Clock Synchronized Transmission of 51.2 GBd Optical Packets for Optically Switched Data Center Interconnects

Zichuan Zhou, *Student Member, IEEE*, Kari Clark, *Member, IEEE*,
Callum Deakin, *Student Member, IEEE*, and Zhixin Liu, *Senior Member, IEEE*

(Top-scored Paper)

***Abstract*— Optical switching has attracted significant attention in recent research on data center networks (DCNs) as it is a promising viable route for the further scaling of hyper scale data centers, so that DCNs can keep pace with the rapid growth of machine-to-machine traffic. It has been shown that optical clock synchronization enables sub-nanosecond clock and data recovery time and is crucial to high performance optically switched DCN. Moreover, the interconnect data rate is expected to increase from the current 100 Gb/s per fiber to scale to 800 Gb/s and beyond, requiring high baud rate signaling at >50 GBd. Thus, future optically switched DCN should support >50 GBd data transmission with optical clock synchronization. Here, we demonstrate the clock-synchronized transmission of 128-byte optical packets at 51.2 GBd and study the impact of reference clock phase noise on system performance, focusing on the tolerance to the clock phase misalignment that affects the system scalability and reliability. By comparing the tolerable sampling clock phase offsets using different reference clocks, we show that a clock phase offset window of about 8 ps could be achieved with a <0.2ps source clock. Furthermore, we model and numerically study the de-correlation of clock phase noise. This allows the total jitter to be estimated, and thereby, the estimation of the transmission performance for future generations of high baud rate, clock synchronized DC interconnects.**

***Index Terms*— Optical Fiber Communication, Clock Synchronization, Optical Switching, Data Centers, Phase Noise**

I. Introduction

THE scaling of data centers (DC) to millions of servers requires scalable, low power consumption and low latency switching for networking. Optical switching is a promising alternative for future DC networks, offering potential advantages including a flat network architecture, transparency to modulation formats, and a potential reduction of power consumption by up to 74% [1], [2]. Previous research has shown that the throughput of optically-switched DC networks is limited by the clock and data recovery (CDR) locking time [3], and proposed optical clock synchronization combined with clock phase caching to overcome this limit, demonstrating sub-nanosecond CDR locking time for achieving >90% network throughput [2], [3]. Clock phase caching has been demonstrated at 25 GBd for both the on-off keying (OOK) [3]-[6] and 4-level pulse amplitude modulation (PAM4) [2], [7] signaling formats. Nevertheless, the drastic growth of data traffic in cloud data centers is driving the signaling baud rate from the current 25 GBd to 50 GBd and possibly 100 GBd in the future [8]-[11]. Therefore, future optically switched data center networks must support >50 GBd optical packet transmission with sub-nanosecond CDR locking time.

The operational principle of clock phase caching shown in [2]-[4], using distribution of an optical clock through control plane fibers combined with clock phase synchronization of transceivers, is compatible with higher baud rates and high-order modulation formats. It can effectively track and compensate for the slow (e.g., <10 Hz) clock phase change due to the temperature dependent fiber time-of-flight and reference clock's random frequency walk [5], [12] and [13], enabling <625 ps CDR locking time for high baud rate signals. However, the timing jitter, or more specifically, the high frequency component (e.g., 10 Hz – 10 MHz) of the clock phase noise, cannot be tracked and compensated by the clock phase caching. As a result, the quadratic increase of clock phase noise with frequency will fundamentally limit the performance of high baud rate, clock synchronized, optical packet transmission systems.

For a given transceiver pair, the timing jitter arises from two sources: a) random temperature variation that changes the light propagation delay in optical fibers [12], and b) the jitter of the clock sources to which the transmitter and receiver are both synchronized. The temperature induced timing jitter is low frequency and, therefore, can be compensated by using clock phase caching [3] or significantly reduced by using low thermal sensitivity fibers [5], [6]. The clock source induced jitter, however, is primarily high frequency and thereby, hard to compensate. Although the impact of the source clock jitter on

Manuscript received xx xx, 2021. The authors acknowledge EPSRC (EP/R035342/1, EP/R041792/1, EP/V051377/1) for the funding support.

The authors are with the Department of Electronic and Electrical Engineering, University College London, Gower Street, Bloomsbury, London WC1E 6BT, U.K. (e-mail: zczlzz0@ucl.ac.uk; zhixin.liu@ucl.ac.uk).

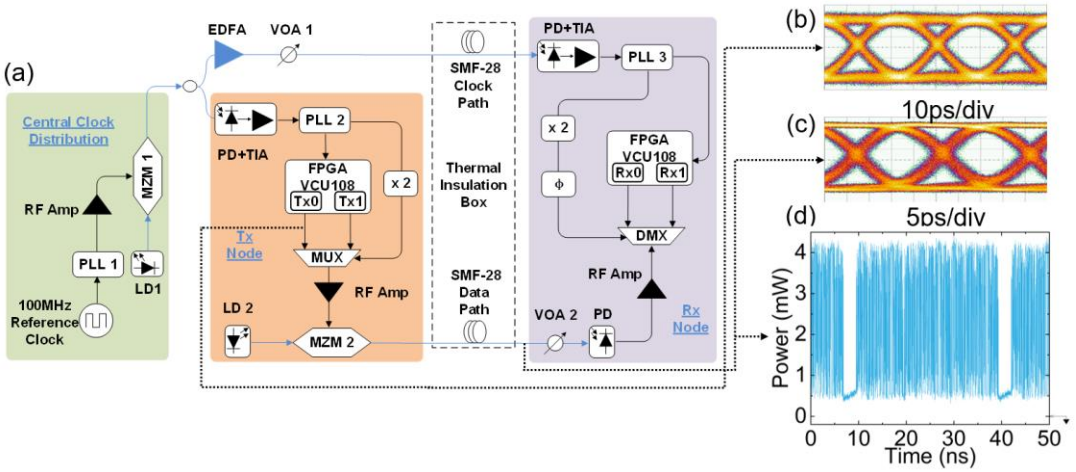

Fig. 1. a) Experimental set-up. PLL: Phase-lock-loop, TIA: transimpedance amplifier, VOA: variable optical attenuator, PD: photodetector, MZM: Mach-Zehnder Modulator, EDFA: Erbium-doped fiber amplifier. b) Eye diagram of 25.6 GBd OOK signal generated from the FPGA transmitter. c) Eye diagram of multiplexed 51.2 GBd optical OOK signal. d) Transmitted optical packets at 51.2 GBd.

system performance was insignificant in low baud rate interconnects (e.g., <25 GBd), it becomes prominent as the baud rate for DC interconnects increases to 50 GBd and even >100 GBd [8]-[11]. It is therefore important to study the jitter and timing tolerance of clock synchronized networks for high baud rate signaling in optically switched DC networks

In this paper, we demonstrate a real-time clock-synchronized optical packet transmission system at 51.2 GBd and study of the source clock jitter's impact on such systems. The symbol rate is doubled compared to previous demonstration in [3]-[6]. It extends our contributed paper at the optical fiber communication conference (OFC) [14] with added experimental details as well as the modeling and simulation of the phase noise de-correlation, based on which one can estimate the system BER performance for a given modulation format. To show the performance of the system with different source clock jitter, four different clock sources were used to study the tolerance to clock timing misalignment and the transmission BER performance. The real-time system allows us to study the system performance using a target BER of $<10^{-10}$, which is preferred for minimizing latency associated with forward error correction (FEC) in DC interconnects, and consequently, more informed results than off-line based studies. Our measurement results indicate that by using a distributed clock with rms jitters of 0.16, 0.23, 3.02 and 6.73 ps, the corresponding phase offset tolerances are 8, 7.7, 6.4 and 1ps, respectively. Furthermore, we model and experimentally demonstrate that phase noise de-correlation due to fiber length mismatch between signal and clock path leads to a reduction of the tolerance to clock phase misalignment by 30%.

The rest of the paper is organized as follows: Section II shows our proof-of-concept experiment for clock-synchronized 51.2 GBd packet transmission. Section III shows the experimental results including measurement of bit error rate (BER), phase noise of the optical clocks and the tolerance to sampling clock phase misalignment using different source clocks. In Section IV we show the modeling and simulation of phase noise decorrelation and the resulted jitter. Section V concludes this paper.

II. EXPERIMENTS

As shown in Fig. 1, our experimental set-up consisted of a transmitter (Tx), a receiver (Rx) and a reference clock that is optically distributed to the Tx and Rx. The clock distribution section started with a 100-MHz reference clock, i.e., the source clock, which seeded a phase lock loop (PLL1, TI LMX2595) with a locking bandwidth of about 460 kHz to generate a 400 MHz digital clock signal. The 400 MHz clock was modulated on to a 1550-nm continuous wave (CW) signal using a $LiNbO_3$ Mach-Zehnder modulator (MZM 1) biased at quadrature to generate intensity modulated optical clock signals. The modulated optical clock was split by a 50:50 coupler and sent to transmitter via a 2-m fiber patchcord. The signal in the other branch was amplified by an erbium-doped fiber amplifier (EDFA) and was attenuated by a variable optical attenuator (VOA 1), which emulated the loss induced by distributing the clock to many nodes, before being launched into a standard single-mode fiber (SMF-28) that connected to the receiver. Different lengths of SMF-28, including 0.26, 0.52, 1, 1.55, 2.07 km, were used to study the impact of the de-correlation of the clock phase noise between the clock and data signal paths.

The received power of the optical clock signals was the same (-5 dBm) for both Tx and Rx. The 400-MHz optical clock was detected by a PIN photodiode followed by a transimpedance amplifier (TIA) that amplified the RF power to -7.5 dBm.

Four inexpensive commercial 100-MHz oscillators with different phase noise were used as source clocks: the Wenzel Sprinter, the Texas Instruments LMK61PD0A2, the Crystek CPRO33-100, and the Agilent E4432B. Their phase noise was measured using an RF phase noise analyzer (R&S FSWP).

At the Tx side, the detected optical clock signal was input to PLL2, which generated a 400 MHz reference clock to drive the Tx-side field programmable gate arrays (FPGA, Xilinx VCU108). PLL2 also generated a 12.8 GHz clock, which was doubled to 25.6 GHz to drive a 2:1 RF multiplexer at half clock rate to interleave the outputs of two FPGA transmitters. The two GTY transmitters on the Tx FPGA outputted 25.6 GBd on-off-

keyed packets containing 768 bits, which were electronically interleaved by the RF multiplexer to create 51.2 GBd packets. Fig. 1b and 1c shows the eye diagram of the 25.6 GBd OOK signal generated from the FPGA transmitter and the multiplexed 51.2 GBd optical OOK signal, respectively. The generated 51.2 GBd packets each contained a 1536-bit (30 ns) payload consisting of three consecutive PRBS-9 sequences each followed by a final 0 bit, with a 96-bit pattern embedded at the beginning of the third sequence for frame alignment. A 2.5 ns gap between consecutive packets was generated by padding 'zeros' between the digital packets. The 2.5 ns gap was chosen to emulate high throughout, small packet dominated DC networks. The electronic packet stream modulated an 18-dBm CW laser through a 40-GHz RF amplifier and a 40 GHz MZM, generating optical packets (with an extinction ratio of 8.7 dB) as shown in Fig. 1d.

To emulate rack-to-rack interconnects, the optical packets were transmitted through 130 m of SMF-28 before detection by a 50 GHz bandwidth photodiode at the Rx side (purple background). A variable optical attenuator (VOA2) was used to attenuate the optical power for BER measurement. The detected signal was amplified by a 65-GHz RF amplifier before entering 1:2 RF demultiplexer driven by a 25.6 GHz clock generated by PLL 3. The demultiplexer converted the 51.2 GBd packet to two streams of 25.6 GBd packets, which were fed into the Rx FPGA for real-time BER measurement. The BER for 51.2 GBd reception was measured by averaging the BER of the two 25.6 GBd packet streams (from the first two 512-bit payloads of the 51.2 GBd packets). To claim a BER of under 10^{-10}, we captured 1.311×10^{11} bits for each BER value. Following the same configuration as the Tx, the Rx FPGA was driven by a 400 MHz reference clock generated by PLL 3, which was synchronized to PLL 2 by optical clock distribution. No DSP is implemented within the transmitter or receiver FPGA, and error detection is performed in real-time [15]. Clock phase caching [3] was not employed due to the use of a manually tuned RF phase shifter for studying the tolerance to clock phase misalignment. To minimize the error caused by the change of light propagation delay due to the temperature variation, we enclosed the fiber spools in a thermal insulation chamber (shown as the dashed line in Fig. 1a). This ensured that the clock phase remained stable over long periods of time without requiring active clock phase tracking. A calibrated 40-GHz RF phase shifter with a resolution of 175 fs was used to shift the phase of the 25.6 GHz clock that drove the demultiplexer shifter. The impact of clock phase noise decorrelation was studied by changing the fiber length in the clock transmission path.

III. EXPERIMENTAL RESULTS AND DISCUSSION

A. Phase noise of synchronization clock

Fig. 2 shows the measured phase noise of the distributed 400 MHz clock signals using different source clocks. The corresponding jitter values are calculated by integrating the phase noise from 10 Hz to 10 MHz. As shown in Fig. 2, when using the Wenzel Sprinter and the TI LMK61PD0A2 clock as source oscillators, the optically distributed clock exhibits similar phase noise from 10 Hz to 10 MHz, while the TI LMK61PD0A2 has slightly worse phase noise in low frequency region (10 Hz to 100 Hz) and total jitter are of 0.16 ps and 0.23 ps, respectively. Similar high frequency phase noise performance was obtained when using the Crystek oscillator (orange curve). However, the converted optical clock showed a higher phase noise within the frequency region of 10 Hz – 2 kHz due to the higher source oscillator noise. The green line shows the phase noise of Agilent E4428B, which has a significantly higher phase noise as well as integrated jitter (6.73 ps) than the other source clocks.

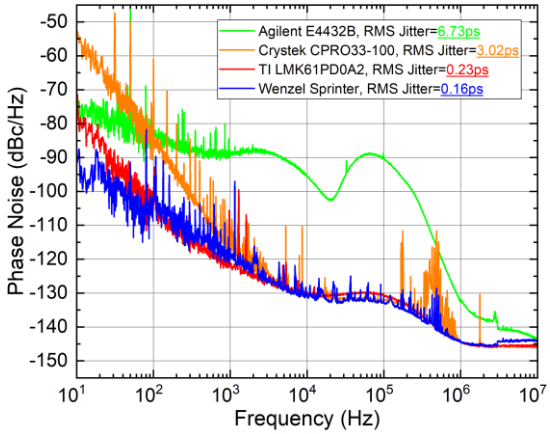

Fig. 2. Phase noise measurement of optically distributed 400 MHz clocks using different source clocks.

B. System performance

Since the Wenzel Sprinter had the lowest jitter of all the source clocks used, we measured the performance of the 51.2 GBd packets using Wenzel Sprinter as the reference clock, with the clock phase of the sampling clock optimized for the lowest BER. As shown in Fig. 3, the optical power for the BER of 10^{-10} at back-to-back was about -2.6 dBm. No penalty was measured after 130-m. The receiver sensitivity is relatively low because the used RF demultiplexer module requires a high RF input power of approximately 0 dBm. The receiver sensitivity

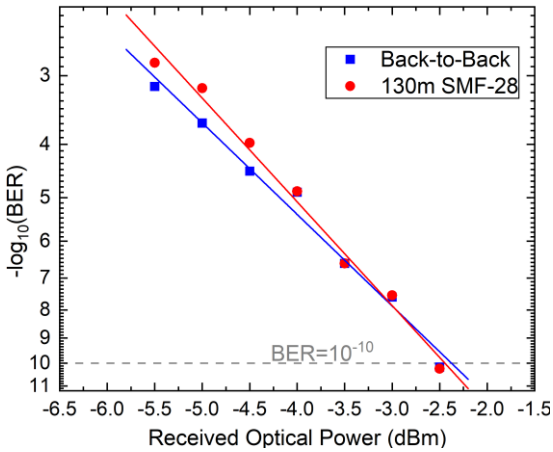

Fig. 3. BER characterization (Red circles: Back-to-Back, Blue squares: after 130 m SMF-28).

can be improved by using transceivers designed and optimized for 50 GBd signaling instead of an RF multiplexer.

Next, we offset the phase of the sampling clock and measure

the BER as shown in Fig. 4. All four source clocks were used and their tolerance to clock phase misalignment are compared, shown as different markers. The received optical power was kept at 4.5 dBm. Compared to 25.6 GBd signaling [3], the doubled baud rate in this work inherently led to a reduced timing window due to the shorter symbol period as well as a longer trialing/falling edge due to limited transceiver bandwidth. Consequently, the impact of jitter on the timing tolerance window is more prominent. An approximately 8 ps window was obtained when using the Wenzel Sprinter and the TI LMK61PD0A2 as the source oscillators. The timing tolerance window decreased to 6.4 ps and 1.0 ps when the Crystek CPRO and the Agilent E4432B respectively were used as the source oscillators, due to their larger jitter. The reduced tolerance to clock phase offset means that the system is more sensitive to optical fiber time-of-flight change due to temperature variation. For instance, assuming a thermal co-efficient of delay of 37 ps/knm/°C [12], 1 °C temperature change would lead to a 4.8 ps time-of-flight variation for 130 m SMF-28. When the clock phase offset drifts beyond the tolerance window, in addition to BER degradation, the CDR locking time will greatly increase (from sub-nanosecond to over 40 ns) as the optimal sampling phase needs to be re-calibrated [3], reducing network throughput. For networks with clock phase caching [3], reducing the timing window from 8 ps to 1 ps leads to an increase of the network overhead (or phase offset updating rate) by a factor of 8, lowering network efficiency. Our results suggest that high quality source clocks with jitter less than 3 ps should be used for clock synchronized

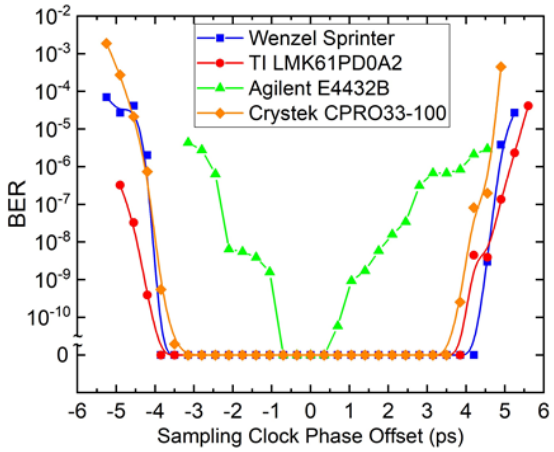

Fig. 4. Tolerable sampling clock phase offsets with different reference clocks (Red circles: LMK61PD0A2, Green triangles: Agilent E4432B, Blue squares: Wenzel Sprinter, Orange diamonds: Crystek CPRO33-100).

data center networks operating at 51.2 GBd. For 100 GBd interconnects, clock sources with sub-1.5 ps jitter are required.

As the path length difference between the clock and data link increases, the phase noise of the distributed 400 MHz clock becomes decorrelated and results in increased timing jitter, and consequently, a reduced tolerance window for clock phase misalignment. This is shown in Fig. 5. As with the previous experiment, the optical power was kept at 4.5 dBm for all the measurements. Using TI LMK61PD0A2 as the source clock, the timing tolerance window was reduced from 7.7 ps to about 5.4 ps when the length difference between the clock and data paths increased from 130 m to about 1 km, and remained the same as the difference further increased up to 1.94 km. Similarly, when the Crystek CPRO was used as the source oscillator, the timing tolerance window decreased from 6.4 ps to 4.4 ps as the fiber length difference increased from 130 m to 1 km, with no further degradation as the clock fiber length increased beyond this length difference. Our results indicate that low phase noise source oscillators offer better distance scalability for clock synchronized DC interconnects. System designers should not only minimize the absolute jitter of the clock or data signals but also the total jitter due to the decorrelated phase noise, for clock synchronized data communications.

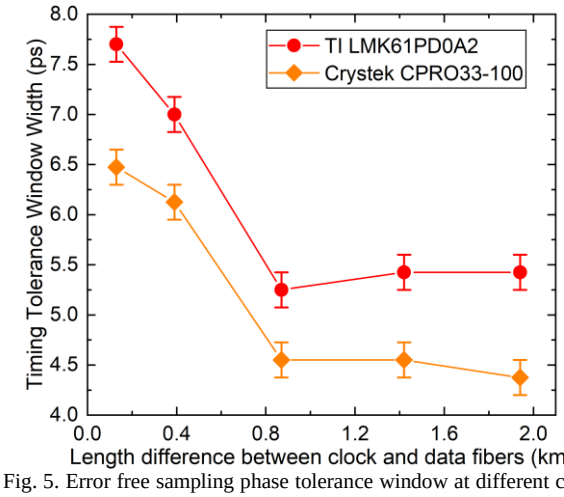

Fig. 5. Error free sampling phase tolerance window at different clock and data path length difference for Crystek and TI LMK oscillators (Red circle: TI LMK61PD0A2, Orange diamond: Crystek CPRO33-100)

IV. Modeling and Simulation of Phase Noise Decorrelation

Assuming a low speed (e.g., <10 Hz) clock phase caching is used to compensate for the initial clock phase difference between different transceiver pairs. The reduction of clock phase misalignment tolerance as the clock path increases is due to phase noise de-correlation. As this is one of the fundamental limits to the scalability of clock synchronized systems, we model and study total jitter due to phase noise decorrelation, starting with the optical clock signal:

$$E_{clk} = A(1 + mcos(2\pi f_m t + \varphi_n(t))e^{-j2\pi f_c t + \varphi_{nc}(t)} \quad (1)$$

where f_m and φ_n are the frequency and phase noise of the optical clock; f_c and φ_{nc} are the optical carrier frequency and phase. A is the field amplitude and m ($m < 1$) is the modulation index. In this case, we assume optical clock signal is detected by an ideal photodetector (i.e., ignoring thermal/shot noise and nonlinearity). Let ΔL be the length difference between clock and data fibers, the difference in the time-of-flight Δt:

$$\Delta t = \frac{\Delta L n_g}{c} \quad (2)$$

where n_g is the group refractive index ($n_g = 1.46$ in this case) and c is the light speed in vacuum. The distributed optical clock is converted to a full rate clock by multiplying by a factor of M (e.g., $M = 64$ in our experiment for multiplying 400 MHz

clock to 25.6 GHz), resulting in a relative phase noise of:

$$\varphi_r(t,\Delta L) = M(\varphi_n(t) - \varphi_n(t - \tfrac{\Delta L n_g}{c})) \quad (3)$$

To obtain the power spectrum density (PSD) of the relative phase noise, we first calculate the autocorrelation of $\varphi_r(t,\Delta L)$:

$$R_{\varphi_r\varphi_r}(\tau,\Delta L) = M^2(2R_{\varphi_n\varphi_n}(\tau) - R_{\varphi_n\varphi_n}\left(\tau - \tfrac{\Delta L n_g}{c}\right) - R_{\varphi_n\varphi_n}\left(\tau + \tfrac{\Delta L n_g}{c}\right)) \quad (4)$$

where $R_{\varphi_n\varphi_n}$ represents the autocorrelation of $\varphi_n(t)$. According to Wiener–Khinchin theorem [16], autocorrelation and power spectrum density is a Fourier Transform pair. Therefore, the Fourier transform of $R_{\varphi_n\varphi_n}$ becomes the PSD of the clock phase noise $\varphi_{PSD}(f)$ (which can be estimated using a phase noise analyzer). By taking the Fourier transform of equation (4), the PSD of the relative phase noise as follows:

$$PSD_{\varphi_r}(f,\Delta L) = M^2 PSD_{\varphi}(f)[2 - 2\cos(2\pi f \tfrac{\Delta L n_g}{c})] \quad (5)$$

With the PSD of the relative phase noise, one can calculate the jitter by integrating the PSD_{φ_r} over the frequency range. Here we use 10 Hz to 10 MHz in consideration of a clock phase cached system where the sub-10Hz phase noise is tracked and compensated. The 10 MHz upper bound assumes a PLL output filter bandwidth of 10 MHz:

$$\sigma_{rms}(\Delta L) = \frac{\sqrt{2\int_{10\,Hz}^{10\,MHz} PSD_{\varphi_r}(f,\Delta L)df}}{2\pi f_m} \quad (6)$$

Equation (5) indicates that the delay between the clock and data path essentially acts as a delay interferometer for the phase noise. Using equation (5) and the measured phase noise of the source clocks, we calculate the resulted PSD of the relative phase noise, as shown in Fig. 6, using Crystek as the source clock. Given that the loop bandwidth of PLL used in this experiment (TI LMX2595) is 460 kHz we assume the PSD_{φ_r} above 460 kHz is completely de-correlated, thus $PSD_{\varphi_r}(f > 460\,kHz)$ is simply $2\varphi_{PSD}(f > 460\,kHz)$. As with a delay interferometer, the high frequency phase noise firstly becomes decorrelated as ΔL increases, resulting in more dips due to destructive interference of the corresponding frequency components due to a phase difference of π between the clock and data path signals. This can be observed from the grey and black curves in Fig. 6, which show the PSD of the relative phase noise with a path length difference of 0.5 and 2 km, respectively. The orange curve shows the calculated phase noise of 25.6 GHz signal by scaling up the phase noise of the optically distributed clock. As shown in Fig. 6, when the fiber length difference is less than 2 km, the low frequency phase noise is largely correlated and therefore, shows a significantly lower phase noise PSD at the <10 kHz frequency region. The high frequency (100 kHz – 10 MHz) are largely uncorrelated, resulting in more dips and increased PSD.

As ΔL tends to infinity the phase noise in the clock and signal paths becomes fully decorrelated and the total phase noise is simply a power summation of two independent phase noises of the same PSD, resulting in a 3 dB higher PSD over the whole frequency range. The lower limit of the rms jitter in this case is determined by the PLL loop bandwidth as phase noise beyond this region is completely de-correlated. In practice, the lower limit of the rms jitter also depends on white thermal noise in the photodetection process, additive phase noise from PLL loop components as well as the amplified and source spontaneous emission noise. Nevertheless, the jitter is dominated by the decorrelated phase noise in this study, therefore, (5) provides a good estimation of the relative phase noise using an equation of low complexity.

Fig. 7a and Fig. 7b show the calculated jitter for different fiber delay length/time using all four source clocks. As the delay increases, the total jitter increases quickly due to the fast de-correlation of the high frequency phase noise. With a delay of a few hundred meters, the total jitter reaches a stable value for clocks with low low-frequency phase noise, e.g., TI and Wenzel cases as shown in Fig. 7a. Generally, lower phase noise clock results in lower total jitter after the decorrelation and they reach the saturated jitter value with a lower delay length (about 200 m) as shown in Fig. 2 because of their significantly lower low-frequency phase noise, i.e., the contribution of the low frequency phase noise to jitter becomes insignificant. With a high phase noise clock, such as the Agilent E4432B case (green curve) shown in Fig. 7b, the total jitter is higher, reaching saturation after about 800 m delay due to the higher phase noise across the whole frequency region.

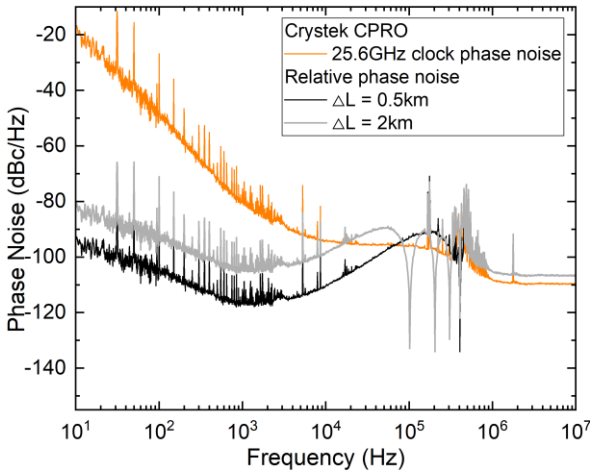

Fig. 6. De-correlated PN of 25.6 GHz clock using Crystek CPRO. Orange: 25.6 GHz clock phase noise (multiplied up by PLL); Black: relative phase noise when $\Delta L = 0.5km$; Grey: relative phase noise when $\Delta L = 2km$

These results agree with our experimental results shown in Fig. 5, where the fast increase of the total jitter resulted a reduction of the timing window tolerance before reaching a stable value. Note that the Wenzel Sprinter and TI LMK61PD0A2 cases show similar performance, which agrees with our sampling phase tolerance results shown in Fig. 4.

The oscillation of the total jitter when using the Crystek CPRO is due to the deterministic jitter of the source clock, which is exhibited as spurs in the phase noise. With 490 m (or 2401 ns) delay, the destructive interference of the clock and data phase noise cancelled the 416.7 kHz spur and resulted in a total jitter of 0.27 ps, and this appears periodically as the fiber delay further increases. Fig. 8 compares the relative phase noise of the Crystek CPRO with delay lengths of 0.26 km and 0.49 km, with a total jitter of 0.33 ps and 0.27 ps. This can also

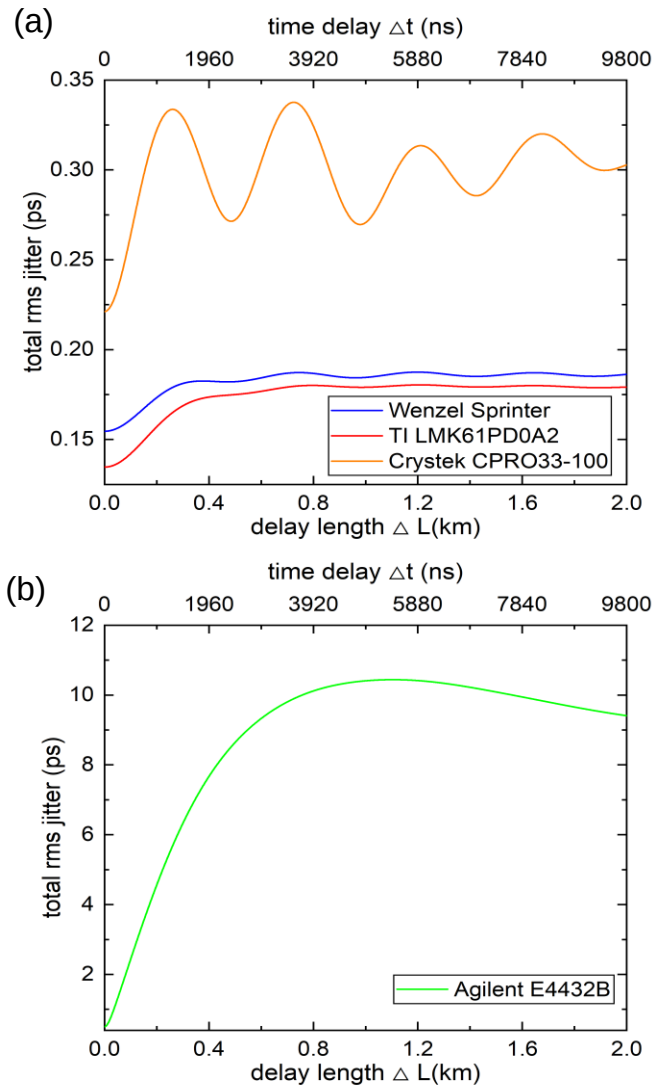

Fig. 7. Calculated total jitter for different delay (ΔL and Δt) for 25.6 GHz clock signal using a) LMK61PD0A2, Wenzel Sprinter and Crystek CPRO33-100. b) using Agilent E4432B.

be explained by looking at equation (5). By setting f to 416.7 kHz, $\varphi_{rPSD}(416.7kHz, \Delta L)$ reaches a minimum when $\Delta L \approx 0.49km$ and reaches maximum when $\Delta L \approx 0.26km$.

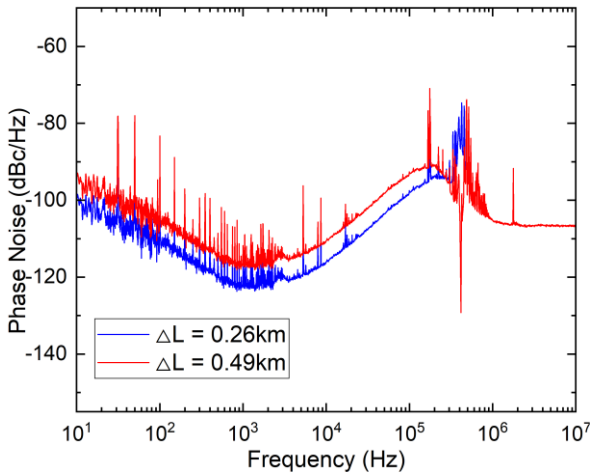

Fig. 8. The relative phase noise measurement when using Crystek CPRO at $\Delta L = 0.26km$ and $\Delta L = 0.49km$.

V. Conclusion

We demonstrated clock synchronous transmission of real-time optical packets at 51.2 GBd and benchmarked the system performance by using different inexpensive commercial oscillators as the system source clock. We demonstrated that a low phase noise source oscillator is crucial in clock synchronized systems as it offers a larger tolerance window for clock phase misalignment (8.0 ps in our experiment), which is necessary for minimizing overhead of an optically-switched DCN with clock phase caching. Our experimental results show that the phase noise de-correlation, which is caused by the path length difference between the clock and data fiber, could lead to a 30% reduction in the clock phase misalignment tolerance with a path length difference of more than 1 km. Based on the model and the experimental results, we point out that the phase noise of the source clock would limit the scalability of clock synchronized DCN for systems operating at >50 GBd. Low phase noise clock synchronization must be achieved to enable high-throughput, sub-nanosecond clock recovery time for optical switched data center interconnection.

DATA AVAILABILITY STATEMENT

The data that supports the figures within this paper is available from the UCL Research Data Repository (https://doi.org/10.5522/04/19074233), which is hosted by FigShare.

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," *Journal of Lightwave Technology*, vol. 33, no. 5, pp. 1117–1125, 2015.

[2] 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*, 2020, pp. 782–797.

[3] 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 centres using clock phase caching," *Nature Electronics*, vol. 3, pp. 426–433, 2020.

[4] K. Clark and Z. Liu, "Analytically Modeling the Clock Phase Caching Approach to Clock and Data Recovery," *Journal of Lightwave Technology*, accepted.

[5] R. S. Sohanpal, K. A. Clark, B. J. Puttnam, Y. Awaji, N. Wada, P. Bayvel, and Z. Liu, "Clock and data recovery-free data communications enabled by multi-core fiber with low thermal sensitivity of skew," *Journal of Lightwave Technology*, vol. 38, no. 7, pp. 1636–1643, 2020.

[6] 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," *Journal of Lightwave Technology*, vol. 38, no. 9, pp. 2703–2709, 2020.

[7] K. Shi, S. Lange, I. Haller, D. Cletheroe, R. Behrendt, B. Thomsen, F. Karinou, K. Jozwik, P. Costa and H. Ballani, "System demonstration of nanosecond wavelength switching with burst-mode PAM4 transceiver," *45th European Conference on Optical Communication (ECOC 2019)*, 2019, pp. 1-4.

[8] IEEE Standard, Media access control parameters, physical layers, and management parameters for 200 Gb/s and 400 Gb/s operation.

[9] J. Wei, Q. Cheng, R. V. Penty, I. H. White and D. G. Cunningham, "400 Gigabit Ethernet using advanced modulation formats:

Performance, complexity, and power dissipation," in *IEEE Communications Magazine*, vol. 53, no. 2, pp. 182-189, 2015.

[10] X. Pang, W. Hu, G. Jacobsen, S. Popov, J. Chen, O. Ozolins, R. Lin, L. Zhang, A. Udalcovs, L. Xue, R. Schatz, U. Westergren, and S. Xiao, "200 gbps/lane im/dd technologies for short reach optical interconnects," *Journal of Lightwave Technology*, vol. 38, no. 2, pp. 492–503, 2020.

[11] Y. Matsui, R. Schatz, D. Che, F. Khan, M. Kwakernaak, and T. Sudo, "Low-chirp isolator-free 65-GHz bandwidth directly modulated lasers," *Nature Photonics*, vol. 15, pp. 59–63, 2021.

[12] R. Slavík, G. Marra, E. N. Fokoua, N. Baddela, N. V. Wheeler, M. Petrovich, F. Poletti, and D. J. Richardson, "Ultralow thermal sensitivity of phase and propagation delay in hollow core optical fibres," *Scientific Reports*, vol. 5, 2015.

[13] A. Saljoghei, H. Yuan, V. Mishra, M. Enrico, N. Parsons, C. Kochis, P. Dobbelaere, D. Theodoropoulos, D. Pnevmatikatos, D. Syrivelis, A. Reale, T. Hayashi, T. Nakanishi and G. Zervas, "MCF-SMF Hybrid Low-Latency Circuit-Switched Optical Network for Disaggregated Data Centers," in *Journal of Lightwave Technology*, vol. 37, no. 16, pp. 4017-4029, 2019.

[14] Z. Zhou, K. Clark, C. Deakin, P. Laccotripes and Z. Liu, "Clock Synchronized Transmission of 51.2 GBd Optical Packets for Optically Switched Data Center Interconnects," *2021 Optical Fiber Communications Conference and Exhibition (OFC)*, 2021, pp. 1-3.

[15] 2019. [Online]. Available: https://www.xilinx.com/content/dam/xilinx/support/documentation/boards_and_kits/vcu108/ug1066-vcu108-eval-bd.pdf

[16] L. Cohen, "Generalization of the Wiener-Khinchin theorem," in *IEEE Signal Processing Letters*, vol. 5, no. 11, pp. 292-294, 1998.