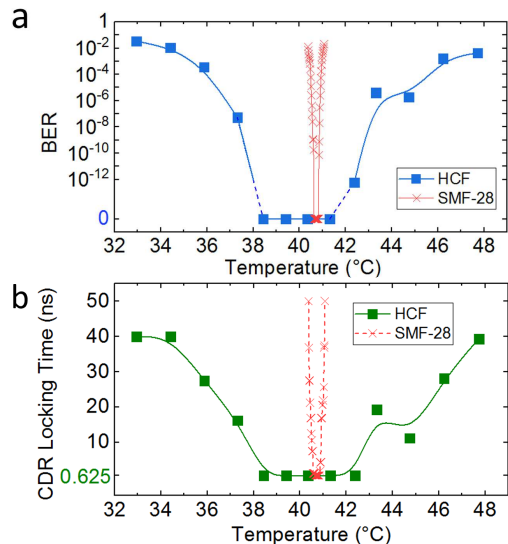


Fig. 5 Impact of temperature variation on the performance of an optically-switched system. Blue marker: BER using HCF; Red cross marker: using SMF-28. Green: CDR locking time

### 4 Conclusion

We demonstrated 20 times greater tolerance to temperature change using HCF for clock-synchronised data transmission systems. Error free transmission with  $<625$ ps clock recovery was achieved without the need to actively track the clock phase of incoming optical packets. HCF provides a robust solution to data centre interconnection and could support the growing need for temperature tolerance and transmission length of large-scale optically-switched data centre networks.

### 5 Acknowledgements

The authors acknowledge support from EPSRC Grants AirGuide Photonics (EP/P030181/1), TRANSNET (EP/R035342/1), PhotoDAC (EP/R041792/1) and the ERC Lightpipe project (682724), in addition to the funding provided by Microsoft Research. The authors also thank Dr Paolo Costa for his discussion.

## 6 References

- [1] Ballani, H., Costa, P., Haller, I., et al.: 'Bridging the last mile for optical switching in data centers'. Proc. Optical Fibre Communication Conference, San Diego, USA, March 2018.
- [2] Clark K., Ballani, H., Bayvel, P., et al.: 'Sub-Nanosecond Clock and Data Recovery in an Optically-Switched Data Centre Network'. Proc. European Conference on Optical Communication, Rome, Italy, September 2018.
- [3] ASHRAE TC9.9: 'Data Center Power Equipment Thermal Guidelines and Best Practices', 2016.
- [4] Slavík, R., Giuseppe, M., Fokoua, E.N., et al.: 'Ultralow thermal sensitivity of phase and propagation delay in hollow core optical fibres', Scientific Reports 5, 15447, 2015.
- [5] Fokoua, E.N., Petrovich, M.N., Bradley, B., et al.: 'How to make the propagation time through an optical fiber fully insensitive to temperature variations', Optica, 2017, 4, (6), pp. 659.
- [6] IEEE 802.3bm-2015: 'IEEE Standard for Ethernet - Amendment 3: Physical Layer Specifications and Management Parameters for 40 Gb/s and 100 Gb/s Operation over Fiber Optic Cables', 2015.
- [7] Bradley, T.D., Hayes, J.R., Chen, Y., et al.: 'Record low-loss 1.3 dB/km data transmitting antiresonant hollow core fibre'. Proc. European Conference on Optical Communication, Rome, Italy, September 2018.